USING ULTRAPOWERS TO CHARACTERIZE ELEMENTARY EQUIVALENCE

MIKAYLA KELLEY

Abstract. This paper will establish that ultrapowers can be used to determine whether or not two models have the same theory. More precisely, assuming the generalized continuum hypothesis and using a special type of ultrafilter, we will show that two models are elementarily equivalent if and only if they have isomorphic ultrapowers. A more general statement holds without assuming GCH, but the proof is beyond the scope of this paper. While presenting the equivalence, we will explore basic model theory, ultrafilters, the ultraproduct construction, types, and saturation.

Contents

1. Introduction

Model theory is the study of formal languages (syntax) and meaningful interpretations of such a language (semantics). The interplay between the syntactic and semantic perspective is a fundamental aspect of the subject. For example, the theory of an ordered field is a formally defined set of properties while $\mathbb{R}$ is a rich and meaningful example of the theory. By building "examples", or *models* for formal *theories*, we come to understand the theory better, and by studying the formal theory of a model, we come to understand the model better as well.

The relationship between syntax and semantics is highlighted in this paper. An *ultraproduct* is a model constructed from a collection of existing models, and an *ultrapower* is an ultraproduct constructed from some number of copies of the same

Date: November 11, 2015.

model. An *ultrafilter* is a tool used in the construction of an ultraproduct. Informally speaking, the main theorem presented is that under certain conditions and the assumption of GCH, two models have the same theory if and only if their ultrapowers constructed with a *countably incomplete good ultrafilter* are isomorphic. In some sense, the theorem states that a syntactic notion of sameness is equivalent to a semantic notion of sameness. It is the goal of this paper to give a relatively complete account of the statement and proof of this equivalence. Note that proving two models are elementarily equivalent given that their ultrapowers are isomorphic is straightforward, while proving the reverse implication requires much more machinery.

The main theorem will be presented in the following steps:

(1) Give a foundational account of both the semantic and syntactic perspective through an introduction to basic model theory. (Section 2)
(2) Describe what ultrafilters and ultrapowers are, and how to construct them. (Sections 3-4)
(3) Describe what a *saturated* model is, and why two saturated models of a given size are isomorphic if and only if they have the same theory. (Section 5)
(4) Show how a countably incomplete good ultrafilter gives rise to a saturated ultrapower. (Section 6)
(5) Use (1)-(4) to prove the main theorem. (Section 6)

2. Background Material

As stated in the introduction, model theory concerns itself both with syntax and semantics, as well as how they relate. The following table depicts analogous ideas that will be introduced in this section. For each pair, the semantic term is the meaningful "translation" of the purely formal, syntactic notion.

Syntax	*Semantics*
formal language	"real" language (e.g. English)
formula	property
sentence	T/F statement
theory	model
proof ($\vdash$)	satisfaction ($\models$)
elementary equivalence	isomorphism

I will not be rigorous at parts, as the foundational theory is not the focus of this paper. See [1] for a detailed introduction to model theory.

2.1. **Languages.**

Definition 2.1. A first order predicate *language* $\mathscr{L}$ is a set of n-placed relation symbols, m-placed function symbols, and constant symbols for $n, m \geq 1$. As in [1], we use capital Latin letters P and F with subscripts to denote relations and functions respectively. We use the letter c with subscripts to denote constants. Let I_P, I_F, and I_c be indexing sets for the relations, functions, and constants respectively. Then we write

$$\mathscr{L} = \{P_i\}_{i \in I_R} \cup \{F_i\}_{i \in I_F} \cup \{c_i\}_{i \in I_c}.$$

Often times we refer to a language by specifying its *power*, denoted $||\mathscr{L}||$ and defined as

$$||\mathscr{L}|| = |\mathscr{L}| + \aleph_0.$$

We say a language is either countably infinite or uncountable, depending on whether $||\mathscr{L}||$ is countable or uncountable.

We use a formal language to formally "say" things. I will not provide a detailed definition of a well formed formula. See [1] for complete definitions. Intuitively, a *formula* describes a potential property of elements in a model while a *sentence* is a true or false statement. More precisely:

Definition 2.2. A *formula* of $\mathscr{L}$ is a finite string of symbols using the symbols of $\mathscr{L}$, logical connectives such as $\{\vee, \wedge, \forall, \neg\}$, and variables with formal restrictions on how strings can be formed. A *sentence* of $\mathscr{L}$ is a formula with no free variables. Coarsely speaking, a *free variable* is not accounted for by a quantifier such as $\forall$ or $\exists$. For example, $\exists x \exists y[x = y]$ is a sentence while $\exists x[x = y]$ is a formula.

Definition 2.3. We say ϕ is a *formula in the (free) variables* $x_1, \ldots, x_n$ if the free variables in ϕ are amongst $\{x_1, \ldots, x_n\}$.

2.2. **Models.**

Definition 2.4. An $\mathscr{L}$*-model* $\mathfrak{A}$ is a non-empty set equipped with an interpretation of every symbol in $\mathscr{L}$. The model consists of the following:

(1) A non-empty set A, called the *universe* of $\mathfrak{A}$.
(2) For each $n \in \mathbb{N}$ and n-placed relation symbol $P_j \in \{P_i\}_{i \in I_P}$, a subset $R_j \subseteq A^n$.
(3) For each $m \in \mathbb{N}$ and m-placed function symbol $F_j \in \{F_i\}_{i \in I_F}$, a function $G_j : A^m \to A$.
(4) An element $a_j \in A$ for each constant $c_j \in \{c_i\}_{i \in I_c}$.

Notation 2.5. We will always denote models using Gothic letters, such as $\mathfrak{A}$ or $\mathfrak{B}$, and their corresponding universes with the corresponding capital letters, such as A or B.

Example 2.6. Let $\mathscr{L}$ be the language with a 2-placed relation symbol $\leq$, 2-placed function symbol $+$, 2-placed function symbol $\times$, and constant symbols $\{0, 1\}$. Then we have the following $\mathscr{L}$-model

$$\mathfrak{R} = \langle \mathbb{R}, \leq_{\mathbb{R}}, +_{\mathbb{R}}, \times_{\mathbb{R}}, 0_{\mathbb{R}}, 1_{\mathbb{R}} \rangle,$$

where the interpretations $\leq_{\mathbb{R}}, +_{\mathbb{R}}$, $\times_{\mathbb{R}}$, $0_{\mathbb{R}}$ and $1_{\mathbb{R}}$ are the usual order, addition operation, multiplication operation, additive unit element, and multiplicative unit element in the real numbers. For example, $(2, 3) \in \leq_{\mathbb{R}}$ and $+_{\mathbb{R}}(2, 3) = 5$.

It is often useful to expand a language and model to include a constant symbol for every element in a subset of the model's universe.

Definition 2.7. If $\mathfrak{A}$ is an $\mathscr{L}$-model and $X \subseteq A$, we may construct a *simple expansion* of $\mathscr{L}$ by adjoining a new constant symbol for each $a \in X$

$$\mathscr{L}_X = \mathscr{L} \cup \{c_a : a \in X\}.$$

We can regard $\mathfrak{A}$ as an $\mathscr{L}_X$ model

$$\mathfrak{A}_X = (\mathfrak{A}, a)_{a \in X}$$

by interpreting each c_a as $a \in X$ and keeping all other interpretations of symbols in $\mathscr{L}$ the same.

There is a natural notion of isomorphism for $\mathscr{L}$-models:

Definition 2.8. Two models $\mathfrak{A}$ and $\mathfrak{A}'$ are *isomorphic* if and only if there is a bijective function f mapping A to A' such that:

(1) For each n-placed relation symbol P, with interpretations R and R' in $\mathfrak{A}$ and $\mathfrak{A}'$ respectively,
$$R(x_1, \dots, x_n) \iff R'(f(x_1), \dots, f(x_n))$$
for any $x_1, \dots, x_n \in A$.

(2) For each m-placed function symbol F, with interpretations G and G' in $\mathfrak{A}$ and $\mathfrak{A}'$ respectively,
$$f(G(x_1, \dots, x_m)) = G'(f(x_1), \dots, f(x_m))$$
for any $x_1, \dots, x_m \in A$.

(3) For each constant symbol c, with interpretations a and a' in $\mathfrak{A}$ and $\mathfrak{A}'$ respectively,
$$f(a) = a'.$$

In short, $\mathfrak{A}$ and $\mathfrak{A}'$ are isomorphic if there is a bijective mapping f such that, for all $x \in A$, $f(x)$ behaves the same way x does with respect to relations, functions, and constants.

2.3. **Satisfaction.** Models are possible worlds realizing $\mathscr{L}$ and we can ask what it is like in such worlds. Specifically, we can ask whether elements in a model have the property described by a formula and whether a sentence is true or false in a model. We define the truth of a formula $\phi(v_0, \dots, v_p)$, for any set of values $x_0, \dots, x_q \in A$ where $q \geq p$, inductively on the complexity of the formula. We define ϕ to be *true* for $x_0, \dots, x_q$ if and only $(x_0, \dots, x_p)$ has the property described by ϕ. Along the way, a sentence σ is defined to be true in a model exactly when the "English translation" of σ is a true description of the model. The truth of a sentence is independent of which $x_0, \dots, x_q \in A$ are considered because a sentence requires no "input" values.

We write
$$\mathfrak{A} \models \phi[x_0, \cdots, x_q]$$
when $\phi(v_0, \dots, v_p)$ is true for $x_0, \dots, x_q \in A$. This notion of *satisfaction* indicates that a specific tuple of elements in $\mathfrak{A}$ has the property ϕ. We say any of the following:
$$x_0, \dots, x_q \text{ satisfies } \phi \text{ in } \mathfrak{A};$$
$$\mathfrak{A} \text{ is a model of } \phi[x_0, \dots, x_p];$$
$$\mathfrak{A} \text{ satisfies } \phi[x_0, \dots, x_p].$$
For the sentence σ, we write
$$\mathfrak{A} \models \sigma$$
to indicate that $\mathfrak{A}$ thinks σ is a true sentence. We say either of the following:
$$\mathfrak{A} \text{ is a model of } \sigma;$$
$$\mathfrak{A} \text{ satisfies } \sigma.$$

Thus, satisfaction in a model can indicate which elements of the model have a certain property defined by a formula, or the truth of a sentence. This is as far as we go into the definition of satisfaction. See [1] for a complete construction.

2.4. **Theories and Consistency.** Until now, we have been mostly focusing on the semantic side of model theory. We now recall the syntactic perspective.

Definition 2.9. A *theory* T in $\mathscr{L}$ is a set of sentences of $\mathscr{L}$. We say that $\mathfrak{A}$ is a *model of* T if $\mathfrak{A} \models \sigma$ for each $\sigma \in T$.

For example, the usual axioms for a field form a theory, and models of this theory are exactly the fields used across mathematics.

Definition 2.10. Let $\mathfrak{A}$ be a model. The *theory of* $\mathfrak{A}$ is the set $Th(\mathfrak{A})$ of all sentences that are true in $\mathfrak{A}$ or, by definition:

$$Th(\mathfrak{A}) = \{\sigma | \sigma \text{ is an } \mathscr{L} \text{ sentence and } \mathfrak{A} \models \sigma\}.$$

Definition 2.11. We say that two $\mathscr{L}$-models $\mathfrak{A}$ and $\mathfrak{B}$ are *elementarily equivalent* if and only if $Th(\mathfrak{A}) = Th(\mathfrak{B})$.

Remark 2.12. It is useful to note that elementary equivalence is strictly weaker than isomorphism. See Example 5.8 for an example of a model that is elementarily equivalent but not isomorphic to $\langle \mathbb{R}, \leq_{\mathbb{R}} \rangle$.

Definition 2.13. If Σ is a set of sentences of $\mathscr{L}$, then a *proof from* Σ is a finite, non-empty sequence of sentences of $\mathscr{L}$, $\phi_0, \dots, \phi_n$, such that, for each i, either $\phi_i \in \Sigma$, ϕ_i is a logical axiom, or ϕ_i follows from a deductive law using any sentences ϕ_j where $j < i$. [2]

Definition 2.14. For a set of sentences Σ and a sentence ϕ, we say that ϕ *is deducible from* Σ, or Σ *proves* ϕ, if there is a proof of ϕ from Σ. We write

$$\Sigma \vdash \phi.$$

Definition 2.15. Let Σ be a set of sentences. Then Σ is *consistent* if it does not prove a contradiction. That is, there is no sentence ϕ such that

$$\Sigma \vdash \phi \wedge \neg\phi.$$

Since proofs are finite, we have the following:

Fact 2.16. Let Σ be a set of sentences. Then Σ is consistent if and only if every finite subset of Σ is consistent.

Definition 2.17. A theory is *complete* if and only if it is *maximally consistent*, or, by definition, is not properly contained in any other consistent theory.

A useful fact about completeness:

Fact 2.18. Let Σ be a theory. The following are equivalent:

(1) Σ is complete.
(2) For every sentence ϕ, exactly one of the following holds:

$$\phi \in \Sigma \text{ or } \neg\phi \in \Sigma.$$

It follows that the theory of a model is complete because, for any sentence, the model says the sentence is either true or false. This is the case if and only if the sentence or its negation is contained in the model's theory.

To conclude, I mention two pivotal theorems of model theory.

Theorem 2.19. *(Extended Completeness Theorem) Let Σ be a set of sentences of $\mathscr{L}$. Then Σ is consistent if and only if Σ has model.*

If Σ is modeled, it is easy to show that Σ cannot prove a contradiction as this will lead to the model satisfying the truth and falsity of some sentence. For the reverse direction, one takes the consistent theory and builds a model for it. See [1] for a full proof.

The Completeness Theorem is a good example of how, in some sense, a theory Σ looks the same both from the syntactic and semantic perspective. A theory is modeled if and only if it is "syntactically sound".

Theorem 2.20. *(Compactness Theorem) A set of sentences Σ in $\mathscr{L}$ has an $\mathscr{L}$-model if and only if every finite subset of Σ has an $\mathscr{L}$-model.*

The theorem follows from the following chain of if and only if statements given by Fact 2.16 and Theorem 2.19: A set of sentences Σ has a model $\iff$ Σ is consistent $\iff$ every finite subset of Σ is consistent $\iff$ every finite subset of Σ has a model. See Theorem 4.8 for an alternative proof.

3. Filters and Ultrafilters

In this section, we introduce and describe filters and ultrafilters, as well as prove the existence of such objects. A filter on a set I makes rigorous a notion of "almost all" or a "majority" of I, which satisfies upward and finite intersection closure properties. An ultrafilter is a maximal filter. Recall that an ultraproduct is constructed from a collection of existing models. We use an ultrafilter in the construction to clearly define a "majority" of the collection. Defining a majority of the collection will, in turn, be used to determine the theory of the ultraproduct.

Definition 3.1. Let I be a non-empty set called the *index set.* Let $\mathcal{P}(I)$ be the set of all subsets of I. Then we say a set of subsets $D \subseteq \mathcal{P}(I)$ is a *filter over I* if and only if:

(1) $I \in D$. [1]
(2) (Closure Under Finite Intersection) If $X, Y \in D$, then $X \cap Y \in D$.
(3) (Upward Closure) If $X \in D$ and $X \subseteq Y \subseteq I$, then $Y \in D$.

Example 3.2. One can check that $\mathcal{P}(I)$ satisfies the properties of a filter. A filter F is *proper* if and only if $F \neq \mathcal{P}(I)$ which, by upward closure, is equivalent to saying $\emptyset \notin F$. We will only be interested in proper filters.

Example 3.3. For any $X \subseteq I$, the filter $D = \{Y \subseteq I : X \subseteq Y\}$ is the *principal filter generated by X.*

Example 3.4. For any subset $E \subseteq \mathcal{P}(I)$, the intersection of all filters containing E is a filter and is the *filter generated by E.* Since $\mathcal{P}(I)$ is a filter, this is well-defined.

Since only proper filters are interesting to us, we'd like to know when the filter generated by E is proper.

[1] This entails that filters are always non-empty.

Definition 3.5. A set E has the *finite intersection property* if and only if any finite and non-empty set of elements of E has a non-empty intersection.

Proposition 3.6. *Let $E \subseteq \mathcal{P}(I)$ and consider D, the filter generated by E. Then D is a proper filter if and only if E has the finite intersection property.*

Proof. First, consider D' the set of $Y \in \mathcal{P}(I)$ such that

$$e_1 \cap e_2 \cap \ldots \cap e_n \subseteq Y$$

for some finite number of elements $e_1, \ldots, e_n \in E$. I claim $D' = D$. To show $D' \subseteq D$, it suffices to show that $D' \subseteq F$ for any filter F containing E. If F is an arbitrary filter containing E, then F contains every finite intersection of elements in E by closure under finite intersection. Therefore, by upward closure, F also contains any subset of I that contains a finite intersection of elements in E. Thus, every element $Y \in D'$ lives in F, which proves $D' \subseteq D$. To show $D \subseteq D'$, we can show that D' is a filter over I such that $E \subseteq D'$ so that the intersection of all such filters (i.e. D) is contained in D'. This is a routine check and so we conclude that $D = D'$.

Lastly, to prove the proposition, assume E generates a proper filter D. Then D cannot contain the empty set, but this entails that finite intersections of elements in E must be non-empty, since D is closed under finite intersection. Thus E has the finite intersection property. If E has the finite intersection property, then $\emptyset$ cannot contain the intersection of any finite set elements of E and is therefore not in D. Thus $D \neq \mathcal{P}(I)$ and is a proper filter. □

Now, we introduce the notion of an ultrafilter over I.

Definition 3.7. $U \subseteq \mathcal{P}(I)$ is an *ultrafilter over I* if and only if U is a proper filter over I and, for all $X \in \mathcal{P}(I)$, exactly one the following is true:

$$X \in U$$

$$I \setminus X \in U$$

Proposition 3.8. *U is an ultrafilter over I if and only if U is a maximal proper filter. That is, U is a proper filter over I such that if V is a proper filter and $U \subseteq V \subseteq \mathcal{P}(I)$, then $U = V$.*

Proof. In the forward direction, assume U is an ultrafilter over I. Let V be a proper filter over I such that $U \subseteq V \subseteq \mathcal{P}(I)$. Assume there is some $X \in V$ with $X \notin U$. Then $I \setminus X \in U$ by definition of an ultrafilter but then $I \setminus X \in V$ as U is contained in V. By closure under finite intersection, $X \cap (I \setminus X) = \emptyset \in V$ and V is not proper. Thus, there is no $X \in V$ with $X \notin U$, so $U = V$ and U is a maximal proper filter.

In the reverse direction, assume U is a maximal proper filter over I. Let $X \in \mathcal{P}(I)$. We know if $X \in U$ then $I \setminus X$ cannot be in U by closure under finite intersection and because U is proper. So if $X \in U$, then $I \setminus X \notin U$. To conclude, we show that if $X \notin U$, then $I \setminus X \in U$. Suppose $X \notin U$. By maximality of U, the filter generated by $U \cup \{X\}$ is improper and therefore, by Proposition 3.6, does not have the finite intersection property. Since every finite intersection of elements in U is non-empty, a finite intersection including X must be empty. That is, there are $Y_1, \ldots, Y_n \in U$ such that $Y_1 \cap \ldots \cap Y_n \cap X = \emptyset$, and therefore $\bigcap Y_i \subseteq I \setminus X$. This entails that $I \setminus X \in U$ by upward closure. By definition, U is an ultrafilter. □

Given any $E \subseteq \mathcal{P}(I)$ with the finite intersection property, we can build a proper filter containing E and we now show that we can always build an ultrafilter containing E, as well.

Proposition 3.9. *If $E \subseteq \mathcal{P}(I)$ has the finite intersection property, then there is an ultrafilter U over I that contains E.*

Proof. By Zorn's Lemma. Consider the non-empty set K of all proper filters over I containing E, ordered under $\subseteq$. The set is non-empty by Proposition 3.6. Let C be a chain in K. Then $\bigcup C$ is an upper bound of C contained in K. That is, $\bigcup C$ is a proper filter over I containing E. Let U be a maximal element of K given by Zorn's Lemma. If U' is a proper filter such that $U \subseteq U'$, then $E \subseteq U'$ so U' is in K. But then, by the maximality of U, $U = U'$. By Proposition 3.8, U is an ultrafilter over I. □

An important corollary is that *any proper filter can be extended to an ultrafilter* because all proper filters have the finite intersection property.

Definition 3.10. A set is *co-finite* if its complement is finite.

We now notice that all co-finite subsets of I are "large" with respect to non-principal ultrafilters over I:

Proposition 3.11. *Any non-principal ultrafilter U over I contains all co-finite sets.*

Proof. First, we show that if U is an ultrafilter such that $\{x\} \in U$ for some $x \in I$, then U is principal. Assume $\{x\} \in U$ and let $Y \in U$ be arbitrary. Then $\{x\} \cap Y \neq \emptyset$ so $\{x\} \subseteq Y$ and U is contained in the principal ultrafilter generated by $\{x\}$. U contains the principal ultrafilter by upward closure, which proves that U is principal.

Thus, since U is non-principal, $\{x\} \notin U$ and $I \setminus \{x\} \in U$ for each $x \in I$. It follows that $\bigcap_{i=1}^{n} I \setminus \{x_i\} \in U$ by closure under finite intersection. That is, U contains all co-finite sets. □

4. Ultraproducts and Ultrapowers

In this section, the ultraproduct construction is introduced. We focus on the ultraproduct's interpretation of some language $\mathscr{L}$ and the nature of its theory. We begin to explore the intricate relationship between the ultraproduct and the ultrafilter used to build it. The Fundamental Theorem of Ultraproducts, which is stated in this section, is an important example of this relationship. The section concludes with a proof of Theorem 2.20 using the ultraproduct construction.

Definition 4.1. Let I be a nonempty set, D an ultrafilter over I, and, for each $i \in I$, let A_i be a set. We let $C = \prod_{i \in I} A_i$ be the Cartesian product of these sets and define an equivalence relation $=_D$ such that for $f, g \in C$

$$f =_D g \iff \{i \in I : f(i) = g(i)\} \in D.$$

Intuitively, we take a "majority vote". If D says the majority of components say $f = g$, then we deem them equal with respect to $=_D$. Denote the equivalence class of f by

$$f_D = \{g \in C : f =_D g\}.$$

The set of equivalence classes of $=_D$ is the *ultraproduct of A_i modulo D*.

Definition 4.2. An ultraproduct of A_i modulo D where $A_i = A$ for all $i \in I$ is called an *ultrapower of A modulo D*.

We wish to extend the ultraproduct construction to a set of $\mathscr{L}$-models $\mathfrak{A}_i$. What we end up building is another $\mathscr{L}$-model, whose interpretation and theory are determined by taking a "vote". Formally:

Definition 4.3. Let I be a non-empty set and D be an ultrafilter over I. For each $i \in I$, let $\mathfrak{A}_i$ be a model for the language $\mathscr{L}$. For each $\mathfrak{A}_i$, we have relation symbols P interpreted as R_i, functions F as G_i, and constants c as a_i. The *ultraproduct* $\prod_D \mathfrak{A}_i$ is a model for $\mathscr{L}$ with the following properties:

(1) The universe set of $\prod_D \mathfrak{A}_i$ is $\prod_{i\in I} A_i$, the ultraproduct defined on sets.
(2) Let P be an n-placed relation symbol in $\mathscr{L}$. Then the interpretation of P in $\prod_D \mathfrak{A}_i$ is the relation S such that
$$S(f^1_D, \ldots, f^n_D) \iff \{i \in I : R_i(f^1(i), \ldots, f^n(i))\} \in D.$$
(3) Let F be an n-placed function symbol of $\mathscr{L}$. Then F is interpreted in $\prod_D \mathfrak{A}_i$ by the function H such that
$$H(f^1_D, \ldots, f^n_D) = g_D$$
where
$$g_D(i) = G_i(f^1(i), \ldots, f^n(i)) \in A_i.$$
(4) Let c be a constant of $\mathscr{L}$. Then c is interpreted as the element $a \in \prod_D \mathfrak{A}_i$ represented by the tuple
$$a = (a_1, a_2, \ldots).$$

A few notes:

Notation 4.4. When defining an element in $\prod_D \mathfrak{A}_i$ coordinate-wise, such as when we define the value of H, we equivalently write
$$g_D = \left\langle G_i(f^1(i), \ldots, f^n(i)) : i \in I \right\rangle_D .$$
So a, the interpretation of c, is equivalently denoted
$$a = \langle a_i : i \in I \rangle_D .$$

Remark 4.5. The ultraproduct is well defined because the interpretations S and H do not depend on which f in each equivalence class is taken as the representative.

There are at least two reasons model theorists find ultraproducts and ultrapowers useful that are discussed in this paper. First, ultraproducts can be built to be saturated models (see Section 6), which means that, in some sense, everything that could possibly happen in the ultraproduct does happen. Thus, if one wants to study certain properties or kinds of elements, an ultraproduct is a fruitful world to be in. Secondly, as mentioned above, ultrapowers characterize elementary equivalence. The main theorem shows that a weak condition of sameness, namely elementary equivalence, can be transformed into a strong condition of sameness, namely isomorphism, via the ultrapower construction.

The following theorem will be stated without proof, but is pivotal in understanding the theory of an ultraproduct. It precisely states that the ultraproduct takes a vote to determine what its theory will be.

Theorem 4.6 (The Fundamental Theorem of Ultraproducts)**.** *Let $\prod_D \mathfrak{A}_i$ be an ultraproduct of $\mathscr{L}$-models $\mathfrak{A}_i$ indexed over I. For any sentence ϕ of $\mathscr{L}$,*

$$\prod_D \mathfrak{A}_i \models \phi \iff \{i \in I : \mathfrak{A}_i \models \phi\} \in D.$$

Corollary 4.7. *Let $\mathfrak{A}$ be a model of $\mathscr{L}$ and D be any ultrafilter. Then*

$$\prod_D \mathfrak{A} \equiv \mathfrak{A}.$$

Theorem 4.6 and Corollary 4.7 are very important facts and will be used extensively throughout the rest of this paper. A few "philosophical" notes to consider. We see yet another way in which the world of an ultraproduct is a democracy [2]. The ultraproduct takes a vote amongst its underlying models on what first order sentences it satisfies, and the majority, as determined by the ultrafilter, gets the final say. This is the first indication of the intricate relationship between the theory of the ultrapower and its related ultrafilter, which leads me to a digression.

Ultrafilters are used to build ultraproducts and the elements in said ultrafilter determine what the ultraproduct says and doesn't say. However, our construction of non-principal ultrafilters is dependent upon the Axiom of Choice, and is not explicitly given. So, in some sense, we know these ultrafilters exist and a few things about their structure, but we largely don't know what specific elements they contain. Thus, we engineer ultrafilters in a special way so that we can "look inside" of the them to obtain a greater understanding of their contents. Ultimately, what we can know about the ultraproduct is restricted by what we know about the ultrafilter used to build it. We will see this at work in theorems and proofs in coming sections.

We close this section with an interesting exercise, namely an alternative proof of the Compactness Theorem using ultraproducts.

Theorem 4.8. *(Compactness Theorem) Let Σ be a set of sentences of $\mathscr{L}$. Then Σ has an $\mathscr{L}$-model if and only if every finite subset has an $\mathscr{L}$-model.*

Proof. The necessary direction is trivial.

Going the other direction, the essence of the proof is to show that an ultraproduct of the models satisfying the finite subsets of Σ can be built to satisfy Σ, by cleverly constructing an ultrafilter. Our index set will be $S_\omega(\Sigma)$, the set of finite subsets of Σ. For each $s \in S_\omega(\Sigma)$, let M_s be the $\mathscr{L}$-model such that $M_s \models s$. We build an ultrafilter U over $S_\omega(\Sigma)$ so that the ultraproduct of $\{M_s\}_{s \in S_\omega(\Sigma)}$ modulo U models all of Σ.

Define for each $\phi \in \Sigma$

$$A_\phi = \{s \in S_\omega(\Sigma) : \phi \in s\}$$

and let

$$E = \{A_\phi : \phi \in \Sigma\}.$$

Then E is a subset of $\mathcal{P}(S_\omega(\Sigma))$ and has the finite intersection property, since $\{\phi_1, \ldots, \phi_n\} \in A_{\phi_1} \cap \ldots \cap A_{\phi_n}$ for any $A_{\phi_1}, \ldots, A_{\phi_n} \in E$. Thus, by Theorem 3.8, there exists an ultrafilter U over $S_\omega(\Sigma)$ containing E.

[2]Or a dictatorship if one uses a principal ultrafilter.

Consider $\prod_U M_s$. Fix $\phi \in \Sigma$. If $\phi \in s$, then $M_s \vDash \phi$. But for every $s \in A_\phi$, we have that $\phi \in s$. So

$$A_\phi \subseteq \{s \in S_\omega(\Sigma) : M_s \vDash \phi\} \in U$$

by the upward closure. By the fundamental theorem $\prod_U M_s \vDash \phi$, which entails $\prod_U M_s \vDash \Sigma$. □

5. Saturation of Models

This section is focused on the notion of a *saturated* model, which can informally be thought of as a model that contains all of its "limit points". In order to understand what a saturated model looks like, *types* and the *realizations of types* are introduced. We conclude with an important theorem that shows how to build an isomorphism between two saturated, elementarily equivalent models of the same size. This uniqueness theorem is invoked to prove the more difficult direction of the main theorem, which states that two ultrapowers built from elementarily equivalent models and a countably incomplete good ultrafilter are isomorphic.

Definition 5.1. Suppose $\Sigma(x_1, \dots, x_n)$ is a set of formulas of $\mathscr{L}$ in the variables $x_1, \dots, x_n$. A tuple $(a_, \dots, a_n) \in A^n$ *realizes* $\Sigma(x_1, \dots, x_n)$ if and only if

$$M \models \sigma[a_1, \cdots, a_n]$$

for every $\sigma \in \Sigma$. A model $\mathfrak{A}$ *realizes* $\Sigma(x_1, \dots, x_n)$ if and only if there is a tuple $(a_1, \dots, a_n) \in A^n$ that realizes Σ.

Definition 5.2. A set of formulas $\Sigma(x_1, \dots, x_n)$ of $\mathscr{L}$ in the variables $x_1, \dots, x_n$ is *consistent* if and only if some model realizes Σ.

Informally, a *type* in one variable is a set of first order formulas that gives a complete "functional description" of an element v that may or may not exist in a model. Rather than explicitly naming v, the element can be "pointed" at by a set of formulas. If the type is realized in a model, then there is an element v of the model's universe that satisfies all of the properties expressed by the type. Formally:

Definition 5.3. A *type* in the variables $x_1, \dots, x_n$ is a maximal consistent set $\Gamma(x_1, \dots, x_n)$ of formulas in the variables $x_1, \dots, x_n$. A *partial type* is a consistent set $\Gamma(x_1, \dots, x_n)$ of formulas in the variables $x_1, \dots, x_n$.

Definition 5.4. A set of formulas $\Sigma(x_1, \dots, x_n)$ is *consistent with a theory T* if and only if there is a model $\mathfrak{A}$ of T that realizes $\Sigma(x_1, \dots, x_n)$.

This proposition will be helpful later on. It states that consistent finite types are always realized.

Proposition 5.5. *A model $\mathfrak{A}$ realizes any finite type $\Sigma(x_1, \dots, x_n)$ that is consistent with its theory.*

Proof. Let $\Sigma(x_1, \dots, x_n)$ be a finite type that is consistent with the model's theory. By definition, there is some model $\mathfrak{B}$ such that $\mathfrak{B} \models Th(\mathfrak{A})$ and realizes $\Sigma(x_1, \dots, x_n)$. Since $Th(\mathfrak{A})$ is complete, $Th(\mathfrak{A}) = Th(\mathfrak{B})$. Consider $\phi = \exists x_1 \dots \exists x_n \bigwedge_{\sigma \in \Sigma[x_1, \dots, x_n]} \sigma[x_1, \dots, x_n]$. Then $\phi \in Th(\mathfrak{B})$, as $\mathfrak{B}$ realizes $\Sigma(x_1, \dots, x_n)$, which entails that $\phi \in Th(\mathfrak{A})$. By definition of ϕ, $\mathfrak{A}$ realizes $\Sigma(x_1, \dots, x_n)$. □

The notion of *saturation* gives a measure of how often or how completely a model realizes its types. Let α be a cardinal. Informally, a model is α-saturated whenever it realizes every type that can be described with less than α many terms in the model's universe. One could also think of a saturated model as one that fills the "holes" in its universe. Formally:

Definition 5.6. A model $\mathfrak{A}$ is *α-saturated* if and only if, for any subset $X \subseteq A$ such that $|X| < \alpha$, the expansion $(\mathfrak{A}, a)_{a\in X}$ realizes every type $\Sigma(v)$ of the language $\mathscr{L} \cup \{c_a : a \in X\}$ that is consistent with the theory of $(\mathfrak{A}, a)_{a\in X}$. A model with universe A is *saturated* if it is $|A|$-saturated.

Example 5.7. Consider the (partial) type

$$\Sigma(v) = \{c_n \leq v : n \in \mathbb{N}\} \cup \{c_n \neq c_m : m \neq n\}$$

in the language $\mathscr{L}' = \mathscr{L} \cup \{c_n : n \in \mathbb{N}\}$ where $\mathscr{L}$ only contains the $\leq$ relation symbol. Further, consider the $\mathscr{L}$-model $\mathfrak{N} = \langle \mathbb{N}, \leq_{\mathbb{N}} \rangle$ with the usual ordering. $\Sigma(v)$ describes an element v that is greater or equal to countably many distinct elements of the universe. It is clear that $\mathfrak{N}$ with the usual ordering cannot be extended to an $\mathscr{L}'$-model in such a way that it realizes Σ. For any element $n \in \mathbb{N}$, $m \leq_{\mathbb{N}} n$ for only finitely many m and so given any interpretation of the constant symbols, no element of the natural numbers satisfies $\Sigma(v)$.

However, if one considers the "reverse" ordering on the natural numbers, I claim the natural numbers do satisfy such a type. Let $\mathfrak{N}' = \langle \mathbb{N}, \leq_{\mathbb{N}'} \rangle$ where $\leq_{\mathbb{N}'}$ is the reverse ordering on the natural numbers. That is, define $\leq_{\mathbb{N}'}$ such that $n \leq_{\mathbb{N}'} m \iff m \leq_{\mathbb{N}} n$. Extend $\mathfrak{N}$ such that n interprets c_n for each $n \in \mathbb{N}$. Notice then that $c_n \leq_{\mathbb{N}'} 0$ for all $n \in \mathbb{N}$. Thus $(\mathfrak{N}', n)_{n\in\mathbb{N}}$ realizes $\Sigma(v)$.

Example 5.8. I claim that $\mathfrak{Q} = \langle \mathbb{Q}, \leq_{\mathbb{Q}} \rangle$ with the usual ordering is saturated. We must consider the existence of any finitely describable element that is described entirely in terms of its place in the order $\leq_{\mathbb{Q}}$. Density of $\leq_{\mathbb{Q}}$ will allow us to find any element in the order consistently described using only finitely many terms of $\mathbb{Q}$.

On the other hand, $\mathfrak{R} = \langle \mathbb{R}, \leq_{\mathbb{R}} \rangle$ is not saturated, because it is not ω_1-saturated. All countably describable positions in the order are not filled. Consider the (partial) type $\Sigma(v)$ that describes an infinitesimally small element. Let

$$X = \{\frac{1}{n} : n \in \mathbb{N}\} \cup \{0\} \subset \mathbb{R}.$$

Consider the expansions $\mathscr{L} \cup \{c_a : a \in X\}$ and $(\mathfrak{R}, a)_{a\in X}$. Then we have the following type:

$$\Sigma(v) = \{c_0 \leq v \leq c_a : a \in X \setminus \{0\}\} \cup \{v \neq c_0\}.$$

Notice that Σ is describing an element v such that $0 < v \leq \frac{1}{n}$ for all $n \in \mathbb{N}$. No such element of $\mathbb{R}$ satisfies this type. But interestingly, if we use a special type of ultrafilter D, the ultrapower of $\mathfrak{R}$ modulo D is ω_1 saturated and does realize an infinitesimally small element. [3]. This is a useful property of ultrapowers: they can be built to be saturated "extensions" of models that are not saturated. Much more on this later!

---

[3]In particular, the ultrapower is elementarily equivalent but not isomorphic to $\mathfrak{R}$.

A saturated model is special in that, in some sense, its first order theory has complete control over the model's properties, up to cardinality. The Löwenheim-Skolem-Tarski Theorem shows that if a theory T has an infinite model, then T has infinite models of any given power $\alpha \geq ||\mathscr{L}||$ (See Section 2.1 of [1]). In this way, a first order theory cannot force a particular infinite cardinality on its models. However, when an infinite model is saturated, the theory of the model controls everything else about the model.

Theorem 5.9 (Uniqueness of Saturated Models)**.** *Let $\mathfrak{A}$ and $\mathfrak{B}$ be elementarily equivalent, saturated models of the same cardinality. Then $\mathfrak{A} \cong \mathfrak{B}$.*

Proof. Essentially, the proof is to build a map from $\mathfrak{A}$ to $\mathfrak{B}$ that sends types in the model $\mathfrak{A}$ to the type's realization in $\mathfrak{B}$, which can be found due to the saturation of the models. We do this systematically, through a "back and forth" method, so that both universes are completely accounted for at the end of the recursion (i.e. the map we construct will be necessarily surjective).

We will only prove the case where α is an infinite cardinal. The finite case is basically the same; one repeats the successor stage until the recursion terminates. Let $|A| = |B| = \alpha$ and let us well order A and B as the cardinal α. We build an isomorphism $f : A \to B$ via transfinite recursion. For each $\beta < \alpha$, we define A_β, B_β and $f_\beta : A_\beta \to B_\beta$ such that:

(1) $f_\beta : A_\beta \to B_\beta$ is a bijection.
(2) $\{a_\zeta\}_{\zeta<\beta} \subseteq A_\beta$ and $\{b_\zeta\}_{\zeta<\beta} \subseteq B_\beta$
(3) $A_\lambda \subseteq A_\beta$ and $B_\lambda \subseteq B_\beta$ for $\lambda \leq \beta$
(4) f_β extends f_λ for each $\lambda \leq \beta$
(5) For each formula $\phi(v_1, \ldots, v_n)$ in $\mathscr{L}$ and $x_1, \ldots, x_n \in A_\beta$
$$\mathfrak{A} \models \phi[x_1, \ldots, x_n] \iff \mathfrak{B} \models \phi[f_\beta(x_1), \ldots, f_\beta(x_n)].$$
(6) $(\mathfrak{A}, A_\beta) \equiv (\mathfrak{B}, B_\beta)$ in the language $\mathscr{L} \cup \{c_a\}_{a \in A_\beta}$. Note that $a \in A_\beta$ and $f(a) \in B_\beta$ interpret c_a in $(\mathfrak{A}, A_\beta)$ and $(\mathfrak{B}, B_\beta)$ respectively.

Base Case: Set $A_0 = \emptyset, B_0 = \emptyset$ and $f_0(\emptyset) = \emptyset$. $(1) - (5)$ hold and (6) holds by the assumption that $\mathfrak{A} \equiv \mathfrak{B}$.

Successor Case: Assume $(1) - (6)$ hold for f_*, g_*, A_*, B_* up to $S(\theta) < \alpha$, a successor ordinal. We begin with the "forth" construction. Let a_θ be the least element in $A \setminus A_\theta$. Trust, for now, that $|A_\theta| < \alpha$. I conclude by showing this is true for all A_β, such that $\beta < \alpha$. Consider the type of a_θ in $(\mathfrak{A}, A_\theta)$, a model in the language $\mathscr{L} \cup \{c_a\}_{a \in A_\theta}$. By (6), we have that $(\mathfrak{A}, A_\theta)$ and $(\mathfrak{B}, B_\theta)$ are elementarily equivalent. Thus the type of a_θ is consistent with $Th(\mathfrak{B}, B_\theta)$, and by α-saturation, is realized in the model. Set $f_{S(\theta)}(a_\theta)$ to some chosen realization.

We continue with the "back" part of the construction. Consider the least element of $B \setminus \{B_\theta \cup \{f_{S(\theta)}(a_\theta)\}\}$, denoted b_θ. We repeat the "forth" argument, using the models $(\mathfrak{A}, A_\theta \cup \{a_\theta\})$ and $(\mathfrak{B}, B_\theta \cup \{f_{S(\theta)}(a_\theta)\})$ in the language $\mathscr{L} \cup \{c_a\}_{a \in A_\theta \cup \{a_\theta\}}$, where a_θ and $f_{S(\theta)}(a_\theta)$ interpret the added constant, and A_θ and B_θ continue to interpret constants as they did in the "forth" direction. By the choice of $f_{S(\theta)}(a_\theta)$, it follows that $(\mathfrak{A}, A_\theta \cup \{a_\theta\}) \equiv (\mathfrak{B}, B_\theta \cup \{f_{S(\theta)}(a_\theta)\})$. Thus, the type of b_θ is consistent with $(\mathfrak{A}, A_\theta \cup \{a_\theta\})$. Because $|A_\theta| < \alpha$, it follows that $|A_\theta| + 1 < \alpha$ and by α-saturation $(\mathfrak{A}, A_\theta \cup \{a_\theta\})$ realizes the type of b_θ via some element a'_θ. Define $f_{S(\theta)}(a'_\theta) = b_\theta$. We set:
$$A_{S(\theta)} = A_\theta \cup \{a_\theta, a'_\theta\}$$

$$B_{S(\theta)} = B_\theta \cup \{f_{S(\theta)}(a_\theta), b_\theta\}$$

Let $f_{S(\theta)}$ be defined as discussed at a_θ and a'_θ and $f_{S(\theta)} = f_\theta$ on A_θ. $(1) - (4)$ hold by construction. By defining $f_{S(\theta)}$ to match types on the added elements, (5) holds and therefore (6) holds.

Limit case: Let $\lambda < \alpha$ be a limit ordinal and assume that f_*, g_*, A_*, B_* are all defined to satisfy $(1) - (6)$ for $\beta < \lambda$. Then we define:

$$A_\lambda = \bigcup_{\beta<\lambda} A_\beta$$

$$B_\lambda = \bigcup_{\beta<\lambda} B_\beta$$

(2) and (3) follow immediately. Further, we can define f_λ by utilizing any applicable f_β for each element in the union since the functions were built compatibly. Then (1) and (4) follow as well. (5) and (6) follow by hypothesis because we can always consider a sentence or formula in $\mathscr{L} \cup \{c_a\}_{a\in A_\lambda}$ as living in some language with $\beta < \lambda$ many constants, where formulas and sentences act as (5) and (6) require.

Continue recursively until $f_\beta : A_\beta \to B_\beta$ is defined for all $\beta < \alpha$. Then implement the limit stage construction one last time to get f_α, an isomorphism from A to B by $(1), (2), (5)$. There is one detail that must be addressed, namely that we are not adding so many elements at each stage that saturation was not sufficient for realizing a type in the extended model. Recall that saturation guarantees realization only if less than α constants are introduced.

We use transfinite recursion to prove $|A_\theta| < \alpha$ for every $\theta < \alpha$ by building a bijection $i_\theta : A_\theta \to \theta \times \{0, 1\}$ at each stage. Note that $|A_{S(\theta)} \setminus A_\theta| = 2$ so given that $|\theta| < \alpha$ and $|\theta \times \{0, 1\}| = 2|\theta| = |\theta| < \alpha$, constructing the above bijection suffices to show $|A_\theta| < \alpha$. At the base case, define $i_0 : A_0 \to \emptyset$ to be the identity on $\emptyset$. For the successor case, assume there is a bijection from A_β to $\beta \times \{0, 1\}$. We want to build a bijection $i_{S(\beta)} : A_{S(\beta)} \to S(\beta) \times \{0, 1\}$. When building $A_{S(\beta)}$ from A_β, we add two elements. Thus we can send A_β bijectively to $\beta \times \{0, 1\}$ by i_β and send the 2 added elements to $(\beta, 0)$ and $(\beta, 1)$. Then $i_{S(\beta)}$ is a bijection from $A_{S(\beta)}$ to $S(\beta) \times \{0, 1\}$.

For the limit case, let $\lambda < \alpha$ be a limit ordinal. Then we are looking to construct a bijection $i_\lambda : \bigcup_{\beta<\lambda} A_\beta \to \lambda \times \{0, 1\}$. A bijection exists because the bijection at each stage was taken to be compatible with bijections at all stages before it. Thus we can define a bijection on the union in terms of the bijections on A_β for $\beta < \lambda$. This completes the proof. □

Corollary 5.10. *Any finite, elementarily equivalent models $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic.*

Proof. Since there is a first order sentence stating that a model has exactly n elements for any finite n, $|A| = |B|$ by elementary equivalence of $\mathfrak{A}$ and $\mathfrak{B}$. Let $X \subseteq A$ such that $|X| < |A|$. We show $\mathfrak{A}$ is saturated by showing that if $\Sigma(v)$, a type in $\mathscr{L} \cup \{c_a\}_{a\in X}$, is not realized then $\Sigma(v)$ is not consistent with $(\mathfrak{A}, a)_{a\in X}$. If $\Sigma(v)$ is not realized in $(\mathfrak{A}, a)_{a\in X}$, then for each $b \in A$, there is some formula $\sigma_b \in \Sigma(v)$ such that $(\mathfrak{A}, a)_{a\in X} \models \neg\sigma_b[b]$. However, we can then show that $\{\sigma_b : b \in A\}$, a finite subset of $\Sigma(v)$, is not consistent with $(\mathfrak{A}, a)_{a\in X}$ by showing every model of the theory of $(\mathfrak{A}, a)_{a\in X}$ does not realize $\{\sigma_b : b \in A\}$. Define the sentence

$$\phi = \neg\exists x \bigwedge_{b\in A} \sigma_b[x].$$

Notice that the finiteness of A entails that ϕ is a first order sentence. Further, $(\mathfrak{A}, a)_{a\in X} \models \phi$ and, for any $\mathfrak{B}$ that models the theory of $(\mathfrak{A}, a)_{a\in X}$, we know $\mathfrak{B} \models \phi$. By definition of ϕ, $\mathfrak{B}$ cannot realize $\{\sigma_b : b \in A\}$, therefore this finite subset of $\Sigma(v)$ is inconsistent with $(\mathfrak{A}, a)_{a\in X}$. It follows that $\Sigma(v)$ is inconsistent with $(\mathfrak{A}, a)_{a\in X}$. Therefore, $\mathfrak{A}$ and $\mathfrak{B}$ are saturated models. By Theorem 5.9, $\mathfrak{A} \cong \mathfrak{B}$. □

We can now sketch how the main theorem of this paper works. As stated in the introduction, the main theorem states that under certain conditions, $\prod_D \mathfrak{A} \cong \prod_D \mathfrak{B}$ if and only if $\mathfrak{A} \equiv \mathfrak{B}$. The forward direction holds because an ultrapower is elementarily equivalent to its underlying model and elementary equivalence is transitive. The reverse direction utilizes the uniqueness of saturated models theorem. Thus, we show that for a certain ultrafilter D:

(1) $\prod_D \mathfrak{A} \equiv \prod_D \mathfrak{B}$: We will use the fundamental theorem of ultrapowers.
(2) Both ultrapowers are saturated: We will restrict the particular ultrafilter used and bound the power of the language to get saturation at the upper bound of their powers.
(3) $|\prod_D \mathfrak{A}| = |\prod_D \mathfrak{B}|$: We will use the generalized continuum hypothesis to bound cardinality from above and saturation to bound cardinality from below.

If conditions are set appropriately, the isomorphism between the two ultrapowers with elementarily equivalent underlying models will follow from Theorem 5.9.

6. Good Ultrafilters and The Main Theorem

This last section is devoted to a discussion of countably incomplete good ultrafilters. We show how such an ultrafilter saturates its ultrapowers, which will be a key step in proving the main theorem. We conclude with a formal statement and proof of the main theorem.

To begin, a *countably incomplete* ultrafilter is used to "keep things finite" in the construction of the ultraproduct:

Definition 6.1. A filter D over I is said to be *countably incomplete* if there exists a countable subset $E \subseteq D$ such that $\bigcap E \notin D$.

Proposition 6.2. *An ultrafilter D is countably incomplete if and only if there exists a countable decreasing chain*

$$I = I_0 \supset I_1 \supset I_2 \supset I_3 \supset \cdots$$

of elements $I_n \in D$ with $\bigcap_n I_n = \emptyset$.

Proof. Assume D is a countably incomplete ultrafilter. Then there exists a countable set $\{I_n \in D : n \in \mathbb{N}\}$ such that $\bigcap I_n \notin D$. Let

$$I_1' = I_1 \setminus \bigcap I_n$$

$$I_{n+1}' = I_n' \cap I_{n+1}$$

Then we see that each $I_n' \in D$, since $(\bigcap I_n)^c \in D$ and $I_n \in D$ for all $n \in \mathbb{N}$. Further, one can check that $I_n' \supset I_{n+1}'$. The intersection is empty, because $\bigcap_n I_n' =$

$\bigcap_n I_n \setminus \bigcap_n I_n = \emptyset$. So we have a countable decreasing chain of $I'_n \in D$ with $\bigcap_n I'_n = \emptyset$.

In the other direction, if there exists such a countable decreasing chain with empty intersection, then $\{I_n : n \in \mathbb{N}\} \subseteq D$ is a countable set whose intersection is the empty set, which is not in D. Thus, D is countably incomplete. □

However, a countably incomplete ultrafilter is far from sufficient to guarantee α-saturation for arbitrary α. The ultrafilter must take a more restrictive form. First, a few necessary definitions about functions:

Definition 6.3. Let I be a nonempty set, β be a cardinal, and consider functions f, g on the set $S_\omega(\beta)$, the set of finite subsets of β, into $\mathcal{P}(I)$. We say that $g \leq f$ if and only if for all $u \in S_\omega(\beta)$, $g(u) \subseteq f(u)$. That is, the value of g at u is contained in the value of f at u for all $u \in S_\omega(\beta)$.

Definition 6.4. Let f be a function like in Definition 6.3. Then f is *monotonic* if and only if for $u, v \in S_\omega(\beta)$, we have

$$u \subseteq v \Longrightarrow f(u) \supseteq f(v).$$

One may think of monotonicity as "anti-monotonicity".

Definition 6.5. Let g be a function like in Definition 6.3. Then g is *additive* on $S_\omega(\beta)$ if and only if for $u, v \in S_\omega(\beta)$, we have

$$g(u \cup v) = g(u) \cap g(v).$$

Once again, one may think of additivity as "anti-additivity".

Because an ultrafilter over I is contained in $\mathcal{P}(I)$, it makes sense to look at functions from $S_\omega(\beta)$ into the ultrafilter. In doing so, we can specify a type of ultrafilter that will allow us to prove the main theorem.

Definition 6.6. An ultrafilter D over I is said to be α-*good* if it has the following property: For every $\beta < \alpha$ and for every monotonic function $f : S_\omega(\beta) \to D$, there exists a function $g : S_\omega(\beta) \to D$ such that $g \leq f$ and g is additive on $S_\omega(\beta)$.

Remark 6.7. It is easy to check that every additive function is monotonic and that if D is α-good then it is β-good for all $\beta < \alpha$

I state the existence of good countably incomplete ultrafilters without proof. See Section 6.1 of [1] for the construction.

Theorem 6.8. *Let I be a set of size α. Then there exists an α^+-good* [4] *countably incomplete ultrafilter D over I.*

In the following theorem and corollary, we show how to build models of arbitrarily large saturation using α-good ultrafilters.

Theorem 6.9. *Let α be an infinite cardinal and let D be a countably incomplete α-good ultrafilter over a set I. Let $||\mathscr{L}|| < \alpha$. Suppose $\{\mathfrak{A}_i\}_{i \in I}$ is a family of $\mathscr{L}$-structures and $\Sigma(x)$ is a set of formulas of $\mathscr{L}$. If every finite subset of $\Sigma(x)$ is realized in $\prod_D \mathfrak{A}_i$, then $\Sigma(x)$ is realized in $\prod_D \mathfrak{A}_i$.*

[4]We use α^+ to denote the successor of the ordinal number α.

Proof. Since D is countably incomplete, there exists a countable, decreasing sequence

$$I = I_0 \supset I_1 \supset I_2 \supset \dots$$

with $I_n \in D$ for all $n \in \mathbb{N}$ and $\bigcap_n I_n = \emptyset$. Since $||\mathscr{L}|| < \alpha$, we know $|\Sigma(x)| < \alpha$. We define the function $f : S_\omega(\Sigma) \to D$ as follows:

$$f(\sigma) = I_{|\sigma|} \cap \{i \in I : \mathfrak{A}_i \models \exists x \bigwedge_{\phi \in \sigma} \phi(x)\}.$$

I claim f is monotonic. If $\tau \subset \sigma$, then

$$I_{|\tau|} \supset I_{|\sigma|}$$

and

$$\{i \in I : \mathfrak{A}_i \models \exists x \bigwedge_{\phi \in \tau} \phi(x)\} \supseteq \{i \in I : \mathfrak{A}_i \models \exists x \bigwedge_{\phi \in \sigma} \phi(x)\},$$

which entails that $f(\sigma) \supseteq f(\tau)$. Since D is α-good and $|\Sigma(x)| < \alpha$, there exists an additive function $g : S_\omega(\Sigma) \to D$ such that for all $\sigma \in S_\omega(\Sigma)$, $g(\sigma) \subseteq f(\sigma)$. Constructing and using such a refinement of f is the crux of the proof, so I will take a moment to consider what is at work here.

If the ultraproduct realizes every $\theta \in \Sigma$ at once, as we'd like, then there is some collection of "majorities" $\{A_\theta\}_{\theta \in \Sigma}$, one for each $\theta \in \Sigma$, that agree there is an element that realizes each formula in Σ. If a model $\mathfrak{A}_i$ is contained in some finite subset of $\{A_\theta\}_{\theta \in \Sigma}$, then there is some $\sigma(i) \subseteq \Sigma$ of formulas such that $\mathfrak{A}_i \models \exists x \theta[x]$ for each $\theta \in \sigma(i)$. However, even stronger, it is also necessarily true that $\mathfrak{A}_i \models \exists x \bigwedge_{\theta \in \sigma(i)} \theta[x]$.

The additive nature of g gives us the power to find such a collection of majorities by setting $A_\theta = g(\{\theta\})$. By inductively applying that $g(u \cup v) = g(u) \cap g(v)$, it follows that for a finite subset $\sigma \subset \Sigma$, $g(\sigma) = \bigwedge_{\theta \in \sigma} g(\{\theta\})$. When applying this to the finite subset $\sigma(i) \subset \Sigma$, additivity ensures that $\bigwedge_{\theta \in \sigma(i)} g(\{\theta\}) \subset g(\sigma(i))$. Translating, we have that any model that belongs to the majority vote for there being an element that realizes a formula, for finitely many formulas, belongs to the majority vote for there being an element that realizes the finite set of formulas at the same time. This is sufficient for the ultraproduct to realize Σ. Note that if $\sigma(i)$ was infinite, additivity would not be strong enough to ensure the model realized $\sigma(i)$. By utilizing a countably incomplete filter, we manage to keep $\sigma(i)$ finite.

Here is the formalization of this commentary. For each $i \in I$, let

$$\sigma(i) = \{\theta \in \Sigma : i \in g(\{\theta\})\}.$$

I claim that $\sigma(i)$ is finite for all $i \in I$. Fix $i \in I$. If $|\sigma(i)| \geq n$, then choose n distinct elements $\theta_1, \dots, \theta_n$ of $\sigma(i)$. We then have

$$i \in g(\{\theta_j\})$$

for $1 \leq j \leq n$, which by additivity gives us that

$$i \in g(\{\theta_1\}) \cap \dots \cap g(\{\theta_n\}) = g(\{\theta_1 \cup \dots \cup \theta_n\}) \subseteq f(\{\theta_1 \cup \dots \cup \theta_n\}) \subseteq I_n.$$

If $\sigma(i)$ was not finite, then $i \in I_n$ for infinitely many n, which is a contradiction. Thus $\sigma(i)$ is finite and i is contained in $g(\{\theta\})$ for finitely many $\theta \in \Sigma$.

We build h_D that will satisfy $\Sigma(x)$ in $\prod_D \mathfrak{A}_i$. If $\sigma(i)$ is empty, then choose an arbitrary $h(i) \in A_i$. Else, for each $i \in I$ with $\sigma(i)$ non-empty, we have by additivity and the definition of $\sigma(i)$ that

$$i \in \bigcap\{g(\{\theta\}) : \theta \in \sigma(i)\} = g(\{\theta \in \Sigma : i \in g(\{\theta\})\}) \subseteq f(\sigma(i)).$$

Thus $i \in \{j \in I : \mathfrak{A}_j \models \exists x \bigwedge_{\phi \in \sigma(i)} \phi(x)\}$.

We choose $h(i)$ to be some element that satisfies $\bigwedge_{\phi \in \sigma(i)} \phi(x)$. That is, we pick $h(i)$ such that

$$\mathfrak{A}_i \models \bigwedge_{\phi \in \sigma(i)} \phi[h(i)].$$

I claim that $\prod_D \mathfrak{A}_i \models \Sigma[h_D]$. Pick any $\theta \in \Sigma$ and consider $i \in g(\{\theta\}) \in D$. We know that $\theta \in \sigma(i)$ so $\mathfrak{A}_i \models \theta[h(i)]$. Thus, since $g(\{\theta\}) \in D$ and $g(\{\theta\}) \subset \{i \in I : \mathfrak{A}_i \models \theta[h(i)]\}$, we have by upward closure that $\{i \in I : \mathfrak{A}_i \models \theta[h(i)]\} \in D$. Theorem 4.6 entails that $\prod_D \mathfrak{A}_i \models \theta[h_D]$. This holds for any $\theta \in \Sigma$ so that $\prod_D \mathfrak{A}_i \models \Sigma[h_D]$ and $\Sigma(x)$ is realized. □

Corollary 6.10. *Let α be an infinite cardinal and let D be a countably incomplete α-good ultrafilter over a set I. Let $||\mathscr{L}|| < \alpha$. Then for any family of $\mathfrak{A}_i, i \in I$, of $\mathscr{L}$-models, the ultraproduct is α-saturated.*

Proof. Let $X \subseteq \prod_{i \in I} A_i$ such that $|X| < \alpha$. Assume $\Sigma(x)$ is a type in the language $\mathscr{L} \cup \{c_a : a \in X\}$ that is consistent with $(\prod_D \mathfrak{A}_i, a)_{a \in X}$. We must show that $\Sigma(x)$ is realized in $(\mathfrak{A}, a)_{a \in X}$. Let $\mathscr{L}' = \mathscr{L} \cup \{c_a : a \in X\}$. Then $||\mathscr{L}'|| = ||\mathscr{L}|| + |X| = \max(||\mathscr{L}||, |X|) < \alpha$. Also note that

$$(\prod_D \mathfrak{A}_i, a)_{a \in X} = \prod_D (\mathfrak{A}_i, a(i))_{a \in X}.$$

Since $\Sigma(x)$ is consistent with $(\prod_D \mathfrak{A}_i, a)_{a \in X}$, every finite subset of $\Sigma(x)$ is consistent with $(\prod_D \mathfrak{A}_i, a)_{a \in X}$ and therefore every finite subset of $\Sigma(x)$ is realized by Proposition 5.5. Considering $(\prod_D \mathfrak{A}_i, a)_{a \in X}$ as the ultraproduct of a family of $\mathscr{L}'$ structures where $||\mathscr{L}'|| < \alpha$, Theorem 6.9 entails that $\Sigma(x)$ is realized in $(\prod_D \mathfrak{A}_i, a)_{a \in X}$. This completes the proof. □

Before proving the main theorem, I state a follow-up question I have:

Question 6.11. *Are there any other kinds of ultrafilters that build α-saturated ultraproducts, other than good ultrafilters?*

Due to the scope of this problem and time constraints, I have not been able to explore this question thoroughly, or to determine if it has been answered already. The general question being asked is what are the necessary, rather than sufficient, conditions under which an ultraproduct is saturated at some given cardinality. On this note, there is a result of Keisler that shows some theories actually *require* a good ultrafilter to construct saturated ultrapowers [3].

To prove the main theorem, we need the following lemma about how the saturation of a model affects its size.

Lemma 6.12. *Let $\mathfrak{A}$ be an $\mathscr{L}$-model with an infinite universe, and α an infinite cardinal. If $\mathfrak{A}$ is α-saturated, then $|\mathfrak{A}| \geq \alpha$.*

Proof. Suppose for contradiction that $|\mathfrak{A}| = \beta < \alpha$. Expand the language to $\mathscr{L} \cup \{c_a\}_{a\in A}$, and expand $\mathfrak{A}$ such that a interprets c_a. Consider the following type:

$$\Sigma(v) = \{v \neq c_a : a \in A\}$$

I claim $\Sigma(v)$ is consistent with $(\mathfrak{A}, a)_{a\in A}$. To show this, one must show that some model of $Th(\mathfrak{A}, a)_{a\in A}$ realizes $\Sigma(v)$. Expand the language again to include an additional constant d. Then consider the following set of sentences in $\mathscr{L} \cup \{c_a\}_{a\in A} \cup \{d\}$:

$$Th(\mathfrak{A}, a)_{a\in A} \cup \Sigma(d)$$

such that d is substituted for v in the type. We know every finite subset of the above set of sentences is modeled in the extended language, namely by $(\mathfrak{A}, a)_{a\in A}$ where d is interpreted by some element not specified by a the finite subset of $\Sigma(d)$. This element exists for any finite subset of $\Sigma(d)$ as A is infinite. Thus, by compactness, there is a model $\mathfrak{B}$ for $Th(\mathfrak{A}, a)_{a\in A} \cup \Sigma(d)$ in the extended language. Considering $\mathfrak{B}$ in $\mathscr{L} \cup \{c_a\}_{a\in A}$ by forgetting about d and its interpretation, $\mathfrak{B}$ continues to model $Th(\mathfrak{A}, a)_{a\in A}$ and there is still an element of B that realizes $\Sigma(v)$. Thus $\Sigma(v)$ is consistent with $(\mathfrak{A}, a)_{a\in A}$, and by α-saturation, must be realized in the model. But then there is an element in the universe not equal to any element in the universe. This is a contradiction. Thus $|\mathfrak{A}| \geq \alpha$. □

Theorem 6.13. *(The Main Theorem) Let $||\mathscr{L}|| \leq \alpha$ and $\mathfrak{A}, \mathfrak{B}$ be $\mathscr{L}$-models with $|A|, |B| \leq \alpha^+$. Assume that $2^\alpha = \alpha^+$. Let D be a countably incomplete α^+-good ultrafilter over a set of power α. Then the following are equivalent:*

(1) $\mathfrak{A} \equiv \mathfrak{B}$
(2) $\prod_D \mathfrak{A} \cong \prod_D \mathfrak{B}$

Proof. To prove (1) $\to$ (2), we use the uniqueness of saturated models. Thus we want to show that the ultrapowers are elementarily equivalent, are saturated, and have the same power. Assume $\mathfrak{A} \equiv \mathfrak{B}$. By Corollary 4.7, we know that $\prod_D \mathfrak{A} \equiv \mathfrak{A} \equiv \mathfrak{B} \equiv \prod_D \mathfrak{B}$, since an ultrapower's theory is entirely determined by the theory of its underlying model. Further, given that D is a countably incomplete α^+-good ultrafilter with $||\mathscr{L}|| < \alpha^+$, Corollary 6.10 gives us that $\prod_D \mathfrak{A}$ and $\prod_D \mathfrak{B}$ are α^+-saturated.

Since $|A|, |B| \leq \alpha^+$, we know $|\prod_D \mathfrak{A}|, |\prod_E \mathfrak{B}| \leq (\alpha^+)^\alpha = 2^{\alpha\cdot\alpha} = \alpha^+$ since we assumed GCH. If either $\prod_D \mathfrak{A}$ or $\prod_D \mathfrak{B}$ is finite, then they must both be finite. If not, only the infinite model would satisfy a sentence describing the existence of n distinct elements for each finite n, and this contradicts the models being elementarily equivalent. Thus, if either model is finite, then both are finite. Finite elementarily equivalent models are always isomorphic by Corollary 5.10 and the proof is complete. Thus, assume both $\prod_D \mathfrak{A}$ and $\prod_D \mathfrak{B}$ are infinite.

By Lemma 6.12, $|\prod_D \mathfrak{A}|$, $|\prod_D \mathfrak{B}| \geq \alpha^+$ and so $|\prod_D \mathfrak{A}| = |\prod_D \mathfrak{B}| = \alpha^+$. It follows from the uniqueness of saturated models, Theorem 5.9, that $\prod_D \mathfrak{A} \cong \prod_D \mathfrak{B}$.

(2) $\to$ (1) follows immediately by Corollary 4.7 and the transitivity of elementary equivalence. □

Whether Theorem 6.13 can be proven without GCH remains an open question. However, we do know the following result holds.

Theorem 6.14. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be models for $\mathscr{L}$. Then $\mathfrak{A} \equiv \mathfrak{B}$ if and only if they have isomorphic ultrapowers.*

There is additional machinery required to prove this result (see Section 6.1 in [1]), and such mathematics is beyond the scope of this paper.

Acknowledgments. I would like to thank my mentor Jonathan Rubin for his invaluable assistance in understanding the above material and for providing thorough lessons in the communication of mathematics. I would also like to thank Professor Maryanthe Malliaris for her guidance and insightful commentary. Lastly, I thank Professor Peter May for the opportunity to spend time grappling with these interesting constructions in such an inspiring and welcoming environment.

References

[1] C.C Chang and H. Jerome Keisler. Model Theory. Dover Publications, 1990.
[2] Kenneth Kunen. The Foundations of Mathematics. 2007.
[3] Jerome Keisler. *Ultraproducts which are not Saturated. The Journal of Symbolic Logic.* 32 (1967), 23-46.