SATURATION BY ULTRAPOWERS AND KEISLER'S ORDER

DIEGO BEJARANO

Abstract. Given a first-order model $\mathfrak{M}$ and a set of formulas $p(x)$ with at most one free variable and parameters from a subset of the universe M of the model, we can ask whether there is some $m \in M$ such that m satisfies all the formulas in $p(x)$. In this case, we say that $p(x)$ is realized. These kinds of sets of formulas, called types, can give us an insight into how complex and expressive the theory of the model is. Keisler's order gives us a way to quantify how complex a complete, countable theory is by looking at saturation, i.e. the realization of types, in ultrapowers of models of the theory. In this paper, we will define Keisler's order and give an exposition of some of the main theorems and some open questions surrounding it.

Contents

1. Introduction

H. Jerome Keisler first defined Keisler's order in his paper "Ultraproducts which are not saturated." Defined as a pre-order on complete countable theories, and therefore a partial order on the equivalence classes. Keisler's order attempts to rank theories according to the difficulty of saturating their regular ultrapowers. Theories that are minimal in Keisler's order tend to be well behaved; for example, the theory of the complex numbers is minimal in Keisler's order. On the other hand theories that are maximal tend to be very complex and puzzling; for instance the theory of the natural numbers is maximal.

When we look at the equivalence classes of complete countable theories, the structure of Keisler's order is quite rich. Keisler showed that the ordering has a minimal and a maximal class and gave some examples of theories that belong to both classes. Further work by Shelah, primarily in his book *Classification Theory*, characterized the minimal class using the *finite cover property* (f.c.p.), and showed that the stable theories with the f.c.p form their own equivalence class. We know very little about what happens with unstable theories, particularly those that are not maximal. Recent work by Malliaris and Shelah showed that Keisler's order has in fact infinitely many classes, and that under some set theoretic assumptions it is not a linear order.

The goal of this paper is to define Keisler's order on complete countable theories, as well as present and discuss some key theorems concerning the overall structure of Keisler's order and also some open questions that still puzzle mathematicians today. First, we will first define the notion of types and saturation, and give a necessary and sufficient condition for realizing types in the ultrapower. Then we will give a proper definition of Keisler's order, show that it has both a minimal and a maximal class, and delve into current research that shows Keisler's order is more complex than initially thought.

1.1. **Notation:** We use the Greek letters κ, λ, θ to denote cardinals, which we assume are infinite unless we state otherwise; ω is the first infinite cardinal. Similarly, the Greek letters α, β will be used to denote ordinals. We use $\mathcal{L}$ to denote a first-order language, and define the size of the language $|\mathcal{L}|$ as the cardinality of $\mathcal{L}$ as a set. The number of formulas of the language $\mathcal{L}$ is written as $\|\mathcal{L}\|$ and it is equal to $|\mathcal{L}| + \aleph_0$. The letters $\mathfrak{M}$ and $\mathfrak{N}$ denote models, with universes M and N respectively. The size $|\mathfrak{M}|$ of a model is defined to be the size of its universe, so for example $|\mathfrak{M}| := |M|$.

We use the notation $[\lambda]^{<\omega}$, where λ is a cardinal, to denote the set of finite subsets of λ. In general, if B is a set, then $[B]^{<\omega}$ denotes the set of finite subsets of B. If $\mathfrak{M}$ is a model and $\mathcal{I}$ is a set, then $\mathfrak{M}^{\mathcal{I}}$ is the set of all function $f : \mathcal{I} \to M$. Note that the target is the universe and not the model.

If $\mathbf{x}$ and $\mathbf{y}$ are vectors, then $\varphi(\mathbf{x}; \mathbf{y})$ is a formula where all free variables that appear in $\varphi(\mathbf{x}; \mathbf{y})$ are elements of the vector $\mathbf{x}$ and all parameters that appear in $\varphi(\mathbf{x}; \mathbf{y})$ are elements of the vector $\mathbf{y}$. In general, our notation regarding models and first-order logic follows the conventions set up in [2].

2. Preliminaries

The purpose of this section is to give a quick introduction to the objects that we study (that is, models, also called structures), as well as to the basics of the ultraproduct construction and some useful properties that we will need later on in the paper. If the reader is comfortable with the basics of model theory, then they should feel free to skip this section and move on to section three.

Definition 2.1. By a *language* in first-order logic we mean a set $\mathcal{L}$ containing the usual logical connectives, quantifiers and parentheses alongside a (possibly infinite) number of constant, variable, function, and relation symbols. We assume that every function and relation symbol takes a particular number of arguments, called the arity of the symbol.

We won't go into much detail in this paper as to how one defines the notions of a term, formula, free variable, or sentence in the context of first-order logic, but we refer the reader to [2] for a detailed construction. However, the reader might intuitively think of a formula $\varphi(x_1, \ldots, x_n)$ as a string of characters, all from the language $\mathcal{L}$, that is syntactically well-defined, and where $x_1, \ldots, x_n$ are the only variables in φ that are not bounded by quantifiers. A formula φ without free variable (i.e. all variables are bounded by quantifiers) is called a *sentence*, and a set of sentences is called a *theory*. We say that a theory T is *consistent* if it doesn't prove a contradiction, and *inconsistent* otherwise. In addition, we say that a theory is *complete* if the set of consequences of T, meaning the set of all sentences that can be proven from T, is a maximal consistent set of sentences. We now define the main object of model theory.

Definition 2.2. By a *model* of a language $\mathcal{L}$, also referred to as an $\mathcal{L}$-*structure*, we mean a tuple $\mathfrak{M} = (M, I)$ where M is a set, also called the domain of $\mathfrak{M}$, and I is an interpretation function that maps each constant symbol c in $\mathcal{L}$ to an element m of M, each n-placed function symbol f to to a function $f^{\mathfrak{M}} : M^n \to M$, and each n-placed relation symbol R to a set $R^{\mathfrak{M}} \subset M^n$, which we interpret as the set of all n-tuples where R holds. In some cases, we might denote the interpretation of a function as $I(f)$ as opposed to $f^{\mathfrak{M}}$ and the interpretation of a relation symbol as $I(R)$ instead of $R^{\mathfrak{M}}$.

We can intuitively ask if an $\mathcal{L}$-sentence is "true" within an $\mathcal{L}$-structure $\mathfrak{M}$. If so we write $\mathfrak{M} \models \varphi$ and read as $\mathfrak{M}$ *models* φ. This is defined rigorously by induction on the complexity of the sentence φ, but the definition matches the usual intuition of what it means for a sentence like $\forall x \exists y (x < y)$ to be true in a model where the ordering symbol has been interpreted. Similarly, if T is a theory and $\mathfrak{M}$ is a model of the same lanugage, then $\mathfrak{M}$ is a model of T, which we write as $\mathfrak{M} \models T$, if $\mathfrak{M} \models \varphi$ for all $\varphi \in T$. Two models $\mathfrak{M}, \mathfrak{N}$ of the same language $\mathcal{L}$ are said to be *elementarily equivalent*, denoted as $\mathfrak{M} \equiv \mathfrak{N}$, if for all sentences φ in the language, $\mathfrak{M} \models \varphi$ iff $\mathfrak{N} \models \varphi$. We can give an alternative characterization of elementary equivalence of models as follows: if $\mathfrak{M}$ is a model, then the theory of $\mathfrak{M}$, denoted as $T(\mathfrak{M})$, is the set of all sentences that are true in $\mathfrak{M}$. With this in mind, $\mathfrak{M} \equiv \mathfrak{N}$ iff $T(\mathfrak{M}) = T(\mathfrak{N})$.

Definition 2.3. Given two models $\mathfrak{M}$ and $\mathfrak{N}$ of the same language $\mathcal{L}$, an $\mathcal{L}$-*homomorphism* is a map $f : \mathfrak{M} \to \mathfrak{N}$ such that

(1) $f(c^{\mathfrak{M}}) = c^{\mathfrak{N}}$ for all constants $c \in \mathcal{L}$;
(2) $f(G^{\mathfrak{M}}(m_1, \ldots, m_k)) = G^{\mathfrak{N}}(f(m_1), \ldots, f(m_k))$ for all function symbols $G \in \mathcal{L}$;
(3) $\mathfrak{N} \models R^{\mathfrak{N}}(f(m_1), \ldots, f(m_k))$ iff $\mathfrak{M} \models R^{\mathfrak{M}}(m_1, \ldots, m_k)$ for all relation symbols $R \in \mathcal{L}$.

Moreover, an $\mathcal{L}$-homomorphism that is also a bijection is called an $\mathcal{L}$-*isomorphism*. Two $\mathcal{L}$-structures are $\mathcal{L}$-*isomorphic* if the exists a $\mathcal{L}$-isomorphism between them.

We leave it to the reader to show that isomorphism implies elementary equivalence but note that the converse is not true in general. Finally, we mention a couple of important theorems in model theory. We omit the proofs, but those can be found in [2].

Theorem 2.4 (Completeness Theorem). *Let T be a theory of a language $\mathcal{L}$. Then T is consistent if and only if it has a model.*

Theorem 2.5 (Compactness Theorem). *Let T be a theory of a language $\mathcal{L}$. Then T has a model if and only if every finite subset of T has a model.*

Theorem 2.6 (Downwards Löwenheim-Skolem-Tarski Theorem). *Let T be a consistent theory of a language $\mathcal{L}$. Then T has a model of size at most $|\mathcal{L}| + \aleph_0$.*

Theorem 2.7 (Upwards Löwenheim-Skolem-Tarski Theorem). *If a theory T of a language $\mathcal{L}$ has infinite models, then it has at least one model of size λ for all $\lambda \geq |\mathcal{L}| + \aleph_0$.*

Theorem 2.8 (Łoś-Vaught Test). *If T is a consistent theory of a language $\mathcal{L}$ with no finite models and only one model of size λ up to isomorphism, for some $\lambda \geq |\mathcal{L}| + \aleph_0$. Then T is complete.*

We now give an introduction to the ultraproduct construction and its basic properties.

Definition 2.9. An *ultrafilter* $\mathcal{D}$ over a set $\mathcal{I}$, with $|\mathcal{I}| = \lambda$, is a set $\mathcal{D} \subset \mathscr{P}(\mathcal{I})$ satisfying all of the following properties:

(1) $\emptyset \notin \mathcal{D}$ and $\mathcal{I} \in \mathcal{D}$;
(2) if $X \in \mathcal{D}$ and $X \subset Y$, then $Y \in \mathcal{D}$;
(3) if $X, Y \in \mathcal{D}$, then $X \cap Y \in \mathcal{D}$;
(4) for all $X \in \mathcal{I}$, either $X \in \mathcal{D}$ or $X^c \in \mathcal{D}$, but not both.

Intuitively, an ultrafilter is a set containing precisely the "big" subsets of $\mathcal{I}$. Notice that $\mathcal{I}$ is big while the empty set is not, a set bigger than a "big" set is also "big," and the finite intersections of "big" sets are also "big." Moreover, an ultrafilter has made a decision for every subset of $\mathcal{I}$, classifying exactly one of X or $\mathcal{I} \setminus X$ as "big".

Definition 2.10. An ultrafilter $\mathcal{D}$ over $\mathcal{I}$ is said to be *principal* if it is of the form $\{X \subset \mathcal{I} \mid x \in X\}$ for some $x \in \mathcal{I}$. Conversely, an ultrafilter is said to be *non-principal* if it is not of the form specified above. Non-principal ultrafilters are also known as *free* ultrafilters.

The existence of non-principal ultrafilters over infinite sets can be proved using Zorn's lemma. In fact, for any set $\mathcal{I}$, a collection of subset $E \subset \mathcal{P}(\mathcal{I})$ can be extended to an ultrafilter $\mathcal{D}$ (meaning that $E \subset \mathcal{D}$) iff every finite intersection of elements of E is nonempty. This property is called the *finite intersection property* (*F.I.P.*). It is also important to note that there are models of Zermelo-Fraenkel set theory (ZF) where no non-principal ultrafilters exist. Therefore, from now one we will work under Zermelo-Fraenkel set theory with the axiom of choice (ZFC).

Example 2.11. If we let $\mathcal{I} = \omega$, then the set $E = \{X \subset \omega \mid \omega \setminus X \text{ is finite}\}$ has the finite intersection property. Therefore, we can extend E to an ultrafilter $\mathcal{D}$ in ω. The reader can check that this ultrafilter is non-principal, and, in fact, that every non-principal ultrafilter on ω extends E. E is called the *Frechet filter* or the *cofinite filter* since it satisfies properties (1)-(3) of the definition above but not (4).

Definition 2.12. Given a collection of sets $\{M_i \mid i \in \mathcal{I}\}$ for some index set $\mathcal{I}$, we define the Cartesian product $\prod_{i\in\mathcal{I}} M_i$ to be the set of all function $f : \mathcal{I} \to \bigcup_{i\in\mathcal{I}} M_i$ such that $f(i) \in M_i$. The reader can verify this generalizes the finite case, where we think of the points in the cartesian product as tuples of elements from each set.

Given a collection of models $\{\mathfrak{M}_i = (M_i, I_i) \mid i \in \mathcal{I}\}$ and an ultrafilter $\mathcal{D}$ on $\mathcal{I}$, we can define an equivalence relation $=_\mathcal{D}$ on the Cartesian product $\prod_{i\in\mathcal{I}} M_i$ as follows.

Definition 2.13. Two elements $f, g \in \prod_{i\in I} M_i$ are said to be *equivalent modulo* $\mathcal{D}$ if $\{i \in I \mid f(i) = g(i)\} \in \mathcal{D}$. We write $f =_\mathcal{D} g$ to indicate this relationship. Furthermore, for some function $f \in \prod_{i\in I} M_i$, we define $f/\mathcal{D}$ (some authors use the notation $[f]_\mathcal{D}$) as the equivalence class of all functions $g \in \prod_{i\in I} M_i$ such that $g =_\mathcal{D} f$. Then the *reduced product* of $\{M_i\}_{i\in I}$ modulo $\mathcal{D}$ is defined as

$$\prod_{i\in I} M_i \,/\mathcal{D} := \{f/\mathcal{D} \mid f \in \prod_{i\in I} M_i\}.$$

Definition 2.14. Let $\mathcal{D}$ be an ultrafilter over some set $\mathcal{I}$ and, for each $i \in \mathcal{I}$, let $\mathfrak{M}_i = (M_i, I_i)$ be an $\mathcal{L}$-structure of some language $\mathcal{L}$. Then the reduced product can be made into a $\mathcal{L}$-structure called the *ultraproduct*, and denoted by $\prod\limits_{i\in I} \mathfrak{M}_i/\mathcal{D} = \left(\prod\limits_{i\in I} M_i/\mathcal{D}, I\right)$, with an interpretation function I defined as follows:

(1) if c is a constant in $\mathcal{L}$, then $I(c) = \langle c^{\mathfrak{M}_i} \mid i \in \mathcal{I}\rangle/\mathcal{D}$;

(2) if f is a function symbol of arity n and $g_1, \dots, g_n \in \prod_{i \in \mathcal{I}} M_i$, then

$$I(f)(g_1/\mathcal{D}, \dots, g_n/\mathcal{D}) = \langle f^{\mathfrak{M}_i}(g_1(i), \dots, g_n(i)) \mid i \in \mathcal{I} \rangle / \mathcal{D};$$

(3) If R is a relation symbol of arity n and $g_1, \dots, g_n \in \prod_{i \in \mathcal{I}} M_i$, then

$$(g_1/\mathcal{D}, \dots, g_n/\mathcal{D}) \in I(R) \quad \textit{iff} \quad \{i \in I \mid (g_1(i), \dots, g_n(i)) \in R^{\mathfrak{M}_i}\} \in \mathcal{D}.$$

Moreover, in the case where all of the models $\mathfrak{M}_i$ are isomorphic, we call the resulting $\mathcal{L}$-structure $\prod_{i \in I} \mathfrak{M}/\mathcal{D}$ the *ultrapower* and use $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ to denote the model and $M^{\mathcal{I}}/\mathcal{D}$ to denote the universe.

From now on, we assume all formulas φ can be expressed in first-order logic.

Theorem 2.15 (Łoś's Theorem). *Let $\mathcal{L}$ be a language, $\mathcal{D}$ be an ultrafilter on some index set $\mathcal{I}$ and $\mathfrak{M}_i$ be an $\mathcal{L}$-structure for all $i \in \mathcal{I}$. Then for all formulas $\varphi(x_1, \dots, x_n)$ of $\mathcal{L}$ and each $f_1, \dots, f_n \in \prod_{i \in \mathcal{I}} M_i$ we have that*

$$\prod_{i \in I} \mathfrak{M}_i/\mathcal{D} \models \varphi(f_1/\mathcal{D}, \dots, f_n/\mathcal{D}) \quad \textit{iff} \quad \{i \in I \mid \mathfrak{M}_i \models \varphi(f_1(i), \dots, f_n(i))\} \in \mathcal{D}.$$

Corollary 2.16 (Transfer Principle). *Let $\mathcal{L}$ be a language, $\mathcal{D}$ be some ultrafilter on some index set $\mathcal{I}$, and $\mathfrak{M}$ be an $\mathcal{L}$-structure. Then for all sentences φ of $\mathcal{L}$ we have that $\mathfrak{M}^{\mathcal{I}}/\mathcal{D} \models \varphi$ iff $M \models \varphi$. In other words, $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ and $\mathfrak{M}$ have the same theory.*

3. Saturation

Consider an $\mathcal{L}$-structure $\mathfrak{M}$ for some language $\mathcal{L}$, and let B be a subset of the universe M of size no more than some cardinal $\lambda < |M|$. We define the expanded language $\mathcal{L}(B)$ by $\mathcal{L}(B) = \mathcal{L} \cup \{c_b \mid b \in B\}$, meaning that we have added a constant symbol c_b for every element $b \in B$, and we expand $\mathfrak{M}$ to be an $\mathcal{L}(B)$-structure by setting $c_b^{\mathfrak{M}} = b$. In other words, we expand the language in such a way that we can consider the elements of B as elements of our language. Within this language there are some special sets of formulas: a *consistent n-type* over B in $\mathfrak{M}$ is a set of formulas $p(x_1, \dots, x_n)$ in the language $\mathcal{L}(B)$ such that every formula has at most n free variables, and such that for every finite subset $p_0(x_1, \dots, x_n)$ it is true that

$$\mathfrak{M} \models (\exists x_1, \dots, x_n) \bigwedge_{\varphi \in p_0} \varphi[x_1, \dots, x_n].$$

By this we mean that there are elements $m_1, \dots, m_n \in M$ that simultaneously satisfy all the conditions imposed by the sentences in $p_0(x_1, \dots, x_n)$. Such an n-type is called *complete* if it is a maximal consistent set of formulas in the language $\mathcal{L}(B)$.

We can think of types as objects that point to the existence of certain elements within the model, since for every finite subset of a type we can actually find element that satisfies all the conditions of the subset simultaneously. Therefore, we say that an n-type $p(x_1, \dots, x_n)$ over B in $\mathfrak{M}$ is *realized* if

$$\mathfrak{M} \models (\exists x_1, \dots, x_n) \bigwedge_{\varphi \in p} \varphi[x_1, \dots, x_n],$$

meaning that the element that type "points to" exists within the model. Similarly, we say that an n-tuple $(m_1, \dots, m_n)$ *realizes a type* $p(x_1, \dots, x_n)$ if $\mathfrak{M} \models p(m_1, \dots, m_n)$.

Definition 3.1. A $\mathcal{L}$-structure is called *κ-saturated* if every complete, consistent n-type over B in $\mathfrak{M}$, where $|B| < \kappa$, in $\mathfrak{M}$ is realized. Further, we say that $\mathfrak{M}$ is *saturated* if it is $|M|$-saturated.

Note that a $\mathcal{L}$-structure cannot be κ-saturated for some $\kappa > |M|$ since we can construct the type $p(x) = \{x \neq c_m \mid m \in M\}$, which cannot be realized in any model. Saturation is a nontrivial property as we will see below.

Example 3.2. $(\mathbb{N}, <)$ is not saturated. For example, in the language $\mathcal{L}(\{0\})$ the type

$$p(x) = \{x > 0, \quad (\exists w_1)\ x > w_1 > 0, \quad (\exists w_1, w_2)\ x > w_2 > w_1 > 0, ...\}$$

is not realized since $\mathbb{N}$ has no maximum element, but it is easy to see that it is consistent and it is also complete.

Example 3.3. $(\mathbb{Q}, <)$ is saturated. One way to see this is to observe that $T(\mathbb{Q}, <)$ has a property called *quantifier elimination*. In this case, this implies that any formula φ in the language of $(\mathbb{Q}, <)$ is equivalent to a formula ψ that has no quantifiers. Moreover, one can show that such a formula ψ is a description of the ordering of all the elements mentioned in ψ. Thus, since $(\mathbb{Q}, <)$ is a dense linear order, no matter what the ordering looks like, we can always find some rational number that satisfies all the formulas at the same time (assuming that the type is consistent).

Example 3.4. $(\mathbb{R}, <)$ is not saturated. The set $B = \{\frac{1}{n} \mid n \in \mathbb{N}\}$ has cardinality strictly less than the cardinality of the continuum, but the type $p(x) = \{0 < x < \frac{1}{n} \mid n \in \mathbb{N}\}$ is a consistent, and actually complete, type over B that can't be realized due to the Archimedean property of $\mathbb{R}$.

Being saturated is actually a strong property; in some sense, it says that the model is complete or full with respect to types. Just like we say that a metric space is complete if every Cauchy sequence converges to a point, a model is saturated if for every type there is a tuple that realizes it. Thus, for every possible set of conditions we can express in the language of the model, with some set of parameters from the model, there is a tuple that actually satisfies all of those conditions. Saturation is also a hard property to check since we are asking that every type, no matter how complicated, is actually realized. However, there is a result that makes this process a bit more manageable.

Theorem 3.5. *A model $\mathfrak{M}$ is κ-saturated if every complete, consistent 1-type over a set B of cardinality less than κ is realized in $\mathfrak{M}$.*

Since we have not given a thorough introduction to first-order logic, we skip the proof of this theorem. However, the proof is not difficult, and a statement of this result can be found early in Section 5.1 of [2].

Suppose now that we have a model $\mathfrak{M}$ that is not saturated. Is it possible to find some elementarily equivalent model $\mathfrak{N}$ that is saturated? What if we only ask for κ-saturation for some cardinal κ? Both questions are quite subtle. Answering the first one is beyond the scope of this paper, but we will tackle the second question in the next few sections. From there we will define Keisler's order, which allows us to use saturation to gauge the complexity of models, and more generally of theories.

Remark 3.6. For the rest of this paper, we assume that all types we consider are complete, consistent 1-types.

4. Regular ultrafilters and hints of saturation

We start this section with the following question: suppose $\mathfrak{M}$ is an $\mathcal{L}$-structure and $\mathcal{D}$ an ultrafilter on an index set $\mathcal{I}$. How saturated is $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$? We know that the ultrapower is elementarily equivalent to $\mathfrak{M}$ by

Łoś's Theorem, but since types allow us to use parameters, this question is still quite nontrivial, and in fact, depends on the ultrafilter. The trivial example occurs when $\mathcal{D}$ is a principal ultrafilter; then the level of saturation of $\mathfrak{N} = \mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ is the same as that of $\mathfrak{M}$ since the two are actually $\mathcal{L}$-isomorphic. For this reason, we will consider non-principal ultrafilters going forward. This is enough to give us our first interesting result.

Theorem 4.1. *Suppose $\mathfrak{M}$ is an $\mathcal{L}$-structure for a countable language $\mathcal{L}$, and $\mathcal{D}$ is a non-principal ultrafilter on ω. Then the ultrapower $\mathfrak{N} = \mathfrak{M}^{\omega}/\mathcal{D}$ is $\aleph_1$-saturated.*

Proof. Fix some set $B \subset N$ with $|B| < \aleph_1$, and consider a 1-type $p(x)$ over B. It suffices to show that $p(x)$ is realized in $\mathfrak{N}$. First, observe that $|B| < \aleph_1$ implies that B is at most a countable set, which implies that $\mathcal{L}(B)$ is countable, so $p(x)$ must be at most countable. Therefore, we may write

$$p(x) = \{\varphi_k(x; \mathbf{b}_k) \mid k \in \omega, \mathbf{b}_k \in [B]^{<\omega}\}.$$

Next, since every finite subset of $p(x)$ is realized in $\mathfrak{N}$, it must be true that for every $n \in \omega$,

$$\mathfrak{N} \models \exists x \ \ (\varphi_0[x; \mathbf{b}_0] \wedge \ldots \wedge \varphi_n[x; \mathbf{b}_n])$$

Then, by Łoś's theorem, we obtain that

$$G_n = \{t \in \omega \mid \mathfrak{M} \models \exists x \ \ (\varphi_0[x; \mathbf{b}_0[t]] \wedge \ldots \wedge \varphi_n[x; \mathbf{b}_n[t]]) \in \mathcal{D}.$$

Note that this doesn't mean that $G_n = \omega$, since the parameters are not necessarily the same for every $t \in \omega$. Indeed, this is not enough to show that every formula is realized on some big set, but we can make a couple of changes so that this holds. Since $\mathcal{D}$ is a non-principal ultrafilter on ω, $\mathcal{D}$ contains the cofinite filter over ω. In fact, if we let $X_n = \{t \in \omega \mid t \geq n\}$, then $X_n \in \mathcal{D}$ and the collection $\{X_n \mid n \in \omega\} \subset \mathcal{D}$ forms what we call a *regularizing family* (we will give a formal definition in a more general context later in this section).

Define $H_n = G_n \cap X_n$. Then every $t \in \omega$ belongs to at most finitely many H_n (which means that the H_n also form a regularizing family), so for every $t \in \omega$ we may define a number q_t as follows. Let q_t be the largest number such that $t \in H_{q_t}$, which is guaranteed to exist if t belongs to some H_n, since t then belongs to finitely many H_n. Let q_t be -1 if t belongs to no H_n. We can now construct an element that will realize $p(x)$. Define a function $f : \omega \to M$ as follows. If q_t is -1 then let $f(t)$ be any element of M (using the axiom of choice if necessary). Otherwise, by the definition of q_t there is some $m \in M$ such that

$$\mathfrak{M} \models \varphi_0[m; \mathbf{b}_0[t]] \wedge \ldots \wedge \varphi_{q_t}[m; \mathbf{b}_{q_t}[t]]$$

and so we define $f(t)$ to be such an m.

It remains to show that $f/\mathcal{D}$ realizes $p(x)$. For this, it suffices to show that $\mathfrak{N} \models \varphi_k[f/\mathcal{D}; \mathbf{b}_k]$ for every $k \in \omega$. Fix one such k and consider the fact that

$$\{t \in \omega \mid \mathfrak{M} \models \varphi_k[f(t); \mathbf{b}_0[t]]\} \supset \{t \in \omega \mid q_t \geq k\} \supset H_k \in \mathcal{D}.$$

By Łoś's theorem $f/\mathcal{D}$ satisfies φ_k, and thus satisfies every formula in the type. Therefore, $f/\mathcal{D}$ realizes $p(x)$ and this completes the proof. $\square$

Looking at the proof above, the reader might wonder why we need the regularizing family. The first time we use the regularizing family is when we claim that the q_t are guaranteed to exist, so what happens if we do not use the regularizing family and attempt to define the q_t using only the G_n? In that case, it is possible that some $t \in \omega$ belongs to all the G_n, which is a priori not a problem. However, the problem comes when trying to define $f(t)$, since we no longer have a good candidate. By completeness, we can satisfy only finitely many number of formulas from $p(x)$ simultaneously, but since we did not assume anything about $\mathfrak{M}$ or $p(x)$, we may not be able to satisfy all the formulas from some infinite subset of $p(x)$ at the same time. Here is where the regularizing family is useful: it makes it so that each factor model is tasked with finding a point that realizes only finite number of formulas. In some sense, the regularizing family spreads the work "evenly" across the factor models so that each model only needs to do a finite amount of work.

Now, we want to generalize this nice property to ultrafilters on sets of size larger than $\aleph_0$. This naturally gives rise to the idea of a regular ultrafilter.

Definition 4.2. An ultrafilter $\mathcal{D}$ on a set $\mathcal{I}$ of size λ is said to be *κ-regular* if there exists a family of sets $\{X_\alpha \mid \alpha < \kappa\} \subset \mathcal{D}$ such that each $i \in \mathcal{I}$ belongs to at most finitely many of the X_α. Equivalently, every intersection of infinitely many of the X_α is empty. We say that $\mathcal{D}$ is *regular* if it is λ-regular. We refer to the family $\{X_\alpha \mid \alpha < \kappa\} \subset \mathcal{D}$ as a *regularizing family.*

Intuitively, the collection $\{X_\alpha \mid \alpha < \kappa\} \subset \mathcal{D}$ has the interesting property that every finite intersection is $\mathcal{D}$-large, but every infinite intersection is empty. One might think of a regularizing family as collection of large sets that are so "spread out" that no $i \in \mathcal{I}$ ends up in infinitely many of them. The next theorem shows that we can find regular ultrafilters on sets of cardinality λ for all λ.

Theorem 4.3 (Proposition 4.3.5 in [2])**.** *For every set $\mathcal{I}$ of size $\lambda \geq \aleph_0$, there exists a regular ultrafilter $\mathcal{D}$ over $\mathcal{I}$.*

Proof. Without loss of generality, we take $\mathcal{I} = [\lambda]^{<\omega}$. For each $\alpha \in \lambda$ we define

$$\hat{\alpha} = \{I \in \mathcal{I} \mid \alpha \in I\}$$

and set

$$A = \{\hat{\alpha} \mid \alpha \in \lambda\}.$$

Notice that A has cardinality λ by construction. Moreover, for each $I \in \mathcal{I}$ we have that $I \in \hat{\alpha}$ iff $\alpha \in I$, and since I is finite, it belongs to at most finitely many $\hat{\alpha}$. Then A is a regularizing family of size λ, and hence any filter containing A is regular. Next, we claim that A has the finite intersection property. To see this note that for any $k \in \omega$ and any $\hat{\alpha}_1, \ldots, \hat{\alpha}_k \in A$ we have $\{\alpha_1, \ldots, \alpha_k\} \in \hat{\alpha}_1 \cap \ldots \cap \hat{\alpha}_k$. Therefore, A can be extended to an ultrafilter $\mathcal{D}$ over $\mathcal{I}$, which is guaranteed to be regular. $\square$

Regular ultrafilters have a number of desirable properties, which we will explore in the following section. One of these properties is that the size of the ultrapower by a regular ultrafilter is equal to the size of the cartesian power. Formally speaking, if $\mathfrak{M}$ is a model and $\mathcal{D}$ is a regular ultrafilter on $\mathcal{I}$, with $|\mathcal{I}| = \lambda$, then

$$|\mathfrak{M}^{\mathcal{I}}/\mathcal{D}| = |\mathfrak{M}|^{|\mathcal{I}|} = |\mathfrak{M}|^{\lambda}.$$

The reader can find a proof of this fact as Proposition 4.3.7 in [2], so we will not give a proof here.

Now that we have generalized the idea of a regularizing family to arbitrarily large cardinals, and shown that ultrafilters with this property exist, we have to wonder if the result of 4.1 also generalizes. Unfortunately, this does not generalize to ultrafilters defined on sets of larger cardinality, but we must restrict our attention to regular ultrafilters to obtain a more general result. In the next section, we focus on finding a necessary and sufficient condition to realize types in the ultrapower, which will help use obtain this result.

5. Realizing types and good regular ultrafilters

In this section we will start by making more assumptions. First, we will assume that all of our languages are countable, as we did in the previous section. Second, we will assume $\mathfrak{M}$ is a model of a complete theory T on a countable language. One reason why we require the language to be countable is simple: if we let the language be arbitrarily large, then our types would also be arbitrarily large, which could cause us to fail to realize types for reasons that do not depend on the model or the ultrafilter. Moreover, a countable language is sufficient to express most interesting theories we consider in model theory. Finally, we will assume that all ultrafilters are defined on infinite sets.

Now, suppose $\mathfrak{M}$ is a model of a complete theory T in a countable language $\mathcal{L}$ and $\mathcal{D}$ is some regular ultrafilter on a set $\mathcal{I}$ of size $\lambda \geq \aleph_0$. As before, let $\mathfrak{N} = \mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ and let us consider some type $p(x)$ with parameters from some set $B \subset N$ of size strictly less than $|N|$. Then we ask what are sufficient and necessary conditions for $p(x)$ to be realized? Since $\mathcal{D}$ is assumed to be regular, we know that $|N| = |M|^{|\mathcal{I}|}$, which is in general much larger than $|M|$, and unfortunately we won't be able to realize types in the ultrapower with arbitrarily large sets of parameters. However, if we further restrict our set of parameters to be of size no more than λ, then we do have some tools that we can use to talk about saturation.

Definition 5.1. Suppose $\mathfrak{M}$ is a model of a complete theory T in a countable language $\mathcal{L}$, $\mathcal{D}$ is a regular ultrafilter on some set $\mathcal{I}$ of size $\lambda \geq \aleph_0$, and $\mathfrak{N} = \mathfrak{M}^{\mathcal{I}}/\mathcal{D}$. Let $p(x)$ be a type with parameters form some set $B \subset N$ of size no more than λ. Then $|p(x)| \leq \lambda + \|\mathcal{L}\| = \lambda$, so we may enumerate the formulas in $p(x)$ as

$$p(x) = \left\{\varphi_\alpha(x; \mathbf{b}_\alpha) \mid \alpha < \lambda, \mathbf{b}_\alpha \in [B]^{<\omega}\right\}.$$

We then define a *distribution* of $p(x)$ as a map $d : [\lambda]^{<\omega} \to \mathcal{D}$ such that

(1) d is anti-monotonic: $d(v) \subset d(u)$ for all $u \subset v$;
(2) d is refinement of the Łoś map for finite subsets, meaning that for all $u \in [\lambda]^{<\omega}$,

$$d(u) \subset \left\{ t \in \mathcal{I} \mid \mathfrak{M} \models \exists x \bigwedge_{\alpha \in u} \varphi_\alpha(x; \mathbf{b}_\alpha[t]) \right\};$$

(3) The image of d is a regularizing family. This means that the set $\{u \in [\lambda]^{<\omega} \mid t \in d(u)\}$ is finite for every $t \in \mathcal{I}$.

In fact, it is also true that the set $\{d(\{\alpha\}) \mid \alpha < \lambda\}$ is also a regularizing family. Moreover, a distribution function d is said to be *accurate* if for every $t \in \mathcal{I}$ and every $u \subset \{\alpha \mid t \in d(\alpha)\}$ we have that

$$t \in d(u) \iff \mathfrak{M} \models \exists x \bigwedge_{\alpha \in u} \varphi_\alpha(x, \mathbf{b}_\alpha[t]).$$

Lemma 5.2. *Suppose $p(x)$ is a type in $\mathfrak{N} = \mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ with the same assumption as in 5.1. Then $p(x)$ has at least one accurate distribution function.*

Proof. As in 5.1, we can express $p(x)$ as $\{\varphi_\alpha(x; \mathbf{b}_\alpha) \mid \alpha < \lambda, \mathbf{b}_\alpha \in [B]^{<\omega}\}$. Since $\mathcal{D}$ is a regular ultrafilter, we know that there is a regularizing family $\{X_\alpha \mid \alpha < \lambda\} \subset \mathcal{D}$. Then, we can define our distribution function $d : [\lambda]^{<\omega} \to \mathcal{D}$ by

$$d(u) = \left\{ t \in \mathcal{I} \mid \mathfrak{M} \models \exists x \bigwedge_{\alpha \in u} \varphi_\alpha(x; \mathbf{b}_\alpha[t]) \right\} \cap \bigcap_{\alpha \in u} X_\alpha.$$

We now verify that d is indeed an accurate distribution for $p(x)$. It should be clear from the definition that d satisfies (2). To see that d satisfies (1), if $u \subset v$, then $d(v)$ intersects more sets from the regularizing family than $d(u)$, and by the nature of first order-logic

$$\left\{t \in \mathcal{I} \mid \mathfrak{M} \models \exists x \bigwedge_{\alpha \in v} \varphi_\alpha(x; \mathbf{b}_\alpha[t])\right\} \subset \left\{t \in \mathcal{I} \mid \mathfrak{M} \models \exists x \bigwedge_{\alpha \in u} \varphi_\alpha(x; \mathbf{b}_\alpha[t])\right\}.$$

Next, (3) follows from the fact that each $d(u)$ is a subset of at least one set in the regularizing family we started with. Finally, to see that d is accurate, notice that $t \in d(u)$ always implies that $\mathfrak{M} \models \exists x \bigwedge_{\alpha \in u} \varphi_\alpha(x, \mathbf{b}_\alpha[t])$ since d is a refinement of the Łoś map, so we only need to show the other direction direction. Suppose u is a subset of $\{\alpha \mid t \in d(\alpha)\}$. Since $\{\alpha \mid t \in d(\alpha)\}$ is necessarily finite, so u. Thus, we can write $u = \{\alpha_1, \ldots, \alpha_k\}$. Now suppose that $\mathfrak{M} \models \exists x \bigwedge_{\alpha \in u} \varphi_\alpha(x, \mathbf{b}_\alpha[t])$. Then $t \in d(\alpha_j)$ for $j = 1, \ldots, k$ by anti-monotonicity, which implies that $t \in X_{\alpha_j}$ for $j = 1, \ldots, k$. Then, it must be the case that $t \in \bigcap_{\alpha \in u} X_\alpha$. Finally, using our assumption we get that $t \in d(u)$ as desired. □

Distribution functions will help us figure out why we might not be able to realize a particular type, and also what conditions are necessary to realize it. Suppose we have a distribution function d of a type $p(x)$. As before, our goal is to construct a function $f : \mathcal{I} \to M$ such that for each formula $\varphi_\alpha(\alpha(x; \mathbf{b}_\alpha) \in p(x)$ the set $\{t \in \mathcal{I} \mid \mathfrak{M} \models \varphi_\alpha[f(t); \mathbf{b}_\alpha[t]]\} \in \mathcal{D}$. The fact that the image of d is a regularizing family guarantees that, for each $t \in \mathcal{I}$, the are only finitely many $\alpha < \lambda$ such that $t \in d(\{\alpha\})$. Therefore, each copy of $\mathfrak{M}$ has to fulfill a type involving only finitely many formulas, which one might think should always be possible. However, each formula comes with a set of parameters that can sometimes conflict and make it impossible to find a value for $f(t)$.

For example, fix some t and let $\alpha_1, \ldots, \alpha_n$ be the $\alpha < \lambda$ such that $t \in d(\{\alpha\})$. Now, for each α_i we can consider the set $S_i = \{m \in M \mid \mathfrak{M} \models \varphi_{\alpha_i}[m; \mathbf{b}_{\alpha_i}[t]]\}$; this is the "solution set" for φ_{α_i}. In general, the issue is that S_i and S_j being nonempty subsets of M does not imply that $S_i \cap S_j$ is nonempty. In other words, it is not true in general that $t \in d(u)$ and $t \in d(v)$ imply that $t \in d(u \cup v)$. A possible solution, then, is to find a refinement of d such that the above property holds and such that we still manage to realize each formula on a large set. In fact, this is exactly the solution. To make this precise, we first give some of relevant definitions:

Definition 5.3. Given a distribution function d of some type $p(x)$, we say that a function $\rho : [\lambda]^{<\omega} \to \mathcal{D}$ is a *multiplicative refinement* of d if $\rho(u) \subset d(u)$ for all $u \in [\lambda]^{<\omega}$ and $\rho(u \cup v) = \rho(u) \cap \rho(v)$ for all $u, v \in [\lambda]^{<\omega}$.

We can now state the main theorem of this section.

Theorem 5.4. *Let $\mathfrak{M}$ be model of a complete theory T in a countable language $\mathcal{L}$, $\mathcal{D}$ a regular ultrafilter on a set $\mathcal{I}$ of size λ, and $p(x)$ a type on $\mathfrak{N} = \mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ with parameters from a set B of size no more that λ. Then the following are equivalent:*

(1) some distribution d of $p(x)$ has a multiplicative refinement.
(2) every accurate distribution d of $p(x)$ has a multiplicative refinement.
(3) $p(x)$ is realized in $\mathfrak{N}$.

Proof. We will show that (1) $\implies$ (3) $\implies$ (2) $\implies$ (1).

(1)⇒(3): Suppose d is a distribution of $p(x)$ with some multiplicative refinement ρ. We will define a function $f : \mathcal{I} \to M$ as follows: for every $t \in \mathcal{I}$, let $\alpha_1, \ldots, \alpha_k$ be the finitely many ordinals less than λ such that $t \in \rho(\{\alpha_i\})$. Since ρ is multiplicative, we have that

$$t \in \rho(\{\alpha_1\}) \cap \ldots \cap \rho(\{\alpha_k\}) = \rho(\{\alpha_1, \ldots, \alpha_k\}),$$

so there must be some $m \in M$ satisfying $\varphi_{\alpha_1}, \ldots, \varphi_{\alpha_k}$, since ρ is a refinement of the Łośmap. We may define $f(t)$ to be one such m. To see that $f/\mathcal{D}$ is the desired element, observe that

$$\{t \mid \mathfrak{M} \models \varphi_\alpha[f(t), \mathbf{b}_\alpha[t]]\} \supset \{t \mid (\exists \sigma \in [\lambda]^{<\omega})\ [(\alpha \in \sigma) \wedge (t \in \rho(\sigma))]\} \supset \rho(\{\alpha\}) \in \mathcal{D},$$

where we obtain the rightmost containment by noticing that $\{\alpha\}$ is one such σ. Therefore, by Łoś's theorem, $f/\mathcal{D}$ realizes $p(x)$.

(3)⇒(2): Suppose now that $p(x)$ is realized by some equivalence class of the form $f/\mathcal{D}$, and let d be some accurate distribution of $p(x)$. We construct a function ρ as follows: define

$$\rho(u) = \left\{ t \in \mathcal{I} \mid \mathfrak{M} \models \bigwedge_{\alpha \in u} \varphi_\alpha[f(t), \mathbf{b}_\alpha[t]] \right\} \cap d(u).$$

We claim that ρ is a multiplicative refinement. Observe that $\rho(u) \subset d(u)$ for all $u \in [\lambda]^{<\omega}$ by definition. Thus, we only need to show that ρ is multiplicative. Next, observe that ρ is also anti-monotonic, so $\rho(u \cup v) \subset \rho(u) \cap \rho(v)$, so all that is left is to show is that $\rho(u) \cap \rho(v) \subset \rho(u \cup v)$. If $t \in \rho(u) \cap \rho(v)$, then $f(t)$ satisfies the formulas indexed by elements of $u \cup v$ simultaneously, Moreover, since d is accurate, and $u \cap v \subset \{\alpha \mid t \in d(\alpha)\}$, we must have that $t \in d(u \cup v)$, which clearly implies that $t \in \rho(u \cup v)$. This completes this part of the proof.

(2)⇒(1): We showed in 5.2 that every type has at least one accurate distribution, so (2) necessarily implies (1) as desired. □

This result provides us with an useful condition that we can place on $\mathcal{D}$ to guarantee some level of saturation.

Definition 5.5. An ultrafilter $\mathcal{D}$ on a set $\mathcal{I}$ of cardinality λ is said to be *κ-good* if for every $\mu < \kappa$ and every monotonic function $f : [\mu]^{<\omega} \to \mathcal{D}$ there is a multiplicative refinement g of f. We say that $\mathcal{D}$ is *good* if it is λ^+-good.

The existence of good ultrafilters on sets of cardinality λ is a nontrivial statement, and their construction requires some level of cleverness. The original proof of their existence is due to Keisler, and assumes the generalized continuum hypothesis (GCH). Later, Kunen gave a proof of their existence without GCH. We are, however, interested in the existence of *good regular ultrafilters* on sets of cardinality λ, since our work above shows that ultraproducts constructed with such ultrafilters will be λ^+- saturated. Thankfully, such ultrafilters exists, even without assuming GCH.

Theorem 5.6. *Let $\mathcal{I}$ be a set of size λ. Then there exists an ω-regular good ultrafilter on $\mathcal{I}$.*

Proof. The construction of such an ultrafilter is detailed in Section 6.1 of [2]. Theorem 6.1.4 in [2] is equivalent to this theorem by Proposition 4.3.4 from [2]. □

Theorem 5.7. *Let $\mathcal{I}$ be a set of size λ, and $\mathcal{D}$ be an ω-regular good ultrafilter on $\mathcal{I}$. Then $\mathcal{D}$ is in fact a good regular ultrafilter on $\mathcal{I}$.*

Proof. Suppose $\mathcal{D}$ is λ^+-good and ω-regular. Then, we can find some countable regularizing family $\{X_n \mid n \in \omega\} \subset \mathcal{D}$. Let $f : [\lambda]^{<\omega} \to \mathcal{D}$ be defined by $f(u) = X_{|u|}$, and let $g : [\lambda]^{<\omega} \to \mathcal{D}$ be the multiplicative refinement of f. We claim that $\{g(\{\alpha\}) \mid \alpha < \lambda\} \subset \mathcal{D}$ is a regularizing family of size λ. Let w be an infinite subset of λ and suppose for contradiction that there is some $t \in \bigcap_{\alpha \in w} g(\{\alpha\})$. Then for each $n < \omega$, we can choose $w_n \subset w$, with $|w_n| = n$, and observe that, for each $n < \omega$, we must have $t \in \bigcap_{\alpha \in w_n} g(\{\alpha\}) = g(w_n) \subset f(w_n) = X_n$ by the multiplicative property of g. Hence, $t \in \bigcap_{n \in \omega} X_n = \emptyset$, which implies that $\bigcap_{\alpha \in w} g(\{\alpha\}) = \emptyset$. This completes the proof. $\square$

Thus we know that there are good regular ultrafilters on sets of every infinite cardinality, so we can use 5.4 to get the following result.

Corollary 5.8. *Let $\mathfrak{M}$ be a model of a complete theory T in a countable language $\mathcal{L}$, and $\mathcal{D}$ be a good regular ultrafilter on a set $\mathcal{I}$ of power λ. Then the ultrapower $\mathfrak{N} = \mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ is λ^+-saturated.*

Thus, given any model $\mathfrak{M}$, we can always find λ^+-saturated ultrapowers of $\mathfrak{M}$. Does this always require a good ultrafilter, or are there models that have λ^+-saturated ultrapowers besides those constructed using good ultrafilters? It turns out such models do exists, and, in fact, looking at how good an ultrafilter has to be to saturate a particular model is an interesting question that gives rise to Keisler's order.

6. Keisler's Order

We want to define an ordering $\trianglelefteq$ on models defined on countable languages. Informally, we want to say that $\mathfrak{M} \trianglelefteq \mathfrak{N}$ if regular ultrapowers of $\mathfrak{N}$ are at least as hard to saturate as regular ultrapowers of $\mathfrak{M}$. Formally speaking, we say that $\mathfrak{M} \trianglelefteq \mathfrak{N}$ if for every infinite cardinal λ, and for every regular ultrafilter $\mathcal{D}$ defined on a set $\mathcal{I}$ of size λ, we have

$$\mathfrak{N}^{\mathcal{I}}/\mathcal{D} \text{ is } \lambda^+\text{-saturated} \quad \Rightarrow \quad \mathfrak{M}^{\mathcal{I}}/\mathcal{D} \text{ is } \lambda^+\text{-saturated.}$$

At first glance, it isn't clear why this ordering is interesting, as it is possible that the only way to guarantee a λ^+-saturated ultrapower is by using good ultrafilters, in which case there is only one equivalence class in the ordering. However, in this section we will show that Keisler's order can be defined in complete countable theories, and that its structure is highly nontrivial.

Theorem 6.1. *Let $\mathcal{D}$ be a regular ultrafilter on some set $\mathcal{I}$ of size λ and suppose $\mathfrak{M} \equiv \mathfrak{N}$. Then $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ is λ^+-saturated iff $\mathfrak{N}^{\mathcal{I}}/\mathcal{D}$ is λ^+-saturated.*

The original proof of this fact is due to Keisler in [1] as Corollary 2.1a. A different approach to this proof can be found in Chapter 6 of [3]. We will give a proof that relies on the existence of accurate distribution and 5.4. As in the previous section, we assume that our language are countable.

Proof of 6.1. Let $\mathfrak{M}$ and $\mathfrak{N}$ be elementarily equivalent $\mathcal{L}$-models, and $\mathcal{D}$ is a regular ultrafilter on $\mathcal{I}$, with $|\mathcal{I}| = \lambda$. Without loss of generality, assume that $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ is λ^+-saturated. We will show that $\mathfrak{N}^{\mathcal{I}}/\mathcal{D}$ is λ^+-saturated.

Suppose $p(x)$ is a type in $\mathfrak{N}^{\mathcal{I}}/\mathcal{D}$ with parameters from $B \subset N^{\mathcal{I}}/\mathcal{D}$, $|B| \leq \lambda$. Our goal is to find a type $q(x)$ in $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ such that $q(x)$ and $p(x)$ have the same accurate distribution. In fact, we will show that if we write $p(x) = \{\varphi_\alpha(x; \mathbf{b}_\alpha) \mid \alpha < \lambda, \mathbf{b}_\alpha \in [B]^{<\omega}\}$, then there is a collection $\{\mathbf{c}_\alpha \mid \alpha \leq \lambda, \mathbf{c}_\alpha \in \left(M^{\mathcal{I}}/\mathcal{D}\right)^{<\omega}\}$ such that $q(x) = \{\varphi_\alpha(x; \mathbf{c}_\alpha) \mid \alpha < \lambda\}$.

Let $F : [\lambda]^{<\omega} \to \mathcal{D}$ is the accurate distribution for $p(x)$ defined in 5.2. Then fix some $t \in \mathcal{I}$ and consider the set $u = \{\alpha \mid t \in F(\{\alpha\})\}$. We know that u is finite, so we may write $u = \{\alpha_1, \ldots, \alpha_k\}$ for $k \in \omega$. Then $\mathfrak{N}$ can satisfy the formulas $\varphi_{\alpha_1}(x; \mathbf{b}_{\alpha_1}[t]), \ldots, \varphi_{\alpha_k}(x; \mathbf{b}_{\alpha_k}[t])$ individually and perhaps some combinations of those formulas at the same time. It might be the case that $\mathfrak{N}$ satisfies all tuples, but only some tripes, and maybe only one group of four is satisfiable. In any case, there is a unique pattern, however complex, which we can exploit as follows. Define

$$R[t] = \left\{ \theta \subset u \mid \mathfrak{N} \models \exists x \bigwedge_{\alpha_j \in \theta} \varphi_{\alpha_j}[x, \mathbf{b}_{\alpha_j}[t]] \right\} \text{ and } S[t] = \left\{ \theta \subset u \mid \mathfrak{N} \models \neg \exists x \bigwedge_{\alpha_j \in \theta} \varphi_{\alpha_j}[x, \mathbf{b}_{\alpha_j}[t]] \right\}.$$

Observe that $R[t]$ and $S[t]$ are disjoint, and that $R[t] \cup S[t] = \mathcal{P}(u)$. More importantly, we have that

$$\mathfrak{N}, \mathfrak{M} \models \exists \mathbf{b}_1, \ldots, \mathbf{b}_k \left[\bigwedge_{\theta \in R[t]} \left(\exists x \bigwedge_{\alpha_j \in \theta} \psi_{\alpha_j}[x, \mathbf{b}_j] \right) \right] \wedge \left[\bigwedge_{\theta \in S[t]} \left(\neg \exists x \bigwedge_{\alpha_j \in \theta} \psi_{\alpha_j}[x, \mathbf{b}_j] \right) \right]$$

since we assumed that $\mathfrak{M} \equiv \mathfrak{N}$. This allows us to define $\mathbf{c}_\alpha[t]$ for all $\alpha < \lambda$. Start by fixing an element $0 \in M$, which require to be the same element for all t. Then, if $\beta \notin \{\alpha_1, \ldots, \alpha_k\}$, define $\mathbf{c}_\beta[t]$ to be a tuple of of the appropriate length with each entry equal to 0. For $\mathbf{c}_{\alpha_1}[t], \ldots, \mathbf{c}_{\alpha_k}[t]$, we use the fact that $\mathfrak{M}$ models the sentence above to find tuples of the appropriate such that

$$\mathfrak{M} \models \left[\bigwedge_{\theta \in R[t]} \left(\exists x \bigwedge_{\alpha_j \in \theta} \psi_{\alpha_j}[x, \mathbf{c}_{\alpha_j}] \right) \right] \wedge \left[\bigwedge_{\theta \in S[t]} \left(\neg \exists x \bigwedge_{\alpha_j \in \theta} \psi_{\alpha_j}[x, \mathbf{c}_j] \right) \right].$$

Finally, define $\mathbf{c}_\alpha = \langle \mathbf{c}_\alpha[t] \mid t \in \mathcal{I} \rangle / \mathcal{D}$.

We want to show show that $q(x)$ is a complete consistent type $\mathfrak{M}$ with parameters from a set of size no more than λ. Fix a unique representative for each $b \in B$, which we will also refer to as b, and define $B[t] = \{b[t] \mid b \in B\}$. Then, we can define a map $f_t : B \to M$ as follows. If $b[t]$ is mentioned in one of the finitely many tuples $\mathbf{b}_{\alpha_1}[t], \ldots, \mathbf{b}_{\alpha_k}[t]$, where $\{\alpha_1, \ldots, \alpha_k\} = \{\alpha \mid t \in F(\{\alpha\})$, then $f(b)$ is the element of M we chose to in the previous paragraph to play the same role; otherwise, we set $f(b) = 0$. This gives us a map $f : B \to M^{\mathcal{I}}/\mathcal{D}$ define by $f(b) = \langle f_t(b) \mid t \in \mathcal{I} \rangle / \mathcal{D}$, and our construction of the maps f_t guarantees that $\mathbf{c}_\alpha \in [f(B)]^{<\omega}$. Since $|f(B)| \leq |B| \leq \lambda$, we see that $q(x)$ has parameters from a set of size at most λ. Therefore, both $p(x)$ and $q(x)$ can be written in the language $\mathcal{L}(\lambda)$ (and are in fact the same in that language), so the fact that $p(x)$ is a complete type implies that $q(x)$ is a maximal consistent set of formulas.

It remains to show that $q(x)$ is finitely satisfiable. Note that we will omit the parameters to ease notation from now on. Let $p_0(x)$ be the corresponding subset of $p(x)$, meaning the same with the same formulas but with the different parameters, and let $w = \{\beta_1, \ldots, \beta_n\}$ be the set of indices of the formulas that appear in $q_0(x)$. Then, suppose that $w \subset \{\beta \mid t \in F(\{\beta\})\}$ for some $t \in \mathcal{I}$. This would imply that $w \subset R[t] \cup S[t]$, so our choice of parameters guarantees that $\mathfrak{M} \models (\exists x)\ q_0[x]$ if and only if $\mathfrak{M} \models (\exists x)\ p_0[x]$ at index t. In particular, this shows that

$$\{t \mid \mathfrak{M} \models (\exists x)\ q_0[x]\} \supset \{t \mid \mathfrak{N} \models (\exists x)\ p_0[x]\} \cap F(\{\beta_1\}) \cap \ldots \cap F(\{\beta_n\}) \in \mathcal{D},$$

which implies that $q_0(x)$ is finitely satisfiable in $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$.

Therefore, $q(x)$ is a complete consistent type in $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ with parameters on a set of size no more than λ. Moreover, we claim that our construction of $q(x)$ guarantees that F is, in fact, an accurate distribution for $q(x)$. To see that it is a distribution, notice that F satisfies conditions (2) and (3) from 5.1 since it is a distribution for $p(x)$; F also satisfies (1) and it is accurate for $q(x)$ thanks to our choice of parameters. Finally, recall that $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ is λ^+ saturated, so $q(x)$ is realized, which in turn gives us that F must have a multiplicative refinement. Therefore, using 5.4 one more time, we can see that $p(x)$ is realized in $\mathfrak{N}^{\mathcal{I}}/\mathcal{D}$. This completes the proof. $\square$

The main implication of this theorem, is that any two models of the same theory are $\trianglelefteq$-equivalent, so it makes sense to define Keisler's order in terms of theories as opposed to models.

Definition 6.2. Let T_1 and T_2 be two complete countable theories on possibly different languages. We say that $T_1 \trianglelefteq T_2$ if for some $\mathfrak{M} \models T_1$ and $\mathfrak{N} \models T_2$, for all $\lambda \geq \aleph_0$, and for all regular ultrafilters $\mathcal{D}$ defined on a set $\mathcal{I}$ of cardinality λ, we have

$$\mathfrak{N}^{\mathcal{I}}/\mathcal{D} \text{ is } \lambda^+\text{-saturated} \quad \Rightarrow \quad \mathfrak{M}^{\mathcal{I}}/\mathcal{D} \text{ is } \lambda^+\text{-saturated.}$$

It is clear from the definition that Keisler's order is both transitive and reflexive (note that this requires 6.1). However, antisymmetry is not obvious from the definition, and it is in fact not true. It is possible to have two complete countable theories T_1, T_2 such that $T_1 \trianglelefteq T_2$ and $T_2 \trianglelefteq T_1$, but $T_1 \neq T_2$ (there is an example of this at the end of this section). Therefore, we will pass to equivalence classes to obtain a partial order on the equivalence classes of complete countable theories. Then the reader might ask what the structure of this order is. While there is still much to discover, most of what we know is due to the work of Keisler, Shelah, and Malliaris. The first result we will discuss is that Keisler's order has a minimal class. To show this, we will show that there is a theory such that all of its ultrapowers by regular ultrafilters on sets of size λ are saturated. One such theory, and probably the most straightforward, is that of algebraically closed field of characteristic zero (ACF_0). First, we need the following lemma.

Lemma 6.3. *For every $\lambda > \aleph_0$ there exists at most one model of ACF_0 of size λ, up to isomorphism.*

Proof. This uses the fact that two algebraically closed fields are isomorphic if they have the same characteristic and transcendence degree. In the case where the cardinality of the field is uncountable, its transcendence degree is precisely its cardinality, so any two fields of characteristic zero and size $\lambda > \aleph_0$ must be isomorphic. $\square$

Thus, by the Łoś-Vaught test (see 2.8), we can see that ACF_0 is a complete theory, and the reader can easily verify that ACF_0 can be written in a countable language. Now we just need to show that ACF_0 is minimal.

Theorem 6.4. *Suppose $\mathfrak{M} \models ACF_0$ and $\mathcal{D}$ is a regular ultrafilter on $\mathcal{I}$ of size λ. Then $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ is λ^+-saturated.*

Proof. Since $\mathcal{D}$ is regular, it follows that $|\mathfrak{M}^{\mathcal{I}}/\mathcal{D}| = |\mathfrak{M}|^{\lambda}$. By 5.6 and 5.7, we know that there is some good, regular ultrafilter $\mathcal{E}$ over $\mathcal{I}$, which implies that the model $\mathfrak{M}^{\mathcal{I}}/\mathcal{E}$ is λ^+-saturated. Moreover, since $\mathcal{E}$ is also regular, $|\mathfrak{M}^{\mathcal{I}}/\mathcal{E}| = |\mathfrak{M}|^{\lambda}$, so by the previous lemma we must have $\mathfrak{M}^{\mathcal{I}}/\mathcal{D} \cong \mathfrak{M}^{\mathcal{I}}/\mathcal{E}$. Finally, since isomorphisms preserve the level of saturation, it must be the case that $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ is λ^+-saturated. $\square$

It is easy to see at this point that the equivalence class of all theories Keisler-equivalent to ACF_0 is minimal with respect to Keisler's order. We call this theory $\mathcal{T}_{min}$. A different characterization of this class can be given in terms of the *finite cover property.*

Definition 6.5. A theory T does not have the *finite cover property* (f.c.p.) if for every formula $\varphi(x; y)$ there exists an integer n depending only on φ such that for every $\mathfrak{M} \models T$, $A \subset M$, and for every $p \subset \{\varphi(x; a), \neg\varphi(x; a) \mid a \in A\}$, we have that if every $q \subset p$ with $|q| < n$ is consistent, then p is consistent.

Theories that do not have the finite cover property seem to have a satisfy version of the compactness theorem for some sets, where consistency for subsets of size at most some n implies the consistency of the whole set. A more extensive study of the f.c.p. can be found in Section 4 of Chapter 2 of [3], and in [5] and [4]. However, we mention the f.c.p primarily because of the following result.

Theorem 6.6. *$\mathcal{T}_{min}$ precisely is the set of complete countable theories without the finite cover property.*

The f.c.p. was originally defined by Keisler in [1] to construct unsaturated ultrapowers. In fact, one of the main results in Section 4 of [1] is that a theory that is minimal with respect to Keisler's order does not have the f.c.p. (see Corollary 4.2a in [1]). The proof of the converse statement, and therefore of 6.6, is due to Shelah and can be found in Chapter 6 of [3] as Theorem 5.8.

Our next goal is to show that there exists a maximal class with respect to Keisler's order, and give a characterization of this class with respect to good regular ultrafilters. Keisler's original proof of this fact relied on the existence of what he called *versatile formulas.* Essentially, models that have versatile formulas are capable of creating types that will be realized if and only if some other type in a different model is realized. Therefore, if a model with a versatile formula has a λ^+-saturated ultrapower, then every other model will also have a λ^+-saturated ultrapower. While this approach shows that a maximal class exists, it is also complex, so we will instead take a simpler approach by finding a theory that is saturated only by good regular ultrafilters.

As we have already shown, if $\mathfrak{M}$ is a model of some theory T and $\mathcal{D}$ is some good regular ultrafilter on some set $\mathcal{I}$ of cardinality λ, then $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ is λ^+-saturated, so good regular ultrafilters always give us the maximum level of saturation in the sense of Keisler's order. It would be nice if we could show that there is some theory T such that if $\mathfrak{M} \models T$ and $\mathcal{D}$ is a regular ultrafilter on a set $\mathcal{I}$ with $|\mathcal{I}| = \lambda$, the ultrapower $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ is λ^+-saturated if and only if $\mathcal{D}$ is good. However, it is also possible that there is no such theory, and that good ultrafilters are just too powerful to distinguish theories within Keisler's order. Thankfully, the next result shows that such theories exist, so Keisler's order has a maximal class.

Theorem 6.7. *Let $\mathcal{L} = \{\subset\}$ and let $M = [\omega]^{<\omega}$. Then M can be made into a $\mathcal{L}$-structure $\mathfrak{M}$ under the usual interpretation of $\subset$. Moreover, if $\mathcal{D}$ is a regular ultrafilter on $\mathcal{I}$, with $|\mathcal{I}| = \lambda$, then $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ is λ^+-saturated if and only if $\mathcal{D}$ is good.*

Proof. The first part of the theorem is trivial, so we will focus on the second part. Suppose $\mathcal{D}$ is not good; we will show that the ultrapower is not λ^+-saturated. Since $\mathcal{D}$ is not good, there exists a function $f : [\lambda]^{<\omega} \to \mathcal{D}$ without a multiplicative refinement. Our goal is to find a elements $\{m_\alpha \mid \alpha < \lambda\} \subset M^{\mathcal{I}}$ such that f is an accurate distribution of $p(x) = \{x \subsetneq m_\alpha/\mathcal{D} \mid \alpha < \lambda\}$. We can assume that $\{f(\{\alpha\}) \mid \alpha < \lambda\}$ is a regularizing family, because otherwise we can define $f'(u) = f(u) \cap \bigcap_{\alpha \in u} X_\alpha$ for $u \in [\lambda]^{<\omega}$ and $\{X_\alpha \mid \alpha < \lambda\} \subset \mathcal{D}$ a regularizing family (here f' has a multiplicative refinement if and only if f does). Now, by regularity, we have that for each $t \in \mathcal{I}$ there is some $n \in \omega$ such that $\alpha_1, \dots, \alpha_n$ are exactly those α such that $t \in f(\{\alpha\})$. Thus, we can define $m_\beta[t] = \emptyset$ for $\beta \notin \{\alpha_1, \dots, \alpha_n\}$ and define $m_{\alpha_1}[t], \dots, m_{\alpha_n}[t]$ so that for $u \subset \{\alpha_1, \dots, \alpha_n\}$ we have

$$t \in f(u) \iff \mathfrak{M} \models \exists x \bigwedge_{\alpha \in u} x \subsetneq m_\alpha[t].$$

For instance, we could construct $m_{\alpha_i}[t]$ in stages, first making sure that we satisfy the singletons, then the pairs, then the triples, and so on. Setting $m_\alpha = \prod_{t \in \mathcal{I}} m_\alpha[t]$, we have that f is an accurate distribution for $p(x)$. Therefore, by 5.4, we see that $p(x)$ is not realized. $\square$

Thus, if we let $T = Th(\mathfrak{M})$, with $\mathfrak{M}$ defined as in the theorem above, it is clear that the equivalence class $\mathcal{T}_{max}$ of T is maximal with respect to Keisler's order.

Keisler used versatile formulas in [1] that the theory of $\mathfrak{N} = (\mathbb{N}; +, \times, 0, 1)$, with the usual interpretations for all the symbols, is also in $\mathcal{T}_{max}$. Therefore, since $T(\mathfrak{N})$ is not equal to the theory of the model in 6.7, this shows that $\trianglelefteq$ is preorder on theories.

7. Linear orders

In this section, we will explore the theory of linear orders, present some recent work by Malliaris and Shelah that has solved a number of long-standing questions in set theory. Moreover, we will use these results to show that the theory of linear orders is maximal in Keisler's order. A proof of this fact can be found in Section 4 of Chapter 6 of [3] and uses a different approach.

Definition 7.1. A *linear order* is a model $\mathfrak{X}$ of the language $\mathcal{L} = \{<\}$ such that $\mathfrak{X}$ satisfies the standard axioms of a linear order.

It should be easy to see that $(\mathbb{N}, <)$, $(\mathbb{Q}, <)$ and $(\mathbb{R}, <)$ are examples of linear orders. The reader should also note that the ultrapower of a linear order is also a linear order by Łoś's theorem. As we have done before, we want to characterize the complexity of linear orders by looking at its ultrapowers. However, in this case we are specially interest in a particular kind of type, which we call a pre-cut.

Definition 7.2. A cardinal λ is said to be *regular* if

$$\lambda = \inf\{\kappa \mid \exists \langle a_\beta \mid \beta < \kappa \rangle \text{ such that } \lim_{\beta \to \kappa} a_\beta = \lambda\}.$$

By taking the sequence of all cardinals less than λ, it is clear that the infimum of this set exists and it is at most λ. If the infimum is less than λ, we say that λ is a *singular* cardinal.

Definition 7.3. Given a linear order $\mathfrak{X}$, with $|\mathfrak{X}| = \lambda$ regular, a (κ_1, κ_2)*-pre-cut in* $\mathfrak{X}$ is a pair of sequences $\langle a_\alpha \in X \mid \alpha < \kappa_1 \rangle$ and $\langle b_\beta \in X \mid \beta < \kappa_2 \rangle$ such that $a_\alpha < a_\gamma < b_\tau < b_\beta$ for all $\alpha < \gamma < \kappa_1$ and $\beta < \tau < \kappa_2$. If there does not exist a $c \in X$ satisfying $a_\alpha < c < b_\beta$ for all $\alpha < \kappa_1$ and $\beta < \kappa_2$, we call this is a (κ_1, κ_2)-*cut* instead.

If we assume that $\kappa_1 + \kappa_2 \leq \lambda$, then a pre-cut can be represented as a type in the ultrapower. Therefore, a cut represents a type that is not realized, so the existence of cuts implies that the ultrapower is not saturated. What can we say about linear orders without cuts? First, let us make our definitions precise by introducing the cut spectrum.

Definition 7.4. Given an regular ultrafilter $\mathcal{D}$ on λ and a liner order $\mathfrak{X}$, we define the *cut spectrum* of $\mathcal{D}$ and $\mathfrak{X}$ as

$$C(\mathcal{D}, \mathfrak{X}) = \{(\kappa_1, \kappa_2) \mid \kappa_1, \kappa_2 \text{ are infinite and regular, } \kappa_1 + \kappa_2 < \lambda \text{ and } \mathfrak{X}^\lambda/\mathcal{D} \text{ has a } (\kappa_1, \kappa_2) \text{ cut}\}.$$

The requirement that κ_1, κ_2 and λ be all regular cardinals prevents some pathological cases. Recall that 6.1 shows that saturation of ultrapowers depends only on the theory of the model and not on the model itself. In fact, this is also the case for the cut spectrum.

Theorem 7.5 (Fact 1.3 in [10])**.** *If $\mathcal{D}$ is a regular ultrafilter and $\mathfrak{N} = (N, <^{\mathfrak{N}})$, $\mathfrak{M} = (M, <^{\mathfrak{M}})$ are any two infinite linearly ordered sets, then $C(\mathcal{D}, \mathfrak{N}) = C(\mathcal{D}, \mathfrak{M})$.*

Proof. Suppose there is a (κ, θ)-cut in $\mathfrak{N}^{\mathcal{I}}/\mathcal{D}$ that is witnessed by two sequences $\langle a_\alpha/\mathcal{D} \mid \alpha < \kappa\rangle$ and $\langle b_\beta/\mathcal{D} \mid \beta < \theta\rangle$ in $N^{\mathcal{I}}/\mathcal{D}$. Then we will show that $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ also has a (κ, θ)-cut.

Define $A = \{a_\alpha \mid \alpha < \kappa\}$ and $B = \{b_\beta \mid \beta < \theta\}$. Since $\mathcal{D}$ is regular there is a regularizing family $\{X_\alpha \mid \alpha < \lambda\}$ and, since $\kappa + \theta \leq \lambda$, we may choose $\kappa + \theta$ distinct sets of the regularizing family and index them as $\{Y_\alpha \mid \alpha < \kappa\}$ and $\{Z_\beta \mid \beta < \theta\}$.

Next, let $f : A \cup B \to \mathcal{D}$ be a map defined by $a_\alpha \mapsto Y_\alpha$ and $b_\beta \mapsto Z_\beta$ for $\alpha < \kappa$ and $\beta < \theta$, and define sets $A[t] = \{a_\alpha[t] \mid t \in f(a_\alpha)\}$ and $B[t] = \{b_\beta[t] \mid t \in f(b_\beta)\}$ for each $t \in \mathcal{I}$. Observe that $A[t]$ and $B[t]$ are finite for all $t \in \mathcal{I}$ by regularity. Moreover, $A[t] \cup B[t]$ is a discrete linearly ordered set since $\mathfrak{N}$ is a linear order. Next, we fix an order preserving injection $\mathbf{i}_t : A[t] \cup B[t] \to M$ such that

$$\mathfrak{N} \models (\exists z)(x < z < y) \implies \mathfrak{M} \models (\exists z)(\mathbf{i}_t(x) < z < \mathbf{i}_t(y)) \tag{1}$$

for all $x, y \in A[t] \cup B[t]$. Observe that since $\mathfrak{M}$ is an infinite linear order, there is at least one such injection for all t, so we fix a particular one.

We now define our sequences in $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ as follows. First, fix some element $0 \in \mathfrak{M}$. Then, for each $\alpha < \kappa$ and for each $t \in \mathcal{I}$, define $c_\alpha[t]$ to be $\mathbf{i}_t(a_\alpha[t])$ if $t \in Y_\alpha$, and zero otherwise. Similarly, for each $\beta < \theta$ and for each $t \in \mathcal{I}$ define $d_\beta[t]$ to be $\mathbf{i}_t(b_\beta[t])$ if $t \in Z_\beta$, and zero otherwise. Finally, let $c_\alpha = \langle c_\alpha[t] \mid t \in \mathcal{I}\rangle/\mathcal{D}$ and $d_\beta = \langle d_\beta[t] \mid t \in \mathcal{I}\rangle/\mathcal{D}$.

We claim that the resulting sequences for a pre-cut in $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$. Fix $\alpha < \gamma < \kappa$ and $\beta < \tau < \theta$, and observe that if $c_\alpha[t] = \mathbf{i}_t(a_\alpha[t]), c_\gamma[t] = \mathbf{i}_t(a_\gamma[t]), d_\tau[t] = \mathbf{i}_t(b_\tau[t]), d_\beta[t] = \mathbf{i}_t(b_\beta[t])$ and $\mathfrak{N} \models a_\alpha[t] <^{\mathfrak{N}} a_\gamma[t] <^{\mathfrak{N}} b_\tau[t] <^{\mathfrak{N}} b_\beta[t]$ then the definition of $\mathbf{i}_t$ guarantees that $\mathfrak{M} \models c_\alpha[t] <^{\mathfrak{M}} c_\gamma[t] <^{\mathfrak{M}} d_\tau[t] <^{\mathfrak{M}} d_\beta[t]$. Therefore, by construction we have that

$$\{t \mid \mathfrak{M} \models c_\alpha[t] <^{\mathfrak{M}} c_\gamma[t] <^{\mathfrak{M}} d_\tau[t] <^{\mathfrak{M}} d_\beta[t]\} \supset Y_\alpha \cap Y_\gamma \cap Z_\beta \cap Z_\tau \cap \{t \mid \mathfrak{N} \models a_\alpha[t] <^{\mathfrak{N}} a_\gamma[t] <^{\mathfrak{N}} b_\tau[t] <^{\mathfrak{N}} b_\beta[t]\}$$

where the intersection on the right is in the ultrafilter. Thus, the result follows by applying Łoś's theorem.

It suffices now to show that this is actually a cut in $\mathfrak{M}^{\mathcal{I}}/\mathcal{D}$. Define a type $p(x)$ with parameters from $\langle c_\alpha/\mathcal{D} \mid \alpha < \kappa\rangle \cup \langle d_\beta/\mathcal{D} \mid \beta < \theta\rangle$ by

$$p(x) = \{c_\alpha <^{\mathfrak{M}} x \mid \alpha < \kappa\} \cup \{x <^{\mathfrak{M}} d_\beta \mid \kappa \leq \beta < \kappa + \theta\} \cup \{x = x \mid \kappa + \theta \leq \gamma < \lambda\}.$$

Notice that this type is realized if, and only if, the sequences $\langle c_\alpha/\mathcal{D} \mid \alpha < \kappa\rangle$ and $\langle d_\beta/\mathcal{D} \mid \beta < \theta\rangle$ from a pre-cut but not a cut. Therefore, assume for the sake of contradiction that that there exists some $m/\mathcal{D} \in \mathfrak{M}^{\mathcal{I}}/\mathcal{D}$ that realizes the type. Then there exists a distribution function d with a multiplicative refinement ρ. We claim that we can use ρ to realize the cut defined by $\langle a_\alpha/\mathcal{D} \mid \alpha < \kappa\rangle$ and $\langle b_\beta/\mathcal{D} \mid \beta < \theta\rangle$ as follows. For each $t \in \mathcal{I}$, there is a set $u_t \in [\lambda]^{<\omega}$ such that u_t is the largest set with $t \in \rho(u_t)$. Since ρ is a distribution for $p(x)$, this implies that there is a chain of the form

$$c_{\alpha_1}[t] < \ldots < c_{\alpha_{k(t)}}[t] < m[t] < d_{\beta_{l(t)}} < \ldots < d_{\beta_1}[t],$$

where $k(t), l(t) < \omega$. Then, by the definition of $\mathbf{i}_t$, we know that

$$a_{\alpha_1}[t] < \ldots < a_{\alpha_{k(t)}}[t] < b_{\beta_{l(t)}}[t] < \ldots < b_{\beta_1}[t],$$

so we define $n[t]$ to be a point above all the a_{α_i} and below the b_{β_j} if one exists, and let $n[t] = a_1[t]$ otherwise. We also define $a[t] = n[t]$ if $n[t] \neq 1$, and otherwise let $a[t]$ be the largest a_{α_i} in the finite chain above. Notice that, since ρ is a distribution function, each a_α and b_β appears in this maximal chain for a large set of $t \in \mathcal{I}$, meaning that that set is in the ultrafilter. Now, if $n/\mathcal{D}$ fills the cut in $\mathfrak{N}^\mathcal{I}/\mathcal{D}$, then we have a contradiction and we are done. Otherwise, consider $a/\mathcal{D}$. Notice that $a/\mathcal{D}$ is strictly smaller than each of the b_β, and there are two possibilities: either $a/\mathcal{D} = a_\alpha/\mathcal{D}$ for some $\alpha < \kappa$ or not. If $a/\mathcal{D} = a_\alpha/\mathcal{D}$, then $a_\alpha/\mathcal{D}$ has to be the largest element of the sequence $\langle a_\alpha/\mathcal{D} \mid \alpha < \kappa\rangle$, but this can't be the case since κ is a limit ordinal and the sequence is assumed to be strictly increasing. Otherwise, $a/\mathcal{D} > a_\alpha/\mathcal{D}$, so $a/\mathcal{D}$ fill the cut in $\mathfrak{N}^\mathcal{I}/\mathcal{D}$, which is also a contradiction. Since all cases lead to a contradiction, it must be the case that the sequences $\langle c_\alpha/\mathcal{D} \mid \alpha < \kappa\rangle$ and $\langle d_\beta/\mathcal{D} \mid \beta < \theta\rangle$ form a cut in $\mathfrak{M}^\mathcal{I}/\mathcal{D}$, which completes the proof. □

Therefore, the cut spectrum depends not on the model but only on the ultrafilter, so we can redefine the cut spectrum as follows.

Definition 7.6. Given an regular ultrafilter $\mathcal{D}$ on λ, we define the *cut spectrum* of $\mathcal{D}$ as

$$C(\mathcal{D}) = \{(\kappa_1, \kappa_2) \mid \kappa_1, \kappa_2 \text{ are regular, } \kappa_1 + \kappa_2 < \lambda \text{ and } (\mathbb{N}, <)^\lambda/\mathcal{D} \text{ has a } (\kappa_1, \kappa_2) \text{ cut}\}.$$

Moreover, if we set $\mathfrak{M} = \mathfrak{N}$ and choose the functions $\mathbf{i}_t$ to be order reversing, we can deduce that $(\kappa, \theta) \in C(\mathcal{D})$ if and only if $(\theta, \kappa) \in C(\mathcal{D})$.

We can now ask specific questions about the cut spectrum. When is the cut spectrum of an ultrafilter empty? What are the implications of having an empty cut spectrum? One might hope that answering this questions can help use place the theory of linear orders within Keisler's. We should point out that there are types in the ultrapower that do not come from pre-cuts, so having an empty cut spectrum might not tell use anything conclusive concerning saturation. Nonetheless, Malliaris and Shelah showed in [9] that the cut spectrum can tell us if the ultrafilter $\mathcal{D}$ is good. Even more surprising is the fact that we only have to look at a particular type of cut to show that the ultrafilter is good.

Theorem 7.7 (Theorem 10.25 in [9]). *Let $\mathcal{D}$ be a regular ultrafilter on $\mathcal{I}$, with $|\mathcal{I}| = \lambda$. Then the following are all equivalent:*

(1) $C(\mathcal{D}) = \emptyset$;
(2) If $\kappa < \lambda$ is a regular cardinal, then $(\kappa, \kappa) \notin C(\mathcal{D})$;
(3) $\mathcal{D}$ is good.

In particular, this implies that the ultrapower of a linear order is saturated if and only $\mathcal{D}$ is a good ultrafilter. Thus, the theory of linear orders is maximal in Keisler's order. While we have only stated this result to show that linear orders are maximal, 7.7 has some other consequences that Malliaris and Shelah were able to use to solve a long standing question in set theory in [9].

8. Open problems and unanswered questions

In the previous section we gave a vague description of the structure of Keisler's order by showing that it has a minimal and a maximal class. One must wonder what else can be said about the structure of Keisler's order. For about fifty years, there wasn't much more to be said, with many believing that it was a simple linear order with four classes (see [7] for more details). Malliaris and Shelah uncovered a bit more of the structure in recent work and answered a couple of long-standing questions. In particular, they show that Keisler's order has many more than four classes. As always, we give a couple of definitions before discussing the main results.

Definition 8.1. A model $\mathfrak{M}$ is said to be λ-*stable* if for every set of parameters $B \subset M$, of size no more than λ, there are no more than λ many complete n-types over B in $\mathfrak{M}$, for all $n \in \omega$. Moreover, a theory T is said to be λ-*stable* if every model $\mathfrak{M} \models T$ is λ-stable, and it is *stable* if is said to be λ-stable for some λ.

Intuitively, if T is a stable theory, then the models of T don't have "too many" types. In fact, its turns out that models of stable theories are easy to saturate using the ultrapower construction. In particular, if T_1 is a stable theory and T_2 is an unstable theory, then $T_1 \trianglelefteq T_2$.

With the above definitions in mind, we state an important result to explain the present picture.

Theorem 8.2. *Let $\mathcal{T}_{min}$ and $\mathcal{T}_{max}$ be as defined in the previous section. Then*

(1) *There is a class of theories in Keisler's order called $\mathcal{T}_{stable}$ that contains all the stable theories with the finite cover property. This class is strictly above $\mathcal{T}_{min}$ and strictly below $\mathcal{T}_{max}$.*
(2) *Let T_C be the theory of a random graph. Then T_C is a countable complete theory which is unstable and it is not an element of the maximal class. Hence, the equivalence class $\mathcal{T}_C$ is distinct from $\mathcal{T}_{min}, \mathcal{T}_{max}$, and $\mathcal{T}_{stable}$. Moreover, it is the minimal class containing an unstable theory.*
(3) *There exists a countable collection $\{T_n \mid n \in \omega\}$ of complete countable theories such that the T_n form a strictly descending chain with respect to Keisler's order. We denote the equivalence class of T_n as $\mathcal{T}_n$.*

In particular, we have that the equivalence classes defined above satisfy the following relation:

$$\mathcal{T}_{min} \lhd \mathcal{T}_{stable} \lhd \mathcal{T}_C \lhd \ldots \lhd \mathcal{T}_{n+1} \lhd \mathcal{T}_n \lhd \ldots \lhd \mathcal{T}_1 \lhd \mathcal{T}_{max}.$$

While (1) and (2) have been known for some time, (3) is a new and exciting result by Malliaris and Shelah in [7]. This theorem gives us two big insights into the structure of Keisler's order. First, it shows that the order has infinitely many classes, and second, it shows that it is not a well-ordering.

However, there are still many open question about the structure of Keisler's order, and about the classes that comprise it. This first question arises naturally from the previous theorem.

Question 8.3. *Is Keisler's order a linear order?*

The question remains open in ZFC. However, under some set theoretic assumptions, the answer is no. If we assume the existence of a super-compact cardinal (we ommit the definition since it requires a strong background in set theory, but can be found in [12]), then there are equivalence classes of theories that are incomparable in Keisler's order. This result was discovered by Malliaris and Shelah, and independently by Ulrich, and it follows from the techniques developed in [8].

Another interesting question concerns the maximal class $\mathcal{T}_{max}$. While we have been able to characterize the maximal class in terms of saturation only by good ultrafilters, one might wonder if there is a set of axioms Σ such that a theory T is in the maximal class if, and only if, every model of T is also a model of Σ. While this question remains unanswered, Malliaris and Shelah have the following conjecture.

Definition 8.4. A theory T has the SOP_n property, for $n \geq 3$, if there there is a formula $\varphi(\mathbf{x}; \mathbf{y})$, with $\mathbf{x}$ and $\mathbf{y}$ of length k, such that for every model $\mathfrak{M} \models T$ there is a sequence $\{\mathbf{a}_i \mid i \in \omega\}$ with each $\mathbf{a}_i \in M^k$ and such that

(1) $M \models \varphi(\mathbf{a}_i; \mathbf{a}_j)$ for all $i < j < \omega$;

(2) $M \models \neg\exists x_1, \ldots, x_n \left(\bigwedge\{\varphi(x_i, x_j) \mid \ i < j < n \text{ and } j = i+1 \ mod\ n\}\right)$.

This means that φ acts like an ordering on infinitely many elements with no cycles of length n.

Definition 8.5. A theory T has the SOP_2 property (also known as TP_1), if there there is a formula $\varphi(x; \mathbf{y})$, with $\mathbf{y}$ of some length k, such that for every model $\mathfrak{M} \models T$ there is a sequence of parameters $\{\mathbf{a}_\eta \mid \eta \in 2^{<\omega}\}$ with each $\mathbf{a}_\eta \in M^k$ such that

(1) for all $\eta, \rho \in 2^{<\omega}$ incomparable (i.e $\neg(\eta \preceq \rho) \wedge \neg(\rho \preceq \eta)$), we have that $\{\varphi(x; \mathbf{a}_\eta), \varphi(x; \mathbf{a}_\rho)\}$ is inconsistent;

(2) for all $\nu \in 2^\omega$, the set $\{\varphi(x; \mathbf{a}_{\nu\restriction n}) \mid n \in \omega\}$ is a consistent partial type.

Conjecture 8.6. *SOP_2 characterizes the maximal class, meaning that $T \in \mathcal{T}_{max}$ iff T has SOP_2.*

While it is know that T having SOP_2 implies that T is maximal in Keisler's order (see [9]), necessity has not been proven so far.

In conclusion, Keisler's order is an interesting, and nontrivial object, that highlights some interesting properties of complete countable theories. 8.2 shows that this ordering is quite complex, and the open questions we have mentioned in this section show that there is still much we don't know about Keisler's order. Therefore there is still much work to be done in uncovering and understanding the structure of Keisler's order.

Acknowledgments. I would like to thank my mentor Sarah Reitzes for her input, comments and thoughts throughout the REU, and Alexander Kastner and Alex Johnson for all the discussion surrounding model theory this summer. I would also like to thank Professor Maryanthe Malliaris for suggesting the topic for this paper and being an outstanding mentor without whom this paper would not have been possible. Finally, I would like to thank Professor Peter May for making this REU possible.

References

[1] H. J. Keisler. *Ultraproducts which are not Saturated.* The Journal of Symbolic Logic, Vol. 32, No. 1, pp. 23-46; 1967.

[2] C. C. Chang and H. J. Keisler. *Model theory,* North-Holland Publishing Co., Amsterdam, Third edition; 1990.

[3] Saharon Shelah. *Classification theory and the number of non-isomorphic models,* Studies in Logic and the Foundation of Mathematics, Volume 92; 1990.

[4] Baldwin, Grossberg and Shelah. *Characterizations of the finite cover property and stability,* http://repository.cmu.edu/cgi/viewcontent.cgi?article=1569&context=math; 1995.

[5] Baldwin, Grossberg and Shelah. *Transfering saturation, the finite cover property, and stability,* The Journal of Symbolic Logic 64, no. 2, 678-684; 1999.

[6] Maryanthe Malliaris. *Persistence and regularity in unstable model theory* University of California, Berkeley; 2009.

[7] M. Malliaris and S. Shelah. *Keisler's order has infinitely many classes,* arXiv:1503.08341; 2015.

[8] M. Malliaris and S. Shelah. *Existence of optimal ultrafilters and the fundamental complexity of simple theories,* arXiv:1404.2919; 2014.

[9] M. Malliaris and S. Shelah. *Cofinality spectrum theorems in model theory, set theory and general topology,* arXiv:1208.5424; 2012

[10] M. Malliaris and S. Shelah. *Open problems on ultrafilters and some connections to the continuum,* Contemporary Mathematics Volume 690; 2017.

[11] Douglas Ulrich. *Keisler's order is not linear, assuming a supercompact,* arXiv:1702.0123; 2017.

[12] Thomas Jech. *Set theory, third millennium edition (revised and expanded),* Springer; 2002.