Available online at www.sciencedirect.com

SciVerse ScienceDirect

ADVANCES IN
Mathematics

Advances in Mathematics 231 (2012) 1731–1772

www.elsevier.com/locate/aim

A measure-theoretic approach to the theory of dense hypergraphs

Gábor Elek [a,*], Balázs Szegedy [b]

[a] Alfred Renyi Institute of the Hungarian Academy of Sciences, POB 127, H-1364, Budapest, Hungary
[b] University of Toronto, Department of Mathematics, St George St. 40, Toronto, ON, M5R 2E4, Canada

Received 21 May 2007; accepted 29 June 2012
Available online 11 August 2012

Communicated by the Managing Editors of AIM

Abstract

In this paper we develop a measure-theoretic method to treat problems in hypergraph theory. Our central theorem is a correspondence principle between three objects: an increasing hypergraph sequence, a measurable set in an ultraproduct space and a measurable set in a finite dimensional Lebesgue space. Using this correspondence principle we build up the theory of dense hypergraphs from scratch. Along these lines we give new proofs for the Hypergraph Removal Lemma, the Hypergraph Regularity Lemma, the Counting Lemma and the Testability of Hereditary Hypergraph Properties. We prove various new results including a strengthening of the Regularity Lemma and an Inverse Counting Lemma. We also prove the equivalence of various notions for convergence of hypergraphs and we construct limit objects for such sequences. We prove that the limit objects are unique up to a certain family of measure preserving transformations. As our main tool we study the integral and measure theory on the ultraproduct of finite measure spaces which is interesting on its own right.

MSC: primary 05C99; secondary 82B99

Keywords: Hypergraphs; Regularity lemma; Limit objects; Property testing

Contents

* Corresponding author.
E-mail address: elek@renyi.hu (G. Elek).

0001-8708/$ - see front matter
doi:10.1016/j.aim.2012.06.022

1. Introduction

The so-called Hypergraph Regularity Lemma (Rödl–Skokan [15], Rödl–Schacht [14], Gowers [5], later generalized by Tao [18]) is one of the most exciting result in modern combinatorics. It exists in many different forms, strength and generality. The main message in all of them is that every k-uniform hypergraph can be approximated by a structure which consists of boundedly many random-looking (quasi-random) parts for any given error ϵ. Another common feature of these theorems is that they all come with a corresponding counting lemma [13] which describes how to estimate the frequency of a given small hypergraph from the quasi-random approximation of a large hypergraph. One of the most important applications of this method is that it implies the Hypergraph Removal Lemma (first proved by Nagle et al. [13]) and by an observation of Solymosi [17] it also implies Szemerédi's celebrated theorem on arithmetic progressions in dense subsets of the integers even in a multidimensional setting.

In this paper we present an analytic approach to the subject. First, for any given sequence of hypergraphs we associate the so-called ultralimit hypergraph, which is a measurable hypergraph in a large (non-separable) probability measure space. The ultralimit method enables us to convert theorems of finite combinatorics to measure theoretic statements on our ultralimit space. In the second step, using separable factors we translate these measure-theoretic theorems to well-known results on the more familiar Lebesgue spaces.

The paper is built up in a way that these two steps are compressed into a **correspondence principle** between the following three objects.

1. An increasing sequence $\{H_i\}_{i=1}^{\infty}$ of k-uniform hypergraphs.
2. The ultraproduct hypergraph $\mathbf{H} \subseteq \mathbf{X}^k$, where $\mathbf{X}$ is the ultraproduct of the vertex sets.
3. A measurable subset $W \subseteq [0, 1]^{2^k-1}$.

Using this single correspondence principle we are able to prove several results in hypergraph theory. The next list is a summary of some of these results.

1. **Removal lemma.** We prove the Hypergraph Removal Lemma directly from Lebesgue's density theorem applied for the set $W \subseteq [0, 1]^{2^k-1}$. In a nutshell, we convert the original removal lemma into the removal of the non-density points from W which is a 0-measure set. (Theorem 1.)
2. **Regularity lemma.** We deduce the Hypergraph Regularity Lemma from a certain finite box approximation of W in L_1. To be more precise, W is approximated by a set which is the disjoint union of finitely many direct product sets in $[0, 1]^{2^k-1}$. (Theorem 2.)
3. **Limit object.** We prove that W serves as a limit object for hypergraph sequences $\{H_i\}_{i=1}^{\infty}$ which are convergent in the sense that the densities of every fixed hypergraph F converge. Limits of k-uniform hypergraphs can also be represented by $2^k - 2$ variable measurable functions $w : [0, 1]^{2^k-2} \to [0, 1]$ such that the coordinates are indexed by the proper non-empty subsets of $\{1, 2, \ldots, k\}$ and w is invariant under the induced action of S_k on the coordinates. This generalizes a theorem by Lovász and Szegedy. (Theorem 7.) Note that a similar limit object was defined by Kallenberg, in the context of exchangeable arrays [8].
4. **Sampling and concentration.** Even tough W is a measurable set, it makes sense to talk about random samples from W which are ordinary hypergraphs. We prove concentration results for this sampling process. The sampling processes give rise to random hypergraph models which are interesting on their own right. (Theorems 11 and 12.)
5. **Testability of hereditary properties.** We give a new proof for the testability of hereditary hypergraph properties. (This was first proved for graphs by Alon-Shapira and later for hypergraphs by Rödl–Schacht.) The key idea is based on a modified sampling process from the limit object W that we call "hyperpartition sampling". This creates an overlay of samples from W and the members of the sequence $\{H_i\}_{i=1}^{\infty}$ such that expected Hamming distance of H_i and the corresponding sample is small. (Theorem 8.)
6. **Regularity as compactness.** We formulate a strengthening of the Hypergraph Regularity Lemma which puts the regularity in the framework of compactness. Roughly speaking this theorem says that every increasing hypergraph sequence has a subsequence which converges in a very strong (structural) sense. Here we introduce the notion of **strong convergence**. (Theorem 4.)
7. **Distance notions.** We introduce several distance notions between hypergraph limit objects (and hypergraphs) and we analyze their relationship. (Theorem 10.)
8. **Uniqueness.** We prove the uniqueness of the limit object up to a family of measure preserving transformations on $[0, 1]^{2^k-1}$. This generalizes a result of Borgs–Chayes–Lovász from graphs to hypergraphs. (Theorem 9.)
9. **Counting Lemma.** We prove that the structure of regular partitions determine the subhypergraph densities. (Theorem 13 and Corollary 4.2.)

10. **Equivalence of convergence notions.** We prove that convergence and strong convergence are equivalent. For technical reasons we introduce a third convergence notion which is a slight variation of strong convergence and we call it **structural convergence**. This is also equivalent with the other two notions. The third notation enables us to speak about structural limit objects which turns out to be the same as the original limit object. (Theorem 14.)
11. **Inverse counting lemma.** Using the equivalence of convergence notions we obtain that if two hypergraphs have similar sub-hypergraph densities then they have similar regular partitions. In other words this means that regular partitions can be tested by sampling small hypergraphs. (Corollary 4.1.)

Remark. In our proofs we use the Axiom of Choice. However, Gödel in his seminal work *The Consistency of the Axiom of Choice and the Generalized Continuum Hypothesis with the Axioms of Set Theory* proved that (see also [4]): if Γ is an arithmetical statement and Γ is provable in **ZF** with the Axiom of Choice then Γ is provable in **ZF**. In fact, Gödel gave an algorithm to convert a formal **ZFC**-proof of an arithmetical statement to a **ZF**-proof. An arithmetical statement is a statement in the form of

$$(\mathbf{Q}_1 x_1 \mathbf{Q}_2 x_2 \dots \mathbf{Q}_k x_k) P(x_1, x_2, \dots, x_k),$$

where the $\mathbf{Q}_i$'s are existential or universal quantifiers and the relation $P(x_1, x_2, \dots, x_k)$ can be checked by a Turing machine in finite time. The reader can convince himself that the Hypergraph Removal Lemma, The Hypergraph Regularity Lemma, the Counting Lemma and the Inverse Counting Lemma are all arithmetical statements.

2. Preliminaries

2.1. Homomorphisms, convergence and completion of hypergraphs

Let $\mathcal{H}_k$ denote the set of isomorphism classes of finite k-uniform hypergraphs. For an element $H \in \mathcal{H}_k$ we denote the vertex set by $V(H)$ and the edge set by $E(H)$. In this paper we view a k-uniform hypergraph H on the vertex set V as a subset of V^k without having repetitions in the coordinates and being invariant under the action of the symmetric group S_k. Let $v_1, v_2, v_3, \dots, v_{|V|}$ be the elements of V. Then an edge $E \in E(H)$ is a subset of k-elements $\{v_{i_1}, v_{i_2}, \dots, v_{i_k}\} \subset V$ such that $(v_{i_1}, v_{i_2}, \dots, v_{i_k}) \in H$. If L is a family of edges in H, then $\hat{L}$ denote the set of elements $(v_{i_1}, v_{i_2}, \dots, v_{i_k}) \in H$ such that $\{v_{i_1}, v_{i_2}, \dots, v_{i_k}\} \in L$.

Definition 2.1. A **homomorphism** between two elements $F, H \in \mathcal{H}_k$ is a map $f : V(F) \mapsto V(H)$ which maps edges of F into edges of H. We denote by $\hom(F, H)$ the number of homomorphisms from F to H and by $\hom^0(F, H)$ the number of injective homomorphisms. An **induced homomorphism** is a map $f : V(G) \mapsto V(H)$ which maps edges to edges and non-edges to non-edges (see [3]).

Note that in the definition of hom the map $V(F) \mapsto V(H)$ does not have to be injective but the definition implies that it is injective if we restrict it to any edge of F. There is a simple inclusion–exclusion type formula which computes $\hom^0$ from hom. To state this we will need some more definitions.

Let $\mathcal{P} = \{P_1, P_2, \dots, P_s\}$ be a partition of $V(F)$ and let $f : V(F) \mapsto \mathcal{P}$ be the function which maps each vertex to its partition set. We define a hypergraph $F(\mathcal{P})$ whose vertex set is $\mathcal{P}$ and the edge set is $f(E(F))$. Note that $F(\mathcal{P})$ is a k-uniform hypergraph if and only if every partition set intersect every edge in at most 1 element. We define the height $h(\mathcal{P})$ of $\mathcal{P}$ as $|V(F)| - |\mathcal{P}|$.

Lemma 2.1. *If F and H are k-uniform hypergraphs then*

$$\hom(F, H) = \sum_{\mathcal{P}} \hom^0(F(\mathcal{P}), H)$$

and

$$\hom^0(F, H) = \sum_{\mathcal{P}} (-1)^{h(\mathcal{P})} \hom(F(\mathcal{P}), H)$$

where $\mathcal{P}$ runs through all partitions of $V(F)$ and $\hom(F(\mathcal{P}), H)$ and $\hom^0(F(\mathcal{P}), H)$ are defined to be 0 if $F(\mathcal{P})$ is not k-uniform.

Proof. The first equation is obvious from the definitions. It implies that for any partition $\mathcal{P}$ we have that

$$\hom(F(\mathcal{P}), H) = \sum_{\mathcal{P}' \leq \mathcal{P}} \hom^0(F(\mathcal{P}'), H)$$

where the sum runs through all partitions $\mathcal{P}'$ such that $\mathcal{P}$ is a refinement of $\mathcal{P}'$. The inversion formula for the partition lattice yields the second equation. □

Now we are ready to prove the next lemma.

Lemma 2.2. *If $H_1, H_2 \in \mathcal{H}_k$ are two hypergraphs such that $\hom(F, H_1) = \hom(F, H_2)$ for every element $F \in \mathcal{H}_k$ then H_1 and H_2 are isomorphic.*

Proof. Lemma 2.1 implies that $\hom^0(F, H_1) = \hom^0(F, H_2)$ for all hypergraphs $F \in \mathcal{H}_k$. In particular $\hom^0(H_1, H_2) = \hom^0(H_1, H_1) > 0$ and $\hom^0(H_2, H_1) = \hom^0(H_2, H_2) > 0$ which implies that $|V(H_1)| = |V(H_2)|$ and $|E(H_1)| = |E(H_2)|$. We obtain that every injective homomorphism from H_1 to H_2 is an isomorphism. Since such a homomorphism exists the proof is complete. □

The next two definitions will be crucial.

Definition 2.2. The **homomorphism density** $t(F, H)$ denotes the probability that a random map $f : V(F) \mapsto V(H)$ is a homomorphism. It can also be defined by the equation

$$t(F, H) = \frac{\hom(F, H)}{|V(H)|^{|V(F)|}}.$$

We also define $t_{\text{ind}}(F, G)$ which is the probability that a random map $f : V(F) \mapsto V(H)$ is an induced homomorphism. Finally $t^0_{\text{ind}}(F, H)$ denotes the probability that a random injective map is an induced homomorphism.

Definition 2.3. A t-fold **equitable blowup** of a hypergraph $H \in H_k$ is a hypergraph H' which is obtained by replacing each vertex of H by t new vertices and each edge of H by a complete k-partite hypergraph on the corresponding new vertex sets.

It is clear that if H' is a t-fold equitable blowup of H then $\hom(F, H') = \hom(F, H)t^{|V(F)|}$ and consequently $t(F, H) = t(F, H')$. The next lemma shows that hypergraphs from $\mathcal{H}_k$ are "essentially" separated by homomorphism densities except that equitable blowups of a hypergraph cannot be separated.

Lemma 2.3. *Let $H_1, H_2 \in \mathcal{H}_k$ be two hypergraphs and assume that $t(F, H_1) = t(F, H_2)$ for every $F \in \mathcal{H}_k$. Then there exists a $H \in \mathcal{H}_k$ which is an equitable blowup of both H_1 and H_2.*

Proof. Let H_1' be the $|V(H_2)|$-fold equitable blowup of H_1 and let H_2' be the $|V(H_1)|$-fold equitable blowup of H_2. Then

$$|V(H_1')| = |V(H_2')| = |V(H_1)||V(H_2)|$$

and $t(F, H_1') = t(F, H_2')$ for every $F \in \mathcal{H}_k$. We obtain that

$$\hom(F, H_1') = t(F, H_1')|V(H_1')|^{|V(F)|} = t(F, H_2')|V(H_2)'|^{|V(F)|} = \hom(F, H_2')$$

for every $F \in \mathcal{H}_k$. By Lemma 2.2 the proof is complete. □

The previous lemma motivates the following definition.

Definition 2.4. Two hypergraphs $H_1, H_2 \in \mathcal{H}_k$ will be called **density equivalent** if there exists $H \in \mathcal{H}_k$ which is an equitable blowup of both H_1, H_2 or equivalently, by Lemma 2.3, $t(F, H_1) = t(F, H_2)$ for every $F \in \mathcal{H}_k$.

Homomorphism densities can be used to define two convergence notions on the set $\mathcal{H}_k$ which are slight variations of each other.

Definition 2.5. A hypergraph sequence $\{H_i\}_{i=1}^{\infty}$ in $\mathcal{H}_k$ is called **convergent** if

$$\lim_{i\to\infty} t(F, H_i)$$

exists for every $F \in \mathcal{H}_k$. We say that $\{H_i\}_{i=1}^{\infty}$ is **increasingly convergent** if it is convergent and

$$\lim_{i\to\infty} |V(H_i)| = \infty.$$

Both convergence notions lead to a completion of the set $\mathcal{H}_k$. We denote the first completion by $\bar{\mathcal{H}}_k$ and the second one by $\hat{\mathcal{H}}_k$. These two spaces are very closely related to each other. It will turn out that $\bar{\mathcal{H}}_k$ is arc-connected whereas $\hat{\mathcal{H}}_k$ is the union of $\mathcal{H}_k$ with the discrete topology and $\bar{\mathcal{H}}_k$. In the space $\hat{\mathcal{H}}_k$ the set $\bar{\mathcal{H}}_k$ behaves as a "boundary" for the set $\mathcal{H}_k$. An advantage of the set $\hat{\mathcal{H}}_k$ is that it directly contains the familiar set $\mathcal{H}_k$ of hypergraphs. A disadvantage of $\hat{\mathcal{H}}_k$ is that it is not connected. On the other hand $\bar{\mathcal{H}}_k$ is connected and k-uniform hypergraphs are represented in it up to dense equivalence. In this paper we focus only on $\bar{\mathcal{H}}_k$ so we give a precise definition only of this space.

Let δ be the following metric on $\mathcal{H}_k$. For two elements $H_1, H_2 \in \mathcal{H}_k$ we define $\delta(H_1, H_2)$ as the infimum of the numbers $\epsilon \geq 0$ for which $|t(F, H_1) - t(F, H_2)| \leq \epsilon$ holds for all $F \in \mathcal{H}_k$ with $|V(F)| \leq 1/\epsilon$. Two hypergraphs have δ-distance zero if and only if they are density equivalent. We denote the completion of this metric space by $\bar{\mathcal{H}}_k$.

The elements of the space $\bar{\mathcal{H}}_k$ have many interesting representations. We give here one which is the most straightforward. Let $\mathcal{M}_k$ denote the compact space $[0, 1]^{\mathcal{H}_k}$. Every graph $H \in \mathcal{H}_k$ can be represented as a point in $\mathcal{M}_k$ by the sequence $T(H) = \{t(F, H)\}_{F\in\mathcal{H}_k}$. By Lemma 2.3 the point set $T(\mathcal{H}_k)$ represents the density equivalence classes of k-uniform hypergraphs. The closure of $T(\mathcal{H}_k)$ in $\mathcal{M}_k$ is a representation of $\bar{\mathcal{H}}_k$. This representation shows immediately that $\bar{\mathcal{H}}_k$ is compact since it is a closed subspace of the compact space $\mathcal{M}_k$. To see that $\bar{\mathcal{H}}_k$ is arc-connected requires some more effort, but it will follow easily from one of our results in this paper. (Theorem 7.)

An important feature of the space $\bar{\mathcal{H}}_k$ is that it makes sense to talk about homomorphism densities of the form $t(F, X)$ if $X \in \bar{\mathcal{H}}_k$ and $F \in \mathcal{H}_k$.

We will denote by $[n]$ the set $\{1, 2, \dots, n\}$. For a subset $B \subset [k]$, $r(B)$ will stand for the non-empty subsets of B. Similarly, $r([n], k)$ will denote the set of all non-empty subsets of $[n]$ having size at most k. If K is a hypergraph on $[n]$ and $H \subset V^{[k]}$ is a k-uniform hypergraph then $T(K, H) \subset V^{[n]}$ denotes the (K, H)-**homomorphism set**, where $(x_1, x_2, \dots, x_n) \in T(K, H)$ if $1 \to x_1, 2 \to x_2, \dots, n \to x_n$ defines a homomorphism. Clearly $|T(K, H)| = hom(K, H)$. For a subset $E \subset [n]$, $|E| = k$ let $P_E : V^{[n]} \to V^E$ be the natural projection and $P_{s_E} : V^{[k]} \to V^E$ be the natural bijection associated to a bijective map $s_E : [k] \to E$. Then it is easy to check that

$$T(K, H) = \bigcap_{E \in E(K)} P_E^{-1}\left(P_{s_E}(H)\right).$$

Similarly, $T_{ind}(K, H) \subset V^{[n]}$ denotes the (K, H)-**induced homomorphism set**, where $(x_1, x_2, \dots, x_n) \in T_{ind}(K, H)$ if $1 \to x_1, 2 \to x_2, \dots, n \to x_n$ defines an induced homomorphism. Then

$$T_{ind}(K, H) = \bigcap_{E \in E(K)} P_E^{-1}\left(P_{s_E}(H)\right) \cap \bigcap_{E' \in E(K)^c} P_{E'}^{-1}\left(P_{s_{E'}}(H^c)\right),$$

where H^c denotes the complement of H in the complete hypergraph on the set V. A simple inclusion–exclusion argument shows that if a hypergraph sequence $\{H_i\}_{i=1}^{\infty}$ is convergent, then for any k-uniform hypergraph F the sequence $\{t_{ind}(F, H_i)\}_{i=1}^{\infty}$ is convergent as well.

2.2. *The Removal and the Regularity Lemmas*

First we state the Removal Lemma.

Theorem 1 (*Hypergraph Removal Lemma*). *For every k-uniform hypergraph K and constant $\epsilon > 0$ there exists a number $\delta = \delta(K, \epsilon)$ such that for any k-uniform hypergraph H on the node set X with $t(K, H) < \delta$ there is a subset L of $E(H)$ with $L \le \epsilon \binom{|X|}{k}$ such that $t(K, H \setminus \hat{L}) = 0$. ([5,7,13,18].)*

Now let us turn to the regularity lemma. Let X be a finite set, then $K_r(X) \subset X^r$ denotes the complete r-uniform hypergraph on X. An l-**hyperpartition** $\mathcal{H}$ is a family of partitions $K_r(X) = \cup_{j=1}^{l} P_r^j$, where P_r^j is an r-uniform hypergraph, for $1 \le r \le k$. We call $\mathcal{H}$ δ-**equitable** if for any $1 \le r \le k$ and $1 \le i < j \le l$:

$$\frac{\| P_r^i| - |P_r^j \|}{|K_r(X)|} < \delta.$$

An l-hyperpartition $\mathcal{H}$ induces a partition on $K_k(X)$ in the following way.

- Two elements $\underline{a}, \underline{b} \in K_k(X)$, $\underline{a} = \{a_1, a_2, \dots, a_k\}$, $\underline{b} = \{b_1, b_2, \dots, b_k\}$ are equivalent if there exists a permutation $\sigma \in S_k$ such that for any subset $A = \{i_1, i_2, \dots, i_{|A|}\} \subset [k]$, $\{a_{i_1}, a_{i_2}, \dots, a_{i_{|A|}}\}$ and $\{b_{\sigma(i_1)}, b_{\sigma(i_2)}, \dots, b_{\sigma(i_{|A|})}\}$ are both in the same $P_{|A|}^j$ for some $1 \le j \le l$.

It is easy to see that this defines an equivalence relation and thus it results in a partition $\cup_{j=1}^{t} C_j$ of $K_k(X)$ into $\mathcal{H}$-**cells**. A **cylinder intersection** $L \subset K_r(X)$ is an r-uniform hypergraph defined in the following way. Let $B_1, B_2, \ldots, B_r$ be $(r-1)$-uniform hypergraphs on X; then an r-edge $\{a_1, a_2, \ldots, a_r\}$ is in L if there exists a permutation $\tau \in S_r$ such that

$$\{a_{\sigma(1)}, a_{\sigma(2)}, \ldots, a_{\sigma(i-1)}, a_{\sigma(i+1)}, \ldots, a_{\sigma(r)}\} \in B_i \quad \text{for any } 1 \le i \le r.$$

As in the graph case, we call an r-uniform hypergraph G ϵ-regular if

$$\left| \frac{|G|}{|K_r(X)|} - \frac{|G \cap L|}{|L|} \right| \le \epsilon,$$

for each cylinder intersection L, where $|L| \ge \epsilon |K_r(X)|$. Now we are ready to state the Hypergraph Regularity Lemma for k-uniform hypergraphs (see [5,7,14,15,18]).

Theorem 2 (*Hypergraph Regularity Lemma*)**.** *Let fix a constant $k > 0$. Then for any $\epsilon > 0$ and function $F : \mathbb{N} \to (0, 1)$ there exist constants $c = c(\epsilon, F)$ and $N_0(\epsilon, F)$ such that if H is a k-uniform hypergraph on a set X, $|X| \ge N_0(\epsilon, F)$, then there exists an $F(l)$-equitable l-hyperpartition $\mathcal{H}$ for some $1 < l \le c$ such that*

- *Each P_j^r is $F(l)$-regular.*
- *$|H \triangle T| \le \epsilon \binom{|X|}{k}$ where T is the union of some $\mathcal{H}$-cells.*

2.3. Combinatorial structures

In this subsection we introduce some further definitions about hyperpartitions. Let $\mathcal{H} = \{P_r^j\}$ be an l-hyperpartition on a set X where $1 \le j \le l$ and $1 \le r \le k$. We shall need the notion of a **directed $\mathcal{H}$-cell**. Let $f : r([k]) \mapsto [l]$ be an arbitrary function. Then the directed cell with coordinate f is the set of ordered k-tuples $(x_1, x_2, \ldots, x_k) \in X^k$ such that $\{x_{i_1}, x_{i_2}, \ldots, x_{i_r}\} \in P_r^{f(S)}$ for every set $S = \{i_1, i_2, \ldots, i_r\} \in r([k])$.

The symmetric group S_k is acting on X^k by permuting the coordinates and this action induces an action on the directed $\mathcal{H}$-cells. Note that a $\mathcal{H}$-cell in the non-directed sense is the union of an orbit of a directed $\mathcal{H}$-cell under the action of S_k.

An abstract (k, l)-cell is a function $c : r([k]) \mapsto [l]$. A (k, l)-cell system $\mathcal{C}$ is a subset of all possible (k, l)-cells. The symmetric group S_k is acting on $r([k])$ and this induces an action on the (k, l)-cells. We say that the system $\mathcal{C}$ is symmetric if it is invariant under the action of S_k. Such a symmetric (k, l)-system shall be called a **combinatorial structure**.

Thus if $\mathcal{H}$ is an l-hyperpartition on $[n]$ and $\mathcal{C}$ is a combinatorial structure then we can define a k-uniform hypergraph $H(\mathcal{H}, \mathcal{C}, [n])$ in the following way. The hypergraph $H(\mathcal{H}, \mathcal{C}, [n])$ is the union of those $\mathcal{H}$-cells in $[n]$ which belong to the coordinates of the combinatorial structure $\mathcal{C}$. If F is a k-uniform hypergraph then we may compute the homomorphism density of F in a combinatorial structure $\mathcal{C}$ as follows. Assume that $V(F) = [n]$ and fix a bijection $s_E : [k] \to E$ for each edge of F. A function $g : r([n], k) \mapsto [l]$ is called a homomorphism of F into $\mathcal{C}$ if for every edge E the restriction $g \circ s_E : r[k] \to [l]$ is a (k, l)-cell of $\mathcal{C}$. The homomorphism density $t(F, \mathcal{C})$ is the probability that a random map $f : r([n], k) \mapsto [l]$ is a homomorphism.

2.4. Regularity lemma as compactness

In this section we state a new type of regularity lemma together with a counting lemma which implies the one stated in the previous section. An interesting feature of this regularity lemma is that arbitrarily decreasing functions (which are common features in “strong” regularity lemmas) are replaced by a sequential compactness type statement.

Theorem 3 (*Hypergraph Sequence Regularity Lemma*). *For every $\epsilon > 0$ and k-uniform increasing hypergraph sequence $\{H_i\}_{i=1}^{\infty}$ there is a natural number $l = l(\epsilon, \{H_i\}_{i=1}^{\infty})$ such that there is a subsequence $\{H_i'\}_{i=1}^{\infty}$ of $\{H_i\}_{i=1}^{\infty}$ together with a sequence of l-hyperpartitions $\{\mathcal{H}_i\}_{i=1}^{\infty}$ satisfying the following properties.*

1. *For every i there is T_i which is the union of some $\mathcal{H}_i$-cells such that $|H_i' \triangle T_i| \le \epsilon \binom{|X_i|}{k}$ where T is the union of some $\mathcal{H}_i$-cells and X_i is the vertex set of H_i'.*
2. *The hyperpartition $\mathcal{H}_i$ is δ_i-equitable and δ_i-regular where* $\lim_{i\to\infty} \delta_i = 0$.
3. *Every T_i has the same combinatorial structure $\mathcal{C}$*
4. $\lim_{i\to\infty} t(F, T_i) = t(F, \mathcal{C})$ *for every k-uniform hypergraph F.*

Note that the value of l depends on the concrete sequence $\{H_i\}_{i=1}^{\infty}$. To see this one can take a large random graph G on n vertices and then take the i-fold equitable blowups G_i of G. The reader can check that in this case (with high probability) $l = n$ for any $\epsilon < 1/2$.

It is quite natural to interpret Theorem 3 in terms of compactness.

Definition 2.6. An increasing hypergraph sequence $\{H_i\}_{i=1}^{\infty}$ is called **strongly convergent** if for every $\epsilon > 0$ there is a number l, hypergraphs T_i on the vertex sets X_i of H_i and l-hyperpartitions $\mathcal{H}_i$ on X_i for every i such that

1. T_i is the union of some $\mathcal{H}_i$-cells
2. $|H_i \triangle T_i| \le \epsilon \binom{|X_i|}{k}$
3. The hyperpartition $\mathcal{H}_i$ is δ_i regular and δ_i equitable where $\lim_{i\to\infty} \delta_i = 0$.
4. Every T_i has the same combinatorial structure.

Using this definition the sequence regularity lemma gets the following simple form.

Theorem 4 (*Regularity as Compactness*). *Every hypergraph sequence has a strongly convergent subsequence.*

2.5. Euclidean hypergraphs

The goal of this subsection is to generalize the notion of k-uniform hypergraphs and homomorphism densities to the Euclidean setting in order to define limit objects for convergent sequences of finite hypergraphs. Seemingly, the appropriate Euclidean analogue of k-uniform hypergraphs would be just the S_k-invariant measurable subsets of $[0, 1]^k$. One could easily define the notion of homomorphisms from finite k-hypergraphs to such Euclidean hypergraphs and even the associated homomorphism densities. The problem with this simple notion of Euclidean hypergraphs is that they could serve as limit objects only for very special finite hypergraph sequences. In order to construct (see Example 1) limit objects to the various random construction of convergent hypergraph sequences one needs a little bit more complicated notion.

Let $k > 0$ and consider $[0, 1]^{2^k-1} = [0, 1]^{r([k])}$, that is the set of points in the form $(x_{A_1}, x_{A_2}, \ldots, x_{A_{2^k-1}})$, where $A_1, A_2, \ldots, A_{2^k-1}$ is a list of the non-empty subsets of $[k]$. Observe that the symmetry group S_k acts on $[0, 1]^{r([k])}$ by

$$\pi((x_{A_1}, x_{A_2}, \ldots, x_{A_{2^k-1}})) = (x_{\pi^{-1}(A_1)}, x_{\pi^{-1}(A_2)}, \ldots, x_{\pi^{-1}(A_{2^k-1})}).$$

We call a measurable $S-k$-invariant subset $\mathcal{H} \subseteq [0, 1]^{2^k-1}$ a k-uniform **Euclidean hypergraph**. Now let K be a finite k-uniform hypergraph and let $\Sigma(K) \subseteq r([n], k)$ be the simplicial complex of K consisting of the non-empty subsets of the k-edges of K. Let $C_1, C_2, \ldots, C_{|\Sigma(K)|}$ be a list of the elements of $\Sigma(K)$.

Definition 2.7 (*Euclidean Hypergraph Homomorphism*). A map $g : r([n], k) \to [0, 1]$ is called a Euclidean hypergraph homomorphism from K to $\mathcal{H}$ if for any edge $E \in E(K)$:

$$(g(s_E(A_1)), g(s_E(A_2)), \ldots, g(s_E(A_{2^k-1}))) \in \mathcal{H},$$

where $s_E : [k] \to E$ is a fixed bijection. The induced Euclidean hypergraph homomorphism is defined accordingly.

Note that the notion of hypergraph homomorphism does not depend on the choice of s_E. Thus the **Euclidean hypergraph homomorphism set** $T(K, \mathcal{H}) \subset [0, 1]^{r([n],k)}$ is the set of points $(y_{B_1}, y_{B_2}, \ldots, y_{B_{|r([n],k)|}})$ such that the map $g :\to [0, 1], g(B_i) = y_{B_i}$ is a homomorphism. One can similarly define the **Euclidean hypergraph induced homomorphism set**. We call $\lambda(T(K, \mathcal{H}))$ the $|\Sigma(K)|$-dimensional Lebesgue-measure of the homomorphism set the **homomorphism density**. We say that the hypergraph $\mathcal{H}$ is the limit of the k-uniform hypergraphs $\{H_n\}_{n=1}^\infty$ if

$$\lim_{n\to\infty} t(K, H_n) = \lambda(T(K, \mathcal{H}))$$

for any finite k-uniform hypergraph K.

Example 1. There are many ways to define random k-uniform hypergraph sequences. The most natural one is the random sequence $\{H_n\}_{n=1}^\infty$, where each edge of the complete hypergraph on n-vertices is chosen with probability $\frac{1}{2}$ to be an edge of H_n. Thus for any k-uniform hypergraph K, $\lim_{n\to\infty} t(K, H_n) = (\frac{1}{2})^{|E(K)|}$ with probability 1. Let us consider the hypergraph

$$\mathcal{H} = \left\{(x_{A_1}, x_{A_2}, \ldots, x_{A_{2^k-1}}) \in [0, 1]^{2^k-1} \mid 0 \le x_{[k]} \le \frac{1}{2}\right\}.$$

An easy calculation shows that $\lambda(T(K, \mathcal{H})) = (\frac{1}{2})^{|E(K)|}$ that is $\mathcal{H}$ is the limit of a random hypergraph sequence $\{H_n\}_{n=1}^\infty$ with probability 1.

Example 2. Now we consider a different notion of randomness. Let the random sequence $\{H'_n\}_{n=1}^\infty$ be constructed in the following way. First choose each $(k-1)$-subset of $[n]$ randomly with probability $\frac{1}{2}$. Then E will be an edge of H'_n if all its $(k-1)$-dimensional hyperedges are chosen. Clearly, $\lim_{n\to\infty} t(K, H'_n) = (\frac{1}{2})^{|K|_{k-1}}$ with probability 1, where $|K|_{k-1}$ is the number of $(k-1)$-hyperedges in $\Sigma(K)$. Now we consider the hypergraph

$$\mathcal{H}' = \Big\{(x_{A_1}, x_{A_2}, \ldots, x_{A_{2^k-1}}) \in [0, 1]^{2^k-1} \mid 0 \le x_{1,2,3,\ldots,k-1} \le \frac{1}{2}, \\ 0 \le x_{1,2,3,\ldots,k-2,k} \le \frac{1}{2}, \ldots, 0 \le x_{2,3,\ldots,k} \le \frac{1}{2}\Big\}.$$

Then $\lambda(T(K, \mathcal{H}')) = (\frac{1}{2})^{|K|_{k-1}}$. Thus $\mathcal{H}'$ is the limit of a random hypergraph sequence $\{H'_n\}_{n=1}^\infty$ with probability 1.

Now let K be a finite k-uniform hypergraph. For any $E \in E(K)$ we fix a bijection $s_E : [k] \to E$ as above. Let $L_{s_E} : [0, 1]^{r([k])} \to [0, 1]^{r([E])}$,

$$L_{s_E}(x_{A_1}, x_{A_2}, \ldots, x_{A_{2^k-1}}) = (x_{s_E(A_1)}, x_{s_E(A_2)}, \ldots, x_{s_E(A_{2^k-1})})$$

be the natural measurable isomorphism associated to the map s_E. Also, let $L_E : [0,1]^{r([n],k)} \to [0,1]^{r(E)}$ be the natural projection. Then for a k-uniform Euclidean hypergraph $\mathcal{H}$ and a finite k-uniform hypergraph K on n vertices

$$T(K,\mathcal{H}) = \bigcap_{E\in E(K)} L_E^{-1}(L_{s_E}(\mathcal{H})). \tag{1}$$

Also,

$$T_{ind}(K,\mathcal{H}) = \bigcap_{E\in E(K)} L_E^{-1}(L_{s_E}(\mathcal{H})) \cap \bigcap_{E'\in E(K)^c} L_{E'}^{-1}(L_{s_{E'}}(\mathcal{H}^c)). \tag{2}$$

We formulate (1) in an integral form as well. Let $W_{\mathcal{H}} : [0,1]^{r([k])} \to \{0,1\}$ be the characteristic function of the Euclidean hypergraph $\mathcal{H}$. We call such an object a **hypergraphon**. Then

$$\lambda(T(K,\mathcal{H})) = \int_0^1 \int_0^1 \dots \int_0^1 \left(\prod_{E\in E(K)} \Psi_E \right) dx_{C_1} dx_{C_2} \dots dx_{C_{\Sigma(K)}},$$

where Ψ_E is the characteristic function of $L_E^{-1}(L_{s_E}(\mathcal{H}))$. Clearly,

$$\Psi_E(x_{C_1}, x_{C_2}, \dots, x_{C_{\Sigma(K)}}) = W_{\mathcal{H}}(x_{s_E(A_1)}, x_{s_E(A_2)}, \dots, x_{s_E(A_{2^k-1})}).$$

Thus, we have the integral formula

$$\lambda(T(K,\mathcal{H})) = \int_0^1 \int_0^1 \dots \int_0^1 \left(\prod_{E\in E(K)} W_{\mathcal{H}}(x_{s_E(A_1)}, x_{s_E(A_2)}, \dots, x_{s_E(A_{2^k-1})}) \right) dx_{C_1} \dots dx_{C_{\Sigma(K)}}.$$

Remark. One can introduce the notion of a **projected hypergraphon** $\widetilde{W}_{\mathcal{H}}$ which is the projection of a hypergraphon to the first 2^k-2 coordinates, where the last coordinate is associated to $[k]$ itself. That is

$$\widetilde{W}_{\mathcal{H}}(x_{A_1}, x_{A_2}, \dots, x_{A_{2^k-2}}) = \int_0^1 W_{\mathcal{H}}(x_{A_1}, x_{A_2}, \dots, x_{A_{2^k-1}}) dx_{A_{2^k-1}}.$$

That is $\widetilde{W}_{\mathcal{H}}$ is a $[0,1]$-valued function which is symmetric under the induced S_k-action of its coordinates. By the classical Fubini theorem we obtain that if $\mathcal{H}$ is the limit of the hypergraphs $\{H_i\}_{i=1}^\infty$ then

$$lim_{i\to\infty} t(K,H_i) = \int_0^1 \int_0^1 \dots \int_0^1 \prod_{E\in E(K)} \widetilde{W}_{\mathcal{H}}(x_{s_E(A_1)}, x_{s_E(A_2)}, \dots, x_{s_E(A_{2^k-2})}) dx_{C_1} dx_{C_2} \dots dx_{C_{|K|_{k-1}}},$$

where the integration is over the variables associated to the simplices of dimension less than k. Note that in the case $k=2$ it is just the graph limit formula of [10].

Note that for a combinatorial structure $\mathcal{C}$ one can define a hypergraphon $W_{\mathcal{C}} \subseteq [0,1]^{2^k-1}$. Recall that an l-box Z in $[0,1]^{2^k-1}$ is a product set in the form

$$\left(\frac{i_1}{l}, \frac{i_1+1}{l}\right) \times \left(\frac{i_2}{l}, \frac{i_2+1}{l}\right) \times \dots \times \left(\frac{i_{2^k-1}}{l}, \frac{i_{2^k-1}+1}{l}\right).$$

The map $f : r[k] \to [l]$, defined by $f(A_j) = i_j$ is the coordinate function of the box Z. Then $W_{\mathcal{C}}$ is the union of the boxes corresponding to the coordinates of the combinatorial structure $\mathcal{C}$. It is easy to check that $t(F, \mathcal{C}) = t(F, W_{\mathcal{C}})$ for any k-uniform hypergraph F.

2.6. *W-random graphs and sampling*

Let us consider the following natural sampling process for k-uniform hypergraphs. We pick n vertices $v_1, v_2, \dots, v_n$ independently and uniformly at random from the vertex set X of H and then we create a hypergraph $\mathbb{G}(H, n)$ with vertex set $[n]$ such that $\{i_1, i_2, \dots, i_k\}$ is an edge in $\mathbb{G}(H, n)$ if and only if $\{v_{i_1}, v_{i_2}, \dots, v_{i_k}\}$ is an edge in H. Thus $\mathbb{G}(H, n)$ is a hypergraph valued random variable. The distribution of $\mathbb{G}(H, n)$ can be described in terms of the homomorphism densities $t_{\text{ind}}(F, H)$ where $|V(F)| \leq n$. The probability that we see a fixed hypergraph F on $[n]$ in $\mathbb{G}(H, n)$ is equal to $t_{\text{ind}}(F, H)$.

Now we generalize sampling for Euclidean hypergraphs $W \subset [0, 1]^{r([k])}$. Let us introduce a random variable X_S for every set $S \in r([n], k)$ which are independent and have uniform distribution in $[0, 1]$. Then $\{i_1, i_2, \dots, i_k\}$ is an edge in $\mathbb{G}(W, n)$ if $W(X_{A_1}, X_{A_2}, \dots, X_{A_{2^k-1}}) = 1$ where $A_1, A_2, \dots, A_{2^k-1}$ are the non-empty subsets of $\{i_1, i_2, \dots, i_k\}$. This again gives a hypergraph valued random variable on $[n]$ which is the infinite analogy of the finite setting.

Another important sampling process from W will be called the **hyperpartition sampling**. Assume that $\mathcal{H} = \{P_j^r\}_{1 \leq j \leq l, 1 \leq r \leq k}$ is an l-hyperpartition on the set $[n]$. We consider the function $g : r([n], k) \to [l]$, which is equal to j if and only if $S \in P_j^{|S|}$. Now we define a sampling process $\mathbb{G}(W, \mathcal{H}, n)$ in the same way as $\mathbb{G}(W, n)$ with the extra restriction that X_S has uniform distribution in the interval $[(g(S) - 1)/l, g(S)/l)$. This sampling process has the property that $t_{\text{ind}}(F, W) = 0$ implies $t_{\text{ind}}(F, \mathbb{G}(W, \mathcal{H}, n)) = 0$ with probability 1.

Finally, we introduce the notion of random coordinate systems. Let Z_n be the random variable which is a random point in $[0, 1]^{r([n],k)}$ with uniform distribution. In other words Z_n is a $r([n], k)$-tuple of independent random variables with uniform distribution $\{f_T\}_{T \in r([n],k)}$. Let $[n]_0^k$ be the set of elements in $[n]^k$ without having repetitions in their coordinates. We introduce the random variables $\tau^n : [n]_0^k \mapsto [0, 1]^{r([k])}$ such that the component $\tau_S^n(x_1, x_2, \dots, x_k)$ corresponding to an element $\{i_1, i_2, \dots, i_t\} = S \in r([k])$ is equal to the value of $f_{x_{i_1}, x_{i_2}, \dots, x_{i_t}}$. We call the random variables τ^n **random coordinate systems** corresponding to $[n]$. An important property of (τ^n) is that for a measurable set $W \subseteq [0, 1]^{r([k])}$ the distribution of the random hypergraph-valued function $(\tau^n)^{-1}(W)$ is exactly the same as the distribution of $\mathbb{G}(W, [n])$.

2.7. *Ultraproducts of finite sets*

First we recall the ultraproduct construction of finite probability measure spaces (see [9]). Let $\{X_i\}_{i=1}^{\infty}$ be finite sets. We always suppose that $|X_1| < |X_2| < |X_3| < \cdots$. Let ω be a nonprincipal ultrafilter and $\lim_\omega : l^\infty(\mathbb{N}) \to \mathbb{R}$ be the corresponding ultralimit. Recall that $\lim_\omega$ is a bounded linear functional such that for any $\epsilon > 0$ and $\{a_n\}_{n=1}^{\infty} \in l^\infty(\mathbb{N})$

$$\{i \in \mathbb{N} \mid a_i \in [\lim_\omega a_n - \epsilon, \lim_\omega a_n + \epsilon]\} \in \omega.$$

The ultraproduct of the sets X_i is defined as follows.

Let $\widetilde{X} = \prod_{i=1}^{\infty} X_i$. We say that $\widetilde{p} = \{p_i\}_{i=1}^{\infty}, \widetilde{q} = \{q_i\}_{i=1}^{\infty} \in \widetilde{X}$ are equivalent, $\widetilde{p} \sim \widetilde{q}$, if

$$\{i \in \mathbb{N} \mid p_i = q_i\} \in \omega.$$

Define $\mathbf{X} := \widetilde{X}/\sim$. Now let $\mathcal{P}(X_i)$ denote the Boolean-algebra of subsets of X_i, with the normalized measure $\mu_i(A) = \frac{|A|}{|X_i|}$. Then let $\widetilde{\mathcal{P}} = \prod_{i=1}^{\infty} \mathcal{P}(X_i)$ and $\mathcal{P} = \widetilde{\mathcal{P}}/I$, where I is the ideal of elements $\{A_i\}_{i=1}^{\infty}$ such that $\{i \in \mathbb{N} \mid A_i = \emptyset\} \in \omega$. Notice that the elements of $\mathcal{P}$ can be identified with certain subsets of $\mathbf{X}$: if

$$\overline{p} = [\{p_i\}_{i=1}^{\infty}] \in \mathbf{X} \quad \text{and} \quad \overline{A} = [\{A_i\}_{i=1}^{\infty}] \in \mathcal{P}$$

then $\overline{p} \in \overline{A}$ if $\{i \in \mathbb{N} \mid p_i \in A_i\} \in \omega$. Clearly, if $\overline{A} = [\{A_i\}_{i=1}^{\infty}]$, $\overline{B} = [\{B_i\}_{i=1}^{\infty}]$ then

- $\overline{A}^c = [\{A_i^c\}_{i=1}^{\infty}]$,
- $\overline{A} \cup \overline{B} = [\{A_i \cup B_i\}_{i=1}^{\infty}]$,
- $\overline{A} \cap \overline{B} = [\{A_i \cap B_i\}_{i=1}^{\infty}]$.

That is $\mathcal{P}$ is a Boolean algebra on $\mathbf{X}$. Now let $\mu(\overline{A}) = \lim_{\omega} \mu_i(A_i)$. Then $\mu : \mathcal{P} \to \mathbb{R}$ is a finitely additive probability measure. We will call $\overline{A} = [\{A_i\}_{i=1}^{\infty}]$ the **ultraproduct** of the sets $\{A_i\}_{i=1}^{\infty}$.

Definition 2.8. $N \subseteq \mathbf{X}$ is a **nullset** if for any $\epsilon > 0$ there exists a set $\overline{A_\epsilon} \in \mathcal{P}$ such that $N \subseteq \overline{A_\epsilon}$ and $\mu(\overline{A_\epsilon}) \leq \epsilon$. The set of nullsets is denoted by $\mathcal{N}$.

Proposition 2.1. *$\mathcal{N}$ satisfies the following properties:*

- *if $N \in \mathcal{N}$ and $M \subseteq N$, then $M \in \mathcal{N}$.*
- *If $\{N_k\}_{k=1}^{\infty}$ are elements of $\mathcal{N}$ then $\cup_{k=1}^{\infty} N_k \in \mathcal{N}$ as well.*

Proof. The first part is obvious, for the second part we need the following lemma.

Lemma 2.4. *If $\{\overline{A_k}\}_{k=1}^{\infty}$ are elements of $\mathcal{P}$ and $\lim_{l\to\infty} \mu(\cup_{k=1}^{l} \overline{A_k}) = t$ then there exists an element $\overline{B} \in \mathcal{P}$ such that $\mu(\overline{B}) = t$ and $\overline{A_k} \subseteq \overline{B}$ for all $k \in \mathbb{N}$.*

Proof. Let $\overline{B_l} = \cup_{k=1}^{l} \overline{A_k}$, $\mu(\overline{B_l}) = t_l$, $\lim_{l\to\infty} t_l = t$. Let

$$T_l = \left\{ i \in \mathbb{N} \mid |\mu_i(\cup_{k=1}^{l} A_k^i) - t_l| \leq \frac{1}{2^l} \right\},$$

where $\overline{A_k} = [\{A_k^i\}_{i=1}^{\infty}]$. Observe that $T_l \in \omega$. If $i \in \cap_{l=1}^{m} T_l$ but $i \notin T_{m+1}$, then let $C_i = \cup_{k=1}^{m} A_k^i$. If $i \in T_l$ for all $l \in \mathbb{N}$, then clearly $\mu_i(\cup_{k=1}^{\infty} A_k^i) = t$ and we set $C_i := \cup_{k=1}^{\infty} A_k^i$. Let $\overline{B} := [\{C_i\}_{i=1}^{\infty}]$. Then $\mu(\overline{B}) = t$ and for any $k \in \mathbb{N}$: $\overline{A_k} \subseteq \overline{B}$. □

Now suppose that for any $j \geq 1$, $\overline{A_j} \in \mathcal{N}$. Let $\overline{B}_j^{\epsilon} \in \mathcal{P}$ such that $\overline{A_j} \subseteq \overline{B}_j^{\epsilon}$ and $\mu(\overline{B}_j^{\epsilon}) < \epsilon \frac{1}{2^j}$. Then by the previous lemma, there exists $\overline{B}^{\epsilon} \in \mathcal{P}$ such that for any $j \geq 1$ $\overline{B}_j^{\epsilon} \subseteq \overline{B}^{\epsilon}$ and $\mu(\overline{B}^{\epsilon}) \leq \epsilon$. Since $\cup_{j=1}^{\infty} \overline{A_j} \subseteq \overline{B}^{\epsilon}$, our proposition follows. □

Definition 2.9. We call $B \subseteq \mathbf{X}$ a **measurable set** if there exists $\widetilde{B} \in \mathcal{P}$ such that $B \triangle \widetilde{B} \in \mathcal{N}$.

Proposition 2.2. *The measurable sets form a σ-algebra $\mathcal{B}_\omega$ and $\mu(B) = \mu(\widetilde{B})$ defines a probability measure on $\mathcal{B}_\omega$.*

Proof. We call two measurable sets B and B' equivalent, $B \cong B'$ if $B \triangle B' \in \mathcal{N}$. Clearly, if $A \cong A'$, $B \cong B'$ then $A^c \cong (A')^c$, $A \cup B \cong A' \cup B'$, $A \cap B \cong A' \cap B'$. Also if $A, B \in \mathcal{P}$ and $A \cong B$, then $\mu(A) = \mu(B)$. That is the measurable sets form a Boolean algebra with a finitely additive measure. Hence it is enough to prove that if $\overline{A_k} \in \mathcal{P}$ are disjoint sets, then there exists

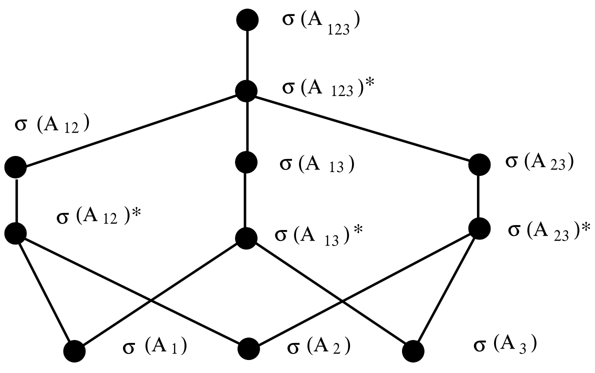

Fig. 1. The σ-algebras.

$\overline{A} \in \mathcal{P}$ such that $\cup_{k=1}^{\infty} \overline{A_k} \cong \overline{A}$ and $\mu(\overline{A}) = \sum_{k=1}^{\infty} \mu(\overline{A_k})$. Note that by Lemma 2.4 there exists $\overline{A} \in \mathcal{P}$ such that $\mu(\overline{A}) = \sum_{k=1}^{\infty} \mu(\overline{A_k})$ and $\overline{A_k} \subseteq \overline{A}$ for all $k \geq 1$. Then for any $j \geq 1$,

$$\overline{A} \setminus \cup_{k=1}^{\infty} \overline{A_i} \subseteq \overline{A} \setminus \cup_{k=1}^{j} \overline{A_k} \in \mathcal{P}.$$

Since $\lim_{j\to\infty} \mu(\overline{A} \setminus \cup_{k=1}^{j} \overline{A_k}) = 0$, $\overline{A} \setminus \cup_{k=1}^{\infty} \overline{A_k} \in \mathcal{N}$ thus $\cup_{k=1}^{\infty} \overline{A_k} \cong \overline{A}$. □

Hence we constructed an atomless probability measure space $(\mathbf{X}, \mathcal{B}_\omega, \mu)$. Note that this space is non-separable, that is it is not measurably equivalent to the interval with the Lebesgue measure.

2.8. σ-algebras and the Total Independence Theorem

We fix a natural number k and we denote by $[k]$ the set $\{1, 2, \dots, k\}$. Let $X_{i,1}, X_{i,2}, \dots, X_{i,k}$ be k copies of the finite set X_i and for a subset $A \subseteq \{1, 2, \dots, k\}$ let $X_{i,A}$ denote the direct product $\bigoplus_{j\in A} X_{i,j}$. Let $\mathbf{X}^A$ denote the ultraproduct of the sets $X_{i,A}$, with a Boolean algebra $\mathcal{P}_A$. There is a natural map $p_A : \mathbf{X}^{[k]} \to \mathbf{X}^A$ (the projection). Let $\mathcal{B}_A$ be the σ-algebra of measurable subsets in $\mathbf{X}^A$ as defined in the previous sections. Define $\sigma(A)$ as $p_A^{-1}(\mathcal{B}_A)$, the σ-algebra of measurable sets depending only on the A-coordinates together with the probability measure μ_A. For a nonempty subset $A \subseteq [k]$ let A^* denote the set system $\{B | B \subseteq A, |B| = |A| - 1\}$ and let $\sigma(A)^*$ denote the σ-algebra $\langle \sigma(B) | B \in A^* \rangle$. An interesting fact is (as it will turn out in Section 6) that $\sigma(A)^*$ is strictly smaller than $\sigma(A)$. The following figure shows how the lattice of the various σ-algebras look like. Recall that if $\mathcal{B} \subset \mathcal{A}$ are σ-algebras on X with a measure μ and g is an $\mathcal{A}$-measurable function on X, then $E(g \mid \mathcal{B})$ is the $\mathcal{B}$-measurable function (unique up to a zero measure perturbation) with the property that

$$\int_Y E(g \mid \mathcal{B})\, d\mu = \int_Y g\, d\mu,$$

for any $Y \in \mathcal{B}$ (see Appendix). If $A \in \mathcal{A}$ we say that A is independent from the σ-algebra $\mathcal{B}$ if $E(\chi_A \mid \mathcal{B})$ is a constant function. One of the main tools in our paper (the proof will be given in Section 5) is the following theorem. (See Fig. 1).

Theorem 5 (*The Total Independence Theorem*)**.** *Let $A_1, A_2, \dots, A_r$ be a list of distinct nonempty subsets of $[k]$, and let $S_1, S_2, \dots, S_r$ be subsets of $\mathbf{X}^{[k]}$ such that $S_i \in \sigma(A_i)$ and $E(S_i | \sigma(A_i)^*)$ is a constant function for every $1 \leq i \leq r$. Then*

$$\mu(S_1 \cap S_2 \cap \dots \cap S_r) = \mu(S_1)\mu(S_2)\dots\mu(S_r).$$

3. Correspondence principles and the proofs of the Removal and Regularity Lemmas

3.1. The ultraproduct method and the correspondence principles

The ultraproduct method for hypergraphs relies on various correspondence principles between the following objects that are infinite variations of the concept of a k-uniform hypergraph.

1. An infinite sequence of hypergraphs $H_1, H_2, \dots$ in $\mathcal{H}_k$.
2. The ultraproduct hypergraph $\mathbf{H}$.
3. A k-uniform Euclidean hypergraph $\mathcal{H} \subseteq [0,1]^{2^k-1}$.

Additionally we will need correspondence principles between homomorphism sets

$$\{T(K, H_i)\}_{i=1}^{\infty}, \qquad T(K, \mathbf{H}) \quad \text{and} \quad T(K, \mathcal{H})$$

for every fixed k-uniform hypergraph K. Let $\{H_i \subset X_i^k\}_{i=1}^{\infty}$ be a sequence of finite k-uniform hypergraphs. Then the **ultraproduct hypergraph** $\mathbf{H} = [\{H_i\}_{i=1}^{\infty}] \subset \mathbf{X}^k$ is well-defined. Clearly, $\mathbf{H}$ is S_k-invariant and has no repetitions in its coordinates. One can formally define the homomorphism set $T(K, \mathbf{H})$ for any finite k-uniform hypergraph K exactly as in Section 2.1. Note that we shall refer to any measurable S_k-invariant set $\mathbf{P} \subset \mathbf{X^k}$ without repetitions in its coordinated k-uniform hypergraph on $\mathbf{X}$.

The following lemma is a trivial consequence of the basic properties of the ultraproduct sets.

Lemma 3.1 (*Homomorphism Correspondence I*). *The homomorphism set $T(F, \mathbf{H})$ is the ultraproduct of the homomorphism sets $T(F, H_i)$. The induced homomorphism set $T_{ind}(F, \mathbf{H})$ is the ultraproduct of the homomorphism sets $T_{ind}(F, H_i)$.*

To state the next theorem we need some notation. For an arbitrary set S let $r(S, m)$ denote the set of non-empty subsets of S of size at most m and let $r(S)$ denote $r(S, |S|)$. The symmetric group S_n is acting on $[n]$ and this action induces an action on $r([n], m)$. Furthermore S_n is acting on $[0,1]^{r([n],m)}$ by permuting the coordinates according to the action on $r([n], m)$. Let X, G, G_2 be sets such that $G_2 \subseteq G$. Then we will denote the projection $X^G \mapsto X^{G_2}$ by P_{G_2}. If a function f takes values in X^G then for an element $a \in G$ we denote the corresponding coordinate function by f_a which is the same as the composition $P_{\{a\}} \circ f$.

Definition 3.1 (*Separable Realization*). For any $k \in \mathbb{N}$ a separable realization is a measure preserving map $\phi : \mathbf{X}^k \mapsto [0,1]^{r([k])}$ such that

1. Any permutation $\pi \in S_k$ commutes with ϕ in the sense that $\phi(\mathbf{x})^{\pi} = \phi(\mathbf{x}^{\pi})$.
2. For any $D \in r([k])$ and measurable set $A \subseteq [0,1]$ the set $\phi_D^{-1}(A)$ is in $\sigma(D)$ and is independent from $\sigma(D)^*$.

Note that the fact that ϕ commutes with the S_k-action means that $\phi_A(\mathbf{x}^{\pi}) = \phi_{A^{\pi^{-1}}}(\mathbf{x})$ for each $\pi \in S_k$. The second condition in the previous definition expresses the fact that the functions ϕ_D of a separable realization depend only on the D-coordinates. Also, by Lemma A.2 of Appendix and the Total Independence Theorem a separable realization ϕ gives a parametrization of $\mathbf{X}^k$ by $|r([k])|$ coordinates such a way that ϕ^{-1} defines an injective measure algebra homomorphism from $\mathcal{M}([0,1]^{r([k])}, \mathcal{B}^k, \lambda^k)$ to a subalgebra of $\mathcal{M}(\mathbf{X}^k, \mathcal{B}_{[k]}, \mu_{[k]})$. The next theorem is the heart of the hypergraph ultraproduct method. The proof of it will be discussed in Section 6.

Theorem 6 (*Euclidean Correspondence*). *Let $\mathcal{A}$ be a separable sub-σ-algebra of $\sigma_{[k]}$ on $\mathbf{X}^k$. Then there is a separable realization $\phi : \mathbf{X}^k \mapsto [0,1]^{r([k])}$ such that for every $A \in \mathcal{A}$ there is a measurable set $B \subseteq [0,1]^{2^k-1}$ with $\mu(\phi^{-1}(B)\triangle A) = 0$.*

Corollary 3.1. *Let $\mathbf{E}$ be an S_k-invariant measurable subset of $\mathbf{X}^k$. Then there is a separable realization ϕ and S_k-invariant measurable set $W \subseteq [0,1]^{r([k])}$ such that $\mu(\phi^{-1}(W)\triangle\mathbf{E}) = 0$.*

The following definition and lemma will be needed to state the main correspondence between homomorphism sets.

Definition 3.2 (*Lifting*). Let $\phi : \mathbf{X}^k \mapsto [0,1]^{r([k])}$ be a separable realization and let $n \geq k$ be an arbitrary natural number. Then a measure preserving map $\psi : \mathbf{X}^n \mapsto [0,1]^{r([n],k)}$ is called a degree n **lifting** of ϕ if $P_{r([k])} \circ \psi$ is equal to $\phi \circ P_{[k]}$ on $\mathbf{X}^n$ and $\psi(\mathbf{x})^\pi = \psi(\mathbf{x}^\pi)$ for all permutations $\pi \in S_n$.

Lemma 3.2 (*Lifting Exists*). *Let $\phi : \mathbf{X}^k \mapsto [0,1]^{r([k])}$ be a separable realization and let $n \geq k$ be an arbitrary natural number. Then there exists a degree n lifting ψ of ϕ.*

Proof. Let $A \in r([n],k)$ be an arbitrary set with t elements and let $\pi \in S_n$ be a permutation such that $A^\pi = [t]$. We define $\psi_A(\mathbf{x})$ to be $\phi_{[t]}(P_{[k]}(\mathbf{x}^\pi))$. Using the fact that ϕ commutes with the S_k action we obtain that $\phi_A \circ P_{[k]} = \psi_A$ for every $A \in r([k])$. Now if π_2 is an arbitrary permutation from S_n then the A-coordinate of $\psi(\mathbf{x})^{\pi_2}$ is the $A^{\pi_2^{-1}}$-coordinate of $\psi(\mathbf{x})$ which is the A-coordinate of $\psi(\mathbf{x}^{\pi_2})$. This proves that ψ commutes with S_n. It remains to show that ψ is measure preserving. The coordinate functions ψ_A are constructed in a way which guarantees that they are measure preserving. Let $I_A \subseteq [0,1]$ be intervals of length l_A for every $A \in r([n],k)$ and let

$$W = \prod_{A\in r([n],k)} I_A$$

be their direct product. Since every measurable set in $[0,1]^{r([n],k)}$ can be approximated by the disjoint union of such cubes it is enough to check that ψ^{-1} preserves the measure of such a set W. The preimage $\psi^{-1}(W)$ is the intersection of the preimages $\psi_A^{-1}(I_A)$ which are in $\sigma(A)$ and are independent from $\sigma(A)^*$. Now the Total Independence Theorem completes the proof. □

Lemma 3.3 (*Homomorphism Correspondence II*). *Let $W \subseteq [0,1]^{r([k])}$ be an S_k-invariant measurable set and let $\mathbf{E}$ be the preimage of W under some separable realization ϕ. Then for an arbitrary finite hypergraph K*

$$\psi^{-1}(T(K,W)) = T(K,\mathbf{E}),$$

where ψ is a $|K|$ degree lifting of ϕ. Similarly,

$$\psi^{-1}(T_{ind}(K,W)) = T_{ind}(K,\mathbf{E}).$$

Proof. Assume that the vertex set of K is defined on $[n]$ and that the edges of K are $\pi_1([k])$, $\pi_2([k]), \ldots, \pi_t([k])$ for some permutations $\pi_1, \pi_2, \ldots, \pi_t$ in S_n. Let $\mathbf{E}_2 \subset \mathbf{X}^n$ be the preimage of $\mathbf{E}$ under the projection $P_{[k]}$ and let $W_2 \subset [0,1]^{r([n])}$ be the preimage of W under the projection $P_{r([k])}$. By definition we have that

$$T(K,\mathbf{E}) = \bigcap_{i=1}^{t} \mathbf{E}_2^{\pi_i}$$

and

$$T(K, W) = \bigcap_{i=1}^{t} W_2^{\pi_i}.$$

Since ψ is a lifting of ϕ the first lifting property shows that $\psi^{-1}(W_2) = \mathbf{E}_2$. Furthermore since ψ commutes with the elements of S_n we get that $\psi^{-1}(W_2^{\pi}) = \mathbf{E}_2^{\pi}$ for every $\pi \in S_n$. This completes the proof. □

3.2. The proof of the Hypergraph Removal lemma

Lemma 3.4 (*Infinite Removal Lemma*). *Let* $\mathbf{H}$ *be the ultraproduct of the k-uniform hypergraphs* $H_1, H_2, \ldots$ *and let* F *be a finite k-uniform hypergraph such that* $T(F, \mathbf{H})$ *has measure 0. Then there is a 0-measure* S_k*-invariant subset* $\mathbf{I}$ *of* $\mathbf{H}$ *such that* $T(F, \mathbf{H} \setminus \mathbf{I})$ *is empty.*

Proof. We use Corollary 3.1 for the set $\mathbf{H}$ and we get a separable realization ϕ and a measurable set $W \subseteq [0, 1]^{r([k])}$ satisfying the statement of the corollary. Let D denote the set density points in W. Lebesgue's density theorem says that $W \setminus D$ has measure 0. Furthermore D will remain symmetric under the action of the symmetric group on $[0, 1]^{r([k])}$. Let $\mathbf{D}$ be the preimage of D under the map ϕ. Using the first property in Definition 3.1 we obtain that $\mathbf{D}$ is S_k- invariant. Furthermore the measure of $\mathbf{H} \triangle \mathbf{D}$ is 0.

Now let F be a k-uniform hypergraph on the vertex set $[n]$ and let ψ be a degree n lifting of ϕ. Lemma 3.3 shows that $T(F, \mathbf{D})$ is the preimage of $T(F, D)$ under ψ^{-1}. On the other hand $T(F, D)$ is the intersection of finitely many sets consisting only of density points. This show that $T(F, D)$ and thus $T(F, \mathbf{D})$ is either empty or has a positive measure. This means that the set $\mathbf{I} = \mathbf{H} \setminus \mathbf{D}$ satisfies the required condition. □

Proof of the hypergraph removal lemma. We proceed by contradiction. Let K be a fixed hypergraph and $\epsilon > 0$ be a fixed number for which the theorem fails. This means that there is a sequence of hypergraphs H_i on the sets X_i such that $lim_{i\to\infty} t(K, H_i) = 0$ but in each H_i there is no set L with the required property. Again let $\mathbf{H} \subseteq \mathbf{X}^k$ denote the ultraproduct hypergraph. Then $\mu(T(K, \mathbf{H})) = \lim_\omega t(K, H_i) = 0$ and thus by the previous lemma there is a zero measure S_k-invariant set $\mathbf{I} \subseteq \mathbf{X}^k$ such that $T(K, \mathbf{H} \setminus \mathbf{I}) = \emptyset$. By the definition of nullsets, for any $\epsilon_1 > 0$ there exists an ultralimit set $\mathbf{J} \subset \mathbf{X}^k$ such that $\mathbf{I} \subset \mathbf{J}$ and $\mu(\mathbf{J}) < \epsilon_1$. We can suppose that $\mathbf{J}$ is S_k-invariant as well. Let $[\{J_i\}_{i=1}^{\infty}] = \mathbf{J}$; then for ω-almost all i, J_i is S_k- invariant, $|J_i| \le \epsilon_1 |X_i|^k$ and $T(K, H_i \setminus L_i) = \emptyset$, where L_i is the set of edges $\{x_1, x_2, \ldots, x_k\}$ such that $(x_1, x_2, \ldots, x_k) \in J_i$. Clearly, $|L_i| \le |J_i|$; hence if ϵ_1 is small enough then $|L_i| \le \epsilon \binom{|X_i|}{k}$ leading to a contradiction.

3.3. The existence of the hypergraph limit object

Proposition 3.1. *Let* $\{H_i\}_{i=1}^{\infty}$ *be a sequence of k-uniform hypergraphs and let* $\mathbf{H}$ *be their ultraproduct hypergraph. Assume furthermore that* $\phi : \mathbf{X}^k \mapsto [0, 1]^{r([k])}$ *is a separable realization such that there is an* S_k*-invariant measurable set* $\mathcal{H} \subseteq [0, 1]^{r([k])}$ *with* $\mu(\phi^{-1}(\mathcal{H}) \triangle \mathbf{H}) = 0$. *Then for every k-uniform hypergraph K we have that*

$$\lim_{\omega} t(K, H_i) = t(K, \mathcal{H}).$$

Proof. Let K be a k uniform hypergraph on n vertices and let ψ be a degree n lifting of ϕ. Lemma 3.1 implies that $t(K, \mathbf{H}) = \lim_\omega T(K, H_i)$ furthermore, using that ψ is measure preserving, Lemma 3.3 implies that $t(K, \mathbf{H}) = t(K, \mathcal{H})$. □

The following theorem is an immediate corollary of the previous one.

Theorem 7 (*Existence of the Limit Object*). *If $\{H_i\}_{i=1}^\infty$ is a convergent sequence of k-uniform hypergraphs then there exists a Euclidean hypergraph $\mathcal{H} \subset [0,1]^{r([k])}$ such that $lim_{i\to\infty} t(K, H_i) = t(K, \mathcal{H})$ for every k-uniform hypergraph K.*

3.4. The proof of the Hypergraph Regularity Lemma

Suppose that the theorem does not hold for some $\epsilon > 0$ and $F : \mathbb{N} \to (0,1)$. That is there exists a sequence of k-uniform hypergraphs H_i without having $F(j)$-equitable j-hyperpartitions for any $1 < j \leq i$ satisfying the conditions of our theorem. Let us consider their ultraproduct $\mathbf{H} \subset \mathbf{X}^k$. Similarly to the proof of the Removal Lemma we formulate an infinite version of the Regularity Lemma as well.

Let $K_r(\mathbf{X})$ denote the complete r-uniform hypergraph on X, that is the set of points $(x_1, x_2, \ldots, x_r) \in \mathbf{X}^r$ such that $x_i \neq x_j$ if $i \neq j$. Clearly $K_r(\mathbf{X}) \subset \mathbf{X}^r$ is measurable and $\mu_{[r]}(K_r(\mathbf{X})) = 1$. An r-uniform hypergraph on $\mathbf{X}$ is an S_r-invariant measurable subset of $K_r(\mathbf{X})$. An l-hyperpartition $\widetilde{\mathcal{H}}$ is a family of partitions $K_r(\mathbf{X}) = \cup_{j=1}^l \mathbf{P_r^j}$, where $\mathbf{P_r^j}$ is an r-uniform hypergraph for $1 \leq r \leq k$. Again, an l-hyperpartition induces a partition of $K_k(\mathbf{X})$ into $\widetilde{\mathcal{H}}$-cells exactly the same way as in the finite case. It is easy to see that each $\widetilde{\mathcal{H}}$-cell is measurable.

Proposition 3.2 (*Hypergraph Regularity Lemma, Infinite Version*). *For any $\epsilon > 0$, there exists a 0-equitable l-hyperpartition (where l depends on $\mathbf{H}$) $\widetilde{\mathcal{H}}$ such that*

- *Each $\mathbf{P_r^j}$ is independent from $\sigma([r])^*$.*
- *$\mu_{[k]}(\mathbf{H} \triangle T) \leq \epsilon$, where T is a union of some $\widetilde{\mathcal{H}}$-cells.*

Proof. Let ϕ be a separable realization for $\mathbf{H}$ that is such a ϕ that there exists an S_k-invariant subset $Q \subseteq [0,1]^{2^k-1}$ such that $\mu_{[k]}(\phi^{-1}(Q) \triangle \mathbf{H}) = 0$. Since Q is a Lebesgue-measurable set, there exists some $l > 0$ such that $Vol_{2^k-1}(Q \triangle Z) < \epsilon$, where Z is a union of l-boxes (see Section 2.3).

By the usual symmetrization argument we may suppose that the set Z is invariant under the S_k-action on the l-boxes. For each $1 \leq r \leq k$ we consider the partition $\mathbf{X}^r = \cup_{j=1}^l \mathbf{P_r^j}$, where $\mathbf{P_r^j} = \phi_{[r]}^{-1}(\frac{j-1}{l}, \frac{j}{l})$. We call the resulting l-hyperpartition $\widetilde{\mathcal{H}}$. Note that by the S_r-invariance of the separable realization each $\mathbf{P_r^j}$ is an r-uniform hypergraph and also $\mathbf{P_r^j}$ is independent from $\sigma([r])^*$.

Now we show that $\mathbf{C}$ is an $\widetilde{\mathcal{H}}$-cell if and only if $\mathbf{C} = \phi^{-1}(\cup_{\pi \in S_k} \pi(D))$, where D is an l-box in $[0,1]^{2^k-1}$. By definition $\mathbf{a} = (a_1, a_2, \ldots, a_k) \in \mathbf{X}^k$ and $\mathbf{b} = (b_1, b_2, \ldots, b_k) \in \mathbf{X}^k$ are in the same $\widetilde{\mathcal{H}}$-cell if and only if there exists $\pi \in S_k$ such that $(a_{i_1}, a_{i_2}, \ldots, a_{i_{|A|}})$ and $(b_{i_{\pi(1)}}, b_{i_{\pi(2)}}, \ldots, b_{i_{\pi(|A|)}})$ are in the same $\mathbf{P_r^j}$ for any $A \subseteq [k]$. That is $\phi(\mathbf{a})$ and $\phi(\mathbf{b}^\pi) = (\phi(\mathbf{b}))^\pi$ are in the same l-box.

Since Z is a union of S_k-orbits of l-boxes the set $T = \phi^{-1}(Z)$ is the union of $\widetilde{\mathcal{H}}$ cells. Using that ϕ is measure preserving the proof is complete. □

Now we return to the proof of the Hypergraph Regularity Lemma. First pick an r-hypergraph $\widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}}$ on $\mathbf{X}$ such that $\mu_{[r]}(\widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}} \triangle \mathbf{P}^{\mathbf{j}}_{\mathbf{r}}) = 0$, $\widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}} \in \mathcal{P}_{[r]}$ and $\cup_{j=1}^{l} \widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}} = K_r(\mathbf{X})$. Let $[\{P^j_{r,i}\}_{i=1}^{\infty}] = \widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}}$. Then for ω-almost all indices $\cup_{j=1}^{l} P^j_{r,i} = K_r(X_i)$ is an $F(l)$-equitable l-partition and $|H_i \triangle \cup_{m=1}^{q} C^i_m| < \epsilon \binom{|X|}{k}$ for the induced $\mathcal{H}$-cell approximation. Here $\cup_{m=1}^{q} \widetilde{\mathbf{C}}_{\mathbf{m}}$ is the $\widetilde{\mathcal{H}}$-cell approximation with respect to the l-hyperpartitions $\cup_{j=1}^{l} \widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}} = K_r(\mathbf{X})$ and $[\{C^i_m\}_{i=1}^{\infty}] = \widetilde{\mathbf{C}}_{\mathbf{m}}$.

The only thing remained to be proved is that for ω-almost all indices i the resulting l-hyperpartitions are $F(l)$-regular. If it does not hold then there exist $1 \le r \le k$ and $1 \le j \le l$ such that for almost all i there exists a cylinder intersection $W_i \subset K_r(X_i)$, $|W_i| \ge F(l)|K_r(X_i)|$, such that

$$\left| \frac{|P^j_{r,i}|}{|K_r(X_i)|} - \frac{|P^j_{r,i} \cap W_i|}{|W_i|} \right| > F(l). \tag{3}$$

Let $\mathbf{W} = [\{W_i\}_{i=1}^{\infty}]$. Then $\mathbf{W} \in \sigma([r])^*$. Hence $\widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}}$ and $\mathbf{W}$ are independent sets. However, by (3)

$$\mu_{[r]}(\widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}})\mu_{[r]}(\mathbf{W}) \neq \mu_{[r]}(\widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}} \cap \mathbf{W}),$$

leading to a contradiction. □

3.5. The proof of the hypergraph sequence regularity lemma

Let us consider the ultralimit $\mathbf{H}$ of the hypergraph sequence $\{H_i\}_{i=1}^{\infty}$ as in the proof of the regularity lemma together with the l-hyperpartition $\widetilde{\mathcal{H}}$ given by the partition $\cup_{j=1}^{l} \widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}} = K_r(\mathbf{X})$, where $[\{P^j_{r,i}\}_{i=1}^{\infty}] = \widetilde{\mathbf{P}}^{\mathbf{j}}_{\mathbf{r}}$. If $s \ge 1$, then for ω-almost all indices

- $\cup_{j=1}^{l} P^j_{r,i} = K_r(X_i)$ is an $\frac{1}{s}$-equitable $\frac{1}{s}$-regular partition
- $|H_i \triangle \cup_{m=1}^{q} C^i_m| < \epsilon \binom{|X_i|}{k}$.
- T_i has combinatorial structure $\mathcal{C}$, where $T_i = \cup_{m=1}^{q} C^i_m$.

Also, by Lemmas 3.1 and 3.3

$$\lim_{\omega} t(F, T_i) = t(F, (\cup_{m=1}^{q} \widetilde{\mathbf{C}}_{\mathbf{m}})) = t(F, \mathcal{C}).$$

Thus for ω-almost all i, $|t(F, T_i) - t(F, \mathcal{C})| < \frac{1}{s}$. Therefore we can pick a subsequence H'_i satisfying the four conditions of the hypergraph sequence regularity lemma. □

3.6. Testability of hereditary properties

We omit here the definition of Property Testing but we state a theorem which is equivalent with the statement that hereditary hypergraph properties are testable.

Theorem 8. *Let $\mathcal{F}$ be a family of k-uniform hypergraphs. Then for every $\epsilon > 0$ there is a $\delta = \delta(\epsilon, \mathcal{F}) > 0$ and a natural number $n = n(\epsilon, \mathcal{F})$ such that if H satisfies $t_{\text{ind}}(F, H) \le \delta$ for every $F \in \mathcal{F}$ with $V(F) \le n$ then there is a hypergraph H' on the vertex set X of H with $|H \triangle H'| \le \epsilon \binom{|X|}{k}$ such that $t^0_{\text{ind}}(F, H') = 0$ for every $F \in \mathcal{F}$. (see also* [16,1,2]*).*

Proof. We proceed by contradiction. Assume that there is a sequence $\{H_i\}_{i=1}^{\infty}$ and $\epsilon > 0$ such that $\lim_{i\to\infty} t_{\text{ind}}(F, H_i) = 0$ for every $F \in \mathcal{F}$; however no member of the sequence can be modified in the way guaranteed by the theorem. Let us repeat the construction used in the proof of the Regularity Lemma again. Let $\mathbf{H}$ be the ultralimit hypergraph of $\{H_i\}_{i=1}^{\infty}$. We use Corollary 3.1 for the set $\mathbf{H}$ in order to obtain a separable realization ϕ and a measurable set $W \subseteq [0,1]^{r([k])}$ satisfying the statement of the corollary. Then $t_{\text{ind}}(F, W) = \lim_{\omega} t_{\text{ind}}(F, H_i) = 0$ for every $F \in \mathcal{F}$.

Thus there is an l-step Euclidean hypergraph (a union of l-boxes) W' such that $Vol(W \triangle W') \le \epsilon/4$. Let $\mathbf{Q}$ be the preimage of W' under ϕ. Denote by $\mathcal{C}$ the combinatorial structure of W'. As in the proof of the regularity lemma for each $1 \le r \le k$ we consider the partition $\mathbf{X}^r = \cup_{j=1}^{l} \mathbf{P}_\mathbf{r}^\mathbf{j}$, where $\mathbf{P}_\mathbf{r}^\mathbf{j} = \phi_{[r]}^{-1}(\frac{j-1}{l}, \frac{j}{l})$. We call the resulting l-hyperpartition $\widetilde{\mathcal{H}}$. The set $\mathbf{Q}$ is the union of some cells in $\widetilde{\mathcal{H}}$. Again, we modify the sets $\mathbf{P}_\mathbf{r}^\mathbf{j}$ to obtain the sets $[\{P_{r,i}^{j}\}_{i=1}^{\infty}] = \widetilde{\mathbf{P}}_\mathbf{r}^\mathbf{j}$. Consider the resulting l-hyperpartitions $\mathcal{H}_i$ on X_i and for every i denote the union of $\mathcal{H}_i$-cells with coordinates in $\mathcal{C}$ by Q_i. That is $Q_i = \mathbb{G}(W', \mathcal{H}_i, [n_i])$, where $[n_i]$ is the vertex set of H_i. Note that $\mathbb{G}(W', \mathcal{H}_i, [n_i])$ is a random hypergraph, nevertheless it always takes the same value. Then of course, $\mu(\mathbf{Q}' \triangle \mathbf{Q}) = 0$ where $\mathbf{Q}'$ is the ultralimit of the hypergraphs $\{Q_i\}_{i=1}^{\infty}$.

Now we consider the random hypergraph model $G_i = \mathbb{G}(W, \mathcal{H}_i, [n_i])$. For an ordered set $S = (i_1, i_2, \ldots, i_k) \in [n_i]^k$ let Y_S denote the random variable which takes 1 if S is in $G_i \triangle Q_i$ and takes 0 elsewhere. One can easily see that the expected value of Y_S is $l^{2^k-1}\text{Vol}((W' \cap B) \triangle (W \cap B))$ where B is the box representing the coordinate of the directed cell containing S. This shows that

$$E(|G_i \triangle Q_i|) = \sum_S E(Y_S) = \sum_C |C| l^{2^k-1} \text{Vol}((W' \cap B(C)) \triangle (W \cap B(C))),$$

where C runs through the directed cells of $\mathcal{H}_i$ and $B(C)$ is the box in $[0,1]^{2^k-1}$ corresponding to the coordinate of C.

Observe that $\lim_{\omega} |C_i^f|/(n_i)^k = l^{2^k-1}$ where C_i^f is the cell in $\mathcal{H}_i$ corresponding to the coordinate f. Indeed, the ultralimit of $\{C_i^f\}_{i=1}^{\infty}$ is a cell in the l-hyperpartition of $\mathbf{X}$. That is

$$\lim_{\omega} \frac{E(|G_i \triangle Q_i|)}{n_i^k} = \text{Vol}(W' \cap W) \le \epsilon/4.$$

On the other hand we know that $\lim_{\omega} \frac{|Q_i \triangle H_i|}{n_i^k} = |\mathbf{Q} \triangle \mathbf{H}| \le \epsilon/4$.

Consequently, $\lim_{\omega} \frac{E(|G_i \triangle H_i|)}{n_i^k} \le \epsilon/2$. Note that by probability 1, $t_{ind}(F, G_i) = 0$ for any $F \in \mathcal{F}$. That is there exists a hypergraph H_i' which is a value of the hypergraph valued random variable G_i such that

- $t_{ind}(F, H_i') = 0$.
- $\lim_{\omega} \frac{|H_i' \triangle H_i|}{n_i^k} < \epsilon$.

This leads to a contradiction. □

4. Uniqueness results and metrics

4.1. Distances of hypergraphs and hypergraphons

Let U and W be two measurable sets in $[0,1]^{r([k])}$. The distance $d_1(U, W)$ is defined as the measure of their symmetric difference $U \triangle W$. Let F be a k uniform hypergraph. It is clear from the definitions that

$$|t(F,U)-t(F,W)| \leq |E(F)| d_1(U,W).$$

We can also introduce a distance using subhypergraph-densities.

Let $\delta = \delta_w(U,W)$ denote the smallest number such that

$$|t(F,U)-t(F,W)| \leq |E(F)|\delta, \quad \text{for any } F.$$

Clearly, $\delta_w(U,W) \leq d_1(U,W)$. It is easy to see that δ_w satisfy the triangle inequality. On the other hand δ_w is only a pseudometric since (as we will see) there are different sets U and W with $\delta_w(U,W) = 0$. Our goal is to understand which two functions have distance 0 in the pseudometric δ_w.

For every set $S \in r([k])$ we denote by $\mathcal{A}_S$ the σ-algebra generated by the projection $[0,1]^{r([k])} \mapsto [0,1]^{r(S)}$. Let $\mathcal{A}^*_S$ denote the σ-algebra generated by all the algebras $\mathcal{A}_T$ where T is a proper subset of S for every $S \in r([k])$. We say that a measurable map $\phi : [0,1]^{r([k])} \mapsto [0,1]^{r([k])}$ is **structure preserving** if

1. ϕ is measure preserving.
2. $\phi^{-1}(\mathcal{A}_S) \subseteq \mathcal{A}_S$.
3. The sets $\phi_S^{-1}(I)$ are independent from $\mathcal{A}^*_S$ for every measurable set $I \subseteq [0,1]$.
4. $\phi \circ \pi = \pi \circ \phi$ for every permutation in S_k.

The following lemma shows that structure preserving maps do not change the homomorphism densities in hypergraphons.

Lemma 4.1. *For any structure preserving map ϕ we have that $\delta_w(U, \phi^{-1}(U)) = 0$.*

Proof. We need to prove that for any finite k-uniform hypergraph F

$$t(F,U) = t(F, \phi^{-1}(U)).$$

Mimicking the proof of Lemma 3.2 we can easily see that there exists a map $\hat{\phi} : [0,1]^{r([n],k)} \to [0,1]^{r([n],k)}$ such that $\hat{\phi}$ commutes with the S_n-action and

$$\phi \circ L_{[k]} = L_{[k]} \circ \hat{\phi}$$

where $L_{[k]}$ is the projection to the $[k]$-coordinates. Therefore, we have the following formula for the homomorphism sets:

$$\hat{\phi}^{-1}\left(\cap_{E\in E(F)} L_E^{-1}(L_{s_E}(U))\right) = \cap_{E\in E(F)} L_E^{-1}(L_{s_E}(\phi^{-1}(U))).$$

Hence the lemma follows. □

Definition 4.1. A structure preserving map $\psi : [0,1]^{r([k])} \mapsto [0,1]^{r([k])}$ is called a structure preserving equivalence if there is a structure preserving map ϕ such that both $\psi \circ \phi$ and $\phi \circ \psi$ are equivalent to the identity map on $[0,1]^{r([k])}$ (recall that equivalence means that two maps define the same measure algebra homomorphism).

Now we introduce the pseudodistance δ_1 by the formula

$$\delta_1(U,W) = \inf_{\phi,\psi} d_1(\phi^{-1}(U), \psi^{-1}(W)),$$

where ϕ and ψ run through all the structure preserving transformations. We will prove the following uniqueness theorem (see [11] for the graph case).

Theorem 9 (*Uniqueness I*). *$\delta_w(U, W) = 0$ if and only if there are two structure preserving measurable maps $\phi, \psi : [0,1]^{r([k])} \mapsto [0,1]^{r([k])}$ such that the measure of $\phi^{-1}(U) \triangle \psi^{-1}(W)$ is zero.*

Theorem 10 (*Uniqueness II*). *$\delta_w(U, W) = 0$ if and only if $\delta_1(U, W) = 0$.*

4.2. Technical lemmas

First we prove a simple real analysis lemma.

Lemma 4.2. *Let $Y \subseteq [0,1]^n$ be a measurable set independent from the σ-algebra $\mathcal{A}_{n-1}$ generated by the projection onto the first $(n-1)$-coordinates. Then there exist measurable subsets $X_k \subseteq [0,1]^n$ in the form*

$$X_k = (A_1^k \times B_1^k) \cup (A_2^k \times B_2^k) \cup \cdots \cup (A_{n_k}^k \times B_{n_k}^k)$$

such that $\lim_{k\to\infty} Vol(X_k \triangle Y) = 0$, where $A_1^k \cup A_2^k \cup \cdots \cup A_{n_k}^k$ is a measurable partition of $[0,1]^{n-1}$ and $\lambda(B_1^k) = \lambda(B_2^k) = \cdots = \lambda(B_{n_k}^k) = Vol(X_k)$. Obviously, the sets X_k are all independent from $\mathcal{A}_{n-1}$.

Proof. Fix a real number $\epsilon > 0$. Let H be a union of l-boxes in $[0,1]^n$ such that $l > \frac{1}{1000\epsilon^2}$ and $Vol(H \triangle Y) < \frac{\epsilon}{1000}$. By Fubini's Theorem, for almost all $z \in [0,1]^{n-1}$, $\lambda(A_z^Y) = Vol(Y)$, where

$$A_z^Y = \{t \in [0,1],\ (z,t) \in Y\}.$$

For each l-box T in $[0,1]^{n-1}$ let

$$H_T = \{s \in [0,1], T \times s \in H\}.$$

Lemma 4.3. *The number of l-boxes in $[0,1]^{n-1}$ for which $|\lambda(H_T) - Vol(Y)| > \frac{\epsilon}{10}$ is less than $\frac{\epsilon}{10} l^{n-1}$.*

Proof. By Fubini's Theorem,

$$\sum_T |\lambda(H_T) - Vol(Y)| \le Vol(H \triangle Y).$$

Hence the lemma follows. □

Now the set X_ϵ is constructed in the following way. Pick an integer m such that $|\frac{m}{l} - Vol(Y)| < \frac{\epsilon}{10}$. If for an l-box T $|\lambda(H_T) - Vol(Y)| < \frac{\epsilon}{10}$ then add or delete less than $\frac{\epsilon}{10} l$ l-boxes of H above T to obtain exactly m boxes. On the other hand if $|\lambda(H_T) - Vol(Y)| \ge \frac{\epsilon}{10}$, then just pick m arbitrary boxes above T. Then X_ϵ is in the right form and $Vol(X_\epsilon \triangle Y) \to 0$ as $\epsilon \to 0$. □

The following lemma establishes the functorality of separable realizations and structure preserving maps.

Lemma 4.4. *Let $\phi : \mathbf{X}^k \to [0,1]^{r([k])}$ be a separable realization and $\rho : [0,1]^{r([k])} \to [0,1]^{r([k])}$ be a structure preserving map. Then $\rho \circ \phi$ is a separable realization as well. Similarly the composition of two structure preserving maps, or the inverse of a structure preserving equivalence is a structure preserving map.*

Proof. For the first part it is enough to prove that if $M \subseteq [0,1]^{r([k])}$, $M \in \mathcal{A}_S$ for some $S \subseteq [k]$ such that M is independent from $\mathcal{A}_S^*$ then $\phi^{-1}(M)$ is independent from $\sigma(S)^*$.

First suppose that M is in block-form that is

$$M = \cup_{i=1}^n (A_i \cap B_i),$$

where for any $1 \leq i \leq n$, $B_i \in \mathcal{B}_S$ and $A_i \in \mathcal{A}_S^*$ so that $\cup_{i=1}^n A_i$ is a measurable partition of $[0,1]^{r([k])}$. Let $\mathbf{I} \in \sigma(S)^*$. Then

$$\phi^{-1}(M) \cap \mathbf{I} = \cup_{i=1}^n (\phi^{-1}(A_i) \cap \mathbf{I}) \cap \phi^{-1}(B_i).$$

Hence

$$\mu(\phi^{-1}(M) \cap \mathbf{I}) = \sum_{i=1}^n \mu(\phi^{-1}(A_i) \cap \mathbf{I}) \mu(\phi^{-1}(B_i)).$$

Note that $\mu(\phi^{-1}(B_i)) = Vol(M)$ and $\sum_{i=1}^n \mu(\phi^{-1}(A_i) \cap \mathbf{I}) = \mu(\mathbf{I})$. Therefore $\phi^{-1}(M)$ is independent from $\sigma(S)^*$. By Lemma 4.2, any set in $\mathcal{A}_S$ which is independent from $\mathcal{A}_S^*$ can be approximated by sets in block-form, thus the proof of the first part of our lemma follows. The second part can be proved completely similarly. □

The following lemma is a baby-version of the Total Independence Lemma.

Lemma 4.5. *For any $S \subseteq [k]$, let $X_S \in \mathcal{A}_S$ such that X_S is independent from $\mathcal{A}_S^*$. Then $\{X_S\}_{S \subseteq [k]}$ is a totally independent system.*

Proof. We need to prove that for any set-system $\{S_i\}_{i=1}^r \subset r([k])$

$$Vol(\cap_{i=1}^r X_{S_i}) = \prod_{i=1}^r Vol(X_{S_i}). \tag{4}$$

Let us proceed by induction. Suppose (4) holds for a certain r. Let $\{S_i\}_{i=1}^{r+1} \subset r([k])$ be a set-system and suppose that S_{r+1} is not a subset of S_j, for $1 \leq j \leq r$. It is enough to see that

$$Vol\left(X_{S_{r+1}} \cap \bigcap_{i=1}^r X_{S_i}\right) = \prod_{i=1}^{r+1} Vol(X_{S_i}). \tag{5}$$

By Lemma 4.2 we may assume that $X_{S_{r+1}}$ is in the block-form $\cup_{i=1}^n (A_i \cap B_i)$, where $\cup_{i=1}^n A_i$ is a partition of $[0,1]^{r([k])}$ such that $\{A_i\}_{i=1}^n$ are in the σ-algebra $\mathcal{C}_S$ generated by $\{\mathcal{A}_S\}_{S \subset [k],\, S \neq S_{r+1}}$ and $\{B_i\}_{i=1}^n \subset \mathcal{B}_S$. Since $\cap_{i=1}^r X_{S_i} \in \mathcal{C}_S$, (5) follows. □

We shall need the auxiliary notion of structure preserving measure algebra embeddings. Let $\mathcal{L}^{r([k])}$ denote the measure algebra associated to $([0,1]^{r([k])}, \mathcal{B}, \lambda)$. For any $S \subseteq [k]$ let $\mathcal{B}_S$ be the subalgebra generated by the S-coordinate, that is for any $S \subset [k]$, $\{B_T\}_{T \subseteq S}$ are jointly independent subalgebras generating $\mathcal{A}_S$. We say that an injective homomorphism $\Phi : \mathcal{L}^{r([k])} \mapsto \mathcal{L}^{r([k])}$ is a **structure preserving embedding** if

1. Φ is measure preserving.
2. $\Phi(\mathcal{B}_S) \subset \mathcal{A}_S$ for any $S \subseteq [k]$.
3. $\Phi(\mathcal{B}_S)$ is independent of $\mathcal{A}_S^*$.
4. $\Phi \circ \pi = \pi \circ \Phi$ for every permutation in S_k.

Lemma 4.6. *Let $\Phi : \mathcal{L}^{r([k])} \mapsto \mathcal{L}^{r([k])}$ be a (measure algebra) structure preserving embedding. Then Φ can be represented (see Lemma A.1) by a structure preserving map $\phi : [0,1]^{r([k])} \mapsto [0,1]^{r([k])}$.*

Proof. Let us consider the map $\Phi_{[i]} : \mathcal{B}_{[i]} \to \mathcal{A}_{[i]}$. By the fourth axiom of structure preserving embeddings the image of $\Phi_{[i]}$ consists of $S_{[i]}$-invariant elements. We claim that we can represent $\Phi_{[i]}$ by maps $\phi_{[i]} : [0,1]^{r([i])} \mapsto [0,1]$ such that $\phi^{-1}(I)$ is $S_{[i]}$ invariant for every measurable set $I \subseteq [0,1]$. First we represent $\Phi_{[i]}$ by a measurable map $\phi'_{[i]}$. Now Lemma 6.8 implies that $S_{[i]}$ acts freely on $[0,1]^{r([i])}$ withe measurable sets $Q_1, Q_2, \dots, Q_{i!}$. Let $G = \cup_i Q_i$. If $x \in G$ then we define $\phi_{[i]}(x)$ as $\phi'[i](\pi(x))$ where $\pi \in S_{[i]}$ is the unique permutation with $\pi(x) \in Q_1$. If $x \in [0,1]^{r([i])} \setminus G$ the $\phi_{[i]}(x)$ is defined to be 0. For a general set $S \in r([k])$ with $|S| = i$ we define $\phi_S(x)$ to be $\phi_{[i]}(\pi(x))$ where $\pi \in S_{[k]}$ is an arbitrary permutation with $\pi(S) = [i]$. The $S_{[i]}$-invariance of $\phi_{[i]}$ guarantees that $\phi_S; [0,1]^{r(S)} \to [0,1]$ is well defined and represents the map $\Phi_S : \mathcal{B}_S \to \mathcal{A}_S$. It is easy to see that the map $\times_{S \in r([k])} \phi_S \circ L_S$ is a structure preserving map which represents Φ. □

Lemma 4.7. *Let $W \subseteq [0,1]^{r([k])}$ be an l-step hypergraphon and let $\phi : [0,1]^{r([k])} \mapsto [0,1]^{r([k])}$ be a structure preserving map with $T = \phi^{-1}(W)$. Then there is a structure preserving equivalence (see Definition 4.1) ψ such that $W \triangle \psi^{-1}(T)$ has measure 0.*

Proof. Let P_t^i denote the set $\phi_{[t]}^{-1}([(i-1)/l, i/l)) \in \mathcal{A}_{[t]}$ for $t = 1, 2, \dots, k$. By the definition of structure preserving maps the set P_t^i is independent from $\mathcal{A}^*_{[t]}$, has measure $1/l$ and is symmetric under $S_{[t]}$. Using Lemma 6.6, for every $t = 1, 2, \dots, k$ we construct a σ-algebra $\mathcal{C}_{[t]} \subseteq \mathcal{A}_{[t]}$ such that

1. $\mathcal{C}_{[t]}$ is an independent complement for $\mathcal{A}^*_{[t]}$ in $\mathcal{A}_{[t]}$
2. $P_t^i \in \mathcal{C}_{[t]}$ for $1 \le i \le l$
3. Every set in $\mathcal{C}_{[t]}$ is invariant under the symmetric group $S_{[t]}$.

In general, for a set $S \in r([k])$, we introduce $\mathcal{C}_S$ as $\pi(\mathcal{C}_{[|S|]})$ where π is an arbitrary permutation taking $\mathcal{A}[|S|]$ to $\mathcal{A}_S$. By the invariance of $\mathcal{C}_{[|S|]}$ this is well defined.

Now the system of σ-algebras $\{\mathcal{C}_S\}_{S \in r([k])}$ satisfies the following properties.

1. The σ-algebras $\mathcal{C}_S$ generate $[0,1]^{r([k])}$ where S runs through the elements in $r([k])$.
2. $\mathcal{C}_S \subset \mathcal{A}_S$ and $\mathcal{C}_S$ is independent from $\mathcal{A}^*_S$. That is by Lemma 4.5 the algebras $\mathcal{C}_S$ are totally independent.

Now let $\rho_{[t]}$ be a measure algebra isomorphism from $[0,1]$ to $\mathcal{C}_{[t]}$ taking $[(i-1)/l, i/l)$ to P_t^i. Using the $S_{[k]}$ action we also define maps ρ_S for every $S \in r([k])$ satisfying $\pi \circ \rho_S = \rho_S \circ \pi$ for every $\pi \in S_{[k]}$. Since the algebras $\mathcal{C}_S$ are totally independent, by Lemma A.2, the product of the maps ρ_S creates a measure algebra equivalence from $[0,1]^{r([k])}$ to itself which is a structure preserving equivalence.

Lemma 4.8. *For every pair $U, W \subseteq [0,1]^{r([k])}$ of hypergraphons and $\epsilon > 0$ there is a structure preserving equivalence $\phi : [0,1]^{r([k])} \mapsto [0,1]^{r([k])}$ such that $d_1(U, \phi^{-1}(W)) \le \delta_1(U, W) + \epsilon$.*

Proof. Let T_1, T_2 be two l-step hypergraphons with $d_1(T_1, W) \le \epsilon/8$ and $d_1(T_2, U) \le \epsilon/8$. We know that there are two structure preserving maps ψ_1 and ψ_2 such that $d_1(\psi_2^{-1}(U), \psi_1^{-1}(W)) \le \delta_1(U, W) + \epsilon/8$. By Lemma 4.7 there are structure preserving equivalences ρ_1 and ρ_2 with $d_1(\rho_1^{-1}(T_1), \psi_1^{-1}(T_1)) = 0$ and $d_1(\rho_2^{-1}(T_2), \psi_2^{-1}(T_2)) = 0$. Now

$$d_1(\rho_1^{-1}(T_1), \rho_2^{-1}(T_2)) \le \delta_1(U, W) + \epsilon/4.$$

By Lemma 4.4, $\rho_2 \circ \rho_1^{-1}$ is a structure preserving equivalence that takes W into a set whose distance from U is at most $\delta_1(U, W) + \epsilon$.

4.3. A concentration result for W-random graphs

Theorem 11 (*Concentration*)**.** *Let $W \subseteq [0,1]^{r([k])}$ be a hypergraphon. Then*

$$\Pr(|t_0(F, \mathbb{G}(W,[n])) - t(F,W)| \geq \epsilon) \leq 2\exp\left(-\frac{\epsilon^2 n}{2|V(F)|^2}\right).$$

The proof of the lemma is identical with the proof of Theorem 2.5 in [10], that was used for the case $k = 2$. For the sake of completeness we repeat the proof.

Proof. Let us consider the system of random hypergraph models $G_1, G_2, \dots, G_n$ such that the distribution of G_n is $\mathbb{G}(W,[n])$ and G_i is the sub-hypergraph in G_n induced by $[i]$. It is clear that the distribution on G_i is the same as $\mathbb{G}(W,[i])$. Let F be a fixed k-uniform hypergraph on the vertex set $[r]$. For any injective map $\psi : [r] \mapsto [n]$ we denote by A_ψ the event that ψ is a homomorphism from F to G_n. Let

$$B_m = \frac{(n-r)!}{n!}\sum_\psi \Pr(A_\psi \mid G_m).$$

The sequence $B_0, B_1, \dots, B_n$ is a martingale, where $B_0 = \Pr(A_\psi) = t(F,W)$ and $B_n = \Pr(A_\psi \mid G_n)$ is 1 if ψ is a homomorphism and 0 elsewhere. This implies that $B_n = t_0(F, G_n)$. Now we have that

$$|B_m - B_{m-1}| \leq \frac{(n-r)!}{n!}\sum_\psi \big|\Pr(A_\psi \mid G_m) - \Pr(A_\psi \mid G_{m-1})\big|.$$

The terms in the sum for which m is not in the range of ψ are 0 and all the other terms are at most one. The number of terms of the second type is $r\frac{(n-r)!}{(n-1)!}$ and so $|B_m - B_{m-1}| \leq r/n$. By applying Azuma's inequality we get that

$$\Pr(|t_0(F,G_n) - t(F,W)| \geq \epsilon) = \Pr(|B_n - B_0| \geq \epsilon) \leq 2\exp\left(\frac{-\epsilon^2}{2n(r/n)^2}\right)$$
$$= 2\exp\left(-\frac{\epsilon^2 n}{2r^2}\right). \quad \square$$

Theorem 12 (*Convergence*)**.** *The sequences $t_0(F, \mathbb{G}(W,[n]))$ and $t(F, \mathbb{G}(W,[n]))$ converge to $t(F,W)$ with probability one as n goes to infinity.*

Proof. The convergence of $t_0(F, \mathbb{G}(W,[n]))$ follows from Theorem 11 and the Borel–Cantelli lemma since for every fixed $\epsilon > 0$ the sum of the right hand side in the inequality is finite. $\square$

4.4. Proof of the uniqueness theorems

Let $\mathbf{X}$ be the ultraproduct of the sets $[n]$. Let Z_n be the random variable which is a random point in $[0,1]^{r([n],k)}$ with uniform distribution as in Section 2.6 and let $\{\tau^n\} : [n]_0^k \to [0,1]^{r([k])}$ be the associated random coordinate systems. The ultraproduct function $\tau = [\{\tau^n\}_{n=1}^\infty]$ on $\mathbf{X}^k$ will also be called random coordinate system.

Lemma 4.9. *The random coordinate system $\tau : \mathbf{X}^k \mapsto [0,1]^{r([k])}$ is a separable realization with probability one.*

Proof. Let $I \subseteq [0,1]$ be a measurable set and $S \in r([k])$. It is clear from the definition that $\tau_S^{-1}(I)$ is in $\sigma(S)$. We show that (with probability one) $\tau_S^{-1}(I)$ is independent from $\sigma(S)^*$ and has measure $\lambda(I)$.

Let $I_{a,b}$ be an open interval with rational endpoints a, b. Let $\hat{\mathbf{I}}_{a,b}$ denote the ultraproduct $[\{(\tau_S^n)^{-1}(I_{a,b})\}_{n=1}^{\infty}]$. By Proposition 6.1 (and the remark after the proposition) we have that almost surely $\hat{\mathbf{I}}_{a,b}$ has measure $b-a$ and is independent from $\sigma(S)^*$. Then we have

$$\hat{\mathbf{I}}_{a+\epsilon,b-\epsilon} \subseteq \tau_S^{-1}(I_{a,b}) \subseteq \hat{\mathbf{I}}_{a,b}$$

for every small enough rational number $\epsilon > 0$. Since there are only countable many rational numbers this holds simultaneously for every rational number with probability 1. This implies that $\tau_S^{-1}(I_{a,b})$ has measure $b-a$ and is independent from $\sigma(S)^*$ with probability 1. Since τ_S is measurable and measure preserving on rational intervals it has to be measure preserving on Lebesgue sets. By approximating an arbitrary measurable sets by unions of disjoint intervals we get the independence from $\sigma(S)^*$.

Now let $B \subseteq [0,1]^{r([k])}$ be a box of the form $\prod_{S\in r([k])} I_S$ where I_S is an interval with rational endpoints. The measure of B is equal to $\prod_{S\in r([k])} \lambda(I_S)$. The set $\tau^{-1}(B)$ is equal to

$$\cap_{S\in r([k])} \tau_S^{-1}(I_S).$$

Therefore using the total independence theorem we obtain that with probability one $\tau^{-1}(B) = \lambda(B)$. Again this holds simultaneously for every rational interval system with probability 1. As a consequence τ is almost surely a measure preserving map.

The symmetry on τ under S_k is clear from its definition. □

Lemma 4.10. *Let W be a hypergraphon. Then with probability one the ultraproduct $\mathbf{H} = [\{\mathbb{G}(W,[n])\}_{n=1}^{\infty}] \subseteq \mathbf{X}^k$ has a separable realization $\phi : \mathbf{X}^k \mapsto [0,1]^{r([k])}$ such that $\mathbf{H}\triangle\phi^{-1}(W)$ has measure 0.*

Proof. We will use that the set $\mathbf{H} = [\{\mathbb{G}(W,[n])\}_{n=1}^{\infty}]$ can be written as the ultraproduct $[\{(\tau^n)^{-1}(W)\}_{n=1}^{\infty}]$. Our goal is to prove that almost surely $\mathbf{H}\triangle\tau^{-1}(W)$ has measure 0. First by applying Theorem 12 to a single hyperedge F we deduce that $\mathbf{H}$ has measure $\lambda(W)$ with probability one.

If W is open then $\tau^{-1}(W)$ is contained in $\mathbf{H}$ and, by Lemma 4.9, has measure $|W|$ with probability 1. This means that with probability 1 the set $\tau^{-1}(W)\triangle\mathbf{H}$ has measure 0.

For an arbitrary measurable set $W \subseteq [0,1]^{r([k])}$ and $\epsilon > 0$ there are open sets O_1 and O_2 in $[0,1]^{r([k])}$ such that $O_1 \setminus O_2 \subseteq W \subseteq O_1$ and $|O_2| \le \epsilon$. We have that

$$(\tau^n)^{-1}(O_1) \setminus (\tau^n)^{-1}(O_2) \subseteq (\tau^n)^{-1}(W) \subseteq (\tau^n)^{-1}(O_1)$$

and thus by taking the ultra product

$$[\{(\tau^n)^{-1}(O_1)\}_{n=1}^{\infty}] \setminus [\{(\tau^n)^{-1}(O_2)\}_{n=1}^{\infty}] \subseteq \mathbf{H} \subseteq [\{(\tau^n)^{-1}(O_1)\}_{n=1}^{\infty}].$$

Using our observation about open sets and that τ is measure preserving with probability 1 we obtain that the measure of $\mathbf{H}\triangle\tau^{-1}(W)$ is at most ϵ. By Lemma 4.9 the proof is complete. □

Proof of Theorem 9. Let U and W be two functions with $\delta_w(U, W) = 0$. This means that $\mathbb{G}(U, [n])$ and $\mathbb{G}(W, [n])$ are equal to the same distribution Z_n. Let $\mathbf{H} = [\{Z_n\}_{n=1}^{\infty}]$. By Lemma 4.10 with probability one there are two separable realizations $\phi_1, \phi_2 : \mathbf{X}^n \mapsto [0, 1]^{r([k])}$ such that $\phi_1^{-1}(U), \phi_2^{-1}(W)$ and $\mathbf{H}$ differ only in a zero measure set. Let $\mathcal{A}$ denote the separable sigma algebra generated by ϕ_1 and ϕ_2. By the Euclidean correspondence (Theorem 6) there is a separable realization $\phi_3 : \mathbf{X}^k \mapsto [0, 1]^{r([k])}$ corresponding to the algebra $\mathcal{A}$. The maps ϕ_1 and ϕ_2 define unique structure preserving maps ψ_1 and ψ_2 on the measure algebra $\mathcal{L}^{r([k])}$ such that $(\phi_3)^{-1}\psi_i(S)$ is equivalent with $\phi_i^{-1}(S)$. This means that $\psi_1(U) = \psi_2(W)$ in the measure algebra $\mathcal{L}^{r([k])}$. Therefore by Lemma 4.6 our theorem follows. □

Proof of Theorem 10. By the previous theorem, if $\delta_w(U, V) = 0$, then $\delta_1(U, V) = 0$. On the other hand, if $\delta_1(U, V) = 0$ then by the fact that $\delta_w(U, V) \leq d_1(U, V)$ and Lemma 4.1, $\delta_w(U, V) = 0$. □

4.5. *The counting lemma*

Let $\mathcal{C} \subseteq [l]^{r([k])}$ be a symmetric combinatorial structure. Let V be a finite set. An (l, k)-map is a function from $r(V, k)$ to $[l]$. If $E \subseteq V$ has k-elements then the restriction of an (l, k)-map f to E is an element x in $[l]^{r(E)}$. By specifying an arbitrary bijection g between E and $[k]$ we can also represent x by an element x' in $[l]^{r([k])}$. The S_k-orbit of x' does not depend on g and so we can talk about the S_k-orbit determined by the restriction of f to E.

Let F be a k-uniform hypergraph on V, let $\mathcal{C} \subseteq [l]^{r([k])}$ be a symmetric combinatorial structure and let f be an (l, k)-map on V. We say that f is a **homomorphism** from F to $\mathcal{C}$ if the restriction of f to any edge of F determines an S_k orbit which is in $\mathcal{C}$.

The **homomorphism density** $t(F, \mathcal{C})$ is the probability that a random (l, k)-map on V is a homomorphism. Note that here we take the uniform probability distribution on all (l, k)-maps. For technical reasons we will also need the number $t(F, \mathcal{C}, P)$ which is the probability that an (l, k)-map chosen with distribution P on $[l]^{r(V,k)}$ (the set of all (l, k)-functions) is a homomorphism.

Let $\mathcal{H}$ be an l-hyperpartition on a finite set U. Every injective map $g : V \mapsto U$ induces an (l, k)-map f_g on V such that for a set $S \in r(V, k)$ the value $f(S)$ is the index i of the partition set $P^i_{|S|}$ containing $g(S)$. Let $D(V, \mathcal{H})$ denote the probability distribution of f_g if g is chosen uniformly at random from all the injective maps $g : V \mapsto U$. Using this notation the following lemma follows immediately from the definitions.

Lemma 4.11. *Let $\mathcal{C} \subseteq [l]^{r([k])}$ be a symmetric combinatorial structure and let H be a hypergraph on the set U which is the union of $\mathcal{H}$-cells with coordinates in $\mathcal{C}$. Then the probability $t^0(F, H)$ that a random injective map $g : V \mapsto U$ is a homomorphism from F to H is equal to $t(F, \mathcal{C}, D(V, \mathcal{H}))$.*

Theorem 13 (*Counting Lemma*)**.** *Let $\{U_i\}_{i=1}^{\infty}$ be increasing finite sets with l-hyperpartitions $\{\mathcal{H}_i\}_{i=1}^{\infty}$ such that $\mathcal{H}_i$ is ϵ_i-regular and δ_i-equitable with $\lim_{i\to\infty} \epsilon_i = \lim_{i\to\infty} \delta_i = 0$. Let furthermore $\mathcal{C} \subseteq [l]^{r([k])}$ be a symmetric combinatorial structure and H_i be the union of $\mathcal{H}_i$-cells with coordinates in $\mathcal{C}$. Then for every finite hypergraph F we have that*

$$\lim_{i\to\infty} t(F, H_i) = t(F, \mathcal{C}).$$

Proof. Let V denote the vertex set of F. Since $\{U_i\}_{i=1}^{\infty}$ is an increasing sequence of sets we have that

$$\lim_{i\to\infty} t^0(F, H_i) = \lim_{i\to\infty} t(F, H_i).$$

Now by Lemma 4.11 it suffices to show that $\lim_{i\to\infty} D(V, \mathcal{H}_i)$ is the uniform distribution on $[l]^{r(V,k)}$. We proceed by contradiction. By choosing an appropriate subsequence of $\{\mathcal{U}_i\}_{i=1}^{\infty}$ we can assume that the limit of $D(V, \mathcal{H}_i)$ exists and it is not uniform. This means that there is a function $f : r(V, k) \mapsto [l]$ such that

$$\lim_{i\to\infty} p_i \neq l^{-|r(V,k)|}, \quad \text{where } p_i = P(f_g = f \mid g : V \mapsto U_i, g \text{ is injective})$$

holds. The set of all injective maps from V to U_i can be represented as the collection of elements in U_i^V with no repetitions in the coordinates. This subset in U_i^V has relative density tending to 1 as i goes to infinity. Now let $T_i \subseteq U_i^V$ defined by

$$T_i := \bigcap_{S\in r(V,k)} \pi_S^{-1}(P_{|S|}^{f(S)}).$$

For $S \subset r(V, k)$, $\pi_S : U_i^V \to U_i^{[|S|]}$ is defined as $L_{\rho_S} \circ L_S$, where $L_S : U_i^V \to U_i^S$ is the natural projection and L_{ρ_S} is given by a bijection $\rho_S : S \to [|S|]$. Here $P_{|S|}^{f(S)}$ denotes the corresponding partition set in $\mathcal{H}_i$. Since $P_{|S|}^{f(S)}$ is symmetric in its coordinates the set T_i is independent of the concrete choice of the bijections ρ_S. Let $\mathbf{X}$ denote the ultraproduct $[\{U_i\}_{i=1}^{\infty}]$ and let $\mathbf{H}^f$ denote the ultraproduct $[\{T_i\}_{i=1}^{\infty}] \subseteq \mathbf{X}^V$. Furthermore for every $S \in r(V, k)$ let $\mathbf{H}_S^f$ denote the ultralimit of the partition sets $\pi_S^{-1}(P_{|S|}^{f(S)})$ from $\mathcal{H}_i$ where i tends to infinity. Then

$$\mathbf{H}^f = \bigcap_{S\in r(V,k)} \mathbf{H}_S^f.$$

Also, the measure of $\mathbf{H}^f$ is equal to $\lim_\omega p_i = \lim_{i\to\infty} p_i$. Now the condition $\lim_{i\to\infty} \epsilon_i = \lim_{i\to\infty} \delta_i = 0$ implies that for every $S \in r(V, k)$ the set $\mathbf{H}_S^f$ has measure l^{-1} and that $\mathbf{H}_S^f \in \sigma(S)^*$. The total independence theorem implies that the measure of $\mathbf{H}$ is $l^{-|r(V,k)|}$ providing a contradiction. □

4.6. Equivalence of convergence notions and the inverse counting lemma

Let $W \subseteq [0, 1]^{r([k])}$ be a hypergraphon. We say that a sequence of hypergraphs $\{H_i\}_{i=1}^{\infty}$ is **structurally converges** to W if for every l-step hypergraphon U with $\delta_1(W, U) \le \epsilon$ and combinatorial structure $\mathcal{C}$ there is a sequence of l-hyperpartitions $\mathcal{H}_i$ on the vertex sets of H_i such that

1. $\mathcal{H}_i$ is δ_i-regular and δ_i-equitable with $\lim_{i\to\infty} \delta_i = 0$.
2. The union T_i of $\mathcal{H}_i$-cells with coordinates in $\mathcal{C}$ satisfies $\limsup_{i\to\infty} d_1(T_i, H_i) \le \epsilon$.

Definition 4.2. We say that an l-step hypergraphon U with combinatorial structure $\mathcal{C}$ is (ϵ, δ)-close to a hypergraph H if there is an l-hyperpartition $\mathcal{H}$ on the vertex set of H such that

1. $\mathcal{H}$ is both δ-regular and δ-equitable.
2. The union T_U of $\mathcal{H}$-cells with combinatorial structure $\mathcal{C}$ satisfies $d_1(H, T_U) \le \epsilon$.

Theorem 14. *For an increasing sequence* $\{H_i\}_{i=1}^\infty$ *of k-uniform hypergraphs the following statements are equivalent:*

1. $\{H_i\}$ *is strongly convergent*
2. $\{H_i\}$ *is weakly convergent*
3. $\{H_i\}$ *structurally converges to a hypergraphon W which is also the weak limit of* $\{H_i\}$.

Proof. Let us start with (2) implies (3). By Theorem 7 we know that there is a hypergraphon W such that $lim_{i\to\infty} t(F, H_i) = t(F, W)$. Assume by contradiction that $\{H_i\}$ is not structurally convergent to W. Then for some $\epsilon > 0$ there is a $\delta > 0$, an l-step hypergraphon U of combinatorial structure $\mathcal{C}$ with $\delta_1(U, W) \le \epsilon$ and an infinite subsequence $\{J_i\}$ of $\{H_i\}$ such that none of the elements of $\{J_i\}$ is $(\delta, \epsilon + \delta)$-close to T. Let $\mathbf{J}$ be the ultraproduct hypergraph $[\{J_i\}_{i=1}^\infty] \subseteq \mathbf{X}^k$ and let $\phi : \mathbf{X}^k \mapsto [0, 1]^{r([k])}$ be a separable realization of $\mathbf{J}$. That is for some $V \subset [0, 1]^{r([k])}$, $\mathbf{J} \triangle \phi^{-1}(V)$ has measure zero. By Proposition 3.1, $\delta_w(V, W) = 0$. By Theorem 10, $\delta_1(V, W) = 0$ and consequently $\delta_1(U, V) \le \epsilon$. By Lemma 4.7 there exists a measure preserving equivalence ρ with $d_1(\rho^{-1}(U), V) \le \epsilon + \delta/2$. This means that $(\rho \circ \phi)^{-1}(U) \triangle \mathbf{J}$ has measure at most $\epsilon + \delta/2$. By Lemma 4.4 $\rho \circ \phi$ is a separable realization; hence $(\rho \circ \phi)^{-1}(U)$ is a cell system with combinatorial structure $\mathcal{C}$ of a 0-regular and 0-equitable hyperpartition on $\mathbf{X}$. This leads to a contradiction.

The implication (3) $\Rightarrow$ (1) is trivial.

The implication (1) $\Rightarrow$ (2) follows from the Counting Lemma (Theorem 13). Let us fix a k-uniform hypergraph F on the vertex set V and with edge set E. According to the definition of strong convergence for every $\epsilon > 0$ there is a fixed combinatorial structure $\mathcal{C}$ and modifications H_i' of H_i with an at most ϵ-density edge set such that every H_i' is the union of the cells with coordinates in $\mathcal{C}$ of some hyperpartition which is getting more and more regular and balanced as i tends to infinity. The Counting Lemma implies that $\lim_{t\to\infty} t(F, H_i') = t(F, \mathcal{C})$. On the other hand $|t(F, H_i) - t(F, H_i')| \le |E|\epsilon$. Using this inequality for every $\epsilon > 0$, we obtain the convergence of $t(F, H_i)$. □

The following immediate corollary states that if two hypergraphs have similar sub-hypergraph densities then they have similar regular partitions.

Corollary 4.1 (*Inverse Counting Lemma*)**.** *Fix* $k > 0$. *Then for any* $\epsilon > 0$ *there exist positive constants* $\delta = \delta(\epsilon), C = C(\epsilon), N = N(\epsilon)$ *such that if* H_1, H_2 *are two k-uniform hypergraphs,* $|V(H_1)| \ge N, |V(H_2)| \ge N$ *and* $\delta_1(H_1, H_2) < \delta$, *then there exists an l-hyperpartition U,* $1 < l \le C$ *so that both hypergraphs are* (ϵ, ϵ)*-close to U.*

We also have a corollary of the Counting Lemma, using the notion of (ϵ, δ)-closeness.

Corollary 4.2 (*Counting Lemma Finitary Version*)**.** *For any finite k-uniform hypergraph F, l-step hypergraphon U and* $\epsilon > 0$ *there is a constant* $\delta = \delta(F, U, \epsilon)$ *such that if a k-uniform hypergraph H is* (δ, δ)*-close to U then*

$$|t(F, U) - t(F, H)| < \epsilon.$$

(see also [13]*).*

5. The proof of the total independence theorem

Let $\{X_i\}_{i=1}^\infty$ be finite sets as in Section 2 and $f_i : X_i \to [-d, d]$ be real functions, where $d > 0$. Then one can define a function $\mathbf{f} : \mathbf{X} \to [-d, d]$ whose value at $\overline{p} = [\{p_i\}_{i=1}^\infty]$ is the ultralimit of $\{f_i(p_i)\}_{i=1}^\infty$. We say that $\mathbf{f}$ is the ultraproduct of the functions $\{f_i\}_{i=1}^\infty$. We shall use the notation $\mathbf{f} = [\{f_i\}_{i=1}^\infty]$. Note that the characteristic function of the ultraproduct of sets is exactly the ultraproduct of their characteristic functions. From now on we call such bounded functions **ultraproduct functions**.

Lemma 5.1. *The ultraproduct functions are measurable on* $\mathbf{X}$ *and*

$$\int_{\mathbf{X}} \mathbf{f}\, d\mu = \lim_\omega \frac{\sum_{p\in X_i} f_i(p)}{|X_i|}.$$

Proof. Let $-d \le a \le b \le d$ be real numbers. It is enough to prove that $\mathbf{f}_{[a,b]} = \{\overline{p} \in \mathbf{X} \mid a \le \mathbf{f}(\overline{p}) \le b\}$ is measurable. Let $f^i_{[a,b]} = \{p \in X_i \mid a \le f_i(p) \le b\}$. Note that $[\{f^i_{[a,b]}\}_{i=1}^\infty]$ is not necessarily equal to $\mathbf{f}_{[a,b]}$. Nevertheless if

$$P_n := \left[\left\{f^i_{[a-\frac{1}{n},b+\frac{1}{n}]}\right\}_{i=1}^\infty\right],$$

then $P_n \in \mathcal{P}$ and $\mathbf{f}_{[a,b]} = \bigcap_{n=1}^\infty P_n$. This shows that $\mathbf{f}_{[a,b]}$ is a measurable set. Hence the function $\mathbf{f}$ is measurable.

Now we prove the integral formula. Let us consider the function g_i on X_i which takes the value $\frac{j}{2^k}$ if f_i takes a value not smaller than $\frac{j}{2^k}$ but less than $\frac{j+1}{2^k}$ for $-N_k \le j \le N_k$, where $N_k = d2^k + 1$. Clearly $|[\{g_i\}_{i=1}^\infty] - \mathbf{f}| \le \frac{1}{2^k}$ on $\mathbf{X}$. Observe that $\mathbf{g} = [\{g_i\}_{i=1}^\infty]$ is a measurable step-function on $\mathbf{X}$ taking the value $\frac{j}{2^k}$ on $C_j = [\{f^i_{[\frac{j}{2^k},\frac{j+1}{2^k})}\}_{i=1}^\infty]$. Hence,

$$\int_X \mathbf{g}\, d\mu = \sum_{-N_k}^{N_k} \frac{j}{2^k}\mu(C_j) = \lim_\omega \left(\sum_{j=-N_k}^{N_k} \frac{j}{2^k} \frac{\left|f^i_{[\frac{j}{2^k},\frac{j+1}{2^k})}\right|}{|X_i|} \right).$$

Also, $|\mathbf{g} - \mathbf{f}| \le \frac{1}{2^k}$ on $\mathbf{X}$ uniformly, that is $|\int_{\mathbf{X}} \mathbf{f}\, d\mu - \int_{\mathbf{X}} g\, d\mu| \le \frac{1}{2^k}$. Notice that for any $i \ge 1$

$$\left| \sum_{j=-N_k}^{N_k} \frac{\left|f^i_{[\frac{j}{2^k},\frac{j+1}{2^k})}\right|}{|X_i|} \frac{j}{2^k} - \frac{\sum_{p\in X_i} f_i(p)}{|X_i|} \right| \le \frac{1}{2^k}.$$

Therefore for each $k \ge 1$,

$$\left| \int_{\mathbf{X}} \mathbf{f}\, d\mu - \lim_\omega \frac{\sum_{p\in X_i} f_i(p)}{|X_i|} \right| \le \frac{1}{2^{k-1}}.$$

Thus our lemma follows. □

Proposition 5.1. *For every measurable function* $\mathbf{f} : \mathbf{X} \to [-d, d]$*, there exists a sequence of functions* $f_i : X_i \to [-d, d]$ *such that the ultraproduct of the sequence* $\{f_i\}_{i=1}^\infty$ *is almost everywhere equal to* $\mathbf{f}$*. That is any element of* $L^\infty(\mathbf{X}, \mathcal{B}_\omega, \mu)$ *can be represented by an ultraproduct function.*

Proof. Recall a standard result of measure theory. If $\mathbf{f}$ is a bounded measurable function on $\mathbf{X}$, then there exists a sequence of bounded stepfunctions $\{h_k\}_{k=1}^\infty$ such that

- $\mathbf{f} = \sum_{k=1}^\infty h_k$
- $|h_k| \leq \frac{1}{2^{k-1}}$, if $k > 1$.
- $h_k = \sum_{n=1}^{n_k} c_n^k \chi_{A_n^k}$, where $\cup_{n=1}^{n_k} A_n^k = \mathbf{X}$ is a measurable partition, $c_n^k \in \mathbb{R}$ if $1 \leq n \leq n_k$.

Now let $B_n^k \in \mathcal{P}$ such that $\mu(A_n^k \triangle B_n^k) = 0$. We can suppose that $\cup_{n=1}^{n_k} B_n^k$ is a partition of $\mathbf{X}$. Let $h'_k = \sum_{n=1}^{n_k} c_n^k \chi_{B_n^k}$ and $\mathbf{f}' = \sum_{k=1}^\infty h'_k$. Then clearly $\mathbf{f}' = \mathbf{f}$ almost everywhere. We show that $\mathbf{f}'$ is an ultraproduct function.

Let $B_n^k = [\{B_{n,i}^k\}_{i=1}^\infty]$. We set $T_k \subset \mathbb{N}$ as the set of integers i for which $\cup_{n=1}^{n_k} B_{n,i}^k$ is a partition of X_i. Then obviously, $T_k \in \omega$. Now we use our diagonalizing trick again. If $i \notin T_1$ let $s_i \equiv 0$. If $i \in T_1, i \in T_2, \dots, i \in T_k, i \notin T_{k+1}$ then define $s_i := \sum_{j=1}^k (\sum_{n=1}^{n_j} c_n^j \chi_{B_{n,i}^j})$. If $i \in T_k$ for each $k \geq 1$ then set $s_i := \sum_{j=1}^i (\sum_{n=1}^{n_j} c_n^i \chi_{B_{n,j}^j})$. Now let $\overline{p} \in B_{j_1}^1 \cap B_{j_2}^2 \cap \dots \cap B_{j_k}^k$. Then

$$|(\lim_\omega s_i)(\overline{p}) - \mathbf{f}'(\overline{p})| \leq \frac{1}{2^{k-1}}.$$

Since this inequality holds for each $k \geq 1$, $\mathbf{f}' \equiv [\{s_i\}_{i=1}^\infty]$. □

Lemma 5.2. *Let $A, B \subseteq [k]$ and let $\mathbf{f} : \mathbf{X}^{[k]} \to \mathbb{R}$ be a $\sigma(B)$-measurable ultraproduct function. Then for all $y \in \mathbf{X}^{A^c}$ the function $\mathbf{f}_y$ is $\sigma(A \cap B)$-measurable, where A^c denotes the complement of A in $[k]$ and $\mathbf{f}_y(x) = \mathbf{f}(x, y)$.*

Proof. Let $\mathbf{f} : \mathbf{X}^{[k]} \to \mathbb{R}$ be a $\sigma(B)$-measurable ultraproduct function. Note that there exist functions $f_i : X_{i,[k]} \to \mathbb{R}$ depending only on the B-coordinates such that $\mathbf{f}$ is the ultraproduct of $\{f_i\}_{i=1}^\infty$. Indeed, let $\mathbf{f}$ be the ultraproduct of the functions g_i. For $x \in X_{i,B}$, let $f_i(x,t) := \frac{\sum_{z \in X_{i,B^c}} g_i(x,z)}{|X_{i,B^c}|}$. Then f_i depends only on the B-coordinates. Also by the integral formula of Lemma 5.1, $\lim_\omega f_i = \mathbf{f}$. Let $y \in \mathbf{X}^{A^c}$, $y = [\{y_i\}_{i=1}^\infty]$. Then $\mathbf{f}_y$ is the ultraproduct of the functions $f_i^{y_i}$. Clearly $f_i^{y_i}$ depends only on the $A \cap B$-coordinates, thus the ultraproduct $\mathbf{f}_y$ is $\sigma(A \cap B)$-measurable. □

Proposition 5.2 (*Fubini's Theorem*). *Let $A \subseteq [k]$ and let $\mathbf{f} : \mathbf{X}^{[k]} \to \mathbb{R}$ be a bounded $\sigma([k])$-measurable function. Then for almost all $y \in \mathbf{X}^{A^c}$, $\mathbf{f}_y(x)$ is a measurable function on $\mathbf{X}^A$ and the function $y \to \int_{\mathbf{X}^A} \mathbf{f}_y(x) d\mu_A(x)$ is $\mathbf{X}^{A^c}$-measurable. Moreover:*

$$\int_{\mathbf{X}^{[k]}} \mathbf{f}(p) d\mu_{[k]}(p) = \int_{\mathbf{X}^{A^c}} \left(\int_{\mathbf{X}^A} \mathbf{f}_y(x) d\mu_A(x) \right) d\mu_{A^c}(y).$$

Proof. First let $\mathbf{f}$ be the ultraproduct of $\{f_i : X_{i,[k]} \to \mathbb{R}\}_{i=1}^\infty$. Define the functions $\overline{f_i} : X_{i,A^c} \to [-d, d]$ by

$$\overline{f_i}(y) = |X_{i,A}|^{-1} \sum_{x \in X_{i,A}} f_i(x, y).$$

By Lemma 5.1

$$\lim_\omega \overline{f_i}(y) = \int_{\mathbf{X}^A} \mathbf{f}(x, y)\, d\mu_A(x).$$

Applying Lemma 5.1 again for the functions $\overline{f_i}$, we obtain that

$$\lim_{\omega} |X_{i,A^c}|^{-1} \sum_{y \in X_{i,A^c}} \overline{f_i}(y) = \int_{\mathbf{X}^{A^c}} \left(\int_{\mathbf{X}^A} \mathbf{f}(x,y) d\mu_A(x) \right) d\mu_{A^c}(y).$$

Then our proposition follows, since

$$|X_{i,A^c}|^{-1} \sum_{y \in X_{i,A^c}} \overline{f_i}(y) = \frac{\sum_{p \in X_i} f_i(p)}{|X_i|}.$$

Now let $\mathbf{f}$ be an arbitrary bounded $\sigma([k])$-measurable function. Since there exists an ultraproduct function $\mathbf{g}$ that is a zero measure perturbation of $\mathbf{f}$ it is enough to prove the following lemma.

Lemma. *Let $Y \subset \mathbf{X}^{[k]}$ be a measurable set of zero measure, then for almost all $y \in \mathbf{X}^{A^c}$,*

$$\{x \in \mathbf{X}^A \mid \mathbf{X}^A \times y \in Y\}$$

has measure zero.

Proof. Since Y is a set of zero measure, there exist sets $Z_n \in \mathcal{P}_{[k]}$ such that

- $\mu_{[k]}(Z_n) \leq \frac{1}{4^n}$
- $Y \subset Z_n$.

Let $L_n \subset \mathbf{X}^{A^c}$ be the set of points y in $\mathbf{X}^{A^c}$ such that

$$\mu_A(\{x \in \mathbf{X}^A \mid \mathbf{X}^A \times y \in Z_n\}) \geq \frac{1}{2^n}.$$

Since Fubini's Theorem holds for ultraproduct functions it is easy to see that $\mu_{A^c}(L_n) \leq \frac{1}{2^n}$. Thus by the Borel–Cantelli Lemma almost all $y \in \mathbf{X}^{A^c}$ is contained only in finitely many sets L_n. Clearly, for those y, $\{x \in \mathbf{X}^A \mid \mathbf{X}^A \times y \in Y\}$ has measure zero. □

Proposition 5.3 (*Integration Rule*)**.** *Let $g_i : \mathbf{X}^{[k]} \to \mathbb{R}$ be bounded $\sigma(A_i)$-measurable functions for $i = 1, 2, \dots, m$. Let B denote the σ-algebra generated by $\sigma(A_1 \cap A_2), \sigma(A_1 \cap A_3), \dots, \sigma(A_1 \cap A_m)$. Then*

$$\int_{\mathbf{X}^{[k]}} g_1 g_2 \dots g_m \, d\mu_{[k]} = \int_{\mathbf{X}^{[k]}} E(g_1|B) g_2 g_3 \dots g_m \, d\mu_{[k]}.$$

Proof. First of all note that $E(g_1 \mid B)$ does not depend on the A_1^c-coordinates. By Fubini's Theorem,

$$\int_{\mathbf{X}^{[k]}} g_1 g_2 g_3 \dots g_m \, d\mu_{[k]} = \int_{\mathbf{X}^{A_1^c}} \left(\int_{\mathbf{X}^{A_1}} g_1(x) g_2(x,y) \dots g_m(x,y) \, d\mu_{A_1}(x) \right) d\mu_{A_1^c}(y).$$

Now we obtain by Lemma 5.2 that for all $y \in \mathbf{X}^{A_1{}^c}$ the function

$$x \to g_2(x,y) g_3(x,y) \dots g_m(x,y) \quad (x \in \mathbf{X}^{A_1})$$

is B-measurable. This means that

$$\int_{\mathbf{X}^{A_1}} g_1(x) g_2(x,y) \dots g_m(x,y) d\mu_{A_1}(x)$$
$$= \int_{\mathbf{X}^{A_1}} E(g_1|B)(x) g_2(x,y) g_3(x,y) \dots g_m(x,y) d\mu_{A_1}(x)$$

for all y in $\mathbf{X}_{A_1^c}$. This completes the proof. □

Now we finish the proof of the Total Independence Theorem. We can assume that $|A_i| \geq |A_j|$ whenever $j > i$. Let χ_i be the characteristic function of S_i. We have that

$$\mu(S_1 \cap S_2 \cap \cdots \cap S_r) = \int_{\mathbf{X}^{[k]}} \chi_1 \chi_2 \ldots \chi_r d\mu_{[k]}.$$

The Integration Rule shows that

$$\int_{\mathbf{X}^{[k]}} \chi_i \chi_{i+1} \ldots \chi_r \, d\mu_{[k]} = \int_{\mathbf{X}^{[k]}} E(\chi_i | \sigma(A_i)^*) \chi_{i+1} \ldots \chi_r \, d\mu_{[k]}$$
$$= \mu(S_i) \int_{\mathbf{X}^{[k]}} \chi_{i+1} \chi_{i+2} \ldots \chi_r \, d\mu_{[k]}.$$

A simple induction finishes the proof. $\square$

6. The proof of the Euclidean correspondence principle

6.1. Random partitions

The goal of this subsection is to prove the following proposition.

Proposition 6.1. *Let $A \subset [k]$ be a subset, then for any $n \geq 1$ there exists a partition $\mathbf{X}^A = S_1 \cup S_2 \cup \cdots \cup S_n$, such that $E(S_i \mid \sigma(A)^*) = \frac{1}{n}$.*

Proof. The idea of the proof is that we consider random partitions of $\mathbf{X}^A$ and show that by probability one these partitions will satisfy the property of our proposition. Let $\Omega = \prod_{i=1}^{\infty} \{1, 2, \ldots, n\}^{X_{i,A}}$ be the set of $\{1, 2, \ldots, n\}$-valued functions on $\cup_{i=1}^{\infty} X_{i,A}$. Each element f of Ω defines a partition of X_A in the following way. Let

$$S_f^{i,j} = \{p \in X_{i,A} \mid f(p) = j\}, \quad 1 \leq j \leq n, \ i \geq 1.$$
$$[\{S_f^{i,j}\}_{i=1}^{\infty}] = S_f^j.$$

Then $\mathbf{X}^A = S_f^1 \cup S_f^2 \cup \cdots \cup S_f^n$ is our partition induced by f.

Note that on Ω one has the usual Bernoulli probability measure P,

$$P(T_{p_1,p_2,\ldots,p_r}(i_1, i_2, \ldots, i_r)) = \frac{1}{n^r},$$

where

$$T_{p_1,p_2,\ldots,p_r}(i_1, i_2, \ldots, i_r) = \{f \in \Omega \mid f(p_s) = i_s \ 1 \leq s \leq r\}.$$

A **cylindric intersection set** T in $X_{i,A}$ is a set $T = \cap_{C, C \subsetneq A} T_C$, where $T_C \subset X_{i,C}$. First of all note that the number of different cylindric intersection sets in $X_{i,A}$ is not greater than

$$\prod_{C, C \subsetneq A} 2^{|X_{i,C}|} \leq 2^{(|X_i|^{|A|-1})2^k}.$$

Let $0 \leq \epsilon \leq \frac{1}{10n}$ be a real number and T be a cylindric intersection set of elements of size at least $\epsilon |X_{i,A}|$. By the Chernoff-inequality the probability that an $f \in \Omega$ takes the value 1 more than $(\frac{1}{n} + \epsilon)|T|$-times or less than $(\frac{1}{n} - \epsilon)|T|$-times on the set T is less than $2\exp(-c_\epsilon |T|)$, where the positive constant c_ϵ depends only on ϵ. Therefore the probability that there exists a cylindric intersection set $T \subset X_{i,A}$ of size at least $\epsilon|X_{i,A}|$ for which $f \in \Omega$ takes the value 1 more than $(\frac{1}{n}+\epsilon)|T|$-times or less than $(\frac{1}{n}-\epsilon)|T|$-times on the set T is less than

$$2^{(|X_i|^{|A|-1})2^k} 2\exp(-c_\epsilon \epsilon |X_i|^{|A|}).$$

Since $|X_1| < |X_2| < \cdots$ by the Borel–Cantelli lemma we have the following lemma.

Lemma 6.1. *For almost all $f \in \Omega$ the following holds. If $\epsilon > 0$, then there exist only finitely many i such that there exists at least one cylindric intersection set $T \subset X_{i,A}$ for which $f \in \Omega$ takes the value 1 more than $(\frac{1}{n}+\epsilon)|T|$-times or less than $(\frac{1}{n}-\epsilon)|T|$-times on the set T.*

Now let us consider a cylindric intersection set $Z \subseteq \mathbf{X}^A$, $Z = \cap_{C, C \subsetneq A} Z_C$, $Z_C \subset \mathbf{X}^C$. By the previous lemma, for almost all $f \in \Omega$,

$$\mu(S_f^1 \cap Z) = \frac{1}{n}\mu(Z).$$

Therefore for almost all $f \in \Omega$:

$$\mu(S_f^1 \cap Z') = \frac{1}{n}(\mu(Z')),$$

where Z' is a finite disjoint union of cylindric intersection sets in $\mathbf{X}^A$. Consequently, for almost all $f \in \Omega$,

$$\mu(S_f^1 \cap Y) = \frac{1}{n}(\mu(Y)),$$

where $Y \in \sigma(A)^*$. This shows immediately that $E(S_f^1 \mid \sigma(A)^*) = \frac{1}{n}$ for almost all $f \in \Omega$. Similarly, $E(S_f^i \mid \sigma(A)^*) = \frac{1}{n}$ for almost all $f \in \Omega$, thus our proposition follows. □

Remark. Later on we need a simple modification of our proposition. Let $\{q_i\}_{i=1}^n$ be non-negative real numbers, such that $\sum_{i=1}^n q_i = 1$. Repeat the construction of the measure on Ω as in Proposition 6.1 with the exception that for any $p \in X_{i,A}$ the probability that $f(p) = i$ is q_i instead of $\frac{1}{n}$. Then with probability one $E(S_f^i \mid \sigma(A)^*) = q_i$.

6.2. *Independent complement in separable σ-algebras*

Let $\mathcal{A}$ be a separable σ-algebra on a set X, and let μ be a probability measure on $\mathcal{A}$. Two sub σ-algebras $\mathcal{B}$ and $\mathcal{C}$ are called independent if $\mu(B \cap C) = \mu(B)\mu(C)$ for every $B \in \mathcal{B}$ and $C \in \mathcal{C}$. We say that $\mathcal{C}$ is an *independent complement* of $\mathcal{B}$ in $\mathcal{A}$ if it is independent from $\mathcal{B}$ and $\langle \mathcal{B}, \mathcal{C} \rangle$ is dense in $\mathcal{A}$.

Definition 6.1. Let $\mathcal{A} \geq \mathcal{B}$ be two σ-algebras on a set X and let μ be a probability measure on $\mathcal{A}$. A $\mathcal{B}$-random k-partition in $\mathcal{A}$ is a partition $A_1, A_2, \ldots, A_k$ of X into $\mathcal{A}$-measurable sets such that $E(A_i|\mathcal{B}) = 1/k$ for every $i = 1, 2, \ldots, k$.

Theorem 15 (*Independent Complement*)**.** *Let $\mathcal{A} \geq \mathcal{B}$ be two separable σ-algebras on a set X and let μ be a probability measure on $\mathcal{A}$. Assume that for every natural number k there exists a $\mathcal{B}$-random k-partition $\{A_{1,k}, A_{2,k}, \ldots, A_{k,k}\}$ in $\mathcal{A}$. Then there is an independent complement $\mathcal{C}$ of $\mathcal{B}$ in $\mathcal{A}$. (Note that this is basically the Maharam-lemma; see* [12]*).*

Proof. Let $S_1, S_2, \dots$ be a countable generating system of $\mathcal{A}$ and let $\mathcal{P}_k$ denote the finite Boolean algebra generated by $S_1, S_2, \dots, S_k$ and $\{A_{i,j} | i \le j \le k\}$. Let $\mathcal{P}_k^*$ denote the atoms of $\mathcal{P}_k$. It is clear that for every atom $R \in \mathcal{P}_k^*$ we have that $E(R|\mathcal{B}) \le 1/k$ because R is contained in one of the sets $A_{1,k}, A_{2,k}, \dots, A_{k,k}$. During the proof we fix one $\mathcal{B}$-measurable version of $E(R|\mathcal{B})$ for every R. The algebra $\mathcal{P}_k$ is a subalgebra of $\mathcal{P}_{k+1}$ for every k. Thus we can define total orderings on the sets $\mathcal{P}_k^*$ such a way that if $R_1, R_2 \in \mathcal{P}_k^*$ with $R_1 < R_2$ and $R_3, R_4 \in \mathcal{P}_{k+1}^*$ with $R_3 \subseteq R_1$, $R_4 \subseteq R_2$ then $R_3 < R_4$. We can assume that $\sum_{R \in \mathcal{P}_k^*} E(R, \mathcal{B})(x) = 1$ for any element $x \in X$. It follows that for $k \in \mathbb{N}$, $x \in X$ and $\lambda \in [0, 1)$ there is a unique element $R(x, \lambda, k) \in \mathcal{P}_k^*$ satisfying

$$\sum_{R < R(x,\lambda,k)} E(R|\mathcal{B})(x) \le \lambda$$

and

$$\sum_{R \le R(x,\lambda,k)} E(R|\mathcal{B})(x) > \lambda.$$

For an element $R \in \mathcal{P}_k^*$ let $T(R, \lambda, k)$ denote the set of those points $x \in X$ for which $R(x, \lambda, k) = R$. It is easy to see that $T(R, \lambda, k)$ is $\mathcal{B}$-measurable. Let us define the $\mathcal{A}$-measurable set $S(\lambda, k)$ by

$$S(\lambda, k) = \bigcup_{R \in \mathcal{P}_k^*} (T(R, \lambda, k) \cap (\cup_{R_2 < R} R_2))$$

and $S'(\lambda, k)$ by

$$S'(\lambda, k) = \bigcup_{R \in \mathcal{P}_k^*} (T(R, \lambda, k) \cap (\cup_{R_2 \le R} R_2)).$$

Note that

$$S(\lambda, k) = \left\{ x \in X \,\middle|\, \sum_{R_2 \le R_k(x)} E(R \mid \mathcal{B})(x) \le \lambda \right\},$$

where $R_k(x)$ is the element of $\mathcal{P}_k^*$ that contains x.

Proposition 6.2. (i) $\lambda - \frac{1}{k} \le E(S(\lambda, k) \mid \mathcal{B})(x) \le \lambda$ *for any* $x \in X$.
(ii) *If* $k < t$, *then* $S(\lambda, k) \subseteq S(\lambda, t) \subseteq S'(\lambda, k)$.
(iii) $E(S'(\lambda, k) \setminus S(\lambda, k) \mid \mathcal{B})(x) \le \frac{1}{k}$ *for any* $x \in X$.

Proof. First observe that

$$\lambda - \frac{1}{k} \le \sum_{R < R(x,\lambda,k)} E(R \mid \mathcal{B})(x) \le \lambda,$$

for any $x \in X$. Also, we have

$$\begin{aligned} S(\lambda, k) &= \bigcup_{R, R_1 \in \mathcal{P}_k^*, R < R_1} (R \cap T(R_1, \lambda, k)), \\ S'(\lambda, k) &= \bigcup_{R, R_1 \in \mathcal{P}_k^*, R \le R_1} (R \cap T(R_1, \lambda, k)). \end{aligned} \tag{6}$$

That is by the basic property of the conditional expectation:

$$E(S(\lambda,k)\mid\mathcal{B})=\sum_{R,R_1\in\mathcal{P}_k^*,R<R_1}E(R\mid\mathcal{B})\chi_{T(R_1,\lambda,k)}.$$

That is

$$E(S(\lambda,k)\mid\mathcal{B})(x)=\sum_{R<R(x,\lambda,k)}E(R\mid\mathcal{B})(x) \tag{7}$$

and similarly

$$E(S'(\lambda,k)\mid\mathcal{B})(x)=\sum_{R\le R(x,\lambda,k)}E(R\mid\mathcal{B})(x). \tag{8}$$

Hence (i) and (iii) follow immediately, using the fact that $E(R'\mid\mathcal{B})\le\frac{1}{k}$ for any $R'\in\mathcal{P}_k^*$.

Observe that for any $R\in\mathcal{P}_k^*$, $T(R,\lambda,k)=\cup_{R'\subseteq R,R'\in\mathcal{P}_t^*}T(R',\lambda,t)$. Hence

$$\begin{aligned}\bigcup_{R,R_1\in\mathcal{P}_k^*,R<R_1}(R\cap T(R_1,\lambda,k))&\subseteq\bigcup_{R',R_1'\in\mathcal{P}_t^*,R'<R_1'}(R'\cap T(R_1',\lambda,t))\\&\subseteq\bigcup_{R,R_1\in\mathcal{P}_k^*,R\le R_1}(R\cap T(R_1,\lambda,k)).\end{aligned}$$

Thus (6) implies (ii). □

Lemma 6.2. *Let $S(\lambda)=\cup_{k=1}^{\infty}S(\lambda,k)$. Then if $\lambda_2<\lambda_1$, then $S(\lambda_2)\subseteq S(\lambda_1)$.*

Proof. Note that $x\in S(\lambda_2,k)$ if and only if $x\in R_2$ for some $R_2<R(x,\lambda_2,k)$. Obviously, $R(x,\lambda_2,k)<R(x,\lambda_1,k)$, thus $x\in S(\lambda_1,k)$. Hence $S(\lambda_2)\subseteq S(\lambda_1)$ □

Lemma 6.3. $E(S(\lambda)\mid\mathcal{B})=\lambda$.

Proof. Since $\chi_{S(\lambda,k)}\overset{L_2(X,\mu)}{\to}\chi_{S(\lambda)}$, we have $E(S(\lambda,k)\mid\mathcal{B})\overset{L_2(X,\mu)}{\to}E(S(\lambda)\mid\mathcal{B})$. That is by (i) of Proposition 6.2 $E(S(\lambda)\mid\mathcal{B})=\lambda$. □

The last two lemmas together imply that the sets $S(\lambda)$ generate a σ-algebra $\mathcal{C}$ which is independent from $\mathcal{B}$.

Now we have to show that $\mathcal{B}$ and $\mathcal{C}$ generate $\mathcal{A}$. Let $S\in\mathcal{P}_k$ for some $k\in\mathbb{N}$. We say that S is an interval if there exists an element $R\in\mathcal{P}_k^*$ such that $S=\cup_{R_1\le R}R_1$. It is enough to show that any interval $S\in\mathcal{P}_k$ can be generated by $\mathcal{B}$ and $\mathcal{C}$.

Suppose that $\{T_t\}_{t=1}^{\infty}$ are sets in $\langle\mathcal{B},\mathcal{C}\rangle$ such that $T_t\subset S$ and $\|E(S\mid\mathcal{B})-E(T_t\mid\mathcal{B})\|$ tends uniformly to 0 as $t\to\infty$. Then $\mu(S\setminus T_t)\to 0$ as $t\to\infty$, that is $\mathcal{B}$ and $\mathcal{C}$ generate S. Indeed,

$$\mu(S\setminus T_t)=\int_X(\chi_S-\chi_{T_t})=\int_X(E(S\mid\mathcal{B})-E(T_t\mid\mathcal{B})).$$

So let $t\ge k$ be an arbitrary natural number. It is clear that S is an interval in $\mathcal{P}_t$. For a natural number $0\le d\le t-1$ let F_d denote the $\mathcal{B}$-measurable set on which $E(S|\mathcal{B})$ is in the interval $(\frac{d}{t},\frac{d+1}{t}]$. Now we approximate S by

$$T_t=\bigcup_{d=0}^{t-1}\left(F_d\cap S\left(\frac{d}{t}\right)\right)\in\langle\mathcal{B},\mathcal{C}\rangle.$$

Lemma 6.4. $T_t \subseteq S$.

Proof. It is enough to prove that $F_d \cap S(\frac{d}{t}, k) \subset S$ for any $0 \le d \le t-1, t < k$. Observe that

$$F_d = \left\{ x \in X \mid \frac{d}{t} < \sum_{R_1 \le R} E(R_1 \mid \mathcal{B})(x) \le \frac{d+1}{t} \right\}$$

and

$$S\left(\frac{d}{t}, k\right) = \left\{ x \in X \mid \sum_{R_2 \le R_k(x)} E(R_2 \mid \mathcal{B})(x) \le \frac{d}{t} \right\}.$$

Thus if $x \in F_d \cap S(\frac{d}{t}, k)$ then $x \in S$. $\square$

Lemma 6.5. *For any* $x \in X$,

$$|E(S \mid \mathcal{B})(x) - E(T_t \mid \mathcal{B})(x)| \le \frac{3}{t}.$$

Proof. First note that by Proposition 6.2 (iii)

$$\left| E\left(S\left(\frac{d}{t}\right) \mid \mathcal{B}\right)(x) - E\left(S\left(\frac{d}{t}, t\right) \mid \mathcal{B}\right)(x) \right| \le \frac{1}{t}. \tag{9}$$

Note that

$$E(T_t \mid \mathcal{B})(x) = \sum_{d=0}^{t-1} \chi_{F_d}(x) E\left(S\left(\frac{d}{t}\right) \mid \mathcal{B}\right)(x).$$

Suppose that $x \in F_d$. Then by (7) and (9),

$$\left| E(T_t \mid \mathcal{B})(x) - \sum_{R' < R(x, \frac{d}{t}, t)} E(R' \mid \mathcal{B})(x) \right| \le \frac{1}{t}.$$

On the other hand $E(S \mid \mathcal{B})(x) = \sum_{R' \le R} E(R' \mid \mathcal{B})(x)$ and $\frac{d}{t} \le \sum_{R' \le R} E(R' \mid \mathcal{B})(x) < \frac{d+1}{t}$. That is

$$|E(S \mid \mathcal{B})(x) - E(T_t \mid \mathcal{B})(x)| \le \frac{3}{t}. \qquad \square$$

The Theorem now follows from Lemma 6.5 immediately. $\square$

Definition 6.2. Let $(X, \mathcal{A}, \mu)$ be a probability space, and assume that a finite group G is acting on X such that $\mathcal{A}$ is G-invariant as a set system. We say that the action of G is free if there is a subset S of X with $\mu(S) = 1/|G|$ such that $S^{g_1} \cap S^{g_2} = \emptyset$ whenever g_1 and g_2 are distinct elements of G.

We will need the following consequence of Theorem 15.

Lemma 6.6. *Let $\mathcal{A} \ge \mathcal{B}$ be two separable σ-algebras on the set X and let μ be a probability measure on $\mathcal{A}$. Assume that a finite group G is acting on X such that $\mathcal{A}$, $\mathcal{B}$ and μ are G invariant. Assume furthermore that the action of G on $(X, \mathcal{B}, \mu)$ is free and for any $k > 1$ there exists a $\mathcal{B}$-random k partition of X in $\mathcal{A}$. Then there is an independent complement $\mathcal{C}$ in $\mathcal{A}$ for $\mathcal{B}$ such that $\mathcal{C}$ is elementwise G-invariant.*

Proof. Let $S \in \mathcal{B}$ be a set showing that G acts freely on $\mathcal{B}$. Let $\mathcal{A}|_S$ and $\mathcal{B}|_S$ denote the restriction of $\mathcal{A}$ and $\mathcal{B}$ to the set S. It is clear that if $\{A_1, A_2, \ldots, A_k\}$ is a $\mathcal{B}$-random k-partition in $\mathcal{A}$ then $\{S \cap A_1, S \cap A_2, \ldots, S \cap A_k\}$ is a $\mathcal{B}|_S$-random k partition in $\mathcal{A}|_S$. Hence by Theorem 15 there exists an independent complement $\mathcal{C}_1$ of $\mathcal{B}|_S$ in $\mathcal{A}|_S$. The set

$$\mathcal{C} = \left\{ \bigcup_{g \in G} H^g | H \in \mathcal{C}_1 \right\}$$

is a σ-algebra because the action of G is free. Note that the elements of $\mathcal{C}$ are G-invariant. Since $E(\cup_{g\in G} H^g|\mathcal{B}) = \sum_{g\in G} E(H|\,\mathcal{B}|_S)^g$ we obtain that the elements of $\mathcal{C}$ are independent from $\mathcal{B}$. It is clear that $\langle \mathcal{C}, \mathcal{B} \rangle$ is dense in $\mathcal{A}$. □

6.3. Separable realization

In this subsection we show how to pass from nonseparable σ-algebras to separable ones. First note that the symmetric group S_k acts on the space $\mathbf{X}^k$ by permuting the coordinates:

$$(x_1, x_2, \ldots, x_k)^\pi = (x_{\pi^{-1}(1)}, x_{\pi^{-1}(2)}, \ldots, x_{\pi^{-1}(k)}).$$

The group also acts on the subsets of $[k]$ and $\sigma(A)^\pi = \sigma(A^\pi)$, where A^π denotes the image of the subset A under $\pi \in S_k$. We will denote by S_A the symmetric group acting on the subset A.

Definition 6.3. A **separable system** on $\mathbf{X}^k$, $r \le k$ is a system of atomless separable σ-algebras $\{l(A) \mid A \in r([k])\}$ and functions $\{F_A : \mathbf{X}^k \to [0, 1] \mid A \in r([k])\}$ with the following properties.

1. $l(A)$ is a subset of $\sigma(A)$ and is independent from $\sigma(A)^*$ for every $\emptyset \neq A \subseteq [k]$.
2. $l(A)^\pi = l(A^\pi)$ for every permutation $\pi \in S_k$.
3. $S^\pi = S$ for every $S \in l(A)$ and $\pi \in S_A$.
4. F_A is an $l(A)$-measurable function which defines a measurable equivalence between the measure algebras of $(\mathbf{X}^k, l(A), \mu^k)$ and $[0, 1]$. (See Appendix.)
5. $F_A(\mathbf{x}) = F_{A^\pi}(\mathbf{x}^\pi)$ for every element $\mathbf{x} \in \mathbf{X}^k$, $\pi \in S_k$ and $A \subseteq [k]$.

The main proposition in this section is the following one.

Proposition 6.3. *For every separable σ-algebra $\mathcal{A}$ in $\sigma([k])$ there exists a separable system such that for every set $M \in \mathcal{A}$ there is a set $Q \in \langle l(A) \mid A \in r([k])\rangle$ with $\mu_{[k]}(M \triangle Q) = 0$.*

This proposition immediately implies Theorem 6 since the map $F : \mathbf{X}^k \mapsto [0, 1]^{r([k])}$ whose coordinate functions are $\{F_A \mid A \in r([k])\}$ constructed in Proposition 6.3 is a separable realization.

We will need the following three lemmas.

Lemma 6.7. *Let $\mathcal{B} \subseteq \mathcal{A}$ be two σ-algebras on a set Y, and let μ be a probability measure on $\mathcal{A}$. Then for any separable sub-σ-algebra $\bar{\mathcal{A}}$ of $\mathcal{A}$ there exists a separable sub σ-algebra $\bar{\mathcal{B}}$ of $\mathcal{B}$ such that $E(A|\mathcal{B}) = E(A|\bar{\mathcal{B}})$ for every $A \in \bar{\mathcal{A}}$.*

Proof. We use the fact that $\bar{\mathcal{A}}$ is a separable metric space with the distance $d(A, B) = \mu(A \triangle B)$. Let $W = \{D_1, D_2, \ldots\}$ be a countable dense subset of $\bar{\mathcal{A}}$ with the previous distance. Let $C^i_{p,q} = E(D_i \mid \mathcal{B})^{-1}(p, q)$, where $p < q$ are rational numbers. Clearly, $E(D_i \mid \mathcal{B})$ is a $\mathcal{B}_i$-measurable function, where $\mathcal{B}_i = \langle C^i_{p,q} \mid p < q \in \mathbb{Q}\rangle$. Obviously, $E(D_i \mid \overline{\mathcal{B}}) = E(D_i \mid \mathcal{B})$

for any $i \geq 1$, where $\overline{\mathcal{B}} = \langle \mathcal{B}_i \mid i = 1, 2, \ldots \rangle$. Now observe that $E(D_i \mid \mathcal{B}) \xrightarrow{L_2} E(D, \mathcal{B})$ and $E(D_i \mid \overline{\mathcal{B}}) \xrightarrow{L_2} E(D, \overline{\mathcal{B}})$ if $D_i \to D$. Hence for any $D \in \overline{\mathcal{A}}$, $E(D \mid \overline{\mathcal{B}}) = E(D \mid \mathcal{B})$. □

Lemma 6.8. *Let $A \subseteq [k]$ be a subset and assume that there are atomless separable σ-algebras $d(\{i\}) \subset \sigma(\{i\})$, $i \in A$ such that $d(\{i\})^\pi = d(\{i^\pi\})$ for every $i \in A$ and $\pi \in S_A$. Then S_A acts freely on $\langle d(\{i\}) | i \in A \rangle$.*

Proof. The permutation invariance implies that there exists a σ-algebra $\mathcal{A}$ on $\mathbf{X}$ such that $P^{-1}_{\{i\}}(\mathcal{A}) = d(\{i\})$ for every $i \in A$. Let $F : \mathbf{X} \to [0, 1]$ be a $\mathcal{A}$-measurable measure preserving map. Now we can define the map $G : \mathbf{X}^A \to [0, 1]^A$ by

$$G(x_{i_1}, x_{i_2}, \ldots, x_{i_{|A|}}) := (F(x_{i_1}), F(x_{i_2}), \ldots, F(x_{i_{|A|}})).$$

Let us introduce $S' := \{(y_1, y_2, \ldots, y_r) | y_1 < y_2 < \ldots < y_r\} \subset [0, 1]^A$ and $S := G^{-1}(S')$. Clearly $\mu^A(S) = 1/|A|!$ and $S^\pi \cap S^\rho = \emptyset$ for every two different elements $\pi \neq \rho$ in S_A. □

Lemma 6.9. *Let k be a natural number and assume that for every $A \subseteq [k]$ there is a separable σ-algebra $c(A)$ in $\sigma(A)$. Then for every $A \subseteq [k]$ there is a separable σ-algebra $d(A)$ in $\sigma(A)$ with $c(A) \subseteq d(A)$ such that*

1. $E(R | \langle d(B) | B \in A^* \rangle) = E(R | \sigma(A)^*)$ *whenever* $R \in d(A)$.
2. $d(A)^\pi = d(A^\pi)$ *for every element* $\pi \in S_k$.
3. $d(B) \subseteq d(A)$ *whenever* $B \subseteq A$

Proof. First we construct algebras $d'(A)$ recursively. Let $d'([k])$ be $\langle c([k])^\pi | \pi \in S_k \rangle$. Assume that we have already constructed the algebras $d'(A)$ for $|A| \geq t$. Let $A \subseteq [k]$ be such that $|A| = t$. By Lemma 6.7 we can see that there exists a separable subalgebra $\widetilde{d'(A)}$ of $\sigma(A)^*$ such that $E(R | \sigma(A)^*) = E(R | \widetilde{d'(A)})$ for every $R \in d'(A)$. Since $\sigma(A)^*$ is generated by the algebras $\{\sigma(B) | B \in A^*\}$ we have that every element of $\sigma(A)^*$ is a countable expression of some sets in these algebras. This implies that any separable sub σ-algebra of $\sigma(A)^*$ is generated by separable sub σ-algebras of the algebras $\sigma(B)$ where $B \in A^*$. In particular we can choose separable σ-algebras $d'(A, B) \supset c(B)$ in $\sigma(B)$ for every $B \in A^*$ such that $\langle d'(A, B) | B \in A^* \rangle \supseteq \widetilde{d'(A)}$. For a set $B \subseteq [k]$ with $|B| = t - 1$ we define $d'(B)$ as the σ-algebra generated by all the algebras in the form of $d'(C, D)^\pi$, where $\pi \in S_k$, $D^\pi = B$, $|C| = |D| + 1$ and $D \subseteq C$. Since $d'(C, D)^\pi \subseteq \sigma(D)^\pi = \sigma(B)$ we have that $d'(B) \subseteq \sigma(B)$. Furthermore we have that $d'(B)^\pi = d'(B^\pi)$ for every $\pi \in S_k$.

Now let $d(A) := \langle d'(B) \mid B \subseteq A \rangle$. The second requirement in the lemma is trivial by definition. We prove the first one. The elements of $d(A)$ can be approximated by finite unions of intersections of the form $\bigcap_{B \subseteq A} T_B$ where $T_B \in d'(B)$ and so it is enough to prove the statement if R is such an intersection. Let $Q = \bigcap_{B \subset A, B \neq A} T_B$. Now

$$E(R | \langle d(B) | B \in A^* \rangle) = E(R | \langle d'(B) | B \subset A, B \neq A \rangle).$$

By the basic property of the conditional expectation (see Appendix) :

$$\begin{aligned} E(R | \langle d'(B) | B \subset A, B \neq A \rangle) &= E(T_A | \langle d'(B) | B \subset A, B \neq A \rangle) \chi_Q = E(T_A | \sigma(A)^*) \chi_Q \\ &= E(R | \sigma(A)^*). \quad \square \end{aligned}$$

Proof of Proposition 6.3. We construct the algebras $l(A)$ in the following steps. For each non-empty subset $A \subseteq [k]$ we choose an atomless separable σ-algebra $c(A) \subseteq \sigma(A)$ containing a $\sigma(A)^*$-random r-partition for every r. We also assume that $\mathcal{A} \subseteq c([k])$. Applying Lemma 6.9 for the previous system of separable σ-algebras $c(A)$ we obtain the σ-algebras $d(A)$. By Lemma 6.8 and the permutation invariance property of the previous lemma, $S_{[r]}$ acts freely on $d([r])^* = \langle d(B) | B \in [r]^* \rangle$. Hence using Lemma 6.6, for every $\emptyset \neq A \in [k]$ we can choose an independent complement $l([r])$ for $d([r])^*$ in $d([r])$ such that $l([r])$ is elementwise invariant under the action of $S_{[r]}$. The algebras $l([r])$ are independent from $\sigma([r])^*$ since $\mu(R) = E(R|d([r])^*) = E(R|\sigma([r])^*)$ for every $R \in l([r])$. Now we define $l(A)$, where $|A| = r$ by $l(A) = l([r])^\pi$ for some $\pi \in S_k$, $\pi([r]) = A$. Note that $l(A)$ does not depend on the choice of π. By Lemma A.1 of Appendix we have maps $F_{[r]} : \mathbf{X}^r \to [0, 1]$ such that F^{-1} defines a measure algebra isomorphism between $\mathcal{M}([0, 1], \mathcal{B}, \lambda)$ and $\mathcal{M}(\mathbf{X}^r, l[r], \mu^r)$. Let $F_A = \pi^{-1} \circ F_{[r]}$, where π maps $[r]$ to A. Again, $F_{[r]}$ does not depend on the particular choice of the permutation π. $\square$

Acknowledgment

We are very indebted to Terence Tao and László Lovász for helpful discussions.

Appendix. Basic measure theory

In this section we collect some of the basic results of measure theory we frequently use in our paper.

Separable measure spaces. Let $(X, \mathcal{A}, \mu)$ be a probability measure space. Then we call $A, A' \in \mathcal{A}$ equivalent if $\mu(A \triangle A') = 0$. The equivalence classes form a complete metric space, where $d([A], [B]) = \mu(A \triangle B)$. These classes form a Boolean-algebra as well, called the **measure algebra** $\mathcal{M}(X, \mathcal{A}, \mu)$. We say that $(X, \mathcal{A}, \mu)$ is a **separable** measure space if $\mathcal{M}(X, \mathcal{A}, \mu)$ is a separable metric space. It is important to note that if $(X, \mathcal{A}, \mu)$ is separable and atomless, then its measure algebra is isomorphic to the measure algebra of the standard Lebesgue space $([0, 1], \mathcal{B}, \lambda)$, where $\mathcal{B}$ is the σ-algebra of Borel sets (see e.g. [6]). We use the following folklore version of this theorem.

Lemma A.1. *If $(X, \mathcal{A}, \mu)$ is a separable and atomless measure space, then there exists a map $f : X \to [0, 1]$ such that $f^{-1}(\mathcal{B}) \subset \mathcal{A}$, $\mu(f^{-1}(U)) = \lambda(U)$ for any $U \in \mathcal{B}$ and for any $L \in \mathcal{A}$ there exists $M \in \mathcal{B}$ such that L is equivalent to $f^{-1}(M)$.*

In other words, if $F : [0, 1] \to X$ is an injective measure preserving measure algebra homomorphism such that the image of the Borel-algebra is just $\mathcal{A}$, then F can be ***represented*** *by the map f. That is for any measurable set $U \subset [0, 1]$, $F(U)$ is the set representing $f^{-1}(U)$.*

Proof. Let I_0 denote the interval $[0, \frac{1}{2}]$, $I_1 = [\frac{1}{2}, 1]$. Then let $I_{0,0} = [0, \frac{1}{4}]$, $I_{0,1} = [\frac{1}{4}, \frac{1}{2}]$, $I_{1,0} = [\frac{1}{2}, \frac{3}{4}]$, $I_{1,1} = [\frac{3}{4}, 1]$. Recursively, we define the dyadic intervals $I_{\alpha_1,\alpha_2,\ldots,\alpha_k}$, where $(\alpha_1, \alpha_2, \ldots, \alpha_k)$ is a $0-1$-string. Let T be the Boolean-algebra isomorphism between the measure algebra of $(X, \mathcal{A}, \mu)$ and the measure algebra of $([0, 1], \mathcal{B}, \lambda)$. Then we have disjoint sets $U_0, U_1 \in \mathcal{A}$ such that $T([U_0]) = [I_0]$, $T([U_1]) = [I_1]$. Clearly $\mu(X \setminus (U_0 \cup U_1)) = 0$. Similarly, we have disjoint subsets of U_0, $U_{0,0}$ and $U_{0,1}$ such that $T([U_{0,0}]) = [I_{0,0}]$ and $T([U_{0,1}]) = [I_{0,1}]$. Recursively, we define $U_{\alpha_1,\alpha_2,\ldots,\alpha_k} \in \mathcal{A}$ such that $U_{\alpha_1,\alpha_2,\ldots,\alpha_{k-1},0}$ and

$U_{\alpha_1,\alpha_2,\ldots,\alpha_{k-1},0}$ are disjoint and $T([U_{\alpha_1,\alpha_2,\ldots,\alpha_k}]) = I_{\alpha_1,\alpha_2,\ldots,\alpha_k}$. For any $k > 0$, the set of points in X which are not included in some $U_{\alpha_1,\alpha_2,\ldots,\alpha_k}$ has measure zero. Now define

$$f(p) := \cap_{k=1}^{\infty} I_{\alpha_1,\alpha_2,\ldots,\alpha_k},$$

where for each $k \geq 1$, $p \in U_{\alpha_1,\alpha_2,\ldots,\alpha_k}$. It is easy to see that f satisfies the conditions of our lemma. □

Generated σ-algebras. Let $(X, \mathcal{C}, \mu)$ be a probability measure space and $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_k$ be sub-σ-algebras. Then we denote by $\langle \mathcal{A}_i \mid 1 \leq i \leq k\rangle$ the generated σ-algebra that is the smallest sub-σ-algebra of $\mathcal{C}$ containing the $\mathcal{A}_i$'s. Then the equivalence classes

$$[\cup_{j=1}^{n}(A_1^j \cap A_2^j \cap \ldots \cap A_k^j)],$$

where $A_i^j \in \mathcal{A}_i$ and $(A_1^s \cap A_2^s \cap \ldots \cap A_k^s) \cap (A_1^t \cap A_2^t \cap \ldots \cap A_k^t) = \emptyset$ if $s \neq t$ form a dense subset in the measure algebra $\mathcal{M}(X, \langle \mathcal{A}_i \mid 1 \leq i \leq k\rangle, \mu)$ with respect to the metric defined above (see [6]).

Independent subalgebras and product measures. The sub-σ-algebras $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_k \subset \mathcal{C}$ are **independent** subalgebras if

$$\mu(A_1)\mu(A_2)\ldots\mu(A_k) = \mu(A_1 \cap A_2 \cap \cdots \cap A_k),$$

if $A_i \in \mathcal{A}_i$.

Lemma A.2. *Let $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_k \subset \mathcal{C}$ be independent subalgebras as above and $f_i : X \to [0, 1]$ be maps such that f_i^{-1} defines isomorphisms between the measure algebras $\mathcal{M}(X, \mathcal{A}_i, \mu)$ and $\mathcal{M}([0, 1], \mathcal{B}, \lambda)$. Then the map F^{-1}, $F = \oplus_{i=1}^k f_i : X \to [0, 1]^k$ defines an isomorphism between the measure algebras $\mathcal{M}(X, \langle \mathcal{A}_i \mid 1 \leq i \leq k\rangle, \mu)$ and $\mathcal{M}([0, 1]^k, \mathcal{B}^k, \lambda^k)$.*

Proof. Observed that

$$\mu(F^{-1}(\cup_{i=1}^{s}[A_1^i \times \ldots \times A_k^i])) = \sum_{i=1}^{s} \lambda^k[A_1^i \times \ldots \times A_k^i]$$

whenever $\{A_1^i \times \ldots \times A_k^i\}_{i=1}^s$ are disjoint product sets. Hence F^{-1} defines an isometry between dense subsets of the two measure algebras. □

Radon–Nikodym Theorem. Let $(X, \mathcal{A}, \mu)$ be a probability measure space and ν be an absolutely continuous measure with respect to μ. That is if $\mu(A) = 0$ then $\nu(A) = 0$ as well. Then there exists an integrable $\mathcal{A}$-measurable function f such that

$$\mu(A) = \int_A f d\mu$$

for any $A \in \mathcal{A}$.

Conditional expectation. Let $(X, \mathcal{A}, \mu)$ be a probability measure space and $\mathcal{B} \subset \mathcal{A}$ be a sub-σ-algebra. Then by the Radon–Nikodym-theorem for any integrable $\mathcal{A}$-measurable function f there exists an integrable $\mathcal{B}$-measurable function $E(f \mid \mathcal{B})$ such that

$$\int_B E(f \mid \mathcal{B}) d\mu = \int_B f d\mu,$$

if $B \in \mathcal{B}$. The function $E(f \mid \mathcal{B})$ is called the conditional expectation of f with respect to $\mathcal{B}$. It is unique up to a zero-measure perturbation. Note that if $a \leq f(x) \leq b$ for almost all $x \in X$, then $a \leq E(f \mid \mathcal{B})(x) \leq b$ for almost all $x \in X$ as well. Also, if g is a bounded $\mathcal{B}$-measurable function, then

$$E(fg \mid \mathcal{B}) = E(f \mid \mathcal{B})g \quad \text{almost everywhere.}$$

The map $f \to E(f, \mathcal{B})$ extends to a Hilbert-space projection $E : L^2(X, \mathcal{A}, \mu) \to L^2(X, \mathcal{B}, \mu)$.

Lebesgue density theorem. Let $A \in \mathbb{R}^n$ be a measurable set. Then almost all points $x \in A$ is a **density point**. The point x is a density point if

$$\lim_{r\to 0} \frac{Vol(B_r(x) \cap A)}{Vol(B_r(x))} = 1,$$

where Vol denotes the n-dimensional Lebesgue-measure.

Coupling. Let A, B be sets. Let X be an A-valued random variable and Y be a B-valued random variable. A **coupling** of X and Y is a $A \times B$-valued random variable Z, such that the first component of Z has the distribution of X and the second component of Z has the distribution of Y.

References

[1] T. Austin, On exchangeable random variables and the statistics of large graphs and hypergraphs, Probab. Surv. 5 (2008) 80–145. electronic.
[2] T. Austin, T. Tao, On the testability and repair of hereditary hypergraph properties, Random Structures Algorithms 36 (4) (2010) 373463.
[3] C. Borgs, J. Chayes, L. Lovasz, V.T. Sós, B. Szegedy, K. Vesztergombi, Graph limits and parameter testing, in: STOC'06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing, ACM, New York, 2006, pp. 261–270.
[4] C.C. Chang, H.J. Keisler, Model theory., in: Studies in Logic and the Foundations of Mathematics, 73, North-Holland Publishing Co, Amsterdam, 1990.
[5] T. Gowers, Quasirandomness, counting and regularity for 3-uniform hypergraphs, Combin. Probab. Comput. 15 (1-2) (2006) 143–184.
[6] P.R. Halmos, Measure Theory, Van Nostrand Company, Inc., New York, N. Y, 1950.
[7] Y. Ishigami, A Simple Regularization of Hypergraphs, (unpublished) http://arxiv.org/abs/math/0612838.
[8] O. Kallenberg, Symmetries on random arrays and set-indexed processes, J. Theoret. Probab. 5 (4) (1992) 727–765.
[9] P.E. Loeb, Conversion from nonstandard to standard measure spaces and applications in probability theory, Trans. Amer. Math. Soc. 211 (1975) 113–122.
[10] L. Lovasz, B. Szegedy, Limits of dense graph sequences, J. Combin. Theory Ser. B 96 (6) (2006) 933–957.
[11] C. Borgs, J. Chayes, L. Lovász, Moments of two-variable functions and the uniqueness of graph limits, Geom. Funct. Anal. 19 (6) (2010) 1597–1619.
[12] D. Maharam, On homogeneous measure algebras, Proc. Nat. Acad. Sci. U. S. A. 28 (1942) 108–111.
[13] B. Nagle, V. Rödl, M. Schacht, The counting lemma for regular k-uniform hypergraphs, Random Structures Algorithms 28 (2) (2006) 113–179.
[14] V. Rödl, M. Schacht, Regular partitions of hypergraphs: regularity lemmas, Combin. Probab. Comput. 16 (6) (2007) 833–885.
[15] V. Rödl, J. Skokan, Regularity lemma for k-uniform hypergraphs, Random Structures Algorithms 25 (1) (2004) 1–42.
[16] V. Rödl, M. Schacht, Generalizations of the removal lemma, Combinatorica 29 (4) (2009) 467–501.
[17] J. Solymosi, A note on a question of Erdös and Graham, Combin. Probab. Comput. 13 (2) (2004) 263–267.
[18] T. Tao, A variant of the hypergraph removal lemma, J. Combin. Theory Ser. A 113 (7) (2006) 1257–1280.