EXCURSIONS IN MODEL THEORY

RAFAEL WINGESTER RIBEIRO DE OLIVEIRA

Abstract. This paper aims to introduce the reader familiar with undergraduate level logic to some fundamental constructions in Model Theory. A introductory course covering the basics of logic and including the completeness theorem, alongside with the fundamentals of ultraproducts, should serve the reader as a solid enough background. The paper follows the exposition style of Chang and Keisler's classic very closely, with minor alterations and the occasional input from Hodges' treatise on the use of games to build models. The paper begins with emphasis on applications of the Henkin construction in the first few sections, passes through the back and forth construction, and ends with a brief introduction to Skolem functions.

Contents

1. Introduction

Rather than presenting a set of tools and combining them to prove a central result, this paper aims to provide an overview of some key model-theoretic ideas and constructions that, although accessible, might not have been presented in an undergraduate introduction to Model Theory. Topics in this paper will include different ways of using the Henkin constrcution, the back and forth constrcution, combinatorial results using ultrafilters, and basic results on indiscernibles and skolem functions. While the next section is mostly concerned with reviewing the Henkin construction, which is perhaps the most central idea of this piece, we use the next section to restate some basic results and definitions that we assume are in the reader's background along with Henkin's Completeness proof (which in fact requires some of them) but might warrant recollection.

Date: August 2016.

2. Recalling Some Key Ideas

We begin with the definition of elementary extensions and elementary embeddings:

Definition 2.1. Let $\mathfrak{A}$ and $\mathfrak{B}$ be models. Then $\mathfrak{B}$ is said to be an *elementary extension* of $\mathfrak{A}$, denoted $\mathfrak{A} \prec \mathfrak{B}$ if

(1) $\mathfrak{B}$ is an extension of $\mathfrak{A}$, that is, $\mathfrak{A} \subset \mathfrak{B}$.
(2) For any formula $\phi(x_1, \ldots, x_n)$ of $\mathcal{L}$, and any $a_1, \ldots, a_n$ in A, the tuple $(a_1, \ldots, a_n)$ satisfies ϕ in $\mathfrak{A}$ if and only if it satisfies ϕ in $\mathfrak{B}$.

Another way of denoting this is by saying $\mathfrak{A}$ in an *elementary submodel* of $\mathfrak{B}$

Definition 2.2. A function $f : A \to B$ is an *elementary embedding* of $\mathfrak{A}$ into $\mathfrak{B}$, denoted $f : \mathfrak{A} \prec \mathfrak{B}$, if, for all formulas $\phi(x_1, \ldots, x_n)$ of the language and n-tuples $(a_1, \ldots, a_n) \in A$, the following holds:

$$\mathfrak{A} \models \phi[a_1, \ldots, a_n] \text{ if and only if } \mathfrak{B} \models \phi[f(a_1), \ldots, f(a_n)].$$

When there exists an elementary embedding from the model $\mathfrak{A}$ to the model $\mathfrak{B}$, we denote it by $\mathfrak{A} \precsim \mathfrak{B}$.

On a similar note, we also define the elementary diagram of a model as follows:

Definition 2.3. Let $\mathcal{L}_A = \mathcal{L} \cup \{c_a : a \in A\}$. The *elementary diagram* of $\mathfrak{A}$ is the set of all sentences of $\mathcal{L}_A$ which are true in the model $\mathfrak{A}_A = (\mathfrak{A}, a)_{a \in A}$. The diagram of $\mathfrak{A}$ is defined in similar fashion as the set of all atomic and negative atomic sentences of $\mathcal{L}_A$ which are true in $\mathfrak{A}_A$.

It can also be useful to know a few results regarding normal and special forms to get to convenient shortcuts in proofs. While Chang and Keisler [1] has a very useful general result in page 50 (T-equivalence), we find stating the following result as found in Ebbinghaus, Flum, and Thomas [4] more convenient for the purposes of this paper.

Definition 2.4. A formula ψ is said to be in *prenex normal form* if it is written as a string of quantifiers Q_n followed by a quantifier-free formula θ, that is, $\psi \equiv Q_1 Q_2 Q_3 \ldots Q_n \theta$.

Theorem 2.5. *Every formula ϕ is logically equivalent to a formula ψ in prenex normal form with $free(\phi) = free(\psi)$.*

Crucial to following the proofs presented here is to recall the Generalization and Generalization on Constants theorems. As an optional shortcut, one might also find helpful to recollect the Deduction Theorem. We state them here as found in pages 116-123 from Enderton [3].

Theorem 2.6 (Generalization Theorem)**.** *If $\Gamma \vdash \phi$ and x do not occur free in any formula in Γ, then $\Gamma \vdash (\forall x)\phi$.*

Theorem 2.7. (Generalization on Constants) *Assume that $\Gamma \vdash \phi$ and that c is a constant that does not occur in Γ. Then there is a variable y (which does not occur in ϕ) such that $\Gamma \vdash (\forall y)\phi^c_y$. Furthermore, there is a deduction of $(\forall y)\phi^c_y$ from Γ in which c does not occur.*

Theorem 2.8 (Deduction Theorem)**.** *If $\Gamma; \gamma \vdash \phi$, then $\Gamma \vdash (\gamma \to \phi)$.*

Finally, it is worthwhile to restate an essential result that will appear in the proof of Theorem (6.4), which will use the back and forth construction.

Proposition 2.9. *Let two models $\mathfrak{A}$ and $\mathfrak{B}$ be given. If $\mathfrak{A} \equiv \mathfrak{B}$ and $\mathfrak{A}$ is finite, then $\mathfrak{A} \cong \mathfrak{B}$.*

And as a last comment:

Definition 2.10. Let a set of formulas $\Sigma(x_1, \dots, x_n)$ of $\mathcal{L}$ be given. A formula $\phi(x_1, \dots, x_n)$ is said to be a *consequence* of Σ, denoted by $\Sigma \models \phi$, if for every model $\mathfrak{A}$ and every n-tuple $(a_1, \dots, a_n)$ of elements from the domain A, if $(a_1, \dots, a_n)$ satisfies Σ, then it satisfies ϕ.

3. Review: The Henkin Construction and the Completeness Theorem

We now begin with the review of Henkin's standard "witness" construction to prove the Completeness Theorem. Since this construction already belongs in the typical undergraduate logic curriculum, we present it here as a review section. The reader who is familiar with this particular construction may just quickly gloss over this part of the paper and move on to section 3. Many of the subsequent proofs will follow a similar method of building models.

Definition 3.1. Let T be a set of sentences of $\mathcal{L}$ and let C be a set of constant symbols of $\mathcal{L}$. C may or may not be the set of all constants in the language. We say that C is a *set of witnesses* of T in $\mathcal{L}$ if and only if for every formula ϕ of the entire language $\mathcal{L}$ in one free variable x there is a constant $c \in C$ such that

$$T \vdash (\exists x)\phi \to \phi(c)$$

We say that T has witnesses in $\mathcal{L}$ if T has some set C of witnesses in $\mathcal{L}$

Lemma 3.2. *Let T be a consistent set of sentences of $\mathcal{L}$. Let C be a set of constants not in $\mathcal{L}$ and the power of C be $|C| = ||\mathcal{L}||$, and let $\overline{\mathcal{L}} = \mathcal{L} \cup C$ be the simple expansion of $\mathcal{L}$ formed by adding C. Then T can be extended to a consistent set of sentences $\overline{T}$ in $\overline{\mathcal{L}}$ which has C as a set of witnesses in $\overline{\mathcal{L}}$.*

Proof. We consider $\kappa = ||\mathcal{L}||$ and enumerate the elements of C by κ in a transfinite sequence $\{c_\beta : \beta < \kappa\}$ such that no single element repeats. Since $\kappa = ||\mathcal{L}|| = |C|$, we naturally have that $\kappa = ||\overline{\mathcal{L}}||$, so that we may also enumerate the formulas in at most one free variable of $\overline{\mathcal{L}}$ by κ in the same fashion. The next step is to define an increasing (under inclusion) sequence of sets of sentences of $\overline{\mathcal{L}}$:

$$T = T_0 \subset T_1 \subset T_2 \subset T_3 ... \subset T_\xi \subset ..., \qquad \text{for } \xi < \kappa$$

and a sequence of constants d_ξ, for $\xi < \kappa$, from C such that:

(i) each T_ξ is consistent in $\overline{\mathcal{L}}$;
(ii) if ξ is a successor ordinal with $\xi = \zeta + 1$, put $T_\xi = T_\zeta \cup \{(\exists x_\xi)\phi_\xi \to \phi_\xi(d_\xi)\}$; x_ξ is the free variable in ϕ_ξ if it has one, otherwise $x_\xi = \nu_0$.
(iii) if ξ is a non-zero limit ordinal, then put $T_\xi = \bigcup_{\zeta<\xi} T_\zeta$.

Suppose that T_ζ has been defined. We remark that by our choice of enumeration by a cardinal we always have that the number of sentences (cardinality-wise) in T_ξ is strictly less than κ. Also given that each sentence contains at most a finite number of constants from C, we can let d_ξ be the first element of C which has not yet occurred in T_ξ. For example, since the original set $T = T_0$ does not contain any sentence with constants from C, then $d_0 = c_0$ from our original enumeration of C. We just have to see that $T_{\zeta+1} \cup (\exists x_\zeta)\phi_\zeta \to \phi_\xi(d_\zeta)$ is consistent given any ordinal $\zeta < \kappa$, which is indeed the case, for if it were not, then it must be that:

$$T_\zeta \vdash \neg(\exists x_\zeta)\phi_\zeta \to \phi_\xi(d_\zeta).$$

Which by the rules of propositional logic implies:

$$T_\zeta \vdash (\exists x_\zeta)\phi_\zeta \wedge \neg\phi_\xi(d_\zeta).$$

And because the constant d_ζ by choice does not occur in T_ζ, predicate logic gives by generalizing on constants (2.7) that:

$$T_\zeta \vdash (\forall x_\zeta)((\exists x_\zeta)\phi_\zeta \wedge \neg\phi_\zeta(x_\zeta))$$

which is basically

$$T_\zeta \vdash (\exists x_\zeta)\phi_\zeta \wedge \neg(\exists x_\zeta)\phi_\xi(x_\zeta)$$

thus giving an obvious problem with the consistency of T_ζ. The limit ordinal case is very similar but instead one should consider the union of the chain of preceding consistent theories indexed by smaller ordinals. Given the finiteness of deductions, one can just pick the least ordinal less than ζ, for which all ingredients for the contradiction proof show up, and we have found an inconsistent precedent countering the assumption. This finishes the induction.

Now let $\overline{T} = \bigcup_{\xi<\kappa} T_\xi$. We are granted by the inductive procedure that it is both consistent and an extension of the original T. Take any $\overline{\mathcal{L}}$ formula ϕ with at most one free variable. Naturally $\phi = \phi_\xi$ and $x = x_\xi$ for some $\xi < \kappa$. Therefore the sentence

$$(\exists x_\xi)\phi_\xi \to \phi_\xi(d_\xi)$$

is in $T_{\xi+1}$ and so in $\overline{T}$. ⊣

We now want to somehow build a model out of the set of constants C we added to the language $\mathcal{L}$. To do that, we start with a natural choice of equivalence relation and possibly an enlargement of T to ensure maximal consistency. Albeit simple, the procedure is very involved, as we see in the proof of the following lemma:

Lemma 3.3. *Let T be a consistent set of sentences and C be a set of witnesses for T in $\mathcal{L}$. Then T has a model $\mathfrak{A}$ such that every element of $\mathfrak{A}$ is an interpretation of a constant $c \in C$*

Proof. We begin with the remark that if T has witnesses C, then C is also a set of witnesses for any extension of T, given that the witness sentences refer to all possible formulas in at most one free variable in the entire language L and the things that T deduces are still deduced by any extension of it. So we just assume that T is maximal consistent, given that we can extend any consistent set of sentences to a maximal consistent one (Lindenbaum's Theorem). Now consider the following equivalence relation:

$$c \sim d \quad \text{iff} \quad (c \equiv d) \in T$$

Note that this defines an equivalence relation because:

(i) $c \sim c$
(ii) $c \sim d$ and $d \sim e$ together imply $c \sim e$
(iii) $c \sim d$ implies $d \sim c$

For each c in C we have the equivalence class $\tilde{c} = \{d \in C : d \sim c\}$. We use these equivalence classes to build our model $\mathfrak{A}$ by the following steps:

(1) We define: $A = \{\tilde{c} : c \in C\}$
(2) For any relation symbol P in $\mathcal{L}$ of n places we define an n-place relation R on the set of constants C so that:

$$R(c_1, c_2, ..., c_n) \text{ if and only if } (P(c_1, c_2, ..., c_n)) \in T$$

By our axioms of identity we are granted that:

$$\vdash P(c_1, c_2, ..., c_n) \wedge c_1 \equiv d_1 \wedge ... \wedge c_n \equiv d_n \rightarrow P(d_1, ..., d_n)$$

So now we can unambiguously interpret P in A by

$$P(\tilde{c_1}, ..., \tilde{c_n}) \text{ if and only if } P(c_1, ..., c_n)$$

independently of the choice of representative. We then found an interpretation for P in our new model $\mathfrak{A}$.

(3) We look at any constant symbol e of $\mathcal{L}$. Predicate logic assures us that:

$$\vdash (\exists v_0)(e \equiv v_0)$$

Hence we have that T certainly includes it, and since T has witnesses there is a constant c in C so that

$$(e \equiv c) \in T$$

The constant c itself may not be unique, but its equivalence class certainly is, because our identity axioms say:

$$\vdash ((e \equiv \bar{c} \wedge e \equiv c) \rightarrow (\bar{c} \equiv c))$$

So in $\mathfrak{A}$ each constant e is uniquely interpreted by the equivalence class of its witness constant. In particular, constants from C are naturally interpreted as their own equivalence classes since $c \equiv c \in T$, clearly.

(4) Now we look at function symbols of $\mathcal{L}$. Take any such n-placed function symbol F in it and we observe that naturally

$$\vdash (\exists v_0)(F(c_1, \ldots, c_n) \equiv v_0)$$

by propositional logic. But since we have witnesses for T, we have

$$(F(c_1, ..., c_n) \equiv c) \in T$$

for some c in C. Uniqueness is then guaranteed by the identity axioms in the same fashion as in the relation symbol procedure by

$$\vdash ((F(c_1, \ldots, c_n) \equiv c) \wedge c_1 \equiv d_1 \wedge \ldots \wedge c_n \equiv d_n \wedge c \equiv d) \rightarrow (F(d_1, \ldots, d_n) \equiv d)$$

So in this fashion we also establish, just like in the relational case, that F can be interpreted uniquely in the equivalence classes by making any choice of representatives in C and asking T through its witnesses to which constant (class) to send the n-tuple. Therefore:

$$F(\tilde{c_1}, ..., \tilde{c_n}) \equiv c \text{ if and only if } (F(c_1, ..., c_n) \equiv c) \in T$$

Now our next items present us with the procedure for showing $\mathfrak{A}$ indeed models T.

(5) Using the function interpretation rule as the first step of an induction on the complexity of terms, we take any term t with no free variable and for any c we end up with

$$\mathfrak{A} \models t \equiv c \text{ if and only if } (t \equiv c) \in T$$

(6) Since C is a set of witnesses, we use the above to see that given any two terms t_1 and t_2 with no free variables

$$\mathfrak{A} \models t_1 \equiv t_2 \text{ if and only if } (t_1 \equiv t_2) \in T$$

(7) Now we proceed to any atomic formula $P(t_1, ..., t_n)$ on terms with no free variables

$$\mathfrak{A} \models P(t_1, ..., t_n) \text{ if and only if } P(t_1, ..., t_n) \in T.$$

(8) Combining items (6) and (7), we can prove that for any sentence ϕ in $\mathcal{L}$:

$$\mathfrak{A} \models \phi \text{ if and only if } \phi \in T.$$

This can be proven by induction on the complexity of formulas. For boolean combinations, the atomic case is ensured by (6) and (7). By maximal consistency and truth in a model, for formulas ϕ and ψ with no quantifiers, we have that:

$$\mathfrak{A} \models \neg\phi \text{ if and only if } \neg\phi \in T$$

and

$$\mathfrak{A} \models \phi \wedge \psi \text{ if and only if } \phi \wedge \psi \in T.$$

Let $\phi \equiv (\exists x)\psi$. On the first step, let ψ be quantifier-free, but then we can build up from it. If $\mathfrak{A} \models \phi$ then for some $\tilde{c}$ in the domain A of $\mathfrak{A}$, it must be that $\mathfrak{A} \models \psi[\tilde{c}]$. But by picking a any representative of the class $\tilde{c}$ and replacing all free occurrences of x in ψ by it, we have then that $\mathfrak{A} \models \psi(c)$. So $\psi(c) \in T$ and since the rules of deduction give:

$$\vdash \psi(c) \rightarrow (\exists x)\psi(x),$$

we have that ϕ is in T. On the other hand, if $\phi \in T$ with $\phi \equiv (\exists x)\psi$, then, because T has witnesses, there is a constant $c \in C$ so that

$$T \vdash (\exists x)\psi \rightarrow \psi(c)$$

So because T is maximal consistent, $\psi(c) \in T$, then $\mathfrak{A} \models \psi(c)$. This then gives $\mathfrak{A} \models \psi[\tilde{c}]$. So $\mathfrak{A} \models \phi$. So we can now mix this quantifier case with the boolean combinations case and get all sentences of T to be exactly $\mathrm{Th}(\mathfrak{A})$.

$\dashv$

Remark 3.4. As a passing remark, a converse of the previous lemma can be very easily proved. If we have a model in which every element is the interpretation of some constant, then its theory is a maximal consistent set with witnesses. They form the set of representatives for elements in the said model.

Without going too much into details, given the review nature of this section, we now state four very important results following from these past two lemmas:

Theorem 3.5 (Completeness)**.** *Let Σ be a set of sentences. Then Σ is consistent if and only if it has a model.*

Proof. If Σ has a model $\mathfrak{M}$, its consistency is easy to show by the definition of truth in a model. The $\models$ sign is defined by induction on the complexity of formulas to give always either true or false values, never both. Note also that the negation sign $\neg$ always changes true and false values, so for any formula ϕ and fixed tuple if there are free variables, exactly one of ϕ and $\neg\phi$ is in the set of formulas (including sentences) modeled by the assumed model $\mathfrak{M}$ of Σ. Thus we can by induction on the length of deductions, using each of the rules of deduction in a text of choice (Chang and Keisler [1] or Enderton [3], for example), see that, if Σ is in $\mathrm{Th}(\mathfrak{M})$, then $\mathrm{Cn}T \subset \mathrm{Th}(\mathfrak{M})$. So $\mathrm{Cn}T$ cannot be inconsistent. So let Σ in $\mathcal{L}$ be consistent. We can use Lemma 3.2 to expand Σ to $\overline{\Sigma}$ in an equally expanded language $\overline{\mathcal{L}}$, so that $\overline{\Sigma}$ has witnesses in $\overline{\mathcal{L}}$. By Lemma 3.3 we have a model $\mathfrak{A}$ for $\overline{\Sigma}$ in $\overline{\mathcal{L}}$. Now we just consider $\mathfrak{A}$ as it is but 'forgetting' all added constant symbols, i.e. coming back to $\mathcal{L}$. We call it now $\overline{\mathfrak{A}}$, and it clearly models Σ as an $\mathcal{L}$-structure once Σ does not mention the added witness constants. ⊣

Theorem 3.6 (Compactness)**.** *Let Σ be a set of sentences. Then Σ has a model if and only if every finite subset of Σ has a model.*

Proof. Σ is consistent if and only if every finite subset of Σ is consistent. ⊣

Theorem 3.7 (Downward Löwenheim-Skolem-Tarski)**.** *Every consistent set of sentences Σ in $\mathcal{L}$ has a model of size at most $||\mathcal{L}||$.*

Proof. The construction carried out in Lemmas 3.2 and 3.3 assures it. ⊣

Theorem 3.8 (Upward Löwenheim-Skolem-Tarski)**.** *If a set of sentences Σ in $\mathcal{L}$ has some infinite model, then it has models of any given power $\kappa \geqslant ||\mathcal{L}||$.*

Sketch: Application of Compactness and the Downward version. ⊣

4. The Omitting Types Theorem

In this section we will concern ourselves with how certain sets of sentences can 'fail' or certain kinds of elements not exist in some particular model. To prove the main result, however, we need to give a few definitions and establish some facts first.

Definition 4.1.
(1) A set of formulas Σ in a language $\mathcal{L}$ is said to *be (free) in the variables* $x_1, ..., x_n$ if and only if $x_1, ..., x_n$ are distinct individual variables and every formula σ in Σ contains at most the variables $x_1, ..., x_n$ free. We introduce the notation $\Sigma(x_1, ..., x_n)$.
(2) $\mathfrak{A} \models \sigma[a_1, ..., a_n]$ if and only if the sequence $a_1, ..., a_n$ of elements in the domain A satisfies σ in $\mathfrak{A}$.
(3) $\mathfrak{A} \models \Sigma[a_1, ..., a_n]$ if and only if for every $\sigma \in \Sigma$ the tuple $a_1, ..., a_n$ satisfies σ in $\mathfrak{A}$. In this case, we say the tuple $a_1, ..., a_n$ *satisfies, or realizes* Σ in $\mathfrak{A}$. One can more plainly say Σ is *satisfiable in*, or *realized by* $\mathfrak{A}$ (Σ is consistent if and only if satisfiable in some model).
(4) We shall say that a formula $\sigma(x_1, \ldots, x_n)$ *is consistent with* T if there is some model $\mathfrak{A}$ of T such that σ is realized by it. Equivalently, we say a set of formulas $\Sigma(x_1, \ldots, x_n)$ is consistent with the theory T if there is some model $\mathfrak{A}$ of T that realizes Σ.

(5) We say Σ is *omitted* in a model $\mathfrak{A}$ whenever it is not realized.
(6) By a type $\Gamma(x_1, ..., x_n)$ in the variables $x_1, ..., x_n$ we mean a maximal consistent set of formulas of $\mathcal{L}$ in these variables. Given any model $\mathfrak{A}$ and n-tuple $a_1, ..., a_n \in A$, the set $\Gamma(x_1, ..., x_n)$ of all formulas $\gamma(x_1, ..., x_n)$ satisfied by $a_1, ..., a_n$ is a type. In fact, this set is the unique type realized by $a_1, ..., a_n$ (maximal consistency). So Γ is called the type of $a_1, ..., a_n$ in $\mathfrak{A}$.

Example 4.2. We show that the model $\mathfrak{A}$ of the ordered field of real numbers has continuum many types. By the density of the rationals, for any two distinct real numbers $a < b$, we have distinct types, since one can with summing 1's, taking multiplicative inverses of the integral denominator, and multiplying by the summed integral numerator, obtain some rational number r between a and b. In one free variable, we would already have $x < r$ satisfied by a and $\neg(x < r)$ satisfied by b. Thus $\mathfrak{A}$ realizes 2^ω many different types in one variable (each element satisfying a unique type).

Proposition 4.3. *Let T be a theory and let $\Sigma = \Sigma(x_1, ..., x_n)$. Then the following are equivalent:*
(1) T has a model which realizes Σ.
(2) Every finite subset of Σ is realized in some model of T.
(3) $T \cup \{(\exists x_1, ..., x_n)(\sigma_1 \wedge ... \wedge \sigma_m) : m < \omega, \sigma_1, ..., \sigma_n \in \Sigma\}$ is consistent.

Proof. Done via direct applications of the compactness theorem in directions (1) to (2), (2) to (3), and (3) to (1). $\dashv$

Remark 4.4. As a consequence of this proposition, all of the statements $(1) - (3)$ are equivalent to the definition of σ being consistent with T.

We now take up the question of how to omit some set of formulas in a model of a given theory. When Σ is finite, we can combine all of its elements to make a new formula to 'summarize' the type. Then add the necessary existential quantifiers to make it a sentence, and, finally, take the negation of that and see consistency with the theory. What we really want, though, is to analyze the case where Σ is infinite. For that, we introduce another crucial notion to get to the Omitting Types Theorem.

Definition 4.5. Let $\Sigma(x_1, \ldots, x_n)$ be a set of $\mathcal{L}$-formulas. A theory T in $\mathcal{L}$ is said to *locally realize* Σ if there is a formula ϕ of $\mathcal{L}$ such that:
(1) The formula ϕ is consistent with T.
(2) For all $\sigma \in \Sigma$, we have $T \models \phi \to \sigma$.

A way of seeing this is that any n-tuple which satisfies ϕ in a model of T must also realize Σ. The opposite of this notion is:

Definition 4.6. We say T *locally omits* Σ if T does not locally realize Σ, meaning that for every formula $\phi(x_1, \ldots x_n)$ which is consistent with T, there exists some $\sigma \in \Sigma$ such that $\phi \wedge \neg\sigma$ is consistent with T.

Proposition 4.7. *Let T be a complete theory in $\mathcal{L}$, and let $\Sigma(x_1, \ldots, x_n) \equiv \Sigma$ be a set of formulas of $\mathcal{L}$. If T has a model that omits Σ, then T locally omits Σ.*

Proof. We take the contrapositive of the proposition statement. That is, we have to prove that if T locally realizes Σ, then every model of T realizes Σ. So let

$\phi(x_1, \ldots, x_n)$ be a formula consistent with T such that $T \models \phi \to \sigma$ for every $\sigma \in \Sigma$. Let $\mathfrak{A}$ be a model of T. Since T is complete, $T \models (\exists x_1 \ldots x_n)\phi(x_1, \ldots, x_n)$. So some n-tuple $a_1, \ldots, a_n$ satisfies ϕ in $\mathfrak{A}$ and thus $\mathfrak{A}$ realizes Σ. ⊣

We now proceed to the main theorem in this section. It is a converse to the proposition we just proved. In fact, it holds for any consistent set of formulas in a countable language.

Theorem 4.8 (Omitting Types)**.** *Let T be a consistent theory in a countable language $\mathcal{L}$, and let $\Sigma(x_1, \ldots x_n)$ be a set of formulas in that same language. If T locally omits Σ, then T has a countable model which omits Σ.*

Proof. For the sake of simplicity, we take $\Sigma(x)$ in one free variable (the proof goes very similarly with more than one free variable). Assume T locally omits Σ. Let $C = \{c_0, c_1, \ldots\}$ be a countable set of constants not in $\mathcal{L}$. Then we create an expanded language $\mathcal{L}' = \mathcal{L} \cup C$ which is still countable. Now we enumerate the sentences of $\mathcal{L}'$ in a sequence $\phi_0, \phi_1, \phi_2, \ldots$. We want to build an increasing sequence of consistent theories

$$T = T_0 \subset T_1 \subset T_2 \subset T_3 \ldots$$

such that the following are satisfied:

1. (1) Each T_m is a consistent theory of $\mathcal{L}'$, which is a finite extension of T.
2. (2) Either $\phi_m \in T_{m+1}$, or $(\neg\phi_m) \in T_{m+1}$.
3. (3) If $\phi_m \equiv (\exists x)\psi(x)$ and $\phi_m \in T_{m+1}$, then $\psi(c_p) \in T_{m+1}$ where c_p is the first constant not occurring in T_m or ϕ_m.
4. (4) There is a formula $\sigma(x) \in \Sigma(x)$ such that $(\neg\sigma(c_m)) \in T_{m+1}$.

Given T_m, we craft T_{m+1} by the following procedure: T_m can be decomposed into $T \cup \{\theta_1, \ldots, \theta_n\}$, with finite $n > 0$. Therefore, we can just combine all θ_k into one single big formula $\theta \equiv \theta_1 \wedge \ldots \wedge \theta_n$. Let $c_0, \ldots, c_m$ be the constants from C showing up in θ. We create formulas $\theta(x_r)$ replacing each c_i with the variable x_i, renaming bound variables from θ if necessary. We make each new formula free in one variable e by adding $\exists x_i$ for each x_i at the beginning for $i \not\equiv r$. We have that each $\theta(x_r)$ is consistent with T, and thus each $\theta(x_r) \wedge \neg\sigma(x_r)$ is also consistent with T for some $\sigma(x)$ in $\Sigma(x)$. We put the sentence $\neg\sigma(c_r)$ in T_{m+1}, and so we satisfy requirement (4).

If ϕ_m is consistent with $T_m \cup \neg\sigma(c_r)$, we add ϕ_m to T_{m+1}. Otherwise, we add $\neg\phi_m$. So (2) is satisfied. If $\phi_m \equiv (\exists x)\psi(x)$ is consistent with $T_m \cup \neg\sigma(c_r)$ we just add $\psi(c_p)$ to T_{m+1}. This takes care of (3). Since T_{m+1} constructed in this way is a consistent finite extension of T, it satisfies $(1) - (4)$.

Let $T_\omega = \bigcup_{n<w} T_n$. By (1) and (2) it is a maximal consistent set of sentences in $\mathcal{L}'$. Consider a given countable model $\overline{\mathfrak{B}} = (\mathfrak{B}, b_0, \ldots)$ of T_ω and take the submodel $\overline{\mathfrak{A}} = (\mathfrak{A}, b_0, \ldots)$ of $\overline{\mathfrak{B}}$ generated by the constants $b_0, b_1, \ldots$. By (3), the domain is $A = b_0, b_1, \ldots$. Furthermore, it follows from completeness of T_ω and (3) by induction of the complexity of sentences that the following are equivalent:

- (i) $\overline{\mathfrak{A}} \models \phi$
- (ii) $\overline{\mathfrak{B}} \models \phi$
- (iii) $T_\omega \models \phi$

Thus $\overline{\mathfrak{A}}$ is a model of T_ω and the reduct $\mathfrak{A}$ is a model of T. By (4), we are thus assured $\mathfrak{A}$ omits Σ. ⊣

We now state a generalization of this result whose proof is the same as the one just presented, save only for a minor alteration in the enumeration of constant tuples.

Theorem 4.9 (Extended Omitting Types Theorem)**.** *Let T be a consistent theory in a countable language $\mathcal{L}$, and let $\Sigma_r(x_1, \ldots x_n)$, for each $r < \omega$, be a set of formulas in that same language. Each of them is free on n_r variables. If T locally omits each Σ_r, then T has a countable model which omits all Σ_r simultaneously.*

Proof. The argument is the same as that for omitting types but the n_r-tuples of new constants are arranged in the following fashion for each r:

$$s^r_r, s^r_{r+1}, s^r_{r+2} \ldots$$

The theory T_m is built up so that for each $r \leq m$ there is a formula $\sigma \in \Sigma_r$ such that $\neg\sigma(s^r_m) \in T_{m+1}$. ⊣

As a concluding comment, it is noticeable that countability of the language and that of the number of Σ's were used in a crucial way in the previous proof. That said, it is worthwhile to bear in mind that, as of a first inspection, removing any of these two hypotheses could make the proof substantially harder.

5. The Craig Interpolation Theorem

This section will provide one more application of the method of constructing models from constants. The setting, however, may at first seem somewhat peculiar and surprising, as the statement of the main result is mostly concerned with syntatics instead of directly mentioning specific models. With no further delay, we proceed to said result.

Theorem 5.1 (Craig Interpolation Theorem)**.** *Let ϕ and ψ be sentences such that $\phi \models \psi$. Then there exists a sentence θ such that:*

(1) $\phi \models \theta$ and $\theta \models \psi$.
(2) except for maybe the identity, every relation, constant, or function symbol occurring in θ, also occurs in both ϕ and ψ.

Before we prove this result, we make a detour to provide some insights on the nature of θ. The sentence θ is called the *Craig interpolant* of ϕ and ψ. One may question the necessity of permitting the identity symbol in θ, regardless of ϕ and ψ. The cases that follow show why this may be necessary.

Examples 5.2. The following pairs of ϕ and ψ are such that $\phi \models \psi$ but have no Craig interpolant without the identity symbol (take M, N, and R to be unary relation symbols):

(1) ϕ is $(\exists x)(M(x) \wedge \neg M(x))$ and ψ is $(\exists x)N(x)$
(2) ϕ is $(\exists x)N(x)$ and ψ is $(\exists x)(M(x) \vee \neg M(x)$
(3) ϕ is $(\forall xy)(x \equiv y)$ and ψ is $(\forall xy)(R(x) \leftrightarrow R(y)$

We now move on to the proof of the theorem.

Proof of the Craig Interpolation Theorem: We assume for the sake of contradiction that there is no Craig interpolant θ of ϕ and ψ and deploy this assumption to build a model of $\phi \wedge \neg\psi$, and thus failing $\phi \models \psi$. The construction is at this point more familiar but some preliminary work most be done. We first begin by calling $\mathcal{L}^1$ the language consisting of the symbols in ϕ and $\mathcal{L}^2$ that consisting of the symbols in ψ. Furthermore, we set:

$$\mathcal{L}^0 = \mathcal{L}^1 \cap \mathcal{L}^2 \qquad \text{and} \qquad \mathcal{L} = \mathcal{L}^1 \cup \mathcal{L}^2$$

We now pick a countable set C of new constants constants and form the following expansions:

(a) $\overline{\mathcal{L}} = \mathcal{L} \cup C$
(b) $\overline{\mathcal{L}^0} = \mathcal{L}^0 \cup C$
(c) $\overline{\mathcal{L}^1} = \mathcal{L}^1 \cup C$
(d) $\overline{\mathcal{L}^2} = \mathcal{L}^2 \cup C$

For the remainder of the proof, it will be necessary to introduce the notion of *separability*. Consider a pair of theories T and V in $\overline{\mathcal{L}^1}$ and $\overline{\mathcal{L}^2}$, respectively. A sentence θ in $\overline{\mathcal{L}^0}$ is said to *separate* T and V if $T \models \theta$ and $V \models \neg\theta$. If no such θ exists, then T and V are said to be *inseparable*.

By assumption, there is no Craig interpolant of ϕ and ψ. Hence, $\{\phi\}$ and $\{\neg\psi\}$ are inseparable. This is the case because if there were $\theta(c_1, \ldots, c_n)$ separating them, then by completeness and generalizing on constants (2.7), for variables $x_1, \ldots, x_n$, the sentence $(\forall x_1, \ldots, x_n)\theta(x_1, \ldots, x_n)$ is a Craig interpolant.

Since formulas have finitely many symbols, and so the languages defined in this proof are all of countable power, we introduce an enumeration of $\overline{\mathcal{L}^1}$ and $\overline{\mathcal{L}^2}$:

($\overline{\mathcal{L}^1}$) $\phi_0, \phi_1, \phi_2, \phi_3, \phi_4, \ldots$
($\overline{\mathcal{L}^2}$) $\psi_0, \psi_1, \psi_2, \psi_3, \psi_4, \ldots$.

The aim is to construct two increasing (under inclusion) sequences of theories, in $\overline{\mathcal{L}^1}$ and $\overline{\mathcal{L}^2}$ respectively, with

$\{\phi\} = T_0 \subset T_1 \subset T_2 \subset T_3 \subset T_4 \ldots$ and
$\{\neg\psi\} = S_0 \subset S_1 \subset S_2 \subset S_3 \subset S_4 \ldots$

such that the following are satisfied:

(1) The theories T_m and S_m are inseparable finite sets of sentences .
(2) If $T_m \cup \{\phi_m\}$ and S_m are inseparable, then $\phi_m \in T_{m+1}$.
If T_{m+1} and $U_m \cup \{\psi_m\}$ are inseparable, then $\psi_m \in S_{m+1}$.
(3) If $\phi_m \equiv (\exists x)\sigma(x)$ and $\phi_m \in T_{m+1}$, then $\sigma(c) \in T_{m+1}$ for some $c \in C$.
If $\psi_m \equiv (\exists x)\eta(x)$ and $\psi_m \in S_{m+1}$, then $\eta(d) \in S_{m+1}$ for some $d \in C$.

So given some S_m and T_m the construction works in the intuitive manner as outlined in (1) – (3). For (3) we have to remember to only pick constants c and d that have not yet occurred in T_m, S_m, ϕ_m, or ψ_m. So inseparability is preserved. Take:

$$T_\omega = \bigcup_{n<\omega} T_m \quad \text{and} \quad S_\omega = \bigcup_{n<\omega} S_m.$$

Then T_ω and S_ω are naturally separable given completeness and the finite length of deductions. Also, each is consistent (for if any of them were inconsistent we can pick the negation of a logical axiom in the language of the other and make it the separating θ).

The aim now is to show that $T_\omega \cup S_\omega$ is consistent. To do that, we begin by showing that T_ω and S_ω are each maximal consistent in $\overline{\mathcal{L}^1}$ and $\overline{\mathcal{L}^1}$ respectively. To see this, suppose there is some ϕ_m such that $\phi_m \notin T_\omega$ and $\neg\phi_m \notin T_\omega$. Since by construction $T_m \cup \{\phi_m\}$ is separable from S_m, there exists θ in $\overline{\mathcal{L}^0}$ such that $T_\omega \models \phi_m \to \theta$ and $S_\omega \models \neg\theta$. The same argument, repeated for $\neg\phi_m$, gives the existence of some ν such that $T_\omega \models \neg\phi_m \to \nu$ and $S_\omega \models \neg\nu$. This leads to $T_\omega \models (\theta \vee \nu)$ and $S_\omega \models \neg(\theta \vee \nu)$. But this violates the inseparability of T_ω and S_ω. The same proof goes for maximal consistency of S_ω.

Now, it is on the agenda to show that that $T_\omega \cap S_\omega$ is a maximal consistent theory in $\overline{\mathcal{L}^0}$. Let some σ of $\overline{\mathcal{L}^0}$ be given. By maximal consistency of T_ω and S_ω in larger languages, either σ or its negation is in T_ω and the same applies for S_ω. Moreover, it cannot be the case that these theories disagree on σ, as that would give σ as their separator. This way, we are given maximal consistency of the intersection $T_\omega \cap S_\omega$ in the intersection of the languages.

Finally, this chunk of information is used to construct a model. We let $\overline{\mathfrak{D}} = (\mathfrak{D}, a_0, a_1, a_2, \ldots)$ be a model of T_ω. By (3) and maximal consistency, we see that the submodel $\overline{\mathfrak{A}} = (\mathfrak{A}, a_0, a_1, a_2, \ldots)$ of $\overline{\mathfrak{D}}$ with domain $A = \{a_0, a_1, a_2, \ldots\}$ is also a model of T_ω (to see this quickly one can imagine the formulas of T_ω in prenex normal form, as outlined in section 2, and proceed inductively on the number of quantifiers, or just induct plainly on the complexity of sentences). Along the same line of argument, it also possible to conclude that S_m has a model $\overline{\mathfrak{B}} = (\mathfrak{B}, b_0, b_1, b_2, \ldots)$ with domain $B = \{b_0, b_1, b_2, \ldots\}$. And by maximal consistency of $T_\omega \cap S_\omega$ in $\overline{\mathcal{L}^0}$ and the fact that all our elements are named by constants in C, we have that the $\overline{\mathcal{L}^0}$ reducts of $\overline{\mathfrak{A}}$ and $\overline{\mathfrak{B}}$ are isomorphic, with a_n corresponding to b_n. Thus, $\mathfrak{A}$ and $\mathfrak{B}$ have the same reduct to $\mathcal{L}^0$. Hence, "merge" $\mathfrak{A}$ and $\mathfrak{B}$ (that is, transfer to one the relation and function symbols of the other along with the interpretation on corresponding elements) to produce a model $\mathfrak{Z}$ in $\mathcal{L}$ whose reduct to $\mathcal{L}^1$ is $\mathfrak{A}$ and whose reduct to $\mathcal{L}^2$ is $\mathfrak{B}$. Since $\phi \in T_\omega$ and $(\neg\psi) \in S_\omega$, and both are in $\mathcal{L}$, then $\mathfrak{Z} \models \phi \wedge (\neg\psi)$. ⊣

There is a particularly interesting corollary to the Craig Interpolation Theorem that sheds light on how consistency and inconsistency cement the relationship between theories in distinct languages. As mentioned in Chang and Keisler's classic treatise [1], the original argument given by Robinson was a different one, but here we present the one which relies on Craig's interpolation result.

Theorem 5.3 (Robinson Consistency Theorem)**.** *Let $\mathcal{L}^1$ and $\mathcal{L}^2$ be two languages and let $\mathcal{L} = \mathcal{L}^1 \cap \mathcal{L}^2$. Suppose T is a complete theory in $\mathcal{L}$, and also that there exist consistent theories T_1 in $\mathcal{L}^1$ and T_2 in $\mathcal{L}^2$ with $T \subset T_1$ and $T \subset T_2$. Then, it must be that $T_1 \cup T_2$ is a consistent theory in $\mathcal{L}^1 \cup \mathcal{L}^2$.*

Proof. Assume $T_1 \cup T_2$ is inconsistent. Then there are finite subsets $\Gamma_1 \subset T_1$ and $\Gamma_2 \subset T_2$ so that $\Gamma_1 \cup \Gamma_2$ is inconsistent. Let η_1 be the conjunction of all elements in Γ_1 and η_2 the conjunction of all elements in Γ_2. It must be then that $\eta_1 \models \neg\eta_2$. By the Craig Interpolation Theorem, there exists θ such that $\eta_1 \models \theta$ and $\theta \models \eta_2$ and every function, relation, or constant symbol in θ (except maybe for identity) occurs in both η_1 and η_2. Thus, θ is a sentence in $\mathcal{L}$, since $\mathcal{L} = \mathcal{L}^1 \cap \mathcal{L}^2$. Observing T_1, it must be that $T_1 \models \theta$, and, because T_1 is consistent, $T_1 \not\models \neg\theta$, and so $T \not\models \neg\theta$. But it must also be that $T_2 \models \neg\theta$, hence $T_2 \not\models \theta$ and therefore $T \not\models \theta$. However, this contradicts the hypothesis of the theorem that T is a complete theory in $\mathcal{L}$. ⊣

6. Countable Models of Complete Theories in Countable Languages

The ideas of realizing and omitting types developed in the earlier sections seem to invite an investigation of the structure of models which may omit or realize many types. In this section, the aim is to present the reader with elementary notions regarding the back and forth construction and countable models of 'small' (atomic) and 'large' (countably saturated) varieties. We also discuss the basic characterizations of ω-categoricity and present a couple of interesting results. We use here the omitting types theorem. Furthermore, in this section, languages are assumed to be of countable power.

Definition 6.1. Given a complete theory T in $\mathcal{L}$, call a formula ϕ *complete (in T)* if for every formula $\psi(x_1 \dots x_n)$ exactly one of the following holds:

(1) $T \models \phi \to \psi$.
(2) $T \models \phi \to \neg\psi$.

Furthermore, a formula θ is said to be *completable (in T)* if there is a complete formula $\phi(x_1 \dots x_n)$ such that $T \models \phi \to \theta$. If θ is not completable, then it is said to be *incompletable*.

Definition 6.2. Call a theory T *atomic* if every $\mathcal{L}$-formula which is consistent with T is also completable in T. Additionally, a model $\mathfrak{A}$ is said to be an *atomic model* if every n-tuple in the domain A satisfies a complete formula in $Th(\mathfrak{A})$.

The next result is an existence one. It discusses the relationship between countable atomic models and atomic theories.

Theorem 6.3 (Existence of Atomic Models)**.** *Let T be a complete theory. Then T has a countable atomic model if and only if T is atomic.*

Proof. We first assume T has an atomic model $\mathfrak{A}$. Let $\phi(x_1, \dots, x_n)$ be consistent with T. Then, since T is complete, we have $T \models (\exists x_1, \dots, x_n)\phi(x_1, \dots, x_n)$. We let the tuple of elements $a_1, \dots, a_n$ in the domain A satisfy ϕ, and let $\psi(x_1, \dots, x_n)$ be a complete formula satisfied by $a_1, \dots, a_n$. Notice that $T \models \psi \to \neg\phi$ is therefore impossible, so it must be the case that $T \models \psi \to \phi$. Thus T is atomic and ϕ is completable.

For the other direction, suppose T is atomic. For each $n < \omega$, let $\Gamma_n(x_1, \dots, x_n)$ be the set of all negations of complete formulas $\psi(x_1, \dots, x_n) \in T$. We know by definition that every $\phi(x_1, \dots, x_n)$ which is consistent with T is completable and therefore $\phi \wedge (\neg\gamma)$ is consistent with T for some $\gamma \in \Gamma_n$. It is then the case that T locally omits each Γ_n. Therefore, by the extended omitting types theorem, the theory T has a countable model $\mathfrak{A}$ that omits each Γ_n. It is then the case that each n-tuple $(a_1, \dots, a_n)$ in the domain A satisfies a complete formula, thus yielding an atomic model. $\dashv$

It is also remarkable that whenever these atomic models of a given complete theory do exist they must actually be isomorphic. For the following, the back and forth construction will be very useful. As a visual guide and preview of the construction, the following diagram may be useful:

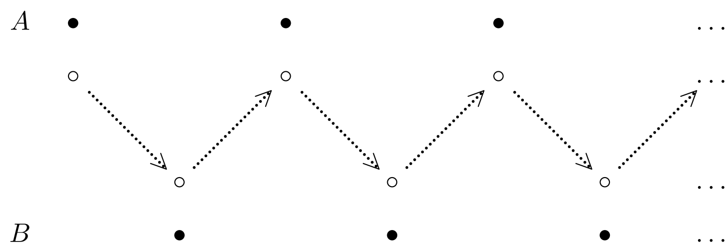

Theorem 6.4. (Uniqueness of Atomic Models) *If $\mathfrak{A}$ and $\mathfrak{B}$ are countable atomic models and $\mathfrak{A} \equiv \mathfrak{B}$, then $\mathfrak{A} \cong \mathfrak{B}$.*

Proof. If $\mathfrak{A}$ (and by elementary equivalence also $\mathfrak{B}$) is finite, then the isomorphism is clear by Proposition (2.9) in the second section. Let $\mathfrak{A}$ be infinite and we induce an enumeration of order type ω on both our countable domains. Now the back and forth construction begins.

Let a_0 be the first element in the ordering of A and we let $\psi_0(x_0)$ be some complete formula satisfied by a_0 in $\mathfrak{A}$. Given elementary equivalence and that $\mathfrak{A} \models (\exists x_0)\psi(x_0)$, we have $\mathfrak{B} \models (\exists x_0)\psi(x_0)$. So we choose $b_0 \in B$ that satisfies $\psi(x_0)$. Now place b_0 in the first position of a newly started ordered list of B. Now take the first element b_1 of $B \setminus \{b_0\}$ in the original ordering of type ω. Take $\psi_1(x_0, x_1)$ some complete formula satisfied by the pair (b_0, b_1) of elements in the domain of $\mathfrak{B}$. Given that ψ_0 is complete, it follows that both $\mathfrak{A}$ and $\mathfrak{B}$ satisfy

$$(\forall x_0)(\psi_0(x_0) \rightarrow (\exists x_1)\psi_1(x_0, x_1).$$

Thus there must exist $a_1 \in A$ such that (a_0, a_1) satisfies ψ_1. Place that element in a newly started ordered list with a_0. In the next step, let a_2 be the first element of $A \setminus \{a_0, a_1\}$ in the original ordering, run the procedure again to choose an element in B, and just move on with repeating the alternation on the ever shrinking orders. For example, the next step involves the following formula:

$$(\forall x_0 x_1)(\psi_1(x_0, x_1) \rightarrow (\exists x_2)\psi_2(x_0, x_1, x_2).$$

By going back and forth ω times, we form two sequences:

$$a_0, a_1, a_2, a_3, \ldots \quad \text{and} \quad b_0, b_1, b_2, b_3, \ldots$$

that list the entirety of A and B. Additionally, for each finite n, the n-tuples $a_0, \ldots, a_{n-1}$ and $b_0, \ldots, b_{n-1}$ satisfy the same complete formula by the procedure of the back and forth construction. Therefore, the map $a_n \mapsto b_n$ gives the desired isomorphism between $\mathfrak{A}$ and $\mathfrak{B}$. $\dashv$

Definition 6.5. A model $\mathfrak{A}$ is called *prime* if $\mathfrak{A}$ is elementarily embedded in every model of $Th(\mathfrak{A})$, that is, $\mathfrak{A}$ is isomorphic to a submodel of and (obviously) elemetarily equivalent to any other model of $Th(\mathfrak{A})$. Also, $\mathfrak{A}$ is said to be *countably prime* if it is elementarily embedded in every countable model of $Th(\mathfrak{A})$.

Theorem 6.6. *The following are equivalent:*

(1) $\mathfrak{A}$ is a countable atomic model.
(2) $\mathfrak{A}$ is prime.
(3) $\mathfrak{A}$ is countably prime.

Proof. We begin by assuming (1) and take the theory $T = Th(\mathfrak{A})$. We are now using a partial variation of the back and forth construction. Let $A = \{a_0, a_1, a_2, \ldots\}$ and let $\mathfrak{B}$ be any model of T. We take $\phi_0(x_0)$ to be some complete formula satisfied by the element a_0 of the domain. Therefore, we have that $T \models (\exists x)\phi_0(x)$. Thus, we may choose b_0 in the domain of $\mathfrak{B}$ such that it satisfies $\phi_0(x_0)$. Now we consider a complete formula $\phi_1(x_0, x_1)$ satisfied by the tuple (a_0, a_1). Given completeness of the formulas, we have that $T \models \phi_0(x_0) \to (\exists x_1)\phi_1(x_0, x_1)$. We choose now b_1 such that the tuple (b_0, b_1) will satisfy $\phi_1(x_0, x_1)$. We proceed in this fashion ω times. We in the end produce an elementary embedding given by $a_k \mapsto b_k$, with $k \in \omega$, from $\mathfrak{A}$ to $\mathfrak{B}$.

We now assume (2). So $\mathfrak{A}$ is embedded, by definition, in every single model of its theory $Th(\mathfrak{A})$, and thus in every one of its countable models. So $\mathfrak{A}$ is countably prime.

Now we assume (3). We let the tuple $(a_1, \ldots, a_n)$ of elements from the domain A of $\mathfrak{A}$ be given. We further let $\Gamma(x_1, \ldots, x_n)$ be the set of $\mathcal{L}$-formulas $\gamma(x_1, \ldots, x_n)$ satisfied by the given tuple. For any countable model $\mathfrak{B}$ of $T = Th(\mathfrak{A})$, there exists some embedding $f : \mathfrak{A} \prec \mathfrak{B}$, where $(f(a_1), \ldots, f(a_n))$ satisfies $\Gamma(x_1, \ldots, x_n)$ in $\mathfrak{B}$. Thus we have that Γ is realized in every countable model of T. Invoking the Omitting Types Theorem, Γ is locally realized by T. Hence, there is a formula $\psi(x_1, \ldots, x_n)$ consistent with T such that $T \models \psi \to \gamma$ for every $\gamma \in \Gamma$. However, for each formula $\phi(x_1, \ldots, x_n)$, either ϕ or $\neg\phi$ is in Γ. Therefore, ψ is complete in T. Note that we cannot have $T \models \psi \to (\neg\psi)$, so we conclude that $\psi \in \Gamma$. Finally, we observe that we obtained a complete formula $\psi(x_1, \ldots, x_n)$ satisfied in $\mathfrak{A}$ by the given tuple $(a_1, \ldots, a_n)$ of elements from the domain A. We conclude that $\mathfrak{A}$ is atomic.

$\dashv$

Definition 6.7. A model $\mathfrak{A}$ is said to be *ω-saturated* if for every finite set $Y \subset A$, every set of formulas $\Gamma(x)$ consistent with $Th(\mathfrak{A}_Y)$ is realized in $\mathfrak{A}_Y$. A model is called *countably saturated* if it is countable and ω-saturated.

We now make a brief pause to introduce the new notation in the definition. Given a model $\mathfrak{C}$ and a subset $Y \subset C$, the expanded model $(\mathfrak{C}, c)_{c \in C}$ will be denoted by $\mathfrak{C}_Y$ and its language by $\mathcal{L}_C$.

As noted before, a type in $x_1, \ldots, x_n$ is a maximal consistent set of formulas $\Gamma(x_1, \ldots, x_n)$. The set T' of all sentences in in Γ (which is a maximal consistent set of sentences), will be called the *theory of* Γ. If $T \subset \Gamma$, we will refer to Γ as a *type of* T. Given a model $\mathfrak{A}$ of T and a type of some n-tuple of elements in the domain, we name that type of the n-tuple in $\mathfrak{A}$ (which is a type of $\mathfrak{A}$) a *type of Th($\mathfrak{A}$).*

One might ask if any power is lost by considering formulas in one free variable in the definition of ω-saturation, or if there is any particular reason for making such a choice. The truth is that it preserves all the power necessary to deal with sets of formulas free in more than one variable. The proposition that follows explains why this is the case.

Proposition 6.8. *Let $\mathfrak{A}$ be an ω-saturated model. Then for each finite $Y \subset A$, each set of formulas $\Gamma(x_1 \ldots x_n)$ of $\mathcal{L}_Y$ consistent with $Th(\mathfrak{A}_Y)$ is realized in $\mathfrak{A}_Y$.*

Proof. The argument is given by induction. The definition of ω-saturation assures it for $n = 1$, so we now assume the result for $n-1$ and take $\Gamma(x_1, \ldots, x_n)$ consistent with $Th(\mathfrak{A}_Y)$. We can further assume that Γ is closed under finite conjunctions. Let the set $\{\Gamma'(x_1, \ldots, x_{n-1})\} = \{(\exists x_n)\gamma(x_1, \ldots, x_n) : \gamma \in \Gamma\}$ be given. Then Γ' is consistent with $Th(\mathfrak{A}_Y)$. By the inductive hypothesis, there is an $(n-1)$-tuple $a_1, \ldots, a_{n-1}$ realizing Γ' in $\mathfrak{A}_Y$. Expand Y to $Y' = Y \cup a_1, \ldots, a_{n-1})$. Clearly Y' remains finite and the set $\Gamma(c_1, \ldots, c_{n-1}, x)$ is consistent with $Th(\mathfrak{A}'_Y)$ because given any finite collection $\gamma_1, \ldots, \gamma_m$ from Γ, we are given that $(\exists x_n)(\gamma_1 \wedge \ldots \wedge \gamma_m)$ is in Γ'. Since $\mathfrak{A}$ is ω-saturated, there is some a_n in the domain such that a_n realizes $\Gamma(c_1, \ldots, c_{n-1}, x_n)$ in $\mathfrak{A}'_Y$. So, taking the reduct, $(a_1, \ldots, a_n)$ realizes Γ in $\mathfrak{A}_Y$. ⊣

Theorem 6.9 (Existence of Countably Saturated Models)**.** *Let T be a complete theory. Then T has a countably saturated model if and only if, for each finite n, the theory T has only countably many types in n variables.*

Proof. Assume that T has a countably saturated model $\mathfrak{A}$. So by the previous result it realizes every type in n variables. Since an n-tuple for each n can realize at most one type, then, given that there are only countably many n-tuples of all $n \in \omega$ in $\mathfrak{A}$, the theory T can only have countably many types. Now we assume T has only countably many types. Add a countable set of new constants $C = c_1, \ldots, c_n$ to $\mathcal{L}$ to form $\mathcal{L}'$. For each finite subset $Y = d_1, \ldots, d_n \subset C$, the types $\Gamma(x)$ of T in $\mathcal{L}_Y$ are in bijective correspondence with the types $\Sigma(x_1, \ldots, x_n, x)$ of T in $\mathcal{L}$. Thus T has only countably many types $\Gamma(x)$ in $\mathcal{L}_Y$. Also taking into account that there are only countably many finite subsets of C, we enumerate the types of T in all expansions $\mathcal{L}_Y$ with Y as defined:

$$\Gamma_1, \Gamma_2, \ldots, \Gamma_n, \ldots.$$

and we enumerate all sentences of $\mathcal{L}'$:

$$\phi_1, \phi_2, \ldots, \phi_n, \ldots.$$

Now we form an increasing chain of theories in $\mathcal{L}'$

$$T = T_0 \subset T_1 \subset T_2 \subset \ldots \subset T_n \subset \ldots$$

such that the following are satisfied for each $m \in \omega$:

1. T_m is a consistent theory which contains only finitely many constants from C.
2. Either ϕ_m, or $\neg\phi_m$ is in T_{m+1}.
3. If $\phi_m = (\exists x)\psi(x)$ is in T_{m+1}, then $\psi(c)$ is in T_{m+1} for some $c \in C$.
4. If $\Gamma_m(x)$ is consistent with T_{m+1}, then $\Gamma_m(b) \subset T_{m+1}$ for some $b \in C$.

The construction of the chain is very similar to what has been done so far, in that the details are left for the reader to check.

We now consider the union $T_\omega = \bigcup_{n<\omega} T_n$, which is by (2) maximal consistent in $\mathcal{L}'$. Using (3) (witnesses), we are given that T_ω has a model $\mathfrak{A}' = (\mathfrak{A}, a_1, \ldots)$ with domain $A = \{a_1, a_2, \ldots, a_n, \ldots\}$. Therefore $\mathfrak{A}$ is a countable model of T which we now show is ω-saturated. Let $Y \subset A$ be finite. Take any $\Sigma(x)$ consistent with $Th(\mathfrak{A}_Y)$. Extend $\Sigma(x)$ to a type $\Gamma(x)$ in $Th(\mathfrak{A}_Y)$. For some m, we have $\Gamma(x) = \Gamma_m(x)$. If $\Gamma_m(x)$ is consistent with T_ω, it is consistent with $T_{m+1} \subset T_\omega$, so it was added with a constant from C plugged in, that is, $\Gamma_m(c_j) \subset T_{m+1}$, as described by (4) for some $c_j \in C$. So we establish the existence of some element in the domain A that realizes the type $\Gamma(x)$, and thus $\Sigma(x)$ in $\mathfrak{A}_Y$. ⊣

Corollary 6.10. *If T is a complete theory with only countably many nonisomorphic countable models, then T has a countably saturated model.*

Proof. This is a cardinality argument. It turns out that each type of T is realized in some countable model of T (completeness and downward LST), and each countable model, as pointed out previously, can only realize countably many types. Moreover, we are given only countably many countable models, thus giving in total only countably many types to realize. The conclusion follows from the theorem above. $\dashv$

Theorem 6.11 (Uniqueness of Countably Saturated Models)**.** *If $\mathfrak{A}$ and $\mathfrak{B}$ are countably saturated models and $\mathfrak{A} \equiv \mathfrak{B}$, then $\mathfrak{A} \cong \mathfrak{B}$.*

Proof. We again encounter the back and forth construction. However, this time, instead of using complete formulas, we will direct our attention to types and use them in the same manner (getting larger n-tuples by increasing n in order to pick the next element from iterated lists of order type ω). By countable saturation, we arrive at two sequences

$$a_0, a_1, a_2, a_3, \ldots \quad \text{and} \quad b_0, b_1, b_2, b_3, \ldots$$

listing the entirety of the domains A and B, respectively. And, for each n, the element a_n realizes the same type in $(\mathfrak{A}, a_0, a_1, \ldots, a_{n-1})$ as b_n realizes in $(\mathfrak{B}, b_0, b_1, \ldots, b_{n-1})$. The map $a_n \mapsto b_n$ gives an isomorphism because for each relation, function, or constant symbol, and each appropriate tuple, one only needs to pick N large enough so that all involved pre-image elements appear in $\{a_n : n < N\}$ and are confirmed by the back and forth construction in $\{b_n : n < N\}$. $\dashv$

Definition 6.12. A model $\mathfrak{A}$ is said to be *countably universal* if $\mathfrak{A}$ is countable and every countable model of $Th(\mathfrak{A})$ can be elementarily embedded in $\mathfrak{A}$.

Theorem 6.13. *Every countably saturated model is countably universal.*

Proof. Let $\mathfrak{B}$ be a countable model with $\mathfrak{A}$ a countably saturated model of $Th\mathfrak{B}$. Further let $B = \{b_0, b_1, \ldots\}$. Using a similar procedure to the back and forth construction (as in the prime model equivalence theorem) and the saturation of $\mathfrak{A}$, we obtain a sequence of elements $a_0, a_1, \ldots$ from A such that

$$(\mathfrak{B}, b_0, b_1, \ldots) \equiv (\mathfrak{A}, a_0, a_1, \ldots).$$

Hence the map $b_k \mapsto a_k$ is an elementary embedding from $\mathfrak{B}$ into $\mathfrak{A}$. $\dashv$

Now we present a characterization of ω-categoricity.

Theorem 6.14 (Characterization of ω-Categorical Theories)**.** *Let T be a complete theory. Then the following are equivalent:*

1. *T is ω-categorical.*
2. *T has a model $\mathfrak{A}$ which is both countably saturated and atomic.*
3. *For each $n \in \omega$, each type $\Gamma(x_1, \ldots, x_n)$ of T contains a complete formula.*
4. *For each $n \in \omega$, the theory T has only finitely many types in $x_1, \ldots, x_n$.*
5. *For each $n \in \omega$, there are only finitely many formulas $\phi(x_1, \ldots, x_n)$, up to equivalence, with respect to T.*
6. *All models of T are atomic.*

Proof. We assume (1) to prove (2). Take $\mathfrak{A}$ to be the unique countable model of T. Then $\mathfrak{A}$ is countably prime, so $\mathfrak{A}$ is atomic. Given that T has only one (and $1 \in \omega$) countable model, it has a countably saturated model. Thus $\mathfrak{A}$ must be that countably saturated model.

We now assume (2) to prove (3). Since $\mathfrak{A}$ is ω-saturated, the type Γ is realized in $\mathfrak{A}$ by some tuple $(a_1, \dots, a_n)$. Since $\mathfrak{A}$ is atomic, the tuple $(a_1, \dots, a_n)$ given satisfies a complete formula $\gamma(x_1, \dots, x_n)$. Since $\neg\gamma$ cannot be in Γ, and Γ is a type, then $\gamma \in \Gamma$.

Let (3) hold and now we prove (4). Let $\Sigma(x_1, \dots, x_n)$ be the set of negations of all complete formulas $\sigma(x_1, \dots, x_n)$ in T. Thus Σ cannot be extended to a type in $x_1, \dots, x_n$ so Σ is inconsistent with T. Hence, some finite subset $\{\neg\sigma_1, \dots, \neg\sigma_k\} \subset \Sigma$ is inconsistent with T. Thus, we get $T \models \neg(\neg\sigma_1 \wedge \dots \wedge \neg\sigma_k)$. Which yields that $T \models \sigma_1 \vee \dots \vee \sigma_n$. Note that for each $i \leq k$, the set $\Gamma_i(x_1, \dots, x_n)$ of all consequences of $T \cup \{\sigma_i\}$ is a type of T. However, given inconsistency of $\neg\sigma_1 \wedge \dots \wedge \neg\sigma_k$ with T, every n-tuple satisfies some σ_i, and thus realizes some Γ_i. Hence, $\Gamma_1, \dots, \Gamma_k$ are the only types of T in $x_1, \dots, x_n$.

Take (4) to hold and we give a proof of (5). We take some formula $\phi(x_1, \dots, x_n)$, and we let $\phi*$ be the set of all types $\Gamma(x_1, \dots, x_n)$ of T which contain ϕ. So $\phi* = \psi*$ implies that $T \models \phi \leftrightarrow \psi$. But we know that there are only $k < \omega$ types of T in $x_1, \dots, x_n$. Therefore there are only 2^k possible sets of types and, consequently, only 2^k formulas up to equivalence in T.

Assume now (5) and we prove (6). Take $\mathfrak{A}$ to be a model of T and let $a_1, \dots, a_n \in A$ be given. We create a finite list $\phi_1(x_1, \dots, x_n), \dots, \phi_k(x_1, \dots, x_n)$ of all formulas satisfied by the n-tuple $(a_1, \dots, a_n)$ up to equivalence in T. Then $\phi_1 \wedge \dots \wedge \phi_k$ is a complete formula in T satisfied by $(a_1, \dots, a_n)$ in $\mathfrak{A}$. Therefore, $\mathfrak{A}$ is atomic.

Now assume (6) and we prove (1). We can observe that any two countable models of T are atomic by assumption and elementarily equivalent since T is complete. The uniqueness of atomic models, then, assures us that these two models are isomorphic. Thus T is ω-categorical. ⊣

Theorem 6.15. *Any complete theory T which has a countably saturated model has a countable atomic model.*

Proof. We begin assuming for the sake of contraposition that T has no countable atomic model. Then T is not atomic. It is, therefore, the case that there exists some incompletable (in T) formula $\psi(x_1, \dots, x_n)$. Thus, for each consistent incompletable formula $\psi(x_1, \dots, x_n)$ of T we can choose from two consistent formulas with T, namely $\psi_0(x_1, \dots, x_n)$ and $\psi_1(x_1, \dots, x_n)$, such that $T \models \psi_0 \to \psi$ and $T \models \psi_1 \to \psi$ hold, and also $T \models \neg(\psi_0 \wedge \psi)$. This gives that ψ_0 and ψ_1 are also incompletable. In this fashion, we obtain a tree of incompletable formulas:

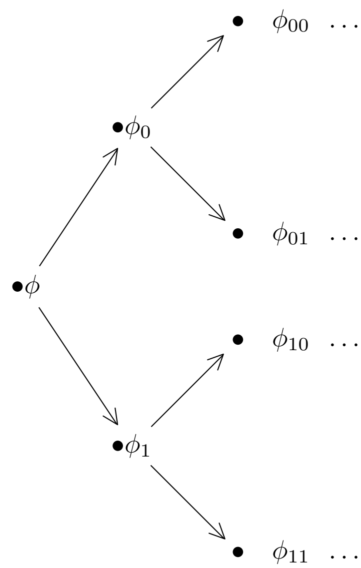

Each infinite binary sequence gives a branch $\Gamma_s = \{\phi, \phi_{s_0}, \phi_{s_0,s_1}, \phi_{s_0,s_1,s_2}, \ldots\}$. There are then naturally 2^ω branches to this tree. By the condition that $T \models \psi_0 \to \psi$ and $T \models \psi_1 \to \psi$ hold, along with $T \models \neg(\psi_0 \wedge \psi)$, each branch is a set of formulas which is consistent with T, and any two branches are inconsistent with each other. Extending each branch Γ_s to a type of T, we get 2^ω distinct types. This violates the existence theorem for countably saturated models and therefore T has no countably saturated model. ⊣

Theorem 6.16 (Vaught's Two Model Theorem)**.** *No complete theory T has exactly two nonisomorphic countable models.*

Proof. We assume for the sake of a contradiction that T has exactly two nonisomorphic countable models. The prior results on existence and on the characterization of ω-categoricity (item (2)) give that T has a countably saturated model $\mathfrak{B}$ and an atomic model $\mathfrak{A}$ such that these two are nonisomorphic (otherwise (2) would be satisfied and there would be only one countable model). Since $\mathfrak{B}$ is not atomic, it has an n-tuple $(b_1, \ldots, b_n)$ which does not satisfy a complete formula. The idea is to construct a countable atomic model $\mathfrak{C}$ as a reduct of some model $(\mathfrak{C}, c_1, \ldots, c_n)$ of the theory $T' = Th(\mathfrak{B}, b_1, \ldots, b_n)$ and show that it fails to be either countably saturated or atomic. Since $\mathfrak{B}$ is countably saturated, the model $(\mathfrak{B}, b_1, \ldots, b_n)$ is countably saturated. Thus we obtain that T' has a countably saturated model, and thus has a countable atomic model $(\mathfrak{C}, c_1, \ldots, c_n)$ by the previous result. The reduct $\mathfrak{C}$ is a model of T, naturally. Note that the model $\mathfrak{C}$ is also not atomic because the tuple $(c_1, \ldots, c_n)$ does not satisfy a complete formula. Now we claim that $\mathfrak{C}$ is also not ω-saturated. Given that T is not ω-categorical, it has infinitely many non-equivalent formulas. Therefore, T' also has infinitely many non-equivalent formulas. The same result also indicates, via item (2), that no model of t' is both atomic and ω-saturated. Also, noting that $(\mathfrak{C}, c_1, \ldots, c_n)$ is atomic, we see it canot be ω-saturated. Therefore, it is the case that $\mathfrak{C}$ is not ω-saturated. So it is neither $\mathfrak{A}$ nor $\mathfrak{B}$. ⊣

7. Elementary Chains and Indiscernibles

This section takes a slightly different turn from the rest of the paper so far. Here, we analyze ways in which theories cannot discern between n-tuples of elements in a model and explore different flavors of model-theoretic constructs. We begin with a combinatorial result, which we will prove with the aid of ultrafilters. First, however, it is necessary to introduce the following shorthand notation: for any set X, we set $[X]^n$ to be the set of all subsets of X with exactly n elements.

Theorem 7.1 (Ramsey's Theorem)**.** *Let I be an infinite set and $n \in \omega$. Consider any partition $[I]^n = A_0 \cup A_1$. Then there is an infinite set $J \subset I$ such that either $[J]^n \subset A_0$, or $[J]^n \subset A_1$*

Proof. We begin by letting I to be countably infinite, since for larger sets any countably infinite subset I will satisfy the hypothesis. We totally order the elements of I in an increasing sequence under some order relation $<$:

$$i_0 < i_1 < i_2 < \ldots < i_k < \ldots$$

We may assume $n > 1$. We take D some nonprincipal ultrafilter over I. Notice that for all k, we have that $\{i \in I : i_k < i\} \in D$ given that the finite-cofinite filter is a subset of any nonprincipal D. For each $r < n$ we define two subsets A_0^{n-r} and A_1^{n-r} of $[I]^{n-r}$ by induction on r in the following manner:

$$A_0^n = A_0 \text{ and } A_1^n = A_1$$

Assume that A_0^{n-r} and A_1^{n-r} have been defined so that $[I]^{n-r} \subset A_0^{n-r} \cup A_1^{n-r}$. Let

$$A_0^{n-r-1} = \{y_1, \ldots, y_{n-r-1} : \{i \in I : y_{n-r-1} < i \text{ and } \{y_1, \ldots, y_{n-r-1}, i\} \in A_0^{n-r}\} \in D\}$$

and

$$A_1^{n-r-1} = \{y_1, \ldots, y_{n-r-1} : \{i \in I : y_{n-r-1} < i \text{ and } \{y_1, \ldots, y_{n-r-1}, i\} \in A_1^{n-r}\} \in D\}.$$

Using the properties of the ultrafilter D we obtain that:

$$[I]^{n-r-1} = A_0^{n-r-1} \cup A_1^{n-r-1}.$$

In this fashion, we increase r until we have that $I = A_0^1 \cup A_1^1$. We now have that either A_0^1 or A_1^1 is in D. Without loss of generality, assume $A_0^1 \in D$. We give an infinite sequence $j_0 < j_1 < j_2 < \ldots < j_k < \ldots$ of elements of I inductively by the following procedure: let $j_0 \in A_0^1$, and assume $j_0 < j_1 < j_2 < \ldots < j_k$ have been defined so that

$$\text{for all } r,\ 1 \le r \le n, \text{ and all } y_1 < \ldots < y_r, \text{ from } \{j_0 \ldots j_k\}, \text{ the set } \{y_1, \ldots, y_r\} \in A_0^r.$$

Now define j_{k+1} as follows. By the inductive hypothesis, given the ascending sequence $y_1 < \ldots < y_r$, from $\{j_0, \ldots, j_k\}$, with $r < n$, the set $X_{y_1,\ldots,y_r}$ has the following property:

$$X_{y_1,\ldots,y_r} = \{i \in I : y_r < i \text{ and } \{y_1, \ldots, y_r, i\} \in A_0^{r+1}\} \in D.$$

Given that there exist at most some finite number of increasing sequences of length less than or equal to $n - 1$ from the set $\{j_0, \ldots, j_k\}$, the number of sets of the $X_{y_1,\ldots,y_r}$ kind is also finite. Moreover, note that their intersection $Y \in D$. Since D is nonprincipal, we can pick an element $\jmath_{k+1} \in Y$ such that $j_k < j_{k+1}$. In this way, the inductive hypothesis will now hold with k replaced by $k + 1$. In this fashion, we notice that the infinite set $J = \{j_0, j_1, \ldots, j_k, \ldots\}$ is constructed as desired. It is clear, also, that any ascending finite segment $j_1 < \ldots < j_n$ from J will satisfy

$\{j_0, \ldots, j_n\} \in A_0^n = A_0$, yielding $[J]^n \subset A_0$. The case with $A_1^1 \in D$ is exactly the same argument. ⊣

We now introduce the defintion of a set of indiscernibles in a model.

Definition 7.2. Let $\mathfrak{A}$ be an $\mathcal{L}$-structure and $X \subset A$ some subset of the domain carrying a simple ordering $<$ (which may not apply at all to the rest of A). X is a *set of indiscernibles* in $\mathfrak{A}$ (with respect to $<$) if, for every $n \in \omega$ and every pair of finite sequences $x_1 < x_2 < \ldots < x_n$ and $y_1 < y_2 < \ldots < y_n$ of elements in X, we have that $(\mathfrak{A}, x_1, x_2, \ldots, x_n) \equiv (\mathfrak{A}, y_1, y_2, \ldots, y_n)$.

Lemma 7.3. *Let $\langle X, < \rangle$ be an ordered subset of a model $\mathfrak{A}$. Suppose that for any two increasing n-tuples $x_1 < \ldots < x_n$ and $y_1 < \ldots < y_n$ of X, there is an automorphism f of $\mathfrak{A}$ such that $f(x_i) = y_i$ for all $1 \leq i \leq n$. Then X is a set of indiscernibles in $\mathfrak{A}$.*

Proof. We have by the definition of automorphism that $f : (\mathfrak{A}, x_1, \ldots, x_n) \overset{\cong}{\to} (\mathfrak{A}, y_1, \ldots, y_n)$. Hence, $(\mathfrak{A}, x_1, \ldots, x_n) \equiv (\mathfrak{A}, y_1, \ldots, y_n)$, and it follows that X is a set of indiscernibles. ⊣

Lemma 7.4. *Let $\overline{\mathcal{L}} = \mathcal{L} \cup \{c_n : n < \omega\}$, where c_n are constants not in $\mathcal{L}$. Let T be a theory in $\mathcal{L}$ with infinite models. Then the following set T' is consistent:*

$$\begin{aligned} T' = T \cup \{&\phi(c_{i_1}, \ldots, c_{i_n}) \leftrightarrow \phi(c_{j_1}, \ldots, c_{j_n}) : \phi(v_1, \ldots, v_n) \text{ is a formula of } \mathcal{L}, \\ &n \in \omega, \text{ and } i_1 < \ldots < i_n, j_1 < \ldots < j_n\} \\ &\cup \{\neg(c_{i_1} \equiv c_{i_2}) : i_1 \neq i_2\} \end{aligned}$$

Proof. Take $\mathfrak{A}$ to be any infinite model of T and let I be a countably infinite subset of A. Let a well-ordering $<$ of I with order type ω be given. Then

$$i_0 < i_1 < i_2 < \ldots < i_k < \ldots$$

is a complete list of the elements of I in ascending order. Now we make the following intermediate claim:

- ➪ Let Δ be any finite subset of T'. Then there is an infinite subset J_Δ of I such that for each infinite subset

$$j_0 < j_1 < \ldots < j_k < \ldots$$

of J_Δ, the expansion $(\mathfrak{A}, j_n)_{n \in \omega}$ satisfies Δ.

The intermediate claim is proven by induction on the number of sentences in Δ. Assume the claim holds for some finite subset Δ of T' and take some $\mathcal{L}$-formula $\phi(v_1, \ldots, v_n)$. Now we partition $[J_\Delta]^n$ into two pieces as follows:

$$A_0 = \{x_1 < \ldots < x_n : x_i \in J_\Delta \text{ and } \mathfrak{A} \models \phi[x_1, \ldots, x_n]\}$$

and

$$A_1 = \{x_1 < \ldots < x_n : x_i \in J_\Delta \text{ and } \mathfrak{A} \models \neg\phi[x_1, \ldots, x_n]\}$$

We first note that $[J_\Delta]^n \subset A_0 \cup A_1$ so we can apply Ramsey's Theorem. So, by said result, we have an infinite subset $K \subset [J_\Delta]^n$ so that either $[K]^n \subset A_0$ or $[K]^n \subset A_1$. We now consider any infinite subset $k_0 < k_1 < k_2 < \ldots$ of K. Then we have that the expansion $(\mathfrak{A}, k_n)_{n \in \omega}$ satisfies $\phi(c_{s_1}, \ldots, c_{s_n}) \leftrightarrow \phi(c_{t_1}, \ldots, c_{t_n})$, where $s_1 < \ldots < s_n$ and $t_1 < \ldots < t_n$. One can notice this by considering the two possible cases, namely, if $[K]^n \subset A_0$ or if $[K]^n \subset A_1$. Thus we proved the intermediate claim holds whenever one extra sentence is added to it, and so the induction is complete.

Moreover, since we just proved the consistency of every finite subset $\Delta \subset T'$, we have that T' must be consistent. ⊣

Theorem 7.5. *Let T be a theory in $\mathcal{L}$ with infinite models, and let $\langle X, < \rangle$ be any simply ordered set. Then there is a model $\mathfrak{A}$ of T with $X \subset A$ and such that X is a set of indiscernibles in $\mathfrak{A}$.*

Proof. We begin by expanding the language $\mathcal{L}$ by adding a set of new constants to obtain $\overline{\mathcal{L}} = \mathcal{L} \cup \{c_x : X \in X\}$ and let:

$$\begin{aligned} T' = T \cup \{&\phi(c_{x_1}, \dots, c_{x_n}) \leftrightarrow \phi(c_{y_1}, \dots, c_{y_n}) : \phi(v_1, \dots, v_n) \text{ is a formula of } \mathcal{L}, \\ &n \in \omega, \text{ and } x_1 < \dots < x_n, y_1 < \dots < y_n \text{ from } X\} \\ &\cup \{\neg(c_{x_1} \equiv c_{x_2}) : x_1 \neq x_2 \text{ in } X\} \end{aligned}$$

Since every finite subset of X will satisfy the conditions outlined in the intermediate claim of Lemma 7.3 (every finite subset of X can be embedded preserving order in ω), we have that T' is consistent in $\overline{\mathcal{L}}$. Take $\overline{\mathfrak{A}}$ to be some model of T' in $\overline{\mathcal{L}}$ and let $\mathfrak{A}$ be its reduct to $\mathcal{L}$. We notice that $\mathfrak{A}$ is a model of T. Naturally, without loss of generality, we can also identify the interpretations of each constant $c_x, x \in X$ with the element of X after which it is named. Also the elements of T' assure that given any $\mathcal{L}$-formula $\phi(v_1, \dots, v_n)$ and $x_1 < \dots < x_n, y_1 < \dots < y_n$ of X, we are given:

$$\mathfrak{A} \models \phi[x_1, \dots, x_n] \text{ if and only if } \mathfrak{A} \models \phi[y_1, \dots, y_n].$$

It then follows that $(\mathfrak{A}, x_1, \dots, x_n) \equiv (\mathfrak{A}, y_1, \dots, y_n)$, thus making X a set of indiscernibles in $\mathfrak{A}$. ⊣

Now, we will begin a discussion on elementary extensions. First, we note a few convenient facts.

Proposition 7.6. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be models*

(1) *If $\mathfrak{A} \prec \mathfrak{B}$, then $\mathfrak{A} \equiv \mathfrak{B}$.*
(2) *$\mathfrak{A} \prec \mathfrak{B}$*
(3) *If $\mathfrak{A} \prec \mathfrak{B}$ and $\mathfrak{B} \prec \mathfrak{C}$, then If $\mathfrak{A} \prec \mathfrak{C}$*
(4) *If $\mathfrak{A} \prec \mathfrak{C}$, $\mathfrak{B} \prec \mathfrak{C}$, and $\mathfrak{A} \subset \mathfrak{B}$, then $\mathfrak{A} \prec \mathfrak{B}$.*

Proof. These items all follow after simple manipulations of the definitions recalled in section 2. ⊣

It also turns out to be very convenient to have a different characterization of elementary submodels, namely the one suggested by the following proposition.

Proposition 7.7. *It is the case that $\mathfrak{A} \prec \mathfrak{B}$ if and only if $\mathfrak{A} \subset \mathfrak{B}$ and for all formulas $(\exists x)\phi(x, x_1, \dots, x_n)$ and all n-tuples $(a_1, \dots, a_n)$ of elements in A, the following occurs: If $\mathfrak{B} \models (\exists x)\phi(x, a_1, \dots, a_n)$, then there exists some $a \in A$ such that $\mathfrak{B} \models \phi(a, a_1, \dots, a_n)$.*

Proof. The 'only if' direction follows *a fortiori* from the definition of an elementary submodel. It is then left to prove the other direction. To do so, we proceed by induction on the complexity of formulas ϕ on the statement: if we have $\phi(v_1, \dots, v_n)$ and some n-tuple $(a_1, \dots, a_n)$ of elements from A:

$$\mathfrak{A} \models \phi[a_1, \dots, a_n] \text{ if and only if } \mathfrak{B} \models \phi[a_1, \dots, a_n].$$

Verifying the inductive procedure at the level of atomic formulas and sentential connectives $\wedge$ and $\neg$ is straightforward. Now the aim is to be able to deal with quantifiers, which amounts to moving from $\phi(v_1, \ldots, v_n)$ to $(\exists v_1)\phi(v_2, \ldots, v_n)$. To make this more tractable, we can induct on the number of existential quantifiers after turning all $\forall$ ones into $\neg\exists\neg$. The base case is quantifier-free formulas is already dealt with early on in this proof (we know how to deal with the atomic case and the sentential connectives). Now take the inductive step. We notice the following: take some $n-1$-tuple $(a_2, \ldots, a_n)$ of elements from A. If $\mathfrak{A} \models (\exists v_1)\phi[a_2, \ldots, a_n]$, then there exists some a_1 in A such that $\mathfrak{A} \models \phi[a_1, \ldots, a_n]$. By the inductive step, we have that $\mathfrak{B} \models \phi[a_1, \ldots, a_n]$, and so $\mathfrak{B} \models (\exists v_1)\phi[a_2, \ldots, a_n]$. Now, assuming $\mathfrak{B} \models (\exists v_1)\phi[a_2, \ldots, a_n]$, the hypothesis of this proposition gives that there is some a_1 *inA* such that $\mathfrak{B} \models \phi[a_1, \ldots, a_n]$. So, by induction, $\mathfrak{A} \models \phi[a_1, \ldots, a_n]$, and hence $\mathfrak{A} \models (\exists v_1)\phi[a_2, \ldots, a_n]$. ⊣

In order to establish our next important result, we need to introduce two definitions to make way for the theorem. The first one just deals with chains under containment.

Definition 7.8. (1) A *chain of models* is an increasing sequence

$$\mathfrak{A}_0 \subset \mathfrak{A}_1 \subset \ldots \subset \mathfrak{A}_\beta \subset \ldots \quad \text{with} \quad \beta < \alpha$$

for some ordinal α.

(2) The *union of the chain* is the model $\mathfrak{A} = \bigcup_{\beta<\alpha} \mathfrak{A}_\beta$ with domain $A = \bigcup_{\beta<\alpha} A_\beta$. Each relation R of $\mathfrak{A}$ in the language is the union of the corresponding relations in the models $\mathfrak{A}_\beta$ of the chain. The same goes for the function symbols, that is, each is the union of the corresponding functions of the models $\mathfrak{A}_\beta$ in the chain.

We are ready to define elementary chains.

Definition 7.9. An *elementary chain* of models is a chain:

$$\mathfrak{A}_0 \prec \mathfrak{A}_1 \prec \ldots \prec \mathfrak{A}_\beta \prec \ldots \quad \text{with} \quad \beta < \alpha$$

so that $\mathfrak{A}_\gamma \prec \mathfrak{A}_\eta$ if $\gamma < \eta < \alpha$.

We should just state a result, which can quickly be fact-checked by finding an isomorphism, as it exhibits a nice property of the union of a chain of models.

Proposition 7.10. *Given a chain of models $\mathfrak{A}_\beta$, with $\beta < \alpha$, the union of the chain $\bigcup_{\beta<\alpha} \mathfrak{A}_\beta$ is the unique model with domain $\bigcup_{\beta<\alpha} A_\beta$ which contains each $\mathfrak{A}_\beta$ as a submodel.*

And now we turn to a very important theorem:

Theorem 7.11 (Elementary Chain Theorem)**.** *Let $\mathfrak{A}_\zeta$, for $\zeta < \alpha$, be an elementary chain of models. Then $\mathfrak{A}_\zeta \prec \bigcup_{\zeta<\alpha} \mathfrak{A}_\zeta$ for all $\zeta < \alpha$.*

Proof. Let $\mathfrak{A} = \bigcup_{\zeta<\alpha} \mathfrak{A}_\zeta$ be given. We prove the following statement by induction on the complexity of formulas ϕ:

➪ For all formulas ϕ in $x_1, \ldots, x_n$, all $\zeta < \alpha$ and all elements $a_1, \ldots, a_n \in A_\zeta$,

$$\mathfrak{A}_\zeta \models \phi[a_1, \ldots, a_n] \text{ if and only if } \mathfrak{A} \models \phi[a_1, \ldots, a_n].$$

The construction of $\mathfrak{A}$ assures that the statement will hold in the case of atomic formulas. Checking for the sentential connectives is also straightforward. So we proceed to the quantifier case in the usual manner. Let $\phi \equiv (\exists x_1)\psi$ be given with ϕ a formula in $x_2, \ldots, x_n$, ordinal $\zeta < \alpha$, and $a_2, \ldots, a_n \in A_\zeta$. If $\mathfrak{A}_\zeta \models \phi[a_2, \ldots, a_n]$, then there is some a_1 in A_ζ such that $\mathfrak{A}_\zeta \models \psi[a_1, \ldots, a_n]$. So by induction $\mathfrak{A} \models \psi[a_1, \ldots, a_n]$ and $\mathfrak{A} \models \phi[a_2, \ldots, a_n]$. Now let $\mathfrak{A} \models \phi[a_2, \ldots, a_n]$, then for some $\eta < \alpha$ and $a_1 \in A$, we are given an $a_1, \ldots, a_n$ all in A_η and $\mathfrak{A} \models \psi[a_1, \ldots, a_n]$. Since $\mathfrak{A}_\zeta$, with $\zeta < \alpha$ is a chain, we can assume $\zeta \leq \eta$. Given that $a_1, \ldots, a_n$ are all in A_η, induction assures that $\mathfrak{A}_\eta \models \psi[a_1, \ldots, a_n]$, and hence $\mathfrak{A}_\eta \models \phi[a_2, \ldots, a_n]$. And since $\mathfrak{A}_\zeta \prec \mathfrak{A}_\eta$ we have that $\mathfrak{A}_\zeta \models \phi[a_2, \ldots, a_n]$. ⊣

8. A quick introduction to Skolem Functions

Skolem functions are a very interesting idea in Model Theory. Here, we define what they are and observe how they interact with the indiscernibles we introduced in the previous section. Our starting point is a language $\mathcal{L}$ which we expand to a language $\mathcal{L}*$ by adding a new collection of function symbols. Let F be mapping from the set of all formulas of the form $\psi \equiv (\exists x)\phi(x)$ to a list of new function symbols F_ψ. We will assume that F is injective and that if ψ has n free variables, then F_ψ is an n-placed (*Skolem*, as we shall see) function symbol.

Definition 8.1. Let some language $\mathcal{L}$ be given.

(1) The expansion $\mathcal{L} \cup \{F_\psi : \psi \equiv (\exists x)\phi(x)$ a formula of $\mathcal{L}\}$ is called a *Skolem expansion.*
(2) We make a list of sentences of $\mathcal{L}*$ and call sentences like the following the axioms of the *Skolem Theory* of $\mathcal{L}$, denoted by $\Sigma_{\mathcal{L}}$:
$$(\forall y_1, \ldots, \forall y_n)(\psi(y_1, \ldots, y_n) \to \phi(F_\psi(y_1, \ldots, y_n), y_1, \ldots, y_n))$$
(3) Let $\mathfrak{A}$ be an $\mathcal{L}$-structure. An expansion $\mathfrak{A}*$ of $\mathfrak{A}$ to the new language $\mathcal{L}*$ is called a *Skolem expansion of* $\mathfrak{A}$ if $\mathfrak{A}* \models \Sigma_{\mathcal{L}}$.
(4) A *Skolem expansion* $T*$ *of a theory* T is the theory $T* = T \cup \Sigma_{\mathcal{L}}$.

Proposition 8.2. *Let some language $\mathcal{L}$, consistent theory T of $\mathcal{L}$, and $\mathcal{L}$-structure $\mathfrak{A}$ be given.*

(1) Every model $\mathfrak{A}$ of $\mathcal{L}$ has a Skolem expansion $\mathfrak{A}*$.
(2) If T is a consistent theory in $\mathcal{L}$, then its Skolem expansion $T*$ is a consistent theory in $\mathcal{L}*$.
(3) Let $\mathfrak{A}$ and $\mathfrak{B}$ be models of $\mathcal{L}$. Consider their Skolem expansions $\mathfrak{A}*$ and $\mathfrak{B}*$. If $\mathfrak{A}* \subset \mathfrak{B}*$, then $\mathfrak{A} \prec \mathfrak{B}$.

Proof. We begin by proving (1). Take $\mathfrak{A}$ to be some $\mathcal{L}$-structure. Let $\psi \equiv (\exists x)\phi$ and assume that ψ has exactly the free variables $x_1, \ldots, x_n$. The idea is to define an interpretation G_ψ of F_ψ in $\mathfrak{A}$. We begin by well-ordering the domain A. Given any n-tuple $(a_1, \ldots, a_n)$ of elements from the domain A, if $\mathfrak{A} \models \psi[a_1, \ldots, a_n]$, then take $G_\psi(a_1, \ldots, a_n)$ to be the first element a of A such that $\mathfrak{A} \models \phi[a, a_1, \ldots, a_n]$. If we have that $\mathfrak{A} \not\models \psi[a_1, \ldots, a_n]$, then $G_\psi(a_1, \ldots, a_n)$ is allowed to be arbitrary. Now it follows that the expansion $\mathfrak{A}* = (\mathfrak{A}, \{G_\psi : \psi \equiv (\exists x)\phi\})$ is a Skolem expansion of $\mathfrak{A}$. We also have that (2) is a direct consequence of (1), whereas for (3) all we have to do is revisit Proposition 7.5. ⊣

Definition 8.3. Let $\mathfrak{A}*$ be a Skolem expansion of $\mathfrak{A}$, and let $X \subset A$. We call the *Skolem hull* of X in $\mathfrak{A}*$ the *smallest* (under containment) set Y that:

(1) is closed under all functions in $\mathcal{L}*$,
(2) contains all constants in $\mathfrak{A}$,
(3) and naturally $X \subset Y \subset A$.

The *Skolem hull* of X is denoted by $H(X)$ and $\mathfrak{H}(X)$, the submodel of $\mathfrak{A}$ generated by $H(X)$, by $\mathfrak{H}(X)$.

Proposition 8.4. *Let $\mathfrak{A}*$ be a Skolem expansion of $\mathfrak{A}$, and let $X \subset A$. Then the Skolem hull $\mathfrak{H}(X)$ is an elementary submodel of $\mathfrak{A}$. Moreover, we have* $|H(X)| \leq |X| \cup ||\mathcal{L}||$.

Proof. We take the model $\mathfrak{H}(\mathrm{X})^*$ to be the model generated by $H(X)$ in $\mathfrak{A}^*$. Clearly, we have $\mathfrak{H}(X)* \subset \mathfrak{A}*$. Once $\mathfrak{H}(\mathrm{X})^*$ is an expansion of $\mathfrak{H}(\mathrm{X})$ to $\mathcal{L}*$, Proposition 8.2.(3) gives the result. ⊣

Definition 8.5. A theory T of $\mathcal{L}$ has *built-in Skolem functions* if for every formula $\psi = (\exists x)\phi$ with exactly the variables $x_1, \ldots, x_n$ occurring free, there is an n-placed term t_ψ of $\mathcal{L}$ such that:

$$T \vdash (\forall y_1, \ldots, \forall y_n)(\psi(y_1, \ldots, y_n) \rightarrow \phi(t_\psi(y_1, \ldots, y_n), y_1, \ldots, y_n))$$

and the variables $y_1, \ldots, y_n$ occur in neither ψ nor t_ψ.

Proposition 8.6. *If a theory T has built-in Skolem functions, then T is* ***model complete****, that is, whenever $\mathfrak{A}$ and $\mathfrak{B}$ are two models of T and $\mathfrak{A} \subset \mathfrak{B}$, we have that $\mathfrak{A} \prec \mathfrak{B}$.*

Proof. Given that $\mathfrak{A}$ is a submodel of $\mathfrak{B}$ (both models of T), it must be the case that $\mathfrak{A}$ is closed under all terms t_ψ of $\mathcal{L}$. The result is thus given by Proposition 7.6. ⊣

Proposition 8.7. *Let T be a theory of a language $\mathcal{L}$. Then there exist expansions $\mathcal{L}'$ of $\mathcal{L}$ and T' of T (T' is a theory of $\mathcal{L}'$) such that T' has built-in Skolem functions. Furthermore, every model of T has an expansion which is a model of T'.*

Proof. We craft an increasing sequence of expansions $\mathcal{L}_n$, $n \in \omega$ by defining $\mathcal{L} = \mathcal{L}_0$ and letting $\mathcal{L}_{n+1} = (\mathcal{L}_n)*$. In this way, for each n, the Skolem theory $\Sigma_{\mathcal{L}_n}$ is a set of $\mathcal{L}_{n+1}$-sentences. Now we take the union $\overline{\mathcal{L}} = \bigcup_{n\in\omega} \mathcal{L}_n$. We construct a theory $\overline{T}$ of $\overline{\mathcal{L}}$ to have the set of axioms $T \cup \bigcup_{n\in\omega} \Sigma_{\mathcal{L}_n}$.

Since every formula of $\overline{\mathcal{L}}$ has at most a finite number of symbols, we are assured that $\overline{T}$ has built-in Skolem functions. For every formula $\psi \equiv (\exists x)\phi$ as in the definition, we pick the least n where all symbols show up in the countable union $\bigcup_{n\in\omega} \mathcal{L}_n$. The language $\mathcal{L}_{n+1}$ will have the desired function symbol for the term. Furthermore, induction on Proposition 8.2.(1) will yield that every model of T has an expansion which is a model of $\overline{T}$. ⊣

After these preliminary results, we proceed to finish this paper with a series of interesting interactions between Skolem functions and indiscernibles. Though the proofs here are going to be generally short, the ideas seemed to be a worthwhile addition. Our starting point is a set X of indiscernibles in a model $\mathfrak{A}$ of a theory T with built-in Skolem functions.

Theorem 8.8 (Subset Theorem)**.** *If $Y \subset X$, then Y is a set of indiscernibles in $\mathfrak{H}(Y)$ with respect to the order inherited from X, and $\mathfrak{H}(Y) \prec \mathfrak{H}(X)$*

Proof. Notice that $\mathfrak{H}(Y)$ is an elementary submodel of $\mathfrak{H}(X)$ by our previous results. Naturally, we also have that increasing sequences of elements from Y will satisfy the same formulas as these satisfied by increasing sequences of elements from X. ⊣

Theorem 8.9 (Stretching Theorem)**.** *Let X and Y be infinite totally ordered sets. Then there exists a model $\mathfrak{B}$ in which Y is a set of indiscernibles and the sets of formulas satisfied by increasing sequences of elements of X in $\mathfrak{A}$ and Y in $\mathfrak{B}$ are the same.*

Proof. Take Σ to be the set of all formulas $\phi(v_1, \ldots, v_n)$ of $\mathcal{L}$ satisfied by increasing sequences $x_1, \ldots, x_n$ of elements from X. For some expansion $\overline{\mathcal{L}} = \mathcal{L} \cup \{c_y : y \in Y\}$ of $\mathcal{L}$, take the set Σ' of all sentences $\phi(c_{y_1}, \ldots, c_{y_n})$, where $\phi(x_1, \ldots, x_n) \in \Sigma$ and $y_1 < \ldots < y_n$ holds in the ordering of Y. Given that X is infinite, the set Σ' is consistent once one can ordermorphically embed finite sequences of Y into X. Therefore, by Theorem 7.4, we can find a model $\mathfrak{B}$ with the set Y of indiscernibles. ⊣

Theorem 8.10 (Elementary Embedding Theorem)**.** *Let Y be a set of indiscernibles in a model $\mathfrak{B}$ with the property that the sets of formulas of $\mathcal{L}$ satisfied by increasing sequences of elements from X and Y are the same. Let f be some injective ordermorphic embedding of X into Y; then f can be extended* ***uniquely*** *to an* ***elementary embedding*** *$\overline{f}$ of $\mathfrak{H}(X)$ into $\mathfrak{H}(Y)$. The range of $\overline{f}$ turns out to be H(range of f).*

Proof. Notice that every element in $H(X)$ is generated by some term $t(v_1, \ldots, v_n)$ and some collection of elements $x_1, \ldots, x_n$ from X. We may assume that said term t and elements x_i from X are chosen so that exactly $v_1, \ldots, v_n$ are free in t, $x_1 < \ldots < x_n$, and $y = t(x_1, \ldots, x_n)$, whereby we mean $y = t[x_1, \ldots, x_n]$ (the result plugging the x_1 in) but choose the parenthesis notation to emphasize the idea of looking at t as a function. This is to be called a *standard representation* of y in $\mathfrak{H}$(X).

We let $y = t(x_1, \ldots, x_n)$ be standard representation of y and define $\overline{f} = t(f(x_1), \ldots, f(x_n))$. We first have to show that $\overline{f}$ is well-defined. Suppose $t'(z_1, \ldots, z_k)$ is another standard representation of y. In $\mathfrak{H}$(X) we have, then, that $t(x_1, \ldots, x_n) = t'(z_1, \ldots, z_k)$. Take the increasing sequence $u_1 < \ldots < u_i$ to be the listing in increasing order of the set $\{x_1, \ldots, x_n, z_1, \ldots, z_k\}$. We obtain a formula ϕ saying exactly $t(x_1, \ldots, x_n) = t'(z_1, \ldots, z_k)$ in terms of $u_1 < \ldots u_i$, and it is the case that $\mathfrak{H}(X) \models \phi[u_1, \ldots, u_i]$. So by the hypothesis of the theorem, since the elements $u_1, \ldots, u_i$ come from X and are placed by f in Y ordermorphically, we have that $\mathfrak{H}(Y) \models \phi[f(u_1), \ldots, f(u_i)]$. This gives that $t(f(x_1), \ldots, f(x_n)) = t'(f(z_1), \ldots, f(z_k))$ in $\mathfrak{H}$(Y).

Now we consider any formula $\phi(v_1, \ldots, v_n)$ of $\mathcal{L}$, and let $y_1, \ldots, y_l$ be given so that $\mathfrak{H}(X) \models \phi[y_1, \ldots, y_l]$. We now produce standard representations of $y_1, \ldots, y_l$ given by $t_1, \ldots, t_l$, and take a finite sequence of generators $x_1 < \ldots < x_n$ from X. Assume that each t_i when applied to an adequate subsequence of $x_1 < \ldots < x_n$ yields y_i. Now we are in position to find a formula ψ containing the terms $t_1, \ldots, t_l$ and the variables $v_1, \ldots, v_n$ so that the following holds:

$$\mathfrak{H}(X) \models \phi[y_1, \ldots, y_l] \text{ if and only if } \mathfrak{H}(X) \models \psi[x_1, \ldots, x_n].$$

Note that we have by hypothesis, once again, that $\mathfrak{H}(Y) \models \psi[f(x_1), \ldots, f(x_n)]$. By taking a look at the form of chosen ψ, it is easily obtained that $\mathfrak{H}(X) \models \phi[\overline{f}(y_1), \ldots, \overline{f}(y_l)]$. Thus $\overline{f}$ is an isomorphism of models.

Taking $z \in H(\text{range of } f)$, there is a standard representation $z = t(y_1, \ldots, y_k)$ with $y_1, \ldots, y_k$ all in the range of f. Take $x_1 < \ldots < x_k$ in X such that $f(x_i) = y_i$.

It follows that $\overline{f}$ maps $t(x_1, \ldots, x_k)$ onto z. Therefore $\overline{f}$ is onto H(range of f) and the Embedding Theorem follows from Theorem 10.8. Uniqueness is assured since X generated $H(X)$. $\dashv$

Theorem 8.11 (Automorphism Theorem)**.** *Let f be any ordermorphism of X into X. Then f can be extended* ***uniquely*** *to an automorphism of $\mathfrak{H}(X)$ onto $\mathfrak{H}(X)$.*

Proof. Follows directly from Theorem 8.10. $\dashv$

Theorem 8.12 (Realizing and Omitting Types Theorem)**.** *Let Y be a set of indiscernibles in a model $\mathfrak{B}$ with the property that the sets of formulas of $\mathcal{L}$ satisfied by increasing sequences of elements from X and Y are the same. Suppose further that both X and Y are infinite. Then, given any type $\Sigma(v_1, \ldots, v_n)$ of $\mathcal{L}$, the model $\mathfrak{H}(X)$ realizes Σ if and only if $\mathfrak{H}(Y)$ realizes Σ.*

Proof. We take X, Y, $\mathfrak{H}(X)$, and $\mathfrak{H}(Y)$ that satisfy the hypothesis of the theorem. Let the tuple $(z_1, \ldots, z_n)$ of elements from $H(X)$ realize Σ in $\mathfrak{H}(X)$. Suppose also that $z_1, \ldots, z_n$ has standard representations in $\mathfrak{H}(X)$, and we take these standard representations to involve at most the generators $x_1 < \ldots < x_k$ from X. We let any map f that sends $x_1 < \ldots < x_k$ onto $y_1 < \ldots < y_k$ preserving order be given. Notice that $\mathfrak{H}(\{x_1, \ldots, x_k\}) \prec \mathfrak{H}(X)$ by the Subset Theorem (8.8). By the Embedding Theorem (8.10), we get $\mathfrak{H}(\{x_1, \ldots, x_k\}) \cong \mathfrak{H}(\{y_1, \ldots, y_k\})$ via the mapping $\overline{f}$. Thus, the n-tuple of elements $(\overline{f}(z_1), \ldots, \overline{f}(z_n))$ of $\mathfrak{H}(\{y_1, \ldots, y_k\})$ realizes Σ. Since also $\mathfrak{H}(\{y_1, \ldots, y_k\}) \prec \mathfrak{H}(Y)$ by the Subset Theorem (8.8), it must be the case that $(\overline{f}(z_1), \ldots, \overline{f}(z_n))$ also realizes Σ in $\mathfrak{H}$(Y). In the same style one establishes the remaining direction. $\dashv$

Acknowledgments. It is a pleasure to thank Professor Maryanthe Malliaris for the opportunity to explore Model Theory under her supervision this summer and for her enthusiasm in sharing with the author a little bit of her extraordinary knowledge in this fascinating area. It was a great honor. I would also like to thank my mentor, Yordanka Kovacheva, for her patience and dedication in the many hours I spent in her office talking about logic and debugging my tex code. Furthermore, I'd like to express my gratitude to Professor Peter May for once again spearheading one more amazing summer of mathematics for undergraduate students.

References

[1] C. C. Chang and H. Jerome Keisler. Model Theory. Dover Publications. 2012.
[2] Wilfrid Hodges. Building Models by Games. Dover Publications. 2006.
[3] Herbert Enderton. A Mathematical Introduction to Logic. Harcourt/Academic Press. 2001.
[4] H. -D. Ebbinghaus, J. Flum, and W. Thomas. Mathematical Logic. Springer. 1996.