Received: 31 March 2023 | Accepted: 23 August 2023

DOI: 10.1112/jlms.12822

RESEARCH ARTICLE

Journal of the London Mathematical Society

Continuous stable regularity

Nicolas Chavarria[1] | **Gabriel Conant**[2] | **Anand Pillay**[1]

[1]Department of Mathematics, University of Notre Dame, Notre Dame, Indiana, USA

[2]Department of Mathematics, The Ohio State University, Columbus, Ohio, USA

Correspondence
Nicolas Chavarria, Department of Mathematics, University of Notre Dame, Notre Dame IN 46656, USA.
Email: nchavarr@nd.edu

Funding information
National Science Foundation, Grant/Award Numbers: DMS-1665025, DMS-1665035, DMS-1760212, DMS-2054271, DMS-1855503; Fields Institute for Research in Mathematical Sciences, Grant/Award Number: SimonsDistinguishedVisitorship

Abstract

We prove an analytic version of the stable graph regularity lemma by Malliaris and Shelah (Trans. Amer. Math. Soc. **366** (2014), no. 3, 1551–1585), which applies to stable functions $f : V \times W \to [0,1]$. Our methods involve continuous model theory and, in particular, results on the structure of local Keisler measures for stable continuous formulas. Along the way, we develop some basic tools around ultraproducts of metric structures and linear functionals on continuous formulas, and we also describe several concrete families of examples of stable functions.

MSC 2020
03C45, 03C66, 05C75 (primary)

INTRODUCTION

Szemerédi's regularity lemma [25] is a structure theorem for arbitrary finite graphs, and has become a fundamental tool in graph theory (see, e.g., [15]). In the setting of a bipartite graph $(V, W; E)$ (which will be our focus), the regularity lemma says roughly that V and W can be partitioned into a small number (depending only on a fixed ε) of sets V_i, W_j such that for all i, j outside of a small number of "irregular pairs," the graph $(V_i, W_j; E \cap (V_i \times W_j))$ is ε-regular, meaning that sufficiently large induced subgraphs have a common edge density up to error at most ε. (See Section 6.4 for discussion of how to reconcile the setting of bipartite graphs with unpartitioned graphs $(V; E)$.)

The regularity lemma for graphs can be recast and generalized as a decomposition theorem for functions $f : V \times W \to [0,1]$, which is sometimes called the *analytic* form of Szemerédi regularity. See [8, 17, 26], for example. In this regime, f is decomposed as the sum of a "structured" part, a "pseudorandom" part, and an "error" part. In the special case where f is $\{0,1\}$-valued, we are

back in the setting of graphs, and it is explained in several places how one recovers the usual statement of Szemerédi regularity from the analytic version (e.g., [26, Lemma 2.11]). See Section 6.3 for further discussion as well as new results in the "stable" case.

It is rather natural to try to improve the conclusion of Szemerédi's regularity lemma by placing restrictions on the class of finite graphs under consideration. In their seminal paper, Malliaris and Shelah [19] considered the restriction to "k-stable graphs," namely, graphs that omit the k-half graph $([k],[k];\leqslant)$. The improvements in the conclusions involved better bounds (as a function of ε), no irregular pairs, and "ε-homogeneity" replacing ε-regularity. Here, by ε-homogeneity of a bipartite graph $(V,W;E)$, we mean that either $|E| \geqslant (1-\varepsilon)|V||W|$ or $|E| \leqslant \varepsilon|V||W|$.

The aim of the current paper is to prove an analytic version of stable graph regularity. We refer to this as *continuous stable regularity* for various reasons to be explained later. In any case, we say that a function $f : V \times W \to [0,1]$ is (k,δ)-**stable** if there do not exist $a_1,\dots,a_k \in V$ and $b_1,\dots,b_k \in W$ such that $|f(a_i,b_j) - f(a_j,b_i)| \geqslant \delta$ for all $i<j$. (This generalizes the previous notion for graphs after possibly changing k; see Remark 1.2.) In Section 1, we will describe several families of examples of stable functions, drawing from the model-theoretic study of Hilbert spaces (see Corollary 1.4).

Among our main results is Theorem 5.1, which we quote now.

Theorem A. *Let V and W be finite sets, and suppose $f : V \times W \to [0,1]$ is a (k,δ)-stable function. Then, for any $\varepsilon > 0$ and any "decay" function $\sigma : \mathbb{N} \to (0,1)$, there are partitions $V = V_0 \cup V_1 \cup \dots \cup V_m$ and $W = W_0 \cup W_1 \cup \dots \cup W_n$, with $m,n \leqslant O_{k,\delta,\varepsilon,\sigma}(1)$, satisfying the following properties.*

- *For all $(i,j) \in [m] \times [n]$, the pair (V_i,W_j) is $(5\delta+\varepsilon;\sigma(mn))$-homogeneous.*
- *$|V_0| \leqslant \varepsilon|V_1|$ and $|W_0| \leqslant \varepsilon|W_1|$.*

Homogeneity is defined in Section 4 (see Definition 4.1). Roughly speaking, the pair (V_i,W_j) is $(\delta;\varepsilon)$-homogenous (with respect to the fixed $f : V \times W \to [0,1]$), if for all but $\varepsilon|W_j|$-many $b \in W_j$, for all but $\varepsilon|V_i|$-many $a \in V_i$, the value $f(a,b)$ is within δ of a fixed number $r \in [0,1]$ that depends only on i and j (along with a dual statement quantifying in the reverse order). In analogy to homogeneity for graphs, this yields the "density" bound $\|f|_{V_i \times W_j} - r\|_1 \leqslant \delta + 2\varepsilon$ (see Remark 4.2). The proof of Theorem A also provides definability conditions on the sets V_i and W_j in terms of the function f (see Remark 5.3). Finally, we note that the bound $O_{k,\delta,\varepsilon,\sigma}(1)$ is ineffective due to the use of pseudofinite methods.

The sets V_0 and W_0 in Theorem A are small "exceptional" sets of vertices, which can be incorporated into V_1 and W_1 to obtain a total partition (with certain costs; see Remark 4.7 and Theorem 5.2). We also prove a version of Theorem A involving balanced *equipartitions* (see Theorem 6.1). Moreover, Theorem 6.8 gives a strong decomposition theorem for stable functions in the sense of "analytic" regularity.

We now discuss our methods, as well as other aspects of this paper (including purely model-theoretic results). In combinatorics, the analytic generalization of regularity is often viewed as being very useful, but only requiring routine variations of existing arguments. So, we make the important remark that this is not the case for stable regularity. This is largely because of the shift from regularity to homogeneity, which yields a much stronger description of the object in question using highly structured ingredients. The adaptation of these ingredients to the setting of functions will require several nontrivial steps and new results.

Szemerédi's original proof of his regularity lemma was direct and finitary. Likewise, the proof of stable graph regularity in [19] was at the finitary level but with tools from stability such as the 2-rank. On the other hand, it is well known that a theorem about all finite objects of a certain kind can be obtained from a theorem about a single "nonstandard finite" (or *pseudofinite*) object (although this method typically does not yield effective bounds, where relevant). Such an approach to graph regularity is described in a blog post by Tao [27], with origins in broader work of Elek and Szegedy [6]. An account can also be found in course notes of the third author [20]. Insofar as stable regularity is concerned, it is very natural to use pseudofinite methods as one can plug into to the existing theory of local stability in model theory. This approach to stable graph regularity was carried out in [18], giving a structure theorem for infinite bipartite graphs $(V, W; E)$ for which the formula $E(x, y)$ is stable, and in the presence of finitely additive (Keisler) measures on the relevant Boolean algebras of subsets of V and W.

This will be our approach to stable regularity in the current paper. However, because we are working with a function $f : V \times W \to [0, 1]$, rather than a relation $E \subseteq V \times W$, the relevant nonstandard environment is in the realm of *continuous logic* where the basic formulas are real-valued rather than Boolean-valued. The change in logic is one reason for our use of the expression "continuous stable regularity." This also fits with combinatorics where the word "continuous" is sometimes used to indicate the passage from graph relations to real-valued functions. Finally, we note that our setting is related to recent work of Chernikov and Towsner [5] on tame regularity for [0,1]-valued $(k + 1)$-ary functions of bounded VC_k-dimension (a higher arity analog of NIP, which refers to 2-ary functions and VC-dimension; for more information and definitions, see the aforementioned citation).

There is some work on local (formula-by-formula) stability in continuous logic which we can and will appeal to (such as [4]); but we will also need to develop new results concerning Keisler measures and stable formulas in the continuous environment. This is done in Section 3. In fact, we will work in large part under the weaker assumption that the given continuous formula $\varphi(x, y)$ is "δ-stable" for a particular $\delta > 0$ (rather than *fully stable*, i.e., δ-stable for all $\delta > 0$; see Definition 3.1). For a given model M, the relevant Keisler measures will be regular Borel (probability) measures on the space $S_\varphi(M)$ of complete φ-types over M.

In the classical first-order context, a basic theorem is that if $\varphi(x, y)$ is stable, then any Keisler measure μ on the space $S_\varphi(M)$ is a countable weighted sum $\mu = \sum \alpha_i p_i$ of Dirac measures, where $p_i \in S_\varphi(M)$ and $\sum \alpha_i = 1$. An account of this result is given by the third author in [22], drawing from earlier work of Keisler [13]. We prove the following analog in the continuous setting (quoting Theorem 3.12).

Theorem B. *Let M be a metric structure. Suppose that $\varphi(x, y)$ is δ-stable, and μ is a Keisler measure on $S_\varphi(M)$. Then there is a countable collection $\{C_i\}_{i\in I}$ of pairwise disjoint closed subsets of $S_\varphi(M)$ each of diameter at most 2δ, such that $\mu = \sum_{i\in I} \alpha_i \mu_i$, where μ_i is a Keisler measure concentrating on C_i and $\sum_{i\in I} \alpha_i = 1$.*

Here, "diameter" refers to the local "d-metric" on $S_\varphi(M)$ (see Definition 3.5), which is simply the discrete metric in the classical setting. When $\varphi(x, y)$ is fully stable, Theorem B is similar to a result of Ben Yaacov from [1] (see Remark 3.13). In Section 4, we will use Theorem B to prove the following model-theoretic regularity statement (paraphrasing Theorem 4.6).

Theorem C. *Let M be an ω-saturated metric structure. Suppose that $\varphi(x,y)$ is δ-stable, and μ and ν are Keisler measures on $S_\varphi(M)$ and $S_{\varphi^*}(M)$, respectively. Then, for any $\varepsilon > 0$, there are $m, n \geqslant 1$ such that for any $\gamma \in (0,1)$, there are partitions $M^x = A_0 \cup A_1 \cup \ldots \cup A_m$ and $M^y = B_0 \cup B_1 \cup \ldots \cup B_n$ satisfying the following properties.*

- *For all $(i,j) \in [m] \times [n]$, the pair (A_i, B_j) is $(5\delta;\gamma)$-homogeneous for μ and ν.*
- *$\mu(A_0) \leqslant \varepsilon\mu(A_1)$ and $\mu(B_0) \leqslant \varepsilon\mu(B_1)$.*

Moreover, A_i and B_j satisfy specific "definability conditions" involving $\varphi(x,y)$.

A more elaborate version of this theorem is Lemma 4.4, which is stated in such a form that when M is an ultraproduct of (continuous) finite structures, then the data can be transferred to the finite to achieve Theorem A. The very delicate aspects of the transfer are carried out in Sections 5.1 and 5.2. When specialized to the classical first-order setting, the proof structure of Theorem C becomes a kind of synthesis of the strategies from [18] and [22], with some additional simplifications.

1 | STABLE FUNCTIONS

Before stating the definition of stability for functions, let us set some general notation, which will be used throughout the paper.

- Given an integer $n \geqslant 1$, let $[n] = \{1, \ldots, n\}$.
- Given a set X and a subset $A \subseteq X$, we let $\mathbf{1}_A : X \to \{0,1\}$ denote the indicator function of A. When X is understood, we write $\mathbf{1}$ for $\mathbf{1}_X$.

Definition 1.1. Let $f : V \times W \to \mathbb{R}$ be a function, where V and W are arbitrary sets. Given a linear order I and some $\delta \geqslant 0$, we say that f is (I,δ)-**stable** if there do not exist $a_i \in V$ and $b_i \in W$, for $i \in I$, such that

$$|f(a_i, b_j) - f(a_j, b_i)| \geqslant \delta \text{ for all } i < j \text{ from } I.$$

Given $k \geqslant 1$, we say that f is (k,δ)-**stable** if it is $([k],\delta)$-stable; and we say that f is δ-**stable** if it is $(\mathbb{N},\delta)$-stable. (Here we use the standard orders on $\mathbb{N}$ and $[k]$.)

It is easy to check that $f : V \times W \to \mathbb{R}$ is δ-stable if and only if it is (I,δ)-stable for all infinite linear orders I. We also note that in certain contexts, f is δ-stable if and only if it is (k,δ)-stable for some $k \geqslant 1$ (specifically, when f has bounded image and (V,W,f) is "saturated" in the model-theoretic sense).

Our definition of stability has been formulated in a particular way so as to agree with previous work in the model-theoretic setting (e.g., [4]), and connect to the appearance of stability in broader mathematics (as we discuss below). However, the definition differs slightly from the usual definition of stability for (bipartite) graphs, and so, we take a moment to clarify this.

Remark 1.2. Let V and W be sets and fix a binary relation $E \subseteq V \times W$. Then, E is k-**stable** if there do not exist $a_1, \ldots, a_k \in V$ and $b_1, \ldots, b_k \in W$ such that $E(a_i, b_j)$ holds if and only if $i \leqslant j$ (i.e., $(V,W;E)$ omits $([k],[k];\leqslant)$ as an induced subgraph). While this does not necessarily coincide with (k,δ)-stability of the indicator function $\mathbf{1}_E$ for some choice of δ, note that if $\mathbf{1}_E$ is $(k,1)$-

stable, then it easily follows that E is k-stable as defined above. (Moreover, $\mathbf{1}_E$ is $(k,1)$-stable if and only if it is (k,δ)-stable for any $\delta > 0$.) Conversely, by a routine Ramsey argument, one can show that if E is k-stable then $\mathbf{1}_E$ is $(O_k(1),1)$-stable. We will comment on a continuous analog of this situation in the Appendix.

Call a function $f : V \times W \to \mathbb{R}$ **stable** if it is δ-stable for all $\delta > 0$. This notion turns out to be quite pervasive in functional analysis. First, it corresponds to Grothendieck's "double-limit" condition from [9], used to characterize relatively weakly compact sets in the Banach space of bounded continuous functions on an arbitrary topological space. This connection has been used to provide analytic proofs of several important theorems from stability theory [2, 21]. Second, Krivine and Maurey [16] defined a Banach space B to be *stable* if the function $f(x,y) = \|x+y\|$ is stable when restricted to the unit ball U in B. They observe that any L^p-space is stable for $1 \leqslant p < \infty$, and their main result is that any infinite-dimensional stable Banach space contains ℓ^p for some $1 \leqslant p < \infty$.

Third, we discuss Hilbert spaces. Given a Hilbert space H, there is a natural interpretation of H as a metric structure in an appropriate (multisorted) language, which includes the vector space structure and the inner product. In this case, the complete theory T of H is stable, meaning that *every* formula is stable in every model of T (see [3, Section 15]). Thus, any Hilbert space satisfies the Krivine–Maurey definition of stability. Another natural formula to consider is the inner product function $f(x,y) = \langle x,y\rangle$, which is stable in any Hilbert space when restricted to the unit ball. It is now understood that the stability of the inner product in Hilbert spaces is largely responsible for the recurring phenomenon of stable formulas arising naturally in several previous settings. A well-known example is [11, Proposition 2.25] from Hrushovski's breakthrough work on the structure of approximate groups. Previous related results include [10, Lemma 6.1], [12, Lemma 5.21], and [14, Lemma 3.4], each of which relates to the study of *simple* theories. The proposition from [11] is also a key ingredient in the model-theoretic proof from [23] of Tao's algebraic regularity lemma for definable sets in finite fields [29].

We now take the opportunity to explain how stability of Hilbert spaces can be used to produce very general examples of (k,δ)-stable functions (such as those underlying [11, Proposition 2.25]). Fix a formula $\varphi(x,y)$ in the language of Hilbert spaces. Assume that $\varphi(x,y)$ is [-1, 1]-valued when restricted to the unit ball (e.g., $\langle x,y\rangle$ or $\frac{1}{2}\|x+y\|$, but x and y could also be tuples of variables). Now let H be a Hilbert space with unit ball U. Given arbitrary sets V and W, and functions $g : V \to U$ and $h : W \to U$, let $\varphi_{g,h} : V \times W \to [-1,1]$ be defined by $\varphi_{g,h}(a,b) = \varphi(g(a),h(b))$.

Theorem 1.3. *For any $\varphi(x,y)$ and $\delta > 0$, there is some $k \geqslant 1$ such that for any H, V, W, g, and h as above, $\varphi_{g,h} : V \times W \to [-1,1]$ is (k,δ)-stable.*

Proof. It suffices to show that for any $\varphi(x,y)$ and $\delta > 0$, there is some $k \geqslant 1$ such that $\varphi(x,y)$ is (k,δ)-stable in any Hilbert space. But this follows from compactness and the fact that any completion of the theory of Hilbert spaces is stable. □

Let us point out some specific cases of the previous theorem.

Corollary 1.4. *For any $\delta > 0$, there is some $k \geqslant 1$ such that, in each case below, the function f is (k,δ)-stable.*

(i) *Let X, V, and W be finite sets, and fix arbitrary functions $g : X \times V \to [-1,1]$ and $h : X \times W \to [-1,1]$. Let $f : V \times W \to [-1,1]$ be defined by*

$$f(a,b) = \frac{1}{|X|} \sum_{x \in X} g(x,a)h(x,b).$$

(ii) *Let G be a locally compact group, and fix continuous functions $g, h : G \to [-1,1]$ with compact support C. Let μ be a left Haar measure on G, normalized so that $\mu(C \cup C^{-1}) \leqslant 1$. The convolution $g * h : G \to [-1,1]$ is the function*

$$(g * h)(x) = \int g(t)h(t^{-1}x)\, d\mu.$$

Let $f : G \times G \to [-1,1]$ be defined by $f(x,y) = (g * h)(xy)$.

(iii) *Let M be a metric structure and let $\mathfrak{f}(x)$ be a Keisler functional over M (see Definition 2.12). Fix $[-1,1]$-valued formulas $\psi_1(x,y)$ and $\psi_2(x,z)$. Let $f : M^y \times M^z \to [0,1]$ be defined by $f(a,b) = \mathfrak{f}(\psi_1(x,a)\psi_2(x,b))$.*

Proof. In each case, we apply Theorem 1.3 with $\varphi(x,y) = \langle x, y \rangle$. For case (i), let $H = \mathbb{R}^X$ with the *normalized* inner product. Then f is $\varphi_{g',h'}$, where $g' : V \to U$ maps a to $g(x,a)$ and $h' : W \to U$ maps b to $h(x,b)$.

For case (ii), let $H = L^2(G,\mu)$ with inner product $\langle u, v \rangle = \int uv\, d\mu$. Let $g', h' : G \to U$ be defined by $g'(a) = g(at)$ and $h'(b) = h(t^{-1}b)$. Then, f is $\varphi_{g',h'}$.

For case (iii), let μ be the Borel measure on $S_x(M)$ corresponding to $\mathfrak{f}$, and let $H = L^2(S_x(M), \mu)$ with inner product as in (ii). Let $g : M^y \to U$ and $h : M^z \to U$ be defined by $g(a) = \psi_1(x,a)$ and $h(b) = \psi_2(x,b)$. Then, f is $\varphi_{g,h}$. □

Linear functionals on metric structures will be discussed in more detail in Section 2.3. Note that case (iii) of Corollary 1.4 generalizes case (i). Moreover, if one applies (iii) to a *classical* first-order structure M, and identifies Boolean formulas with their indicator functions, then $f(a,b) = \mu(\psi_1(x,a) \wedge \psi_2(x,b))$ where μ is a Keisler measure on $\mathrm{Def}_x(M)$. If M is sufficiently saturated and μ is invariant (over some small set), then stability of f is [11, Proposition 2.25]. This last point, namely, the use of stability of the inner product in Hilbert spaces to deduce [11, Proposition 2.25], was communicated to the third author by Remi Jaoui, after seeing it in a course by Hrushovski in Paris.

Remark 1.5. In the context of Theorem 1.3, suppose that H is $L^2(X,\mu)$ for some probability space (X,μ), $\varphi(x,y)$ is the inner product, and g and h take values in the set of $v \in H$ such that $\|v\|_\infty = 1$ (which is a subset of the unit ball). Then, in this case, one can show $k = \exp^2(O(\delta^{-1}))$ by an elementary argument due to Tao [28]. Note that this situation covers most of the examples in Corollary 1.4.

2 | PRELIMINARIES ON CONTINUOUS LOGIC

We assume familiarity with the foundations of continuous logic and continuous model theory. See [3] for an introduction to this subject.

2.1 | Basic notions

Let $\mathcal{L}$ be a continuous language. We work in the setting where formulas may take values in a bounded subset of $\mathbb{R}$. Specifically, each predicate symbol P in $\mathcal{L}$ comes with a distinguished closed bounded interval $I_P \subseteq \mathbb{R}$ (in addition to the modulus of uniform continuity Δ_P), and the definition of an $\mathcal{L}$-structure includes the requirement that P is I_P-valued.

Let T be a complete $\mathcal{L}$-theory, and fix an $\mathcal{L}$-formula $\varphi(x, y)$. For simplicity, we assume that $\varphi(x, y)$ is [0,1]-valued (but this is not crucial). Throughout this section, we work with a fixed model $M \models T$.

Let $S_\varphi(M)$ denote the space of local φ-types over M. Recall that a type $p \in S_\varphi(M)$ is uniquely determined by the function $b \mapsto \varphi(p, b)$ from M^y to [0,1]. Moreover, $S_\varphi(M)$ is a compact Hausdorff space under the natural quotient topology inherited from $S_x(M)$. Further details can be found in [4, Section 6].

Definition 2.1.

(1) A **φ-formula over** M is a continuous function $\psi : S_\varphi(M) \to \mathbb{R}$.
(2) A **φ-generated formula** is a uniformly continuous combination of $\varphi(x, y_i)$ for $i < \omega$, that is, a formula of the form $\zeta(x, \bar{y}) = \alpha(\varphi(x, y_i)_{i<\omega})$, where $\alpha : [0, 1]^\omega \to \mathbb{R}$ is a continuous function.

Remark 2.2. In [4], continuous functions on $S_\varphi(M)$ are referred to as *φ-predicates* (over M), and the word "formula" is reserved for the smaller class of syntactic or *finitary* formulas. For our purposes, this distinction will not be significant because we will either be working with the general class of formulas as defined above, or with very specific families of finitary formulas.

The next result, which is part of [4, Fact 6.4], says that φ-formulas (over M) coincide with "instances" of φ-generated formulas.

Fact 2.3. *A function $\psi : S_\varphi(M) \to \mathbb{R}$ is continuous if and only if there is a φ-generated formula $\zeta(x, \bar{y})$ and some $\bar{b} \in M^{\bar{y}}$ such that $\psi(x) = \zeta(x, \bar{b})$.*

Given a φ-formula $\psi(x)$ over M and a set $B \subseteq \mathbb{R}$, define

$$[\psi(x) \in B] := \{p \in S_\varphi(M) : \psi(p) \in B\}.$$

Given φ-formulas $\psi_1(x)$, $\psi_2(x)$ over M, and Borel sets $B_1, B_2 \subseteq \mathbb{R}$, we let

$$[\psi_1(x) \in B_1 \wedge \psi_2(x) \in B_2] := [\psi_1(x) \in B_1] \cap [\psi_2(x) \in B_2],$$

and similarly for $\vee$.

Definition 2.4.

(1) A **subbasic open set** in $S_\varphi(M)$ is a set of the form $[\varphi(x, b) \in U]$ where $b \in M^y$ and $U \subseteq \mathbb{R}$ is a bounded open interval.
(2) A **basic open set in** $S_\varphi(M)$ is a finite intersection of subbasic open sets.
(3) A subset of $S_\varphi(M)$ is **explicitly open** if it is a finite union of basic open sets.

Fact 2.5. *The basic open sets in $S_\varphi(M)$ (as defined above) are a basis for the topology on $S_\varphi(M)$.*

Definition 2.6. A **Keisler measure on** $S_\varphi(M)$ is a regular Borel probability measure on $S_\varphi(M)$.

Let $\psi(x)$ be a φ-formula over M, and fix $B \subseteq \mathbb{R}$. If B is Borel (resp., open, closed, etc.), then $[\psi(x) \in B]$ is Borel (resp., open, closed, etc.). In this case, if μ is a Keisler measure on $S_\varphi(M)$, then we write $\mu(\psi(x) \in B)$ to denote $\mu([\psi(x) \in B])$.

We will also use the previous notation in the global setting where μ is a Keisler measure on $S_x(M)$ and $\psi(x)$ is an $\mathcal{L}_M$-formula. In this case, we have the pushforward measure $\tilde{\mu}$ of μ to $S_\varphi(M)$ and, if $\psi(x)$ is a φ-formula and $B \subseteq \mathbb{R}$ is Borel, then $\mu(\psi(x) \in B) = \tilde{\mu}(\psi(x) \in B)$.

Definition 2.7. Given a set $U \subseteq \mathbb{R}$, let $\alpha_U : \mathbb{R} \to [0,1]$ be defined by

$$\alpha_U(x) = \min\{d(x, \mathbb{R}\backslash U), 1\}.$$

Note that α_U is uniformly continuous, and so, in particular, it is a logical connective. Moreover, if U is open, then $\alpha_U(x) > 0$ if and only if $x \in U$.

Suppose $V \subseteq S_\varphi(M)$ is an explicitly open set. Then, we can write

$$V = \bigcup_{i=1}^{m} \bigcap_{j=1}^{n_i} [\varphi(x, b_{i,j}) \in U_{i,j}],$$

where each $b_{i,j}$ is from M^y and each $U_{i,j}$ is a bounded open interval. Set $\bar{y} = (y_{i,j})$ and define the φ-generated formula

$$\psi(x, \bar{y}) := \max_{1 \leqslant i \leqslant m} \min_{1 \leqslant j \leqslant n_i} \alpha_{U_{i,j}}(\varphi(x, y_{i,j})).$$

Then, $\psi(x, \bar{y}) > 0$ is logically equivalent to $\bigvee_{i=1}^{m} \bigwedge_{j=1}^{n_i} \varphi(x, y_{i,j}) \in U_{i,j}$. Therefore, $V = [\psi(x, \bar{b}) > 0]$. A formula of the form $\psi(x, \bar{b})$ is called an **explicit φ-formula over** M. So, we have shown that $V \subseteq S_\varphi(M)$ is an explicitly open set if and only if it is of the form $[\psi(x, \bar{b}) > 0]$ for some explicit φ-formula $\psi(x, \bar{b})$ over M.

Remark 2.8. One can further assume that the intervals $U_{i,j}$ have rational endpoints. However, we will not make this assumption in general.

Note that if $\psi(x)$ is a φ-formula over M, and $D \subseteq \mathbb{R}$ is closed, then the expression $\psi(x) \in D$ is logically equivalent to the *$\mathcal{L}$-condition* $\alpha_{\mathbb{R}\backslash D}(\psi(x)) = 0$. In particular, $\psi(x) \in D$ defines a zeroset in M^x.

Definition 2.9. An **explicit φ-zeroset** is a subset of M^x in the lattice[†] generated by zerosets defined by $\varphi(x, b) \in D$, where $b \in M^y$ and $D \subseteq \mathbb{R}$ is closed.

For example, if $\psi(x)$ is an explicit φ-formula over M, and $\eta \in [0,1]$, then $\psi(x) \geqslant \eta$ defines an explicit φ-zeroset in M^x.

[†] Recall that a lattice of subsets of some set X is a collection of subsets closed under (finite) unions and intersections.

2.2 | Linear functionals

Let V be an ordered real vector space with a norm $\|\cdot\|$. A **linear functional** on V is map $\mathfrak{f}: V \to \mathbb{R}$ such that $\mathfrak{f}(rv + sw) = r\mathfrak{f}(v) + s\mathfrak{f}(w)$ for all $v, w \in V$ and $r, s \in \mathbb{R}$. A linear functional $\mathfrak{f}$ is **positive** if $\mathfrak{f}(v) \geqslant 0$ for all $v \in V$ such that $v \geqslant 0_V$. A linear functional $\mathfrak{f}$ is **bounded** if there is some $c \geqslant 0$ such that $|\mathfrak{f}(v)| \leqslant c\|v\|$ for all $v \in V$. The **operator norm** $\|\mathfrak{f}\|$ of a bounded linear functional $\mathfrak{f}$ is the infimum of all $c \geqslant 0$ satisfying the previous condition.

Given a compact Hausdorff space X, the space $C(X, \mathbb{R})$ of all continuous functions from X to $\mathbb{R}$ is a real ordered normed vector space (in fact, a Banach space) under the pointwise partial order and the uniform norm $\|f\|_\infty = \sup_{x \in X} |f(x)|$.

Fact 2.10. *If X is a compact Hausdorff space and $\mathfrak{f}$ is a positive linear functional on $C(X, \mathbb{R})$, then $\mathfrak{f}$ is bounded and $\|\mathfrak{f}\| = \mathfrak{f}(\mathbf{1})$.*

Proof. This is a basic exercise, but we will include the proof for later reference. We first fix $\varphi \in C(X, \mathbb{R})$, and show $|\mathfrak{f}(\varphi)| \leqslant \mathfrak{f}(\mathbf{1})\|\varphi\|_\infty$. Note that $|\varphi| \leqslant \|\varphi\|_\infty \mathbf{1}$, and so, $\|\varphi\|_\infty \mathbf{1} - \varphi \geqslant 0$ and $\|\varphi\|_\infty \mathbf{1} + \varphi \geqslant 0$. Thus, $\|\varphi\|_\infty \mathfrak{f}(\mathbf{1}) - \mathfrak{f}(\varphi) \geqslant 0$ and $\|\varphi\|_\infty \mathfrak{f}(\mathbf{1}) + \mathfrak{f}(\varphi) \geqslant 0$, as desired. This shows that $\mathfrak{f}$ is bounded and $\|\mathfrak{f}\| \leqslant \mathfrak{f}(\mathbf{1})$. Conversely, $\mathfrak{f}(\mathbf{1}) = |\mathfrak{f}(\mathbf{1})| \leqslant \|\mathfrak{f}\|\|\mathbf{1}\|_\infty = \|\mathfrak{f}\|$. □

Recall that if X is a compact Hausdorff space, then a **Radon measure (on** X**)** is a regular Borel measure μ on X such that $\mu(X) < \infty$. Given a Radon measure μ on a compact Hausdorff space X, one obtains a positive linear functional $\mathfrak{f}_\mu$ on $C(X, \mathbb{R})$ such that $\mathfrak{f}_\mu(f) = \int_X f \, d\mu$. Note that $\|\mathfrak{f}_\mu\| = \mu(X)$. The *Riesz–Markov–Kakutani theorem* states that the map $\mu \mapsto \mathfrak{f}_\mu$ is a bijection between Radon measures on X and positive linear functionals on $C(X, \mathbb{R})$. Therefore, regular Borel *probability* measures correspond to positive linear functionals of operator norm 1.

2.3 | Linear functionals on metric structures

Let $\mathcal{L}$ be a continuous language. Given an $\mathcal{L}$-structure M and some sort x, we work with the vector space $C(S_x(M), \mathbb{R})$ of *$\mathcal{L}$-formulas over M* in x, with the pointwise partial order and uniform norm as in the previous subsection. We may also decorate the norm as $\|\varphi\|_\infty^M$ for emphasis. Note that the set of finitary (syntactic) $\mathcal{L}$-formulas in x over M forms a subspace of $C(S_x(M), \mathbb{R})$, which is dense by [4, Fact 6.4]. Moreover, if $\varphi(x)$ is such a formula, then $\|\varphi\|_\infty^M = \sup_{a \in M^x} |\varphi(a)|$.

Remark 2.11. Suppose that T is a complete $\mathcal{L}$-theory. Then, the ordered normed vector space structure on $\mathcal{L}$-formulas (over $\varnothing$) in some fixed sort x is part of the theory of T. In other words, if $M \equiv N$, then for any $\mathcal{L}$-formula $\varphi(x)$, we have $\|\varphi\|_\infty^M = \|\varphi\|_\infty^N$, and given $\mathcal{L}$-formulas $\varphi(x)$ and $\psi(x)$, we have $\varphi^M \leqslant \psi^M$ if and only if $\varphi^N \leqslant \psi^N$.

Definition 2.12. Given an $\mathcal{L}$-structure M, a **Keisler functional (in** x**) over** M is a positive linear functional $\mathfrak{f}$ on $C(S_x(M), \mathbb{R})$ such that $\|\mathfrak{f}\| = 1$.

Note that any nonzero positive linear functional on $C(S_x(M), \mathbb{R})$ can be normalized to a Keisler functional. Let us restate the Reisz–Markov–Kakutani theorem in this context.

Fact 2.13. *If M is an $\mathcal{L}$-structure, then the map $\mu \mapsto \int d\mu$ is a bijection between Keisler measures on $S_x(M)$ and Keisler functionals in x over M.*

We also recall that a Keisler functional is determined entirely by its behavior on the subspace of finitary formulas.

Proposition 2.14. *Suppose that M is an $\mathcal{L}$-structure, and $\mathfrak{f}_0$ is a positive linear functional on finitary $\mathcal{L}$-formulas over M in x. Then $\mathfrak{f}_0$ extends uniquely to a positive linear functional $\mathfrak{f}$ on $C(S_x(M), \mathbb{R})$ with $\|\mathfrak{f}\| = \|\mathfrak{f}_0\|$.*

Proof. Using the same steps as in the proof of Fact 2.10, one sees that $\mathfrak{f}_0$ is bounded, and hence continuous. Thus, the claim follows from the fact that the linear subspace of finitary $\mathcal{L}$-formulas is dense in $C(S_x(M), \mathbb{R})$ (see [4, Fact 6.4]), together with basic facts in functional analysis (see [24, Exercise I.I.19]). □

In light of the previous result, we will sometimes view linear functionals as maps on $C(S_x(M), \mathbb{R})$, while in other cases as maps on the subspace of finitary $\mathcal{L}$-formulas over M, depending on the relevant context.

We end this section by examining ultraproducts of functionals. Fix an infinite index set Σ, a collection $(M_s)_{s\in\Sigma}$ of $\mathcal{L}$-structures, and an ultrafilter $\mathcal{U}$ on Σ. Let M be the metric ultraproduct $\prod_{\mathcal{U}} M_s$. Let x be a fixed sort, and suppose that for all $s \in \Sigma$, we have a positive linear functional $\mathfrak{f}_s$ on $C(S_x(M_s), \mathbb{R})$. Assume further that the set $\{\|\mathfrak{f}_s\| : s \in \Sigma\}$ of norms is bounded. Given a finitary $\mathcal{L}$-formula $\varphi(x, b)$ over M, define $\mathfrak{f}(\varphi(x, b)) = \lim_{\mathcal{U}} \mathfrak{f}_s(\varphi(x, b^s))$ where $(b^s)_{s\in\Sigma}$ is a representative of b.

Proposition 2.15. *$\mathfrak{f}$ induces a well-defined positive linear functional on $C(S_x(M), \mathbb{R})$. Moreover, $\|\mathfrak{f}\| = \lim_{\mathcal{U}} \|\mathfrak{f}_s\|$.*

Proof. We first show $\mathfrak{f}$ is well defined. Fix a finitary $\mathcal{L}$-formula $\varphi(x, y)$ and metrically $\mathcal{U}$-equivalent sequences $(b^s)_{s\in\Sigma}$ and $(c^s)_{s\in\Sigma}$. We verify $\lim_{\mathcal{U}} \mathfrak{f}_s(\varphi(x, b^s)) = \lim_{\mathcal{U}} \mathfrak{f}_s(\varphi(x, c^s))$. Toward this end, we fix $\varepsilon > 0$ and show that the set

$$X := \{s \in \Sigma : |\mathfrak{f}_s(\varphi(x, b^s)) - \mathfrak{f}_s(\varphi(x, c^s))| \leqslant \varepsilon\}$$

is in $\mathcal{U}$. Let $R > 0$ be a bound on $\|\mathfrak{f}_s\|$ for all $s \in \Sigma$. By assumption, the set

$$Y := \{s \in \Sigma : d(b^s, c^s) < \Delta_{\varphi(x,y)}(\varepsilon/R)\}$$

is in $\mathcal{U}$. Moreover, if $s \in Y$, then $\|\varphi(x, b^s) - \varphi(x, c^s)\|_\infty \leqslant \varepsilon/R$, and so,

$$|\mathfrak{f}_s(\varphi(x, b^s)) - \mathfrak{f}_s(\varphi(x, c^s))| = |\mathfrak{f}_s(\varphi(x, b^s) - \varphi(x, c^s))| \leqslant R\|\varphi(x, b^s) - \varphi(x, c^s)\|_\infty \leqslant \varepsilon.$$

Therefore, $Y \subseteq X$, whence $X \in \mathcal{U}$.

Now, since ultralimits preserve vector space operations, it follows that $\mathfrak{f}$ is a linear functional. We show next that $\mathfrak{f}$ is positive. Fix an $\mathcal{L}$-formula $\varphi(x, b)$ over M such that $\varphi(x, b) \geqslant 0$. By Łoś's theorem, $\lim_{\mathcal{U}} \inf_x \varphi(x, b^s) \geqslant 0$. Now fix $\varepsilon > 0$. Then the set $Z := \{s \in \Sigma : \varphi(x, b^s) \geqslant -\varepsilon\}$ is in $\mathcal{U}$.

If $s \in Z$ then $\varphi(x, b^s) + \varepsilon \geqslant 0$, and so, $\mathfrak{f}_s(\varphi(x, b^s) + \varepsilon) \geqslant 0$, which implies

$$\mathfrak{f}_s(\varphi(x, b^s)) \geqslant -\mathfrak{f}_s(\varepsilon) = -\varepsilon \mathfrak{f}_s(\mathbf{1}) = -\varepsilon \|\mathfrak{f}_s\|.$$

Therefore, $\mathfrak{f}(\varphi(x, b)) = \lim_{\mathcal{U}} \mathfrak{f}_s(\varphi(x, b^s)) \geqslant -\varepsilon \lim_{\mathcal{U}} \|\mathfrak{f}_s\| \geqslant -\varepsilon R$. Since $\varepsilon > 0$ was arbitrary and R is fixed, it follows that $\mathfrak{f}(\varphi(x, b)) \geqslant 0$. So, we have shown that $\mathfrak{f}$ is positive. Finally, note that $\|\mathfrak{f}\| = \mathfrak{f}(\mathbf{1}) = \lim_{\mathcal{U}} \mathfrak{f}_s(\mathbf{1}) = \lim_{\mathcal{U}} \|\mathfrak{f}_s\|$. □

Example 2.16. Working in the above setting, suppose also that each M_s is finite. Then, for a given sort x, we can define the "average value functional" $\mathfrak{f}_s$ such that if $\varphi(x)$ is a finitary L-formula over M_s, then $\mathfrak{f}_s(\varphi) = \frac{1}{|M_s^x|} \sum_{a \in M_s^x} \varphi(a)$. Note that each $\mathfrak{f}_s$ is a Keisler functional in x over M_s. By the previous proposition, it follows that $\mathfrak{f} := \lim_{\mathcal{U}} \mathfrak{f}_s$ is a Keisler functional on M, which we refer to as the **pseudofinite average value functional on** M **in sort** x.

We will also need the following standard fact (see [3, Proposition 7.6]).

Fact 2.17. *Assume $\mathcal{L}$ is countable and let $\{M_s : s \in \Sigma\}$ be a countable family of $\mathcal{L}$-structures. Then, $\prod_{\mathcal{U}} M_s$ is ω_1-saturated for any nonprincipal ultrafilter $\mathcal{U}$ on Σ.*

3 | KEISLER MEASURES ON STABLE FORMULAS

Let T be a complete $\mathcal{L}$-theory, and fix a [0,1]-valued $\mathcal{L}$-formula $\varphi(x, y)$. Since we will be working locally around $\varphi(x, y)$, there is no harm in assuming that $\mathcal{L}$ is countable.

Given $r, s \in \mathbb{R}$, and some $\varepsilon > 0$, we write $r \approx_\varepsilon s$ to denote $|r - s| \leqslant \varepsilon$. We also let $r \dot{-} s = \max\{r - s, 0\}$. So, $r \dot{-} s = 0$ if and only if $r \leqslant s$.

Definition 3.1. Given $\delta \in [0, 1]$, we say that $\varphi(x, y)$ is δ**-stable (in** T) if for every $M \models T$, the function $\varphi : M^x \times M^y \to [0, 1]$ is δ-stable (as defined in Section 1).

We say that $\varphi(x, y)$ is **stable** if it is δ-stable for all $\delta > 0$.

It is not hard to show that when checking δ-stability of $\varphi(x, y)$ with respect to T, it suffices to consider a single ω-saturated model $M \models T$. In the next lemma, we further note that δ-stability is an "open condition."

Lemma 3.2. *There is some $\delta_\varphi \in [0, 1]$ such that for any $\delta \in [0, 1]$, $\varphi(x, y)$ is δ-stable if and only if $\delta > \delta_\varphi$.*

Proof. Suppose that $\varphi(x, y)$ is δ-stable for some $\delta > 0$. We find some $\varepsilon > 0$ such that $\varphi(x, y)$ is δ'-stable for all $\delta' \in (\delta - \varepsilon, \delta)$. This suffices to prove the lemma (take $\delta_\varphi = 1$ if $\varphi(x, y)$ is not δ-stable for any $\delta \in [0, 1]$). By compactness, there is some $k \geqslant 1$ such that φ^M is (k, δ)-stable for all $M \models T$. Define the formula

$$\theta(x_1, \ldots, x_k, y_1, \ldots, y_k) := \max_{i<j} (\delta \dot{-} |\varphi(x_i, y_j) - \varphi(x_j, y_i)|).$$

Then for any $M \models T$ and $a_1, \dots, a_k \in M^x$, $b_1 \dots, b_k \in M^y$, we have $\theta(\bar{a}, \bar{b}) > 0$. By compactness, there is some $\varepsilon > 0$ such that $\inf_{\bar{x},\bar{y}}^{M} \theta(\bar{x}, \bar{y}) = \varepsilon$ for all $M \models T$. Unpacking this, it follows immediately that φ^M is (k, δ')-stable for any $M \models T$ and $\delta' \in (\delta - \varepsilon, \delta)$. Therefore, $\varphi(x, y)$ is δ'-stable for any $\delta' \in (\delta - \varepsilon, \delta)$. □

Let $\varphi^*(y, x)$ denote the same formula $\varphi(x, y)$, but with the roles of object and parameter variables exchanged.

Definition 3.3. A **min-max φ^*-generated formula** is a formula of the form

$$\zeta(y, \bar{x}) = \min_{1 \leqslant i \leqslant m} \max_{1 \leqslant j \leqslant n_i} \varphi(x_{i,j}, y)$$

for some $m, n_1, \dots, n_m \geqslant 1$.

The next result says that if $\varphi(x, y)$ is δ-stable, then φ-types are "uniformly approximately φ^*-definable." See [4, Lemma 7.4] for details.

Lemma 3.4. *Assume that $\varphi(x, y)$ is δ-stable. Then there is a min-max φ^*-generated formula $\zeta_\varphi^\delta(y, \bar{x})$ such that, for any $M \models T$ and $p \in S_\varphi(M)$, there is some $\bar{c} \in M^{\bar{x}}$ such that for all $b \in M^y$, $\varphi(p, b) \approx_\delta \zeta_\varphi^\delta(b, \bar{c})$.*

It is easy to see that if $\varphi(x, y)$ is δ-stable, then so is $\varphi^*(y, x)$, and thus, by the previous lemma, we obtain a min-max φ-generated formula $\zeta_{\varphi^*}^\delta(x, \bar{y})$ satisfying the analogous conclusion for all types $q \in S_{\varphi^*}(M)$.

For the rest of this section, we fix a model $M \models T$. We now recall the topometric space structure on $S_\varphi(M)$, as well as the associated Cantor–Bendixson ranks, as defined in [1, 4]. Recall that a *topometric space* is a pair (X, d) where X is a Hausdorff space and d is a metric on X satisfying the following properties.

(i) The metric refines the topology, that is, for every open $V \subseteq X$ and every $p \in V$, there is some $\varepsilon > 0$ such that $\{q \in X : d(p, q) < \varepsilon\} \subseteq V$.
(ii) The metric function is lower semicontinuous, that is, for any $\varepsilon > 0$, the set $\{(p, q) \in X \times X : d(p, q) \leqslant \varepsilon\}$ is closed in $X \times X$.

Definition 3.5. The **d-metric** on $S_\varphi(M)$ is $d(p, q) = \sup_{b \in M^y} |\varphi(p, b) - \varphi(q, b)|$.

Despite our use of the terminology "d-metric," we note that this metric on $S_\varphi(M)$ can be quite different from what is usually called the d-metric on the space of *complete* types (see [4, Section 4.3]). For example, if the metric on M is discrete, then so is the d-metric on complete types, although what we call the d-metric on $S_\varphi(M)$ may not be discrete (for a given φ).

Fact 3.6. *$(S_\varphi(M), d)$ is a (compact) topometric space.*

Proof. See [4, Fact 6.2] and [1] (especially Definition 1.2 and the remarks after Lemma 1.9). Given the definition of the metric on $S_\varphi(M)$, properties (i) and (ii) are also easy to check directly. □

The following is a straightforward exercise that we will use later.

Proposition 3.7. *Suppose that $\zeta(y,\bar{x})$ is a min-max φ^*-generated formula. Then $|\zeta(q,\bar{c}) - \zeta(q',\bar{c})| \leqslant d(q,q')$ for any $\bar{c} \in M^{\bar{x}}$ and $q, q' \in S_{\varphi^*}(M)$.*

Given $C \subseteq S_\varphi(M)$, let $\operatorname{diam}(C) = \sup\{d(p,q) : p,q \in C\}$.

Definition 3.8. Fix $\delta > 0$.

(1) We define closed sets $X_{\delta,\alpha} \subseteq S_\varphi(M)$, for α an ordinal. Let $X_{\delta,0} = S_\varphi(M)$ and, for a limit ordinal α, set $X_{\delta,\alpha} = \bigcap_{\beta<\alpha} X_{\delta,\beta}$. Finally, set

$$X_{\delta,\alpha+1} = \bigcap \{F \subseteq X_{\delta,\alpha} : F \text{ is closed and } \operatorname{diam}(X_{\delta,\alpha} \setminus F) \leqslant \delta\}.$$

(2) Fix a nonempty closed set $C \subseteq X$. The **δ-Cantor–Bendixson rank of C** is

$$\mathrm{CB}_\delta(C) := \sup\{\alpha : X_{\delta,\alpha} \cap C \neq \emptyset\} \in \mathrm{Ord} \cup \{\infty\}.$$

If $\mathrm{CB}_\delta(C) = \alpha < \infty$, then define $\mathrm{CBm}_\delta(C) := C \cap X_{\delta,\alpha}$.

The previous definition also makes sense when $\delta = 0$, and yields the usual Cantor–Bendixson rank of a topological space. However, we will not use this case.

Proposition 3.9. *Fix $\delta > 0$ and suppose $C \subseteq S_\varphi(M)$ is closed and nonempty.*

(a) $\mathrm{CBm}_\delta(C)$ *is closed and nonempty.*
(b) *If $\mathrm{CB}_\delta(C) < \infty$, and $D \subseteq C$ is a nonempty closed set disjoint from $\mathrm{CBm}_\delta(C)$, then $\mathrm{CB}_\delta(D) < \mathrm{CB}_\delta(C)$.*
(c) *If $\mathrm{CB}_\delta(C) < \infty$, then $\mathrm{CBm}_\delta(C)$ admits a finite open cover $U_1, \ldots, U_n$ such that $\operatorname{diam}(U_i \cap \mathrm{CBm}_\delta(C)) \leqslant \delta$ for all $1 \leqslant i \leqslant n$.*

Proof. These follow easily from the definitions. See also the remarks following [4, Definition 7.9]. □

The following is a "δ-local" analog of [4, Proposition 7.11].

Proposition 3.10. *If $\varphi(x,y)$ is δ-stable, then $\mathrm{CB}_{2\delta}(S_\varphi(M)) < \infty$.*

Proof. This is more or less implicit in [4]. But we will provide a sketch. Toward a contradiction, suppose $\mathrm{CB}_{2\delta}(S_\varphi(M)) = \infty$. Arguing inductively as in [4, Proposition 7.11], we may construct $M_0 \preceq M$ of countable density character, and types $\{p_\eta : \eta \in 2^\omega\}$ in $S_\varphi(M_0)$ such that $d(p_\eta, p_\lambda) > 2\delta$ for all distinct $\eta, \lambda \in 2^\omega$. By Lemma 3.2, we can fix $\delta' < \delta$ such that $\varphi(x,y)$ is δ'-stable. For each $\eta \in 2^\omega$, apply Lemma 3.4 to find $\bar{c}_\eta \in M_0^{\bar{x}}$ such that for all $b \in M_0^y$, $\varphi(p_\eta, b) \approx_{\delta'} \zeta_\varphi^{\delta'}(b, \bar{c}_\eta)$. Since $M_0^{\bar{x}}$ has countable density character, there are distinct $\eta, \lambda \in 2^\omega$ such that $\|\zeta_\varphi^{\delta'}(y,\bar{c}_\eta) - \zeta_\varphi^{\delta'}(y,\bar{c}_\lambda)\|_\infty < 2(\delta - \delta')$. But then $d(p_\eta, p_\lambda) < 2\delta$ by the triangle inequality, which is a contradiction. □

Our next main goal is a description of local Keisler measures for δ-stable formulas. We will obtain a continuous analog of the well-known result that in a classical discrete theory, if $\varphi(x,y)$ is stable, then any Keisler measure on $S_\varphi(M)$ can be written as an infinite weighted sum of types (see [22, Fact 1.1]).

Lemma 3.11. *Let μ be a Keisler measure on $S_\varphi(M)$, and fix $\delta > 0$. Then for any closed set $C \subseteq S_\varphi(M)$, if $\mu(C) > 0$ and $\mathrm{CB}_\delta(C) < \infty$, then there is a closed set $C' \subseteq C$ such that $\mathrm{diam}(C') \leqslant \delta$ and $\mu(C') > 0$.*

Proof. We first prove the lemma in the special case that $\mu(\mathrm{CBm}_\delta(C)) > 0$. So, assume this is the case, and let $D = \mathrm{CBm}_\delta(C)$. By Proposition 3.9($c$), there are open sets $U_1, \ldots, U_n \subseteq S_\varphi(M)$ such that $D \subseteq \bigcup_{i=1}^n U_i$ and $\mathrm{diam}(U_i \cap D) \leqslant \delta$. Let $Y_i = U_i \cap D$, and note that each Y_i is Borel. Since $\mu(D) > 0$, there must be some $1 \leqslant i \leqslant n$ such that $\mu(Y_i) > 0$. Let $C' = \overline{Y_i}$ and note that $C' \subseteq D \subseteq C$ and $\mu(C') \geqslant \mu(Y_i) > 0$. Finally, by lower semicontinuity of the metric, and since $\mathrm{diam}(Y_i) \leqslant \delta$, it follows that $\mathrm{diam}(C') \leqslant \delta$.

Now we prove the lemma in the general case. Let $C \subseteq S_\varphi(M)$ be closed with $\mu(C) > 0$ and $\mathrm{CB}_\delta(C) < \infty$. We proceed by induction on CB_δ-rank. If $\mathrm{CB}_\delta(C) = 0$, then $\mathrm{CBm}_\delta(C) = C$, and so, we can apply the special case above. So, assume the result for ranks strictly less than $\mathrm{CB}_\delta(C)$.

Let $C_1 = \mathrm{CBm}_\delta(C)$. By the special case above, we can assume $\mu(C_1) = 0$. Set $X = C \backslash C_1$. Then X is Borel and $\mu(X) = \mu(C) > 0$. By regularity, there is a closed set $D \subseteq X$ such that $\mu(D) > 0$. By Proposition 3.9(b), we have $\mathrm{CB}_\delta(D) < \mathrm{CB}_\delta(C)$. So, we can apply the induction hypothesis to find a closed set $C' \subseteq D \subseteq C$ such that $\mathrm{diam}(C') \leqslant \delta$ and $\mu(C') > 0$, as desired. □

Given a Keisler measure μ on $S_\varphi(M)$ and a Borel set $X \subseteq S_\varphi(M)$, with $\mu(X) > 0$, the *localization of μ at X* is the Keisler measure μ_X on $S_\varphi(M)$ such that $\mu_X(B) = \mu(B \cap X)/\mu(X)$ for any Borel $B \subseteq S_\varphi(M)$. We now restate and prove Theorem B from the introduction.

Theorem 3.12. *Assume that $\varphi(x,y)$ is δ-stable, and let μ be a Keisler measure on $S_\varphi(M)$. Then there is a countable collection $(C_i)_{i\in I}$ of pairwise disjoint closed sets in $S_\varphi(M)$ such that:*

(i) *for all $i \in I$, $\mathrm{diam}(C_i) \leqslant 2\delta$ and $\mu(C_i) > 0$, and*
(ii) *$\mu = \sum_{i\in I} \alpha_i \mu_{C_i}$, where $\alpha_i = \mu(C_i)$.*

Proof. Let $\{C_i : i \in I\}$ be a maximal family of pairwise disjoint closed sets in $S_\varphi(M)$ with positive measure and diameter at most 2δ. Then I must be countable by countable additivity of μ. Let $B = \bigcup_{i\in I} C_i$, and note that B is Borel. We claim that $\mu(B) = 1$. Indeed, if not then by regularity, there is a closed set $C \subseteq S_\varphi(M)$ such that $\mu(C) > 0$ and $C \cap B = \emptyset$. By Proposition 3.10, we may apply Lemma 3.11 to C to obtain a closed set $C' \subseteq C$ with $\mu(C') > 0$ and $\mathrm{diam}(C') \leqslant 2\delta$. Then C is disjoint from C_i for all $i \in I$, contradicting maximality of the family.

Now let $\alpha_i := \mu(C_i)$. Then, for any Borel set $X \subseteq S_\varphi(M)$, we have

$$\mu(X) = \mu(B \cap X) = \sum_{i\in I} \mu(C_i \cap X) = \sum_{i\in I} \alpha_i \mu_{C_i}(X).$$

Therefore, $\mu = \sum_{i\in I} \alpha_i \mu_{C_i}$, as desired. □

Remark 3.13. Call a subset of $S_\varphi(M)$ *ε-finite* if it can be covered by finitely many sets of diameter at most ε. It follows from Theorem 3.12 that if $\varphi(x,y)$ is δ-stable and μ is a Keisler measure on $S_\varphi(M)$, then for any $\varepsilon > 0$, there is some closed 2δ-finite set $C \subseteq S_\varphi(M)$ such that $\mu(C) > 1-\varepsilon$. This is closely related to [1, Theorem 3.31] which (in our setting) says that if $\varphi(x,y)$ is stable, then for any $\varepsilon > 0$, there is a metrically compact set $C \subseteq S_\varphi(M)$ such that $\mu(C) > 1-\varepsilon$. One can deduce this from Theorem 3.12 by choosing a closed $\frac{1}{n}$-finite set C_n such that $\mu(C_n) > 1 - 2^{-n}\varepsilon$, and setting $C = \bigcap_{n>0} C_n$.

Before continuing with the main theme of this paper, we state some corollaries connecting Theorem 3.12 with the idea of approximating measures by types.

Corollary 3.14. *Suppose that $\varphi(x,y)$ is δ-stable, and let μ be a Keisler measure on $S_\varphi(M)$. Then, there is a finitely supported Keisler measure μ' on $S_\varphi(M)$ such that for each $b \in M^y$,*

$$\int_{S_\varphi(M)} \varphi(x,b)\,d\mu \approx_{2\delta} \int_{S_\varphi(M)} \varphi(x,b)\,d\mu'.$$

Proof. By Lemma 3.2, we can fix some $\delta' < \delta$ such that $\varphi(x,y)$ is δ'-stable. By Theorem 3.12, we can write $\mu = \sum_{i\in I} \alpha_i \mu_{C_i}$ where I is countable and the C_i's are pairwise disjoint closed sets of diameter at most $2\delta'$. For each $i \in I$, pick some $p_i \in C_i$. So, for any $q \in C_i$ and $b \in M^y$, we have $\varphi(q,b) \approx_{2\delta'} \varphi(p_i,b)$. It follows that for any $i \in I$ and any $b \in M^y$,

$$\int_{C_i} \varphi(x,b)\,d\mu_{C_i} \approx_{2\delta'} \int_{C_i} \varphi(p_i,b)\,d\mu_{C_i} = \varphi(p_i,b). \tag{$\dagger$}$$

Let $\varepsilon = \delta - \delta'$, and choose a finite set $I_0 \subseteq I$ such that $\alpha := \sum_{i\in I_0} \alpha_i \geqslant 1-\varepsilon$. Consider the finitely supported Keisler measure $\mu' = \alpha^{-1} \sum_{i\in I_0} \alpha_i p_i$, where p_i is identified with its Dirac measure. Then, by ($\dagger$), we have that for any $b \in M^y$,

$$\begin{aligned}\int_{S_\varphi(M)} \varphi(x,b)\,d\mu &= \sum_{i\in I} \alpha_i \int_{C_i} \varphi(x,b)\,d\mu_{C_i} \approx_{2\delta'} \sum_{i\in I} \alpha_i \varphi(p_i,b)\\ &\approx_\varepsilon \sum_{i\in I_0} \alpha_i \varphi(p_i,b) = \alpha \int_{S_\varphi(M)} \varphi(x,b)\,d\mu' \approx_\varepsilon \int_{S_\varphi(M)} \varphi(x,b)\,d\mu',\end{aligned}$$

as desired. □

When $\varphi(x,y)$ is stable we obtain the following conclusion, which also follows from [1, Corollary 3.32].

Corollary 3.15. *Suppose that $\varphi(x,y)$ is stable, and let μ be a Keisler measure on $S_\varphi(M)$. Then there is a sequence $(\mu_n)_{n=0}^\infty$ of finitely supported Keisler measures on $S_\varphi(M)$ such that for all $\varepsilon > 0$, there is an N such that for all $n \geqslant N$ and $b \in M^y$,*

$$\int_{S_y(M)} \varphi(x,b)\,d\mu \approx_\varepsilon \int_{S_y(M)} \varphi(x,b)\,d\mu_n.$$

We point out that both of the previous corollaries can be rephrased as statements about linear functionals. Let us also note that the conclusion of Corollary 3.15 is a continuous generalization of what is known in the classical first-order case for NIP formulas, and thus, we expect that it extends to the setting where the continuous formula $\varphi(x, y)$ is NIP. But in any case, the approximation by finitely supported measures in Corollary 3.15 is much weaker than the statement of Theorem 3.12. In the classical first-order case, this is analogous to property (b) below of a Keisler φ-measure μ over M being substantially stronger than property (a) below.

(a) For each $\varepsilon > 0$, there are $p_1, \dots, p_k \in S_\varphi(M)$ such that for any $b \in M$, $\mu(\varphi(x, b))$ is within ε of the average value of $\varphi(x, b)$ at the p_i's.
(b) For each $\varepsilon > 0$, there are $p_1, \dots, p_k \in S_\varphi(M)$ such that $\mu(\{p_1, \dots, p_k\}) \geqslant 1 - \varepsilon$.

In the discrete setting, $\varphi(x, y)$ is stable if and only if property (b) holds for all Keisler φ-measures μ; and $\varphi(x, y)$ is NIP if and only if property (a) holds for all μ and all φ-generated formulas $\theta(x, \bar{y})$.

4 | MODEL-THEORETIC STRUCTURE OF STABLE FORMULAS

The goal of this section is to prove Theorem C, which provides a structure theorem for stable continuous formulas in terms of "homogeneous pairs." We first define our notion of homogeneity precisely, and in a general setting.

Let V and W be nonempty sets, and fix a function $f : V \times W \to [0, 1]$. Let $\mathcal{A}$ and $\mathcal{B}$ be Boolean algebras of subsets of V and W, respectively. Assume that $\mathcal{A}$ contains any set of the form $\{x \in V : f(x, b) \in D\}$ where $b \in W$ is fixed and $D \subseteq [0, 1]$ is a closed interval. Also assume the analogous condition for $\mathcal{B}$. Let μ and ν be finitely additive probability measures on $\mathcal{A}$ and $\mathcal{B}$, respectively.

Definition 4.1. Fix real numbers $\delta, \gamma, \varepsilon > 0$. Given $V_* \in \mathcal{A}$ and $W_* \in \mathcal{B}$, we say that the pair (V_*, W_*) is $(\delta; \gamma, \varepsilon)$-**homogeneous for μ and ν** if there are $V' \subseteq V_*$ and $W' \subseteq W_*$ satisfying the following properties.

(i) $V' \in \mathcal{A}$ and $\mu(V') \geqslant (1 - \gamma)\mu(V_*)$.
(ii) $W' \in \mathcal{B}$ and $\nu(W') \geqslant (1 - \gamma)\nu(W_*)$.
(iii) There is some $r \in [0, 1]$ such that for all $b \in W'$,

$$\mu(\{a \in V_* : f(a, b) \approx_\delta r\}) \geqslant (1 - \varepsilon)\mu(V_*).$$

(iv) There is some $s \in [0, 1]$ such that for all $a \in V'$,

$$\nu(\{b \in W_* : f(a, b) \approx_\delta s\}) \geqslant (1 - \varepsilon)\nu(W_*).$$

If $\gamma = \varepsilon$, then we say that (V_*, W_*) is $(\delta; \gamma)$-**homogeneous for μ and ν**.

Note that the fixed function f has been suppressed from the terminology. When the ambient measures are fixed, and there is no possibility of confusion, we will also often omit "for μ and ν." For example, if V and W are finite, then we will generally be in the situation where μ and ν are the normalized counting measures on $\mathcal{P}(V)$ and $\mathcal{P}(W)$. Let us now comment more specifically on the finite case. Given a (nonempty) finite set X and a function $f : X \to \mathbb{R}$, we define the *normalized* ℓ^1-norm $\|f\|_1 = \frac{1}{|X|} \sum_{x \in X} |f(x)|$.

Remark 4.2. Assume that V and W are finite, and suppose that (V_*, W_*) is $(\delta;\gamma,\varepsilon)$-homogeneous, witnessed by $r, s \in [0,1]$. Then:

- for all but at most $\gamma|W_*|$-many $b \in W_*$, for all but at most $\varepsilon|V_*|$-many $a \in V_*$, we have $f(a,b) \approx_\delta r$, *and dually*,
- for all but at most $\gamma|V_*|$-many $a \in V_*$, for all but at most $\varepsilon|W_*|$-many $b \in W_*$, we have $f(a,b) \approx_\delta s$.

In this case, it should intuitively follow that r and s are not much different. Indeed, if $f' = f|_{V_* \times W_*}$, then by a direct computation, $\|f'(x,y) - r\|_1$ and $\|f'(x,y) - s\|_1$ are both bounded by $\delta + \gamma + \varepsilon$. So $|r - s| \leqslant 2(\delta + \gamma + \varepsilon)$ by the triangle inequality.

We now return to the previous setting where T is a complete continuous $\mathcal{L}$-theory, with $\mathcal{L}$ countable. Fix $M \vDash T$ and a [0,1]-valued $\mathcal{L}$-formula $\varphi(x,y)$. The first step toward our main result is the following corollary of Theorem 3.12.

Corollary 4.3. *Assume that $\varphi(x,y)$ is δ-stable, and let μ be a Keisler measure on $S_\varphi(M)$. Then for any $\varepsilon > 0$, there are pairwise disjoint closed sets $C_1, \ldots, C_n \subseteq S_\varphi(M)$ satisfying the following properties.*

(i) $\operatorname{diam}(C_i) \leqslant 2\delta$ *for all* $1 \leqslant i \leqslant n$.
(ii) $\mu(\bigcup_{i=1}^n C_i) > 1 - \varepsilon\mu(C_1)$.
(iii) *For any $\gamma \in (0,1)$, there are pairwise disjoint explicitly open sets $V_1, \ldots, V_n \subseteq S_\varphi(M)$ such that, for all $1 \leqslant i \leqslant n$, $C_i \subseteq V_i$ and $\mu(V_i \backslash C_i) < \gamma\mu(V_i)$.*

Proof. Let $\{C_i : i \in I\}$ be as in Theorem 3.12, and set $\alpha_i := \mu(C_i)$. Identify I with an initial segment of $\mathbb{Z}^+$, and choose $n \geqslant 1$ large enough so that $\sum_{i=1}^n \alpha_i > 1 - \varepsilon\alpha_1$. This yields $C_1, \ldots, C_n$ satisfying (i) and (ii). Toward (iii), fix $\gamma \in (0,1)$.

For all distinct $i, j \leqslant n$, we can choose disjoint open sets $V_{i,j}$ and $W_{i,j}$ such that $C_i \subseteq V_{i,j}$ and $C_j \subseteq W_{i,j}$. Set $V_i' = \bigcap_{j \neq i}(V_{i,j} \cap W_{j,i})$. Then V_i' is open, $C_i \subseteq V_i'$, and if $i, j \leqslant n$ are distinct, then $V_i' \cap V_j' \subseteq V_{i,j} \cap W_{i,j} = \varnothing$.

Now, by regularity of μ, there is an open set $U_i \supseteq C_i$ such that $\mu(U_i) < \alpha_i/(1-\gamma)$. Let $V_i'' = V_i' \cap U_i$. Then V_i'' is open, $C_i \subseteq V_i''$, $\mu(V_i'') < \alpha_i/(1-\gamma)$, and $V_1'', \ldots, V_n''$ are pairwise disjoint. Since C_i is compact, there is an explicitly open set V_i such that $C_i \subseteq V_i \subseteq V_i''$. Note that $V_1, \ldots, V_n$ are pairwise disjoint and, for all $1 \leqslant i \leqslant n$, we have

$$\mu(V_i \backslash C_i) = \mu(V_i) - \alpha_i < \gamma\mu(V_i),$$

where the final inequality follows from $\mu(V_i) \leqslant \mu(V_i'') \leqslant \mu(U_i) < \alpha_i/(1-\gamma)$. □

We now prove a rather technical lemma that will be used to obtain both Theorems A and C from the introduction.

Lemma 4.4. *Assume $\varphi(x,y)$ is δ-stable, and let μ and ν be Keisler measures on $S_\varphi(M)$ and $S_{\varphi^*}(M)$, respectively. Fix some $\varepsilon > 0$. Then there are $m, n \geqslant 1$ such that for any $\gamma \in (0,1)$, there are:*

- *explicit φ-formulas $\psi_1(x), \ldots, \psi_m(x)$ over M,*
- *explicit φ^*-formulas $\theta_1(y), \ldots, \theta_n(y)$ over M,*
- *finite tuples $\bar{c}_1, \ldots, \bar{c}_m, \bar{d}_1, \ldots, \bar{d}_n$ from M, and*
- *some $\eta > 0$,*

satisfying the following properties.

(1) $[\psi_1(x) > 0], \dots, [\psi_m(x) > 0]$ *are pairwise disjoint.*
(2) $[\theta_1(y) > 0], \dots, [\theta_n(y) > 0]$ *are pairwise disjoint.*
(3) $\mu(\bigvee_{i=1}^m \psi_i(x) \geqslant \eta) > 1 - \varepsilon\mu(\psi_1(x) \geqslant \eta)$.
(4) $\nu(\bigvee_{j=1}^n \theta_j(y) \geqslant \eta) > 1 - \varepsilon\nu(\theta_1(y) \geqslant \eta)$.
(5) *If* $1 \leqslant i \leqslant m$, *then for all* $b \in M^y$,

$$\mu\Big(\psi_i(x) \geqslant \eta \wedge \varphi(x,b) \approx_{3\delta} \zeta_\varphi^\delta(b, \bar{c}_i)\Big) > (1-\gamma)\mu(\psi_i(x) > 0).$$

(6) *If* $1 \leqslant j \leqslant n$, *then for all* $a \in M^x$,

$$\nu\Big(\theta_j(y) \geqslant \eta \wedge \varphi(a,y) \approx_{3\delta} \zeta_{\varphi^*}^\delta(a, \bar{d}_j)\Big) > (1-\gamma)\nu(\theta_j(y) > 0).$$

(7) *For all* $(i,j) \in [m] \times [n]$, *there is some* $r \in [0,1]$ *such that*

$$\nu\Big(\theta_j(y) \geqslant \eta \wedge \zeta_\varphi^\delta(y, \bar{c}_i) \approx_{2\delta} r\Big) > (1-\gamma)\nu(\theta_j(y) > 0).$$

(8) *For all* $(i,j) \in [m] \times [n]$, *there is some* $s \in [0,1]$ *such that*

$$\mu\Big(\psi_i(x) \geqslant \eta \wedge \zeta_{\varphi^*}^\delta(x, \bar{d}_j) \approx_{2\delta} s\Big) > (1-\gamma)\mu(\psi_i(x) > 0).$$

Remark 4.5. Before starting the proof of Lemma 4.4, we take a moment to analyze properties (5)–(8), and explain the connection to homogeneous pairs.

Fix $(i,j) \in [m] \times [n]$, and consider the sets $V_i := [\psi_i(x) \geqslant \eta] \subseteq S_\varphi(M)$ and $W_j := [\theta_j(y) \geqslant \eta] \subseteq S_{\varphi^*}(M)$. Then (5) says that for any $b \in M^y$, the function $\varphi(x,b)$ is within 3δ of the *fixed value* $\zeta_\varphi^\delta(b, \bar{c}_i)$ on almost all of V_i. For readers familiar with [19], this is a functional analog of the notion of a "good set" in stable graph regularity. By itself, this does not give homogeneity, because the value $\zeta_\varphi^\delta(b, \bar{c}_i)$ depends on the choice of b. However, property (7) says that there is some single $r \in [0,1]$ (depending only on i and j) such that the function $\zeta_\varphi^\delta(y, \bar{c}_i)$ is within 2δ of r on almost all of W_j. Applying the triangle inequality, we conclude that for almost all $b \in W_j(M)$, $\varphi(x,b)$ is within 5δ of r on almost all of V_i. By symmetric arguments with properties (6) and (8), we also have some $s \in [0,1]$ such that for almost all $a \in V_i(M)$, $\varphi(a,y)$ is within 5δ of s on almost all of W_j. Altogether, (V_i, W_j) resembles a homogeneous pair.

Proof of Lemma 4.4. Apply Corollary 4.3 to both μ and ν to obtain pairwise disjoint closed sets $C_1, \dots, C_m \subseteq S_\varphi(M)$ and pairwise disjoint closed sets $D_1, \dots, D_n \subseteq S_{\varphi^*}(M)$ satisfying the following properties.

(i) For all $(i,j) \in [m] \times [n]$, $\operatorname{diam}(C_i) \leqslant 2\delta$ and $\operatorname{diam}(D_j) \leqslant 2\delta$.
(ii) $\mu(\bigcup_{i=1}^m C_i) > 1 - \varepsilon\mu(C_1)$ and $\nu(\bigcup_{j=1}^n D_j) > 1 - \varepsilon\nu(D_1)$.
(iii) Corollary 4.3(*iii*) holds for for both $C_1, \dots, C_m$ and $D_1, \dots, D_n$.

Now fix $\gamma \in (0,1)$. By (*iii*), there are pairwise disjoint explicitly open sets $V_1, \dots, V_m \subseteq S_\varphi(M)$, and pairwise disjoint explicitly opens sets $W_1, \dots, W_n \subseteq S_{\varphi^*}(M)$, such that if $(i,j) \in [m] \times [n]$, then $C_i \subseteq V_i$, $D_j \subseteq W_j$, $\mu(V_i \backslash C_i) < \gamma\mu(V_i)$, and $\nu(W_j \backslash D_j) < \gamma\nu(D_j)$.

Let $\psi_i(x)$ be an explicitly open φ-formula over M such that $V_i = [\psi_i(x) > 0]$. Let $\theta_j(y)$ be an explicitly open φ^*-formula over M such that $W_j = [\theta_j(y) > 0]$. Note that properties (1) and (2) of the lemma hold. For the rest of the properties, we need to define the appropriate $\bar{c}_i$, $\bar{d}_j$, and $\eta > 0$.

For each $1 \leqslant i \leqslant m$, let p_i be a fixed type in C_i and, using Lemma 3.4, choose a finite tuple $\bar{c}_i$ from M such that for all $b \in M^y$, $\zeta^\delta_\varphi(b, \bar{c}_i) \approx_\delta \varphi(p_i, b)$. For each $1 \leqslant j \leqslant n$, let q_j be a fixed type in D_j, and let $\bar{d}_j$ be chosen similarly (using Lemma 3.4 applied to $\varphi^*(x, y)$ and $\zeta^\delta_{\varphi^*}(x, \bar{y})$).

Next, let $\tau > 0$ be small enough so that $\mu(V_i \backslash C_i) < (\gamma - \tau)\mu(V_i)$ for all $1 \leqslant i \leqslant m$, and $\nu(W_j \backslash D_j) < (\gamma - \tau)\nu(W_j)$ for all $1 \leqslant j \leqslant n$. Note also that, by (ii) above,

$$\mu(\bigcup_{i=1}^{m} V_i) > 1 - \varepsilon\mu(V_1)[.4in]and\nu(\bigcup_{j=1}^{n} W_j) > 1 - \varepsilon\nu(W_1).$$

Altogether, by countable additivity, we may choose some $\eta > 0$ satisfying the following conditions:

* $\mu(\psi_i(x) \geqslant \eta) \geqslant (1 - \tau)\mu(V_i)$ for all $1 \leqslant i \leqslant m$,
* $\nu(\theta_j(y) \geqslant \eta) \geqslant (1 - \tau)\nu(W_j)$ for all $1 \leqslant j \leqslant n$,
* $\mu(\bigvee_{i=1}^{m} \psi_i(x) \geqslant \eta) > 1 - \varepsilon\mu(\psi_1(x) \geqslant \eta)$, and
* $\nu(\bigvee_{j=1}^{n} \theta_j(y) \geqslant \eta) > 1 - \varepsilon\nu(\theta_1(y) \geqslant \eta)$.

In particular, we have properties (3) and (4) of the lemma. As for the remaining properties, we show (5) and (7). The arguments for (6) and (8) are nearly identical.

For (5), fix $1 \leqslant i \leqslant m$ and $b \in M^y$, and define

$$V' = [\psi_i(x) \geqslant \eta \wedge \varphi(x, b) \approx_{3\delta} \zeta^\delta_\varphi(b, \bar{c}_i)].$$

We want to show $\mu(V') > (1 - \gamma)\mu(V_i)$. Note that if $p \in C_i$, then $d(p, p_i) \leqslant 2\delta$ by (i), and so,

$$\varphi(p, b) \approx_{2\delta} \varphi(p_i, b) \approx_\delta \zeta^\delta_\varphi(b, \bar{c}_i).$$

Thus, we have $[\psi_i(x) \geqslant \eta] \backslash V' \subseteq V_i \backslash C_i$. Therefore,

$$\mu([\psi_i(x) \geqslant \eta] \backslash V') \leqslant \mu(V_i \backslash C_i) < (\gamma - \tau)\mu(V_i) \leqslant$$

$$\frac{\gamma - \tau}{1 - \tau}\mu(\psi_i(x) \geqslant \eta) = \left(1 - \frac{1 - \gamma}{1 - \tau}\right)\mu(\psi_i(x) \geqslant \eta).$$

So,

$$\mu(V') > \frac{1 - \gamma}{1 - \tau}\mu(\psi_i(x) \geqslant \eta) \geqslant (1 - \gamma)\mu(V_i),$$

as desired.

Finally, for (7), fix $(i, j) \in [m] \times [n]$ and set $r = \zeta^\delta_\varphi(q_j, \bar{c}_i)$. Define

$$W' = \left[\theta_j(y) \geqslant \eta \wedge \zeta^\delta_\varphi(y, \bar{c}_i) \approx_{2\delta} r\right].$$

We want to show $\nu(W') > (1-\gamma)\nu(W_j)$. If $q \in D_j$, then $d(q, q_j) \leqslant 2\delta$ by (i), and so, $\zeta^{\delta}_{\varphi}(q, \bar{c}_i) \approx_{2\delta} r$ by Proposition 3.7. Thus, we have $[\theta_j(y) \geqslant \eta] \backslash W' \subseteq W_j \backslash D_j$. By choice of η and similar steps as above, it follows that $\nu(W') > (1-\gamma)\nu(W_j)$. □

In order to derive a statement about subsets of M (rather than $S_{\varphi}(M)$) from the previous lemma, we now assume that M is ω-saturated. In this case, any Keisler measure μ on $S_{\varphi}(M)$ induces a well-defined finitely additive probability measure on the Boolean algebra generated by explicit φ-zerosets in M^x (see Definition 2.9). We now restate and prove Theorem C from the introduction.

Theorem 4.6. *Suppose that M is ω-saturated. Assume that $\varphi(x, y)$ is δ-stable, and let μ and ν be Keisler measures on $S_{\varphi}(M)$ and $S_{\varphi^*}(M)$, respectively. Fix some $\varepsilon > 0$. Then there are $m, n \geqslant 1$ such that for any $\gamma \in (0, 1)$, there are partitions $M^x = A_0 \cup A_1 \cup \ldots \cup A_m$ and $M^y = B_0 \cup B_1 \cup \ldots \cup B_n$ satisfying the following properties.*

- *$A_1, \ldots, A_m$ are explicit φ-zerosets, and $B_1, \ldots, B_n$ are explicit φ^*-zerosets.*
- *If $(i, j) \in [m] \times [n]$, then (A_i, B_j) is $(5\delta; \gamma)$-homogeneous for μ and ν.*
- *$\mu(A_0) \leqslant \varepsilon\mu(A_1)$ and $\nu(B_0) \leqslant \varepsilon\nu(B_1)$.*

Proof. Apply Lemma 4.4. Let A_i be defined by $\psi_i(x) \geqslant \eta$, and B_j be defined by $\theta_j(y) \geqslant \eta$. Then, $A_1, \ldots, A_m$ are pairwise disjoint explicit φ-zerosets, and $B_1, \ldots, B_n$ are pairwise disjoint explicit φ^*-zerosets. Let $A_0 = M^x \backslash \bigcup_{i=1}^{m} A_i$ and $B_0 = M^y \backslash \bigcup_{j=1}^{n} B_j$. Then $\mu(A_0) \leqslant \varepsilon\mu(A_1)$ and $\nu(B_0) \leqslant \varepsilon\nu(B_1)$ by parts (3) and (4) of the lemma. As outlined in Remark 4.5, it follows from parts (5)–(8) and the triangle inequality that for all $(i, j) \in [m] \times [n]$, the pair (A_i, B_j) is $(5\delta; \gamma)$-homogeneous for μ and ν. □

Remark 4.7. In the previous theorem, A_0 and B_0 serve as small "exceptional" sets, and are *complements* of zerosets. These sets can be removed at the cost of weaker control of the error in some of the homogeneous pairs. More specifically, in the context of the theorem, let $A'_1 = A_0 \cup A_1$. Then for any $1 \leqslant j \leqslant n$, since (A_1, B_j) is $(5\delta; \gamma)$-homogeneous and $\mu(A_0) \leqslant \varepsilon\mu(A_1)$, it follows that (A'_1, B_j) is $(5\delta; \gamma, \varepsilon + \gamma)$-homogeneous. Similarly, if $B'_1 = B_0 \cup B_1$, then (A_i, B'_1) is $(5\delta; \gamma, \varepsilon + \gamma)$-homogeneous for all $1 \leqslant i \leqslant m$. Moreover, (A'_1, B'_1) is $(5\delta; \varepsilon + \gamma)$-homogeneous. On the other hand, A'_1 and B'_1 are no longer zerosets. In the continuous setting, we expect that is not generally possible to obtain partitions into zerosets with no exceptional sets. However, one could replace A_0 with an explicit φ-zeroset A'_0, which contains A_0 and satisfies $\mu(A'_0) \leqslant (1 + \gamma)\mu(A_0)$. After a similar adjustment to B_0, this would result in coverings (rather than partitions) of M^x and M^y by finitely many zerosets such that all pairs are homogeneous (with parameters as above), and there is arbitrarily small overlap between the pieces in each partition.

Remark 4.8. By applying Theorem 4.6 in the setting that T is a classical discrete theory, and $\varphi(x, y)$ is a stable formula, we recover the model-theoretic version of the Malliaris–Shelah stable regularity lemma [19], proved by Malliaris and Pillay in [18]. However, we have introduced a finer control on the error in the homogeneous pairs by means of the parameter γ, *which is allowed to depend on the size of the partition*. This will be reflected in our main result below for stable functions on finite sets (Theorem 5.1), where the degree of homogeneity is controlled by an arbitrarily chosen "decay function."

5 | STABLE REGULARITY FOR FUNCTIONS ON FINITE SETS

We now restate Theorem A, which is the main finite result. The proof is given in Section 5.2 below.

Theorem 5.1. *Let V and W be finite sets, and suppose $f : V \times W \to [0,1]$ is a (k,δ)-stable function. Then for any $\varepsilon > 0$ and any function $\sigma : \mathbb{N} \to (0,1)$, there are partitions $V = V_0 \cup V_1 \cup \dots \cup V_m$ and $W = W_0 \cup W_1 \cup \dots \cup W_n$, with $m, n \leqslant O_{k,\delta,\varepsilon,\sigma}(1)$, satisfying the following properties.*

- *For all $(i,j) \in [m] \times [n]$, the pair (V_i, W_j) is $(5\delta + \varepsilon; \sigma(mn))$-homogeneous.*
- *$|V_0| \leqslant \varepsilon |V_1|$ and $|W_0| \leqslant \varepsilon |W_1|$.*

Before proving this result, we make a few remarks. First, using the same calculations as in Remark 4.7, one can remove the exceptional sets V_0 and W_0 at the cost of weaker error in *some* of the homogeneous pairs. For example, by choosing $\sigma(n) = \frac{1}{2}\varepsilon$ in the previous result, we obtain the following simpler (but weaker) version of Theorem 5.1 without the decay function σ or the exceptional sets.

Theorem 5.2. *Let V and W be finite sets and suppose $f : V \times W \to [0,1]$ is a (k,δ)-stable function. Then for any $\varepsilon > 0$, there are partitions $V = V_1 \cup \dots \cup V_m$ and $W = W_1 \cup \dots \cup W_n$, with $m, n \leqslant O_{k,\delta,\varepsilon}(1)$, such that (V_i, W_j) is $(5\delta + \varepsilon; \varepsilon)$-homogeneous for all $(i,j) \in [m] \times [n]$.*

The proof of Theorem 5.1 will also provide strong definability conditions on the sets V_i and W_j. In order to give a precise formulation, we first define some terminology. Given a rational number $\alpha = \frac{r}{s}$, with $r, s \in \mathbb{Z}$, $s > 0$, and $\gcd(r,s) = 1$, define the *complexity* of α to be $\max\{|r|, s\}$.[†] The *complexity* of a rational interval is the maximum complexity of its endpoints.

Now let $f : V \times W \to [0,1]$ be a function and fix an integer $N \geqslant 1$. Then, we say that a subset $V' \subseteq V$ is *f-definable of complexity N* if

$$V' = \bigcup_{i=1}^{m} \bigcap_{j=1}^{n_i} \{a \in V : f(a, b_{i,j}) \in D_{i,j}\}$$

for some $b_{i,j} \in W$, some $m, n_1, \dots, n_m \leqslant N$, and some closed rational intervals $D_{i,j}$ of complexity at most N. We analogously define f-definable subsets of W. Finally, we define a *min-max f-function* on V of complexity N to be a function of the form

$$\min_{1 \leqslant i \leqslant m} \max_{1 \leqslant j \leqslant n_i} f(x, y_{i,j})$$

for some $m, n_1, \dots, n_m \leqslant N$. Min-max f-functions on W are defined analogously.

Remark 5.3. With the above terminology in hand, we can now elaborate on Theorem 5.1. In particular, in the conclusion of the theorem, we also have the following.

(1) $V_1, \dots, V_m, W_1, \dots, W_n$ are f-definable of complexity $O_{k,\delta,\varepsilon,\sigma}(1)$.
(2) There are min-max f-functions $\zeta_1(y, \bar{x})$ on W and $\zeta_2(x, \bar{y})$ on V of complexity $O_{k,\delta,\varepsilon,\sigma}(1)$, and tuples $\bar{c}_1, \dots, \bar{c}_n \in V^{|\bar{x}|}$ and $\bar{d}_1, \dots, \bar{d}_m \in W^{|\bar{y}|}$ satisfying the following properties.

[†] This is the standard number-theoretic "height" function. However, no special properties of this function will be used other than that it defines a map from $\mathbb{Q}$ to $\mathbb{N}$.

- If $1 \leqslant i \leqslant m$, then for all $b \in W_s$,

$$|\{a \in V_i : f(a,b) \approx_{3\delta+\frac{1}{2}\varepsilon} \zeta_1(b, \bar{c}_i)\}| > (1-\sigma(mn))|V_i|.$$

- If $1 \leqslant j \leqslant n$, then for all $a \in V_s$,

$$|\{b \in W_j : f(a,b) \approx_{3\delta+\frac{1}{2}\varepsilon} \zeta_2(a, \bar{d}_j)\}| > (1-\sigma(mn))|W_j|.$$

- For all $(i,j) \in [m] \times [n]$, there are rational numbers $r_{i,j}, s_{i,j} \in [0,1]$ of complexity $O_{k,\delta,\varepsilon,\sigma}(1)$ such that, if we define

$$W'_{i,j} := \{b \in W_j : \zeta_1(b, \bar{c}_i) \approx_{2\delta+\frac{1}{2}\varepsilon} r_{i,j}\} \text{ and}$$

$$V'_{i,j} := \{a \in V_i : \zeta_2(a, \bar{d}_j) \approx_{2\delta+\frac{1}{2}\varepsilon} s_{i,j}\},$$

then $|W'_{i,j}| > (1-\sigma(mn))|W_j|$ and $|V'_{i,j}| > (1-\sigma(mn))|V_i|$.

In particular, it follows from the above properties that for all $(i,j) \in [m] \times [n]$, (V_i, W_j) is $(5\delta+\varepsilon, \sigma(mn))$-homogeneous, witnessed by $r_{i,j}$, $s_{i,j}$, $W'_{i,j}$, and $V'_{i,j}$.

See Remark 5.8 for a summary of how the proof yields these extra details.

The decay function σ in Theorem 5.1 leads to strong control of the error in the homogeneous pairs. This will be used to match our work to the setting of analytic regularity for functions (see Section 6.3). Another application of the decay function will appear in Section 6.1, where we modify Theorem 5.1 so that it yields equipartitions (at the cost of the definability described in Remark 5.3).

We now start toward the proof of Theorem 5.1. Despite the similarity between this result and Theorem 4.6, the proof will not be as straightforward as corresponding results in discrete logic. There are essentially three reasons for this. The first is that the zerosets in Theorem 4.6 are not necessarily definable (in the strict sense of continuous logic), and so, we need to argue directly with the underlying formulas used to construct these zerosets. The second complication has to do with the small discrepancy that exists between an abstract ultralimit of normalized counting measures on finite sets when compared to the pseudofinite normalized average value functional (see Lemma 5.4 below). Finally, rather than working directly with homogeneous pairs (which involve an "if-then" statement with measures), it will be much cleaner to instead transfer the individual underlying components that control homogeneity, as given by Lemma 4.4.

5.1 | Approximating the counting measure

Let $\mathrm{Def}_{\mathbb{Q}}(\mathbb{R}, <)$ be the set of $D \subseteq \mathbb{R}$ that are first-order definable over $\mathbb{Q}$ in $(\mathbb{R}, <)$.

Fix a countable continuous language $\mathcal{L}$. Let $\mathcal{F}$ be a countable set of (finitary) $\mathcal{L}$-formulas, which is closed under the connectives α_D for $D \in \mathrm{Def}_{\mathbb{Q}}(\mathbb{R}, <)$ (see Definition 2.7). We define an expanded language $\mathcal{L}^+$ consisting of $\mathcal{L}$ together with the following new symbols.

(i) For each $\psi(\bar{x}) \in \mathcal{F}$, add a new [0,1]-valued predicate symbol $\mathbf{1}^+_\psi(\bar{x})$.
(ii) For each rational $r \in [0,1]$ and $\theta(x, \bar{y}) \in \mathcal{F}$, with x a singleton, add new [0,1]-valued predicate symbols $P_{\theta,r}(\bar{y})$ and $Q_{\theta,r}(\bar{y})$.

Each new predicate is given a trivial modulus of uniform continuity.

Let M be a finite $\mathcal{L}$-structure with the discrete metric. We expand M to an $\mathcal{L}^+$-structure M^+ as follows. For each $\psi(\bar{x}) \in \mathcal{F}$, interpret $\mathbf{1}^+_\psi(\bar{x})$ in M^+ as the indicator function of $\psi^M(\bar{x}) > 0$. For each rational $r \in [0,1]$ and $\theta(x,\bar{y}) \in \mathcal{F}$, define

$$P^{M^+}_{\theta,r} : M^{\bar{y}} \to [0,1] \text{ so that } \bar{b} \mapsto |\theta^M(x,\bar{b}) \leqslant r|/|M^x|, \text{ and}$$

$$Q^{M^+}_{\theta,r} : M^{\bar{y}} \to [0,1] \text{ so that } \bar{b} \mapsto |\theta^M(x,\bar{b}) \geqslant r|/|M^x|.$$

Now let $(M_s)_{s\in\mathbb{N}}$ be a collection of finite $\mathcal{L}$-structures with discrete metrics. Fix an ultrafilter $\mathcal{U}$ on $\mathbb{N}$ and let $M^+ = \prod_{\mathcal{U}} M^+_s$. Let $M = \prod_{\mathcal{U}} M_s$ be the reduct of M^+ to $\mathcal{L}$. Note that $\mathcal{L}$ and $\mathcal{L}^+$ are both countable, and so, M and M^+ are both ω_1-saturated by Fact 2.17. We identify M with its underlying universe, and use M^+ when emphasis on the expanded language is necessary.

Given a sort x, let μ^s_x and $\mathfrak{f}^s_x$ denote the normalized counting measure and average value functional on $(M^+_s)^x$. Let $\mathfrak{f}_x = \lim_{\mathcal{U}} \mathfrak{f}^s_x$ be the pseudofinite average value functional, and let μ_x denote the Keisler measure induced by $\mathfrak{f}_x$. When there is no possibility of confusion, we will omit the subscript x in the previous notation.

Given an $\mathcal{L}^+$-formula $\psi(x,\bar{b})$ over M and a set $B \subseteq \mathbb{R}$, define

$$\mu^*(\psi(x,\bar{b}) \in B) := \lim_{\mathcal{U}} \mu^s(\psi^{M_s}(x,\bar{b}^s) \in B),$$

where $(\bar{b}^s)_{s\geqslant 0}$ is a choice of representative for $\bar{b}$. Note that this is well defined because the metric on M is discrete.

Lemma 5.4. *Let $\psi(x,\bar{b})$ be an $\mathcal{L}$-formula over M, with $\psi(x,\bar{y}) \in \mathcal{F}$.*

(a) *If $U \in \mathrm{Def}_{\mathbb{Q}}(\mathbb{R},<)$ is open, then $\mu(\psi(x,\bar{b}) \in U) \leqslant \mu^*(\psi(x,\bar{b}) \in U)$.*
(b) *If $C \in \mathrm{Def}_{\mathbb{Q}}(\mathbb{R},<)$ is closed, then $\mu^*(\psi(x,\bar{b}) \in C) \leqslant \mu(\psi(x,\bar{b}) \in C)$.*

Proof. Note that if $C \subseteq \mathbb{R}$ is closed, then $\mu(\psi(x,\bar{b}) \in C) = 1 - \mu(\psi(x,\bar{b}) \in \mathbb{R}\backslash C)$ (and similarly for μ^*). Moreover, if $U \subseteq \mathbb{R}$ is open, then $\psi(x,\bar{y}) \in U$ is logically equivalent to $\alpha_U(\psi(x,\bar{y})) > 0$. So, in light of the assumptions on $\mathcal{F}$, it suffices to just prove part (a), and only consider the case of $\psi(x,\bar{b}) > 0$ (i.e., $U = (0,\infty)$).

The predicate $\mathbf{1}^+_\psi(x,\bar{b})$ induces a continuous $\{0,1\}$-valued function on $S_x(M^+)$. Let $X \subseteq S_x(M^+)$ denote the support of this function. Then

$$\begin{aligned}\mu(X) = \int_{S_x(M^+)} \mathbf{1}^+_\psi(x,\bar{b})\ d\mu &= \mathfrak{f}(\mathbf{1}^+_\psi(x,\bar{b})) = \lim_{\mathcal{U}} \mathfrak{f}^s(\mathbf{1}^+_\psi(x,\bar{b}^s)) \\ &= \lim_{\mathcal{U}} \mu^s(\psi^{M_s}(x,\bar{b}^s) > 0) = \mu^*(\psi(x,\bar{b}) > 0).\end{aligned}$$

So, to prove the result, it suffices to show $[\psi(x,\bar{b}) > 0] \subseteq X$.

Fix $p \in [\psi(x,\bar{b}) > 0]$, and set $r = \psi(p,\bar{b})$. Suppose $a \in M^x$ is such that $|\psi(a,\bar{b}) - r| \leqslant \frac{r}{2}$. Since $\psi(a,\bar{b}) = \lim_{\mathcal{U}} \psi(a^s,\bar{b}^s)$, it follows that for $\mathcal{U}$-many $s \geqslant 0$, we have $|\psi(a^s,\bar{b}^s) - r| < r$. So, for $\mathcal{U}$-many $s \geqslant 0$, we have $\psi(a^s,\bar{b}^s) > 0$, that is, $\mathbf{1}^+_\psi(a^s,\bar{b}^s) = 1$. Therefore, $\mathbf{1}^+_\psi(a,\bar{b}) = \lim_{\mathcal{U}} \mathbf{1}^+_\psi(a^s,\bar{b}^s) = 1$. Altogether, for any $a \in M^x$ if $|\psi(a,\bar{b}) - r| \leqslant \frac{r}{2}$, then $\mathbf{1}^+_\psi(a,\bar{b}) = 1$.

Now, since M^+ is ω-saturated, it follows (e.g., via [3, Proposition 7.14]) that for any $q \in S_x(M^+)$, if $|\psi(q,\bar{b}) - r| \leqslant \frac{r}{2}$, then $\mathbf{1}^+_\psi(q,\bar{b}) = 1$. Since $\psi(p,\bar{b}) = r$, we therefore have $\mathbf{1}_\psi(p,\bar{b}) = 1$, that is, $p \in X$, as desired. □

Note that the previous lemma did not involve the extra predicates introduced above in (*ii*). These will be used in the next subsection.

5.2 | Proof of Theorem 5.1

Suppose that the theorem fails for some fixed $k \geqslant 1, \delta, \varepsilon > 0$, and $\sigma : \mathbb{N} \to (0,1)$. Then, for all $s \geqslant 0$, we have a (k,δ)-stable function $f_s : V_s \times W_s \to [0,1]$ that admits no partition as in the statement of Theorem 5.1, with $m, n \leqslant s$.

Let $\mathcal{L}$ be a continuous language with two sorts V and W, along with a [0,1]-valued binary predicate symbol f on $V \times W$ with trivial modulus of uniform continuity. For each $s \geqslant 0$, define an $\mathcal{L}$-structure M_s such that $V(M_s) = V_s$, $W(M_s) = W_s$, and $f^{M_s} = f_s$. We equip M_s with the discrete metric.

Let $\mathcal{F}$ be a countable set of $\mathcal{L}$-formulas which contains $f(x,y)$ and is closed under variable substitution as well as the connectives max, min, $|\psi - \theta|$, $\psi \dot{-} r$ and $r \dot{-} \psi$ for rational r, and α_D for $D \in \mathrm{Def}_{\mathbb{Q}}(\mathbb{R},<)$ (see Definition 2.7). Let M_s^+ be the expansion of M_s to an $\mathcal{L}^+$-structure as described in Section 5.1. Fix a nonprincipal ultrafilter $\mathcal{U}$ on $\mathbb{N}$, and let $M^+ = \prod_{\mathcal{U}} M_s^+$. Let $M = \prod_{\mathcal{U}} M_s$ be the reduct of M^+ to $\mathcal{L}$, and let μ_x^s, μ_x, and μ_x^* be as defined in Section 5.1. We will only use the case where x is a singleton (in either V or W), and so, we will just write μ, μ^s, and μ^* (the relevant sort V or W will be clear from context). Set $T = \mathrm{Th}(M)$.

Lemma 5.5. *$f(x,y)$ is δ'-stable in T for any $\delta' > \delta$.*

Proof. Set $\bar{x} = (x_1, \ldots, x_k)$ and $\bar{y} = (y_1, \ldots, y_k)$. Fix $\delta' > \delta$ and define the formula

$$\theta(\bar{x},\bar{y}) := \max_{i<j}(\delta' \dot{-} |f(x_i,y_j) - f(x_j,y_i)|).$$

Then, for any $s \geqslant 0$ and $\bar{a} \in V_s^k, \bar{b} \in W_s^k$, we have $\theta(\bar{a},\bar{b}) > \delta' - \delta > 0$. Therefore, $\inf_{\bar{x},\bar{y}}^M \theta(\bar{x},\bar{y}) \geqslant \delta' - \delta > 0$. So, $f(x,y)$ is δ'-stable in T (via Fact 2.17). □

We now view μ as determining local Keisler measures on $S_f(M)$ and on $S_{f^*}(M)$ (specifically, we will work with the pushforward of μ to these local type spaces, while still using the symbol μ). In the next lemma, we transfer the properties of Lemma 4.4 to obtain suitable statements in the finite setting.

Lemma 5.6. *Let $\psi(x,\bar{b})$ be an explicit f-formula over M, and let $\theta(y,\bar{a})$ and $\zeta(y,\bar{c})$ be explicit f^*-formulas over M. Fix $\delta', \gamma \in (0,1)$ and $\eta > \lambda > 0$. Set $\rho = \eta - \lambda$. For $s \in \mathbb{N}$, define*

$$V_{s,*} = \{a \in V_s : \psi(a,\bar{b}^s) \geqslant \lambda\}[.4in] and W_{s,*} = \{b \in W_s : \theta(b,\bar{a}^s) \geqslant \lambda\}.$$

(a) *Suppose that for all $b \in M^y$,*

$$\mu\big(\psi(x,\bar{b}) \geqslant \eta \wedge f(x,b) \approx_{\delta'} \zeta(b,\bar{c})\big) \;>\; (1-\gamma)\mu(\psi(x,\bar{b}) > 0).$$

Then, for $\mathcal{U}$-many $s \in \mathbb{N}$, we have that for all $b \in W_s$,

$$|\{a \in V_{s,*} : f(a,b) \approx_{\delta'+\rho} \zeta(b,\bar{c}^s)\}| > (1-2\gamma)|V_{s,*}|.$$

(b) *Suppose that there is some $r \in [0,1]$ such that*

$$\mu(\theta(y,\bar{a}) \geqslant \eta \wedge \zeta(y,\bar{c}) \approx_{\delta'} r) \ > \ (1-\gamma)\mu(\theta(y,\bar{a}) > 0).$$

Then, for $\mathcal{U}$-many $s \in \mathbb{N}$, we have

$$|\{b \in W_{s,*} : \zeta(b,\bar{c}^s) \approx_{\delta'+\rho} r\}| > (1-\gamma)|W_{s,*}|.$$

Proof. Without loss of generality, we may assume that δ', η, and λ are rational.

Part (a). First, note that our assumptions imply $\mu(\psi(x,\bar{b}) \geqslant \eta) > 0$. So, by Lemma 5.4, we have $\mu^*(\psi(x,\bar{b}) \geqslant \lambda) > 0$. Let $\tau = \frac{1}{2}\mu^*(\psi(x,\bar{b}) \geqslant \lambda)$. Then, there is some $X_1 \in \mathcal{U}$ such that $\mu^s(\psi(x,\bar{b}^s) \geqslant \lambda) > \tau$ for all $s \in X_1$.

Now define the $\mathcal{L}$-formula

$$\xi(x,y,\bar{y},\bar{z}) := \max\{\eta \dot{-} \psi(x,\bar{y}), |f(x,y) - \zeta(y,\bar{z})| \dot{-} \delta'\},$$

which is in $\mathcal{F}$. By assumption, for all $b \in M^y$, we have

$$\mu(\xi(x,b,\bar{b},\bar{c}) = 0) \ > \ (1-\gamma)\mu(\psi(x,\bar{b}) > 0).$$

By Lemma 5.4, for all $b \in M^y$, we have

$$\mu^*(\xi(x,b,\bar{b},\bar{c}) \leqslant \rho) \ > \ (1-\gamma)\mu^*(\psi(x,\bar{b}) \geqslant \lambda).$$

So, if we define the $\mathcal{L}^+$-formula

$$\pi(y,\bar{y},\bar{z}) := (1-\gamma)Q_{\psi,\lambda}(\bar{y}) \dot{-} P_{\xi,\rho}(y,\bar{y},\bar{z}),$$

then $\sup_y^{M^+} \pi(y,\bar{b},\bar{c}) = 0$. By Łoś's theorem, there is some $X_2 \in \mathcal{U}$ such that if $s \in X_2$, then $\sup_y^{M_s^+} \pi(y,\bar{b}^s,\bar{c}^s) \leqslant \gamma\tau$. Therefore, if $s \in X_1 \cap X_2$, then, for all $b \in W_s$,

$$\mu^s(\xi(x,b,\bar{b}^s,\bar{c}^s) \leqslant \rho) \ \geqslant \ (1-\gamma)\mu^s(\psi(x,\bar{b}^s) \geqslant \lambda) - \gamma\tau \ > \ (1-2\gamma)\mu^s(\psi(x,\bar{b}^s) \geqslant \lambda),$$

and so, $|\{a \in V_{s,*} : f(a,b) \approx_{\delta'+\rho} \zeta(b,\bar{c}^s)\}| > (1-2\gamma)|V_{s,*}|$.

Part (b). Define the $\mathcal{L}$-formula

$$\chi(y,\bar{x},\bar{z}) := \max\{\eta \dot{-} \theta(y,\bar{x}), |\zeta(y,\bar{z}) - r| \dot{-} \delta'\}.$$

Then, $\mu(\chi(y,\bar{a},\bar{c}) = 0) > (1-\gamma)\mu(\theta(y,\bar{a}) > 0)$ by assumption. By Lemma 5.4,

$$\mu^*(\chi(y,\bar{a},\bar{c}) \leqslant \rho) \ \geqslant \ \mu(\chi(y,\bar{a},\bar{c}) = 0) \ > \ (1-\gamma)\mu^*(\theta(y,\bar{a}) \geqslant \lambda).$$

So, for $\mathcal{U}$-many $s \in \mathbb{N}$, we have

$$\mu^s(\chi(y, \bar{a}^s, \bar{c}^s) \leqslant \rho) \ > \ (1-\gamma)\mu^s(\theta(y, \bar{a}^s) \geqslant \lambda),$$

that is, $|\{b \in W_{s,*} : \zeta(b, \bar{c}^s) \approx_{\delta'+\rho} r\}| > (1-\gamma)|W_{s,*}|$. □

Choose some $\delta' \in (\delta, \delta + \frac{1}{6}\varepsilon)$. By Lemma 5.5, $f(x,y)$ is (k, δ')-stable in T. We apply Lemma 4.4 (with our fixed $\varepsilon > 0$) to obtain some $m, n \geqslant 1$, and then choose $\gamma := \frac{1}{2}\sigma(mn)$ in the conclusion. For the reader's convenience, we reiterate the full statement. Set $\zeta_1(y, \bar{x}) := \zeta_f^{\delta'}(y, \bar{x})$ and $\zeta_2(x, \bar{y}) := \zeta_{f^*}^{\delta'}(x, \bar{y})$. Then there are:

* explicit f-formulas $\psi_1(x, \bar{b}_1), \ldots, \psi_m(x, \bar{b}_m)$ over M,
* explicit f^*-formulas $\theta_1(y, \bar{a}_1), \ldots, \theta_n(y, \bar{a}_n)$ over M,
* finite tuples $\bar{c}_1, \ldots, \bar{c}_m, \bar{d}_1, \ldots, \bar{d}_n$ from M, and
* some $\eta > 0$,

satisfying the following properties.

(1) $[\psi_1(x, \bar{b}_1) > 0], \ldots, [\psi_m(x, \bar{b}_m) > 0]$ are pairwise disjoint.
(2) $[\theta_1(y, \bar{a}_1) > 0], \ldots, [\theta_n(y, \bar{a}_n) > 0]$ are pairwise disjoint.
(3) $\mu(\bigvee_{i=1}^m \psi_i(x, \bar{b}_i) \geqslant \eta) \ > \ 1 - \varepsilon\mu(\psi_1(x, \bar{b}_1) \geqslant \eta)$.
(4) $\mu(\bigvee_{j=1}^n \theta_j(y, \bar{a}_j) \geqslant \eta) \ > \ 1 - \varepsilon\mu(\theta_1(y, \bar{a}_1) \geqslant \eta)$.
(5) If $1 \leqslant i \leqslant m$, then for all $b \in M^y$,

$$\mu\big(\psi_i(x, \bar{b}_i) \geqslant \eta \wedge f(x,b) \approx_{3\delta'} \zeta_1(b, \bar{c}_i)\big) \ > \ (1-\gamma)\mu(\psi_i(x, \bar{b}_i) > 0).$$

(6) If $1 \leqslant j \leqslant n$, then for all $a \in M^x$,

$$\mu\big(\theta_j(y, \bar{c}_j) \geqslant \eta \wedge f(a,y) \approx_{3\delta'} \zeta_2(a, \bar{d}_j)\big) \ > \ (1-\gamma)\mu(\theta_j(y, \bar{c}_j) > 0).$$

(7) For all $(i,j) \in [m] \times [n]$, there is some $r_{i,j} \in [0,1]$ such that

$$\mu\big(\theta_j(y, \bar{c}_j) \geqslant \eta \wedge \zeta_1(y, \bar{c}_i) \approx_{2\delta'} r_{i,j}\big) \ > \ (1-\gamma)\mu(\theta_j(y) > 0).$$

(8) For all $(i,j) \in [m] \times [n]$, there is some $s_{i,j} \in [0,1]$ such that

$$\mu\big(\psi_i(x, \bar{b}_i) \geqslant \eta \wedge \zeta_2(x, \bar{d}_j) \approx_{2\delta'} s_{i,j}\big) \ > \ (1-\gamma)\mu(\psi_i(x, \bar{b}_i) > 0).$$

Note that properties (3)–(8) remain true if η is replaced by something smaller. So, we may assume $\eta \leqslant 6\delta - 6\delta' + \varepsilon$ (the right side is positive by choice of δ'). Let $\bar{a} = (\bar{a}_1, \ldots, \bar{a}_n)$ and $\bar{b} = (\bar{b}_1, \ldots, \bar{b}_m)$, and define

$$\psi(x, \bar{b}) = \min_{1 \leqslant i \leqslant m} \psi_i(x, \bar{b}_i)[.4in]and\theta(y, \bar{a}) = \min_{1 \leqslant j \leqslant n} \theta_j(y, \bar{a}_j).$$

Given $s \geqslant 0$, $1 \leqslant i \leqslant m$, and $1 \leqslant j \leqslant n$, set $V_{s,i} = \{a \in V_s : \psi_i(a, \bar{b}_i^s) \geqslant \frac{1}{2}\eta\}$ and $W_{s,j} = \{b \in W_s : \theta_j(b, \bar{a}_j^s) \geqslant \frac{1}{2}\eta\}$.

Claim 5.7. The following properties hold for $\mathcal{U}$-many $s \in \mathbb{N}$.

(a) $V_{s,1}, \dots, V_{s,m}$ are pairwise disjoint and $W_{s,1}, \dots, W_{s,n}$ are pairwise disjoint.
(b) $|\bigcup_{i=1}^{m} V_{s,i}| > |V_s| - \varepsilon|V_{s,1}|$ and $|\bigcup_{j=1}^{n} W_{s,j}| > |W_s| - \varepsilon|W_{s,1}|$.
(c) If $1 \leqslant i \leqslant m$, then for all $b \in W_s$,

$$|\{a \in V_{s,i} : f(a,b) \approx_{3\delta+\frac{1}{2}\varepsilon} \zeta_1(b, \bar{c}_i^s)\}| > (1 - \sigma(mn))|V_{s,i}|.$$

(d) If $1 \leqslant j \leqslant n$, then for all $a \in V_s$,

$$|\{b \in W_{s,j} : f(a,b) \approx_{3\delta+\frac{1}{2}\varepsilon} \zeta_2(a, \bar{d}_j^s)\}| > (1 - \sigma(mn))|W_{s,j}|.$$

(e) For all $(i,j) \in [m] \times [n]$, $|\{b \in W_{s,j} : \zeta_1(b, \bar{c}_i^s) \approx_{2\delta+\frac{1}{2}\varepsilon} r_{i,j}\}| > (1 - \sigma(mn))|W_{s,j}|$.
(f) For all $(i,j) \in [m] \times [n]$, $|\{a \in V_{s,i} : \zeta_2(a, \bar{d}_j^s) \approx_{2\delta+\frac{1}{2}\varepsilon} s_{i,j}\}| > (1 - \sigma(mn))|V_{s,i}|$.

Proof. Since the metric on M is discrete, part (a) follows easily from Łoś's theorem and properties (1) and (2) above. For part (b), note that by (3) and Lemma 5.4, we have $\mu^*(\bigvee_{i=1}^{m} \psi_i(x, \bar{b}_i) \geqslant \frac{1}{2}\eta) > 1 - \varepsilon\mu^*(\psi_1(x, \bar{b}_i) \geqslant \frac{1}{2}\eta)$, and so, $|\bigcup_{i=1}^{m} V_{s,i}| > |V_s| - \varepsilon|V_{s,1}|$ holds for $\mathcal{U}$-many $s \in \mathbb{N}$. We similarly get $|\bigcup_{j=1}^{n} W_{s,j}| > |W_s| - \varepsilon|W_{s,1}|$ for $\mathcal{U}$-many $s \in \mathbb{N}$ from (4) and Lemma 5.4. For parts (c) through (f), apply Lemma 5.6 to (5) through (8), while choosing $\lambda = \frac{1}{2}\eta$. □

Since $\mathcal{U}$ is nonprincipal, we may choose $s \geqslant m, n$ satisfying the properties in the previous claim. Set $V_{s,0} = V_s \setminus \bigcup_{i=1}^{m} V_{s,i}$ and $W_{s,0} = W_s \setminus \bigcup_{j=1}^{n} W_{s,j}$. Then $|V_{s,0}| < \varepsilon|V_{s,1}|$ and $|W_{s,0}| < \varepsilon|W_{s,1}|$ by part (b) of Claim 5.7. By parts (c), (d), (e), and (f) of Claim 5.7, and the triangle inequality, it follows that for all $(i,j) \in [m] \times [n]$, $(V_{s,i}, W_{s,j})$ is $(5\delta + \varepsilon; \sigma(mn))$-homogeneous. Altogether, this contradicts the choice of $f_s : V_s \times W_s \to [0,1]$, and we have finished the proof of Theorem 5.1.

Remark 5.8. In order to obtain the extra definability conditions described in Remark 5.3, one only needs to further assume that f_s admits no such partition in which the complexity of the ingredients is bounded by s. Then, at the end of the proof, choose s to be larger than the complexity of the objects constructed. This requires one to also assume that various parameters are rational, in particular, η, $r_{i,j}$, $s_{i,j}$, and the endpoints of the intervals involved in each ψ_i and θ_j. For η, this is easy, and for ψ_i and θ_j use Remark 2.8. Finally, replace $r_{i,j}$ and $s_{i,j}$ by a sufficiently close rational number, and use ε to absorb the difference.

6 | FURTHER RESULTS

6.1 | Equipartitions

A common tension between regularity lemmas proved using finitary methods versus those proved using model-theoretic methods is that the model-theoretic proofs typically do not provide equipartitions without further work. In the case of regularity for *arbitrary* graphs or functions, a partition can be turned into an equipartition using a number of standard methods. However, for stable regularity, which involves homogeneity and *and no irregular pairs*, more care is required to build an equipartition. In this section, we will demonstrate how the decay function in The-

orem 5.1 makes this relatively easy.[†] On the other hand, we note that the process of turning a partition into an equipartition usually results in a loss of "definability" of the pieces (as in Remark 5.3).

Theorem 6.1. *Let V and W be finite sets, and suppose $f : V \times W \to [0,1]$ is a (k,δ)-stable function. Then for any $\varepsilon > 0$ and any function $\sigma : \mathbb{N} \to (0,1)$, there are partitions $V = V_0 \cup V_1 \cup \ldots \cup V_m$ and $W = W_0 \cup W_1 \cup \ldots \cup W_n$, with $m, n \leqslant O_{k,\delta,\varepsilon,\sigma}(1)$, satisfying the following properties.*

(i) *For all $(i,j) \in [m] \times [n]$, the pair (V_i, W_j) is $(5\delta + \varepsilon; \sigma(mn))$-homogeneous.*
(ii) *$|V_i| = |V_j|$ for all $1 \leqslant i, j \leqslant m$; and $|W_i| = |W_j|$ for all $1 \leqslant i, j \leqslant n$.*
(iii) *$|V_0| \leqslant \varepsilon |V|$ and $|W_0| \leqslant \varepsilon |W|$.*

Proof. Without loss of generality, assume that σ is decreasing. Let $\tau : \mathbb{N} \to (0,1)$ be defined by $\tau(n) = \frac{\varepsilon}{2n}\sigma(4n^2\lceil \varepsilon^{-1} \rceil^2)$, and let N be the bound $O_{k,\delta,\varepsilon/2,\tau}(1)$ from Theorem 5.1. Now let $f : V \times W \to [0,1]$ be (k,δ)-stable. By Theorem 5.1, there are partitions $V = V'_0 \cup V'_1 \cup \ldots \cup V'_{m'}$ and $W = W'_0 \cup W'_1 \cup \ldots \cup W'_{n'}$, with $m', n' \leqslant N$, such that $|V'_0| \leqslant \varepsilon |V'_1|/2$, $|W'_0| \leqslant \varepsilon |W'_1|/2$, and (V'_i, W'_j) is $(5\delta + \varepsilon; \tau(m'n'))$-homogeneous for all $(i,j) \in [m'] \times [n']$. Set $N^* = 2N^2\varepsilon^{-1} = O_{k,\delta,\varepsilon,\sigma}(1)$.

For each $1 \leqslant i \leqslant m'$, partition $V'_i = V'_{i,1} \cup \ldots \cup V'_{i,t_i} \cup X_i$ so that $|V'_{i,p}| = \lceil \frac{\varepsilon}{2m'}|V| \rceil$ for all $1 \leqslant p \leqslant t_i$, and $|X_i| \leqslant \frac{\varepsilon}{2m'}|V|$ (note that we allow $t_i = 0$). Similarly, for each $1 \leqslant j \leqslant n'$, partition $W'_j = W'_{j,1} \cup \ldots \cup W'_{j,u_j} \cup Y_j$ so that $|W'_{i,q}| = \lceil \frac{\varepsilon}{2n'}|W| \rceil$ and $|Y_j| \leqslant \frac{\varepsilon}{2n'}|W|$. Let $V_1, \ldots, V_m$ enumerate $\{V'_{i,p} : 1 \leqslant i \leqslant m',\ 1 \leqslant p \leqslant t_i\}$; and let $W_1, \ldots, W_n$ enumerate $\{W'_{j,q} : 1 \leqslant j \leqslant n',\ 1 \leqslant q \leqslant u_j\}$. Then (ii) holds by construction. For each $1 \leqslant i \leqslant m'$, we have $t_i \frac{\varepsilon}{2m'}|V| \leqslant |V'_i| \leqslant |V|$, and so, $t_i \leqslant 2m'\varepsilon^{-1}$. Thus, $m \leqslant 2(m')^2\varepsilon^{-1} \leqslant N^*$. Similarly, $n \leqslant 2(n')^2\varepsilon^{-1} \leqslant N^*$.

To show (i), we fix $V'_{i,p}$ and $W'_{j,q}$, and show that $(V'_{i,p}, W'_{j,q})$ is $(5\delta + \varepsilon; \sigma(mn))$-homogeneous. By construction, there are $r \in [0,1]$ and $W' \subseteq W'_j$ such that $|W'| \geqslant (1 - \tau(m'n'))|W'_j|$ and, for all $b \in W'$,

$$|\{a \in V'_i : f(a,b) \approx_{5\delta+\varepsilon} r\}| \geqslant (1 - \tau(m'n'))|V'_i|. \qquad (\dagger)$$

Let $W'' = W' \cap W'_{j,q}$. Then,

$$|W''| \geqslant |W'_{j,q}| - \tau(m'n')|W'_j| \geqslant |W'_{j,q}| - \tau(m'n')|W| \geqslant |W'_{j,q}| - \tau(m'n')\tfrac{2n'}{\varepsilon}|W'_{i,q}|.$$

Recall that $mn \leqslant 4(m'n')^2\varepsilon^{-2}$, and so $\tau(m'n')\frac{2n'}{\varepsilon} \leqslant \sigma(mn)$ by choice of τ. So, $|W''| \geqslant (1 - \sigma(mn))|W'_{i,q}|$. Moreover, if $b \in W''$, then by (†),

$$|\{a \in V'_{i,p} : f(a,b) \approx_{5\delta+\varepsilon} r\}| \geqslant |V'_{i,p}| - \tau(m'n')|V'_i| \geqslant (1 - \sigma(mn))|V'_{i,p}|,$$

where the final inequality follows by similar calculations. By a symmetric argument, we obtain the desired homogeneity for $(V'_{i,p}, W'_{j,q})$.

Finally, set $V_0 = V'_0 \cup \bigcup_{i=1}^{m'} X_i$ and $W_0 = W'_0 \cup \bigcup_{j=1}^{n'} Y_j$. Note that we now have partitions $V = V_0 \cup V_1 \cup \ldots \cup V_m$ and $W = W_0 \cup W_1 \cup \ldots \cup W_n$ satisfying (i) and (ii). So, it remains to prove the

[†] This was also observed by Terry in the setting of stable graph regularity.

bounds in (iii). For this, we have

$$|V_0| = |V_0'| + \sum_{i=1}^{m'} |X_i| \leqslant \tfrac{\varepsilon}{2}|V_1'| + m'\left(\tfrac{\varepsilon}{2m'}|V|\right) \leqslant \varepsilon|V|.$$

By a similar argument, we get $|W_0| \leqslant \varepsilon|W|$. □

Remark 6.2. As in Theorem 5.1, the sets V_0 and W_0 are exceptional sets of vertices, which are used to ensure strong homogeneity with a decay function (c.f. Theorem 5.2) *and* to achieve perfectly balanced equipartitions. The use of an exceptional set to achieve the latter feature is typical in general regularity as well (see, e.g., [15, Theorem 1.7]). A standard alternate approach is to evenly distribute the exceptional set among the remaining pieces of the partition, which yields a new partition satisfying $||V_i| - |V_j|| \leqslant 1$ for all i, j (see, e.g., [15, Theorem 1.8]). In our situation, one could do this to remove V_0 and W_0, but it would again result in homogeneity controlled only by ε rather than σ (as in Theorem 5.2).

6.2 | The case of graphs

Continuing with Remarks 1.2 and 4.8, we note that our previous results yield a qualitative version of stable graph regularity [19] with homogeneity controlled by a decay function. Our formulation is in terms of bipartitioned graph relations, which differs from [19] (see Section 6.4 for further discussion).

Corollary 6.3. *Let V and W be finite sets and suppose $E \subseteq V \times W$ is k-stable. Then for any $\varepsilon > 0$ and any function $\sigma : \mathbb{N} \to (0, 1)$, there are partitions $V = V_0 \cup V_1 \cup \ldots \cup V_m$ and $W = W_0 \cup W_1 \cup \ldots \cup W_n$, with $m, n \leqslant O_{k,\varepsilon,\sigma}(1)$, such that for all $(i, j) \in [m] \times [n]$,*

$$|E \cap (V_i \times W_j)| \leqslant \sigma(mn)|V_i||W_j| \text{ or } |E \cap (V_i \times W_j)| \geqslant (1 - \sigma(mn))|V_i||W_j|.$$

Moreover, one of the following cases holds.

(i) $|V_0| \leqslant \varepsilon|V_1|$, $|W_0| \leqslant \varepsilon|W_1|$, *and* $V_1, \ldots, V_m, W_1, \ldots, W_n$ *are E-definable of complexity* $O_{k,\varepsilon,\sigma}(1)$.
(ii) $|V_0| \leqslant \varepsilon|V|$, $|W_0| \leqslant \varepsilon|W|$, $|V_i| = |V_j|$ *for all* $1 \leqslant i, j \leqslant m$, *and* $|W_i| = |W_j|$ *for all* $1 \leqslant i, j \leqslant n$.

Proof. Let $f = \mathbf{1}_E$. As explained in Remark 1.2, f is $(O_k(1), \delta)$-stable for any $\delta > 0$. So, we can apply Theorem 5.1 with $\delta < \frac{1}{10} - \varepsilon$ (without loss of generality, assume $\varepsilon < \frac{1}{10}$). It follows that in a $(5\delta + \varepsilon; \sigma(mn))$-homogeneous pair (V_i, W_j), we may take the uniform value $r_{i,j}$ to be either 0 or 1. This yields the main claim with case (i) (via Remark 5.3). To instead obtain case (ii), replace Theorem 5.1 with Theorem 6.1 in the previous argument. □

A quantitative proof of the previous result for graphs (stated with case (ii)) was recently given by Terry and Wolf [30, Theorem 4.7].

6.3 | Function decomposition

In this section, we use Theorem 5.1 to derive a stable analog of the analytic form of Szemerédi's regularity lemma for functions. We follow the formalism described by Tao in [26]. Since the analytic setting involves decompositions of functions into *sums* of various components, we will need to develop terminology for functions that are not necessarily [0,1]-valued.

Given a (nonempty) finite set X and functions $f, g : X \to \mathbb{R}$, define the *normalized* inner product

$$\langle f, g\rangle = \frac{1}{|X|}\sum_{x\in X} f(x)g(x).$$

The *normalized* ℓ^2-norm of f is $\|f\|_2 = \sqrt{\langle f, f\rangle}$. Note that $\|f\|_1 = \langle |f|, \mathbf{1}\rangle$, and recall also that $\|f\|_\infty = \max_{x\in X} |f(x)|$. So, the inequalities $\|f\|_1 \leqslant \|f\|_2 \leqslant \|f\|_\infty$ hold for any f (the first inequality follows from Cauchy–Schwarz).

Definition 6.4. Let V and W be finite sets, and fix $f : V \times W \to [-1, 1]$.

(1) f is (m, n)-**structured** it is of the form $\sum_{i,j} r_{i,j}\mathbf{1}_{V_i\times W_j}$ for some partitions $V = V_0 \cup V_1 \cup \ldots \cup V_m$ and $W = W_0 \cup W_1 \cup \ldots \cup W_n$, with $r_{i,j} \in [-1, 1]$.
(2) f is ε-**pseudorandom** if $|\langle f, \mathbf{1}_{A\times B}\rangle| \leqslant \varepsilon$ for all $A \subseteq V$ and $B \subseteq W$.

The previous definitions are adapted from [26], although we remark that the "structured" notion there applies to $\mathbb{R}$-valued functions, and thus includes an additional parameter K bounding $|r_{i,j}|$ (which we do not require). The notion of pseudorandomness is a special case of the definition from [26], which involves a more general schematic. We have specialized to the "product structure" setting described in [26, Example 2.3]. We can now state the analytic regularity lemma for functions (quoting [8, Lemma 1.1], with some clarification to follow).

Lemma 6.5 (Szemerédi's regularity lemma, analytic form). *Let V and W be finite sets, and let $f : V \times W \to [0, 1]$ be a function. Then for any $\varepsilon > 0$ and $\sigma : \mathbb{N} \to (0, 1)$, there are $m, n \leqslant O_{\varepsilon,\sigma}(1)$ and a decomposition $f = f_{\mathrm{str}} + f_{\mathrm{psd}} + f_{\mathrm{err}}$ such that f_{str} is (m, n)-structured, f_{psd} is $\sigma(mn)$-pseudorandom, and $\|f_{\mathrm{err}}\|_2 \leqslant \varepsilon$. Moreover, f_{psd} is $[-1, 1]$-valued, whereas f_{str} and $f_{\mathrm{str}} + f_{\mathrm{err}}$ are [0,1]-valued.*

In [8], pseudorandomness is formulated using the "box norm," and the equivalence with the definition above is a fundamental result in graph theory (see, e.g., [7, Theorem 2.4] for a precise statement in the context of functions). We also note that the previous lemma is stated in a "bipartite form," which differs from [8] (see Section 6.4 for further discussion).

As a segue to stability, let us discuss how classical graph regularity can be used to obtain a prototype of Lemma 6.5 for the indicator function of a bipartite graph. Fix some $E \subseteq V \times W$ and $\varepsilon > 0$. Given (nonempty) $A \subseteq V$ and $B \subseteq W$, set $\alpha_{A,B} := |E \cap (A \times B)|/|A \times B|$ (the density of E on $A \times B$). Then, Szemerédi's regularity lemma provides partitions $V = V_1 \cup \ldots \cup V_m$ and $W = W_1 \cup \ldots \cup W_n$ (with $m, n \leqslant O_\varepsilon(1)$), and a set Σ of pairs (i, j), such that if $Z = \bigcup_{(i,j)\in\Sigma} V_i \times W_j$, then $|Z| \leqslant \varepsilon|V \times W|$, and if $(i, j) \notin \Sigma$, then (V_i, W_j) is ε-regular. So, for $(i, j) \notin \Sigma$, $\alpha_{i,j} := \alpha_{V_i,W_j} \approx_\varepsilon \alpha_{A,B}$ for any $A \subseteq V_i$ and $B \subseteq W_j$ with $|A| \geqslant \varepsilon|V_i|$ and $|B| \geqslant \varepsilon|W_j|$. An easy calculation then shows that $\mathbf{1}_{E\cap(V_i\times W_j)} - \alpha_{i,j}$ is 2ε-pseudorandom as a function on $V_i \times W_j$. Thus, if we set $f_{\mathrm{str}} = \sum_{(i,j)\notin\Sigma} \alpha_{i,j}\mathbf{1}_{V_i\times W_j}$, $f_{\mathrm{err}} = \mathbf{1}_{E\cap Z}$, and $f_{\mathrm{psd}} = \mathbf{1}_E - f_{\mathrm{str}} - f_{\mathrm{err}}$, then $\mathbf{1}_E = f_{\mathrm{str}} + f_{\mathrm{psd}} + f_{\mathrm{err}}$, and we

have that f_{str} is (m,n)-structured, f_{psd} is 2ε-pseudorandom, and $\|f_{\mathrm{err}}\|_2 \leqslant \sqrt{\varepsilon}$. Also, f_{str} and f_{err} are [0,1]-valued, while f_{psd} is [-1, 1]-valued.

In the setting of stable *graphs*, the previous situation is qualitatively strengthened in two ways. First, the set Σ of irregular pairs does not appear, which removes the f_{err} term. Second, ε-regularity is replaced by ε-homogeneity, which is to say that each density $\alpha_{i,j}$ is within ε of some $\hat{\alpha}_{i,j} \in \{0,1\}$.[†] Thus, $(V_i, W_j; E)$ is almost complete or empty, and $\hat{\alpha}_{i,j}$ is the "generic" value of $\mathbf{1}_E$ on $V_i \times W_j$. Consequently, if one *redefines* f_{str} above so that $\alpha_{i,j}$ is replaced by $\hat{\alpha}_{i,j}$, then f_{str} is still (m,n)-structured (in fact, it is the indicator function of the union of all $V_i \times W_j$ on which E is almost complete), and $f_{\mathrm{psd}} = f - f_{\mathrm{str}}$ is a {-1, 0, 1}-valued function whose support has size at most $\varepsilon|V \times W|$. This motivates the following remark.

Remark 6.6. Suppose $f : V \times W \to [-1,1]$ is an arbitrary function such that $|\mathrm{supp}(f)| \leqslant \varepsilon|V \times W|$. Then $\|f\|_1 \leqslant \varepsilon$, which implies that *for any other function* $g : V \times W \to [-1,1]$, we have $|\langle f, g\rangle| \leqslant \|f\|_1 \leqslant \varepsilon$. So, f is automatically ε-pseudorandom. In other words, having small support can be viewed as a very strong form of pseudorandomness.

In the setting of stable functions, we will see the same picture emerge, in the sense that f_{str} will be determined by the "generic value" of f on each homogeneous pair, and f_{psd} will have small support. At this level, one could obtain from Theorem 5.2 a decomposition for stable functions that is very much like the one described above for stable graphs (see Remark 6.11). However, note that the above discussion of stable graphs does not include the stronger control on pseudorandomness using a decay function, as in Lemma 6.5. In order to reintroduce this aspect, we will need to deal with the issues of the exceptional sets V_0 and W_0, and the appearance of δ and ε in the homogeneous pairs. To handle the latter issue, we will relax the "structured" component f_{str} by allowing for some uniformly bounded fluctuation.

Definition 6.7. A function $f : V \times W \to [0,1]$ is $(\delta; m, n)$-**structured** if there is some (m,n)-structured function $g : V \times W \to [0,1]$ such that $\|f - g\|_\infty \leqslant \delta$.

As for the exceptional sets, we will deal with those in the same way as one deals with irregular pairs, which is to put them in an error term. This will yield an error term supported on a set of the form $Z' = (V_0 \times W) \cup (V \times W_0)$, where $|V_0| \leqslant \varepsilon|V|$ and $|W_0| \leqslant \varepsilon|W|$. Note that $|Z'| \leqslant 2\varepsilon|V \times W|$. So, Z' is comparable to the error set Z above in terms of size, and we similarly get bounded ℓ^2-norm for any function supported on Z'. On the other hand, it is important to note that Z and Z' are *qualitatively* different in a way undetected by norms. Indeed, Z' is built from essentially unary ingredients, whereas Z is necessarily binary. More precisely, if a function $f : V \times W \to [0,1]$ is decomposed as $g + h$, where h is supported on Z' (and g has some desirable properties), then after removing a small amount of V and W individually, one can assume $f = g$. On the other hand, if h were supported on Z, then one would need to remove some possibly complicated subset of $V \times W$, which represents a more drastic change to the nature of f.

To close this discussion, we will say that a function $f : V \times W \to [-1,1]$ has ε-**structured support** if its support is contained in a set of the form Z' above. Note that if f has ε-structured support, then $|\mathrm{supp}(f)| \leqslant 2\varepsilon|V \times W|$, and so, $\|f\|_2 \leqslant \sqrt{2\varepsilon}$.

[†] This is essentially equivalent to a suitable graph-theoretic analog of functional homogeneity (as in Definition 4.1), up to uniform change in ε.

Theorem 6.8. *Let V and W be finite sets, and suppose $f : V \times W \to [0,1]$ is a (k,δ)-stable function. Then for any $\varepsilon > 0$ and any function $\sigma : \mathbb{N} \to (0,1)$, there are $m, n \leqslant O_{k,\delta,\varepsilon,\sigma}(1)$ and a decomposition $f = f_{\text{str}} + f_{\text{psd}} + f_{\text{err}}$ such that:*

(i) *f_{str} is $(5\delta+\varepsilon; m, n)$-structured,*
(ii) *$|\text{supp}(f_{\text{psd}})| \leqslant \sigma(mn)|V \times W|$, and*
(iii) *f_{err} has ε-structured support.*

Moreover, f_{psd} is [-1,1]-valued, f_{str} and f_{err} are [0,1]-valued, and $f_{\text{err}} \in \{f_{\text{str}}, |f_{\text{psd}}|\}^{\perp}$.

Proof. Apply Theorem 5.1 to obtain partitions $V = V_0 \cup V_1 \cup \ldots \cup V_m$ and $W = W_0 \cup W_1 \cup \ldots \cup W_n$, with $m, n \leqslant O_{k,\delta,\varepsilon,\sigma}(1)$, so that $|V_0| \leqslant \varepsilon|V_1|$, $|W_0| \leqslant \varepsilon|W_1|$, and for all $(i,j) \in [m] \times [n]$, (V_i, W_j) is $(5\delta+\varepsilon; \frac{1}{2}\sigma(mn))$-homogeneous. For each $(i,j) \in [m] \times [n]$, choose $r_{i,j} \in [0,1]$ and $W_{i,j} \subseteq W_j$ such that $|W_{i,j}| \geqslant (1 - \frac{1}{2}\sigma(mn))|W_j|$ and, for all $b \in W_{i,j}$, $|\{a \in V_i : f(a,b) \approx_{5\delta+\varepsilon} r_{i,j}\}| \geqslant (1 - \frac{1}{2}\sigma(mn))|V_i|$.

Set $Z = (V_0 \times W) \cup (V \times W_0)$. Let $f_{\text{str}} : V \times W \to [0,1]$ be defined by

$$f_{\text{str}}(a,b) = \begin{cases} f(a,b) & \text{if } (a,b) \in V_i \times W_j \text{ and } f(a,b) \approx_{5\delta+\varepsilon} r_{i,j}, \\ r_{i,j} & \text{if } (a,b) \in V_i \times W_j \text{ and } f(a,b) \not\approx_{5\delta+\varepsilon} r_{i,j}, \text{ and} \\ 0 & \text{if } (a,b) \in Z. \end{cases}$$

Define $f_{\text{err}} := f\mathbf{1}_Z$ and $f_{\text{psd}} := f - f_{\text{str}} - f_{\text{err}}$. Then, other than (ii), the claims in the theorem follow easily by construction. So, we show (ii).

Set $\gamma = \frac{1}{2}\sigma(mn)$. Given $(i,j) \in [m] \times [n]$ and $b \in W$, partition $V_i = X^b_{i,j} \cup Y^b_{i,j}$ so that $a \in X^b_{i,j}$ if and only if $f(a,b) \approx_{5\delta+\varepsilon} r_{i,j}$. So, if $b \in W_{i,j}$, then $|Y^b_{i,j}| \leqslant \gamma|V_i|$ and $f_{\text{psd}}(a,b) = 0$ for all $a \in X^b_{i,j}$. It follows that

$$\text{supp}(f_{\text{psd}}) \subseteq \bigcup_{(i,j)\in[m]\times[n]} \Big((V_i \times (W_j \backslash W_{i,j})) \cup \bigcup_{b \in W_{i,j}} (Y^b_{i,j} \times \{b\}) \Big).$$

So, $|\text{supp}(f_{\text{psd}})| \leqslant \sum_{i,j} (\gamma|V_i||W_j| + \gamma|V_i||W_{i,j}|) \leqslant 2\gamma|V \times W| = \sigma(mn)|V \times W|$. □

Remark 6.9. The previous proof also implies that the underlying partition of f_{str} involves f-definable pieces of bounded complexity. We could instead use Theorem 6.1 and obtain an equipartition. (Note that we only need $|V_0| \leqslant \varepsilon|V|$ and $|W_0| \leqslant \varepsilon|W|$ to know that f_{err} has ε-structured support.)

Remark 6.10. Call a function $f : V \times W \to [0,1]$ *k-stable* if it is (k,δ)-stable for some $\delta > 0$ such that for all $x, y \in V \times W$, if $f(x) \neq f(y)$ then $|f(x) - f(y)| > 10\delta$. For example, this includes the setting of k-stable bipartite graphs (after changing k as discussed in Remark 1.2). Suppose $f : V \times W \to [0,1]$ is k-stable, and fix $V' \subseteq V$ and $W' \subseteq W$ such that (V', W') is $(5\delta+\varepsilon; \varepsilon)$-homogeneous for sufficiently small ε. Then, it follows that there is some $r \in \text{Im}(f)$ such that for all but at most $\varepsilon|W'|$-many $b \in W'$, for all but at most $\varepsilon|V'|$-many $a \in V'$, $f(a,b) = r$, *and dually*, for all but at most $\varepsilon|V'|$-many $a \in V'$, for all but at most $\varepsilon|W'|$-many $b \in W'$, $f(a,b) = r$. Altogether, one obtains a version of Theorem 6.8 for k-stable functions in which f_{str} is (m,n)-structured.

Remark 6.11. One can obtain a decomposition of (k,δ)-stable functions involving *no error term*, at the cost of the decay function. Indeed, in Theorem 6.8, if we replace f_{psd} with $f_{\text{psd}} + f_{\text{err}}$, then we have $|\text{supp}(f_{\text{psd}})| \leqslant (\sigma(mn) + 2\varepsilon)|V \times W|$. Note that we can also write $f_{\text{str}} = f'_{\text{str}} + h$ where f'_{str} is (m,n)-structured and $\|h\|_\infty \leqslant 5\delta + \varepsilon$. So, if one is willing to allow f_{psd} to involve δ, then f_{str} can be made perfectly structured by replacing f_{psd} with $f_{\text{psd}} + h$. In this case, f_{psd} no longer has small support, but one still has a bound on $\|f_{\text{psd}}\|_1$ in terms of δ, ε, and $\sigma(mn)$. Therefore, f_{psd} is still pseudorandom in a strong qualitative sense (see Remark 6.6).

6.4 | Bipartite versus symmetric

In this section, we clarify some subtleties regarding our bipartite viewpoint on graphs and functions. First, recall that any graph $(V;E)$ can be "coded" as bipartite graph $(V,V;E)$ (sometimes called the *bipartite double cover* of $(V;E)$). Similarly, a [0,1]-valued binary function f on a set V can be viewed as a bipartitioned function $f : V \times V \to [0,1]$. From this perspective, the bipartite setting is more flexible because it allows one to distinguish between the two sets. This is also a very natural setting in which to apply model-theoretic tools. On the other hand, there is one issue with the bipartite approach, which is usually not discussed in the model-theoretic literature on regularity. In particular, given a graph $(V;E)$ if one applies a bipartite regularity lemma to $(V,V;E)$, then this results in two potentially different partitions of V.

One way to remedy this issue is to again exploit the decay function. In particular, suppose that we have partitions $V_1 \cup \ldots \cup V_m$ and $W_1 \cup \ldots \cup W_n$ of the same finite set V, in which each pair (V_i, W_j) satisfies a desired homogeneity property with respect to some $f : V \times V \to [0,1]$ (we ignore the issue of exceptional sets for the moment). Then, we have a common refinement, which partitions V into at most mn sets. Thus, given a target decay function $\sigma : \mathbb{N} \to (0,1)$, one can define a modified function τ (similar to the proof of Theorem 6.1) so that if the initial partitions are homogeneous with respect to τ, then any sufficiently large piece of the common refinement maintains homogeneity using σ, while the remaining small pieces can be put into an exceptional set. Moreover, if we also start out with two exceptional sets V_0 and W_0, then these can be added to the larger exceptional set, along with any W_j that intersects V_0 in a large set (and vice versa).

As a final remark, we note that when $f : V \times V \to [0,1]$ is a *symmetric* function, one can also address this issue "upstairs" at the model-theoretic stage. Indeed, given a continuous structure M and a symmetric [0,1]-valued formula $\varphi(x,y)$, with x and y variables of the same sort, the type spaces $S_\varphi(M)$ and $S_{\varphi^*}(M)$ can be naturally identified. Therefore, in the setting of Lemma 4.4, if the measures μ and ν are the same (e.g., if they both arise from the pseudofinite average value functional), then one can prove the lemma using *the same* φ-formulas on both sides. Carrying this through the rest of the steps, we obtain a single partition of V.

APPENDIX

We discuss a natural variation of (k,δ)-stability (as defined in Definition 1.1), which is closer to the definition of k-stability for bipartite graphs.

Definition A.1. A function $f : V \times W \to [0,1]$ is $^*(k,\delta)$-**stable** if there do not exist $a_1, \ldots, a_k \in V$, $b_1, \ldots, b_k \in W$, and $r \in [0,1]$ such that $f(a_i, b_j) \geqslant r + \delta$ if $i \leqslant j$, and $f(a_i, b_j) \leqslant r$ if $i > j$.

Note that a binary relation $E \subseteq V \times W$ is k-stable (as defined in Remark 1.2) if and only if $\mathbf{1}_E$ is $^*(k,1)$-stable (if and only if $\mathbf{1}_E$ is $^*(k,\delta)$-stable for all $\delta > 0$). As we previously observed, for "discretely valued" functions such as $\mathbf{1}_E$, $^*(k,\delta)$-stability implies (ℓ,δ)-stability for suitable ℓ. On the other hand, there is a small discrepancy for arbitrary functions.

Proposition A.2.

(a) *Any (k,δ)-stable [0,1]-valued function is $^*(k,\delta)$-stable.*
(b) *For any $k \geqslant 1$ and $\delta' > \delta > 0$, there is some $\ell \geqslant 1$ such that any $^*(k,\delta)$-stable [0,1]-valued function is (ℓ,δ')-stable.*

Proof. Part (a). This is straightforward to check.

Part (b). Given $m, n \geqslant 1$, let $R_m(n)$ be an integer such that any m-coloring of $\binom{[R_m(n)]}{2}$ admits a monochromatic subset of size n. Fix $k \geqslant 1$ and $\delta' > \delta > 0$. Set $\varepsilon = \delta' - \delta$ and fix $m \geqslant \varepsilon^{-1}$. Set $\ell = R_2(R_m(2k+1))$. Suppose $f : V \times W \to [0,1]$ is not (ℓ,δ')-stable. We will show that f is not $^*(k,\delta)$-stable.

Fix $a_1, \ldots, a_\ell \in V$ and $b_1, \ldots, b_\ell \in W$ such that $|f(a_i,b_j) - f(a_j,b_i)| \geqslant \delta'$ for all $i < j$. Consider a 2-coloring of $\binom{[\ell]}{2}$ according to whether $f(a_i,b_j) \geqslant f(a_j,b_i) + \delta'$ or $f(a_i,b_j) \leqslant f(a_j,b_i) - \delta'$. Set $n = R_m(2k+1)$. By choice of ℓ, and after relabeling and reversing the order (if necessary), we may assume that there are $a_1, \ldots, a_n \in V$ and $b_1, \ldots, b_n \in W$ such that $f(a_i,b_j) \geqslant f(a_j,b_i) + \delta'$ for all $i < j$.

For $1 \leqslant t \leqslant m$, set $I_t = [\frac{t-1}{m}, \frac{t}{m}]$. Consider an m-coloring of $\binom{[n]}{2}$ by the minimal t such that $f(a_j,b_i) \in I_t$ (for $i < j$). By choice of n, and after relabeling, we may assume that we have $1 \leqslant t \leqslant m$, $a_1, \ldots, a_{2k+1} \in V$, and $b_1, \ldots, b_{2k+1} \in W$ such that for all $i < j$, $f(a_i,b_j) \geqslant f(a_j,b_i) + \delta'$ and $f(a_j,b_i) \in I_t$. Set $r = \frac{t}{m}$. Then, for all $i < j$, we have $f(a_j,b_i) \leqslant r$ and $f(a_i,b_j) \geqslant r - \frac{1}{m} + \delta' \geqslant r + \delta$. Altogether, the sequences $(a_{2i})_{i=1}^k$ and $(b_{2i+1})_{i=1}^k$ witness that f is not $^*(k,\delta)$-stable. □

In the context of a complete theory, however, the discrepancy can be removed.

Corollary A.3. *Let T be a complete theory in a continuous language $\mathcal{L}$, and fix a [0,1]-valued $\mathcal{L}$-formula $\varphi(x,y)$. Fix $\delta > 0$, and assume that $\varphi(x,y)$ is $^*(k,\delta)$-stable in T for some $k \geqslant 1$. Then $\varphi(x,y)$ is (ℓ,δ)-stable in T for some $\ell \geqslant 1$.*

Proof. First, one checks that $^*(k,\delta)$-stability in T is an open condition in the sense of Lemma 3.2. So, there is some $\delta_0 < \delta$ such that $\varphi(x,y)$ is $^*(k,\delta_0)$-stable in T. By Proposition A.2, $\varphi(x,y)$ is (ℓ,δ)-stable in T for some $\ell \geqslant 1$. □

The previous corollary also follows from [4, Lemma 7.2], whose proof involves similar Ramsey arguments.

ACKNOWLEDGEMENT

The authors were supported by NSF grants DMS-1665025, DMS-1665035, DMS-1760212, and DMS-2054271 (Pillay), and DMS-1855503 (Conant). Much of the work in this paper was carried out during the 2021 Thematic Program on Trends in Pure and Applied Model Theory at the Fields Institute, where Pillay was appointed as a Simons Distinguished Visitor. Research funds from this

appointment were also used to support Chavarria and Conant at Fields. Chavarria's Notre Dame stipend for Fall 2021 was provided through DMS-1665025.

We thank the Fields Institute for their hospitality. The second author would also like to thank J. Hanson for helpful conversations. Special thanks are due to C. Terry for a number of invaluable discussions, and for suggesting Definition 6.7.

JOURNAL INFORMATION

The *Journal of the London Mathematical Society* is wholly owned and managed by the London Mathematical Society, a not-for-profit Charity registered with the UK Charity Commission. All surplus income from its publishing programme is used to support mathematicians and mathematics research in the form of research grants, conference grants, prizes, initiatives for early career researchers and the promotion of mathematics.

ORCID

Nicolas Chavarria https://orcid.org/0000-0001-8888-6296

REFERENCES

1. I. Ben Yaacov, *Topometric spaces and perturbations of metric structures*, Log. Anal. **1** (2008), no. 3–4, 235–272. MR 2448260.
2. I. Ben Yaacov, *Model theoretic stability and definability of types, after A. Grothendieck*, Bull. Symbolic Logic **20** (2014), no. 4, 491–496. MR 3294276.
3. I. Ben Yaacov, A. Berenstein, C. W. Henson, and A. Usvyatsov, *Model theory for metric structures*, Model theory with applications to algebra and analysis. Vol. 2, London Math. Soc. Lecture Note Ser., vol. 350, Cambridge University Press, Cambridge, 2008, pp. 315–427. MR 2436146 (2009j:03061).
4. I. Ben Yaacov and A. Usvyatsov, *Continuous first order logic and local stability*, Trans. Amer. Math. Soc. **362** (2010), no. 10, 5213–5259. MR 2657678.
5. A. Chernikov and H. Towsner, *Hypergraph regularity and higher arity VC-dimension*, arXiv:2010.00726, 2020.
6. G. Elek and B. Szegedy, *A measure-theoretic approach to the theory of dense hypergraphs*, Adv. Math. **231** (2012), no. 3–4, 1731–1772. MR 2964622.
7. W. T. Gowers, *Quasirandom groups*, Combin. Probab. Comput. **17** (2008), no. 3, 363–387. MR 2410393.
8. B. Green and T. Tao, *An arithmetic regularity lemma, an associated counting lemma, and applications*, An irregular mind, Bolyai Soc. Math. Stud., vol. 21, János Bolyai Math. Soc., Budapest, 2010, pp. 261–334. MR 2815606.
9. A. Grothendieck, *Critères de compacité dans les espaces fonctionnels généraux*, Amer. J. Math. **74** (1952), 168–186. MR 0047313.
10. E. Hrushovski, *Pseudo-finite fields and related structures*, Model theory and applications, Quad. Mat., vol. 11, Aracne, Rome, 2002, pp. 151–212. MR 2159717 (2006d:03059).
11. E. Hrushovski, *Stable group theory and approximate subgroups*, J. Amer. Math. Soc. **25** (2012), no. 1, 189–243. MR 2833482.
12. E. Hrushovski and A. Pillay, *Groups definable in local fields and pseudo-finite fields*, Israel J. Math. **85** (1994), no. 1–3, 203–262. MR 1264346.
13. H. J. Keisler, *Measures and forking*, Ann. Pure Appl. Logic **34** (1987), no. 2, 119–169. MR 890599.
14. B. Kim and A. Pillay, *Simple theories*, Ann. Pure Appl. Logic **88** (1997), no. 2–3, 149–164, Joint AILA-KGS Model Theory Meeting (Florence, 1995). MR 1600895 (99b:03049).
15. J. Komlós and M. Simonovits, *Szemerédi's regularity lemma and its applications in graph theory*, Combinatorics, Paul Erdős is eighty, Vol. 2 (Keszthely, 1993), Bolyai Soc. Math. Stud., vol. 2, János Bolyai Math. Soc., Budapest, 1996, pp. 295–352. MR 1395865
16. J.-L. Krivine and B. Maurey, *Espaces de Banach stables*, Israel J. Math. **39** (1981), no. 4, 273–295. MR 636897.
17. L. Lovász and B. Szegedy, *Szemerédi's lemma for the analyst*, Geom. Funct. Anal. **17** (2007), no. 1, 252–270. MR 2306658

18. M. Malliaris and A. Pillay, *The stable regularity lemma revisited*, Proc. Amer. Math. Soc. **144** (2016), no. 4, 1761–1765. MR 3451251.
19. M. Malliaris and S. Shelah, *Regularity lemmas for stable graphs*, Trans. Amer. Math. Soc. **366** (2014), no. 3, 1551–1585. MR 3145742.
20. A. Pillay, *Pseudofinite model theory*, Lecture notes, Fall 2014, https://www3.nd.edu/~apillay/notes_greg-final.pdf.
21. A. Pillay, *Generic stability and Grothendieck*, South Amer. J. Log. **2** (2016), no. 2, 437–442. MR 3671044.
22. A. Pillay, *Domination and regularity*, Bull. Symbolic Logic **26** (2020), no. 2, 103–117. MR 4222408.
23. A. Pillay and S. Starchenko, *Remarks on Tao's algebraic regularity lemma*, arXiv:1310.7538, 2013, research note.
24. W. Rudin, *Functional analysis*, 2nd ed., International Series in Pure and Applied Mathematics, McGraw-Hill, Inc., New York, 1991. MR 1157815.
25. E. Szemerédi, *Regular partitions of graphs*, Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976), Colloq. Internat. CNRS, vol. 260, CNRS, Paris, 1978, pp. 399–401. MR 540024.
26. T. Tao, *Structure and randomness in combinatorics*, FOCS 2007 lecture notes, arXiv:0707.4269.
27. T. Tao, *Nonstandard analysis as a completion of standard analysis*, What's new (blog), November 27, 2010, http://terrytao.wordpress.com/2010/11/27/nonstandard-analysis-as-a-completion-of-standard-analysis/.
28. T. Tao, *A spectral theory proof of the algebraic regularity lemma*, What's new (blog), October 29, 2013, https://terrytao.wordpress.com/2013/10/29/a-spectral-theory-proof-of-the-algebraic-regularity-lemma/.
29. T. Tao, *Expanding polynomials over finite fields of large characteristic, and a regularity lemma for definable sets*, Contrib. Discrete Math. **10** (2015), no. 1, 22–98. MR 3386249.
30. C. Terry and J. Wolf, *Irregular triads in 3-uniform hypergraphs*, arXiv:2111.01737, 2021.