A DISCUSSION OF KEISLER'S ORDER

ADELE PADGETT

Abstract. In model theory, the complex numbers as an algebraically closed field are often given as an example of a simple, well-behaved mathematical structure, while the reals as a dense linear order without endpoints are more complicated and behave more strangely in many ways. Can we compare the complexity of *other* mathematical structures in a meaningful and informative way? Keisler's order seeks to address this question by moving from the realm of structures to the theories that describe them. The complexities of two theories can be compared by measuring the relative saturation of certain ultrapowers of each theory's models. The details of how this comparison works will be discussed in detail and justified at each step.

Contents

1. Introduction

One of the reasons model theory is powerful is that it allows us to manipulate general mathematical structures in terms of sets of statements that are true about them, called theories. An example of this is Morley's categoricity theorem, which states that if there is only one structure up to isomorphism satisfying a theory in some uncountable cardinality, then there is only one such structure up to isomorphism in *any* uncountable cardinality. Now, $\mathbb{C}$ is the only algebraically closed field of characteristic zero and cardinality $\mathfrak{c}$ up to isomorphism, while there are many dense linear orders without endpoints of cardinality $\mathfrak{c}$ which are not isomorphic to $\mathbb{R}$. This is one of the reasons $\mathbb{C}$ as an algebraically closed field is relatively simple and easy to deal with, and $\mathbb{R}$ as a dense linear order is not. However, Morley's theorem tells us that there is only one algebraically closed field of characteristic zero but at least two (and in fact many, by later work) dense linear orders without

Date: September 23, 2016.

endpoints in *any* uncountable cardinality, indicating that this difference between $\mathbb{C}$ and $\mathbb{R}$ is actually a specific instance of a difference in the theories describing them. Indeed, the complexities of many other mathematical structures can be compared by looking at the theories that they *model*. Keisler's order is a preorder on theories given by comparing the *saturation* of special *ultrapowers* of the models of each theory.

Before going into detail about the pieces needed to construct Keisler's order, it helps to have some idea of what is known about it. Considering equivalence classes in the preorder gives a partial order, which can be depicted as the following:

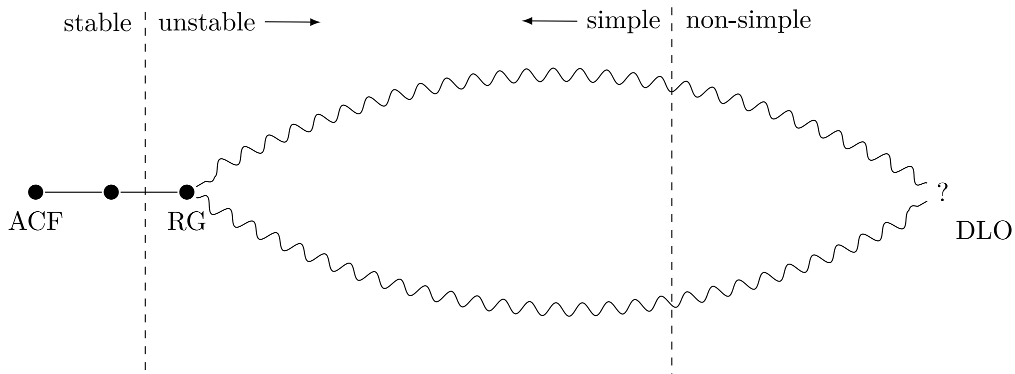

Keisler's order was originally believed to have only finitely many equivalence classes, but it is actually much richer and more complicated than was first thought. The minimum equivalence class contains the theory of algebraically closed fields (ACF), and the minimum two classes make up the *stable* theories. The third equivalence class contains the minimal *unstable* theories, including the theory of random graphs (RG). Beyond this, the order starts to get complicated. There is an infinite descending chain of simple theories, so the order is infinite and not a well-order. At some point, there is a division between the *simple* and *non-simple* theories (under a set-theoretic hypothesis). There is a maximum class containing the the theory of dense linear orders (DLO), but finding a meaningful characterization of the theories that form the maximum class is an open problem. It is even possible that all non-simple theories form the maximum class. Also, the middle section of the picture is wide to leave open the possibility of incomparable classes of theories, but the order may actually be linear.

The rest of this section contains basic definitions and results in model theory that are necessary later on in this paper. More detailed information can be found in the first chapter of Chang and Keisler's *Model Theory* [1].

1.1. **Preliminaries.** A *first order language* $\mathscr{L}$ consists of specified relation, function, and constant symbols along with parentheses, variables, logical connectives, and quantifiers. Every language assumes an identity relation $\equiv$, and for this paper, it is assumed that all languages are at most countable. A model M of a language $\mathscr{L}$ consists of a non-empty set $\text{dom}(M)$ called the *domain* along with an *interpretation* that maps each n-placed relation symbol R in $\mathscr{L}$ to a set $R' \subset \text{dom}(M)^n$, each m-placed function symbol F to function $F' : \text{dom}(M)^m \to \text{dom}(M)$, and each constant symbol in $\mathscr{L}$ to a constant in $\text{dom}(M)$. In this paper, it will cause no confusion to conflate the notations for a model M and its underlying set $\text{dom}(M)$, so M will be used for both.

A *term* in a language $\mathscr{L}$ is some expression that evaluates to an element of a model when values are substituted for its free variables. A term can be a variable, a constant, or if $t_1, \ldots, t_n$ are terms and F is an n-placed function symbol, then $F(t_1, \ldots, t_n)$ is a term. *Formulas* are expressions that evaluate to true or false when values are substituted for their free variables. *Atomic formulas* are built by applying an n-placed relation symbol to n terms, and formulas are built from atomic formulas by adding logical connectives and quantifying over variables. A formula with no free variables is called a *sentence*. The notations $t(v_1, \ldots, v_n)$ and $\varphi(v_1, \ldots, v_n)$ are used for a term or formula with free variables among $v_1, \ldots, v_n$, and $t[a_1, \ldots, a_n]$ and $\varphi[a_1, \ldots, a_n]$ are used for a term or formula evaluated at elements $a_1, \ldots, a_n$ of a given model.

Example 1.1. $\mathscr{L} = \{+, \cdot, 0, 1\}$ is often called the language of rings because it contains all the symbols needed to express the ring axioms. In $\mathscr{L}$, $t(v, w) = v \cdot w$ is a term, while

$$\varphi(v) = \exists w((v \cdot w = 1) \wedge (w \cdot v = 1))$$

which states that v has a 2-sided multiplicative inverse, is a formula. In the model $(\mathbb{Q}, +, \cdot, 0, 1)$, $\varphi[a]$ is true for all non-zero a in M, but in $(\mathbb{Z}, +, \cdot, 0, 1)$, $\varphi[a]$ is only true when $a = \pm 1$.

Implicit in this example is what it means for a sentence φ to be "true" or "not true" in a model M; notationally, $M \vDash \varphi$. Sentences of the form $\varphi \wedge \psi$ and $\neg\varphi$ are evaluated inductively based on the truth of φ and ψ, and sentences of the form $(\forall v)\sigma$ are evaluated by substituting each $a \in M$ for v and then asking whether there are other elements of M that can be substituted for the remaining free and bound variables of σ to make σ true. If the same sentences are true in two models M and N, we say these models are *elementarily equivalent* and write $M \equiv N$[1].

Unlike sentences, the truth of most formulas depends on what values are substituted for their free variables. One way to think about when formulas are true is in terms of their *solution sets*. The solution set of a formula ψ is the set of elements, or tuples of elements, on which ψ evaluates to true. In the example above, $\{\pm 1\}$ is the solution set of φ. To satisfy multiple formulas simultaneously, pick an element from the intersection of their solution sets. If there is some model in which such an element exists, the formulas are said to be *consistent*.

A set of sentences Σ is consistent if not every formula φ in the underlying language can be deduced from it. Any formula can be deduced from a formula of the form $\varphi \wedge \neg\varphi$, so deducing this is sufficient to prove inconsistency. A deduced inconsistency in a set of sentences must come from a finite proof using only finitely many sentences of Σ, so a set Σ of sentences is consistent if and only if every finite subset of Σ is consistent. Note that if a sentence is thought of as a formula with n free variables whose value is independent of the sentence's truth, then the definition of a consistent set of sentences is in accordance with the definition of a consistent set of formulas: In a model M, if a sentence is false, it corresponds to $\varnothing$, and if it is true, it corresponds to M^n, so for any formula φ, $\varphi \wedge \neg\varphi$ corresponds to the empty set. The following important theorem in model theory establishes a connection between the *semantic* notion of truth in a model and the *syntactic* notion of consistency.

Theorem 1.2 (Completeness)**.** *A set Σ of sentences is consistent if and only if Σ has a model.*

Although this theorem only guarantees one model, a consistent set of sentences can have many different models. In the other direction, all sentences true in a model M form a consistent set. A consistent set of sentences whose set of consequences is maximal (meaning that for every sentence

[1]This relation between two models should not be confused with the identity relation in a language, which holds between two elements of a model

φ, either φ or $\neg\varphi$ can be deduced), is called *complete*. The set of all sentences true in a model M is complete, and so is any smaller set from which all sentences true in M can be deduced. Keiser's Order deals with consistent sets of sentences, also called *theories*, that are complete. The following examples explain the theories mentioned in the earlier discussion of Keisler's Order.

Example 1.3. The theory of algebraically closed fields in the language $\mathscr{L} = \{+, \cdot, 0, 1\}$ consists of the field axioms along with the following statement for each $n \in \mathbb{N}$:

$$\forall a_0 \dots \forall a_n \exists x (a_n \cdot x^n + \dots + a_1 \cdot x + a_0 = 0)$$

where x^n is shorthand for $x \cdot \ldots \cdot x$ n-times. Together, these statements say that every polynomial has a root. For a prime p, adding the statement $0 = 1 + \dots + 1$ (p-times) gives the theory of algebraically closed fields of characteristic p. Adding the statements $0 \neq 1 + \dots + 1$ (n-times) for all $n \in \mathbb{N}$ gives the theory of algebraically closed fields of characteristic 0. The theories of algebraically closed fields of specified characteristic are *complete*.

Example 1.4. The theory of random graphs uses the language containing one binary relation R. The underlying set of any model is the vertex set of a graph, and the interpretation of R gives the edge set. The theory consists of a statement for each $n \in \mathbb{N}$ saying there are at least n distinct vertices (together these say the graph is infinite), and the following statement for each n.

$$(\forall y_1 \dots \forall y_n \forall z_1 \dots \forall z_n) \bigvee_{i,j=1}^{n} (y_i = z_j) \vee \exists x \left(\bigwedge_{i \leq n} x R y_i \wedge \bigwedge_{j \leq n} \neg x R z_j \right)$$

These axioms come from a result about random graphs: Fix $0 < p < 1$ and form a graph G with an infinite vertex set by letting an edge occur between each pair of vertices with probability p. Then almost surely for any $2n$ vertices $y_1, \dots, y_n, z_1, \dots, z_n$, there is another vertex adjacent to $y_1, \dots, y_n$ but not $z_1, \dots, z_n$.

This theory is also complete.

Example 1.5. The theory of dense linear orders consists of the axioms for a linear order (transitivity, antisymmetry, reflexivity, and comparability), along with the statement

$$\forall v \forall w \big((v < w) \Rightarrow \exists x (v < x) \wedge (x < w) \big)$$

where $v < w$ is shorthand for $(v \leq w) \wedge (v \neq w)$, and a statement saying there are at least two distinct elements. The statement $\exists x \forall y (x \leq y)$ says there is a least element, and $\exists x \forall y (y \leq x)$ says there is a greatest element. Adding in both of these statements, both of their negations, or one positive statement and one negation gives four different theories, all of which are complete.

A formula $\psi(v_1, \dots, v_n)$ is consistent with a theory T if there is a model M of T and elements $a_1, \dots, a_n \in M$ such that $M \vDash \psi[a_1, \dots, a_n]$. Given a model M of a theory T, the *definable sets* of M are the solution sets to all formulas in a given language. All the formulas not consistent with T will correspond to the empty set. Some very special theories have few distinct definable sets, and to these theories we can apply the method of *quantifier elimination* to simplify every formula to a Boolean combination of appropriately selected *basic formulas* with no quantifiers.

Theorem 1.6. *The method of quantifier elimination can be applied to the theories of algebraically closed fields and random graphs.*

2. Ultrafilters

Ultrafilters are integral to the construction of Keisler's order. They are used to build *ultrapowers*, which simultaneously amplify structures and even them out in a way that shrinks quirks, and makes size easy to control. Ultrapowers can be endowed with other desirable properties by placing additional constraints on the ultrafilters used to construct them. This section will define ultrafilters and ultrapowers and give some basic results about them. After a discussion of *saturation* in the next section, we will return to discuss the ways in which restrictions on ultrafilters affect ultrapowers.

2.1. **Filters.** Intuitively, a filter is a collection of "big" subsets of a set. We want the whole set to be big, finite intersections of big sets to be big, and sets containing other big sets to be big. Formally, this gives the following definition.

Definition 2.1. A *filter* over a set I is a family $\mathcal{F}$ of subsets of I such that

i) $\varnothing \notin \mathcal{F}$,
ii) if $A \in \mathcal{F}$ and $A \subset B \subset I$, then $B \in \mathcal{F}$, and
iii) if $A \in \mathcal{F}$ and $B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$.

Example 2.2. For any point a in a set I, the collection $\mathcal{F} = \{A : a \in A\}$ is a filter on I. This is the *principal* filter generated by a.

Example 2.3. We can construct a non-principal filter $\mathcal{F}$ on an infinite set I by taking $\mathcal{F}$ to be the collection of all cofinite sets (sets with finite complement). This is called the *Fréchet filter* on I.

In the first example, every single subset of I either contains the point a or doesn't contain a, so every set is either big (if it contains a) or small (if its complement contains a). In the second example, however, if we consider $I = \mathbb{N}$, then the sets of evens and odds are complements, but neither is in the Fréchet filter because neither is cofinite. This means the evens and odds are not accounted for in the Fréchet filter's notion of largeness.

The additional condition required for a filter to be an ultrafilter resolves this issue by requiring an ultrafilter to decide whether each subset is big or small:

Definition 2.4. An *ultrafilter* is a filter $\mathcal{F}$ on I which also satisfies the condition that for all $A \subset I$, either $A \in \mathcal{F}$ or $I \setminus A \in \mathcal{F}$.

The principal filter in example 2.2 is an ultrafilter, but most filters, like the filter in example 2.3, do not have an opinion about every single set. A natural question is whether we can extend a filter to an ultrafilter by adding in each undecided set or its complement in such a way that the collection of sets remains a filter. The answer is yes, and actually a set need only have the property that the intersection of any finitely many of its elements is nonempty to be extended to a filter, and thus to an ultrafilter. This is called the *finite intersection property.* Thus, an ultrafilter can be built on any infinite set by extending its Fréchet filter.

The following quick fact will be of use in a later example.

Proposition 2.5. *Let $\mathcal{D}$ be a non-principal ultrafilter on a set I. Then any cofinite subset of I is $\mathcal{D}$-large.*

Proof. Suppose a finite set $A = \{i_i, \ldots, i_n\}$ is $\mathcal{D}$-large. Since $\mathcal{D}$ is non-principal, there exist sets $X_1, \ldots, X_n$ that exclude $i_1, \ldots, i_n$ respectively. Then $A \cap X_1 \cap \cdots \cap X_n = \varnothing$, and this is a contradiction because the intersection of finitely many large sets should still be in $\mathcal{D}$, but $\varnothing \notin \mathcal{D}$. □

2.2. **Ultraproducts.** One important use of ultrafilters is in constructing *ultraproducts*. The ultraproduct of a set of objects indexed by some set I, modulo some ultrafilter $\mathcal{D}$ on I, has as its underlying set the Cartesian product of the terms in the sequence under the equivalence relation that two elements of this product are identified if the set of indices on which their coordinates agree is in $\mathcal{D}$. As stated formally in definition 2.6 below, additional structure on the ultraproduct arises from the functions, relations, and constants that hold on a large set of the objects in the sequence, where the meaning of "large" is determined by the ultrafilter.

Formally, let $\langle M_i : i \in I \rangle$ be a set of structures and $\mathcal{D}$ an ultrafilter on I. Define a relation $=_{\mathcal{D}}$ for two elements $a = (a[i])$ and $b = (b[i])$ in $\prod_{i \in I} M_i$ by $a =_{\mathcal{D}} b$ if and only if $\{i \in I : a[i] = b[i]\} \in \mathcal{D}$. The relation is trivially reflexive and symmetric, and transitivity comes from the finite intersection and upward closure properties of filters (if $a =_{\mathcal{D}} b$ and $b =_{\mathcal{D}} c$, then $\{i \in I : a[i] = b[i]\} \cap \{i \in I : b[i] = c[i]\} \subset \{i \in I : a[i] = c[i]\}$, so $\{i \in I : a[i] = c[i]\} \in \mathcal{D}$). Therefore, $=_{\mathcal{D}}$ is an equivalence relation. Write $\prod_{i \in I} M_i/\mathcal{D}$ to denote the Cartesian product modulo this equivalence relation, and write $=$ in place of $=_{\mathcal{D}}$ for equality in the ultraproduct. When considering this set, it is often useful to fix a set theoretic section $\prod_{i \in I} M_i/\mathcal{D} \to \prod_{i \in I} M_i$ so that we can unambiguously consider the coordinates $a[i]$ of any element $a \in \prod_{i \in I} M_i/\mathcal{D}$. We can get such a section by picking an element from each $\mathcal{D}$-equivalence class using the axiom of choice.

Definition 2.6. Let $\mathscr{L}$ be a language, $\langle M_i : i \in I \rangle$ a sequence of models of $\mathscr{L}$, and $\mathfrak{D}$ an ultrafilter on I. The *ultraproduct* of this sequence is the model with underlying set $N = \prod_{i \in I} M_i/\mathcal{D}$ and additional structure as follows:

i) If R is a relation symbol of arity n interpreted as R_i in M_i, then the ultraproduct interprets R as R', where
$$(a_1, \ldots, a_n) \in R' \iff \big\{i \in I : (a_1[i], \ldots, a_n[i]) \in R_i\big\} \in \mathcal{D}$$

ii) If f is a function of arity n interpreted as f_i in M_i, then the ultraproduct interprets f as a function f', where
$$f'(a_1, \ldots, a_n) = a_* \in N \iff \big\{i \in I : f_i(a_1[i], \ldots, a_n[i]) = a_*[i]\big\} \in \mathcal{D}$$
In words, $f'(a_1, \ldots, a_n)$ is defined to be the unique $\mathcal{D}$-equivalence class containing $\prod_{i \in I} f_i(a_1[i], \ldots, a_n[i])$.

iii) If c is a constant symbol interpreted as c_i in M_i, then the ultraproduct interprets c as a constant c', where
$$c' = a \iff \big\{i \in I : c_i = a[i]\big\} \in \mathcal{D}$$
Again, in words, c' is defined to be the unique $\mathcal{D}$-equivalence class containing $\prod_{i \in I} c_i$.

If $M_i = M$ for all $i \in I$, we write $M^I/\mathcal{D}$ and call this the *ultrapower* of M modulo $\mathcal{D}$.

This definition states that the functions, relations, and constants in an ultraproduct come from the functions, relations, and constants on large sets of index models, giving the ultraproduct the "average" model structure of the elements in the sequence. In fact, much more structure is transferred from the elements in the sequence to the ultraproduct than is immediately apparent. The following theorem states that if a formula of first order logic is true in M_i for a large set of indices i, then it is true in the ultraproduct, and vice versa.

Theorem 2.7 (Fundamental Theorem of Ultraproducts). *Let $\mathcal{D}$ be an ultrafilter on I, $\langle M_i : i \in I\rangle$ a sequence of models, and $N = \prod_{i\in I} M_i/\mathcal{D}$ their ultraproduct. For any formula of first order logic $\varphi(x_1, \cdots, x_n)$ and any $a_1, \cdots, a_n \in N$, $\varphi[a_1, \cdots, a_n]$ is true in N if and only if $\{i \in I : \varphi[a_1[i], \cdots, a_n[i]$ is true in $M_i\} \in \mathcal{D}$.*

A proof of this theorem can be found in Chang and Keisler's *Model Theory* [1, p. 217-219]. The fundamental theory of ultraproducts is an important indication of the deep connection between ultrafilters and first order logic.

The importance of the ultrapower construction to Keisler's order is its close relationship to saturation. Before giving more details about ultrapowers, we will introduce saturation and discuss why it is a good way to measure the complexity of models.

3. Saturation

In section 1.1, it was mentioned that formulas can be thought of in terms of their solution sets, and the elements or tuples that satisfy multiple formulas simultaneously are exactly the elements in the intersection of the corresponding solution sets. In a given model, as long as there is some element satisfying a set of formulas, its corresponding set of solution sets has the finite intersection property, which means it can be extended to a filter. Thus, a *maximal* consistent set of formulas works like an ultrafilter. There is a name for such sets of formulas.

Definition 3.1. A maximal consistent set of formulas in a fixed language is called a *type*.

We will mostly talk about partial types (another term for sets of formulas that are not necessarily maximal) and types generated by positive and negative instances of a specified formula φ called *φ-types*.

Now, if an ultrafilter is principal, the intersection of all of its sets will be the single value that generates it; if it is non-principal the intersection of all its sets will be empty. This illustrates how a type describes one specific kind of element, or one specific kind of tuple, that may or may not be in a given model, depending on whether its solution sets in the model correspond to a principal or non-principal ultrafilter.

Definition 3.2. When there are elements $a_1, \ldots, a_n$ of a model that satisfy all formulas in a type $p(v_1, \ldots, v_n)$, we say that the type is *realized* in the model by $a_1, \ldots, a_n$.

When a type is realized, the corresponding sets in the ultrafilter correspond to formulas that describe how $a_1, \ldots, a_n$ relate to all other tuples of elements in M. We can also start with a tuple of elements in a model and build a type from every formula that those elements satisfy. This is exactly the way principal ultrafilters are constructed.

The following example describes a situation in which a type is *not* realized.

Example 3.3. Let I be an infinite set and $\mathscr{L} = \{c_i : i \in I\}$ be a language with $|I|$-many constants. Let $p = \{\varphi_i(v) = (v \not\equiv c_i) : i \in I\}$ be a partial type and let M be a set of size $|I|$. Give M a model structure by interpreting each constant c_i as a different element m_i of M. The type p is consistent because any finitely many formulas are satisfied in M by any of the infinitely many non-excluded elements of M, but the type cannot be realized in M.

One way to formalize the idea of "complexity" of models is in terms of definable sets: Models that have more complicated definable sets should be more complex. Simply counting definable sets is not enough to gauge complexity, but a more subtle measure of the differences between definable sets in different models is how often types are realized. If the definable sets in a model are relatively simple, then it should not be too difficult to realize types in the model. However, if the definable sets in a model are complicated, then types will be harder to realize. The notion of the ability to realize types is formalized in the following definition.

Definition 3.4. Let λ be an infinite cardinal. A model M in a language $\mathscr{L}$ is λ-*saturated* if for every subset $A \subset M$ with strictly fewer than λ elements, the expansion $(M, a)_{a \in A}$ realizes every type in the expanded language $\mathscr{L} \cup \{c_a : a \in A\}$.

If a model is saturated in its own cardinality, we say simply that it is *saturated*. Notice that no model M is saturated in any cardinal larger than its own size because a constant corresponding to each element of M can be added to the language, and then the type in example 3.3 is not realized. We will now give several examples of saturated and unsaturated models.

Example 3.5. $(\mathbb{N}, <)$ is not saturated because we can find a type in the expanded language $\{<, 1\}$ that is not realized in $(\mathbb{N}, <, 1)$. One such type $p(v)$ is generated by

$$\begin{aligned}
&\{(1 < v),\\
&\ \exists w_1((1 < w_1) \wedge (w_1 < v)),\\
&\ \exists w_1 \exists w_2((1 < w_1) \wedge (w_1 < w_2) \wedge (w_2 < v)), \dots\}
\end{aligned}$$

Any finitely many of these formulas can be satisfied by a large enough element of $(\mathbb{N}, <, 1)$, so this partial type is consistent, but no element simultaneously satisfies all the formulas.

Example 3.6. $(\mathbb{Q}, <)$ is saturated, in part because the rationals are dense, so the issue in the previous example is not a problem.

Example 3.7. $(\mathbb{R}, <)$ is not saturated. Let $\mathscr{L} = \{<, c_n : n \in \mathbb{N}\}$ be the expanded language where c_n corresponds to $\frac{1}{n}$ in $(\mathbb{R}, <)$, and consider the type $p(v)$ generated by $\{(v < \frac{1}{n} : n \in \mathbb{N}\} \cup \{(v > 0)\}$. This partial type is consistent, but not all of its formulas can be satisfied simultaneously. We can think of this type as describing some infinitesimal larger than zero that is not in $(\mathbb{R}, <)$.

In fact, there are many other types that are not realized in $(\mathbb{R}, <)$, and we can think of these types as pointing out holes where infinitesimals or other kinds of numbers are missing. In some sense, saturation in different cardinalities reveals holes of different sizes.

Example 3.8. $(\mathbb{C}, +, \cdot, 0, 1)$ is saturated. In the next section, a proof will be given that certain ultrapowers of algebraically closed fields are saturated.

Saturation is a good candidate for measuring complexity, but the capacity for models to be λ-saturated for some cardinal λ depends partially on their size. Using saturation alone, it's difficult to meaningfully comparing the complexity of models of different sizes, but taking ultrapowers gives more control. However, to order theories based on saturation of their ultrapowers, we would need ultrapowers of different models of a theory to have the same degree of saturation. Otherwise, the order would not be well defined. However, this is not always the case:

Example 3.9. Let $\mathcal{D}_1$ be a principal ultrafilter on $\mathbb{N}$ and $\mathcal{D}_2$ be an ultrafilter derived from the Frèchet filter on $\mathbb{N}$. Let $M = (\mathbb{Q}, <)$. Then $M^{\mathbb{N}}/\mathcal{D}_1 \cong M$, but $M^{\mathbb{N}}/\mathcal{D}_2$ is strictly larger than M because any order preserving map $M \to M^{\mathbb{N}}/\mathcal{D}_2$ will miss the element $a = (1, 2, 3, \dots) \in M^{\mathbb{N}}/\mathcal{D}_2$.

This example illustrates how ultrapowers depend heavily on the ultrafilters used to construct them. It also raises the question of how to construct a well-defined order using saturation of ultrapowers. Clearly, we must ask for non-principal ultrafilters, but this is not quite enough.

4. Regular Ultrapowers

Definition 4.1. An ultrafilter $\mathcal{D}$ on a set I is *regular* if there is a family of sets $\{A_\alpha : \alpha < |I|\} \subset \mathcal{D}$ such that the intersection of any infinitely many sets in this family is empty. Such a collection of sets is called a *regularizing family.*

We can think of this condition as requiring that the sets be "spread out" enough that any infinitely many of them have empty intersection. Equivalently, an ultrafilter $\mathcal{D}$ is regular if for every ultraproduct $N = \prod_{i \in I} M_i/\mathcal{D}$ and every subset $A \subseteq N$ of size $\leq |I|$, there is a sequence $\langle A_i : i \in I \rangle$ where each $A_i \subset M_i$ is finite and $A \subseteq \prod_{i \in I} A_i/\mathcal{D}$. Another useful equivalent statement is that an ultrafilter $\mathcal{D}$ on a set I is regular if every element $i \in I$ is contained in only finitely many sets in $\mathcal{D}$.

Example 4.2. Let $I = \{i_0, i_1, i_2, \dots\}$ be countable and $\mathcal{D}$ be a non-principal ultrafilter on I. Pick countably many sets $\{U_n : n \in \mathbb{N}\}$ from $\mathcal{D}$. Define $X_n = I \setminus \{i_0, \dots, x_n\}$. Note that $X_n \in \mathcal{D}$ because any cofinite set is $\mathcal{D}$-large by proposition 2.5. Let $A_n = X_n \cap U_n \in \mathcal{D}$. Then the intersection of any infinitely many A_n is empty, so we have found a regularizing family of sets. This means that any ultrafilter on a countable set is either principal or regular.

For regular ultrafilters to be of use, we must know that we can construct them on any index set.

Proposition 4.3. *Regular ultrafilters exist on sets of any infinite size κ.*

For a proof of this fact, see [1, p. 249]. The following result ensures that ordering theories by saturation of their ultrapowers is well-defined when we use regular ultrapowers.

Theorem 4.4. *If $M \equiv N$ and $\mathcal{D}$ is regular, then $M^\lambda/\mathcal{D}$ is saturated if and only if $N^\lambda/\mathcal{D}$ is saturated.*

This was first proved by Keisler in his paper introducing the order [2, p. 29-30]. Another useful property of regular ultrafilters is the control they give over the size of ultrapowers.

Theorem 4.5. *If M is an infinite model and $\mathcal{D}$ a regular ultrafilter, then $|M^I/\mathcal{D}| = |M|^{|I|}$.*

This is proved in [1, p. 250]. If M is finite, then $|M^I/\mathcal{D}| = |M|$ because for any element $a \in M^I/\mathcal{D}$, there is some $m \in M$ such that $\{i : a[i] = m\} \in \mathcal{D}$, since there are only finitely many possibilities for $a[i]$ in each index model. Thus, every element of $M^I/\mathcal{D}$ corresponds to an element of M. This theorem allows us to prove an interesting fact about cardinalities of models.

Corollary 4.6. *If M is an infinite model, then there are arbitrarily large ultrapowers of M.*

With theorem 4.4, we have all we need to be able to define the order Keisler introduced in [2]. We use the notation λ^+ to denote the smallest cardinal number strictly greater than λ.

Definition 4.7. (Keisler's Order) Let T_1 and T_2 be countable complete theories. Then $T_1 \trianglelefteq T_2$ if for every infinite set I, every regular ultrafilter $\mathcal{D}$ on I, and every model $M_1 \vDash T_1$ and model $M_2 \vDash T_2$, we have that $(M_2)^I/\mathcal{D}$ being $|I|^+$-saturated implies $(M_1)^I/\mathcal{D}$ is $|I|^+$-saturated.

This defines a preorder on theories because it is possible to have two theories T_1 and T_2 such that $T_1 \trianglelefteq T_2$ and $T_2 \trianglelefteq T_1$, but $T_1 \neq T_2$, as the following example illustrates.

Example 4.8. The theory of algebraically closed fields is minimal in Keisler's order, which will be shown later, so the theories of algebraically closed fields of characteristic 0 and of characteristic p are also minimal in the order. However, these two theories are not the same and have no models in common.

As is indicated by how little is known about Keisler's order, the definition as it is given can be difficult to work with. To simplify the issue of determining the degree of saturation of regular ultrapowers, we must take a closer look at how saturation works in ultrapowers. One useful simplification is the following:

Theorem 4.9. *For a countable theory T, a regular ultrafilter $\mathcal{D}$ on a set I, and a model $M \vDash T$, $M^I/\mathcal{D}$ is $|I|^+$-saturated if and only if for each formula φ, $M^I/\mathcal{D}$ is $|I|^+$ saturated for φ-types.*

Introduced in [7], this is very useful because types generated by positive and negative instances of a single formula are easier to describe and work with than general types. Another way to simplify the problem is by examining the mechanics of how types are actually realized in regular ultrapowers. This leads to the idea of *distributions*.

4.1. **Distributions.** For a fixed φ-type and ultrapower, a distribution is roughly a way of splitting a type into finite sets of formulas by looking for where (i.e. in which index models) the formulas are satisfied. This allows us to pick an element in each index model satisfying finitely many formulas in such a way that the image of these elements in the ultrapower satisfies the type. A distribution exists for every *small* type (a type over a set of cardinality less than or equal to the index set). The distinguishing characteristic of types that are realized is that their distributions have *multiplicative refinements*.

The definition of a distribution seems technical, but in fact, it directly leads into the way types are realized in regular ultrapowers. The notation $\mathcal{P}_{\aleph_0}(X)$ denotes the set of finite subsets of X.

Definition 4.10. Fix a theory T, a model M of T, an infinite cardinal λ, a regular ultrafilter $\mathcal{D}$ on λ. Define $N = M^\lambda/\mathcal{D}$, choose any section $N \to M^\lambda$, let $A \subset N$ be such that $|A| \leq \lambda$, and let $p(x)$ be a type over A. A distribution $d : \mathcal{P}_{\aleph_0}(p) \to \mathcal{D}$

(1) is monotonic: For all $u \subset v \in \mathcal{P}_{\aleph_0}(p)$, $d(v) \subset d(u)$
(2) is a refinement of the Łoś map: for each $u \in \mathcal{P}_{\aleph_0}(p)$

$$d(u) \subset \left\{ t < \lambda : M[t] \vDash \exists x \left(\bigwedge_{\varphi(x;a) \in u} \varphi(x; a[t]) \right) \right\}$$

(3) has a regularizing set as its image: for each $t < \lambda$, $|\{u : t \in d(u)\}|$ is finite.

Proposition 4.11. *A distribution exists on any small φ-type in a regular ultrapower.*

Proof. To construct a distribution for a φ-type $\{\varphi_i(v; a_i) : i < \lambda\}$ and an ultrapower $M^I/\mathcal{D}$, start with the map $d_0 : p \to \mathcal{D}$ sending each formula to the set of indices t for which

$$M[t] \vDash \exists x(\varphi_i(x; a_i[t]))$$

Enumerate a regularizing set of $\mathcal{D}$ by $\langle A_i : i < \lambda \rangle$, and define $d : p \to \mathcal{D}$ by $d(\{\varphi_i\}) = d_0(\varphi_i) \cap A_i$. This will ensure that the image of the distribution is a regularizing set. Finally, extend d to finite subsets of p by

$$d(\{\varphi_{i_1},\dots,\varphi_{i_n}\}) = \bigcap_{k\leq n} d(\{\varphi_{i_k}\}) \cap \left\{ t : M[t] \vDash \exists x \bigwedge_{k\leq n} \varphi_{i_k}(x; a_{i_k}[t]) \right\}$$

The first part of this function definition, $\bigcap_{k\leq n} d(\{\varphi_{i_k}\})$, ensures that the distribution is monotonic, and this intersection is $\mathcal{D}$-large by the finite intersection property. The second part of the function definition, $\left\{t : M[t] \vDash \exists x \bigwedge_{k\leq n} \varphi_{i_k}(x; a_{i_k}[t])\right\}$, ensures that the distribution satisfies condition (2). This second set is large by theorem 2.7, so each value of d is a $\mathcal{D}$-large set, as desired. □

The proof of the following proposition shows a distribution in action.

Proposition 4.12. *Ultrapowers of algebraically closed fields are always saturated.*

Proof. Let M be a model of the theory of algebraically closed fields, λ an infinite cardinal, and $\mathcal{D}$ a regular ultrafilter on λ. Let $N = M^\lambda/\mathcal{D}$.

Let $A \subset N$ be such that $|A| \leq \lambda$, and let $p(x)$ be a type over A. Because we have quantifier elimination, $p(x)$ is of the form $\{f_i(x;\bar{a}_i) : i < \lambda\}$ where each f_i is a finite conjunction of polynomial equations and $\neg$-polynomial equations with coefficients in A. If the type is realized in M, then it will be realized in N by theorem 2.7, so suppose p describes an element that doesn't satisfy any nontrivial polynomials with coefficients in $\bar{a}_i$, a finite subset of A. This means $p(x)$ is of the form $\{\neg g_i(x,\bar{a}_i) : i < \lambda\}$, where each g_i is a finite conjunction of polynomial equations with coefficients in $\bar{a}_i$.

Now, let d be a distribution. Then for each t, d sends finitely many $\neg g_i$ to sets containing t, so there are only finitely many formulas associated with each index model $M[t]$. In each $M[t]$, there is some element $b[t]$ that satisfies the finitely many $\neg g_i(x;\bar{a}_i)$, so we can define $b = (\prod_{t<\lambda} b[t])/\mathcal{D}$. b will always satisfy p because it satisfies each $\neg g_i$ on a large set, by the definition of a distribution. Thus, any type over a set of size less than or equal to λ is realized in N, and this means that N is saturated. □

The above proof fails to show the essential barrier to realization of types: If $\{\varphi_{i_1}\},\dots,\{\varphi_{i_n}\}$ are elements of $\mathcal{P}_{\aleph_0}(p)$ whose images under d all contain some index t, then for each $k = 1,\dots,n$, we have $M[t] \vDash \exists x(\phi_{i_k}(x; a_{i_k}[t]))$. However, it is not necessarily true that any element of $M[t]$ satisfies more than one of these formulas simultaneously. Only if $t \in d(\{\varphi_{j_i},\dots,\varphi_{j_l}\})$ can $\varphi_{j_i},\dots,\varphi_{j_l}$ be satisfied simultaneously in $M[t]$. In general, if $u,v \in \mathcal{P}_{\aleph_0}(p)$, then $d(u\cup v) \subseteq d(u)\cap d(v)$ by monotonicity, but we will show that it is only when we have equality for some distribution that p is realized.

Before proving this fact, we introduce a couple of new definitions.

Definition 4.13. A function $f : \mathcal{P}_{\aleph_0}(I) \to \mathcal{D}$ is *multiplicative* if $d(u\cup v) = d(u)\cap d(v)$ for all finite subsets u and v of I.

Having the image of a distribution be a regularizing set ensures that only finitely many formulas are mapped to each index model. Adding the multiplicativity condition ensures that each finite subset of a type projects to a set of formulas that can be realized in each index model. These two conditions ensure that given a type on a subset of an ultrapower, we can first associate each formula with a large set in such a way that there are only finitely many formulas associated to any particular set by regularity; secondly, we can pick an element in each index model that simultaneously satisfies the finitely many formulas associated to it, by multiplicativity; and finally, we take the element in the ultrapower given by the sequence of elements we've picked from each index model, and this satisfies our type.

Definition 4.14. A distribution $d : \mathcal{P}_{\aleph_0}(p) \to \mathcal{D}$ is *accurate* if for each $t \in I$ and each finite subset $\{\varphi_{i_1}, \dots, \varphi_{i_n}\} \subset \{\varphi_j : t \in d(\{\varphi_j\})\}$, we have $t \in d(\{\varphi_{i_1}, \dots, \varphi_{i_n}\})$ if and only if all the formulas are satisfied simultaneously in $M[t]$.

Notice that the distribution constructed in the proof of proposition 4.11 is accurate. The main importance of accurate distributions is that they are easier to work with than arbitrary distributions, and that they are used in an intermediate step of the proof that a type is realized if and only if some distribution on it has a multiplicative refinement.

Theorem 4.15. *For a theory T, a model $M \vDash T$, a set I, and an ultrafilter $\mathcal{D}$ on I, let $N = M^I/\mathcal{D}$, $A \subset N$ be small, and p be a type over A. Then the following are equivalent:*

(1) Some distribution of p has a multiplicative refinement.
(2) Every accurate distribution of p has a multiplicative refinement.
(3) p is realized in N.

Proof. First, suppose some distribution d on $p(x) = \{\varphi_i(x; a_i) : i \in I\}$ has a multiplicative refinement d'. We will build an element of N that realizes p. There are only finitely many formulas $\varphi_1, \dots, \varphi_n$ assigned to each index model $M[t]$ by regularity, and we can find an element $b[t]$ that satisfies all of them simultaneously because if $t \in d'(\{\varphi_i\})$ for $i = 1, \dots, n$, then $t \in d'(\{\varphi_1, \dots, \varphi_n\})$ by multiplicativity. Define $b = \prod_{t\in I} b[t]/\mathcal{D}$. Then $\varphi_i(b[t], a_i[t])$ is true in $M[t]$ for a $\mathcal{D}$-large set of t, so by Łòs's theorem, $N \vDash \varphi_i(b, a_i)$. Since this is true for every formula φ_i, N realizes p at b.

Now suppose p is realized by some element $b \in N$, and d is an accurate distribution. Define a function $d' : \mathcal{P}_{\aleph_0}(p) \to \mathcal{D}$ by

$$d'(\{\varphi_{i_1}, \dots, \varphi_{i_n}\}) = \left\{ t : M[t] \vDash \bigwedge_{k\leq n} \varphi_{i_n}\big(b[t]; a_i[t]\big) \right\} \cap d(\{\varphi_{i_1}, \dots, \varphi_{i_n}\})$$

The first set in the equation above is large because p is realized. Also, d' is monotonic, refines the Łòs map, and has a regularizing set as its image, so d' is a distribution and a refinement of d. Now, to show that d' is multiplicative, $d'(u \cup v) \subseteq d'(u) \cap d'(v)$ comes from monotonicity, so we need that $d'(u \cup v) \supseteq d'(u) \cap d'(v)$. If $s \in d'(u) \cap d'(v)$, then $b[s]$ satisfies the formulas in u simultaneously and the formulas in v simultaneously, so it must satisfy the formulas in $u \cup v$ simultaneously. Note that $s \in d(u) \cap d(v)$ because d' is a refinement. Since d is accurate and $M[s] \vDash \bigwedge_{\varphi_i \in u\cup v} \varphi_i\big(b[s]; a_i[s]\big)$, we have that $s \in d(u \cup v)$. Thus,

$$s \in \left\{ t : M[t] \vDash \bigwedge_{\varphi_i\in u\cup v} \varphi_i\big(b[t]; a_i[t]\big) \right\} \cap d(u \cup v) = d'(u \cup v)$$

so d' is multiplicative, and we've shown (3) ⇒ (2).

Since any small type in a regular ultrapower has an accurate distribution by proposition 4.11, (2) ⇒ (1). □

We have shown that types are realized in a fixed ultrapower only when some distribution has a multiplicative refinement. Naturally, the next step is to explore how different ultrapowers allow or prevent the existence of multiplicative refinements. Specifically, what effects do different ultrafilters have on distributions? One way to explore this question is to come up with a useful criterion and then define a class of ultrafilters as the ones that satisfy that criterion.

Definition 4.16. If every monotonic function $f : \mathcal{P}_{\aleph_0}(\lambda) \to \mathcal{D}$ has a multiplicative refinement, then $\mathcal{D}$ is λ^+-good.

Theorem 4.17. *λ^+-good ultrafilters exist on sets of size λ.*

For a proof, see [1, p. 391-392]. Good ultrafilters were first proved to exist by Keisler assuming the generalized continuum hypothesis, and Kunen later gave a proof of their existence in full generality in [3].

Notice that if every monotonic function has a multiplicative refinement, then in particular, every distribution has a multiplicative refinement. This means that for any model M, when an $|I|^+$-good ultrafilter $\mathcal{D}$ is used to build the ultrapower $M^I/\mathcal{D}$, every small type in this ultrapower is realized. Thus, any theory whose ultrapowers are *only* saturated by good ultrafilters would be maximal in Keisler's order if such a theory were to exist. There is, in fact, such a theory.

Proposition 4.18. *Let M be the model whose elements are finite subsets of $\mathbb{N}$ in the language $\{\subset\}$, and let λ be an infinite cardinal. If $\mathcal{D}$ is not a λ^+-good ultrafilter, then $N = M^\lambda/\mathcal{D}$ is not saturated.*

Proof. Let T be the theory of M. Define $\varphi(x; y) = x \subset y$. We will show there is a small φ-type omitted in N by showing that there is an accurate distribution with no multiplicative refinement.

Because $\mathcal{D}$ is not λ^+-good, there is some monotonic function $f : \mathcal{P}_{\aleph_0}(\lambda) \to \mathcal{D}$ with no multiplicative refinement. For each $t < \lambda$, consider the set of singletons $\{i\}$ such that f maps $\{i\}$ to a set containing t. We would like to find $\{a_i : i < \lambda\} \subset N$ such that f is an accurate distribution of a consistent partial type $\{x \subseteq a_i : i < \lambda\}$. This means we want there to be an incidence of finitely many $a_{i_j}[t]$ exactly when f maps the indices i_j to t. However, there are problems: f is not a distribution, so the set of $\{i\}$ such that f maps $\{i\}$ to a set contain t may be infinite, which makes it harder to pick $a_{i_j}[t]$ that make the formulas $x \subseteq a_{i_j}$ consistent in $M[t]$ when they should be (according to f).

To fix the problem of having to choose each $a_i[t]$ so that it satisfies possibly infinitely many incidence relations, we can replace M by a larger model $M' \equiv M$ that is saturated in a sufficiently high cardinality. Types are realized in ultrapowers of M' iff they are realized in ultrapowers of M by theorem 4.4, so switching to M' won't affect the realization of p. If all λ-small types are realized in M', then for the $\{i\}$ that f maps to sets containing t, we can pick $a_i[t]$ so that

$$M'[t] \models \exists x \left(\bigcap_{j \le k} x \subset a_{i_j}[t] \right) \iff t \in f(\{i_i, \dots, i_k\})$$

i.e. we're creating an incidence in $M'[t]$ of finitely many formulas exactly when f maps the indices of the formulas to t. If f does not map $\{j\}$ to a set containing t, set $a_j[t] = \varnothing$. Let $a_i = \prod a_i[t]/\mathcal{D}$.

Now we would like to build a distribution f' from f using the sets $\{a_i : i < \lambda\}$. Let

$$p = \{x \subset a_i : i < \lambda\}$$

and let $\varphi_i(x) = x \subset a_i$. Consider $f : \mathcal{P}_{\aleph_0}(\lambda) \to \mathcal{D}$ as mapping from $\mathcal{P}_{\aleph_0}(p)$ to $\mathcal{D}$. Let $\{X_i : i < \lambda\}$ be a regularizing set in $\mathcal{D}$. Define f' by

$$f'(\{\varphi_i\}) = f(\{\varphi_i\}) \cap X_i$$

Extend this definition to $\mathcal{P}_{\aleph_0}(p)$ by

$$f'(\{\varphi_{i_i},\ldots,\varphi_{i_k}\}) = \left\{t : \bigwedge_{k\le n} t\in f'(\{\varphi_{i_k}\})\right\} \cap \left\{t : M[t]\vDash \exists x \bigwedge_{k\le n} \varphi_{i_k}(x;a_{i_k}[t])\right\} \cap f(\{\varphi_{i_i},\ldots,\varphi_{i_k}\})$$

This is a refinement of f by construction, so it is not multiplicative by hypothesis. It is also an accurate distribution by construction, so we have found an accurate distribution with no multiplicative refinement, and by the previous theorem, this means that p is not realized. □

The above proof is an expanded version of a proof in [6, p. 28-29]. By this result, the ultrapowers of models of T are only saturated by good ultrafilters, so T is maximal in Keisler's order. With the framework that we currently have, we cannot prove that the theory of random graphs is in the third equivalence class, but it is still worth discussing as an intermediate example where some types are realized and some are not.

Example 4.19. (Random Graphs) Let RG be the theory of random graphs, M a model of RG, I an infinite set, and $\mathcal{D}$ a regular ultrafilter on I. RG has quantifier elimination, so every formula can be reduced to boolean combinations of formulas of the form $xRa\wedge\neg xRb$. Let

$$p(x) = \{\varphi_i(x;a_i;b_i) : i\in I\}$$

be a small type where $\varphi_i(x;a_i;b_i) = xRa_i\wedge\neg xRb_i$, and let d be a distribution on p with multiplicative refinement d'. Suppose $\varphi_{i_1},\ldots,\varphi_{i_n}$ are the formulas d' maps to sets containing some index t. Then these formulas are simultaneously satisfied in $M[t]$ by multiplicativity, i.e.

$$M[t]\vDash\exists x\left(\bigwedge_{i\le n} xRa_{i_1}\wedge\neg xRb_{i_n}\right)$$

However, this cannot happen if $a_{i_j}[t] = b_{i_k}[t]$ for some j and k. Let u_t be the finite set of all formulas d maps to sets containing t. Define $A[t] = \{a_i[t] : \varphi_i(x;a_i;b_i)\in u_t\}$ and $B[t] = \{b_i[t] : \varphi_i(x;a_i;b_i)\in u_t\}$. Then a necessary condition for d to have a multiplicative refinement is that $\{t : A[t]\cap B[t]=\varnothing\}\in\mathcal{D}$. There are indeed ultrafilters that do not satisfy this property, so some types in models of RG are not realized.

This example is an expanded version of a discussion that can be found in [6, p. 24-25]. For more information about the ways in which ultrafilters can affect the saturation of ultrapowers, see [6].

5. Keisler's Order

This final section will sum up the information about Keisler's order described in this paper. Recall, $T_1\trianglelefteq T_2$ in Keisler's order if for every infinite set I, every regular ultrafilter $\mathcal{D}$ on I, and every model $M_1\vDash T_1$ and model $M_2\vDash T_2$, we have that $(M_2)^I/\mathcal{D}$ is $|I|^+$-saturated implies $(M_1)^I/\mathcal{D}$ is $|I|^+$-saturated.

Proposition 5.1. *The theory of algebraically closed fields is minimal in Keisler's order.*

Proof. Let $M_1\vDash$ ACF, let T be a complete countable theory, let $M_2\vDash T$, and let $\mathcal{D}$ be a regular ultrafilter on a set I for which $(M_2)^I/\mathcal{D}$ is $|I|^+$-saturated. By proposition 4.12, $(M_1)^I/\mathcal{D}$ is also $|I|^+$-saturated. Since the ultrafilter $\mathcal{D}$ saturating M_2, the set I, and the models M_1 and M_2 were arbitrary, ACF $\trianglelefteq T$. □

Proposition 5.2. *Keisler's order has a maximum equivalence class of theories. A necessary and sufficient condition for a theory T to be in this class is that for any set I and model $M \models T$, $M^I/\mathcal{D}$ is $|I|^+$-saturated if and only if $\mathcal{D}$ is $|I|^+$-good.*

Proof. Let T be a theory satisfying this condition, M any model of T, T' any countable complete theory, $M' \models T'$, I an infinite set, and $\mathcal{D}$ a $|I|^+$-good ultrafilter on I. Then $(T')^I/\mathcal{D}$ is $|I|^+$-saturated, so $T' \trianglelefteq T$. This means the condition is sufficient. Since proposition 4.18 shows that there does exist a theory whose ultrapowers are saturated only by good ultrafilters, the condition is also necessary. Thus, there is maximum class of theories in Keisler's order. □

In the paper introducing the order, Keisler proved that there were minimum and maximum classes of theories, but he did not characterize them. The maximum class of theories still has not been characterized. Shelah later characterized the first two classes of the order and proved that the maximum class contained the theory of dense linear orders in [9]. More recently, Malliaris showed that the theory of random graphs is in the minimum class of unstable theories. Additional results, including the facts that the order is infinite and not a well order, can be found in [8].

Acknowledgments. It is a pleasure to thank my mentors, Yun Cheng and Sean Howe, for all their help. Special thanks are also due to Maryanthe Malliaris, without whom this project would not have been possible. Her guidance, patience, and encouragement have all been very much appreciated. I would also like to thank Peter May for three interesting, enjoyable, and math-filled summers in the UChicago REU.

References

[1] C. C. Chang and H. Jerome Keisler. Model Theory. Amsterdam and New York: North-Holland Press, 1990.
[2] H. Jerome Keisler. Ultraproducts which are not saturated. The Journal of Symbolic Logic. Vol. 32, No. 1 (Mar., 1967), pp. 23-46.
[3] Kenneth Kunen. Ultrafilters and independent sets. Transactions of the American Mathematical Society. Vol. 172, No. 1 (1972), pp. 299-306.
[4] Maryanthe Malliaris. The Clique Covering Problem and Other Questions.
[5] Maryanthe Malliaris. Hypergraph sequences as a tool for saturation of ultrapowers. Journal of Symbolic Logic. Vol. 157, No. 1, (2012) pp. 195-223.
[6] Maryanthe Malliaris. Persistence and Regularity in Unstable Model Theory. 2009.
[7] Maryanthe Malliaris. Realization of phi-types and Keisler's order. Annals of Pure and Applied Logic. Vol. 77, No. 2-3, (2009) pp. 220-224.
[8] Maryanthe Malliaris and Saharon Shelah. Keisler's order has infinitely many classes. Arxiv. 28 Mar 2015.
[9] Saharon Shelah. Classification Theory: and the Number of Non-Isomorphic Models. Amsterdam and New York: North-Holland Press, 1978.