

INDEPENDENCE, ORDER, AND THE INTERACTION OF ULTRAFILTERS AND THEORIES

M. E. MALLIARIS

ABSTRACT. We consider the question, of longstanding interest, of realizing types in regular ultrapowers. In particular, this is a question about the interaction of ultrafilters and theories, which is both coarse and subtle. By our prior work it suffices to consider types given by instances of a single formula. In this article, we analyze a class of formulas φ whose associated characteristic sequence of hypergraphs can be seen as describing realization of first- and second-order types in ultrapowers on one hand, and properties of the corresponding ultrafilters on the other. These formulas act, via the characteristic sequence, as points of contact with the ultrafilter \mathcal{D} , in the sense that they translate structural properties of ultrafilters into model-theoretically meaningful properties and vice versa. Such formulas characterize saturation for various key theories (e.g. T_{rg} , T_{feq}), yet their scope in Keisler's order does not extend beyond T_{feq} . The proof applies Shelah's classification of second-order quantifiers.

1. INTRODUCTION

Regular ultrafilters and countable first-order theories are both, a priori, quite complicated objects. And yet the mystery is that their interaction is often quite coarse, for reasons that have to do both with model theory and with the structure of ultrafilters. In this article, we build and investigate a framework in which this interaction is visible: namely, certain sequences of hypergraphs whose solution (in a sense defined below) alternately describes realization of first-order types in ultrapowers, existence of second-order structure in those same ultrapowers and structural properties of the corresponding ultrafilters.

To begin, in §2, we motivate these investigations by looking at regular ultrapowers of countable stable theories. We then define Keisler's order and give the known results. We explain how saturation of ultrapowers can be analyzed in terms of characteristic sequences, i.e. sequences of hypergraphs defined on the parameter space of first-order formulas, following our prior work [6], [8]. We recall two relevant classical “dichotomies” between order and independence in unstable theories, namely, independence/strict order and TP_1/TP_2 . In [8] we showed that the “independent” half in each of these cases has a Keisler-minimal theory of a certain simple form; namely, a theory in which key formulas capture a certain interaction between types and ultrafilters which, in this paper, we abstract and investigate.

Motivation 2.20 explains the objects of study in this paper, a basic but rich class of formulas (or, more generally, their associated hypergraph sequences) called “fundamental” because they have model-theoretic significance on one hand and capture properties of ultrafilters on the other. §3 gives an example to show that analyzing such formulas depends on understanding a certain interaction between first- and second-order structure in ultrapowers. §4 gives a formal correspondence: each fundamental formula φ can be associated to a “second-order quantifier” in the sense of [11] via its characteristic sequence. That is, for any regular ultrafilter \mathcal{D} , \mathcal{D} realizes all φ -types over small sets if and only if \mathcal{D} solves its associated quantifier, in the sense defined there. However, the real interest of connecting our investigations to [11] is its proof that upto interpretability, there are very few such quantifiers. In §5, we show that with some care, many of the interpretability arguments

Partially supported by NSF grant DMS-1001666.

of [11] can be translated to ultrapowers. Applying these results, we show that any fundamental formula is dominated (in the sense of Keisler’s order) either by the empty theory, by the random graph, or by the minimal TP_2 theory. In Theorem 5.21 we prove, among other things, that the scope of the second-order quantifiers (and therefore the fundamental formulas) in Keisler’s order does not extend beyond TP_2 . In §6 we apply the prior analysis to prove Theorem 6.1, which points towards a possible gap in complexity between independence and strict order.

Throughout this paper, variables and parameters written without an overline need not necessarily have length 1.

CONTENTS

1. Introduction	1
2. Background and Context	2
2.1. Regularity of ultrafilters	3
2.2. Keisler’s order on unstable theories	5
2.3. Results of our prior work and motivation for fundamental formulas	7
3. The example of the random graph	8
3.1. Analysis of the example	9
4. Second-order quantifiers	10
5. Interpretability	13
6. = and <, or independence and order	18
7. Appendix: Construction of theories	20
References	23

2. BACKGROUND AND CONTEXT

A regular ultrapower is a regular ultraproduct, i.e. a reduced product where equivalence is computed modulo some regular ultrafilter \mathcal{D} , in which the index models M_i are all taken to be the same. (“Regular” is Definition 2.4 below, but may not be necessary for these first remarks.) A model is λ^+ -saturated if it realizes all types over sets of size λ .

As mentioned in the introduction, the interaction of regular ultrafilters and countable first-order theories is often quite coarse; it is interesting to examine why and how. An illustrative case is the following. Suppose we are given two regular ultrafilters $\mathcal{D}_1, \mathcal{D}_2$ on λ and two countable stable theories T_1, T_2 . Let $M_1 \models T_1, M_2 \models T_2$. Can it happen that saturation of their ultrapowers is independent, i.e. that if we ask whether $M_i^\lambda/\mathcal{D}_j$ ($i, j \leq 2$) is λ^+ -saturated, the answer is

	\mathcal{D}_1	\mathcal{D}_2	
M_1	yes	no	
M_2	no	yes	?

Surprisingly, it cannot ([12].VI.5). In fact, saturation of regular ultrapowers of stable theories depends on one parameter: the minimal size (modulo \mathcal{D}) of a pseudofinite set. More precisely, it is a theorem of classification theory that a model of a countable stable theory is κ^+ -saturated if and only if it is \aleph_1 -saturated and every maximal indiscernible set has size greater than κ . (This relies heavily on uniqueness of nonforking extensions, so fails in unstable theories.) Since any nonprincipal ultrapower, in particular a regular ultrapower, is \aleph_1 -saturated, it suffices to show that every maximal indiscernible set is sufficiently large. One can show that this will be true precisely when the minimal size, modulo \mathcal{D} , of the product of an unbounded sequence of natural numbers is strictly greater than λ . Call this minimal size $\mu(\mathcal{D})$.

On the other hand, suppose M is a model of an equivalence relation E with a class of size n for each $n \in \mathbb{N}$. Suppose that for some sequence of natural numbers n_t , we have that $\prod_{t < \lambda} n_t / \mathcal{D} = n_*$ where $n_* \leq \lambda$. Then the ultrapower $N = M^\lambda / \mathcal{D}$ contains an equivalence class of size n_* . Letting $\langle a_i : i < n_* \rangle$ list the elements of this class, the type $\{E(x, a_i) \wedge x \neq a_i : i < n_*\}$ has size $\leq \lambda$ but is omitted in N . In general, say that a formula φ has the *finite cover property* if for arbitrarily large n there is a set $\{\varphi(x; \bar{a}_0), \dots, \varphi(x; \bar{a}_n)\}$ of instances of φ which is inconsistent but whose n -element subsets are all consistent. Saturation of regular ultrapowers of countable stable theories can be described as follows:

Theorem A. (Shelah [12] VI.5) *Let T be a countable stable theory and λ an infinite cardinal.*

- *If T does not contain a formula with the finite cover property and \mathcal{D} is any regular ultrafilter on λ , then for any $M \models T$, M^λ / \mathcal{D} is λ^+ -saturated.*
- *If T does contain a formula with the finite cover property and \mathcal{D} is any regular ultrafilter on λ , then for any $M \models T$, we have that M^λ / \mathcal{D} is λ^+ -saturated if and only if the minimum product of an unbounded sequence of natural numbers modulo \mathcal{D} , i.e. $\mu(\mathcal{D})$, is at least λ^+ .*

This theorem gives the only two known classes in Keisler's order. More precisely, define:

Definition 2.1. (Keisler 1967 [2]) *Let T_1, T_2 be countable theories and λ be an infinite cardinal.*

- (1) $T_1 \leq_\lambda T_2$ means: for every $M_1 \models T_1, M_2 \models T_2$, and for every regular ultrafilter \mathcal{D} on λ , if $(M_2)^\lambda / \mathcal{D}$ is λ^+ -saturated, then $(M_1)^\lambda / \mathcal{D}$ is λ^+ -saturated.
- (2) (Keisler's order) $T_1 \leq T_2$ means that $T_1 \leq_\lambda T_2$ for all $\lambda \geq \aleph_0$.

The hypothesis of regularity justifies the quantification over all models: if $M \equiv N$ and \mathcal{D} is a regular ultrafilter on λ , then M^λ / \mathcal{D} is λ^+ -saturated iff N^λ / \mathcal{D} is, Fact 2.5 below. Definition 2.1(2) describes a preorder on the countable first-order theories, often thought of as a partial order on the \leq -equivalence classes. The major question is then:

Problem 2.2. [2] *Determine the structure of Keisler's order.*

Convention 2.3. *In the discussion of ultrapowers which follows, we will generally use M for the index model, identify the index set with its cardinality λ , and use N for the ultrapower M^λ / \mathcal{D} . We fix in advance a canonical representative $a \in M^\lambda$ of each \mathcal{D} -equivalence class, so that the projection of a given element $c \in N$ to its value in the t -th index model, denoted $c[t]$, is well defined. Since we primarily consider ultrapowers, not ultraproducts, in proofs we write e.g. " $M \models \exists x \varphi(x; c[t])$ " rather than using $M[t]$ to distinguish the t -th copy of M .*

2.1. Regularity of ultrafilters. Before continuing our discussion of Keisler's order, we consider the hypothesis "regular."

Definition 2.4. *An ultrafilter \mathcal{D} on I , $|I| = \lambda \geq \aleph_0$ is said to be regular if there exists $X = \{X_j : j < \lambda\} \subseteq \mathcal{D}$, called a regularizing family, with the property that for any $\sigma \subset \lambda$,*

$$\bigcap_{j \in \sigma} X_j \neq \emptyset \quad \text{iff} \quad |\sigma| < \aleph_0$$

Equivalently, for any index $t \in I$, t belongs to only finitely many of the X_j . For an extensive discussion of regular ultrapowers and Keisler's order, with many examples, the interested reader is referred to the paper [8], as well as the foundational sources [2] and [12] Chapter VI. However, for completeness, we summarize several well known properties here.

Regular ultrafilters exist on any infinite cardinal (given $|I| = \lambda$ and any bijection $f : \mathcal{P}_{\aleph_0}(\lambda) \rightarrow I$, notice that $\{\{t \in I : \eta \in f^{-1}(t)\} : \eta \in \lambda\}$ can be extended to an ultrafilter). Moreover, the degree of saturation of regular ultrapowers is a property of the theory and not of the index model:

Fact 2.5. (Keisler [2] Corollary 2.1 p. 30; see also Shelah [12].VI.1) *If $M \equiv M'$ and \mathcal{D} is a regular ultrafilter on λ , then M^λ/\mathcal{D} is λ^+ -saturated iff $(M')^\lambda/\mathcal{D}$ is λ^+ -saturated.*

Thus the quantification over all models in Definition 2.1 makes sense. A related fact is that saturation in regular ultrapowers reduces to satisfying a growing, but *almost everywhere finite*, set of conditions in each index model:

Fact 2.6. *Consider a type in a regular ultrapower. That is, suppose we are given $M \models T$, \mathcal{D} a regular ultrafilter on λ , $N = M^\lambda/\mathcal{D}$, $A \subseteq N$, $|A| \leq \lambda$, $p \in S(A)$. Let $\mathcal{P}_{\aleph_0}(p)$ denote the set of finite subsets of p . Then:*

- (a) *There exists a map $d : \mathcal{P}_{\aleph_0}(p) \rightarrow \mathcal{D}$, whose image is a regularizing family, such that:*
- *d is monotonic, i.e. $u \subseteq v$ implies $d(v) \subseteq d(u)$*
 - *d refines the Loś map, i.e. if $t \in d(u)$ then $M \models \exists x \bigwedge_{\varphi(x;a) \in u} \varphi(x; a[t])$*

Moreover for any such d , called a distribution, the following are equivalent:

- (b) *There exists a distribution d' refining d which is multiplicative, i.e. $d'(u) \cap d'(v) = d'(u \cup v)$.*
(c) *The type p is realized in N .*

Proof. (a) Let $\langle X_i : i < \lambda \rangle$ be a regularizing family and for each $u \in \mathcal{P}_{\aleph_0}(p)$, let

$$f_1(u) = \{t < \lambda : M \models \exists x \bigwedge_{\varphi(x;a) \in u} \varphi(x; a[t])\}$$

be the Loś map. Fix a bijection $g : \lambda \rightarrow \mathcal{P}_{\aleph_0}(p)$, and let $f_2(u) = f_1(u) \cap X_i$, where $i = g^{-1}(u)$. Finally, to ensure monotonicity, define (by induction on the size of $|u|$) the desired distribution $d(u) = f_2(u) \cap \bigcap \{d(v) : v \subsetneq u\}$, which remains \mathcal{D} -large since u is finite.

For the “moreover” clause, notice that realizing the type depends simply on whether or not it is almost everywhere true under some distribution that the projections $\varphi_{i_1}(x; a_{i_1}[t]), \dots, \varphi_{i_n}(x; a_{i_n}[t])$ of the finitely many formulas $\varphi_{i_1}(x; a_{i_1}), \dots, \varphi_{i_n}(x; a_{i_n})$ assigned to the index t have a *common* solution, that is, whether or not the distribution can be chosen to be multiplicative. More precisely, if (c) holds, let c realize the type and let f_1 send any finite subset of p to the set of indices t on which it is realized by $c[t]$. This map is multiplicative, but its image may not be a regularizing set. To obtain a distribution, first fix an enumeration $g : \lambda \rightarrow \mathcal{P}_{\aleph_0}(p)$. Let $f_2(u) = f_1(u) \cap X_{g^{-1}(u)}$, so the image of f_2 is a regularizing set. We now recover multiplicativity as follows: for each $u \in \mathcal{P}_{\aleph_0}(p)$, $|u| \geq 1$, define

$$f(u) = \{t < \lambda : \text{for each } \varphi(x; a) \in u, t \in f_2(\{\varphi(x; a)\})\}$$

Since $t \in f_2(\{\varphi(x; a)\})$ implies $M \models \varphi(c[t]; a[t])$, f refines the Loś map, and it is clear that f assigns only finitely many subsets u to each index model. So f is a multiplicative distribution as desired.

If (b) holds, choose a common solution c_t in each index model; by Loś’ theorem and the definition of distribution, $\prod_{t < \lambda} c_t/\mathcal{D}$ will realize the type. \square

In [5] we proved that the reductions just explained for regular ultrapowers – namely that first, saturation depends on the theory and not the model chosen, and second, that realizing types depends on almost everywhere finite projections to the index models – can be further reduced to the study of types in finite fragments of the language:

Theorem 2.7. (Malliaris [5]) *Let \mathcal{D} be a regular ultrafilter on $\lambda \geq \aleph_0$ and let T be a countable theory, $M \models T$. Then M^λ/\mathcal{D} is λ^+ -saturated iff, for all formulas φ , M^λ/\mathcal{D} realizes all φ -types over sets of size $\leq \lambda$. That is, local saturation implies saturation.*

For types in a single formula, the complexity of distributions (see Fact 2.6 above) can be usefully abstracted by the following simple representation of a *characteristic sequence*, which we now explain. These hypergraphs were introduced and developed in our prior work [6], [7], [8] as a context where graph-theoretic arguments could be applied to give model-theoretic information, but we will introduce ideas from those papers only as they are needed.

Definition 2.8. (Characteristic sequences, Malliaris [6]) *Let T be a first-order theory and φ a formula of the language of T .*

- For $n < \omega$, $P_n(z_1, \dots, z_n) := \exists x \bigwedge_{i \leq n} \varphi(x; z_i)$.
- The characteristic sequence of φ in T is $\langle P_n : n < \omega \rangle$.
- Write $(T, \varphi) \mapsto \langle P_n \rangle$ for this association.

Without loss of generality, we will identify the predicates P_n with their interpretation in the monster model. When it is important to specify the model, write P_n^M for interpretation in some given model M .

Remark 2.9. *In practice, when computing the characteristic sequence, we will often choose formulas of the form $\varphi(x; y) \wedge \neg \varphi(x; z)$, or $\theta(x; y, z, w) = ((z = w) \wedge \varphi(x; y)) \vee ((z \neq w) \wedge \neg \varphi(x; y))$. For instance, in the random graph, it is the characteristic sequence of $\varphi(x; y, z) = xRy \wedge \neg xRz$, not that of xRy , which captures the essential complexity. The characteristic sequence accurately describes “positive” partial types, as the next remark shows; so if the formula chosen can code negation, we can describe consistent partial types.*

Remark 2.10. *Note that for any characteristic sequence, and any $A \subset P_1$, the following are equivalent:*

- (1) $A^n \subset P_n$ for all n
- (2) $\{\varphi(x; a) : a \in A\}$ is a consistent partial type.

We will call a set $A \subset P_1$ satisfying either of these conditions a P_∞ -complete graph, or equivalently a positive base set to emphasize its connection with types.

In the context of ultrapowers, the analogue of Fact 2.6 is simply:

Fact 2.11. (Malliaris [8] Lemma 4.8) *Let \mathcal{D} be a regular ultrafilter on λ and $N = M^\lambda/\mathcal{D}$ an ultrapower. The following are equivalent for any characteristic sequence and any positive base set $A \subseteq N$:*

- (1) The type $p(x) = \{\varphi(x; a) : a \in A\}$ corresponding to A is realized.
- (2) There exists $d : A \rightarrow \mathcal{D}$ whose image is a regularizing family and such that writing $A[t]$ for $\{a[t] : t \in d(a)\}$, we have that almost everywhere, $A[t]$ is a $P_{|A[t]|}$ -complete graph in M .

Without loss of generality, for any positive base set A , if $d : A \rightarrow \mathcal{D}$ then we may assume d refines $a \mapsto \{t < \lambda : a[t] \in P_1^M\}$.

Definition 2.12. *Let \mathcal{D} be a regular ultrafilter on λ , $\langle P_n \rangle$ a characteristic sequence, and $N = M^\lambda/\mathcal{D}$ an ultrapower.*

- (1) Call $A \subseteq N$ small if $|A| \leq \lambda$.
- (2) If Fact 2.11(2) holds for every small positive base set $A \subset N$, say that \mathcal{D} solves $\langle P_n \rangle$.

2.2. Keisler’s order on unstable theories. By Theorem A, stable theories fall into precisely two classes in Keisler’s order. However, despite much progress in the model-theoretic understanding of unstable theories, the structure of Keisler’s order on unstable theories has remained elusive. Before stating one further result, Theorem C, we recall a classic structure/randomness tradeoff for unstable theories.

Definition 2.13. (Stability and the order property) *A theory is said to be stable if it does not contain a formula with the order property. We say that: (note that $\ell(x), \ell(y)$ need not be 1)*

- $\varphi(x; y)$ has the order property with respect to the background theory T if it is consistent with T that there exist elements $\langle a_i : i < \omega \rangle$ such that for all $j < \omega$, the following partial type is consistent:

$$\{\neg\varphi(x; a_i) : i < j\} \cup \{\varphi(x; a_i) : i \geq j\}$$

This is often stated as: there exist elements $\langle c_i, a_i : i < \omega \rangle$ in some sufficiently saturated $M \models T$ such that $M \models \varphi(c_i, a_j)$ iff $i < j$.

Note that “ φ has the order property with respect to T ” can be expressed by countably many sentences.

Remark 2.14. *The definition of order property remains agnostic as to whether or not types which do not correspond to cuts over the a_i are consistent. A fundamental structural property of unstable theories is the following “dichotomy” (strictly speaking, the two possibilities are not mutually exclusive, but we use this word as they represent different ends of a spectrum).*

Theorem B. (A dichotomy above stable theories, [12] Chapter II §2,4) *Suppose T is unstable. Then either:*

- T contains a formula with the independence property, i.e. for some $\varphi(x; y)$ there exists a sequence $\langle a_i : i < \omega \rangle$ such that for any two disjoint finite $\sigma, \tau \subseteq \omega$, the following partial type is consistent:

$$\{\neg\varphi(x; a_i) : i \in \sigma\} \cup \{\varphi(x; a_i) : i \in \tau\}$$

- T contains a formula with the strict order property, i.e. for some $\psi(x; y)$ there exists a sequence $\langle a_i : i < \omega \rangle$ such that $\{\psi(x; a_i)\} \cup \{\neg\psi(x; a_j)\}$ is consistent if and only if $j < i$.

We can now state:

Theorem C. (Summary of known results on the structure of Keisler’s order, from introduction to [8])

- (1) *The theories without the finite cover property (FCP) are minimal in Keisler’s order. [2] (necessary), [12] (sufficient)*
- (2) *There is a dividing line between theories with and without FCP. [2], [12]*
- (3) *The stable theories with FCP are an equivalence class in Keisler’s order. [12]*
- (4) *There is a dividing line between stable and unstable theories. [12]*
- ⋮
- (5) *There is a maximum class, namely, the theories which are λ^+ -saturated iff the ultrafilter is λ^+ -good. The strict order property is sufficient for maximality. In fact, SOP_3 is sufficient for maximality; however, the model-theoretic identity of the maximal class is not known. [2], [12], [16], [1], [15]*

Theorem C explains our interest in the independence property in this article. Namely, among unstable theories, those with enough rigidity (i.e. those containing the strict order property of Theorem B) are viewed by ultrafilters as maximally complex. Thus, speaking informally, they sink to the bottom of the classification, whereas those whose complexity comes from many degrees of freedom as in the independence property do not. Before giving some of our results, we record one higher level of the dichotomy of Theorem B. Note that stable implies simple but the reverse does not hold. (Also, any theory with the strict order property or SOP_3 is not simple.)

Definition 2.15. *A theory is said to be simple if it does not contain a formula with the tree property, where this means: (note that $\ell(x), \ell(y)$ need not be 1)*

- $\varphi(x; y)$ has the tree property (strictly speaking, the 2-tree property) if it is consistent with T that there exist elements $\{a_\eta : \eta \in {}^\omega > \omega\}$ such that for any $\eta \in {}^\omega > \omega$, the elements of $\{\varphi(x; a_{\eta \restriction i}) : i < \omega\}$ are pairwise inconsistent but for every $\eta \in {}^\omega \omega$, the branch $\{\varphi(x; a_{\eta \restriction n}) : n < \omega\}$ is consistent.

That is, there is a tree of instances of φ such that paths correspond to consistent partial types, and the successors of any given node are pairwise inconsistent. Once again, notice that this definition remains agnostic as to the consistency of instances from incomparable nodes without a common immediate predecessor.

Here too, there is a very useful “dichotomy.”

Theorem D. (A dichotomy above simple theories, [13] Theorem 0.2 p. 177)

Suppose T is not simple. Then either:

- there are φ, a_η witnessing the tree property, such that furthermore, for any $\eta, \nu \in {}^\omega > \omega$ such that neither of η, ν is an initial segment of the other, we have that $\varphi(x; a_\eta), \varphi(x; a_\nu)$ are contradictory. In this case, say that φ has SOP_2 , sometimes called TP_1 . The strict order property implies SOP_2 .
- there are φ, a_n^ℓ ($\ell, n < \omega$) such that the elements of $\Gamma_\ell = \{\varphi(x; a_n^\ell) : n < \omega\}$ are pairwise inconsistent, but for any $\eta \in {}^\omega \omega$, the set $\{\varphi(x; a_{\eta \restriction \ell}^\ell) : \ell < \omega\}$ is consistent. In this case, say that φ has TP_2 .

Note that the second item is (a priori) stronger than simply adding a compatibility condition onto the tree property for nodes not already determined to be inconsistent. It says, roughly speaking, that we have countably many distinct sets of pairwise inconsistent choices, but these choices are independent.

2.3. Results of our prior work and motivation for fundamental formulas. We now connect the two structural “dichotomies” between independence and rigidity, Theorem B and Theorem D, to our prior work and our aims in this article.

Definition 2.16. T_{rg} is the theory of the Rado graph, here informally called the random graph, in the language with a single binary relation R . T_{feq} , studied in [14] 2.1, [1], [16], [8] is the model completion of the following theory: there are two sorts X, Y and a three-place relation $E(x, y, z)$ on $X \times Y \times Y$ such that for each $x \in X$, $E(x, y, z)$ is an equivalence relation on Y with infinitely many infinite classes. Write $E_x(y, z)$ to indicate that y, z are E_x -equivalent.

Theorem 2.17. (Malliaris [8]) There is a minimal unstable theory in Keisler’s order, namely the theory T_{rg} of the random (Rado) graph. There is a minimal TP_2 theory, namely the theory T_{feq} of a parametrized family of independent equivalence relations.

Remark 2.18. Theorem 2.17 shows that in both of the “independent” halves of the dichotomy results Theorem B and Theorem D, there is a Keisler-minimal theory which, moreover, has a particularly simple form. That is, both T_{rg} and T_{feq} contain formulas which assert the existence of certain partitions in the ultrapower. In ultrapowers of the random graph, realizing 1-types corresponds to finding elements R -related to all elements in some set A and not to any elements of B , which can be done exactly when the sets A, B can be almost everywhere separated, Example 3.3 below. Realizing 1-types in the formula $\varphi(x; y, z) = E_x(y, z)$ asks that the ultrafilter be able to separate larger families of sets, namely, the proposed equivalence classes, almost everywhere. This observation suggests the following definition.

Definition 2.19. (Fundamental formulas) Let φ be a formula and $\langle P_n : n < \omega \rangle$ its characteristic sequence with respect to a background theory T . Say that φ is fundamental if the following both hold:

- (1) $\langle P_n \rangle$ is $=$ -definable, i.e. for each P_n there is a formula $\nu_n(y_1, \dots, y_n)$ in the language of equality such that $\models P_n(a_1, \dots, a_n)$ iff $\models \nu_n(a_1, \dots, a_n)$.
- (2) $\langle P_n \rangle$ has finite support, i.e. there is $k < \omega$ such that for all $n > k$, $P_n(y_1, \dots, y_n)$ iff $P_k(y_{i_1}, \dots, y_{i_k})$ holds on all k -element subsets $\{i_1, \dots, i_k\} \subseteq \{1, \dots, n\}$.

We will also use the word *fundamental* to refer simply to characteristic sequences satisfying these two conditions.

Motivation 2.20. The apparent nature of Keisler’s order, in which certain paradigmatic configurations serve as points of contact between theories (that is, they correspond to realization of certain φ -types) and ultrafilters (that is, they make clear demands on certain regularizing families), suggest that it is potentially very useful to define and classify the formulas exhibiting the “particularly simple form” of Remark 2.18. Definition 2.19 proposes a formal description of this class. Informally speaking, we look for formulas which, (α) on the one hand, have model-theoretic significance: e.g. $\varphi(x; y, z) = xRy \wedge \neg xRz$ in the random graph captures the independence property, whereas $E_x(y, z) \wedge \neg E_u(v, w)$ captures TP_2 , which can be thought of as independence in the presence of dividing. On the other hand, (β) we ask that realization of φ -types be expressible in terms of conditions on the filters: e.g. partitioning or separating certain families of sets in the ultrapower. What motivates this work is the observation that classifying these fundamental formulas is likely to be useful both for establishing equivalence classes and for determining dividing lines in Keisler’s order:

- (1) (Equivalence classes) *Prior results, e.g. Theorem A suggest that such paradigmatic examples may bring many other theories with them. Such paradigmatic examples tend, informally speaking, to provide the ultrafilter with clear invariants, like partition properties or growth rates of nonstandard integers, which it may then manipulate.*
- (2) (Dividing lines) *Any equivalence established between realization of certain types and properties of ultrafilters helps the construction problem by isolating structural properties of ultrafilters with model-theoretic significance. It is in principle easier to build a filter which, say, can separate pairs of sets but not families of sets, than to try, before this reduction, simply to saturate models of the random graph while not saturating some non-simple theory.*

Remark 2.21. *We remark that Definition 2.19 satisfies (α) and (β) from Motivation 2.20 above in the following sense:*

- (α) *The model-theoretic significance of the fundamental formulas is to give a series of “independence properties” of potentially increasing complexity. That is, they assert that sets of instances of φ always have a common witness provided that there is no explicit contradiction in the parameters as given by the formulas ν_n .*
- (β) *Each fundamental formula corresponds to an assertion about induced structure in the ultrapower by Claim 4.10 below.*

3. THE EXAMPLE OF THE RANDOM GRAPH

A first key observation is that what is at stake in studying the fundamental formulas is the relation of first-order and second-order structure in regular ultrapowers, which the following example illustrates. Suppose that M is the countable model of the theory T_{rg} of the random (i.e. Rado) graph in the language $\mathcal{L} = \{=, R(x, y)\}$. Let \mathcal{D} be a regular ultrafilter on λ . Recall that ultrapowers commute with reducts:

Fact 3.1. *Let M be a model of signature \mathcal{L} , $\mathcal{L}_0 \subset \mathcal{L}$ and \mathcal{D} an ultrafilter on λ . Then*

$$\left(M^\lambda / \mathcal{D} \right) |_{\mathcal{L}_0} = (M |_{\mathcal{L}_0})^\lambda / \mathcal{D}$$

Thus any ultrapower $N = M^\lambda/\mathcal{D}$ can naturally be expanded to a model of the full theory of M (the theory in a language of size $2^{|M|}$ in which all possible relations, functions and constants are named). In particular, the following class of sets will play an important role in reflecting structure between M and N :

Definition 3.2. (Induced) *Let $N = M^\lambda/\mathcal{D}$. A set $Q \subset N^k$ is said to be induced if it is equal to the ultraproduct (modulo \mathcal{D}) of its projections to the index models.*

Example 3.3. *The following are equivalent for $M \models T_{rg}$ and a regular ultrafilter \mathcal{D} on λ :*

- (1) $N := M^\lambda/\mathcal{D}$ is λ^+ -saturated.
- (2) For any two disjoint $A, B \subset N$ with $|A| + |B| \leq \lambda$, there exists an induced unary predicate Q such that $A \subset Q$ and $B \cap Q = \emptyset$.

Proof. First, by quantifier elimination, N is λ^+ -saturated iff it realizes all types of the form: $p(x, A, B) := \{xRa : a \in A\} \cup \{\neg xRb : b \in B\}$ for all disjoint sets $A, B \subset N$ with $|A| + |B| \leq \lambda$.

(1) \rightarrow (2): Let A, B be given and let c realize the type $p(x, A, B)$ in N . Let $Q[t] := \{y \in M : M \models c[t]Ry\}$. Then $\prod_{t < \lambda} Q[t]/\mathcal{D}$ is the desired set.

(2) \rightarrow (1): Let Q be given and, as above, write $Q[t]$ for its trace in the index model M . Let $d : A \cup B \rightarrow \mathcal{D}$ be a map whose image is a regularizing family. Refining this map by Loś' theorem, we may assume that almost everywhere, $A[t] \subset Q[t]$ and $B[t] \cap Q[t] = \emptyset$. Define $c[t]$ to be an element satisfying $p(x, A[t], B[t])$ in M , if it exists; it will almost everywhere, by the axioms of the random graph and the fact that the traces of A, B are almost everywhere finite and disjoint. Then $c := \prod_{t < \lambda} c[t]/\mathcal{D}$ will realize $p(x, A, B)$ in N . \square

In other words, in ultrapowers of the random graph, we can find an induced unary predicate which separates two given sets of size λ (i.e. we can realize a certain second-order type) if and only if a certain type in the formula $xRy \wedge \neg xRz$ is realized over those sets.

3.1. Analysis of the example. We make some observations which will then generalize. First, we observe that characteristic sequences (Definition 2.8 above), introduced in [6] to analyze first-order types, have a larger scope.

Remark 3.4. *By analogy to Definition 2.8, we can describe the existence of a unary predicate Q in Example 3.3 by writing down the (a priori non-first-order) sequence $\langle P_n^Q : n < \omega \rangle$ given by*

$$P_n^Q((y_1, z_1), \dots, (y_n, z_n)) = \exists Q \left(\bigwedge_{i \leq n} \psi(Q, y_i) \wedge \neg \psi(Q, z_i) \right)$$

where $\psi(Q, y) = Q(y)$ has one second-order and one first-order variable.

That is, in any model M (not necessarily an ultrapower) the interpretation of P_k^Q in M holds on a k -tuple of pairs of elements of M whenever it is possible to expand the model by a unary predicate whose interpretation is as required.

Notice, however, that in Example 3.3 the P_n^Q associated to the induced predicate are first-order definable because their truth depends only on collisions between the parameters y_i, z_j . The predicate P_k^Q holds on a k -tuple of pairs of elements of some model M exactly when the set of first coordinates of the pairs does not intersect the set of second coordinates. More generally, a possibly infinite subset $A \subset M^2$ is a P_∞^Q -complete graph precisely when it is possible to expand the model by a unary predicate which separates first and second coordinates of A . Thus in Remark 3.4, the second-order structure simply records the existence of a possible partition.

Returning to the random graph:

Observation 3.5. Let $\varphi(x; y, z) = xRy \wedge \neg xRz$ and $T = T_{rg}$, and $\langle P_n : n < \omega \rangle$ be its characteristic sequence. Then φ is fundamental in the sense of Definition 2.19 above.

Proof. We check the conditions of Definition 2.19. (1) It suffices to check that, writing $y_i = (z_i, w_i)$, we have that $\{z_1, \dots, z_n\} \cap \{w_1, \dots, w_n\} = \emptyset$. (2) In fact, $k = 2$ suffices. \square

Conclusion 3.6. Example 3.3 can be explained via Fact 2.11 by noting that the characteristic sequences $\langle P_n \rangle$ for φ and $\langle P_n^Q \rangle$ for Q are identical and, moreover, both sequences are fundamental in the sense of Definition 2.19.

4. SECOND-ORDER QUANTIFIERS

We now generalize the analysis of Example 3.3 to all fundamental formulas.

In his paper “There are just four second-order quantifiers” [11], Shelah gives the definition (suggested by Stavi):

Definition 4.1. A second order quantifier is of the form Q_ψ , where $\psi = \psi(r)$ is a first-order sentence with the single predicate r , and $(Q_\psi(r))\phi$ means “There is a relation r satisfying ψ such that $\phi\dots$ ”

For instance, Q_ψ may be a unary predicate, a bijection, a linear order, or an equivalence relation all of whose classes have size 2. We will need the following notation:

Definition 4.2. [11] Let Q_ψ be a second-order quantifier as in Definition 4.1. Then for any set B , $R_\psi(B) = \{R : R \text{ is an } \ell(\psi)\text{-ary relation over } B \text{ and } B \models \psi[R]\}$.

Definition 4.3. (\mathcal{D} solves Q_ψ) Let \mathcal{D} be a regular ultrafilter on λ and Q_ψ a second-order quantifier.

- (1) Given a set A and $Q \in Q_\psi$ with $\ell(\psi) = n$, we write $A \models Q = R$ to abbreviate: $Q \cap A^n = R \cap A^n$.
- (2) Say that \mathcal{D} solves Q_ψ if for any infinite M , any small set $A \subseteq N = M^\lambda / \mathcal{D}$, and any $Q \in R_\psi(A)$, there exists an induced predicate R such that $R \in R_\psi(N)$ and $A \models Q = R$.

As in Remark 3.4, we may naturally assign characteristic sequences to assertions of the form $Q_\psi\phi$:

Definition 4.4. To any second-order quantifier Q_ψ and associated formula $\phi(Q, \bar{y})$ from Definition 4.1, we may associate the characteristic sequence $\langle P_n : n < \omega \rangle$ given by:

$$P_n(\bar{y}_1, \dots, \bar{y}_n) \iff Q_\psi(R) \left(\bigwedge_{i \leq n} \phi(R, \bar{y}_i) \right)$$

For concreteness, we will be most interested in formulas ϕ which allow us to determine the “type” of the predicate Q_ψ over some set of size $\leq \lambda$. The use of the word “type” looks towards Claim 4.11 below.

Definition 4.5. Let Q_ψ be a second-order quantifier of arity n . Let $\phi(R, y_1, \dots, y_n, z, w) = ((R(y_1, \dots, y_n) \wedge z = w) \vee (R(y_1, \dots, y_n) \wedge z \neq w))$. Then ϕ is a formula in the language of equality with a single second-order variable R (alternately, a first-order formula in the language with a symbol for R). We say that ϕ is a true description of Q_ψ .

Informally, a true description allows us to determine how R partitions some given set of parameters. In general, of course, one can do this with many different formulas in the language of equality coding the assertions “ $R(\bar{x})$ ” and “ $\neg R(\bar{x})$ ”.

Convention 4.6. *Below, when we write “ $\langle P_n \rangle$ is ‘a’ characteristic sequence of Q_ψ ” this will always mean that it is the characteristic sequence of $Q_\psi\phi(r, \bar{y})$ computed with respect to the true description ϕ just given in the language of equality (Definition 4.5). However, it will not matter much in our context if another true description is chosen; “a” rather than “the” reflects this fact.*

We now briefly discuss the hypothesis of “finite support” in Definition 2.19. Recall that a characteristic sequence has *finite support* if for some $n < \omega$ and for all $m > n$, $P_m(y_1, \dots, y_m)$ holds iff P_n holds on all n -element subsets of $\{y_1, \dots, y_m\}$. In the first-order case, the characteristic sequence of φ has finite support iff φ does not have the finite cover property, see [6].

Remark 4.7. *If ψ asserts that R is an equivalence relation, and $\phi(R, x, y, z, w) = (R(x, y) \leftrightarrow z = w) \wedge (\neg R(x, y) \leftrightarrow z \neq w)$ is a true description, then the associated characteristic sequence has finite support, in fact support 2.*

Likewise, in the Example 3.3 of the random graph above, saturation depends on a formula $\varphi(x; y, z) = xRy \wedge \neg xRz$ of finite support. Still, any unstable theory contains a formula with the finite cover property (recall definition on page 3):

Fact 4.8. ([12] Theorem 4.2 p. 62) *If $\varphi(\bar{x}; \bar{y})$ has the order property, then the formula*

$$\theta(\bar{x}; \bar{y}_1, \bar{y}_2, \bar{y}_3, \bar{y}_4) = (\varphi(\bar{x}; \bar{y}_1) \leftrightarrow \neg\varphi(\bar{x}; \bar{y}_2)) \wedge (\varphi(\bar{x}; \bar{y}_3) \leftrightarrow \varphi(\bar{x}; \bar{y}_4))$$

has the finite cover property.

In particular, this can be reflected in second-order quantifiers.

Example 4.9. *Let Q_M be the monadic quantifier (which corresponds to the unstable formula xRy in the random graph) and consider the related quantifier Q_ψ where*

$$\psi(r) = \forall x(r(x) = r(x)) \wedge \exists x\neg r(x) \wedge \exists xr(x)$$

Then the formula $(r(x) \leftrightarrow \neg r(y)) \wedge (r(z) \leftrightarrow r(w))$ is a true description of Q_ψ whose associated characteristic sequence does not have finite support.

However, this is not an essential loss, as the fundamental formulas are covered by the second order quantifiers:

Claim 4.10. (Fundamental formulas and quantifiers)

- (1) *Each characteristic sequence $\langle P_n : n < \omega \rangle$ of finite support which is definable by formulas $\langle \theta_1, \dots, \theta_k \rangle$ in the language of equality is a characteristic sequence of some second-order quantifier Q_ψ .*
- (2) *Each second-order quantifier Q_ψ in the sense of Definition 4.1 has a characteristic sequence $\langle P_n^\psi : n < \omega \rangle$ which is definable in the language of equality.*

Proof. (1) Let $\langle P_n : n < \omega \rangle$ be such a characteristic sequence; suppose it has support k , and let $\theta_1, \dots, \theta_k$ be the defining formulas for P_1, \dots, P_k respectively. Let m be the arity of a tuple from P_1 . Then let $\psi(R)$ be the formula

$$\forall \bar{y}(r(\bar{y}) \equiv r(\bar{y})) \wedge \bigwedge_{i \leq k} \forall \bar{y}_1 \dots \bar{y}_i \left(\bigwedge_{j \leq i} (r(\bar{y}_j) \implies \theta_i(\bar{y}_1, \dots, \bar{y}_i)) \right)$$

(2) Suppose the arity of Q_ψ is m . Let $\langle P_n \rangle$ be a characteristic sequence of Q_ψ . Note first that for any (sufficiently large) set B , $R \in R_\phi(B)$ and $f : B \rightarrow B$ a bijection, if we identify R with its interpretation in B then f carries R to some other $R' \in R_\phi(B)$ and thus preserves the quantifier. So for each $k < \omega$, $P_k(y_1, \dots, y_k)$ depends only on the (principal) type of y_1, \dots, y_k in the language of equality, and we may define it by taking the disjunction of the finitely many such types on which it is satisfied. \square

The next result says that all such second-order quantifiers “descend” to some first-order representative of equivalent complexity, i.e. with the same characteristic sequence; we delay the (straightforward but lengthy) proof to the Appendix, page 20.

Lemma 7.1. (see p. 20 below) *Suppose Q_ψ is a second-order quantifier of arity m . Let*

$$\varphi(r, y_1, \dots, y_{m+2}) = (r(y_1, \dots, y_m) \wedge y_{m+1} = y_{m+2}) \vee (\neg r(y_1, \dots, y_m) \wedge y_{m+1} \neq y_{m+2})$$

be a true description and $\langle P_n \rangle$ its associated characteristic sequence.

Then there exist a first-order theory T in the language $\mathcal{L} = \{=, X, Y, \rho\}$ where X, Y are unary predicates and ρ is an $(m+1)$ -ary relation such that $\langle P_n \rangle$ is also the characteristic sequence of (T, ξ) , where

$$\xi(x, y_1, \dots, y_m, z, w) = \begin{cases} \rho(x, y_1 \dots y_m) & \text{if } z = w \\ \neg \rho(x, y_1 \dots y_m) & \text{if } z \neq w \end{cases}$$

We have the following equivalence between any fundamental formula and second-order quantifier which share a characteristic sequence:

Claim 4.11. *The following are equivalent, for a regular ultrafilter \mathcal{D} on λ :*

- (1) \mathcal{D} solves the second-order quantifier Q_ψ in the sense of Definition 4.3.
- (2) \mathcal{D} solves $\langle P_n \rangle$ in the sense of Definition 2.12, where $\langle P_n \rangle$ is a characteristic sequence of Q_ψ .
- (3) Any \mathcal{D} -ultrapower of some model of T_ρ realizes all ρ -types over sets of size $\leq \lambda$, where T_ρ, ρ are the theory and formula constructed in §7 such that the characteristic sequence of ρ modulo T_ρ is a characteristic sequence of Q_ψ .
- (4) Any \mathcal{D} -ultrapower of some model of T realizes all φ -types over sets of size $\leq \lambda$, where T, φ are such that the characteristic sequence of φ w.r.t. T is the same as that of Q_ψ .

Proof. By construction, characteristic sequences work in precisely the same way for the first- and second-order types: namely, Fact 2.11 applies. In other words, the following are equivalent: (1) \mathcal{D} solves $\langle P_n \rangle$, (2) for any small positive base set $A \subset P_1^N$ in any regular ultrapower $N = M^\lambda/\mathcal{D}$, there is a distribution $d : A \rightarrow \mathcal{D}$ which is a.e. a P_∞ -complete graph as required. But this corresponds to (3), (4) as well just by the fact that Q_ψ and φ share the same characteristic sequence. \square

Given this correspondence between solution of second-order quantifiers and realization of certain first-order φ -types, we now consider how this identification can help with the original goal of evaluating saturation for these types.

As the title of his paper suggests, Shelah [11] proved the following remarkable result:

Theorem E. ([11] Theorem 2 p. 285) *Each Q_ψ is equivalent to exactly one of the following quantifiers:*

- A) Q_{FO} , the trivial quantifier, i.e. Q_{ψ_1} , $\psi_1 = r$, $n(\psi_1) = 0$, so the language with this additional quantifier is just first order logic.
- B) Q_M , the monadic second-order quantifier, i.e. Q_{ψ_M} , $\psi_M = (\forall x)[r(x) \equiv r(x)]$, and $n(\psi_M) = 1$.
- C) Q_σ , the permutational second-order quantifier, ranging over permutations of the universe of order 2, i.e. Q_{ψ_σ} where $\psi_\sigma = (\forall x)[f(f(x)) = x]$.
- D) Q_{II} , the (full) second-order quantifier, i.e. $Q_{\psi_{II}}$, $\psi_{II} = (\forall xy)[r(x, y) \equiv r(x, y)]$, $n(\psi_{II}) = 2$.

Note that Q_{FO} and Q_σ are called Q_I and Q_P , respectively, in [11]. “Equivalent” in Theorem E means up to interpretability, Definition 5.1.

5. INTERPRETABILITY

We now consider whether Theorem E remains true when the relevant definition of “interpretation” is considered in the context of ultrapowers.

Definition 5.1. ([11], p. 282) *The quantifier Q_{ψ_1} is interpretable in Q_{ψ_2} if there is a first-order formula $\theta(\bar{x}, y_1, \dots, r_1, \dots)$ such that for any infinite set A , and relation R over it such that $A \models \psi_1[R]$, there are elements $a_1, \dots \in A$ and relations S_1, \dots over A , with $A \models \psi_2[S_i]$ for each i , such that $A \models (\forall \bar{x} (R(\bar{x}) \equiv \theta(\bar{x}, a_1, \dots, S_1, \dots)))$.*

Remark 5.2. *To illustrate the subtlety of translation to ultrapowers, consider a simple example. We might argue “in the real world” that Q_{II} interprets Q_σ as follows. First, if we have available any equivalence relation with infinitely many infinite classes, then we certainly can name any infinite set: simply let this be one of the classes. Thus, to interpret an equivalence relation with infinitely many classes of size 2 on some set $A \subset B$, we might first choose an equivalence relation on B with infinitely many infinite classes (whose restriction to A is as desired), and then intersect this with a monadic predicate naming A . This works well on some fixed infinite set. However, in an ultrapower, we do not have access to every second-order predicate, but only the induced ones. Induced sets are very large and they interact with each other in “large” (i.e. coarse) ways. While it is certainly possible to find an equivalence relation E and a monadic predicate P which have the desired type over some small $A \subset N$, this is no guarantee that $E \cap P$ will be an equivalence relation with classes of size 2 on the rest of the model (or, what amounts to the same thing, in most index models). So this attempted “interpretation” has no purchase since the proposed copy of Q_σ over A which we obtain may not be itself induced.*

However, by appealing to the precise reductions made in the course of Shelah’s proof, we now show that interpretability works in our context as well (with one twist). Throughout, we will make repeated use of Fact 3.1 combined with Łoś’ theorem, as well as induced sets, Definition 3.2.

Definition 5.3. *Let Q_ϕ, Q_ψ be quantifiers. Say that $Q_\phi \trianglelefteq Q_\psi$ if for all $\lambda \geq \aleph_0$ and all regular ultrafilters \mathcal{D} on λ , if \mathcal{D} solves Q_ψ then \mathcal{D} solves Q_ϕ . If $Q_\psi \trianglelefteq Q_\phi \trianglelefteq Q_\psi$, write $Q_\phi \sqsupseteq Q_\psi$.*

Remark 5.4. *The notation \trianglelefteq reflects an analogy to Keisler’s order: if φ_1 (wrt T_1), φ_2 (wrt T_2) are first-order formulas whose characteristic sequences coincide with those of Q_{ψ_1}, Q_{ψ_2} respectively, then by Claim 4.11, $Q_{\psi_1} \trianglelefteq Q_{\psi_2}$ iff any ultrafilter which realizes all small ψ_2 types in ultrapowers of models of T_2 also realizes all small ψ_1 types in ultrapowers of models of T_1 .*

Observation 5.5. $Q_M \trianglelefteq Q_\sigma$.

Proof. Let A, B be disjoint small subsets of some given ultrapower N . To solve Q_M in this instance, it is sufficient to find an induced permutation σ of order 2 which fixes A pointwise and does not fix any of the elements of B pointwise. That is, given σ as described, the set $X = \{x : \sigma \text{ is a permutation of order 2 and } \sigma(x) = x\}$ is also induced and separates A, B almost everywhere by Łoś’ theorem. \square

Claim 5.6. *The following are equivalent for a regular ultrafilter \mathcal{D} on λ .*

- (1) \mathcal{D} solves Q_σ , the quantifier asserting the existence of a permutation of order 2.
- (2) \mathcal{D} solves Q_β , the quantifier asserting the existence of a bijection f (without loss of generality we will choose the sequences A, B enumerating the domain and range to be disjoint)

Proof. 1 \rightarrow 2: Clear.

2 \rightarrow 1: Let us check that we can allow for the possibility that the permutation has fixed points. Let $a_i \mapsto b_i$ (for $i < \lambda$) describe the desired permutation, and let $X \subseteq \lambda$ be the set on which $a_i = b_i$. Now we know that there exists *some* infinite, coinfinite induced predicate P which does not contain

$\{b_i : i \notin X\}$. For each $i \in X$, choose $c_i \in N$ which does not equal any of the elements a_i (easy as any infinite induced predicate will have size at least λ^+). Now we ask for a bijection f such that $f(a_i) = c_i$ if $i \in X$ and $f(a_i) = b_i$ otherwise. Let $\varphi(f)$ say that f is a bijection in the language with a symbol for f (see Fact 3.1). Now we distribute the elements via:

- for $i \in X$, set $d(a_i) = d(c_i) = \{t < \lambda : P(c_i) \wedge \varphi(f) \wedge f(a_i) = c_i\}$
- for $i \notin X$, set $d(a_i) = d(b_i) = \{t < \lambda : \neg P(c_i) \wedge \varphi(f) \wedge f(a_i) = c_i\}$

Now for each index $t < \lambda$, we define a function g_t as follows: (1) if $f(x) = y$, $y \notin P$ then set $g(x) = y, g(y) = x$. If $f(x) = y$ and $y \in P$, then set $f(x) = x, f(y) = y$. Now $g = \prod_t g_t / \mathcal{D}$ is as desired. \square

Observation 5.7. *Likewise, Q_E (the quantifier asserting the existence of an equivalence relation) is implied by the quantifier Q_f asserting the existence of a many-to-one function.*

Proof. Let the preimages be the desired equivalence classes. \square

Claim 5.8. $Q_\sigma \leq Q_f$.

Proof. Work in some ultrapower M^λ / \mathcal{D} . Suppose we want to realize a type which asks that $E(a_i, b_i)$ (for $i < \lambda$, and $E \models Q_\sigma$, and parameters from the ultrapower). By hypothesis we may find an induced many-to-one function f which satisfies $f(a_i) = b_i$ for all $i < \lambda$, and likewise find an induced many-to-one function g which satisfies $g(b_i) = a_i$ for all $i < \lambda$. Now distribute the elements a_i, b_i by sending any element with subscript i to the set $h(i) := \{t < \lambda : M[t] \models (f(a_i[t]) = b_i[t]) \wedge (g(b_i[t]) = a_i[t]) \wedge \text{“}f, g \text{ are functions”}\} \in \mathcal{D}$. Since f, g are (almost everywhere) functions, we have that if $t \in h(i), h(j)$ for $i \neq j$, then $|\{a_i[t], b_i[t]\} \cap \{a_j[t], b_j[t]\}| \neq 1$. So we can construct the desired permutation in each index model. \square

Conclusion 5.9. $Q_{FO} \leq Q_M \leq Q_\sigma \leq Q_f$.

Notice, however, that for now this is a statement about these four specifically, not about the class they belong to up to interpretability.

Recall the definition of TP_2 from Theorem D page 7 above. It was shown in [8] §6 that any ultrafilter which saturates models of some theory with TP_2 must have the following property, which we will make extensive use of. (More precisely, as mentioned above, it was shown that there is a Keisler-minimal TP_2 -theory and that this next property suffices for its saturation). Recall that “small” for M^λ / \mathcal{D} means $\leq \lambda$.

Definition 5.10. (\mathcal{D} solves (ω, ω) , Malliaris [8])

- (1) Let $\langle P_n \rangle$ be a characteristic sequence. An (ω, ω) -array is an infinite set $C = \{c_i^t : t < \omega, i < \omega\}$ such that first, $P_2(c_i^t, c_j^s)$ iff $(i \neq j) \vee (t = s)$, i.e. elements in the same column are pairwise inconsistent, and second, the sequence restricted to C has support 2, i.e. any subset of C which contains no more than one element from each column is P_∞ -complete.
- (2) Let \mathcal{D} be a regular ultrafilter on λ , $N = M^\lambda / \mathcal{D}$, $\langle P_n \rangle$ a characteristic sequence and $C \subseteq M$ an (ω, ω) -array for $\langle P_n \rangle$. Without loss of generality, by Fact 3.1, suppose the language has a predicate for C , so C^N is its interpretation in N . We say that \mathcal{D} solves (ω, ω) if \mathcal{D} solves any small positive base set $A \subset N$ such that $A \subseteq C^N$.

Equivalently, by Łoś’ theorem, \mathcal{D} solves (ω, ω) if it solves any small positive base set A such that for almost all indices t , $A[t] \subseteq C^M$. In other words, a particular kind of type (which can be seen from the configuration as arising from many pairwise compatible instances of dividing) is known to be solvable.

Now the first-order theory intervenes. Q_E , the quantifier asserting the existence of an equivalence relation, and (ω, ω) represent a priori different conditions on ultrafilters: informally speaking, Q_E asks us to separate infinitely many infinite sets, while (ω, ω) and Q_f ask us to choose a single element out of each of these sets. However, from the first-order point of view, Q_E implies (ω, ω) .

Claim 5.11. *Let $\langle P_n \rangle$ be a characteristic sequence for Q_E . Let φ be a first-order formula with the same characteristic sequence modulo some background theory T . Then for any regular ultrafilter \mathcal{D} on λ and $M \models T$, if M^λ/\mathcal{D} is λ^+ -saturated, then \mathcal{D} solves (ω, ω) .*

Proof. Let φ be the formula given (see Lemma 7.1 for an example). So φ parametrizes a family of equivalence relations which, since the characteristic sequence is $=$ -definable, are independent. But if $E(x; y, z)$ is a x -indexed family of independent equivalence relations, then $\psi(y; x, z) = E(x, y, z)$ has TP_2 . By [6] Claim 3.8 (or see [8] §6), any regular ultrafilter \mathcal{D} which solves some formula with TP_2 solves (ω, ω) . \square

Below, this will be enough to show any such filter will solve Q_f (c.f. Observation 5.7).

Corollary 5.12. *Let Q_{2e} be the quantifier asserting the existence of an equivalence relation with classes of size 2. Let $\langle P_n \rangle$ be a characteristic sequence for Q_{2e} . Let φ be a first-order formula with the same characteristic sequence modulo some background theory T . Then for any regular ultrafilter \mathcal{D} on λ and $M \models T$, if M^λ/\mathcal{D} is λ^+ -saturated, then \mathcal{D} solves (ω, ω) .*

Proof. By the analogous argument; for dividing it is enough to have classes of size 2. \square

Lemma 5.13. *Suppose Q_ψ is interpretable by Q_{FO} . Then any regular ultrafilter which solves Q_{FO} solves Q_ψ .*

Proof. By Shelah's classification, Q_ψ does not interpret Q_M . In the notation of that paper, say that two sequences \bar{a}_1, \bar{a}_2 are similar over a set C if they satisfy the same quantifier-free type over C in the language of equality. Then by Claim 4B p. 286, for every formula $\phi(\bar{x}, \bar{y}, \bar{r})$, so in particular the formula $\phi(\bar{x}, r) = r(\bar{x})$, we have that for any set A and $R \in R_\psi(A)$ there exists a finite set $C_R \subseteq A$ (uniformly definable in terms of R) such that if \bar{a}_1, \bar{a}_2 are similar over C_R then $A \models R(\bar{a}_1) \equiv R(\bar{a}_2)$.

Let $A \subset N = M^\lambda/\mathcal{D}$ be a small subset of an infinite regular ultrapower, and fix $R \in R_\psi(A)$. We would like to find an induced $Q \in R_\psi(N)$ such that $A \models R = Q$. Let C_R be the finite set given by the previous paragraph. Since C_R is finite, we may assume that on a \mathcal{D} -large set $X \subseteq \lambda$ its projection to the index $s \in X$ has the same quantifier-free type in the language of equality as does C_R in N . Let k be the arity of R . Now each k -tuple of elements of A will have one of the finitely many possible similarity types of k -tuples over C_R (in the language of equality), and by the quoted Claim 4B, this will entirely determine whether or not it belongs to R . Let $d : A^k \rightarrow \mathcal{D}$ be a map whose image is a regularizing family and which takes each tuple \bar{a} to a subset Y of X such that for each $s \in Y$, $\bar{a}[s]$ satisfies the same similarity type over $C_R[s]$ as it does in N . By Łoś' theorem, since the characteristic sequence is definable in the language of equality, we will have that for almost all indices s , there is $Q^s \in R_\psi(M)$ such that for all $\bar{a} \in A^k$ with $s \in d(\bar{a})$, $M \models Q^s(\bar{a}[s])$ iff $A \models R(\bar{a})$ in N . But then letting Q be the induced predicate whose projection to index s is Q^s satisfies our demands. \square

Remark 5.14. *Note that this proof shows more, namely, that Q_ψ is solved by any infinite regular ultrafilter. In particular, this implies by Theorem A above that any formula which shares a characteristic sequence with Q_ψ must be stable without the finite cover property.*

Fact 5.15. *We will use the following two steps from [11] in the analysis of quantifiers Q_ψ which do not interpret Q_σ .*

- (1) ([11] Claim 5H p. 292) *If Q_σ is not interpretable by Q_ψ , then for every $A, R \in R_\psi(A)$, $e^+(R, A) = \{ \langle a, b \rangle : a, b \in A, \text{ and the permutation } f \text{ such that } f(a) = b, f(b) = a, f(c) = c \text{ for } c \neq a, b \text{ is an automorphism of } (A, R) \}$ is an equivalence relation with finitely many equivalence classes.*

- (2) ([11], proof of Lemma 5, bottom of p. 293, in the context of any infinite set) *If Q_σ is not interpretable by Q_ψ then there is some $n_5 < \omega$ such that for any $A, R \in R_\psi(A)$, $e^+(R, A)$ has $\leq n_5$ equivalence classes. Let us show that this implies Q_ψ is interpretable by Q_M . This implies that for every $A, R \in R_\psi(A)$ there are sets $B_1, \dots, B_\ell \subseteq A$, the $e^+(R, A)$ -equivalence classes, such that for any tuple of elements of A , the truth value of $R(a_1, \dots, a_{n(\psi)})$ depends only on the truth values of $a_i = a_j$ and $a_i \in B_k$. Hence there is a quantifier free formula ϕ such that*

$$A \models (\forall \bar{x})[R(\bar{x}) \equiv \phi(\bar{x}, B_1, \dots, B_\ell)]$$

Lemma 5.16. *Suppose Q_ψ is interpretable by Q_M . Then any regular ultrafilter which solves Q_M solves Q_ψ .*

Proof. Let $A \subset N = M^\lambda/\mathcal{D}$ be a small subset of some regular ultrapower, and let $R \in R_\psi(A)$ be given. If Q_ψ is interpretable by Q_M , then by Theorem E it does not interpret Q_σ . By Fact 5.15, there is a partition of A into finitely many sets B_i such that for any tuple of elements of A , the truth value of $R(a_1, \dots, a_{n(\psi)})$ depends only on the truth values of $a_i = a_j$ and $a_i \in B_k$. By hypothesis, \mathcal{D} solves Q_M , i.e. we are guaranteed an induced monadic predicate which has the desired type over any given small set. Define, by induction on $i \leq \ell$, induced sets C_1, \dots, C_ℓ such that

- (1) C_1, \dots, C_ℓ partition N
- (2) $A \models C_i = B_i$ for each $i \leq \ell$
- (3) if $|B_i| < \aleph_0$, then $|C_i| = |B_i|$, so no new elements outside A are added to finite sets

Let $Q(\bar{x}) = \phi(\bar{x}, C_1, \dots, C_\ell)$. Then Q is also induced, since it is definable in terms of induced sets. By compactness (since the characteristic sequence is $=$ -definable), $Q \in R_\psi(N)$. By construction, $A \models R = Q$, which completes the proof. \square

Remark 5.17. *In particular, any regular ultrafilter which saturates models of the random graph solves any Q_ψ which is interpretable by Q_M , by the analysis of Example 3.3. Since the random graph is minimal among the unstable theories in Keisler's order [8], we can equivalently say that any regular ultrafilter which saturates some unstable theory solves any such Q_ψ .*

However above monadic there is a surprise. A first indication of this was the analysis of [8] §6, given in a slightly different language, proving the existence of a minimal TP_2 theory in Keisler's order.

Definition 5.18. [8] *Let \mathcal{D} be a regular ultrafilter on λ , $N = M^\lambda/\mathcal{D}$ and $A \subset N$ with $|A| \leq \lambda$. A true distribution of A is a map $d : A \rightarrow \mathcal{D}$ whose image is a regularizing family such that for almost all $t < \lambda$ and for all $a \neq b \in A$, if $t \in d(a) \cap d(b)$ then $a[t] \neq b[t]$.*

In particular, it was shown that any regular ultrafilter which could saturate some theory with TP_2 must be able to give a true distribution of any small set, and therefore certain induced predicates would always be available. Here a slight generalization of that argument gives more power.

Claim 5.19. *Let \mathcal{D} be a regular ultrafilter on λ and suppose that in any ultrapower $N = M^\lambda/\mathcal{D}$, if $A \subset N$, $|A| \leq \lambda$ then there exists a true distribution of A . Then for any $k < \omega$ and any set $C = \{a_1^i, \dots, a_k^i : i < \lambda\} \subseteq N^k$, there is a map $d : C \rightarrow \mathcal{D}$ whose image is a regularizing family and such that for any two k -tuples $a_1^i, \dots, a_k^i, a_1^j, \dots, a_k^j$ from C ,*

$$\text{if } t \in d(a_1^i, \dots, a_k^i) \cap d(a_1^j, \dots, a_k^j) \text{ then for all } \ell \leq k, M \models a_\ell^i[t] = a_\ell^j[t] \text{ iff } N \models a_\ell^i = a_\ell^j$$

Proof. Write C_j for the set $\{a_j^i : i < \lambda\}$ of j th coordinates of elements of C_j . For each pair $\ell, j \leq k$ (not necessarily distinct) let $d_{\ell,j} : C_\ell \cup C_j \rightarrow \mathcal{D}$ be a true distribution. Now define the desired d by $a_1^i \dots a_k^i \mapsto \bigcap_{\ell, j \leq k} d_{\ell,j}(a_j^i)$. As the quantifier-free type in the language of equality depends on checking all pairs of elements, this is sufficient. \square

Let us give this sharper notion a name (which refers to goodness of an ultrafilter).

Definition 5.20. *Let \mathcal{D} be a regular ultrafilter on λ . Say that \mathcal{D} is good for equality if the following holds: for any infinite M , if $N = M^\lambda/\mathcal{D}$, $k < \omega$ and $A \subset N^k$ with $|A| \leq \lambda$, then there exists $d : A \rightarrow \mathcal{D}$ whose image is a regularizing family such that for almost all $t < \lambda$ and for all $a \neq b \in A$, writing $q_{a,b}(x,y)$ for the type of $\{a,b\}$ in the language of equality, if $t \in d(a) \cap d(b)$ then $M \models q_{a,b}(a[t], b[t])$.*

This has the following consequences for quantifiers. We discuss its significance in Remark 5.22.

Theorem 5.21. *The following are equivalent for a regular ultrafilter \mathcal{D} on $\lambda \geq \aleph_0$.*

- 0. \mathcal{D} solves (ω, ω) .
- (1) \mathcal{D} is good for equality.
- (2) \mathcal{D} solves Q_σ .
- (3) \mathcal{D} solves every second-order quantifier.

Proof. (0) \rightarrow (1) It was proved in [8] Lemma 6.8 that (0) implies that any small set $A \subseteq N$ has a true distribution modulo \mathcal{D} . Briefly, given $A \subset N$ of size λ , let $d : A \rightarrow \mathcal{D}$ be any map whose image is a regularizing family. Then the trouble is that we may have many pairs $a \neq b \in A$ such that $a[s] = b[s]$ on a given index $s \in d(a) \cap d(b)$; we would like to refine d so this does not happen. This can be coded into a combinatorial problem visible to (ω, ω) as follows. Assign to each element $a \in A$ an element b in the model of (ω, ω) such that if a, a' are distinct in A but $a[t] = a'[t]$, then $b[t] \neq b'[t]$ however $b[t], b'[t]$ are in the same column of the (ω, ω) array. By Łoś' theorem the elements b push forward to a set B of elements in distinct (ω, ω) columns. Thus B is a positive base set in the (ω, ω) context, and any solution of B corresponds to a true distribution for A . Applying Claim 5.19 gives (1).

(1) \rightarrow (3) follows from Claim 5.19: it amounts to saying we can distribute any positive base set in a characteristic sequence of Q_ψ so that it is almost everywhere a positive base set (since each of the predicates P_n is definable in the language of equality).

(3) \rightarrow (2) is immediate.

Finally, let us show (2) \rightarrow (0). Suppose that we are given an (ω, ω) -array in M whose image in the ultrapower N is an array $W \subset N$, and assume, without loss of generality by Fact 3.1, that we have a unary relation symbol for W , a binary relation symbol E interpreted in N as the induced equivalence relation on W in which P_2 -incomparable elements are E -equivalent, and a unary relation symbol C which chooses one canonical representative from each W -equivalence class. Now the combinatorial problem is the following: we are given a small set $A \subset W$ with the property that for $a, a' \in A$, $\neg E(a, a')$. We would like to find a distribution $d : A \rightarrow \mathcal{D}$ such that for each $s \in \lambda$ and each $a, a' \in A$, if $s \in d(a) \cap d(a')$ then $a[s] \neq a'[s] \implies \neg E^M(a[s], a'[s])$. This will solve (ω, ω) .

Define for each $a \in A$ and element $c_a \in C^N$ as follows: $c_a = \prod_s c_a[s]/\mathcal{D}$ where $c_a[s]$ is the unique element of C^M such that $E^M(c_a[s], a[s])$. By hypothesis (2), there exists an induced bijection f which takes c_a to a for each $a \in A$. Now consider the distribution $d : A \rightarrow \mathcal{D}$ given by:

$$a \mapsto \{s < \lambda : f \text{ is a function and } f(c_a) = a \text{ and } E(c_a, a)\} \cap d_0(a)$$

where $d_0 : A \rightarrow \mathcal{D}$ is any map whose image is a regularizing family. By Łoś' theorem, this distribution is as required. \square

Remark 5.22. *How to interpret Theorem 5.21, i.e. why is the existence of a bijection enough to give goodness for equality, and thus guarantee the solution of any $=$ -definable second-order quantifier? One answer is to appeal to the characterization of Morley-Vaught: the saturated models are exactly the homogeneous-universal ones [9]. It is not difficult to see that if \mathcal{D} is a regular ultrafilter on λ , then M^λ/\mathcal{D} is universal for elementarily equivalent models of cardinality $\leq \lambda$. Thus failures*

of saturation come from failures of homogeneity. Theorem 5.21 may then be understood as saying that to require the existence of an induced bijection between any two sequences of size λ ensures homogeneity on the level of equality.

In the following definition, we remember the fact that realizing φ -types always occurs in the context of a theory, and that we care about the realization of all types in such a theory. Thus certain configurations which are not a priori of the same strength (in ultrapowers) when considered as as quantifiers may nonetheless be related when they are represented by some first-order formula.

Definition 5.23. (Equivalent at the level of theories)

- (1) Let Q_ψ be a second-order quantifier. Let T be a first-order theory. Say that T represents Q_ψ if there exists a formula φ in the language of T such that $(T, \varphi) \mapsto \langle P_n \rangle$ for some characteristic sequence $\langle P_n \rangle$ of Q_ψ .
- (2) Let Q_ψ, Q_ρ be second-order quantifiers. Say that Q_ψ implies Q_ρ at the level of theories if for any infinite cardinal λ , any regular ultrafilter \mathcal{D} on λ , any T which represents Q_ψ and any $M \models T$, if M^λ/\mathcal{D} is λ^+ -saturated then \mathcal{D} solves Q_ρ .
- (3) Say that Q_ψ, Q_ρ are equivalent at the level of theories if each implies the other in the sense of (2).

Claim 5.24. (1) Q_{2e} and Q_σ are equivalent at the level of theories.

(2) Q_E and Q_f are equivalent at the level of theories.

Proof. By Claim 5.11 and Corollary 5.12, any ultrafilter which saturates models of some theory representing either Q_{2e} or Q_E will solve (ω, ω) . By Theorem 5.21, it will therefore solve both Q_σ and Q_f .

The other direction is given by Theorem 5.21 (2) \rightarrow (3) in the first case, and Observation 5.7 in the second. \square

Remark 5.25. To complete this classification, it would be enough to show, analogously to Claim 5.11 and Corollary 5.12, that any theory which represents a quantifier which interprets Q_M has the independence property, and that any theory which represents a quantifier which interprets Q_σ has TP_2 . By the analysis of [6], this appears feasible as dividing and infinite D -rank are visible in the characteristic sequence.

Conclusion 5.26. From our analysis above the fundamental formulas can be characterized as follows. (Each comes with some background theory T .) “Interprets” means in the sense of [11].

- (1) $\varphi \mapsto \langle P_n \rangle \mapsto Q_\psi$ where Q_ψ does not interpret Q_M . Then $\langle P_n \rangle$ is solved by any regular ultrafilter.
- (2) $\varphi \mapsto \langle P_n \rangle \mapsto Q_\psi$ where Q_ψ interprets Q_M but does not interpret Q_σ . Then any ultrafilter \mathcal{D} which solves Q_M , or equivalently, saturates models of the random graph will solve $\langle P_n \rangle$.
- (3) $\varphi \mapsto \langle P_n \rangle \mapsto Q_\psi$ where Q_ψ interprets Q_σ . Then any ultrafilter \mathcal{D} which solves Q_σ , or, equivalently, saturates models of T_{feq} will solve $\langle P_n \rangle$.

Proof. By Theorem E, these are the only possible cases. The conclusion of Case (1) is Lemma 5.13 and the conclusion of Case (2) is Lemma 5.16. Case (3) follows from Theorem 5.21. \square

6. = AND <, OR INDEPENDENCE AND ORDER

We began, in §2.2 above, with a discussion of various tradeoffs in unstable theories between independence and order. It appears this “dichotomy” (again, quotes as the possibilities are not mutually exclusive for a theory or formula but are in some sense structurally opposite) is of basic interest in understanding Keisler’s order. It was argued in Motivation 2.20 and Remark 2.21 that the fundamental formulas express various levels of independence. By Theorem 5.21, their scope

does not extend past TP_2 , which captures “goodness for equality”. On the other hand, by Theorem C, we know that $(\omega, <)$ is maximal in Keisler’s order. Thus the distance to maximality is already covered by ultrafilters whose distributions almost everywhere respect another binary relation, linear order $<$.

It would be very interesting, both for Keisler’s order and for model-theoretic analysis more generally, to determine whether, and if so where, there exists a dividing line among the non simple theories. The theories at stake, those with the tree property but without SOP_3 , are increasingly an object of model-theoretic study. Still, as far as we are aware, despite recent results on SOP_3 (see [15], [16], [1]; also, [8] Theorem 7.11, connects SOP_3 to high density in the characteristic sequence in the sense of Szemerédi regularity) it is generally considered that finding a dividing line is not tractable until the structure of such theories is model-theoretically more understood.

Our methods in this paper, however, bring to light a potentially useful gap between two non-simple theories, to do with the distinction between independence and order. Recall that “small” for M^λ/\mathcal{D} means of size no more than λ .

Theorem 6.1. *Let $M = (\omega, <)$. Let \mathcal{D} be a regular ultrafilter on λ which solves T_{rg} (equivalently, Q_M) and $N = M^\lambda/\mathcal{D}$. Let $\varphi(x; y, z) = y < x < z$. Then for any $A \subset N^2$ and any small φ -type p over A , there exist $B \subseteq N^2$ with $|A| = |B|$ and a small φ -type q over B , such that:*

- (1) N realizes p if and only if N realizes q .
- (2) There is a true distribution of B modulo \mathcal{D} .

In other words, provided \mathcal{D} saturates the minimum unstable theory, the difficulty of realizing a cut in N , i.e. finding a distribution of its base set which almost everywhere respects $<$, is not made easier if there is a distribution of its base set which almost everywhere respects $=$.

Note that by quantifier elimination, types of this form determine saturation of ultrapowers of $(\omega, <)$.

The following observation will be useful for the proof. For clarity of notation, we distinguish the index set: let \mathcal{D} be an ultrafilter on I , where $|I| = \lambda$.

Observation 6.2. *Let p be a small type in $N = (\omega, <)^I/\mathcal{D}$, given by $\{a_i^0 < x < a_i^1 : i < \lambda\}$. Let r be another small type in N , given by $\{c_i^0 < x < c_i^1 : i < \lambda\}$. Write a_i for the pair (a_i^0, a_i^1) and c_i for (c_i^0, c_i^1) , so we have that $A = \{a_i : i < \lambda\}$ and $C = \{c_i : i < \lambda\}$ are positive base sets for the characteristic sequence $\langle P_n \rangle$ of φ . Note that this sequence has support 2.*

Suppose that for all $i < j < \lambda$,

$$X_{i,j} := \{s \in I : M \models P_2(a_i[s], a_j[s])\} = \{s \in I : M \models P_2(c_i[s], c_j[s])\}$$

Then N realizes p if and only if N realizes r .

Proof. The conclusion is symmetric, so let us show that if N realizes r then N realizes p . Let $d : C \rightarrow \mathcal{D}$ be a distribution which is almost everywhere P_∞ -complete, given by Fact 2.11. Define the analogous distribution $d' : A \rightarrow \mathcal{D}$ by $a_i \mapsto d(c_i)$. By choice of d , if $i < j$ then $d(c_i) \cap d(c_j) \subseteq X_{i,j}$ and so by choice of d' , $d'(a_i) \cap d'(a_j) = d(c_i) \cap d(c_j) \subseteq X_{i,j}$. By hypothesis and definition of $X_{i,j}$, this means the distribution d' is almost everywhere a P_2 -complete graph. As noted, this characteristic sequence has support 2. So d' is almost everywhere P_∞ -complete, thus p is realized, as desired. \square

We now return to the proof of the theorem.

Proof. (of Theorem 6.1) Let p be a small type in $N = (\omega, <)^I/\mathcal{D}$, given by $\{a_i^0 < x < a_i^1 : i < \lambda\}$, where $i < j < \lambda$ implies $a_i^0 < a_j^0 < a_j^1 < a_i^1$. Once again, write a_i for (a_i^0, a_i^1) and let $A = \{a_i : i < \lambda\}$, so A is a positive base set.

Let $d : A \rightarrow \mathcal{D}$ be a.e. in P_1^M . By the assumption that \mathcal{D} solves Q_M , we may further assume that under d the sets $L_p = \{a_i^0 : i < \lambda\}$ and $R_p = \{a_i^1 : i < \lambda\}$, describing the left and right sides

of the cut p , are disjoint under the distribution d . If d is a true distribution of both L_p and R_p , we are done: let $B = A$, $q = p$. If not, for any $s \in I$, denote by Y_s the set $\{i < \lambda : s \in d(a_i)\}$ of indices of elements represented at index s . For each $s \in I$, we may define equivalence relations E_ℓ^s, E_r^s on Y_s by $E_\ell^s(i, j) \iff (a_i^0[s] = a_j^0[s])$, and likewise $E_r^s(i, j) \iff (a_i^1[s] = a_j^1[s])$, which record the collisions. (By construction, the only collisions will be between elements on the same side of the cut.)

Define a linear order $<_s$ on the set of pairs $\{(t, i) : t < 2, i \in Y_s\}$ by: $(t, i) <_s (t', i')$ if:

- $t = t' = 0$ and $E_\ell^s(i, i')$ and $i < i'$
- $t = t' = 1$ and $E_r^s(i, i')$ and $i' < i$
- $t \neq t'$ and $a_i^t[s] < a_{i'}^{t'}[s]$

By the hypothesis that L_p and R_p are disjoint at s , this is well defined and is indeed linear. Essentially, what we have done is the following. At index s the projected elements cluster into blocks (elements whose projections to s are equal). Each block consists entirely of elements from either the left or the right side of the cut, though of course the blocks may alternate in the linear order $<$ on $M = (\omega, <)$. We define a second order $<_s$ which refines the given order on the blocks by ordering the elements within each block according to their natural order as seen in N . We will then want to choose witnesses B whose order-type in M with respect to $<$ is the same as the order-type of the projected parameter set $A[s]$ with respect to $<_s$. This may require “spreading out” the elements along the linear order to accommodate the expansion. Thus we ensure goodness for equality of this second set, while preserving the configuration from A with respect to $<$.

More formally, we construct B as follows. At each index s , let $f : Y_s \rightarrow M \times M$ choose a set of pairs of elements $\{b_i[s] = (b_i^0[s], b_i^1[s]) : i \in Y_s\}$ which satisfies $b_i^t[s] < b_{i'}^{t'}[s]$ in M if and only if $(t, i) <_s (t', i')$. For each $i < \lambda$, let $b_i = \prod_{s \in I} b_i[s] / \mathcal{D}$.

The set $B = \{b_i : i < \lambda\}$ has a true distribution, i.e. $d(b_i) := d(a_i)$, by construction. Now let $i < j$ and let a_i, a_j and b_i, b_j be the corresponding elements of A, B respectively. For any $s \in I$, $P_2(a_i[s], a_j[s])$ iff $a_j^0[s] < a_i^1[s]$ iff $b_j^0[s] < b_i^1[s]$ iff $P_2(b_i[s], b_j[s])$, by definition of $<_s$ and the fact that the blocks were uniformly composed of left or right elements. In particular, since the characteristic sequence of $\varphi(x; y, z) = y < x < z$ has support 2, we have that B is a positive base set, and we may call its corresponding type q . By Observation 6.2, p is realized iff q is realized, which completes the proof. \square

7. APPENDIX: CONSTRUCTION OF THEORIES

In this appendix, we show that the characteristic sequences of second-order quantifiers coincide with $=$ -definable characteristic sequences of certain first-order formulas, as mentioned on page 12 above. The construction will require some intermediate definitions.

Lemma 7.1. *Suppose Q_ψ is a second-order quantifier of arity m . Let*

$$\varphi(r, y_1, \dots, y_{m+2}) = (r(y_1, \dots, y_m) \wedge y_{m+1} = y_{m+2}) \vee (\neg r(y_1, \dots, y_m) \wedge y_{m+1} \neq y_{m+2})$$

be a true description and $\langle P_n \rangle$ its associated characteristic sequence. Let ν_n be the formula in the language of equality which defines the predicate P_n , given by Claim 4.10.2.

Then there exist a first-order theory T in the language $\mathcal{L} = \{=, X, Y, \rho\}$ where X, Y are unary predicates and ρ is an $(m + 1)$ -ary relation such that $\langle P_n \rangle$ is also the characteristic sequence of (T, ξ) , where

$$\xi(x, y_1, \dots, y_m, z, w) = \begin{cases} \rho(x, y_1 \dots y_m) & \text{if } z = w \\ \neg \rho(x, y_1 \dots y_m) & \text{if } z \neq w \end{cases}$$

Remark 7.2. *For the complexity of this theory, as compared to that of the formula ρ , see Remark 7.6.*

First part of the proof. Consider the language \mathcal{L} with equality, two disjoint unary predicates (or sorts) X and Y which will partition the domain, and an $(m + 1)$ -ary relation $\rho(x, y_1, \dots, y_m)$ with domain $X \times Y^m$. We will build M as the union of an increasing chain of \mathcal{L} -structures M_i .

At stage 0: X_0 is empty, Y_0 has m distinct elements, ρ is empty.

At odd stages $i + 1$: Say that $C \subseteq (Y_i)^m$ is a maximal P_∞ -complete graph if (1) for each $\ell < \omega$ and $\bar{c}_1, \dots, \bar{c}_\ell \in C$, we have $P_\ell(\bar{c}_1, \dots, \bar{c}_\ell)$ and (2) no $C', C \subsetneq C' \subseteq (Y_i)^m$ satisfies this condition. For each maximal P_∞ -complete graph $C \subseteq (Y_i)^m$, add, if no such element already exists, a new element a to X_{i+1} and set $\rho(a, \bar{c})$ to hold iff $\bar{c} \in C$. The structure with all such witnesses added is M_i . At odd stages, Y does not change; let $Y_{i+1} = Y_i$.

End of first part; proof continues on page 22. □

Before describing the even stages, we need one further definition.

To simplify the exposition, let us assume that z takes values in the natural numbers, though this is understood to be shorthand for an expression with many more variables definable in the language of equality as long as there are enough elements in the model (e.g. consider inputs z_1, \dots, z_ℓ and divide cases according to what combination of $z_i = z_j$ are true).

Definition 7.3. Recall that $\varphi(X, y_1, \dots, y_m)$ is a true description of the second-order quantifier Q_ψ , and that $\ell(\psi) = m$. Recall that the formula ψ defining the quantifier can be thought of as a first-order sentence with a symbol for the predicate r . Without loss of generality, the variable x does not appear in ψ . Let ψ' be the first-order sentence constructed from ψ as follows: for any m -tuple of variables y_1, \dots, y_m of \mathcal{L} , if the string $r(y_1, \dots, y_m)$ appears in ψ replace this with $\rho(x, y_1, \dots, y_m)$. Note that since the result will be a sentence, the variables do not need to coincide with those of the formula θ we now define.

Define the reverse description $\theta(y; x, y_1, \dots, y_m, z)$ as follows. We assume $x \in X$, $y_i \in Y$.

$$\theta(y; x, y_1, \dots, y_m, z) = \begin{cases} \psi'(x) \wedge \rho(x, y_1, \dots, y_m) & \text{if } z = 0, z > 2m \\ \psi'(x) \wedge \rho(x, y_1, \dots, y_{z-1}, y, y_{z+1}, \dots, y_m) & \text{if } 1 \leq z \leq m \\ y \neq y_z & \text{if } m + 1 \leq z \leq 2m \end{cases}$$

In other words, θ allows us to describe all possible types in this context which an element y of a given set A may have relative to other elements y_i and to the element x , which is just the first-order counterpart of the predicate ρ in the next construction. For instance, if x induces an equivalence relation on A and we have decided which x -classes the elements of A fall into, the consistent partial θ -types over $\{x\} \cup A$ would include the type of an element x -equivalent to some $\{a_1, \dots, a_n\}$ but not equal to any of the a_i ; and a consistent partial θ -type over $\{x_1, x_2\} \cup A$ would include the type of an element x_1 -equivalent to some a but not x_2 -equivalent to a .

Remark 7.4. Let θ be a reverse description of the quantifier Q_ψ , in the notation of Definition 7.3. Let $\langle P_n^\theta : n < \omega \rangle$ be the characteristic sequence of $\theta(y; x, y_1, \dots, y_m, z)$. As the characteristic sequence $\langle P_n \rangle$ of $Q_\psi \varphi$ is $=$ -definable, the sequence $\langle P_n^\theta : n < \omega \rangle$ is $\{=, \rho\}$ -definable.

Proof. Let $\pi = \{\theta(y; a^i, b_1^i, \dots, b_m^i, t_i) : i < k, t_i \leq 2m\}$ be a finite partial type. It suffices to determine consistency in the case where the a^i are all equal to a single a (such cases are independent), so assume this is the case for π . First consider the subtype $\pi_0 = \{\theta(y, a, b_1^i, \dots, b_m^i, t_i) : i < k, m + 1 \leq t_i \leq 2m\}$, which is definable in the language of equality. If π_0 is consistent, let c realize it and let $\pi_1 = tp_{=}(b_1^1, \dots, b_m^k, c)$ be the type of such a tuple in the language of equality.

Let $\pi_r = \{\theta(c; x, b_1^i, \dots, b_m^i, t_i) : i < k, t_i \leq m\} \subseteq \pi$ be the relevant formulas from π .

Now π is consistent just in case the type $\pi_r(x; c, \bar{b}) \wedge \pi_1(c, \bar{b})$ is consistent, i.e. we can add a new element c provided that requiring a to relate to it in the specified way will not violate the condition $\psi'(x)$. But each P_m in the characteristic sequence of Q_ψ is $=$ -definable by a formula ζ_m , so this can

be definably checked. Note that the predicate symbol ρ appears in ψ' , which (usually) prevents the characteristic sequence of θ from being $=$ -definable. \square

We can now complete the construction for Lemma 7.1.

Continuation of proof of Lemma 7.1. At even stages $i > 0$: Let $\langle P_n^\theta : n < \omega \rangle$ be the characteristic sequence of the reverse description θ . Since this sequence is definable in the language $\{=, \rho\}$, we can evaluate it on B . For every P_∞^θ -complete graph $B_0 \subseteq Y^{m+1}$, add, if one does not already exist, a new element b to Y_i and set ρ to hold following the template of this partial type. The structure with all such witnesses added is M_{i+1} . Note that at even stages, $X_{i+1} = X_i$.

Finally, let $M = \bigcup_{i < \omega} M_i$, so $X = \bigcup X_i, Y = \bigcup Y_i$. Let $T = Th(M)$.

This completes the construction of the theory T and the formula ρ . It remains to check, in the following Claim, that this construction works. \square

Claim 7.5. *The theory of the structure just built eliminates quantifiers and ξ behaves as intended.*

Proof. It suffices to show that any formula $\exists z \phi(z, \dots)$ is equivalent to a quantifier-free formula, where ϕ is a boolean combination of atomic and negative atomic formulas. In our context, this means either

$$\exists x \left(x \in X \wedge \bigwedge_i \rho(x; y_1^i, \dots, y_m^i) \wedge \bigwedge_j \neg \rho(x; y_1^j, \dots, y_m^j) \wedge \bigwedge_k x = x_k \wedge \bigwedge_\ell x \neq x_\ell \right)$$

or

$$\exists y \left(y \in Y \wedge \bigwedge_i \rho(x^i; y_1^i, \dots, y, \dots, y_m^i) \wedge \bigwedge_j \neg \rho(x_j; y_1^j, \dots, y, \dots, y_m^j) \wedge \bigwedge_k y = y_k \wedge \bigwedge_\ell y \neq y_\ell \right)$$

which correspond to statements expressible in terms of the predicates P_n and P_n^θ , respectively. Since each of these are definable by construction, there is a quantifier-free equivalent in the language $\{=\}$ and $\{=, \rho\}$, respectively.

As for the statement that “ ξ behaves as intended,” this means simply that the characteristic sequence of ξ in T is indeed a characteristic sequence of Q_ψ . This follows from the fact that, by construction, the same definition schema works for both sequences, as the predicates P_n are just names for sets definable modulo the theory T . \square

This completes the correspondence between predicates and formulas described in Definition 2.19 above. To conclude, we note that the second-order nature of these quantifiers is somehow essential to their fineness; though the combinatorial problems they present accurately reflect, and can be accurately reflected by, first-order formulas, this cannot necessarily be done on the level of theories without an increase in complexity:

Remark 7.6. *Let Q_ψ assert the existence of a total linear order. Then by Theorem 5.21 the complexity of Q_ψ is simply that of saturating some theory with TP_2 . On the other hand, the reverse description is maximally complex, because any theory with the strict order property is maximal in Keisler’s order (see Theorem C above). As earlier sections in this paper have suggested, given indications of a large gap between goodness for equality and goodness for order, it would be very surprising if the minimal TP_2 theory were itself already maximal. But then it appears unlikely that one could modify the construction of a “canonical theory” given above to construct some first-order theory associated to Q_ψ in which the quantifiers “descend” to first-order objects and yet the complexity is not increased. For this “descent” seems to necessarily involve using elements to name the instances of Q_ψ , and therefore to name at least one definable linear order.*

REFERENCES

- [1] M. Džamonja and S. Shelah, “On \triangleleft^* -maximality.” APAL 125 (2004) 119–158.
- [2] H. J. Keisler, “Ultraproducts which are not saturated.” JSL, 32 (1967) 23–46.
- [3] S. Kochen, “Ultraproducts in the Theory of Models.” Ann. of Math. (2) Vol. 74, No. 2 (Sep., 1961), pp. 221-261
- [4] K. Kunen, “Ultrafilters and independent sets.” Trans. AMS 172 (1972), 299-306.
- [5] M. Malliaris, “Realization of φ -types and Keisler’s order.” APAL 157 (2009) 220–224.
- [6] M. Malliaris, “The characteristic sequence of a first-order formula.” JSL, 75, 4 (2010) pp. 1415-1440.
- [7] M. Malliaris, “Edge distribution and density in the characteristic sequence.” APAL 162, 1 (2010) 1-19.
- [8] M. Malliaris, “Hypergraph sequences as a tool for saturation of ultrapowers.” (2010)
Submitted. Available at <http://math.uchicago.edu/~mem/>
- [9] M. Morley and R. L. Vaught, “Homogeneous universal models.” Math. Scand. 11 (1962) 37-57.
- [10] B. Nagle, V. Rödl, and M. Schacht. “The counting lemma for regular k-uniform hypergraphs.” Random Structures and Algorithms (2006) 28(2), 113-179.
- [11] S. Shelah, “There are just four second-order quantifiers.” Israel J Math 15 (1973) 282-300.
- [12] S. Shelah, *Classification Theory and the number of non-isomorphic models*, rev. ed. North-Holland, 1990.
- [13] S. Shelah, “Simple unstable theories.” Annals of Mathematical Logic 19 (1980) 177-203.
- [14] S. Shelah, “The Universality Spectrum: Consistency for more classes.” *Combinatorics, Paul Erdős is Eighty* (1993) 403–420
- [15] S. Shelah, “Toward classifying unstable theories.” Annals of Pure and Applied Logic 80 (1996) 229–255.
- [16] S. Shelah and A. Usvyatsov, “More on SOP_1 and SOP_2 .” APAL 155 (2008), no. 1, 16–31.
- [17] E. Szemerédi, “On sets of integers containing no k elements in arithmetic progression.” Acta Arithmetica 27 (1975), 199–245.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CHICAGO, 5734 S. UNIVERSITY AVENUE, CHICAGO, IL 60637