

AN APPROXIMATE LOGIC FOR MEASURES

ISAAC GOLDBRING AND HENRY TOWNSNER

ABSTRACT. We present a logical framework for formalizing connections between finitary combinatorics and measure theory or ergodic theory that have appeared in various places throughout the literature. We develop the basic syntax and semantics of this logic and give applications, showing that the method can express the classic Furstenberg correspondence as well as provide short proofs of the Szemerédi Regularity Lemma and the hypergraph removal lemma. We also derive some connections between the model-theoretic notion of stability and the Gowers uniformity norms from combinatorics.

1. INTRODUCTION

Since the 1970's, diagonalization arguments (and their generalization to ultrafilters) have been used to connect results in finitary combinatorics with results in measure theory or ergodic theory [14]. Although it has long been known that this connection can be described with first-order logic, a recent striking paper by Hrushovski [22] demonstrated that modern developments in model theory can be used to give new substantive results in this area as well.

Papers on various aspects of this interaction have been published from a variety of perspectives, with almost as wide a variety of terminology [30, 10, 9, 1, 32, 13, 4]. Our goal in this paper is to present an assortment of these techniques in a common framework. We hope this will make the entire area more accessible.

We are typically concerned with the situation where we wish to prove a statement 1) about sufficiently large finite structures which 2) concerns the finite analogs of measure-theoretic notions such as density or integrals. The classic example of such a statement is Szemerédi's Theorem (other examples can be found later in the paper and in the references):

Theorem 1.1 (Szemerédi's Theorem). *For each $\epsilon > 0$ and each $k \in \mathbb{N}^{>0}$, there is an N such that whenever $n > N$ and $A \subseteq [1, n] := \{1, \dots, n\}$ with $\frac{|A|}{n} \geq \epsilon$, there exist a, d such that $a, a + d, a + 2d, \dots, a + (k - 1)d \in A$.*

Suppose this statement were false, and therefore that for some particular ϵ , some k , and infinitely many n , there were sets $A_n \subseteq [1, n]$ serving as counterexamples to the theorem. For each n , we could view the set $[1, n]$, together with the subset A_n and the function $+ \pmod n$, as a *structure* \mathfrak{M}_n of first-order logic. Using the ultraproduct construction, these finite models can be assembled into a single *infinite* model \mathfrak{M} satisfying the Łoś Theorem, which, phrased informally, states:

Theorem 1.2. *A first-order sentence is true in \mathfrak{M} if and only the sentence is true in \mathfrak{M}_n for “almost every” n .*

2010 *Mathematics Subject Classification.* Primary 03B48; Secondary 03C20, 60A10.

Key words and phrases. first-order logic, ultraproducts, measure theory.

Goldbring's work was partially supported by NSF grant DMS-1007144 and Townsner's work was partially supported by NSF grant DMS-1001528.

The existence of a and d satisfying the theorem can be expressed by a first-order sentence, so if we could show that this sentence was true in the infinite structure \mathfrak{M} , then the same sentence would be true in many of the finite structures \mathfrak{M}_n , which would contradict our assumption.

In each of the finite models, we have a natural measure on $[1, n]$, namely the (normalized) counting measure μ_n , and the assumption that $\frac{|A_n|}{n} \geq \epsilon$ is simply the statement that $\mu_n(A_n) \geq \epsilon$. We would like to produce a measure on \mathfrak{M} which satisfies some analog of the Loś theorem.

One way to accomplish this is to expand the language of our models with many new predicates for measures, one for each definable set S and rational δ , and specify that the formula $m_{S,\delta}$ holds in \mathfrak{M}_n exactly if $\mu_n(S) < \delta$. Then, for instance, the failure of $m_{A,\epsilon}$ in the finite structures \mathfrak{M}_n implies the failure of the same formula in \mathfrak{M} . With a bit of work, these new predicates are enough to ensure that we can construct a genuine measure, the Loeb measure (see [15]), on the structure \mathfrak{M} whose properties are related to those of the finite structures via the Loś theorem. The extension of a language of first-order logic by these predicates has been used in [22], but the details of the construction have not appeared in print.

The applications of these predicates for measures are sufficiently compelling to justify the study of this extension of first-order logic in its own right. It would be quite interesting to see certain results about first-order logic proven for this extension; for example, some of the results in [7] can be viewed as studying the addition of a “random” predicate P to a language of first-order logic, a notion which has natural application in the presence of measures if the result can be extended to include AML.

In this paper we define *approximate measure logic*, or *AML*, to be an extension of first-order logic in which these new predicates are always present. This serves two purposes. First, to give the full details of extending first-order logic by measure predicates in a fairly general way. Second, we hope that providing a precise formulation will spur the investigation of the distinctive properties of AML.

In general, a logic which intends to talk about measure has to make some kind of compromise, since there is a tension between talking about measures in a natural way and talking about the usual first-order notions (which, for instance, impose the presence of projections of sets). AML is not the first attempt at such a combination (see [24, 2]), but differs because it stays much closer to conventional first-order logic. The main oddity, relative to what a naive attempt at a logic for measures might look like, is that the description of measures in the logic is *approximate*: there can be a slight mismatch between which sentences are true in a structure and what is true about the actual measure. For instance, it is possible to have a structure \mathfrak{M} together with a measure μ so that \mathfrak{M} satisfies a sentence saying the measure of a set B is strictly less than ϵ while in fact the measure of B is precisely ϵ . This is necessary to accommodate the ultraproduct construction.

The next three sections lay out the basic definitions of AML and describe some properties of structures of AML, culminating in the construction of ultraproducts which have well-behaved measures. In Section 5, we present two examples of simple applications, giving a proof of the Furstenberg correspondence between finite sets and dynamical systems, and a proof of the Szemerédi Regularity Lemma.

In Section 6 we prove a Downward Löwenheim-Skolem theorem for AML. In Section 7, we illustrate several useful model-theoretic techniques in AML and

apply these to show a connection between definability, the model-theoretic notion of stability, and combinatorial quantities known as the Gowers uniformity norms. We intend this section to be a survey of methods that have proven useful. These methods are either more fully developed in the references, or await a fuller development.

In this article, m and n range over $\mathbb{N} := \{0, 1, 2, \dots\}$. For a relation $R \subseteq X \times Y$, $x \in X$, and $y \in Y$, we set $R^x := \{y \in Y : (x, y) \in R\}$ and $R_y := \{x \in X : (x, y) \in R\}$.

The authors thank Ehud Hrushovski, Lou van den Dries, and the members of the UCLA reading seminar on approximate groups: Matthias Aschenbrenner, Greg Hjorth, Terence Tao, and Anush Tserunyan, for helpful conversations on the topic of this paper.

2. SYNTAX AND SEMANTICS

A *signature for approximate measure logic* (AML) is nothing more than a first-order signature. Suppose that \mathcal{L} is a first-order signature. One forms AML \mathcal{L} -terms as usual, by applying function symbols to constant symbols and variables. We let \mathcal{T} denote the set of \mathcal{L} -terms.

Definition 2.1. We define the \mathcal{L} -formulae of AML as follows:

- Any first-order atomic \mathcal{L} -formula is an AML \mathcal{L} -formula.
- The AML \mathcal{L} -formulae are closed under the use of the usual connectives \neg and \wedge .
- If φ is an AML \mathcal{L} -formula and x is a variable, then $\exists x\varphi$ and $\forall x\varphi$ are AML \mathcal{L} -formulae.
- If $\vec{x} = (x_1, \dots, x_n)$ is a sequence of distinct variables, q is a non-negative rational, $\bowtie \in \{<, \leq\}$, and φ is an \mathcal{L} -formula, then $m_{\vec{x}} \bowtie q.\varphi$ is an \mathcal{L} -formula.

The last constructor, which we will refer to as the *measure constructor*, in the above definition behaves like a quantifier in the sense that any of the variables x_1, \dots, x_n appearing free in φ become bound in the resulting formula.

From now on, \bowtie will always denote $<$ or \leq .

Convention 2.2. It will become convenient to introduce the following notations:

- $m_{\vec{x}} \geq q.\varphi$ is an abbreviation for $\neg m_{\vec{x}} < q.\varphi$
- $m_{\vec{x}} > q.\varphi$ is an abbreviation for $\neg m_{\vec{x}} \leq q.\varphi$

Remark 2.3. It is not necessary that AML contain precisely one measure for each tuple length. A more general definition could allow multiple measures at each arity. In this case, one would also want to track which measures on pairs, for instance, correspond to the products of which measures on singletons, and so on. There are even situations where it is natural to have a measure on pairs which does not correspond to any product of measures on singletons.¹ Further generalizations could involve such things as measures valued in ordered rings. A fully general definition of AML would include in the signature \mathcal{L} , for each arity, a list of measures together with the ordered ring in which they are valued. In this

¹A typical example is the proofs in extremal graph theory of sparse analogs of density statements, as in the regularity lemma for sparse graphs [26], where the counting measure on pairs is replaced by counting the proportion of pairs from some fixed (typically small) set of pairs. If the fixed set of pairs satisfies certain randomness properties, this measure cannot be viewed as any product measure.

paper we eschew these complications and focus only on the simplest case, with a single real-valued measure; the changes needed to handle more general cases are largely routine.

Particular examples of AML structures are given by finite structures with the normalized counting measure and by structures with definable Keisler measures. However we would like to define an *a priori* notion of semantics for AML: analogously to the definition of a structure of first-order logic, we would like to first state that a certain object is an AML structure and then interpret formulae and sentences within that object. The heart of such a definition should be that an AML structure consists of a first-order structure together with a measure such that all definable sets are measurable. Since identifying the definable sets requires interpreting the logic, this makes it difficult to recognize a valid AML structure without simultaneously interpreting the logic.

We adopt a compromise: an AML *quasistructure* is an object which would be an AML structure except that definable sets may not be measurable. We define the rank of a formula to be the level of nesting of measure quantifiers, and we then, by simultaneous recursion, identify the definable sets of rank n and the quasistructures in which all sets defined by formulae of rank n are measurable; as long as all sets defined by formulae of rank n are measurable, it will be possible to interpret formulae of rank $n+1$. An *AML structure* will then be a quasistructure in which this recursion continues through all ranks.

We face one additional obstacle: a first-order structure together with a measure does not provide enough information to completely specify the semantics. In order to satisfy compactness, we need to allow the possibility that

$$\mathfrak{M} \models m_{\vec{x}} < r + \epsilon \cdot \varphi(\vec{x})$$

for all $\epsilon > 0$, but

$$\mathfrak{M} \not\models m_{\vec{x}} \leq r \cdot \varphi(\vec{x}).$$

In such a structure, the measure of the set defined by φ should be r , but this structure must have a different semantics from a structure \mathfrak{N} where the measure of the set defined by φ is r and

$$\mathfrak{N} \models m_{\vec{x}} \leq r \cdot \varphi(\vec{x}).$$

To accomodate this, we pair the measure portion of a structure with a function v which assigns to each set X a value \odot , \oplus , or \ominus . The values \oplus and \ominus indicate that the structure can only identify the measure of the set X as it is approximated from above or from below, respectively, while the value \odot indicates that the structure actually identifies the measure of the set X exactly. This issue can only occur when the value of the measure is rational, so if the measure of X is irrational, we always set $v(X) = \odot$.

Definition 2.4. Let \mathcal{L} be a first-order signature. An AML \mathcal{L} -*quasistructure* consists of:

- a first-order \mathcal{L} -structure \mathfrak{M} ;
- for each $n \geq 1$, an algebra $\mathcal{B}_n^{\mathfrak{M}}$ of subsets of M^n and a finitely additive measure $\mu_n^{\mathfrak{M}}$ on $\mathcal{B}_n^{\mathfrak{M}}$ such that, for any m, n , $\mathcal{B}_{m+n}^{\mathfrak{M}}$ extends $\mathcal{B}_m^{\mathfrak{M}} \otimes \mathcal{B}_n^{\mathfrak{M}}$ and $\mu_{m+n}^{\mathfrak{M}}$ extends $\mu_m^{\mathfrak{M}} \times \mu_n^{\mathfrak{M}}$;
- for each set $X \in \mathcal{B}_n^{\mathfrak{M}}$, an assignment of a value $v_n^{\mathfrak{M}}(X) \in \{\oplus, \ominus, \odot\}$ such that $v_n^{\mathfrak{M}}(X) = \odot$ whenever $\mu_n^{\mathfrak{M}}(X) \notin \mathbb{Q}$.

If the structure \mathfrak{M} and the arity n are clear from context, we will write μ and v instead of $\mu_n^{\mathfrak{M}}$ and $v_n^{\mathfrak{M}}$.

Suppose that \mathfrak{M} is an AML \mathcal{L} -quasistructure. Let V be the set of variables and suppose also that $s : V \rightarrow M$ is a function; we call such an s a *valuation on \mathfrak{M}* . Given any sequence of distinct variables $\vec{v} = (v_1, \dots, v_n)$ and any sequence $\vec{c} = (c_1, \dots, c_n)$ from M , we let $s(\vec{c}/\vec{v})$ be the valuation obtained from s by redefining (if necessary) s at v_i to take the value c_i .

We would now like to define the satisfaction relation $\mathfrak{M} \models \varphi[s]$ by recursion on complexity of formulae in such a way that this extends the satisfaction relation for first-order logic. However, in order to specify the semantics of the measure constructor, one needs to know that the set defined by the formula following the measure constructor is measurable. Thus, we need to define the semantics in stages while ensuring that the sets defined at each stage are measurable.

Towards this end, we define the *rank* of an AML \mathcal{L} -formula φ , denoted $\text{rk}(\varphi)$, by induction:

- If φ is a classical \mathcal{L} -formula, then $\text{rk}(\varphi) = 0$;
- $\text{rk}(\neg\varphi) = \text{rk}(\forall x\varphi) = \text{rk}(\varphi)$;
- $\text{rk}(\varphi \vee \psi) = \max(\text{rk}(\varphi), \text{rk}(\psi))$;
- $\text{rk}(m_{\vec{x}} \bowtie r.\varphi(\vec{x}, \vec{y})) = \text{rk}(\varphi) + 1$.

Suppose that the satisfaction relation $\mathfrak{M} \models \varphi(\vec{x})[s]$ has been defined for formulae of rank $\leq n$. (Note that this is automatic for $n = 0$.) As is customary, we often write $\mathfrak{M} \models \varphi(\vec{a})$ if $\varphi(\vec{x})$ is a formula with free variables among \vec{x} , s is a valuation for which $s(x_i) = a_i$, and $\mathfrak{M} \models \varphi[s]$. If $\varphi(\vec{x}, \vec{y})$ is a formula of rank $\leq n$ and $b \in M^{|\vec{y}|}$, we then set $\varphi(M, \vec{b}) := \{a \in M^{|\vec{x}|} : \mathfrak{M} \models \varphi(\vec{a}, \vec{b})\}$. We also use the notation $v(\varphi, \vec{b})$ to denote $v(\varphi(M, \vec{b}))$.

We say that the AML \mathcal{L} -quasistructure \mathfrak{M} satisfies the *rank n measurability condition* if, for all formulae $\varphi(\vec{x}, \vec{y})$ of rank $\leq n$ and all $b \in M^{|\vec{y}|}$, we have $\varphi(M, b) \in \mathcal{B}_{|\vec{x}|}^{\mathfrak{M}}$.

Assume that \mathfrak{M} satisfies the rank n measurability condition. We now define the satisfaction relation \models for rank $n + 1$ formulae. We treat the connective and quantifier cases as in first-order logic and only specify how to deal with the semantics of the measure constructor. Suppose that $\varphi(\vec{x}, \vec{y})$ is a formula of rank n and s is a valuation such that $s(\vec{y}) = \vec{b}$. We then declare:

- $\mathfrak{M} \models m_{\vec{x}} < r.\varphi(\vec{x}, \vec{b})$ if and only if:
 - $\mu(\varphi(M, \vec{b})) < r$, or
 - $\mu(\varphi(M, \vec{b})) = r$ and $v(\varphi, \vec{b}) = \ominus$.
- $\mathfrak{M} \models m_{\vec{x}} \leq r.\varphi(\vec{x}, \vec{b})$ if and only if:
 - $\mu(\varphi(M, \vec{b})) < r$, or
 - $\mu(\varphi(M, \vec{b})) = r$ and $v(\varphi, \vec{b}) \neq \oplus$.

Definition 2.5. We say that the AML quasistructure \mathfrak{M} is a *structure* if it satisfies the rank n measurability condition for all n .

If φ has no free variables, then φ is called a *sentence*. If T is a set of sentences and \mathfrak{M} is a structure, we write $\mathfrak{M} \models T$ if and only if $\mathfrak{M} \models \varphi$ for all $\varphi \in T$; in this case, we say that \mathfrak{M} is a *model* of T .

Example 2.6. Let \mathcal{L} be the signature consisting of a single binary function symbol \cdot and a single constant symbol e . Let G be a group with $|G| = m$. We view G as a first-order \mathcal{L} -structure \mathfrak{G} by interpreting \cdot as multiplication in G and

e as the identity of G . In order to view \mathfrak{G} as an AML \mathcal{L} -quasistructure, we set $\mathcal{B}_n^{\mathfrak{G}} := \mathcal{P}(G^n)$ and $\mu_n(X) := \frac{|X|}{m^n}$, that is, we equip each cartesian power G^n with the (normalized) counting measure. It is clear that the AML quasistructure \mathfrak{G} is actually a structure.

Now fix $g \in G$ and let $\varphi(x, y)$ be the first-order \mathcal{L} -formula $x \cdot y = y \cdot x$. Then $\mathfrak{G} \models m_x < r \cdot \varphi(g)$ if and only if the cardinality of $C_G(g)$, the centralizer of g in G , is $< r \cdot m$. In particular, if $r < 1$, then $\mathfrak{G} \not\models \forall y m_x < r \cdot \varphi$ as the centralizer of the identity has cardinality m .

As in classical logic, we say that $D \subseteq M^n$ is *definable (in \mathfrak{M})* if there is a formula $\varphi(\vec{x}, \vec{y})$, with $|\vec{x}| = n$, and $\vec{a} \in M^{|\vec{y}|}$ such that $D = \varphi(M, \vec{a})$. If $A \subseteq M$ and the tuple \vec{a} as above lies in A , we also say that D is *A-definable (in \mathfrak{M})* or *definable (in \mathfrak{M}) over A*. We let $\text{Def}_n(M)$ denote the Boolean algebra of definable subsets of M^n . A function $h : M^m \rightarrow M^n$ is *definable (in \mathfrak{M})* if the graph of h is a definable subsets of M^{m+n} .

Convention 2.7. From now on, when we say that \mathfrak{M} is an \mathcal{L} -(quasi)structure, we mean that \mathfrak{M} is an \mathcal{L} -(quasi)structure of AML. We use fraktur letters \mathfrak{M} and \mathfrak{N} (sometimes decorated with subscripts) to denote quasistructures. The corresponding roman letter denotes the underlying universe of the quasistructure. (So, for example, the underlying universe of \mathfrak{M} will be M and the underlying universe of \mathfrak{N}_i will be N_i .) Likewise, \mathcal{L} -formulae will always mean AML \mathcal{L} -formulae. If we wish to consider a structure from classical first-order logic, then we will speak of a *classical* structure and denote it using calligraphic letters such as \mathcal{M} . Likewise, we will speak of *classical* formulae.

Suppose \mathcal{L}' is signature extending the signature \mathcal{L} . Suppose that \mathfrak{M} is an \mathcal{L}' -quasistructure. Then the \mathcal{L} -reduct of \mathfrak{M} , denoted $\mathfrak{M}|_{\mathcal{L}}$, is the \mathcal{L} -quasistructure obtained from \mathfrak{M} by “forgetting” to interpret symbols from $\mathcal{L}' \setminus \mathcal{L}$. We also refer to \mathfrak{M} as an *expansion* of $\mathfrak{M}|_{\mathcal{L}}$.

Definition 2.8. Suppose that \mathfrak{M} and \mathfrak{N} are \mathcal{L} -structures. We say that \mathfrak{M} is a *substructure* of \mathfrak{N} if:

- (1) the first-order part of \mathfrak{M} is a substructure of the first-order part of \mathfrak{N} , and
- (2) for each $B \in \mathcal{B}_n^{\mathfrak{N}}$, we have $B \cap M^n \in \mathcal{B}_n^{\mathfrak{M}}$.

We further say that \mathfrak{M} is an *elementary substructure* of \mathfrak{N} if, for all formulae $\varphi(\vec{x})$ and $a \in M^{|\vec{x}|}$, we have $\mathfrak{M} \models \varphi(\vec{a})$ if and only if $\mathfrak{N} \models \varphi(\vec{a})$.

It may seem a bit strange that the notion of substructure does not mention the measure part of the structures in question. However, given that the constructor $m_{\vec{x}} < r$ acts as a quantifier and arbitrary substructures need not respect the truth of quantified formulae, this lack of a requirement need not be so surprising.

We define the notions of *embedding* and *elementary embedding* in the obvious way so as to align with our notion of substructure and elementary substructure.

3. THE CANONICAL MEASURES ON AN \aleph_1 -SATURATED AML STRUCTURE

An extremely important source of AML structures are those that are equipped with an actual measure.

Definition 3.1. A *measured \mathcal{L} -structure* is an AML \mathcal{L} -structure \mathfrak{M} such that there is a measure ν^n on a σ -algebra \mathcal{B}^n of subsets of M^n extending $\text{Def}_n(M)$ such that:

- (1) For definable $B \subseteq M^n$, we have $\nu^n(B) = \mu_n^{\mathfrak{M}}(B)$.
- (2) For each $m, n \geq 1$, the measure ν^{m+n} is an extension of the product measure $\nu^m \otimes \nu^n$.
- (3) (Fubini property) For every $m, n \geq 1$ and $B \in \mathcal{B}^{m+n}$, we have
 - (a) for almost all $x \in M^m$, $B^x \in \mathcal{B}^n$;
 - (b) for almost all $y \in M^n$, $B_y \in \mathcal{B}^m$;
 - (c) the functions $x \mapsto \nu^n(B^x)$ are $y \mapsto \nu^m(B_y)$ are ν^m - and ν^n -measurable, respectively, and

$$\nu^{m+n}(B) = \int \nu^n(B^x) d\nu^m(x) = \int \nu^m(B_y) d\nu^n(y).$$

For example, any finite \mathcal{L} -structure equipped with its normalized counting measure is a measured structure; in fact, these structures (and their ultraproducts) are our primary examples of measured structures.

Even if an AML structure does not come equipped with a measure, often there is still a “natural” measure that can be placed on the structure, a process that we now describe. First, we say that an \mathcal{L} -structure \mathfrak{M} is \aleph_1 -compact if whenever $n \geq 1$ and $(D_m : m \in \mathbb{N})$ is a family of definable subsets of M^n with the finite intersection property, then $\bigcap_{m \in \mathbb{N}} D_m \neq \emptyset$. (If the signature \mathcal{L} is countable, then this coincides with the notion of an \aleph_1 -saturated structure more commonly encountered in model theory; when the signature is countable, we will often use the term saturated rather than compact.) For example, any finite \mathcal{L} -structure is \aleph_1 -compact; many ultraproducts are also \aleph_1 -compact. In this section, we will show that certain \aleph_1 -compact AML structures can be equipped with a family of canonical measures, the so-called *Loeb measures*.

We write $\sigma(\mathcal{B})$ for the σ -algebra generated by the algebra \mathcal{B} .

Suppose that \mathfrak{M} is an \aleph_1 -compact AML \mathcal{L} -structure and $(A_m : m \in \mathbb{N})$ is a family of definable subsets of M^n for which $\bigcup_{m \in \mathbb{N}} A_m$ is also definable. Then, by \aleph_1 -compactness, we have that $\bigcup_{m \in \mathbb{N}} A_m = \bigcup_{m=0}^k A_m$ for some $k \in \mathbb{N}$. Consequently, $\mu_n^{\mathfrak{M}}|_{\text{Def}_n(M)}$ is a pre-measure. Thus, by the Caratheodory extension theorem, there is a measure, which we again call $\mu_n^{\mathfrak{M}}$ or simply μ_n , on $\sigma(\text{Def}_n(M))$, extending the original $\mu_n^{\mathfrak{M}}$ given by

$$\mu_n(B) := \inf \left\{ \sum_m \mu_n^{\mathfrak{M}}(D_m) : B \subseteq \bigcup_m D_m, D_m \in \text{Def}_n(M) \text{ for all } m \right\}.$$

Moreover, if $\mu_n^{\mathfrak{M}}|_{\sigma(\text{Def}_n(M))}$ is σ -finite, then μ_n is the unique measure on $\sigma(\text{Def}_n(M))$ extending the original $\mu_n^{\mathfrak{M}}$. Observe, however, that by \aleph_1 -compactness again, $\mu_n^{\mathfrak{M}}|_{\text{Def}_n(M)}$ is σ -finite if and only if $\mu_n^{\mathfrak{M}}|_{\text{Def}_n(M)}$ is finite (which occurs if and only if $\mu_1^{\mathfrak{M}}(M)$ is finite by the product property). We refer to the family of measures $(\mu_n^{\mathfrak{M}})$ as the *canonical measures* on \mathfrak{M} . Although they are defined for all \aleph_1 -compact \mathcal{L} -structures, we mainly consider them in the case when $\mu_1^{\mathfrak{M}}(M)$ is finite, for then they are truly “canonical.” (Indeed, when a pre-measure on an algebra \mathcal{A} is not σ -finite, there are many extensions of it to a measure on $\sigma(\mathcal{A})$; see Exercise 1.7.8 in [29].)

Until further notice, we assume that $\mu_1^{\mathfrak{M}}|_{\text{Def}_1(M)}$ is a finite pre-measure on $\text{Def}_1(M)$ (and hence $\mu_n^{\mathfrak{M}}|_{\text{Def}_n(M)}$ is a finite pre-measure on $\text{Def}_n(M)$ for all $n \geq 1$). Let $\sigma(\text{Def}_m(M)) \otimes \sigma(\text{Def}_n(M))$ denote the product σ -algebra and $\mu_m \otimes \mu_n$ denote the product measure. Observe that $\sigma(\text{Def}_m(M)) \otimes \sigma(\text{Def}_n(M))$ is generated by sets of the form $A \times B$, with $A \in \text{Def}_m(M)$ and $B \in \text{Def}_n(M)$; since $A \times B \in \text{Def}_{m+n}(M)$ for such A, B , we have $\sigma(\text{Def}_m(M)) \otimes \sigma(\text{Def}_n(M)) \subseteq$

$\sigma(\text{Def}_{m+n}(M))$. Now observe that, if $A \in \text{Def}_m(M)$ and $B \in \text{Def}_n(M)$, then $\mu_{m+n}(A \times B) = \mu_m(A) \cdot \mu_n(B)$ by the Product property. Thus

$$E := \{D \in \sigma(\text{Def}_m(M)) \otimes \sigma(\text{Def}_n(M)) : \mu_{m+n}(D) = (\mu_m \otimes \mu_n)(D)\}$$

contains all sets of the form $A \times B$, where $A \in \text{Def}_m(M)$ and $B \in \text{Def}_n(M)$. One can easily check that E is a λ -class. Since the sets of the form $A \times B$, where $A \in \text{Def}_m(M)$ and $B \in \text{Def}_n(M)$, form a π -class, we have, by the π - λ Theorem, that $\sigma(\text{Def}_m(M)) \otimes \sigma(\text{Def}_n(M)) \subseteq E$. (For the definitions of λ - and π -classes as well as the statement of the π - λ -theorem, see Section 17.A of [23].) Consequently, we see that $\mu_m \otimes \mu_n = \mu_{m+n}|_{\sigma(\text{Def}_m(M)) \otimes \sigma(\text{Def}_n(M))}$. This proves:

Lemma 3.2. *If \mathfrak{M} is an \aleph_1 -compact \mathcal{L} -structure such that $\mu^{\mathfrak{M}}(M)$ is finite, then \mathfrak{M} , equipped with its canonical measures, satisfies the first two axioms in the definition of a measured structure.*

The third axiom, requiring that Fubini's Theorem hold, might still fail for sets which do not belong to $\sigma(\text{Def}_m(M)) \otimes \sigma(\text{Def}_n(M))$. We call an \aleph_1 -compact \mathcal{L} -structure whose canonical measures satisfy the Fubini property a *Fubini structure*. (The measurability of the functions in (3) always holds for the canonical measures, so really a Fubini structure is one where the displayed equation in the statement of the Fubini property is required to hold.) Thus, by the preceding lemma, a Fubini structure, equipped with its canonical measures, is a measured \mathcal{L} -structure.

The following situation is also important:

Definition 3.3. Suppose that \mathcal{M} is a classical \mathcal{L} -structure which is equipped with a family of measures ν^n on a σ -algebra \mathcal{B}^n of subsets of \mathfrak{M}^n extending the algebra of definable sets and such that ν^{m+n} is an extension of $\nu^m \otimes \nu^n$ for all $m, n \geq 1$. We call $(\mathcal{M}, (\nu^n))$ a *classical measured structure*. Then we can turn \mathcal{M} into an \mathcal{L} -quasistructure \mathfrak{M} by declaring, for $B \subseteq M^n$, that $v_n^{\mathfrak{M}}(B) = \odot$. We call \mathfrak{M} the *AML quasistructure associated to $(\mathcal{M}, (\nu^n))$* .

Remark 3.4. Usually, the AML quasistructure associated to a classical measured structure is not a structure: due to the added expressivity power, there are new definable sets which may not be measurable. (However this quasistructure will satisfy the rank 0 measurability condition, and so satisfaction for formulae of rank 1 can be defined.) For finite classical measured structures (which are the structures most important for our combinatorial applications), the associated AML quasistructure is in fact a structure because all sets are measurable.

Remark 3.5. Consider \mathfrak{M} , an \aleph_1 -compact \mathcal{L} -structure with $\mu_1^{\mathfrak{M}}(M)$ finite, equipped with its family $(\mu_n^{\mathfrak{M}})$ of canonical measures. Let \mathcal{M} be the classical \mathcal{L} -structure obtained from \mathfrak{M} by considering only the interpretations of the symbols in \mathcal{L} . Then $(\mathcal{M}, (\mu_n^{\mathfrak{M}}))$ is a classical measured structure. Let \mathfrak{M}' be the AML quasistructure associated to $(\mathcal{M}, (\mu_n^{\mathfrak{M}}))$. How do \mathfrak{M} and \mathfrak{M}' compare? Suppose that $\varphi(x)$ is a classical \mathcal{L} -formula and $r \in \mathbb{Q}^{\geq 0}$. It is straightforward to see that $\mathfrak{M} \models m_x \leq r \cdot \varphi(x)$ implies $\mathfrak{M}' \models m_x \leq r \cdot \varphi(x)$. However, the converse need not hold. For example, suppose that $\mathfrak{M}' \models m_x \leq 0 \cdot \varphi(x)$, whence $\mu(\varphi(M)) = 0$. However, if $v^{\mathfrak{M}}(\varphi) = \oplus$, then $\mathfrak{M} \models \neg m_x \leq 0 \cdot \varphi(x)$. (Such phenomena can, and will, happen in ultraproducts, as we will see in later sections.) Thus, in some sense, \mathfrak{M}' is a “completion” of \mathfrak{M} .

We now briefly turn to the question of what our logic can say about these canonical measures. We present a couple of examples. Suppose that (X, \mathcal{S}) is a measurable space such that $\{x\} \in \mathcal{S}$ for every $x \in X$. (This happens, for example, when

(X, \mathcal{S}) = ($M^n, \sigma(\text{Def}_n(M))$) for \mathfrak{M} an \mathcal{L} -structure.) Then a measure μ on X is said to be *continuous* if $\mu(\{x\}) = 0$ for each $x \in X$. The following is immediate.

Proposition 3.6. *Suppose that \mathfrak{M} is an \aleph_1 -compact \mathcal{L} -structure. Then we have $\mathfrak{M} \models \{\forall x m_y \leq q(x = y) : q \in \mathbb{Q}^{>0}\}$ if and only if μ^1 is a continuous measure.*

More generally, there is a set $\Gamma(x)$ consisting of \mathcal{L} -formulae containing a single free variable x such that, for every \mathcal{L} -structure \mathfrak{M} and every $a \in M$,

$$\mu^1(\{a\}) = 0 \Leftrightarrow \mathfrak{M} \models \varphi(a) \text{ for all } \varphi \in \Gamma.$$

Indeed, take $\Gamma(x) := \{m_y < q.(x = y) : q \in \mathbb{Q}^{>0}\}$. In model-theoretic terms, the set of elements of μ^1 -measure 0 is uniformly type-definable in all \aleph_1 -compact \mathcal{L} -structures. Below, as a consequence of the ultraproduct construction, we will see that this set is not uniformly definable in all \mathcal{L} -structures.

Here is another proposition along the same lines.

Proposition 3.7. *There is a set T of AML sentences such that, for any \aleph_1 -compact \mathcal{L} -structure \mathfrak{M} with $\mu_1^{\mathfrak{M}}(M)$ finite, $\mathfrak{M} \models T$ if and only if μ^1 is a probability measure.*

Proof. Take $T := \{m_x \leq 1(x = x)\} \cup \{\neg m_x \leq q(x = x) : q \in (0, 1) \cap \mathbb{Q}\}$. \square

Proposition 3.8. *There is a set T of AML sentences such that, for any \aleph_1 -compact \mathcal{L} -structure \mathfrak{M} with $\mu_1^{\mathfrak{M}}(M)$ finite, $\mathfrak{M} \models T$ if and only if \mathfrak{M} is a Fubini structure.*

Proof. We take T to consist of the following two schemes:

- (1) For every $\varphi(\vec{x}, \vec{y}, \vec{z}), \psi(\vec{x}, \vec{z})$ and q, r, t with $t < qr$,

$$\forall \vec{z} [(\forall \vec{x}(\psi \rightarrow m_{\vec{y}} \geq r.\varphi) \wedge m_{\vec{x}} \geq q.\psi) \rightarrow m_{\vec{x}, \vec{y}} > t.(\varphi \wedge \psi)],$$

- (2) For every $\varphi(\vec{x}, \vec{y}, \vec{z}), \psi(\vec{x}, \vec{z})$ and q, r, t with $qr < t$,

$$\forall \vec{z} [(\forall \vec{x}(\psi \rightarrow m_{\vec{y}} \leq r.\varphi) \wedge m_{\vec{x}} \leq q.\psi) \rightarrow m_{\vec{x}, \vec{y}} < t.(\varphi \wedge \psi)].$$

Note that ψ may not contain variables from \vec{y} .

Suppose $\mathfrak{M} \models T$. It suffices to show that for any definable set $B \subseteq M^{m+n}$, we have $\mu^{m+n}(B) = \int \mu^n(B^{\vec{x}}) d\mu^m(\vec{x})$. Fix a formula $\varphi(\vec{x}, \vec{y}, \vec{z})$ and $\vec{c} \in M^{|\vec{z}|}$ so that $B = \varphi(M, \vec{c})$. Given rationals r, r' , we set $B_{r, r'} = \{\vec{x} : \mathfrak{M} \models m_{\vec{y}} \geq r.\varphi(\vec{x}, \vec{y}, \vec{c}) \wedge m_{\vec{y}} < r'.\varphi(\vec{x}, \vec{y}, \vec{c})\}$; we will also write $B_{r, r'}$ as an abbreviation for its defining formula.

Observe that we can approximate the function $\vec{x} \mapsto \mu^n(B^{\vec{x}})$ by step functions of the form $\sum_i r_i \chi_{B_{r_i, r_{i+1}}}$ for sequences of rationals $r_1 < \dots < r_k$. Let $r_1 < \dots < r_k$ be such a sequence of rationals. For any rational $q_i < \mu^m(B_{r_i, r_{i+1}})$, we must have $\mathfrak{M} \models m_{\vec{x}} \geq q_i.B_{r_i, r_{i+1}}$. Since $\mathfrak{M} \models T$, we must have

$$\mathfrak{M} \models m_{\vec{x}, \vec{y}} > t.(\varphi \wedge B_{r_i, r_{i+1}})$$

for every $t < q_i r_i$, and so $\mu^{m+n}(\{(\vec{x}, \vec{y}) \in B : \vec{x} \in B_{r_i, r_{i+1}}\}) \geq q_i r_i$. Since the $B_{r_i, r_{i+1}}$ are pairwise disjoint, it follows that $\mu^{m+n}(B) \geq \sum_i q_i r_i$. Since we may choose q_i, r_i as above so that $\sum_i q_i r_i$ is arbitrarily close to $\int \mu^n(B^{\vec{x}}) d\mu^m(\vec{x})$, we have $\mu^{m+n}(B) \geq \int \mu^n(B^{\vec{x}}) d\mu^m(\vec{x})$.

Similarly, for any $q_i > \mu^m(B_{r_i, r_{i+1}})$, we must have $\mathfrak{M} \models m_{\vec{x}} \leq q_i.B_{r_i, r_{i+1}}$, and since $\mathfrak{M} \models T$, we also have

$$\mathfrak{M} \models m_{\vec{x}, \vec{y}} < t.(\varphi \wedge B_{r_i, r_{i+1}})$$

for any $t > q_i r_{i+1}$. Therefore $\mu^{m+n}(\{(\vec{x}, \vec{y}) \in B : \vec{x} \in B_{r_i, r_{i+1}}\}) \leq q_i r_{i+1}$, and so $\mu^{m+n}(B) \leq \sum_i q_i r_{i+1}$. Again, since $\sum_i q_i r_{i+1}$ can be made arbitrarily close to $\int \mu^n(B^{\vec{x}}) d\mu^m(\vec{x})$, we have $\mu^{m+n}(B) \leq \int \mu^n(B^{\vec{x}}) d\mu^m(\vec{x})$.

Now suppose \mathfrak{M} is Fubini; we must show that every sentence in T holds. For the first family of sentences, suppose $\mathfrak{M} \models \forall \vec{x} (\psi \rightarrow m_{\vec{y}} \geq r \cdot \varphi) \wedge m_{\vec{x}} \geq q \cdot \psi$ (with parameters $\vec{c} \in M^{|\vec{z}|}$). Setting $C = \psi(M, \vec{c})$, we see that $\mu(C) \geq q$. Let $B = \varphi(M, \vec{c})$. For each $\vec{a} \in C$, $\mu(B^{\vec{a}}) \geq r$. Since μ is Fubini, $\mu^{m+n}(B) = \int_C \mu^n(B^{\vec{x}}) d\mu^m(\vec{x}) \geq qr$, and therefore for each $t < qr$ we have $\mathfrak{M} \models m_{\vec{x}} > t \cdot (\varphi \wedge \psi)$.

For the second family of sentences, suppose $\mathfrak{M} \models \forall \vec{x} (\psi \rightarrow m_{\vec{y}} \leq r \cdot \varphi) \wedge m_{\vec{x}} \leq q \cdot \psi$. Again setting $C = \psi(M, \vec{c})$ and $B = \varphi(M, \vec{c})$, for each $\vec{a} \in C$, $\mu(B^{\vec{a}}) \leq r$. Since μ is Fubini, $\mu^{m+n}(B) = \int_C \mu^n(B^{\vec{x}}) d\mu^m(\vec{x}) \leq qr$, and therefore for each $t > qr$ we have $\mathfrak{M} \models m_{\vec{x}} < t \cdot (\varphi \wedge \psi)[s]$. \square

Remark 3.9. Fix T as in the proof of the previous proposition. The proof we have given above shows that $\mathfrak{M} \models T$ for any measured structure \mathfrak{M} .

4. ULTRAPRODUCTS

We suppose that the reader is familiar with the standard ultraproduct construction from first-order logic. We sometimes use the phrase “ $P(i)$ holds a.e.” to mean that the set of i 's for which $P(i)$ holds belongs to the ultrafilter.

Suppose that $(\mathfrak{M}_i : i \in I)$ is a family of \mathcal{L} -structures and suppose that \mathcal{U} is an ultrafilter on I . We let \mathfrak{M} denote the classical ultraproduct of the family (\mathfrak{M}_i) with respect to \mathcal{U} . We declare $A \in \mathcal{B}_n^{\mathfrak{M}}$ if and only if $A = \prod_{\mathcal{U}} A_i$, where each $A_i \in \mathcal{B}_n^{\mathfrak{M}_i}$; in this case, we declare $\mu_n^{\mathfrak{M}}(A) = \lim_{\mathcal{U}} \mu_n^{\mathfrak{M}_i}(A_i)$ (the *ultralimit* of $\mu_n^{\mathfrak{M}_i}(A_i)$). Observe that $\lim_{\mathcal{U}} (\mu_n^{\mathfrak{M}_i}(A_i))$ always exists and is unique since $[0, \infty]$ is a compact hausdorff space.

If $\mu_n^{\mathfrak{M}}(A) = r \in \mathbb{Q}$, we set

$$v_n^{\mathfrak{M}}(A) = \begin{cases} \oplus & \text{if } \mu_n^{\mathfrak{M}_i}(A_i) > r \text{ a.e.} \\ \ominus & \text{if } \mu_n^{\mathfrak{M}_i}(A_i) < r \text{ a.e.} \\ \oplus & \text{if } \mu_n^{\mathfrak{M}_i}(A_i) = r \text{ and } v_n^{\mathfrak{M}_i}(A_i) = \oplus \text{ a.e.} \\ \ominus & \text{if } \mu_n^{\mathfrak{M}_i}(A_i) = r \text{ and } v_n^{\mathfrak{M}_i}(A_i) = \ominus \text{ a.e.} \\ \odot & \text{if } \mu_n^{\mathfrak{M}_i}(A_i) = r \text{ and } v_n^{\mathfrak{M}_i}(A_i) = \odot \text{ a.e.} \end{cases}$$

It is reasonably straightforward to check that each $\mathcal{B}_n^{\mathfrak{M}}$ is an algebra and $\mu_n^{\mathfrak{M}}$ is a finitely additive measure on $\mathcal{B}_n^{\mathfrak{M}}$, whence the ultraproduct \mathfrak{M} is a quasistructure.

Note that the assignment of $v_n^{\mathfrak{M}}(A)$ is consistent with the interpretations of \oplus and \ominus as indicating approximation from above or below: if $\mu_n^{\mathfrak{M}_i}(A_i) > \mu_n^{\mathfrak{M}}(A)$ a.e., we have obtained the measure $\mu_n^{\mathfrak{M}}(A)$ by approximations from above, thus (irrespective of whether the values of $\mu_n^{\mathfrak{M}_i}(A_i)$ were obtained exactly or obtained by approximation themselves) $v_n^{\mathfrak{M}}(A) = \oplus$; if $\mu_n^{\mathfrak{M}_i}(A_i) = r$ and $v_n^{\mathfrak{M}_i}(A_i) = \oplus$ a.e. then almost every $\mu_n^{\mathfrak{M}_i}(A_i)$ was found by approximations from above, so the measure of $\mu_n^{\mathfrak{M}}(A)$ is found the same way.

The following theorem is of fundamental importance for our applications of AML.

Theorem 4.1 (Łoś Theorem for AML). *Suppose that $(\mathfrak{M}_i : i \in I)$ is a family of \mathcal{L} -structures, \mathcal{U} is a ultrafilter on I , and $\mathfrak{M} = \prod_{\mathcal{U}} \mathfrak{M}_i$ is the ultraproduct of the family (\mathfrak{M}_i) . Let φ be an \mathcal{L} -formula and s a valuation on \mathfrak{M} . For each $i \in I$, let s_i be a valuation on \mathfrak{M}_i such that, for each variable v , we have $s(v) = [s_i(v)]_{\mathcal{U}}$. Then $\mathfrak{M} \models \varphi[s] \Leftrightarrow \mathfrak{M}_i \models \varphi[s_i]$ for \mathcal{U} -almost all i .*

Proof. As in the proof of the classical Loś theorem, we proceed by induction on the complexity of formulae. We only need to explain how to deal with the measure constructor cases, that is, we consider a formula $\varphi(\vec{x}, \vec{y})$ such that the theorem holds true for φ and we show that it holds for $m_{\vec{x}} \boxtimes r.\varphi(\vec{x}, \vec{y})$, where $\boxtimes \in \{<, \leq\}$. We fix valuations s_i and s as in the statement of the theorem and set $\vec{b}_i := s_i(\vec{y})$ and $\vec{b} := s(\vec{y})$. We will repeatedly use the inductive hypothesis that $\varphi(M, \vec{b}) = \prod_{\mathcal{U}} \varphi(M_i, \vec{b}_i)$.

If $\mu^{\mathfrak{M}_i}(\varphi(M_i, \vec{b}_i)) < r$ a.e. then $\mathfrak{M}_i \models m_{\vec{x}} < r.\varphi(\vec{x}, \vec{b}_i)$ and $\mathfrak{M}_i \models m_{\vec{x}} \leq r.\varphi(\vec{x}, \vec{b}_i)$ a.e.; also either $\mu^{\mathfrak{M}}(\varphi(M, \vec{b})) < r$ or both $\mu^{\mathfrak{M}}(\varphi(M, \vec{b})) \leq r$ and $v^{\mathfrak{M}}(\varphi, \vec{b}) = \ominus$, and in either case $\mathfrak{M} \models m_{\vec{x}} < r.\varphi(\vec{x}, \vec{b})$ and $\mathfrak{M} \models m_{\vec{x}} \leq r.\varphi(\vec{x}, \vec{b})$.

If $\mu^{\mathfrak{M}_i}(\varphi(M_i, \vec{b}_i)) > r$ a.e. then, symmetrically, we have $\mathfrak{M}_i \models m_{\vec{x}} > r.\varphi(\vec{x}, \vec{b}_i)$ and $\mathfrak{M}_i \models m_{\vec{x}} \geq r.\varphi(\vec{x}, \vec{b}_i)$ a.e., and also $\mathfrak{M} \models m_{\vec{x}} > r.\varphi(\vec{x}, \vec{b})$ and $\mathfrak{M} \models m_{\vec{x}} \geq r.\varphi(\vec{x}, \vec{b})$.

Finally, if $\mu^{\mathfrak{M}_i}(\varphi(M_i, \vec{b}_i)) = r$ a.e. then $v^{\mathfrak{M}}(\varphi(M, \vec{b})) = v^{\mathfrak{M}_i}(\varphi(M_i, \vec{b}_i))$ a.e., so $\mathfrak{M} \models m_{\vec{x}} < r.\varphi(\vec{x}, \vec{b})$ iff $\mathfrak{M}_i \models m_{\vec{x}} < r.\varphi(\vec{x}, \vec{b}_i)$ a.e. and $\mathfrak{M} \models m_{\vec{x}} \leq r.\varphi(\vec{x}, \vec{b})$ iff $\mathfrak{M}_i \models m_{\vec{x}} \leq r.\varphi(\vec{x}, \vec{b}_i)$. \square

Corollary 4.2. *The ultraproduct of structures is a structure.*

Proof. Definable subsets of the ultraproduct are ultraproducts of definable sets in the factor structures. \square

Corollary 4.3. *With the same hypotheses as in the previous theorem, if σ is an \mathcal{L} -sentence, then*

$$\mathfrak{M} \models \sigma \Leftrightarrow \mathfrak{M}_i \models \sigma \text{ for } \mathcal{U}\text{-almost all } i.$$

Corollary 4.4 (Compactness Theorem for AML). *If Σ is a set of \mathcal{L} -sentences such that each finite subset of Σ has a model, then Σ has a model.*

Proof. The usual ultraproduct proof of the Compactness theorem from the Loś theorem applies; see, for example, Corollary 4.1.11 of [6]. \square

Corollary 4.5. *For any \mathcal{L} -structure \mathfrak{M} , any set I , and any nonprincipal ultrafilter \mathcal{U} on I , the diagonal embedding $j : \mathfrak{M} \rightarrow \mathfrak{M}^{\mathcal{U}}$ given by $j(a) = [(a)]_{\mathcal{U}}$ is an elementary embedding.*

The following result is standard in the classical context; this proof carries over immediately to the framework of AML.

Proposition 4.6. *Suppose that \mathcal{U} is a countably incomplete ultrafilter on a set I and suppose that \mathcal{L} is countable. Suppose that $(\mathfrak{M}_i : i \in I)$ is a family of \mathcal{L} -structures. Then $\prod_{\mathcal{U}} \mathfrak{M}_i$ is \aleph_1 -saturated. In particular, if $I = \mathbb{N}$ and \mathcal{U} is any nonprincipal ultrafilter on \mathbb{N} , then $\prod_{\mathcal{U}} \mathfrak{M}_i$ is \aleph_1 -saturated.*

Consequently, an ultraproduct \mathfrak{M} for which $\mu_1^{\mathfrak{M}}$ is finite can be equipped with its canonical family of measures. (Observe that, by the Loś theorem, $\mu_1^{\mathfrak{M}}$ is finite if and only if the sequence $(\mu_1^{\mathfrak{M}_i})$ is “essentially bounded.”) These measures are instrumental for our applications of AML to problems in finite combinatorics and ergodic theory.

We can also use the Loś theorem to show that properties of the canonical measures on the factor models are inherited by the canonical measures on the ultraproduct.

Corollary 4.7. *Suppose that $(\mathfrak{M}_i : i \in I)$ is a family of Fubini structures, \mathcal{U} is a nonprincipal ultrafilter on I , and $\mathfrak{M} = \prod_{\mathcal{U}} \mathfrak{M}_i$. Suppose further that there is a $c \in \mathbb{R}^{>0}$ such that $\mu_1^{\mathfrak{M}}(M_i) \leq c$ for \mathcal{U} -almost all i . Then \mathfrak{M} is a Fubini structure.*

Proof. This is immediate from Proposition 3.8 and the Łoś theorem. \square

Corollary 4.8. *Suppose that $(\mathfrak{M}_i : i \in I)$ is a family of \aleph_1 -compact \mathcal{L} -structures such that $\mu_1^{\mathfrak{M}_i}$ is a continuous measure for almost all i . Suppose that \mathcal{U} is a nonprincipal ultrafilter on I and $\mathfrak{M} = \prod_{\mathcal{U}} \mathfrak{M}_i$. Then $\mu_1^{\mathfrak{M}}$ is also a continuous measure.*

Proof. This is immediate from Proposition 3.6 and the Łoś theorem. \square

Remark 4.9. Under some set-theoretic assumptions, the converse of the above corollary holds. A nonprincipal ultrafilter \mathcal{U} on a (necessarily uncountable) index set I is said to be *countably complete* if whenever $(D_m : m \in \mathbb{N})$ is a sequence of sets from \mathcal{U} such that $D_m \supseteq D_{m+1}$ for all m , we have $\bigcap_{m \in \mathbb{N}} D_m \in \mathcal{U}$. With the same hypotheses as in the previous corollary, if we further assume that \mathcal{U} is countably complete, then continuity of $\mu_1^{\mathfrak{M}}$ implies continuity of almost all $\mu_1^{\mathfrak{M}_i}$. The existence of a countably complete ultrafilter is tantamount to the existence of a *measurable cardinal*, which is a fairly mild set-theoretic assumption that most set theorists are willing to assume when necessary. However, it is straightforward to check that if \mathcal{U} is countably complete and each \mathfrak{M}_i is finite, then $\prod_{\mathcal{U}} \mathfrak{M}_i$ will also be finite, making such ultraproducts unusable for many purposes.

Corollary 4.10. *There does not exist an \mathcal{L} -formula $\varphi(x)$ such that for every \aleph_1 -compact \mathcal{L} -structure \mathfrak{M} , we have*

$$\mu_1^{\mathfrak{M}}(\{a\}) = 0 \Leftrightarrow \mathfrak{M} \models \varphi(a).$$

Proof. Suppose, towards a contradiction, that such an \mathcal{L} -formula $\varphi(x)$ exists. For $i \geq 1$, let \mathfrak{M}_i be the classical measured \mathcal{L} -structure whose universe is $\{1, \dots, i\}$, whose measures are given by the normalized counting measures, and which interprets the \mathcal{L} -structure in an arbitrary fashion. Let \mathfrak{M}_i be the AML \mathcal{L} -structure associated with \mathfrak{M}_i . Let \mathcal{U} be a nonprincipal ultrafilter on $\mathbb{N}^{>0}$ and let $\mathfrak{M} = \prod_{\mathcal{U}} \mathfrak{M}_i$. It is easy to see that $\mu_1^{\mathfrak{M}}(M) \leq 1$ and $\mu_1^{\mathfrak{M}}(\{a\}) = 0$ for any $a \in M$. It follows that $\mathfrak{M} \models \forall x \varphi(x)$. By Łoś, we have that $\mathfrak{M}_i \models \forall x \varphi(x)$ for some $i \in \mathbb{N}^{>0}$. Since \mathfrak{M}_i is \aleph_1 -compact, it follows that each $m \in M_i$ has measure 0, which is a contradiction. \square

Corollary 4.11. *There does not exist an \mathcal{L} -sentence φ such that, whenever \mathfrak{M} is an \aleph_1 -compact structure with $\mu_1^{\mathfrak{M}}(M)$ finite, then $\mathfrak{M} \models \varphi$ if and only if $\mu_1^{\mathfrak{M}}$ is a probability measure.*

Proof. Suppose, towards a contradiction, that such a sentence φ exists. For $n \geq 1$, let \mathfrak{M}_n be a finite \mathcal{L} -structure such that $\mu_1^{\mathfrak{M}_n}(M_n) = 1 - \frac{1}{n}$. Let $\mathfrak{M} := \prod_{\mathcal{U}} \mathfrak{M}_n$, where \mathcal{U} is a nonprincipal ultrafilter on $\mathbb{N}^{>0}$. Observe that $\mu_1^{\mathfrak{M}}(M) = 1$, whence $\mathfrak{M} \models \varphi$. Thus, $\mathfrak{M}_n \models \varphi$ for some n , contradicting the fact that $\mu^{\mathfrak{M}_n} 1(M_n) < 1$. \square

Corollary 4.12. *There does not exist a set T of \mathcal{L} -sentences such that, for any \mathcal{L} -structure \mathfrak{M} , $\mathfrak{M} \models T$ if and only if $\mu_1^{\mathfrak{M}}$ is σ -finite.*

Proof. Suppose, towards a contradiction, that there is a set T of \mathcal{L} -sentences such that $\mathfrak{M} \models T$ if and only if $\mu_1^{\mathfrak{M}}$ is σ -finite. Let M be a countably infinite set equipped with the unnormalized counting measure $\nu(B) = |B|$ and

turn M into an AML \mathcal{L} -structure \mathfrak{M} by interpreting the symbols in \mathcal{L} in an arbitrary fashion. Note that $\mathfrak{M} \models T$ because \mathfrak{M} is the union of its singletons, which are definable and have measure 1. Let \mathcal{U} be a nonprincipal ultrafilter on \mathbb{N} and let $\mathfrak{N} := \mathfrak{M}^{\mathcal{U}}$. Then, by the Łoś theorem, we have that $\mathfrak{N} \models T$, so $\mu_1^{\mathfrak{N}}$ is σ -finite. Since \mathfrak{N} is \aleph_1 -compact, we have that $\mu_1^{\mathfrak{N}}$ is a finite measure. Thus, for some $n \in \mathbb{N}$, $\mathfrak{N} \models \neg \exists x_1 \cdots \exists x_n (\bigwedge_{i=1}^n \neg m_x < 1.(x = x_i))$. Consequently, $\mathfrak{M} \models \neg \exists x_1 \cdots \exists x_n (\bigwedge_{i=1}^n \neg m_x < 1.(x = x_i))$, contradicting the fact that M contains infinitely many singletons of measure 1. \square

We now turn to the connection between integration and ultraproducts. Suppose that $(\mathfrak{M}_i : i \in I)$ is a family of \mathcal{L} -structures and $h_i : M_i^m \rightarrow M_i^n$ are functions. We will say that the family $(h_i : i \in I)$ is *uniformly definable* if there exists a formula $\varphi(\vec{x}, \vec{y}, \vec{z})$, where $|\vec{x}| = m$ and $|\vec{z}| = n$, and tuples $\vec{c}_i \in M_i^{|\vec{z}|}$ such that, for all $i \in I$, $\vec{a} \in M_i^m$ and $\vec{b} \in M_i^n$, we have $h_i(\vec{a}) = \vec{b}$ if and only if $\mathfrak{M}_i \models \varphi(\vec{a}, \vec{b}, \vec{c}_i)$. If $\mathfrak{M} = \prod_{\mathcal{U}} \mathfrak{M}_i$ and $\vec{c} := [\vec{c}_i]_{\mathcal{U}}$, then by Łoś, we have that $\varphi(\vec{x}, \vec{y}, \vec{c})$ defines a function $h : M^m \rightarrow M^n$, which we call the *ultraproduct* of the uniformly definable family (h_i) .

We will also need to consider families of functions. Suppose $(\mathfrak{M}_i : i \in I)$ is a family of measured \mathcal{L} -structures, and, for each i , $g_i : M_i^m \rightarrow \mathbb{R}$ is a function. We then say the family $(g_i : i \in I)$ is *uniformly layerwise definable*² if, for each $q \in \mathbb{Q}$, there is a formula $\varphi_q(\vec{x}, \vec{z})$ and tuples $\vec{c}_i \in M_i^{|\vec{z}|}$ such that, for $\vec{a} \in M_i^m$, $g_i(\vec{a}) < q$ if and only if $\mathfrak{M}_i \models \varphi_q(\vec{a}, \vec{c}_i)$.

In the following theorem, the bound of 1 can easily be replaced by any positive real number.

Theorem 4.13. *Suppose that $(\mathfrak{M}_i : i \in I)$ is a family of measured \mathcal{L} -structures with $\mu_i^1(M_i) \leq 1$ for each $i \in I$. Suppose that \mathcal{U} is a nonprincipal ultrafilter on I and $\mathfrak{M} := \prod_{\mathcal{U}} \mathfrak{M}_i$ is equipped with its canonical measures (μ^m) . For each $i \in I$, let $g_i : M_i^n \rightarrow [-1, 1]$ be a function and suppose that the family $(g_i : i \in I)$ is uniformly layerwise definable. Then there is an expansion \mathfrak{M}' of \mathfrak{M} to a structure and a measurable (with respect to $\mathcal{B}^{\mathfrak{M}', n}$) function $g : M^n \rightarrow [-1, 1]$ such that, for every uniformly definable family $(h_i : i \in I)$ of functions of m variables with ultraproduct h , we have $\int (g \circ h) d\mu^m = \lim_{\mathcal{U}} \int (g_i \circ h_i) d\mu_i^m$.*

Proof. Consider the extension \mathcal{L}' of \mathcal{L} obtained by adding, for each $q \in \mathbb{Q} \cap [-1, 1]$, a new unary predicate G_q . We define expansions \mathfrak{M}'_i of \mathfrak{M}_i by interpreting $G_q^{\mathfrak{M}'_i} = \{\vec{x} \in M_i^n : g_i(\vec{x}) < q\}$; since this is a definitional expansion, \mathfrak{M}'_i is indeed an \mathcal{L}' -structure. Let $\mathfrak{M}' := \prod_{\mathcal{U}} \mathfrak{M}'_i$, and define $g(\vec{x}) = \inf\{q \in \mathbb{Q}^{>0} : \mathfrak{M}' \models G_q(\vec{x})\}$. It is clear that g is $\mathcal{B}^{\mathfrak{M}', n}$ -measurable.

Now suppose that $(h_i : i \in I)$ is a uniformly definable family of functions, say defined by $\varphi(\vec{x}, \vec{y})$ (and some parameters, which we suppress for sake of exposition), with ultraproduct h . Set $m := |\vec{x}|$ and $n := |\vec{y}|$. Fix $c \in \mathbb{Q}^{>0}$ and choose a partition $-1 = q_0 < q_1 < \cdots < q_k = 1$ of $[-1, 1]$ by rational numbers such that that $|q_j - q_{j+1}| < 1/c$ for each $j < k$. Fix also some rational $q_{k+1} > 1$. For $0 \leq j \leq k$, set

$$S_i^{q_j, q_{j+1}} := \{\vec{x} \in M_i^m : q_j \leq (g_i \circ h_i)(\vec{x}) < q_{j+1}\}$$

and take $S^{q_j, q_{j+1}} = \prod_{\mathcal{U}} S_i^{q_j, q_{j+1}}$, whence $\mu^m(S^{q_j, q_{j+1}}) = \lim_{\mathcal{U}} \mu_i^m(S_i^{q_j, q_{j+1}})$.

²This notion of definability of a function into \mathbb{R} differs from the usual notion in the literature, which requires the inverse image of a closed set be type-definable.

We have by definition that

$$\sum_j q_{j+1} \mu_i^m(S_i^{q_j, q_{j+1}}) - \frac{1}{c} \leq \int (g_i \circ h_i) d\mu_i^m \leq \sum_j q_{j+1} \mu_i^m(S_i^{q_j, q_{j+1}}).$$

On the other hand, the sets $S^{q_j, q_{j+1}}$ partition M^m and if $\vec{x} \in S^{q_j, q_{j+1}}$ then, since $g \circ h(\vec{x}) = \lim_{\mathcal{U}} g_i \circ h_i(\vec{x}_i)$, $q_j \leq g \circ h(\vec{x}) \leq q_{j+1}$, so also

$$\sum_j q_{j+1} \mu^m(S^{q_j, q_{j+1}}) - \frac{1}{c} \leq \int (g \circ h) d\mu^m \leq \sum_j q_{j+1} \mu^m(S^{q_j, q_{j+1}}).$$

Therefore we have $|\lim_{\mathcal{U}} \int (g_i \circ h_i) d\mu_i^m - \int (g \circ h) d\mu^m| \leq \frac{1}{c}$. Letting c go to ∞ , we have the desired result. \square

We call the function g as in the conclusion of the previous theorem the *ultra-product* of the family (g_i) . Observe that $g(\vec{x}) = \lim_{\mathcal{U}} g_i(\vec{x}_i)$ for $\vec{x} = [\vec{x}_i]_{\mathcal{U}}$. We will particularly be interested in the case that $I = \mathbb{N}$, in which case, using the same notation as in the statement of the previous theorem, if $\lim_{i \rightarrow \infty} \int (g \circ h_i) d\mu_i^m = r$, then $\int (g \circ h) d\mu^m = r$.

Remark 4.14. In the previous theorem, one could drop the requirement that the family (g_i) be uniformly layerwise definable and instead only require that each g_i be measurable. However, the \mathfrak{M}'_i and, consequently, \mathfrak{M}' , would only be quasistructures rather than structures.

5. CORRESPONDENCE THEOREMS

5.1. The Furstenberg Correspondence Principle. One application of AML is to give a general framework for *correspondence principles* between finite structures and dynamical systems. We will only present a proof of the original correspondence argument given by Furstenberg [14]; since then, many variations and generalizations have been produced (for instance, [5]). The method given extends immediately to these generalizations.

Definition 5.1. Let $E \subseteq \mathbb{Z}$. The *upper Banach density* of E , $\bar{d}(E)$, is

$$\limsup_{m-n \rightarrow \infty} \frac{|[n, m] \cap E|}{m-n}.$$

Definition 5.2. A *dynamical system* is a tuple (Y, \mathcal{B}, μ, T) , where (Y, \mathcal{B}, μ) is a probability space and $T : Y \rightarrow Y$ is a bimeasurable bijection for which $\mu(A) = \mu(TA)$ for all $A \in \mathcal{B}$.

Furstenberg's original correspondence can be stated as follows:

Theorem 5.3 (Furstenberg). *Let $E \subseteq \mathbb{Z}$ with positive upper Banach density be given. Then there is a dynamical system (Y, \mathcal{B}, μ, T) and a set $A \in \mathcal{B}$ with $\mu(A) = \bar{d}(E)$ such that for any finite set of integers U ,*

$$\bar{d}\left(\bigcap_{i \in U} (E - i)\right) \geq \mu\left(\bigcap_{i \in U} T^{-i} A\right).$$

Proof. Let $(\epsilon_N : N \in \mathbb{N})$ be an increasing sequence of positive rational numbers such that $\sup_N \epsilon_N = \bar{d}(E)$. Then for each N , we may find n, m such that $m - n > N$ and $\frac{|[n, m] \cap E|}{m-n} \geq \epsilon_N$. Let $f_{n, m} : [n, m] \rightarrow [n, m]$ be the function such that $f_{n, m}(x) = x+1$ when $x < m$ and $f_{n, m}(m) = n$. We consider the classical measured structure $([n, m], E \cap [n, m], f_{n, m})$ equipped with the normalized counting measure and consider the corresponding AML structure $\mathfrak{M}_{n, m}$.

Let (Y, A, T) be the ultraproduct of the $\mathfrak{M}_{n,m}$'s with respect to some nonprincipal ultrafilter on \mathbb{N} . Let $\mathcal{B} := \sigma(\text{Def}_1(Y))$ and let μ be the canonical measure on \mathcal{B} . It is clear from Łoś' theorem that (Y, \mathcal{B}, μ, T) is a dynamical system. Next observe that, by our construction, $\mu(A) = \bar{d}(E)$. Now notice that, for any finite set of integers U , any rational $\delta < \mu(\bigcap_{i \in U} T^{-i}A)$, and any N , we may find m, n such that $m - n > N$ and

$$\frac{|[n, m] \cap \bigcap_{i \in U} f_{n,m}^{-i}(E \cap [n, m])|}{m - n} > \delta.$$

Fix $\gamma > 0$. Then for sufficiently large N , we also have that

$$\frac{|[n, m] \cap \bigcap_{i \in U}(E - i)|}{m - n} > \delta - \gamma.$$

Consequently, $\bar{d}(\bigcap_{i \in U}(E - i)) \geq \mu(\bigcap_{i \in U} T^i A)$. \square

5.2. The Regularity Lemma.

Definition 5.4. Let \mathfrak{M} be an AML structure, let $A \subseteq M$ be a set, let n be a positive integer, and let $I \subseteq [1, n]$ be given. We define $\mathcal{B}_{n,I}^0(A)$ to be the Boolean algebra of subsets of M^n generated by sets of the form

$$\{(x_1, \dots, x_n) \in M^n : \mathfrak{M} \models \varphi(x_1, \dots, x_n)\},$$

where φ is a formula with parameters from A whose free variables belong to I .

When $k \leq n$, we define $\mathcal{B}_{n,k}^0(A)$ to be the Boolean algebra generated by the algebras $\mathcal{B}_{n,I}^0(A)$ as I ranges over subsets of $[1, n]$ of cardinality k .

In all cases, we drop 0 to indicate the σ -algebra generated by the algebra. When $A = \emptyset$, we omit it and write $\mathcal{B}_{n,k}$.

These algebras were introduced, in a somewhat different context, by Tao in [30, 31].

Definition 5.5. Suppose that (G, E) is a finite (undirected) graph, $U, U' \subseteq G$ are nonempty, and ϵ is a positive real number.

- (1) We set $d(U, U') := \frac{|E \cap (U \times U')|}{|U \times U'|}$.
- (2) We say that U and U' are ϵ -regular if whenever $V \subseteq U, V' \subseteq U'$ are such that $|V| \geq \epsilon|U|$ and $|V'| \geq \epsilon|U'|$, then we have

$$|d(U, U') - d(V, V')| < \epsilon.$$

Theorem 5.6 (Szemerédi's Regularity Lemma, [28]). *For any k and any $\epsilon > 0$, there is a $K \geq k$ such that whenever (G, E) is a finite graph, there is an $n \in [k, K]$ and a partition $G = U_1 \cup \dots \cup U_n$ such that, setting $B = \{(i, j) : U_i \text{ and } U_j \text{ are not } \epsilon\text{-regular}\}$, we have*

$$\left| \bigcup_{(i,j) \in B} U_i \times U_j \right| \leq \epsilon |G|^2.$$

Proof. For a contradiction, suppose the conclusion fails. Then we may find k, ϵ such that for every $K \geq k$ there is a finite graph (G_K, E_K) with no partition into at least k but at most K components satisfying the theorem. The partition of a graph into singleton elements consists entirely of ϵ -regular pairs, so in particular $|G_K| > K$ for all K .

Let $\mathcal{L}_0 := \{E\}$, where E is a binary predicate symbol. Inductively assume that AML signatures $\mathcal{L}_0 \subseteq \mathcal{L}_1 \subseteq \mathcal{L}_2 \subseteq \dots$ have been constructed. Then for any two AML \mathcal{L}_n -formulae $\varphi(x)$ and $\psi(x)$, at least one of which does not belong to $\bigcup_{i=1}^{n-1} \mathcal{L}_i$, with only the displayed free variable, we add new unary predicates $R_{\varphi,\psi}(x)$ and $S_{\varphi,\psi}(x)$ to \mathcal{L}_{n+1} . We let $\mathcal{L} := \bigcup_{n=1}^{\infty} \mathcal{L}_n$. We define finite measured \mathcal{L} -structures \mathfrak{G}_K as follows:

- The universe of \mathfrak{G}_K is G_K ;
- The measures on \mathfrak{G}_K are given by the normalized counting measure;
- E is interpreted in \mathfrak{G}_K by E_K ; and
- $R_{\varphi,\psi}, S_{\varphi,\psi}$ are interpreted by induction on formulae. Having interpreted φ, ψ , we may set $U = \varphi(G_K)$ and $U' = \psi(G_K)$. If U, U' are not ϵ -regular then set $R_{\varphi,\psi}^{\mathfrak{G}_K} = S_{\varphi,\psi}^{\mathfrak{G}_K} = \emptyset$. Otherwise, choose $V \subseteq U, V' \subseteq U'$ witnessing the failure of ϵ -regularity and set $R_{\varphi,\psi}^{\mathfrak{G}_K} = V$ and $S_{\varphi,\psi}^{\mathfrak{G}_K} = V'$.

Let \mathfrak{G} be the ultraproduct of the \mathfrak{G}_K . We work in the canonical measure space on G^2 . From here on, we will identify formulae φ with their interpretation $\varphi(G)$. Let $h := \mathbb{E}(\chi_E \mid \mathcal{B}_{2,1})$, the conditional expectation of χ_E with respect to the σ -algebra $\mathcal{B}_{2,1}$. Since h can be approximated by simple functions, we may write

$$h = \sum_{i \leq K_0} \alpha'_i \chi_{C_i \times D_i} + h'$$

where $\|h'\|_{L^2} < \epsilon^4/4$ and C_i, D_i are definable. Let U_1, U_2, \dots, U_n be the atoms of the finite algebra generated by $\{C_i\} \cup \{D_i\}$. Then we have

$$h = \sum_{i,j \leq n} \alpha_{i,j} \chi_{U_i \times U_j} + h''$$

with $\|h''\|_{L^2} \leq \|h'\|_{L^2}$ and $\{U_i\}_{i \leq n}$ is a finite partition of G into definable sets.

Let B be the collection of (i,j) such that R_{U_i, U_j} (and therefore S_{U_i, U_j}) are non-empty, and for each $(i,j) \in B$, define $\beta_{i,j} = \frac{\mu(E \cap (R_{U_i, U_j} \times S_{U_i, U_j}))}{\mu(R_{U_i, U_j})\mu(S_{U_i, U_j})}$. By Loś' theorem, for each $(i,j) \in B$, we have $|\beta_{i,j} - \alpha_{i,j}| \geq \epsilon$. Let $B^+ \subseteq B$ be those (i,j) such that $\alpha_{i,j} - \beta_{i,j} \geq \epsilon$. Suppose $\mu(\bigcup_{(i,j) \in B} U_i \times U_j) \geq \epsilon/2$. We may assume $\mu(\bigcup_{(i,j) \in B^+} U_i \times U_j) \geq \epsilon/4$ (for otherwise we carry out a similar argument with $B \setminus B^+$).

Again by Loś' theorem, for each $(i,j) \in B^+$, $\mu(R_{U_i, U_j}) \geq \epsilon \mu(U_i)$ and $\mu(S_{U_i, U_j}) \geq \epsilon \mu(U_j)$. Set $Z = \bigcup_{(i,j) \in B^+} (R_{U_i, U_j} \times S_{U_i, U_j})$. Note that $Z \in \mathcal{B}_{2,1}$ and

$$0 = \int \chi_E \chi_Z - h \chi_Z d\mu = \int \chi_E \chi_Z - \sum_{i,j \leq n} \alpha_{i,j} \chi_{U_i \times U_j} \chi_Z d\mu + \int h' \chi_Z d\mu$$

so we have

$$\begin{aligned} \left| \int h' \chi_Z d\mu \right| &= \left| \int \sum_{i,j \leq n} \alpha_{i,j} \chi_{U_i \times U_j} \chi_Z - \chi_E \chi_Z d\mu \right| \\ &= \left| \sum_{(i,j) \in B^+} \int \alpha_{i,j} \mu(R_{U_i, U_j}) \mu(S_{U_i, U_j}) - \chi_E \chi_{R_{U_i, U_j} \times S_{U_i, U_j}} d\mu \right| \\ &\geq \sum_{(i,j) \in B^+} \alpha_{i,j} \mu(R_{U_i, U_j}) \mu(S_{U_i, U_j}) - (\alpha_{i,j} - \epsilon) \mu(R_{U_i, U_j}) \mu(S_{U_i, U_j}) \end{aligned}$$

$$\begin{aligned}
&= \sum_{(i,j) \in B^+} \epsilon \mu(R_{U_i, U_j}) \mu(S_{U_i, U_j}) \\
&\geq \sum_{(i,j) \in B^+} \epsilon^3 \mu(U_i) \mu(U_j) \\
&\geq \epsilon^4 / 4.
\end{aligned}$$

But this is a contradiction, since $\int h' \chi_Z d\mu \leq \|h'\|_{L^2} \|\chi_Z\|_{L^2} < \epsilon^4 / 4 \cdot 1$. It follows that $\mu(\bigcup_{(i,j) \in B} U_i \times U_j) \geq \epsilon / 2$.

Note that what we have constructed is a partition of G which roughly satisfies the regularity conditions, but in the infinite model. Our final step is to pull this partition down to give a finite model, giving a contradiction. By choosing $K \geq n$ large enough, we may ensure that the following hold:

- Whenever $(i, j) \notin B$, $R_{U_i, U_j}^{\mathfrak{G}_K} = \emptyset$, and therefore $U_i(G_K)$ and $U_j(G_K)$ are ϵ -regular,
- $\mu(U_i(G_K)) \leq \sqrt{2} \mu(U_i(G))$ for each i .

This implies that

$$\mu \left(\bigcup_{(i,j) \in B} U_i(G_K) \times U_j(G_K) \right) \leq 2\mu \left(\bigcup_{(i,j) \in B} U_i(G) \times U_j(G) \right) \leq \epsilon$$

as desired. \square

Remark 5.7. The regularity lemma is usually stated with an additional requirement that the cardinalities of the partition pieces be nearly equal. This could be obtained with the following changes: for each integer n , include new unary predicates $P_{n,1}, \dots, P_{n,n}$ in the language, and in the finite models interpret these predicates as a partition into nearly equal pieces. Using these predicates, refine the partition $\{U_i\}$ into one where the components have nearly the same measure. When we pull this down to the finite model, the sets will differ in size by a very small fraction (choosing parameters correctly, this fraction can be arbitrarily small, say ϵ^2), so we may redistribute a small number of points to make these pieces have exactly the same size without making much change to the edge density of large subsets.

Remark 5.8. The method in this subsection extends very naturally to hypergraphs, using the algebras $\mathcal{B}_{k,k-1}$ in place of $\mathcal{B}_{2,1}$. An example of a related proof for hypergraphs is given below in Section 7.3.

6. THE DOWNWARD LÖWENHEIM-SKOLEM THEOREM

In this section, we prove the Downward Löwenheim-Skolem Theorem for AML using an appropriate version of the Tarski-Vaught test.

Proposition 6.1 (Tarski-Vaught Test). *Suppose that \mathfrak{M} is a substructure of \mathfrak{N} . Then \mathfrak{M} is an elementary substructure of \mathfrak{N} if and only if:*

- (1) *for all formulae $\varphi(x; \vec{y})$ and \vec{a} from M , if there is $b \in N$ with $\mathfrak{N} \models \varphi(b, \vec{a})$, then there is $c \in M$ with $\mathfrak{N} \models \varphi(c, \vec{a})$; and*
- (2) *for all formulae $\varphi(x; \vec{y})$ and \vec{a} from M , we have $\mu^{\mathfrak{M}}(\varphi(N, \vec{a}) \cap M) = \mu^{\mathfrak{N}}(\varphi(N, \vec{a}))$ and $v^{\mathfrak{M}}(\varphi(N, \vec{a}) \cap M) = v^{\mathfrak{N}}(\varphi(N, \vec{a}))$.*

Observe that, assuming (1), $\varphi(N, a) \cap M = \varphi(M, a)$, whence it makes sense to speak of $\mu^{\mathfrak{M}}(\varphi(N, a) \cap M)$ in (2).

Proof. First suppose that \mathfrak{M} is an elementary substructure of \mathfrak{N} . Then (1) is immediate. To prove (2), first notice that $\varphi(N, \vec{a}) \cap M = \varphi(M)$ by elementarity. Suppose, towards a contradiction, that $\mu(\varphi(M, \vec{a})) = r$ while $\mu(\varphi(N, \vec{a})) = s$ with $r \neq s$. Without loss of generality, suppose that $r < s$. Fix $q \in \mathbb{Q}$ with $r < q < s$. Then $\mathfrak{M} \models m_x < q \cdot \varphi(x, \vec{y})[\vec{a}]$ while $\mathfrak{N} \not\models m_x < q \cdot \varphi(x, \vec{y})[\vec{a}]$, a contradiction.

We now must show that $v(\varphi(N, \vec{a}) \cap M) = v(\varphi(N, \vec{a}))$. Set $X = \varphi(N, \vec{a})$, $Y = \varphi(M, \vec{a})$, and $r := \mu(\varphi(N, \vec{a})) = \mu(\varphi(M, \vec{a}))$. Without loss of generality, we may assume that $r \in \mathbb{Q}$. First suppose that $v(X) = \oplus$. If $v(Y) = \ominus$, then $\mathfrak{M} \models m_{\vec{x}} < r \cdot \varphi(\vec{x}; \vec{a})$ while $\mathfrak{N} \not\models m_{\vec{x}} < r \cdot \varphi(\vec{x}; \vec{a})$; if $v(Y) = \odot$, then $\mathfrak{M} \models m_{\vec{x}} \leq r \cdot \varphi(\vec{x}; \vec{a})$ while $\mathfrak{N} \not\models m_{\vec{x}} \leq r \cdot \varphi(\vec{x}; \vec{a})$. Now suppose that $v(X) = \ominus$. If $v(Y) \neq \ominus$, then $\mathfrak{M} \not\models m_{\vec{x}} < r \cdot \varphi(\vec{x}; \vec{a})$ while $\mathfrak{N} \models m_{\vec{x}} < r \cdot \varphi(\vec{x}; \vec{a})$. Finally suppose that $v(X) = \odot$. If $v(Y) = \ominus$, then $\mathfrak{M} \models m_{\vec{x}} < r \cdot \varphi(\vec{x}; \vec{a})$ while $\mathfrak{N} \not\models m_{\vec{x}} < r \cdot \varphi(\vec{x}; \vec{a})$; if $v(Y) = \oplus$, then $\mathfrak{M} \not\models m_{\vec{x}} \leq r \cdot \varphi(\vec{x}; \vec{a})$ while $\mathfrak{N} \models m_{\vec{x}} \leq r \cdot \varphi(\vec{x}; \vec{a})$.

Conversely, suppose that (1) and (2) hold for every formula; we prove that \mathfrak{M} is an elementary substructure of \mathfrak{N} by induction on complexity of formulae. The quantifier case is handled as usual using (1) while the measure constructor case is handled using (2). \square

Theorem 6.2 (Downward Löwenheim-Skolem). *Suppose that \mathfrak{N} is a structure and X is a subset of N . Then there is an elementary substructure \mathfrak{M} of \mathfrak{N} containing X with $|M| \leq \max(|X|, |\mathcal{L}|, \aleph_0)$.*

Proof. We define a sequence $(X_i)_{i < \omega}$ of subsets of N as follows. Set $X_0 := X$. Assume now that X_i has already been defined. For each formula $\varphi(x)$ with parameters from X_i that is satisfiable in \mathfrak{N} , we fix a witness b_φ . We then let X_{i+1} be the closure of X_i and all the b_φ 's under the function symbols in \mathcal{L} . Observe that, since we only added measure constructors $m_{\vec{x}} \bowtie r$ for rational r , we have $|X_i| \leq \max(|X|, |\mathcal{L}|, \aleph_0)$ for each $i < \omega$.

We now define an elementary substructure \mathfrak{M} as follows. We take $M := \bigcup_{i < \omega} X_i$ and $\mathcal{B}_{\mathfrak{M}} := \mathcal{B}_{\mathfrak{N}} \cap M$. By the usual Tarski-Vaught test, the first-order part of \mathfrak{M} is an elementary substructure of the first-order part of \mathfrak{N} . This is the beginning of an inductive argument that allows one to define, for $a \in M$, $\mu^{\mathfrak{M}}(\varphi(N, a) \cap M) := \mu^{\mathfrak{N}}(\varphi(N, a))$ and $v^{\mathfrak{M}}(\varphi(N, a) \cap M) := v^{\mathfrak{N}}(\varphi(N, a))$. Another inductive argument allows one to show that the quasistructure \mathfrak{M} is actually a structure. Now the AML Tarski-Vaught Test (Proposition 6.1) yields that \mathfrak{M} is an elementary substructure of \mathfrak{N} . \square

7. MODEL THEORY AND APPLICATIONS

7.1. More Model Theory. We introduce some standard (and less standard) notions from model theory. We often discuss the case where the language is countable, \mathfrak{M} is \aleph_1 -saturated, and we consider a countable substructure, since this is the most interesting case for our purposes. Most of these results would still hold as long as the saturation of \mathfrak{M} is greater than the cardinality of either the signature or the submodels we consider.

From now on, we fix a countable signature \mathcal{L} and a measured \mathcal{L} -structure \mathfrak{M} (although Lemmas 7.11, 7.10, and 7.13 below make no mention of the measure and thus these Lemmas go through for an arbitrary \mathcal{L} -structure). For example, \mathfrak{M} could be a Fubini structure equipped with its canonical measures. Also, for sake of readability, we will write μ for the measure, regardless of the cartesian power of M^n under discussion. Finally, we will often identify a formula $\varphi(x)$ with parameters from M with the subset of M^n that it defines; in this way, we may

write $\mu(\varphi)$ to denote the measure of the definable subset of M^n defined by φ . Likewise, we may write $\neg Y$ to indicate the complement of the definable set Y .

Definition 7.1. Suppose $A \subseteq M$.

- (1) A *partial type* of n -tuples (or an *n -type*) over A is a collection $p(\vec{x})$ of formulae with parameters from A , where $\vec{x} = (x_1, \dots, x_n)$, such that for every $\varphi_1, \dots, \varphi_n \in p$, there is $\vec{a} \in M^n$ such that $\mathfrak{M} \models \varphi_i(\vec{a})$ for every $i \leq n$.
- (2) A partial type p is said to be a *type* over A if for every formula φ with suitable free variables and parameters from A , either $\varphi \in p$ or $\neg\varphi \in p$.
- (3) A *global type* is a type over M .
- (4) If \vec{a} is a sequence of elements, we write $\text{tp}(\vec{a}/A)$ for the set of formulae with parameters from A satisfied by \vec{a} .
- (5) If \mathfrak{M} is measured, then the partial type p is *wide* if for every $\varphi \in p$, $\mu(\varphi) > 0$.

Remark 7.2. An \mathcal{L} -structure \mathfrak{M} is \aleph_1 -saturated if and only if, whenever $A \subseteq M$ is countable and $p(\vec{x})$ is a partial type over A , there is $a \in M^{|\vec{x}|}$ such that $\mathfrak{M} \models \varphi(\vec{a})$ for all $\varphi \in p$. In this case, we write $\vec{a} \models p$.

Lemma 7.3. *Suppose that $A \subseteq M$ is countable. Then for almost every \vec{a} , $\text{tp}(\vec{a}/A)$ is wide.*

Proof. Any \vec{a} such that $\text{tp}(\vec{a}/A)$ is not wide must belong to some set of measure 0 definable with parameters from A . There are countably many such sets, so their union has measure 0. \square

Lemma 7.4. *Suppose that $A \subseteq M$ is countable. Then for almost every (\vec{a}, \vec{b}) , $\text{tp}(\vec{a}/A \cup \{\vec{b}\})$ and $\text{tp}(\vec{b}/A \cup \{\vec{a}\})$ are both wide.*

Proof. By the previous lemma, for each \vec{b} the set

$$T_{\vec{b}} = \{\vec{a} : \text{tp}(\vec{a}/A \cup \{\vec{b}\}) \text{ is not wide}\}$$

has measure 0. Therefore $\mu(\{(\vec{a}, \vec{b}) : \vec{a} \in T_{\vec{b}}\}) = \int \mu(T_{\vec{b}}) d\mu^{|\vec{b}|}(\vec{b}) = 0$, so the set of (\vec{a}, \vec{b}) such that $\text{tp}(\vec{a}/A \cup \{\vec{b}\})$ is not wide has measure 0. By a symmetric argument, also the set of (\vec{a}, \vec{b}) such that $\text{tp}(\vec{b}/A \cup \{\vec{a}\})$ is not wide has measure 0. Therefore almost every (\vec{a}, \vec{b}) is in neither of these sets; for such (\vec{a}, \vec{b}) , $\text{tp}(\vec{a}/A \cup \{\vec{b}\})$ and $\text{tp}(\vec{b}/A \cup \{\vec{a}\})$ are both wide. \square

Definition 7.5. Suppose that $(\vec{b}_i : i \in I)$ is a sequence of r -tuples of elements of M , where I is an initial segment of \mathbb{N} , and $A \subseteq M$. Then (\vec{b}_i) is a *sequence of indiscernibles over A* if whenever $\varphi(\vec{x}_0, \dots, \vec{x}_n)$ is a formula with parameters from A and the displayed variables and $m_0 < \dots < m_n$, we have

$$\mathfrak{M} \models \varphi(\vec{b}_0, \dots, \vec{b}_n) \Leftrightarrow \varphi(\vec{b}_{m_0}, \dots, \vec{b}_{m_n}).$$

That is, for every n there is a type p_n of $r \cdot (n+1)$ -tuples such that $\text{tp}(\vec{b}_{m_0}, \dots, \vec{b}_{m_n}/A) = p_n$ whenever $m_0 < \dots < m_n$ is an increasing sequence.

7.2. Amalgamation. The arguments in this section are essentially derived from [22].

Lemma 7.6. *Suppose that $\mu(M)$ is finite, (\vec{b}_i) is a sequence of indiscernibles over $A \subseteq M$, X is a definable set over A , and $\mu(X_{\vec{b}_0}) > 0$ holds. Then for any n , $\mu(\bigcap_{i \leq n} X_{\vec{b}_i}) > 0$.*

Proof. Suppose the conclusion of the lemma is false and take $k > 0$ minimal so that $\mu(\bigcap_{i \leq k} X_{\vec{b}_i}) = 0$. For $n \geq k - 1$, let $C_n = \bigcap_{i < k-1} X_{\vec{b}_i} \cap X_{\vec{b}_n}$. Then $\mu(C_n) = \mu(\bigcap_{i \leq k-1} X_{\vec{b}_i}) > 0$ by indiscernibility. However, when $k - 1 \leq n < m$, by indiscernibility again, $\mu(C_n \cap C_m) = \mu(\bigcap_{i \leq k} X_{\vec{b}_i}) = 0$. This is a contradiction since μ is a finite measure. \square

Lemma 7.7. *Suppose that \mathfrak{M} is \aleph_1 -saturated, $\mu(M)$ is finite, and $A \subseteq M$ is countable. If (\vec{b}_i) is a sequence of indiscernibles over A and $\text{tp}(\vec{a}/A \cup \{\vec{b}_0\})$ is wide, then for any n , there is an \vec{a}' with $\text{tp}(\vec{a}', \vec{b}_i/A) = \text{tp}(\vec{a}, \vec{b}_0/A)$ for all $i \leq n$.*

Proof. Let X be a set definable over A and suppose $(\vec{a}, \vec{b}_0) \in X$. Then $\vec{a} \in X_{\vec{b}_0}$ and, since $\text{tp}(\vec{a}/A \cup \{\vec{b}_0\})$ is wide, $\mu(X_{\vec{b}_0}) > 0$. By the previous lemma, $\mu(\bigcap_{i \leq n} X_{\vec{b}_i}) > 0$, and is therefore non-empty. By \aleph_1 -saturation, there is an \vec{a}' such that $\vec{a}' \in \bigcap_{i \leq n} X_{\vec{b}_i}$ simultaneously for all $X \ni (\vec{a}, \vec{b}_0)$ that are definable over A . Therefore $\text{tp}(\vec{a}', \vec{b}_i/A) = \text{tp}(\vec{a}, \vec{b}_0/A)$ for all $i \leq n$. \square

Remark 7.8. The previous lemma suggests a connection between wideness and the model-theoretic notion of *nonforking*, as we now explain. First, a formula $\varphi(x, b)$ *k-divides* over a set $A \subseteq M$ if there is a sequence $(b_i : i \in \mathbb{N})$ which is indiscernible over A for which $b_0 = b$ and such that $\bigcap_{i \in I} \varphi(M, b_i) = \emptyset$ for any $I \subseteq \mathbb{N}$ with $|I| = k$. We say that $\varphi(x, b)$ *divides* over A if it *k-divides* over A for some $k \geq 2$. We say that $\varphi(x, b)$ *forks* over A if there are formulae $\psi_1(x, b_1), \dots, \psi_n(x, b_n)$ such that $\varphi(M, b) \subseteq \bigcup_{i=1}^n \psi_i(M, b_i)$ and such that each $\psi_i(x, b_i)$ divides over A . Finally, we say that a type $p(x)$ forks over A if it contains a formula which forks over A . In model theory, when $\text{tp}(a/A \cup \{b\})$ does not fork over A , this implies that a is in some sense independent from b over A . For example, when working in an algebraically closed field K and k is a subfield of K , then $\text{tp}(a/k \cup \{b\})$ does not fork over k if and only if, whenever a is a generic point for a Zariski closed set V defined over k , then a does not lie in any Zariski closed set V' defined over $k(b)$ with $\dim(V') < \dim(V)$. Lemma 7.7 shows that if $A \subseteq M$ is countable and $\text{tp}(a/A \cup \{b\})$ is wide, then $\text{tp}(a/A \cup \{b\})$ does not fork over A .

Definition 7.9. Suppose that $A \subseteq M$ and p is a global type.

- (1) p is *A-invariant* if whenever $\text{tp}(\vec{a}_1, \dots, \vec{a}_n/A) = \text{tp}(\vec{a}'_1, \dots, \vec{a}'_n/A)$, we have $X_{\vec{a}_1, \dots, \vec{a}_n} \in p \Leftrightarrow X_{\vec{a}'_1, \dots, \vec{a}'_n} \in p$.
- (2) p is *A-finitely satisfiable* if whenever $X \in p$, $X \cap A \neq \emptyset$.

Lemma 7.10. *Suppose that p is a global type and $A \subseteq M$. If p is A-finitely satisfiable then p is A-invariant.*

Proof. Suppose not; then there are definable X and $\vec{a}_1, \dots, \vec{a}_n, \vec{a}'_1, \dots, \vec{a}'_n$ so that $X_{\vec{a}_1, \dots, \vec{a}_n} \in p$ and $X_{\vec{a}'_1, \dots, \vec{a}'_n} \notin p$. Since p is a global type, this means $\neg X_{\vec{a}'_1, \dots, \vec{a}'_n} \in p$, and since p is closed under intersections, $X_{\vec{a}_1, \dots, \vec{a}_n} \setminus X_{\vec{a}'_1, \dots, \vec{a}'_n} \in p$. But then there is an element $m \in A$ such that $m \in X_{\vec{a}_1, \dots, \vec{a}_n} \setminus X_{\vec{a}'_1, \dots, \vec{a}'_n}$, contradicting the fact that $\text{tp}(\vec{a}_1, \dots, \vec{a}_n/A) = \text{tp}(\vec{a}'_1, \dots, \vec{a}'_n/A)$. \square

Lemma 7.11. *Let \mathfrak{N} be an elementary substructure of \mathfrak{M} . If p is a type over N , then there is an extension of p to a global N -finitely satisfiable (and therefore N -invariant) type.*

Proof. Note that, since \mathfrak{M} is an elementary submodel of \mathfrak{M} , each $X \in p$ has the property that $X \cap N^n \neq \emptyset$. Let

$$p' = p \cup \{\neg Y : Y \in \text{Def}_n(M) \text{ and } Y \cap N^n = \emptyset\}.$$

We will show that every finite subset of p' is satisfied by an element of N . If not, there is an $X \in p$ and Y_1, \dots, Y_n with $Y_i \cap N = \emptyset$ such that $X \cap \bigcap_i \neg Y_i = \emptyset$. So $X \cap N \cap \bigcap_i (Y_i \cup \neg Y_i) \neq \emptyset$, and since each $Y_i \cap N = \emptyset$, $X \cap M \cap \bigcap_i Y_i \neq \emptyset$.

Let q be an arbitrary extension of p' to a (global) type. Suppose q is not finitely satisfiable in M ; then there is a $Y \in q$ with $Y \cap N = \emptyset$, which is impossible since then $\neg Y \in p' \subseteq q$. \square

Definition 7.12. If S is a set and p a global type, write $p \upharpoonright S$ for the restriction of p to sets definable over S .

Lemma 7.13. *Suppose p is a global A -invariant type, and recursively choose $\vec{b}_n \models p \upharpoonright A \cup \{\vec{b}_i\}_{i < n}$. Then (\vec{b}_i) (if it exists) is a sequence of indiscernibles over A . (This sequence always exists if \mathfrak{M} is \aleph_1 -saturated.)*

Proof. By induction on n , we show $\text{tp}(\vec{b}_{m_0}, \dots, \vec{b}_{m_n}/A)$ is constant whenever $m_0 < \dots < m_n$. When $n = 1$, this follows since each $\text{tp}(\vec{b}_i/A) = p \upharpoonright A$. Suppose now that $(\vec{b}_{n+1}, \vec{b}_n, \dots, \vec{b}_0) \in X$. Then also $(\vec{b}_{m_{n+1}}, \vec{b}_n, \dots, \vec{b}_0) \in X$ holds because $\text{tp}(\vec{b}_{m_{n+1}}/A \cup \{\vec{b}_0, \dots, \vec{b}_n\}) = \text{tp}(\vec{b}_{n+1}/A \cup \{\vec{b}_0, \dots, \vec{b}_n\}) = p \upharpoonright A \cup \{\vec{b}_0, \dots, \vec{b}_n\}$. By the induction hypothesis, we have that $\text{tp}(\vec{b}_0, \dots, \vec{b}_n/A) = \text{tp}(\vec{b}_{m_0}, \dots, \vec{b}_{m_n}/A)$. Since $X_{\vec{b}_n, \dots, \vec{b}_0} \in p$, also $X_{\vec{b}_{m_n}, \dots, \vec{b}_{m_0}} \in p$, whence we have $(\vec{b}_{m_{n+1}}, \vec{b}_{m_n}, \dots, \vec{b}_{m_0}) \in X$. \square

Lemma 7.14. *Suppose that \mathfrak{M} is \aleph_1 -saturated and $\mu(M)$ is finite. Suppose that $A \subseteq M$ is countable, q is an A -invariant type, $\vec{b} \models q \upharpoonright A \cup \{\vec{a}\}$, $\text{tp}(\vec{b}'/A) = \text{tp}(\vec{b}/A)$, $\text{tp}(\vec{a}'/A) = \text{tp}(\vec{a}/A)$, $\text{tp}(\vec{a}'/A \cup \{\vec{b}'\})$ is wide, X, Y are definable with parameters from A , and $\mu(X_{\vec{a}} \cap Y_{\vec{b}}) > 0$. Then $\mu(X_{\vec{a}'} \cap Y_{\vec{b}'}) > 0$.*

Proof. Suppose the claim fails, so $\mu(X_{\vec{a}'} \cap Y_{\vec{b}'}) = 0$. By \aleph_1 -saturation and the fact that $\text{tp}(\vec{b}'/A) = \text{tp}(\vec{b}/A)$, there an \vec{a}'' with $\text{tp}(\vec{a}'', \vec{b}/A) = \text{tp}(\vec{a}', \vec{b}/A)$; in this way, we may assume $\vec{b} = \vec{b}'$.

Since q is A -invariant, whenever $\vec{a}^* \models \text{tp}(\vec{a}/A)$ and $\vec{b}^* \models q \upharpoonright A \cup \{\vec{a}^*\}$ we have $\text{tp}(\vec{a}^*, \vec{b}^*/A) = \text{tp}(\vec{a}, \vec{b}/A)$, and so in particular $\mu(X_{\vec{a}^*} \cap Y_{\vec{b}^*}) = \mu(X_{\vec{a}} \cap Y_{\vec{b}}) > 0$.

Choose $\vec{a}_0 \models p$ and $\vec{b}_0 \models q \upharpoonright A \cup \{\vec{a}_0\}$. Let $\{\vec{a}_i, \vec{b}_i\}_{i \leq n}$ be given so that (\vec{b}_i) is indiscernible over A . For each $\delta > 0$ and each $C \in \text{tp}(\vec{a}/A)$, by the wideness of $\text{tp}(\vec{a}'/A \cup \{\vec{b}\})$, we have

$$\mu(C \cap \{\vec{x} : \mathfrak{M} \models (m_{\vec{z}} < \delta \cdot X(\vec{z}, \vec{x}) \wedge Y(\vec{z}, \vec{b}))\}) > 0,$$

and therefore by Lemma 7.7 and \aleph_1 -saturation we may find an \vec{a}_{n+1} with $\text{tp}(\vec{a}_{n+1}/A) = \text{tp}(\vec{a}/A)$ such that $\mu(X_{\vec{a}_{n+1}} \cap Y_{\vec{b}_i}) = 0$ for $i \leq n$. Let $\vec{b}_{n+1} \models q \upharpoonright (A \cup \{\vec{a}_i, \vec{b}_i\}_{i \leq n} \cup \{\vec{a}_{n+1}\})$; by Lemma 7.13, we have ensured that (\vec{b}_i) remains indiscernible over A . For each i , let $C_i = X_{\vec{a}_i} \cap Y_{\vec{b}_i}$. Since each $\vec{b}_n \models q \upharpoonright \vec{a}_n$, $\mu(C_i) \geq \epsilon$; while for $i \neq j$, $\mu(C_i \cap C_j) = 0$. This is a contradiction to the finiteness of $\mu(M)$. \square

Lemma 7.15. *Suppose that \mathfrak{M} is \aleph_1 -saturated and $\mu(M)$ is finite. Suppose that $A \subseteq M$ is countable, q is an A -invariant type, $\vec{b} \models q \upharpoonright A \cup \{\vec{a}\}$, $\text{tp}(\vec{b}'/A) = \text{tp}(\vec{b}/A)$, $\text{tp}(\vec{a}'/A) = \text{tp}(\vec{a}/A)$, $\text{tp}(\vec{a}'/A \cup \{\vec{b}'\})$ is wide, X, Y are definable with parameters from A , and $\mu(X_{\vec{a}} \cap Y_{\vec{b}}) = 0$. Then $\mu(X_{\vec{a}'} \cap Y_{\vec{b}'}) = 0$.*

Proof. Suppose the claim fails, so $\mu(X_{\vec{a}'} \cap Y_{\vec{b}'}) > 0$. As in the proof of the previous lemma, we may assume that $\vec{b} = \vec{b}'$. Suppose that $\{\vec{a}_i, \vec{b}_i\}_{i < n}$ has been constructed (including the empty case where $n = 0$) so that $\vec{a}_i \models \text{tp}(\vec{a}/A)$ for each i and so that (\vec{b}_i) is indiscernible over A . Choose $\vec{b}_n \models q \mid A \cup \{\vec{a}_i, \vec{b}_i\}_{i < n}$. There is some $\epsilon \in \mathbb{Q}^{>0}$ such that for each $C \in \text{tp}(\vec{a}/A)$,

$$\mu(C \cap \{\vec{x} : \mathfrak{M} \models (m_{\vec{z}} \geq \epsilon \cdot X(\vec{z}, \vec{x}) \wedge Y(\vec{z}, \vec{b}))\}) > 0,$$

and therefore by Lemma 7.7 and \aleph_1 -saturation we may find an \vec{a}_n with $\text{tp}(\vec{a}_n/A) = \text{tp}(\vec{a}/A)$ such that $\mu(X_{\vec{a}_n} \cap Y_{\vec{b}_i}) \geq \epsilon$ for $i \leq n$. For each i , let $C_i = X_{\vec{a}_i} \cap Y_{\vec{b}_i}$. Then $\mu(C_i) \geq \epsilon$ for each i and, when $i < j$, $\mu(C_i \cap C_j) = 0$ since $\vec{b}_j \models q \mid A \cup \{\vec{a}_i\}$. Thus $\mu(X_{\vec{a}_i} \cap Y_{\vec{b}_j}) = 0$ by A -invariance. This contradicts the fact that $\mu(M)$ is finite. \square

Theorem 7.16. *Suppose \mathfrak{M} is \aleph_1 -saturated and $\mu(M)$ is finite. Let \mathfrak{N} be a countable elementary substructure of \mathfrak{M} . Let p, q be types over N and suppose $\text{tp}(\vec{a}/N) = \text{tp}(\vec{a}'/N) = p$, $\text{tp}(\vec{b}/N) = \text{tp}(\vec{b}'/N) = q$ with $\text{tp}(\vec{a}/N \cup \{\vec{b}\})$, $\text{tp}(\vec{a}'/N \cup \{\vec{b}'\})$ wide. Then for any definable sets X and Y over N , we have $\mu(X_{\vec{a}} \cap Y_{\vec{b}}) > 0$ iff $\mu(X_{\vec{a}'} \cap Y_{\vec{b}'}) > 0$.*

Proof. Fix an extension q' of q to an N -invariant global type; this is possible by Lemma 7.11. Let $\vec{a}^* \models p$ and $\vec{b}^* \models q' \mid N \cup \{\vec{a}^*\}$. If $\mu(X_{\vec{a}^*} \cap Y_{\vec{b}^*}) > 0$ then by Lemma 7.14, both $\mu(X_{\vec{a}} \cap Y_{\vec{b}}) > 0$ and $\mu(X_{\vec{a}'} \cap Y_{\vec{b}'}) > 0$. Otherwise, $\mu(X_{\vec{a}^*} \cap Y_{\vec{b}^*}) = 0$, and by Lemma 7.15, both $\mu(X_{\vec{a}} \cap Y_{\vec{b}}) = 0$ and $\mu(X_{\vec{a}'} \cap Y_{\vec{b}'}) = 0$. \square

Definition 7.17. Let \mathfrak{N} be an elementary substructure of \mathfrak{M} . A set $R \subseteq M^{2m}$ is *wide-stable* over \mathfrak{N} if R is N -invariant and whenever $\text{tp}(\vec{a}/N) = \text{tp}(\vec{a}'/N)$, $\text{tp}(\vec{b}/N) = \text{tp}(\vec{b}'/N)$, both $\text{tp}(\vec{a}/N \cup \{\vec{b}\})$ and $\text{tp}(\vec{a}'/N \cup \{\vec{b}'\})$ are wide, and $(\vec{a}, \vec{b}) \in R$, then also $(\vec{a}', \vec{b}') \in R$.

R is *stable* over \mathfrak{N} if whenever (\vec{a}_i, \vec{b}_i) is a sequence of indiscernibles over N and $R(\vec{a}_i, \vec{b}_j)$ holds for every $i < j$, also $R(\vec{a}_j, \vec{b}_i)$ for some $i < j$.

We say $R \in \mathcal{B}_2(N)$ is *approximated by definable stable sets (ADS)* (respectively, *approximated by definable wide-stable sets (ADWS)*) over N if for every $\epsilon > 0$, there is a stable (respectively, wide-stable) set S definable over N such that $\mu(R \Delta S) < \epsilon$.

Observe that the wide-stable subsets of M^{2m} form a σ -algebra. It can be shown that stable subsets of M^{2m} form a Boolean algebra. Note that if we replace wideness in the definition of wide-stability with non-forking, we obtain a standard consequence of stability (see Lemma 3.3 of [25]). In particular, since wide types are non-forking, stable implies wide-stable.

Theorem 7.18. *Let \mathfrak{M} be \aleph_1 -saturated, let \mathfrak{N} be a countable elementary substructure of \mathfrak{M} , and suppose μ is finite. If $R \subseteq M^2$ is in $\mathcal{B}_2(N)$, then the following are equivalent:*

- (1) *There is $U \in \mathcal{B}_{2,1}(N)$ such that $\mu(R \Delta U) = 0$;*
- (2) *R is ADS;*
- (3) *R is ADWS.*

Proof. (1) \Rightarrow (2): First observe that every set of the form $A \times B$ with A, B definable over N , is stable: if $(a_i, b_i : i \in I)$ is a sequence of indiscernibles over N and $(a_i, b_j) \in A \times B$ for $i < j$ then $(a_j, b_i) \in A \times B$ for all $i < j$. Consequently,

the algebra generated by the sets of the form $A \times B$, with A, B definable over N , which is dense in $\mathcal{B}_{2,1}(N)$, is contained in the algebra of stable subsets of M^2 . Thus, if there is $U \in \mathcal{B}_{2,1}(N)$ such that $\mu(R \Delta U) = 0$, then for every $\epsilon > 0$, there is $S \in \mathcal{B}_{2,1}^0(N)$ such that $\mu(R \Delta S) < \epsilon$; such an S is stable.

(2) \Rightarrow (3): Immediate since stable definable sets are wide-stable and definable.

(3) \Rightarrow (1): Let $R \subseteq M^2$ be definable over N and wide-stable over N . Since $\mathcal{B}_{2,1}(N)$ is complete with respect to the pseudometric $\mu(\cdot \Delta \cdot)$, it suffices to prove (1) for such an R . For any N -definable A and B and any $\epsilon > 0$, define $A_\epsilon^B = \{a \in A : \mathfrak{M} \models m_y < \epsilon \cdot B \setminus R_a\}$. Set

$$\mathcal{U} = \{A \times B : \mu((A \times B) \setminus R) = 0, A, B \in \mathcal{B}_1(N)\} \cup \left\{ \bigcap_n (A_{1/n}^B \times B) : A, B \in \mathcal{B}_1(N) \right\}.$$

Clearly $U = \bigcup \mathcal{U}$ belongs to $\mathcal{B}_{2,1}(N)$ and satisfies $\mu(U \setminus R) = 0$.

Suppose, towards a contradiction, that $\mu(R \setminus U) > 0$. Then we may find an $(a, b) \in R \setminus U$ such that $\text{tp}(a/N)$ and $\text{tp}(b/N \cup \{a\})$ are wide. For each $A \in \text{tp}(a/N)$ and $B \in \text{tp}(b/N)$, we have $A \times B \notin \mathcal{U}$, and therefore $\mu((A \times B) \setminus R) > 0$. Therefore $p = (\text{tp}(a/N) \times \text{tp}(b/N)) \cup \{\neg R(x, y)\}$ is a wide partial type. Set

$$p' = p \cup \{\neg S(x, y) : (\forall \epsilon \in \mathbb{Q}^{>0})(m_y < \epsilon \cdot S(x, y) \in p)\}.$$

We claim this is also a wide partial type. Towards this end, fix A and B definable over N such that $a \in A$ and $b \in B$. Also fix $S(x, y)$ definable over N so that, for each $\epsilon \in \mathbb{Q}^{>0}$, $m_y < \epsilon \cdot S(x, y) \in p$. Since $(a, b) \notin U$, we have some $\epsilon \in \mathbb{Q}^{>0}$ such that $A_\epsilon^B \notin \text{tp}(a/N)$. Set

$$A' = \{a' \in A : \mathfrak{M} \models (m_y \geq \epsilon \cdot (B \setminus R_{a'})) \wedge (m_y < \epsilon \cdot S(a', y))\}.$$

Observe that $A' \in \text{tp}(a/N)$, whence $\mu(A') > 0$. It is enough to show that $\mu((A' \times B) \cap (\neg R \cap \neg S)) > 0$. Suppose this is not the case. Then

$$\mu((A' \times B) \setminus R) = \mu((A' \times B) \cap (S \setminus R)).$$

However,

$$\mu((A' \times B) \setminus R) = \int_{a' \in A'} \mu(B \setminus R_{a'}) d\mu \geq \epsilon \mu(A')$$

while

$$\mu((A' \times B) \cap (S \setminus R)) = \int_{a' \in A'} \mu(B \cap (S_{a'} \setminus R_{a'})) d\mu < \epsilon \mu(A'),$$

yielding a contradiction.

Now take any extension of p' to a wide type q and suppose that (a', b') realizes q . It follows immediately that $\text{tp}(b'/N \cup \{a'\})$ is wide, $\text{tp}(a'/N) = \text{tp}(a/N)$, and $\text{tp}(b'/N) = \text{tp}(b/N)$. Since $\text{tp}(b/N \cup \{a\})$ is wide and $\text{tp}(a'/N) = \text{tp}(a/N)$, whenever $B_a \in \text{tp}(b/N \cup \{a\})$, $\mu(B_{a'}) > 0$. By \aleph_1 -saturation, there is a b'' with $\text{tp}(b''/N \cup \{a'\}) = \text{tp}(b/N \cup \{a\})$. But now we have $(a', b') \notin R$ and $(a', b'') \in R$, contradicting wide-stability. \square

We could generalize these notions to stable relations on pairs of n -tuples, by using the two algebras $\mathcal{B}_{2n, [1, n]}$ and $\mathcal{B}_{2n, [n+1, 2n]}$ (the algebras of sets of $2n$ -tuples depending only on the first n or only on the second n coordinates, respectively); the arguments are exactly analogous, replacing the singletons with n -tuples.

Combining Theorem 7.16 and the proof of the previous theorem, we see:

Corollary 7.19. *Suppose \mathfrak{M} is \aleph_1 -saturated and $\mu(M)$ is finite. Suppose that \mathfrak{N} is a countable elementary substructure of \mathfrak{M} and that X and Y are definable over N . Then $\{(a, b) : \mu(X_a \cap Y_b) > 0\}$ differs from an element of $\mathcal{B}_{2,1}$ by a null set.*

7.3. Hypergraph Removal. We now give a proof of the hypergraph removal lemma [11, 17, 27]; this proof is closely related to the infinitary proof given by Tao [30] in a different framework.

Suppose \mathcal{B} is the σ -algebra generated by some collection of formulae with parameters from an elementary substructure \mathfrak{M} of \mathfrak{N} . Two natural ways to extend \mathcal{B} would be by extending the collection of formulae and by extending the set M . The following lemma states that these two methods are orthogonal in a certain sense. This lemma will be applied in the particular case of the algebras $\mathcal{B}_{n,<I}(M)$, $\mathcal{B}_{n,<I}(M \cup \{a\})$ (to be defined below), and $\mathcal{B}_{n,I}(M)$ (as defined in Definition 5.4), and the reader will not be misled assuming these are the algebras used. (Recall our notation that \mathcal{B}^0 is the Boolean algebra of definable sets generating the algebra \mathcal{B} .)

Lemma 7.20. *Let Φ be a set of formulae, let \mathfrak{M} be an elementary substructure of \mathfrak{N} , and let n be an integer. Suppose \mathcal{A}^0 is the collection of sets of the form $\varphi(N^n, \vec{a})$ where $\varphi(\vec{x}, \vec{y}) \in \Phi$ and $\vec{a} \in M^{|\vec{y}|}$, \mathcal{B}^0 is the set of M -definable sets of n -tuples in \mathfrak{N} , and \mathcal{C}^0 is the collection of sets of the form $\varphi(N^n, \vec{a})$ where $\varphi(\vec{x}, \vec{y}) \in \Phi$ and $\vec{a} \in N^{|\vec{y}|}$.*

Then for any $f \in L^2(\mathcal{B})$, $\|\mathbb{E}(f \mid \mathcal{A}) - \mathbb{E}(f \mid \mathcal{C})\|_{L^2} = 0$.

*Proof.*³ Suppose not. Then setting $\epsilon := \|f - \mathbb{E}(f \mid \mathcal{A})\|_{L^2}$ and $\delta := \|f - \mathbb{E}(f \mid \mathcal{C})\|_{L^2}$, we must have $\delta < \epsilon$. Since f is \mathcal{B} -measurable, for some β_1, \dots, β_m , some ψ_1, \dots, ψ_m , and some $\vec{b}_1, \dots, \vec{b}_m$ in M , we have $\|f - \sum_{i \leq m} \beta_i \chi_{\psi_i}(\vec{x}, \vec{b}_i)\|_{L^2} < (\epsilon - \delta)/4$.

Since $\|f - \mathbb{E}(f \mid \mathcal{C})\|_{L^2} = \delta$, there are $\alpha_1, \dots, \alpha_k$, formulae $\varphi_1, \dots, \varphi_k \in \Phi$, and $\vec{d}_1, \dots, \vec{d}_k$ in N such that $\|f - \sum_{i \leq k} \alpha_i \chi_{\varphi_i}(\vec{x}, \vec{d}_i)\|_{L^2} < \epsilon - 3(\epsilon - \delta)/4$, and therefore

$$\left\| \sum_{i \leq m} \beta_i \chi_{\psi_i}(\vec{x}, \vec{b}_i) - \sum_{i \leq k} \alpha_i \chi_{\varphi_i}(\vec{x}, \vec{d}_i) \right\|_{L^2} < \epsilon - (\epsilon - \delta)/2.$$

This is not itself a formula, but we can construct a formula which corresponds to it.

Squaring both sides and expanding the product, we have

$$\begin{aligned} (\epsilon - (\epsilon - \delta)/2)^2 &> \int \sum_{i,j \leq m} \beta_i \beta_j \chi_{\psi_i}(\vec{x}, \vec{b}_i) \chi_{\psi_j}(\vec{x}, \vec{b}_j) \\ &\quad - 2 \int \sum_{i \leq m, j \leq k} \beta_i \alpha_j \chi_{\psi_i}(\vec{x}, \vec{b}_i) \chi_{\varphi_j}(\vec{x}, \vec{d}_j) \\ &\quad + \int \sum_{i,j \leq k} \alpha_i \alpha_j \chi_{\varphi_i}(\vec{x}, \vec{d}_i) \chi_{\varphi_j}(\vec{x}, \vec{d}_j). \end{aligned}$$

For each $i \leq m, j \leq k$, we may choose an r_{ij} such that $\int \chi_{\psi_i}(\vec{x}, \vec{b}_i) \chi_{\varphi_j}(\vec{x}, \vec{d}_j) d\mu < r_{ij}$ and for each $i, j \leq k$ we may choose an s_{ij} such that $\int \chi_{\varphi_i}(\vec{x}, \vec{d}_i) \chi_{\varphi_j}(\vec{x}, \vec{d}_j) d\mu < s_{ij}$, and we may choose these so that

$$(\epsilon - (\epsilon - \delta)/2)^2 > \int \sum_{i,j \leq m} \beta_i \beta_j \chi_{\psi_i}(\vec{x}, \vec{b}_i) \chi_{\psi_j}(\vec{x}, \vec{b}_j) - 2 \sum_{i \leq m, j \leq k} r_{ij} \beta_i \alpha_j + \sum_{i,j \leq k} s_{ij} \alpha_i \alpha_j.$$

³Hrushovski has pointed out that this proof is reminiscent of early proofs in the stability theory of algebraically closed fields. The full connection between this argument and the amalgamation methods from the previous section is not completely understood.

Now let $\rho(\vec{b}_1, \dots, \vec{b}_m, \vec{d}_1, \dots, \vec{d}_k)$ be the formula

$$\bigwedge_{i \leq m, j \leq k} m_{\vec{x}} < r_{ij} \cdot \varphi_i(\vec{x}, \vec{b}_i) \wedge \psi_j(\vec{x}, \vec{d}_j) \wedge \bigwedge_{i, j \leq k} m_{\vec{x}} < r_{ij} \cdot \psi_i(\vec{x}, \vec{d}_i) \wedge \psi_j(\vec{x}, \vec{d}_j).$$

Since $\mathfrak{N} \models \rho(\vec{b}_1, \dots, \vec{b}_m, \vec{d}_1, \dots, \vec{d}_k)$, we have $\mathfrak{N} \models \exists \vec{y}_1, \dots, \vec{y}_k \rho(\vec{b}_1, \dots, \vec{b}_m, \vec{y}_1, \dots, \vec{y}_k)$.

By the elementarity of \mathfrak{M} , there exist $\vec{d}_1^i, \dots, \vec{d}_k^i$ in M such that $\mathfrak{M} \models \rho(\vec{b}_1, \dots, \vec{b}_m, \vec{d}_1^i, \dots, \vec{d}_k^i)$.

Consequently, for each $i \leq m, j \leq k$, we have $\int \chi_{\psi_i}(\vec{x}, \vec{b}_i) \chi_{\varphi_j}(\vec{x}, \vec{d}_j^i) d\mu \leq r_{ij}$ and

for each $i, j \leq k$, we have $\int \chi_{\varphi_i}(\vec{x}, \vec{d}_i^i) \chi_{\varphi_j}(\vec{x}, \vec{d}_j^i) d\mu \leq s_{ij}$. Therefore

$$\left\| \sum_{i \leq m} \beta_i \chi_{\psi_i}(\vec{x}, \vec{b}_i) - \sum_{i \leq k} \alpha_i \chi_{\varphi_i}(\vec{x}, \vec{d}_i^i) \right\|_{L^2} < \epsilon - (\epsilon - \delta)/2$$

and

$$\left\| f - \sum_{i \leq k} \alpha_i \chi_{\varphi_i}(\vec{x}, \vec{d}_i^i) \right\|_{L^2} < \epsilon - (\epsilon - \delta)/4.$$

Since $\sum_{i \leq k} \alpha_i \chi_{\varphi_i}(\vec{x}, \vec{d}_i^i)$ is measurable with respect to \mathcal{A} , this contradicts the assumption that $\|f - E(f \mid \mathcal{A})\|_{L^2} = \epsilon$. \square

The following theorem is essentially the infinitary version of hypergraph removal:

Theorem 7.21. *Let \mathfrak{M} be an elementary substructure of \mathfrak{N} . Let $n \geq k$, $\mathcal{I} \subseteq \binom{[1, n]}{k}$, and suppose that for each $I \in \mathcal{I}$ with $|I| = k$, we have a set $A_I \in \mathcal{B}_{n, I}(M)$, and suppose there is a $\delta > 0$ such that whenever $B_I \in \mathcal{B}_{n, I}^0(M)$ and $\mu^n(A_I \setminus B_I) < \delta$ for all $I \in \mathcal{I}$, $\bigcap_{I \in \mathcal{I}} B_I$ is non-empty. Then $\mu^n(\bigcap_{I \in \mathcal{I}} A_I) > 0$.*

Proof. We proceed by main induction on k . When $k = 1$, the claim is trivial: we must have $\mu(A_I) > 0$ for all I , since otherwise we could take $B_I = \emptyset$; then $\mu^n(\bigcap A_I) = \prod \mu(A_I) > 0$. So we assume that $k > 1$ and that whenever $B_I \in \mathcal{B}_{n, I}(M)$ and $\mu^n(A_I \setminus B_I) < \delta$ for all I , $\bigcap_{I \in \mathcal{I}} B_I$ is non-empty. Throughout this proof, the variables I and I_0 range over elements of \mathcal{I} .

We write $\mathcal{B}_{n, < I}(M)$ for $\bigcup_{J \subsetneq I} \mathcal{B}_{n, J}(M)$.

Claim 1. For any I_0 ,

$$\int (\chi_{A_{I_0}} - E(\chi_{A_{I_0}} \mid \mathcal{B}_{n, < I_0}(M))) \prod_{I \neq I_0} \chi_{A_I} d\mu^n = 0.$$

Proof. When $k = n$, this is trivial since $\prod_{I \neq I_0} \chi_{A_I}$ is an empty product, and therefore equal to 1.

If $k < n$, we have

$$\begin{aligned} & \int (\chi_{A_{I_0}} - E(\chi_{A_{I_0}} \mid \mathcal{B}_{n, < I_0}^\sigma(M))) \prod_{I \neq I_0} \chi_{A_I} d\mu^n \\ &= \int (\chi_{A_{I_0}} - E(\chi_{A_{I_0}} \mid \mathcal{B}_{n, < I_0}^\sigma(M))) \prod_{I \neq I_0} \chi_{A_I} d\mu^k(\{x_i\}_{i \in I_0}) d\mu^{n-k}(\{x_i\}_{i \notin I_0}). \end{aligned}$$

Observe that for any choice of $\{a_i\}_{i \notin I_0}$, $\mathcal{B}_{n, < I_0}(M)$, $\mathcal{B}_{n, < I_0}(M \cup \{a_i\}_{i \notin I_0})$, $\mathcal{B}_{n, I_0}(M)$ satisfy the preceding lemma, so

$$\|E(\chi_{A_{I_0}} \mid \mathcal{B}_{n, < I_0}^\sigma(M)) - E(\chi_{A_{I_0}} \mid \mathcal{B}_{n, < I_0}^\sigma(M \cup \{a_i\}_{i \notin I_0}))\|_{L^2} = 0.$$

The function $\prod_{I \neq I_0} \chi_{A_I}(\{x_i\}_{i \in I_0}, \{a_i\}_{i \notin I_0})$ is measurable with respect to $\mathcal{B}_{n, < I_0}^\sigma(M \cup \{\vec{a}_i\})$. Combining these two facts, we have

$$\begin{aligned} & \int (\chi_{A_{I_0}} - E(\chi_{A_{I_0}} \mid \mathcal{B}_{n, < I_0}^\sigma(M))) \prod_{I \neq I_0} \chi_{A_I} d\mu^k(\{x_i\}_{i \in I_0}) \\ &= \int (\chi_{A_{I_0}} - E(\chi_{A_{I_0}} \mid \mathcal{B}_{n, < I_0}^\sigma(M \cup \{a_i\}_{i \notin I_0}))) \prod_{I \neq I_0} \chi_{A_I} d\mu^k(\{x_i\}) \\ &= 0 \end{aligned}$$

Since this holds for any $\{a_i\}_{i \notin I_0}$, the claim follows by integrating over all choices of $\{a_i\}$. \dashv

We next show that, without loss of generality, we may assume each A_{I_0} belongs to $\mathcal{B}_{n, < I_0}(M)$ by showing that for each $I_0 \in \mathcal{I}$, there is some set $A'_{I_0} \in \mathcal{B}_{n, < I}(M)$ with the property that, if we replace A_{I_0} by A'_{I_0} , the assumptions of the theorem all hold, and such that if we show the conclusion for the modified family of sets, the conclusion also holds for the original family.

Claim 2. For any I_0 , there is an $A'_{I_0} \in \mathcal{B}_{n, < I_0}(M)$ such that:

- Whenever $B_I \in \mathcal{B}_{n, I}^0(M)$ for each I , $\mu^n(A_I \setminus B_I) < \delta$ for each $I \neq I_0$, and $\mu^n(A'_{I_0} \setminus B_{I_0}) < \delta$, $\bigcap_{I \in \mathcal{I}} B_I$ is non-empty, and
- If $\mu^n(A'_{I_0} \cap \bigcap_{I \neq I_0} A_I) > 0$, $\mu^n(\bigcap_{I \in \mathcal{I}} A_I) > 0$.

Proof. Define $A'_{I_0} := \{x_{I_0} \mid \mathbb{E}(\chi_{A_{I_0}} \mid \mathcal{B}_{n, < I_0}(M))(x_{I_0}) > 0\}$. If $\mu^n(A'_{I_0} \cap \bigcap_{I \neq I_0} A_I) > 0$ then we have

$$\int \mathbb{E}(\chi_{A_{I_0}} \mid \mathcal{B}_{n, < I_0}(M)) \prod_{I \neq I_0} \chi_{A_I} d\mu^n > 0,$$

and by the previous claim, this implies that $\mu^n(\bigcap_{I \in \mathcal{I}} A_I) > 0$.

Suppose that for each I , $B_I \in \mathcal{B}_{n, I}^0(M)$ with $\mu^n(A_I \setminus B_I) < \delta$ for $I \neq I_0$ and $\mu^n(A'_{I_0} \setminus B_{I_0}) < \delta$. Since

$$\mu^n(A_{I_0} \setminus A'_{I_0}) = \int \chi_{A_{I_0}} (1 - \chi_{A'_{I_0}}) d\mu^n = \int \mathbb{E}(\chi_{A_{I_0}} \mid \mathcal{B}_{n, < I_0}(M)) (1 - \chi_{A'_{I_0}}) d\mu^n = 0,$$

we have $\mu^n(A_{I_0} \setminus B_{I_0}) < \delta$ as well, and therefore $\bigcap_{I \in \mathcal{I}} B_I$ is non-empty. \dashv

By applying the previous claim to each $I \in \mathcal{I}$, we may assume for the rest of the proof that for each I , $A_I \in \mathcal{B}_{n, < I}(M)$.

Fix some finite algebra $\mathcal{B} \subseteq \mathcal{B}_{n, k-1}^0(M)$ so that for every I , $\|\chi_{A_I} - \mathbb{E}(\chi_{A_I} \mid \mathcal{B})\|_{L^2(\mu^n)} < \frac{\sqrt{\delta}}{\sqrt{2}(|\mathcal{I}|+1)}$ (such a \mathcal{B} exists because there are finitely many I and each A_I is $\mathcal{B}_{n, k-1}(M)$ -measurable). For each I , set $A_I^* = \{a_I \mid \mathbb{E}(\chi_{A_I} \mid \mathcal{B})(a_I) > \frac{|\mathcal{I}|}{|\mathcal{I}|+1}\}$.

Claim 3. For each I , $\mu^n(A_I \setminus A_I^*) \leq \delta/2$

Proof. $A_I \setminus A_I^*$ is the set of points such that $(\chi_{A_I} - \mathbb{E}(\chi_{A_I} \mid \mathcal{B}))(\vec{a}) \geq \frac{1}{|\mathcal{I}|+1}$. By Chebyshev's inequality, the measure of this set is at most

$$(|\mathcal{I}| + 1)^2 \int (\chi_{A_I} - \mathbb{E}(\chi_{A_I} \mid \mathcal{B}))^2 d\mu^n = (|\mathcal{I}| + 1)^2 \|\chi_{A_I} - \mathbb{E}(\chi_{A_I} \mid \mathcal{B})\|_{L^2(\mu^n)}^2 \leq \frac{\delta}{2}.$$

\dashv

Claim 4. $\mu^n(\bigcap_{I \in \mathcal{I}} A_I) \geq \mu^n(\bigcap_{I \in \mathcal{I}} A_I^*) / (|\mathcal{I}| + 1)$.

Proof. For each I_0 ,

$$\begin{aligned} \mu^n((A_{I_0}^* \setminus A_{I_0}) \cap \bigcap_{I \neq I_0} A_I^*) &= \int \chi_{A_{I_0}^*} (1 - \chi_{A_{I_0}}) \prod_{I \neq I_0} \chi_{A_I^*} d\mu^n \\ &= \int \chi_{A_{I_0}^*} (1 - \mathbb{E}(\chi_{A_{I_0}} \mid \mathcal{B})) \prod_{I \neq I_0} \chi_{A_I^*} d\mu^n \\ &\leq \frac{1}{|\mathcal{I}| + 1} \int \prod_{I \in \mathcal{I}} \chi_{A_I^*} d\mu^n \\ &= \frac{1}{|\mathcal{I}| + 1} \mu^n\left(\bigcap_{I \in \mathcal{I}} A_I^*\right) \end{aligned}$$

But then

$$\mu^n\left(\bigcap_{I \in \mathcal{I}} A_I^* \setminus \bigcap_{I \in \mathcal{I}} A_I\right) \leq \sum_{I_0} \mu^n((A_{I_0}^* \setminus A_{I_0}) \cap \bigcap_{I \neq I_0} A_I^*) \leq \frac{|\mathcal{I}|}{|\mathcal{I}| + 1} \mu^n\left(\bigcap_{I \in \mathcal{I}} A_I^*\right).$$

—

Each A_I^* may be written in the form $\bigcup_{i \leq r_I} A_{I,i}^*$ where $A_{I,i}^* = \bigcap_{J \in \binom{I}{k-1}} A_{I,i,J}^*$ and $A_{I,i,J}^*$ is an element of $\mathcal{B}_{n,J}^0(M)$. We may assume that if $i \neq i'$ then $A_{I,i}^* \cap A_{I,i'}^* = \emptyset$.

We have

$$\mu^n\left(\bigcap_I A_I^*\right) = \mu^n\left(\bigcup_{\vec{i} \in \prod_I [1, r_I]} \bigcap_I \bigcap_{J \in \binom{I}{k-1}} A_{I,i_I,J}^*\right).$$

For each $\vec{i} \in \prod_I [1, r_I]$, let $D_{\vec{i}} = \bigcap_I \bigcap_{J \in \binom{I}{k-1}} A_{I,i_I,J}^*$. Each $A_{I,i_I,J}^*$ is an element of $\mathcal{B}_{n,J}^0(M)$, so we may group the components and write $D_{\vec{i}} = \bigcap_{J \in \binom{[1,n]}{k-1}} D_{\vec{i},J}$ where $D_{\vec{i},J} = \bigcap_{I \supset J} A_{I,i_I,J}^*$.

Suppose, for a contradiction, that $\mu^n(\bigcap_I A_I^*) = 0$. Then for every $\vec{i} \in \prod_I [1, r_I]$, $\mu^n(D_{\vec{i}}) = \mu^n(\bigcap_J D_{\vec{i},J}) = 0$. By the inductive hypothesis, for each $\gamma > 0$, there is a collection $B_{\vec{i},J} \in \mathcal{B}_{n,J}^0(M)$ such that $\mu^n(D_{\vec{i},J} \setminus B_{\vec{i},J}) < \gamma$ and $\bigcap_J B_{\vec{i},J} = \emptyset$. In particular, this holds with $\gamma = \frac{\delta}{2 \binom{k}{k-1} (\prod_I r_I) (\max_I r_I)}$.

For each $I, i \leq r_I, J \subset I$, define

$$B_{I,i,J}^* = A_{I,i,J}^* \cap \bigcap_{\vec{i}, i_I=i} \left[B_{\vec{i},J} \cup \bigcup_{I' \supseteq J, I' \neq I} \overline{A_{I',i_{I'},J}^*} \right].$$

Claim 5. $\mu^n(A_{I,i,J}^* \setminus B_{I,i,J}^*) \leq \frac{\delta}{2 \binom{k}{k-1} (\max_I r_I)}$.

Proof. Observe that if $x \in A_{I,i,J}^* \setminus B_{I,i,J}^*$ then for some \vec{i} with $i_I = i$, $x \notin B_{\vec{i},J} \cup \bigcup_{I' \supseteq J, I' \neq I} \overline{A_{I',i_{I'},J}^*}$. This means $x \notin B_{\vec{i},J}$ and $x \in \bigcap_{I' \supseteq J} A_{I',i_{I'},J}^* = D_{\vec{i},J}$. So

$$\mu^n(A_{I,i,J}^* \setminus B_{I,i,J}^*) \leq \sum_{\vec{i} \in \prod_I [1, r_I]} \mu^n(D_{\vec{i},J} \setminus B_{\vec{i},J}) \leq \frac{\delta}{2 \binom{k}{k-1} (\max_I r_I)}.$$

—

Define $B_I^* = \bigcup_{i \leq r_I} \bigcap_J B_{I,i,J}^*$.

Claim 6. $\mu^n(A_I \setminus B_I^*) \leq \delta$.

Proof. Since $\mu^n(A_I \setminus A_I^*) \leq \delta/2$, it suffices to show that $\mu^n(A_I^* \setminus B_I^*) \leq \delta/2$.

$$\begin{aligned}
\mu^n(A_I^* \setminus \bigcup_i \bigcap_J B_{I,i,J}^*) &= \mu^n\left(\bigcup_i \bigcap_J A_{I,i,J}^* \setminus \bigcup_i \bigcap_J B_{I,i,J}^*\right) \\
&\leq \mu^n\left(\bigcup_i \left(\bigcap_J A_{I,i,J}^* \setminus \bigcap_J B_{I,i,J}^*\right)\right) \\
&\leq \sum_{i \leq r_I} \mu^n\left(\bigcap_J A_{I,i,J}^* \setminus \bigcap_J B_{I,i,J}^*\right) \\
&\leq \sum_{i \leq r_I} \sum_J \mu^n(A_{I,i,J}^* \setminus B_{I,i,J}^*) \\
&\leq r_I \cdot \binom{k}{k-1} \cdot \frac{\delta}{2 \binom{k}{k-1} (\max_I r_I)} \\
&\leq \delta/2.
\end{aligned}$$

+

Note that the sets B_I^* satisfy the assumption, and therefore $\bigcap_I B_I^* \neq \emptyset$.

Claim 7.

$$\bigcap_I B_I^* \subseteq \bigcup_{\vec{i}} \bigcap_J B_{\vec{i},J}^*.$$

Proof. Suppose $x \in \bigcap_I B_I^* = \bigcap_I \bigcup_{i \leq r_I} \bigcap_J B_{I,i,J}^*$. Then for each I , there is an $i_I \leq r_I$ such that $x \in \bigcap_J B_{I,i_I,J}^*$. Since $B_{I,i_I,J}^* \subseteq A_{I,i_I,J}^*$, for each I and $J \subset I$, $x \in A_{I,i_I,J}^*$.

For any J , let $I \supset J$. Then

$$x \in B_{I,i_I,J}^* = A_{I,i_I,J}^* \cap \bigcap_{\vec{i}', i'_I = i_I} \left[B_{\vec{i}',J}^* \cup \bigcup_{I' \supseteq J, I' \neq I} \overline{A_{I',i'_I,J}^*} \right].$$

In particular, $x \in \left[B_{\vec{i},J}^* \cup \bigcup_{I' \supseteq J, I' \neq I} \overline{A_{I',i'_I,J}^*} \right]$ for the particular \vec{i} we have chosen. Since $x \in A_{I,i_I,J}^*$ for each $I' \supset J$, it must be that $x \in B_{\vec{i},J}^*$. This holds for any J , so $x \in \bigcap_J B_{\vec{i},J}^*$. +

Since $\bigcap_I B_I^*$ is non-empty, there is some \vec{i} such that $\bigcap_J B_{\vec{i},J}^* \neq \emptyset$. But this leads to a contradiction, so it must be that $\mu^n(\bigcap_I A_I^*) > 0$, and therefore, as we have shown, $\mu^n(\bigcap_{I \in \mathcal{I}} A_I) \geq \frac{1}{|\mathcal{I}|+1} \mu^n(\bigcap_{I \in \mathcal{I}} A_I^*) > 0$. □

Corollary 7.22. *For each $\epsilon > 0, k$ and each k -regular hypergraph (W, F) there is a $\delta > 0$ such that whenever (V, E) is a k -regular hypergraph such that there are at most $\delta|V|^{|W|}$ copies of (W, F) , it is possible to remove at most $\epsilon|V|^k$ edges to obtain a hypergraph with no copies (W, F) .*

Proof. Suppose not. Fix $\epsilon > 0, k, (W, F)$ so that for each K , there is an k -regular hypergraph (V_K, E_K) with at most $\delta|V|^{|W|}$ copies of (W, F) and such that every sub-hypergraph with at least $|E_K| - \epsilon|V_K|^n$ edges contains a copy.

We consider a signature of AML consisting of an k -ary predicate E . We consider finite measured structures \mathfrak{N}_K defined as follows:

- The universe of \mathfrak{N}_K is V_K ,
- The measures on \mathfrak{N}_K are given by the normalized counting measure,
- E is interpreted by E_K .

Let \mathfrak{N} be a (nonprincipal) ultraproduct of these structures and fix a countable elementary submodel \mathfrak{M} . We may assume $W = [1, n]$. For each $I \in \binom{[1, n]}{k}$, define $E^I \in \mathcal{B}_{n, I}(M)$ to be the copy of E on the coordinates I . Observe that if $\vec{x} \in N^n$, \vec{x} is a copy of W exactly if $W \in \bigcap_{I \in F} E^I$. In particular, $\mu^n(\bigcap_{I \in F} E^I) = 0$.

By the preceding theorem, we can find definable sets B_I such that $\mu^n(E^I \setminus B_I) < \epsilon/|\mathcal{I}|$ for each $I \in \binom{[1, n]}{k}$ so that $\bigcap_I B_I = \emptyset$. Note that saying this intersection of definable sets is empty is expressed by a single formula, and is therefore true for almost every structure (V_K, E_K) . But then for some K , $(V_K, E_K \cap \bigcap B_I(G_K))$ is a sub-hypergraph obtained by removing fewer than $\epsilon|V_K|^k$ edges and containing no copies of (W, F) , yielding a contradiction. \square

The second author has extended this method [33] to the setting where (V, E) is a dense sub-hypergraph of a sparse random hypergraph. In this setting one replaces the normalized counting measure with a new counting measure normalized by the ambient random hypergraph, causing substantial additional complications.

Remark 7.23. Szemerédi's Theorem follows from hypergraph removal using the following, now standard, encoding. Given $A \subseteq [1, n]$, to find an arithmetic progression of length $k + 1$, we define a $k + 1$ -partite k -regular hypergraph. For each $i \in [1, k]$, we take the part X_i to be a copy of $[1, n]$, and we take X_{k+1} to be a copy of $[1, k^2 n]$. Given $x_i \in X_i$ for all i , we say (x_1, \dots, x_k) is an edge iff $\sum_{i \leq k} i \cdot x_i \in A$, and when $1 \leq i \leq k$, $(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{k+1})$ is an edge iff $\sum_{j \leq k, j \neq i} j \cdot x_j + i(x_{k+1} - \sum_{j \leq k, j \neq i} x_j) \in A$.

Suppose (x_1, \dots, x_{k+1}) is a copy of the complete k -regular hypergraph on $k + 1$ vertices (that is, every k -tuple is an edge) with $x_{k+1} \neq \sum_{i \leq k} x_i$. Then set $a = \sum_{i \leq k} i \cdot x_i$ and $d = x_{k+1} - \sum_{i \leq k} x_i \neq 0$. Then we have $a \in A$, and for each $i \leq k$, we have

$$a + id = \sum_{j \leq k} j \cdot x_j + i(x_{k+1} - \sum_{j \leq k} x_j) = \sum_{j \leq k, j \neq i} j \cdot x_j + i(x_{k+1} - \sum_{j \leq k, j \neq i} x_j) \in A.$$

On the other hand, for any $a \in A$ and any sequence (x_1, \dots, x_k) with $a = \sum_{i \leq k} i \cdot x_i$, the sequence $(x_1, \dots, x_k, \sum_{i \leq k} x_i)$ is also a copy of the complete k -regular hypergraph on $k + 1$ vertices. It is not possible to remove all such sequences by removing a small number of edges, so hypergraph removal implies that there must be many copies of the complete k -regular hypergraph on $k + 1$ vertices; but there are only a small number of sequences $(x_1, \dots, x_k, \sum_{i \leq k} x_i)$, so the remaining copies must correspond to genuine arithmetic progressions.

7.4. Gowers Norms.

Definition 7.24. Let G be a finite abelian group and consider $g : G \rightarrow \mathbb{R}$. We define the k -th Gowers uniformity norm, $\| \cdot \|_{U^k}$, by

$$\|g\|_{U^k}^{2^k} = \frac{1}{|G|^{k+1}} \sum_{x \in G} \sum_{\vec{h} \in G^k} \prod_{\omega \in \{0, 1\}^k} g(x + \omega \cdot \vec{h}).$$

It is not immediate that the right-hand side of the above display is nonnegative; however, this can be derived using an argument similar to that in the proof of Lemma 7.27 below.

The Gowers norms were introduced by Gowers in his proof of Szemerédi's Theorem [16]. Since then, the norms have proven to be powerful tools in combinatorics; for instance, these norms have been used to give a proof of the hypergraph regularity lemma [17] and in the proof of the Green-Tao theorem on arithmetic progressions in the primes [19]. Related norms for dynamical systems, the Gowers-Host-Kra norms, we introduced by Host and Kra [21], and have similarly been used to prove recurrence theorems in dynamical systems (for a few examples, see [3, 8]; [12] describes the general method and cites more than twenty examples).

Roughly speaking, the U^2 norm is large when a function has a large correlation with some group character. The U^k norms for larger k are large when a function correlates with a canonical function of a more general type. The exact characterization of these canonical functions was a substantial project [18, 20]. Here we show that, in the setting of AML, these canonical functions can be taken to be the simple functions generated by algebras $\mathcal{B}_{k,k-1}$.

In the infinite setting, we give a slightly more general definition:

Definition 7.25. Let \mathfrak{M} be a measured AML structure and let $f : M^k \rightarrow \mathbb{R}$ be bounded and \mathcal{B}_k -measurable. Define $\| \cdot \|_{U_\infty^k}$ by:

$$\|f\|_{U_\infty^k}^{2^k} = \int \prod_{\omega \in \{0,1\}^k} f(h_1^{\omega(1)}, \dots, h_k^{\omega(k)}) d\mu^{2^k}(\vec{h}^0, \vec{h}^1).$$

Suppose further that $+$ is a definable group operation on M and $g : M \rightarrow \mathbb{R}$ is bounded. Then define $f : M^k \rightarrow \mathbb{R}$ by $f(h_1, \dots, h_k) = g(\sum_{i=1}^k h_i)$ and define $\|g\|_{U_\infty^k} := \|f\|_{U_\infty^k}$.

Note that in the infinite setting, these are seminorms (and not even a seminorm for $k = 1$).

Lemma 7.26. *Let $(G_i : i \in \mathbb{N})$ be a sequence of finite abelian groups and, for each i , let $g_i : G_i \rightarrow [-1, 1]$ be a function such that $\lim_{i \rightarrow \infty} \|g_i\|_{U^k}$ exists. Let \mathcal{L} be the signature obtained by adding to the signature for abelian groups unary predicates P_q for $q \in \mathbb{Q} \cap [-1, 1]$. Let \mathfrak{G}_i be the AML \mathcal{L} -structure associated to G_i , equipped with its normalized counting measure, by interpreting $P_q^{\mathfrak{G}_i} := \{a \in G_i : g_i(a) < q\}$. Let \mathcal{U} be a nonprincipal ultrafilter on \mathbb{N} , let $\mathfrak{G} := \prod_{\mathcal{U}} \mathfrak{G}_i$, and let $g : G \rightarrow [-1, 1]$ be the ultraproduct of (g_i) . Then $\|g\|_{U_\infty^k} = \lim_{i \rightarrow \infty} \|g_i\|_{U^k}$.*

Proof. First observe that

$$\|g_i\|_{U^k}^{2^k} = \frac{1}{|G_i|^{2^k}} \sum_{\vec{h}^0, \vec{h}^1 \in G_i^k} \prod_{\omega \in \{0,1\}^k} g_i(\sum_i h_i^{\omega(i)}),$$

which is easily seen by making the substitutions $x = \sum_i h_i^0$ and $h_i = h_i^1 - h_i^0$. Define $j_i : G_i^k \rightarrow G_i^{2^k}$ by $j_i(\vec{h}) = (\sum h_p^{\omega(p)} : \omega \in \{0,1\}^k)$; observe that (j_i) is uniformly definable in \mathfrak{G}_i . Now define $g'_i : G_i^{2^k} \rightarrow [-1, 1]$ by $g'_i((x_\omega)) = \prod_{\omega \in \{0,1\}^k} g_i(x_\omega)$. Then the lemma follows from Theorem 4.13 (applied to (g'_i) and (j_i)). \square

For the rest of this subsection, we fix $f : M \rightarrow \mathbb{R}$ which is bounded and \mathcal{B}_k -measurable. We further assume that each μ^k is a probability measure on \mathcal{B}_k .

Lemma 7.27.

$$\left| \int f d\mu^k \right| \leq \|f\|_{U_\infty^k}.$$

Proof. Using Fubini's theorem, we have

$$\left| \int f(\vec{h}) d\mu^k(\vec{h}) \right|^{2^k} = \left| \int \left(\int f(\vec{h}) d\mu(h_k) \right) d\mu^{k-1}(h_1, \dots, h_{k-1}) \right|^{2^k}.$$

By Cauchy-Schwarz and Fubini again, we have

$$\left| \int f(\vec{h}) d\mu^k(\vec{h}) \right|^{2^k} \leq \left(\int f(h_1, \dots, h_{k-1}, h_k^0) f(h_1, \dots, h_{k-1}, h_k^1) d\mu^{k+1} \right)^{2^{k-1}}.$$

Repeating this process, we arrive at

$$\begin{aligned} \left| \int f(\vec{h}) d\mu^k(\vec{h}) \right|^{2^k} &\leq \left(\int \prod_{\omega \in \{0,1\}^k} f(h_1^{\omega(1)}, \dots, h_k^{\omega(k)}) d\mu^{2k}(\vec{h}^0, \vec{h}^1) \right) \\ &= \|f\|_{U_\infty^k}^{2^k}. \end{aligned}$$

□

Lemma 7.28. *For each $I \subseteq [1, k]$ with $|I| = k - 1$, let B_I be a definable set in $\mathcal{B}_{k,I}$. Then $0 \leq \|f \prod_I \chi_{B_I}\|_{U_\infty^k} \leq \|f\|_{U_\infty^k}$.*

Proof. It suffices to show that $0 \leq \|f \chi_{B_I}\|_{U_\infty^k} \leq \|f\|_{U_\infty^k}$ for a single I . Without loss of generality, we assume $I = [1, k - 1]$. Since $f = f \chi_{B_I} + f \chi_{\overline{B_I}}$, we have

$$\|f\|_{U_\infty^k}^{2^k} = \|f \chi_{B_I} + f \chi_{\overline{B_I}}\|_{U_\infty^k}^{2^k}$$

which in turn expands into a sum of 2^{2^k} terms of the form

$$\int \prod_{\omega \in \{0,1\}^k} (f \chi_{S_\omega})(h_1^{\omega(1)}, \dots, h_k^{\omega(k)}) d\mu^{2k}(\vec{h}^0, \vec{h}^1),$$

where each S_ω is either χ_{B_I} or $\chi_{\overline{B_I}}$. Observe that $\|f \chi_{B_I}\|_{U_\infty^k}^{2^k}$ corresponds to the term when each $S_\omega = \chi_{B_I}$. Thus, it suffices to show that all of the 2^{2^k} terms are non-negative.

Observe that if $\omega(i) = \omega'(i)$ for each $i \in I$ but $S_\omega \neq S_{\omega'}$ then for any $\chi_{S_\omega} \chi_{S_{\omega'}} = 0$, whence the corresponding integral is 0. We thus restrict ourselves to the case where, whenever $\omega \upharpoonright I = \omega' \upharpoonright I$, $S_\omega = S_{\omega'}$. In this case, we have

$$\begin{aligned} &\int \prod_{\omega \in \{0,1\}^k} (f \chi_{S_\omega})(h_1^{\omega(1)}, \dots, h_k^{\omega(k)}) d\mu^{2k}(\vec{h}^0, \vec{h}^1) \\ &= \int \int \prod_{\omega \in \{0,1\}^k} (f \chi_{S_\omega})(h_1^{\omega(1)}, \dots, h_k^{\omega(k)}) d\mu^2(h_k^0, h_k^1) d\mu^{2(k-1)} \\ &= \int \left[\int \prod_{\omega \in \{0,1\}^k} (f \chi_{S_\omega})(h_1^{\omega(1)}, \dots, h_k^{\omega(k)}) d\mu(h_k) \right]^2 d\mu^{2(k-1)}. \end{aligned}$$

Since the inside of the integral is always non-negative, this term is non-negative. \square

We set $D(f)(\vec{h}^0) := \int \prod_{\omega \in \{0,1\}^k, \omega \neq \vec{0}} f(h_1^{\omega(1)}, \dots, h_k^{\omega(k)}) d\mu^k(\vec{h}^1)$. Observe that $\|f\|_{U_\infty^k}^{2^k} = \int f \cdot D(f) d\mu^k(\vec{h}^0)$.

Lemma 7.29. *$D(f)$ is measurable with respect to $\mathcal{B}_{k,k-1}$.*

Proof. It suffices to show that if $\|\mathbb{E}(f \mid \mathcal{B}_{k,k-1})\| = 0$ then

$$\int f(\vec{h}^0) D(f)(\vec{h}^0) d\mu^k(\vec{h}^0) = 0.$$

We have

$$\begin{aligned} \int f(\vec{h}^0) D(f)(\vec{h}^0) d\mu^k(\vec{h}^0) &= \int f(\vec{h}^0) \prod_{\omega \in \{0,1\}^k, \omega \neq \vec{0}} f(h_1^{\omega(1)}, \dots, h_k^{\omega(k)}) d\mu^{2^k}(\vec{h}^0, \vec{h}^1) \\ &= \iint f(\vec{h}^0) \prod_{\omega \in \{0,1\}^k, \omega \neq \vec{0}} f(h_1^{\omega(1)}, \dots, h_k^{\omega(k)}) d\mu^k(\vec{h}^0) d\mu^k(\vec{h}^1). \end{aligned}$$

Observe that, for a given \vec{h}^1 , we have

$$\int f(\vec{h}^0) \prod_{\omega \in \{0,1\}^k, \omega \neq \vec{0}} f(h_1^{\omega(1)}, \dots, h_k^{\omega(k)}) d\mu^k(\vec{h}^0) = 0$$

since $\prod_{\omega \in \{0,1\}^k, \omega \neq \vec{0}} f(h_1^{\omega(1)}, \dots, h_k^{\omega(k)})$ is $\mathcal{B}_{k,k-1}$ -measurable. The lemma now follows. \square

Theorem 7.30. *$\|f\|_{U_\infty^k} > 0$ if and only if $\|\mathbb{E}(f \mid \mathcal{B}_{k,k-1})\| > 0$.*

Proof. If $\|f\|_{U_\infty^k} > 0$, then $\int f D(f) d\mu^k(\vec{h}^0) > 0$; since $D(f)$ is $\mathcal{B}_{k,k-1}$ -measurable, $\|\mathbb{E}(f \mid \mathcal{B}_{k,k-1})\| > 0$.

For the other direction, if $\|\mathbb{E}(f \mid \mathcal{B}_{k,k-1})\| > 0$, we may find \mathcal{B}_I -measurable B_I so that $\int f \prod \chi_{B_I} d\mu \neq 0$, whence

$$0 < \left| \int f \prod_I \chi_{B_I} d\mu \right| \leq \|f \prod_I \chi_{B_I}\|_{U_\infty^k} \leq \|f\|_{U_\infty^k}.$$

\square

REFERENCES

- [1] Tim Austin. Deducing the multidimensional Szemerédi theorem from an infinitary removal lemma. *J. Anal. Math.*, 111:131–150, 2010.
- [2] Seyed-Mohammad Bagheri and Massoud Pourmahdian. The logic of integration. *Arch. Math. Logic*, 48(5):465–492, 2009.
- [3] V. Bergelson, B. Host, B. Kra, and I. Ruzsa. Multiple recurrence and nilsequences. *Inventiones Mathematicae*, 160(2):261–303, May 2005.
- [4] Vitaly Bergelson. Ergodic Ramsey theory. In *Logic and combinatorics (Arcata, Calif., 1985)*, volume 65 of *Contemp. Math.*, pages 63–87. Amer. Math. Soc., Providence, RI, 1987.
- [5] Vitaly Bergelson. Ergodic theory and Diophantine problems. In *Topics in symbolic dynamics and applications (Temuco, 1997)*, volume 279 of *London Math. Soc. Lecture Note Ser.*, pages 167–205. Cambridge Univ. Press, Cambridge, 2000.
- [6] C.C. Chang and J. Keisler. *Model Theory*, volume 73 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, 1990.
- [7] Z. Chatzidakis and A. Pillay. Generic structures and simple theories. *Ann. Pure Appl. Logic*, 95(1-3):71–92, 1998.

- [8] Qing Chu. Convergence of weighted polynomial multiple ergodic averages. *Proc. Amer. Math. Soc.*, 137(4):1363–1369, 2009.
- [9] G. Elek. Weak convergence of finite graphs, integrated density of states and a cheeger type inequality. *Journal of Combinatorial Theory, Series B*, 98(1):62 – 68, 2008. 10.1016/j.jctb.2007.03.004.
- [10] G. Elek and B. Szegedy. Limits of hypergraphs, removal and regularity lemmas. a non-standard approach. unpublished, 2007.
- [11] Peter Frankl and Vojtech Rödl. Extremal problems on set systems. *Random Structures Algorithms*, 20(2):131–164, 2002.
- [12] N. Frantzikinakis. Some open problems on multiple ergodic averages. *ArXiv e-prints*, March 2011.
- [13] H. Furstenberg and Y. Katznelson. Idempotents in compact semigroups and Ramsey theory. *Israel J. Math.*, 68(3):257–270, 1989.
- [14] Harry Furstenberg. Ergodic behavior of diagonal measures and a theorem of szemerdi on arithmetic progressions. *Journal d'Analyse Mathématique*, 31:204–256, 1977. 10.1007/BF02813304.
- [15] R. Goldblatt. *Lectures on the Hyperreals: An introduction to nonstandard analysis*, volume 188 of *Graduate Texts in Mathematics*. Springer-Verlag, 1998.
- [16] W. T. Gowers. A new proof of Szemerédi's theorem. *Geometric And Functional Analysis*, 11(3):465–588, Aug 2001. 10.1007/s00039-001-0332-9.
- [17] W. T. Gowers. Hypergraph regularity and the multidimensional Szemerédi theorem. *Ann. of Math. (2)*, 166(3):897–946, 2007.
- [18] Ben Green and Terence Tao. An inverse theorem for the Gowers U^3 norm. *Proc. Edinb. Math. Soc. (2)*, 51(1):73–153, 2008.
- [19] Ben Green and Terence Tao. The primes contain arbitrarily long arithmetic progressions. *Ann. of Math. (2)*, 167(2):481–547, 2008.
- [20] Ben Green, Terence Tao, and Tamar Ziegler. An inverse theorem for the Gowers U^4 -norm. *Glasg. Math. J.*, 53(1):1–50, 2011.
- [21] B. Host and B. Kra. Nonconventional ergodic averages and nilmanifolds. *Ann. of Math. (2)*, 161(1):397–488, 2005. 10.4007/annals.2005.161.397.
- [22] Ehud Hrushovski. Stable group theory and approximate subgroups. <http://www.math.huji.ac.il/~ehud/NQF/nqf.pdf>, December 2009.
- [23] A. Kechris. *Classical Descriptive Set Theory*, volume 156 of *Graduate Texts in Mathematics*. Springer-Verlag, 1995.
- [24] H. J. Keisler. Probability quantifiers. In *Model-theoretic logics*, *Perspect. Math. Logic*, pages 509–556. Springer, New York, 1985.
- [25] B. Kim and A. Pillay. Simple theories. *Ann. Pure Appl. Logic*, 88(2-3):149–164, 1997.
- [26] Y. Kohayakawa. Szemerédi's regularity lemma for sparse graphs. In *Foundations of computational mathematics (Rio de Janeiro, 1997)*, pages 216–230. Springer, Berlin, 1997.
- [27] Brendan Nagle, Vojtech Rödl, and Mathias Schacht. The counting lemma for regular k -uniform hypergraphs. *Random Structures Algorithms*, 28(2):113–179, 2006.
- [28] Endre Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, volume 260 of *Colloq. Internat. CNRS*, pages 399–401. CNRS, Paris, 1978.
- [29] T. Tao. *An Introduction to Measure Theory*, volume 126 of *Graduate Studies in Mathematics*. American Mathematical Society, 2011.
- [30] Terence Tao. A correspondence principle between (hyper)graph theory and probability theory, and the (hyper)graph removal lemma. *Journal d'Analyse Mathématique*, 103:1–45, 2007. 10.1007/s11854-008-0001-0.
- [31] Terence Tao. Norm convergence of multiple ergodic averages for commuting transformations. *Ergodic Theory Dynam. Systems*, 28(2):657–688, 2008.
- [32] Henry Towsner. Convergence of diagonal ergodic averages. *Ergodic Theory Dynam. Systems*, 29(4):1309–1326, 2009. 10.1017/S0143385708000722.
- [33] Henry Towsner. An analytic approach to sparse hypergraphs: Hypergraph removal. submitted, 2012.

UNIVERSITY OF ILLINOIS AT CHICAGO, DEPARTMENT OF MATHEMATICS, STATISTICS, AND
COMPUTER SCIENCE, SCIENCE AND ENGINEERING OFFICES (M/C 249), 851 S. MORGAN ST.,
CHICAGO, IL 60607-7045, USA

E-mail address: isaac@math.uic.edu

URL: www.math.uic.edu/~isaac

UNIVERSITY OF PENNSYLVANIA, DEPARTMENT OF MATHEMATICS, 209 S. 33RD STREET,
PHILADELPHIA, PA 19104-6395

E-mail address: htowsner@math.upenn.edu

URL: www.sas.upenn.edu/~htowsner