An Even Shorter Model Theory

Justin Bledin

June 4, 2007

These notes were prepared for the UC Berkeley Group in Logic & the Methodology of Science foundations exam. Material is adapted from Chang and Keisler's (1973) *Model Theory* and Wilfrid Hodges' (1997) *A Shorter Model Theory*. Sections 1-2 review basic semantic/syntactic concepts; Section 3 introduces Robinson diagrams and canonical interpretations; Section 4 presents back-and-forth techniques; Section 5 introduces the method of quantifier elimination; Section 6 discusses Compactness and types; Section 7 examines atomic and saturated models; Section 8 presents the theorems of Löwenheim-Skolem; Section 9 examines preservation phenomena; Section 10 presents amalgamations and Fraïsse limits; Section 11 discusses interpolation; Section 12 discusses model-completeness; Section 13 covers further miscellaneous topics.

1 Mathematical Structures & Mappings

Definition 1 *a structure A is an object with these ingredients:*
(1) domain with cardinality $|A| = |dom(A)|$;
(2) constants c naming c^A, elements in A;
(3) n-ary relations R naming R^A, subsets of $(dom(A))^n$;
(4) n-ary functions F naming F^A, mappings from $(dom(A))^n \mapsto dom(A)$.

Examples: Klein 4-group with signature $\mathcal{L} = \langle \{a, b, c\}, \cdot, e\rangle$; ordered naturals with signature $\mathcal{L} = \langle \mathbb{N}, <\rangle$.

Definition 2 *A is a substructure of B if $dom(A) \subseteq dom(B)$ and the identity i:$dom(A) \mapsto dom(B)$ is an embedding (see below). Then (i) $c^A = c^B$; (ii) $R^A = R^B \cap (dom(A))^n$; (iii) $F^A = F^B|(dom(A))^n$.*

$A = \langle Y\rangle_B$ is the substructure of B generated by the hull Y, where Y is said to be a set of generators for A ($|A| \leq |Y| + |L|$; see Hodges, p.8 for proof)

Definition 3 *a homomorphism f is a mapping such that:*
(1) $f(c^A) = c^B$;
(2) $\overline{a} \in R^A \rightarrow f\overline{a} \in R^B$; (e.g. $x <^A y \rightarrow f(x) <^B f(y)$)
(3) $f(F^A(\overline{a})) = F^B(f\overline{a})$. (e.g. $x = y \cdot^A z \rightarrow f(x) = f(y) \cdot^B f(z)$)

Definition 4 *an embedding f has a stronger (2*) $\overline{a} \in R^A \leftrightarrow f\overline{a} \in R^B$.*

Definition 5 *an isomorphism f is a surjective embedding.*

Definition 6 *a homomorphism $f : A \mapsto A$ is an endomorphism of A.*

Definition 7 *an isomorphism $f : A \mapsto A$ is an automorphism of A.*

$$
\begin{array}{c|cccc}
\cdot & e & a & b & c \\
\hline
e & e & a & b & c \\
a & a & e & c & b \\
b & b & c & e & a \\
c & c & b & a & e
\end{array}
\quad \xrightarrow{f} \quad
\begin{array}{c|cccc}
+ & 0 & x & y & z \\
\hline
0 & 0 & x & y & z \\
x & x & 0 & z & y \\
y & y & z & 0 & x \\
z & z & y & x & 0
\end{array}
$$

Figure 1: *$\langle A, \cdot, e\rangle$ and $\langle B, +, 0\rangle$ are isomorphic*

$$
\begin{array}{ccccccc}
0 < & 1 < & 2 < & 3 < & 4 < & 5 < & \ldots \\
\downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow \\
0 \leq & 1 \leq & 2 \leq & 3 \leq & 4 \leq & 5 \leq & \ldots
\end{array}
$$

Figure 2: *$\langle \mathbb{N}, < \rangle$ is homomorphic but not isomorphic to $\langle \mathbb{N}, \leq \rangle$*

Intuitively, if A is homomorphic to B, then the structural properties of A carry over to B. If A is isomorphic to B (we write $A \cong B$), they are essentially the same structure. Here are some nice properties of homomorphisms/embeddings:

(i) if $f : A \mapsto B$ and $g : B \mapsto C$ are homomorphisms/embeddings, then gf is a homomorphism/embedding from $A \mapsto C$;

(ii) The identity map is an automorphism of A;

(iii) if f is an isomorphism from $A \mapsto B$, $g = f^{-1}$ exists and is an isomorphism from $B \mapsto A$;

(iv) if $f : A \mapsto B$ is a homomorphism and $g : B \mapsto A$ is a homomorphism with $g = f^{-1}$, then f and g are both isomorphisms;

(v) every homomorphism $f : A \mapsto C$ can be factored as $f = hg$ for some surjective homomorphism $g : A \mapsto B$ and extension $h : B \mapsto C$;

(vi) every embedding $f : A \mapsto C$ can be factored as $f = hg$ where $g : A \mapsto B$ is an extension and $h : B \mapsto C$ is an isomorphism (so we can assume A is a substructure of C).

2 Language & Interpretations

A formal language $\mathcal{L}$ can be described in stages: first come the *terms*; from these, the basic *atomic formulae* are built; then after considering the whole array of logical symbols, more complex formulae can be expressed in $\mathcal{L}$.

Definition 8 *a term is a variable, constant or function $F(t_1, ..., t_n)$ whose inputs $\{t_i\}$ are terms. If $t(\overline{x})$ is a term in the language $\mathcal{L}$, $t^A(\overline{a})$ is the element in A named by t with $\overline{x}$ interpreted as a name of $\overline{a} \in A$.*

It is sometimes useful to consider an absolutely free structure in which terms describe themselves. If $\mathcal{L}$ is a language and X a set of variables, the *term algebra of $\mathcal{L}$ with basis X* is the $\mathcal{L}$-structure A with $dom(A)$={set of all terms of $\mathcal{L}$ whose variables are in X}, $c^A = c$ for constants in $\mathcal{L}$, and $F^A(t) = F(t)$ for functions in $\mathcal{L}$ and elements t in $dom(A)$ (R^A is empty). For example, if $\mathcal{L}$ is the signature of additive groups with identity element 0 and $X = \{x, y, z\}$, $dom(A) = \{0, 0 + x, x + y, (x + y) + z, ...\}$. Note that elements of the term algebra A are not elements of a particular group but rather terms in the language of groups.

Definition 9 *atomic formulae are strings of symbols built from terms as follows:*
(1) if s, t are terms, $s = t$ is an atomic formulae of $\mathcal{L}$;
(2) if R is an n-ary symbol and $t_1, ..., t_n$ are terms, $R(t_1, ..., t_n)$ is an atomic formulae of $\mathcal{L}$.

The class of all formulae of $\mathcal{L}$ is then generated from the atomic formulae by introducing the logical connectives. If ϕ, ψ are atomic formulae, then $\neg\phi$, $\phi\vee\psi$, $\phi\wedge\psi$, $\forall y\phi$ and $\exists y\phi$ are all formulae of $\mathcal{L}$ and so on. Formulae with only bound occurrences of variables (*i.e.*, containing only closed terms) are called *sentences*.

Languages are typically denoted with two subscripts $\mathcal{L}_{xy}$. The first x says how many formulae we can join together by $\bigvee$ and $\bigwedge$. The second y says how many quantifiers we can put together in a row. A first-order language $\mathcal{L}_{\omega\omega}$ is one in which all formulae are finite and only finitely many quantifiers appear in each formula.

The cardinality of a language $||\mathcal{L}||$ is the least infinite cardinal $\geq$ the number of symbols $|\mathcal{L}|$ in the language (*i.e.*, $||\mathcal{L}|| = \omega \cup |\mathcal{L}|$). This definition has two virtues: (a) $||\mathcal{L}||$ coincides with the number of non-equivalent (under the relation of being variants) formulae expressible in $\mathcal{L}$ (to see this, note that if $|\mathcal{L}|$ is countable, then # terms is countable $\Rightarrow$ # atomic formulae is countable $\Rightarrow$ # formulae is countable; otherwise, if # functions $\{F_i\}$ is uncountable $\Rightarrow$ # of terms is uncountable and if # relations $\{R_i\}$ is uncountable $\Rightarrow$ # atomic formulae is uncountable); (b) what we often really care about is whether a language is countable or not as this will be relevant in certain constructions where we wish to build countable models (such as in the proof of the Omitting Types theorem). If $|\mathcal{L}|$ is finite, it is still common to call $\mathcal{L}$ a finite signature.

In model theory, quantifiers are especially important (see Section 5 on quantifier elimination and Section 9 on preservation phenomena) and we will work with a hierarchy of first-order formulae:

(i) Π^0_0 and Σ^0_0 formulae have all quantifiers bounded where the quantifiers $(\forall x < y)$ and $(\exists x < y)$ are bounded (alternatively, we use the notation $\forall_0$ and $\exists_0$ when no distinction is made between bounded/unbounded quantification so the formulae are quantifier-free);

(ii) a Π^0_{k+1} formula is of the form $\forall x\psi$ where ψ is a Σ^0_k formula (or $\forall_{k+1}$, $\exists_k$ respectively);

(iii) a Σ^0_{k+1} formula is of the form $\exists x\psi$ where ψ is a Π^0_k formula (or $\exists_{k+1}$, $\forall_k$ respectively);

(as x in (ii) and (iii) can be empty, every Π^0_k/Σ^0_k formula is also both a Π^0_{k+1} and Σ^0_{k+1} formula)

Through interpretation of its symbols, a language allows us to talk about mathematical structures and the relations between them. Specifically, a formula $\phi(a)$ in $\mathcal{L}_A$ talks about an element a in the structure A; a sentence in $\mathcal{L}_A$ talks about global structural features of A (*e.g.*, $\forall x\neg(x < 0)$ says there are no negative elements in a structure). A formula ϕ may speak truthfully but it might also talk rubbish. In the former case, $A \models \phi$ and ϕ is true in A; otherwise, $A \not\models \phi$ and ϕ is false in A (for full truth conditions, see Hodges: 12/24). If $A \models T$ where T is a theory (*i.e.*, set of sentences), we say A is a *model* of T.

We write $Mod(T)$ for all models of the theory T in $\mathcal{L}$. If $K = Mod(T)$, we say T axiomatizes the class K of $\mathcal{L}$-structures. Theories T and T' are equivalent if $Mod(T) = Mod(T')$. We write $Th(K)$ for the set of all first-order sentences true in every structure in K. A theory often aims to describe a particular structure or class of structures. For example, here are the first-order Peano axioms:

(i) $\forall x\neg(x + 1 = 0)$

(ii) $\forall x(x + 0 = x)$

(iii) $\forall xy(x + 1 = y + 1 \rightarrow x = y)$

(iv) $\forall xy(x + (y + 1) = (x + y) + 1)$

(v) $\forall x(x \cdot 0 = 0)$

(vi) $\forall xy(x \cdot (y + 1) = x \cdot y + x)$

(vii) $[\phi(0) \wedge \forall x(\phi(x) \rightarrow \phi(x + 1))] \rightarrow \forall x\phi(x)$ (induction schema for all definable sets)

The conditions (i) and (iii) define succession; (ii) and (iv) recursively define addition; (v) and (vi) recursively define multiplication; (vii) is the first-order induction axiom schema (so (vii) fathers an infinite number of axioms). While the intended model of first-order Peano arithmetic (PA) is $\langle\mathbb{N}, 0, 1, +, \cdot\rangle$, we shall see in Section 6 that there are nonstandard models as well (as implied by Gödel's proof of the incompleteness of arithmetic).

Given a language $\mathcal{L}$, certain relations are definable on a structure A:

Definition 10 *a relation R is first-order definable if $R^A = \phi(A^n, \overline{b}) = \{\overline{a} : A \models \phi(\overline{a}, \overline{b})\}$ for some first-order formula ϕ in $\mathcal{L}$ and parameters $\overline{b} \in A$.*

Example: the formula $\phi(p) = \forall x(x|p \rightarrow (x = 1 \vee x = p))$ defines the prime numbers (where '$x|p$' says x divides p). The ordering $x < y$ can also be explicitly defined in PA with $\phi(x, y) = \exists z(x + (z + 1) = y)$. More generally, a set of natural numbers is a Π^0_k (resp. Σ^0_k) relation if it is definable by some Π^0_k (Σ^0_k) formula. R is a Δ^0_k relation if it is both a Π^0_k and Σ^0_k relation.

Application: a structure A is *minimal* if A is infinite but the only first-order definable subsets of $dom(A)$ with parameters are either finite or cofinite. Likewise, a first-order definable set $X \subseteq dom(A)$ is minimal if X is infinite and for every other definable set Y, $X \cap Y$ or $X \setminus Y$ is finite. A is said to be *O-minimal* if $\mathcal{L}$ contains $<$ which linearly orders $dom(A)$ in such a way that any first-order definable subset is a union of finitely many intervals of the form $(a, b), (a), (-\infty, b), (a, \infty)$ for $a, b \in A$. The following lemma is often helpful in analyzing the definable relations on a structure:

Lemma 1 *suppose Y is a definable set on A with parameters from X, then every automorphism σ of A which fixes X pointwise (i.e., $\sigma(a) = a$ for all $a \in X$) fixes Y setwise (i.e., for all $a \in A$, $a \in Y \leftrightarrow \sigma(a) \in Y$).*

Having a language to talk about structures also provides syntactic characterizations of mappings between structures (more of this in Section 5):

Theorem 1 *$f : A \mapsto B$ is a homomorphism $\leftrightarrow$ $A \models \phi(\overline{a}) \rightarrow B \models \phi(f\overline{a})$ for all atomic formulae $\phi \in \mathcal{L}$ and tuples $\overline{a} \in A$.*

Theorem 2 *$f : A \mapsto B$ is an embedding $\leftrightarrow$ $A \models \phi(\overline{a}) \rightarrow B \models \phi(f\overline{a})$ for all literals $\phi \in \mathcal{L}$ and tuples $\overline{a} \in A$.*

This agrees with our earlier definitions of homomorphism and embedding as formula preservation between structures $\Leftrightarrow$ structural equivalence. If A is homomorphic to B, anything an atomic ϕ truthfully says about A (and its elements) will also hold of B as the structural properties of A carry over to B (and conversely if the structures are isomorphic). There is also the purely syntactic condition of elementary equivalence where a first-order language cannot distinguish between structures:

Definition 11 *an elementary embedding f is a mapping which preserves all first-order formulae. Two structures A and B are elementary equivalent, $A \equiv B$, if they agree on all first-order sentences.*

A hierarchy of mappings has now emerged: homomorphisms preserve atomic formulae; embeddings preserve literals; elementary embeddings preserve all formulae. Clearly, $A \cong B \rightarrow A \equiv B$ but the converse does not hold (unless A, B are finite) as first-order languages cannot distinguish between models of differing infinite cardinalities (see Section 8).

Definition 12 *A is an elementary substructure of B, denoted $A \preccurlyeq B$, if $A \subseteq B$ and the inclusion map $i : A \mapsto B$ is an elementary embedding (B is said to be an elementary extension of A). We write $A \prec B$ when A is a proper elementary substructure of B.*

Note that $A \subseteq B$ and $A \equiv B$ does not necessarily imply that $A \preccurlyeq B$ (though the converse clearly holds). Consider when $A = \langle [1, \omega), < \rangle$ and $B = \langle [0, \omega), < \rangle$. Though $f(n) = n - 1$ is an elementary embedding from $A \mapsto B$, the inclusion map $i : A \mapsto B$ is *not* an elementary embedding as $A \models \forall y(1 \leq y)$ while this formula $\phi(1)$ does not hold in B. Thus $A \preccurlyeq B$ is a stronger notion than $A \equiv B$ as the elements common to A and B must have exactly the same first-order properties (with respect to A,B). The following theorems are useful for constructing/working with elementary substructures:

Theorem 3 *(Tarski-Vaught criterion for elementary substructures) $A \preccurlyeq B \Leftrightarrow$ for every formula $\psi \in \mathcal{L}$ and $\overline{a} \in A$, $B \models \exists y \psi(\overline{a}, y) \rightarrow B \models \psi(\overline{a}, d)$ for some $d \in A$.*

Theorem 4 *(Tarski-Vaught theorem on unions of elementary chains) Let $A_0 \preccurlyeq A_1 \preccurlyeq \ldots \preccurlyeq A_n \preccurlyeq \ldots$ be an elementary chain of $\mathcal{L}$-structures. Then for all A_i, $A_i \preccurlyeq \bigcup_i A_i$.*

To form an elementary substructure of A, we can start with some subset $X \subset A$ and close off X under existential claims true in A. Constructing elementary chains is also a powerful method for iterating the notion of elementary equivalence into the transfinite.

3 From Structures to Language and Back Again

(1) Structures $\Rightarrow$ Language (diagrams)

A powerful tool in model theory, a diagram of a structure is formed by naming elements of the structure by new constants (called *parameters*). The language obtained from $\mathcal{L}$ by adding the constants $\overline{c}$ is denoted $\mathcal{L}(\overline{c})$ and $\mathcal{L}(\overline{c}) \supset \mathcal{L}$. We often do not need to name every element in A (though this would do the trick) as if $A = \langle \overline{a} \rangle_A$ for some $\overline{a} \subset A$, then introducing parameters for $\overline{a}$ ensures every element in A is named in $\mathcal{L}(\overline{c})$ (the structure $(A, \overline{a})$ in $\mathcal{L}(\overline{c})$ is considered an *extension* of A). Diagrams can be regarded as generalizations of the multiplication table for a group (see the Klein 4-group table on p.2 above).

Definition 13 *the Robinson diagram of A, $diag(A)$, is the set of all closed literals of $\mathcal{L}(\overline{c})$ which are true in $(A, \overline{a})$ (i.e., the set $\{\phi \in \mathcal{L}(\overline{c}) | (A, \overline{a}) \models \phi\}$). The set of atomic sentence of $\mathcal{L}(\overline{c})$ true in $(A, \overline{a})$ is called the positive diagram of A and denoted $diag^+(A)$. The set of all sentences of $\mathcal{L}(\overline{c})$ true in $(A, \overline{a})$ (i.e., $Th(A, \overline{a})$) is called the elementary diagram of A, $eldiag(A)$.*

Lemma 2 *(Diagram Lemma) Let $(A, \overline{a})$ and $(B, \overline{b})$ be structures in $\mathcal{L}(\overline{c})$, then the following are equivalent:*

(a) for every atomic sentence $\phi \in \mathcal{L}(\overline{c})$, $(A, \overline{a}) \models \phi \to (B, \overline{b}) \models \phi$;

(b) there exists a unique homomorphism $f : \langle \overline{a} \rangle_A \mapsto B$ such that $f\overline{a} = \overline{b}$.

(if (a) holds for all literals in $\mathcal{L}(\overline{c})$, f is an embedding; if (a) holds for all sentences in $\mathcal{L}(\overline{c})$, f is an elementary embedding)

Note that when $A = \langle \overline{a} \rangle_A$ the Diagram Lemma essentially says that A is homomorphic to a reduct of B if $B \models diag^+(A)$. Once we know the diagram of a structure, we know the structure up to isomorphism. As we will see, diagrams are particularly useful in constructive proofs where we want to find extensions/reductions of structures with additional properties. More generally, by converting a structure to language, diagrams give us a way to work with structures in syntactic arguments.

(2) Language $\Rightarrow$ Structures (canonical models)

One of the essential results in logic is the Completeness of first-order logic; namely, that for every set of consistent sentences we can always find a model in which the sentences hold true. Here we explore the basis of one route to this result (Henkin's famous Completeness/Compactness proof; presented in detail in Section 6), constructing the *canonical interpretation* of a set of atomic sentences.

Let T be an *=-closed* set of atomic sentences in $\mathcal{L}$ (so $t = t \in T$ for all closed terms $t \in \mathcal{L}$ and $s = t \in T$ implies $\phi(s) \in T \leftrightarrow \phi(t) \in T$). Define the equivalence relation $\sim$ by: $s \sim t \leftrightarrow s = t \in T$. The relation $\sim$ thus creates equivalence classes $t \backslash \sim$ of closed terms in $\mathcal{L}$ and for the canonical model A, we set $dom(A) = \{t^{\sim} \in t \backslash \sim\}$. Every element in A is named by a closed term of $\mathcal{L}$.

Constants, relations and functions are then defined as:

(i) $c^A = c^{\sim}$;

(ii) $(t_1^{\sim}, ..., t_n^{\sim}) \in R^A \leftrightarrow R(t_1, ..., t_n) \in T$;

(iii) $F^A(t_1^{\sim}, ..., t_n^{\sim}) = (F(t_1, ..., t_n))^{\sim}$.

It follows that T is the set of all atomic sentences true in A as we have:

$$A \models s = t \Leftrightarrow s^A = t^A \Leftrightarrow s^{\sim} = t^{\sim} \Leftrightarrow s = t \in T$$
$$A \models R(t_1, ..., t_n) \Leftrightarrow (t_1^A, ..., t_n^A) \in R^A \Leftrightarrow (t_1^{\sim}, ..., t_n^{\sim}) \in R^A \Leftrightarrow R(t_1, ..., t_n) \in T$$

For example, in the language $\mathcal{L}_G = \langle G, \cdot, e \rangle$ the only atomic sentences are of the form $e = e$, $e \cdot e = e$, $(e \cdot e) \cdot e = e$, etc. so in the canonical model A of $T = \{$all such atomic formulae in $\mathcal{L}_G\}$, $dom(A) = \{e^{\sim}\}$, $c^A = e^{\sim}$ and $e^{\sim} \cdot^A e^{\sim} = e^{\sim}$. Note that if T is a set of atomic formulae that is not =-closed, we can always add sentences to T to form its *=-closure* T' and if A is the canonical model of T', we then have $A \models T$ as $T \subset T'$. We will later see that this construction can also be generalized to more complex theories (providing certain additional conditions are still met).

4 Back and Forth between Structures

Back-and-forth techniques are useful for showing two structures are isomorphic. The origin of this proof strategy allegedly goes back to Cantor who used back-and-forth methods to show that the elements of any two countable dense linear orderings without endpoints could be mapped from one structure to the other and back again. The modern incarnation of the back-and-forth argument is Fraïsse's notion of 'back-and-forth equivalence', introduced in the 1950s and sometimes presented, as in Hodges (1997: 74-81), in terms of Ehrenfeucht-Fraïsse games.

The players are ∀belard and ∃loise, a twelfth-century Parisian logician and his student/lover, the niece of a Notre Dame canon. Given two structures A and B, ∀belard wants to prove A different from B while ∃loise wants to prove them identical. They take turns choosing elements a_i from A and b_i from B at each step of the game, ∀belard freely choosing from either A or B and ∃loise choosing from the opposite structure. At the end of the game with a (countably) infinite number of steps, denoted $EF_\omega(A, B)$, sequences of elements $\overline{a} = \{a_i\}$ and $\overline{b} = \{b_i\}$ have been chosen from A and B respectively with the pair $(\overline{a}, \overline{b})$ known as the final 'play'. The play $(\overline{a}, \overline{b})$ is a win for ∃loise if there exists an isomorphism $f : \langle\overline{a}\rangle_A \mapsto \langle\overline{b}\rangle_B$ between substructures $\langle\overline{a}\rangle_A \subseteq A$ and $\langle\overline{b}\rangle_B \subseteq B$ such that $f\overline{a} = \overline{b}$. In other words, the substructures of A and B generated from the elements $\{a_i\}$ and $\{b_i\}$ must be structurally equivalent.

In this context, Fraïsse's 'back-and-forth equivalence' can now be defined as follows:

Definition 14 *Two structures A and B are back-and-forth equivalent, $A \sim_\omega B$, if ∃loise can always win the game $EF_\omega(A, B)$.*

When ∃loise knows of an isomorphism $f : A \mapsto B$, $EF_\omega(A, B)$ is easily won as she can choose $f(a_i)$ from B when ∀belard chooses the corresponding element a_i from A and $f^{-1}(b_i)$ when ∀belard picks b_i. The game is more interesting when ∃loise knows of no such isomorphism but fortunately there is a useful criterion for determining exactly when two structures are back-and-forth equivalent (and in the countable case, isomorphic):

Definition 15 *A back-and-forth system from A to B is a set of functions J satisfying the following conditions:*

(BF1) each $f \in J$ is an isomorphism $f : \langle\overline{a}\rangle_A \mapsto \langle\overline{b}\rangle_B$
(BF2) J is non-empty
(BF3) $\forall f \in J$ and $c \in A$, there is $g \supseteq f$ such that $g \in J$ and $c \in dom(g)$
(BF4) $\forall f \in J$ and $d \in B$, there is $g \supseteq f$ such that $g \in J$ and $d \in im(g)$

Lemma 3 *A,B are back-and-forth equivalent ↔ there is a back-and-forth system from A to B.*

Lemma 4 *Let A,B be countable structures. Then A,B are isomorphic ↔ they are back-and-forth equivalent.*

Two examples are provided below.

Theorem 5 *If A,B are countable dense linear orderings without endpoints, A,B are isomorphic.*

Proof : In the base case, $\langle\emptyset\rangle_A = \emptyset$ and $\langle\emptyset\rangle_B = \emptyset$ so (BF2) clearly holds. Now assume that $\langle\overline{a}\rangle_A = a_1 < ... < a_n$, $\langle\overline{b}\rangle_B = b_1 < ... < b_n$ and there exists an isomorphism $f : \langle\overline{a}\rangle_A \mapsto \langle\overline{b}\rangle_B$ with $f(\overline{a}) = \overline{b}$. We must show that the 'forth' step holds. Choose $c \in A$ not in $\{a_1, ..., a_n\}$. Either $c < a_i$, $a_i < c\ \forall a_i$, or $a_i < c < a_j$ for some $a_i, a_j \in \{a_1, ..., a_n\}$. In the latter case, the density of B ensures we can find some $d \in B$ such that $f(a_i) < d < f(a_j)$ - let $g = f \cup \{(c, d)\}$. In the previous cases, the infinite extension of B in both directions ensures we can find the required d. An analogous argument works for the 'back' step. $\square$

Theorem 6 *If A,B are countable atomless Boolean algebras, A,B are isomorphic.*

Proof : In the base case, where $\langle\emptyset\rangle_A$ and $\langle\emptyset\rangle_B$ are Boolean algebras with only zero and one elements, let $f(0_A) = 0_B$ and $f(1_A) = 1_B$; so (BF2) holds. Now consider $\langle\overline{a}\rangle_A$ with atoms $a_1, ..., a_n$ and $\langle\overline{b}\rangle_B$ with atoms $b_1, ..., b_n$ and assume there exists an isomorphism $f : \langle\overline{a}\rangle_A \mapsto \langle\overline{b}\rangle_B$ with $f(\overline{a}) = \overline{b}$. We must show that the 'forth' step holds. Choose $c \in A$ not in $\{a_1, ..., a_n\}$. Then the structural identity (what Hodges calls 'isomorphism type') of c over $\langle\overline{a}\rangle_A$ is determined by, for each atom a_i, whether $c \wedge a_i$ equals a_i, 0 or neither. $c \wedge a_i$ is the greatest lower bound of c and a_i so, intuitively, the split concerns whether c lies above a_i, beside a_i or below a_i. Now as B is atomless, we can find some $d \in B$ such that $d \wedge f(a_i) = f(a_i), 0$ or neither $\leftrightarrow c \wedge a_i = a_i, 0$ or neither - let $g = f \cup \{(c, d)\}$. An analogous argument works for the 'back' step. $\square$

5 Quantifier Elimination

This section takes us back to the early days of model theory, the method of quantifier elimination originating in Tarski's Warsaw seminar in the late 1920s. The idea is simple but useful: show that relative to a theory T, all formulae in a first-order language $\mathcal{L}$ are equivalent to Boolean combinations of formulae in an *elimination set* Φ. If Φ contains only quantifier-free formulae, then all formulae in $\mathcal{L}$ are equivalent (mod T) to quantifier-free formulae and T is said to have the property of quantifier elimination. It is important to distinguish between the *method* of quantifier elimination, *i.e.*, finding an elimination set Φ or reducing a formula to a Boolean combination of formulae in Φ, from the *property* of quantifier elimination. In the special cases where the method of quantifier elimination is successful, the result is a condensed description of all complete extensions of T, simplifying the study of definable sets on models of T and usually leading to completeness and decidability proofs (see Hodges 1997: 60-1 for details).

The traditional approach to quantifier elimination (used here) is a heavily syntactic approach (model theorists such as Abraham Robinson later encouraged the use of good structural information about the models of T, rather than syntax, to show T admits quantifier elimination and we will later see an alternative route to quantifier elimination in Section 12). Chang and Kiesler (1997: 49) write: "the method may be thought of as a direct attack on a theory." Though these 'attacks' are not very difficult, they can be quite tedious. Fortunately, the following lemma eases the burden tremendously:

Lemma 5 *Suppose that:*

(i) every atomic formula $\phi \in \mathcal{L}$ is in Φ;

(ii) for every $\phi = \exists x\psi(x, \overline{y})$ in $\mathcal{L}$ with ψ a Boolean combination of formulae in Φ, ϕ is equivalent to a Boolean combination of formulae in Φ with respect to every structure in a class K;

Then Φ is an elimination set for K.

So finding an elimination set reduces to eliminating the quantifier $\exists x$. One thorough example is provided here from the theory of dense linear orderings without endpoints (DLO). To tighten the below proof, I take the completeness of DLO for granted at the onset. An alternative route, taken by Chang and Kiesler (50-4), is to prove DLO admits quantifier elimination and have completeness fall out as an easy consequence.

Theorem 7 *DLO admits quantifier elimination*

Proof : It must be shown that for every formula ϕ in L, there exists a quantifier-free ψ such that DLO$\vdash \phi \leftrightarrow \psi$ (ϕ is equivalent to ψ mod DLO). First consider when ϕ is a sentence so as DLO is complete, either DLO$\vdash \phi$ or DLO$\vdash\sim \phi$. If DLO$\vdash \phi$, then DLO$\vdash \phi \leftrightarrow x_1 = x_1$; if DLO$\vdash\sim \phi$, then DLO$\vdash \phi \leftrightarrow x_1 < x_1$ (note that $x_1 = x_1$ and $x_1 < x_1$ could be replaced here by $\top$ and $\bot$ if our language allowed it).

So suppose ϕ is a formula with free variables $x_1, ..., x_n$. We show the set of atomic formulae Φ: $\{x_i = x_j, x_i < x_j\}$ forms an elimination set for the class of all models of DLO (*i.e.*, ϕ is DLO-equivalent to a Boolean combination of the formulae in Φ). Define an *arrangement* of $x_1, ..., x_n$ to be the finite conjunction of formulae $\Theta = \theta_1 \wedge \theta_2 \wedge ... \wedge \theta_n$ where each θ_i is of the form $x'_i < x'_{i+1}$ or $x'_i = x'_{i+1}$ for a renumbering of $x_1, ..., x_n$ as $x'_1, ..., x'_n$. Each arrangement gives us an exact ordering of the n free variables. Now when $n = 1$, an open formula ϕ is built from Φ and DLO$\vdash x_1 = x_1$ and DLO$\nvdash x_1 < x_1$ so we must have DLO$\vdash \phi$ (in which case DLO$\vdash \phi \leftrightarrow x_1 = x_1$) or DLO$\vdash\sim \phi$ (in which case DLO$\vdash \phi \leftrightarrow x_1 < x_1$). So consider when $n > 1$. If DLO$\cup\{\phi\}$ is inconsistent, DLO$\vdash \phi \leftrightarrow x_1 < x_1$. If DLO$\cup\{\phi\}$ is consistent with ϕ open, ϕ must be DLO-equivalent to a disjunction of finitely many arrangements of $x_1, ..., x_n$ as this exhausts what we can say about the variables in ϕ in terms of Boolean combinations of the atomic formulae in Φ (*i.e.*, DLO$\vdash \phi \leftrightarrow \bigvee \Theta$). As an example, the formula $\phi(x_1, x_2, x_3) = (x_1 < x_3) \wedge (x_2 < x_3)$ is equivalent (mod DLO) to $\bigvee \Theta = ((x_1 = x_2) \wedge (x_2 < x_3)) \vee ((x_1 < x_2) \wedge (x_2 < x_3)) \vee ((x_2 < x_1) \wedge (x_1 < x_3))$ as only these three possible arrangements of x_1, x_2, x_3 are compatible with ϕ.

We can now use this condition to show that *every* formula in L is DLO-equivalent to an open formula. Here, it suffices to show that if ϕ is an open formula, $\exists x_n \phi$ is DLO-equivalent to an open formula. By the above, DLO$\vdash \exists x_n \phi \leftrightarrow \exists x_n \bigvee \Theta$ so DLO$\vdash \exists x_n \phi \leftrightarrow \bigvee \exists x_n \Theta$. For each arrangment Θ, let Θ^* be the restricted arrangement of $x_1, ..., x_{n-1}$ obtained by omitting the literal involving x_n. Now the crucial step is this: given the extension axioms ((iv) and (v) above) for DLO, telling us that an element can be found in any possible interval over the ordering, $\exists x_n \Theta$ is equivalent (mod DLO) to Θ^*. Thus, DLO$\vdash \exists x_n \phi \leftrightarrow \bigvee \Theta^*$ where $\bigvee \Theta^*$ is a Boolean combination of the formulas in Φ so DLO admits quantifier elimination. $\square$

In the first part of this proof, the completeness of T (*i.e.*, that $T \vdash \phi$ or $T \vdash\sim \phi$ for every sentence $\phi \in \mathcal{L}$, or alternatively, that any models of T are elementarily equivalent) allowed us to swiftly deal with the case where ϕ is a sentence. While complete theories do not necessarily have the property of quantifier elimination (*e.g.*, the theory of dense linear orderings *with* endpoints), it is generally easy to show that an incomplete theory does not have this property by finding a sentence with quantifiers that is not a consequence of the theory. Slight modifications of this proof (in particular, regarding the elimination of the final existential quantifier) can also be used to show that other complete theories admit quantifier elimination, such as the theory of atomless Boolean algebras and the theory of an equivalence relation with infinitely many infinite classes.

6 Into the Infinite

6.1 The Compactness Theorem

Earlier in Section 3, we constructed the canonical interpretation of a set of atomic sentences. This model-existence argument can be generalized to an arbitrary set of sentences satisfying certain properties:

Definition 16 *A Hintikka set for $\mathcal{L}$ is a theory with the following properties:*

(H1) for every atomic sentence $\phi \in \mathcal{L}$, $\phi \in T \rightarrow\sim \phi \notin T$;

(H2) for every closed term $t \in \mathcal{L}$, the sentence $t = t \in T$;

(H3) for closed terms $s, t \in \mathcal{L}$ and atomic formula ϕ, $s = t \in \mathcal{L} \rightarrow (\phi(s) \in T \leftrightarrow \phi(t) \in T)$;

(H4) $\sim\sim \phi \in T \rightarrow \phi \in T$;

(H5) $\bigwedge \Phi \in T \rightarrow \Phi \subseteq T$; $\sim \bigwedge \Phi \in T \rightarrow \exists \psi \in \Phi$ such that $\sim \psi \in T$;

(H6) $\bigvee \Phi \in T \rightarrow \exists \psi \in \Phi$ such that $\psi \in T$; $\sim \bigvee \Phi \in T \rightarrow\sim \Phi \subseteq T$;

(H7) $\exists x \phi(x) \in T \rightarrow \phi(t) \in T$ for some closed term $t \in \mathcal{L}$; $\sim \exists x \phi(x) \in T \rightarrow\sim \phi(t) \in T$ for every closed term $t \in \mathcal{L}$.

Theorem 8 *If T is a Hintikka set for $\mathcal{L}$, then the canonical model A of the set of atomic sentences in T is a model of T.*

The conditions (H1)-(H7) ensure that for every sentence $\phi \in \mathcal{L}$, $A \models \phi \leftrightarrow \phi \in T$ (so $A \models T$). (H2) and (H3) are the conditions from Section 3 which ensure T is =-closed and the closed terms in $\mathcal{L}$ can be divided into equivalence classes based on whether they are equal relative to T, allowing us to ignore redundant closed terms when constructing A. (H1)/(H4) ensure that negation does not lead to any inconsistencies in T and (H5)/(H6) analogously ensure that conjunction and disjunction do not lead to problems. (H7) is known as the 'Henkin property' and the closed terms t *witness* the existential claims.

Theorem 9 *Let T be a theory in $\mathcal{L}$ such that:*

(i) T is consistent;

(ii) for every sentence $\phi \in \mathcal{L}$, either $\phi \in T$ or $\sim \phi \in T$;

(iii) for every sentence $\exists x \psi(x) \in T$ there is a witness $t \in \mathcal{L}$ such that $\psi(t) \in T$;

Then T is a Hintikka set for $\mathcal{L}$.

While we had that for a Hintikka set T, $A \models T$ where A is the canonical interpretation of the set of atomic sentences in T, we now have the reverse: $Th(A)$, the complete theory of A which clearly satisfies (i)-(iii) above, is a Hintikka set. But we now also have a blueprint for constructing Hintikka sets as given a consistent set of sentences in $\mathcal{L}$, we must only expand this set to a maximally consistent set with the Henkin property. Such a construction underlies Henkin's (1949) proof of the Completeness/Compactness of first-order logic:

Theorem 10 *(Completeness Theorem for first-order logic) Let T be a set of sentences in $\mathcal{L}$. Then T is consistent $\leftrightarrow$ T has a model.*

Proof : The direction ($\Leftarrow$) is trivial. To prove the direction ($\Rightarrow$), we expand T to a Hintikka set T^+ and consider the canonical interpretation of T^+. First add new constants $\{c_i | i < ||\mathcal{L}||\}$ to $\mathcal{L}$ to act as witnesses, forming $\mathcal{L}^+$, and enumerate the sentences in $\mathcal{L}^+$ as $\{\phi_i | i < ||\mathcal{L}||\}$. We then define an increasing chain of theories $\{T_i | i < ||\mathcal{L}||\}$ in $\mathcal{L}^+$ as follows:

Set $T_0 = T$. Then let $T'_{i+1} = T_i \cup \{\phi_{i+1}\}$ if this set is consistent; otherwise, set $T'_{i+1} = T_i$. If ϕ_{i+1} is of the form $\exists x \psi(x)$, then take the earliest witness c_j not already used in T'_{i+1} and set $T_{i+1} = T'_{i+1} \cup \{\psi(c_j)\}$; otherwise, put $T_{i+1} = T'_{i+1}$. At limit ordinals δ, take $T_\delta = \bigcup_{i<\delta} T_i$. If we then take the union $T^+ = \bigcup_{i<||\mathcal{L}||} T_i$, then by construction T^+ is consistent, maximal and has the Henkin property so T^+ is a Hintikka set. It follows that the canonical interpretation A of the set of atomic sentences in T^+ models T^+ so as $T \subseteq T^+$, $A|\mathcal{L} \models T$. $\Box$

Theorem 11 *(Compactness Theorem for first-order logic) Let T be a set of sentences in $\mathcal{L}$. Then T is satisfiable $\leftrightarrow$ T is finitely satisfiable.*

Proof : If T is finitely satisfiable, then by Completeness every finite subset of T is consistent so T is consistent. By Completeness again, T has a model. The other direction is trivial. $\Box$

Henkin's proof is our first example of constructing models from constants. Similar constructions are used in the proofs of the Omitting Types theorem (next subsection) and Craig Interpolation theorem (Section 9). To see that Compactness fails in infinite languages, consider the theory $T = \{\bigvee_{0<i<\omega} c_0 = c_i, c_0 \neq c_1, c_0 \neq c_2, c_0 \neq c_3, ...\}$. While T is finitely satisfiable, T does not have a model.

We end this subsection with various easy corollaries/applications of Compactness (though we leave the Upward Löwenheim-Skolem theorem for Section 8):

Theorem 12 *If T is a first-order theory and $T \vdash \phi$ for some sentence ϕ, then $U \vdash \phi$ for some finite subset $U \subseteq T$.*

Proof : Suppose $U \not\vdash \phi$ for every finite subset $U \subseteq T$. Then $T \cup \{\sim \phi\}$ is finitely satisfiable so by Compactness is satisfiable. Thus $T \not\vdash \phi$. $\Box$

Theorem 13 *If T is a first-order theory with arbitrary large finite models, then T has an infinite model. If ϕ is a formula in $\mathcal{L}$ such that for every $n < \omega$, T has a model A with $|\phi(A)| \geq n$, then T has a model B for which $|\phi(B)|$ is infinite.*

Proof : For the first part, add new constants to $\mathcal{L}$ and consider $T' = T \cup \{c_i \neq c_j | i < j < \omega\}$. T' is finitely satisfiable so by Compactness is satisfiable. Thus there exists an infinite model $A \models T'$ and $A|\mathcal{L} \models T$. For the second part, consider $T' = T \cup \{\phi(c_i) \wedge \phi(c_j) \wedge c_i \neq c_j | i < j < \omega\}$. $\Box$

Theorem 14 *The class of infinite sets is not first-order definable.*

Proof : Assume it was by some sentence $\phi \in \mathcal{L}$ so $\sim \phi$ defines the class of finite sets. Now consider $T = \{\sim \phi\} \cup \{\exists_{\geq n} x(x = x) | n \leq \omega\}$ which is finitely satisfiable so by Compactness is satisfiable, a contradiction. $\Box$

Theorem 15 *Let T be a theory in the language of fields with models of arbitrarily high finite characteristics. Then T has a model which is a field of characteristic 0.*

Proof : Let T_f be the theory of fields and consider $T' = T \cup T_f \cup \{p1 \neq 0 | \text{all primes } p\}$. $\Box$

6.2 Realizing and Omitting Types

While our discussion in the previous subsection was restricted to theories, or sets of sentences in $\mathcal{L}$, similar results hold for formulae with free variables. Given a set $\Sigma(\overline{x})$ of open formulae, we can ask whether there is a model A and tuple $\overline{a} \in A$ such that $A \models \Sigma(\overline{a})$. And so, we have types.

Definition 17 *Let $\mathcal{L}$ be a first-order language, A an $\mathcal{L}$-structure with $B \subseteq A$, $\overline{b}$ a sequence listing the elements of B and $\overline{a} \in A$. The complete type of $\overline{a}$ over B w.r.t A, denoted $tp_A(\overline{a}/B)$, is the set of all formulae $\phi(\overline{x}, \overline{b}) \in \mathcal{L}(\overline{b})$ such that $A \models \phi(\overline{a}, \overline{b})$. More generally, a set of formulae $p(\overline{x})$ is a complete type w.r.t A if it is the complete type of some tuple w.r.t some elementary extension of A.*

When we talk of the 'complete type of $\overline{a}$ over X', we refer to everything that can be said about the particular element $\overline{a}$ in terms of X. When we talk of a 'complete type over X', full stop, we refer to everything we can say about some *possible* tuple in terms of X, whether the tuple lies in A or only in some elementary extension of A.

Definition 18 *A type over X w.r.t A is a subset of a complete type over X. A type is an n-type if it has just n free variables. The sets of complete n-types over X w.r.t A are denoted $S_n(X; A)$ and are known as the Stone spaces of A.*

Definition 19 *A type of a theory T is a set Φ of formulae such that $T \cup \{\exists\overline{x} \bigwedge \Psi(\overline{x})\}$ is consistent for every finite $\Psi \subseteq \Phi$ (if T is complete, $T \vdash \exists\overline{x} \bigwedge \Psi(\overline{x})$ for all Ψ). A complete type of T is a maximal type of T.*

For example, in $\mathcal{L}_{\mathbb{N},+,<}$ the type $\Phi(x) = \{0 < x, 1 < x, x = 1+1, x < 1+1+1\}$ is a 1-type in $A = \langle \mathbb{N}, +, < \rangle$ of 'the number 2' (if Φ were expanded to include all formulae in $\mathcal{L}_{\mathbb{N},+,<}$ which hold of '2' in A, then Φ would be a complete type w.r.t A). If a type $\Phi \subseteq tp_A(\overline{a}/X)$, then Φ is *realized* by the tuple $\overline{a} \in A$. If Φ is not realized by any tuple in A, then A *omits* Φ. Now consider Peano arithmetic and let $\Phi(x) = \{x \neq 0, x \neq 1, x \neq 1+1, x \neq 1+1+1, ...\}$. Clearly, Φ is not realized in the standard model $\mathbb{N}$ of PA but is Φ realized in other models of PA? To answer this, we have the following corollary from Compactness:

Theorem 16 *Φ is a type w.r.t $A \leftrightarrow \Phi$ is finitely realized in A.*

Proof : ($\Rightarrow$) Let $\Psi \subseteq \Phi$ be a finite subset of Φ. Then as Φ is a type w.r.t A, there is an elementary extension $B \succcurlyeq A$ and tuple $\overline{b} \in B$ such that $B \models \Psi(\overline{b})$. But now $B \models \exists x\Psi(\overline{x})$ so $A \models \exists x\Psi(\overline{x})$ and Φ is finitely realized in A.
($\Leftarrow$) Now suppose Φ is an n-type finitely realized in A and consider $T = eldiag(A) \cup \Phi(\overline{c})$. T is finitely satisfiable so by Compactness is satisfiable by some model B. As $B \models eldiag(A)$, $A \preccurlyeq B|\mathcal{L}$ by the elementary diagram lemma and the tuple $\overline{b} = (c_0^B, ..., c_n^B)$ realizes Φ. Thus Φ is a type w.r.t A. $\square$

It can also be shown that if Φ is finitely realized in A, then Φ can be extended to a complete type w.r.t A. But returning to our example, $\Phi(x) = \{x \neq 0, x \neq 1, x \neq 1+1, x \neq 1+1+1, ...\}$ is finitely realized in the standard model $\mathbb{N}$ of PA so by the above theorem, there must be some elementary extension $M \succcurlyeq \mathbb{N}$ in which Φ is realized. The model M is called a *nonstandard* model of Peano arithmetic.

Here are some more examples/applications of realizing types:

Theorem 17 *There exist nonstandard primes.*

Proof : Let $\phi(x)$ be a formula in $\mathcal{L}_{\mathbb{Z}}$ expressing that x is a prime and consider the set of formulae $\Phi = \{\phi(x), x > 0, x > 1, x > 2, ...\}$. As Φ is finitely realized in $A = \langle \mathbb{Z}, +, \cdot, 0, 1, < \rangle$, Φ is a type w.r.t A and is realized by some element $\notin \mathbb{Z}$ in some elementary extension of A. $\square$

Example: (Dedekind cuts) A nice historic example, let us describe the complete 1-types of the ordered rationals, *i.e.*, the space $S_1(\mathbb{Q}; A)$ for $A = \langle \mathbb{Q}, < \rangle$. Consider a 1-type $p(v) \in S_1(\mathbb{Q}; A)$. $p(v)$ is uniquely identified with a *cut* in the rationals: $p = \bigwedge_{q \in L_p}[q < v] \wedge \bigwedge_{q' \in U_p}[v < q']$ where

$L_p = \{q \in \mathbb{Q} : q < v \in p\}$ and $U_p = \{q' \in \mathbb{Q} : v < q' \in p\}$ (if v is a rational, we add $[v = q'']$ for some $q'' \in \mathbb{Q}$). The identification is unique since a cut is completely determined using only atomic formulae and A has quantifier elimination. $S_1(A)$ then contains the following 1-types:

(i) for all $q \in \mathbb{Q}$, the unique types p_q containing $v = q$;

(ii) $p_{+\infty}$ where $L_p = \mathbb{Q}$ and $U_p = \emptyset$

$p_{-\infty}$ where $L_p = \emptyset$ and $U_p = \mathbb{Q}$;

(iii) for all $r \in \mathbb{R}\backslash\mathbb{Q}$, the unique types p_r where $L_p = \{q \in \mathbb{Q} : q < r\}$ and $U_p = \{q' \in \mathbb{Q} : r < q'\}$.

As there are as many types $p \in S_1(A)$ as there are real numbers, $|S_1(A)| = 2^\omega$.

Omitting types is more difficult. As Professor Leo Harrington remarked in lecture: "any idiot can realize types but it takes a *real* model theorist to omit them". In analyzing when a theory T has a model which omits particular types, the central idea will be that of T *locally realizing* a type.

Definition 20 *Given a set of formulae $\Phi \in \mathcal{L}$, T locally realizes Φ iff there exists a formula $\theta \in \mathcal{L}$ such that $T \cup \{\exists\overline{x}\theta\}$ has a model (or if T is complete, $T \vdash \exists\overline{x}\theta$) and for every formula $\phi \in \Phi$, $T \vdash \forall\overline{x}(\theta \rightarrow \phi)$. We call θ a support of Φ over T (or when $\theta \in \Phi$, we say θ generates Φ) and call Φ a supported type (or principal type when Φ has a generator). A formula θ is complete (for T) if it generates a complete type.*

The main results here explicate the relationship between local realization of a set of formulae Φ by T and the realization of Φ in every model of T. One direction is simple (when T is complete) while the other, characterized in the Omitting Types theorem, requires a construction similar to that used by Henkin in proving Completeness/Compactness:

Theorem 18 *Let T be a complete theory. T locally realizes $\Phi \rightarrow$ for every model $A \models T$, $A \models \Phi$.*

Proof : As T is complete and locally realizes Φ, there is some formula $\theta \in \mathcal{L}$ such that $T \vdash \exists\overline{x}\theta$ and $T \vdash \forall\overline{x}(\theta \rightarrow \phi)$ for all $\phi \in \Phi$. Now if $A \models T$, $A \models \exists\overline{x}\theta$ so $A \models \theta(\overline{a})$ for some $\overline{a} \in A$. But we also have $A \models \forall\overline{x}(\theta \rightarrow \phi)$ so $A \models \theta(\overline{a}) \rightarrow \phi(\overline{a})$ for all $\phi \in \Phi$. Thus $A \models \Phi(\overline{a})$ as desired. $\square$

Theorem 19 *(Countable Omitting Types theorem) Let T be a consistent theory in a countable language $\mathcal{L}$ and let Φ be a set of formulae in $\mathcal{L}$. T locally omits $\Phi \rightarrow T$ has a countable model which omits Φ.*

Proof : As in Henkin's Completeness proof, add new constants $\{c_i | i < \omega\}$ to $\mathcal{L}$, forming $\mathcal{L}^+$, and enumerate the sentences in $\mathcal{L}^+$ as $\{\phi_i | i < \omega\}$. Similar to the previous construction, we then define an increasing chain of theories $\{T \cup T_i | i < \omega\}$ in $\mathcal{L}^+$ as follows:

Set $T_0 = \emptyset$ so $T \cup T_0 = T$. As before, let $T'_{i+1} = T_i \cup \{\phi_{i+1}\}$ if $T \cup T_i \cup \{\phi_{i+1}\}$ is consistent; otherwise, set $T'_{i+1} = T_i$. If ϕ_{i+1} is of the form $\exists x\psi(x)$, then take the earliest witness c_j not already used in T'_{i+1} and set $T_{i+1} = T'_{i+1} \cup \{\psi(c_j)\}$; otherwise, put $T_{i+1} = T'_{i+1}$. But there is now an additional crucial step: given T_i, write $\bigwedge T_i$ as a sentence $\chi(\overline{c}, \overline{d})$ with $\chi \in \mathcal{L}$ and $\overline{d}$ the distinct witnesses that occur in T_i but not in $\{\overline{c}\}$. Then $T \cup \{\chi\}$ is consistent so $T \cup \{\exists\overline{x}\overline{y}\chi\}$ has a model and as T locally omits Φ, there is some formula $\phi \in \Phi$ such that $T \cup \{\chi\} \cup \{\sim \phi\}$ is consistent. Take $\sim \phi(\bar{c})$ and put it in T_{i+1}.

If we then take the union $T^+ = T \cup \bigcup_{i<\omega} T_i$, then by construction T^+ is consistent, maximal and has the Henkin property so T^+ is a Hintikka set. It follows that the canonical interpretation A of the set of atomic sentences in T^+ models T^+ so as $T \subseteq T^+$, $A|\mathcal{L} \models T$. Further, for each tuple of distinct witnesses $\bar{c}$, there is a formula $\phi \in \Phi$ such that $\sim \phi(\bar{c}) \in T^+$. As every tuple in A is named by these witnesses (consider sentences $\exists x(x = t)$ where t is a closed term), $A|\mathcal{L}$ omits Φ (note also that we constructed A by adding only countably many new constants $\{c_i\}$ to $\mathcal{L}$ so $A|\mathcal{L}$ is a countable model). $\square$

While the Omitting Types theorem was presented with only one type Φ, it is easily extended to countably many types $\{\Phi_i | i < n\}$. We simply run the same argument but at each incremental step, we witness a formula $\sim \phi_i$ for each Φ_i. However, the Omitting Types theorem fails for sets of formulae with infinitely many free variables as consider the complete theory DLO and $\Phi = \{x_0 < x_1, x_1 < x_2, x_2 < x_3, ...\}$. Though DLO has no model which omits Φ, DLO *does* locally omit Φ as if DLO $\cup\ \theta(x_0, ..., x_n)$ is consistent, then DLO $\cup\ \theta(x_0, ..., x_n) \cup \sim (x_{n+1} < x_{n+2})$ is consistent as well. Moreover, the Omitting Types theorem fails for uncountable languages as let $T = \{a_i \neq a_j | i < j < \omega_1\}$ be a theory in the language $\mathcal{L}(c, d)$ with constants $\{c_i | i < \omega_1\} \cup \{d_i | i < \omega\}$ and consider $\Phi = \{x \neq d_i | i < \omega\}$. Though T locally omits Φ, no model of T omits Φ as every model of T has uncountably many elements.

But back to the case where we are working with countable language $\mathcal{L}$ and type Φ with only finitely many free variables, we can now combine the above theorems to give a necessary and sufficient condition for T to have a model omitting Φ:

Theorem 20 *Let T be a consistent theory in a countable language $\mathcal{L}$. T has a countable model which omits $\Phi \leftrightarrow T$ has a complete extension which locally omits Φ.*

And so we now know how to realize and omit types. As we will see in the following section, this will be particularly useful for systematically investigating the countable models of a complete theory. For the time being, we will assume that $\mathcal{L}$ is a countable language.

7 Big and Small Countable Models

7.1 Atomic Models

We begin with the small countable models:

Definition 21 *A is an atomic model if for every $\overline{a} \in A$, $tp_A(\overline{a})$ is principal (i.e., there exists some formula $\theta \in \mathcal{L}$ which generates $tp_A(\overline{a})$ relative to $Th(A)$ and $A \models \theta(\overline{a})$). A is a prime model if A is elementarily embedded in every model of $Th(A)$.*

Definition 22 *a theory T is atomic iff for every $\phi \in \mathcal{L}$ such that $T \cup \{\phi\}$ is consistent, ϕ is generated by a complete formula in T (i.e., ϕ lies in some principal type Φ of T).*

Examples: finite models, the standard model of PA, DLO.

Theorem 21 *(Existence Theorem for Atomic Models) Let T be a complete theory.*
T has a countable atomic model $\leftrightarrow$ T is atomic.

Proof : ($\Rightarrow$) Let $A \models T$ be a countable atomic model and consider a formula $\phi \in \mathcal{L}$ such that $T \cup \{\phi\}$ is consistent. As T is complete, $T \vdash \exists\overline{x}\phi$ so $A \models \phi(\overline{a})$ for some $\overline{a} \in A$. But A is atomic so $tp_A(\overline{a})$ is principal and $A \models \theta(\overline{a}) \to \phi(\overline{a})$ for a complete formula θ. Thus $T \vdash \theta \to \phi$ so T is atomic. ($\Leftarrow$) Assume T is atomic and consider the sets $\Phi_n = \{\sim \theta_i(x_1, ..., x_n) | i < \omega\}$ for $n < \omega$ where $\{\theta_i\}$ are the complete formulae w.r.t T. If some formula $\phi(x_1, ..., x_n) \in \mathcal{L}$ is consistent with T, then as T is atomic, $T \cup \{\phi\} \cup \{\sim\sim \theta_i\}$ is consistent for some complete θ_i. But then T locally omits Φ_n so by the Omitting Types theorem, there is a countable model $A \models T$ which omits Φ_n. But then we have that for each $\overline{a} \in A$, $A \models \theta_i(\overline{a})$ for some complete formulae θ_i. Thus T has a countable atomic model A. [*note:* instead of requiring that T is atomic, we could have here used the equivalent condition that for every $n < \omega$, $|S_n(T)|$ is at most countable (recall that $S_n(T)$ is the set of complete n-types of T) as by the Omitting Types theorem, the at most countably many non-principal complete types can be omitted] $\square$

Theorem 22 *(Uniqueness Theorem for Atomic Models)*
A, B countable atomic models and $A \equiv B \to A \cong B$.

Proof : We show a back-and-forth system exists from A to B. In the base case, as $A \equiv B$, there is an isomorphism $f : \langle\emptyset\rangle_A \mapsto \langle\emptyset\rangle_B$. Now let $\langle\overline{a}\rangle_A$ and $\langle\overline{b}\rangle_B$ be finitely generated substructures of the countable atomic models A, B respectively and assume that there exists an isomorphism $f : \langle\overline{a}\rangle_A \mapsto \langle\overline{b}\rangle_B$ with $f(\overline{a}) = \overline{b}$. This implies $(A, \overline{a}) \equiv (B, \overline{b})$ so $\overline{a}$ and $\overline{b}$ realize the same complete n-type in A, B. We must show the 'forth' step holds so choose $c \in A$ not in $\overline{a}$. As A is atomic, $tp_A(\overline{a}c)$ is principal so has a generator $\theta(\overline{x}, y) \in \mathcal{L}$. But now $tp_A(\overline{a}) = tp_B(\overline{b})$ and $A \models \exists y\theta(\overline{a}, y)$ so $B \models \exists y\theta(\overline{b}, y)$ and for some $d \in B$, $B \models \theta(\overline{b}, d)$. The tuples $\overline{a}c$ and $\overline{b}d$ thus realize the same complete $n+1$-type in A, B so $(A, \overline{a}c) \equiv (B, \overline{b}d)$ and there is an isomorphism $f : \langle\overline{a}c\rangle_A \mapsto \langle\overline{b}d\rangle_B$ with $f(\overline{a}c) = \overline{b}d$ as desired. An analogous argument works for the 'back' step so as $A \sim_\omega B$, $A \cong B$. $\square$

Theorem 23 *A is a countable atomic model $\leftrightarrow$ A is a prime model.*

Proof : ($\Rightarrow$) Let A be a countable atomic model and let $B \models Th(A)$. Add new constants to the language $\mathcal{L}$ and list the elements of A as $\{a_i | i < \omega\}$. Using the argument in the previous proof, we can find elements $\{b_i | i < \omega\} \in B$ such that $(A, a_i) \equiv (B, b_i)$ for all $i < \omega$. Then by the elementary diagram lemma, the mapping $f(a_i) = b_i$ is an elementary embedding of A into B.
($\Leftarrow$) Assume A is not atomic so there is some $\overline{a} \in A$ such that $tp_A(\overline{a})$ is not principal. By the Omitting Types theorem, there is $B \models Th(A)$ which omits $tp_A(\overline{a})$. But then A cannot be elementarily embedded in B as for some formula $\phi \in tp_A(\overline{a})$, $A \models \phi(\overline{a})$ for some $\overline{a} \in A$ while $B \not\models \exists \overline{x} \phi(\overline{x})$. Thus A is not a prime model. [*note:* a prime model must be countable as it is elementarily embedded in every countable model of T] $\square$

Let us summarize: we are interested in the countable models of a complete theory T and the first appearance is made by the small atomic models, structures in which each element satisfies a principal complete type. As T is complete, none of these types can be omitted in any countable model of T so they are here to stay. The results stated above should then come as no surprise as, intuitively, the unique atomic model of T - the model containing the bare minimum of what must be included in any model of T - must be elementarily embedded in every countable model of the theory. In a sense, prime models are like prime numbers: they cannot be factored into pieces but are rather the raw, indestructible components of all countable models.

7.2 ω-Saturated Models

We now turn to the big countable models:

Definition 23 *A is ω-saturated if for every finite subset $X \subset A$, all complete 1-types over X w.r.t A in $\mathcal{L}(X)$ consistent with $Th(A)$ are realized by elements in A [note: this immediately extends to all complete n-types]. A is a countable universal model if A is countable and every countable model $B \equiv A$ is elementarily embedded in A.*

The restriction to finite subsets $X \subset A$ is important here as if X were allowed to be countable, we have the type $\Phi = \{x \neq a_1, x \neq a_2, x \neq a_3, ...\}$ where $\{a_i | i < \omega\}$ lists the elements in A. In this case, the very notion of a countable ω-saturated model becomes impossible.

Examples: finite models, the ordered rationals. For a better example clearly distinguishing atomic and ω-saturated models, consider a language with only constants $\{c_i | i < \omega\}$ and the theory $T = \{c_i \neq c_j | i < j < \omega\}$. There are (up to isomorphism) countably many models of T as for each $n < \omega$, there is a model with exactly n elements that are non-constants. The atomic model has zero non-constants while the saturated model has ω non-constants.

Theorem 24 *(Existence Theorem for ω-Saturated Models) Let T be a complete theory. T has a countable ω-saturated model $\leftrightarrow$ for each $n < \omega$, $|S_n(T)|$ is countable.*

Proof : ($\Rightarrow$) Let $A \models T$ be a countable ω-saturated model so all complete n-types of T in $\mathcal{L}$ are realized in A. As each tuple $\overline{a} \in A$ realizes only one complete type and A is countable, $|S_n(T)|$ must be at most countable for all $n < \omega$.
($\Leftarrow$) The argument is similar to Henkin's Completeness/Compactness proof. Without getting into the details, we can construct a chain of theories $T = T_0 \subset T_1 \subset T_2 \subset ...$ in an expanded language $\mathcal{L}(\overline{c})$ such that the union of the chain $T^+ = \bigcup_{i<\omega} T_i$ is a Hintikka set and the canonical model $A|\mathcal{L} \models T$ realizes all complete types over finite subsets of $A|\mathcal{L}$ (see Chang & Keisler: 98-9). The construction of a countable canonical model is possible as $|S_n(T)|$ is at most countable for each $n < \omega$ (so if $\mathcal{L}(\overline{d})$ is the language obtained by adding $m < \omega$ new constants $\overline{d} \subset \overline{c}$, then since the types $\Phi_{\mathcal{L}(\overline{d})}(x_1, ..., x_n, d_1, ..., d_m)$ of T in $\mathcal{L}(\overline{d})$ are in one-to-one correspondence with the types $\Phi_{\mathcal{L}}(x_1, ..., x_{n+m})$ of T in $\mathcal{L}$, there are at most countably many types $\Phi_{\mathcal{L}(\overline{d})}$ as well). $\Box$

Theorem 25 *Let T is a complete theory with countably many nonisomorphic countable models. Then T has a countable ω-saturated model.*

Proof : Each complete n-type of T is realized in a countable model of T (the countability of $\mathcal{L}$ is essential here) and there are only countably many nonisomorphic countable models of T so T has only countably many complete n-types. By the above theorem, T has a countable ω-saturated model. $\Box$

Theorem 26 *(Uniqueness Theorem for ω-Saturated Models) A, B countable ω-saturated models and $A \equiv B \rightarrow A \cong B$.*

Proof : The proof is similar to the uniqueness proof for countable atomic models. The previous proof was driven by the result that for A an atomic model and B an $\mathcal{L}$-structure, if $(A, \overline{a}) \equiv (B, \overline{b})$ and $c \in A$, there is an element $d \in B$ such that $(A, \overline{a}c) \equiv (B, \overline{b}d)$. Analogously, we now have that for A an ω-saturated model and B an $\mathcal{L}$-structure, if $(A, \overline{a}) \equiv (B, \overline{b})$ and $d \in B$, there is an element $c \in A$ such that $(A, \overline{a}c) \equiv (B, \overline{b}d)$. This holds as $tp_A(\overline{a}) = tp_B(\overline{b})$ and the complete type $tp_B(d; \overline{b})$ is also realized in A (though with parameter $\overline{a}$) as A is ω-saturated. We can thus construct a back-and-forth system from B to A. $\Box$

Theorem 27 *A is a countable ω-saturated model $\rightarrow$ A is a countable universal model.*

The proof is again similar to the earlier proof that all countable atomic models are prime. However, the converse of this theorem now fails (see Chang & Keisler: 101 for a counterexample). Intuitively, if A is a countable ω-saturated model, all countable models $B \equiv A$ must be elementarily embedded in A as A realizes all the complete types we can reasonably expect to be satisfied in a countable model. Note that saturation is a very strong condition as not only does A realize all n-types in the base language $\mathcal{L}$ (in which case we say A is *weakly saturated*) but also all n-types with parameters defined on finite subsets of A. Also, if a complete theory T has a countable ω-saturated model, then T has a countable atomic model, though the converse is not true.

7.3 ω-Categoricity

Now that we have examined atomic and ω-saturated countable models, we might also ask when a complete theory T has (up to isomorphism) a unique countable model. In this section, we give several characterizations for when this is the case. The exact theorem given below is part of the Theorem of Engeler, Ryll-Nardzewski and Svenonius (see Hodges: 171-2 and Chang & Keisler: 101-3) but ignores the bits on automorphism groups.

Definition 24 *A theory T is ω-categorical if T has exactly one countable model up to isomorphism [note: such a theory must be complete]. A structure A is ω-categorical if $Th(A)$ is ω-categorical.*

Theorem 28 *Let T be a complete theory in a countable language $\mathcal{L}$. Then the following are equivalent:*

(a) T is ω-categorical;

(b) T has a countable model which is both atomic and ω-saturated;

(c) for each $n < \omega$, every type in $S_n(T)$ is principal;

(d) for each $n < \omega$, $|S_n(T)|$ is finite;

(e) for each $n < \omega$, there are only finitely many pairwise non-equivalent formulae $\phi(x_1, ...x_n)$ in the language $\mathcal{L}$ mod T;

(f) all models of T are atomic.

Proof : Following Chang & Keisler, we prove (a) $\Rightarrow$ (b) $\Rightarrow$ (c) $\Rightarrow$ (d) $\Rightarrow$ (e) $\Rightarrow$ (f) $\Rightarrow$ (a).
(a) $\Rightarrow$ (b): Let $A \models T$ be the unique countable model of T (up to isomorphism). Then A is trivially prime so A is atomic. Further, T has only countably many n-types so A is ω-saturated.
(b) $\Rightarrow$ (c): Since A is ω-saturated, every type in $S_n(T)$ is realized in A. But A is also atomic so every type in $S_n(T)$ is principal.
(c) $\Rightarrow$ (d): Consider $\Phi = \{\sim \phi(x_1, ..., x_n) |\, \phi$ is a complete formula in $T\}$ for $n < \omega$. As every type in $S_n(T)$ is principal, $T \cup \Phi$ is inconsistent, so $T \cup \Psi$ is inconsistent for some finite subset $\Psi = \{\sim \phi_j | 1 < j < m\} \subset \Phi$ by Compactness. Thus $T \vdash \sim (\sim \phi_1 \wedge ... \wedge \sim \phi_m)$ so $T \vdash \phi_1 \vee ... \vee \phi_m$. But $\{\phi_j\}$ are complete formulae (so generate principal types in T) and every element in every model of T must satisfy one of them so $|S_n(T)| \leq m$ for some finite m.
(d) $\Rightarrow$ (e): Let Φ^*/Ψ^* be the set of all types of T containing the formulas $\phi(x_1, ..., x_n)$ and $\psi(x_1, ..., x_n)$ respectively. Then $\Phi^* = \Psi^*$ implies $T \vdash \phi \leftrightarrow \psi$ so as there are at most $m < \omega$ types in $S_n(T)$, there are at most 2^m pairwise non-equivalent formulae with n free variables in $\mathcal{L}$ mod T.
(e) $\Rightarrow$ (f): Let A be a model of T and consider $tp_A(\overline{a})$ for some $\overline{a} \in A$. As for each $n < \omega$, there are only finite non-equivalent formulae (mod T), the conjunction of the non-equivalent formulae in $tp_A(\overline{a})$ (*i.e.*, $\phi^{\overline{a}} = \phi_1(\overline{a}) \wedge \phi_2(\overline{a}) \wedge ... \wedge \phi_k(\overline{a})$) is itself a formula in $\mathcal{L}$. But $\phi^{\overline{a}}$ generates $tp_A(\overline{a})$ so A is atomic.
(f) $\Rightarrow$ (a): If all models of T are atomic, then since the countable atomic model of T is unique, T is ω-categorical. $\Box$

Examples of ω-categorical theories include DLO, atomless Boolean algebras, infinite Abelian groups with all elements of prime order, and the theory of an equivalence relation with infinitely many infinite classes.

Interestingly, no complete theory has exactly two nonisomorphic countable models:

Theorem 29 *(Vaught) Let T be a complete theory. Then either T is ω-categorical or T has ≥ 3 nonisomorphic countable models.*

Proof : Assume T has exactly two nonisomorphic countable models A, B so A is atomic and B is ω-saturated. Choose some tuple $\bar{b} \in B$ that does not realize a principal type in T and let $B' = (B, \bar{b})$. Since B is ω-saturated and there is a one-to-one correspondence between the types in $\mathcal{L}$ and $\mathcal{L}(\bar{b})$, B' is also ω-saturated so letting $T' = Th(B')$, T' has an ω-saturated model (B') and hence an atomic model $(C, \bar{b}) \models T'$. Now consider the reduct $C = (C, \bar{b})|\mathcal{L} \models T$. First, C is not atomic as the tuple in C named by $\bar{b}$ in $\mathcal{L}(\bar{b})$ does not realize a principal type. Second, C is not ω-saturated: as T is not ω-categorical, neither is T' (both theories have infinitely many nonequivalent formulae) and since no model of T' is both atomic and ω-saturated, $(C, \bar{b})$ is not ω-saturated so neither is C. Thus $C \models T$ is neither atomic nor ω-saturated so $C \not\cong A$ and $C \not\cong B$ respectively. $\square$

8 The Infinite Elevator of Löwenheim-Skolem

In the previous section, we explored what happens in the countable realm, categorizing the models of a complete theory T with countable elements. We now broaden our analysis to the uncountable case. Fortunately, here we have the theorems of Löwenheim-Skolem which tell us we can move up and down with ease through the higher reaches of the infinite. The Downward Löwenheim-Skolem (DLS) theorem is presented first and involves the construction of models by adding *Skolem functions* to the language. The Upward Löwenheim-Skolem (ULS) theorem is then an easy consequence of DLS and Compactness.

Disliking uncountable structures, Skolem proved that for every infinite structure A of countable signature, there is an elementarily equivalent countable substructure $B \prec A$. Hence, uncountable structures are redundant (and sometimes lead to counterintuitive results). An easy way to show this is to start with some countable subdomain of $dom(A)$ and close it off by adding new constants so that the Tarski-Vaught criterion for elementary substructures holds. But instead of adding constants, Skolem added functions:

Definition 25 *A Skolemization of T is a theory $T^+ \supseteq T$ in $\mathcal{L}^+ \supseteq \mathcal{L}$ such that:*
(a) if $A \models T$, there exists an expansion $A^+ \supseteq A$ such that $A^+ \models T^+$;
(b) for every $\phi(\bar{x}, y) \in \mathcal{L}^+$, $T^+ \vdash \forall\bar{x}(\exists y \phi(\bar{x}, y) \to \phi(\bar{x}, t(\bar{x})))$ for some term $t \in \mathcal{L}^+$.
The terms t are called Skolem functions for T^+.

Definition 26 *A theory T is a Skolem theory (or has Skolem functions) if T is a Skolemization of* *itself. A structure A has Skolem functions if $Th(A)$ is a Skolem theory. If $A \models T$ with T a Skolem theory, $\langle X \rangle_A$ is called the Skolem hull of X.*

Theorem 30 *T is a Skolem theory $\rightarrow$ T has the property of quantifier elimination.*

Proof : follows immediately from definitions ($T \vdash \forall \overline{x}(\exists y \phi(\overline{x}, y) \leftrightarrow \phi(\overline{x}, t(\overline{x})))$. $\square$

Theorem 31 *Let $A \models T$ with T a Skolem theory and $X \subset dom(A)$ such that $\langle X \rangle_A$ is non-empty. Then $\langle X \rangle_A \preccurlyeq A$.*

Proof : Let $B = \langle X \rangle_A$ and consider some $\psi \in \mathcal{L}$ and $\bar{b} \in B$ such that $A \models \exists y \psi(\bar{b}, y)$. As T is a Skolem theory, $A \models \psi(\bar{b}, t(\bar{b}))$ for some term $t \in \mathcal{L}$ and element $t^A(\bar{b}) \in A$. Thus by the Tarski-Vaught criterion for elementary substructures, $B \preccurlyeq A$ as desired. $\square$

By this last theorem, Skolem theories have the nice property that every non-empty substructure of their models is an elementary substructure. But there are few natural Skolem theorems. Like Hintikka sets (recall Henkin's Completeness/Compactness proof), Skolem theories are generally artificially constructed:

Theorem 32 *Let $\mathcal{L}$ be a first-order language. Then there are a first-order language $\mathcal{L}^\Sigma \supseteq \mathcal{L}$ and set of sentences $\Sigma \in \mathcal{L}^\Sigma$ such that:*
(a) every $\mathcal{L}$-structure A can be expanded to a model $A^\Sigma \models \Sigma$;
(b) Σ is a Skolem theory in $\mathcal{L}^\Sigma$;
(c) $|\mathcal{L}^\Sigma| = |\mathcal{L}|$.

Proof : For each formula $\phi(x, y) \in \mathcal{L}$, introduce the new function $F_{\phi,x}$ and take Σ_1 to be the set of all sentences of the form $\forall x(\exists y \phi(x, y) \rightarrow \phi(x, F_{\phi,x}))$ in the new language $\mathcal{L}_1$. Now consider an arbitrary non-empty $\mathcal{L}$-structure A. Given $\phi(x, y) \in \mathcal{L}$ and $a \in A$, if there is a $b \in A$ such that $A \models \phi(a, b)$, put $F^{A'}_{\phi,x}(a) = b$; otherwise, set $F^{A'}_{\phi,x}(a) = a$ (it does not matter). The expanded model A^{Σ_1} is then a model of Σ_1. Iterating this expansion, we form the increasing chains $(L_n : n < \omega)$ and $(\Sigma_n : n < \omega)$. Now letting L^Σ and Σ be their unions, (a) and (b) clearly hold. $|\mathcal{L}^\Sigma| \leq |\mathcal{L}| + |\mathcal{L}| \cdot \omega = |\mathcal{L}|$ so (c) holds as well. $\square$

Theorem 33 *Every theory $T \in \mathcal{L}$ has a Skolemization $T^+ \in \mathcal{L}^+$ with $|\mathcal{L}^+| = |\mathcal{L}|$.*

As all the work has already been done in constructing the set Σ, simply let $T^+ = T \cup \Sigma$. We are now ready to prove DLS:

Theorem 34 *(Downward Löwenheim-Skolem theorem) Let $\mathcal{L}$ be first-order, A an $\mathcal{L}$-structure, $X \subseteq Dom(A)$ and λ a cardinal such that $|\mathcal{L}| + |X| \leq \lambda \leq |A|$. Then there exists B with $B \preccurlyeq A$, $|B| = \lambda$, and $X \subseteq dom(B)$.*

Proof : Consider $Y \supseteq X$ in $dom(A)$ with $|Y| = \lambda$ and expand A to a model $A^\Sigma \models \Sigma$ in $\mathcal{L}^\Sigma$. Let $B = \langle Y \rangle_{A^\Sigma} | \mathcal{L}$. Then $|B| \leq |Y| + |\mathcal{L}^\Sigma| = \lambda + |\mathcal{L}| = \lambda = |Y| \leq |B|$. Also, $\langle Y \rangle_{A^\Sigma} \preccurlyeq A^\Sigma$ since Σ is a Skolem theory so $B \preccurlyeq A$. $\square$

DLS thus tells us that we can move downward through the transfinite hierarchy, bounded from below by only the cardinality of the language $\mathcal{L}$. It is now easy to show that we can freely move upward as well:

Theorem 35 *(Upward Löwenheim-Skolem theorem) Let $\mathcal{L}$ be first-order, A an infinite $\mathcal{L}$-structure, both of cardinality $\leq \lambda$. Then there exists B with $A \preccurlyeq B$ and $|B| = \lambda$.*

Proof : Introduce λ new constants and let $T = eldiag(A) \cup \{c_i \neq c_j | i < j < \lambda\}$ in $\mathcal{L}(\overline{c})$. As A is infinite, T is finitely satisfiable so by Compactness, there is a model $B' \models T$. Let $B = B'|\mathcal{L}$. Since $B' \models eldiag(A)$, by the elementary diagram lemma we have an elementary embedding $e : A \mapsto B$ so $A \preccurlyeq B$. Further, B must have at least λ elements so we can invoke DLS to ensure $|B| = \lambda$. $\square$

9 Preservation Theorems

We have already seen some examples of preservation in our discussion of different mappings in Section 2: f preserves ϕ if $A \models \phi(\overline{a}) \rightarrow B \models \phi(f\overline{a})$ for every $\overline{a} \in A$. In this section, we explore several more important preservation results, *i.e.*, that certain mappings preserve all formulae with certain syntactic features. But preservation phenomena (applied to sets of sentences) also work in the other direction: if a theory is preserved under certain mappings, we can also infer the syntactic form of the theory (up to equivalence).

Theorem 36 *embeddings preserve $\exists_1$ formulae.*

Proof : From before, we have that an embedding f preserves all literals. By induction on the complexity of formulae, this easily extends to all Boolean combinations of atomic formulae. Now consider when $\phi = \exists x \psi$ and assume f preserves ψ, *i.e.*, $A \models \psi(\overline{a}) \rightarrow B \models \psi(f\overline{a})$. Now if $A \models \phi$, then $A \models \psi(\overline{c})$ for some $\overline{c}$ so $B \models \psi(f\overline{c})$ and $B \models \phi$ as desired. $\square$

Theorem 37 *a theory T is preserved under submodels $\leftrightarrow$ T is an $\forall_1$ theory.*

Proof : ($\Leftarrow$) Let B be a submodel of A. Then there exists an embedding $f : B \mapsto A$ so from the above theorem, $B \models \phi \rightarrow A \models \phi$ where ϕ is an $\exists_1$ formula. But now $A \models \neg\phi \rightarrow B \models \neg\phi$ and as $\neg\phi$ is an $\forall_1$ formula, $\forall_1$ formulae are preserved in substructures.
($\Rightarrow$) Assume T is preserved under submodels and let $\Delta = \{$set of sentences in $\mathcal{L}$ equivalent to $\forall_1$ sentences$\}$. Consider $A \models T$ and B such that $A \models \phi \rightarrow B \models \phi$ where $\phi \in \Delta$. We must show $B \models T$ (as then T is equivalent to an $\forall_1$ theory). Set $T' = T \cup diag(B)$ in $\mathcal{L}(b)$ and let $\theta_1(b), ..., \theta_n(b)$ be a finite set of formulae from $diag(B)$. Then $B \models \exists x(\theta_1(x), ..., \theta_n(x))$ so given the contrapositive of the above conditional, $A \models \exists x(\theta_1(x), ..., \theta_n(x))$. As $A \models T$, $T \cup \theta_1, ..., \theta_n$ is consistent so as T' is

finitely satisfiable, T' is satisfiable by Compactness. Let $C \models T'$. Then $C \models T$ and $B \subseteq C|\mathcal{L}$ so as T is preserved under submodels, $B \models T$. $\square$

Theorem 38 *a theory T is preserved in unions of chains $\leftrightarrow$ T is an $\forall_2$ theory.*

Proof : ($\Leftarrow$) Let $(A_i : i < \gamma)$ be a chain of $\mathcal{L}$-structures (*i.e.*, $A_0 \subseteq ... \subseteq A_n \subseteq ...$) and $a \in A_0$ such that $A_i \models \forall\overline{y}\exists\overline{x}\psi(\overline{y}, \overline{x}, \overline{a})$ for all $i < \gamma$. Put $B = \bigcup_{i<\gamma} A_i$. We must show $B \models \forall\overline{y}\exists\overline{x}\psi(\overline{y}, \overline{x}, \overline{a})$. If $\overline{b} \in B$, then $\overline{b} \in A_i$ for some $i < \gamma$ so by assumption, $A_i \models \exists\overline{x}\psi(\overline{b}, \overline{x}, \overline{a})$. Now as $\exists_1$ formulas are preserved under embeddings and $A_i \subseteq B$, $B \models \forall\overline{y}\exists\overline{x}\psi(\overline{y}, \overline{x}, \overline{a})$ as desired.
($\Rightarrow$) Assume T is preserved in unions of chains and let $\Delta = \{$set of sentences in $\mathcal{L}$ equivalent to $\forall_2$ sentences$\}$. Consider $A \models T$ and B such that $A \models \phi \rightarrow B \models \phi$ where $\phi \in \Delta$. Set $T' = Th(A) \cup diag^{\forall_1}(B)$ where $Th(A)$ is the complete theory of A in $\mathcal{L}$ and $diag^{\forall_1}(B)$ is the set of universal sentences of $\mathcal{L}(b)$ in $eldiag(B)$. By the contrapositive of the above conditional, $B \models \phi \rightarrow A \models \phi$ where ϕ is an $\exists_2$ sentence in $\mathcal{L}$ so T' is finitely satisfiable and by Compactness, there exists $A' \models T'$ (note that $A'|\mathcal{L} \equiv A$ and $B \subseteq A'|\mathcal{L}$). Now add new constants and set $T'' = diag(A') \cup eldiag(B)$. As we also have $A' \models diag^{\forall_1}(B)$, every existential sentence in $\mathcal{L}(b)$ true in A' is true in $Mod(diag(B))$ so T'' is satisfiable by Compactness. Let $B' \models T''$ so $B \preccurlyeq B'|\mathcal{L}$ and $A'|\mathcal{L} \subseteq B'|\mathcal{L}$.

Iterating this construction, we can form the chain (we write A' for $A'|\mathcal{L}$, B' for $B'|\mathcal{L}$, etc.): $B \subseteq A' \subseteq B' \subseteq A'' \subseteq B'' \subseteq ...$ and let C be the union of this chain. Now C is the union of $(A^i : i < \gamma)$ and as each $A^i \models T$ and T is preserved in unions of chains, $C \models T$. But C is also the union of the elementary chain $B \preccurlyeq B' \preccurlyeq B'' \preccurlyeq ...$ so by the Tarski-Vaught theorem on unions of elementary chains, $B \models T$ as desired. $\square$

The latter part of this proof combines several powerful model-theoretic techniques. First, note that by deliberately considering particular kinds of theories, we can ensure the models given by Compactness have specific features (we have already used this technique extensively in Section 6 and in the proof of ULS). Given a finitely consistent set of sentences, Compactness simply tells us that a model exists in which all the sentences are satisfied. But by using diagrams (or other carefully chosen sentences), we can exercise some control. In T', the component $Th(A)$ ensures that the model $A'|\mathcal{L}$ is elementarily equivalent to A; the component $diag^{\forall_1}(B)$ ensures that $B \subseteq A'|\mathcal{L}$ (and a bit more); using $eldiag(B)$ in T'' ensures the stronger correspondence $B \preccurlyeq B'|\mathcal{L}$. Second, the proof also uses the neat trick of alternating chains - we create an increasing chain of models in which alternating components have special properties. As these properties (namely, elementary equivalence and modeling T) are preserved in the union of the chain, transfinite induction allows us to find a model in which both properties hold.

Other preservation results not proven here include:
(i) $\exists_1^+$ formulae are preserved under homomorphisms
(an $\exists_1^+$ formula is an $\exists_1$ formula in which negation does not occur);
(ii) all positive formulae are preserved under surjective homomorphisms
(a positive formula is one in which negation does not occur);

(iii) all formulae are preserved under isomorphisms (surjective embeddings).

Aside from some of our other preservation results in this section, there is now the following preservation hierarchy: homomorphisms preserve $\exists_1^+$ formulae and if they are surjective, all positive formulae; embeddings preserve $\exists_1$ formulae and if they are elementary, all formulae.

10 Amalgamation Theorems

10.1 Merging Models

In many instances, it will be useful to *amalgamate*, or join, several models together into a larger structure. There are many ways to do this and many interesting proofs that use such constructions. We begin with the father theorem:

Theorem 39 *(Elementary Amalgamation theorem) Let $\mathcal{L}$ be first-order and let B, C be $\mathcal{L}$-structures. If there exists $\overline{a} \in B$ and $\overline{c} \in C$ with $(B, \overline{a}) \equiv (C, \overline{c})$, there is an elementary extension $D \succcurlyeq B$ and an elementary embedding $g : C \mapsto D$ such that $g\overline{c} = \overline{a}$.*

Proof : Without loss of generality, assume $\overline{a} = \overline{c}$ (we consider an isomorphic copy of C if necessary) and consider $T = eldiag(B) \cup eldiag(C)$. It must be shown that T is finitely satisfiable. As in any finite subset $T' \subseteq T$ there are only finitely many sentences in $eldiag(C)$, we can take their conjunction $\phi(\overline{a}, \overline{d})$ with $\phi \in \mathcal{L}$ and $\overline{d}$ elements in C which are not in $\overline{a}$. If T' is unsatisfiable, $eldiag(B) \vdash\sim \phi(\overline{a}, \overline{d})$ so as the elements $\overline{d}$ are not in B, $eldiag(B) \vdash \forall\overline{x} \sim \phi(\overline{a}, \overline{x})$. But now $(B, \overline{a}) \models \forall\overline{x} \sim \phi(\overline{a}, \overline{x})$ so by assumption, $(C, \overline{a}) \models \forall\overline{x} \sim \phi(\overline{a}, \overline{x})$, a contradiction. By Compactness then, there exists a model $D' \models T$ and by two applications of the elementary diagram lemma, the reduct $D = D'|\mathcal{L}$ is an elementary extension of both B and C. $\square$

$$\begin{tikzcd} & D & \\ B \arrow[ur, "\preccurlyeq"] & & C \arrow[ul, "g"'] \\ & \langle\overline{a}\rangle_B \arrow[ul, "\subseteq"] \arrow[ur, "f"'] & \end{tikzcd}$$

Figure 3: Elementary Amalgamation (Hodges: 135)

The elementary amalgamation theorem is depicted in the above diagram. On the right side, the condition $(B, \overline{a}) \equiv (C, \overline{c})$ ensures, by the diagram lemma, that there is a unique embedding $f : \langle\overline{a}\rangle_B \mapsto C$ with $f\overline{a} = \overline{c}$. The right side of the diagram is essentially still symmetric to the left as there is a $C' \cong C$ with $\langle\overline{a}\rangle_B \subseteq C'$ and $C' \preccurlyeq D$ (so there is really not that much extra commotion). The idea here is that if a substructure is contained in two (or more) distinct elementarily equivalent structures, the structures can be elementarily embedded into a larger structure under mappings that agree on the common substructure. In the special case where $\overline{a}$ is empty, the elementary amalgamation theorem tells us that elementarily equivalent structures can be elementarily embedded in a larger structure.

A nice application of the elementary amalgamation theorem is that all of the complete types with respect to a structure A can be realized in some elementary extension of A:

Theorem 40 *Let $\mathcal{L}$ be first-order and let A be an $\mathcal{L}$-structure. There is an elementary extension $B \succcurlyeq A$ such that every type over $dom(A)$ w.r.t A is realized in B.*

Proof : Consider the set of complete types $\{p_i : i < \lambda\}$ over $dom(A)$ w.r.t A. With each $p_k = tp_{A_k}(\overline{a_k}/dom(A))$ for some elementary extension $A_k \succcurlyeq A$ and $\overline{a_k} \in A_k$, we can construct an increasing chain of models $B_0 \preccurlyeq B_1 \preccurlyeq B_2 \preccurlyeq ...$ as follows: let $B_0 = A$ and form B_i by elementarily amalgamating B_{i-1} and A_i. Take unions at limit ordinals and let $B = \bigcup_{i<\lambda} B_i$. $\Box$

An amalgam is said to be *strong* if it has the minimum-overlap property, *i.e.*, the overlap of B and $C' = gC$ in D is precisely $\langle \overline{a} \rangle_B$. It can be shown (see Hodges: 139-140) that a necessary condition for a strong amalgam is that the common substructure $\langle \overline{a} \rangle_B$ is *algebraically closed* in B (or C by symmetry). Note that if $X \subseteq domB$, an element $b \in B$ is *algebraic over* X if there is a first-order formula $\phi(x, \overline{y})$ and $\overline{a} \in X$ such that $B \models \phi(b, \overline{a}) \wedge \exists_{\leq n} x \phi(x, \overline{a})$ for some finite n. Letting $acl_B(X)$ represent the set of all elements in B algebraic over X, a strong amalgam is then possible when $\langle \overline{a} \rangle_B = acl_B(\langle \overline{a} \rangle_B)$.

We now give two of the many variations on amalgamations:

Theorem 41 *(Existential Amalgamation theorem) Let B, C be $\mathcal{L}$-structures. If there exists $\overline{a} \in B$ and a homomorphism $f : \langle \overline{a} \rangle_B \to C$ such that $(C, f\overline{a}) \Rrightarrow_1 (B, \overline{a})$ (i.e., $(C, f\overline{a}) \models \phi \to (B, \overline{a}) \models \phi$ for all first-order $\exists_1$ sentences $\phi \in \mathcal{L}$), there is an elementary extension $D \succcurlyeq B$ and an elementary embedding $g : C \mapsto D$ such that $gf\overline{a} = \overline{a}$.*

Proof : Note that the condition $(C, f\overline{a}) \Rrightarrow_1 (B, \overline{a})$ implies that f is an embedding so without loss of generality, assume f is the identity on $\langle \overline{a} \rangle_B$. The proof of the elementary amalgamation theorem now goes through as before as at the point where $(B, \overline{a}) \models \forall \overline{x} \sim \phi(\overline{a}, \overline{x})$, this new condition also implies that $(C, \overline{a}) \models \forall \overline{x} \sim \phi(\overline{a}, \overline{x})$. $\Box$

When $\overline{a}$ is empty, the existential amalgamation theorem says that if $C \Rrightarrow_1 B$, C is embeddable in some elementary extension of B. Existential amalgamation thus trims the fat from the elementary case, showing us that the weaker condition $\Rrightarrow_1$, rather than full-blown elementary equivalence, ensures that the structures B and C can be amalgamated. There are numerous interesting corollaries/applications of the existential amalgamation theorem and related amalgamation theorems, such as preservation results and interpolation theorems (see Hodges 5.4/5.5).

Our second amalgamation theorem will be useful in proving the Craig interpolation theorem in Section 11:

Theorem 42 *Let B, C be $\mathcal{L}_1, \mathcal{L}_2$-structures respectively and let $\mathcal{L}_0 = \mathcal{L}_1 \cap \mathcal{L}_2$ and $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2$. If there exists $\overline{a} \in B, C$ such that $(B|\mathcal{L}_0, \overline{a}) \equiv (C|\mathcal{L}_0, \overline{a})$, there is an $\mathcal{L}$-structure D such that $B \preccurlyeq D|\mathcal{L}_1$ and there is an elementary embedding $g : C \mapsto D|\mathcal{L}_2$ with $g\overline{a} = \overline{a}$.*

Proof : Using the elementary amalgamation proof with $T = eldiag(B) \cup eldiag(C|\mathcal{L}_0)$, we obtain the weaker result that there is an $\mathcal{L}_1$-structure $D \succcurlyeq B$ and an elementary embedding $g : C|\mathcal{L}_0 \mapsto D|\mathcal{L}_0$. Though this is insufficient to prove the theorem at hand, it allows us to build chains of elementary extensions of B and C such that the $\mathcal{L}_0$-reducts of the unions are isomorphic, *i.e.*, $\bigcup B_i|\mathcal{L}_0 \cong \bigcup C_i|\mathcal{L}_0$ (see Hodges: 147-8 for details). So as the union $\bigcup B_i$ is an $\mathcal{L}_1$-structure and $\bigcup C_i$ is an $\mathcal{L}_2$-structure, we can expand $\bigcup B_i$ (or $\bigcup C_i$ by symmetry) to an $\mathcal{L}$-structure D [*note: expanding* a structure involves adding new properties and introducing new names on a fixed domain of elements; this differs from *extending* a structure where the universe itself is enlarged]. The Tarski-Vaught theorem on unions of elementary substructures does the rest. $\square$

10.2 Fraïsse Limits

One of the more remarkable instances of amalgamation is the *Fraïsse limit*. Fraïsse's ingenious idea was that given a class of finite structures having various properties, we can amalgamate them together to form a 'limit' structure. For example (and one we will use throughout this subsection), the limit of the class of finite linear orderings is the ordered rationals $\langle \mathbb{Q}, < \rangle$. Formal proofs of the existence and uniqueness of Fraïsse limits can be found in Hodges (161-4).

The starting point is a class K of finitely generated structures. K is called the *age* of some structure if K is non-empty and has the following properties:

- *Hereditary property* (HP): If $A \in K$ and B is a finitely generated substructure of A (*i.e.*, $B = \langle a \rangle_A$ for finite a), then B is isomorphic to some structure in K
- *Joint embedding property* (JEP): If A,$B \in K$, then there exists a $C \in K$ and embeddings $f : A \mapsto C$ and $g : B \mapsto C$ (an embedding from $A \mapsto C$ is an isomophism from A to some substructure of C).

In addition, the class of finite linear orderings has the following important property:

- *Amalgamation property* (AP): If A,B,$C \in K$ and $e : A \mapsto B$ and $f : A \mapsto C$ are embeddings, then there is a $D \in K$ and embeddings $g : B \mapsto D$ and $h : C \mapsto D$ such that $ge = hf$

By the AP, the countable limit of the class K of finite orderings must be a dense linear ordering without endpoints so K tends to the rationals rather than, say, the integers or natural numbers (though K is still the age of both $\langle \mathbb{Q}, < \rangle$ and $\langle \mathbb{Z}, < \rangle$).

Theorem 43 *(Fraïsse's Theorem) Let L be a countable language, K a non-empty finite/countable collection of finitely generated L-structures which has HP, JEP and AP. Then there is a unique (up to isomorphism) L-structure D of cardinality $\leq \omega$ such that K is the age of D and D is homogeneous.*

Definition 27 *A structure D is homogeneous if every isomorphism between finitely generated substructures of D extends to an isomorphism from $D \mapsto D$ (*i.e.*, an automorphism of D).*

The structure D in Fraïsse's Theorem is what I, following Hodges, have been calling the Fraïsse limit (it is also known as the 'universal homogeneous structure of age K'). As already mentioned, when K is the class of finite orderings, D is the ordered rationals. Fraïsse's Theorem thus says that the theory of dense linear orderings without endpoints is ω-categorical. Other interesting examples of Fraïsse limits are the countable atomless Boolean algebra (where K is the class of finite Boolean algebras) and the celebrated 'random graph' (where K is the class of all finite graphs).

But what exactly does it mean to say $\langle \mathbb{Q}, < \rangle$ is 'homogeneous'? Well if every isomorphism between substructures of the ordered rationals extends to an automorphism, then when we take a peek at some local region of $\langle \mathbb{Q}, < \rangle$ and it looks the same, it really is. By contrast, consider an ordering of the natural numbers $\langle \mathbb{N}, < \rangle$. As only the zero element is a lower bound of all other elements and the naturals are not dense (this creates a problem as the substructures $A' = 0 < 1$ and $A'' = 0 < 2$ are isomorphic but any isomorphism between them can clearly not be extended to an automorphism of $\langle \mathbb{N}, < \rangle$), the ordered naturals are not homogeneous. Homogeneity is discussed further in Section 13.

11 Syntactic Theorems

This section presents several interesting theorems dealing with language.

Theorem 44 *(Craig Interpolation theorem) Let ϕ, ψ be sentences in $\mathcal{L}$ such that $\phi \vdash \psi$. Then there is a sentence θ such that:*

(a) $\phi \vdash \theta$ and $\theta \vdash \psi$;

(b) every symbol (excluding identity) that occurs in θ also occurs in both ϕ and ψ.

The sentence θ is the Craig interpolant of ϕ and ψ.

Proof : We give a constructive argument similar to Henkin's Completeness/Compactness proof and the proof of the Omitting Types theorem but the focus is now on an inseparable pair of theories rather than on a consistent theory. Given $T_1 \in \mathcal{L}_1$ and $T_2 \in \mathcal{L}_2$, a sentence $\theta \in \mathcal{L}_0 = \mathcal{L}_1 \cap \mathcal{L}_2$ *separates* T_1 and T_2 if $T_1 \vdash \theta$ while $T_2 \vdash \sim \theta$. If no such θ exists, T_1, T_2 are *inseparable*.

Now assuming there is no Craig interpolant θ of $\phi \in \mathcal{L}_1$ and $\psi \in \mathcal{L}_2$, we prove $\phi \not\vdash \psi$ by constructing a model of $\phi \wedge \sim \psi$. Let $\mathcal{L}_0 = \mathcal{L}_1 \cap \mathcal{L}_2$ and $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2$ and add new constants C so $\mathcal{L}'_i = \mathcal{L}_i \cup C$. We construct a pair of maximally consistent inseparable theories T and U in $\mathcal{L}'_1, \mathcal{L}'_2$ respectively as follows:

Enumerate the sentences in $\mathcal{L}_1'$ as $\{\phi_i\}$ with $\phi_1 = \phi$ and the sentences in $\mathcal{L}_2'$ as $\{\psi_i\}$ with $\psi_1 =\sim \psi$. Letting $T_1 = \phi$ and $U_1 =\sim \psi$, the theories T_1, U_1 are inseparable as if not, *i.e.*, $\phi \vdash \theta$ and $\sim \psi \vdash\sim \theta$ for some $\theta \in \mathcal{L}_0'$, the sentence $\forall \overline{x}\theta(\overline{x}) \in \mathcal{L}_0$ is a Craig interpolant of ϕ and ψ. Define increasing chains of theories $\{T_i\}$ and $\{U_i\}$ where we now only add ϕ_n to T_n if $T_n \cup \{\phi_n\}$ and U_n are inseparable (and similarly for ψ_n and U_n). As usual, we witness existential claims. Finally, put T and U as the unions of the respective chains $\{T_i\}$ and $\{U_i\}$ so T, U are inseparable.

T is clearly consistent as if not, everything follows from T so T is separable from U. It can also be shown that T is maximally consistent and similarly for U (see Chang & Keisler: 86) so consider the canonical interpretations $B \models T$ and $C \models U$. We must have $B|\mathcal{L}_0 \equiv C|\mathcal{L}_0$ as if not, there is a sentence $\delta \in \mathcal{L}_0$ such that $B \models \delta$ while $C \models\sim \delta$, contradicting the separability and maximal consistency of T and U. So by the syntactic amalgamation theorem in Section 10, there is an $\mathcal{L}$-structure D such that $B \preccurlyeq D|\mathcal{L}_1$ and there is an elementary embedding $g : C \mapsto D|\mathcal{L}_2$. Thus $D \models T \cup U$ so $D \models \phi \wedge \sim \psi$ as desired. $\square$

The notion of inseparable theories provides us with another valuable construction technique. In such previous arguments as Henkin's Completeness proof, we restricted our attention to a single chain of theories, adding constants until we arrived at a maximally consistent theory (the union of the chain) with certain built-in properties. In the proof of Craig's theorem, we not only build maximally consistent theories T and U (as before) but by ensuring that T, U are inseparable, the construction also ensures that the theories are consistent with each other, *i.e.*, $T \cup U$ is consistent.

When the language $\mathcal{L}$ has no constant or function symbols, there is the following strengthening of Craig's result:

Theorem 45 *(Lyndon Interpolation theorem) Let ϕ, ψ be sentences in a language $\mathcal{L}$ with no function or constant symbols such that $\phi \vdash \psi$. Then there is a sentence θ such that:*

(a) $\phi \vdash \theta$ and $\theta \vdash \psi$;

(b) every relation symbol (excluding identity) that occurs positively in θ also occurs positively in both ϕ and ψ;

(c) every relation symbol (excluding identity) that occurs negatively in θ also occurs negatively in both ϕ and ψ.

Finally, here is a nice application of Craig's Theorem:

Theorem 46 *(Robinson Consistency Theorem) Let $\mathcal{L}_1, \mathcal{L}_2$ be languages and let $\mathcal{L} = \mathcal{L}_1 \cap \mathcal{L}_2$. If T is a complete theory in $\mathcal{L}$ and $T_1 \supset T$, $T_2 \supset T$ are consistent theories in $\mathcal{L}_1, \mathcal{L}_2$ respectively, then $T_1 \cup T_2$ is consistent in $\mathcal{L}_1 \cup \mathcal{L}_2$.*

Proof : Assume $T_1 \cup T_2$ is inconsistent. Then by Compactness, there are finite subsets $\Sigma_1 \subset T_1$, $\Sigma_2 \subset T_2$ with $\Sigma_1 \cup \Sigma_2$ inconsistent. Let χ_1, χ_2 be the conjunctions of Σ_1, Σ_2 respectively. Then $\chi_1 \vdash\sim \chi_2$ so by Craig's theorem, there is a sentence $\theta \in \mathcal{L}$ such that $\chi_1 \vdash \theta$ and $\chi_2 \vdash\sim \theta$. Now since $T_1 \vdash \theta$ and T_1 is consistent, $T_1 \not\vdash\sim \theta$ so $T \not\vdash\sim \theta$. By analogous reasoning with T_2, $T \not\vdash \theta$, contradicting the completeness of T. $\square$

12 Model Completeness

In this section, we discuss existentially-closed structures and Abraham Robinson's related notions of model-completeness and model companions. While model-completeness/companions are important tools for the model theory of algebra, we will only define these notions and present some basic results, staying away from their algebraic applications.

Definition 28 *$A \models T$ is existentially closed (e.c.) iff for every model $B \models T$ with $A \subseteq B$, $A \preccurlyeq_1 B$ (i.e., any existential formula $\exists x \phi(x, \overline{a})$ which holds in B with $\overline{a} \in A$ also holds in A).*

Examples of e.c. structures are the algebraically closed fields and Fraïsse limits. Given a $\forall_2$ theory T, there is also a general method for constructing an e.c. extension of a model of T:

Theorem 47 *let K be the class of all models of $\forall_2$ theory T and let $A \models T$. Then there exists an e.c. structure $B \in K$ s.t. $A \subseteq B$.*

Proof : Enumerate as $(\phi_i, \overline{a}_i)_{i<\lambda}$ all pairs $(\phi, \overline{a})$ where ϕ is an $\exists_1$ formula in $\mathcal{L}$ and $\overline{a} \in A$. Now define a chain of structures $A_0 \subseteq A_1 \subseteq ...$ where $A = A_0$ and if there exists $C \supseteq A_i$ with $C \in K$ and $C \models \phi_i(\overline{a}_i)$, then $A_{i+1} = C$; else $A_{i+1} = A_i$. Letting $A^* = \bigcup_{i<\lambda} A_i$, then $A^* \in K$ (since T is $\forall_2$, T is preserved in unions of chains) and all $\exists_1$ formulas in $\mathcal{L}$ with parameters from A are realized in A^*. Now define a second chain $A \subseteq A^* \subseteq (A^*)^* \subseteq ...$ (*i.e.*, repeat the first chain construction ω times). Then $B = \bigcup_{n<\omega} A^{(n)} \in K$ and B is existentially-closed. To see this, suppose ϕ is $\exists_1$, $\overline{b} \in B$ and $C \models \phi(\overline{b})$ for some $C \in K$ which extends B. Since $\overline{b}$ is finite, $\overline{b} \in A^{(n)}$ for some $n < \omega$. But then $A^{(n+1)} \models \phi(\overline{b})$ so as $\exists_1$ formulae are preserved under embeddings, $B \models \phi(\overline{b})$ as required. $\square$

Note that if $||\mathcal{L}|| \leq \lambda$, we can ensure that the e.c. structure B found in the above proof has cardinality $|B| \leq \lambda$ (in constructing the first chain, we use DLS to find extensions C of cardinality $\leq \lambda$). Also (as we prove below), in cases where a theory T is model-complete (such as when T is a Skolem theory), all models of T turn out to be e.c. structures.

Definition 29 *T is model-complete iff for all models $A, B \models T$, if $A \subseteq B$ then $A \preccurlyeq B$ (i.e., all embeddings between models of T are elementary).*

Neither completeness nor model-completeness implies the other. The theory DLO with endpoints and $Th(\mathbb{N}, <)$ are both complete but not model-complete. The theory of algebraically closed fields is model-complete but not complete. Nonetheless, model-completeness combined with some additional properties of a theory *does* imply completeness, as in the following theorem:

Theorem 48 *Let T be model-complete.*

(i) if any two models of T are isomorphically embedded in a third model, then T is complete.

(ii) if T has a model which is isomorphically embedded in every model of T, then T is complete.

Proof : both (i) and (ii) follow from the fact that isomorphic embeddings are elementary when T is model-complete [*note:* (ii) equivalently says that a model-complete theory with a prime model is complete]. $\square$

Theorem 49 *If T is model-complete, then T is equivalent to a $\forall_2$ theory.*

Proof : By Tarski-Vaught theorem on unions of chains, T is preserved in unions of chains. $\square$

Theorem 50 *The following are equivalent:*
(a) T is model complete;
(b) for every $A \models T$, $T \cup diag(A)$ is complete in $\mathcal{L}_A$;
(c) every model of T is existentially-closed (Robinson's test);
(d) for every $\phi \in \exists_1$, there is $\psi \in \forall_1$ s.t. $T \vdash \phi \leftrightarrow \psi$.

Proof : We prove (a) $\Rightarrow$ (b) $\Rightarrow$ (c) $\Rightarrow$ (d) $\Rightarrow$ (a).
(a) $\Rightarrow$ (b): $T' = T \cup diag(A)$ has the same models as $eldiag(A)$ so all models of T' are elementarily equivalent in $\mathcal{L}_A$.
(b) $\Rightarrow$ (c): if $A \subseteq B$ and $B \models T$ then $B \models T'$ and since T' is complete in $\mathcal{L}_A$, $A \preccurlyeq_1 B$.
(c) $\Rightarrow$ (d): consider $\phi(x_1, ..., x_n) \in \exists_1$. Add $c_1, ..., c_n$ to $\mathcal{L}$ and let $\Gamma = \{\forall_1$ sentences θ in $\mathcal{L}_{\overline{c}}$ s.t. $T \vdash \phi \rightarrow \theta\}$. Let $A \models T \cup \Gamma$ and consider $T' = T \cup \{\phi\} \cup diag(A)$ in $\mathcal{L}_A$. If T' is inconsistent, $T \cup \{\phi\} \rightarrow \neg\Phi$ for some finite conjunction of formulae $\Phi \in diag(A)$. But then $T \cup \{\phi\} \rightarrow \neg\exists\overline{x}\Phi(\overline{x})$ in $\mathcal{L}_{\overline{c}}$ so $T \cup \{\phi\} \rightarrow \forall\overline{x}\neg\Phi(\overline{x})$, a contradiction since $\forall\overline{x}\neg\Phi(\overline{x}) \in \Gamma$ and $A \models \Gamma$. Let $B \models T'$ so $A \subset B$ and $B \models T$. By assumption, $A \preccurlyeq_1 B$ so ϕ holds in A and since A was arbitrary, $T \cup \Gamma \vdash \phi$. But now by Compactness, there is a finite conjunction $\psi \in \Gamma$ s.t. $T \cup \{\psi\} \vdash \phi$ so $T \vdash \phi \leftrightarrow \psi$ as desired.
(d) $\Rightarrow$ (a): first note that (d) implies that for every formula $\phi \in \mathcal{L}$, $T \vdash \phi \leftrightarrow \psi$ for some $\psi \in \forall_1$. Now consider $A, B \models T$ with $A \subset B$ and $B \models \phi(\overline{a})$ for some $\overline{a} \in A$. Then $B \models \psi(\overline{a})$ for some $\psi \in \forall_1$ and since $\forall_1$ formulae are preserved in substructures, $A \models \psi(\overline{a})$ so $A \models \phi(\overline{a})$ (and similarly for $\neg\phi$). $\square$

We can actually get the stronger result: if T has only infinite models, $\alpha \geq ||\mathcal{L}||$ and $A \preccurlyeq_1 B$ for any $A, B \models T$ of power α with $A \subset B$ (*i.e.*, all models of power α are e.c.), then T is model-complete. This gives us a nice criterion for model-completeness:

Theorem 51 *(Lindström's theorem) Let T be a $\forall_2$ theory in countable $\mathcal{L}$ which has only infinite models and is α-categorical for some infinite α. Then T is model-complete.*

Proof : starting with some model $A \models T$ of power α, we can construct an e.c. model of T of power α as above. But now since T is α-categorical, all models of power α are e.c. so T is model-complete. $\square$

Alternatively, to show T is model-complete we need only prove that T admits quantifier elimination as then condition (d) in Theorem 50 is clearly satisfied. But the reverse argument also holds. A particularly useful feature of model-completeness is that it provides an alternative route (recall Section 5) to showing a theory has quantifier elimination:

Theorem 52 *T has quantifier elimination iff T is model complete and $T_\forall$ - the set of all $\forall_1$ sentences of $\mathcal{L}$ s.t. $T \vdash \phi$ - has the amalgamation property (see Section 10.2).*

Finally, a few words on companionship:

Definition 30 *a theory U in $\mathcal{L}$ is a model companion of theory T if: (i) U is model-complete, (ii) every model of T has an extension which is a model of U, (iii) every model of U has an extension which is a model of T. We say T is companionable.*

Theorem 53 *Let T be an $\forall_2$ theory in $\mathcal{L}$.*
(i) T is companionable iff the class of e.c. models of T is axiomatizable by a theory in $\mathcal{L}$.
(ii) If T is companionable, then up to equivalence of theories, its model companion U is unique and is the theory of the class of e.c. models of T.

13 Further Topics

13.1 Saturation & Homogeneity

In Section 7.2, we discussed the ω-saturated models of a theory (models which realize all types over finite subsets of their domain) and countable universal models (models in which every other countable elementarily equivalent model is elementarily embeddable). Here, saturation and universality are generalized from the countable case and we also introduce the notion of homogeneity.

Definition 31 *A is λ-saturated iff for every set X of elements of A, if $|X| < \lambda$ then all complete 1-types over X w.r.t. A are realized in A. A is saturated iff A is $|A|$-saturated.*

Definition 32 *A is λ-homogeneous iff for every pair of sequences $\overline{a}, \overline{b} \in A$ of length $< \lambda$ and $c \in A$, if $(A, \overline{a}) \equiv (A, \overline{b})$ then there is some $d \in A$ s.t. $(A, \overline{a}c) \equiv (A, \overline{b}d)$. A is homogeneous iff A is $|A|$-homogeneous.*

Definition 33 *A is λ-universal iff any structure B of cardinality $< \lambda$ with $B \equiv A$ is elementarily embeddable in A.*

Our first two theorems can be proven using back-and-forth constructions akin to those from Theorems 22/26 above:

Theorem 54 *A is countable and atomic $\rightarrow$ A is countable and ω-homogeneous.*

Theorem 55 *A is λ-saturated $\rightarrow$ A is λ-homogeneous. [note: in fact, it can be shown that λ-saturation = λ-homogeneity + λ-universality.]*

As with building e.c. structures in Section 12, it is a fairly simple matter to construct countable ω-homogeneous (*i.e.*, countably homogeneous) elementary extensions of countable models:

Theorem 56 *every countable model has a countably homogeneous elementary extension.*

Proof : Consider a countable model A and for any two finite sequence $a_1, ..., a_{n+1}$ and $b_1, ..., b_n$ in A s.t. $(A, a_1...a_n) \equiv (A, b_1...b_n)$, introduce a new constant c_i and type $\Sigma_i(b_1...b_n, x)$ in $\mathcal{L}_{\overline{b}}$ corresponding to the set of formulae s.t. $A \models \Sigma_i(a_1...a_{n+1})$. Let $T = eldiag(A) \cup \Sigma_{i<\omega}(b_1...b_n, c)$. Then if $A' \models T$ (we use a standard Compactness argument here), $A'|\mathcal{L} \succcurlyeq A$ and for all finite sequences $a_1, ..., a_{n+1}$ and $b_1, ..., b_n$ in A s.t. $(A, a_1...a_n) \equiv (A, b_1...b_n)$, there exists a $b_{n+1} \in A'$ s.t. $(A', a_1...a_{n+1}) \equiv (A', b_1...b_{n+1})$. Repeating this procedure, we form the elementary chain $A \prec A' \prec A'' \prec ...$ and letting $B = \bigcup_{n\in\omega} A^{(n)}$, it is easy to verify that B is a countably homogeneous elementary extension of A. $\square$

Now that we know how to construct them, here are some simple properties of countably homogeneous models (the first two theorems can be proven with a standard back-and-forth argument):

Theorem 57 *If $A \equiv B$ and A, B are both (countably) homogeneous, of the same cardinality, and realize the same n-types over $\emptyset$ for each $n < \omega$, then $A \cong B$.*

Theorem 58 *Let A be a countably homogeneous and let $a_1, ..., a_n$ and $b_1, ..., b_n$ be sequences s.t. $(A, a_1, ..., a_n) \equiv (A, b_1, ..., b_n)$. Then there is an automorphism $f : A \mapsto A$ with $f(a_i) = b_i$.*

Theorem 59 *Let $(A_i : i < \omega)$ be an elementary chain of countably homogeneous models.*
(i) $B = \bigcup_{i<\omega} A_i$ is countably homogeneous.
(ii) If $A_i \cong A_j$ for all $1 < i < j < \omega$, $B \cong A_i$ for each i.

13.2 Two-Cardinal Theorems

As implied by the Löwenheim-Skolem theorems, a theory T with an infinite model in a countable language $\mathcal{L}$ cannot distinguish between infinite cardinals as T will have a model in each infinite power. Nonetheless, model theorists have been interested in whether theories can distinguish between pairs of infinite cardinals in another sense: given a theory T, is there a model $A \models T$ with a one-place relation V defined on A s.t. $|A| = \alpha$ and $|V| = \beta$? If so, we say A is a (α, β) model and T *admits the pair of cardinals* (α, β). Alternatively, a *two-cardinal theorem* tells us that under certain conditions, there exists a model $A \models T$ in which the $\emptyset$-definable subsets $\phi(A)$,$\psi(A)$ defined by ϕ, ψ have different cardinalities (*note:* if ϕ is $x = x$ and ψ is V, this collapses to our first definition with $\alpha \neq \beta$).

Theorem 60 *Let T be a theory in a countable language and let α, β, γ be infinite cardinals.*
(i) If T admits (α, β), then T admits (γ, β) for all γ s.t. $\alpha \geq \gamma \geq \beta$.
(ii) If T admits (α, β), then T admits all (γ, γ).

Proof : (i) follows immediately from DLS; (ii) follows from (i) since T admits (β, β) and we can introduce a new function F in $\mathcal{L}$ and extend T by a single sentence which says F is a one-to-one mapping from $A \mapsto V$. $\Box$

Theorem 61 *(Vaught's Two-Cardinal theorem) Let $\mathcal{L}$ be countable, $\phi(x), \psi(x) \in \mathcal{L}$ and T a theory. Then the following are equivalent:*

(a) T has a model A in which $|\phi(A)| \leq \omega$ but $|\psi(A)| = \omega_1$;

(b) T has a model A in which $|\phi(A)| < |\psi(A)| \geq \omega$;

(c) T has models A, B s.t. $B \preccurlyeq A$ and $\phi(A) = \phi(B)$ but $\psi(A) \neq \psi(B)$.

Definition 34 *a Vaught pair for ϕ is a pair of structures A, B s.t. $B \preccurlyeq A$, $B \neq A$ and $|\phi(A)| = |\phi(B)| = \lambda$ for infinite λ.*

A formula is two-cardinal for T if there is $A \models T$ s.t. $|A| \neq |\phi(A)|$; otherwise it is one-cardinal for T. A theory T is two-cardinal if there is a two-cardinal formula ϕ for T s.t. $|\phi(A)|$ is infinite in every model of T; otherwise T is one-cardinal. For a countable and complete theory T, the direction (c) $\Rightarrow$ (b) in Vaught's theorem says that if ϕ has a Vaught pair of models of T, then ϕ is two-cardinal for T and T is a two-cardinal theory.

13.3 Indiscernibles

Given a set X, let $[X]^n$ denote the set of all subsets of X with n elements. If X is simply ordered by $<$, we can take $[X]^n$ to be the set of increasing sequences $x_1 < ... < x_n$ from X. We will need the following combinatorial result:

Theorem 62 *(Ramsey's theorem) Let I be an infinite set and let $n \in \omega$. Suppose that $[I]^n = A_0 \cup A_1$. Then there is an infinite subset $J \subset I$ s.t. either $[J]^n \subset A_0$ or $[J]^n \subset A_i$.*

In the case where $n = 1$, Ramsey's theorem amounts to the pigeon-hole principle with two holes. When $n = 2$, Ramsey's theorem tells us that any infinite graph whose edges have been colored with two colors has an infinite monochromatic subgraph. And so on.

Definition 35 *Let A be a model in $\mathcal{L}$ and let $X \subset dom(A)$ be simply ordered by $<$. Then X is a set of indiscernibles in A if for all n and finite sequences $x_1 < ... < x_n$ and $y_1 < ... < y_n$ from X, $(A, x_1 ... x_n) \equiv (A, y_1 ... y_n)$.*

In other words, any finite sequences $x_1 < ... < x_n$ and $y_1 < ... < y_n$ from X cannot be distinguished by any formula $\phi(v_1, ..., v_n) \in \mathcal{L}$. Here is a sufficient condition for a set to be indiscernible:

Theorem 63 *Let $\langle X, < \rangle \subset dom(A)$ and for any sequences $x_1 < ... < x_n$ and $y_1 < ... < y_n$ from X, suppose there is an automorphism $f : A \mapsto A$ s.t. $f(x_i) = y_i$ for all i. Then X is a set of indiscernibles in A.*

Proof : follows immediately from the elementary diagram lemma. $\Box$

For example, if A is an algebraically closed field of characteristic zero, then $X = \{$algebraically independent elements in $X\}$ ordered by $<$ is a set of indiscernibles in A. If A is a Boolean algebra, then the set of all atoms $X \subset dom(A)$ is a set of indiscernibles. More generally, indiscernibles of any order type must exist in models of any theory T with infinite models, as outlined in the following two results:

Lemma 6 *Let T have infinite models and let $\mathcal{L}' = \mathcal{L} \cup \{c_n : n \in \omega\}$.
Then $T' = T \cup \{\phi(c_{i_1}, ..., c_{i_n}) \leftrightarrow \phi(c_{j_1}, ..., c_{j_n}) : \phi \in \mathcal{L}, n \in \omega, i_1 < ... < i_n, j_1 < ... < j_n\} \cup \{\neg c_1 = c_2\}$ is consistent in $\mathcal{L}'$.*

Proof : let A be an infinite model of T and I a countable infinite well-ordered subset of $dom(A)$. We show by induction that for any finite subset $\Delta \subset T'$, there is an infinite subset $J_\Delta \subset I$ s.t. $(A, j_i)_{i\in\omega} \models \Delta$ for $j_i \in J_\Delta$. Assume this holds for some Δ and consider some new formula $\psi(v_1, ..., v_m) \in \mathcal{L}$. We divide $[J_\Delta]^m$ into two parts: $A_0 = \{j_1 < ... < j_m : j_i \in J_\Delta \wedge A \models \psi(j_1, ..., j_m)\}$ and $A_1 = \{j_1 < ... < j_m : j_i \in J_\Delta \wedge A \not\models \psi(j_1, ..., j_m)\}$. Then $[J_\Delta]^m = A_0 \cup A_1$ so by Ramsey's theorem, there is an infinite subset $K \subset J_\Delta$ s.t. either $[K]^m \subset A_0$ or $[K]^m \subset A_1$. In either case, $(A, k_i)_{i\in\omega} \models \Delta \cup \{\psi(c_{i_1}, ..., c_{i_m}) \leftrightarrow \psi(c_{j_1}, ..., c_{j_m})\}$ so the induction is complete. $\square$

Theorem 64 *Let T have infinite models and let $\langle X, < \rangle$ be any simply ordered set. Then there exists $A \models T$ with $X \subset dom(A)$ s.t. X is a set of indiscernibles in A.*

Proof : let $\mathcal{L}' = \mathcal{L} \cup \{c_x : x \in X\}$ and $T' = T \cup \{\phi(c_{x_1}, ..., c_{x_n}) \leftrightarrow \phi(c_{y_1}, ..., c_{y_n}) : \phi \in \mathcal{L}, n \in \omega,$ $x_1 < ... < x_n, y_1 < ... < y_n \in X\} \cup \{\neg c_{x_1} = c_{x_2} : x_1 \neq x_2 \in X\}$. Then T' is finitely consistent as we can ordermorphically embed every finite subset of X into the set K from the above proof. So by Compactness, let $A' \models T'$. Then interpreting c_x as $x \in X$, $A = A'|\mathcal{L} \models T$, $X \subset A$ and X is a set of indiscernibles in A. $\square$