UNCOUNTABLY-CATEGORICAL THEORIES

JACOB PARISH

Abstract. Morley's Categoricity Theorem states that if a theory has a unique model of size κ *some* uncountable cardinal κ, then it has a unique model of size κ for *every* uncountable κ. While the statement is interesting in its own right, the theorem is all the more interesting due to the ideas in the proof, which stimulated the development of stability theory in the following decades. This paper provides an exposition of the theorem, following a later proof due to Baldwin and Lachlan.

Contents

1. Introduction

A first-order theory T is called *κ-categorical* if T has a unique model of cardinality κ, up to isomorphism. For finite κ, κ-categoricity is rather uninteresting. If $\mathcal{M}$ is a finite structure of size κ, then the complete theory of all sentences satisfied by $\mathcal{M}$ is κ-categorical, and has no models of any other size. The question of κ-categoricity for infinite κ is far more fruitful. The complete theory of an infinite structure need not be κ-categorical, and by the Löwenheim–Skolem Theorem, it must have models of every infinite size $\geq |T|$. Łos, investigating categoricity during the 1950s, identified three ways in which a theory in a countable language $\mathcal{L}$ can be κ-categorical for infinite κ:

(1) κ-categorical for every infinite κ.
(2) κ-categorical just for $\kappa > \aleph_0$.

Date: August 2022.

(3) κ-categorical just for $\kappa = \aleph_0$.

There are many theories demonstrating each of these possibilities, for instance:

(a) Algebraically closed fields of a fixed characteristic. Such a field is determined up to isomorphism by its transcendence degree over the prime subfield. The fields of finite or countable transcendence degree are all countable; thus there are countably many countable models. On the other hand, if a field has transcendence degree $\kappa > \aleph_0$, then the field itself has size κ, so the theory is κ-categorical.

(b) Infinite vector spaces over a countable field K. This is a theory in the language $\mathcal{L} = \{0, +, -, (k(x) : k \in K)\}$, where each $k(x)$ is a unary function symbol interpreted as scalar multiplication by k. A basic linear algebra fact is that a vector space is determined by its dimension. Since K is countable, the vector spaces of finite or countable dimension are all countable, so there are countably many countable models. On the other hand, if a vector space has dimension $\kappa > \aleph_0$, then the vector space itself has size κ, so the theory is κ-categorical.

If instead K is a finite field, then the theory is $\aleph_0$-categorical as well, since the finite-dimensional vector spaces are no longer infinite.

(c) Infinite sets. This is a theory in an empty language, with an axiom for each $n < \omega$ stating that there are at least n distinct elements. Evidently this theory is κ-categorical for every infinite κ.

(d) Dense unbounded linear orders. The rationals $(\mathbb{Q}, <)$ are the unique countable model of this theory. However, for every uncountable κ, there are 2^κ distinct models.

(e) The random graph. This is a theory of a graph (i.e., a symmetric binary relation) with an axiom for each $n < \omega$ stating that for every two disjoint sets A, B of n vertices, there is a vertex adjacent to each element of A and no element of B. Like the previous theory, this theory is $\aleph_0$-categorical, but there are 2^κ distinct models for every uncountable κ.

Łos asked if possibilities (1)–(3) are the only ways in which a countable theory can be categorical in an infinite cardinal; which is to say, if a theory is κ-categorical for *some* $\kappa > \aleph_0$, must it be κ-categorical for *every* $\kappa > \aleph_0$? Morley, in his 1965 thesis [6], answered this question in the affirmative:

Theorem 1.1 (Morley's Theorem)**.** *Let T be a countable theory. If T is κ-categorical for some $\kappa > \aleph_0$, then T is κ-categorical for all $\kappa > \aleph_0$.*

The goal of this paper is to prove Morley's Theorem. We will not follow Morley's original proof, but rather a 1971 proof by Baldwin and Lachlan in [1]; we will pull from several accounts of the proof, namely [2, Chapter 3.1], [4, Chapter 9], and [5, Chapter 6.1].

From now on T is assumed to be a countable and complete theory with no finite models. When investigating categoricity, completeness is no restriction at all; a simple application of the Löwenheim–Skolem theorem called the Łos–Vaught test states that if T has no finite models and is κ-categorical for some infinite κ, then T is complete.

Examples (a) and (b) above are the motivating examples of uncountably-categorical theories to keep in mind throughout the proof of the theorem. In both examples, models of the theory are determined up to isomorphism by a cardinal invariant: the transcendence degree for algebraically closed fields, and the dimension for vector spaces. An essential point of this proof of the theorem is that we can develop a notion of dimension for *arbitrary* uncountably-categorical theories.

In Section 2, we introduce *pregeometries*, the structure for which this notion of dimension will be developed. Section 3 introduces types and algebraic closure; Sections 4 and 5 apply these concepts, first to a subclass of uncountably-categorical theories called *strongly minimal theories*, and then to strongly minimal *subsets* of theories. Sections 6, 8, and 9 complete the proof of Morley's Theorem, by showing that uncountably-categorical theories are characterized by two properties: $\aleph_0$-stability (defined in Section 6) and the non-existence of Vaughtian pairs (defined in Section 8). Finally, Section 10 briefly examines the number of countable models of an uncountably-categorical theory.

1.1. **Background and Notation.**

This paper assumes a basic knowledge of model theory (for example, the first quarter of the logic sequence at UChicago). This includes the definition of languages and structures, the compactness and Löwenheim–Skolem theorems, and elementary maps and embeddings. For all relevant definitions and theorems, see [3].

Recall, we are assuming that T is a countable complete theory with only infinite models. The letter $\mathcal{L}$ denotes the language of T. The letters $\mathcal{M}, \mathcal{N}$ are used to denote models of T, and the letters M, N are used to denote the underlying sets. If $A \subset M$ is a subset of the domain of a model, then $\mathcal{L}(A)$ is the expanded language obtained by adding constant symbols for each element of A. Abusing notation, we also use the symbol $\mathcal{L}$ to denote the formulas in the language $\mathcal{L}$, so for instance we will write $\varphi(x) \in \mathcal{L}$ for a formula in one free variable over $\mathcal{L}$. For a formula $\varphi(\bar{x}) \in \mathcal{L}$ in n free variables, $\varphi(\mathcal{M})$ denotes the set $\{\bar{a} \in M^n : \mathcal{M} \models \varphi(\bar{a})\}$. As in the previous sentence, $\bar{a}$ typically denotes a finite tuple of elements of a model. We will often write $\bar{a} \in M$ to denote a finite tuple of elements of an unspecified length, rather than explicitly specifying the length and writing $\bar{a} \in M^n$. For brevity, concatenation is often used to denote the union of many sets or elements, so for instance Aa denotes $A \cup \{a\}$.

2. Pregeometries, Independence, Dimension

As stated, a central part of the proof of Morley's Theorem is the following: to every model of an uncountably-categorical theory, we can assign a cardinal invariant called its *dimension*. This notion will generalize the transcendence degree of an algebraically closed field, as well as the dimension of a vector space. To that end, we first define a *pregeometry*, an abstract setting in which one can talk about closure, independence, bases, and dimension.

Definition 2.1. A *pregeometry* (X, cl) is a set X together with a *closure operator* $\mathrm{cl} : \mathcal{P}(X) \to \mathcal{P}(X)$ defined on subsets of X, which satisfies the following for all $A \subset X$:

1. (enlargement) $A \subset \mathrm{cl}(A)$.
2. (idempotency) $\mathrm{cl}(\mathrm{cl}(A)) = \mathrm{cl}(A)$.
3. (finite character) $\mathrm{cl}(A) = \bigcup_{A_0 \subset A \text{ finite}} \mathrm{cl}(A_0)$.
4. (exchange property) If $a \in \mathrm{cl}(Ab) \setminus \mathrm{cl}(A)$, then $b \in \mathrm{cl}(Aa)$.

A set $A \subset X$ is *closed* if $\mathrm{cl}(A) = A$.

If V is a K-vector space, then (V, span) is a pregeometry, whose closed subsets are precisely the linear subspaces.

Definition 2.2. Let (X, cl) be a pregeometry. A subset $I \subset X$ is called *independent* if for each $i \in I$, $i \notin \mathrm{cl}(I \setminus \{i\})$. A subset $I \subset X$ is a *basis for* X if I is independent and $\mathrm{cl}(I) = X$.

For K-vector spaces, these agree with the usual notions of linear independence and bases. In linear algebra, in order to define dimension, one shows that (1) every vector space has a basis, and (2) any two bases of a vector space have the same cardinality. The proof of (1) is via Zorn's Lemma, and the proof of (2) is by a "replacement" lemma, showing that any basis can be transformed into another, one vector at a time. In the case of pregeometries, very similar proofs will show the analogues of both (1) and (2).

Proposition 2.3. *Let (X, cl) be a pregeometry.*

(1) There is a basis I for X.
(2) If J is another basis for X, then $|I| = |J|$.

Proof.

(1) Let $P = \{I \subset X : I \text{ is independent}\}$, and consider the partially-ordered set $(P, \subset)$. If $C \subset P$ is a chain, then $J = \bigcup_{I \in C} I$ is an independent set which is an upper bound for C. Indeed, if J fails to be independent, then by finite character J fails to be independent on a finite subset $J_0 \subset J$, and this J_0 must be contained in some element of C.

By Zorn's Lemma, P contains a maximal element I. If I is not a basis, then there must exist some $i \in \mathrm{cl}(X) \setminus \mathrm{cl}(I)$. If $I \cup \{i\}$ is independent, then this contradicts the maximality of I. So there must exist some $j \in I$ for which $j \in \mathrm{cl}((I \setminus \{j\}) \cup \{i\})$. The exchange property then implies that $i \in \mathrm{cl}(I)$, a contradiction.

(2) We first prove the following claim.

Claim. Suppose $I_0 \subset I$ and $J_0 \subset J$ are subsets such that $I_0 \cup J_0$ is a basis for X. Then if $i \in I \setminus I_0$, there is some $j \in J \setminus J_0$ so that $I_0 \cup \{i\} \cup J_0 \setminus \{j\}$ is a basis for X.

Let $J_1 \subset J_0$ be a minimal subset with $i \in \mathrm{cl}(I_0 \cup J_1)$ (by finite character, there is some finite J_1 for which this holds, so we can take it to be minimal). Since I is independent, there is at least one element $j \in J_1$, any such j will satisfy the claim. We have $i \in \mathrm{cl}(I_0 \cup J_1) \setminus \mathrm{cl}(I_0 \cup J_1 \setminus \{j\})$, so by the exchange property, $j \in \mathrm{cl}(I_0 \cup \{i\} \cup J_1 \setminus \{j\})$. Then $I_0 \cup \{i\} \cup J_0 \setminus \{j\}$ is a generating set: $\mathrm{cl}(I_0 \cup \{i\} \cup J_0 \setminus \{j\})$ contains j, hence $I_0 \cup J_0$, hence all of X. Moreover, $I_0 \cup \{i\} \cup J_0 \setminus \{j\}$ is independent: suppose $i \in \mathrm{cl}(I_0 \cup J_0 \setminus \{j\})$. By minimality of J_1, $i \notin \mathrm{cl}(I_0 \cup J_1 \setminus \{j\})$. By exchange, $j \in \mathrm{cl}(I_0 \cup J_0 \setminus \{j\})$, contradicting the fact that $I_0 \cup J_0$ is independent, and completing the proof of the claim.

If J is finite, then applying the claim inductively, starting with $I_0 = \emptyset$ and $J_0 = J$, we can show $|I| \leq |J|$. Swapping the roles of I and J, we can also show $|J| \leq |I|$.

If J is infinite, then there are $|J|$ many finite subsets $J_0 \subset J$, and by the finite case, each $I \cap \mathrm{cl}(J_0)$ is finite. Since

$$I = I \cap \mathrm{cl}(J) = \bigcup_{J_0 \subset J, J_0 \text{ finite}} I \cap \mathrm{cl}(J_0),$$

we find $|I| \leq |J|$. Swapping the roles of I and J we obtain $|J| \leq |I|$.

□

Thus, the following definition makes sense, and agrees with the notion of dimension in vector spaces.

Definition 2.4. For (X, cl) a pregeometry, the *dimension of* X, denoted $\dim(X)$, is the cardinality of any basis for X.

3. Types and Algebraicity

Our goal now is to give models of an uncountably-categorical theory T the structure of a pregeometry. First, however, we have to introduce a fundamental concept in model theory: types. Roughly, a type describes an element (or set of elements) that might exist in a model of T.

Definition 3.1. Let T be a theory, $\mathcal{M} \models T$, and $A \subset M$. An *n-type over* A is a set p of $\mathcal{L}(A)$-formulas in n free variables which is consistent with T.

An n-type is *complete* if for every formula $\varphi(\bar{x}) \in \mathcal{L}(A)$ in n free variables, either $\varphi(\bar{x}) \in p$ or $\neg\varphi(\bar{x}) \in p$.

If $\mathcal{M}$ contains a tuple $\bar{a}$ such that $\mathcal{M} \models \varphi(\bar{a})$ for all $\varphi(\bar{x}) \in p$, then we say $\bar{a}$ *realizes* p, or that $\mathcal{M}$ realizes p. In general, it need not be the case that $\mathcal{M}$ realizes p, as we show below. However, by the compactness theorem, there is always *some* model $\mathcal{N} \succ \mathcal{M}$ which realizes p.

Note that every tuple $\bar{a} \in M$ realizes a unique complete type over a parameter set A, which we denote by $\mathrm{tp}(\bar{a}/A)$.

Example 3.2. In the theory of algebraically closed fields of characteristic 0, let p be the set of formulas

$$\{q(x) \neq 0 : q(x) \text{ a nonzero polynomial with coefficients in } \mathbb{Q}\ \}.$$

Then p is a 1-type over $\mathbb{Q}$. Note that p is indeed consistent, as any finite set of polynomials has only finitely many solutions. Since p is exactly the type of an element which is transcendental over $\mathbb{Q}$, p is not realized in $\overline{\mathbb{Q}}$, but is realized in every other model of the theory.

Definition 3.3. For T a theory, $\mathcal{M} \models T$, and $A \subset M$, $S_n^T(A)$ denotes the set of complete n-types over A. We write $S^T(A)$ for the union $\bigcup_{n<\omega} S_n^T(A)$. When T is clear from context, we omit it and write $S_n(A)$ and $S(A)$.

Definition 3.4. A type $p \in S_n(A)$ is *isolated* if there is a formula $\varphi(\bar{x}) \in \mathcal{L}(A)$ in n free variables so that for all $\psi \in p$,

$$T \models \forall \bar{x}(\varphi(\bar{x}) \to \psi(\bar{x})).$$

The formula φ is said to *isolate* p.

Remark 3.5. For $p \in S_n(A)$ and $\varphi(\bar{x}) \in \mathcal{L}(A)$ a formula in n free variables, define

$$U_\varphi = \{p \in S_n(A) : \varphi \in p\}.$$

The sets U_φ give a basis for a topology on $S_n(A)$. Then p is an isolated type iff it is an isolated point in this topology, and a formula isolating p is a formula φ such that $U_\varphi = \{p\}$.

In the world of a complete theory T, isolated types are those types that *must* be realized in every model of T. Indeed, a formula φ isolating a type p must be consistent, so $T \models \exists \bar{x} \varphi(\bar{x})$, and any witness is a realization of p. Conversely, if a type is not isolated, then we have the following fundamental result. For a proof, see any model theory text, e.g., [5, Theorem 4.2.3].

Theorem 3.6 (Omitting Types). *Let $\mathcal{M}$ be a model, $A \subset M$, and $p \in S(A)$ a non-isolated type. Then there is an elementary submodel $\mathcal{N} \prec \mathcal{M}$ in which p is not realized.*

Using types, we can define a closure operator that will give our theories the structure of a pregeometry.

Definition 3.7. Let $\mathcal{M}$ be a model, $A \subset M$, and $\varphi(\bar{x}) \in \mathcal{L}(A)$. We say φ is *algebraic* if $\varphi(\mathcal{M})$ is finite. A type $p \in S(A)$ is *algebraic* if it contains an algebraic formula. A tuple $\bar{a} \in M$ is *algebraic over* A if its type $\mathrm{tp}(\bar{a}/A)$ is algebraic. The *algebraic closure of* A, denoted $\mathrm{acl}(A)$, is the set of elements $a \in M$ which are algebraic over A.

Note that whether a formula is algebraic does not depend on the ambient model. Indeed, for some k

$$\mathcal{M} \models \text{“there are exactly } k \text{ elements satisfying } \varphi\text{”},$$

and thus the same is true in any elementary extension of $\mathcal{M}$, or indeed any elementarily equivalent model containing the parameters of φ. Similarly, whether or not an element is algebraic over a given set does not depend on the model, and hence neither does the algebraic closure.

Example 3.8.

(1) In the theory of algebraically closed fields of a fixed characteristic (or in fact, just in the theory of fields), algebraic closure exactly corresponds with algebraic closure in the field-theoretic sense.
(2) In the theory of K-vector spaces, algebraic closure is exactly the linear span.
(3) In the theory of infinite sets, the algebraic closure is trivial, i.e., $\mathrm{acl}(A) = A$ for every set A.

For an arbitrary model $\mathcal{M}$ of a theory T, (M, acl) need not be a pregeometry. However, properties (1)–(3) of 2.1 always hold.

Lemma 3.9. *For any model $\mathcal{M}$ and any $X \subset M$, the restriction of* acl *to X satisfies properties (1)–(3) of a pregeometry.*

Proof.

(1) (enlargement) If $a \in A$, then $\mathrm{tp}(a/A)$ contains the algebraic formula $x = a$, so $a \in \mathrm{acl}(A)$.

(2) (containment) Let $a \in \operatorname{acl}(\operatorname{acl}(A))$. Then $\operatorname{tp}(a/\operatorname{acl}(A))$ contains an algebraic formula $\varphi(x, b_1, \ldots, b_n)$ with parameters $b_i \in \operatorname{acl}(A)$. For each b_i, $\operatorname{tp}(b_i/A)$ contains an algebraic formula $\psi_i(y) \in \mathcal{L}(A)$. The formula

$$\exists y_1, \ldots, y_n \varphi(x, y_1, \ldots, y_n) \wedge \psi_1(y_1) \wedge \cdots \wedge \psi_n(y_n)$$

is an algebraic $\mathcal{L}(A)$-formula contained in $\operatorname{tp}(a/A)$, and hence $a \in \operatorname{acl}(A)$.

(3) (finite character) Let $a \in \operatorname{acl}(A)$. There is an algebraic formula $\varphi(x) \in \mathcal{L}(A)$ which is satisfied by a. Then for any finite $A_0 \subset A$ which contains the parameters of $\varphi(x)$, we have $a \in \operatorname{acl}(A_0)$.

□

4. Strongly Minimal Theories

Unfortunately, the exchange property will not generally hold even for models of uncountably-categorical theories. However, in this section we will investigate a special class of theories for which exchange holds in all models. Later on, we will see that uncountably-categorical theories are in some sense built out of these special theories.

Definition 4.1. A theory T is *strongly minimal* if for every model $\mathcal{M} \models T$, and every formula $\psi(x) \in \mathcal{L}(M)$ in one free variable, either $\psi(\mathcal{M})$ is finite or $\neg\psi(\mathcal{M})$ is finite. In other words, every definable (with parameters) subset of M is either finite or cofinite.

Example 4.2.

(1) If $\mathcal{M}$ is an infinite set, then every formula is equivalent to a boolean combination of formulas $x = a$ for $a \in M$. So a formula can either say "x is one of these finitely many elements", or "x is not one of these finitely many elements". Hence this theory is strongly minimal.

(2) If $\mathcal{M}$ is a K-vector space, then every formula with parameters $a_1, \ldots, a_m \in M$ is equivalent to boolean combination of formulas of the form $k_1(x_1) + \cdots + k_n(x_n) = l_1(a_1) + \cdots + l_m(a_m)$ where $k_i, l_j \in K$ for all $i \leq n, j \leq m$. In case there is only one free variable, i.e., $n = 1$, this formula reduces to saying that x_1 is equal to some vector. Thus this theory is also strongly minimal.

(3) If $\mathcal{M}$ is an algebraically closed field of a fixed characteristic, then every formula over M is equivalent to a boolean combination of polynomial equations (see for instance [Theorem 3.2.2][5]). Since any polynomial has finitely many solutions, all definable subsets of M are finite or cofinite.

Lemma 4.3. *Let $\mathcal{M}$ be a model of a strongly minimal theory. Then (M, acl) satisfies the exchange property, and hence is a pregeometry.*

Proof. Suppose that $a \in \operatorname{acl}(Ab) \setminus \operatorname{acl}(A)$. Then there is an $\mathcal{L}(A)$-formula $\psi(x, y)$ so that $\mathcal{M} \models \psi(a, b)$ and $\psi(\mathcal{M}, b)$ is finite, say of size n. Let $\chi(y)$ be the $\mathcal{L}(A)$-formula stating

$$|\psi(\mathcal{M}, y)| = n.$$

Thus $\mathcal{M} \models \chi(b)$. If $\chi(\mathcal{M})$ is finite, then $b \in \operatorname{acl}(A)$, hence $\operatorname{acl}(Ab) = \operatorname{acl}(A)$, which is a contradiction. Thus, by strong minimality, $\chi(\mathcal{M})$ is cofinite.

If $\psi(a, \mathcal{M})$ is finite, then it witnesses that $b \in \mathrm{acl}(Aa)$ and we are done. So assume $\psi(a, \mathcal{M})$ is cofinite, and let k be the size of the complement. Let $\xi(x)$ be the $\mathcal{L}(A)$-formula stating

$$|\neg\psi(x, \mathcal{M})| = k.$$

Then $\mathcal{M} \models \xi(a)$. Since $a \notin \mathrm{acl}(A)$, $\xi(\mathcal{M})$ must be cofinite, and we may choose distinct solutions $a_1, \dots, a_{n+1}$. For each such solution, $\psi(a_i, \mathcal{M})$ is cofinite. Since $\chi(\mathcal{M})$ is also cofinite, we may choose $b' \in \chi(\mathcal{M}) \cap \psi(a_1, \mathcal{M}) \cap \cdots \cap \psi(a_n, \mathcal{M})$. Then $\chi(b')$ states that $|\psi(\mathcal{M}, b')| = n$, but $\psi(a_i, b')$ holds for each a_i. □

Lemma 4.4. *Let T be a strongly minimal theory, $\mathcal{M} \models T$, and $A \subset M$. There is a unique non-algebraic type in $S_1(A)$.*

Proof. By strong minimality, for any $\psi(x) \in \mathcal{L}(A)$, either $\psi(\mathcal{M})$ is finite or $\psi(\mathcal{M})$ is cofinite. That is, either ψ is algebraic or $\neg\psi$ is algebraic. So there is at most one non-algebraic complete 1-type over A. If $\psi_1, \dots, \psi_n$ are non-algebraic, then each $\psi_i(\mathcal{M})$ is cofinite, so $\psi_1(\mathcal{M}) \cap \cdots \cap \psi_n(\mathcal{M}) \neq \emptyset$. Thus the set of non-algebraic formulas is consistent. □

Lemma 4.5. *Let $\mathcal{M}, \mathcal{N} \models T$ where T is strongly minimal. Let $I \subset M$ and $J \subset N$ be independent sets. If $f : I \to J$ is any bijection, then f is a (partial) elementary map.*

Proof. If f fails to be elementary, then it fails to be so on a finite subset of its domain, so it suffices to consider the case when I and J are finite. In this case we prove the result by induction on $n = |I| = |J|$.

When $n = 1$, let $I = \{a\}$ and $J = \{b\}$. By independence, $a, b \notin \mathrm{acl}(\emptyset)$, and thus a, b both realize the unique non-algebraic type over $\emptyset$. Thus f is elementary.

In the general case, suppose $I = \{a_1, \dots, a_n, a_{n+1}\}$ and $J = \{b_1, \dots, b_n, b_{n+1}\}$, and assume that the map $a_i \mapsto b_i$ for $i \leq n$ is elementary. By 4.4, there is a unique non-algebraic type over $a_1, \dots, a_n$ as well as over $b_1, \dots, b_n$. By independence, a_{n+1}, b_{n+1} must realize these non-algebraic types, and thus the map $a_i \mapsto b_i$ for $i \leq n+1$ is also elementary. □

Lemma 4.6. *Let $\mathcal{M}$ be a model and $A \subset M$. Then $|\mathrm{acl}(A)| \leq |A| + \aleph_0$.*

Proof. There are $|A| + \aleph_0$ many $\mathcal{L}(A)$-formulas, and each can only determine finitely many elements of $\mathrm{acl}(A)$. □

Proposition 4.7. *Let $A \subset M$, $B \subset N$, and let $f : A \to B$ be an elementary map. Then f extends to an elementary map $\hat{f} : \mathrm{acl}(A) \to \mathrm{acl}(B)$.*

Proof. Consider the set

$$P = \{g \supset f : g \text{ is an elementary map, } \mathrm{dom}(g) \subset \mathrm{acl}(A),\ \mathrm{ran}(g) \subset \mathrm{acl}(B)\},$$

partially ordered by inclusion. A chain of maps in P has an upper bound given by the union, because a union of an increasing chain of elementary maps is elementary. Let $\hat{f} \in P$ be a maximal element, $\hat{A} = \mathrm{dom}(\hat{f})$ and $\hat{B} = \mathrm{ran}(\hat{f})$. Suppose $\hat{A} \neq \mathrm{acl}(A)$, and let $c \in \mathrm{acl}(A) \setminus \hat{A}$. Let $\varphi(x, \bar{a})$ be an algebraic formula isolating $\mathrm{tp}(c/\hat{A})$, which is necessarily algebraic. Then $\varphi(x, f(\bar{a}))$ is also algebraic, and isolates a type p over $\hat{B}$. Since p is algebraic it is realized by some $d \in \mathrm{acl}(B)$. We may extend $\hat{f}$ by setting $\hat{f}(c) = d$. This is elementary because both c and d satisfy φ, and hence both have the same types over the domain/range of $\hat{f}$.

The same argument shows $\hat{B} = \operatorname{acl}(B)$. $\square$

We can now prove a preliminary result relating to uncountable categoricity.

Proposition 4.8. *Let T be a strongly minimal theory. Then T is κ-categorical for every $\kappa > \aleph_0$.*

Proof. Let $\mathcal{M}, \mathcal{N} \models T$ be models of cardinality $\kappa > \aleph_0$. Let I be a basis in M, and let J be a basis in N. By 4.6, we must have $|I| = |J| = \kappa$. Let $f : I \to J$ be a bijection. By 4.5, f is an elementary map. By 4.7, f extends to an elementary map $\hat{f} : \operatorname{acl}(I) \to \operatorname{acl}(J)$. But $\operatorname{acl}(I) = M$ and $\operatorname{acl}(J) = N$, and thus $\hat{f}$ is an isomorphism. $\square$

5. Strongly Minimal Sets

As mentioned, it is not the case that all uncountably-categorical theories are strongly minimal, although each example we have seen so far has been strongly minimal. Consider, for instance, the following theories.

Example 5.1.
1. Let T be the complete theory of $(\mathbb{Z}/4\mathbb{Z})^\omega$. This is an uncountably-categorical theory (in fact, totally categorical; a model is determined up to isomorphism by its rank as a $\mathbb{Z}/4\mathbb{Z}$-module). However, the formula $x + x = 0$ defines an infinite and co-infinite subset of any model, so T is not strongly minimal. Also note that no model satisfies the exchange property. For example, if $a \in \mathbb{Z}/4\mathbb{Z}$ is an element of order 4, then $a + a \in \operatorname{acl}(a) \setminus \operatorname{acl}(\emptyset)$, while $a \notin \operatorname{acl}(a + a)$.
2. Let T be the theory of an equivalence relation E with two infinite classes, and a function f interpreted as a bijection between the two classes, with $f^2 = \mathrm{id}$. This is again totally categorical; a model is determined up to isomorphism by the size of either of the classes. However, for any model $\mathcal{M} \models T$ and any element $a \in M$, the formula $E(x, a)$ which singles out one of the two classes defines an infinite and co-infinite subset of M.

However, in the first example, the set *defined* by $x + x = 0$ is just an infinite $\mathbb{F}_2$-vector space, and we know $\mathbb{F}_2$-vector spaces *are* strongly minimal. Similarly, in the second example, for any element a, the formula $E(x, a)$ singling out one of the two classes defines an infinite set, and there is no way to pick out an infinite and co-infinite subset of this equivalence class. This leads to the following idea.

Definition 5.2. Let T be a theory and $\mathcal{M} \models T$. Let $\varphi \in \mathcal{L}(M)$ be a formula in one free variable. Then φ is a *minimal formula* if $\varphi(\mathcal{M})$ is infinite, and every definable (with parameters) subset of $\varphi(\mathcal{M})$ is either finite or cofinite. If φ is also a minimal formula in every elementary extension $\mathcal{N} \succ \mathcal{M}$, then φ is called a *strongly minimal formula.*

Note that a theory is strongly minimal if and only if the formula $x = x$ is strongly minimal. It will turn out that every uncountably-categorical theory *does* have a strongly minimal formula.

In the same manner as 4.3, one can show the following lemma.

Lemma 5.3. *Let T be a theory, $\mathcal{M} \models T$, and let $\varphi \in \mathcal{L}(M)$ be a strongly minimal formula. Then the restriction of* acl *to the set $\varphi(\mathcal{M})$ satisfies the exchange property, and thus $(\varphi(\mathcal{M}), \operatorname{acl})$ is a pregeometry.*

Proof. See 4.3. □

We can also prove a modification of 4.5, with a slight adjustment to deal with parameters.

Lemma 5.4. *Let $\mathcal{M}, \mathcal{N} \models T$, and let $\mathcal{M}_0 \prec \mathcal{M}, \mathcal{N}$ be a common elementary submodel. Let $\varphi \in \mathcal{L}(M_0)$ be a strongly minimal formula, where $\bar{a} \in M_0$ is the parameter required by φ. Let $I \subset \varphi(M)$ and $J \subset \varphi(N)$ be independent sets, and let $f : I \cup \bar{a} \to J \cup \bar{a}$ be a bijection which fixes $\bar{a}$. Then f is an elementary map.*

Proof. See 4.5. □

In Sections 7 and 8, we will see the condition that the parameters of φ lie in some $\mathcal{M}_0 \prec \mathcal{M}, \mathcal{N}$ is no restriction at all, as (1) there is a model of T which embeds into every other model, and (2) the parameters can be taken to lie in this model.

The next order of business is to show that an uncountably-categorical theory has a strongly minimal formula. To that end, we will show that uncountably-categorical theories are characterized by two properties which place very strong restrictions on the definable subsets of any model of the theory.

6. Stability

Definition 6.1. Let T be a theory and κ an infinite cardinal. We say T is *κ-stable* if for all models $\mathcal{M} \models T$, and all $A \subset M$ with $|A| \leq \kappa$, we also have $|S(A)| \leq \kappa$.

The first property characterizing uncountably-categorical theories is $\aleph_0$-stability.

Proposition 6.2. *Let T be a theory which is κ-categorical for some $\kappa > \aleph_0$. Then T is $\aleph_0$-stable.*

In order to prove this proposition, we introduce indiscernible sequences.

Definition 6.3. Let $\mathcal{M}$ be a model, and let $(I, <)$ be a linear ordering. Let $A = \{a_i : i \in I\} \subset M$ be any subset of M indexed by I. We say A is an *indiscernible sequence* in $\mathcal{M}$ (of order-type $(I, <)$) if for every $n < \omega$ and every $i_1, \ldots, i_n, j_1, \ldots, j_n \in I$ with $i_1 < \cdots < i_n$ and $j_1 < \cdots < j_n$, we have $\operatorname{tp}(a_{i_1}, \ldots, a_{i_n}) = \operatorname{tp}(a_{j_1}, \ldots, a_{j_n})$.

For $a_i, a_j \in A$, we may write $a_i < a_j$ to mean that $i < j$ in I, even though the linear ordering is on I and not on A, and even though the linear ordering need not be definable in $\mathcal{M}$.

Note that an indiscernible sequence A has a unique n-type realized by any sequence of n increasing elements from A. The main result about indiscernible sequences is that we can always find models where they exist. Its proof is an application of Ramsey's Theorem, a combinatorial result concerning edge-colorings of hypergraphs.

Lemma 6.4 (Ramsey's Theorem)**.** *Let X be an infinite set, let $n, k < \omega$, and let $[X]^n$ denote the collection of subsets of X of size n. For every coloring of $[X]^n$ using k different colors, there is an infinite set $S \subset X$ for which all elements of $[S]^n$ are the same color.*

Theorem 6.5. *Let T be a theory, and let $(I, <)$ be any linearly ordered set. Then there is a model $\mathcal{M} \models T$ containing an indiscernible sequence of order-type $(I, <)$.*

Proof. Let $\{x_i : i \in I\}$ be a set of variables indexed by I. Let Γ be the set of formulas

(1) $x_i \neq x_j$, for each $i \neq j$ in I.
(2) $\varphi(x_{i_1}, \ldots, x_{i_n}) \leftrightarrow \varphi(x_{j_1}, \ldots, x_{j_n})$, for each $\varphi \in \mathcal{L}$ and $i_1 < \cdots < i_n$ and $j_1 < \cdots < j_n$ in I.

We show $T \cup \Gamma_0$ is consistent for all finite $\Gamma_0 \subset \Gamma$. There is a finite $I_0 \subset I$ so that formulas in Γ_0 only reference variables x_i with $i \in I_0$. There is also a finite set of formulas Φ_0 so that Γ_0 includes only formulas of type (2) when $\varphi \in \Phi_0$. We may assume all formulas in Φ_0 are in n variables.

Let $\mathcal{M} \models T$ be infinite, and let $X = \{a_k : k < \omega\} \subset M$ be a countable subset. Since X is linearly ordered, elements of $[X]^n$ correspond to increasing sequences of m elements of X. We may view the formulas in Φ_0 as determining a $2^{|\Phi_0|}$-coloring of $[X]^n$. That is, if $k_1 < \cdots < k_n$ and $l_1 < \cdots < l_n$ in ω, then $\{a_{k_1}, \ldots, a_{k_n}\}$ and $\{a_{l_1}, \ldots, a_{l_n}\}$ have the same color iff $\mathcal{M} \models \varphi(a_{k_1}, \ldots, a_{k_n}) \leftrightarrow \varphi(a_{l_1}, \ldots, a_{l_n})$ for each $\varphi \in \Phi_0$. By Ramsey's Theorem, there is an infinite $S \subset X$ so that all elements of $[S]^n$ are of the same color, i.e., have the same behavior on each $\varphi \in \Phi_0$. Assigning the variables $\{x_i : i \in I_0\}$ to elements of S, we find Γ_0 is consistent.

By compactness, there is a model $\mathcal{M} \models T \cup \Gamma$. The interpretations of the variables x_i in $\mathcal{M}$ give an indiscernible sequence of order-type $(I, <)$. □

Definition 6.6. A theory T is *model complete* if for every $\mathcal{M} \models T$, any substructure $\mathcal{N} \subset \mathcal{M}$ is actually elementary: $\mathcal{N} \prec \mathcal{M}$.

In fact, one can expand any theory T to an essentially equivalent theory T^* in a larger language which is model complete[1]. We refer to [2, Chapter 2.5] for details on the construction; the basic idea is to add new function symbols (called "Skolem functions") which give a "canonical" solution to every consistent formula. This fact, along with indiscernible sequences, are used to prove that in any theory, we can always find models which realize very few types.

Proposition 6.7. *Let T be a theory. For every infinite cardinal κ, there is a model $\mathcal{M} \models T$ of size κ, so that for every $B \subset M$, $\mathcal{M}$ realizes at most $|B| + \aleph_0$ types in $S(B)$.*

Proof. We assume without loss of generality that T is model complete. By 6.5, there is a model $\mathcal{M} \models T$ containing an indiscernible sequence A of order-type $(\kappa, <)$. Passing to an elementary submodel if necessary, we may assume $\mathcal{M}$ is generated by A.

If for any $B \subset M$, $\mathcal{M}$ realizes at most $|B| + \aleph_0$ types in $S_1(B)$, then by induction one can show the same is true for any $S_n(B)$. Thus, it suffices to show the result for $S_1(B)$.

Fix $B \subset M$. Since $\mathcal{M}$ is generated by I, every $b \in B$ is of the form $t_b(\bar{a}_b)$ for some term t_b, and some $\bar{a}_b \in A$. Let $C = \{a \in A : a \in \bar{a}_b \text{ for some } b \in B\}$. Note that $|C| \leq |B| + \aleph_0$.

Since A is linearly ordered, elements of $[A]^m$ correspond to increasing sequences of m elements in A. Define an equivalence relation $\sim_C$ on $[A]^m$ as follows: for $\bar{a}, \bar{a}' \in [A]^m$, $\bar{a} \sim_C \bar{a}'$ iff $\bar{a}, \bar{a}'$ have the same order-type with respect to C, meaning that for each $c \in C$ and $i \leq m$, we have $a_i < c$ iff $a'_i < c$, and $a_i = c$ iff $a'_i = c$. Let

[1]We caution that this is different from the *model completion* of a theory, which is a theory in the same language as T, and does not always exist.

$[A]^{<\omega} = \bigcup_{m<\omega}[A]^m$, and extend the relation $\sim_C$ to $[A]^{<\omega}$ so that sequences can only be equivalent if they have the same length.

When $\bar{a} \sim_C \bar{a}'$, indiscernibility of A implies that $\mathrm{tp}(\bar{a}/C) = \mathrm{tp}(\bar{a}'/C)$. Then it is also the case that $\mathrm{tp}(t(\bar{a})/C) = \mathrm{tp}(t(\bar{a}')/C)$ for any term t, and since every element of B is a term with parameters in C, we have $\mathrm{tp}(t(\bar{a})/B) = \mathrm{tp}(t(\bar{a}')/B)$. Since every element of $\mathcal{M}$ is a term with parameters in A, the number of types in $S_1(B)$ which are realized in $\mathcal{M}$ is at most $|[A]^{<\omega}/\sim_C| + \aleph_0$.

Now the fact that A has order-type κ becomes relevant. For $a \in A \setminus C$, let $C_a = \{c \in C : c < a\}$ be the initial segment of C determined by a. Note that $\bar{a} \sim_C \bar{a}'$ iff for each $i \leq m$, either $a_i = a_i' \in C$, or $a_i, a_i' \notin C$ and $C_{a_i} = C_{a_i'}$. Since the induced ordering on C is a well-ordering, the number of initial segments of $(C, <)$ is $|C| + 1$. Thus, $|[A]^m/\sim_C| \leq (|C| + |C| + 1)^m \leq |C| + \aleph_0$. It follows that

$$|[A]^{<\omega}/\sim_C| = \sum_{m<\omega} |[A]^m/\sim_C| \leq \aleph_0 \cdot (|C| + \aleph_0) \leq |C| + \aleph_0.$$

Since $|C| + \aleph_0 \leq |B| + \aleph_0$, the result follows. □

Proof of 6.2. Assuming T is not $\aleph_0$-stable, we can construct two non-isomorphic models of size κ.

(1) Let $\mathcal{M} \models T$ and let $A \subset M$ be such that $|A| \leq \aleph_0$ and $|S(A)| > \aleph_0$. We can assume $|M| = \kappa$. By compactness and Löwenheim–Skolem, we can pass to an elementary extension, also of size κ, realizing uncountably many types in $S(A)$.
(2) By 6.7, there is a model $\mathcal{M}$ of size κ which realizes at most $|B| + \aleph_0$ types over any $B \subset M$. In particular it cannot be isomorphic to the first model, which realizes uncountably many types over the countable set A.

□

7. Prime Models in $\aleph_0$-Stable Theories

Definition 7.1. Let T be a theory, $\mathcal{M} \models T$, and $A \subset M$. The model $\mathcal{M}$ is *prime over* A if whenever $\mathcal{N} \models T$ and $f : A \to \mathcal{N}$ is an elementary map, there is an elementary embedding $\mathcal{M} \prec \mathcal{N}$ which extends f. A model of T is *prime* if it is prime over $\emptyset$, meaning it embeds into every model of T.

The key theorem in this section is that prime models over parameter sets always exist for $\aleph_0$-stable theories.

Lemma 7.2. *Let T be $\aleph_0$-stable, $\mathcal{M} \models T$, $A \subset M$, and $\varphi \in \mathcal{L}(A)$ a consistent formula. Then there is some isolated type $p \in S(A)$ which contains φ.*

Proof. Supposing otherwise, we will construct $2^{\aleph_0}$ distinct types over a countable $A' \subset A$ to contradict $\aleph_0$-stability. Let $2^{<\omega}$ be the set of finite $\{0,1\}$-valued sequences. Inductively, for each $s \in 2^{<\omega}$, we can find consistent formulas $\varphi_s \in \mathcal{L}(A)$ so that

- $\varphi_\emptyset = \varphi$.
- If $s \subset t$, then $\mathcal{M} \models \forall \bar{x}(\varphi_t(\bar{x}) \to \varphi_s(\bar{x}))$.
- If $s \not\subset t$, and $t \not\subset s$, then $\mathcal{M} \models \neg\exists \bar{x}(\varphi_s(\bar{x}) \wedge \varphi_t(\bar{x}))$.

Having chosen φ_s, there must be a formula $\psi \in \mathcal{L}(A)$ so that $\varphi_s \wedge \psi$, $\varphi_s \wedge \neg\psi$ are consistent. Indeed, were this not the case, φ_s would isolate a complete type containing φ. We then set $\varphi_{s,0} = \varphi_s \wedge \psi$ and $\varphi_{s,1} = \varphi_s \wedge \neg\psi$.

Since we used countably many formulas to build the tree, we can pass to a countable $A' \subset A$ which contains the parameters of each φ_s. Let 2^ω be the set of countable $\{0,1\}$-valued sequences. For each $f \in 2^\omega$, the set $\{\varphi_s : s \subset f\}$ is a consistent set of $\mathcal{L}(A')$-formulas, so let $p_f \in S(A')$ be a complete type extending $\{\varphi_s : s \subset f\}$. Then we obtain $2^{\aleph_0}$ distinct elements of $S(A')$, contradicting $\aleph_0$-stability. □

Proposition 7.3. *Let T be $\aleph_0$-stable, $\mathcal{M} \models T$, and $A \subset M$. Then there is $\mathcal{N} \prec \mathcal{M}$ which is a prime model over A.*

Proof. Let $\lambda = |A| + \aleph_0$. Passing to an elementary submodel if necessary, we may assume $|M| = \lambda$. Inductively, for $\alpha < \lambda$, we choose elements $c_\alpha \in M$ and sets $C_\alpha \subset M$ so that

- $C_\alpha = \{c_\beta : \beta < \alpha\}$.
- $\text{tp}(c_\alpha/A \cup C_\alpha)$ is isolated.
- If $\psi(x) \in \mathcal{L}(C_\alpha)$ is consistent, then $\mathcal{M} \models \psi(c_\beta)$ for some $\beta < \lambda$.

Let $\{\psi_\alpha(x) : \alpha < \lambda\}$ be an enumeration of the consistent $\mathcal{L}(M)$-formulas in one free variable. At step α, let

$$\Gamma = \{\gamma < \lambda : \psi_\gamma(x) \in \mathcal{L}(A \cup C_\alpha), \text{ and for all } c_\beta \in C_\alpha,\ \mathcal{M} \not\models \psi_\gamma(c_\beta)\ \}.$$

If $\Gamma \neq \emptyset$, let γ be the least element of Γ, and let $\psi = \psi_\gamma$. Otherwise, let ψ be $x = x$. By 7.2, there is an isolated type $p_\gamma \in S_1(A \cup C_\alpha)$ which contains ψ. Let $c_\alpha \in M$ be a realization of p_γ.

By the Tarski–Vaught test, $A \cup C_\lambda$ is the universe of an elementary submodel $\mathcal{N} \prec \mathcal{M}$. We have to show that $\mathcal{N}$ is a prime model over A. Let $\mathcal{M}'$ be another model and $f : A \to \mathcal{M}'$ an elementary map. Inductively, we extend f to the elements of C_λ. Suppose we have extended the domain of f to $A \cup C_\alpha$, so that it remains an elementary map. Since $\text{tp}(c_\alpha/A \cup C_\alpha)$ is isolated in $\mathcal{N}$, $f(\text{tp}(c_\alpha/A \cup C_\alpha))$ is isolated in $\mathcal{M}'$. Thus it is realized by some $d_\alpha \in M'$, and we may define $f(c_\alpha) = d_\alpha$. When we are done, f is an elementary embedding of $\mathcal{N}$ into $\mathcal{M}'$. □

8. Vaughtian Pairs

The second property characterizing uncountably-categorical theories is the (lack of) the following.

Definition 8.1. Let $\mathcal{M}, \mathcal{N} \models T$. The pair $(\mathcal{M}, \mathcal{N})$ is called a *Vaughtian pair* (for T) if $\mathcal{N} \prec \mathcal{M}$, and if there is a formula $\varphi \in \mathcal{L}(N)$ in one free variable for which $\varphi(\mathcal{N})$ is infinite and $\varphi(\mathcal{M}) = \varphi(\mathcal{N})$.

This places another significant restriction on the definable subsets. It means that given a model of T, there is no way to extend the model without increasing the size of *every (infinite) definable set*.

Proposition 8.2. *Let T be κ-categorical for some $\kappa > \aleph_0$. Then T has no Vaughtian pairs.*

The proof of this proposition is perhaps the most difficult step of the theorem, requiring some additional machinery from stability theory. Its proof will be delayed until later. However, we will prove the following in this section, which assuming 8.2, will complete the proof of Morley's Theorem.

Proposition 8.3. *Let T be an $\aleph_0$-stable theory which has no Vaughtian pairs. Then T is κ-categorical for every $\kappa > \aleph_0$.*

The most important remaining step is to show that we can find a strongly minimal formula.

Lemma 8.4. *Let T be a theory with no Vaughtian pairs, and let $\psi(x, \bar{y}) \in \mathcal{L}$. There is some n, depending on ψ, so that for all $\mathcal{M} \models T$ and $\bar{a} \in M$,*

$$|\psi(\mathcal{M}, \bar{a})| > n \implies |\psi(\mathcal{M}, \bar{a})| \geq \aleph_0.$$

Proof. Let $\mathcal{L}' = \mathcal{L} \cup \{U, c_1, \ldots, c_m\}$ be an expanded language, where U is a unary relation and $c_1, \ldots, c_m$ are constants (m is the length of the $\bar{y}$ parameter in $\psi(x, \bar{y})$). Let T' be the $\mathcal{L}'$-theory which says

- $U(\mathcal{M})$ is the domain of a proper $\mathcal{L}$-elementary submodel of $\mathcal{M}$.
- For each c_i, $U(c_i)$ holds.
- $\psi(\mathcal{M}, c_1, \ldots, c_m) \subset U(\mathcal{M})$.

Suppose towards contradiction that for each $n < \omega$, we have models $\mathcal{N}_n \models T$ where $n < |\psi(\mathcal{N}_n, \bar{c})| < \aleph_0$. Let $\mathcal{M}_n \succ \mathcal{N}_n$ be a proper elementary extension. Then also $n < |\psi(\mathcal{M}_n, \bar{c})| < \aleph_0$, since finiteness is preserved under elementary extension. Evidently, each $\mathcal{M}_n$ can be expanded to a model $\mathcal{M}'_n \models T'$. Then by compactness, there is a model $\mathcal{M}' \models T$ where $|\psi(\mathcal{M}', \bar{c})|$ is infinite. But this gives us a Vaughtian pair for T. □

The preceding lemma can be thought of as showing that in theories with no Vaughtian pairs, "there exist infinitely many" can be re-written, in first-order, as "there exist more than n". This allows minimal formulas to be upgraded to strongly minimal formulas, as shown in the following corollary.

Corollary 8.5. *Let T be a theory with no Vaughtian pairs, and let $\varphi \in \mathcal{L}(M)$ be a minimal formula over some model $\mathcal{M} \models T$. Then φ is strongly minimal.*

Proof. Fix a formula $\psi(x, \bar{y}) \in \mathcal{L}$. Let n_1 and n_2 be as given by 8.4 for $\varphi \wedge \psi$ and $\varphi \wedge \neg\psi$, respectively. By minimality of φ,

$$\mathcal{M} \models \forall \bar{y}(|\varphi(M) \cap \psi(M, \bar{y})| \leq n_1 \vee |\varphi(M) \cap \neg\psi(M, \bar{y})| \leq n_2).$$

Since this holds in every elementary extension of $\mathcal{M}$, φ is strongly minimal. □

All that remains is to find a minimal formula over the prime model of T.

Lemma 8.6. *If T is $\aleph_0$-stable and $\mathcal{M} \models T$, then there exists a minimal formula $\varphi \in \mathcal{L}(M)$.*

Proof. Suppose there is no minimal formula over $\mathcal{M}$. In a similar manner to 7.2, for each $s \in 2^{<\omega}$, we will find consistent formulas $\varphi_s \in \mathcal{L}(M)$ so that

- $\varphi_\emptyset$ is $x = x$.
- If $s \subset t$, then $\mathcal{M} \models \forall x(\varphi_t(x) \to \varphi_s(x))$.
- If $s \not\subset t$, and $t \not\subset s$, then $\mathcal{M} \models \neg\exists \bar{x}(\varphi_s(\bar{x}) \wedge \varphi_t(\bar{x}))$.
- $\varphi_s(\mathcal{M})$ is infinite.

Having chosen φ_s, since $\varphi_s(\mathcal{M})$ is infinite and φ_s is not minimal, there is some formula $\psi \in \mathcal{L}(M)$ partitioning $\varphi_s(M)$ into two infinite subsets. Let $\psi_{s,0} = \psi_s \wedge \psi$, and let $\psi_{s,1} = \psi_s \wedge \neg\psi$. Concluding in the same way as 7.2, we contradict $\aleph_0$-stability. □

Combining 7.3, 8.6, and 8.5 quickly yields the following proposition.

Proposition 8.7. *Let T be an $\aleph_0$-stable theory with no Vaughtian pairs. Then there is a strongly minimal formula over the prime model of T.*

Proof of 8.3. Let $\mathcal{M}, \mathcal{N} \models T$ be models of size κ, and let $\mathcal{M}_0 \prec \mathcal{M}, \mathcal{N}$ be the prime model of T. By 8.7, T has a strongly minimal formula $\varphi \in \mathcal{L}(M_0)$. Let $\bar{a} \in M_0$ be the parameter required by φ. Since T has no Vaughtian pairs, we must have $|\varphi(\mathcal{M})| = |\varphi(\mathcal{N})| = \kappa$. Choose bases $I \subset \varphi(\mathcal{M})$, $J \subset \varphi(\mathcal{N})$. By 4.6, we must have $|I| = |J| = \kappa$. Let $f : I \cup \bar{a} \to J \cup \bar{a}$ be a bijection fixing $\bar{a}$, which is an elementary map by 5.4. By 4.7, f can be extended to an elementary map $\hat{f} : \varphi(\mathcal{M}) \cup \bar{a} \to \varphi(\mathcal{N}) \cup \bar{a}$. By 7.3, there is a model $\mathcal{M}' \prec \mathcal{M}$ which is prime over $\varphi(\mathcal{M}) \cup \bar{a}$. But since T has no Vaughtian pairs, it must be the case that $\mathcal{M}' = \mathcal{M}$. Thus $\hat{f}$ extends further to an elementary embedding $\mathcal{M} \prec \mathcal{N}$. If this embedding is proper, then $(\mathcal{N}, \mathcal{M})$ is a Vaughtian pair as witnessed by φ. So in fact, the embedding is an isomorphism $\mathcal{M} \cong \mathcal{N}$. □

Combining everything, we obtain

Theorem 8.8. *Let T be a theory, and $\kappa > \aleph_0$. Then T is κ-categorical if and only if T is $\aleph_0$-stable and has no Vaughtian pairs.*

Proof. The first implication is by 6.2 and 8.2. The other implication is by 8.3. □

Corollary 8.9 (Theorem 1.1, Morley's Theorem)**.** *If T is κ-categorical for some $\kappa > \aleph_0$, then T is κ-categorical for all $\kappa > \aleph_0$.*

Proof. This follows immediately from 8.8, as one side of the bidirectional implication does not depend on κ. □

9. Non-existence of Vaughtian Pairs

In this section we fill in the missing piece by proving 8.2. Actually, we will prove the following proposition.

Proposition 9.1. *Suppose T is $\aleph_0$-stable and has a Vaughtian pair. Then for any infinite κ, T has a Vaughtian pair $(\mathcal{M}, \mathcal{N})$ where $|M| = \kappa$ and $|N| = \aleph_0$.*

Then 8.2 follows quickly: let $\varphi \in \mathcal{L}(N)$ be a formula witnessing that T has such a Vaughtian pair. By compactness we can produce a model $\mathcal{M}'$ where $|\varphi(\mathcal{M}')| = \kappa$. Thus we have two distinct models of size κ for all $\kappa > \aleph_0$.

This section, in particular the proof of 9.8, will use a number of basic facts about saturated models, which can be found in the appendix to this paper. A first lemma shows we can reduce any Vaughtian pair to a pair of countable, saturated models.

Lemma 9.2. *Suppose T is $\aleph_0$-stable and has a Vaughtian pair. Then T has a Vaughtian pair consisting of countable saturated models.*

Proof. Suppose we have a Vaughtian pair for T, witnessed by the formula φ. In a similar manner to 8.4, we form an expanded theory T' in the langauge $\mathcal{L}' = \mathcal{L} \cup \{U, c_1, \ldots, c_n\}$, which says

- $U(\mathcal{M})$ is the domain of a proper $\mathcal{L}$-elementary submodel of $\mathcal{M}$.
- For each c_i, $U(c_i)$ holds.
- $\varphi(\mathcal{M}, c_1, \ldots, c_n) \subset U(\mathcal{M})$.

Since we have a Vaughtian pair for T, we have a model $\mathcal{M}' \models T'$, which we may take to be a countable model of T'. Now we form a chain of countable models $\mathcal{M}'_n \models T'$ for $n < \omega$, such that

- $\mathcal{M}'_0 = \mathcal{M}'$.
- $\mathcal{M}'_n \prec \mathcal{M}'_{n+1}$.
- if $A \subset \mathcal{M}'_n$ is finite, and $p \in S^T(A)$, then p is realized in $\mathcal{M}'_{n+1}$; if in fact $A \subset U(\mathcal{M}'_n)$, then p is realized in $U(\mathcal{M}'_{n+1})$.

Having constructed $\mathcal{M}'_n$, we construct $\mathcal{M}'_{n+1}$ as follows. Enumerate as $\{A_i : i < \omega\}$ the finite subsets of M'_n. For each A_i, since T is $\aleph_0$-stable, there are at most countably many types in $S^T(A_i)$. If in fact $A_i \subset U(\mathcal{M}'_n)$, then for each type $p \in S^T_k(A_i)$, the set $p \cup \{U(x_1), \ldots, U(x_k)\}$ is consistent, since $U(\mathcal{M}'_n)$ is an elementary submodel of $\mathcal{M}'_n$. Thus we can take a countable extension of $\mathcal{M}'_{n+1} \succ \mathcal{M}'_n$ where each such p is realized, and if the parameters of p are in U, then p is realized in U.

The union $\bigcup_{n<\omega} \mathcal{M}'_n$ is a countable model of T', and it induces the desired Vaughtian pair of countable, saturated models of T. □

From here, we prove 9.1 by showing that we can "stretch" the larger model in a Vaughtian pair $(\mathcal{M}, \mathcal{N})$ to a larger model $\mathcal{M}'$, so that $(\mathcal{M}', \mathcal{N})$ remains a Vaughtian pair. Repeating this step κ times, we can extend the larger model to have size κ, as desired.

The difficulty is ensuring that in an extension, we still have $\varphi(\mathcal{M}') = \varphi(\mathcal{N})$. First of all, we should choose the extension to be as small as possible. Since we still want it to be a proper extension, one way to enforce that it is as small as possible is to take it to be *prime* over $M \cup \{b\}$ for some choice of $b \in M' \setminus M$. However, not only is it the case that an arbitrary choice of b might satisfy φ, but an arbitrary choice of b might introduce some other c into M' which satisfies φ. Thus b must be chosen carefully.

Any element $a \in M \setminus N$ is an element that does not satisfy φ, and which did not introduce any new elements into M that satisfy φ. So we might choose $b \in M' \setminus M$ which "looks like" the element $a \in M \setminus N$. That is, we pick $b \in M' \setminus M$ for which $\mathrm{tp}(b/N) = \mathrm{tp}(a/N)$. Now b does not satisfy φ, since it has the same type as a. But suppose that b introduced a new element c into M' which satisfies φ. For such a c, $\mathrm{tp}(c/Mb)$ must be isolated by a formula $\psi(x, b) \in \mathcal{L}(M)$, because $\mathcal{M}'$ is prime over $M \cup \{b\}$. In particular, $\psi(x, b)$ implies $\varphi(x)$. We would like to transfer this knowledge back to a, reaching a contradiction by finding an element of $M \setminus N$ satisfying φ.

We might argue as follows: the formulas $\exists x(\psi(x, y))$, $\forall x(\psi(x, y) \to \varphi(x))$, and $\forall x(\psi(x, y) \to x \neq n)$ for $n \in N$ are all satisfied by b. Since a has the same type as b over N, a also satisfies them, and hence there is a $c_0 \in M \setminus N$ which satisfies φ.

Unfortunately, this is incorrect, because *$\psi(x, b)$ is parametrized in M, not in N*. We do not know that a and b behave alike with respect to other elements of M. (In fact, they obviously do not; $x = a$ is an $\mathcal{L}(M)$-formula satisfied by a and not by b). What we want is to choose $b \in M' \setminus M$ which looks like a "generic" element of $M \setminus N$. To formalize this notion, we need to introduce some ideas from stability theory.

Definition 9.3. Let $p \in S(B)$ be a type. We say p *does not split* over $A \subset B$ if for all $\varphi(\bar{x}, \bar{y}) \in \mathcal{L}$, and for all $\bar{b}, \bar{b}' \in B$ with $\mathrm{tp}(\bar{b}/A) = \mathrm{tp}(\bar{b}'/A)$, the following holds:

$$\varphi(\bar{x}, \bar{b}) \in p \iff \varphi(\bar{x}, \bar{b}') \in p.$$

That is, membership of a formula in p depends only on the parameter type over A. In case $q \in S(A)$ and $q \subset p$, we say that p is a *strong heir* of q.

If p does not split over A, this somehow means A already contains enough information to specify p. Strong heirs are the most "free" or "generic" extensions, i.e., the ones that add as little information as possible from the larger parameter set. To resolve the issue described above, what we will end up doing is to taking b to be a realization of a strong heir of $\mathrm{tp}(a/N)$ in $S_1(M)$.

Example 9.4. Let $T = ACF_0$. Let $p \in S_1(\mathbb{Q})$ be the unique non-algebraic 1-type. Let t be a realization of p in some model, i.e., t is a transcendental element, and consider the larger parameter set $\mathbb{Q}(t)$. If $q \in S_1(\mathbb{Q}(t))$ is the unique non-algebraic type, then q is a strong heir of p. On the other hand, if $q \in S_1(\mathbb{Q}(t))$ is the algebraic type isolated by $x = t$, then even though q extends p, q is not a strong heir of p since it does not split over $\mathbb{Q}$; for instance, $\mathrm{tp}(t/\mathbb{Q}) = \mathrm{tp}(t+1/\mathbb{Q})$, but $x = t \in q$ while $x = t + 1 \notin q$. Adding the formula "$x = t$" was too specific of a choice.

Now let $p \in S_1(\mathbb{Q})$ be the algebraic type isolated by $x^2 = 2$. Let $q \in S_1(\mathbb{Q}(\sqrt{2}))$ be the extension of p isolated by $x = \sqrt{2}$. Then q splits over $\mathbb{Q}$, i.e., is not a strong heir of p, since $\mathrm{tp}(\sqrt{2}/\mathbb{Q}) = \mathrm{tp}(-\sqrt{2}/\mathbb{Q})$, but q does not contain $x = -\sqrt{2}$. Unlike the previous case, p has *no* strong heir over $\mathbb{Q}(\sqrt{2})$; any extension of p *must* pick between "$x = \sqrt{2}$" $\in p$ and "$x = -\sqrt{2}$" $\in p$.

In $\aleph_0$-stable theories, the splitting relation has particularly nice properties. We do not prove them here, but collect them in the following lemma.

Lemma 9.5. *Let T be $\aleph_0$-stable.*

(1) Let $\mathcal{M} \models T$, $A \subset M$, and $p \in S(A)$. There is a finite subset $A_0 \subset A$ over which p does not split.

(2) Let $\mathcal{N} \models T$ be an $\aleph_0$-saturated model. For any $p \in S(N)$ and any $\mathcal{M} \succ \mathcal{N}$, p has a strong heir $q \in S(M)$.

Proof. See [2, Lemma 3.1.9] for a proof of (1) and [2, Lemma 3.1.10] for a proof of (2). □

In the informal argument described above, we required $\mathcal{M}'$ to be prime over $M \cup \{b\}$. Since we require $\aleph_0$-saturation in order to guarantee the existence of strong heirs, we will actually need a different notion of "primeness" which the model $\mathcal{M}'$ should satisfy, which can be thought of as "prime relative to the $\aleph_0$-saturated models".

If $p \in S(A)$ and $B \subset A$, then by $p \restriction B$ we denote the type $\{\varphi(\bar{x}, \bar{a}) \in p : \bar{a} \in B\} \in S(B)$.

Definition 9.6. Let $\mathcal{M} \models T$. We say $\mathcal{M}$ is $\aleph_0$-*prime* (over $A \subset M$) if (1) it is $\aleph_0$-saturated, and (2) for any $\aleph_0$-saturated model $\mathcal{N} \models T$ (containing A), there is an elementary embedding $\mathcal{M} \prec \mathcal{N}$ (fixing A).

A type $p \in S(M)$ is $\aleph_0$-*isolated* if there is a finite $B \subset M$ so that p is the only complete extension of $p \restriction B$ to M. The model $\mathcal{M}$ is $\aleph_0$-*atomic* if $\mathrm{tp}(\bar{a}/M)$ is $\aleph_0$-isolated for every $\bar{a} \in M$.

Just as isolated types are the types that *must* be realized in every model, $\aleph_0$-isolated types are the types that *must* be realized in every $\aleph_0$-saturated model.

The analogue of 7.3 holds for $\aleph_0$-prime models of $\aleph_0$-stable theories as well; the proof strategy is similar and we omit it.

Lemma 9.7. *Let T be $\aleph_0$-stable, $\mathcal{M} \models T$ an $\aleph_0$-saturated model, and $A \subset M$. There is a model $\mathcal{N} \prec \mathcal{M}$ containing A which is both $\aleph_0$-prime over A and $\aleph_0$-atomic over A.*

Proof. See [2, Lemma 3.1.6]. □

Now we can state and prove the "stretching" lemma.

Lemma 9.8. *Suppose T is $\aleph_0$-stable and has a Vaughtian pair. Whenever $(\mathcal{M},\mathcal{N})$ is a Vaughtian pair of $\aleph_0$-saturated models of T and $\mathcal{N}$ is countable, there is an $\aleph_0$-saturated model $\mathcal{M}' \gneqq \mathcal{M}$, so that $(\mathcal{M}',\mathcal{N})$ is a Vaughtian pair.*

Proof. Let $a \in M \setminus N$, and $p = \operatorname{tp}(a/N)$. By 9.5, as $\mathcal{N}$ is $\aleph_0$-saturated, p has a strong heir $q \in S(M)$. Moreover, there is a finite set $A \subset N$ over which q, hence p, does not split. Extending A if necessary, assume A contains the parameters of φ. Let $\mathcal{M}' \succ \mathcal{M}$ be an $\aleph_0$-saturated model containing some $b \in M'$ realizing q. By 9.7, we may take $\mathcal{M}'$ to be $\aleph_0$-prime and $\aleph_0$-atomic over $M \cup \{b\}$.

Suppose there is some $c \in M' \setminus M$ with $\mathcal{M}' \models \varphi(c)$. Since $\mathcal{M}'$ is $\aleph_0$-atomic over $M \cup \{b\}$, $\operatorname{tp}(c/Mb)$ is $\aleph_0$-isolated over $B \cup \{b\}$ for some finite $B \subset M$. Extending B if necessary, we assume $B \supset A$. Since $\mathcal{M}$ is $\aleph_0$-saturated, there is a countable saturated model $\mathcal{N}_0 \prec \mathcal{M}$ containing B. Uniqueness of countable saturated models implies there is an isomorphism $f_0 : \mathcal{N}_0 \cong \mathcal{N}$ which fixes A. We claim f_0 extends to an elementary map $f_1 : N_0 \cup \{b\} \to N \cup \{a\}$. Indeed, for $\bar{d} \in N_0$,

$$\begin{aligned}\varphi(x,\bar{d}) \in \operatorname{tp}(b/N_0) &\iff \varphi(x,\bar{d}') \in \operatorname{tp}(b/N_0) \text{ whenever } \operatorname{tp}(\bar{d}'/A) = \operatorname{tp}(\bar{d}/A)\\ &\iff \varphi(x,f_0(\bar{d}')) \in \operatorname{tp}(a/N) \text{ whenever } \operatorname{tp}(f_0(\bar{d}')/A) = \operatorname{tp}(f_0(\bar{d})/A)\\ &\iff \varphi(a,f_0(\bar{d})) \in \operatorname{tp}(a/N),\end{aligned}$$

where the first implication is because $\operatorname{tp}(b/N_0)$ does not split over A, the second is because f_0 is an isomorphism $N_0 \cong N$ which fixes A, and the third is because $\operatorname{tp}(a/N)$ does not split over A.

Let $\mathcal{M}_0' \prec \mathcal{M}'$ be $\aleph_0$-prime over $N_0 \cup \{b\}$. Since $\operatorname{tp}(c/N_0 b)$ is $\aleph_0$-isolated, $\mathcal{M}_0'$ contains c. Since $\mathcal{M}$ is $\aleph_0$-saturated and contains $N \cup \{a\}$, f_1 extends to an embedding f of $\mathcal{M}_0'$ into $\mathcal{M}$. But then $f(c) \in M \setminus N$ is an element satisfying φ, contradicting the fact that $(\mathcal{M},\mathcal{N})$ is a Vaughtian pair. □

With the above lemma, we can prove 9.1, which will conclude the proof of Morley's Theorem.

Proof of 9.1. By 9.2, T has a Vaughtian pair $(\mathcal{M},\mathcal{N})$ of countable, saturated models. Applying 9.8, we may construct a sequence of models $\{\mathcal{M}_\alpha : \alpha < \kappa\}$ so that

- $\mathcal{M}_0 = \mathcal{M}$.
- For each α, $(\mathcal{M}_\alpha,\mathcal{N})$ is a Vaughtian pair, $\mathcal{M}_\alpha$ is $\aleph_0$-saturated, and $\mathcal{M}_\alpha \precneqq \mathcal{M}_{\alpha+1}$.
- At limits, $\mathcal{M}_\alpha = \bigcup_{\beta<\alpha} \mathcal{M}_\beta$.

Let $\mathcal{M}' = \bigcup_{\alpha<\kappa} \mathcal{M}_\alpha$. Then $(\mathcal{M}',\mathcal{N})$ is a Vaughtian pair, and since a new element is added at each step, $|\mathcal{M}'| \geq \kappa$. By Löwenheim–Skolem we can reduce to a model of size precisely κ. □

10. Countable Models of Uncountably-Categorical Theories

The Baldwin–Lachlan proof of Morley's Theorem yields the following corollary.

Corollary 10.1. *Let T be κ-categorical for some $\kappa > \aleph_0$. Then T has at most countably many countable models.*

Proof. The strongly minimal part of a countable model of T has either finite or countable dimension; in proving 8.3, we showed two models of T whose strongly minimal parts have the same dimension are isomorphic. □

Each of the examples of uncountably-categorical theories that we have seen fits into one of two cases:

(1) The only countable model is of dimension $\aleph_0$ (e.g., infinite sets, vector spaces over a finite field, the two examples of 5.1).
(2) There are countably many countable models (e.g., algebraically closed fields, vector spaces over a countable field).

This leads to the following question: is it possible that an uncountably-categorical theory has more than one but only finitely many countable models? Baldwin and Lachlan proved in [1] that the answer is no: an uncountably-categorical theory is either totally categorical, or has exactly $\aleph_0$ countable models. While the full Baldwin–Lachlan result is beyond the scope of the paper, we will be able to show that it holds assuming there is a strongly minimal formula *without parameters*.

Proposition 10.2. *Suppose that T is an uncountably-categorical theory, and $\varphi \in \mathcal{L}$ is a strongly minimal formula. Then T has either one or $\aleph_0$ countable models.*

Proof. Since φ requires no parameters, we can refer unambiguously to the dimension of any model of T (more specifically, the dimension of the strongly minimal part).

Suppose T is not countably-categorical. Then T has a model $\mathcal{M}_k$ of some finite dimension k. Let $I_k \subset \varphi(\mathcal{M}_k)$ be a basis. Let $p \in S_1(I_k)$ be the unique nonalgebraic type containing φ. Let $\mathcal{M} \succ \mathcal{M}_k$ be an elementary extension containing an element a which realizes p, and let $I_{k+1} = I_k \cup \{a\}$. Let $\mathcal{M}_{k+1} \prec \mathcal{M}$ be a model which is prime over I_{k+1}. It is clear that $\mathcal{M}_{k+1}$ has dimension at least $k+1$. To show it has dimension exactly $k+1$, suppose $b \in \varphi(\mathcal{M}_{k+1})$ is independent from I_{k+1}. Since $\mathcal{M}_{k+1}$ is prime over I_{k+1}, $\mathrm{tp}(b/I_{k+1})$ is isolated by some $\psi \in \mathcal{L}(I_{k+1})$. Since $\mathcal{M}_{k+1} \models \psi(b)$, and $b \notin \mathrm{acl}(I_{k+1})$, it must be the case that $\psi(\mathcal{M}_{k+1})$ is cofinite in $\varphi(\mathcal{M}_{k+1})$. Since $\mathrm{acl}(I_{k+1})$ is infinite (it contains $\mathrm{acl}(I_k) = \varphi(\mathcal{M}_k)$), there is some element of $\mathrm{acl}(I_{k+1})$ satisfying ψ. But this is absurd, as ψ isolates a non-algebraic type over I_{k+1}.

Repeating this argument, we obtain a model of every finite dimension $\geq k$. □

As mentioned above, the difficulty arises when the strongly minimal formula φ requires a parameter $\bar{a}$ in the prime model $\mathcal{M}_0$. In this case, given an arbitrary model of T, there may be many realizations of $\mathrm{tp}(\bar{a})$, and each one could, a priori, give rise to a different dimension. So there may be one model which, viewed from many different perspectives, has many possible finite dimensions. The extra work Baldwin and Lachlan did is to show this is not the case: different realizations of $\mathrm{tp}(\bar{a})$ in the same model always give rise to the same dimension. In fact, the extra work in this case is necessary: there are uncountably-categorical theories in which there is no strongly minimal $\mathcal{L}$-formula.

Example 10.3. Recall from 5.1 the theory of an equivalence relation E with two infinite classes and a bijection f between the two classes so that $f^2 = \mathrm{id}$. If $\varphi \in \mathcal{L}$ is a formula in one free variable, then $\varphi(\mathcal{M})$ is invariant under the automorphism defined by f. So if $\varphi(\mathcal{M})$ is infinite, then $E(x, a)$ defines an infinite and co-infinite subset of $\varphi(\mathcal{M})$, and φ is not strongly minimal. That is, a strongly minimal formula must pick out one of the two classes, but there is no way of doing so without choosing a parameter.

Appendix: Saturated Models

This appendix includes the definition and some basic facts about saturated models. For proofs, see [3, Chapter 5].

Definition 10.4. Let κ be an infinite cardinal. A model $\mathcal{M} \models T$ is *κ-saturated* if for all $A \subset M$ with $|A| < \kappa$, $\mathcal{M}$ realizes every type in $S(A)$. We say $\mathcal{M}$ is *saturated* if it is $|M|$-saturated.

Proposition 10.5 (Existence of κ-Saturated Models)**.** *For any $\mathcal{M} \models T$ and any κ, there exists an elementary extension $\mathcal{N} \succ \mathcal{M}$ which is κ-saturated.*

Proof Sketch. We will find a κ^+-saturated extension, which clearly is also κ-saturated. Construct an elementary chain of models $\{\mathcal{M}_\alpha : \alpha < \kappa^+\}$ satisfying the following:

- $\mathcal{M}_0 = \mathcal{M}$.
- For each $A \subset M_\alpha$ with $|A| \leq \kappa$, every type in $S(A)$ is realized in $\mathcal{M}_{\alpha+1}$ (apply compactness for each A and each $p \in S(A)$).
- If α is a limit ordinal, then $\mathcal{M}_\alpha = \bigcup_{\beta<\alpha} \mathcal{M}_\beta$.

Let $\mathcal{N} = \bigcup_{\alpha<\kappa^+} \mathcal{M}_\alpha$. Then regularity of κ^+ implies that $\mathcal{N}$ is κ^+-saturated. □

Proposition 10.6 (Uniqueness of Saturated Models)**.** *Let $\mathcal{M}, \mathcal{N} \models T$ be saturated models with $|M| = |N| = \kappa$. Let $A \subset M$, $|A| < \kappa$, and $f : A \to N$ an elementary map. Then f can be extended to an isomorphism $\mathcal{M} \cong \mathcal{N}$. In particular, any two saturated models of the same cardinality are isomorphic.*

Proof Sketch. Let $\{a_\alpha : \alpha < \kappa\}$ and $\{b_\alpha : \alpha < \kappa\}$ be enumerations of M and N, respectively. By induction on $\alpha < \kappa$, we extend f as follows:

1. If a_α is not already in $\operatorname{dom} f$, use saturation of $\mathcal{N}$ to find $b \in N$ realizing the type $\{\varphi(x, f(\bar{c})) : \varphi(x, \bar{c}) \in \operatorname{tp}(a_\alpha / \operatorname{dom} f)\}$, and set $f(a_\alpha) = b$.
2. If b_α is not already in $\operatorname{ran} f$, use saturation of $\mathcal{M}$ to find $a \in M$ realizing $\{\varphi(x, f^{-1}(\bar{c})) : \varphi(x, \bar{c}) \in \operatorname{tp}(b_\alpha / \operatorname{ran} f)\}$, and set $f(a) = b_\alpha$.

After κ steps, we have an isomorphism $\mathcal{M} \cong \mathcal{N}$. □

We caution that in general, saturated models of a given cardinality do not always exist, and κ-saturated models of the same cardinality need not be isomorphic.

Acknowledgments

I am deeply grateful to Isabella Scott for her mentorship and guidance throughout the summer, for her detailed comments on drafts of this paper, and for suggesting this topic. I would also like to thank Professor Maryanthe Malliaris for meeting with me to discuss project ideas, as well as Olga Medrano Martin del Campo for organizing the logic group meetings throughout the summer. Finally, I would like to thank Peter May for organizing the University of Chicago REU, along with all of the speakers for giving excellent talks.

References

[1] J. T. Baldwin and A. H. Lachlan. "On Strongly Minimal Sets" in *The Journal of Symbolic Logic.* Association for Symbolic Logic. 1971.
[2] Steven Buechler. *Essential Stability Theory.* Springer. 1996.
[3] C.C. Chang and H. J. Kiesler. *Model Theory.* Elsevier Science Publishers. 1990.
[4] Artem Chernikov. *Lecture Notes on Stability Theory.* AMS Open Math Notes. 2019.
[5] David Marker. *Model Theory.* Springer. 2002.
[6] Michael Morley. "Categoricity in Power" in *Transactions of the American Mathematical Society.* American Mathematical Society. 1965.