MODEL THEORY PREPARATION

REWORDING PROOFS AND COMPILING INFORMATION

James Holland

math.rutgers.edu/~jch258

jch258@math.rutgers.edu

2019-05-31

Contents

Section 1. Basic Logic 1
Section 2. Quantifier Elimination 6
Section 3. Realizing and Omitting Types 12
Section 4. Prime and Saturated Models 16
Section 5. Homogeneous Models 24
Section 6. Indiscernibles 29
Section 7. ω-Stable Theories 32
Section 8. Categoricity 35
Section 9. Morley Rank and Degree 44

Section 1. *Basic Logic*

This will be about first order logic, and its structures. In principle, for each language $\mathscr{L}$, we get a different logic system to deal with, but all of these can be collected in a coherent manner. There are three central results about logic which are used in a variety of ways. Even the proofs of these theorems have variants which are used to prove variant, stronger results.

First we introduce some notation and terminology to be used throughout the chapter. The first big chunk are about languages and theories. We will often interchangeably use "model" and "structure".

1 • 1. Definition

Let $\mathscr{L}$ be a language, and $\mathbf{A}$ be an $\mathscr{L}$-structure.

1. The $\mathscr{L}$-theory of $\mathbf{A}$—denoted $\mathrm{Th}(\mathbf{A})$—is the set of all $\mathscr{L}$-sentences modeled by $\mathbf{A}$.
2. For $X \subseteq A$, the language $\mathscr{L}_X$ denotes $\mathscr{L} \cup X$, adding a constant symbol for each element of X.
   a. $\mathbf{A}_X$ denotes the expanded model of $\mathbf{A}$ interpreting each symbol $x \in \mathscr{L}_X \cap X$ as the element $x \in X$.
   b. The theory $\mathrm{Th}_X(\mathbf{A})$ is the $\mathscr{L}_X$-theory of $\mathbf{A}_X$.
3. An $\mathscr{L}$-literal is an atomic $\mathscr{L}$-formula, or negation of an atomic $\mathscr{L}$-formula.
4. The diagram of $\mathbf{A}$—denoted $\mathrm{Diag}(\mathbf{A})$—is the set of $\mathscr{L}_A$-literals modeled by $\mathbf{A}_A$.
5. The elementary diagram of $\mathbf{A}$—denoted $\mathrm{Eldiag}(\mathbf{A})$—is the set $\mathrm{Th}_A(\mathbf{A})$.

This next big chunk of definitions is about relationships between structures.

1 • 2. Definition

Let $\mathcal{L}$ be a language, and A and B $\mathcal{L}$-structures.

1. A is a substructure of B—in symbols $\mathsf{A} \subseteq \mathsf{B}$—iff $\sigma^{\mathsf{A}} = \sigma^{\mathsf{B}} \restriction A$ for all symbols $\sigma \in \mathcal{L}$.
2. A is elementarily equivalent to B—in symbols $\mathsf{A} \equiv \mathsf{B}$—iff $\mathrm{Th}(\mathsf{A}) = \mathrm{Th}(\mathsf{B})$.
3. If $\mathsf{A} \subseteq \mathsf{B}$, then A is an elementarily substructure of B—in symbols $\mathsf{A} \preccurlyeq \mathsf{B}$—iff $\mathsf{B}_A \vDash \mathrm{Eldiag}(\mathsf{A})$.

This last chunk is about maps between structures.

1 • 3. Definition

Let $\mathcal{L}$ be a language, and A and B $\mathcal{L}$-structures. Let $f : A \to B$ be a function.

1. $f``\mathsf{A}$ is the $\mathcal{L}$-structure interpreting $\sigma \in \mathcal{L}$ as $f``\sigma^{\mathsf{A}}$, where $f(\langle a_i : i < n \rangle)$ is taken to mean $\langle f(a_i) : i < n \rangle$.
2. f is an embedding iff $f``\mathsf{A} \subseteq \mathsf{B}$.
3. f is an elementary embedding iff $f``\mathsf{A} \preccurlyeq \mathsf{B}$.
4. f is an isomorphism iff $f``\mathsf{A} = \mathsf{B}$. If there is an isomorphism between A and B, we write $\mathsf{A} \cong \mathsf{B}$.
5. f is an automorphism iff f is isomorphism, and $\mathsf{A} = \mathsf{B}$.

There are of course other definitions that will play a role, but those will be more local, or at least used only after their first introduction. Most of these will be considered throughout this chapter, and in fact this document.

The two other big things left undefined are '$\vDash$' and '$\vdash$'. Both of these are assumed to be understood. The proof system is not laid out explicitly, but is taken to be a sequence of formulas following certain rules. This makes it easier to reason about proofs for our purposes than something like Gentzen's system. Regardless, we still have the following defintions used in basic logic.

1 • 4. Definition

Let $\mathcal{L}$ be a language. Let T be an $\mathcal{L}$-theory. Let κ be a cardinal.

1. T is inconsistent iff $T \vdash \varphi \wedge \neg\varphi$ for any $\mathcal{L}$-formula φ. T is consistent iff T isn't inconsistent.
2. T is satsifiable iff there is an $\mathcal{L}$-model $\mathsf{A} \vDash T$.
3. T is complete iff for every $\mathcal{L}$-sentence φ, either $T \vdash \varphi$ or $T \vdash \neg\varphi$.
4. T is κ-categorical iff $\mathsf{A} \cong \mathsf{B}$ for all $\mathcal{L}$-structures $\mathsf{A} \vDash T$ and $\mathsf{B} \vDash T$ such that $|A| = |B| = \kappa$.

§ 1.A. Completeness and compactness

Without introducing the proof system, the statement of compactness is ambiguous. There are many proofs systems which are equivalent, so the reader can choose whatever is preferred. Just the idea of the proof will be given here rather than a completely full and rigorous proof. Note that the converse is clear from whatever reasonable definition of '$\vdash$' we give.

1.A • 1. Theorem (Completeness)

Let $\mathcal{L}$ be a language. Let T be an $\mathcal{L}$-theory with φ an $\mathcal{L}$-formula. Suppose $T \vDash \varphi$. Therefore $T \vdash \varphi$.

Proof .:.

Suppose $T \nvdash \varphi$. Thus $T \cup \{\neg\varphi\}$ is consistent. If $T \cup \{\neg\varphi\}$ has a model, then $T \nvDash \varphi$. We will construct a model of $T \cup \{\neg\varphi\}$ out of syntax. We do this just by well-ordering the $\mathcal{L}$-sentences, and expanding $T \cup \{\neg\varphi\}$ to a $\mathcal{L}$-theory T_0 which is consistent and complete.

Now by well-ordering T_0, for each existential statement φ being $\exists x \psi(x) \in T_0$, associate a unique constant c_φ, and add in the statement $\psi(c_\varphi)$ to a new theory T_1 in an expanded language $\mathcal{L}_1$. Also expand to make sure T_1 is consistent and complete in $\mathcal{L}_1$. Repeat this process ω times, and you end up with a theory T_ω in an expanded language $\mathcal{L}_\omega$ such that if $\exists x \psi(x)$ is in T_ω, then $\psi(c) \in T_\omega$ for some constant symbol $c \in \mathcal{L}_\omega$ and T_ω is consistent and complete.

For constant symbol $c \in \mathcal{L}_\omega$, consider the equivalence class $[c] = \{d \in \mathcal{L}_\omega : T_\omega \vdash d = c\}$. This will be an equivalence class as T_ω is complete. Consider the structure M with universe $M = \{[c] : c \in \mathcal{L}_\omega \text{ a constant symbol}\}$ and with relation interpretations

$$f^{\mathsf{M}}([d_0], \cdots, [d_{n-1}]) = [d_n] \text{ iff } T_\omega \vdash f(d_0, \cdots, d_{n-1}) = d_n,$$

and similarly for relations. It's not difficult to see that this definition is well-defined, and that the resulting structure $\mathsf{M} \vDash T_\omega$ and hence the reduct to $\mathcal{L}$ models $T \cup \{\neg\varphi\}$. ⊣

Following easily from completeness is compactness. Compactness will be significantly more useful for us than completeness, since one is a statement about the existence of proofs, and the other is a statement about the existence of models. In a chapter called "Model Theory", it seems obvious that one will be preferred.

1.A • 2. Corollary (Compactness)

Let $\mathcal{L}$ be a language. Let T be an $\mathcal{L}$-theory. Therefore T has a model iff each finite subset $\Delta \subseteq T$ has a model.

Proof .:.

If T has a model, clearly every finite subset does. So suppose T does not have a model. Therefore $T \vDash \varphi \wedge \neg\varphi$. By Completeness (1.A • 1), $T \vdash \varphi \wedge \neg\varphi$. Proofs are taken to be finite sequences of formulas. Hence the set of formulas of T which occur in such a proof is a finite subset of T, Δ. But then $\Delta \vdash \varphi \wedge \neg\varphi$ so that $\Delta \vDash \varphi \wedge \neg\varphi$, meaning Δ has no model. ⊣

The two results Completeness (1.A • 1) and Compactness (1.A • 2) are actually equivalent, although the proof of this fact is quite round-about. As a small side note, for finite languages, we don't need the axiom of choice, because we can do these well-orders lexiographically through whatever fixed order on the finite elements. For infinite languages in general, we need choice, and this is what is meant by the two theorems are equivalent: ZF proves the equivalence.

§ 1.B. Löwenheim–Skolem

Before we prove the very useful result, we must prove a lemma which reduces being an elementary substructure to satisfying certain formulas.

1.B • 1. Lemma (The Tarksi–Vaught Test)

Let $\mathcal{L}$ be a language. Let M be an $\mathcal{L}$-model with N a proper $\mathcal{L}$-submodel. Therefore $\mathsf{N} \preccurlyeq \mathsf{M}$ iff for every $\mathcal{L}$-formula $\varphi(x, \vec{w})$ and $\vec{n} \in N^{<\omega}$,

$$\mathsf{M} \vDash \exists x \varphi(x, \vec{n}) \quad \text{iff} \quad \mathsf{N} \vDash \exists x \varphi(x, \vec{n}).$$

Proof .:.

This is really just a proof by structural induction where we just assume the quantifier case goes through. Clearly if $\mathsf{N} \preccurlyeq \mathsf{M}$, then we get the existential result. Let $\vec{n}_0$ and $\vec{n}_1$ be sequences of parameters from N. Let φ and ψ be $\mathcal{L}$-formulas with parameters in N. Clearly if φ or ψ is atomic, then the result holds, since N and M have the same interpretations of the relation and function symbols. If φ is $\neg\psi$, then the result clearly holds by the inductive hypothesis as a result of the "iff". If we consider $\varphi \wedge \psi$, then if $\mathsf{M} \vDash \varphi$ and $\mathsf{M} \vDash \psi$, then by the inductive hypothesis $\mathsf{N} \vDash \varphi$ and $\mathsf{N} \vDash \psi$. So the result holds here too. The existential case is given to us, and thus the result holds. ⊣

1.B • 2. Theorem (Löwenheim–Skolem)

Let $\mathcal{L}$ be a language. Let $\kappa \geq |\mathcal{L}|$ be an infinite cardinal. Let T be an $\mathcal{L}$-theory with an infinite model. Therefore T has a model of size κ.

Proof .:.

There are two parts to this proof: an upward, and a downward version. We have a model $\mathsf{M} \vDash T$, but the trouble is finding a larger, and a smaller model depending on how κ relates to this model. If $\kappa = |M|$ we're done.

For $\kappa > |M|$, consider the language $\mathcal{L}' = \mathcal{L} \cup \{c_\alpha : \alpha < \kappa\}$ adding new constant symbols. Consider the expanded theory $T' = T \cup \{c_\alpha \neq c_\beta : \alpha < \beta < \kappa\}$. It's easy to see that every finite subset of this has a model as an expansion of M, and thus T' has a model. But such a model must have at least $\kappa = |\mathcal{L}'|$ many elements. We want to then find a model with exactly κ many elements. So to finish off the upward version, it suffices to show the downward direction.

Without loss of generality, let $\kappa = |\mathcal{L}|$ just by expanding the language as desired. Suppose $|M| > \kappa$. For each $\mathcal{L}$-formula $\psi(x, \vec{w})$, associate a new function symbol f_ψ and $\mathcal{L} \cup \{f_\psi\}$-sentence

$$\psi' := \forall \vec{w} (\exists x \psi(x, \vec{w}) \to \psi(f_\psi(\vec{w}), \vec{w}).$$

The theory $\{\exists x \psi(x), \psi' : \mathsf{M} \vDash \exists x \psi(x)\}$ has a model which is just an expansion of M—such f_ψ^{M}s exist by choice applied to M. Note that we are working in an expanded language $\mathcal{L}'$. Note further that $|\mathcal{L}'| = |\mathcal{L}| + \aleph_0 = \kappa$.

Consider a subset of M of size κ, say $A \subseteq M$. By recursively applying the functions of $\mathcal{L}'$, we can take the closure of A under these functions, say $B \subseteq M$. It's a simple exercise to show $|B| = |A| \cdot |\mathcal{L}'| \cdot \aleph_0 = \kappa$. But then the submodel $\mathsf{B} \subseteq \mathsf{M}$ has that for $\psi(x, \vec{w})$ an $\mathcal{L}$-formula and $\vec{b} \in B^{<\omega}$, if $\mathsf{M} \vDash \exists x \psi(x, \vec{b})$ then $\mathsf{B} \vDash \exists x \psi(x, \vec{b})$. By The Tarski–Vaught Test (1.B • 1), it follows that $\mathsf{B} \preccurlyeq \mathsf{M}$, and so we get a model $\mathsf{B} \vDash T$ with $|B| = \kappa$. ⊣

The proof of the downward direction actually shows that if $X \subseteq M$ with $|M| \geq \aleph_0$, then there is an elementary substructure $\mathsf{N} \preceq \mathsf{M}$ such that $X \subseteq N$ and $|N| = \kappa$. Such an idea is very useful for all sorts of applications. We will return to the idea of this proof later with the ideas of skolem functions.

1.B • 3. Corollary (Löwenheim–Skolem)

Let $\mathcal{L}$ be a language with M an $\mathcal{L}$-model with $|M| \geq \aleph_0$, and $X \subseteq M$.
Let $\kappa \geq |\mathcal{L}| + |X| + \aleph_0$ be a cardinal.
Therefore there is an elementary substructure $\mathsf{N} \preccurlyeq \mathsf{M}$ such that $X \subseteq N$, and $|N| = \kappa$.

Such a structures is often referred to as a *skolem hull*, and will be denoted here by $\mathsf{Hull}(X) \preccurlyeq \mathsf{M}$ for $X \subseteq M$, having universe $\mathrm{Hull}(X)$. If we want to be clear where we are taking the hull, we might write $\mathsf{Hull}_{\mathsf{M}}(X)$.

This also gives the easy corollary about complete theories. We will assume that T has infinite models to ensure that T isn't κ-categorical vacuously. We also want to avoid finite models, which can be categorical just by virtue of their number of elements, e.g. if T has a model with just one element.

1.B • 4. Corollary (The Łoś–Vaught Test)

Let $\mathcal{L}$ be a language. Let T be a $\mathcal{L}$-theory with infinite models. If T is κ-categorical for any infinite cardinal κ, then T is complete.

Proof .:.

If $\kappa < \omega$, the result is clear. If κ is infinite, but T isn't complete, say $T \cup \{\varphi\}$ and $T \cup \{\neg\varphi\}$ are consistent, use Löwenheim–Skolem (1.B • 3) to get a model of size κ of each of them: $\mathsf{A} \vDash T + \varphi$, $\mathsf{B} \vDash T + \neg\varphi$. This contradicts κ-categoricity, since isomorphic structures are elementarily equivalent. ⊣

When dealing with chains of elementary submodels, we get the following theorem—not to be confused with The Tarksi–Vaught Test (1.B • 1).

1.B • 5. Theorem (The Tarski–Vaught Theorem)

Let $\mathcal{L}$ be a language. Let M_α be an $\mathcal{L}$-model for all $\alpha < \gamma \in \mathrm{Ord}$ with γ a limit ordinal. Suppose $\mathsf{M}_\alpha \preccurlyeq \mathsf{M}_\beta$ for all $\alpha < \beta < \gamma$. Therefore there is a model $\bigcup_{\alpha\in\gamma} \mathsf{M}_\alpha$ where $\mathsf{M}_\beta \preccurlyeq \bigcup_{\alpha\in\gamma} \mathsf{M}_\alpha$ for all $\beta < \gamma$.

Proof .:.

The direct limit $\bigcup_{\alpha\in\gamma} \mathsf{M}_\alpha$ is just given by the union of the corresponding models: the universe is the union of the universes, the relations are the unions of the relations, and the functions are the unions of the functions. The constants are necessarily the constants as interpreted by M_0.

Using this, it's a simple exercise to see that any $\mathcal{L}_{\bigcup_{\alpha<\gamma} M_\alpha}$-formula modeled by $\bigcup_{\alpha<\gamma} \mathsf{M}_\alpha$ will use only finitely many parameters, and so will be modeled by some M_β, $\beta < \gamma$, containing all of them; and vice versa. $\dashv$

From now on, $\mathcal{L}$ will always be a language, and always the default language we're working in.

Section 2. Quantifier Elimination

The process of quantifier elimination is exactly what the name suggests: making formulas quantifier free. To fully explain this, however, we need to include some new definitions about syntax.

2 • 1. Definition

An $\mathcal{L}$-formula φ is *quantifier free* iff it does not contain the symbols $\exists$ and $\forall$, or φ is $\exists x(x = x)$ or its negation, denoted tt and ff respectively.

The benefit of allowing $\exists x(x = x)$ and $\neg\exists x(x = x)$ is the ability to reduce a statement to just "true" or "false" without introducing any more official symbols into the vocabulary of first-order logic. This is especially useful if the language doesn't have any constant symbols: any sentence will need to reduce to either tt or ff.

2 • 2. Definition

Let T be an $\mathcal{L}$-theory. T admits quantifier elimination iff for every $\mathcal{L}$-formula $\varphi(\vec{x})$, there is a quantifier-free formula $\psi(\vec{x})$ such that $T \vdash \varphi(\vec{x}) \leftrightarrow \psi(\vec{x})$.
An $\mathcal{L}$-structure A admits quantifier elimination iff $\mathrm{Th}(\mathsf{A})$ does.

In general, most theories do not admit quantifier elimination. An easy example of this would be ZFC. If $\varphi(x, y) = |\mathcal{P}(x)| = \aleph_y$ were equivalent to some quantifier free formula, then it could be put in the form of the disjunction of a bunch of conjunctions of literals. Such statements would be absolute between transitive models, however, which would imply, for example, $|\mathcal{P}(\omega)| = \aleph_1$ as a statement true in $L \subseteq V$.

An easy example of a theory which *does* admit quantifier elimination would be the models of pure equality, with T saying "there are infinitely many elements". To see how easy this example is, however, it will be useful to build enough theory to apply some tests.

§ 2.A. Theory and tests

In some sense, quantifier elimination is a measure of how simple a structure is in that models that admit quantifier elimination are very basic. We have the following test which allows us to see whether a theory/structure admits quantifier elimination.

2.A • 1. Theorem (Quantifier Elimination Test)

Let A be an $\mathcal{L}$-structure.
Suppose every formula of the form $\exists x(\chi_0 \wedge \cdots \wedge \chi_n)$ for $\mathcal{L}$-literals χ_i, $i \leq n$, is equivalent under A to a quantifier free formula. Therefore A admits quantifier elimination.

Proof .:.

In essence, this is just a proof by structural induction where we remove the existential case. Let $\varphi(\vec{x})$ be an arbitrary $\mathcal{L}$-formula. If $\varphi(\vec{x})$ is atomic, then φ is already quantifier free. If $\varphi(\vec{x})$ is the conjunction or negation of some formulas, then by the inductive hypothesis, the result holds. So all that remains is the existential case: $\varphi(\vec{x})$ is $\exists y\, \psi(y, \vec{x})$. By the inductive hypothesis, $\psi(y, \vec{x})$ can be assumed to be quantifier free. But then we can write $\psi(y, \vec{x})$ in disjunctive normal form, and so get that $\varphi(\vec{x})$ is equivalent under A to the disjunction of formulas as in the statement of the theorem. By hypothesis, these are equivalent to quantifier free formulas, and hence so is $\varphi(\vec{x})$. $\dashv$

Note that if we can effectively eliminate quantifiers from these formulas, then A admits effective quantifier elimination, meaning that there is a computable procedure for eliminating them from arbitrary formulas.

There is another way of understanding quantifier elimination in terms of structures. Namely, if structures of a theory always agree on their common substructures, then their theories admit quantifier elimination.

2.A • 2. Result

Let $\mathcal{L}$ be a language and T an $\mathcal{L}$-theory. Suppose $\mathrm{Th}_M(\mathbf{A}) = \mathrm{Th}_M(\mathbf{B})$ for all $\mathcal{L}$-models $\mathbf{A} \vDash T$ and $\mathbf{B} \vDash T$ with $\mathbf{M}$ a common substructure: $\mathbf{M} \subseteq \mathbf{A}, \mathbf{B}$. Therefore T admits quantifier elimination.

Proof .:.

To prove this, let $\varphi(\vec{x})$ be an arbitrary $\mathcal{L}$-formula. Without loss of generality, assume both $T \cup \{\exists x \varphi(\vec{x})\}$ and $T \cup \{\exists x \neg\varphi(\vec{x})\}$ are consistent. This is because $T \vdash \varphi(\vec{x})$ iff $T \vdash \varphi(\vec{x}) \leftrightarrow \mathrm{tt}$. Consider the set of quantifier free formulas that φ implies relative to T:

$$C(\vec{x}) = \{\psi(\vec{x}) : \psi \text{ is quantifier free} \wedge T \vdash \varphi(\vec{x}) \to \psi(\vec{x})\}.$$

Consider fresh constant symbols $\{c_i : i \leq n\}$—enough to fill up $\vec{x}$—in the expanded language $\mathcal{L}' = \mathcal{L} \cup \{c_i : i \leq n\}$.

If we can show that $T \cup C(\vec{c}) \vDash \varphi(\vec{c})$, then the result holds. To see this, by Compactness (1.A • 2), there would be some finite subset $\{\psi_i(\vec{c}) : i \leq m\} \subseteq C(\vec{c})$ where $T \vdash \bigwedge_{i \leq m} \psi_i(\vec{c}) \to \varphi(\vec{c})$. Since the converse holds by construction of C, and the constant symbols were new, we would get that T would prove the desired equivalence. So the rest of the proof will show that $T \cup C(\vec{c}) \vDash \varphi(\vec{c})$, which is where the hypotheses of the result come into play.

So suppose not: get a model $\mathbf{A} \vDash T \cup C(\vec{c}) \cup \{\neg\varphi(\vec{c})\}$. Consider the substructure $\mathbf{M} \subseteq \mathbf{A}$ generated by $\vec{c}^{\mathbf{A}}$, meaning the structure whose universe is the closure of $\vec{c}$ under the functions of $\mathcal{L}'$, with relations interpreted as in $\mathbf{A}$ restricted to this universe. Note that $\mathbf{A}$ and $\mathbf{M}$ agree on quantifier free formulas so that $\mathbf{M} \vDash C(\vec{c})$.

In the expanded language $\mathcal{L}'_M$, we can also get an elementary extension $\mathbf{B} \vDash T \cup \mathrm{Eldiag}(\mathbf{M}) \cup \{\varphi(\vec{c})\}$. To see this, otherwise $T \cup \mathrm{Eldiag}(\mathbf{M}) \vDash \neg\varphi(\vec{c})$, whence by Compactness (1.A • 2), we get quantifier free $\mathcal{L}'_M$-formulas $\psi_i \in \mathrm{Eldiag}(\mathbf{M})$ where $T \vdash \bigwedge_{i \leq m} \psi_i(\vec{c}) \to \neg\varphi(\vec{c})$, and thus $T \vdash \varphi(\vec{c}) \to \bigvee_{i \leq m} \neg\psi_i(\vec{c})$. But then this big disjunction is in $C(\vec{c})$ so that $\mathbf{M} \vDash \bigvee_{i \leq m} \neg\psi_i(\vec{c})$, contradicting that $\psi_i(\vec{c}) \in \mathrm{Eldiag}(\mathbf{M})$ for every $i \leq m$.

But as an elementary extension, $\mathbf{B} \equiv \mathbf{M}_M$, which means that $\mathbf{M}_M \vDash \varphi(\vec{c})$. Because $\mathrm{Th}_M(\mathbf{A}) = \mathrm{Th}_M(\mathbf{B})$, it follows that $\mathbf{A}_M \vDash \varphi(\vec{c})$, contradicting the hypothesis on $\mathbf{A}$. Hence no such $\mathbf{A}$ can exist, and thus $T \cup C(\vec{c}) \vdash \varphi(\vec{c})$. By the argument above, this gives the result. $\dashv$

The above result is actually equivalent to quantifier elimination, since if T admits quantifier elimination, then for φ an $\mathcal{L}$-formula with $\vec{m} \in M^{<\omega}$, φ as a quantifier free T-equivalent in $\mathbf{A}$: ψ. Hence $\mathbf{A} \vDash \varphi(\vec{m})$ iff $\mathbf{A} \vDash \psi(\vec{m})$. Since substructures and superstructures agree on quantifier free formulas, this happens iff $\mathbf{M} \vDash \psi(\vec{m})$, which holds iff $\mathbf{B} \vDash \psi(\vec{m})$. iff $\mathbf{B} \vDash \varphi(\vec{m})$. Hence $\mathrm{Th}_M(\mathbf{A}) = \mathrm{Th}_M(\mathbf{B})$.

The two tests above actually combine to give a third test for quantifier elimination.

2.A • 3. Corollary

Let T be an $\mathcal{L}$-theory. Suppose for arbitrary $\mathcal{L}$-models $\mathbf{A}, \mathbf{B} \vDash T$ with common substructure $\mathbf{M} \subseteq \mathbf{A}, \mathbf{B}$, and $\vec{m} \in M^{<\omega}$ and quantifier-free $\mathcal{L}$-formula φ that if $\mathbf{A} \vDash \exists x \varphi(\vec{m}, x)$ then $\mathbf{B} \vDash \exists x \varphi(\vec{m}, x)$.
Therefore T admits quantifier elimination.

These are the most basic techniques for testing whether a theory admits quantifier elimination. The usefulness of quantifier elimination comes from a couple places. The definable relations over a model that admits quantifier elimination can be reduced to the quantifier free relations in the same number of variables or fewer. In this sense, the definable sets are quite simple.

One of the issues with these techniques, however, is that they are almost brute force in a way: one must slog through all the different combinations of literals, and work with them; or else work through all the kinds of substructures a model can have. These methods can be done quickly enough for some simple theories in small enough languages, but quickly become unmanageable with more complicated theories. The interaction of model completeness with quantifier elimination provides some help with this with some algebraically motivated examples.

§ 2.B. Association with model completeness

A consequence of quantifier elimination is model completeness, a notion that produces elementary substructures from substructures. This notion can be generalized to theories being *model completions* of other theories.

2.B • 1. Definition

Let T and T_0 be $\mathcal{L}$-theories.
T is model complete iff for all $\mathcal{L}$-structures A, if $\mathsf{B} \subseteq \mathsf{A}$ and $\mathsf{B} \vDash T$ then $\mathsf{B} \preccurlyeq \mathsf{A}$.
T is the model completion of T_0 iff $T_0 \subseteq T$ and $T \cup \mathrm{Diag}(\mathsf{A})$ is a complete $\mathcal{L}_A$-theory for every $\mathsf{A} \vDash T_0$.

In another sense, a theory is model complete iff all embeddings between models of T are elementary embeddings. Note that the model completion T of a theory is necessarily model complete. To see this, suppose $\mathsf{A} \vDash T$ with $\mathsf{M} \subseteq \mathsf{A}$ and $\mathsf{M} \vDash T$. It's clear that then both $\mathsf{M}_M, \mathsf{A}_M \vDash T \cup \mathrm{Diag}(\mathsf{M})$, which is a complete $\mathcal{L}_M$-theory. In particular, $T \cup \mathrm{Diag}(\mathsf{M}) \vDash \mathrm{Eldiag}(\mathsf{M})$ so that $\mathsf{M} \preccurlyeq \mathsf{A}$.

Note that not all theories have model completions—e.g. the group axioms—but when they do, the completions are unique in a certain sense. While it's not quite correct to say that T is *the* model completion of another theory T_0, the point is that we won't have two model completions T, T' which disagree: $T \vDash T'$ and $T' \vDash T$.

2.B • 2. Result

Let T_0 be a $\mathcal{L}$-theory. Let T and T' be $\mathcal{L}$-theories which are model completions of T_0. Therefore $T \vDash T'$ and vice versa.

Proof .:.

Let $\mathsf{A} \vDash T$ be arbitrary. As a model of T_0, we get that the theory $T' \cup \mathrm{Diag}(\mathsf{A})$ is complete, and consistent in particular. Any $\mathcal{L}_A$-model $\mathsf{B} \vDash T' \cup \mathrm{Diag}(\mathsf{A})$ contains $\mathsf{A} \subseteq \mathsf{B}$ as a submodel. So we can take the $\mathcal{L}$-reduct $\mathsf{B}' \vDash T'$ with $\mathsf{A} \subseteq \mathsf{B}'$.

Using this repeatedly, by symmetry, we get a chain of $\mathcal{L}$-models
$$\mathsf{A} = \mathsf{A}_0 \subseteq \mathsf{A}_1 \subseteq \mathsf{A}_2 \subseteq \cdots,$$
where $\mathsf{A}_n \vDash T$ if n is even, and $\mathsf{A}_n \vDash T'$ if n is odd. As model completions, we then have $\mathsf{A}_{2n} \preccurlyeq \mathsf{A}_{2n+2}$ for all $n \in \omega$. Hence by The Tarski–Vaught Theorem (1.B • 5),
$$\mathsf{A}_{2m} \preccurlyeq \bigcup_{n\in\omega} \mathsf{A}_{2n} = \bigcup_{n\in\omega} \mathsf{A}_n$$
for all $m \in \omega$. Yet the same argument applied to $2n+1$ and T' yields that $\mathsf{A}_{2m+1} \preccurlyeq \bigcup_{n\in\omega} \mathsf{A}_n$ and thus that the union models T'. Hence $\mathsf{A} \preccurlyeq \bigcup_{n\in\omega} \mathsf{A}_n \vDash T'$ so that $\mathsf{A} \vDash T'$. ⊣

Despite model completions being somewhat uncommon, for theories which do have model completions, we get lots of nice properties. One way to get model complete theories—which aren't necessarily model completions—is through quantifier elimination.

2.B • 3. Result

Let T be an $\mathcal{L}$-theory. Suppose T admits quantifier elimination. Therefore T is model complete.

Proof .:.

Let $\mathsf{A} \vDash T$ be an arbitrary model with submodel $\mathsf{B} \vDash T$. We must show that $\mathsf{B} \preccurlyeq \mathsf{A}$. To do this, let $\varphi(\vec{x})$ be an $\mathcal{L}$-formula with $\vec{b} \in B^{<\omega}$. Note that $\mathsf{A} \vDash \varphi(\vec{b})$ iff $\mathsf{A} \vDash \psi(\vec{b})$ for some quantifier free $\mathcal{L}$-formula ψ. But as substructures and superstructures agree on quantifier free formulas, this is equivalent to $\mathsf{B} \vDash \psi(\vec{b})$. Since $\mathsf{B} \vDash T$, which proves the equivalence of ψ and φ, this holds iff $\mathsf{B} \vDash \varphi(\vec{b})$, and thus the result. ⊣

Model completeness is not equivalent to quantifier elimination, but we do have an equivalence for some kinds of theories. For example, the theory of real closed fields in the language of rings is model complete, but doesn't admit elimination of quantifiers. The only barrier to elimination of quantifiers is the ordering: the theory of real closed fields in the language of orderings *does* admit elimination of quantifiers, and this is how we can show the model completeness of the theory in the language of rings. In particular, for theories which can be axiomatized with only universal quantifiers—so-called $\forall$-theories—model completions admit elimination of quantifiers. Examples of $\forall$-theories include the axioms for integral domains and ordered rings in the proper languages.

2.B • 4. Theorem

Let T_0 be a $\forall$-theory—an $\mathcal{L}$-theory with an axiomatization using only universal quantifiers. Let T be the model completion of T_0. Therefore T admits quantifier elimination.

Proof .:.

This is really a consequence of Result 2.A • 2. Note that any substructure of a $\forall$-theory is itself a model of that theory: $\mathsf{M} \subseteq \mathsf{A}$ with $\mathsf{A} \vDash T$ implies $\mathsf{M} \vDash T_0$.

So suppose we have to models with a common substructure $\mathsf{A}, \mathsf{B} \vDash T$, and $\mathsf{M} \subseteq \mathsf{A}, \mathsf{B}$. Thus $\mathsf{M} \vDash T_0$. As the model completion, $T_0 \cup \text{Diag}(\mathsf{M})$ is a complete $\mathcal{L}_M$ theory. But this theory is modeled by both A and B, meaning that $\text{Th}_M(\mathsf{A}) = \text{Th}_M(\mathsf{B})$. Hence by Result 2.A • 2, T admits quantifier elimination. ⊣

§ 2.C. Examples and non-examples

The simplest example of a theory which admits quantifier elimination is the theory of dense linear orders without endpoints (DLO). More complicated examples include non-trivial, torsion-free, divisible, abelian groups (DAG); as well as non-trivial, divisible, ordered, abelian groups (ODAG), and algebraically closed fields (ACF).

Some non-examples include presburger arithmetic, which is the $\{+, -, <, 0, 1\}$-theory of $\mathbb{Z}$ as an ordered group of integers. A perhaps unexpected result is that the theory of real closed fields (RCF) does not admit quantifier elimination. This is partly because models of ACF are the only infinite fields admitting quantifier elimination in $\mathcal{L} = \{0, 1, +, -, \cdot\}$. The primary—and in fact *only*—obstruction is ordering: RCF admits quantifier elimination in $\mathcal{L} = \{0, 1, +, -, \cdot, <\}$.

2.C • 1. Example (DLO)

The theory of dense linear orders without endpoints (DLO) in the language $\mathcal{L} = \{<\}$ admits quantifier elimination.

Proof .:.

Proceed as in Quantifier Elimination Test (2.A • 1). We want formulas φ of the form $\exists x(\chi_0(x, \vec{y}) \wedge \cdots \wedge \chi_n(x, \vec{y}))$ to be equivalent to quantifier free formulas where χ_i is a literal. Note that the literals of $\mathcal{L}$ will be of the form

$$x < y \qquad x = y \qquad x \neq y \qquad x \nless y,$$

which reduce just to $x < y$ and $x = y$. The reason why we can reduce to this is just that $x \neq y$ and $x \nless y$ are equivalent to $x < y \vee y < x$ and $y < x \vee x = y$ respectively. Distributing over the conjunctions and existential quantifier gives the result if we can show the reduced case.

To show the reduced case, we will put the χ_is in to blocks. If $x = v$ occurs as one of the literals, then we can replace x with v, and get the equivalence of φ and $\chi(v, \vec{y}) \wedge \cdots \wedge \chi(v, \vec{y})$. Otherwise, φ doesn't have any equalities with x in them. In fact, just by replacing variables, we can assume none of the literals are statements of equality. So assume $x < y_i$ for various $y_i \in Y$, and $v_i < x$ for various $v_i \in V$ for some set Y and V of variables. If either is $\emptyset$, we just reduce to the relations between the variables which aren't x, since DLO states there are no endpoints. As there are only finitely many such y_i and v_is, and the order under DLO is dense, φ is equivalent to the statement $V < Y \wedge \chi$ where χ is the conjunction of all the literals of φ which don't include x, and $V < Y$ just means $v_i < y_j$ for all appropriate i, j.

To see this, one direction is clear: if φ holds, then $V < Y \wedge \chi$. For the converse, if $V < Y \wedge \chi$, then by DLO, there is some element between all of V and all of Y. This means $\exists x(V < x \wedge x < Y \wedge \chi)$. But this is just φ. Hence φ is equivalent to the quantifier free formula $V < Y \wedge \chi$, and so by Quantifier Elimination Test (2.A • 1), the result holds. ⊣

Before proceeding with the proof that ACF admits quantifier elimination, we should explicitly give the axioms we're using: in the language $\mathcal{L}_{\text{rings}} = \{0, 1, +, -, \cdot\}$, ACF consists of the following axioms.

2.C • 2. Definition

ACF is the set of the following axioms:

1. $1 \neq 0$;
2. $\forall x \forall y(x \cdot y = y \cdot x \wedge x + y = y + x)$;
3. $\forall x \forall y \forall z((x \cdot y) \cdot z = x \cdot (y \cdot z) \wedge (x + y) + z = x + (y + z))$;
4. $\forall x \forall y \forall z(x - y = z \leftrightarrow x = z + y)$.
5. $\forall x \forall y \forall z(x \cdot (y + z) = x \cdot y + x \cdot z)$;
6. $\forall x(1 \cdot x = x \wedge 0 + x = x)$;
7. $\forall x(x - x = 0)$;
8. $\forall x \forall y(x \cdot y = 0 \to x = 0 \vee y = 0)$;
9. $\forall x \exists y(x \neq 0 \to x \cdot y = 1)$;
10. For each $n < \omega$, the axiom $\forall a_0 \cdots \forall a_n \exists x(a_0 + a_1 \cdot x + \cdots + a_n(\underbrace{x \cdot \ldots \cdot x}_{n}) = 0)$.

Note that the theory of commutative rings is given by (1)–(6). The theory of integral domains is given by (1)–(7). (1)–(8) gives the theory of fields, and the addition of (9) gives algebraic completeness. If we consider only the universal sentences of ACF—axioms (1)–(7) in Definition 2.C • 2, hereafter called $\mathsf{ACF}_\forall$—we get a $\forall$-theory of integral domains

2.C • 3. Example (ACF admits Quantifier Elimination)

The theory of algebraically closed fields (ACF), in the language of rings, $\mathcal{L} = \{0, 1, +, -, \cdot\}$, admits quantifier elimination.

Proof .:.

We want to show that ACF is the model completion of the axioms for integral domains. Note that the use of '$-$' in the language here is crucial to this characterization, but not to the end result as it's a conservative extension of the axioms in the more natural language of $\{0, 1, +, \cdot\}$. To do this, note that $\mathsf{ACF}_\forall$ is a $\forall$-theory in $\mathcal{L}$. Hence it suffices by Theorem 2.B • 4 to show that ACF is the model completion of $\mathsf{ACF}_\forall$. To do this, we proceed in two parts, first showing that every integral domain embeds into an algebraically closed field. Second, we show that $\mathsf{ACF} \cup \mathrm{Diag}(\mathbf{D})$ is complete for every integral domain $\mathbf{D}$. So let $\mathbf{D} \vDash \mathsf{ACF}_\forall$ be an integral domain.

To embed $\mathbf{D}$ into an algebraically closed field, consider the field of fractions $\mathbf{F}$, an $\mathcal{L}$-model with universe $D \times (D \setminus \{0^{\mathbf{D}}\})$, representing $\langle a, b\rangle$ instead by a/b, and with $\mathbf{F}_D$ interpretting $\mathcal{L}_D$ as follows:

$$\frac{a}{b} \pm^{\mathbf{F}} \frac{a'}{b'} = \frac{a \cdot^{\mathbf{D}} b' \pm a' \cdot^{\mathbf{D}} b}{b \cdot^{\mathbf{D}} b'}, \qquad \frac{a}{b} \cdot^{\mathbf{F}} \frac{a'}{b'} = \frac{a \cdot^{\mathbf{D}} a'}{b \cdot^{\mathbf{D}} b'},$$

and for $d \in D$, $d^{\mathbf{F}} = d/1^{\mathbf{D}}$. It's not hard to see $\mathbf{F}_D \vDash \mathrm{Diag}(D)$ along with (1)–(8) in Definition 2.C • 2. Note that $\mathbf{F}$ is then a field. Taking its algebraic closure $\mathbf{K}$ then yields that $\mathbf{D} \subseteq \mathbf{F} \subseteq \mathbf{K}$. This shows that every integral domain embeds into an algebraically closed field.

To show that $\mathsf{ACF} \cup \mathrm{Diag}(\mathbf{D})$ is complete, we can further appeal to some knowledge of algebra, which is appropriate given that this is an example motivated by algebra. In particular, for each characteristic $p > 0$, consider

$$\mathsf{ACF}_p := \mathsf{ACF} \cup \Big\{\underbrace{1 + \cdots + 1}_{p} = 0\Big\}.$$

This last sentence—or its negation—will be given by the diagram of $\mathbf{D}$. Similarly, if $p = 0$,

$$\mathsf{ACF}_0 := \mathsf{ACF} \cup \Big\{\underbrace{1 + \cdots + 1}_{n} \neq 0 : n \in \omega\Big\}$$

will again have the non-ACF part determined by $\mathrm{Diag}(\mathbf{D})$. Thus it suffices to show that ACF_p is complete for $p \geq 0$. To do this, note that two algebraically closed fields of the same characteristic p and transcendence degree λ (over their ring of integers) are isomorphic. If $\mathbf{K} \vDash \mathsf{ACF}_p$ has transcendence degree is just λ, then $|K| = \lambda + \aleph_0$. Hence if $\mathbf{K} \vDash \mathsf{ACF}_p$ is uncountable, its cardinality is its transcendence degree. Thus any two uncountable, algebraically closed fields of the same characteristic and cardinality are isomorphic, and so ACF_p is κ-categorical for every $\kappa > \aleph_0$. Therefore by The Tarksi–Vaught Test (1.B • 1), ACF_p is complete. $\dashv$

Section 3. *Realizing and Omitting Types*

This section will be about types in models and theories. Realizing or omitting a type is similar to what can be expressed by higher order and infinitary logics. For example, being well-founded cannot be expressed by a theory in first order logic as a result of compactness. This can be expressed in second order logic quite easily though: for R our relation, R is ill-founded iff

$$\exists A\ \forall x\ \exists y\ (A(x) \wedge A(y) \wedge y\ R\ x) \wedge x \neq y).$$

First order logic can still be expanded in some sense to allow for this expression through the idea of omitting a type: there is no assignment to the infinite number of variables $\langle x_n : n \in \omega \rangle$ such that all the formulas of $\{x_{n+1}\ R\ x_n : n \in \omega\}$ are true in your expanded model. Let's now formally introduce the notion of a type.

3 • 1. Definition

Let A be an $\mathcal{L}$-model. Let $X \subseteq A$. A type of A over X is a set of $\mathcal{L}_X$-formulas $\Sigma(\vec{x})$ such that

1. the free variables of $\varphi \in \Sigma(\vec{x})$ are among $\vec{x}$;
2. $\mathrm{Th}_X(\mathsf{A}) \cup \Sigma(\vec{x})$ is *consistent* in the sense that replacing variables with fresh constant symbols C yields a consistent $\mathcal{L} \cup C$-theory.

$\Sigma(\vec{x})$ is an n-type iff $\Sigma(\vec{x})$ is a type, and $\vec{x}$ has at most n variables.
A type $\Sigma(\vec{x})$ over X is complete iff $\varphi \in \Sigma(\vec{x})$ or $\neg\varphi \in \Sigma(\vec{x})$ for every $\mathcal{L}_X$-formula $\varphi(\vec{x})$.
The set of all complete n-types of A over X is denoted $S_n^{\mathsf{A}}(X)$.

We will usually refer just to 1-types, but will occasionally refer to larger n-types, or even κ-types for $\kappa \geq \aleph_0$. Note that in the future, by *type* we just mean a *finitary type*, meaning n-type for $n \in \omega$. For non-finitary types, we will write "infinitary type". Despite being consistent, it may not be the case that the theory is compatible with the model in question. This motivates the idea of realizing versus omitting a type.

3 • 2. Definition

Let A be an $\mathcal{L}$-model with $X \subseteq A$. Let $\Sigma(\vec{x})$ be a type of A over X.
A realizes $\Sigma(\vec{x})$ iff there is some assignment $\vec{a} \in A^{|\vec{x}|}$ such that $\mathsf{A}_X \vDash \Sigma(\vec{a})$. Otherwise A omits $\Sigma(\vec{x})$.

Omitting a type in some sense requires an infinite amount of information—every element doesn't have this property—while realizing a type only requires a small amount of information—this element has these properties. Of course, confirming that either holds is often a difficult procedure. But there are some easy examples to illustrate the concept.

In particular, consider the ordering $\langle \mathbb{Q}, < \rangle$. The 1-type $\{x > n : n \in \mathbb{N}\}$ over $\mathbb{N} \subseteq \mathbb{Q}$ is omitted by $\mathbb{Q}$. The 1-type $\Sigma(x) = \{x < 1/n : n \in \mathbb{N}\}$ over $\mathbb{Q}$ is realized by $\mathbb{Q}$: $\langle \mathbb{Q}, < \rangle_{\mathbb{Q}} \vDash \Sigma(0)$, for example.

In general, for a model A with $X \subseteq A$ and $a \in A$, we can consider the *type* of a over X just by considering the $\mathcal{L}_X$-formulas which are true of a:

$$\mathrm{tp}^{\mathsf{A}}(a/X) = \{\varphi(v) : \mathsf{A}_X \vDash \varphi(a)\}.$$

Such a type is realized in A, and is a complete type, i.e. $\mathrm{tp}^{\mathsf{A}}(a/X) \in S_1^{\mathsf{A}}(X)$. And this idea can be generalized to larger n-types as well. The only thing to consider is that the order of the variables is important: $x < y$ and $y < x$ might each be realized, but not with the same assignment. This idea will eventually lead to the idea of *indiscernibles*. Also, for the sake of space, $\mathrm{tp}^{\mathsf{A}}(a/\emptyset)$ will often just be written $\mathrm{tp}^{\mathsf{A}}(a)$.

§ 3.A. Realizing types

Contained in the idea that a type is consistent is that it can be realized. In fact, it can be realized in an elementary extension. This shows that while a structure might omit a type, even its elementary diagram isn't enough to prevent the type from being realized, again highlighting some of the deficiencies of first order logic.

3.A • 1. Theorem (Realizing Types Theorem)

Let $\mathbf{A}$ be an $\mathcal{L}$-model with $X \subseteq A$. Let $\Sigma(\vec{x})$ be an n-type, $n < \omega$, of $\mathbf{A}$ over X.
Therefore there is an elementary extension $\mathbf{A} \preccurlyeq \mathbf{B}$ such that $\mathbf{B}$ realizes $\Sigma(\vec{x})$ and $|B| = |A|$.

Proof .:.

The result clearly holds of finite models, so assume $\mathbf{A}$ is infinite. Take $T = \mathrm{Eldiag}(\mathbf{A}) \cup \Sigma(\vec{x})$. To be more formal, we could instead replace $\vec{x}$ by some fresh constants $\vec{c}$. To show that T is satisfiable, take a finite subset $\Delta \subseteq T$, and a model $\mathbf{M}$ of $\mathrm{Th}_X(\mathbf{A}) \cup \Sigma(\vec{x})$. The finitely many formulas of $\Sigma(\vec{x}) \cap \Delta$ are realized in $\mathbf{M}$. The sentences of $\mathrm{Eldiag}(\mathbf{A}) \cap \Delta$ might have parameters from $A \setminus X$. As a conjunction $\varphi(\vec{a})$ with parameters in A, we can quantify them out with existential quantifiers: $\exists \vec{x}\varphi(\vec{x})$. The resulting sentence is in $\mathrm{Th}_X(\mathbf{A})$ and hence has witnesses in $\mathbf{M}$. So if we expand $\mathbf{M}$ to an $\mathcal{L}_A$ model, we can get that $\mathbf{M} \vDash \Delta$. Hence the theory is consistent, and so by Subsection 1.B, we get a model of size $|\mathcal{L}_A| + \aleph_0 = |A|$. ⊣

One useful property of this is that we don't expand the types over a subset when passing to an elementary extension, because the theories over the subset are the same: $\mathrm{Th}_X(\mathbf{A}) = \mathrm{Th}_X(\mathbf{B})$. So for all $\mathbf{A} \preccurlyeq \mathbf{B}$ with $n < \omega$ and $X \subseteq A \subseteq B$, $S_n^{\mathbf{A}}(X) = S_n^{\mathbf{B}}(X)$. As a result, the complete types of $\mathbf{A}$ are precisely the types given by elements in larger models: $\mathrm{tp}^{\mathbf{B}}(\vec{b}/X)$ for $\mathbf{B} \succcurlyeq \mathbf{A}$.

Another way to think about types is through automorphisms: having the same type is equivalent—modulo an elementary extension—to having an automorphism moving one to the other. To establish this, we first have some lemmas about partial elementary maps.

3.A • 2. Definition

Let $\mathbf{A}$ and $\mathbf{B}$ be $\mathcal{L}$-models. Let $f : A \rightharpoonup B$ be a map. f is a <u>partial elementary</u> map iff for all $\vec{a} \in \mathrm{dom}(f)^{<\omega}$ and $\mathcal{L}$-formulas $\varphi(\vec{x})$, $\mathbf{A} \vDash \varphi(\vec{a})$ iff $\mathbf{B} \vDash \varphi(f(\vec{a}))$. In other words, $\mathrm{Th}_{\mathrm{dom}(f)}(\mathbf{A}) = \mathrm{Th}_{\mathrm{im}(f)}(\mathbf{B})$.

Contained in this is the idea that f is injective and a *partial embedding*, since we can just take $\varphi(\vec{x})$ to be $\vec{x} \neq \vec{y}$, $R(\vec{x})$, or $F(\vec{x}) = y$. If we consider sentences instead of formulas in general, we can see that we're assuming $\mathbf{A} \equiv \mathbf{B}$. The concept of a partial elementary embedding will be useful, as they provide part of a back and forth argument when elements have the same type. The idea is then that we extend these partial elementary maps to automorphisms of the resulting structures, using Realizing Types Theorem (3.A • 1).

3.A • 3. Lemma

Let $\mathbf{A}$ and $\mathbf{B}$ be $\mathcal{L}$-models. Let $f : A \rightharpoonup B$ be a partial elementary map. Therefore, there is an elementary extension $\mathbf{B}' \succcurlyeq \mathbf{B}$ and elementary embedding $f \subseteq f' : A \to B'$.

Proof .:.

First we show that f can be extended point by point.

Claim 1

For any $\mathbf{A}$, $\mathbf{B}$, f as above and $e \in A$, there is an $\mathbf{B}_+ \succcurlyeq \mathbf{B}$ with partial elementary map $f \subseteq f_+ : A \rightharpoonup B_+$ where $e \in \mathrm{dom}(f_+)$.

Proof .:.

Consider the 1-type over $\mathbf{B}$

$$\Sigma(x) = \mathrm{Eldiag}(\mathbf{B}) \cup \{\varphi(x, f(\vec{a})) : \vec{a} \in \mathrm{dom}(f)^{<\omega} \wedge \mathbf{A} \vDash \varphi(e, \vec{a})\}.$$

The consistency of this is given just by compactness, and that $\mathbf{A} \equiv \mathbf{B}$. Since $\mathbf{A} \vDash \exists x\varphi(x, \vec{a})$ for all $\varphi(x, \vec{a}) \in \Sigma(x)$, it follows that $\mathbf{B} \vDash \exists x\varphi(x, f(\vec{a}))$. Thus by Realizing Types Theorem (3.A • 1), there's an elementary extension $\mathbf{B}_+ \succcurlyeq \mathbf{B}$ realizing $\Sigma(x)$ by some $b \in B_+$. Moreover, setting $f_+ = f \cup \{\langle e, b\rangle\}$ yields the desired partial elementary map. ⊣

Enumerate $A = \{a_\alpha : \alpha < \kappa\}$. We will inductively build an elementary chain $\langle \mathsf{B}_\alpha : \alpha < \kappa\rangle$ and partial elementary maps $\langle f_\alpha : A \rightharpoonup B : \alpha < \kappa\rangle$ where $\{a_\alpha : \alpha < \beta\} \subseteq \mathrm{dom}(f_\beta)$.

To do this, just recursively use Claim 1: for 0, set $\mathsf{B}_0 = \mathsf{B}$, and $f_0 = f$. At successor stages, let $\mathsf{B}_{\alpha+1}$ be the elementary extension from the claim applied to B_α, $e = a_\alpha$, and f_α. Let $f_{\alpha+1}$ be the f_+ from the claim. At limit stages take unions.

In the end, taking $\mathsf{B}' = \bigcup_{\alpha<\kappa} \mathsf{B}_\alpha$ and $f' = \bigcup_{\alpha<\kappa} f_\alpha$ yields an elementary embedding from all of A to B'. By The Tarski–Vaught Theorem (1.B • 5), $\mathsf{B} \preccurlyeq \mathsf{B}'$. ⊣

3.A • 4. Theorem

Let A be an $\mathcal{L}$-model with $X \subseteq A$. Let $a, e \in A^n \setminus X^n$ be such that $\mathrm{tp}^{\mathsf{A}}(a/X) = \mathrm{tp}^{\mathsf{A}}(e/X)$.
Therefore there is an automorphism f of an elementary extension of A such that $f \restriction X = \mathrm{id}$, and $f(a) = e$.

Proof .:.

Proceed by a back and forth argument, elementarily extending our model each time as with Lemma 3.A • 3. Start with the partial elementary embedding $f_0 : X \cup \{a\} \to X \cup \{e\}$ where $f_0 \restriction X = \mathrm{id}$, and $f_0(a) = e$. This is a partial elementary embedding, since a and e share the same type over X.

We can then take elementary extensions $\mathsf{A}_n \succcurlyeq \mathsf{A}$ through extensions of f_n so that $A_n \subseteq \mathrm{dom}(f_{n+1})$. When we then take the union, we get $\mathsf{A} \preccurlyeq \mathsf{A}_\omega$, and an elementary embedding $f = f_\omega : A_\omega \to A_\omega$. ⊣

In particular, for sufficiently saturated models A, elements with the same type over an $X \subseteq A$ of size $|X| < |A|$ can be moved to another by an automorphism. Such models in general are called *homogeneous*, and will be investigated in a later section.

§ 3.B. Omitting types

We've seen that we can realize types just by using Compactness (1.A • 2). It can often be useful to get more information by getting models which omit types. A major result about this is the omitting types theorem, showing that theories—in countable languages—are unable to ensure realization of types that aren't outright proven to be realized. The proof of the theorem itself in essence is a more careful proof of Compactness (1.A • 2).

The countability of the language is necessary, since we could easily just have uncountably many constant symbols, and take the countable type $\Sigma(x) = \{x \neq c_n : n < \omega\}$. So any model of the theory $\{c_\alpha \neq c_\beta : \alpha \neq \beta\}$ will realize this type, since it will necessarily be uncountable.

To start the proof, however, we require some preliminaries about *isolated* types. In particular, it's possible for a complete type to just be the result of a single formula: $\Sigma(\vec{x}) = \{\psi(\vec{x}) :\vDash \varphi(\vec{x}) \to \psi(\vec{x})\}$. In this case, we can't rule out omitting the type, since it's realized iff a single formula, $\exists\vec{x}\varphi(\vec{x})$, holds.

3.B • 1. Definition

Let T be an $\mathcal{L}$-theory. Let $\Sigma(x)$ be a set of $\mathcal{L}$-formulas.
$\Sigma(x)$ is a <u>type</u> of T iff $\Sigma(x) \cup T$ is consistent.
$\Sigma(x)$ is <u>isolated</u> iff there is some formula $\varphi(x)$ where $\Sigma(x) = \{\psi(x) : T \vDash \varphi(x) \to \psi(x)\}$.

And of course these notions can be generalized to allow for more variables.

3.B • 2. Theorem (Omitting Types Theorem)

Let $\mathcal{L}$ be a countable language. Let T be an $\mathcal{L}$-theory. Let $\Sigma(x)$ be a non-isolated set of $\mathcal{L}$-formulas consistent with T. Therefore there is a countable $\mathcal{L}$-structure $\mathsf{A} \vDash T$ omitting $\Sigma(x)$.

Proof .:.

Add countably many constant symbols $C \cap \mathcal{L} = \emptyset$. In the still countable language $\mathcal{L}' = \mathcal{L} \cup C$, enumerate the $\mathcal{L}'$-sentences $\{\varphi_n : n \in \omega\}$. We will construct a chain of $\mathcal{L}'$-theories $T_n \subseteq T_{n+1}$ with $T_0 = T$.

Assume $T_n = T \cup \{\psi_1, \cdots, \psi_m\}$ has already been constructed. Let ψ be the conjunction of the ψ_is. Replace the constants in C of ψ with fresh variables, ending up with ψ as $\psi'(\vec{c})$ for $\psi'(\vec{x})$ an $\mathcal{L}$-formula. Now let $\psi^n(x_n)$ be $\exists \vec{y}, \vec{z}\, \psi'(\vec{y}, x_n, \vec{z})$. Note that ψ^n is consistent with T. As $\Sigma(x)$ is non-isolated, there must be some $\theta(x) \in \Sigma(x)$ where $\psi^n(x_n) \wedge \neg\theta(x_n)$ is consistent with T. In this case, put $\psi^n(c_n) \wedge \neg\theta(c_n)$ into T_{n+1}.

Now put either φ_n or $\neg\varphi_n$ into T_{n+1} depending on which is consistent with $T_n \cup \{\neg\theta(c_n)\}$. If the formula we put in is $\exists x \sigma(x)$, then put $\sigma(c)$ for some fresh constant symbol c.

The resulting theory $T_\omega = \bigcup_{n\in\omega} T_n$ is consistent and complete as an $\mathcal{L}'$-theory. So let $\mathsf{A} \vDash T_\omega$ be a $\mathcal{L}'$-model. Take $\mathsf{M} \subseteq \mathsf{A}$ to be the submodel generated by the constant symbols of $\mathcal{L}$. Note that the interpretations C^{A}—along with the constants of $\mathcal{L}$—are closed under the operations of $\mathcal{L}$, since $\exists x (f(\vec{c}) = x) \in T_n$ for some n requires us to put a witness $f(\vec{c}) = c' \in T_\omega$. Hence $M = C^{\mathsf{A}}$.

Moreover, $\mathsf{M} \vDash T_\omega$ by The Tarksi–Vaught Test (1.B • 1). Thus M omits $\Sigma(x)$, since each constant symbol has $T \vDash \neg\theta(c)$ for some $\theta(x) \in \Sigma(x)$, and any witness to $\Sigma(x)$ will be a constant symbol by construction of M and the argument above. Thus the $\mathcal{L}$-reduct of M is a countable $\mathcal{L}$-model omitting $\Sigma(x)$. $\dashv$

There are various extensions of the omitting types theorem, as seen in [1], but they will not be so important for us. Another thing to note is that the same sort of procedure allows us to omit any countable collection of non-isolated types using $|\omega \cdot \omega| = \aleph_0$.

Section 4. Prime and Saturated Models

The use of omitting types is usually to get certain special kinds of models, but it is usually seen as a kind of weakness of the model, since something which is "possible" doesn't "happen". This is something which makes saturated models nice to work with, since everything which is "possible" does "happen" in the sense that every type consistent with the theory—over a subset of smaller than the whole model—is realized in a saturated model.

Saturated models are then seen as quite large, an intuition backed-up by the result that models of inaccessible cardinality are always saturated. Another sense that makes them large is their *universality* in that all smaller models of the same theory are elementarily embedded in it. But large models are not the only kinds of models of interest. In particular, a kind of dual notion of saturation is being atomic, or prime, a notion of being very small. Both of these notions build heavily on the previous section about realizing and omitting types.

Throughout this section, we will be dealing with countable languages, and complete theories with infinite models[i]. This is both for the ease of arguments, but also because of the requirements of Omitting Types Theorem (3.B • 2).

§ 4.A. Prime and atomic models

First we will introduce the actual definition of a prime model. This is motivated by the notion perhaps already familiar from algebra of a *prime field* like $\mathbb{Q}$, or $\mathbb{F}_3$, which has no proper subfield, and which embeds into every other field of its characteristic.

4.A • 1. Definition

Let T be a theory. An $\mathcal{L}$-model A is a prime model of T—a T-prime model—iff A elementarily embeds into every model of T.

For a slightly more complicated example extending the above, the algebraic closure of $\mathbb{Q}$ as a field, $\overline{\mathbf{Q}}$, is then a prime model of ACF_0. And similarly, the algebraic closure $\overline{\mathbb{F}}_3 \vDash \mathsf{ACF}_3$ is a prime model.

For a more logic-motivated example, we can consider the standard model of natural numbers, $\mathbf{N}$ in the language of PA. As a model of PA, this embeds into every other model. As a model of $\mathrm{Th}(\mathbf{N})$, it elementarily embeds in every model[ii]. As a result, $\mathbf{N}$ is a prime model of $\mathrm{Th}(\mathbf{N})$.

Related to prime models are those of *atomic* models, motivated by isolated types.

4.A • 2. Definition

Let T be a theory. An $\mathcal{L}$-model A is an atomic model of T—a T-atomic model—iff $\mathrm{tp}^{\mathsf{A}}(\vec{a})$ is isolated for all $\vec{a} \in A^{<\omega}$.

As a result, all the information about each element (or sequence of elements) can be understood as the result of a single formula and the theory T. Of course, different elements may have different such formulas, but the idea is that each element can be characterized by such a formula, reducing the infinite amount of information about the element by the finite amount of information given by a formula.

By a not-so-difficult argument, we can see that prime models over countable languages are always atomic, although the converse only holds under restricted conditions. In particular, for countable models, the two properties are equivalent.

[i] And thus the theories necessarily have no finite models: each $\exists \vec{x} \forall y (y$ is one of the $\vec{x})$ or its negation is in the complete theory. Since the theory has infinite models, none of them can be in the theory, and thus all of their negations must be in there.

[ii] Note that every element of $\mathbb{N}$ can be written out explicitly as a sequence of "$+1$"s, so we don't need the expand the model to the language $\mathcal{L} \cup \mathbb{N}$.

4.A • 3. Lemma

Let $\mathcal{L}$ be a countable language. Let T be a complete theory with infinite models. Let $\mathsf{A} \vDash T$ be T-prime. Therefore A is T-atomic.

Proof .:.

Let $\vec{a} \in A^{<\omega}$ be arbitrary. If $\mathrm{tp}^{\mathsf{A}}(\vec{a})$ were non-isolated, then by Omitting Types Theorem (3.B • 2), we would get an elementary embedding into a B which omits $\mathrm{tp}^{\mathsf{A}}(\vec{a})$. For $j : \mathsf{A} \to \mathsf{B}$ such an elementary embedding, we necessarily have that $\mathrm{tp}^{\mathsf{A}}(\vec{a}) = \mathrm{tp}^{\mathsf{B}}(j(\vec{a}))$, meaning $\mathrm{tp}^{\mathsf{A}}(\vec{a})$ is realized in B by $j(\vec{a})$, contradicting that B omits this type. ⊣

Extending this lemma is the equivalence for countable models stated earlier. Note that the requirement for countability is required here, since any prime model is necessarily countable. There are examples, however of atomic models which are not prime.

4.A • 4. Theorem

Let $\mathcal{L}$ be a countable language. Let T be a complete theory with infinite models. Let $\mathsf{A} \vDash T$ be a countable model. Therefore A is T-prime iff it is T-atomic.

Proof .:.

Lemma 4.A • 3 shows that if A is prime, then it is atomic, regardless of the cardinality of A. To show the other direction, we need to make use of the countability of A. In particular, we will proceed by building an elementary embedding from A into any $\mathsf{B} \vDash T$ assuming that A is atomic.

So assume A is atomic, and $\mathsf{B} \vDash T$. Enumerate $A = \{a_n : n \in \omega\}$. For each $n < \omega$, let $\varphi_n(x_0, \cdots, x_n)$ witness that the type of the first $n+1$ elements, $\mathrm{tp}^{\mathsf{A}}(a_0, \cdots, a_n)$, is isolated. We will construct partial elementary embeddings $f_n : A \rightharpoonup B$ where $\{a_i : i < n\} \subseteq \mathrm{dom}(f_n)$. Note that $f_0 = \emptyset$ is elementary since $\mathsf{A} \equiv \mathsf{B}$ as T is complete.

Suppose f_n has been given. Set $b_i = f_n(a_i)$ for $i < n$. By elementarity, $\mathsf{B} \vDash \exists x\, \varphi_n(b_0, \cdots, b_{n-1}, x)$. So there is some witness $b_n \in B$ with $\mathsf{B} \vDash \varphi_n(b_0, \cdots, b_n)$. But as φ_n isolates $\mathrm{tp}^{\mathsf{A}}(a_0, \cdots, a_n)$, the corresponding type $\mathrm{tp}^{\mathsf{B}}(b_0, \cdots, b_n)$ must be the same set of formulas. Hence $f_{n+1} = f_n \cup \{\langle a_n, b_n \rangle\}$ is still partial elementary.

Taking $f = \bigcup_{n \in \omega} f_n$ then yields an elementary embedding from all of A to B. ⊣

As a result, the existence of prime models for a theory is equivalent to the existence of atomic ones for it. Furthermore, the uniqueness of prime models then tells us that countable atomic models are also equivalent.

4.A • 5. Theorem (Uniqueness of Prime Models)

Let $\mathcal{L}$ be a countable language. Let T be a complete theory with infinite models. Let $\mathsf{A}, \mathsf{B} \vDash T$ be countable, T-atomic models. Therefore $\mathsf{A} \cong \mathsf{B}$.

Proof .:.

Clearly each is elementarily embedded in the other by definition. We will then build an isomorphism by a back and forth argument in the same way as with Theorem 4.A • 4: order A and B of type ω, and proceed as follows. Note that although we will make gaps in how we define the isomorphism, we will fill in the gaps as we progress.

Let $f_0 = \emptyset$, a partial elementary embedding. Let a_0 be the first element of A, and let φ_0 be the isolating formula of $\mathrm{tp}^{\mathsf{A}}(a_0)$. As $\mathsf{A} \vDash \exists x \varphi_0(x)$, let b_0 be an arbitrary element of B satisfying $\varphi_0(x_0)$. Let $f_1 = f_0 \cup \{\langle a_0, b_0 \rangle\}$.

Now let b_1 be the first element of B which isn't b_0. Now we can consider the 2-type $\mathrm{tp}^{\mathsf{B}}(b_0, b_1)$ which is isolated by $\varphi_1(x_0, x_1)$. By the same sort of argument as before, we can get an $a_1 \in A$ which witnesses $\varphi_1(a_0, x_1)$, setting $f_2 = f_1 \cup \{\langle a_1, b_1 \rangle\}$.

Then we can pick the least element a_2 of $A \setminus \{a_0, a_1\}$, and continue back and forth, getting two sequences $\langle a_n : n \in \omega \rangle$ and $\langle b_n : n \in \omega \rangle$ which cover A and B with f_n sending $a_i \mapsto b_i$ for $i < n$. It's not difficult to see that each f_n will be a partial elementary embedding, and that $\bigcup_{n\in\omega} f_n$ will be a full elementary embedding from A to B which is both injective and surjective, and thus an isomorphism. ⊣

Note further that although prime models are stated as embedding in *all* models of the theory, it suffices to consider just countable ones, which means that $\aleph_0$-categorical theories necessarily have atomic and prime models.

4.A • 6. Result

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Let $\mathsf{A} \models T$ elementarily embed into every countable model of T. Therefore A is T-prime.

Proof .:.

For every $\vec{a} \in A^{<\omega}$, $\mathrm{tp}^{\mathsf{A}}(\vec{a})$ is then realized in every countable model of T. By the (contrapositive of the) omitting types theorem, $\mathrm{tp}^{\mathsf{A}}(\vec{a})$ cannot be non-isolated, and thus must be isolated. As a result, A is countable and T-atomic, whence by Theorem 4.A • 4, T-prime. ⊣

§ 4.B. Stone spaces and atomic models

We can put a topology on the set of complete types of a model $S_n^{\mathsf{A}}(X)$ for each $X \subseteq A$ and $n \in \omega$. In particular, the basic open sets will be of the form $[\varphi] = \{p \in S_n^{\mathsf{A}}(X) : \varphi \in p\}$ for φ a formula. There are a lot of things to say about these spaces which will be unproven here. In particular, the space is always compact and totally disconnected. This also motivates some of the terminology used before: $\Sigma(x)$ is isolated iff $\Sigma(x) = [\varphi]$ for some φ. We can also do this relative to a theory T. In particular, $S_n(T)$ will be the set of complete sets of formulas consistent with T. The same topological definitons apply to make this a topology.

4.B • 1. Result

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Therefore $S_n(T)$ is compact for all $n \in \omega$.

Proof .:.

Let $C = \{[\varphi_i] : i \in I\}$ be a cover of $S_n(T)$. If C has no finite subcover, then $\Sigma = \{\neg\varphi_i : i \in I\}$ will be consistent with T. This is just a result of Compactness (1.A • 2): for each finite subset $\Delta \subseteq \Sigma$, there is a type $p \notin \bigcup_{\neg\varphi\in\Delta}[\varphi]$. But then any $\mathsf{B} \succcurlyeq \mathsf{A} \models T$, where p is realized in B, has $\mathsf{B} \models T$ and realize Δ. Hence Σ is finitely consistent with T and so Σ is consistent with T. But then C isn't a cover of $S_n(T)$. ⊣

For now, we want to think about what the stone spaces of atomic models look like to get some conditions and information about atomic models, allowing us to easily assert their existence or non-existence. In particular, we have the following theorem.

4.B • 2. Theorem

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Let $S_n(T)$ be the resulting stone space. Therefore T has an atomic model iff the isolated types of $S_n(T)$ are dense in $S_n(T)$ for all $n \in \omega$.

A more important use of this theorem is recognizing that a theory has a prime model. This will be dual to the same sort of fact with countably-saturated models.

4.B • 3. Corollary

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Suppose there are less than continuum many types in $S_n(T)$ for each $n \in \omega$. Therefore T has an atomic model.

Proof .:.

Otherwise, by Theorem 4.B • 2, there is some formula φ where $[\varphi]$ is disjoint from all isolated types. Hence there is a ψ where $[\varphi \wedge \neg\psi]$ and $[\varphi \wedge \psi]$ are also disjoint from all isolated types. Continuing in this way, we can build an injection from $2^{<\omega}$ into the types of $S_n(T)$, associating each such ψ with a point so that branches are sets of formulas. But then each branch in this $2^{<\omega}$ space is a subset of a type in $S_n(T)$. Moreover, each branch disagrees on some formula, meaning all of these types are distinct. Since there are $2^{\aleph_0}$ branches, this means there are continuum many types in $S_n(T)$, a contradiction. $\dashv$

This of course isn't necessary for T to have an atomic model, since PA has uncountably many types, but has the standard model of arithmetic, **N**, as a prime model. To see that PA has uncountably many types, we can just consider the set of primes $\{p_n : n \in \omega\}$, and for each $S \subseteq \{p_n : n \in \omega\} = P$, take the type

$$\Sigma_S(x) = \{\exists y(y \cdot p = x), \neg\exists y(y \cdot q = x) : p \in S \wedge q \notin S\},$$

where p, q are replaced by their syntactic $1 + \cdots + 1$ form. Each will be consistent with PA by compactness. Since each $\Sigma_S(x)$, $\Sigma_Z(x)$ disagree about whether x is divisible by $p \in S \triangle Z$, the complete types containing $\Sigma_S(X)$ and $\Sigma_Z(X)$ are distinct, and thus there are at least $|\mathcal{P}(\omega)| = 2^{\aleph_0}$ complete types of PA. In fact, there are exactly $2^{\aleph_0}$.

§ 4.C. Examples of prime, and atomic models

The start of the discussion about prime and atomic modelsr began with several examples which will not be reiterated here. Instead, we can consider several simple examples as well as non-examples.

Firstly, we noted above that being atomic does not imply being prime unless the model is countable. To see this, we have the following example of an uncountable, atomic, non-prime model.

4.C • 1. Example

Let $\mathcal{L} = \emptyset$, the language of pure identity. Let T be the theory of infinite sets. Let $\mathbf{A} \vDash T$. Thus $\mathbf{A}$ is T-atomic, and if $|A| > \aleph_0$, $\mathbf{A}$ is not prime.

Proof .:.

Clearly the type $\text{tp}^{\mathbf{A}}(\vec{a})$ for any $\vec{a} \in A^{<\omega}$ is the same, and is isolated just by the formula saying which entries in $\vec{a}$ are distinct and which are equal.

To see that $\mathbf{A}$ is not prime if $\mathbf{A}$ is uncountable, just note that $\langle\omega\rangle \vDash T$, but there can be no elementary embedding, indeed no injection, from A to ω. $\dashv$

Another easy example is $\mathbf{A} \vDash$ DLO in the language $\mathcal{L} = \{<\}$. Note that DLO is complete with infinite models. The type of $\vec{a} \in A^{<\omega}$ is isolated just by the order relations on the entries of $\vec{a}$.

Of course, both of these examples are somewhat trivial, but the point remains. Of course, if a theory is $\aleph_0$-categorical, then the countable model is prime and atomic.

§ 4.D. Saturated models

Whereas prime and atomic models realize relatively few types, saturated models will realize quite a lot. In fact, it will realize everything that it can. Saturated models of theories in countable languages aren't guaranteed to exist just under ZFC. There will always be if GCH holds, or if there is an inaccessible cardinal.

4.D • 1. Definition

Let κ be an infinite cardinal. Let $\mathbf{A}$ be an $\mathcal{L}$-model.
$\mathbf{A}$ is $\underline{\kappa\text{-saturated}}$ iff $\mathbf{A}$ realizes all elements of $S_n^{\mathbf{A}}(X)$ for $X \subseteq A$ with $|X| < \kappa$.
$\mathbf{A}$ is <u>saturated</u> iff $\mathbf{A}$ is $|A|$-saturated.

We of course can't strengthen the requirements to realizing all types over subsets of size $\leq |A|$, since $\{x \neq a : a \in A\}$ would be an (incomplete) type which can't be realized.

Note that in general, if **A** is saturated, then the subsets definable with parameters have size $|A|$ or else finite. This provides a nice test of saturation. To see this, let $S \subseteq A$ be infinite. Let S be definable from $\varphi(x)$ with parameters in $X \subseteq A$, a finite subset. Thus $\Sigma(x) = \{\varphi(x)\} \cup \{x \neq a : a \in S\}$ will be a type over $S \cup X$. But $\Sigma(x)$ can't be realized in **A** by definition of S so that by saturation, $S \cup X$ must have size $\kappa = |A|$.

Note that by the same sort of back and forth argument as in Uniqueness of Prime Models (4.A • 5), we get the uniqueness of elementarily equivalent, saturated models of a given cardinality.

4.D • 2. Theorem (Uniqueness of Saturated Models)

Let κ be an infinite cardinal. Let $\mathcal{L}$ be a countable language. Let $\mathbf{A} \equiv \mathbf{B}$ be saturated $\mathcal{L}$-models of size $|A| = |B| = \kappa$. Therefore $\mathbf{A} \cong \mathbf{B}$.

Proof .:.

Proceed by a back-and-forth argument. For the "forth" part, fixing an $a_\alpha \in A$ for $\alpha < \kappa$, let b_α be an element with the same type over $\{b_\beta : \beta < \alpha\}$ as a_α has over $\{a_\beta : \beta < \alpha\}$. Such elements exist since $\mathbf{A} \equiv \mathbf{B}$—meaning the two models have the same complete types—and **B** is saturated. For the "back" part, we do the same, switching the as and bs. If we ensure we're choosing the least (for some fixed well-orders) every time, we hit every element of A and B.

This process gives a sequence of partial elementary maps from A to B where $a_\beta \mapsto b_\beta$ for $\beta < \alpha < \kappa$. Denote these maps by f_α for $\alpha < \kappa$. Taking the union $f = \bigcup_{\alpha<\kappa} f_\alpha$ yields a full elementary embedding from A to B which is surjective by construction. Hence f is an isomorphism. $\dashv$

This is related to the fact that saturated models of theories are *universal*, a dual notion to being *prime*: every model (of at most the same size) elementarily embeds into it. Note that this means any saturated model properly contains infinitely many copies of itself.

4.D • 3. Theorem

Let κ be an infinite cardinal. Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models.
Let $\mathbf{A} \vDash T$ be a saturated $\mathcal{L}$-model with $|A| = \kappa$.
Let $\mathbf{B} \vDash T$ be an arbitrary $\mathcal{L}$-model of size $|B| \leq \kappa$.
Therefore **B** elementarily embeds into **A**. In other words, **A** is a κ^+-universal model of T.

Proof .:.

Similar to Uniqueness of Saturated Models (4.D • 2), enumerating B we build partial elementary maps $f_\alpha \subseteq f_\beta$ for $\alpha < \beta < |B|$ such that $\{b_\alpha \in B : \alpha < \beta\} \subseteq \mathrm{dom}(f_\beta)$ for each $\beta < \kappa$. To do this, start with $f_0 = \emptyset$, and at limit stages take unions: $f_\gamma = \bigcup_{\alpha<\gamma} f_\alpha$. At the successor stage $f_{\alpha+1}$, by saturation, we get a witness a_α to the type of b_α over what's been defined thus far: a_α realizes $\mathrm{tp}^{\mathbf{B}}(b_\alpha / \operatorname{im} f_\alpha)$. Setting $f_{\alpha+1} = f_\alpha \cup \{\langle b_\alpha, a_\alpha\rangle\}$ works so that $\bigcup_{\alpha<|B|} f_\alpha$ is a full elementary embedding from B to A. $\dashv$

For now, we will turn our attention to countably saturated models—i.e. $\aleph_0$-saturated models of size $\aleph_0$. Note that by Uniqueness of Saturated Models (4.D • 2) and Theorem 4.D • 3, such models are unique, and countably universal. This will allow us to dip our feet into ideas about the number of models a theory has up to isomorphism, although this isn't explored too deeply here.

Firstly, we look at when exactly countably saturated models exist.

4.D • 4. Theorem

Let $\mathcal{L}$ be a countable language. Let T be a complete theory with infinite models. Therefore T has a countably saturated model iff there are only countably many complete types of T.

Proof .:.

Suppose that T has a countably saturated model A. Therefore, the complete types of T are just the types of finite sequences of elements from A. Since $|A| = \aleph_0$, there are only $\aleph_0^{<\omega} = \aleph_0$ such complete types, and thus only countably many complete types of T.

If there are only countably many complete types of T, then we can just continually expand a countable model of T by witnesses via Realizing Types Theorem (3.A • 1). A little more explicitly, proceed as follows. Add countably many fresh constants $C = \{c_n : n \in \omega\}$. For each finite subset $\Delta \subseteq C$, there are still only countably many complete types of T as an $\mathcal{L}_\Delta = \mathcal{L} \cup \Delta$-theory. Since there are only countably many finite subsets of C, there are only countably many such types. Thus when we build a countable model realizing all these types with universe C, it will be countably saturated. ⊣

This also gives a nice corollary about the number of non-isomorphic models of complete theories. In particular, if a theory doesn't have a lot of countable models, then it has a saturated one.

4.D • 5. Corollary

Let $\mathcal{L}$ be a countable language. Let T be a complete theory with infinite models. Suppose T has only countably many countable models up to isomorphism. Therefore T has a countably saturated model.

Proof .:.

Each countable model realizes only countably many types. But each complete type of T is realized in some countable model. Thus there are only $\aleph_0 \cdot \aleph_0 = \aleph_0$ such types, so by Theorem 4.D • 4, there is a countably saturated model. ⊣

If we strengthen the hypotheses of Theorem 4.D • 4 considerably, then we can get stronger results. In particular, if there are only finitely many n-types for each $n \in \omega$, then T not only has a countably saturated model, but it only has *one* countable model: T is $\aleph_0$-categorical. This will be a major idea for thinking about the number of countable models a theory can have. Most surprisingly, the answer is never 2, although it can be any other cardinality $\leq 2^{\aleph_0}$.

4.D • 6. Theorem

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory. Therefore, up to isomorphism, T does not have exactly two countable models.

Proof .:.

Clearly if T has finite models, then T has exactly one model up to isomorphism. So assume that T has countably infinite models. If T has exactly two models, then by Corollary 4.D • 5, T has a saturated model.

Claim 1

If T has a saturated model, then T has a prime model.

Proof .:.

If T has a saturated model, then there are only countably many types of T. By Corollary 4.B • 3, T has an atomic model. If this model is countable, we're done since Theorem 4.A • 4 tells us that countable atomic models are prime. Otherwise, since there are only countably many types, we can choose a suitable submodel which will be an atomic model, but now of size $\aleph_0$ so that by Theorem 4.A • 4, we have a prime model of T. ⊣

Note that the prime model of T can't also be saturated without T being $\aleph_0$-categorical by Theorem 4.D • 3 and Theorem 4.A • 4: If $\mathbf{A} \vDash T$ is prime and saturated, then the only types of T are isolated, which means any countable model of T must be prime. Hence we get a contradiction: there are two non-isomorphic models of T which, by Uniqueness of Prime Models (4.A • 5), are isomorphic.

So let $\mathbf{A} \vDash T$ be atomic, and $\mathbf{S} \vDash T$ be saturated. Since $\mathbf{S}$ is not atomic, there are some $\vec{s} \in S^{<\omega}$ with $\text{tp}^{\mathbf{S}}(\vec{s})$ non-isolated. Note that $\mathbf{S}_{\vec{s}}$ is still countably saturated since $\mathbf{S}$ is. Thus the $\mathcal{L}' = \mathcal{L}_{\vec{s}}$ theory $T' = \text{Th}_{\vec{s}}(\mathbf{S})$ has a prime model $\mathbf{M}' \vDash T'$ by the reasoning in Claim 1.

The $\mathcal{L}$-reduct $\mathbf{M}$ is not prime, since it's not atomic: $\vec{s}^{\mathbf{M}'}$ realizes a non-isolated type. $\mathbf{M}$ is not saturated, since $\mathbf{M}'$ is not saturated as an $\mathcal{L}'$-model. This is more easily seen by considering the contrapositive: if $\mathbf{M}$ is saturated, then $\mathbf{M}'$ is. This is clear since if $X \subseteq M'$, then any complete type in $S^{\mathbf{M}'}(X)$ can be regarded as a subset of some complete type of $S^{\mathbf{M}}(X \cup \vec{s})$, which is realized in $\mathbf{M}$. Hence $\mathbf{M}$ is a non-atomic, non-saturated, countable model of T, meaning T has at least three non-isomorphic models: $\mathbf{A}$, $\mathbf{S}$, and $\mathbf{M}$. ⊣

This subsection will end with a result about the existence of saturated models. First we prove two necessary lemmas, both are relatively easy results about elementary chains.

4.D • 7. Lemma

Let $\mathcal{L}$ be a countable language. Let $\mathbf{A}$ be an $\mathcal{L}$-model of size $|A| \geq \aleph_0$. Let $\kappa \geq \aleph_0$ be a cardinal. Therefore there is a $\mathbf{B} \succcurlyeq \mathbf{A}$ with $|B| \leq |A|^\kappa$ such that B realizes all complete types of $\mathbf{A}$ over each $X \in [A]^{\leq\kappa}$.

Proof .:.

This is just using Realizing Types Theorem (3.A • 1) in addition to elementary chains. In particular, first note that there are at most $|A|^\kappa$ subsets of A of size $\leq \kappa$. Note further that for each $X \in [A]^{\leq\kappa}$, there are at most $|\mathcal{P}(\mathcal{L}_X \cup \aleph_0)| = 2^{|\mathcal{L}|+|X|+\aleph_0} \leq 2^\kappa$ complete types of $\mathbf{A}$ over X. Hence there are at most $|A|^\kappa$ complete types of $\mathbf{A}$ over subsets of size at most κ.

So we can successively realize the αth such type using Realizing Types Theorem (3.A • 1), and take limits at limit stages to get an elementary chain $\mathbf{A} = \mathbf{A}_0 \preccurlyeq \mathbf{A}_1 \preccurlyeq \cdots$ with limit $\mathbf{B} = \bigcup_{\alpha<|A|^\kappa} \mathbf{A}_\alpha$. This $\mathbf{B}$ has the desired properties. ⊣

4.D • 8. Lemma

Let $\mathcal{L}$ be a countable language. Let $\kappa \geq \aleph_0$ be a cardinal. Let $\mathbf{A}$ be an $\mathcal{L}$-model. Therefore there is a κ^+-saturated model $\mathbf{B} \succcurlyeq \mathbf{A}$ with $|B| \leq |A|^\kappa$.

Proof .:.

Successively using Lemma 4.D • 7, and that $\kappa^+ \cdot (\lambda^\kappa)^\kappa = \lambda^\kappa$ for $\lambda > 1$, we will realize in $\mathbf{A}_{\alpha+1}$ all types of $\mathbf{A}_\alpha$ over all subsets of A_α of size at most κ. So we can ensure $|A_{\alpha+1}| \leq |A|^\kappa$ for all $\alpha < \kappa^+ \leq |A|^\kappa$. The resulting $\mathbf{B} = \bigcup_{\alpha<\kappa^+} \mathbf{A}_\alpha$ will have size at most $\kappa^+ \cdot |A|^\kappa = |A|^\kappa$, and will then be κ^+ saturated by construction: κ^+ is regular, so that any subset of size κ will appear by some stage $\mathbf{A}_\alpha$, and thus be realized in $\mathbf{A}_{\alpha+1}$. ⊣

Using these two theories, we get the easy corollary of when a theory will have a saturated model.

4.D • 9. Theorem

Let $\mathcal{L}$ be a countable language. Let $\kappa > \aleph_0$ be a regular cardinal with $2^{<\kappa} = \kappa$. Let T be a complete theory with infinite models. Therefore T has a saturated model of size κ.

Proof .:.

Let $\mathbf{A} \vDash T$ be of size $|A| = \kappa$. If $\kappa = \lambda^+$ is a successor, then we are done by Lemma 4.D • 8. To see this, note that $\lambda^+ = \kappa = \kappa^{<\kappa}$ by the hypotheses on κ. There is then a λ^+-saturated model $\mathbf{B} \succcurlyeq \mathbf{A}$ of size $|B| \leq |A|^\lambda = \kappa^\lambda = \kappa$, which means $\mathbf{B}$ is saturated.

So assume κ is a limit, i.e. κ is inaccessible. We now just deal with the successor stages as before, and sup up to κ. To do this, build an elementary chain

$$\mathbf{A} \preccurlyeq \mathbf{A}_1 \preccurlyeq \cdots \preccurlyeq \bigcup_{\alpha<\kappa} \mathbf{A}_\alpha = \mathbf{B},$$

where we take unions at limit stages, and at successors, for $\mathbf{A}_\alpha$ already defined, we take $\mathbf{A}_{\alpha+1}$ to be an $\aleph_{\alpha+1}$-saturated model with $|A_{\alpha+1}| \leq |A_\alpha|^{\aleph_{\alpha+1}} \leq \kappa^{<\kappa} = \kappa$. As in Lemma 4.D • 8, this results in a saturated model of size $\aleph_\kappa = \kappa$. Since κ is regular, any subset of B of size $\aleph_\beta < \kappa$ will appear as a subset by some stage $\mathbf{A}_\alpha$, and thus be realized in $\mathbf{A}_{\alpha+\beta+1}$. ⊣

In contrast to Theorem 4.D • 9, note that some theories provably have no saturated models of given cardinalities. For example, number theory $\mathrm{Th}(\mathbf{N})$ has no countably saturated models, since it has $2^{\aleph_0}$ complete types and we can apply Theorem 4.D • 4. To see that $\mathrm{Th}(\mathbf{N})$ has continuum complete types, we can just consider the infinite set of prime numbers $P \subseteq \mathbb{N}$. For each subset $X \subseteq P$, we can consider the incomplete type

$$\Sigma_X(x) = \{p \mid x, q \nmid x : p \in X \wedge q \in P \setminus X\}.$$

Each $\Sigma_X(x)$ is finitely satisfiable, and so is contained in a complete type of $S_1(\mathrm{Th}(\mathbf{N}))$. Moreover, any two distinct subsets $X, Y \subseteq P$ will have $\Sigma_X(x)$ and $\Sigma_Y(x)$ disagree on some prime dividing x: $p \in X \triangle Y$ has $p \mid x, p \nmid x \in \Sigma_X(x) \triangle \Sigma_Y(x)$. So $\mathrm{Th}(\mathbf{N})$ has uncountably many complete types and so no countably saturated model. Of course, this doesn't rule out saturated models of other cardinalities, like Theorem 4.D • 9 tells us. This theorem does, however, rely on ideas unprovable from ZFC, namely the existence of inaccessibles, or else an instance of $\mathsf{CH}(\kappa)$ for some κ.

§ 4.E. Examples and applications of saturated models

It's clear that the κ-sized models of κ-categorical theories are saturated: if $\Sigma(x)$ is a type omitted in our model, we can take an elementary extension of the same cardinality in which it's realized by Realizing Types Theorem (3.A • 1): Take a skolem hull of the original model with the new witness of size κ, which then must be isomorphic to the original, contradicting that it omitted the type. Hence some easy examples of saturated models come from categorical theories: DLO is $\aleph_0$-categorical with model $\mathbf{Q} = \langle \mathbb{Q}, < \rangle$, which is then saturated. ACF_p is categorical in every uncountable cardinality for each characteristic $p \in \omega$, with $2^{\aleph_0}$-sized model $\mathbf{C} = \langle \mathbb{C}, 0, 1, +, -, \cdot \rangle$ for ACF_0.

A non-trivial example of a saturated model is the countable, random graph in the language with just the edge relation. Again, most of these come from theories categorical in some cardinality: the countable, random graph is $\aleph_0$-categorical. ACF_p is κ-categorical for all $\kappa > \aleph_0$. Divisible, torsion-free, abelian groups are categorical for all $\kappa > \aleph_0$. Obviously pure-identity theory is categorical for all cardinalities.

To investigate a cute example of saturation, we can consider $\aleph_1$-saturated models of DLO. Such models cannot have countable sequences with a limit.

4.E • 1. Example

Let $\mathbf{A} \vDash \mathsf{DLO}$. Let $\langle x_n \in A : n \in \omega \rangle$ be a countable, increasing sequence. Therefore $\lim_{n<\omega} x_n$ does not exist in A.

Proof .:.

Suppose not, and let $c \in A$ be the limit of the sequence. Consider the (partial) type $\Sigma(x) = \{x_n < x < c : n \in \omega\}$. Using only countably many parameters, $\Sigma(x)$ is realized by some $a \in A$. But this a is strictly between the limit c and the sequence $\langle x_n : n \in \omega \rangle$. Hence c is not the limit. ⊣

Section 5. Homogeneous Models

Many of the uniqueness results about prime and saturated models can be restated as results about *homogeneous* models. The terminology is motivated by the idea that two elements are basically the same if there is an automorphism that moves one to the other. For homogeneous models, so long as the two elements have the same type—that is, are at least as similar as they need to be—we have such an automorphism.

5 • 1. Definition

> Let **A** be an infinite $\mathcal{L}$-model. Let κ be an infinite cardinal.
> **A** is κ-homogeneous iff for every $X \subseteq A$ of size $|X| < \kappa$, $f : X \to A$ partial elementary, and $a \in A$, there is a partial elementary map f' where $f \subseteq f' : X \cup \{a\} \to A$.
> **A** is homogeneous iff **A** is $|A|$-homogeneous.

This kind of homogeneity can allow for some back and forth arguments in certain cases. In particular, it turns out that prime and saturated models will have various kinds of homogeneity. This will allow us to condense the proofs of Uniqueness of Prime Models (4.A • 5) and Uniqueness of Saturated Models (4.D • 2) in terms of homogeneity.

A useful property, which isn't used much here, is that for an elementary chain of homogeneous models, the union is also homogeneous.

§ 5.A. Tests and theory

As stated before, if two elements have the same type, then there is an automorphism moving one to the other. Of course, the converse holds in every model, so homogeneity provides the nice property of the two being equivalent.

5.A • 1. Result

> Let **A** be a homogeneous $\mathcal{L}$-model with $X \subseteq A$ of size $|X| < |A|$. Let $f : X \to A$ be a partial elementary map. Therefore there is an automorphism $\pi : A \to A$ such that $f \subseteq \pi$.

Proof .:.

Enumerate $A = \{a_\alpha : \alpha < \kappa\}$ for $\kappa = |A|$. Now we build up to an automorphism π. Let $\pi_0 = f$, and at limit stages take unions. At odd successor stage $\pi_{\alpha+1}$, add a_α to the domain of π_α by Definition 5 • 1, and then add a_α to the image of π_α—i.e. adding it to the domain of the inverse. As a result, $\pi_{\alpha+1}$ has $a_\alpha \in \mathrm{dom}(\pi_{\alpha+1}) \cap \mathrm{im}(\pi_{\alpha+1})$ and so taking the union results in $\pi = \bigcup_{\alpha<\kappa} \pi_\alpha$ as a full automorphism. ⊣

So an immediate corollary is when $X = \{a, b\}$ and $f = \{\langle a, b\rangle\}$. Such an f is partial elementary if $\mathrm{tp}^{\mathsf{A}}(a) = \mathrm{tp}^{\mathsf{B}}(b)$, and thus if two elements have the same type, there's an automorphism that moves one to the other. Of course, it might not be that the two are exchanged, e.g. if there is an order.

The major result of this subsection is that homogeneous models of a given cardinality are determined by the types they realize. This also has the nice property of showing how many homogeneous models of a given cardinality there are. First we have an (unproven) lemma allowing us to do one half of a back and forth argument.

5.A • 2. Lemma

> Let $\mathcal{L}$ be a countable language. Let κ be an infinite cardinal.
> Let T be a complete $\mathcal{L}$-theory.
> Let **A** $\equiv$ **B** $\vDash T$ be κ-homogeneous $\mathcal{L}$-models. Suppose
> 1. $|B| \le \kappa$, and $X \subseteq A$ where $|X| \le \kappa$; and
> 2. every type of T realized in **B** is realized in **A**.
>
> Therefore there is a partial elementary map $f : X \to B$.

The assumption that every type realized in **B** is realized in **A** is important. It allows us to conclude that if they two homogeneous models realized the same types, then they are isomorphic.

5.A • 3. Theorem (Uniqueness of Homogeneous Models)

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory. Let $\mathbf{A} \equiv \mathbf{B} \vDash T$ be two homogeneous models realizing the same types in $S_n(T)$ for $n < \omega$, and with $|A| = |B|$. Therefore $\mathbf{A} \cong \mathbf{B}$.

Proof .:.

We will build an isomorphism $f : A \to B$ be a back-and-forth argument. Enumerate $A = \{a_\alpha : \alpha < \kappa\}$ and $B = \{b_\alpha : \alpha < \kappa\}$ for $|A| = |B| = \kappa$. First set $f_0 = \emptyset$, and at limit stages take unions.

Now suppose f_α has been defined with $\mathrm{dom}(f_\alpha) = X$. By Lemma 5.A • 2, there is a partial elementary $g : X \cup \{a_\alpha\} \to B$. Write $Y = \mathrm{im}(g) \subseteq B$. Taking $f_\alpha \circ g^{-1}$ then yields a partial elementary function from Y to B. As **B** is homogeneous, this can be extended to a partial elementary embedding to one $h : Y \cup \{x\} \to B$. Taking $h(x) = b$ yields a partial elementary map $f_\alpha \cup \{\langle a_\alpha, b\rangle\}$.

This provides one half of the back-and-forth argument, and the same argument can extend this to a partial elementary map with b_α in its range, yielding $f_{\alpha+1}$ with a_α in its domain, and b_α in its range. ⊣

There are several neat consequences of this, mostly in thinking about the number of (non-isomorphic) homogeneous models of a theory. It's obvious that for a cardinal κ, the number of models of size κ is 2^κ when the language is countable. We can get a better bound than this, however, for homogeneous models, and in the case of having a countable saturated model, the bound decreases substantially for all κ.

5.A • 4. Corollary

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory. Let κ be an infinite cardinal. Therefore the number of non-isomorphic, homogeneous models of T of size κ is at most $2^{\mathfrak{c}}$.

If T has a countably saturated model, then the number is at most $\mathfrak{c}$.

Proof .:.

By Uniqueness of Homogeneous Models (5.A • 3), homogeneous models of size κ are determined by the types realized. Because T is countable, there are at most $2^{\aleph_0} = \mathfrak{c}$ types. As a result, there are at most $2^{\mathfrak{c}}$ subsets of types, and so at most that many non-isomorphic, homogeneous models.

If there is a countably saturated model, then by Theorem 4.D • 4, there are only $\aleph_0$ many types, and thus $2^{\aleph_0}$ non-isomorphic, homogeneous models by the same idea as before. ⊣

Another useful result about homogeneous models concerns $\aleph_0$-homogeneity: we can extend to such models without an increase in cardinality.

5.A • 5. Result

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Let $\mathbf{A} \vDash T$. Therefore there is a $\mathbf{B} \succcurlyeq \mathbf{A}$ where **B** is $\aleph_0$-homogeneous, and $|B| = |A|$.

We will see that both prime models and saturated models are homogeneous, and so this provides us with some basic examples of homogeneous models. Obviously κ-categorical theories have their κ-sized models as homogeneous: they are saturated and thus homogeneous. In fact, saturation can be characterized in terms of universality and homogeneity. Before working with saturated models, we will work with prime ones.

We will reframe the results about prime models and saturated models as results about kinds of homogeneous structures, or at least κ-homogeneous of certain cardinals κ.

§ 5.B. Exploring prime models with homogeneity

The primary result for this subsection is that prime models are homogeneous. More generally, atomic models are $\aleph_0$-homogeneous.

5.B • 1. Lemma

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with $\mathsf{A} \vDash T$ an atomic $\mathcal{L}$-model. Therefore A is $\aleph_0$-homogeneous.

Proof .:.

Suppose $f : A \rightharpoonup A$ is partial elementary with $|f| < \omega$. Let $a \in A$ be arbitrary. Ordering the domain, view f as the map sending $d_i \mapsto r_i$ for $i < |\mathrm{dom}(f)| < \omega$. Note that the type $\mathrm{tp}^{\mathsf{A}}(\vec{d}, a)$ is isolated by some $\varphi(\vec{x}, y)$. As a partial elementary map,

$$\mathsf{A} \vDash \exists y \varphi(\vec{d}, y) \quad \text{iff} \quad \mathsf{A} \vDash \exists y \varphi(\vec{r}, y).$$

So if we set b to be a witness to $\varphi(\vec{r}, y)$ in A, then $f' = f \cup \{\langle a, b\rangle\}$ is partial elementary. $\dashv$

5.B • 2. Corollary

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with $\mathsf{A} \vDash T$ a prime $\mathcal{L}$-model.
Therefore A is homogeneous.

With the theory of homogeneous models, this gives an alternative way to view Uniqueness of Prime Models (4.A • 5): such models are unique by Uniqueness of Homogeneous Models (5.A • 3). This new proof is really just the same proof with some new machinary. If we examine the proof of Lemma 5.B • 1, the kind of back and forth argument that we can produce is precisely the same idea[iii] as in the proof of Uniqueness of Prime Models (4.A • 5).

5.B • 3. Corollary

Let $\mathcal{L}$ be a countable language. Let T be a complete theory with infinite models. Let $\mathsf{A}, \mathsf{B} \vDash T$ be countable, T-atomic models. Therefore $\mathsf{A} \cong \mathsf{B}$.

Proof .:.

A and B are prime, and thus homogeneous. Yet as the types realized in A and B are precisely the isolated types, we have two homogeneous models of T of the same cardinality that realize the same types, and thus are isomorphic by Uniqueness of Homogeneous Models (5.A • 3). $\dashv$

§ 5.C. Exploring saturation with homogeneity

We've noted in Subsection 4.D, in particular Theorem 4.D • 3, that saturated models are κ^+-universal for κ the size of the model, although we haven't formally introduced the term yet. A remarkable fact is that saturated models are precisely the universal, homogeneous models.

5.C • 1. Definition

Let T be a complete $\mathcal{L}$-theory, and $\mathsf{A} \vDash T$ an $\mathcal{L}$-model. Let κ be a cardinal.
Therefore A is <u>κ-universal</u> iff for all $\mathsf{B} \vDash T$ with $|B| < \kappa$, B elementarily embeds into A.

Combining Theorem 4.D • 3, we get that saturated models are homogeneous. In fact, homogeneity can be seen as a kind of weak form of saturation.

[iii] albeit applied only between one model and itself

5.C • 2. Lemma

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Let $\mathbf{A} \vDash T$ be a saturated $\mathcal{L}$-model. Therefore $\mathbf{A}$ is homogeneous and $|A|^+$-universal.

Proof .:.

Again, Theorem 4.D • 3 gives $|A|^+$-universality. To show homogeneity, let $X \subseteq A$ with $|X| < |A| = \kappa$. Suppose $f : X \to A$ is partial elementary. For $a \in A \setminus X$, just consider the type of a over X modified by f:

$$\Sigma(x) = \{\varphi(x, f(\vec{m})) : \vec{m} \in X^{<\omega} \wedge \mathbf{A} \vDash \varphi(a, \vec{m})\}.$$

Note that $\Sigma(x)$ is consistent with T, since f is partial elementary: if $\varphi(x, f(\vec{m})) \in \Sigma(x)$, then $\mathbf{A} \vDash \exists x\varphi(x, \vec{m})$ and applying f yields that $\mathbf{A} \vDash \exists x\varphi(x, f(\vec{m}))$. Using conjugation yields that $\Sigma(x)$ is finitely realizable in $\mathbf{A}$ and so consistent with T.

By κ-saturation, there is a witness $b \in A$ with $\mathbf{A} \vDash \Sigma(b)$. The map $f' = f \cup \{\langle a, b\rangle\}$ is then still partial elementary. Hence A is homogeneous. ⊣

This gives one direction of the following theorem which states that the two are equivalent.

5.C • 3. Theorem

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Let $\mathbf{A} \vDash T$ be an $\mathcal{L}$-model. Therefore $\mathbf{A}$ is saturated iff $\mathbf{A}$ is homogeneous and a $|A|^+$-universal model of T.

Proof .:.

Lemma 5.C • 2 gives one direction so it suffices to show that if $\mathbf{A}$ is homogeneous and $|A|^+$-universal then it is saturated.

So suppose $X \subseteq A$ with $|X| < |A| = \kappa$. Let $\Sigma(x) \in S_1^{\mathbf{A}}(X)$. Note that we can realize this type in some elementary extension: $\mathbf{B} \vDash \mathrm{Th}_X(\mathbf{A})$ and for $b \in B$, $\mathbf{B} \vDash \Sigma(b)$. Note that we're assuming here then that $X \subseteq B$. And in particular, we can ensure $|B| \le \kappa$. By universality, there's an elementary embedding $f : B \to A$. But then $f(b)$ must realize the same type as in $\mathbf{B}$. Explicitly, $f \restriction X : X \to A$ is a partial elementary map so by homogeneity, there is an $a \in A$ such that

$$\Sigma(x) = \mathrm{tp}^{\mathbf{B}}(b/X) = \mathrm{tp}^{\mathbf{A}}(f(b)/f"X) = \mathrm{tp}^{\mathbf{A}}(a/X).$$

Hence $\mathbf{A}$ is saturated. ⊣

In fact, we actually get a slightly stronger result for the "if" direction.

5.C • 4. Result

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Let $\mathbf{A} \vDash T$ be a homogeneous $\mathcal{L}$-model. Therefore if $\mathbf{A}$ realizes all complete types of T, then $\mathbf{A}$ is saturated.

Such knowledge gives an alternative proof of the uniqueness of saturated models through the uniqueness of homogeneous models.

5.C • 5. Corollary (Uniqueness of Saturated Models)

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Let $\mathbf{A} \vDash T$ and $\mathbf{B} \vDash T$ be two $\mathcal{L}$-models with $|A| = |B|$. Therefore $\mathbf{A} \cong \mathbf{B}$.

Proof .:.

To see this, by Uniqueness of Homogeneous Models (5.A • 3), it suffices to show that $\mathbf{A}$ and $\mathbf{B}$ realize the same types in $S_n(T)$ for $n < \omega$. Yet as saturated models, they both realize *all* such types, and so have the same types. Hence the two are isomorphic. ⊣

Just as in Subsection 4.D, let's turn our attention to countably saturated models. By the results above, it's clear that such models are $\aleph_0$-homogeneous, a property shared by prime models. The distinguishing property then is about what complete types of T are realized. In particular, while prime models realize just the isolated types, countably saturated ones, unsurprisingly, realize all of them.

Note that for countable models in general, realizing all types of T isn't sufficient for being countably saturated. So the homogeneity here is necessary.

5.C • 6. Result

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Let $\mathsf{A} \vDash T$ be an $\mathcal{L}$-model. Therefore A is $\aleph_0$-saturated iff A is $\aleph_0$-homogeneous and A realizes all types of T.

Proof .:.

The "only if" direction is clear from Lemma 5.C • 2. To show the "if" direction, suppose $\vec{a} \in A^{<\omega}$, and $\Sigma(\vec{x}) \in S_n^{\mathsf{A}}(\vec{a})$. Let $m = \mathrm{lh}(\vec{a})$, and consider the $n+m$-type replacing $\vec{a}$ with variables:

$$\sigma(\vec{x}, \vec{y}) = \{\varphi(\vec{x}, \vec{y}) : \varphi(\vec{x}, \vec{a}) \in \Sigma(\vec{x})\} \in S_{n+m}(T).$$

There are then $\vec{e}_1, \vec{a}_1 \in A^{<\omega}$ realizing this by assumption: $\mathsf{A} \vDash \sigma(\vec{e}_1, \vec{a}_1)$. Since it must follow that $\mathrm{tp}^{\mathsf{A}}(\vec{a}_1) = \mathrm{tp}^{\mathsf{A}}(\vec{a})$, by homogeneity, we can then find a $\vec{e} \in A^{<\omega}$ such that $\mathrm{tp}^{\mathsf{A}}(\vec{e}, \vec{a}) = \mathrm{tp}^{\mathsf{A}}(\vec{e}_1, \vec{a}_1)$. But then $\vec{e}$ realizes $\Sigma(x)$ in A. Therefore A is $\aleph_0$-saturated. $\dashv$

Section 6. Indiscernibles

Indiscernibles are useful for a variety of reasons, most notably for exactly what the name suggests: they are indistinguishable from the model's perspective. Before getting into the theory proper, we will need to introduce a variety of concepts. Of course, this includes what it means to be a set of indiscernibles, but it also includes a fair amount of infinitary combinatorics, and notably Ramsey theory.

6 • 1. Definition

Let $\mathcal{L}$ be a language, and $\mathbf{A}$ an $\mathcal{L}$-model. $X \subseteq A$ with an order $\leqslant$ is a set of indiscernibles iff for all increasing sequences $\vec{a}, \vec{b} \in X^n$, $n \in \omega$, and $\mathcal{L}$-formulas φ, $\mathbf{A} \vDash \varphi(\vec{a}) \leftrightarrow \varphi(\vec{b})$.

First we have some notation to introduce. Firstly, a partition of $[X]^\kappa$, for κ a cardinal and X any set, is just a function $f : [X]^\kappa \to \lambda$ for λ an ordinal. Ore often, we will call this a *coloring* of $[X]^\kappa$. A homogeneous subset of X is just any $Y \subseteq X$ where $f \restriction [Y]^\kappa$ is constant. We have the downright awkward notation $\kappa \to (\theta)^\mu_\lambda$ to represent the statement that for any $f : [\kappa]^\mu \to \lambda$, there is a homogeneous $Y \subseteq \kappa$ with $|Y| \geq \theta$. Clearly if $\kappa \to (\theta)^\mu_\lambda$ and $\eta \geq \kappa$, then $\eta \to (\theta)^\mu_\lambda$. Similarly, if $\kappa \to (\theta)^\mu_\lambda$, then the statement holds if we reduce any cardinal to the right of the arrow.

6 • 2. Theorem (Ramsey's Theorem)

For $m, n < \omega$, $\aleph_0 \to (\aleph_0)^n_m$. In other words, for any coloring $f : [\omega]^n \to m$, there is an infinite, homogeneous subset $X \subseteq \omega$: $1 = |f``[X]^n|$.

The use of these concepts will be for generating indiscernibles.

§ 6.A. Consistency and useful properties

The existence of elements totally indistinguishable in general isn't guaranteed for infinite models, but the existence of order-indiscernibles is. The proof of Löwenheim–Skolem (1.B • 2) used skolem functions without first introducing them. Their use is found throughout set theory and model theory, often in restricted usage, such as only considering Σ_1-skolem functions. They will play a useful role in showing the consistency of indiscernibles below.

6.A • 1. Theorem

Let T be a $\mathcal{L}$-theory with infinite models. Let $\langle X, \leqslant \rangle$ be an infinite linear order. Therefore, there is a model $\mathbf{M} \vDash T$ with $X \subseteq M$ a set of order-indiscernibles.

Proof .:.

Expand the language to $\mathcal{L}_X = \mathcal{L} \cup X$. Consider the $\mathcal{L}_X$-theory

$$T' = T \cup \{x \neq y : x, y \in X\}$$
$$\cup \{\varphi(\vec{x}) \leftrightarrow \varphi(\vec{y}) : \varphi \text{ is an } \mathcal{L}\text{-formula, and } \vec{x}, \vec{y} \in X^{<\omega} \text{ are } \leqslant\text{-increasing}\}.$$

If T' has a model, then we're done: take the $\mathcal{L}$ reduct. To get that T' is satisfiable, use compactness. Suppose that $\Delta \subseteq T'$ is finite. Consider the subset $Y \subseteq X$ of elements which occur in the formulas of Δ. Let Φ be the set of $\mathcal{L}$-formulas for which Δ asserts the indiscernibility of Y. Note that Y, Φ, and Δ are all finite. In particular, Φ contains only the $n + 1 < \omega$ free variables $v_0, \cdots, v_n$.

Let $\mathbf{M}$ be an infinite model of T. Without loss of generality, take $X \subseteq M$, and extend $\leqslant$ to all of M. Define the function $F : [M]^n \to \wp(\Phi)$. If $\{a_i : i \in n\} \in [M]^n$ is an increasing enumeration, define

$$F(A) = \{\varphi \in \Phi : \mathbf{M} \vDash \varphi(a_0, \cdots, a_n)\}.$$

Note that F colors n-subsets of the infinite M into at most $2^{|\Phi|} < \aleph_0$ colors. So by Ramsey's Theorem (6 • 2), there is a F-homogeneous set $Z \subseteq M$. But then $M \vDash \Delta$ in the expanded language interpreting Y as Z, and

$X \setminus Y$ as anything else. As Δ was arbitrary, by Compactness (1.A • 2), T' has a model. $\dashv$

The proof of this is really just a more careful proof of the upward Löwenheim–Skolem theorem: we add a bunch of elements, and extend our model to one where these are indiscernible in the sense of Definition 6 • 1. The obstruction to doing this in general is that certain theories might be incompatible with this, e.g. linear orders: we can't say $x < y$ and $y < x$ are equivalent for any x, y supposedly indistinguishable.

This is really the only obstruction, but it is still why we need to consider increasing sequences of elements rather than just in an arbitrary order. Although the order on X is arbitrary, that order prevents us from adding the statement $x < y \leftrightarrow y < x$ for $x, y \in X$ to our theory.

A very useful idea in working with indiscernibles is working with theories T that have built-in skolem functions, i.e. for every formula $\varphi(x, y)$, there is a function symbol f_φ such that $T \vdash \forall y\big(\exists x \varphi(x, y) \to \varphi(f_\varphi(y))\big)$. That we can always extend a theory to one with built-in skolem functions, and similarly expand a model interpreting these functions. To denote these expansions, we let T^{sk} and $\mathbf{A}^{\mathrm{sk}}$ denote the expansion of T and $\mathbf{A} \vDash T$ to those with built-in skolem functions.

Once we have built in skolem functions, it's useful to consider the skolem hull generated by a set of indiscernibles X: $\mathsf{Hull}_{\mathbf{A}}(X) \prec \mathbf{A}$. Such hulls have some nice, important properties. The proofs of these are relatively straightforward.

6.A • 2. Theorem

Let $\mathcal{L}$ be a countable language. Let T be an $\mathcal{L}$-theory with built-in skolem functions.
Let X, with order $\leqslant_X$, be a set of indiscernibles for an $\mathcal{L}$-model $\mathbf{A} \vDash T$.
Let Y, with order $\leqslant_Y$, be a set of indiscernibles for an $\mathcal{L}$-model $\mathbf{B} \vDash T$. Therefore

1. If $Z \subseteq X$, then $\mathsf{Hull}_{\mathbf{A}}(Z) \preccurlyeq \mathsf{Hull}_{\mathbf{A}}(X)$.
2. Any automorphism of $\langle X, \leqslant_X \rangle$ extends uniquely to an automorphism of $\mathsf{Hull}(X)$.
3. If $\mathrm{tp}^{\mathbf{B}}(Y) = \mathrm{tp}^{\mathbf{A}}(X)$ and as orders X embeds into Y, then $\mathsf{Hull}_{\mathbf{B}}(Y)$ elementarily embedds into $\mathsf{Hull}_{\mathbf{A}}(X)$.
4. If $\mathrm{tp}^{\mathbf{B}}(Y) = \mathrm{tp}^{\mathbf{A}}(X))$ and X, Y are infinite, then $\mathsf{Hull}_{\mathbf{A}}(X)$ and $\mathsf{Hull}_{\mathbf{B}}(Y)$ realize the same types of T.

Note that although tp is defined only for finite sequences of elements, it makes sense to say take tp of a set of indiscernibles, as just the union of the type of the finite sequences of length $n < \omega$.

6.A • 3. Corollary

Let $\mathcal{L}$ be a countable language. Let T be an $\mathcal{L}$-theory with infinite models. Let $X \subseteq A$ be an infinite set of indiscernibles for an $\mathcal{L}^{\mathrm{sk}}$-model $\mathbf{A} \vDash T^{\mathrm{sk}}$.
Suppose $\mathbf{A}$ omits a type $\Sigma(\vec{x}) \in S_n^{\mathbf{A}}(\emptyset)$. Therefore there are arbitrarily large models of T^{sk} omitting $\Sigma(\vec{x})$.

Proof .:.

Let $\kappa \geq \aleph_0$ be arbitrary. Just by compactness, we can find a model $\mathbf{B} \vDash T^{\mathrm{sk}}$ with indiscernibles $Y \subseteq B$ of cardinality $|Y| \geq \kappa$ and $\mathrm{tp}^{\mathbf{B}}(Y) = \mathrm{tp}^{\mathbf{A}}(X)$. Note that then $\mathsf{Hull}_{\mathbf{B}}(Y)$ has size at least κ, and this will omit $\Sigma(\vec{x})$ if $\mathbf{A}$ does. To see this, just note that any realization $\vec{w}$ can be represented by elements $\vec{v}$ of X through built in skolem functions f, but then for some $\vec{u} \in X^{<\omega}$,

$$\mathsf{Hull}_{\mathbf{B}}(Y) \vDash \varphi(\vec{w}) \iff \mathbf{B} \vDash \varphi(f(\vec{v})) \iff \mathbf{A} \vDash \varphi(f(\vec{u})).$$

Thus $\mathsf{Hull}_{\mathbf{B}} \vDash \Sigma(\vec{w})$ iff $\mathbf{A} \vDash \Sigma(f(\vec{u}))$, and since $\mathbf{A}$ omits the type, $\mathsf{Hull}_{\mathbf{B}}$ must as well. $\dashv$

Depending on the order-type, as in Theorem 6.A • 2, we can ensure that we realize relatively few types overall. The idea is just that almost everything becomes indiscernible, and so the different types realized over X are just how the representatives for X cut the other indiscernibles, and how the elements are represented. For $n \in \omega$, there are just $|X| \cdot \aleph_0$ such cuts on n-tuples, and using built-in skolem functions, we still just get $|X| \cdot \aleph_0$ types.

6.A • 4. Lemma

Let $\mathcal{L}$ be a countable language. Let T be an $\mathcal{L}$-theory with infinite models. Let κ be an infinite cardinal. Therefore there is an $\mathcal{L}^{\text{sk}}$-model $\mathsf{A} \vDash T^{\text{sk}}$ of size $|A| = \kappa$ such that if $X \subseteq A$, then A realizes at most $|X| \cdot \aleph_0$ types in $S_n^{\mathsf{A}}(X)$.

Proof .:.

By Theorem 6.A • 1, take a model $\mathsf{M} \vDash T$ with A as in the theorem statement played by $\mathsf{H} = \mathsf{Hull}_{\mathsf{M}}(K) \vDash T$, the skolem hull of a set of indiscernibles K of order-type κ. Note that then $|H| = \kappa$. Now let $X \subseteq H$ be arbitrary.

Consider the set of representative elements used from K to get the elements of X:

$$R = \{k \in K : k \text{ is used in the representation of some element of } X\}.$$

Of course, this requires fixing some canonical representations of H beforehand, but this isn't important for the proof. But anyway, from the definition, we get that $|R| \leq |X| \cdot \aleph_0$.

Now we introduce an equivalence relation that tells us precisely when two elements of H have the same type. In particular, for sequences $\vec{x}, \vec{y} \in K^{<\omega}$, write that $\vec{x} \approx_R \vec{y}$ iff $\vec{x}$ and $\vec{y}$ have the same cuts from R. In symbols, $\vec{x} \approx_R \vec{y}$ iff for all $r \in R$ and $i \leq \text{lh}(\vec{x}) = \text{lh}(\vec{y})$,

$$y_i < r \leftrightarrow x_i < r \text{ and } y_i = x \leftrightarrow x_i = r.$$

Note that if $\vec{x} \approx_R \vec{y}$, then the terms given by $\vec{x}$ and those given by $\vec{y}$ in the same way have the same type: for t a skolem term, $\text{tp}^{\mathsf{H}}(t(\vec{x})) = \text{tp}^{\mathsf{H}}(t(\vec{y}))$ for $\vec{x} \approx_R \vec{y}$.

To see this, just by indiscernibility, for $\vec{m} \in X^{<\omega}$ represented by $t'(\vec{z})$ for $\vec{z} \in R^{<\omega}$,

$$\begin{aligned} \mathsf{H} \vDash \varphi(t(\vec{x}), \vec{m}) &\iff \mathsf{H} \vDash \varphi(t(\vec{x}), t'(\vec{z})) \\ &\iff \mathsf{H} \vDash \varphi(t(\vec{y}), t'(\vec{z})) \\ &\iff \mathsf{H} \vDash \varphi(t(\vec{y}), \vec{m}). \end{aligned}$$

So the number of types realized is at most the number of equivalence classes modulo $\approx_R$. But of course, this is just determined by the number of cuts made by R in n-tuples. But there are at most $|R| \cdot \aleph_0 = |X| \cdot \aleph_0$ such cuts, meaning that H realizes at most $|X| \cdot \aleph_0$ types of $S_n^{\mathsf{H}}(X)$, $n < \omega$. $\dashv$

Section 7. ω-Stable Theories

The notion of stability is an important one that will become more important later in the proof of the categoricity theorem. In general, we have the concept of κ-stable theories for $\kappa \geq \aleph_0$ a cardinal. But for the most part we will be concerned only with $\aleph_0 = \omega$-stable theories, since ω-stable theories are κ-stable for all $\kappa \geq \aleph_0$.

7 • 1. Definition

Let $\mathcal{L}$ be a countable langauge. Let T be a complete $\mathcal{L}$-theory with infinite models. Let κ be an infinite cardinal. T is <u>κ-stable</u> iff for all $\mathsf{A} \models T$ and $X \subseteq A$ with $|X| = \kappa$, there are κ complete types of A over X: $|S_n^{\mathsf{A}}(X)| = \kappa$.

In particular, T is <u>ω-stable</u> iff for $\mathsf{A} \models T$, there are only countably many types of A over countable subsets.
An $\mathcal{L}$-model A is called <u>κ-stable</u> iff $\mathrm{Th}(\mathsf{A})$ is κ-stable.

§ 7.A. Connection with saturation

Firstly note that by using a kind of binary tree (i.e. $2^{<\omega}$) of formulas, we get that ω-stable theories are κ-stable for all $\kappa \geq \aleph_0$. The main idea behind the technique is to identify types as branches of $2^{<\omega}$. If we can do this, then there would be $2^{\aleph_0}$ such types, contradicting ω-stability. So our intended construction hinging on the failure of κ-stability can't be carried out.

7.A • 1. Result

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Suppose T is ω-stable. Therefore T is κ-stable for all $\kappa \geq \aleph_0$.

Proof .:.

Suppose $\mathsf{A} \models T$ with $X \subseteq A$ of size $|X| = \kappa \geq \aleph_0$. If $|S_n^{\mathsf{A}}(X)| > \kappa$, then since there are only $\aleph_0 \cdot \kappa = \kappa$ $\mathcal{L}_X$-formulas, there must be some $\mathcal{L}_X$-formula φ_0 where $[\varphi_0] = \{p \in S_n^{\mathsf{A}}(X) : \varphi_0 \in p\}$ has size $|[\varphi_0]| > \kappa$. Repeating this argument, there must be a ψ where both $|[\varphi_0 \wedge \psi]| > \kappa$ and $|[\varphi_0 \wedge \neg\psi]| > \kappa$. So take φ_{00} to be $\varphi_0 \wedge \neg\psi$, and φ_{01} to be $\varphi_0 \wedge \psi$.

Continuing this idea gives a binary tree of formulas $\{\varphi_\tau : \tau \in 2^{<\omega}\}$ such that any two incomparable elements of $2^{<\omega}$ yield contradictory formulas. Taking X' as the set of parameters in these $\mathcal{L}_X$ formulas yields a countable set, and yet $2^{\aleph_0}$ branches or types, meaning $|S_n^{\mathsf{A}}(X')| = 2^{\aleph_0}$, contradicting ω-stability. ⊣

The next theorem gives a nice test that relates categoricity with ω-stability. This connection will be explored further when proving Morley's categoricity theorem.

The general idea of the result is that for a non-ω-stable theory, we can extend in two ways. Firstly, we can expand to realize uncountably many types over some subset X. Secondly, we can take a skolem hull with a lot of indiscernibles that realizes relatively few types. Expanding these to size κ will contradict κ-categoricity.

7.A • 2. Lemma

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Suppose T is κ-categorical for some $\kappa > \aleph_0$. Therefore T is ω-stable.

Proof .:.

Here is where skolem functions come into play. If T isn't ω-stable, then there is some $\mathsf{A} \models T$ which has a countable subset $X \subseteq A$ with $S_n^{\mathsf{A}}(X) > \aleph_0$. Without loss of generality, A is countable. Note that by compactness, we can find an extension $\mathsf{A} \preccurlyeq \mathsf{B}$ with $|B| = \kappa$ realizing uncountably many of these types in $S_n^{\mathsf{A}}(X)$.

But using a previous result of indiscernibles, Lemma 6.A • 4, there is another model $\mathsf{H} \vDash T$ such that for all $Y \subseteq H$, H realizes at most $|Y| \cdot \aleph_0$ types in $S_n^{\mathsf{A}}(Y)$. Note that because $|X| = \aleph_0$, this means H realizes $\aleph_0$ types in $S_n^{\mathsf{A}}(X)$. Hence $\mathsf{H} \not\cong \mathsf{B}$, contradicting κ-categoricity. ⊣

A useful bit about ω-stable theories is that they have saturated models of size κ for all regular cardinals κ. Compared to the relative difficulty of finding saturated models in general, this should immediately be an indication that ω-stable theories are somewhat special.

7.A • 3. Theorem

Let $\mathcal{L}$ be a countable language. Let T be an $\mathcal{L}$-theory. Let $\kappa \geq \aleph_0$ be a regular cardinal. Suppose T is ω-stable. Therefore T has a saturated model of size κ.

Proof .:.

We begin by building an elementary chain of models each of size κ. In particular, let $\mathsf{A}_0 \vDash T$ with $|A_0| = \kappa$. At limit stages we take unions. At successor stages, we let $\mathsf{A}_{\alpha+1}$ realize all the types of A_α: if $\Sigma \in S^{\mathsf{A}_\alpha}(A_\alpha)$, then Σ is realized in $\mathsf{A}_{\alpha+1}$. As T is ω-stable, we can ensure $\mathsf{A}_{\alpha+1}$ has size just $|A_\alpha| + \aleph_0 = \kappa$. If we consider $\mathsf{A} = \bigcup_{\alpha<\kappa} \mathsf{A}_\alpha$, we get that A is saturated.

To see this, note that $|A| = \kappa \cdot \kappa = \kappa$. So for $X \subseteq A$ of size $|X| < \kappa$, all the elements will appear by some stage A_α, $\alpha < \kappa$, and thus any type of $S_n^{\mathsf{A}}(X)$ will be in $S^{\mathsf{A}}(A_\alpha) = S^{\mathsf{A}_\alpha}(A_\alpha)$, which is realized in $\mathsf{A}_{\alpha+1}$, and thus in A. ⊣

The proof here really only requires κ-stability to conclude a saturated model of size κ. Really this is just the same idea as in Theorem 4.D • 9, is just that ω-stability allows us to bound the number of types that we want to realize. Given that this bound really only needs to be κ, we only need κ-stability for the proof.

A result of this is that theories without countably saturated models are not ω-stable. This includes the theory of arithmetic, for example. This shouldn't be so surprising, however, since Theorem 4.D • 4 tells us that the non-existence of a countably saturated model is equivalent to having continuum many types over just the theory.

§ 7.B. Connection with prime model extensions

There's another notion of being a prime model which is slightly more relative. We know that being a prime model in the sense of Definition 4.A • 1 implies that the model must be countable—just take a countable model of the theory, and the prime model must embed in that. We can relativize this notion to models. The key role ω-stability plays is guaranteeing existence and uniqueness to such extensions.

7.B • 1. Definition

Let T be an $\mathcal{L}$-theory. Let $\mathsf{A} \vDash T$ and $X \subseteq A$. A is prime over X or a prime model extension of X iff whenever $\mathsf{B} \vDash T$ with partial elementary $f : X \to B$, there is an elementary embedding $f' : A \to B$ with $f \subseteq f'$.

7.B • 2. Theorem

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Let $\mathsf{A} \vDash T$ with $X \subseteq A$. Suppose T is ω-stable. Therefore there is a prime model extension $\mathsf{A}_0 \preccurlyeq \mathsf{A}$ of X.

Proof .:.

We continually build up X in a sequence $\langle X_\alpha : \alpha < \gamma \rangle$ for some sufficiently large γ. In particular, we set $X_0 = X$, and at limit stages take unions. At successor stages, we take out an element $a_\alpha \in A \setminus X_\alpha$ that realizes an isolated type over X_α, and let $X_{\alpha+1} = X_\alpha \cup \{a_\alpha\}$. Eventually this process stops at some stage γ. Set $A_0 = X_\gamma$.

Firstly, note that A_0 makes sense: if A_0 weren't closed under the functions of A, there would be the isolated type given by "x is the image of these elements of X_γ", which would then need to be realized in $X_{\gamma+1}$, contradicting that the process had stopped.

Now note that $\mathsf{A}_0 \preccurlyeq \mathsf{A}$ just by The Tarksi–Vaught Test (1.B • 1), given that we've added in all the isolated types, and T being ω-stable implies the isolated types in $S_n^{\mathsf{A}}(X)$ are dense—just building a binary tree of formulas as in Corollary 4.B • 3.

To see that A_0 is a prime model extension of X, we expand whatever partial elementary $f : X \to B$, where $\mathsf{B} \vDash T$, to $f' : A_0 \to B$ just by how the isolated types were added: a_α maps to the element in B that realizes the same type. ⊣

Although not stated, we are ensuring that all the types realized in A_0 are isolated over X. This shows the existence of prime model extensions. But we also get uniqueness from ω-stable theories. The proof of this will be delayed until we have the notions of vaughtian pairs.

7.B • 3. Theorem

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models.
Let $X \subseteq A, B$ where $\mathsf{A}, \mathsf{B} \vDash T$ are two prime model extensions of X such that $\mathsf{A}_X \equiv \mathsf{B}_X$.
Suppose T is ω-stable. Therefore $\mathsf{A} \cong \mathsf{B}$, and in fact, there is an isomorphism fixing X.

§ 7.C. Examples and non-examples

ACF is ω-stable. DLO is not ω-stable (consider the types of $\mathsf{Q} = \langle \mathbb{Q}, < \rangle$ over $\mathbb{Q}$, having one for each real number). Anything theory in a countable language that is categorical in an uncountable cardinal is ω-stable.

For examples of prime model extensions, we can again consider DLO. Although it is not ω-stable, we can build a prime extension of any linear order $\mathsf{A} \vDash \mathsf{LO}$ over A just by adding a copy of $\mathbb{Q}$ around the gaps of A. It's easy to see that the resulting $\mathsf{A}' \vDash \mathsf{DLO}$. Moreover, A' is prime over A. If $f : A \to B$ is partial elementary between A' and $\mathsf{B} \vDash \mathsf{DLO}$, we can then clearly extend it to an $f' : A' \to B$. By quantifier elimination, we can view $\mathsf{A}' \subseteq \mathsf{B}$ via f'. By model completeness of DLO, $\mathsf{A}' \preccurlyeq \mathsf{B}$ via f'.

Section 8. Categoricity

This will primarily be working towards Morley's categoricity theorem, but we begin with a short characterization of $\aleph_0$-categorical theories. Although much of the background of Morley's theorem has been considered in Section 7 and Section 6, we will still need to introduce more concepts.

The main route to proving Morley's categoricity theorem is as follows. Although we have not introduced the terminology, a road map is helpful. Firstly, we have a characterization of κ-categorical theories as ω-stable with no vaughtian pairs. This has two directions: the first is as follows.

1. κ-categorical implies ω-stable;
2. κ-categorical implies no $(\kappa, \aleph_0)$-models;
3. ω-stability and $(\aleph_1, \aleph_0)$-models imply $(\kappa, \aleph_0)$-models; and
4. vaughtian pairs imply $(\aleph_1, \aleph_0)$-models.

This tells us that κ-categorical implies ω-stability and the lack of vaughtian pairs. For the other direction,

1. ω-stability implies prime models;
2. prime models of ω-stable theories without vaughtian pairs yield strongly minimal formulas;
3. strongly minimal sets admit partial elementary maps;
4. without vaughtian pairs, these maps extend to isomorphisms; and
5. such isomorphisms yield κ-categoricity for all $\kappa \geq \aleph_1$.

From this, categoricity is easy, as having no vaughtian pairs and being ω-stable doesn't depend on $\kappa \geq \aleph_1$.

§ 8.A. Countable categoricity

Countable categoricity is just have exactly one countably infinite model up to isomorphism. For complete theories in countable languages, we already know that this is equivalent to having both a countably saturated, and prime model. But there are also some other conditions which are of interest.

8.A • 1. Theorem (Ryll–Nardzewski)

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory. Therefore the following are equivalent:

1. T is $\aleph_0$-categorical;
2. Every type in $S_n(T)$ is isolated for $n < \omega$;
3. $|S_n(T)|$ is finite for all $n < \omega$; and
4. For $n < \omega$ variables, there are finitely many $\mathcal{L}$-formulas with these variables up to equivalence in T.

Proof .:.

(1 $\to$ 2) Clearly if there is a non-isolated type $\Sigma(x) \in S_n(T)$, then there are two countable models: one omitting $\Sigma(x)$ and one realizing $\Sigma(x)$ by Omitting Types Theorem (3.B • 2) and Realizing Types Theorem (3.A • 1). These two can't be isomorphic, contradicting $\aleph_0$-categoricity, meaning that all types of $S_n(T)$ are isolated.

(2 $\to$ 3) Suppose $S_n(T)$ were infinite. So let φ_Σ isolate $\Sigma(x) \in S_n(T)$. Since $[\varphi_\Sigma] = \{\Sigma\}$ is then open, we can reduce the cover $\bigcup_{\Sigma \in S_n(T)} \{\Sigma\} = S_n(T)$ to a finite subcover since $S_n(T)$ is compact by Result 4.B • 1. But this means that $S_n(T)$ is finite.

(3 $\to$ 4) Associate a $\varphi_\Sigma \in \Sigma$ for $\Sigma \in S_n(T)$ such that $\varphi_p \notin \Sigma$ for $p \neq \Sigma \in S_n(T)$. For any formula ψ, as there are only $N < \omega$ such types, $T \vdash \psi \leftrightarrow \bigvee_{\Sigma \in S_n(T)} \varphi_\Sigma$. Thus there are only 2^N such formulas up to equivalence.

(4 $\to$ 1) Let $\mathbf{A} \vDash T$ is countable. For $n < \omega$, enumerate the formulas of n variables $\{\varphi_i^n : i \leq N_n\}$. For $\vec{a} \in A^n$,

$n < \omega$, let $\Phi(\vec{a}) = \{i \leq N_n : \mathsf{A} \vDash \varphi_i^n(\vec{a})\}$. Note that $\mathrm{tp}^{\mathsf{A}}(\vec{a})$ is isolated by all the information of $\Phi(\vec{a})$ and the negation of the information outside $\Phi(\vec{a})$, meaning among the formulas φ_i^n for $i \leq N_n$. Hence A is atomic and thus prime. But as A was arbitrary, and all prime models are isomorphic, T is $\aleph_0$-categorical. ⊣

Note that these provide some nice tests of countable categoricity. In particular, the theories of infinite fields are not $\aleph_0$-categorical.

8.A • 2. Example

Let $\mathcal{L}$ be the language of rings. Let $\mathbb{F}$ be an infinite field. Therefore $\mathrm{Th}(\mathbb{F})$ is not $\aleph_0$-categorical.

§ 8.B. Vaughtian pairs and the two-cardinal theorem

To introduce some notation, for φ a formula and A a model with $X \subseteq A$, let $\varphi(X)$ denote the set $\{x \in X : \mathsf{A} \vDash \varphi(x)\}$. In this sense $\varphi(A)$ is just the set defined by φ. This subsection will deal with definable sets and their cardinalities, so this notation will simplify things tremendously.

Now as usually defined, a vaughtian pair is a certain pair of models $\mathsf{A} \preccurlyeq \mathsf{B}$. Instead of thinking of these as two models, however, we will often just consider one model in the expanded language $\mathcal{L}^U = \mathcal{L} \cup \{U\}$ with a predicate U interpreted as membership in A.

8.B • 1. Definition

Let T be a complete $\mathcal{L}$-theory.
A vaughtian pair is a pair of $\mathcal{L}$-models $\mathsf{A} \precneqq \mathsf{B} \vDash T$ such that $\varphi(A) = \varphi(B)$ is infinite for some $\mathcal{L}_A$-formula φ.

Reframed in $\mathcal{L}^U$, if $\mathsf{A} \prec \mathsf{B}$ then $\langle \mathsf{B}, A\rangle \vDash \varphi^U(\vec{a})$ iff $\mathsf{A} \vDash \varphi(\vec{a})$, for all $\mathcal{L}$-formulas φ and $\vec{a} \in A^{<\omega}$.

Important for this concept is the case when $|A| < |B|$, in which case we now get into the ideas of the size of definable sets. For example, a saturated model B can't have an infinite $\varphi(A) = \varphi(B)$ with $|A| < |B|$, because saturated models need to have definable sets as being finite, or else cofinite.

8.B • 2. Definition

Let $\aleph_0 \leq \lambda < \kappa$ be cardinals. A (κ, λ)-model is a model A of size κ with an $\mathcal{L}$-definable subset of size λ.

Hence these notions should not be surprising to appear in discussion of categoricity. If T has a (κ, λ)-model, then we violate saturation as we have a definable subset of size λ for $\aleph_0 \leq \lambda < \kappa$ which is then neither finite nor cofinite. In this section, we will prove one direction of the Baldwin-Lachlan characterization in this section:

1. κ-categorical implies ω-stable;
2. κ-categorical implies no $(\kappa, \aleph_0)$-models;
3. ω-stability and $(\aleph_1, \aleph_0)$-models imply $(\kappa, \aleph_0)$-models; and
4. vaughtian pairs imply $(\aleph_1, \aleph_0)$-models.

We already have (1) through Lemma 7.A • 2. Because κ-sized models of κ-categorical theories are saturated, and thus have definable subsets as finite or cofinite, there can be no $(\kappa, \aleph_0)$-models. Hence (2) is also done. To do (3), we have the following idea.

8.B • 3. Theorem

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models. Let $\kappa \geq \aleph_1$ be a cardinal. Suppose T is ω-stable with an $(\aleph_1, \aleph_0)$-model. Therefore there is a $(\kappa, \aleph_0)$-model of T.

Proof .:.

Let $\mathsf{A} \vDash T$ with $|A| = \kappa$ with an $\mathcal{L}$-formula Φ such that $|\Phi(A)| = \aleph_0$. Now we will show that we can find an extension $\mathsf{A}_1 \succcurlyeq \mathsf{A}_0 = \mathsf{A}$ such that A_1 doesn't realize any more countable types over A.

To do this, by building a binary tree of $\mathcal{L}_A$-formulas otherwise, note that there is an $\mathcal{L}_A$-formula φ such that $|[\varphi]| \geq \aleph_1$ with $|[\varphi \wedge \neg\psi]| \leq \aleph_0$ or else $|[\varphi \wedge \psi]| \leq \aleph_0$ for all $\mathcal{L}_A$-formulas ψ. Now consider the type $\Sigma = \{\psi : |[\varphi \wedge \psi]| \geq \aleph_1\}$. Σ is a complete type over A by the hypothesis on φ. Realizing Σ with some element σ in an elementary extension, take a prime model extension of $A \cup \{\sigma\}$, B. This model works.

Iteratively using this construction, extend the model A κ-many times, taking unions at limit stages. The end result B doesn't add any elements defined by Φ, since $\Gamma = \{\Phi(x), x \neq a : a \in A\}$ is a countable, omitted type in A. Hence B is a $(\kappa, \aleph_0)$-model of T. $\dashv$

(4) above will pop out of the proof of Vaught's two-cardinal theorem. So this is what we will work out of for the rest of the subsection.

8.B • 4. Corollary

Let T be a complete $\mathcal{L}$-theory. Let $\kappa > \lambda \geq \aleph_0$ be cardinals. Suppose T has a (κ, λ)-model as witnessed by φ. Therefore there is a vaughtian pair of models of T of size κ and λ as witnessed by φ.

Proof .:.

Let $\mathsf{B} \vDash T$ be a (κ, λ)-model with $X = \varphi(B)$ of size λ. Taking the skolem hull $\mathsf{A} = \mathsf{Hull}_{\mathsf{B}}(X)$ yields that $\varphi(A) = \varphi(B) = X$ is infinite with $\mathsf{A} \prec \mathsf{B} \vDash T$. Hence these form a vaughtian pair. $\dashv$

The reverse isn't necessarily true, however. Now rather than restricting ourselves to particular, unknown cardinalities, we can always find vaughtian pairs where the models are countable.

8.B • 5. Lemma

Let T be a complete $\mathcal{L}$-theory. Suppose $\langle \mathsf{B}_0, A_0\rangle \equiv \langle \mathsf{B}, A\rangle$ as $\mathcal{L}^U$-models.
Therefore $\mathsf{A} \prec \mathsf{B}$ are a vaughtian pair iff $\mathsf{A}_0 \prec \mathsf{B}_0$ are.

Proof .:.

Suppose $\langle \mathsf{B}, A\rangle$ form a vaughtian pair. Let $\varphi(A) = \varphi(B)$ be infinite. so that $\mathsf{B}_0 \vDash T$. As $\mathsf{A} \prec \mathsf{B}$, for all $\mathcal{L}^U$-formulas ψ,

$$\langle \mathsf{B}, A\rangle \vDash \forall \vec{v}\left(\left(\psi(\vec{v}) \wedge \bigwedge_{i<\mathrm{lh}(\vec{v})} U(v_i)\right) \to \psi^U(\vec{v})\right),$$

and thus in $\langle \mathsf{B}_0, A_0\rangle$ models this as well. Thus $\mathsf{A}_0 \prec \mathsf{B}_0$. Note that $A_0 \neq B_0$ since $\langle \mathsf{B}, A\rangle \vDash \exists x \neg U(x)$. To show that $\varphi(A_0) = \varphi(B_0)$ is infinite, just note that the following sentences hold in $\langle \mathsf{B}, A\rangle$ and thus $\langle \mathsf{B}_0, A_0\rangle$:

1. For each $n < \omega$, the sentence "there are at least n witnesses to φ"; and
2. $\forall \vec{v}(\varphi(\vec{v}) \to \bigwedge_{i<\mathrm{lh}(\vec{v})} U(v_i))$.

But with the conditions we have established before, this means $\langle \mathsf{B}_0, A_0\rangle$ is a vaughtian pair. By symmetry, we get the reverse direction. $\dashv$

A simple application of this with Löwenheim–Skolem (1.B • 3) gives the following.

8.B • 6. Corollary

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory.
Suppose $\mathsf{A} \prec \mathsf{B} \vDash T$ is a vaughtian pair. Therefore there is a vaughtian pair $\langle \mathsf{B}_0, A_0\rangle \preccurlyeq \langle \mathsf{B}, A\rangle$ such that $|B_0| = \aleph_0$.

This allows us to go down with vaughtian pairs. The next result allows us to go back up: realizing a type. By taking enough types, we can ensure that the expanded models are homogeneous, and use this in showing Vaught's two cardinal theorem, a crucial result in proving the downward version of Morley's categoricity theorem.

8.B • 7. Lemma

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory. Suppose $\mathbf{A}_0 \prec \mathbf{B}_0 \models T$ are countable. Let $\vec{a}_0 \in A_0^{<\omega}$ and $\vec{b}_0 \in B_0^{<\omega}$. Therefore,

1. If a type $\Sigma \in S_n^{\mathbf{A}_0}(\vec{a}_0)$ is realized in $\mathbf{B}_0$, then there are countable $\langle \mathbf{B}, A\rangle \succ \langle \mathbf{B}_0, A_0\rangle$ such that $\mathbf{A}$ realizes Σ.
2. If a type $\Sigma \in S_n^{\mathbf{B}_0}(\vec{b}_0)$, then there are countable $\langle \mathbf{B}, A\rangle \succ \langle \mathbf{B}_0, A_0\rangle$ such that $\mathbf{B}$ realizes Σ.

Proof .:.

For (1), just consider the type $\Sigma' = \{\varphi^U : \varphi \in \Sigma\} \cup \mathrm{Eldiag}(\mathbf{B}_0, A_0)$. Note that this is going to be consistent with $\mathrm{Th}(\mathbf{B}_0, A_0)$ since $\mathbf{B}_0$ realizes the type, and $\mathbf{A}_0 \prec \mathbf{B}_0$:

$$\mathbf{B}_0 \models \exists x \left(\textstyle\bigwedge_{i \le n < \omega} \varphi_i(x)\right) \quad \text{implies} \quad \mathbf{B}_0 \models \exists x \left(\textstyle\bigwedge_{i \le n < \omega} \varphi_i^U(x)\right)$$

Realizing this type in an elementary extension yields the result. For (2), we do the same thing, going to an elementary extension which is still then a vaughtian pair. ⊣

Iteratively using this lemma to witness the same types of $S_n(T)$, we get

8.B • 8. Corollary

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory. Suppose $\mathbf{A}_0 \prec \mathbf{B}_0 \models T$ are countable. Therefore there are $\langle \mathbf{B}, A\rangle \succcurlyeq \langle \mathbf{B}_0, A_0\rangle$ such that $\mathbf{A} \cong \mathbf{B}$ are countable, and homogeneous.

This gives a major result in reducing (κ, λ)-models to just $(\aleph_1, \aleph_0)$-models.

8.B • 9. Theorem (Vaught's Two-Cardinal Theorem)

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory. Let $\aleph_0 \le \lambda < \kappa$ be cardinals. Suppose T has a (κ, λ)-model. Therefore T has a $(\aleph_1, \aleph_0)$-model.

Proof .:.

Let φ be an $\mathcal{L}$-formula that witnesses that T has a (κ, λ)-model. By Corollary 8.B • 4, T has a vaughtian pair $\langle \mathbf{B}, A\rangle$ where $|B| = \kappa$ and $|A| = \lambda$, and $\varphi(B) = \varphi(A)$. By Corollary 8.B • 6, there's a vaughtian pair $\langle \mathbf{B}_0, A_0\rangle \prec \langle \mathbf{B}, A\rangle$ where $|B_0| = \aleph_0$. By Corollary 8.B • 8, we can assume $\mathbf{B}_0 \cong \mathbf{A}_0$ are both homogeneous. As an $\mathcal{L}$-elementary substructure, we still have that $\varphi(B_0) = \varphi(A_0)$ with $|B_0 \setminus \varphi(B_0)| = |\varphi(B_0)| = \aleph_0$.

Now we build an elementary chain $\langle \mathbf{B}_\alpha : \alpha \in \omega_1\rangle$ such that $\varphi(B_{\alpha+1}) = \varphi(B_\alpha)$ and

$$\langle \mathbf{B}_{\alpha+1}, B_\alpha\rangle \cong \langle \mathbf{B}_0, A_0\rangle.$$

Of course, we start with $\mathbf{B}_0$ already defined. At limit stages, we take unions. Note that as a union of homogeneous models, $\mathbf{B}_\alpha = \bigcup_{\xi<\alpha} \mathbf{B}_\xi$ is homogeneous. $\mathbf{B}_\alpha$ still realizes the same types as $\mathbf{B}_0$, and so the two are isomorphic by Uniqueness of Homogeneous Models (5.A • 3). Clearly $\varphi(B_\alpha) = \varphi(B_0)$ as well.

At successor stages, since $\mathbf{B}_\alpha \cong \mathbf{B}_0 \cong \mathbf{A}_0$, let $\mathbf{B}_{\alpha+1} \cong \mathbf{B}_0$ be the extension of $\mathbf{B}_\alpha$ as $\mathbf{B}_0$ is to $\mathbf{A}_0$. There are thus no new elements of $B_{\alpha+1}$ defined by φ, and we have all the relevant conditions we want.

Taking $\mathbf{B}_{\omega_1} = \bigcup_{\alpha<\omega_1} \mathbf{B}_\alpha$ yields that $|B_{\omega_1}| = \aleph_1$, and $\varphi(B_{\omega_1}) = \varphi(A_0)$ with $A_0 \subseteq B_{\omega_1}$. Hence $\mathbf{B}_{\omega_1}$ is an $(\aleph_1, \aleph_0)$-model. ⊣

To expand this proof and result to higher cardinalities, we require ω-stability to ensure that we don't add too many types in our construction. But regardless, we have the following result relating categoricity with vaughtian pairs.

8.B • 10. Corollary

Let $\mathcal{L}$ be a countable language. Let T be an $\mathcal{L}$-theory. Suppose T is κ-categorical for some $\kappa \ge \aleph_1$. Therefore T is ω-stable with no vaughtian pairs.

Proof .:.

To show that T has no vaughtian pairs, we can do the same sort of proof as in Vaught's Two-Cardinal Theorem (8.B • 9) to show that otherwise T has an $(\aleph_1, \aleph_0)$-model. So suppose T has a vaughtian pair. T then has countable, homogeneous ones $\mathbf{A} \prec \mathbf{B}$. Producing the same elementary chain as there, we expand $\mathbf{B}$ to size $\aleph_1$ without adding elements defined by whatever φ we care about. In this case, $\mathbf{B}$ has a definable subset of size $\aleph_0$, meaning $\mathbf{B}$ is an $(\aleph_1, \aleph_0)$-model.

Because uncountably categorical theories are ω-stable, it follows by Theorem 8.B • 3 that T has a $(\kappa, \aleph_0)$-model. But this is impossible in a categorical theory, as all definable subsets must be either finite or cofinite. Hence there can be no vaughtian pairs. $\dashv$

§ 8.C. Minimal sets and stability

Moving away from vaughtian pairs and (κ, λ)-models, we will still be focusing on the kinds of definable sets of a theory.

8.C • 1. Definition

Let $\mathbf{A}$ be an $\mathcal{L}$-model. Let $D \subseteq A$ be an infinite set defined by an $\mathcal{L}$-formula φ.
D is minimal in $\mathbf{A}$ iff all subsets of D definable with parameters are finite or cofinite relative to D.
φ is minimal iff $\varphi(A)$ is. D or φ is strongly minimal iff φ is minimal in any extension $\mathbf{B}$ of $\mathbf{A}$.
An $\mathcal{L}$-theory T is strongly minimal iff $x = x$ is strongly minimal, i.e. all subsets of A definable with parameters, where $\mathbf{A} \vDash T$, are either finite or cofinite.

Strongly minimal theories pop up in a variety of circumstances, in particular, ACF_p is strongly minimal for $p \geq 0$. It should be any surprise that in talking about definable subsets, we will be talking about the *algebraic* elements.

8.C • 2. Definition

Let $\mathbf{A}$ be an $\mathcal{L}$-model, and $X \subseteq A$. Let $D \subseteq A$ be strongly minimal.

- $a \in A$ is algebraic over X iff there if is an $\mathcal{L}_X$-formula φ where $a \in \varphi(A)$ and $|\varphi(A)| < \aleph_0$.
- Define $\mathrm{acl}(X) := \{a \in A : a \text{ is algebraic over } X\}$.
- For $X \subseteq D$, define $\mathrm{acl}_D(X) := \mathrm{acl}(X) \cap D$.
- $X \subseteq D$ is independent iff $x \notin \mathrm{acl}(X \setminus \{x\})$ for all $x \in X$.
- $X \subseteq D$ is independent over $Y \subseteq D$ iff $x \notin \mathrm{acl}(Y \cup (X \setminus \{x\}))$ for all $x \in X$.
- X is a basis for a subset $Y \subseteq D$ iff $X \subseteq Y$, X is independent, and $\mathrm{acl}_D(X) = \mathrm{acl}_D(Y)$.
- The dimension of $Y \subseteq D$, $\dim(Y)$, is the cardinality of any basis for Y.

In other words, being algebraic means being defined modulo finite sets. Being independent just means you can't algebraically define any of the elements in terms of the others, a notion similar to linear independence or algebraic independence in the context of fields. Being independent over some Y just means you still can't define any of the elements in terms of the others, but even using Y doesn't help define you modulo a finite set.

The dimension is a similar property as in linear algebra, as the "minimum" number of elements needed to define all of Y modulo finite sets around each element. We will see that the dimension of $Y \subseteq D$ is well defined in that all bases share the same cardinality. Note that if $\mathcal{L}$ is countable, then the cardinality for any basis of Y is just $|Y|$ if $|Y| \geq \aleph_1$.

8.C • 3. Result

Let $X \subseteq Y \subseteq A$ for **A** an $\mathcal{L}$-model. Therefore
1. $\text{acl}(\text{acl}(X)) = \text{acl}(X) \supseteq X$;
2. $\text{acl}(X) \subseteq \text{acl}(Y)$;
3. $a \in \text{acl}(X)$ implies $a \in \text{acl}(X_0)$ for some finite $X_0 \subseteq X$.

8.C • 4. Result

Let **A** be an $\mathcal{L}$-model with $X \subseteq A$, and $a \in \text{acl}(X)$. Therefore $\text{tp}^{\mathsf{A}}(a/X)$ is isolated.

Proof .:.

Let $\varphi(x)$ be an $\mathcal{L}_X$-formula with $|\varphi(A)|$ minimal. We will show that $\varphi(x)$ isolates $\text{tp}^{\mathsf{A}}(a/X)$. To do this, suppose $\mathsf{A} \vDash \varphi(b) \wedge \neg\psi(b)$ for some $b \in A$ and $\mathcal{L}_X$-formula $\psi \in \text{tp}^{\mathsf{A}}(a/X)$. Thus $|\varphi(A) \wedge \psi(A)| < |\varphi(A)|$, contradicting minimality. ⊣

Note that independent sets are totally indiscernible, meaning that we don't even need the order condition of Definition 6 • 1. In this way, distinguishing different independent sets from the perspective of the model can only be done through their cardinalities.

8.C • 5. Result

Let **A**, **B** be two $\mathcal{L}$-theories. Let $X \subseteq M$ where $\mathsf{M} \preccurlyeq \mathsf{A}, \mathsf{B}$.
Let φ be a strongly minimal $\mathcal{L}_X$-formula.
Suppose $\vec{a} \in \varphi(A)^{<\omega}$ and $\vec{b} \in \varphi(B)^{<\omega}$ are independent over X. Therefore $\text{tp}^{\mathsf{A}}(\vec{a}/X) = \text{tp}^{\mathsf{B}}(\vec{b}/X)$.

Proof .:.

Proceed by induction on $\text{lh}(\vec{a}) = n$. For $n = 1$, let $a \in \varphi(A) \setminus \text{acl}(X)$, and $b \in \varphi(B) \setminus \text{acl}(X)$. Suppose $\mathsf{A} \vDash \psi(a)$ for ψ an $\mathcal{L}_X$-formula. We will show $\mathsf{B} \vDash \psi(b)$. To do this, since $a \notin \text{acl}(X)$, $\varphi(A) \cap \psi(A)$ is infinite, which means by strong minimality that $\varphi(A) \setminus \psi(A)$ is finite, and thus

$$|\{x \in A : \mathsf{A} \vDash \varphi(x) \wedge \neg\psi(x)\}| = n$$

for some $n \in \omega$. But this can be represented by an $\mathcal{L}_X$-sentence ψ_n so that by elementarity, $\mathsf{B} \vDash \psi_n$ and since $b \notin \text{acl}(X)$, $\mathsf{B} \vDash \neg\varphi(b) \vee \psi(b)$. Since $\mathsf{B} \vDash \varphi(B)$, it follows that $\mathsf{B} \vDash \psi(b)$, as desired. The inductive case proceeds just like this case, just involving more parameters. ⊣

A simpler result is when $\mathsf{A} = \mathsf{B} = \mathsf{M}$, so that for any strongly minimal $\mathcal{L}_X$-formula φ, the independent elements of $\varphi(A)$ are totally indiscernible. Much like with a basis for linear algebra, because cardinality is the only way to distinguish sets of independent elements. To show that dimension still makes sense as in linear algebra, we need the following principle, which is where strongly minimal sets come into play.

8.C • 6. Lemma (Exchange Principle)

Let **A** be an $\mathcal{L}$-model with strongly minimal set $D \subseteq A$ with $d, \eth \in D$. Suppose $X \subseteq D$ and $d \in \text{acl}_D(X \cup \{\eth\}) \setminus \text{acl}_D(X)$. Therefore $\eth \in \text{acl}_D(X \cup \{d\})$.

Proof .:.

Since we need $\eth$ to define d, let $d \in \varphi(D, \eth)$ for some $\mathcal{L}_X$-formula $\varphi(x, y)$. Suppose $|\varphi(D, \eth)| = n < \omega$. Let $\psi_n(y)$ be the formula stating $|\{x \in D : \varphi(x, y)\}| = n$ so that $\mathsf{A} \vDash \psi_n(\eth)$. If $\psi_n(y)$ defines a finite subset of D, then $\eth$ is algebraic over X, and hence so is d, a contradiction. Hence $\psi_n(D)$ is infinite, and as D is minimal, this means ψ_n defines a cofinite subset of D.

Now consider the subset definable over $\{d\}$, $F = \{y \in D : \varphi(d, y) \wedge \psi_n(y)\}$. If F is finite, we're done: $F \ni \eth$ witnessses $\eth \in \text{acl}_D(X \cup \{d\})$. Thus we may assume F is infinite. Again, minimality tells us F is cofinite. So

let $\theta(x)$ be the $\mathcal{L}_X$-formula stating this: $|E_x| = |D \setminus F| < \omega$ where
$$E_x = \{y \in D : \neg\varphi(x, y) \vee \neg\psi_n(y)\}.$$
Again, we need θ to define a cofinite subset, as otherwise $\mathbf{A} \vDash \theta(d)$ yields $d \in \mathrm{acl}_D(X)$. Hence $\theta(x)$ defines a cofinite subset of D. So choose $m > n$ elements $t_i \in \theta(D)$. By definition of θ, for each $i < m$,
$$G_i = \{y \in D : \varphi(t_i, y) \wedge \psi_n(y)\}$$
is cofinite. So let $g \in \bigcap_{i<m} G_i$. Since $\varphi(t_i, g)$ for each $i < m$, we must have at least $m > n$ elements of $\varphi(D, g)$, contradicting $\mathbf{A} \vDash \psi_n(g)$. ⊣

We now get some immediate results about bases, mirroring basic linear algebra, and using Exchange Principle (8.C • 6).

8.C • 7. Result

Let $\mathbf{A}$ be an $\mathcal{L}$-model. Let $D \subseteq A$ be strongly minimal. Let $X, Y \subseteq D$ be independent with $X \subseteq \mathrm{acl}_D(Y)$.
1. If X and Y are bases for $Z \subseteq D$, then $|X| = |Y|$.
2. Suppose $X_0 \subseteq X$ and $Y_0 \subseteq Y$ have $X_0 \cup Y_0$ as a basis for $\mathrm{acl}(Y)$, and that $x \in X \setminus X_0$. Therefore there is a $y \in Y_0$ where $X_0 \cup \{x\} \cup (Y_0 \setminus \{y\})$ is still a basis for $\mathrm{acl}_D(Y)$.

Now recall the steps to proving the other direction of the Baldwin–Lachlan characterization of uncountably categorical theories. We now have all the background needed to start the process of proving this direction: that ω-stability without vaughtian pairs implies κ-categoricity for all $\kappa \geq \aleph_1$.
1. ω-stability implies prime models;
2. prime models of ω-stable theories without vaughtian pairs yield strongly minimal formulas;
3. strongly minimal sets admit partial elementary maps;
4. without vaughtian pairs, these maps extend to isomorphisms; and
5. such isomorphisms yield κ-categoricity for all $\kappa \geq \aleph_1$.

(1) was accomplished in Theorem 7.B • 2. Of course, (4) implying (5) is immediate, and so now we only need to prove (2), (3), and (4).

We begin to prove (2). This has two steps, first showing that ω-stable theories yield minimal formulas, and if the theory has no vaughtian pairs, then any minimal formula is strongly minimal.

8.C • 8. Lemma

Let T be ω-stable. Let $\mathbf{A} \vDash T$ be an infinite $\mathcal{L}$-model. Therefore $\mathbf{A}$ has a minimal subset. Moreover, if $\mathbf{A}$ is $\aleph_0$-saturated, then all minimal sets are strongly minimal.

Proof .:.

First we prove that $\mathbf{A}$ has a minimal subset. To do this, suppose not. To contradict ω-stability as always, we build a binary tree of formulas with the branches as distinct types. Start with $\varphi_\emptyset(x)$ as $x = x$. At each node φ_τ, since $\varphi_\tau(A)$ isn't minimal, there is a formula ψ where $\varphi_\tau \wedge \psi$ and $\varphi_\tau \wedge \neg\psi$ both define infinite subsets. Hence continuing in this pattern, we get a binary tree of formulas, and contradict ω-stability. ⊣

8.C • 9. Lemma

Let T be ω-stable. Suppose T has no vaughtian pairs. Therefore any model $\mathbf{A} \vDash T$ has a strongly minimal formula. In particular, there is one from the prime model of T.

Proof .:.

We know by Lemma 8.C • 8 that T has a minimal formula. So it suffices to show that any minimal formulas are strongly minimal. First we note that subsets definable with parameters can only get so large in models of T before becoming infinite. This threshold may not be uniform with respect to all $\mathcal{L}$-formulas.

Note that for any formula $\psi(x, y)$ and model $\mathbf{A}$, there is a threshold $n \in \omega$ where if $\vec{a} \in A^{<\omega}$ and $\mathbf{A} \vDash$

$|\psi(A,\vec{a})| > n$, then $\mathsf{A} \vDash \psi(A,\vec{a})$ is infinite. To see this, just elementarily extend A to a model where $\psi(x,\vec{a})$ defines an infinite subset by using a certain type. Additionally, in the language $\mathcal{L}^U$, we may make the type say that U is an elementary submodel, and the set defined by ψ is a subset of U. Since A has arbitrarily large numbers of witnesses to ψ, an elementary extension realizing this type is possible, but yields a vaughtian pair.

For the proof of the actual lemma, suppose φ is a minimal formula over $\mathsf{A} \vDash T$. If φ isn't strongly minimal, we get an extension $\mathsf{A} \preccurlyeq \mathsf{B}$ with $\psi(B,\vec{b})$ infinite and co-infinite in $\varphi(B,\vec{a})$. But then there is a threshold has above that B violates: $\psi(B,\vec{b})$ is infinite and co-infinite in $\varphi(B,\vec{a})$ iff they both have more than $n < \omega$ members. But A says these have less than n members always. Hence by elementarity, we get a contradiction in B. ⊣

Now we begin work on (3): that strongly minimal sets admit partial elementary maps. More specifically, this deals with strongly minimal theories and their dimension: two models of a strongly minimal theory are isomorphic iff they have the same dimension. Although this is the simpler result, what's more useful is that if we share a common elementary substructure and φ is strongly minimal there, then there is a partial elementary map between us so long as $\dim(\varphi)$ is the same for us both.

8.C • 10. Lemma

Let T be an $\mathcal{L}$-theory with no vaughtian pairs. Let $\mathsf{A}, \mathsf{B} \vDash T$ be two $\mathcal{L}$-theories. Let $X \subseteq M$ where $\mathsf{M} \preccurlyeq \mathsf{A}, \mathsf{B}$. Let φ be a strongly minimal $\mathcal{L}_X$-formula such that $\dim(\varphi(A)) = \dim(\varphi(B))$.
Therefore there is partial elementary map $f : \varphi(A) \to \varphi(B)$ a bijection.

Proof .:.

Let A_0 be a basis for $\varphi(A)$, and B_0 a basis for $\varphi(B)$. Since $|A_0| = |B_0|$, we can find a bijection $f_0 : A_0 \to B_0$. Since these are totally indiscernible by Result 8.C • 5, f_0 is partial elementary. Now we must extend f_0 to a partial elementary map between all of $\varphi(A)$ and $\varphi(B)$.

To do this, consider a Zorn's lemma argument on the space

$$\{f_1 : A_1 \to B_1 : A_0 \subseteq A_1 \subseteq \varphi(A) \land B_0 \subseteq B_1 \subseteq \varphi(B) \land f_0 \subseteq f_1\}.$$

Ordered by inclusion and taking unions, Zorn's lemma gives a maximal element $f_1 : A_1 \to B_1$. Now we will show $f_1 : \varphi(A) \to \varphi(B)$ is surjective, and thus a bijection.

To show that $\mathrm{im}(f_1) = B_1 = \varphi(B)$, for $b \in \varphi(B) \setminus B_1$, we know that $\mathrm{tp}^{\mathsf{B}}(b/B_0)$ is isolated by Result 8.C • 4. So take ψ isolating this type. As f_1 is partial elementary, we get a corresponding realization $a \in \varphi(A)$. But then we could extend f_1 to a partial elementary map $f_2 = f_1 \cup \{\langle a, b\rangle\}$, contradicting maximality. The same idea applies to show $\mathrm{dom}(f_1) = \varphi(A)$. ⊣

This allows us to conclude (3), and thus the only remaining part of Baldwin–Lachlan the characterization of κ-categoricity is (4) above. Really the characterization follows the larger steps of

(i) Get strongly minimal formulas which characterize a kind of dimension of models.
(ii) Ensure this dimension characterizes the models up to isomorphism.
(iii) Get κ-categoricity as a result.

To do (ii) above, we need to do (4) from before: extend the partial elementary maps in Lemma 8.C • 10 to full elementary maps. We do this using prime model extensions, which is where we use ω-stability.

8.C • 11. Theorem (Morley's Categoricity Theorem)

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory.
Therefore T is κ-categorical for some cardinal $\kappa \geq \aleph_1$ iff T is κ-categorical for all cardinals $\kappa \geq \aleph_1$.

Proof .:.

We instead prove that T is κ-categorical iff T is ω-stable with no vaughtian pairs. The "only if" direction was proven in Corollary 8.B • 10. For the "if" direction, suppose T is ω-stable with no vaughtian pairs, and let $\kappa \geq \aleph_1$ be arbitrary.

By Theorem 7.B • 2, T has a prime model $\mathbf{P} \vDash T$, and by Lemma 8.C • 9, $\mathbf{P}$ has a strongly minimal $\mathcal{L}_P$-formula $\varphi(x)$. Let $\mathbf{A}, \mathbf{B} \vDash T$ be two models of size $|A| = |B| = \kappa$. As $\mathbf{P} \preccurlyeq \mathbf{A}, \mathbf{B}$ and $\kappa > |P| = \aleph_0$, $\dim(\varphi(A)) = \dim(\varphi(B)) = \kappa$. Hence by Lemma 8.C • 10, there is a partial elementary map $f : \varphi(A) \to \varphi(B)$.

Now we will show that $\mathbf{A}$ is prime over $\varphi(A)$. To see this, note that any elementary submodel $\mathbf{A}_0 \prec \mathbf{A}$ with $\varphi(A) \subseteq A_0$ yields that $\varphi(A) = \varphi(A_0)$ and hence $(\mathbf{A}, A_0)$ is a vaughtian pair, contradicting the hypothesis of T. Now since T is ω-stable, by Theorem 7.B • 2, there is a prime extension of $\varphi(A)$, which then must be $\mathbf{A}$.

Because $\mathbf{A}$ is prime over $\varphi(A)$, we can extend $f : \varphi(A) \to \varphi(B)$ to an elementary embedding $f \subseteq \pi : A \to B$. But the same reasoning as above yields that $\mathbf{B}$ is prime over $\varphi(B)$. Hence $\operatorname{im}\pi$ can't be a proper subset of B, meaning that π is a bijection and thus an isomorphism. As $\mathbf{A}$ and $\mathbf{B}$ were arbitrary, this implies T is κ-categorical. As κ was arbitrary, we have the "if" direction, and thus the result. ⊣

Some consequences of uncountable categoricity include a limit on the possible number of *countable* models.

8.C • 12. Result (Baldwin–Lachlan Theorem)

Let $\mathcal{L}$ be a countable language. Let T be uncountably categorical.
Therefore either T is countably categorical, or T has $\aleph_0$ countable models.

The proof of this is not some easy corollary of Morley's Categoricity Theorem (8.C • 11), but it is still an interesting result of Baldwin and Lachlan's work.

Section 9. *Morley Rank and Degree*

T is a complete theory with infinite models.

Something developed out of Morley's analysis of uncountably categorical theories is the notion of a kind of rank of formulas according to how they decompose the universe of the model. In some sense, this is a refinement of the idea of minimal and strongly minimal formulas. The degree thus comes in with just how much more one can decompose when given a certain rank.

§ 9.A. Morley Rank

The definition of morley rank is difficult enough to parse, and still more difficult to really understand. To help with this, the definition will be presented, and it will be applied to the idea of vector spaces.

9.A • 1. Definition

Let A be an $\mathcal{L}$-model, and φ an $\mathcal{L}_A$-formula.
The morley rank of φ in A, $\mathfrak{R}^{\mathsf{A}}(\varphi)$, is the greatest ordinal α such that $\mathfrak{R}^{\mathsf{A}}(\varphi) \geq \alpha$, which is defined inductively below:

1. $\mathfrak{R}^{\mathsf{A}}(\varphi) \geq 0$ iff $\varphi(A) \neq \emptyset$;
2. $\mathfrak{R}^{\mathsf{A}}(\varphi) \geq \alpha$ for α a limit iff $\mathfrak{R}^{\mathsf{A}}(\varphi) \geq \beta$ for all $\beta < \alpha$;
3. $\mathfrak{R}^{\mathsf{A}}(\varphi) \geq \beta + 1$ iff there are $\mathcal{L}_A$-formulas $\{\psi_n(\vec{x}) : n \in \omega\}$ such that $\mathfrak{R}^{\mathsf{A}}(\psi_n) \geq \beta$ for all $n \in \omega$, and $$\bigsqcup_{n \in \omega} \psi_n(A) \subseteq \varphi(A).$$

If the above hold for all ordinals α, then write $\mathfrak{R}^{\mathsf{A}}(\varphi) = \text{Ord}$. If they hold for no ordinals, write $\mathfrak{R}^{\mathsf{A}}(\varphi) = -1$.
Define $\mathfrak{R}(\varphi)$ as $\mathfrak{R}^{\mathsf{B}}(\varphi)$ for any $\aleph_0$-saturated $\mathsf{B} \succcurlyeq \mathsf{A}$.

The use of (2) tells us that if $\mathfrak{R}^{\mathsf{A}}(\varphi)$ isn't -1 or Ord, then there is an ordinal α with $\mathfrak{R}^{\mathsf{A}}(\varphi) = \alpha$. We don't yet know that the definition of $\mathfrak{R}(\varphi)$ is independent of our choice of $\aleph_0$-saturated extension. So we will aim to prove that after some examples.

For a motivating example, consider the vector space $V = \mathbb{R}^n$. Let $f \in V^*$ be a (non-constant) linear function from V to $\mathbb{R}$. We can then decompose V by preimages: $f^{-1}"\{r\}$ for each $r \in \mathbb{R}$. Note that each of these preimages has dimension (in the sense of linear algebra) $n-1$. So we have decomposed V by a bunch of disjoint subsets each of "rank" $n-1$, and so V should have "rank" n.

For a more concrete and illustrative example, any finite subset definable with parameters $X = \varphi(A)$ has rank $\mathfrak{R}^{\mathsf{A}}(\varphi) = 0$. Clearly $\mathfrak{R}^{\mathsf{A}}(\varphi) \geq 0$. If $\mathfrak{R}^{\mathsf{A}}(\varphi) \geq 1$, then there are infinitely many non-empty sets whose disjoint union is X, contradicting that X is finite. We will also see that with strongly minimal theories, rank is equal to dimension, and in fact that strongly minimal formulas are those with rank 1.

To show that $\mathfrak{R}(\varphi)$ is well defined, we first show that $\aleph_0$-saturated models don't care about parameters that look the same.

9.A • 2. Lemma

Let A be $\aleph_0$-saturated with $\vec{a}, \vec{e} \in A^{<\omega}$. Let $\varphi(\vec{x}, \vec{y})$ be an $\mathcal{L}$-formula.
Suppose $\text{tp}^{\mathsf{A}}(\vec{a}) = \text{tp}^{\mathsf{A}}(\vec{e})$. Therefore $\mathfrak{R}^{\mathsf{A}}(\varphi(\vec{x}, \vec{a})) = \mathfrak{R}^{\mathsf{A}}(\varphi(\vec{x}, \vec{e}))$.

Proof .:.

Write φ_a for $\varphi(\vec{x}, \vec{a})$ and φ_e for $\varphi(\vec{x}, \vec{e})$. Proceed by induction on α that $\Re^{\mathsf{A}}(\varphi_a) \geq \alpha$ iff $\Re^{\mathsf{A}}(\varphi_e) \geq \alpha$. Of course, if $\varphi_a(A) = \emptyset$, then so too is $\varphi_e(A)$, and thus $\Re^{\mathsf{A}}(\varphi_a) \geq 0$ iff $\Re^{\mathsf{A}}(\varphi_e) \geq 0$. For α a limit, the result holds by the inductive hypothesis.

For $\alpha + 1$, suppose $\Re^{\mathsf{A}}(\varphi_a) \geq \alpha + 1$. Thus there are $\mathcal{L}$-formulas $\psi_n(\vec{x}, \vec{u}_n)$ with parameters $\vec{u}_n \in A^{<\omega}$ witnessing that $\Re^{\mathsf{A}}(\varphi_a) \geq \alpha + 1$.

As A is $\aleph_0$-saturated, we can do a back-and-forth argument to see that there are $\vec{w}_n \in A^{<\omega}$ such that $\mathrm{tp}^{\mathsf{A}}(\vec{a}, \vec{u}_n) = \mathrm{tp}^{\mathsf{A}}(\vec{e}, \vec{w}_n)$ for each $n < \omega$. But then by the inductive hypothesis, $\Re^{\mathsf{A}}(\psi_n(\vec{x}, \vec{w}_n)) \geq \alpha$ and so the $\psi_n(\vec{x}, \vec{w}_n)$ witness that $\Re^{\mathsf{A}}(\varphi_e) \geq \alpha + 1$. By symmetry, we get both directions, and so complete the induction. ⊣

Using this, we can show that any two $\aleph_0$-saturated models will agree on the rank so long as one is elementarily embedded in the other.

9.A • 3. Lemma

Let $\mathsf{A} \preccurlyeq \mathsf{B}$ be two $\aleph_0$-saturated models. Let φ be an $\mathcal{L}_A$-formula. Therefore $\Re^{\mathsf{A}}(\varphi) = \Re^{\mathsf{B}}(\varphi)$.

Proof .:.

Again proceed by induction on α to show $\Re^{\mathsf{A}}(\varphi) \geq \alpha$ iff $\Re^{\mathsf{A}}(\varphi) \geq \alpha$. Clearly for α a limit or $\alpha = 0$, the result holds. So it suffices to show the successor case.

So suppose $\Re^{\mathsf{A}}(\varphi) \geq \alpha + 1$. This clearly implies $\Re^{\mathsf{B}}(\varphi) \geq \alpha + 1$ by elementarity: A thinks all the ψ_is are disjoint, and imply φ. By the inductive hypothesis, they all have rank at least α. So they still witness that $\Re^{\mathsf{B}}(\varphi) \geq \alpha + 1$.

Now suppose $\Re^{\mathsf{B}}(\varphi) \geq \alpha + 1$ as witnessed by $\mathcal{L}$-formulas $\psi_n(\vec{x}, \vec{b}_n)$ with parameters $\vec{b}_n \in B^{<\omega}$, and $n < \omega$. Since A is $\aleph_0$-saturated, it must have some parameters $\vec{a}_n \in A^{<\omega}$ where

$$\mathrm{tp}^{\mathsf{B}}(\vec{m}, \vec{b}_n) = \mathrm{tp}^{\mathsf{A}}(\vec{m}, \vec{a}_n),$$

where $\vec{m}$ are the parameters in φ. By elementarity, $\mathrm{tp}^{\mathsf{A}}(\vec{m}, \vec{a}_n) = \mathrm{tp}^{\mathsf{B}}(\vec{m}, \vec{a}_n)$. Hence by Lemma 9.A • 2, changing the parameters doesn't change the rank for B: $\Re^{\mathsf{B}}(\psi_n(\vec{x}, \vec{b}_n)) = \Re^{\mathsf{B}}(\psi_n(\vec{x}, \vec{a}_n)) \geq \alpha$. Yet by the inductive hypothesis, $\Re^{\mathsf{B}}(\psi_n(\vec{x}, \vec{a}_n)) = \Re^{\mathsf{A}}(\psi_n(\vec{x}, \vec{a}_n)) \geq \alpha$. Hence, using elementarity to ensure the other conditions, $\Re^{\mathsf{A}}(\varphi) \geq \alpha + 1$, and so this completes the induction. ⊣

And now we can justify the definition of $\Re(\varphi)$ as in Definition 9.A • 1 just by elementarily extending any two $\mathsf{B}_0, \mathsf{B}_1 \succcurlyeq \mathsf{A}$ to an $\aleph_0$-saturated $\mathsf{C} \succcurlyeq \mathsf{B}_0, \mathsf{B}_1$. The previous lemma says all three then agree on rank when parameters come from A.

9.A • 4. Corollary

Let A be an $\mathcal{L}$-model, and φ an $\mathcal{L}_A$-formula.
Therefore, $\Re^{\mathsf{B}_0}(\varphi) = \Re^{\mathsf{B}_1}(\varphi)$ for any two $\aleph_0$-saturated elementary extensions $\mathsf{A} \preccurlyeq \mathsf{B}_0, \mathsf{B}_1$.

Of course, this doesn't reduce the dependence of the definition on the model the parameters come from, but it does reduce some dependence in that it doesn't matter what the model might *lack* in omitting types.

9.A • 5. Definition

Let A be an $\mathcal{L}$-model. Let φ be an $\mathcal{L}_A$-formula. For $X = \varphi(A) \subseteq A^{<\omega}$, the morley rank of X is $\Re(X) = \Re(\varphi) = \Re^{\mathsf{B}}(\varphi)$ for any $\aleph_0$-saturated $\mathsf{B} \preccurlyeq \mathsf{A}$.

As a result, if X is definable with parameters, its morley rank is $\geq \alpha + 1$ iff there are infinitely many pairwise disjoint Y_n, $n < \omega$, of rank $\geq \alpha$. This is supposed to be a kind of notion of dimension as in other areas, and so it's useful to have the following easy results.

9.A • 6. Result

Let $\mathbf{A}$ be an $\mathcal{L}$-model. Let $X, Y \subseteq A^n$, $n < \omega$, be definable with parameters. Therefore
1. $X \subseteq Y$ implies $\mathfrak{R}(X) \leq \mathfrak{R}(Y)$;
2. $\mathfrak{R}(X \cup Y) = \max(\mathfrak{R}(X), \mathfrak{R}(Y))$;
3. $\mathfrak{R}(X) = 0$ iff X is finite.

It's of course not true that every formula has a rank. For example, any model of DLO has formulas of rank Ord. The property of having *every* formula of ordinal rank is called being totally transcendental, and for countable languages, this turns out to be equivalent to ω-stability.

9.A • 7. Definition

A complete theory T is totally transcendental iff for all $\mathbf{A} \vDash T$, $\mathfrak{R}(\varphi) \in \text{Ord} \cup \{-1\}$ for all $\mathcal{L}_A$-formulas φ.

Now adopting the approach of [2], we will let $\underset{\sim}{\mathbb{M}}$ denote a "monster model" with universe $\mathbb{M}$. In practice, we might as well assume that $\underset{\sim}{\mathbb{M}}$ is just a model of inaccessible cardinality $|\mathbb{M}|$ and much larger than any fixed thing we're working with, and which elementarily contains the models we're considering. In principle, for any particular result we will use this in, the approach could be eliminated. But it makes the arguments simpler to understand. The same approach is used in forcing over the actual universe of sets V. Philosophically speaking, this isn't really possible, but the worry can be eliminated for the sake of consistency results just by slightly modifying the proof relative to a countable $V_\alpha \preccurlyeq V$[iv].

So without too much worry, we'll just assume for this section that there's a proper class of inaccessibles and $\underset{\sim}{\mathbb{M}}$ is a saturated model of one of these cardinals that is sufficiently large relative to the rest of the objects in the proof.

§ 9.B. Morley degree

To further refine morley rank, we can consider the degree of a formula. Note that a formula φ with $\mathfrak{R}(\varphi) = \alpha$ cannot be decomposed into infinitely many disjoint, definable subsets of rank α. The *degree* is then the number φ *can* be partitioned into using subsets of rank α. A curious result is that this is well-defined—or at least there is a maximum number.

9.B • 1. Definition

Let $\mathbf{A}$ be an $\aleph_0$-saturated model. Let φ be an $\mathcal{L}_A$-formula.
The morley degree of φ, $\mathfrak{D}(\varphi)$, is the maximum number $n < \omega$ of $\mathcal{L}_A$-formulas $\psi_0, ..., \psi_{n-1}$ such that $\mathfrak{R}(\psi_i) = \mathfrak{R}(\varphi)$ for all $i < n$, and $\bigsqcup_{i<n} \psi_i(A) \subseteq \varphi(A)$.

As a reminder, we let $\underset{\sim}{\mathbb{M}}$ be a monster model with universe $\mathbb{M}$, which is then $\aleph_0$-saturated.

9.B • 2. Result

Let $\underset{\sim}{\mathbb{M}}$ be a monster model. Let φ be an $\mathcal{L}_{\mathbb{M}}$-formula. Therefore $\mathfrak{D}(\varphi)$ exists.

Proof .:.

Let $\mathfrak{R}(\varphi) = \alpha$ for some ordinal α. Build a binary tree of formulas of morley rank α, T. The sucessor nodes of T will take the following form. For φ_τ already in the tree, suppose for some ψ,

$$\mathfrak{R}(\varphi_\tau \wedge \psi) = \mathfrak{R}(\varphi_\tau \wedge \neg\psi) = \alpha.$$

In this case, let $\varphi_{\tau^\frown 1}$ be $\varphi_\tau \wedge \psi$, and let $\varphi_{\tau^\frown 0}$ be $\varphi_\tau \wedge \neg\psi$. If there are no such ψ, then φ_τ has no successor. So we don't necessarily have $2^{\aleph_0}$ branches in this tree of formulas T. In fact, we'll have only finitely many.

To see this, suppose T is infinite. There is then an infinite branch. This of course yields an infinite antichain,

[iv] Of course, these might not exist in all models of ZFC, but consistently there are relative to Con(ZFC).

which then yield disjoint subsets of $\varphi(\mathbb{M})$, yielding that $\mathfrak{R}(\varphi) \geq \alpha + 1$.

Thus T is finite, and so we can consider the finite branches of T as nodes of T. Suppose there are $d < \omega$ of these branches $\psi_0, ..., \psi_{d-1}$. We claim that $d = \mathfrak{D}(\varphi)$. Note that by construction $\varphi(\mathbb{M}) = \bigsqcup_{i<d} \psi_i(\mathbb{M})$. Now suppose that the result fails: we have $\theta_0, \cdots, \theta_d$ disjoint $\mathcal{L}_{\mathbb{M}}$-formulas of rank α such that $\bigsqcup_{i<d+1} \theta_i(\mathbb{M}) \subseteq \varphi(\mathbb{M})$.

As branches of T, for any $i < d$ and $j < d+1$, we can't have both $\mathfrak{R}(\psi_i \wedge \theta_j) = \alpha$ and $\mathfrak{R}(\psi_i \wedge \neg\theta_j) = \alpha$. Since the θ_j are disjoint, and we have more of them than we do of the ψ_i, it follows that we must have some $j < d+1$, without loss of generality $j = d$, where $\mathfrak{R}(\psi_i \wedge \theta_d) < \alpha$ for all $i < d$. But because $\varphi(\mathbb{M}) = \bigsqcup_{i<d} \psi_i(\mathbb{M})$, we also have $\theta_d(\mathbb{M}) = \bigsqcup_{i<d} \psi_i \wedge \theta_d$, and so this means $\mathfrak{R}(\theta_d) < \max(\{\alpha : i < d\}) = \alpha$ by (2) of Result 9.A • 6. But this contradicts $\mathfrak{R}(\theta_d) = \alpha$ by hypothesis. Therefore, no such formulas $\theta_0, ..., \theta_d$ can exist. $\dashv$

After introducing the concepts of Morley rank and degree, we immediately get some results about degree and some other, previously introduced concepts.

9.B • 3. Result

Let $\underset{\sim}{\mathbb{M}}$ be a monster model. Let φ and ψ be an $\mathcal{L}_{\mathbb{M}}$-formulas such that $\mathfrak{R}(\varphi) = \mathfrak{R}(\varphi \wedge \psi) = \mathfrak{R}(\varphi \wedge \neg\psi)$. Therefore $\mathfrak{D}(\varphi \wedge \psi) < \mathfrak{D}(\varphi)$. Moreover, φ is strongly minimal iff $\mathfrak{R}(\varphi) = \mathfrak{D}(\varphi) = 1$.

Proof .:.

The first result is immediate by Definition 9.B • 1: decomposing $\varphi(\mathbb{M}) \wedge \psi(\mathbb{M})$ into n pieces yields a decomposition of $\varphi(\mathbb{M})$ into $n + 1$ pieces: $\varphi(\mathbb{M}) \wedge \neg\psi(\mathbb{M})$ and the n pieces of $\varphi(\mathbb{M}) \wedge \psi(\mathbb{M})$.

Now suppose φ is strongly minimal. Because $\varphi(\mathbb{M})$ can't be decomposed into two infinite parts, φ has rank $\mathfrak{R}(\varphi) \leq 1$. Since $\varphi(\mathbb{M})$ is infinite, $\mathfrak{R}(\varphi) \geq 1$, and hence equality. Again, since φ can't be decomposed into two infinite parts, $\mathfrak{D}(\varphi) = 1$.

Now suppose $\mathfrak{R}(\varphi) = \mathfrak{D}(\varphi) = 1$. Thus $\varphi(\mathbb{M})$ is infinite and can't be partitioned into two infinite parts, and hence is strongly minimal. $\dashv$

Note that the reverse direction actually made use of the monster model idea for the strength of the minimality.

§ 9.C. Morley rank for types and tuples

Note that types can be said to have a morley rank, just as the minimum of its elements. Its degree can also be considered as the minimum of the degrees—of those with the minimal rank. The benefit of extending it in this way is just to get a handle on types by a single element, a formula witnessing the rank and degree of the type. It will turn out that these formulas determine the type completely.

9.C • 1. Definition

Let **A** be an $\mathcal{L}$-model. Let $\Sigma \in S_n(X)$ be a complete type over $X \subseteq A$.
The morley rank of Σ, $\mathfrak{R}(\Sigma)$, is $\min\{\mathfrak{R}(\varphi) : \varphi \in \Sigma\}$.
If $\mathfrak{R}(\Sigma) \in \mathrm{Ord}$, then the morley degree of Σ, $\mathfrak{D}(\Sigma)$, is $\min\{\mathfrak{D}(\varphi) : \varphi \in \Sigma \wedge \mathfrak{R}(\varphi) = \mathfrak{R}(\Sigma)\}$.

As a result, for each $\Sigma \in S_n(X)$, there is a formula $\Phi_\Sigma \in \Sigma$ witnessing both of these properties: $\mathfrak{R}(\Phi_\Sigma) = \mathfrak{R}(\Sigma)$, and $\mathfrak{D}(\Phi_\Sigma) = \mathfrak{D}(\Sigma)$. As a bit of notation, we will write Φ_Σ given Σ for such a formula.

9.C • 2. Lemma

Let **A** be an $\mathcal{L}$-model with $X \subseteq A$. Let $\Sigma, \Gamma \in S_n(X)$ be distinct, complete types over X.
Suppose $\mathfrak{R}(\Sigma), \mathfrak{R}(\Gamma) \in \mathrm{Ord}$. Therefore $\Phi_\Sigma \neq \Phi_\Gamma$.

Proof .:.

Let $\varphi \in \Sigma \setminus \Gamma$ so that $\neg\varphi \in \Gamma \setminus \Sigma$. As complete types, and by definition of $\mathfrak{R}(\Sigma)$ and $\mathfrak{R}(\Gamma)$, it follows that

$$\mathfrak{R}(\Phi_\Sigma \wedge \varphi) = \mathfrak{R}(\Phi_\Sigma) = \mathfrak{R}(\Sigma)$$
$$\mathfrak{R}(\Phi_\Gamma \wedge \neg\varphi) = \mathfrak{R}(\Phi_\Gamma) = \mathfrak{R}(\Gamma).$$

Yet if $\Phi_\Sigma = \Phi_\Gamma$, then $\mathfrak{R}(\Phi_\Sigma \wedge \varphi) = \mathfrak{R}(\Phi_\Sigma \wedge \neg\varphi) = \mathfrak{R}(\Phi_\Sigma)$. But by Result 9.B • 3, this means $\mathfrak{D}(\Phi_\Sigma \wedge \varphi) < \mathfrak{D}(\Phi_\Sigma) = \mathfrak{D}(\Sigma)$, contradicting the definition of $\mathfrak{D}(\Sigma)$. $\dashv$

This is a useful counting result, allowing us to calculuate the number of types with ordinal rank as bounded by the number of such formulas. This will be important in the following result, which tells us that the concept of totally transcendental theories is unnecessary for countable languages, being equivalent to ω-stability.

9.C • 3. Theorem

Let $\mathcal{L}$ be a countable language. Let T be a complete $\mathcal{L}$-theory with infinite models.
Therefore T is ω-stable iff T is totally transcendental.

Proof .:.

Suppose T is totally transcendental. Let $|X| \leq \aleph_0$. For each complete type Σ of $\underset{\sim}{\mathbb{M}}$ over X, it follows that $\mathfrak{R}(\Sigma) \in \mathrm{Ord}$, and so there is some Φ_Σ witnessing this. By Lemma 9.C • 2, these Φ_Σs are unique, meaning there are as many complete types over X as there are $\mathcal{L}_X$-formulas, being just countably many. Hence T is ω-stable.

Now suppose T is ω-stable. Suppose T isn't totally transcendental, as witnessed by an $\mathcal{L}_{\mathbb{M}}$-formula φ with $\mathfrak{R}(\varphi) = \mathrm{Ord}$. As a set of ordinals, we can take the supremum

$$\rho = \sup\{\mathfrak{R}(\psi) \in \mathrm{Ord} : \psi \text{ is an } \mathcal{L}_{\mathbb{M}}\text{-formula}\}.$$

But as $\mathfrak{R}(\varphi) = \mathrm{Ord}$, we can find formulas ψ with $\mathfrak{R}(\varphi \wedge \psi) \geq \rho + 1$ and $\mathfrak{R}(\varphi \wedge \neg\psi) \geq \rho + 1$ so that $\mathfrak{R}(\varphi \wedge \psi) = \mathfrak{R}(\varphi \wedge \neg\psi) = \mathrm{Ord}$. Continuing in this way allows us to build a binary tree $2^{<\omega}$ of formulas. Taking the countable set of parameters X that these formulas use yields uncountably many types in $S_n(X)$, contradicting ω-stability. $\dashv$

We actually only need a countable language for the "if" direction: ω-stability implies total transcendence for any sized language.

Now we may further extend the notion of morley rank to that of tuples just by the rank of the corresponding type.

9.C • 4. Definition

Let $\underset{\sim}{\mathbb{M}}$ be a monster model. Let $X \subseteq \mathbb{M}$ and $\vec{a} \in \mathbb{M}^{<\omega}$.
The morley rank of $\vec{a}$ over X, $\mathfrak{R}(\vec{a}/X)$, is $\mathfrak{R}(\mathrm{tp}^{\underset{\sim}{\mathbb{M}}}(\vec{a}/X))$. Write $\mathfrak{R}(\vec{a})$ for $\mathfrak{R}(\vec{a}/\emptyset)$.

We then get the following not-super-difficult-but-not-exactly-immediate results below.

9.C • 5. Lemma

Let $\underset{\sim}{\mathbb{M}}$ be a monster model. Let $X \subseteq \mathbb{M}^{<\omega}$ be definable. Therefore

1. If $\alpha < \mathfrak{R}(X)$, then there is a definable $Y \subseteq X$ with $\alpha = \mathfrak{R}(Y)$;
2. $\mathfrak{R}(X) = \sup\{\mathfrak{R}(\vec{x}/Y) : \vec{x} \in X \wedge Y \subseteq \mathbb{M} \wedge |Y| < |\mathbb{M}| \wedge X \text{ is definable over } Y\}$.
3. For $\vec{m} \in \mathbb{M}^{<\omega}$ and $a \in \mathbb{M}$, if a is algebraic over $X \cup \{\vec{m}\}$, then $\mathfrak{R}(a, \vec{m}/X) = \mathfrak{R}(\vec{m}/X)$.

A useful result is that ω-stable theories have a kind of preservation theorem with respect to morley rank. In particular if $f : X \to Y$ is a definable bijection then $\mathfrak{R}(X) = \mathfrak{R}(Y)$.

9.C • 6. Corollary

Let T be a complete ω-stable $\mathcal{L}$-theory. Let $\mathsf{A} \vDash T$ be infinite with definable $X, Y \subseteq A^{<\omega}$. Suppose $f : X \to Y$ is a definable finite-to-one surjection. Therefore $\mathfrak{R}(X) = \mathfrak{R}(Y)$.

Proof .:.

Suppose $M \subseteq \mathbb{M}$ with X and Y definable over M. Now if $f(\vec{x}) = \vec{y}$, then $\vec{y}$ is definable over M and $\vec{x}$. Since f is finite-to-one, $\vec{a}$ is algebraic over M and $\vec{y}$. So $\vec{x}$ and $\vec{y}$ are algebraic over each other relative to M. So by (3) of Lemma 9.C • 5, they have the same rank over M: $\mathfrak{R}(\vec{x}/M) = \mathfrak{R}(\vec{x}, \vec{y}/M) = \mathfrak{R}(\vec{y}/M)$.

Now let $\vec{x} \in X$ be such that $\mathfrak{R}(\vec{x}/M) = \mathfrak{R}(X)$. We thus have by the above reasoning that

$$\mathfrak{R}(Y) \geq \mathfrak{R}(f(\vec{x})/M) = \mathfrak{R}(\vec{x}/M) = \mathfrak{R}(X).$$

But the same reasoning for $\vec{y} \in Y$ with $\mathfrak{R}(\vec{y}/M) = \mathfrak{R}(Y)$ yields the reverse inequality—noting surjection to be able to pull back to such an $\vec{x} \in X$. Hence $\mathfrak{R}(X) = \mathfrak{R}(Y)$. ⊣

Recall Result 9.B • 3, that the morley rank and degree of strongly minimal sets is just 1. In all strongly minimal theories, we can calculate the morley rank of a set just by its dimension. Note that we're taking Ord + Ord to just be Ord here.

9.C • 7. Theorem

Let T be a strongly minimal $\mathcal{L}$-theory with $\underset{\sim}{\mathbb{M}} \vDash T$ a monster model. Let $X \subseteq \mathbb{M}$ with $\vec{m} \in \mathbb{M}^{<\omega}$. Therefore

1. $\mathfrak{R}(\vec{m}/X) = \dim(\vec{m}/X)$.
2. For $A, B \subseteq \mathbb{M}^{<\omega}$ disjoint, and definable, $\mathfrak{R}(A \times B) = \mathfrak{R}(A) + \mathfrak{R}(B)$.

Continuing this idea with strongly minimal theories, we have the ability to more easily get ranks between 1 and Ord.

§ 9.D. Examples and non-examples

First we note that all sorts of ranks can occur. We've already seen strongly minimal sets as having rank and degree 1. This equality $\mathfrak{R}(\varphi) = \mathfrak{D}(\varphi)$ occurs in all strongly minimal theories, like ACF_p for $p \geq 0$.

9.D • 1. Example

Let $p \geq 0$, and suppose $\mathsf{ACF}_p \vDash \exists \vec{x} \varphi(\vec{x})$ for φ an formula in the language of rings. Therefore $\mathfrak{R}(\varphi) \in \{0, 1\}$.

Proof .:.

Clearly $\mathfrak{R}(\varphi) \neq -1$. Note that ACF_p admits quantifier elimination by ACF admits Quantifier Elimination (2.C • 3). What this means is that φ is equivalent to a boolean combination of polynomial equations, each of which has only finitely many solutions, and cofinitely many non-solutions. Hence $\mathfrak{R}(\varphi) \leq 1$. As we're assuming $\mathfrak{R}(\varphi) \geq 0$, this means $\mathfrak{R}(\varphi) \in \{0, 1\}$. ⊣

Note that this is really just a result of strong minimality: every definable subset is either finite or cofinite in any algebraically closed field. Of course, this doesn't extend to formulas with parameters. Pursuing this example further, for $\mathsf{K} \vDash \mathsf{ACF}_p$, $\mathfrak{R}(\vec{k}/X)$ is just the transcendence degree of $\vec{k} \in K^{<\omega}$ over $X \subseteq K$, being the dimension in the model-theoretic sense of Definition 8.C • 2

We also have the following results allowing us to generate more examples of different morley ranks and degrees. Firstly we have Theorem 9.C • 7, allowing us to take the sum of two ranks. The following allows us to take the sum of two degrees as well.

9.D • 2. Result

Let $X, Y \subseteq \mathbb{M}$ be disjoint definable sets with $\mathfrak{R}(X), \mathfrak{R}(Y) \in \mathrm{Ord}$. Therefore

$$\mathfrak{D}(X \sqcup Y) = \begin{cases} \mathfrak{D}(X) + \mathfrak{D}(Y) & \text{if } \mathfrak{R}(X) = \mathfrak{R}(Y) \\ \max(\mathfrak{D}(X), \mathfrak{D}(Y)) & \text{otherwise.} \end{cases}$$

Bibliography

[1] Chen Chung Chang and Jerome Keisler, *Model Theory*, 3rd ed., Studies in logic and the foundations of mathematics, vol. 73, North-Holland Press, 1990. Republished by Dover Publications, 2012.

[2] David Marker, *Model Theory*: *An Introduction*, Graduate Texts in Mathematics, vol. 217, Springer-Verlag, 2002.

[3] Yiannis Moschovakis, *The Lower Predicate Calculus with identity*, 2015, unpublished. Class notes for UCLA Math 114L.

[4] ———, *Lecture Notes in Logic*, 2014, unpublished. Class notes for UCLA Math 220A, 220B, and 220C.

[5] Wolfgang Rautenberg, *A Concise Introduction to Mathematical Logic*, 3rd ed., Universitext, Springer-Verlag, 2010.