

Set Theory and Model Theory

VERSION 3

Rahim Moosa, University of Waterloo

January 2015

Contents

Part 1. Set Theory	1
Chapter 1. Ordinals	3
1.1. Well-orderings	3
1.2. Zermelo-Fraenkel Axioms	4
1.3. Cartesian products and functions	7
1.4. Well-ordering the natural numbers	7
1.5. Introducing ordinals	9
1.6. Transfinite induction and recursion	11
1.7. Ordinal arithmetic	13
1.8. Well-orderings and ordinals	15
Chapter 2. Axiom of Choice	17
2.1. A first look at cardinals	17
2.2. The axiom of choice and its equivalents	18
2.3. Axiom of choice and cardinals	20
Chapter 3. Cardinals	23
3.1. Hierarchy of infinite cardinals	23
3.2. A word on the continuum hypothesis	24
3.3. Cardinal arithmetic	25
3.4. Regularity and cofinality	29
Part 2. Model Theory	33
Chapter 4. First-order Logic	35
4.1. Structures	35
4.2. Languages	36
4.3. Some syntax	38
4.4. Truth and satisfaction	40
4.5. Elementary embeddings	43
4.6. Definable sets and parameters	45
4.7. Theories and their models	48
Chapter 5. Compactness and Consequences	53
5.1. A proof of compactness using ultraproducts	53
5.2. Applications of compactness, Löwenheim-Skolem, Vaught	56
5.3. An application to geometry: the Lefschetz Principle	59
Chapter 6. Quantifier Elimination	61

6.1. Preliminaries on substructures	62
6.2. A criterion for quantifier elimination	63
6.3. Model completeness	67
Part 3. More Model Theory	69
Chapter 7. Types and Saturation	71
7.1. Types	71
7.2. Some examples	74
7.3. Saturation and homogeneity	77
Chapter 8. Totally Transcendental Theories	85
8.1. Morley Rank	85
8.2. Total transcendentality and ω -stability	91
8.3. The strongly minimal case	93
Chapter 9. Nonforking in Totally Transcendental Theories	99
9.1. First properties	99
9.2. Definability of types and stable formulas	101
9.3. Further properties of nonforking	105
Chapter 10. An Introduction to Totally Transcendental Groups	109
10.1. Definable groups	109
10.2. Descending chain condition and connected components	110
10.3. Generics	111
10.4. Totally transcendental fields	114

The set theory part of these lecture notes were originally based on those of Tuna Altinel (Université Lyon 1). The faithfulness to that text is variable; certain sections are essentially translations, while other topics are substantially reworked. For the model theory I have used Dave Marker's book "Model Theory: An Introduction" (Springer 2002) as a source for examples and ideas about exposition. I am grateful to both of these authors.

Part 1

Set Theory

CHAPTER 1

Ordinals

1.1. Well-orderings

DEFINITION 1.1. Suppose R is a binary relation on a set E . The pair (E, R) is called a *partially ordered set* (or *poset*) if R satisfies the following properties

Reflexivity: aRa for all $a \in E$

Antisymmetry: if aRb and bRa then $a = b$, for all $a, b \in E$

Transitivity: if aRb and bRc then aRc , for all $a, b, c \in E$.

A poset (E, R) is *linearly ordered* (or *totally ordered*) if in addition it satisfies

Linearity: for all $a, b \in E$, either aRb or bRa .

A linear ordering is *well-ordered* if every non-empty subset has a least element.

DEFINITION 1.2. A *strict partially ordered set* is a pair (E, R) where E is a set and R is a binary relation satisfying antisymmetry and transitivity, plus

Antireflexivity: $\neg(aRa)$ for all $a \in E$.

The notions of *strict linear order* and *strict well-order* are similarly defined by replacing reflexivity with antireflexivity.

Note that to any poset (E, R) we can canonically associate the strict poset (E, R_{\neq}) where $aR_{\neq}b$ if and only if aRb and $a \neq b$. It is not hard to see that this association preserves linearity and well-orderedness. Similarly, to any strict poset there is an associated poset.

We often use \leq to denote the ordering on a poset and $<$ to denote the corresponding strict ordering.

LEMMA 1.3. *Well-orderings are rigid; that is, the identity map is the only automorphism of a well-ordering.*

PROOF. Of course, by an automorphism of a poset (E, \leq) we mean a bijective map $f : E \rightarrow E$ such that $a \leq b$ if and only if $f(a) \leq f(b)$, for all $a, b \in E$. Suppose f is such an automorphism of a well-ordering (E, \leq) , and suppose, toward a contradiction, that $f \neq \text{id}$. Let $D := \{x \in E : f(x) \neq x\}$ be the (nonempty) set of points that are moved by f . Since $f(x) = x$ if and only if $x = f^{-1}(x)$, D is also the set of elements moved by f^{-1} . Now D must have a least element, say $a \in E$. Either $f(a) < a$ or $a < f(a)$. In the first case $f(a) \notin D$, so f^{-1} fixes $f(a)$, which means that $a = f^{-1}f(a) = f(a)$, contradicting that $a \in D$. In the second case, applying f^{-1} to both sides, we get that $f^{-1}(a) < a$ and hence $f^{-1}(a) \notin D$, which means it is fixed by f and so $a = ff^{-1}(a) = f^{-1}(a)$, again contradicting $a \in D$. \square

COROLLARY 1.4. *Suppose (E, \leq) and (F, \preceq) are isomorphic well-orderings. Then there exists a unique isomorphism from (E, \leq) to (F, \preceq) .*

PROOF. Of course, an isomorphism from (E, \leq) to (F, \preceq) is a bijection $f : E \rightarrow F$ such that $a \leq b$ if and only if $f(a) \preceq f(b)$. Suppose f and g are two such isomorphisms. Then $g^{-1} \circ f$ is an automorphism of (E, \leq) and hence, by Lemma 1.3, $g^{-1} \circ f = \text{id}$. So $f = g$. \square

LEMMA 1.5. *Suppose $(E, <)$ is a strict well-ordering. For any $b \in E$, $(E, <)$ and $(\{x \in E : x < b\}, <)$ are not isomorphic.*

PROOF. Suppose, toward a contradiction, that $f : E \rightarrow \{x \in E : x < b\}$ is an isomorphism. Let $D = \{x \in E : f(x) \neq x\}$ be the set of elements moved by f . Since $f(b) \neq b$ (as $b \notin \{x \in E : x < b\}$), D is nonempty. We finish as in the proof of Lemma 1.3 by considering a least element a of D , and showing that both $a < f(a)$ and $a > f(a)$ lead to a contradiction. \square

1.2. Zermelo-Fraenkel Axioms

The prototypical example of a well-ordering is the set of natural numbers. Our goal in this chapter is to develop an infinite analogue of natural number that will allow us to “count” beyond finite sets. We begin however by recasting the natural numbers themselves, together with their ordering, in purely set-theoretic terms. This is consistent with the point of view that considers only sets and the membership relation \in as basic notions from which everything else is developed. To this end we define

$$\begin{aligned} 0 &:= \emptyset \\ 1 &:= \{\emptyset\} = \{0\} \\ 2 &:= \{\emptyset, \{\emptyset\}\} = \{0, 1\} \\ 3 &:= \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\} = \{0, 1, 2\} \\ &\vdots \end{aligned}$$

so that $S(n) := n \cup \{n\}$ for all n . Here S denotes the successor function. Each of these sets is a (finite) strict well-ordering: you should check that the membership relation on each fixed n is antireflexive, antisymmetric, transitive, and linear. That n is well-ordered is clear as it is finite.

Before continuing let us pause to consider what set-theoretical axioms we have used to construct natural numbers. To start with we needed a set that contained no elements.

AXIOM 1 (Emptyset). *There is a set, denoted by \emptyset , which has no members.*

Next, given a set x we needed to consider the set $\{x\}$. It is convenient to do this using the following axiom.

AXIOM 2 (Pairset). *For any two sets x and y there is a set p with the property that $t \in p$ if and only if $t = x$ or $t = y$. This set p is usually denoted by $\{x, y\}$.*

Note that we get the singleton $\{x\}$ by using the pairset axiom with $x = y$. To be more precise, this yields the set $\{x, x\}$ which we recognise as being just $\{x\}$. In fact, that “recognition” is another axiom:

AXIOM 3 (Extensionality). *For any two sets x and y , $x = y$ if and only if x and y have exactly the same members.*

Next, to form $S(n)$ from n we needed to be able to take the union of two sets (namely n and $\{n\}$). This can be done using the pairset axiom to form $\{n, \{n\}\}$ and then taking the union as follows.

AXIOM 4 (Unionset). *For any set x there exists a set denoted by $\bigcup x$ whose members are exactly the members of the members of x .*

With the above axioms in place we have the existence of each natural number as a set. However we do not yet have the existence of the set of all natural numbers; indeed, our axioms so far do not ensure the existence of any infinite set. One might hope to simply add an axiom asserting its existence, something like “there exists a set whose members are the natural numbers”. The problem is that the property “whose members are the natural numbers” is not clearly expressible in a finitary way. In order to make this objection more precise, I need to define a *definite condition*. They are conditions on sets that are built up in a finitary way from assertions about membership and equality using the usual logical operations. More precisely, they are conditions that can be expressed by a *first-order formula in the language of set theory*. We will be studying first-order formulas (in arbitrary languages) in the model theory part of this course. For now let us give a quick and informal definition:

DEFINITION 1.6 (Definite conditions and operations). The two basic types of definite conditions are $x \in y$ and $x = y$, where x and y are either sets or indeterminates standing for sets. Given definite conditions P and Q , the following are also definite conditions:

- “not P ”, denoted by $\neg P$
- “ P and Q ”, denoted by $P \wedge Q$
- “ P or Q ”, denoted by $P \vee Q$.

By combining the above connectives we see that “if P then Q ” (denoted by $P \rightarrow Q$) is also definite – it is equivalent to “(not P) or Q ”. If x is an indeterminate, then

- “for all x , P ”, denoted by $\forall xP$, and
- “there exists x , P ”, denoted by $\exists xP$,

are definite conditions. Only conditions obtained by the above operations in finitely many steps are definite. Given a definite condition P we sometimes write $P(x)$ to mean that “ x satisfies P ”. By a *definite operation* we mean an operation H on sets such that the condition $y = H(x)$ is definite.

We require that the defining property of the set whose existence our axiom is asserting be a definite condition. For example, the unionset axiom can be expressed as

$$\text{Given a set } x \text{ there exists a set } y \text{ satisfying } \forall z(z \in y \leftrightarrow \exists u(u \in x \wedge z \in u))$$

Try to write “there exists a set whose members are the natural numbers” in this way and you will see what goes wrong. We can get around this by noticing that the set of natural numbers could also be described as the smallest set containing 0 and closed under the successor function. Now, the successor function is definite since $y = S(x)$ is expressed as

$$\forall z(z \in y \leftrightarrow (z \in x \vee z = x))$$

and hence “containing 0 and closed under the successor function” is definite. So we may introduce:

AXIOM 5 (Infinity). *There exists a set I which contains 0 as well as the successor of each of its members; that is, if $x \in I$ then $x \cup \{x\} \in I$.*

But we still don't have the set of all natural numbers; if I is a set whose existence is asserted by the infinity axiom then certainly each natural number is in I , but there is no reason why I should not contain other elements. We could remedy this by considering

$$\bigcap \{J \subset I : 0 \text{ is in } J \text{ and } J \text{ is closed under successor}\}$$

This somehow captures our intuition about what “the set of natural numbers” is. But for the right hand side to exist we need more axioms: we need to be able to consider all subsets of a given set, we need to be able to pick out of a set the subset of elements satisfying some property, and we need to be able to take intersections of families of sets.

AXIOM 6 (Powerset). *For each set A there exists a set $\mathcal{P}(A)$ whose members are the subsets of A .*

AXIOM 7 (Separation). *Suppose P is a definite condition. For each set A there exists a set B whose members are exactly the members of A that satisfy P .*

EXERCISE 1.7. Express “whose members are the subsets of A ” and “whose members are exactly the members of A that satisfy P ” as a definite condition.

As it turns out, we can prove the existence of intersections without more axioms:

EXERCISE 1.8. Suppose that A is a nonempty set of sets. Using the above axioms show that there exists a set denoted by $\bigcap A$ such that $x \in \bigcap A$ if and only if $x \in y$ for all $y \in A$.

We finally have enough axiomatic power to define the set of all natural numbers.

DEFINITION 1.9. The *set of natural numbers* is the set

$$\omega := \bigcap \{J \in \mathcal{P}(I) : 0 \in J \wedge \forall x(x \in J \rightarrow S(x) \in J)\}$$

where I is any set given by the infinity axiom.

EXERCISE 1.10. Show that the above definition of ω does not depend on the choice of I .

There is one remaining axiom which we did not need to construct ω but whose usefulness we will see later.

AXIOM 8 (Replacement). *If H is a definite operation and A is a set then there exists a set B which is the image of A under H . That is, $y \in B$ if and only if there exists $x \in A$ such that $y = H(x)$.*

These eight axioms together form the axiom system we call *Zermelo-Fraenkel set theory* and denote by **ZF**. Note that all the axioms except extensionality assert the existence of sets. The emptyset and infinity axioms assert the existence of particular sets, while the others allow us to make new sets out of old ones.

1.3. Cartesian products and functions

Partly as an exercise in using the above axioms, and partly for future use, let us with some care discuss cartesian products of sets and functions between sets.

Given sets x, y , define the ordered pair

$$(x, y) := \{\{x\}, \{x, y\}\}$$

This exists by the pairset axiom, and is an element of $\mathcal{P}(\mathcal{P}(X) \cup \mathcal{P}(X \cup Y))$. The point of this definition is that $(x, y) = (x', y')$ if and only if $x = x'$ and $y = y'$. I leave it to you to check this.

Now, given sets X and Y there exists a set denoted by $X \times Y$ and called the *cartesian product of X and Y* whose elements are precisely the ordered pairs (x, y) where $x \in X$ and $y \in Y$. To see that $X \times Y$ exists, note first of all that “ $p = \{x, y\}$ ” is a definite condition, expressible by

$$\exists x, y(x \in p \wedge y \in p \wedge \forall z(z \in p \rightarrow z = x \vee z = y))$$

Hence

$$\text{Pairs}(X, Y) := \{p \in \mathcal{P}(X \cup Y) : \exists x, y(x \in X \wedge y \in Y \wedge p = \{x, y\})\}$$

exists by separation. Similarly

$$\text{Singletons}(X) := \{s \in \mathcal{P}(X) : s = \{x\}\}$$

exists. Hence $X \times Y := \text{Pairs}(\text{Singletons}(X), \text{Pairs}(X, Y))$ exists.

I also leave it to you to extend these definitions to longer (finite) ordered tuples and cartesian products.

Formally, given sets X and Y , by a *function from X to Y* we will mean a subset $\Gamma \subseteq X \times Y$ with the property that for all $x \in X$, there is a unique element $y \in Y$ such that $(x, y) \in \Gamma$. In practice, and more naturally, we will view a function as an operation on X , denoted by $f : X \rightarrow Y$, which assigns to $x \in X$ the unique element $f(x) \in Y$ such that $(x, f(x)) \in \Gamma$. Under this latter formulation, Γ is called the *graph of f* .

EXERCISE 1.11. A function is precisely a definite operation whose domain is a set.

1.4. Well-ordering the natural numbers

Our construction of the natural numbers has the following consequence: If $J \subseteq \omega$ contains 0 and is such that $x \in J$ implies $S(x) \in J$, then $J = \omega$. This is just by Definition 1.9. It is referred to as the *Induction Principle* and is very useful for proving statements about ω .

LEMMA 1.12. *Suppose n is a natural number.*

- (a) *Every element of n is a natural number. (So every element of ω is a subset of ω .)*
- (b) *Every element of n is a subset of n .*
- (c) *$n \notin n$*
- (d) *Either $n = 0$ or $0 \in n$.*
- (e) *If $y \in n$ then either $S(y) \in n$ or $S(y) = n$.*

PROOF. Let $J := \{n \in \omega : n \subseteq \omega\}$. Part (a) is claiming that $J = \omega$. Clearly $0 \in J$ as the empty set is a subset of every set. If $n \in J$ then $n \subseteq \omega$, and hence $S(n) = n \cup \{n\} \subseteq \omega$. So by the induction principle, $J = \omega$.

Next, let J be the set of all elements $n \in \omega$ satisfying the property that the elements of n are all subsets of n . Then $0 \in J$ vacuously. Suppose $n \in J$ and consider $S(n) = n \cup \{n\}$. We wish to show that every element of $S(n)$ is a subset of $S(n)$. This is certainly true of n which is both an element and a subset of $S(n)$. If x is any other element of $S(n)$, then $x \in n$, and hence by assumption $x \subseteq n$, so that $x \subseteq S(n)$. We have shown that $S(n) \in J$, and hence by the induction principle, $J = \omega$. This proves (b)

For (c), let $J = \{n \in \omega : n \notin n\}$. Then $0 \in J$. Suppose $n \in J$ but $S(n) \in S(n)$. This would mean that either $S(n) \in n$ or $S(n) = n$. The former would imply by part (b) that $S(n) \subseteq n$, and hence $n \in n$, which contradicts $n \in J$. The latter directly implies the contradiction $n \in n$. Hence it must be that $n \in J$ implies $S(n) \in J$, and so $J = \omega$.

For (d) we need to show that $\omega = \{0\} \cup \{n \in \omega : 0 \in n\}$. Letting J be the right-hand-side, we have $0 \in J$ explicitly. Suppose $n \in J$. If $n = 0$ then $0 \in S(n)$ and hence $S(n) \in J$. If $n \neq 0$ then $0 \in n$ by assumption, and hence as $n \subseteq S(n)$ we have $0 \in S(n)$ too. So again, $S(n) \in J$. Hence $J = \omega$, as desired

Finally, let J be the set of natural numbers satisfying (e). Vacuosly, $0 \in J$. Suppose $n \in J$. Consider $S(n) = n \cup \{n\}$, and $y \in S(n)$. If $y \in n$ then either $S(y) \in n$ or $S(y) = n$, by definition of J , and in either case we have $S(y) \in S(n)$. Otherwise, $y = n$ and so $S(y) = S(n)$. We have shown that $S(n)$ is in J , so that $J = \omega$. \square

PROPOSITION 1.13. (ω, \in) is a strict well-ordering. So is (n, \in) for every $n \in \omega$.

PROOF. Antireflexivity is Lemma 1.12(c). If $x, y \in \omega$ then, by 1.12(a), x and y are subsets of ω , so antisymmetry is clear. For transitivity, suppose $x, y, z \in \omega$, with $x \in y$ and $y \in z$. Then by 1.12(b) applied to z , we have $y \subseteq z$, and hence $x \in z$ as desired.

Next we prove linearity. Fix $x \in \omega$ and let $J = x \cup \{x\} \cup \{y \in \omega : x \in y\}$. We need to show that $J = \omega$. Note that if $x = 0$ then this is exactly what 1.12(d) says. So we may assume $x \neq 0$. It follows by 1.12(d) applied to x , that $0 \in x$, and hence $0 \in J$. Suppose $y \in J$ and consider the three possibilities for y given by the definition of J . If $y \in x$ then, by 1.12(e), either $S(y) \in x$ or $S(y) = x$. In either case we have $S(y) \in J$. If $y = x$ then $x \in S(y)$ and so again $S(y) \in J$. Finally, if $x \in y$ then $x \in S(y)$ and hence $S(y) \in J$. So $J = \omega$ by the induction principle.

It remains to prove that (ω, \in) is well-ordered. Suppose $X \subseteq \omega$ has no least element. Let J be the set of natural numbers x such that $S(x) \cap X = \emptyset$. Since 0 is a member of every nonzero natural number – that’s Lemma 1.12(d) – it is least in (ω, \in) and hence by assumption cannot be in X . It follows that $S(0) \cap X = \emptyset$ and so $0 \in J$. Now suppose $x \in J$. So $S(x) \cap X = \emptyset$. It follows that if $S(x)$ were in X it would have to be the least element of X – so $S(x) \notin X$. Hence $S(S(x)) \cap X = \emptyset$, showing that $S(x) \in J$. Therefore, $J = \omega$. So for every $x \in \omega$ we have $x \notin X$. That is, $X = \emptyset$.

That each (n, \in) is also well-ordered follows immediately as $n \subseteq \omega$ by Lemma 1.12(a), and hence (n, \in) is the induced ordering. \square

1.5. Introducing ordinals

Inspired by the results of the previous section, we generalise the natural numbers as follows:

DEFINITION 1.14. An *ordinal* is a set α satisfying the following two properties:

1. Every element of α is a subset of α .
2. The membership relation induces a strict well-ordering on α .

We equip an ordinal with the ordering induced by \in . So for $x, y \in \alpha$, $x < y$ will be synonymous with $x \in y$.

We saw in the previous section that ω , as well as each natural number, is an ordinal.

LEMMA 1.15. *If α and β are ordinals with $\alpha \subset \beta$ and $\alpha \neq \beta$, then $\alpha \in \beta$.*

PROOF. Let $D := \beta \setminus \alpha$ be the (nonempty) set of elements in β that are not in α , and let γ be the least element of D in β . (That D is a set uses the axiom of separation.) Note that as β is an ordinal, γ being a member of β implies that it is a subset of β also. We will show that $\gamma = \alpha$. This will imply in particular that $\alpha \in \beta$, as desired.

Suppose there exists an element $\delta \in \gamma \setminus \alpha$. But then $\delta \in D$ and $\delta < \gamma$, which contradicts the minimal choice of γ . So $\gamma \subseteq \alpha$.

For the converse, suppose $\delta \in \alpha$. By linearity of β , of which δ and γ are members, either $\gamma = \delta$ or $\gamma < \delta$ or $\delta < \gamma$. As γ is in D , $\gamma \notin \alpha$, and hence it cannot be that $\gamma = \delta$. Now as α is an ordinal, each member of α is a subset of α , and so $\delta \subseteq \alpha$. Hence $\gamma < \delta$ would also imply the impossible $\gamma \in \alpha$. Hence $\delta < \gamma$, i.e., $\delta \in \gamma$. We have shown that $\alpha \subseteq \gamma$, as desired. \square

- PROPOSITION 1.16.
- (a) *Every member of an ordinal is an ordinal.*
 - (b) *No ordinal is a member of itself.*
 - (c) *If α is an ordinal then so is its successor $S(\alpha) := \alpha \cup \{\alpha\}$.*
 - (d) *The intersection of two ordinals is an ordinal.*

PROOF. For part (a), let α be a member of an ordinal β . So $\alpha \subseteq \beta$ and hence the membership relation on α is just the membership relation on β restricted to α . Now it is not hard to see that every subset of a strict well-ordering is again a strict well-ordering. Hence \in induces a strict well-ordering on α . So it remains to show that every member of α is a subset of α . Suppose $x \in \alpha$ and $y \in x$. Since $\alpha \subseteq \beta$ we have that $x \in \beta$, and hence, using again that β is an ordinal, we have that $x \subseteq \beta$. So $y \in \beta$. We have that $x, y, \alpha \in \beta$, $y < x$, and $x < \alpha$. Hence by transitivity (in β) we have that $y < \alpha$, i.e., $y \in \alpha$. We have shown that $x \subseteq \alpha$, as desired.

For part (b) we point out that for α an ordinal, $\alpha \notin \alpha$. First, by antireflexivity of \in on α we have that for all $x \in \alpha$, $x \notin x$. Hence, if $\alpha \in \alpha$, then $\alpha \notin \alpha$, as desired!

We leave parts (c) and part (d) as exercises. \square

- PROPOSITION 1.17.
- (a) *If α and β are ordinals then either $\alpha \in \beta$ or $\alpha = \beta$ or $\beta \in \alpha$.*
 - (b) *Every set of ordinals is strictly well-ordered by \in .*
 - (c) *If E is a set of ordinals then its supremum $\sup E := \bigcup E$ is an ordinal.*
 - (d) *There does not exist a set whose elements are all the ordinals.*

PROOF. If α and β are ordinals then $\alpha \cap \beta$ is an ordinal by Proposition 1.16(d), and it is a subset of both α and β . Hence by Lemma 1.15, if it is neither α nor β , then $\alpha \cap \beta \in \alpha$ and

$\alpha \cap \beta \in \beta$. That is, $\alpha \cap \beta \in \alpha \cap \beta$, which contradicts the antireflexivity of \in on α (of which $\alpha \cap \beta$ is a member). Hence, either $\alpha \cap \beta = \alpha$ or $\alpha \cap \beta = \beta$. In the first case we get $\alpha \subseteq \beta$, which by Lemma 1.15 again, implies that either $\alpha \in \beta$ or $\alpha = \beta$. Similarly, if $\alpha \cap \beta = \beta$ then either $\beta \in \alpha$ or $\beta = \alpha$. This proves part (a).

For part (b), suppose that E is a set of ordinals.

- Antireflexivity: This is Proposition 1.16(b).
- Antisymmetry: Suppose $\alpha, \beta \in E$. Note that under antireflexivity, antisymmetry is equivalent to saying that it is not the case that both $\alpha \in \beta$ and $\beta \in \alpha$. Indeed, if this were the case, then $\alpha \subseteq \beta$ (as every member of β is a subset of β) and so $\beta \in \beta$ which we know is impossible.
- Transitivity: Suppose $\alpha, \beta, \gamma \in E$ and $\alpha \in \beta$ and $\beta \in \gamma$. Then $\beta \subseteq \gamma$ (as every member of γ is a subset of γ) and so $\alpha \in \gamma$, as desired.
- Linearity: This is just part (a) of this Proposition.
- Well-orderedness: Suppose A is a nonempty subset of E . Let $\alpha \in A$. If $\alpha \cap A = \emptyset$ then α is the least element of A . Otherwise $\alpha \cap A$ is a nonempty subset of α and hence it contains a least element β in α . If $\gamma \in \beta \cap A$, then by transitivity of \in on α , $\gamma \in \alpha \cap A$ and $\gamma < \beta$, contradicting the minimal choice of β . So $\beta \cap A = \emptyset$, which implies that β is the least element of A .

This shows that every set of ordinals is strictly well-ordered by membership.

Now, for (c), suppose that E is a set of ordinals and consider $\sup E$. By Proposition 1.16(a), $\sup E$ is itself a set of ordinals and hence strictly well-ordered by \in . So it remains to show that if $\alpha \in \sup E$ and $\beta \in \alpha$, then $\beta \in \sup E$. But, by definition of $\sup E$, there exists a $\gamma \in E$ with $\alpha \in \gamma$. Now by transitivity, $\beta \in \gamma$, and hence $\beta \in \sup E$.

For part (d), suppose that E is a set consisting of all ordinals. Then by part (b), membership induces a strict well-ordering on E . Moreover, Proposition 1.16(a) tells us that every member of E is a subset of E . Hence E is itself an ordinal. But then $E \in E$, which contradicts E being an ordinal. \square

Note that 1.17(d) is not inconsistent with the axiom of separation. While being an ordinal is a definite condition, it cannot be used with the axiom of separation to obtain “the set of all ordinals”, since that axiom only applies when restricted to a set. So, if we start with a set A , then we can, using separation, form the subset of all elements *in* A that are ordinals. In fact, Proposition 1.17(d) shows that we cannot drop this boundedness restriction on the axiom of separation without running into “paradoxes”.

Despite 1.17(b), not every set of ordinals is itself an ordinal: $\{1\}$ is not an ordinal because $1 = \{\emptyset\}$ is not a subset of $\{1\} = \{\{\emptyset\}\}$.

NOTATION 1.18. As is somewhat justified by Proposition 1.17, from now on, given ordinals α and β , the notation $\alpha < \beta$ will be synonymous with $\alpha \in \beta$.

- LEMMA 1.19. (a) *There is no ordinal strictly between an ordinal and its successor.*
 (b) *If E is a nonempty set of ordinals then $\sup E$ is the least upper bound for E .*
 (c) *If E is a set of ordinals then there is a least ordinal not in E .*

PROOF. Suppose $\alpha \leq \gamma \leq S(\alpha)$. If $\gamma \neq S(\alpha)$ then $\gamma \in S(\alpha) = \alpha \cup \{\alpha\}$. Since $\alpha \leq \gamma$, it cannot be that $\gamma \in \alpha$ (by antireflexivity and transitivity). Hence $\gamma \in \{\alpha\}$, that is, $\gamma = \alpha$.

To see that $\sup E$ is an upper bound for E it suffices, by the totality of the ordering of ordinals (Proposition 1.17(a)), to show that no element of E is greater than $\sup E$. But $\alpha \in E$ and $\sup E \in \alpha$ would imply the impossible $\sup E \in \sup E$. Now suppose $\alpha < \sup E$. Then there is a $\beta \in E$ with $\alpha \in \beta$. Hence $\alpha < \beta$ and so α is not an upper bound of E . This shows that $\sup E$ is the least upper bound of E .

Finally we prove (c). Suppose E is a set of ordinals. Let $\alpha := S(S(\sup E))$. Note that $E \subseteq \alpha$. Now, for any ordinal β , $\sup S(\beta) = \beta$, by part (b) for example. Hence if $E = \alpha$ then $\sup E = S(\sup E)$ which is a contradiction. So $E \subsetneq \alpha$. Let μ be least in $\alpha \setminus E$. If $\beta \notin \alpha$ then $\beta \geq \alpha > \mu$. If $\beta \in \alpha \setminus E$ then $\beta \geq \mu$ by choice. So μ is the least ordinal not in E . \square

DEFINITION 1.20. A *successor ordinal* is an ordinal of the form $S(\alpha)$ for some ordinal α . A *limit ordinal* is an ordinal that is not a successor.

EXERCISE 1.21. Show that ω is a limit ordinal.

1.6. Transfinite induction and recursion

THEOREM 1.22 (Transfinite Induction). *Suppose P is a definite condition satisfying the following property: if α is an ordinal and $P(\beta)$ is true of all $\beta < \alpha$, then $P(\alpha)$ is true. Then P is true of all ordinals.*

PROOF. Note that, vacuously, the hypothesis of the Theorem implies that $P(0)$ is true. Suppose, toward a contradiction, that there exists an ordinal α such that $P(\alpha)$ is false. Then $D := \{\beta \leq \alpha : P(\beta) \text{ is false}\}$ is a nonempty set of ordinals. (We use separation to see that it is a set at all.) Let α_0 be its least element, which exists by Proposition 1.17(b). So for every $\beta < \alpha_0$, $P(\beta)$ is true. Hence $P(\alpha_0)$ is true, which contradicts α_0 being in D . \square

The following corollary is just another form of transfinite induction that is often useful. Its proof is left as an exercise.

COROLLARY 1.23 (Transfinite Induction – Second Form). *Suppose P is a definite condition satisfying the following:*

1. $P(0)$ is true.
2. For all ordinals α , $P(\alpha)$ implies $P(S(\alpha))$.
3. For all limit ordinals $\alpha > 0$, if $P(\beta)$ is true for all $\beta < \alpha$ then $P(\alpha)$ is true.

Then P is true of all ordinals.

While functions always have set domains, an *operation* (even a definite one) may be defined on “larger” classes of objects. For example, one can have an operation defined on all ordinals, like the successor operation (which is definite, see page 5). If F is an operation on ordinals and α is an ordinal, then $F \upharpoonright \alpha$ denotes the restriction of F to α .

THEOREM 1.24 (Transfinite Recursion). *Suppose G is a definite operation which takes as input any function whose domain is an ordinal. Then there is a unique definite operation on ordinals satisfying $F(\alpha) = G(F \upharpoonright \alpha)$ for all ordinals α .*

PROOF. First of all, to prove that this condition uniquely determines F is a straightforward transfinite induction that we leave to the reader.

To prove the existence of F , we begin by introducing some terminology: we say that a function t with domain an ordinal α is an α -function defined by G if for all $\beta < \alpha$, $t(\beta) = G(t \upharpoonright \beta)$. So such functions are approximations to the operation F that we are after. The first thing to observe (using transfinite induction) is that an α -function defined by G is necessarily unique, and that for all $\beta \leq \alpha$, the α -function defined by G extends the β -function defined by G . That they exist at all is the following claim.

CLAIM 1.25. *For each ordinal α there exists an α -function defined by G .*

PROOF OF CLAIM 1.25. We use (the second form of) transfinite induction. For $\alpha = 0$, the empty function vacuously satisfies the condition. Suppose that $\alpha \neq 0$ and that there is an α -function t_α defined by G . Then define

$$t := t_\alpha \cup \{(\alpha, G(t_\alpha))\}$$

where the union here is really by identifying the function with its graph. Clearly t is a function whose domain is $S(\alpha)$. To see that it is the $S(\alpha)$ -function defined by G , note that $t(\alpha) = G(t_\alpha) = G(t \upharpoonright \alpha)$ while for $\beta < \alpha$ we have $t(\beta) = t_\alpha(\beta) = G(t_\alpha \upharpoonright \beta) = G(t \upharpoonright \beta)$.

Finally, suppose $\alpha \neq 0$ is a limit ordinal and that for each $\beta < \alpha$ there is a β -function, say t_β , that is defined by G . By the proof of Claim 1.26 below, the operation $\beta \mapsto t_\beta$ is definite, and hence by replacement $T := \{t_\beta : \beta < \alpha\}$ is a set. (This is our first use of replacement.) Let $t = \bigcup T$. To verify that t is the α -function defined by G we need to verify that:

- t is a function on α . This follows from the above observation that the t_β s form a chain of extensions.
- t is the α -function defined by G : for all $\beta < \alpha$, note that $S(\beta) < \alpha$ as α is a limit ordinal, and $t(\beta) = t_{S(\beta)}(\beta) = G(t_{S(\beta)} \upharpoonright \beta) = G(t \upharpoonright \beta)$.

This proves Claim 1.25. □

Let us denote by t_α the α -function defined by G .

CLAIM 1.26. *The operation $\alpha \mapsto t_\alpha$ is definite.*

PROOF OF CLAIM 1.26. Indeed, $y = t_\alpha$ if and only if y is a function whose domain is α and for all $x < \alpha$, $y(x) = G(y \upharpoonright x)$. In order to verify that this is a definite condition it suffices to prove that being a function with domain α is definite, and that given a function y the operation $x \mapsto y \upharpoonright x$ is definite. Both are straightforward and left to the reader. □

Define $F(\alpha) := t_{S(\alpha)}(\alpha)$, for all ordinals α . By Claim 1.26 and the fact that the successor operation is definite, we see that F is a definite operation. Note that for any $\beta < \alpha$, since $t_{S(\alpha)}$ extends $t_{S(\beta)}$ we have that $t_{S(\alpha)}(\beta) = t_{S(\beta)}(\beta) = F(\beta)$. That is, $F \upharpoonright \alpha = t_{S(\alpha)} \upharpoonright \alpha$ for all ordinals α . Hence,

$$F(\alpha) = t_{S(\alpha)}(\alpha) = G(t_{S(\alpha)} \upharpoonright \alpha) = G(F \upharpoonright \alpha)$$

as desired. □

The following two other versions of transfinite recursion are left as exercises.

COROLLARY 1.27 (Transfinite Recursion – Second Form). *Suppose G_1 is a set, G_2 is an operation on sets, and G_3 is an operation on partial functions on ordinals. There exists a unique operation F on ordinals, such that*

1. $F(0) = G_1$,
2. $F(S(\alpha)) = G_2(F(\alpha))$ for all ordinals α ,
3. $F(\alpha) = G_3(F \upharpoonright \alpha)$ for all limit ordinals α .

THEOREM 1.28 (Parametric Transfinite Recursion). *Suppose G is a definite binary operation. Then there exists a unique definite binary operation F such that*

$$F(z, \alpha) = G(z, F(z, -) \upharpoonright \alpha)$$

for all ordinals z and α .

1.7. Ordinal arithmetic

In Section 1.2 we constructed the natural numbers. We can continue that construction:

$$\begin{aligned} \omega + 1 &:= S(\omega) = \omega \cup \{\omega\} \\ \omega + 2 &:= S(\omega + 1) = \omega \cup \{\omega\} \cup \{\omega \cup \{\omega\}\} \\ &\vdots \\ \omega \cdot 2 &:= \sup\{\omega + n : n \in \omega\} = \omega \cup \{\omega, \omega + 1, \omega + 2, \dots\} \\ \omega \cdot 3 &:= \sup\{\omega \cdot 2 + n : n \in \omega\} = \omega \cdot 2 \cup \{\omega \cdot 2, \omega \cdot 2 + 1, \omega \cdot 2 + 2, \dots\} \\ &\vdots \\ \omega^2 &:= \sup\{\omega \cdot n : n \in \omega\} = \{\omega \cdot n + m : n, m \in \omega\} \\ &\vdots \end{aligned}$$

To see how $\omega \cdot 2$, for example, is obtained, prove first that $n \mapsto \omega + n$ is definite by transfinite recursion, conclude that $\{\omega + n : n \in \omega\}$ is a set by the replacement axiom, and then apply 1.17(c). The arithmetic operations suggested by the above notation can be extended to all ordinals, and that is the purpose of this section.

DEFINITION 1.29 (Ordinal addition). For all ordinals β we define the definite operation $\beta + \alpha$ recursively as follows:

$$\begin{aligned} \beta + 0 &= \beta \\ \beta + S(\alpha) &= S(\beta + \alpha) \text{ for all ordinals } \alpha \\ \beta + \alpha &= \sup\{\beta + \gamma : \gamma < \alpha\} \text{ for all limit ordinals } \alpha \neq 0. \end{aligned}$$

To see that the above definition is well-founded one has to express it as an instance of the recursion theorem. This can be done as follows: Fix an ordinal $G_1 := \beta$, let G_2 be the successor operation, and let $G_3(y) = \sup(\text{Im } y)$ be the operation which takes as input functions with ordinal domains. Then the definite operation $F(x)$ determined by G_1, G_2, G_3 according to the second form of the transfinite recursion theorem (cf. 1.27) is exactly the operation $x \mapsto \beta + x$.

EXERCISE 1.30. Show that $(\beta, \alpha) \mapsto \beta + \alpha$ is a definite *binary* operation on ordinals.

Note that $\alpha + 1 = S(\alpha)$, and we will from now on tend to use the former notation.

Note that ordinal addition is not commutative: $1 + \omega = \omega \neq \omega + 1$.

DEFINITION 1.31 (Ordinal product). For all ordinals β we define the definite operation $\beta \cdot \alpha$ recursively as follows:

$$\begin{aligned}\beta \cdot 0 &= 0 \\ \beta \cdot S(\alpha) &= \beta \cdot \alpha + \beta \text{ for all ordinals } \alpha \\ \beta \cdot \alpha &= \sup\{\beta \cdot \gamma : \gamma < \alpha\} \text{ for all limit ordinals } \alpha \neq 0.\end{aligned}$$

DEFINITION 1.32 (Ordinal exponentiation). For all ordinals β we define the definite operation β^α recursively as follows:

$$\begin{aligned}\beta^0 &= 1 \\ \beta^{S(\alpha)} &= \beta^\alpha \cdot \beta \text{ for all ordinals } \alpha \\ \beta^\alpha &= \sup\{\beta^\gamma : \gamma < \alpha\} \text{ for all limit ordinals } \alpha \neq 0.\end{aligned}$$

We leave it to you to express these precisely by transfinite recursion. Note that ordinal product is also not commutative as $2 \cdot \omega = \omega \neq \omega \cdot 2$.

Here are some basic properties of this arithmetic:

PROPOSITION 1.33. *Suppose α, β, δ are ordinals.*

- (a) $\alpha < \beta$ iff $\delta + \alpha < \delta + \beta$
- (b) $\alpha = \beta$ iff $\delta + \alpha = \delta + \beta$
- (c) $(\alpha + \beta) + \delta = \alpha + (\beta + \delta)$
- (d) For $\delta \neq 0$, $\alpha < \beta$ iff $\delta \cdot \alpha < \delta \cdot \beta$
- (e) For $\delta \neq 0$, $\alpha = \beta$ iff $\delta \cdot \alpha = \delta \cdot \beta$
- (f) $(\alpha \cdot \beta) \cdot \delta = \alpha \cdot (\beta \cdot \delta)$

PROOF. These are essentially an exercise in transfinite induction; we prove some and leave others to you.

Fix α and δ and consider the (definite) property $P(x)$ which says that “if $\alpha < x$ the $\delta + \alpha < \delta + x$ ”. Now $P(0)$ is vacuously correct. Suppose $P(\beta)$ is correct and consider $P(\beta + 1)$. If $\alpha < \beta + 1$ then either $\alpha < \beta$ or $\alpha = \beta$. In the former case we have $\delta + \alpha < \delta + \beta$ by the truth of $P(\beta)$, and since $\delta + \beta < \delta + \beta + 1$, we get that $\delta + \alpha < \delta + \beta + 1$. In the latter case, $\delta + \alpha = \delta + \beta < \delta + \beta + 1$ also. So $P(\beta + 1)$ is true. If β is a limit ordinal and $\alpha < \beta$, then $\delta + \alpha < \sup\{\delta + \gamma : \gamma < \beta\} = \delta + \beta$ where the first inequality is by the fact that $\delta + \alpha \in \sup\{\delta + \gamma : \gamma < \beta\}$ since $\delta + \alpha \in S(\delta + \alpha) = \delta + S(\alpha)$ and $S(\alpha) < \beta$. Hence $P(\beta)$ is true. By the second form of transfinite induction, we have $P(\beta)$ for all ordinals β , which is the left-to-right implication of part (a).

For the converse, suppose $\delta + \alpha < \delta + \beta$. Then $\alpha \neq \beta$ by the well-definedness of ordinal addition, and $\neg(\beta < \alpha)$ by the left-to-right implication proved above. So it must be that $\alpha < \beta$. This completes the roof of part (a).

The left-to-right implication of part (b) is just the well-definedness of ordinal addition. For the converse, suppose $\delta + \alpha = \delta + \beta$. Then part (a) rules out $\alpha < \beta$ and $\beta < \alpha$, hence forcing $\alpha = \beta$.

Toward part (c), fix α and β and consider the condition $P(x)$ which says that $(\alpha + \beta) + x = \alpha + (\beta + x)$. It is clear that $P(0)$ is true. Suppose $P(\delta)$ is true. Then

$$\begin{aligned} (\alpha + \beta) + (\delta + 1) &= ((\alpha + \beta) + \delta) + 1 \text{ by the definition of ordinal addition} \\ &= (\alpha + (\beta + \delta)) + 1 \text{ by the truth of } P(\delta) \\ &= \alpha + ((\beta + \delta) + 1) \text{ by the definition of ordinal addition} \\ &= \alpha + (\beta + (\delta + 1)) \text{ by the definition of ordinal addition} \end{aligned}$$

So $P(\delta)$ implies $P(\delta + 1)$.

Now suppose that δ is a limit ordinal.

CLAIM 1.34. $\sup\{\alpha + (\beta + \gamma) : \gamma < \delta\} = \sup\{\alpha + \zeta : \zeta < \beta + \delta\}$

PROOF OF CLAIM 1.34. Indeed the left-to-right containment is just by part (a). For the converse suppose $x \in \alpha + \zeta$ for some $\zeta < \beta + \delta$. As δ is a limit ordinal, $\beta + \delta = \sup\{\beta + \gamma : \gamma < \delta\}$. Hence $\zeta < \beta + \gamma$ for some $\gamma < \delta$. So $\alpha + \zeta < \alpha + (\beta + \gamma)$ by part (a), and by transitivity $x \in \alpha + (\beta + \gamma)$ as desired \square

CLAIM 1.35. $\beta + \delta$ is again a limit ordinal.

PROOF OF CLAIM 1.35. Assume toward a contradiction that for some ordinal ξ , $S(\xi) = \beta + \delta$. As $\xi < S(\xi)$, and δ is a limit ordinal, we get that $\xi < \beta + \gamma$ for some $\gamma < \delta$. But then $\beta + \delta = S(\xi) < S(\beta + \gamma) = \beta + S(\gamma)$. So by part (a), $\delta < S(\gamma)$. But $\gamma < \delta$ by choice, which is a contradiction (cf. Lemma 1.19(a)). \square

If we now assume that $P(\gamma)$ is true of all $\gamma < \delta$, then we get

$$\begin{aligned} \alpha + (\beta + \delta) &= \sup\{\alpha + \zeta : \zeta < \beta + \delta\} \text{ by Claim 1.35} \\ &= \sup\{\alpha + (\beta + \gamma) : \gamma < \delta\} \text{ by Claim 1.34} \\ &= \sup\{(\alpha + \beta) + \gamma : \gamma < \delta\} \text{ by the truth of } P(\gamma) \\ &= (\alpha + \beta) + \delta \text{ as } \delta \text{ is a limit ordinal.} \end{aligned}$$

So $P(\delta)$ is true. This completes the proof of part (c).

Parts (d)–(f) are proved in a similar fashion, and left as an exercise. \square

1.8. Well-orderings and ordinals

At this point one might have the impression that ordinals are a very special class of well-orderings. In fact the opposite is true: ordinals are all there is.

THEOREM 1.36. *Every strict well-ordering is isomorphic to an ordinal. In fact, if (E, \prec) is a strict well-ordering, then there exists a unique ordinal α and a unique isomorphism between (E, \prec) and $(\alpha, <)$.*

PROOF. The fact that if there exists such an isomorphism between (E, \prec) and an ordinal α , then it must be unique, is just Corollary 1.4.

Suppose (E, \prec) is isomorphic to ordinals $\alpha < \beta$. Then $(\alpha, <) = (\{\gamma \in \beta : \gamma < \alpha\}, <)$. But, by Lemma 1.5, there can be no isomorphism between this latter well-ordering and $(\beta, <)$, which contradicts that both are isomorphic to (E, \prec) . Hence $\alpha = \beta$.

So it remains to find an ordinal α isomorphic to (E, \prec) . We may assume that E is nonempty, else the desired ordinal is 0. Given $x \in E$ let us denote by E_x the initial segment $\{e \in E : e \prec x\}$. Consider

$$A = \{x \in E : (E_x, \prec) \text{ is isomorphic to an ordinal}\}$$

In order to verify that such a set A exists one would need to show that the displayed condition is definite and then use separation. We leave that to the reader. Note that A is not empty since if $x \in E$ is the least element of E , then E_x is isomorphic to 0. By the uniqueness proved above we can consider the function f on A where $f(x)$ is the unique ordinal isomorphic to (E_x, \prec) . By replacement $\text{Im}(f)$ is a set of ordinals. Let α be the least ordinal which is not in $\text{Im}(f)$; this exists by Lemma 1.19(c). We prove that $f : A \rightarrow \text{Im}(f)$ is in fact an isomorphism between (E, \prec) and α .

1. f preserves the ordering: in fact, if $x, y \in E$ with $y \in A$ and $x \prec y$, then $x \in A$ and $f(x) < f(y)$. Indeed, since $x \prec y$, E_x is an initial segment of E_y . Now, if h is the isomorphism between (E_y, \prec) and $f(y)$, then

$$h(E_x) = \{\alpha \in f(y) : \alpha < h(x)\} = \{\alpha \in f(y) : \alpha \in h(x)\} = h(x)$$

So h restricts to an isomorphism between E_x and the ordinal $h(x)$. It follows that $x \in A$, and by uniqueness that $h(x) = f(x)$. So $f(x) < f(y)$.

2. $\alpha = \text{Im}(f)$. Suppose $\beta \in \alpha$. Then by choice of α , $\beta \in \text{Im}(f)$. Conversely, suppose $\beta \in \text{Im}(f)$, and h is the isomorphism between E_x and β for some $x \in A$. Then $\beta \neq \alpha$. If $\alpha < \beta$, then $\alpha = h(y)$ for some $y \prec x$. By the proof of (1) above, $y \in A$ and $f(y) = h(y) = \alpha$, which contradicts $\alpha \notin \text{Im}(f)$. Hence $\beta < \alpha$, as desired.
3. f is injective. Suppose $f(x) = f(y)$. If $x \prec y$ then E_x is a proper initial segment of E_y that is isomorphic to E_y – but this is forbidden by Lemma 1.5. Similarly, $y \prec x$ is impossible. Hence $x = y$.
4. $A = E$. Toward a contradiction assume that $E \setminus A$ is nonempty and let x be a least element of this set. By (1) above, no element greater than x is in A . On the other hand, by minimality, every element less than x is in A . That is, $A = E_x$. Since we have already proved that f is an isomorphism between (A, \prec) and the ordinal α , this implies that $x \in A$, a contradiction.

This proves that f is an isomorphism between (E, \prec) and $(\alpha, <)$, as desired. □

CHAPTER 2

Axiom of Choice

2.1. A first look at cardinals

The ordinals were introduced as a transfinite generalisation of the natural numbers. One aspect of the natural numbers is that they can be used to measure the size of finite sets; a set is of size n iff there exists a bijection between it and the finite ordinal n . Infinite ordinals do not serve the same purpose simply because there are many distinct ordinals that are of the “same size” in this sense – for example there is a bijection between ω and $\omega + 1$. Let us first formalise this notion of “same size”.

DEFINITION 2.1. Two sets A and B are said to be *equinumerous* if there exists a bijection between them.

Let us record for future use the following well-known fact.

PROPOSITION 2.2 (Schröder-Bernstein). *Sets A and B are equinumerous if and only if there are injective functions from A to B and from B to A .*

PROOF. Only the right-to-left implication requires proof. We begin with a claim.

CLAIM 2.3. *Suppose X is a set and $f : X \rightarrow X$ is injective. Then for any $Y \subseteq X$ with $f(X) \subseteq Y$, Y is equinumerous with X .*

PROOF OF CLAIM 2.3. Let $X_n := f^n(X)$, $Y_n := f^n(Y)$, $Z_n := X_n \setminus Y_n$. Let $Z := \bigcup_{n \geq 0} Z_n$ and $W := X \setminus Z$. Since $Z_0 = X \setminus Y$ and $f(X) \subseteq Y$, it is clear that $f(Z) \cup W \subseteq Y$. For the converse consider $y \in Y \setminus W$. Then $y \in Z$. But $y \notin Z_0$, so $y \in Z_n$ for some $n > 0$. That is, $y = f^n(x)$ for some $x \in X$. Since $y \notin Y_n$, $f^{n-1}(x) \in X_{n-1} \setminus Y_{n-1} = Z_{n-1} \subseteq Z$. Hence $y \in f(Z)$. We have shown that $Y = f(Z) \cup W$. Hence, $g : X \rightarrow Y$ defined by

$$g(x) = \begin{cases} f(x) & \text{if } x \in Z \\ x & \text{if } x \in W \end{cases}$$

is clearly surjective. It is also injective because it is injective on Z and on W , and $g(Z) = f(Z) \subseteq Z$ is disjoint from $g(W) = W = X \setminus Z$. This proves Claim 2.3. \square

Now suppose we have injective maps $i : A \rightarrow B$ and $j : B \rightarrow A$. Apply the claim to $X = A$, $Y = j(B)$, and $f = j \circ i$. This tells us that A and $j(B)$ are equinumerous. But of course B and $j(B)$ are equinumerous. Hence A and B are equinumerous. \square

LEMMA 2.4. *For every infinite ordinal α , α and $\alpha + 1$ are equinumerous.*

PROOF. Define $f : \alpha + 1 \rightarrow \alpha$ as follows:

$$f(x) = \begin{cases} x + 1 & \text{if } x \in \omega \\ 0 & \text{if } x = \alpha \\ x & \text{otherwise} \end{cases}$$

Since α is infinite, $\alpha \notin \omega$, and hence f is well-defined on all of $\alpha + 1$. It is surjective since its image clearly contains ω and if $x \in \alpha \setminus \omega$ then $f(x) = x$. It is also clearly injective. \square

This problem of equinumerous distinct infinite ordinals can be overcome by choosing canonically one ordinal from each collection of equinumerous ordinals:

DEFINITION 2.5 (Cardinals). An ordinal α is called a *cardinal* if α is not equinumerous to any ordinal $\beta < \alpha$.

Note that by Lemma 2.4 every infinite cardinal is a limit ordinal.

For example, every finite ordinal is a cardinal. It is also clear that ω is a cardinal. Are there other cardinals? The answer is yes.

PROPOSITION 2.6. For every set E there exists a unique cardinal $h(E)$ which is the least ordinal that is not equinumerous with any subset of E .

PROOF. It suffices to prove the existence of such an ordinal $h(E)$; it will be a cardinal as all ordinals less than it are equinumerous with subsets of E but it is not. To prove the existence of $h(E)$ we need to show that the collection of ordinals that are equinumerous with a subset of E form a set. Consider

$$W := \{(A, \prec) : A \subseteq E, \prec \text{ is a well-ordering of } A\}.$$

Note that W is a nonempty subset of $\mathcal{P}(E) \times \mathcal{P}(E \times E)$. By Theorem 1.36 we have an operation F on W which assigns to each member of W the unique ordinal that is order-isomorphic to it. It is not hard to see that F is definite. So $\text{Im}(F)$ is a set of ordinals by replacement. But

$$\alpha \in \text{Im}(F) \iff \alpha \text{ is an ordinal equinumerous with a subset of } E.$$

Hence the least ordinal not in $\text{Im}(F)$, which exists by 1.19(c), is our desired $h(E)$. \square

In particular, there exist uncountable cardinals, for example $h(\omega)$.

But we want much more of our cardinals. If they are to measure the size of arbitrary sets then every set must be equinumerous with a cardinal. But that would in particular imply that every set is well-orderable, something that cannot be proved from the axioms we have so far. We need the axiom of choice.

2.2. The axiom of choice and its equivalents

What makes the axiom of choice different from the other axioms we have used so far is that it asserts the existence of sets that are not “definable” from \in and $=$. As a result, the axiom of choice has many surprising consequences. Nevertheless, this axiom is indispensable to much of contemporary mathematics.

DEFINITION 2.7 (Choice functions). Suppose \mathcal{F} is a set of sets. A *choice function* on \mathcal{F} is a function from \mathcal{F} to $\bigcup \mathcal{F}$ which assigns to each $F \in \mathcal{F}$ a member of F .

AXIOM 9 (Choice). *Every set of nonempty sets has a choice function.*

THEOREM 2.8. *The following are equivalent:*

1. *The axiom of choice.*
2. *The well-ordering principle: every set is well-orderable.*
3. *Zorn's lemma: If (E, \prec) is a nonempty strict poset with the property that every totally ordered subset of E has an upper bound in E , then E has a maximal element.*

PROOF. Let us assume the axiom of choice and derive from it the well-ordering principle. Let A be an arbitrary set and let c be a choice function on $\mathcal{P}(A) \setminus \{\emptyset\}$. Consider the operation on ordinals defined (using transfinite recursion) as follows:

$$F(\alpha) = \begin{cases} c(A \setminus \text{Im}(F \upharpoonright \alpha)) & \text{if } A \setminus \text{Im}(F \upharpoonright \alpha) \neq \emptyset \\ \zeta & \text{else} \end{cases}$$

where ζ is some fixed ordinal not in A that is used as a default value. To show that A is well-ordered it will suffice to show that

- F halts, in the sense that it takes on the value ζ at some ordinal, and hence at every future ordinal.
- If α is the least ordinal at which F halts, then $F \upharpoonright \alpha$ is a bijection onto A .

To see that F halts, let $h(A)$ be the cardinal whose existence is guaranteed by Proposition 2.6; so $h(A)$ is the least ordinal which is not equinumerous with any subset of A . Given ordinals $\alpha < \beta$ in $h(A)$, note that if $F(\beta) \neq \zeta$ then $F(\beta) = c(A \setminus \text{Im}(F \upharpoonright \beta))$ and so $F(\beta) \notin \text{Im}(F \upharpoonright \beta)$. But $F(\alpha) \in \text{Im}(F \upharpoonright \beta)$. Hence $F(\alpha) \neq F(\beta)$. So, if $F(\beta) \neq \zeta$ for every $\beta \in h(A)$, then we have shown that F restricts to an embedding of $h(A)$ in A . But by definition of $h(A)$, this is not possible. Hence, for some $\beta \in h(A)$, $F(\beta) = \zeta$. Moreover, if we let α be the least such β , then we have shown that $F \upharpoonright \alpha$ is an injection from α to A . It remains to show that this map is surjective. But this is immediate since $F(\alpha) = \zeta$ implies that $A = \text{Im}(F \upharpoonright \alpha)$.

Now let us assume the well-ordering principle and prove Zorn's lemma. Suppose (E, \prec) is a poset satisfying the stated condition on totally ordered subsets. By the well-ordering principle there exists a well-ordering, say $<$, on E . Without loss of generality we may assume that both orderings are strict. By Theorem 1.36 we may as well assume that E is an ordinal and $<$ is \in . Let $h(E)$ be the cardinal given by Proposition 2.6, that is $h(E)$ is the least ordinal that is not equinumerous with any subset of E . We will verify Zorn's lemma as follows: we will assume toward a contradiction that E has no \prec -maximal element and use that to construct an (impossible) embedding of $h(E)$ into E . Let $F : h(E) \rightarrow E$ be defined by transfinite recursion (cf. Corollary 1.27), as follows:

$$F(0) = 0$$

For all ordinals $\alpha < h(E)$,

$$F(\alpha + 1) = \text{the } <\text{-least } \beta \in E \text{ such that } F(\alpha) \prec \beta$$

For all limit ordinals $\alpha < h(E)$,

$$F(\alpha) = \begin{cases} < \text{-least } \beta \text{ such that } F(\gamma) \prec \beta \text{ for all } \gamma < \alpha & \text{if such a } \beta \text{ exists} \\ 0 & \text{else} \end{cases}$$

This function is well-defined because E has no \prec -maximal element and $<$ is a well-ordering on E . Since $h(E)$ is a limit ordinal (as its a cardinal), to show that F is injective it will suffice to prove the following:

CLAIM 2.9. *For all $\alpha < h(E)$, $F \upharpoonright \alpha$ is strictly order preserving; that is, for all $x < y < \alpha$, $F(x) \prec F(y)$.*

PROOF OF CLAIM 2.9. We prove this by transfinite induction on $\alpha < h(E)$. For $\alpha = 0$ there is nothing to prove. If α is a limit ordinal then $\alpha = \bigcup_{\beta < \alpha} \beta$ and so the claim follows

by the induction hypothesis. Finally, suppose α is any ordinal and consider $F \upharpoonright (\alpha + 1)$. There are two possibilities: either α is itself a successor or a limit. If $\alpha = \beta + 1$ then $F \upharpoonright (\alpha + 1) = F \upharpoonright \{\gamma : \gamma \leq \beta + 1\}$, and by definition $F(\beta) \prec F(\beta + 1)$. This, together with the induction hypothesis, shows that $F \upharpoonright (\alpha + 1)$ is strictly order preserving. So suppose α is a limit ordinal. By the induction hypothesis, $F \upharpoonright \alpha$ is strictly order preserving, and hence $\text{Im}(F \upharpoonright \alpha)$ forms a totally \prec -ordered set. It follows that there exists an upper bound to this set, and hence by definition $F(\alpha)$ is such a bound. So, $F(\gamma) \prec F(\alpha)$ for all $\gamma < \alpha$. Hence, $F \upharpoonright (\alpha + 1)$ is strictly order preserving, as desired. \square

Finally, let us assume Zorn's lemma and derive the axiom of choice from it. Suppose \mathcal{F} is a set of nonempty sets, and let us consider the set Λ of all *partial* choice functions on \mathcal{F} , identified with their graphs. That is, the elements of Λ are sets of the form

$$\{(G, x) : G \in \mathcal{G}, x \in G\}$$

where \mathcal{G} is a subset of \mathcal{F} . Note that Λ is non-empty, it contains for example $\{(F, x)\}$ for each $F \in \mathcal{F}$ and $x \in F$. Now Λ forms a poset under \subset . Moreover, if Θ is a totally ordered subset of Λ , then $\bigcup \Theta$ is quite easily seen to be a partial choice function on \mathcal{F} , and hence an upper bound for Θ in Λ . So the assumptions of Zorn's lemma are satisfied, and Λ must have a maximal element, say f_∞ . I claim that f_∞ is a (total) choice function on \mathcal{F} . If not, then there must exist some $F \in \mathcal{F}$, such that $F \notin \text{Dom}(f_\infty)$. But then, fixing $x \in F$, $f_\infty \cup \{(F, x)\}$ is a strictly larger partial choice function on \mathcal{F} , contradicting the maximality of f_∞ in Λ . Hence f_∞ is a choice function on \mathcal{F} , and we have proven the axiom of choice. \square

There are many other interesting equivalents to the axiom of choice, which I leave to you to investigate.

2.3. Axiom of choice and cardinals

Recall that cardinals were defined as those ordinals which are not equinumerous with any strictly lesser ordinals (cf. Section 2.1). In the presence of the axiom of choice, cardinals become a robust measure of the "size" of sets. We collect together in the following proposition several consequences of the axiom of choice to the theory of cardinals.

PROPOSITION 2.10. *Assume the Axiom of Choice.*

- (a) Every set is equinumerous with a unique cardinal. Given a set A we denote the unique cardinal equinumerous to it by $|A|$ and call this the cardinality of A . A set A is finite if $|A| < \omega$ and countable if $|A| \leq \omega$.
- (b) Given sets A and B , $|A| \leq |B|$ if and only if there is an injective function from A to B .
- (c) Given sets A and B , either there is an injective map from A to B or there is an injective function from B to A .
- (d) Suppose f is a function on a set A . Then $|\text{Im}(f)| \leq |A|$.
- (e) A countable union of countable sets is countable.

PROOF. (a) Suppose X is a set. By the axiom of choice, and Theorem 2.8, X is well-orderable. It follows by Theorem 1.36 that X is in bijection with an ordinal. Let α be the least ordinal equinumerous with X . Clearly α is a cardinal, and it is the unique cardinal equinumerous with X .

(b) Let $|A| = \kappa$ and $|B| = \lambda$ and fix bijections $f : A \rightarrow \kappa$ and $g : B \rightarrow \lambda$. If $\kappa \leq \lambda$ then $\kappa \subseteq \lambda$ and hence $g^{-1} \circ f$ is an injective function from A to B . Conversely, if $h : A \rightarrow B$ is injective then $g \circ h \circ f^{-1}$ is an embedding of κ in λ . If $\lambda < \kappa$ then there is also an embedding of λ in κ (by containment). But then Schröder-Bernstein (Proposition 2.2) implies that κ and λ are equinumerous, contradicting that the cardinal κ cannot be equinumerous with a lesser ordinal. Hence, $\kappa \leq \lambda$, as desired.

Part (c) is an immediate consequences of (b).

(d) Let c be a choice function on $\mathcal{F} = \{f^{-1}(y) : y \in \text{Im}(f)\}$, the set of fibres of f . Set $\iota : \text{Im}(f) \rightarrow A$ to be $\iota(y) = c(f^{-1}(y))$. Now if x and y are distinct in $\text{Im}(f)$ then $f^{-1}(x)$ and $f^{-1}(y)$ are disjoint, and hence $\iota(x) \neq \iota(y)$. So ι is injective and $|\text{Im}(f)| \leq |A|$.

(e) First we leave it as an exercise for you to see that $\omega \times \omega$ is countable. Now suppose A is a countable set with the property that every member of A is countable. Since A is countable we have an embedding $g : A \rightarrow \omega$. Moreover, for each $x \in A$ the set of embeddings from x into ω is nonempty and hence by choice there is a function $x \mapsto f_x$ on A such that $f_x : x \rightarrow \omega$ is an embedding. Now set $I := \{(i, j) \in \omega \times \omega : i \in \text{Im}(g), j \in \text{Im}(f_{g^{-1}(i)})\}$ and define $h : I \rightarrow \bigcup A$ by $(i, j) \mapsto f_{g^{-1}(i)}^{-1}(j)$. We show that h is surjective. Indeed, if $a \in \bigcup A$ then $a \in x$ for some $x \in A$, and setting $i = g(x)$ and $j = f_x(a)$ we see that $h(i, j) = a$. Hence $|\bigcup A| \leq |I| \leq |\omega \times \omega| \leq \omega$, as desired. \square

In particular we have now defined the *cardinality* of any set. Note that if X is a set of cardinality κ then we can enumerate X as $X = \{x_i : i < \kappa\}$. Indeed, let $f : \kappa \rightarrow X$ be a bijection and set $x_i := f(i)$.

EXERCISE 2.11. Suppose α is an ordinal and κ is a cardinal. Show that $|\alpha| \leq \alpha$, and that if $|\alpha| < \kappa$ then $\alpha < \kappa$.

In the rest of this course we will assume the axiom of choice, along with the Zermelo-Fraenkel axioms of set theory, without explicitly saying so. That is we work in ZFC; Zermelo-Fraenkel set theory with choice.

CHAPTER 3

Cardinals

In the last chapter we introduced cardinals and showed that in the presence of the axiom of choice cardinality is a robust measure of the size of a set, and that our intuition about the relative sizes of sets corresponds exactly to the ordering of the cardinals by \in . In this chapter we pursue further the structure of infinite cardinals. We use the axiom of choice freely.

3.1. Hierarchy of infinite cardinals

Our goal here is to describe a complete hierarchy of cardinals.

NOTATION 3.1. For any cardinal κ , $h(\kappa)$, which by definition is the least ordinal not equinumerous with any subset of κ (see Proposition 2.6), is also the least cardinal strictly bigger than κ . For that reason, we will denote $h(\kappa)$ by κ^+ .

DEFINITION 3.2. We define, using transfinite recursion, the following ordinal-enumerated collection of cardinals.

- $\aleph_0 = \omega$.
- For all ordinals α , $\aleph_{\alpha+1} = \aleph_\alpha^+$.
- For all limit ordinals α , $\aleph_\alpha = \sup\{\aleph_\beta : \beta < \alpha\}$.

LEMMA 3.3. *For all ordinals α , \aleph_α is a cardinal.*

PROOF. By transfinite induction on α . This is clear for $\alpha = 0$. At successor stages we use the fact that $h(E)$ is a cardinal for any set E , by Proposition 2.6. Finally suppose that α is a limit ordinal, and $\beta < \aleph_\alpha$. By definition of \aleph_α we have that for some $\gamma < \alpha$, $\beta < \aleph_\gamma$. So we have $\beta < \aleph_\gamma \leq \aleph_\alpha$ and hence $|\beta| < |\aleph_\gamma| \leq |\aleph_\alpha|$, where the first inequality uses the fact that \aleph_γ is a cardinal by the inductive hypothesis. So \aleph_α is not equinumerous with any lesser ordinal – that is, it is a cardinal. \square

LEMMA 3.4. *For all ordinals $\alpha < \beta$, $\aleph_\alpha < \aleph_\beta$.*

PROOF. By transfinite induction on β . There is nothing to prove for $\beta = 0$. If $\beta = \gamma + 1$ then $\aleph_\beta = \aleph_\gamma^+ > \aleph_\gamma$ and the result follows by the inductive hypothesis. If β is a limit ordinal then there exists $\gamma < \beta$ with $\alpha < \gamma$. So $\aleph_\alpha < \aleph_\gamma$ by the inductive hypothesis, and $\aleph_\gamma \leq \aleph_\beta$ since $\aleph_\gamma \subseteq \aleph_\beta$. Hence $\aleph_\alpha < \aleph_\beta$, as desired. \square

LEMMA 3.5. *For all ordinals α , $\alpha \leq \aleph_\alpha$. The inequality is strict if α is a successor ordinal.*

PROOF. By transfinite induction on α . This is clear for $\alpha = 0$. If $\alpha = \beta + 1$ then by the induction hypothesis $\beta \leq \aleph_\beta$, and so $\alpha = \beta + 1 \leq \aleph_\beta + 1$. On the other hand,

$\aleph_\beta + 1 < |\aleph_\beta|^+ = \aleph_\alpha$. Hence $\alpha < \aleph_\alpha$. Finally suppose α is a limit ordinal. For every $\beta < \alpha$, $\beta \leq \aleph_\beta < \aleph_\alpha$, where the first inequality is by the induction hypothesis and the second is Lemma 3.4. Hence, $\alpha = \sup\{\beta : \beta < \alpha\} \leq \aleph_\alpha$, as desired. \square

EXERCISE 3.6. The inequality in Lemma 3.5 may not be strict. For example, consider the sequence of ordinals defined recursively by $\alpha_0 = 0$ and $\alpha_{n+1} = \aleph_{\alpha_n}$. Now let $\alpha = \sup\{\alpha_n : n \in \mathbb{N}\}$. Verify that $\alpha = \aleph_\alpha$. In fact, this works if we start with any ordinal α_0 , not just 0.

PROPOSITION 3.7. *Every infinite cardinal is of the form \aleph_α for some ordinal α .*

PROOF. Suppose κ is a cardinal. By Lemma 3.5 $\kappa \leq \aleph_\kappa < \aleph_{\kappa+1}$. Hence it will suffice to show that *for every ordinal β and every infinite cardinal $\kappa < \aleph_\beta$, there exists an ordinal $\alpha < \beta$ such that $\kappa = \aleph_\alpha$* . We prove that statement by transfinite induction on β . Since there are no infinite cardinals strictly below ω , there is nothing to prove in the case $\beta = 0$. Suppose $\beta = \gamma + 1$ and let $\kappa < \aleph_\beta$. Then, as there is no cardinal strictly between \aleph_γ and $\aleph_\gamma^+ = \aleph_\beta$, $\kappa \leq \aleph_\gamma$. Hence either $\kappa = \aleph_\gamma$ or by the induction hypothesis $\kappa = \aleph_\alpha$ for some $\alpha < \gamma$. This deals with the successor stage. Suppose now that β is a limit ordinal and $\kappa < \aleph_\beta$. Then by definition $\kappa < \aleph_\gamma$ for some $\gamma < \beta$. Hence by the induction hypothesis $\kappa = \aleph_\alpha$ for some $\alpha < \gamma$, as desired. \square

We call a cardinal of the form $\aleph_{\alpha+1}$ a *successor cardinal* and one of the form \aleph_β for some limit ordinal β a *limit cardinal*. Note that by Proposition 3.7, successor cardinals are exactly those of the form κ^+ for some cardinal κ , and limit cardinals are exactly those that are not successors. It is important to not get confused between successor/limit cardinals and successor/limit ordinals. Indeed, every cardinal is a limit ordinal by Lemma 2.4.

3.2. A word on the continuum hypothesis

THEOREM 3.8 (Cantor's diagonalisation). *For every set E , $|E| < |\mathcal{P}(E)|$.*

PROOF. Since $x \mapsto \{x\}$ is an embedding of E in $\mathcal{P}(E)$, we have $|E| \leq |\mathcal{P}(E)|$. Suppose, toward a contradiction that there exists a bijective function $f : E \rightarrow \mathcal{P}(E)$. Let $\Delta := \{x \in E : x \notin f(x)\}$. So $\Delta \in \mathcal{P}(E)$ and hence $\Delta = f(x)$ for some $x \in E$. If $x \in \Delta$ then by definition $x \notin f(x) = \Delta$, a contradiction. Hence $x \notin \Delta$. But then $x \notin f(x)$ and so by definition $x \in \Delta$, again a contradiction. Hence no such f exists and $|E| < |\mathcal{P}(E)|$. \square

The question arises as to what cardinal $|\mathcal{P}(\aleph_0)|$ is in terms of the hierarchy described in the previous section. One might expect that $|\mathcal{P}(\aleph_0)| = \aleph_1$. In fact, that statement, which is called the *continuum hypothesis* (CH), cannot be proved from the set-theoretic axioms we have introduced so far, nor can its negation be proved. This is also the case for the stronger statement, called the *generalised continuum hypothesis* (GCH), which states that $|\mathcal{P}(\kappa)| = \kappa^+$ for all cardinals κ . Unlike the axiom of choice, the (generalised) continuum hypothesis is not indispensable to most of contemporary mathematics. So we do not add a set-theoretic axiom determining the value of $|\mathcal{P}(\aleph_0)|$, but rather work independently of its status.

3.3. Cardinal arithmetic

LEMMA 3.9. Suppose X_1, X_2, Y_1, Y_2 are such that $|X_i| = |Y_i|$, for each $i = 1, 2$.

- (a) If $X_1 \cap X_2 = Y_1 \cap Y_2 = \emptyset$, then $|X_1 \cup X_2| = |Y_1 \cup Y_2|$.
- (b) $|X_1 \times X_2| = |Y_1 \times Y_2|$.

PROOF. Let $f_i : X_i \rightarrow Y_i$ be a bijection witnessing the equinumerosity of X_i and Y_i , for $i = 1$ and 2 . Then $g : X_1 \cup X_2 \rightarrow Y_1 \cup Y_2$ given by $g(x) = \begin{cases} f_1(x) & \text{if } x \in X_1 \\ f_2(x) & \text{if } x \in X_2 \end{cases}$ and $h : X_1 \times X_2 \rightarrow Y_1 \times Y_2$ given by $h(a, b) = (f_1(a), f_2(b))$, are easily seen to be bijections. \square

DEFINITION 3.10 (Cardinal sum and product). Given cardinals κ_1, κ_2 we define their *sum* to be

$$\kappa_1 + \kappa_2 := |X_1 \cup X_2|$$

where X_1, X_2 are disjoint sets of cardinality κ_1 and κ_2 , respectively. Similarly, we define the *product* to be

$$\kappa_1 \cdot \kappa_2 := |X_1 \times X_2|$$

where X_1, X_2 are sets of cardinality $\kappa_1, \dots, \kappa_n$. By Lemma 3.9 these operations are well-defined; they do not depend on the choice of X_1, X_2 .

REMARK 3.11. It is very important here to distinguish between cardinal sum/product and ordinal sum/product. For example, viewing ω and 1 as cardinals we have that their sum is the cardinal ω (see 3.12(d) below), while viewing them as ordinals their sum is the ordinal $\omega + 1 \neq \omega$. We hope that the context will always make clear which is meant. One way to do so is to use the \aleph notation for cardinals. So for example, in the above case, for ordinal sum we would write $\omega + 1$, while for cardinal sum we would write $\aleph_0 + 1$ (which is equal to \aleph_0).

EXERCISE 3.12. (a) Sum and product are commutative and associative.

(b) Product distributes over sum: $\kappa_1 \cdot (\kappa_2 + \kappa_3) = \kappa_1 \cdot \kappa_2 + \kappa_1 \cdot \kappa_3$.

(c) For all cardinals κ , $0 + \kappa = \kappa$ and $0 \cdot \kappa = 0$.

(d) For all $n \in \omega$, $\aleph_0 + n = \aleph_0 \cdot n = \aleph_0$.

(d) $\aleph_0 + \aleph_0 = \aleph_0$.

(e) $\aleph_0 \cdot \aleph_0 = \aleph_0$.

THEOREM 3.13. For all infinite cardinals κ , $\kappa \cdot \kappa = \kappa$.

PROOF. First, given a cardinal κ , we describe a well-ordering on $\kappa \times \kappa$. This ordering denoted by \prec will be essentially the lexicographic one. We define $(x, y) \prec (x', y')$ if $\max\{x, y\} < \max\{x', y'\}$, or if $\max\{x, y\} = \max\{x', y'\}$ and $x < x'$, or if $\max\{x, y\} = \max\{x', y'\}$ and $x = x'$ and $y < y'$. We leave it to you to check that this is a strict linear ordering on $\kappa \times \kappa$. That it is well-ordered can be seen as follows: If X is a nonempty subset of $\kappa \times \kappa$ we let δ be the least ordinal such that $\delta = \max\{x, y\}$ for some $(x, y) \in X$, and we let D be the set of $(x, y) \in X$ with $\delta = \max\{x, y\}$. So D is a nonempty subset of X and a \prec -least element of D will, by definition, be a \prec -least element of X . Now set x to be least such that $(x, y) \in D$ for some y , and let E be the set of all such $y \in \kappa$. Then E is a nonempty subset of κ and if y is the least element in E , then (x, y) is the \prec -least element of D and hence of X .

Now to the proof that $\kappa \cdot \kappa = \kappa$. The diagonal map embeds κ into $\kappa \times \kappa$, so it remains to show that $\kappa \cdot \kappa \leq \kappa$. Writing $\kappa = \aleph_\alpha$ for some ordinal α (Proposition 3.7), we prove by induction on α that $\aleph_\alpha \cdot \aleph_\alpha \leq \aleph_\alpha$. The case of $\alpha = 0$ is Exercise 3.12(e) which is the fact that a countable union of countable sets is countable (Proposition 2.10(e)). Suppose $\aleph_\beta \cdot \aleph_\beta \leq \aleph_\beta$ for all $\beta < \alpha$.

Let γ be the ordinal to which $(\aleph_\alpha \times \aleph_\alpha, \prec)$ is order-isomorphic (Theorem 1.36). Now, if $\aleph_\alpha < \aleph_\alpha \cdot \aleph_\alpha$ then $\aleph_\alpha < |\aleph_\alpha \times \aleph_\alpha| = \gamma$. It follows that \aleph_α is a proper initial segment of γ , and hence order-isomorphic to a proper initial segment of $(\aleph_\alpha \times \aleph_\alpha, \prec)$. So it will suffice to show that every proper initial segment of $(\aleph_\alpha \times \aleph_\alpha, \prec)$ is of cardinality strictly less than \aleph_α . This is what we now do.

Suppose S is a proper initial segment of $(\aleph_\alpha \times \aleph_\alpha, \prec)$. That is, $S = \{(x, y) : (x, y) \prec (x_0, y_0)\}$ for some $x_0, y_0 \in \aleph_\alpha$. Let $\zeta = \max\{x_0, y_0\} + 1$. Note that $S \subseteq \zeta \times \zeta$. Now, as every cardinal is a limit ordinal, $\zeta < \aleph_\alpha$, and so $|\zeta| < \aleph_\alpha$ (see Exercise 2.11). By Theorem 3.7, there is a $\beta < \alpha$ such that $|\zeta| = \aleph_\beta$. Hence $|S| \leq |\zeta \times \zeta| = \aleph_\beta \cdot \aleph_\beta \leq \aleph_\beta < \aleph_\alpha$, where the penultimate inequality is the induction hypothesis and the ultimate one is Lemma 3.4. So $|S| < \aleph_\alpha$ as desired. \square

It follows that cardinal arithmetic for infinite cardinals trivialises to computing maxima:

COROLLARY 3.14. *For all cardinals $\kappa_1 \leq \kappa_2$ with κ_2 infinite we have:*

- (a) *If $\kappa_1 \neq 0$, then $\kappa_1 \cdot \kappa_2 = \kappa_2$.*
- (b) *$\kappa_1 + \kappa_2 = \kappa_2$.*

PROOF. As κ_2 is an infinite cardinal it is of the form \aleph_α for some ordinal α . Fixing $x \in \kappa_1$ we have that $y \mapsto (x, y)$ is an embedding of \aleph_α into $\kappa_1 \times \aleph_\alpha$. Hence

$$\begin{aligned} \aleph_\alpha &\leq \kappa_1 \cdot \aleph_\alpha \\ &\leq \aleph_\alpha \cdot \aleph_\alpha \quad \text{as } \kappa_1 \subseteq \aleph_\alpha \text{ and hence } \kappa_1 \times \aleph_\alpha \subseteq \aleph_\alpha \times \aleph_\alpha \\ &= \aleph_\alpha \quad \text{by Theorem 3.13.} \end{aligned}$$

This proves part (a). For part (b) we have

$$\begin{aligned} \aleph_\alpha &\leq \kappa_1 + \aleph_\alpha \\ &\leq \aleph_\alpha + \aleph_\alpha \quad \text{as } \kappa_1 \subseteq \aleph_\alpha \\ &= 2 \cdot \aleph_\alpha \\ &= \aleph_\alpha \quad \text{by part (a)} \end{aligned}$$

where the penultimate equality is because the disjoint union of two \aleph_α -cardinality sets is equinumerous with $\{0, 1\} \times \aleph_\alpha$. \square

We now extend the notions of cardinal sum and product to arbitrary (that is, possibly infinite) sets of cardinals. But first we need some preliminaries on arbitrary sequences of sets.

DEFINITION 3.15. Suppose X and I are sets.

- (a) By an *I-sequence* in X we mean simply a function $f : I \rightarrow X$. We often use the notation $(x_i : i \in I)$ for sequences, where $x_i := f(i)$.
- (b) Suppose $(X_i : i \in I)$ is a sequence of sets in X . By the *cartesian product* $\times_{i \in I} X_i$ we mean the set of all function $f : I \rightarrow \bigcup X$ such that $f(i) \in X_i$.

DEFINITION 3.16. Suppose $(\kappa_i : i \in I)$ is a sequence of cardinals. Then we can define the sum

$$\sum_{i \in I} \kappa_i = \left| \bigcup_{i \in I} X_i \right|$$

where $(X_i : i \in I)$ is a sequence of pairwise disjoint sets with $|X_i| = \kappa_i$ for all $i \in I$. Similarly we define the product

$$\prod_{i \in I} \kappa_i = \left| \times_{i \in I} X_i \right|$$

where $(X_i : i \in I)$ is such that $|X_i| = \kappa_i$ for all $i \in I$.

LEMMA 3.17. *The sums and products of Definition 3.16 are well-defined; they do not depend on the choice of $(X_i : i \in I)$.*

PROOF. This is exactly as in the finite case (cf. the proof of Lemma 3.9), except that we need to use the axiom of choice. Suppose $(X_i : i \in I)$ and $(X'_i : i \in I)$ are sequences of pairwise disjoint sets such that $|X_i| = |X'_i|$ for all $i \in I$. Using the axiom of choice, consider the function $i \mapsto f_i$ where $f_i : X_i \rightarrow X'_i$ is a bijection for each $i \in I$. We obtain the bijection $f : \bigcup_{i \in I} X_i \rightarrow \bigcup_{i \in I} X'_i$ given by $f(x) = f_i(x)$ where $i \in I$ is the (unique) index such that $x \in X_i$. Similarly we get the bijection $g : \times_{i \in I} X_i \rightarrow \times_{i \in I} X'_i$ given by $g(x_i : i \in I) = (f_i(x_i) : i \in I)$. \square

Infinite cardinal sums also reduce to computing suprema, as Proposition 3.20 below demonstrates. But we begin with some easy lemmas.

LEMMA 3.18. *If κ and λ are cardinals, then $\sum_{i < \lambda} \kappa = \lambda \cdot \kappa$.*

PROOF. It suffices to show that if X_i , for $i < \lambda$, are pairwise disjoint sets with each $|X_i| = \kappa$, then there is a bijection between $\bigcup_{i < \lambda} X_i$ and $\lambda \times \kappa$. But this is clear since each X_i is equinumerous with $\{i\} \times \kappa$. \square

LEMMA 3.19. *If $(\kappa_i : i \in I)$ is a sequence of cardinals then $\sup_{i \in I} \kappa_i$ is again a cardinal.*

Moreover, $\sup_{i \in I} \kappa_i \leq \sum_{i < \lambda} \kappa_i$.

PROOF. Recall that by definition $\sup_{i \in I} \kappa_i = \bigcup_{i \in I} \kappa_i$. Now $\left| \bigcup_{i \in I} \kappa_i \right| \leq \sum_{i \in I} \kappa_i$. If this inequality were strict, then as the right hand side is the least upper bound (cf. Lemma 1.19(b)), there exists $j \in I$ such that $\left| \bigcup_{i \in I} \kappa_i \right| < \kappa_j$, which contradicts the fact that κ_j is a cardinal and $\kappa_j \subseteq \bigcup_{i \in I} \kappa_i$. Hence, $\left| \bigcup_{i \in I} \kappa_i \right| = \sum_{i \in I} \kappa_i$, proving that the latter is a cardinal.

For the moreover clause observe that if $(X_i : i \in I)$ is a pairwise disjoint sequence of sets with bijections $f_i : X_i \rightarrow \kappa_i$, then $\bigcup_{i \in I} f_i : \bigcup_{i \in I} X_i \rightarrow \bigcup_{i \in I} \kappa_i$ is surjective. \square

PROPOSITION 3.20. *If λ is a nonzero cardinal and $(\kappa_i : i < \lambda)$ is a sequence of nonzero cardinals, then*

$$\sum_{i < \lambda} \kappa_i = \lambda \cdot \sup_{i < \lambda} \kappa_i = \max \left(\lambda, \sup_{i < \lambda} \kappa_i \right).$$

PROOF. The second equality is just by Corollary 3.14(a). For the first equality, let $\kappa := \sup_{i < \lambda} \kappa_i$. Then it is easy to see that $\sum_{i < \lambda} \kappa_i \leq \sum_{i < \lambda} \kappa$, and the latter is $\lambda \cdot \kappa$ by Lemma 3.18. Hence $\sum_{i < \lambda} \kappa_i \leq \lambda \cdot \kappa$ and we need only show the opposite inequality. Now $\lambda = \lambda \cdot 1 = \sum_{i < \lambda} 1 \leq \sum_{i < \lambda} \kappa_i$, where the second equality is by 3.18 and the inequality is by the fact that each κ_i is nonzero and hence ≥ 1 . On the other hand, $\kappa = \sup_{i < \lambda} \kappa_i \leq \sum_{i < \lambda} \kappa_i$ by Lemma 3.19. Taking products we get

$$\lambda \cdot \kappa \leq \left(\sum_{i < \lambda} \kappa_i \right) \cdot \left(\sum_{i < \lambda} \kappa_i \right) = \sum_{i < \lambda} \kappa_i$$

where the final equality uses the fact that $\sum_{i < \lambda} \kappa_i$ is an infinite cardinal (since λ is and each κ_i is nonzero). \square

Infinite cardinal products, on the other hand, do not reduce to computing suprema: Given any cardinal $\lambda \geq 2$, let $\kappa_i = 2$ for all $i < \lambda$. The correspondence which assigns to each $a = (a_i : i < \lambda) \in \times_{i < \lambda} \kappa_i$ the set $\{i \in \lambda : a_i = 1\}$, is a bijection between $\times_{i < \lambda} \kappa_i$ and $\mathcal{P}(\lambda)$. Hence $\prod_{i < \lambda} \kappa_i = |\mathcal{P}(\lambda)|$. But this is strictly larger than λ by Cantor's diagonalisation (Theorem 3.8), and also strictly larger than each $\kappa_i = 2$.

The above example suggests the connection between infinite products and exponentiation.

DEFINITION 3.21 (Cardinal exponentiation). If κ and λ are cardinals then κ^λ is defined to be the cardinality of the set of all functions from λ to κ .

LEMMA 3.22. *If κ and λ are cardinals, then $\prod_{i < \lambda} \kappa = \kappa^\lambda$.*

PROOF. This follows from the fact that the λ th cartesian power of a set is exactly the set of functions from λ to that set. \square

LEMMA 3.23. *Suppose κ, λ, μ are cardinals. Then*

- (a) *If $\lambda \leq \mu$ then $\kappa^\lambda \leq \kappa^\mu$ and $\lambda^\kappa \leq \mu^\kappa$.*
- (b) *$\kappa^{\lambda+\mu} = \kappa^\lambda \cdot \kappa^\mu$, and*
- (c) *$(\kappa^\lambda)^\mu = \kappa^{\lambda \cdot \mu}$.*

PROOF. Part (a) follows rather easily from the definition.

For part (b) let X be a set of cardinality λ and Y a set, disjoint from X , of cardinality μ . Let $\text{Fun}(X, \kappa)$ denote the set of all functions from X to κ , and similarly for Y . Then $\kappa^\lambda \cdot \kappa^\mu$ is equinumerous with $\kappa^\lambda \times \kappa^\mu$ which is in turn equinumerous with $\text{Fun}(X, \kappa) \times \text{Fun}(Y, \kappa)$. But this latter set is equinumerous with $\text{Fun}(X \cup Y, \kappa)$ via the association $(f, g) \mapsto f \cup g$. Finally, as $X \cup Y$ is of cardinality $\lambda + \mu$, $\text{Fun}(X \cup Y, \kappa)$ is of cardinality $\kappa^{\lambda+\mu}$, as desired.

For part (c) it suffices to show that $\text{Fun}(\lambda \times \mu, \kappa)$ is equinumerous with $\text{Fun}(\mu, \text{Fun}(\lambda, \kappa))$. To see that this is the case suppose $f : \lambda \times \mu \rightarrow \kappa$. Then for each $x \in \mu$ let $g_x : \lambda \rightarrow \kappa$ be given by $g_x(y) = f(y, x)$ for all $y \in \lambda$. The association $f \mapsto (x \mapsto g_x : x \in \mu)$ is a bijection between $\text{Fun}(\lambda \times \mu, \kappa)$ and $\text{Fun}(\mu, \text{Fun}(\lambda, \kappa))$. \square

EXERCISE 3.24. (a) Show that for any infinite cardinal κ , $\kappa^\kappa = 2^\kappa$.

(b) Show that $\prod_{i < \aleph_0, i \neq 0} i = 2^{\aleph_0}$ and so $\sum_{i < \aleph_0} i < \prod_{i < \aleph_0, i \neq 0} i$.

The following is a generalisation of Cantor's diagonalisation.

THEOREM 3.25 (König's Theorem). *Suppose $(\kappa_i : i \in I)$ and $(\lambda_i : i \in I)$ are sequences of cardinals with $\kappa_i < \lambda_i$ for all $i \in I$. Then $\sum_{i \in I} \kappa_i < \prod_{i \in I} \lambda_i$.*

PROOF. Let $(X_i : i \in I)$ be a sequence of pairwise disjoint sets with $|X_i| = \kappa_i$ for all $i \in I$, and let $(Y_i : i \in I)$ be such that $|Y_i| = \lambda_i$ for all $i \in I$. We will show that there is an injection from $\bigcup_{i \in I} X_i$ to $\times_{i \in I} Y_i$, but no surjection.

By choice let $i \mapsto \pi_i$ be a function where $\pi_i : X_i \rightarrow Y_i$ is an embedding for all $i \in I$. Also by choice, fix an element $c \in \times_{i \in I} Y_i$ such that for each $i \in I$, $c_i \in Y_i \setminus \pi_i(X_i)$. For each $x \in \bigcup_{i \in I} X_i$ let

$$f(x)_i = \begin{cases} \pi_i(x) & \text{if } x \in X_i \\ c_i & \text{else.} \end{cases}$$

Then $f(x) \in \times_{i \in I} Y_i$ and we have constructed a function $f : \bigcup_{i \in I} X_i \rightarrow \times_{i \in I} Y_i$. Note that $f(x)$ differs from c at exactly the (unique) i th co-ordinate where $x \in X_i$, and then for that i , $f(x)_i = \pi_i(x)$. So by the injectivity of the π_i 's we have that f is injective.

Suppose, toward a contradiction, that there is a surjection $h : \bigcup_{i \in I} X_i \rightarrow \times_{i \in I} Y_i$. For each $i \in I$, let $p_i : \times_{i \in I} Y_i \rightarrow Y_i$ be the projection onto the i th coordinate. We obtain $h_i := p_i \circ (h \upharpoonright X_i) : X_i \rightarrow Y_i$ which cannot be surjective since $\kappa_i < \lambda_i$ (cf. Proposition 2.10(d)). Using choice, fix $c \in \times_{i \in I} Y_i$ such that for each $i \in I$, $c_i \in Y_i \setminus h_i(X_i)$. By the assumption that h is surjective there must exist some $j \in I$ and some $x \in X_j$ with $h(x) = c$. But then $h_j(x) = c_j$ which contradicts our choice of c . \square

3.4. Regularity and cofinality

LEMMA 3.26. *Suppose κ is a cardinal. The following are equivalent.*

- (i) *For every subset X of κ , of cardinality strictly less than κ , $\sup X < \kappa$.*
- (ii) *Suppose $(X_i : i < \lambda)$ is a sequence of subsets of κ of length $\lambda < \kappa$, and each $|X_i| < \kappa$. Then $|\bigcup_{i < \lambda} X_i| < \kappa$.*

PROOF. Suppose (i) holds. Then for each $i < \lambda$, for X_i as in the statement of (ii) we have $\lambda_i := \sup X_i < \kappa$. So $\Lambda := \{\lambda_i : i < \lambda\}$ is a subset of κ of cardinality $\leq \lambda < \kappa$. Applying (i) again, $\sup \Lambda < \kappa$. Now every member of X_i is $\leq \lambda_i < \sup \Lambda + 1$. Hence $\bigcup_{i < \lambda} X_i \subseteq \sup \Lambda + 1$.

Hence, $|\bigcup_{i < \lambda} X_i| \leq |\sup \Lambda + 1| = |\sup \Lambda| \leq \sup \Lambda < \kappa$, as desired.

For the converse let $\lambda := |X|$ and enumerate $X = (x_i : i < \lambda)$. Then each x_i is an ordinal less than κ and hence $x_i \subseteq \kappa$ with $|x_i| \leq x_i < \kappa$. Applying (ii) we get that $|\sup X| = |\bigcup X| = |\bigcup_{i < \lambda} x_i| < \kappa$. So $\sup X < \kappa$ by Exercise 2.11. \square

DEFINITION 3.27 (Regular/Singular Cardinals). A cardinal is called *regular* if it satisfies the equivalent conditions of Lemma 3.26. A cardinal is *singular* if it is not regular.

For example every finite cardinal is regular, as is \aleph_0 . An example of a singular cardinal is \aleph_ω (which is by definition $= \bigcup_{i < \omega} \aleph_i$). Indeed, for each $i < \omega$ fix $x_i \in \aleph_{i+1} \setminus \aleph_i$. Then $X := \{x_i : i < \omega\}$ is a countable subset of \aleph_ω with $\sup X = \aleph_\omega$. This example motivates the following notion:

DEFINITION 3.28 (Cofinality). The *cofinality* of an ordinal α , denoted by $\text{cof}(\alpha)$, is the least ordinal β for which there exists a strictly increasing unbounded function $f : \beta \rightarrow \alpha$. That is, if $x < y < \beta$ then $f(x) < f(y)$, and for every $z < \alpha$ there is an $x < \beta$ with $f(x) \geq z$.

Note that since the identity function $\text{id} : \alpha \rightarrow \alpha$ is strictly increasing and unbounded, $\text{cof}(\alpha)$ always exists and is less than or equal to α .

LEMMA 3.29. *Suppose α is an ordinal. Then*

- (a) $\text{cof}(\alpha)$ is a cardinal, and
- (b) $\text{cof}(\text{cof}(\alpha)) = \text{cof}(\alpha)$.

PROOF. We will need the following claim.

CLAIM 3.30. *If $\beta < \alpha$ are equinumerous ordinals, then there exists $\beta' \leq \beta$ and a strictly increasing unbounded function $f : \beta' \rightarrow \alpha$.*

PROOF OF CLAIM 3.30. Suppose $g : \beta \rightarrow \alpha$ is a bijection. Define $g' : \beta \rightarrow \alpha$ by transfinite recursion as follows:

$$g'(0) = g(0),$$

for all ordinals $\gamma < \beta$ with $\gamma + 1 < \beta$

$$g'(\gamma + 1) = \begin{cases} g(x) \text{ where } x \text{ is least such that } g'(\gamma) < g(x) & \text{if such } x \text{ exists} \\ g'(\gamma) & \text{else} \end{cases}$$

and for all limit ordinals $\gamma < \beta$,

$$g'(\gamma) = \sup\{g(\gamma), g'(x) : x < \gamma\}.$$

Then it is not hard to see that g' is an increasing (though not necessarily strictly increasing) function with the property that for all $y < \beta$, $g'(y) \geq g(x)$ for all $x \leq y$. In particular, $\text{Im } g'$ has no strict upper bound in α . If g' is strictly increasing then set $\beta' = \beta$ and $f = g'$ and

we are done. If g' is not strictly increasing then let $\gamma < \beta$ be least such there exists η with $\gamma < \eta < \beta$ and $g'(\eta) = g(\gamma)$. Since g' is increasing it must be that $g'(\gamma + 1) = g'(\gamma)$. But then by definition $g'(\gamma) = \sup \alpha$. Hence setting $\beta' = \gamma + 1$ and $f = g' \upharpoonright (\gamma + 1)$ satisfies the claim. \square

Now to the proof of part (a) of Lemma 3.29. Suppose toward a contradiction that there exists a $\beta < \text{cof}(\alpha)$ equinumerous with $\text{cof}(\alpha)$. Applying the claim to this pair of ordinals we obtain $\beta' \leq \beta$ and a strictly increasing unbounded function $f : \beta' \rightarrow \text{cof}(\alpha)$. Now by definition we have a strictly increasing unbounded $g : \text{cof}(\alpha) \rightarrow \alpha$. But then $g \circ f : \beta' \rightarrow \alpha$ is strictly increasing and unbounded. As $\beta' < \text{cof}(\alpha)$ this contradicts the definition of $\text{cof}(\alpha)$. So no such β can exist and $\text{cof}(\alpha)$ is shown to be a cardinal.

For part (b) we have already observed that $\text{cof}(\text{cof}(\alpha)) \leq \text{cof}(\alpha)$. For the converse, by definition there is a strictly increasing unbounded function $f : \text{cof}(\text{cof}(\alpha)) \rightarrow \text{cof}(\alpha)$, and a strictly increasing unbounded function $g : \text{cof}(\alpha) \rightarrow \alpha$. But then $g \circ f : \text{cof}(\text{cof}(\alpha)) \rightarrow \alpha$ is strictly increasing and unbounded. Hence $\text{cof}(\text{cof}(\alpha)) \geq \text{cof}(\alpha)$, as desired. \square

Our explanation above of why \aleph_ω is singular also showed that $\text{cof}(\aleph_\omega) \leq \omega$.

PROPOSITION 3.31. *A cardinal κ is regular if and only if $\text{cof}(\kappa) = \kappa$.*

PROOF. Suppose $\text{cof}(\kappa) = \kappa$ and X is a subset of κ of cardinality $\lambda < \kappa$. Now the induced ordering $(X, <)$ is strictly well-ordered and hence is order-isomorphic to an ordinal α . That is, we have a strictly increasing bijection $f : \alpha \rightarrow X$. Since $|\alpha| = \lambda < \kappa$, $\alpha < \kappa$. As $\text{cof}(\kappa) = \kappa > \alpha$, it follows that f is bounded in κ . That is, there exists $z < \kappa$ such that z is a strict upper bound for $\text{Im } f = X$. Hence $\sup X < \kappa$. We have shown that κ is regular.

Now suppose $\text{cof}(\kappa) = \alpha < \kappa$, so that there exists a strictly increasing unbounded $f : \alpha \rightarrow \kappa$. This means that $\sup(\text{Im } f) = \sup \kappa$. As κ is a limit ordinal (it is a cardinal), $\sup \kappa = \kappa$. So we have a subset $\text{Im } f \subseteq \kappa$ whose cardinality is $|\alpha| \leq \alpha < \kappa$ but with supremum equal to κ . It follows that κ is not regular. \square

EXERCISE 3.32. Every successor cardinal (i.e., cardinal of the form $\aleph_{\alpha+1}$) is regular.

Note that cof is not an increasing operation on cardinals. For example $\text{cof}(\aleph_n) = \aleph_n$ for all $n < \omega$ by Exercise 3.32 and Proposition 3.31, but $\text{cof}(\aleph_\omega) = \aleph_0$.

The following is a strengthening of Cantor's diagonalisation, and it is the principal corollary of König's Theorem.

PROPOSITION 3.33. *For all infinite cardinals κ , $\text{cof}(2^\kappa) > \kappa$.*

PROOF. This is immediate for κ finite, so we assume κ is infinite. Let $\lambda \leq \kappa$ and suppose $f : \lambda \rightarrow 2^\kappa$ is a strictly increasing function. Then $\sup(\text{Im } f) = \bigcup_{i < \lambda} f(i)$. But we can compute

$$\begin{aligned}
 \left| \bigcup_{i < \lambda} f(i) \right| &\leq \sum_{i < \lambda} |f(i)| && \text{as unions embed into disjoint unions (exercise)} \\
 &< \prod_{i < \lambda} 2^\kappa && \text{by 3.25 (König's Theorem) since each } |f(i)| < 2^\kappa \\
 &= (2^\kappa)^\lambda && \text{by Lemma 3.22} \\
 &= 2^{\kappa \cdot \lambda} && \text{by Lemma 3.23(c)} \\
 &= 2^\kappa && \text{by Corollary 3.14(a) as } \lambda \leq \kappa \text{ and } \kappa \text{ is infinite.}
 \end{aligned}$$

Hence $\sup(\text{Im } f) < 2^\kappa$. That is, f must be bounded in 2^κ . We have shown that λ cannot be the cofinality of 2^κ . So $\text{cof}(2^\kappa) > \kappa$, as desired. \square

This gives us a little bit more information on the continuum:

COROLLARY 3.34. $2^{\aleph_0} \neq \aleph_\omega$ \square

Part 2

Model Theory

CHAPTER 4

First-order Logic

4.1. Structures

DEFINITION 4.1. A *structure* \mathcal{M} consists of a nonempty underlying set M , called the *universe* of \mathcal{M} , together with

- a sequence $(c_i : i \in I_{\text{con}})$ of distinguished elements of M , called *constants*;
- a sequence $(f_i : M^{n_i} \rightarrow M : i \in I_{\text{fun}})$ of distinguished maps from various cartesian powers of M to M itself, called *basic functions*; and,
- a sequence $(R_i \subseteq M^{k_i} : i \in I_{\text{rel}})$ of distinguished subsets of various cartesian powers of M , called *basic relations*.

Each of the natural numbers n_i and k_i that appear above are assumed to be nonzero and are called the *arity* of the corresponding function or relation. The constants, basic functions and basic relations together make up the *signature* of \mathcal{M} .

- REMARK 4.2.
- (a) While the universe of a structure is never empty, any of the sets of constants, relations or functions may be empty. So, for example, a nonempty set M by itself forms a structure.
 - (b) Note that constants are nothing other than 0-ary functions, under the convention that $M^0 = \{\emptyset\}$. So if we had allowed 0-ary functions we could have done without constants.
 - (c) Some treatments of this material ask that the equality relation on M always be included among the relations in a structure. We will not do so here, only because our point of view will be that equality is an inherent part of the set M itself, and hence need not be distinguished as a named relation.

The notion of structure here is very natural, it is nothing other than a set M equipped with a signature which dictates what “structure” on M we are interested in studying. For example, consider the set of real numbers \mathbb{R} . If we wish to view \mathbb{R} purely as a *set* (with equality), then we can consider the structure whose universe is \mathbb{R} and whose signature is empty. On the other hand, to study the structure $(\mathbb{R}, 0, +, -)$ whose signature consists of 0 (a constant), addition (a binary function) and additive inverse (a unary function), is to study the reals as a *group* under addition. If we wish to study the reals as a *ring* then we can consider the structure $(\mathbb{R}, 0, 1, +, -, \times)$. The structure $(\mathbb{R}, 0, 1, +, -, \times, <)$ where $<$ is the usual ordering on the reals (as a binary relation) corresponds to the *ordered ring* of real numbers. Each of the examples above have the same underlying set, just equipped with ever expanding signatures. A natural question, very much at the heart of model theory, would be to ask whether, for example, the ordered ring of reals can be recovered from just the ring structure, or whether the ring of reals can be recovered from the additive group of reals. Of course we need to make precise what we mean by “recovered”, but once we have done so we

will see that the answer to the former is yes (this is easy), and to the latter is no (this is a little harder). We are, however, getting ahead of ourselves. For now, let us just record this notion of expansion we have been discussing.

DEFINITION 4.3 (Expansion and Reduct). Suppose \mathcal{M} and \mathcal{N} are structures. We say that \mathcal{N} is an *expansion* of \mathcal{M} , or that \mathcal{M} is a *reduct* of \mathcal{N} , if they have the same universe and the signature of \mathcal{N} contains the signature of \mathcal{M} .

4.2. Languages

Consider the structures $(\mathbb{R}, 0, 1, +, -, \times)$ and $(\mathbb{F}_5, 0, 1, +, -, \times)$, where \mathbb{F}_5 is the set of integers mod 5. They do not actually have the same signature, since, for example, the $+$ in the first case denotes addition on the reals while the $+$ in the second denotes addition on \mathbb{F}_5 . There is however a sense in which both $+$ s are interpretations of some “additive structure”. This is made precise by the notion of a common *language*.

DEFINITION 4.4 (Language and L -structure). A *language* L is determined by specifying the following three sets of symbols:

- (i) a set of *constant symbols* L^{con} ;
- (ii) a set of *function symbols* L^{fun} , together with a positive integer n_f for every $f \in L^{\text{fun}}$ called the *arity* of f ; and,
- (iii) a set of *relation symbols* L^{rel} , together with a positive integer k_R for every $R \in L^{\text{rel}}$ called the *arity* of R .

An L -*structure* is then a structure \mathcal{M} together with bijective correspondences between L^{con} and I_{con} , L^{fun} and I_{fun} , and L^{rel} and I_{rel} , that preserve arity. So to each constant symbol $c \in L^{\text{con}}$ is associated a constant $c^{\mathcal{M}} \in M$ of \mathcal{M} , to each function symbol $f \in L^{\text{fun}}$ of arity n is associated a basic n -ary function $f^{\mathcal{M}} : M^n \rightarrow M$ of \mathcal{M} , and to each relation symbol $R \in L^{\text{rel}}$ of arity n is associated a basic n -ary relation $R^{\mathcal{M}} \subseteq M^n$ of \mathcal{M} . These constants, functions, and relations, $c^{\mathcal{M}}, f^{\mathcal{M}}, R^{\mathcal{M}}$, are the *interpretations* in \mathcal{M} of the corresponding symbols.

So, for example, $(\mathbb{R}, 0, 1, +, -, \times)$ and $(\mathbb{F}_5, 0, 1, +, -, \times)$ are both naturally L -structures where $L := \{0, 1, +, -, \times\}$ is made up of two constant symbols, two binary function symbols, one unary function symbol, and no relation symbols. This particular language is often called the *language of rings* for obvious reasons.

REMARK 4.5. Note that we have already begun to abuse the notation: we have used the same notation, $+$ for example, to denote both the function symbol in L and its interpretations in the L -structures $\mathcal{R} := (\mathbb{R}, 0, 1, +, -, \times)$ and $\mathcal{F} := (\mathbb{F}_5, 0, 1, +, -, \times)$. To be correct and unambiguous, if we used $+$ for the symbol in the language then we should use $+^{\mathcal{R}}$ for addition on \mathbb{R} and $+^{\mathcal{F}}$ for addition on \mathbb{F} . This would indeed become unwieldy and it is common to not make this notational distinction, at least when the context makes clear whether we mean a symbol in a language or its interpretation in some particular structure.

Let us consider some further examples.

EXAMPLE 4.6 (Orderings). In the first part of this course we studied ordinals. The structure we were interested in was the ordering on the ordinal as given by the membership relation. In model-theoretic terms we were interested in structures of the form $\mathcal{A} := (\alpha, \in)$

where α is an ordinal. A natural language for this structure is the *language of orderings*, $L := \{<\}$, consisting of a single binary relation symbol (and no constant or function symbols). Thus \mathcal{A} is an L -structure with $<^{\mathcal{A}} = \in$. The abuse of notation described in Remark 4.5 is thus consistent with (and explains) our convention from the last chapter that $<$ and \in are synonymous when dealing with ordinals. However there are other L -structures that look very different than ordinals. For example, the usual ordering on the rational numbers, $(\mathbb{Q}, <)$, is an L -structure. This linear ordering is not an ordinal, it is dense which is very far from well-ordered. In fact, it is important to note that an L -structure need not be a linear ordering at all; an L -structure is just a nonempty set together with a binary relation. For example, consider the L -structure \mathcal{C} whose universe is the complex numbers and where $<$ is so interpreted that $a <^{\mathcal{C}} b$ if and only if $a^2 + b^2 = 1$. This L -structure is not even a poset.

EXAMPLE 4.7 (Vector spaces). Suppose F is a field. The *language of F -vector spaces* usually refers to the language $L := \{0, +, -, \lambda_a : a \in F\}$ consisting of one constant symbol 0 , one binary function symbol $+$, a unary function symbol $-$, and a set of unary function symbols $\{\lambda_a : a \in F\}$, indexed by F . Any F -vector space, V , is made into an L -structure by interpreting 0 as the zero vector, $+$ as vector addition, $-$ as the operation which takes the negative of a vector, and for each $a \in F$, interprets λ_a as scalar multiplication by a . As in the previous example the converse is not true; not every L -structure is a vector space. Only those L -structures *satisfying certain axioms* about how the interpretations behave will be F -vector spaces. We will study this notion of satisfaction later.

We ended the previous section by defining expansions and reducts: when the universe of a structure is unchanged but the signature is expanded or reduced. One can also vary structures in the opposite way: let the universe increase or decrease but leave the signature, or rather the language, constant.

DEFINITION 4.8 (Embedding). Suppose L is a language and \mathcal{M} and \mathcal{N} are L -structures with universes M and N respectively. An *L -embedding* of \mathcal{M} in \mathcal{N} is an injective map $j : M \rightarrow N$ such that

- (i) for all $c \in L^{\text{con}}$, $j(c^{\mathcal{M}}) = c^{\mathcal{N}}$;
- (ii) for all $f \in L^{\text{fun}}$ and all $a \in M^{n_f}$, $j(f^{\mathcal{M}}(a)) = f^{\mathcal{N}}(j(a))$; and,
- (iii) for all $R \in L^{\text{rel}}$ and all $a \in M^{k_R}$, $a \in R^{\mathcal{M}}$ if and only if $j(a) \in R^{\mathcal{N}}$.

A surjective L -embedding is called an *L -isomorphism*.

If $M \subseteq N$ and the inclusion map is an L -embedding then we say that \mathcal{M} is a *L -substructure* of \mathcal{N} , or that \mathcal{N} is an *L -extension* of \mathcal{M} . In this case we write $\mathcal{M} \subseteq \mathcal{N}$.

- REMARK 4.9.**
- (a) It is important that in (iii) we have an “if and only if” and not just an “only if”. That, together with the injectivity of ι , distinguish L -embeddings from *L -homomorphisms*, which we will not discuss here. The injectivity of ι is itself just the “if” part of (iii) applied to the implicit equality relation.
 - (b) Note that $\mathcal{M} \subseteq \mathcal{N}$ if and only if $M \subseteq N$, $c^{\mathcal{M}} = c^{\mathcal{N}}$ for all $c \in L^{\text{con}}$, $f^{\mathcal{M}} = f^{\mathcal{N}} \upharpoonright M^{n_f}$ for all $f \in L^{\text{fun}}$, and $R^{\mathcal{M}} = R^{\mathcal{N}} \cap M^{k_R}$ for all $R \in L^{\text{rel}}$.

The notion of substructure is sensitive to the choice of language, and indeed often informs what language we choose. For example, consider the languages $L_1 = \{0, 1, +, -, \times\}$, $L_2 = \{0, +, -\}$, and $L_3 = \emptyset$. Then, under the natural interpretations, the L_1 -substructures of \mathbb{R}

are its *subrings*, the L_2 -substructures are the *subgroups* of the additive group of reals, and the L_3 -substructures are the nonempty *subsets* of \mathbb{R} .

EXERCISE 4.10. Suppose \mathcal{N} is an L -structure and M is a nonempty subset of the universe of \mathcal{N} . Show that M is the universe of a substructure of \mathcal{N} if and only if M contains all the constants of \mathcal{N} and is preserved by all the functions in \mathcal{N} . Moreover, if this is the case, then there is a unique L -substructure of \mathcal{N} that has M as its underlying set.

Note that $j : M \rightarrow N$ is an embedding of \mathcal{M} in \mathcal{N} if and only if it induces an isomorphism between \mathcal{M} and a substructure of \mathcal{N} . For convenience we may therefore sometimes replace L -embeddings by L -substructures.

4.3. Some syntax

In this section we will describe the formal rules for writing down *first-order formulas* in a given language L . These formulas will be used, in later sections, on the one hand to express properties that tuples from a particular L -structure may or may not possess (giving rise to *definable sets*), and on the other hand to express axioms that L -structures may or may not satisfy (giving rise to the class of *models of a theory*). The study of these two kinds of objects, the sets definable in L -structures and the class of models of a theory, is essentially what model theory is about.

The starting point is the notion of term which makes precise which functions can be defined, in a finitary manner, from the function and constant symbols in the language. We will make use of a fixed countably infinite set of distinct *variable symbols* $\text{Var} = \{v_0, v_1, \dots\}$.

Let us fix a language L .

DEFINITION 4.11 (Term). The set of L -terms is the smallest set of strings of symbols satisfying:

- (i) A variable symbol is an L -term.
- (ii) A constant symbol is an L -term.
- (iii) If $f \in L^{\text{fun}}$ and t_1, \dots, t_{n_f} are L -terms then $f(t_1, \dots, t_{n_f})$ is an L -term.

We sometimes write a term t as $t(x_1, \dots, x_n)$ to mean that the variable symbols appearing in t come from the list x_1, \dots, x_n . (Though not all of these variables need appear in t .)

For the sake of readability we may often diverge from the above rules for writing terms according to the natural use of the symbols in everyday mathematics. For example we will write $(v_0 + v_1) - 1$ for the more accurate but rather inscrutable $+(+(v_0, v_1), -(1))$. Another example is $(1 + v_1)^{-1}$ rather than the official $^{-1}(+(1, v_1))$.

The next step is the atomic formula, which is the simplest kind of formula and is obtained by equating or relating terms.

DEFINITION 4.12 (Atomic formula). An *atomic L -formula* is a string of symbols of the form

- (i) $(t = s)$ where t and s are L -terms, or
- (ii) $R(t_1, \dots, t_{k_R})$ where $R \in L^{\text{rel}}$ and t_1, \dots, t_{k_R} are L -terms.

Again we may use common abbreviations for the sake of readability. For example, we write $v_2 < -v_1$ instead of $< (v_2, -(v_1))$.

Finally, we can build recursively on atomic formulas, using the logical *connectives* $\{\neg, \wedge, \vee\}$ and the *quantifiers* $\{\forall, \exists\}$, to define the set of formulas.

DEFINITION 4.13 (Formula). The set of *L-formulas* is the smallest set of strings of symbols satisfying:

- (i) Every atomic formula is a formula.
- (ii) If ϕ and ψ are formulas then $\neg\phi$, $(\phi \wedge \psi)$, and $(\phi \vee \psi)$ are all formulas.
- (iii) If ϕ is a formula and x is a variable symbol then $\forall x\phi$ and $\exists x\phi$ are formulas.

We also make the following two abbreviations: For any *L-formulas* ϕ and ψ ,

- $(\phi \rightarrow \psi)$ abbreviates the formula $(\neg\phi \vee \psi)$, and,
- $(\phi \leftrightarrow \psi)$ abbreviates the formula $((\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi))$.

EXAMPLE 4.14 (Axioms of ZF). In the first part of this course we talked a little about Zermelo-Fraenkel set theory. The axioms of ZF can all be viewed as formulas in the *language of set theory*, namely $L := \{\in\}$. For example, let $\phi(x, y)$ be the formula saying that y is the successor of x . That is, $\phi(x, y)$ is the formula

$$(x \in y) \wedge \forall z(z \in y \rightarrow (z \in x \vee z = x)).$$

Note that for the sake of readability we have dropped some of the parantheses. So actually this formula should be written $((x \in y) \wedge \forall z((z \in y) \rightarrow ((z \in x) \vee (z = x))))$. Now the Axiom of Infinity, which says there is a set containing \emptyset and closed under the successor function, can be expressed using $\phi(x, y)$ as follows:

$$\exists w (\forall v(\forall z(z \notin v) \rightarrow (v \in w)) \wedge \forall x\forall y((x \in w \wedge \phi(x, y)) \rightarrow y \in w)).$$

Note that we use x, y, z , etc., for variable symbols. This will be common.

This is a good time to point out that a definite condition in set theory is precisely one that can be expressed by a formula in the language of set theory (just compare Definition 1.6 with Definition 4.13). Thus the Separation Axiom is in fact an axiom *scheme*; for each formula $\psi(z)$ we have the axiom $\forall x\exists y\forall z((z \in y) \leftrightarrow ((z \in x) \wedge \psi(z)))$.

DEFINITION 4.15. Suppose ϕ is an *L-formula*. An occurrence of a variable x in ϕ is said to be *bound* if it appears inside the scope of a quantifier $\exists x$ or $\forall x$. If the occurrence is not bound then it is said to be *free*. Sometimes we write ϕ as $\phi(x_1, \dots, x_n)$ to mean that the variables which occur freely in ϕ all come from the list x_1, \dots, x_n . (Though not all of these variable need appear freely in ϕ .) A formula in which all variable symbols always occur bound is called a *sentence*.

So in the formula $\phi(x, y)$ of Example 4.14 all the occurrences of x and y are free while all the occurrences of z are bound. Still working in Example 4.14, the formula expressing the Axiom of Infinity is a sentence. In the formula $((x > 0) \vee \exists x(x \times x = 1))$ the first occurrence of x is free while the other occurrences are all bound.

REMARK 4.16. As a simplification we will assume that no variable occurs both free and bound in the same formula; so the formula $((x > 0) \vee \exists x(x \times x = 1))$ discussed above would be implicitly replaced by $((x > 0) \vee \exists z(z \times z = 1))$.

4.4. Truth and satisfaction

Until now terms and formulas are officially just certain strings of symbols; they do not *mean* anything. In this section we describe how, given an L -structure \mathcal{M} , it makes sense to interpret what L -terms and L -formulas mean in \mathcal{M} . We begin with terms. These will be interpreted as functions on the universe of \mathcal{M} .

DEFINITION 4.17 (Interpreting terms). Suppose \mathcal{M} is an L -structure with universe M , and $t = t(x_1, \dots, x_n)$ is an L -term. Then the *interpretation of $t(x_1, \dots, x_n)$ in \mathcal{M}* is the function $t^{\mathcal{M}} : M^n \rightarrow M$ defined inductively as follows:

- (i) If t is a variable symbol x_i for some $1 \leq i \leq n$, then $t^{\mathcal{M}} : M^n \rightarrow M$ is the function $(a_1, \dots, a_n) \mapsto a_i$.
- (ii) If t is a constant symbol c , then $t^{\mathcal{M}} : M^n \rightarrow M$ is the function $(a_1, \dots, a_n) \mapsto c^{\mathcal{M}}$.
- (iii) If $t = f(t_1, \dots, t_{n_f})$ where $f \in L^{\text{fun}}$ and $t_1(x_1, \dots, x_n), \dots, t_{n_f}(x_1, \dots, x_n)$ are L -terms, then $t^{\mathcal{M}} : M^n \rightarrow M$ is the function

$$(a_1, \dots, a_n) \mapsto f^{\mathcal{M}}(t_1^{\mathcal{M}}(a_1, \dots, a_n), \dots, t_{n_f}^{\mathcal{M}}(a_1, \dots, a_n)).$$

Note that the function $t^{\mathcal{M}}$ depends not only on the term t but on its presentation as $t = t(x_1, \dots, x_n)$.

For example, the interpretation of the term $(v_0 + v_1) - 1$, presented as $t(v_1, v_2)$, in the structure $(\mathbb{R}, 0, 1, +, -, \times)$, is the function $(a, b) \mapsto a + b - 1$.

EXERCISE 4.18. Suppose $L = \{0, 1, +, -, \times\}$ is the language of rings. Show that the L -terms are the polynomials with integer coefficients. More precisely, for every L -term $t = t(x_1, \dots, x_n)$ there exists a polynomial $P_t \in \mathbb{Z}[X_1, \dots, X_n]$ such that, for every ring \mathcal{R} , viewed in the natural way as an L -structure, $t^{\mathcal{R}} = P_t$ as functions on R^n . Conversely, every polynomial in $\mathbb{Z}[X_1, \dots, X_n]$ is of the form P_t for some L -term $t = t(x_1, \dots, x_n)$.

EXERCISE 4.19. Suppose F is a field and $L = \{0, +, -, \lambda_a : a \in F\}$ is the language of F -vector spaces. What, in the spirit of Exercise 4.18, do the L -terms look like?

DEFINITION 4.20 (Satisfaction). Suppose \mathcal{M} is an L -structure with universe M , $\phi = \phi(x)$ is an L -formula where $x = (x_1, \dots, x_n)$, and $a = (a_1, \dots, a_n) \in M^n$. We define $\mathcal{M} \models \phi(a)$ inductively as follows:

- (i) If ϕ is $(t_1 = t_2)$ where $t_1(x)$ and $t_2(x)$ are L -terms, then $\mathcal{M} \models \phi(a)$ means that $t_1^{\mathcal{M}}(a) = t_2^{\mathcal{M}}(a)$.
- (ii) If ϕ is $R(t_1, \dots, t_{k_R})$ where $R \in L^{\text{rel}}$ and $t_1(x), \dots, t_{k_R}(x)$ are L -terms, then $\mathcal{M} \models \phi(a)$ means that $(t_1^{\mathcal{M}}(a), \dots, t_{k_R}^{\mathcal{M}}(a)) \in R^{\mathcal{M}}$.
- (iii) If ϕ is $\neg\psi$ where $\psi(x)$ is an L -formula, then $\mathcal{M} \models \phi(a)$ means that $\mathcal{M} \not\models \psi(a)$.
- (iv) If ϕ is $(\psi \wedge \theta)$ where $\psi(x)$ and $\theta(x)$ are L -formulas, then $\mathcal{M} \models \phi(a)$ means that $\mathcal{M} \models \psi(a)$ and $\mathcal{M} \models \theta(a)$.
- (v) If ϕ is $(\psi \vee \theta)$ where $\psi(x)$ and $\theta(x)$ are L -formulas, then $\mathcal{M} \models \phi(a)$ means that $\mathcal{M} \models \psi(a)$ or $\mathcal{M} \models \theta(a)$.
- (vi) If ϕ is $\exists z\psi$ where z is a variable symbol and $\psi(x, z)$ is an L -formula, then $\mathcal{M} \models \phi(a)$ means that there exists $b \in M$ such that $\mathcal{M} \models \psi(a, b)$.
- (vii) If ϕ is $\forall z\psi$ where z is a variable symbol and $\psi(x, z)$ is an L -formula, then $\mathcal{M} \models \phi(a)$ means that $\mathcal{M} \models \psi(a, b)$ for all $b \in M$.

If $\mathcal{M} \models \phi(a)$ then we say that \mathcal{M} *satisfies* $\phi(a)$ or that $\phi(a)$ *is true in* \mathcal{M} , or that a *realises* $\phi(x)$ *in* \mathcal{M} . The set of all realisation of ϕ in \mathcal{M} , $\{a \in M^n : \mathcal{M} \models \phi(a)\}$, is denoted by $\phi^{\mathcal{M}}$ and is called *the set defined by ϕ in \mathcal{M}* .

Let us consider the case when when $n = 0$, that is when ϕ is a sentence. Since, by convention, $M^0 = \{\emptyset\}$, the only question in this case is whether \emptyset realises ϕ or not. If it does then we say that ϕ *is true in* \mathcal{M} and write $\mathcal{M} \models \phi$, otherwise ϕ *is false in* \mathcal{M} . Note that a sentence is false in \mathcal{M} if and only if its negation is true.

EXAMPLE 4.21. Suppose $\phi(y)$ is $\exists x(x^2 = y)$, where x^2 is an abbreviation for $x \times x$. Then $(\mathbb{R}, 0, 1, +, -, \times) \models \phi(2)$ while $(\mathbb{Q}, 0, 1, +, -, \times) \models \neg\phi(2)$ and $(\mathbb{C}, 0, 1, +, -, \times) \models \forall y\phi(y)$.

The following proposition shows that satisfaction for formulas without quantifiers – we call them *quantifier-free formulas* – is inherited by substructures and extensions.

PROPOSITION 4.22. *Suppose $\mathcal{M} \subseteq \mathcal{N}$ are L -structures with universes $M \subseteq N$, $\phi = \phi(x_1, \dots, x_n)$ is an L -formula, and $a \in M^n$.*

- (a) *If ϕ is quantifier-free then $\mathcal{M} \models \phi(a)$ if and only if $\mathcal{N} \models \phi(a)$.*
- (b) *If ϕ of the form $\exists y_1 \dots \exists y_m \psi$ for some quantifier-free L -formula ψ , then $\mathcal{M} \models \phi(a)$ implies $\mathcal{N} \models \phi(a)$.*
- (c) *If ϕ of the form $\forall y_1 \dots \forall y_m \psi$ for some quantifier-free L -formula ψ , then $\mathcal{N} \models \phi(a)$ implies $\mathcal{M} \models \phi(a)$.*

PROOF. This proposition has a very typical proof. In order to prove something about all formulas one usually has to begin by proving something about terms and then proceeding by induction on the complexity of the formula. The result about terms is itself usually proved by induction on the complexity of the term.

Let $x = (x_1, \dots, x_n)$. We first observe that *if $t = t(x)$ is any L -term then $t^{\mathcal{M}} = t^{\mathcal{N}} \upharpoonright M^n$* . Indeed, we prove this by induction on the complexity of t . If t is a constant or variable symbol then this is clear. Suppose $t = f(t_1, \dots, t_{n_f})$, where $f \in L^{\text{fun}}$ and $t_1(x), \dots, t_{n_f}(x)$ are L -terms for which the result is known. Then for any $a \in M^n$,

$$\begin{aligned}
t^{\mathcal{M}}(a) &= f^{\mathcal{M}}(t_1^{\mathcal{M}}(a), \dots, t_{n_f}^{\mathcal{M}}(a)) \quad \text{by definition of the interpretation of terms} \\
&= f^{\mathcal{M}}(t_1^{\mathcal{N}}(a), \dots, t_{n_f}^{\mathcal{N}}(a)) \quad \text{by the induction hypothesis} \\
&= f^{\mathcal{N}}(t_1^{\mathcal{N}}(a), \dots, t_{n_f}^{\mathcal{N}}(a)) \quad \text{since } f^{\mathcal{M}} = f^{\mathcal{N}} \upharpoonright M^{n_f}, \text{ see Remark 4.9(b)} \\
&= t^{\mathcal{N}}(a)
\end{aligned}$$

as desired.

We now prove part(a) by induction on the complexity of the quantifier-free formula ϕ . If ϕ is of the form $(t_1 = t_2)$ for some L -terms $t_1(x)$ and $t_2(x)$, then

$$\begin{aligned}
\mathcal{M} \models \phi(a) &\iff t_1^{\mathcal{M}}(a) = t_2^{\mathcal{M}}(a) \\
&\iff t_1^{\mathcal{N}}(a) = t_2^{\mathcal{N}}(a) \quad \text{since } t_i^{\mathcal{M}} = t_i^{\mathcal{N}} \upharpoonright M^n, \text{ for } i = 1, 2 \\
&\iff \mathcal{N} \models \phi(a).
\end{aligned}$$

If ϕ is of the form $R(t_1, \dots, t_{k_R})$ for some $R \in L^{\text{rel}}$ and $t_1(x), \dots, t_{k_R}(x)$ L -terms, then

$$\begin{aligned}
\mathcal{M} \models \phi(a) &\iff (t_1^{\mathcal{M}}(a), \dots, t_{k_R}^{\mathcal{M}}(a)) \in R^{\mathcal{M}} \\
&\iff (t_1^{\mathcal{N}}(a), \dots, t_{k_R}^{\mathcal{N}}(a)) \in R^{\mathcal{M}} \text{ since } t_i^{\mathcal{M}} = t_i^{\mathcal{N}} \upharpoonright M^n, \text{ for } i = 1, \dots, k_R \\
&\iff (t_1^{\mathcal{N}}(a), \dots, t_{k_R}^{\mathcal{N}}(a)) \in R^{\mathcal{N}} \text{ since } R^{\mathcal{M}} = R^{\mathcal{N}} \cap M^{k_r}, \text{ see Remark 4.9(b)} \\
&\iff \mathcal{N} \models \phi(a).
\end{aligned}$$

Now suppose $\psi(x)$ and $\theta(x)$ are quantifier-free L -formulas for which the result is known. If ϕ is $\neg\psi$ then

$$\begin{aligned}
\mathcal{M} \models \phi(a) &\iff \mathcal{M} \not\models \psi(a) \\
&\iff \mathcal{N} \not\models \psi(a) \text{ by the induction hypothesis} \\
&\iff \mathcal{N} \models \phi(a).
\end{aligned}$$

If ϕ is $(\psi \wedge \theta)$ then

$$\begin{aligned}
\mathcal{M} \models \phi(a) &\iff \mathcal{M} \models \psi(a) \text{ and } \mathcal{M} \models \theta(a) \\
&\iff \mathcal{N} \models \psi(a) \text{ and } \mathcal{N} \models \theta(a) \text{ by the induction hypothesis} \\
&\iff \mathcal{N} \models \phi(a).
\end{aligned}$$

If ϕ is $(\psi \vee \theta)$ then

$$\begin{aligned}
\mathcal{M} \models \phi(a) &\iff \mathcal{M} \models \psi(a) \text{ or } \mathcal{M} \models \theta(a) \\
&\iff \mathcal{N} \models \psi(a) \text{ or } \mathcal{N} \models \theta(a) \text{ by the induction hypothesis} \\
&\iff \mathcal{N} \models \phi(a).
\end{aligned}$$

Since ϕ is a quantifier-free formula, this completes the induction.

To prove part (b) we write $\phi(x) = \exists y\psi(x, y)$ where $y = (y_1, \dots, y_m)$ and ψ is quantifier-free. Then

$$\begin{aligned}
\mathcal{M} \models \phi(a) &\implies \text{there exists } b \in M^m \text{ such that } \mathcal{M} \models \psi(a, b) \\
&\implies \text{there exists } b \in M^m \text{ such that } \mathcal{N} \models \psi(a, b) \text{ by part (a)} \\
&\implies \text{there exists } b \in N^m \text{ such that } \mathcal{N} \models \psi(a, b) \text{ as } M \subseteq N \\
&\implies \mathcal{N} \models \phi(a).
\end{aligned}$$

To prove part (c), write $\phi(x) = \forall y\psi(x, y)$. Then

$$\begin{aligned}
\mathcal{N} \models \phi(a) &\implies \text{for all } b \in N^m, \mathcal{N} \models \psi(a, b) \\
&\implies \text{for all } b \in M^m, \mathcal{N} \models \psi(a, b) \text{ as } M \subseteq N \\
&\implies \text{for all } b \in M^m, \mathcal{M} \models \psi(a, b) \text{ by part (a)} \\
&\implies \mathcal{M} \models \phi(a).
\end{aligned}$$

This completes the proof of the proposition. \square

The formulas of the form $\exists y_1 \cdots \exists y_m \psi$ where ψ is quantifier-free are called *existential formulas*, while those of the form $\forall y_1 \cdots \forall y_m \psi$ are called *universal formulas*.

4.5. Elementary embeddings

While Proposition 4.22 is about $\mathcal{M} \subseteq \mathcal{N}$, the proof clearly goes through for any L -embedding $j : \mathcal{M} \rightarrow \mathcal{N}$. We can strengthen the notion of embedding to force 4.22(a) to hold for all formulas ϕ , and not just quantifier-free ones.

DEFINITION 4.23 (Elementary embedding). Suppose \mathcal{M} and \mathcal{N} are L -structures with universes M and N , respectively. An L -embedding $j : \mathcal{M} \rightarrow \mathcal{N}$ is called an *elementary embedding* if for all L -formulas $\phi(x_1, \dots, x_n)$ and all n -tuples $a \in M^n$, $\mathcal{M} \models \phi(a)$ if and only if $\mathcal{N} \models \phi(j(a))$.

If $M \subseteq N$ and the containment map is an elementary embedding, then we say that \mathcal{M} is an *elementary substructure* of \mathcal{N} , or that \mathcal{N} is an *elementary extension* of \mathcal{M} ; and we denote this by $\mathcal{M} \preceq \mathcal{N}$.

COROLLARY 4.24. *Every isomorphism is an elementary embedding.*

PROOF. Suppose $f : \mathcal{M} \rightarrow \mathcal{N}$ is an L -isomorphism between L -structures. We need to show that for any formula $\phi(x)$ where $x = (x_1, \dots, x_n)$ and any tuple $a \in M^n$, $\mathcal{M} \models \phi(a)$ if and only if $\mathcal{N} \models \phi(f(a))$. This is done by induction on the complexity of ϕ . As usual, to deal with the atomic case we need to prove something about terms. Namely: *For all L -terms $t(x)$, $f(t^{\mathcal{M}}(a)) = t^{\mathcal{N}}(f(a))$.* This in turn is done by induction on the complexity of the term t , and uses only the fact that f is an L -embedding. We leave it to the reader.¹ The case of ϕ atomic follows easily, as well as the inductive steps corresponding to \neg, \wedge and \vee ; see the proof of Proposition 4.22(a). Since we can write \forall as $\neg\exists\neg$, it remains to consider the case when $\phi(x) = \exists y\psi(x, y)$. In that case

$$\begin{aligned} \mathcal{M} \models \phi(a) &\iff \mathcal{M} \models \psi(a, b) \text{ for some } b \in M \\ &\iff \mathcal{N} \models \psi(f(a), f(b)) \text{ for some } b \in M && \text{by the induction hypothesis} \\ &\iff \mathcal{N} \models \psi(f(a), c) \text{ for some } c \in N && \text{as } f \text{ is surjective} \\ &\iff \mathcal{N} \models \phi(f(a)) \end{aligned}$$

as desired. □

EXERCISE 4.25. Show that $j : \mathcal{M} \rightarrow \mathcal{N}$ is an elementary embedding of \mathcal{M} in \mathcal{N} if and only if it induces an isomorphism between \mathcal{M} and an elementary substructure of \mathcal{N} . The study of elementary L -embeddings thereby reduces to that of elementary L -substructures.

By virtue of Proposition 4.22(a) the difference between substructures and elementary substructures can only be seen by considering formulas with quantifiers. For example, $(\mathbb{Z}, 0, +, -) \subseteq (\mathbb{Q}, 0, +, -)$ as the integers form a subgroup of the rationals, but $(\mathbb{Z}, 0, +, -) \not\preceq (\mathbb{Q}, 0, +, -)$ since if $\phi(x)$ is $\exists y(y + y = x)$ then $(\mathbb{Z}, 0, +, -) \models \neg\phi(1)$ while $(\mathbb{Q}, 0, +, -) \models \phi(1)$. This example can be restated as saying that the equation $y + y = 1$ has a solution in \mathbb{Q} but not in \mathbb{Z} . In fact, this is the typical way in which substructures fail to be elementary substructures, as the following proposition explains:

PROPOSITION 4.26 (Tarski-Vaught Test). *Suppose $\mathcal{M} \subseteq \mathcal{N}$ with universes $M \subseteq N$. Then the following are equivalent:*

¹In fact, when f is the inclusion map we have already verified this claim about terms in the proof of Proposition 4.22(a).

- (i) $\mathcal{M} \preceq \mathcal{N}$
- (ii) For every L -formula $\phi(x_1, \dots, x_n, y)$ and all n -tuples $a \in M^n$, if $\mathcal{N} \models \exists y\phi(a, y)$ then there exists $b \in M$ such that $\mathcal{N} \models \phi(a, b)$.

PROOF. Suppose $\mathcal{M} \preceq \mathcal{N}$, and $\mathcal{N} \models \exists y\phi(a, y)$ as in the statement of (ii). Let $\psi(x) = \exists y\phi(x, y)$. Since $\mathcal{N} \models \psi(a)$, $\mathcal{M} \models \psi(a)$. The latter means that there is a $b \in M$ such that $\mathcal{M} \models \phi(a, b)$. Applying the definition of an elementary substructure again, we have that $\mathcal{N} \models \phi(a, b)$, as desired.

For the converse we assume that (ii) holds and show by induction on the complexity of formulas $\psi(x_1, \dots, x_n)$ that for any $a \in M^n$, $\mathcal{M} \models \psi(a)$ if and only if $\mathcal{N} \models \psi(a)$. Since $\mathcal{M} \subseteq \mathcal{N}$, this is true of all atomic (indeed quantifier-free) formulas ψ by Proposition 4.22(a). The induction steps with regard to \wedge, \vee, \neg are straightforward. Since \forall can be written as $\neg\exists\neg$, it suffices to consider the case when ψ is of the form $\exists y\phi(x_1, \dots, x_n, y)$, and the result is known for ϕ . Now if $\mathcal{M} \models \psi(a)$ then there exists $b \in M$ such that $\mathcal{M} \models \phi(a, b)$ and so by the inductive hypothesis $\mathcal{N} \models \phi(a, b)$, and so $\mathcal{N} \models \psi(a)$. Now (ii) tells us that there is a $b \in M$ with $\mathcal{N} \models \phi(a, b)$, and so by the inductive hypothesis again we have $\mathcal{M} \models \phi(a, b)$, which implies that $\mathcal{M} \models \psi(a)$, as desired. \square

I leave it to you to work out the Tarski-Vaught test for elementary embeddings (rather than just elementary substructures as above).

The following is maybe one of the first theorems in model theory (circa 1915).

THEOREM 4.27 (Downward Löwenheim-Skolem). *Suppose \mathcal{M} is an L -structure and $A \subseteq M$. Then there exists an elementary substructure of \mathcal{M} that contains A and is of cardinality at most $\max\{|A|, |L|, \aleph_0\}$.*

In particular, if L is countable then every L -structure has a countable elementary substructure.

PROOF. Let $\kappa = \max\{|A|, |L|, \aleph_0\}$. We define recursively a countable chain of subsets of M cardinality $\leq \kappa$, $A = A_0 \subseteq A_1 \subseteq \dots$ such that: for each $n \geq 0$, if $\phi(\bar{x}, y)$ is any L -formula and \bar{a} is any tuple from A_n with $\mathcal{M} \models \exists y\phi(\bar{a}, y)$, then there exists $b \in A_{n+1}$ with $\mathcal{M} \models \phi(\bar{a}, b)$. Given A_n we show how to construct A_{n+1} . First observe that an L -formula is a finite string from a set of symbols of size $|L| + \aleph_0 \leq \kappa$, and there are only at most κ -many such finite strings since κ is infinite (see Corollary 3.14 and Proposition 3.20). Similarly there are at most κ -many finite tuples from A_n since $|A_n| \leq \kappa$. Hence the set of all pairs $(\phi(\bar{x}, y), \bar{a})$ where $\phi(\bar{x}, y)$ is an L -formula and \bar{a} is a tuple from A_n , has cardinality at most κ . For each such pair with $\mathcal{M} \models \exists y\phi(\bar{a}, y)$, we choose a realisation and include it in A_{n+1} . Note that $A_n \subseteq A_{n+1}$ (consider the pairs $(y = x, a)$ for each $a \in A_n$), $|A_{n+1}| \leq \kappa$, and A_{n+1} satisfies the desired property.

Having constructed the chain, let $B := \bigcup_{n \in \omega} A_n$. Then $|B| \leq \kappa$ too. Considering the formulas $y = c$ for each constant symbol, we see that A_1 , and hence B , contains all the constants of \mathcal{M} . For every function symbol f of arity m , and m -tuple \bar{a} from B , taking n big enough so that \bar{a} comes from A_n and considering the pair $(y = f(\bar{x}), \bar{a})$ shows that $f(\bar{a}) \in A_{n+1} \subseteq B$. It follows by Exercise 4.10 that B is the universe of a substructure $\mathcal{N} \subseteq \mathcal{M}$. To see that $\mathcal{N} \preceq \mathcal{M}$ we apply the Tarski-Vaught test: if $\mathcal{M} \models \exists y\phi(\bar{a}, y)$ for some \bar{a} from B , then taking n large enough so that \bar{a} comes from A_n , we get a realisation in $A_{n+1} \subseteq B$, so there exists $b \in B$ such that $\mathcal{M} \models \phi(\bar{a}, b)$. \square

4.6. Definable sets and parameters

Suppose \mathcal{M} is an L -structure with universe M , and $B \subseteq M$. Then by L_B we mean the language obtained by adding to L a new constant symbol, \underline{b} , for each $b \in B$. We can extend \mathcal{M} to an L_B -structure in a natural way, denoted sometimes by \mathcal{M}_B , by interpreting $\underline{b}^{\mathcal{M}_B} = b$. This process is called “naming constants.” Often we drop the underscore and rely on context to distinguish between $b \in B$ and $b \in L_B^{\text{con}}$. Also, as the construction of \mathcal{M}_B from \mathcal{M} is canonical, we sometimes drop the subscript and rely on context to determine whether \mathcal{M} is being viewed as an L -structure or an L_B -structure.

EXERCISE 4.28. Suppose $B \subseteq M$ is a nonempty subset of an L -structure \mathcal{M} . Show that B is the universe of an elementary L -substructure of \mathcal{M} if and only if every L_B -formula in one variable that is realised in \mathcal{M} has a realisation in the set B .

DEFINITION 4.29 (Definable set). A set $X \subseteq M^n$ is *definable over B* (or *B -definable*) in \mathcal{M} if there exists an L_B -formula $\phi(x_1, \dots, x_n)$, such that

$$X = \{a \in M^n : \mathcal{M}_B \models \phi(a)\}.$$

In this case we write $X = \phi^{\mathcal{M}}$ and say that ϕ *defines* X . We say that X is *definable* if it is M -definable and that it is *0-definable* if it is \emptyset -definable. We say that X is *quantifier-free* (respectively *existentially* or *universally*) definable if there is a quantifier-free (respectively existential or universal) formula ϕ such that $X = \phi^{\mathcal{M}}$.

REMARK 4.30. Note that if $\phi(x_1, \dots, x_n)$ is an L_B -formula then there exists an L -formula $\psi(x_1, \dots, x_n, y_1, \dots, y_m)$ and a tuple $b \in B^m$, such that $\phi(x_1, \dots, x_n) = \psi(x_1, \dots, x_n, b)$. Hence X is B -definable if and only if $X = \{a \in M^n : \mathcal{M} \models \psi(a, b)\}$ for some L -formula $\psi(x_1, \dots, x_n, y_1, \dots, y_m)$ and some $b \in B^m$.

Let us consider a few examples.

EXAMPLE 4.31 (Zero sets of polynomials). Suppose $\mathcal{M} = (F, 0, 1, +, -, \times)$ where F is a field. Then for any finite set of polynomials $p_1, \dots, p_\ell \in F[X_1, \dots, X_n]$, the zero set,

$$V(p_1, \dots, p_\ell) := \{a \in F^n : p_1(a) = \dots = p_\ell(a) = 0\}$$

is a quantifier-free definable set in \mathcal{M} . Such zero sets are called *algebraic sets* and are the main objects of study in algebraic geometry. They form the closed sets of a topology on F^n called the *Zariski topology*. A subset of F^n is *Zariski-constructible* if it is a finite boolean combination of algebraic sets – that is, if it is obtained in finitely many steps from algebraic sets by taking unions, intersections and complements. Given that the terms in the language of rings are just polynomials (cf. Exercise 4.18) it is not hard to see that the quantifier-free definable sets in \mathcal{M} are precisely the Zariski-constructible sets. It is a much harder problem to describe the definable sets of \mathcal{M} in general. Indeed, without additional “tameness” assumption on the field F this is in some precise sense impossible. On the other hand, we will see in Section 6.2 that if F is an algebraically closed field then *every* definable set in \mathcal{M} is Zariski-constructible (Tarski’s Theorem).²

²By a theorem of Angus Macintyre from the seventies the converse is also true, if F is a field such that every definable set in $\mathcal{M} = (F, 0, 1, +, -, \times)$ is Zariski-constructible, then F is an algebraically closed field.

EXAMPLE 4.32 (Ordering in the ring of reals). Suppose $\mathcal{M} = (\mathbb{R}, 0, 1, +, -, \times)$ and let $\phi(x, y)$ be the formula

$$\exists z((z \neq 0) \wedge (y = x + z^2)).$$

Note that $a < b$ if and only if $\mathcal{M} \models \phi(a, b)$. So $\phi^{\mathcal{M}} \subset \mathbb{R}^2$ is the ordering on \mathbb{R} . This answers (just by making precise) a question we posed at the beginning of the chapter: the ordered ring of reals is recoverable from the ring structure in the sense that the ordering is definable in the ring structure. But note that the ordering is *not* quantifier-free definable: otherwise the set of positive reals would be Zariski-constructible (cf. Example 4.31), but as nonzero polynomials in one variable can have only finitely many zeros all Zariski-constructible subsets of the real line are either finite or cofinite.

EXAMPLE 4.33 (Ordering in the ring of integers). Suppose $\mathcal{M} = (\mathbb{Z}, 0, 1, +, -, \times)$ and let $\phi(x, y)$ be the formula

$$\exists z_1 \exists z_2 \exists z_3 \exists z_4 ((z_1 \neq 0) \wedge (y = x + z_1^2 + z_2^2 + z_3^2 + z_4^2)).$$

By Lagrange's theorem an integer is positive if and only if it is the sum of four squares. Hence, $m < n$ if and only if $\mathcal{M} \models \phi(m, n)$. So the ordering is also 0-definable in the integers.

EXAMPLE 4.34 (Constants in the ring of polynomials). Suppose $\mathcal{M} = (K[X], 0, 1, +, -, \times)$ where K is a field and $X = (X_1, \dots, X_m)$ is a sequence of indeterminates. So we are dealing with the ring of polynomials in m variables over the field K . Then K is definable in \mathcal{M} . Indeed, K is the set of units in $K[X]$, and hence is defined by $(x \neq 0) \vee (\exists z(xz = 1))$.

EXAMPLE 4.35. Suppose $\mathcal{M} = (M, \in)$ where M is a set. Let $\phi(x, y)$ be the formula of Example 4.14, namely $(x \in y) \wedge \forall z(z \in y \rightarrow (z \in x \vee z = x))$. Let $\psi(y)$ be $\exists x \phi(x, y)$. What set does $\psi(y)$ define? At first glance it might seem that ψ defines the set of elements in M that are successors of elements in M . But since the variables range only over the universe M , this is not exactly correct. Looking closer we see that $\mathcal{M} \models \psi(a)$ if and only if $a \cap M = (b \cap M) \cup \{b\}$ for some $b \in M$. If every member of M were a subset of M (for example, if M were an ordinal) then indeed $\psi^{\mathcal{M}} = \{a \in M : a = S(b) \text{ for some } b \in M\}$.

How can we show that some set is not definable?

EXERCISE 4.36. Suppose \mathcal{M} is an L -structure and $\kappa = \max\{\aleph_0, |L|, |M|\}$. Show that there are at most κ -many definable sets in \mathcal{M} . Deduce that if \mathcal{M} is infinite and of cardinality at least that of the language, then there exist non-definable subsets of M^n for all $n > 0$.

The above exercise is just counting, but the following lemma gives us one tool for actually showing that a set is not definable.

LEMMA 4.37. *If $X \subset M^n$ is B -definable and f is an L -automorphism of \mathcal{M} fixing B pointwise, then $f(X) = X$.*

PROOF. Suppose $X = \phi^{\mathcal{M}}$ where $\phi(x_1, \dots, x_n)$ is an L_B -formula. Then $\phi(x) = \psi(x, b)$ for some tuple $b \in B^m$ and L -formula ψ . Now

$$\begin{aligned} a \in X &\iff \mathcal{M} \models \psi(a, b) \\ &\iff \mathcal{M} \models \psi(f(a), f(b)) \quad \text{by Corollary 4.24} \\ &\iff \mathcal{M} \models \psi(f(a), b) \quad \text{as } f(b) = b \\ &\iff f(a) \in X \end{aligned}$$

as desired. \square

Together with some knowledge about the automorphisms of a structure, this can be used to produce non-definable sets.

COROLLARY 4.38. *The real numbers are not definable in $(\mathbb{C}, 0, 1, +, -, \times)$.*

PROOF. This requires a little bit of field theory. Suppose toward a contradiction that \mathbb{R} is definable in \mathbb{C} . Then it would be definable over some finite set B . Since the transcendence degree of \mathbb{R} over \mathbb{Q} is infinite, there exists $r \in \mathbb{R} \setminus \mathbb{Q}(B)^{\text{alg}}$. Now choose $s \in \mathbb{C} \setminus \mathbb{Q}(B, r)^{\text{alg}}$ such that s is not real (check that this is possible!). Since s and r are algebraically independent over $\mathbb{Q}(B)$, there is an L -isomorphism $f : \mathbb{Q}(B, r) \rightarrow \mathbb{Q}(B, s)$ that takes r to s and fixes $\mathbb{Q}(B)$ pointwise. This isomorphism can be extended to an automorphism of the complex field (again something to check). But then $f(\mathbb{R}) \neq \mathbb{R}$ contradicting Lemma 4.37. \square

EXERCISE 4.39. Show that the addition is not definable in $(\mathbb{R}, <)$.

This method of proving that something is not definable relies on the existence of many automorphisms. When a structure does not have many automorphisms one has to work harder to understand which sets are definable and which are not. For example, $(\mathbb{R}, 0, 1, +, -, \times)$ has no automorphisms except the identity (why?). So to prove, for example, that the integers are not definable in this structure (and they aren't) requires a more thorough understanding of what the definable sets look like. One has to first prove a “quantifier elimination theorem”. We will return to this important theme in Chapter 6.

We conclude this section by characterising the definable sets of a structure in a way that avoids logic entirely; that is, without making any reference to formulas and satisfaction.

PROPOSITION 4.40. *Suppose \mathcal{M} is an L -structure, let $\text{Def}_n(\mathcal{M})$ be the collection of all definable subsets of M^n in \mathcal{M} , and let $\text{Def}(\mathcal{M}) := \bigcup_n \text{Def}_n(\mathcal{M})$. Then $\text{Def}(\mathcal{M})$ satisfies the following closure properties:*

- (i) For each n -ary function symbol f , the graph of $f^{\mathcal{M}}$ is in $\text{Def}_{n+1}(\mathcal{M})$.
- (ii) For each n -ary relation symbol R , $R^{\mathcal{M}}$ is in $\text{Def}_n(\mathcal{M})$.
- (iii) For all $i, j \leq n$, $\Delta_{i,j}^{(n)} := \{(a_1, \dots, a_n) \in M^n : a_i = a_j\} \in \text{Def}_n(\mathcal{M})$.
- (iv) If $X \in \text{Def}_n(\mathcal{M})$ then $X \times M \in \text{Def}_{n+1}(\mathcal{M})$.
- (v) Each $\text{Def}_n(\mathcal{M})$ is closed under complements, unions, and intersections.
- (vi) If $X \in \text{Def}_{n+1}(\mathcal{M})$ and $\pi : M^{n+1} \rightarrow M^n$ is the projection onto the first n coordinates, then $\pi(X) \in \text{Def}_n(\mathcal{M})$.
- (vii) If $X \in \text{Def}_{n+m}(\mathcal{M})$ and $b \in M^m$ then $X_b := \{a \in M^n : (a, b) \in X\} \in \text{Def}_n(\mathcal{M})$.

Moreover, $\text{Def}(\mathcal{M})$ is the smallest collection of subsets of cartesian powers of M satisfying (i) through (vii).

PROOF. We first show that $\text{Def}(\mathcal{M})$ satisfies (i) through (vii):

- (i) The graph of $f^{\mathcal{M}}$ is defined by $(f(x_1, \dots, x_n) = x_{n+1})$.
- (ii) $R^{\mathcal{M}}$ is defined by $R(x_1, \dots, x_n)$.
- (iii) $\Delta_{i,j}^{(n)} = \phi^{\mathcal{M}}$ where $\phi(x_1, \dots, x_n)$ is the formula $(x_i = x_j)$.
- (iv) If X is defined by $\phi(x_1, \dots, x_n)$ then $X \times M$ is defined by $\psi(x_1, \dots, x_n, x_{n+1})$, where $\psi = \phi$.

- (v) If X is defined by $\phi(x_1, \dots, x_n)$ and Y is defined by $\psi(x_1, \dots, x_n)$ then $M^n \setminus X$ is defined by $\neg\phi$, $X \cup Y$ is defined by $\phi \vee \psi$, and $X \cap Y$ is defined by $\phi \wedge \psi$.
- (vi) If X is defined by $\phi(x_1, \dots, x_{n+1})$ then $\pi(X)$ is defined by $\exists x_{n+1}\phi(x_1, \dots, x_n, x_{n+1})$.
- (vii) If X is defined by $\phi(x_1, \dots, x_n, y_1, \dots, y_m)$ then X_b is defined by $\phi(x_1, \dots, x_n, \underline{b})$.

Next we show that if $D = \bigcup_n D_n$ with $D_n \subseteq \mathcal{P}(M^n)$ satisfies (i) through (vii), then all definable sets are in D . One thing that will be useful is to note that (iii), (iv) and (vi) imply D is preserved by co-ordinate permutations. We begin with the claim

CLAIM 4.41. *For each L -term $t(x_1, \dots, x_n)$, the graph of $t^{\mathcal{M}}$ is in D_{n+1} .*

PROOF OF CLAIM 4.41. This is by induction on the complexity of the term t . If t is x_i then the graph of $t^{\mathcal{M}}$ is $\Delta_{i,n+1}^{(n+1)}$. If t is a constant c , then the graph of $t^{\mathcal{M}}$ is $M^n \times \{c^{\mathcal{M}}\} = (M^{n+1})_{c^{\mathcal{M}}} = (\Delta_{1,1}^{(n+1)})_{c^{\mathcal{M}}}$. Finally suppose t is $f(t_1, \dots, t_m)$ where f is an m -ary function symbol and t_1, \dots, t_m are L -terms for which we know the claim. So by repeated use of (iv) we have that $\Gamma(t_i^{\mathcal{M}}) \times M^m \in D_{n+m+1}$. Hence, for each $i \leq m$, by permuting co-ordinates,

$$\{(a_1, \dots, a_n, b_1, \dots, b_m, c) : t_i^{\mathcal{M}}(a_1, \dots, a_n) = b_i\} \in D_{n+m+1}.$$

Also, by (i) and (iv) and permuting co-ordinates again, we have

$$\{(a_1, \dots, a_n, b_1, \dots, b_m, c) : f^{\mathcal{M}}(b_1, \dots, b_m) = c\} \in D_{n+m+1}.$$

Taking intersections and projecting onto the appropriate $n+1$ co-ordinates – by repeated applications of (vi) after a co-ordinate permutation – we get that

$$\Gamma(t^{\mathcal{M}}) = \{(a_1, \dots, a_n, c) : f^{\mathcal{M}}(t_1^{\mathcal{M}}(a_1, \dots, a_n), \dots, t_m^{\mathcal{M}}(a_1, \dots, a_n)) = c\} \in D_{n+1}$$

as desired. □

Now, using properties (i) through (vi) as we have been doing, together with Claim 4.41, it is not hard to show by induction on L -formulas that all 0-definable sets are in D . By (vii), it follows that every definable set is in D . □

4.7. Theories and their models

Logical syntax appears in model theory in two ways. On the one hand it gives us the notion of a definable set in a given structure (via formulas and interpretation). On the other hand, and this is the topic of the present section, formulas (or rather sentences) can be used to axiomatise classes of structures. We will see later that these two aspects of model theory are intimately related and complementary.

DEFINITION 4.42. An L -theory is a set of L -sentences. A *model* of a theory T is an L -structure \mathcal{M} such that $\mathcal{M} \models \sigma$ for each $\sigma \in T$. This is written $\mathcal{M} \models T$. If T has a model then it is said to be *consistent*.

We denote the class of all models of a theory T by $\text{Mod}(T)$. A class \mathcal{K} of L -structures is called *elementary* or *axiomatisable* if there exists an L -theory T such that $\mathcal{K} = \text{Mod}(T)$. In this case we say that T *axiomatises* \mathcal{K} .

Given an L -structure \mathcal{M} , by the *theory of \mathcal{M}* , denoted by $\text{Th}(\mathcal{M})$, we mean the set of all L -sentences true in \mathcal{M} . Two structures \mathcal{M} and \mathcal{N} are *elementarily equivalent*, denoted by $\mathcal{M} \equiv \mathcal{N}$, if $\text{Th}(\mathcal{M}) = \text{Th}(\mathcal{N})$.

Given a theory T and a sentence σ , we write $T \models \sigma$, and say that T *entails* σ , or that σ is a *consequence of T* , if $\mathcal{M} \models \sigma$ for every model \mathcal{M} of T . We say that T is *complete* if for every L -sentence σ , either $T \models \sigma$ or $T \models \neg\sigma$.

Let's consider some examples of theories and their models.

EXAMPLE 4.43. (a) Suppose $L = \{e, \cdot, {}^{-1}\}$ is the language of (multiplicatively written) groups made up of a constant symbol e , a binary function symbol \cdot , and a unary function symbol ${}^{-1}$. The following classes of L -structures are axiomatisable: groups, abelian groups, torsion-free groups, groups of order N (for any fixed N), divisible groups. I leave it to you to write down the L -theories. Let me only point out that the torsion-free groups require infinitely many sentences, including the sentence $\forall x((x = e) \vee \neg(x^n = e))$ for each n , where x^n is just an abbreviation for the product of x with itself n -times. Divisible groups also require infinitely many axioms. We will see later that the class of torsion groups, that is groups where every element is a torsion element, is *not* elementary.

(b) Suppose $L = \{0, 1, +, -, \times\}$ is the language of rings. The following classes of L -structures are elementary: rings, integral domains, fields, algebraically closed fields, algebraically closed fields of characteristic p (for any fixed prime p or $p = 0$). Again I leave it to you to write the appropriate L -sentences, mentioning only that being algebraically closed is expressed by the sentences

$$\forall a_0 \dots \forall a_n \exists x \left(x^n + \sum_{i=0}^{n-1} a_i x^i = 0 \right),$$

one for each $n \geq 1$. Characteristic p for a prime p is expressed by $\forall x(px = 0)$, while characteristic zero requires the negation of all the above sentences as p ranges over all primes. As we will see later, the class of all finite fields is *not* axiomatisable.

- (c) Suppose $L = \{R\}$ is the language consisting of a single binary relation symbol. The following classes of L -structures are elementary: graphs (antireflexivity and symmetry), posets, linear orders, dense linear orders, discrete linear orders (where every element has a least greater element), equivalence relations.
- (d) For any fixed field K the class of K -vector spaces is an elementary class of L -structures where $L = \{0, +, -, \lambda_a : a \in K\}$.
- (e) Suppose $L = \emptyset$. The class of sets and the class of infinite sets are both axiomatisable as L -structures. The former is given by the empty theory, the latter by the infinite theory with sentences saying that there are at least n elements, one for each $n \geq 1$. The class of all finite sets is *not* elementary, as we shall see later.
- (f) Let $L = \{\in\}$. The class of all set of sets (or *universes* of sets) that satisfy the Zermelo-Fraenkel axioms is axiomatisable. Indeed, the Zermelo-Fraenkel axioms can all be expressed as L -sentences. It is not known (and indeed is in some sense impossible to prove) that ZF is consistent.

LEMMA 4.44. (a) *For any L -structure \mathcal{M} , $\text{Th}(\mathcal{M})$ is a complete consistent theory which contains all its consequences.*

- (b) *Suppose T is a consistent theory. The following are equivalent:*
- (i) *T is complete*
 - (ii) *The set of consequences of T is a maximally consistent theory.*
 - (iii) *The set of consequences of T is of the form $\text{Th}(\mathcal{M})$ for some (equivalently any) $\mathcal{M} \models T$.*
 - (iv) *Any two models of T are elementarily equivalent.*
- (c) *If $j : \mathcal{M} \rightarrow \mathcal{N}$ is an elementary embedding then $\mathcal{M} \equiv \mathcal{N}$. In particular, isomorphic structures are elementarily equivalent.*

PROOF. The completeness of $\text{Th}(\mathcal{M})$ follows from the fact that for every σ , either $\mathcal{M} \models \sigma$ or $\mathcal{M} \models \neg\sigma$. Consistency, as well as the fact that $\text{Th}(\mathcal{M})$ contains all its consequences, is immediate from the definitions.

For part (b) let T' be the set of consequences of T . Suppose T is complete. If $T' \subsetneq S$ then there is some $\sigma \in S \setminus T'$. Hence $T \models \neg\sigma$ and so $\sigma, \neg\sigma \in S$. This implies that S has no models. So T' is maximally consistent, as desired. Now assume the maximality of T' and let $\mathcal{M} \models T$. Then $T' \subseteq \text{Th}(\mathcal{M})$ and so by the consistency of $\text{Th}(\mathcal{M})$ we must have $T' = \text{Th}(\mathcal{M})$. We have shown that if T' is maximally consistent then $T' = \text{Th}(\mathcal{M})$ for some $\mathcal{M} \models T$. Now assume $T' = \text{Th}(\mathcal{M})$ for some $\mathcal{M} \models T$. Suppose $\mathcal{N} \models T$. Then $\mathcal{N} \models \text{Th}(\mathcal{M})$, and so $\text{Th}(\mathcal{M}) \subseteq \text{Th}(\mathcal{N})$. So by part (a), $\text{Th}(\mathcal{M}) = \text{Th}(\mathcal{N})$. Finally, assume that every two models of T are elementarily equivalent. Suppose σ is a sentence such that $T \not\models \sigma$. Then there exists $\mathcal{M} \models T$ such that $\mathcal{M} \models \neg\sigma$. As every other model of T is elementarily equivalent to \mathcal{M} , $\neg\sigma$ is true in every model of T . That is $T \models \neg\sigma$. Hence T is complete.

Part (c) follows from the $n = 0$ case of the definition of an elementary embedding (Definition 4.23). Since isomorphisms are elementary embeddings (Corollary 4.24), isomorphic structures are elementarily equivalent. \square

Note that the converse to part (c) is false: there exist elementarily equivalent substructures that are not elementary substructures. For example $(\mathbb{N} \setminus \{0\}, <) \subseteq (\mathbb{N}, <)$ and the map $n \mapsto n - 1$ shows that $(\mathbb{N} \setminus \{0\}, <)$ is isomorphic to, and hence elementarily equivalent to, $(\mathbb{N}, <)$. But $(\mathbb{N} \setminus \{0\}, <)$ is not an elementary substructure since the formula $y < 1$ has a realisation in $(\mathbb{N}, <)$ but not in $(\mathbb{N} \setminus \{0\}, <)$.

In fact, it is easy to see that for $\mathcal{M} \subseteq \mathcal{N}$, we have $\mathcal{M} \preceq \mathcal{N}$ if and only if $\mathcal{M}_M \equiv \mathcal{N}_M$. More generally, the following characterises when a structure can be elementarily embedded in another.

PROPOSITION 4.45. *There exists an elementary embedding of \mathcal{M} into \mathcal{N} if and only if there exists an expansion of \mathcal{N} to an L_M -structure that is a model of $\text{Th}(\mathcal{M}_M)$.*

PROOF. Suppose we have a map $j : M \rightarrow N$. Let \mathcal{N}' be the L_M -structure expanding \mathcal{N} given by $\underline{a}^{\mathcal{N}'} = j(a)$ for all $a \in M$. To say that j is an elementary embedding is precisely to say that $\mathcal{N}' \models \text{Th}(\mathcal{M}_M)$. Conversely, suppose we have an L_M -structure \mathcal{N}' expanding \mathcal{N} . Define $j : \mathcal{M} \rightarrow \mathcal{N}$ by $j(a) := \underline{a}^{\mathcal{N}'}$. To say that $\mathcal{N}' \models \text{Th}(\mathcal{M}_M)$ is precisely to say that j is an elementary embedding. \square

We conclude the chapter by pointing out that for finite structures there is no difference between elementary equivalence and isomorphism. In fact it will follow from the upward Löwenheim-Skolem Theorem (5.12 below) that this is only true of finite structures.

PROPOSITION 4.46. *Suppose \mathcal{M} and \mathcal{N} are finite structures. That is, their universes are finite sets. Then $\mathcal{M} \equiv \mathcal{N}$ if and only if \mathcal{M} and \mathcal{N} are isomorphic.*

The proof of this proposition is a good chance to introduce the following generalisation of the notions of an elementary embedding:

DEFINITION 4.47 (Partial elementary map). Suppose $f : A \rightarrow N$ where \mathcal{M} and \mathcal{N} are L -structures and $A \subseteq M$. We say that f is a *partial elementary map* (or *p.e.m.* for short) if for all L -formulas $\phi(x_1, \dots, x_n)$ and all $a_1, \dots, a_n \in A$,

$$\mathcal{M} \models \phi(a_1, \dots, a_n) \iff \mathcal{N} \models \phi(f(a_1), \dots, f(a_n)).$$

Note that the empty function is a p.e.m. if and only if $\mathcal{M} \equiv \mathcal{N}$, and that a function with domain all of M is a p.e.m. if and only if it is an elementary embedding.

PROOF OF PROPOSITION 4.46. By Lemma 4.44(c), isomorphic structures are elementarily equivalent (without the finiteness assumption). For the converse, enumerate $M = \{a_1, \dots, a_n\}$. The fact that $\mathcal{M} \equiv \mathcal{N}$ tells us that the empty function (from the empty subset of M to N) is a p.e.m. We will inductively extend the empty function to an elementary embedding of \mathcal{M} into \mathcal{N} , which as $|N| = |M|$ (why?) will have to be an isomorphism. Suppose we have built a p.e.m. $f : \{a_1, \dots, a_k\} \rightarrow N$, where $0 \leq k < n$. Let $X_{k+1,1}, \dots, X_{k+1,m_{k+1}}$ be all the $\{a_1, \dots, a_k\}$ -definable subsets of M containing a_{k+1} . (There are only finitely many as $\mathcal{P}(M)$ is finite.) Choose and fix formulas $\phi_{k+1,1}(a_1, \dots, a_k, x), \dots, \phi_{k+1,m_{k+1}}(a_1, \dots, a_k, x)$

defining $X_{k+1,1}, \dots, X_{k+1,m_{k+1}}$, respectively. Now $\mathcal{M} \models \bigwedge_{i=1}^{m_{k+1}} \phi_{k+1,i}(a_1, \dots, a_k, a_{k+1})$. Hence

$\mathcal{M} \models \exists x \bigwedge_{i=1}^{m_{k+1}} \phi_{k+1,i}(a_1, \dots, a_k, x)$, and so $\mathcal{N} \models \exists x \bigwedge_{i=1}^{m_{k+1}} \phi_{k+1,i}(f(a_1), \dots, f(a_k), x)$ by the in-

ductive hypothesis that f is a p.e.m. Let $b \in N$ realise this formula and define f' to be the extension of f that takes a_{k+1} to b . It is not hard to see that f' is also a p.e.m., as desired. \square

So if a complete theory has a finite model then it has only one model up to isomorphism. We are not interested in such theories.

CHAPTER 5

Compactness and Consequences

The compactness theorem for first-order logic is of fundamental importance and is the starting point for model theory and its applications.

THEOREM 5.1 (Compactness Theorem). *Suppose L is a language and T is an L -theory. T is consistent if and only if every finite subset of T is consistent.*

Here is an equivalent formulation.

COROLLARY 5.2. *Suppose T is an L -theory and σ is an L -structure. Then $T \models \sigma$ if and only if there exists a finite subset $\Sigma \subseteq T$ such that $\Sigma \models \sigma$.*

PROOF. Observe that $T \models \sigma$ if and only if $T \cup \{\neg\sigma\}$ is inconsistent. \square

In many texts the compactness theorem is seen as an immediate consequence of Gödel's completeness theorem, which says that T is consistent if and only if there is no "formal derivation" of a contradiction using the sentences in T as assumptions. The finite character of derivations then implies the Compactness Theorem. This approach however takes us into the realm of proof theory, which it is our intention to avoid. We will present an attractive "algebraic" proof of the compactness theorem using ultraproducts; a subject of interest in its own right.

5.1. A proof of compactness using ultraproducts

DEFINITION 5.3 (Filter). Suppose I is a nonempty set. A subset $\mathcal{F} \subseteq \mathcal{P}(I)$ is called a *filter on I* if the following conditions hold:

- (i) $I \in \mathcal{F}$ and $\emptyset \notin \mathcal{F}$.
- (ii) If $A, B \in \mathcal{F}$ then $A \cap B \in \mathcal{F}$.
- (iii) If $A \subseteq B$ and $A \in \mathcal{F}$ then $B \in \mathcal{F}$.

It is reasonable to think of a filter as giving a notion of "largeness" for subsets of I . Here are some examples:

- $\{\mathbb{R} \setminus X : X \text{ has Lebesgue measure } 0\}$ is a filter on \mathbb{R} .
- Suppose I is any infinite set and κ is an infinite cardinal such that $\kappa \leq |I|$. Then $\{I \setminus X : |X| < \kappa\}$ is a filter on I . In particular the set of all cofinite subsets of I forms a filter, called the *Frechet filter*.
- Suppose I is any nonempty set. A *principal filter* on I is a filter of the form $\{A \subseteq I : x \in A\}$ for some fixed $x \in I$.

An *ultrafilter* on I is a maximal filter; that is, a filter not properly contained in any other filter on I . It is not hard to see that every principal filter is an ultrafilter – any larger filter would contain a set that does not contain x and hence its intersection with $\{x\}$, which must

also be in the filter, would be \emptyset , contradicting the fact that \emptyset is not a member of any filter. While it is rather difficult to describe any particular nonprincipal ultrafilters, they do exist by Zorn's Lemma (and hence in ZFC). Indeed, start with any filter \mathcal{F} and consider all the filters on I extending \mathcal{F} . As the union of a chain of filters is a filter, there exists an ultrafilter \mathcal{U} containing \mathcal{F} . If \mathcal{F} was the Frechet filter on I (the set of all cofinite subsets of I), for example, then \mathcal{U} cannot be principal (as for any $x \in I$, $I \setminus \{x\} \in \mathcal{F} \subseteq \mathcal{U}$).

LEMMA 5.4. *A filter \mathcal{U} is an ultrafilter if and only if for every $A \subseteq I$, either $A \in \mathcal{U}$ or $I \setminus A \in \mathcal{U}$.*

PROOF. The right-to-left direction is clear: if $\mathcal{U} \subsetneq \mathcal{F}$ then there is $A \in \mathcal{F} \setminus \mathcal{U}$. Since \mathcal{F} is closed under intersections and does not contain \emptyset , $I \setminus A \notin \mathcal{F}$. Hence neither A nor $I \setminus A$ are in \mathcal{U} . For the left-to-right direction, suppose \mathcal{U} is an ultrafilter and $A \notin \mathcal{U}$. Consider $\mathcal{F} = \{X \subseteq I : Y \setminus A \subseteq X \text{ for some } Y \in \mathcal{U}\}$. Check that \mathcal{F} is a filter containing \mathcal{U} . Hence $\mathcal{U} = \mathcal{F}$. But $I \setminus A \in \mathcal{F}$. \square

DEFINITION 5.5 (Ultraproduct). Suppose L is a language, I is an infinite set, and \mathcal{M}_i is an L -structure for each $i \in I$. Let \mathcal{U} be an ultrafilter on I . The *ultraproduct of $(\mathcal{M}_i : i \in I)$ with respect to \mathcal{U}* will be the L -structure \mathcal{M} given as follows:

- The universe of \mathcal{M} is $M := \left(\prod_{i \in I} M_i \right) / E$ where E is the equivalence relation given by: $f E g$ if $\{i \in I : f(i) = g(i)\} \in \mathcal{U}$. That is, $f E g$ if the indices on which sequences f and g agree is a member of the ultrafilter \mathcal{U} .
- For every constant symbol $c \in L^{\text{con}}$, set $c^{\mathcal{M}}$ to be the E -class of the sequence given by $f(i) = c^{\mathcal{M}_i}$ for each $i \in I$.
- For every function symbol $f \in L^{\text{fun}}$, and all $g_1, \dots, g_{n_f} \in M$, define $f^{\mathcal{M}}(g_1/E, \dots, g_{n_f}/E)$ to be g/E where $g : I \rightarrow \bigcup_{i \in I} M_i$ is the sequence $g(i) = f^{\mathcal{M}_i}(g_1(i), \dots, g_{n_f}(i))$. Check that this definition does not depend on the choice of representative g_1, \dots, g_{n_f} .
- For every relation symbol $R \in L^{\text{rel}}$, define $R^{\mathcal{M}} \subseteq M^{k_R}$ as follows: for all $g_1, \dots, g_{k_R} \in M$, $(g_1/E, \dots, g_{k_R}/E) \in R^{\mathcal{M}}$ if and only if $\{i \in I : (g_1(i), \dots, g_{k_R}(i)) \in R^{\mathcal{M}_i}\} \in \mathcal{U}$. Check that this does not depend on the choice of representatives g_1, \dots, g_{k_R} .

We often denote the ultraproduct by $\prod_{\mathcal{U}} \mathcal{M}_i$.

The following lemma points out that if we start with a principal ultrafilter then this construction does not produce anything new.

LEMMA 5.6. *Suppose $\{\mathcal{M}_i : i \in I\}$ is a set of L -structures, $j \in I$, and \mathcal{U}_j is the principal ultrafilter $\{X \subseteq I : j \in X\}$. Then the j th projection map, $\pi : \times_{i \in I} \mathcal{M}_i \rightarrow \mathcal{M}_j$, induces an isomorphism from $\prod_{\mathcal{U}_j} \mathcal{M}_i$ to \mathcal{M}_j .*

PROOF. Note that a set of indices is in the ultrafilter if and only if it contains j . So if $g E g'$ then $g(j) = g'(j)$. Hence π does induce a map from the universe of $\prod_{\mathcal{U}_j} \mathcal{M}_i$ to \mathcal{M}_j . That it is surjective is clear, and that it is an L -embedding is not hard to check from the definitions. \square

PROPOSITION 5.7 (Łoś' Theorem). *Suppose $\mathcal{M} = \prod_{\mathcal{U}} \mathcal{M}_i$ where $\{\mathcal{M}_i : i \in I\}$ is a set of L -structures and \mathcal{U} is an ultrafilter on I . Suppose $\phi(x_1, \dots, x_n)$ is an L -formula, and*

$g_1, \dots, g_n \in \prod_{i \in I} M_i$. Then $\mathcal{M} \models \phi(g_1/E, \dots, g_n/E)$ if and only if

$$\{i \in I : \mathcal{M}_i \models \phi(g_1(i), \dots, g_n(i))\} \in \mathcal{U}.$$

In particular, for an L -sentence σ , $\mathcal{M} \models \sigma$ if and only if $\{i \in I : \mathcal{M}_i \models \sigma\} \in \mathcal{U}$.

PROOF. First of all, one observes that if $t(x_1, \dots, x_n)$ is an L -term and $g_1, \dots, g_n \in \times_{i \in I} M_i$, then

$$t^{\mathcal{M}}(g_1/E, \dots, g_n/E)(i) = t^{\mathcal{M}_i}(g_1(i), \dots, g_n(i))$$

for each $i \in I$. Indeed, this claim follows very easily by induction on the complexity of t using the fact that the corresponding statement is by definition true of function symbols.

The proof of Loś' Theorem is now a straightforward induction on the complexity of ϕ . The atomic case is an immediate consequence of the above claim on terms (and the definition of the ultraproduct). The inductive steps of \neg , \wedge , and \vee are easily verified using properties of ultrafilters; namely Lemma 5.4 (for \neg), that ultrafilters are closed under intersections (for \wedge), and that they are closed under unions (for \vee). As usual the case of the universal quantifier is reduced to that of the existential quantifier. So it remains to prove that if we know Loś theorem for $\psi(x_1, \dots, x_n, y)$ then it follows for $\exists y\psi$. By the inductive hypothesis

$$\mathcal{M} \models \exists y\psi(g_1/E, \dots, g_n/E, y)$$

if and only if

$$\text{there exists } f \in \times_{i \in I} M_i \text{ such that } X_f := \{i \in I : \mathcal{M}_i \models \psi(g_1(i), \dots, g_n(i), f(i))\} \in \mathcal{U}$$

Now X_f is contained in $Y := \{i \in I : \mathcal{M}_i \models \exists y\psi(g_1(i), \dots, g_n(i), y)\}$, and since ultrafilters are preserved under supersets, we see that if $\mathcal{M} \models \exists y\psi(g_1/E, \dots, g_n/E, y)$ then $Y \in \mathcal{U}$. Conversely, suppose $Y \in \mathcal{U}$. For each $i \in Y$ choose $a_i \in M_i$ such that $\mathcal{M}_i \models \psi(g_1(i), \dots, g_n(i), a_i)$. Define $f : I \rightarrow \bigcup_{i \in I} M_i$ by setting $f(i) = a_i$ for $i \in Y$, and anything otherwise. Then, X_f will contain Y and hence $X_f \in \mathcal{U}$. By the above equivalence, this means that $\mathcal{M} \models \exists y\psi(g_1/E, \dots, g_n/E, y)$. \square

As a corollary we can prove the compactness theorem. Let T be an L -theory each of whose finite subsets are consistent. Let I be the set of finite subsets of T . For each $i \in I$, let \mathcal{M}_i be a model of i . Also for each $i \in I$, let $X_i \subseteq I$ be the set of finite subsets of T extending i . Note that $\{X_i : i \in I\}$ does not include the empty set, does include I (indeed, $I = X_\emptyset$), and is closed under intersections (as $X_i \cap X_j = X_{i \cup j}$). Hence $\mathcal{F} := \{X \subseteq I : X_i \subseteq X \text{ for some } i \in I\}$ is a filter on I . Extend \mathcal{F} to an ultrafilter \mathcal{U} on I . Let $\mathcal{M} := \prod_{\mathcal{U}} \mathcal{M}_i$. I claim that $\mathcal{M} \models T$. Indeed, let $\sigma \in T$. If $j \in X_{\{\sigma\}}$ – that is, if j is a finite subset of T containing σ – then $\mathcal{M}_j \models \sigma$ since $\mathcal{M} \models j$. Hence, $\{i \in I : \mathcal{M}_i \models \sigma\}$ contains $X_{\{\sigma\}}$, and the latter is a member of \mathcal{F} (and hence \mathcal{U}) by construction. So $\{i \in I : \mathcal{M}_i \models \sigma\} \in \mathcal{U}$. By Loś' theorem, $\mathcal{M} \models \sigma$. So $\mathcal{M} \models T$. This completes the proof of the compactness theorem. \square

In fact, Loś' Theorem tells us more:

COROLLARY 5.8. *Suppose \mathcal{M} is an L -structure, I is a nonempty set, and \mathcal{U} is an ultrafilter on I . Then the diagonal induces an elementary embedding $\mathcal{M} \rightarrow \prod_{\mathcal{U}} \mathcal{M}$.*

PROOF. For each $a \in M$ let $f_a \in \times_{i \in I} M$ be given by $f_a(i) = a$ for all $i \in I$. Then the map induced by the diagonal is the one that takes a to f_a/E . Suppose $\phi(x_1, \dots, x_n)$ is an L -formula and $a_1, \dots, a_n \in M$. Then

$$\begin{aligned} \mathcal{M} \models \phi(a_1, \dots, a_n) &\implies \mathcal{M} \models \phi(f_{a_1}(i), \dots, f_{a_n}(i)) \text{ for each } i \\ &\implies \{i \in I : \mathcal{M} \models \phi(f_{a_1}(i), \dots, f_{a_n}(i))\} = I \in \mathcal{U} \\ &\implies \prod_{\mathcal{U}} \mathcal{M} \models \phi(f_{a_1}/E, \dots, f_{a_n}/E) \text{ by Łoś' Theorem} \end{aligned}$$

Conversely if $\{i \in I : \mathcal{M} \models \phi(f_{a_1}(i), \dots, f_{a_n}(i))\} \in \mathcal{U}$, then this set is nonempty and hence $\mathcal{M} \models \phi(f_{a_1}(i), \dots, f_{a_n}(i))$ for some $i \in I$. As each $f_{a_j}(i) = a_j$, $\mathcal{M} \models \phi(a_1, \dots, a_n)$, as desired. \square

Such an ultraproduct, $\times_{i \in I} \mathcal{M}/\mathcal{U}$ is called an *ultrapower* and often denoted by $\mathcal{M}^I/\mathcal{U}$. So we have that every structure elementarily embeds in every ultrapower of itself. Identifying \mathcal{M} with its image under the diagonal, we see that this gives us a very useful technique for producing elementary extensions.

For example, suppose $\mathcal{R} = (\mathbb{R}, 0, 1, +, -, \times, <)$ and let \mathcal{U} be an ultrafilter on ω that extends the Frechet filter. Let $\epsilon := (\frac{1}{n+1} : n \in \omega)/E$, which is an element of the universe of the elementary extension $\mathcal{R}^* := \mathcal{R}^\omega/\mathcal{U}$. Łoś' theorem implies that $\mathcal{R}^* \models (0 < \epsilon)$ but $\mathcal{R}^* \models (\epsilon < r)$ for every positive real number r . Indeed, $\{n : \frac{1}{n+1} < r\}$ is cofinite, hence in the Frechet filter and hence in \mathcal{U} . Such an ϵ is called an *infinitesimal*. So we have shown that there exist elementary extensions of the reals with infinitesimals. This is the beginning of *nonstandard analysis*. (Note that we are viewing $\mathcal{R} \preceq \mathcal{R}^*$ by identifying \mathcal{R} with its image under the diagonal.)

5.2. Applications of compactness, Löwenheim-Skolem, Vaught

We begin with some typical application of compactness to particular classes of structures.

PROPOSITION 5.9. *Let $L = \{0, +, -\}$ be the language of (additively written) groups. The class of torsion groups is not axiomatisable.*

PROOF. Suppose T was an axiomatisation of the torsion groups. Consider the language $L' := L \cup \{c\}$ where we have added a new constant symbol. Let T' be the L' -theory $T \cup \{nc \neq 0 : n > 0\}$. We claim that T' is consistent. Indeed, if Σ is a finite subset of T' then, for some $\ell > 0$, Σ is contained in $T \cup \{nc \neq 0 : n = 1, \dots, \ell\}$. Now the L -structure $(\mathbb{Z}/(\ell+1)\mathbb{Z}, 0, +, -)$ is a torsion group, and if we make it into an L' -structure by interpreting c as 1 then we get a model of $T \cup \{nc \neq 0 : n = 1, \dots, \ell\}$. Hence Σ is consistent. So by compactness, T' is consistent. But this is absurd since the reduct to L of any model must be a torsion group, while the interpretation of c will be torsion-free. \square

The compactness theorem can also be used to show that certain elementary classes are not *finitely* axiomatisable.

PROPOSITION 5.10. *Let $L = \{0, +, -\}$ be the language of (additively written) groups. The class of torsion-free groups is not finitely axiomatisable.*

PROOF. Suppose it were and seek a contradiction. Let σ be the conjunction of the (finitely many) sentences in this finite axiomatisation of the torsion-free groups. Let T be the natural axiomatisation of the class of torsion-free groups. That is, T is made up of the axioms for groups, G , together with sentences τ_n , saying $\forall x(x \neq 0 \rightarrow nx \neq 0)$, for each $n > 0$ (see Example 4.43(a)). So $T \models \sigma$. Hence, by compactness, there exists an $\ell > 0$ such that $G \cup \{\tau_1, \dots, \tau_\ell\} \models \sigma$. But then $\mathbb{Z}/p\mathbb{Z} \models \sigma$, if p is chosen to be a prime greater than ℓ . But this contradicts the assumption that σ axiomatises torsion-free groups. \square

Here is another typical use of the compactness theorem, similar to that of Proposition 5.9 above, though the same result could be obtained more explicitly using ultrapowers (how?).

PROPOSITION 5.11. *There exists an elementary extension of $\mathcal{Q} = (\mathbb{Q}^{\text{alg}}, 0, 1, +, -, \times)$ with a transcendental element.*

PROOF. Let L be the language of rings and let $L' = L_{\mathbb{Q}^{\text{alg}}} \cup \{c\}$ where c is a new constant symbol. Consider the L' -theory

$$T = \text{Th}(\mathcal{Q}_{\mathbb{Q}^{\text{alg}}}) \cup \{p(c) \neq 0 : p \in \mathbb{Q}[X], p \neq 0\}$$

By compactness, T is consistent: \mathcal{Q} itself can be made into a model of any finite subset of T by interpreting c to be an algebraic number whose minimal polynomial over \mathbb{Q} is of sufficiently large degree. Suppose $\mathcal{M}' \models T$ and let \mathcal{M} be the reduct of \mathcal{M}' to L . By Proposition 4.45, there is an elementary embedding of \mathcal{Q} in \mathcal{M} ; this is just the map that takes $q \in \mathbb{Q}^{\text{alg}}$ to the interpretation of \underline{q} in \mathcal{M}' . Finally, the interpretation of c in \mathcal{M}' gives us a transcendental element in \mathcal{M} . \square

What Proposition 5.11 tells us is that the algebraicity of \mathbb{Q}^{alg} is not part of the first-order theory of the structure, even after naming all parameters.

We now turn our attention to more general and theoretical consequences of compactness.

THEOREM 5.12 (Upward Löwenheim-Skolem). *Suppose \mathcal{M} is an infinite L -structure and $\kappa \geq \max\{|\mathcal{M}|, |L|\}$. Then there exists an elementary extension of \mathcal{M} of cardinality κ .*

In particular, if L is countable and T is an L -theory with an infinite model, then T has models of any infinite cardinality.

PROOF. We first find an elementary extension of \mathcal{M} of size at least κ . We use compactness: Let $L' = L_{\mathcal{M}} \cup \{c_\lambda : \lambda < \kappa\}$ where the c_λ s are new constant symbols. Let

$$T = \text{Th}(\mathcal{M}_{\mathcal{M}}) \cup \{c_\lambda \neq c_\gamma : \lambda < \gamma < \kappa\}$$

We claim that T is consistent. Indeed, \mathcal{M} itself can be made into a model of any finite subset of T by interpreting the c_λ s which appear as distinct elements (possible as \mathcal{M} is infinite). Hence by the compactness theorem, T is consistent. Let $\mathcal{N}' \models T$ and let \mathcal{N} be the reduct of \mathcal{N}' to L . Then \mathcal{N} is of cardinality at least κ (witnessed by the interpretation of the c_λ s in \mathcal{N}') and by Proposition 4.45 there is an elementary embedding of \mathcal{M} in \mathcal{N} . Identifying \mathcal{M} with its image we may assume that $\mathcal{M} \preceq \mathcal{N}$.

Let $A \subseteq N$ contain M and be of cardinality κ . Applying the downward Löwenheim-Skolem theorem, namely Theorem 4.27, we get $\mathcal{N}_1 \preceq \mathcal{N}$ with $A \subseteq N_1$ and $|N_1| = \kappa$. Moreover, we have $\mathcal{M} \subseteq \mathcal{N}_1 \preceq \mathcal{N}$ and $\mathcal{M} \preceq \mathcal{N}$. It follows that $\mathcal{M} \preceq \mathcal{N}_1$. Indeed, for any L -formula $\phi(x_1, \dots, x_n)$ and any $a \in M^n$, $\mathcal{M} \models \phi(a) \iff \mathcal{N} \models \phi(a) \iff \mathcal{N}_1 \models \phi(a)$. So \mathcal{N}_1 is an elementary extension of \mathcal{M} of size κ . \square

COROLLARY 5.13 (Vaught's Test). *Suppose T is an L -theory with only infinite models. If there exists an infinite cardinal $\kappa \geq |L|$ such that all models of T of size κ are isomorphic (we say that T is κ -categorical) then T is complete.*

PROOF. Suppose $\mathcal{M}_1, \mathcal{M}_2$ are models of T . If $|M_i| \leq \kappa$ let \mathcal{M}'_i be an elementary extension of size κ , which exists by the upward Löwenheim-Skolem theorem. If $|M_i| > \kappa$ then use the downward Löwenheim-Skolem theorem (with some $A \subseteq M_i$ of cardinality κ) to obtain an elementary substructure \mathcal{M}'_i of size κ . In either case, $\mathcal{M}'_i \equiv \mathcal{M}_i$ and so $\mathcal{M}'_i \models T$. By κ -categoricity, $\mathcal{M}'_1 \approx \mathcal{M}'_2$ and hence $\mathcal{M}'_1 \equiv \mathcal{M}'_2$. So $\mathcal{M}_1 \equiv \mathcal{M}'_1 \equiv \mathcal{M}'_2 \equiv \mathcal{M}_2$. Hence T is complete by Lemma 4.44(b). \square

Let us apply this to a few examples.

EXAMPLE 5.14 (DLO is complete). Let $L = \{<\}$ be the language of orderings and let DLO be the theory of dense linear orderings without endpoints. We show that DLO is \aleph_0 -categorical. The method we apply is ubiquitous in model theory and is called *back-and-forth*. Suppose (E_1, \leq_1) and (E_2, \leq_2) are (infinite) countable models of DLO. We enumerate them as $E_1 = \{a_i : i < \omega\}$ and $E_2 = \{b_i : i < \omega\}$. Note that these enumerations are not in any way related to the given orderings. We build a chain of partial order-preserving bijections $f_i : A_i \rightarrow B_i$ where $A_0 \subseteq A_1 \subseteq \dots$ and $B_0 \subseteq B_1 \subseteq \dots$ are finite subsets, such that $E_1 = \bigcup_i A_i$ and $E_2 = \bigcup_i B_i$. The union $f = \bigcup_i f_i$ will be the desired isomorphism.

Stage 0. Let $A_0 = B_0 = f_0 = \emptyset$.

Stage $n + 1 = 2m + 1$. At this stage we ensure that $a_m \in A_{n+1}$. If $a_m \in A_n$ then set $A_{n+1} = A_n$, $B_{n+1} = B_n$, and $f_{n+1} = f_n$. Suppose $a_m \notin A_n$. Then set $A_{n+1} = A_n \cup \{a_m\}$. Now exactly one of the following three cases is possible:

- (i) a_m is greater than every element of A_n , or
- (ii) a_m is less than every element of A_n , or
- (iii) there exists $\alpha <_1 \beta$ in A_n without any elements of A_n between them and $\alpha <_1 a_m <_1 \beta$.

In case (i) let $b \in E_2$ be greater than every element of B_n (which is possible as B_n is finite and E_2 has no endpoints). In case (ii) let $b \in E_2$ be less than every element of B_n . In case (iii) let $b \in E_2$ be such that $f_n(\alpha) <_2 b <_2 f_n(\beta)$, which is possible since $\alpha <_1 \beta$ and so $f_n(\alpha) <_2 f_n(\beta)$ and E_2 is dense. In any case, set $B_{n+1} = B_n \cup \{b\}$ and $f_{n+1} = f_n \cup \{(a_m, b)\}$. Then we have that $f_{n+1} : A_{n+1} \rightarrow B_{n+1}$ is an order preserving bijection.

Stage $n + 1 = 2m + 2$. At this stage we ensure that $b_m \in B_{n+1}$. If $b_m \in B_n$ then we do nothing and set $B_{n+1} = B_n$, $A_{n+1} = A_n$, and $f_{n+1} = f_n$. Otherwise let $B_{n+1} = B_n \cup \{b_m\}$. Now b_m sits with respect to the elements of B_n in three possible ways analogously to (i), (ii), and (iii) above. Then just as in the odd stage we choose $a \in E_1$ according to which case b_m satisfies. Finally we let $A_{n+1} = A_n \cup \{a\}$ and $f_{n+1} = f_n \cup \{(a, b_m)\}$.

Then $f = \bigcup_i f_i : E_1 = \bigcup_i A_i \rightarrow \bigcup_i B_i = E_2$ is an isomorphism. We have shown that DLO is \aleph_0 -categorical, and hence, as dense linear orderings without endpoints are all infinite, Vaught's test implies that DLO is complete. What this says is that no dense linear ordering without endpoints can be distinguished from $(\mathbb{Q}, <)$ by a sentence in the language of orderings.

EXAMPLE 5.15 (ACF_p is complete). We denote by ACF_p the set of sentences in the language of rings that make up the usual axioms for algebraically closed fields of characteristic

p , where p is a prime or zero (see Example 4.43(b)). We can use the Vaught test to see that ACF_p is complete. Indeed, let $K \models \text{ACF}_p$ and let $\mathbb{F} \subseteq K$ be the prime field. (So $\mathbb{F} = \mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$ if p is prime and $\mathbb{F} = \mathbb{Q}$ if $p = 0$.) Let B be a transcendence basis for K over \mathbb{F} . So $K = \mathbb{F}(B)^{\text{alg}}$. I leave it as an exercise for you to check that either B is finite and $|\mathbb{F}(B)^{\text{alg}}| = \aleph_0$, or $|\mathbb{F}(B)^{\text{alg}}| = |B|$. Hence, if K is uncountable then $|B| = |K|$. Assume this is the case. Now if $L \models \text{ACF}_p$ and $|L| = |K|$ then L has a transcendence basis C over \mathbb{F} with $|C| = |B|$. We obtain therefore that L and K are isomorphic fields (any bijection between C and B will extend to an isomorphism). We have shown that ACF_p is κ -categorical for any uncountable cardinal κ . Also, every algebraically closed field is infinite. Hence, by Vaught's test (which actually only requires κ -categoricity for *some* infinite κ), ACF_p is complete. So for example, no sentence in the language of rings can distinguish between \mathbb{Q}^{alg} and \mathbb{C} .

EXERCISE 5.16. Let F be any fixed field and consider T the theory of infinite F -vector spaces. Show that T is complete.

5.3. An application to geometry: the Lefschetz Principle

The completeness of ACF_0 , together with the compactness theorem, can be used to make precise the following “metatheorem” in algebraic geometry: statements true in almost all positive characteristics are true in characteristic zero. In this section we establish this principle and then apply it to prove something about regular endomorphisms of algebraic varieties.

For any prime p , let \mathbb{F}_p denote the prime field of p elements, $\mathbb{Z}/p\mathbb{Z}$.

THEOREM 5.17 (Lefschetz Principle). *Let σ be a sentence in the language of rings. The following are equivalent:*

- (i) $(\mathbb{C}, 0, 1, +, -, \times) \models \sigma$
- (ii) For all but finitely many primes p , $(\mathbb{F}_p^{\text{alg}}, 0, 1, +, -, \times) \models \sigma$.
- (iii) For infinitely many primes p , $(\mathbb{F}_p^{\text{alg}}, 0, 1, +, -, \times) \models \sigma$.

In fact, in all of the above one can replace \mathbb{C} by any algebraically closed field of characteristic zero, and one can replace $\mathbb{F}_p^{\text{alg}}$ by any algebraically closed field of characteristic p .

PROOF. (i) implies (ii). Since ACF_0 is complete, $\text{Th}(\mathbb{C}, 0, 1, +, -, \times)$ is equal to the set of consequences of ACF_0 , and so $\text{ACF}_0 \models \sigma$. By the compactness theorem, there is a finite set of sentences $\Sigma \subseteq \text{ACF}_0$ which implies σ . Now ACF_0 is made up of the axioms for algebraically closed fields together with axioms of the form $\exists x(px \neq 0)$ for all primes p . So for some $N > 0$, $\Sigma \subseteq \text{ACF} \cup \{\exists x(px \neq 0) : p \text{ prime} \leq N\}$. Hence, $\text{ACF} \cup \{\exists x(px \neq 0) : p \text{ prime} \leq N\} \models \sigma$. But then for all primes p bigger than N , $(\mathbb{F}_p^{\text{alg}}, 0, 1, +, -, \times) \models \sigma$.

(ii) implies (iii). Clear.

(iii) implies (i). If $(\mathbb{C}, 0, 1, +, -, \times) \not\models \sigma$ then $(\mathbb{C}, 0, 1, +, -, \times) \models \neg\sigma$ and so by the (i) implies (ii) direction we get that for all but finitely many primes p , $(\mathbb{F}_p^{\text{alg}}, 0, 1, +, -, \times) \models \neg\sigma$. Hence, for only finitely many primes p does $(\mathbb{F}_p^{\text{alg}}, 0, 1, +, -, \times) \models \sigma$. \square

Let us give an example of how the Lefschetz principle can be used in algebraic geometry. We will work with the field of complex numbers, but we could equally work with any algebraically closed field of characteristic zero. Recall the *Zariski topology* on \mathbb{C}^n from Example 4.31; the closed sets are the zero sets of finite collections of polynomials

with complex coefficients in n variables. If $V, W \subseteq \mathbb{C}^n$ are Zariski closed sets, a function $f : V \rightarrow W$ is called a *polynomial map* if there exist $Q_1, \dots, Q_m \in \mathbb{C}[X_1, \dots, X_n]$ such that $f(a) = (Q_1(a), \dots, Q_m(a))$ for all $a = (a_1, \dots, a_n) \in V$.

PROPOSITION 5.18. *Suppose V is a Zariski closed subset of \mathbb{C}^n . Every injective polynomial map $f : V \rightarrow V$ is surjective.*

PROOF. Here is the strategy: convert what we are trying to prove into a sentence in the language of rings, apply the Lefschetz Principle to reduce to positive characteristic, use the pro-finiteness of $\mathbb{F}_p^{\text{alg}}$ to further reduce the problem to Zariski closed subsets of k^n where k is a *finite* field, and the result is now just the pigeonhole principle.

Let V be the zero set of $P_1, \dots, P_\ell \in \mathbb{C}[X_1, \dots, X_n]$ and $f = (Q_1, \dots, Q_n) : V \rightarrow V$ an injective but *not* surjective polynomial map, where $Q_1, \dots, Q_n \in \mathbb{C}[X_1, \dots, X_n]$. Replacing the coefficients of P_i by variables we can write each $P_i = P'_i(X_1, \dots, X_n, b_i)$, where $P'_i \in \mathbb{Z}[X_1, \dots, X_n, Y_1, \dots, Y_{r_i}]$ and $b_i \in \mathbb{C}^{r_i}$, for $i = 1, \dots, \ell$. Similarly, for each $j = 1, \dots, n$, we write $Q_j = Q'_j(X_1, \dots, X_n, c_j)$, where $Q'_j \in \mathbb{Z}[X_1, \dots, X_n, Z_1, \dots, Z_{s_j}]$ and $c_j \in \mathbb{C}^{s_j}$.

Let $L := \{0, 1, +, -, \times\}$ and σ the L -sentence which expresses that “there exist tuples $y_1, \dots, y_\ell, z_1, \dots, z_n$ such that $(Q'_1(X_1, \dots, X_n, z_1), \dots, Q'_n(X_1, \dots, X_n, z_n))$ is an injective but not surjective map from the zero set of $\{P'_1(X_1, \dots, X_n, y_1), \dots, P'_\ell(X_1, \dots, X_n, y_\ell)\}$ to itself”. Note that this an L -sentence as P'_i and Q'_i are L -terms and injectivity and surjectivity are clearly definable properties. By assumption $(\mathbb{C}, 0, 1, +, -, \times) \models \sigma$, witnessed by $b_1, \dots, b_\ell, c_1, \dots, c_n$. By the Lefschetz principle, for some (indeed infinitely many) prime $p > 0$, $(\mathbb{F}_p^{\text{alg}}, 0, 1, +, -, \times) \models \sigma$.

Let $d_1, \dots, d_\ell, e_1, \dots, e_n$ be tuples from $\mathbb{F}_p^{\text{alg}}$ that witness the truth of σ . That is, if $W \subseteq (\mathbb{F}_p^{\text{alg}})^n$ is the zero set of $\{P'_i(X_1, \dots, X_n, d_i) : i = 1, \dots, \ell\}$, and $g : W \rightarrow W$ is given by $(Q'_1(X_1, \dots, X_n, e_1), \dots, Q'_n(X_1, \dots, X_n, e_n))$, then g is injective but not surjective. Let $w \in W \setminus g(W)$, and let $k \leq \mathbb{F}_p^{\text{alg}}$ be the subfield generated by e_1, \dots, e_n, w . As $\mathbb{F}_p^{\text{alg}}$ is the countable increasing union of finite fields, k is a *finite* field. Note that $g(W \cap k^n) \subseteq W \cap k^n$ as the coefficients of the polynomials in g (namely e_1, \dots, e_n) are all in k . Also, $W \cap k^n \setminus g(W \cap k^n) \neq \emptyset$, as it contains w . Hence g restricts to an injective but not surjective map from $W \cap k^n$ to itself. But $W \cap k^n$ is finite! \square

CHAPTER 6

Quantifier Elimination

In the last chapter we saw that ACF_p , where p is a prime or zero, is a complete theory. In particular, $(\mathbb{Q}^{\text{alg}}, 0, 1, +, -, \times) \equiv (\mathbb{C}, 0, 1, +, -, \times)$. But is \mathbb{Q}^{alg} an elementary substructure of \mathbb{C} ? Equivalently, do they satisfy the same $L_{\mathbb{Q}^{\text{alg}}}$ -sentences? Now an $L_{\mathbb{Q}^{\text{alg}}}$ -sentence is a sentence of the form $\phi(a_1, \dots, a_n)$ where $\phi(x_1, \dots, x_n)$ is an L -formula and $a_1, \dots, a_n \in \mathbb{Q}^{\text{alg}}$. As \mathbb{Q}^{alg} is a substructure of \mathbb{C} , if ϕ happens to be quantifier-free then the answer is “yes” by Proposition 4.22(a). To see that the answer is unconditionally “yes” we will show that every formula is equivalent to a quantifier-free formula.

DEFINITION 6.1 (Quantifier Elimination). An L -theory T admits *quantifier elimination* (QE) if for every L -formula $\phi(x_1, \dots, x_n)$, where $n > 0$, there exists a quantifier-free L -formula $\psi(x_1, \dots, x_n)$ such that $T \models \forall x_1 \cdots x_n (\phi(x_1, \dots, x_n) \leftrightarrow \psi(x_1, \dots, x_n))$.

We say that $\phi(x_1, \dots, x_n)$ is *equivalent to $\psi(x_1, \dots, x_n)$ modulo T* to mean that $T \models \forall x_1 \cdots x_n (\phi(x_1, \dots, x_n) \leftrightarrow \psi(x_1, \dots, x_n))$. Equivalently, for all $\mathcal{M} \models T$, $\phi^{\mathcal{M}} = \psi^{\mathcal{M}}$. So QE says that every L -formula in at least one free variable is equivalent modulo T to a quantifier-free formula in the same free variables. Note that we ask for $n > 0$, so the definition does not seem to apply to L -sentences. Of course if σ is an L -sentence then we can write $\phi(x) := \sigma$, and quantifier elimination would produce a quantifier-free formula $\psi(x)$ that is equivalent to σ modulo T . But in $\psi(x)$ the variable x may in fact appear free, even though it did not in $\phi(x)$. So, assuming QE, we can eliminate the quantifiers from σ but at the cost of (possibly) introducing a free variable.

EXERCISE 6.2. Show that if L contains a constant symbol and T admits QE then every L -sentence is equivalent modulo T to a quantifier-free L -sentence.

In this chapter we will develop a criterion for quantifier-elimination and then use it to prove that various theories (including ACF) admit quantifier elimination. The model-theoretic study of almost any structure begins with proving quantifier elimination in some appropriate language. Without knowing QE it is very difficult to get a grasp on what the definable sets look like. However, it is important to realise that quantifier elimination in *some* language can always be forced. For example, suppose \mathcal{M} is any L -structure, and consider the *Skolemisation* of \mathcal{M} : Let L' be the expanded language where there is a new n -ary relation symbol P_A for every 0-definable subset $A \subseteq M^n$, for all $n < \omega$. We can make \mathcal{M} into an L' -structure canonically by interpreting P_A as A . Note that the definable subsets of \mathcal{M} viewed as an L -structure and as an L' -structure are the same – because the expansion we are considering does not introduce any new definable sets – it just makes all the old definable sets atomic. But, as an L' -theory, $\text{Th}(\mathcal{M})$ has quantifier elimination (exercise). While this can always be done we gain nothing by it; the language L' has become very complicated. The point is that quantifier elimination is only useful if we have some control over the language. And when we do, it can be very useful.

6.1. Preliminaries on substructures

We introduce a couple more basic notions that we will need to prove our desired criterion for quantifier elimination, and that are of independent interest.

DEFINITION 6.3. Suppose \mathcal{M} is an L -structure and $A \subseteq M$. Then the *substructure generated by A* is the smallest substructure of \mathcal{M} whose universe contains A . If this happens to be \mathcal{M} itself, then we say that A *generates* \mathcal{M} .

LEMMA 6.4. *Suppose \mathcal{M} is an L -structure and $A \subseteq M$. Assume either that A is nonempty or that L has a constant symbol. The universe of the substructure generated by A is $\{t^{\mathcal{M}}(a_1, \dots, a_n) : n \in \omega, a_1, \dots, a_n \in A, t(x_1, \dots, x_n) \text{ an } L\text{-term}\}$.*

PROOF. By Exercise 4.10 we know that a nonempty subset is the universe of a substructure if and only if it contains all the constants and is preserved by all the basic functions. We first show this is the case for

$$N = \{t^{\mathcal{M}}(a_1, \dots, a_n) : n \in \omega, a_1, \dots, a_n \in A, t(x_1, \dots, x_n) \text{ an } L\text{-term}\}.$$

Considering the term x we see that $A \subseteq N$. If $c \in L^{\text{con}}$ then c is itself a term and we have that $c^{\mathcal{M}} \in N$. That is, N contains A and all the constants of \mathcal{M} . In particular, by assumption $N \neq \emptyset$. If $F \in L^{\text{fun}}$ is ℓ -ary, and $a_i = t_i^{\mathcal{M}}(a_{i,1}, \dots, a_{i,n_\ell})$ are elements of N for $i = 1, \dots, \ell$, then

$$F^{\mathcal{M}}(a_1, \dots, a_\ell) = t^{\mathcal{M}}(a_{1,1}, \dots, a_{1,n_1}, \dots, a_{\ell,1}, \dots, a_{\ell,n_\ell})$$

where $t = F(t_1, \dots, t_\ell)$. Hence $F^{\mathcal{M}}(a_1, \dots, a_\ell) \in N$. So N is the universe of a substructure of \mathcal{M} , say \mathcal{N} .

To see that \mathcal{N} is the substructure generated by A , we need to show that if $\mathcal{N}' \subseteq \mathcal{M}$ and $A \subseteq \mathcal{N}'$, then $\mathcal{N} \subseteq \mathcal{N}'$. But this is the case since for every L -term t and $a_1, \dots, a_n \in A \subseteq \mathcal{N}'$, $t^{\mathcal{M}}(a_1, \dots, a_n) = t^{\mathcal{N}'}(a_1, \dots, a_n)$ since $\mathcal{N}' \subseteq \mathcal{M}$. Hence $t^{\mathcal{M}}(a_1, \dots, a_n) \in \mathcal{N}'$. \square

In Proposition 4.45 we saw that there is an elementary embedding from \mathcal{M} to \mathcal{N} if and only if \mathcal{N} can be expanded to a model of $\text{Th}(\mathcal{M}_M)$. We now give a similar criterion for the existence of an embedding, but we refine it a little to take into account a generating set.

LEMMA 6.5. *Suppose \mathcal{M} is an L -structure generated by $A \subseteq M$. Assume either that A is nonempty or that L has a constant symbol. Consider the L_A -theory*

$$D := \{\phi(\underline{a}) : n \in \omega, \underline{a} = (a_1, \dots, a_n) \in A^n, \phi \text{ a quantifier-free } L\text{-formula, and } \mathcal{M} \models \phi(\underline{a})\}.$$

Suppose \mathcal{N} is an L -structure. Then there exists an L -embedding $j : \mathcal{M} \rightarrow \mathcal{N}$ if and only if \mathcal{N} can be expanded to an L_A -structure \mathcal{N}' such that $\mathcal{N}' \models D$.

PROOF. If such an embedding j exists then expand \mathcal{N} to an L_A -structure \mathcal{N}' by $\underline{a}^{\mathcal{N}'} = j(\underline{a})$ for each $\underline{a} \in A$. Then $\mathcal{N}' \models D$ by Proposition 4.22(a).

For the converse, let $\mathcal{N}' \models D$ be an expansion of \mathcal{N} . We define $j : \mathcal{M} \rightarrow \mathcal{N}$ as follows. Suppose $b \in M$. Then by Lemma 6.4 there is an L -term $t(x_1, \dots, x_n)$, and $\underline{a} \in A^n$, such that $b = t^{\mathcal{M}}(\underline{a})$. Now $t(\underline{a})$ is an L_A -term. Set $j(b) := t(\underline{a})^{\mathcal{N}'}$. This map is injective: if $b \neq b'$ are elements of M with $b = t^{\mathcal{M}}(\underline{a})$ and $b' = s^{\mathcal{M}}(\underline{a}')$, where t and s are L -terms and $\underline{a} \in A^n$ and $\underline{a}' \in A^m$, then $(t(\underline{a}) \neq s(\underline{a}')) \in D$, and so $j(b) \neq j(b')$. For $c \in L^{\text{con}}$, $j(c^{\mathcal{M}}) = c^{\mathcal{N}'}$ by definition and $c^{\mathcal{N}'} = c^{\mathcal{N}}$ as \mathcal{N}' expands \mathcal{N} . If $F \in L^{\text{fun}}$ is n -ary and $b_1, \dots, b_n \in M$, then

writing each $b_i = t_i^{\mathcal{M}}(a_i)$ where t_i is an L -term and $a_i \in A^{n_i}$, we have that $j(b_i) = t_i(\underline{a_i})^{\mathcal{N}'}$. Moreover, $F^{\mathcal{M}}(b_1, \dots, b_n) = F^{\mathcal{M}}(t_1^{\mathcal{M}}(a_1), \dots, t_n^{\mathcal{M}}(a_n))$. Hence,

$$\begin{aligned} j(F(b_1, \dots, b_n)) &= F(t_1(\underline{a_1}), \dots, t_n(\underline{a_n}))^{\mathcal{N}'} \\ &= F^{\mathcal{N}'}(t_1(\underline{a_1})^{\mathcal{N}'}, \dots, t_n(\underline{a_n})^{\mathcal{N}'}) \\ &= F^{\mathcal{N}}(j(b_1), \dots, j(b_n)) \end{aligned}$$

where the last equality again uses that \mathcal{N}' is an expansion of \mathcal{N} . Finally, suppose $R \in L^{\text{rel}}$ is n -ary and $b_1, \dots, b_n \in M$. Again writing each $b_i = t_i^{\mathcal{M}}(a_i)$ where t_i is an L -term and $a_i \in A^{n_i}$ we have

$$\begin{aligned} (b_1, \dots, b_n) \in R^{\mathcal{M}} &\iff (t_1^{\mathcal{M}}(a_1), \dots, t_n^{\mathcal{M}}(a_n)) \in R^{\mathcal{M}} \\ &\iff R(t_1(\underline{a_1}), \dots, t_n(\underline{a_n})) \in D \\ &\iff (t_1(\underline{a_1})^{\mathcal{N}'}, \dots, t_n(\underline{a_n})^{\mathcal{N}'}) \in R^{\mathcal{N}'} \\ &\iff (j(b_1), \dots, j(b_n)) \in R^{\mathcal{N}}. \end{aligned}$$

Hence j is an L -embedding. □

6.2. A criterion for quantifier elimination

We begin with a criterion for eliminating quantifiers from a given formula.

THEOREM 6.6. *Suppose T is an L -theory and $\phi(x)$ is an L -formula where $x = (x_1, \dots, x_n)$. Assume that either L has a constant symbol or $n > 0$. The following are equivalent.*

- (i) $\phi(x)$ is equivalent to a quantifier-free formula $\psi(x)$ modulo T .
- (ii) Suppose $\mathcal{M} \models T$, $\mathcal{N} \models T$, and \mathcal{A} is an L -substructure of both \mathcal{M} and \mathcal{N} . Then for all $a \in A^n$, $\mathcal{M} \models \phi(a)$ if and only if $\mathcal{N} \models \phi(a)$.

PROOF. Suppose $\phi(x)$ is equivalent to a quantifier-free L -formula $\psi(x)$ modulo T . Let $\mathcal{M}, \mathcal{N}, \mathcal{A}$ be as in (ii). Then for any $a \in A^n$,

$$\begin{aligned} \mathcal{M} \models \phi(a) &\iff \mathcal{M} \models \psi(a) \text{ as } \psi \text{ is equivalent to } \psi \text{ modulo } T \text{ and } \mathcal{M} \models T \\ &\iff \mathcal{A} \models \psi(a) \text{ as } \mathcal{A} \subseteq \mathcal{M}, \text{ and by Proposition 4.22(a)} \\ &\iff \mathcal{N} \models \psi(a) \text{ as } \mathcal{A} \subseteq \mathcal{N}, \text{ and by Proposition 4.22(a)} \\ &\iff \mathcal{N} \models \phi(a) \text{ as } \psi \text{ is equivalent to } \psi \text{ modulo } T \text{ and } \mathcal{N} \models T \end{aligned}$$

as desired.

For the converse, suppose (ii) holds. We look for a quantifier-free formula equivalent to $\phi(x)$ modulo T . First consider

$$\Psi(x) := \{\psi(x) : \psi \text{ is quantifier-free and } T \models \forall x(\phi(x) \rightarrow \psi(x))\}.$$

This is the set of all quantifier-free consequences of ϕ . Let c_1, \dots, c_n be new constant symbols and let $L' = L \cup \{c_1, \dots, c_n\}$. Let $c = (c_1, \dots, c_n)$ and denote by $\Psi(c)$ the set of L' -sentences $\{\psi(c) : \psi \in \Psi\}$.

CLAIM 6.7. $T \cup \Psi(c) \models \phi(c)$.

PROOF OF CLAIM 6.7. Suppose not. Then there is an L' -structure $\mathcal{M}' \models T \cup \Psi(c) \cup \{\neg\phi(c)\}$, with universe M . Let \mathcal{A}' be the L' -substructure of \mathcal{M}' generated by \emptyset . If \mathcal{M} and \mathcal{A} denote their reducts to L , then we have $\mathcal{M} \models \neg\phi(b)$ where $b = (b_1, \dots, b_n)$ and each $b_i := c_i^{\mathcal{M}'}$. Note that \mathcal{A} is the substructure of \mathcal{M} generated by $B := \{b_1, \dots, b_n\}$. Our goal is to construct a model of T which is an extension of \mathcal{A} and in which $\phi(b)$ is true. This will contradict (ii).

Consider the L_B -theory $D = \{\theta(\underline{b}) : \theta(x) \text{ is a quantifier-free } L\text{-formula and } \mathcal{A} \models \theta(\underline{b})\}$.

SUBCLAIM 6.8. $\Sigma := T \cup D \cup \phi(\underline{b})$ is consistent.

PROOF OF SUBCLAIM 6.8. By compactness we need only consider a finite subset $\Sigma_0 \subseteq \Sigma$. Only finitely many quantifier-free L_B -sentences from D appear in Σ_0 ; say $\theta_1(\underline{b}), \dots, \theta_\ell(\underline{b})$.

If Σ_0 is not consistent then working in L' we have $T \models \left(\bigwedge_{i=1}^{\ell} \theta(\underline{b}) \rightarrow \neg\phi(\underline{b})\right)$, and hence

$T \models \left(\phi(\underline{b}) \rightarrow \bigvee_{i=1}^{\ell} \neg\theta_i(\underline{b})\right)$. Now, as \underline{b} is a tuple of new constant symbols not in L , an L' -structure that is a model of T is just an L -structure that is a model of T with an arbitrary choice of n -tuple interpreting \underline{b} . Hence, working just in L , $T \models \forall x(\phi(x) \rightarrow \bigvee_{i=1}^{\ell} \neg\theta(x))$. By

definition, it follows that $\bigvee_{i=1}^{\ell} \neg\theta(x) \in \Psi(x)$ and hence $\mathcal{M} \models \bigvee_{i=1}^{\ell} \neg\theta(b)$. As this is a quantifier-free L' -sentence, we get $\mathcal{A} \models \bigvee_{i=1}^{\ell} \neg\theta(b)$. Which is absurd as $\theta_1(\underline{b}), \dots, \theta_\ell(\underline{b}) \in D$. Hence Σ_0 must be consistent. So Σ is consistent. This completes the proof Subclaim 6.8. \square

Let $\tilde{\mathcal{N}} \models \Sigma$ with universe N and let \mathcal{N} be the reduct of $\tilde{\mathcal{N}}$ to L . So $\mathcal{N} \models T$. By Lemma 6.5, since B generates \mathcal{A} and \mathcal{N} expands to the L_B -structure $\tilde{\mathcal{N}} \models D$, we get an L -embedding $j : \mathcal{A} \rightarrow \mathcal{N}$. (Note that either L has a constant symbol or $B \neq \emptyset$, so that Lemma 6.5 does indeed apply.) Under this mapping b maps to $\underline{b}^{\tilde{\mathcal{N}}}$. We can identify \mathcal{A} with its image and thereby view $\mathcal{A} \subseteq \mathcal{N}$ with $b = \underline{b}^{\tilde{\mathcal{N}}}$. So $\mathcal{N} \models \phi(b)$. Since $\mathcal{M} \models \neg\phi(b)$, we have contradicted (ii). This completes the proof of Claim 6.7. \square

We have shown that $T \cup \Psi(c) \models \phi(c)$. By compactness there exists $\psi_1, \dots, \psi_\ell \in \Psi$ such that $T \cup \{\psi_1(c), \dots, \psi_\ell(c)\} \models \phi(c)$. Let $\psi(x) := \bigwedge_{i=1}^{\ell} \psi_i(x)$. We have that $T \models \psi(c) \rightarrow \phi(c)$. Again, as c does not appear in T this means that $T \models \forall x(\psi(x) \rightarrow \phi(x))$. On the other hand, $T \models \forall x(\phi(x) \rightarrow \psi(x))$ since $T \models \forall x(\phi(x) \rightarrow \psi_i(x))$ for each $i = 1, \dots, \ell$. So $T \models \forall x(\phi(x) \leftrightarrow \psi(x))$, and we have found a quantifier-free formula that is equivalent to ϕ modulo T . This completes the proof of Theorem 6.6. \square

While the last theorem tells us how to test whether a given formula is equivalent to a quantifier-free formula, the following, rather easier proposition, tells us which formulas we need to consider in order to obtain quantifier elimination.

PROPOSITION 6.9. *Suppose T is an L -theory. The following are equivalent,*

- (i) *T admits quantifier-elimination.*
- (ii) *For all $n > 0$ and all quantifier-free formulas $\theta(x, y)$ where $x = (x_1, \dots, x_n)$, the formula $\exists y\theta(x, y)$ is equivalent modulo T to a quantifier-free formula $\psi(x)$.*

If L has a constant symbol then we can allow $n = 0$ in (ii).

PROOF. That (i) implies (ii) is immediate. For the converse, we assume (ii) and prove by induction on complexity that every formula $\phi(x)$ is equivalent modulo T to a quantifier-free formula. For ϕ atomic we can take ϕ itself. If ϕ is $\xi_1 \wedge \xi_2$ and ξ_i is equivalent to a quantifier-free ξ'_i modulo T , for $i = 1, 2$, then ϕ is equivalent modulo T to the quantifier-free $\xi'_1 \wedge \xi'_2$. We can deal similarly with the case of \vee and \neg . So assume that $\phi(x)$ is $\exists y\xi(x, y)$. By the induction hypothesis $T \models \forall xy(\xi(x, y) \leftrightarrow \theta(x, y))$, where $\theta(x, y)$ is quantifier-free. Hence $T \models \forall x(\phi(x) \leftrightarrow \exists y\theta(x, y))$. By (ii), $T \models \forall x(\exists y\theta(x, y) \leftrightarrow \psi(x))$ for some quantifier-free $\psi(x)$. But then we have $T \models \forall x(\phi(x) \leftrightarrow \psi(x))$, as desired. (As usual the case of \forall reduces to the cases of \neg and \exists .) \square

Putting Theorem 6.6 and Proposition 6.9 together we obtain:

COROLLARY 6.10. *Suppose T is an L -theory. The following are equivalent:*

- (i) *T admits quantifier elimination.*
- (ii) *Suppose $\mathcal{M} \models T$, $\mathcal{N} \models T$, \mathcal{A} is an L -substructure of both \mathcal{M} and \mathcal{N} , $n > 0$, $a \in A^n$, and $\theta(x_1, \dots, x_n, y)$ is a quantifier-free L -formula. If $\theta(a, y)$ has a realisation in \mathcal{M} , then $\theta(a, y)$ has a realisation in \mathcal{N} .*

If L has a constant symbol then we can allow $n = 0$ in (ii).

PROOF. By Proposition 6.9, (i) is equivalent to showing that every formula of the form $\exists y\theta(x_1, \dots, x_n, y)$, where θ is quantifier-free, is equivalent modulo T to a quantifier-free formula. Applying Theorem 6.6 we get that this is equivalent to showing that: (*) *whenever $\mathcal{M} \models T$, $\mathcal{N} \models T$, \mathcal{A} is an L -substructure of both \mathcal{M} and \mathcal{N} , and $a \in A^n$; $\mathcal{M} \models \exists y\theta(a, y)$ if and only if $\mathcal{N} \models \exists y\theta(a, y)$.* The statement (*) clearly implies (ii). On the other hand, (ii) applied twice, once with the roles of \mathcal{M} and \mathcal{N} reversed, gives us (*). \square

Let us refine the criterion a little further. By a *literal* we mean an atomic or negated atomic formula.

COROLLARY 6.11 (Criterion for QE). *Suppose T is an L -theory satisfying the following condition:*

- (*) *Whenever \mathcal{M} and \mathcal{N} are models of T with a common substructure \mathcal{A} , and $\psi(y)$ is a conjunction of $L_{\mathcal{A}}$ -literals in a single variable, if $\psi(y)$ has a realisation in \mathcal{M} then it has a realisation in \mathcal{N} .*

Then T admits QE.

PROOF. We assume (*) and show that (ii) of Corollary 6.10 holds. Indeed, by De Morgan's laws about how negation interacts with conjunctions and disjunctions, if $\theta(x_1, \dots, x_n, y)$ is a quantifier-free formula with $n > 0$ and $a \in A^n$ then $\theta(a, y)$ is a quantifier-free $L_{\mathcal{A}}$ -formula that is equivalent (modulo the empty theory) to a disjunction of conjunctions of $L_{\mathcal{A}}$ -literals, say $\phi(y)$. If $\theta(a, y)$ has a solution on \mathcal{M} then so does some disjunct, say $\psi(y)$ of ϕ . Now by

($*$), $\psi(y)$ has a solution in \mathcal{N} . Hence so does $\phi(y)$, and thus $\theta(a, y)$. We have shown that condition (ii) of Corollary 6.10 is satisfied, so T admits QE. \square

Condition ($*$) can be used to show that a number of theories admit QE.

EXAMPLE 6.12 (The theory of infinite sets admits QE). Consider $L = \emptyset$ and T the theory of infinite sets. Suppose M and N are infinite sets with as common nonempty subset A . Note that an atomic L_A -formula in the variable y is of the form $y = y$ or $y = a$ for some $a \in A$. Of course $y = y$ is realised by everything while $y \neq y$ is realised by nothing, so the former adds no new information to a conjunction while the latter can never be a conjunct of a formula that has a solution. Similarly, any atomic L_A -formula not involving y is either true or false in A (and hence also in M and N) so either does not add any new information as a conjunct or cannot be present as a conjunct in a formula that has a solution in M . Hence a conjunction of L_A -literals in y , that has a realisation in M , is equivalent to one of the form:

$$\bigwedge_{i=1}^k (y = a_i) \wedge \bigwedge_{j=1}^{\ell} (y \neq b_j)$$

where the a_i and b_j come from A . If any of the $(y = a_i)$ actually appears in the formula then we have a solution in A and hence in N . If not, then as N is infinite we can find a solution in N . So we have shown ($*$) for T – so the theory of infinite sets admits QE.

EXAMPLE 6.13 (ACF admits QE). Suppose M and N are algebraically closed fields, R is a common subring, $\psi(y)$ is a conjunction of L_R -literals such that $\psi(y)$ has a realisation in M . We need to show that $\psi(y)$ has a solution in N . First observe that R is an integral domain (as it is a subring of a field) and hence has a unique field of fractions F , which in turn has a unique algebraic closure F^{alg} . That is, there exists an isomorphism between the algebraic closure of the fraction field of R in M and the algebraic closure of the fraction field of R in N , that fixes R pointwise. Identifying these two fields and extending R , we may as well assume that R is an algebraically closed subfield of both M and N . Now $\psi(y)$ is of the form

$$\bigwedge_{i=1}^k p_i(y) = 0 \wedge \bigwedge_{j=1}^{\ell} q_j(y) \neq 0$$

where the p_i and q_j are polynomials in one variable over R . Let b be a realisation of this in M . If any one of the p_i s are nonzero then b is algebraic over R , and hence in R , and hence in N , and we would be done. So we can assume that all the p_i s are zero polynomials. Now each q_j has only finitely many roots in N , but N is infinite as it is algebraically closed. So if we choose $b' \in N$ not equal to any of these roots, then b' will realise $\psi(y)$, as desired.

EXERCISE 6.14. Let $L = \{E\}$ where E is a binary relation symbol. Let T be the theory which says that there are infinitely many E -classes and each class is infinite. Show that T is complete and admits quantifier elimination.

EXERCISE 6.15. Let F be any fixed field and consider T the theory of infinite F -vector spaces. Show that T admits quantifier-elimination.

Similar arguments can be used to show QE for a number of theories, including the theory of torsion-free divisible abelian groups and the theory of dense linear orderings without endpoints. See chapter 3 of Marker’s “Model Theory: An introduction” (Springer 2002).

EXERCISE 6.16. Following the strategy of Theorem 6.6 prove that given an L -theory T and an L -formula $\phi(x)$ where $x = (x_1, \dots, x_n)$, $n \geq 0$, the following are equivalent:

- (i) $\phi(x)$ is equivalent to a universal formula $\psi(x)$ modulo T .
- (ii) Suppose $\mathcal{M} \models T$, $\mathcal{N} \models T$, and $\mathcal{M} \subseteq \mathcal{N}$. Then for all $a \in M^n$, $\mathcal{M} \models \phi(a)$ if $\mathcal{N} \models \phi(a)$.

6.3. Model completeness

DEFINITION 6.17. A theory T is said to be *model-complete* if whenever \mathcal{M} and \mathcal{N} are models of T , if $\mathcal{M} \subseteq \mathcal{N}$ then $\mathcal{M} \preceq \mathcal{N}$.

PROPOSITION 6.18. *Quantifier elimination implies model-completeness.*

PROOF. Suppose T admits QE, $\mathcal{M} \models T$, $\mathcal{N} \models T$, and $\mathcal{M} \subseteq \mathcal{N}$. Given a formula $\phi(x_1, \dots, x_n)$ with $n > 0$ we have a quantifier-free formula $\psi(x_1, \dots, x_n)$ that is equivalent to ϕ modulo T . Hence, for any $a \in M^n$ we have

$$\begin{aligned} \mathcal{M} \models \phi(a) &\iff \mathcal{M} \models \psi(a) \text{ as } \mathcal{M} \models T \\ &\iff \mathcal{N} \models \psi(a) \text{ as } \psi \text{ is quantifier-free and } \mathcal{M} \subseteq \mathcal{N} \\ &\iff \mathcal{N} \models \phi(a) \text{ as } \mathcal{N} \models T. \end{aligned}$$

We have shown that $\mathcal{M} \preceq \mathcal{N}$. Actually, this is not strictly speaking correct as we do still have to deal with case when $n = 0$. But then we can still write the sentence ϕ as $\phi(x)$, and the above argument still shows that for all $a \in M$, $\mathcal{M} \models \phi(a)$ if and only if $\mathcal{N} \models \phi(a)$. Note that as x does not in fact appear in ϕ , fixing $a \in M$, $\mathcal{M} \models \phi$ if and only if $\mathcal{M} \models \phi(a)$. And similarly for \mathcal{N} . So $\mathcal{M} \models \phi$ if and only if $\mathcal{N} \models \phi$. \square

Model-completeness is strictly weaker than quantifier-elimination. An example is $T = \text{Th}(\mathbb{R}, 0, 1, +, -, \times)$ which we have already seen does not admit QE (as the ordering is definable but not quantifier-free definable, cf. Example 4.32). However, it is a fact (requiring some more work, including Tarski's theorem that $\text{Th}(\mathbb{R}, 0, 1, +, -, \times, <)$ does admit QE) that T is model-complete.

One application of model-completeness is another technique for proving completeness.

PROPOSITION 6.19. *If T is model-complete and there exists a model \mathcal{M} of T which embeds into every other model of T , then T is complete.*

PROOF. Suppose $\mathcal{N}_1 \models T$ and $\mathcal{N}_2 \models T$. By assumption we have $\mathcal{M} \subseteq \mathcal{N}_i$ for $i = 1, 2$. By model-completeness, $\mathcal{M} \preceq \mathcal{N}_i$ and hence $\mathcal{M} \equiv \mathcal{N}_i$. So $\mathcal{N}_1 \equiv \mathcal{N}_2$, as desired. \square

It follows that the theory of infinite sets and the theory of an equivalence relation with infinitely many classes all infinite, are complete. Indeed, we (or rather you) have already pointed out that they admit QE and it is not hard to see that both of these theories have a model that embeds into every other model (exercise). Of course in both of these cases one can also get completeness more directly by \aleph_0 -categoricity and Vaught's test. Another example is that the theory of torsion-free divisible abelian groups is complete: one shows it admits QE by applying Corollary 6.11, and then one observes that the additive group of rational numbers embeds in every model. Once again however, completeness could have been achieved by Vaught's test, since the theory of torsion-free divisible abelian groups is

κ -categorical for any uncountable κ . To see an example where Proposition 6.19 gives us a proof of completeness while Vaught's test does not, consider the structure $(\mathbb{R}, 0, 1, +, -, \times)$. This is a model of the theory of *real closed fields* (RCF), which we have not (and will not) define. Nevertheless, let me just say that RCF is model-complete (for reasons more or less hinted at above), but not categorical in any infinite cardinal (the latter also requires proof, and is by no means obvious). See Marker's "Model Theory: An introduction" for details.

From Proposition 6.18 we get that ACF is model-complete. Here is a corollary of this fact, which is a nice note to end on.

COROLLARY 6.20 (Hilbert's Nullstellensatz). *Suppose K is an algebraically closed field and I is a proper ideal in the polynomial ring $K[X_1, \dots, X_n]$. Then there exists a tuple $a \in K^n$ such that $P(a) = 0$ for all $P \in I$.*

PROOF. Extending I to a prime ideal (for example a maximal ideal by Zorn's lemma) we may as well assume that I is prime. By Noetherianity of $K[X_1, \dots, X_n]$ we know that I is generated by a finite set of polynomials, say P_1, \dots, P_ℓ . To find a common zero for all polynomials in I it will suffice to find a common zero for P_1, \dots, P_ℓ . This is what we will do.

Consider the integral domain $K[X_1, \dots, X_n]/I$. Let F be its fraction field and let $L = F^{\text{alg}}$ be its algebraic closure. Then K is a subfield (and hence a substructure) of L . As both are models of ACF, and ACF is model-complete, we get $K \preceq L$. Now, view L as a $K[X_1, \dots, X_n]$ -algebra in the natural way. That is, let $\pi : K[X_1, \dots, X_n] \rightarrow L$ denote the quotient map $K[X_1, \dots, X_n] \rightarrow K[X_1, \dots, X_n]/I$ composed with the inclusion $K[X_1, \dots, X_n]/I \subseteq L$. For all polynomials $P \in I$, $P(\pi X_1, \dots, \pi X_n) = \pi(P(X_1, \dots, X_n)) = 0$. That is, the n -tuple $(\pi X_1, \dots, \pi X_n)$ from L is a root of all the polynomials in I . Hence

$$L \models \exists x_1 \cdots x_n \left(\bigwedge_{i=1}^{\ell} P_i(x) = 0 \right).$$

Since the polynomials P_1, \dots, P_ℓ have coefficients in K this is a formula over K . So

$$K \models \exists x_1 \cdots x_n \left(\bigwedge_{i=1}^{\ell} P_i(x) = 0 \right),$$

as desired. □

EXERCISE 6.21. Using Exercise 6.16, prove that a theory T is model-complete if and only if every formula is equivalent modulo T to a universal (equivalently existential) formula. Use this, together with Tarski's theorem that $\text{Th}(\mathbb{R}, 0, 1, +, -, \times, <)$ admits QE, to conclude that $\text{Th}(\mathbb{R}, 0, 1, +, -, \times)$ is model-complete.

Part 3

More Model Theory

CHAPTER 7

Types and Saturation

7.1. Types

Fix a language L and suppose Φ is a set of L -formulas. We write $\Phi(x_1, \dots, x_n)$ to mean that each formula in Φ has free variables from among $\{x_1, \dots, x_n\}$.

DEFINITION 7.1. Suppose $\Phi(x)$ is a set of L -formulas where $x = (x_1, \dots, x_n)$. An L -structure \mathcal{M} *realises* $\Phi(x)$ if there exists a tuple $a \in M^n$ such that $\mathcal{M} \models \phi(a)$ for all $\phi \in \Phi$ – that is, if \mathcal{M} realises all the formulas in Φ simultaneously. We say that a *realises* (or *satisfies*) $\Phi(x)$ in \mathcal{M} , and we denote it by $\mathcal{M} \models \Phi(a)$.

We say that $\Phi(x)$ is *satisfiable* (or *consistent*) if it is realised in some L -structure.

PROPOSITION 7.2. *Suppose T is an L -theory and $\Phi(x)$ is a set of L -formulas where $x = (x_1, \dots, x_n)$. The following are equivalent:*

- (i) $\Phi(x)$ is realised in some model of T .
- (ii) Every finite subset of $\Phi(x)$ is realised in some model of T .
- (iii) There is a model of T in which every finite subset of $\Phi(x)$ is realised.

PROOF. We first show that (i) and (ii) are equivalent. Let $L' = L \cup \{c\}$ where $c = (c_1, \dots, c_n)$ is an n -tuple of new constant symbols. Consider the L' -theory $T \cup \Phi(c)$.

CLAIM 7.3. *(i) is equivalent to the consistency $T \cup \Phi(c)$.*

PROOF OF CLAIM 7.3. If $\mathcal{M} \models T$ and $a \in M^n$ with $\mathcal{M} \models \Phi(a)$, then set \mathcal{M}' to be the expansion of \mathcal{M} where we set $c^{\mathcal{M}'} = a$. Then $\mathcal{M}' \models (T \cup \Phi(c))$. Conversely, if $\mathcal{M}' \models (T \cup \Phi(c))$ then set \mathcal{M} to be the reduct of \mathcal{M}' to L , and set $a := c^{\mathcal{M}'}$. Then $\mathcal{M} \models \Phi(a)$ and \mathcal{M} is a model of T . □

CLAIM 7.4. *(ii) is equivalent to the consistency of every finite subset of $T \cup \Phi(c)$.*

PROOF OF 7.4. This is exactly as in the proof Claim 7.3, but working with the finite subsets of Φ . □

Hence, by compactness, (i) and (ii) are equivalent.

Now for the equivalence of (ii) and (iii). Clearly (iii) implies (ii). So assume (ii), and let $\Sigma := T \cup \{\exists x(\phi_1(x) \wedge \dots \wedge \phi_m(x)) : m \geq 1, \phi_1, \dots, \phi_m \in \Phi(x)\}$. We need to show that Σ is consistent. We do so by compactness: let $\Delta \subseteq \Sigma$ be a finite subset. Then Δ involves only finitely many ϕ from Φ . Putting them all together we see that for some $\phi_1, \dots, \phi_N \in \Phi$, Δ is entailed by $T \cup \{\exists x(\phi_1(x) \wedge \dots \wedge \phi_N(x))\}$. By (ii), the latter is consistent, hence Δ is consistent. □

DEFINITION 7.5 (Types of a theory). Suppose T is an L -theory and $n \in \omega$. An (*partial*) n -*type* of T is a set of L -formulas $\Phi(x_1, \dots, x_n)$ for which any of the equivalent conditions

of Proposition 7.2 hold. We call Φ *complete* if in addition for every L -formula $\phi(x)$, either $\phi \in \Phi$ or $\neg\phi \in \Phi$. The set of all complete n -types of T is denoted $S_n(T)$.

REMARK 7.6. Note that if $n = 0$ then types are just consistent theories: the 0-types of T are exactly the consistent theories extending T and $S_0(T)$ is the set of complete consistent extensions of T .

As usual, completeness here corresponds to maximality:

LEMMA 7.7. *An n -type of T is complete if and only if it is a maximal n -type of T .*

PROOF. Suppose $\Phi(x) \in S_n(T)$ and $\Phi(x) \subsetneq \Phi'(x)$. Let $\phi \in \Phi' \setminus \Phi$. Then by completeness, $\neg\phi \in \Phi$. Hence Φ' contains both ϕ and $\neg\phi$, and thus can have no realisations. So Φ is a maximal n -type of T .

Conversely, suppose $\Phi(x)$ is a maximal n -type of T and let $\phi(x)$ be an L -formula not in Φ . We need to show that $\neg\phi$ is in Φ . Consider the extension $\Phi'(x) = \Phi(x) \cup \{\neg\phi(x)\}$ of $\Phi(x)$. We need to show that $\Phi' = \Phi$, for which it suffices by maximality to show that Φ' is an n -type of T . Assume it isn't and seek a contradiction. By Proposition 7.2 there is a finite subset of Φ' has no realisation in any model of T . It follows that $T \models \forall x((\phi_1(x) \wedge \cdots \wedge \phi_m(x)) \rightarrow \phi(x))$, for some $\phi_1, \dots, \phi_m \in \Phi$. As Φ is an n -type of T there is $\mathcal{M} \models T$ with $a \in M^n$ such that $\mathcal{M} \models \Phi(a)$. Hence $\mathcal{M} \models \phi(a)$. So \mathcal{M} and a witness that $\Phi(x) \cup \{\phi(x)\}$ is an n -type of T . But this contradicts the maximality of Φ since by assumption $\phi \notin \Phi$. This contradiction proves that $\Phi(x) \cup \{\neg\phi(x)\}$ is an n -type of T and hence $\neg\phi \in \Phi$ as desired. \square

COROLLARY 7.8. *Every n -type of T is contained in a complete n -type of T .*

PROOF. Suppose $\Phi(x)$ is an n -type of T . Consider the poset $(P, <)$ of all n -types of T containing Φ , ordered by containment. By Proposition 7.2, the union of a chain of n -types of T is again an n -type of T . Hence chains in P have upper bounds in P , and hence by Zorn's Lemma P has a maximal element, say $p(x)$. Then $p(x)$ is a maximal n -type of T , and hence a complete n -type of T by Lemma 7.7, that contains $\Phi(x)$. \square

We are usually going to be interested in types of theories of a structure. In that case we will also want to consider types that refer to parameters from that structure.

DEFINITION 7.9 (Types over parameters). Suppose \mathcal{M} is an L -structure and $A \subseteq M$. A (*complete* or *partial*) n -type in \mathcal{M} over A is a (complete or partial) n -type of the L_A -theory $\text{Th}(\mathcal{M}_A)$. The set of complete n -types in \mathcal{M} over A is denoted by $S_n^{\mathcal{M}}(A)$.

Suppose $b \in M^n$. The *type of b over A in \mathcal{M}* is

$$\text{tp}_{\mathcal{M}}(b/A) := \{\phi(x) : \text{where } \phi(x) \text{ is an } L_A\text{-formula with } \mathcal{M} \models \phi(b)\}.$$

When $A = \emptyset$ we often simply write $\text{tp}_{\mathcal{M}}(b)$.

REMARK 7.10. (a) An n -type in \mathcal{M} over A is thus a set of L_A -formulas in n free variables that is realised in *some* model of $\text{Th}(\mathcal{M}_A)$.

(b) By Proposition 7.2 and the completeness of $\text{Th}(\mathcal{M}_A)$ we have that a set of L_A -formulas $\Phi(x)$ is an n -type in \mathcal{M} over A if and only if every finite subset of $\Phi(x)$ is realised in \mathcal{M} .

(c) As $\text{tp}_{\mathcal{M}}(b/A)$ is realised in \mathcal{M} itself (by $b \in M^n$), we have that $\text{tp}_{\mathcal{M}}(b/A)$ is an n -type over A . Moreover it is clearly complete, so $\text{tp}_{\mathcal{M}}(b/A) \in S_n^{\mathcal{M}}(A)$.

But not all types in $S_n^{\mathcal{M}}(A)$ need be realised here. The following lemma gives us a way of obtaining other types in $S_n^{\mathcal{M}}(A)$.

LEMMA 7.11. *Suppose $\mathcal{M} \preceq \mathcal{N}$ and $A \subseteq M$. The n -types in \mathcal{M} over A agree with the n -types in \mathcal{N} over A . In particular, $S_n^{\mathcal{M}}(A) = S_n^{\mathcal{N}}(A)$.*

PROOF. A set of L_A -formulas is an n -type in \mathcal{M} over A if and only if it is realised in a model of $\text{Th}(\mathcal{M}_A)$. A set of L_A -formulas is an n -type in \mathcal{N} over A if and only if it is realised in a model of $\text{Th}(\mathcal{N}_A)$. So it suffices to observe that $\text{Th}(\mathcal{M}_A) = \text{Th}(\mathcal{N}_A)$. Indeed if $\mathcal{M}_A \models \sigma$ then $\sigma = \phi(a_1, \dots, a_m)$ where $\phi(y)$ is an L -formula, $a \in A^m$, and $\mathcal{M} \models \phi(a)$. As $\mathcal{M} \preceq \mathcal{N}$, $\mathcal{N} \models \phi(a)$ also. Hence $\mathcal{N}_A \models \sigma$. So $\text{Th}(\mathcal{M}_A) = \text{Th}(\mathcal{N}_A)$. \square

So we can get elements of $S_n^{\mathcal{M}}(A)$ that may not be realised in \mathcal{M} as follows: let \mathcal{N} be an elementary extension of \mathcal{M} and let $b \in N^n$. Then $\text{tp}_{\mathcal{N}}(b/A) \in S_n^{\mathcal{N}}(A) = S_n^{\mathcal{M}}(A)$. For example, we know that $\mathcal{Q} := (\mathbb{Q}^{\text{alg}}, 0, 1, +, -, \times) \preceq (\mathbb{C}, 0, 1, +, -, \times) =: \mathcal{C}$. Hence, $p(x) = \text{tp}_{\mathcal{C}}(\pi)$ is a complete 1-type in \mathcal{Q} (over the emptyset). Since $Q(x) \neq 0$ is in $p(x)$ for every nonzero polynomial Q with integer coefficients, $p(x)$ is not realised in \mathcal{Q} .

In fact, this is how all complete types come about:

PROPOSITION 7.12. *Suppose \mathcal{M} is an L -structure, $A \subseteq M$, and $p(x)$ is a set of L_A -formulas in free variables $x = (x_1, \dots, x_n)$. Then $p(x) \in S_n^{\mathcal{M}}(A)$ if and only if $p(x) = \text{tp}_{\mathcal{N}}(b/A)$ for some $\mathcal{N} \succeq \mathcal{M}$ and some $b \in N^n$.*

PROOF. The right-to-left implication is by Lemma 7.11, as explained in the previous paragraph. For the converse, suppose $p(x) \in S_n^{\mathcal{M}}(A)$.

CLAIM 7.13. $\Phi(x) := p(x) \cup \text{Th}(\mathcal{M}_M)$ is satisfiable.

PROOF OF 7.13. By Proposition 7.2 (talking T to be the empty theory), it suffices to show that every finite subset of $\Phi(x)$ is satisfiable. Since both $p(x)$ and $\text{Th}(\mathcal{M}_M)$ are closed under conjunctions we reduce to showing the following: given $\phi(x, \underline{a}) \in p(x)$ and $\psi(\underline{e}) \in \text{Th}(\mathcal{M}_M)$, there exists an L_M -structure in which $\psi(\underline{e})$ is true and $\phi(x, \underline{a})$ is realised. Here ϕ and ψ are L -formulas, a is a tuple from A , and e is a tuple from M . Let us write $e = (f, g)$ where f is a tuple from A and g is a tuple from $M \setminus A$.

Let \mathcal{R} be a model of $\text{Th}(\mathcal{M}_A)$ that realises $p(x)$, say by $r \in R^n$ – possible as $p(x)$ is a type over A . So $\mathcal{R} \models \phi(r, \underline{a})$. Now $\text{Th}(\mathcal{M}_A) \models \exists w \psi(\underline{f}, w)$. Hence, $\mathcal{R} \models \exists w \psi(\underline{f}, w)$. Let g' be a tuple from R witnessing this. Expand \mathcal{R} from an L_A -structure to an L_M -structure, \mathcal{R}' , in such a way that the interpretation of \underline{g} in \mathcal{R}' is g' . Then $\mathcal{R}' \models \psi(\underline{e})$ and $\mathcal{R}' \models \phi(r, \underline{a})$, as desired. \square

Let \mathcal{N}' be an L_M -structure in which $\Phi(x)$ is realised, say by $b \in N^n$. Let \mathcal{N} be the reduct of \mathcal{N}' to L . Since \mathcal{N} can be expanded to a model of $\text{Th}(\mathcal{M}_M)$, we have an elementary embedding of \mathcal{M} in \mathcal{N} (cf. Proposition 4.45). On the other hand, $\mathcal{N} \models p(b)$ and so $p(x) \subseteq \text{tp}_{\mathcal{N}}(b/A)$. By maximality of $p(x)$ (Lemma 7.7), $p(x) = \text{tp}_{\mathcal{N}}(b/A)$, as desired. \square

EXERCISE 7.14. Suppose T is a complete consistent theory. Then $S_n(T) = S_n^{\mathcal{M}}(\emptyset)$ for any $\mathcal{M} \models T$.

EXERCISE 7.15. Suppose $\mathcal{M} \preceq \mathcal{N}$, $A \subseteq M$, and $b \in M^n$. Then $\text{tp}_{\mathcal{M}}(b/A) = \text{tp}_{\mathcal{N}}(b/A)$.

EXERCISE 7.16. Given any finite set of types $p_1, \dots, p_\ell \in S_n^{\mathcal{M}}(A)$ there is an elementary extension of \mathcal{M} in which they are all realised.

When it is clear what the ambient structure is we may drop the subscript \mathcal{M} in $\text{tp}_{\mathcal{M}}$. By Exercise 7.15, we might even do so when the ambient structure is replaced by an elementary extension – since then it does not matter in which of the two structures we are computing the type.

7.2. Some examples

By Proposition 7.12, complete types in \mathcal{M} tell us what first-order properties tuples in elementary extensions of \mathcal{M} can have. In this section we describe the complete types for two examples: dense linear orderings without endpoint and algebraically closed fields.

The following easy lemma will be useful.

LEMMA 7.17. *Suppose T is a theory that admits quantifier elimination, $\mathcal{M} \models T$, $A \subseteq M$, and $b, b' \in M^n$. If b and b' realise the same atomic L_A -formulas then $\text{tp}(b/A) = \text{tp}(b'/A)$.*

PROOF. Note that as b and b' realise the same atomic L_A -formulas they also realise the same quantifier-free L_A -formulas. Indeed, one shows by a very straightforward induction on the quantifier-free L_A -formula $\theta(x)$, that $\mathcal{M} \models \theta(b) \iff \mathcal{M} \models \theta(b')$.

Now suppose $\psi(x, \underline{a})$ is an arbitrary L_A -formula, where $\psi(x, y)$ is an L -formula and $a \in A^\ell$. Let $\theta(x, y)$ be a quantifier-free L -formula that is equivalent to ψ modulo T . Then

$$\begin{aligned} \psi(x, \underline{a}) \in \text{tp}(b/A) &\iff \mathcal{M} \models \psi(b, a) \\ &\iff \mathcal{M} \models \theta(b, a) \quad \text{as } \mathcal{M} \models T \text{ and } \psi \text{ is } T\text{-equivalent to } \theta \\ &\iff \mathcal{M} \models \theta(b', a) \quad \text{as } \theta(x, \underline{a}) \text{ quantifier-free } L_A\text{-formula} \\ &\iff \mathcal{M} \models \psi(b', a) \\ &\iff \psi(x, \underline{a}) \in \text{tp}(b'/A) \end{aligned}$$

as desired. □

7.2.1. Dense linear orderings without endpoints. Consider DLO, the theory of dense linear orderings without endpoints, in the language $L = \{<\}$. Note that the only atomic L -formulas are $x < x$ and $x = x$. Hence, all elements in all models of DLO satisfy the same atomic formulas, namely the single atomic formula $(x = x)$. By QE for DLO and Lemma 7.17, this means that all elements in all models of DLO satisfy the same type (over the emptyset). That is, there is a unique complete 1-type of DLO.

Now suppose that $\mathcal{M} \models \text{DLO}$, $A \subseteq M$ is nonempty, and $p(x) \in S_1^{\mathcal{M}}(A)$. If p is realised by some $a \in A$ then $(x = a) \in p(x)$ and a is the only realisation of $p(x)$ (in any elementary extension of \mathcal{M}). So in this case we know exactly what $p(x)$ is, it carries the same content as the formula $(x = a)$. (We say that $p(x)$ is *isolated* by the formula $(x = a)$.)

So let us assume that $p(x) \in S_1^{\mathcal{M}}(A)$ is not realised in A . Then let $L_p = \{a \in A : (a < x) \in p(x)\}$ and $U_p = \{a \in A : (x < a) \in p(x)\}$. Since p is satisfiable we must have that $L_p < U_p$. So p determines a *cut* in $(A, <)$ – a decomposition of A into two disjoint sets, every element in the first below every element in the second. We show that this is a bijective correspondence between complete 1-types and cuts.

Suppose that two complete 1-types, $p(x)$ and $q(x)$, both not realised in A , determine the same cut, (L, U) . Let $\mathcal{N} \succeq \mathcal{M}$ be such that both p and q are realised in \mathcal{N} . So there are $b, b' \in N$ with $p = \text{tp}_{\mathcal{N}}(b/A)$, $q = \text{tp}_{\mathcal{N}}(b'/A)$. Now the atomic L_A -formulas are exactly those

of the form $(x = a)$, $(x < a)$, or $(a < x)$, for $a \in A$. Since p and q are not realised in A , b and b' can only realise formulas of the latter two forms. That p and q determine the same cut in A therefore means that b and b' realise exactly the same atomic L_A -formulas. By QE for DLO and Lemma 7.17, $p = q$.

Finally, given a cut (L, U) of $(A, <)$, let $\Phi(x) = \{(a < x) : a \in L\} \cup \{(x < a) : a \in U\}$. the axioms of DLO ensure that every finite subset of $\Phi(x)$ is realised in \mathcal{M} . Hence $\Phi(x)$ is a 1-type in $\text{Th}(\mathcal{M}_A)$ and so, by Corollary 7.8, can be extended to a complete 1-type $p(x) \in S_{\mathcal{M}}(A)$. Clearly the cut determined by p must be (L, U) .

We have shown:

PROPOSITION 7.18. *There is a unique complete 1-type of DLO. As for 1-types over paramaters, if $\mathcal{M} \models \text{DLO}$ and $A \subseteq M$ is nonempty, then the types in $S_1^{\mathcal{M}}(A)$ that are not realised in A are in bijective correspondence with cuts in $(A, <)$. This correspondence is given by $p \mapsto (L_p, U_p)$ where $L_p := \{a \in A : (a < x) \in p(x)\}$ and $U_p = \{a \in A : (x < a) \in p(x)\}$. \square*

Note in particular that $|S_1^{\mathcal{Q}}(\mathbb{Q})| = 2^{\aleph_0}$ where $\mathcal{Q} = (\mathbb{Q}, <) \models \text{DLO}$. Indeed, this is because every real number determines a unique cut in $(\mathbb{Q}, <)$ and so there must be at least 2^{\aleph_0} -many complete 1-types, while an easy cardinal arithmetic argument shows that 2^{\aleph_0} is always the maximum possible number of complete types over countably many parameters in a countable language.

7.2.2. Algebraically closed fields. Next let us consider ACF the theory of algebraically closed fields in the language $L = \{0, 1, +, -, \times\}$. Suppose $K \models \text{ACF}$ and $A \subseteq K$.¹ We wish to describe $S_n^K(A)$. The first step is to see that if k denotes the subfield of K generated by A then we can replace A by k . Note that if $\text{char}(K) = 0$ then k is a field extension of \mathbb{Q} and that if $\text{char}(K) = p > 0$ then k is an extension of \mathbb{F}_p . First some notation:

DEFINITION 7.19 (Restriction of a type). Suppose $p(x) \in S_n^{\mathcal{M}}(B)$ for some L -structure \mathcal{M} and $B \subseteq M$. If $A \subseteq B$ then $p \upharpoonright A := \{\phi(x) \in p(x) : \phi \text{ is an } L_A\text{-formula}\}$. Note that $p \upharpoonright B \in S_n^{\mathcal{M}}(A)$.

LEMMA 7.20. *The map $p \mapsto p \upharpoonright A$ is a bijection between $S_n^K(k)$ and $S_n^K(A)$.*

PROOF. Note that every element of $S_n^K(A)$ is a partial n -type over k , and hence can be extended to a complete n -type over k (by Lemma 7.8). So the restriction map is surjective. What we need to show is that if p and q are complete types over k that agree over A , then $p = q$. This will follow from the following:

CLAIM 7.21. *Every L_k -formula is equivalent in K to an L_A -formula.*

PROOF OF CLAIM 7.21. Suppose $\phi(x)$ is an L_k -formula where $x = (x_1, \dots, x_n)$. By quantifier elimination for ACF we may assume that ϕ is quantifier-free. By a straightforward induction we reduce to the case when ϕ is an atomic L_k -formula, that is, when $\phi(x)$ is (equivalent to a formula) of the form $P(x) = 0$ where P is a polynomial in n variables over k . In multi-index notation we can write $P(X) = \sum_{\alpha} b_{\alpha} X^{\alpha}$ where $b_{\alpha} \in k$. As k is generated

¹We will start using this abuse of notation more frequently now: we do not notationally distinguish between the structure and its universe unless we need to.

by A we have $a \in A^m$ and $g_\alpha \in \mathbb{Z}(Y_1, \dots, Y_m)$ such that $b_\alpha = g_\alpha(a)$ for each α . Let Q be the product of all the denominators of the g_α s, so that $g'_\alpha := Q(Y)g_\alpha(Y) \in \mathbb{Z}[Y]$. Note that $P(x)$ and $Q(a)P(x) = \sum_\alpha g'_\alpha(a)X^\alpha$ have the same roots. Hence $P(x) = 0$ defines the same set in K^n as $\sum_\alpha g'_\alpha(a)x^\alpha = 0$, while the latter is an L_A -formula. \square

Given any L_k -formula $\phi(x)$ let us denote by $\phi'(x)$ an L_A -formula that is equivalent to ϕ in K . Let $L \succeq K$ and $b, c \in L^n$ such that $p(x) = \text{tp}_L(b/k)$ and $q(x) = \text{tp}_L(c/k)$. Then

$$\begin{aligned} \phi(x) \in p(x) &\iff L \models \phi(b) \\ &\iff L \models \phi'(b) \quad \text{as } \phi \text{ and } \phi' \text{ are equivalent in } K \text{ and hence in } L \succeq K \\ &\iff \phi' \in \text{tp}_L(b/A) = p \upharpoonright A = q \upharpoonright A = \text{tp}_L(c/A) \\ &\iff L \models \phi'(c) \\ &\iff L \models \phi(b) \\ &\iff \phi(x) \in q(x) \end{aligned}$$

so that $p = q$, as desired. This completes the proof of Lemma 7.20. \square

EXERCISE 7.22. Suppose $\mathcal{B} \subseteq \mathcal{M}$ is a substructure of some L -structure generated by a nonempty subset $A \subseteq M$. Then the map $p \mapsto p \upharpoonright A$ is a bijection between $S_n^{\mathcal{M}}(B)$ and $S_n^{\mathcal{M}}(A)$. (*Hint: This is a generalisation of part of 7.20 and can be proved analogously, using Lemma 6.4.*)

Now for our characterisation of complete types in algebraically closed field.

PROPOSITION 7.23. *Suppose k is a subfield of $K \models \text{ACF}$. There is a bijective correspondence between $S_n^K(k)$ and primed ideals in $k[X]$, where $X = (X_1, \dots, X_n)$. This correspondence is given by $p(x) \mapsto \{f(X) \in k[X] : (f(x) = 0) \in p(x)\}$.*

PROOF. Given $p(x) \in S_n^K(k)$ let $I_p := \{f(X) \in k[X] : (f(x) = 0) \in p(x)\}$. Let $L \succeq K$ be an elementary extension where p is realised, say by $b \in L^n$. Then $I_p = \{f(X) \in k[X] : f(b) = 0\}$. From this description of I_p it follows immediately that I_p is a prime ideal.

Next we show that this association $p \mapsto I_p$ is injective. This will follow from quantifier elimination since the atomic L_k -formulas are exactly those of the form $f(x) = 0$ where f is a polynomial over k . So, if $p, q \in S_n^K(k)$ with $I_p = I_q$, and $L \succeq K$ with $p = \text{tp}_L(b/k)$ and $q = \text{tp}_L(c/k)$, then b and c realise exactly the same atomic L_k -formulas. Hence, by QE and Lemma 7.17, $p = q$.

Finally, given a prime ideal $P \subset k[X]$ we need to find $p \in S_n^K(k)$ with $I_p = P$. Now, by a standard fact from commutative algebra (the fully-faithfulness of $K[X]$ over $k[X]$), there exists a prime ideal in $K[X]$ lying over P . That is, there is prime Q in $K[X]$ with $Q \cap k[X] = P$. Consider the quotient ring $R = K[X]/Q$, and let $b = (b_1, \dots, b_n)$ denote the image of $X = (X_1, \dots, X_n)$ in R . As Q is prime, R is an integral domain and is thus contained in its fraction field F . Let L be the algebraic closure of F . Then we have $K \subseteq L$, and so by model-completeness of ACF, as both K and L are algebraically closed, we have that $K \preceq L$. On the other hand, for any $f(X) \in k[X]$, $f(b) = 0$ if and only if $f \in Q$ if and only if $f \in P$. Hence, setting $p(x) := \text{tp}_L(b/k) \in S_n^L(k) = S_n^K(k)$, we have that $I_p = P$, as desired. \square

So the complete types over parameters in algebraically closed fields correspond to prime ideals. In particular this tells us that if $K \models \text{ACF}$ and $A \subseteq K$ with $|A| = \kappa$ an infinite cardinal, then $|S_n^K(A)| = \kappa$. Indeed, it suffices to count the number of prime ideals in $k[X_1, \dots, X_n]$, where k is the field generated by A . But $k[X_1, \dots, X_n]$ is of size κ , and so, as every ideal is finitely generated, there are only κ -many ideals in $k[X_1, \dots, X_n]$. So, unlike the situation with DLO where one had the maximal possible number of types, in ACF we have the minimum possible number of types.

In algebraic geometry, there is a correspondence between prime ideals $P \subset k[X]$ and k -irreducible Zariski closed subsets $V \subseteq K^n$ – the correspondence is by

$$P \longmapsto V(P) := \{b \in K^n : f(b) = 0 \text{ for all } f \in P\}.$$

Hence we get a bijective correspondence between $S_n^K(k)$ and the k -irreducible Zariski closed subsets of K^n , namely $p \mapsto V(I_p)$. We say that p is the *generic type of $V(I_p)$ over k* and that $V(I_p)$ is the *k -locus of p* .

7.3. Saturation and homogeneity

DEFINITION 7.24 (Saturation). Suppose κ is an infinite cardinal. An L -structure \mathcal{M} is κ -saturated if for all $A \subseteq M$ with $|A| < \kappa$, and all $n \geq 1$, every type $p(x) \in S_n^{\mathcal{M}}(A)$ is realised in \mathcal{M} . We say that \mathcal{M} is *saturated* if it is $|M|$ -saturated.

- REMARK 7.25.** (a) As every partial type can be extended to a complete type, κ -saturation implies that every n -type over $< \kappa$ -many parameters is realised.
 (b) If \mathcal{M} is infinite and κ -saturated then $|M| \geq \kappa$. Indeed, consider $\{(x \neq a) : a \in M\}$.

In fact, κ -saturation can be checked by considering only 1-types:

PROPOSITION 7.26. *Suppose κ is an infinite cardinal. If every 1-type in \mathcal{M} over fewer than κ -many parameters is realised in \mathcal{M} , then \mathcal{M} is κ -saturated.*

PROOF. We show by induction on n that if $p(x) \in S_n^{\mathcal{M}}(A)$ where $|A| < \kappa$, then $p(x)$ is realised in \mathcal{M} . The case of $n = 1$ is just the assumption. Suppose $n > 1$ and write $x = (x_1, \dots, x_n)$. Let $q \subseteq p$ be the set of those L_A -formulas in p whose free variables are from among (x_1, \dots, x_{n-1}) . By the induction hypothesis, q has a realisation, say $b \in M^{n-1}$.

Let $r(x_n) := \{\phi(b, x_n) : \phi(x) \in p(x)\}$. We claim that r is a 1-type over $A \cup \{b_1, \dots, b_{n-1}\}$. As r is closed under conjunctions, it suffices to show that each $\phi(b, x_n) \in r$ is realised in \mathcal{M} . As $\{\phi(x), \neg \exists x_n \phi(x)\}$ is not satisfiable, and as p is a complete type containing ϕ , it must be that $\exists x_n \phi(x)$ is in p . Hence $\exists x_n \phi(x)$ is in q . So $\mathcal{M} \models \exists x_n \phi(b, x_n)$, as desired.

By assumption, since $r(x_n)$ is a 1-type still over $< \kappa$ -many parameters, it is realised in \mathcal{M} , say by c . Then, by the definition of $r(x_n)$, (b, c) realises $p(x)$. \square

COROLLARY 7.27. *Suppose κ is an infinite cardinal, $K \models \text{ACF}$, and $\mathbb{F} \subseteq K$ is the prime subfield (so \mathbb{F}_p in characteristic p and \mathbb{Q} in characteristic zero). Then K is κ -saturated if and only if $\text{trdeg}(K/\mathbb{F}) \geq \kappa$*

PROOF. Suppose $\text{trdeg}(K/\mathbb{F}) < \kappa$ and let A be a transcendence basis for K over \mathbb{F} . So $K = \mathbb{F}(A)^{\text{alg}}$. Hence the 1-type

$$\Phi(x) := \{P(x) \neq 0 : P \in \mathbb{F}(A)[X], P \neq 0\}$$

is not realised in K . But $\Phi(x)$ is over A and hence over $< \kappa$ -many parameters. So K is not κ -saturated.

Conversely, suppose $\text{trdeg}(K/\mathbb{F}) \geq \kappa$ and let $A \subset K$ be a set with $|A| < \kappa$. So $K \neq \mathbb{F}(A)^{\text{alg}}$. Let $k := \mathbb{F}(A)$ and suppose $p(x) \in S_1^K(k)$. It suffices to show that p is realised in K . Let $L \succeq K$ with $p = \text{tp}_L(b/k)$. If $b \in k^{\text{alg}}$ then $b \in K$ and we are done. So assume that $b \notin k^{\text{alg}}$. Hence $I_p = \{f(X) \in k[X] : f(b) = 0\} = (0)$. Let $b' \in K \setminus k^{\text{alg}}$. Then $I_{\text{tp}_K(b'/k)} = (0)$ also. Hence $p = \text{tp}_K(b'/k)$, as desired. \square

Let us now consider the question of the existence of κ -saturated models in general.

PROPOSITION 7.28. *Suppose κ is an infinite cardinal. Every L -structure \mathcal{M} has a κ -saturated elementary extension.*

PROOF. We will actually construct a κ^+ -saturated model, as the regularity of κ^+ will be helpful. We begin with the following claim.

CLAIM 7.29. *Every L -structure \mathcal{M} has an elementary extension \mathcal{N} such that every type in \mathcal{M} is realised in \mathcal{N} .*

PROOF OF CLAIM 7.29. It suffices to find \mathcal{N} realising all complete n -types over M , for all n . Enumerate all the complete types in \mathcal{M} over M , $\{p_\alpha : \alpha < \lambda\}$. We define, recursively, an elementary chain of L -structures $(\mathcal{M}_\alpha : \alpha < \lambda)$ as follows:

- $\mathcal{M}_0 := \mathcal{M}$
- For all $\alpha < \lambda$, $\mathcal{M}_{\alpha+1}$ is an elementary extension of \mathcal{M}_α realising p_α . This is possible because \mathcal{M}_α is assumed to be an elementary extension of \mathcal{M} , and hence p_α is a type in \mathcal{M}_α also, and hence realisable in some elementary extension.
- For limit $\alpha < \lambda$, $\mathcal{M}_\alpha := \bigcup_{\beta < \alpha} \mathcal{M}_\beta$. I leave it to you to check that $\mathcal{M} \preceq \mathcal{M}_\alpha$.

Then, setting $\mathcal{N} := \bigcup_{\alpha < \lambda} \mathcal{M}_\alpha$, we have again that $\mathcal{M} \preceq \mathcal{N}$, but now every type in \mathcal{M} is realised in \mathcal{N} . \square

Now we iterate the claim κ^+ -many times. That is, we define recursively an elementary chain of L -structures $(\mathcal{M}_\alpha : \alpha < \kappa^+)$ as follows:

- $\mathcal{M}_0 := \mathcal{M}$
- For all ordinals $\alpha < \kappa^+$, $\mathcal{M}_{\alpha+1}$ is an elementary extension of \mathcal{M}_α realising all the types in \mathcal{M}_α . This is possible by the claim.
- For limit $\alpha < \kappa^+$, $\mathcal{M}_\alpha := \bigcup_{\beta < \alpha} \mathcal{M}_\beta$.

Then set $\mathcal{N} := \bigcup_{\alpha < \kappa^+} \mathcal{M}_\alpha$. We have again that $\mathcal{M} \preceq \mathcal{N}$. Suppose $A \subseteq N$ with $|A| < \kappa^+$ and $p(x) \in S_1^{\mathcal{N}}(A)$. By the regularity of κ^+ we have that $A \subseteq M_\alpha$ for some $\alpha < \kappa^+$. As $\mathcal{M}_\alpha \preceq \mathcal{N}$, $p \in S_1^{\mathcal{M}_\alpha}(A)$. Hence, by construction, $p = \text{tp}_{\mathcal{M}_{\alpha+1}}(b/A)$ for some $b \in M_{\alpha+1}$. As $\mathcal{M}_{\alpha+1} \preceq \mathcal{N}$, $p = \text{tp}_{\mathcal{N}}(b/A)$. So p is realised in \mathcal{N} . We have shown that \mathcal{N} is κ^+ -saturated (and hence κ -saturated). This completes the proof of Proposition 7.28. \square

In fact, if we paid more attention to the proof and made some effort to produce as small an \mathcal{N} as possible², then we would get that *every* L -structure \mathcal{M} , with $|M| < \kappa$, has a κ^+ -saturated elementary extension of size 2^κ . The reason this is interesting is that, if we further assume the Generalised Continuum Hypothesis, then we would have a κ^+ -saturated elementary extension of size κ^+ . So, assuming GCH, and this more refined version of 7.28 mentioned, we obtain that every L -structure has arbitrarily large saturated elementary extensions. Saturated structures satisfy some additional properties that we will discuss later. Of course, we do *not* want to assume GCH.

We move now to the notion of homogeneity, which begins with a discussion of *partial* elementary maps.

DEFINITION 7.30 (Partial elementary map). Suppose $f : A \rightarrow N$ where \mathcal{M} and \mathcal{N} are L -structures and $A \subseteq M$. We say that f is a *partial elementary map* (or *p.e.m.* for short) if for all L -formulas $\phi(x_1, \dots, x_n)$ and all $a_1, \dots, a_n \in A$,

$$\mathcal{M} \models \phi(a_1, \dots, a_n) \iff \mathcal{N} \models \phi(f(a_1), \dots, f(a_n)).$$

REMARK 7.31. (a) The empty function is a p.e.m. if and only if $\mathcal{M} \equiv \mathcal{N}$.
(b) If $A = M$ then f is a p.e.m. if and only if it is an elementary embedding.

Suppose $\Phi(x)$ is an n -type in \mathcal{M} over A and $f : A \rightarrow N$ is a p.e.m. Then we define

$$f(\Phi) := \{\phi(x, f(a)) : \phi(x, a) \in \Phi\}.$$

LEMMA 7.32. *If Φ is an n -type in \mathcal{M} over A and $f : A \rightarrow N$ is a p.e.m., then $f(\Phi)$ is an n -type in \mathcal{N} over $f(A)$.*

PROOF. Indeed, if $\phi_1(x, f(a_1)), \dots, \phi_\ell(x, f(a_\ell)) \in f(\Phi)$ then $\phi_1(x, a_1), \dots, \phi_\ell(x, a_\ell) \in \Phi$ and so $\mathcal{M} \models \exists x \bigwedge_{i=1}^{\ell} \phi_i(x, a_i)$ as Φ is a type in \mathcal{M} , and hence $\mathcal{N} \models \exists x \bigwedge_{i=1}^{\ell} \phi_i(x, f(a_i))$ as f is a p.e.m. So every finite subset of $f(\Phi)$ is satisfiable in \mathcal{M} , as desired. \square

LEMMA 7.33. *Suppose $f : A \rightarrow N$ is a p.e.m. where \mathcal{M} and \mathcal{N} are L -structures and $A \subseteq M$. Suppose $b \in M$ and $b' \in N$. We can extend f to a p.e.m. taking b to b' if and only if b' realises $f(\text{tp}_{\mathcal{M}}(b/A))$.*

PROOF. Note that, by Lemma 7.32, $f(\text{tp}_{\mathcal{M}}(b/A))$ is a 1-type in \mathcal{N} over $f(A)$. We do the right-to-left direction; the converse, which is obtained by reversing the steps, is left as an exercise.

Assume b' realises $f(\text{tp}_{\mathcal{M}}(b/A))$ and let f' be the extension of f such that $f'(b) = b'$. Suppose $\phi(x_1, \dots, x_n)$ is an L -formula. Given $a_1, \dots, a_n \in A \cup \{b\}$ we need to show that $\mathcal{M} \models \phi(a_1, \dots, a_n)$ if and only if $\mathcal{N} \models \phi(f'(a_1), \dots, f'(a_n))$. If none of the a_i are b , then this follows from the fact that f is a p.e.m. If more than one of the a_i are b , say $a_1 = a_2 = b$, then by considering $\psi(x_1, \dots, x_{n-1})$ given by $\phi(x_1, x_1, x_2, \dots, x_{n-1})$ we can finish by induction on the number of times that b appears among the a_i s. So we may assume that b appears exactly

²I leave it to you to look up the proof of this more refined result. Besides some cardinal arithmetic one uses the downward Löwenheim-Skolem Theorem.

once. Permuting the coordinates we may assume that $a_1 = b$. But then,

$$\begin{aligned}
\mathcal{M} \models \phi(b, a_2, \dots, a_n) &\iff \phi(x_1, a_2, \dots, a_n) \in \text{tp}(b/A) \\
&\iff \phi(x_1, f(a_2), \dots, f(a_n)) \in f(\text{tp}(b/A)) \\
&\iff \mathcal{M} \models \phi(b', f(a_2), \dots, f(a_n)) \\
&\iff \mathcal{M} \models \phi(f'(b), f'(a_2), \dots, f'(a_n))
\end{aligned}$$

as desired. \square

DEFINITION 7.34 (Homogeneity). Suppose κ is an infinite cardinal. An L -structure \mathcal{M} is κ -homogeneous if whenever $f : A \rightarrow M$ is a p.e.m. where $A \subseteq M$ with $|A| < \kappa$, and $b \in M$, then f can be extended to a p.e.m. whose domain is $A \cup \{b\}$. We say that \mathcal{M} is homogeneous if it is $|M|$ -homogeneous

PROPOSITION 7.35. κ -saturation implies κ -homogeneity.

PROOF. Given such a p.e.m. $f : A \rightarrow M$ and $b \in M$, consider the 1-type $f(\text{tp}(b/A))$. By κ -saturation, and the fact that $|f(A)| = |A| < \kappa$, $f(\text{tp}(b/A))$ has a realisation $b' \in M$. Applying Lemma 7.33 we get our desired extension of f to a p.e.m., f' , with $f'(b) = b'$. \square

A much more useful property than homogeneity is:

DEFINITION 7.36 (Strong homogeneity). An L -structure \mathcal{M} is strongly κ -homogeneous if whenever $f : A \rightarrow M$ is a p.e.m. with $A \subseteq M$ of cardinality less than κ , f extends to an automorphism of \mathcal{M} . We say \mathcal{M} is strongly homogeneous if it is strongly $|M|$ -homogeneous.

It is not the case that κ -saturation implies strong κ -homogeneity. However,

PROPOSITION 7.37. Saturation implies strong homogeneity.

PROOF. Suppose \mathcal{M} is a saturated structure of size κ , and $f : A \rightarrow M$ is a p.e.m. where $A \subseteq M$ is of cardinality less than κ . We can iterate κ -homogeneity to obtain an elementary embedding of \mathcal{M} into itself that extends f . But this is not the same thing as an automorphism, which is what we seek. In order to ensure that our extension is surjective one must pursue a back-and-forth construction. Enumerate $M \setminus A = \{b_\alpha : \alpha < \kappa\}$, and build a chain of p.e.m.s, $(f_\alpha : \alpha < \kappa)$, each with domain of cardinality less than κ , recursively as follows:

- $f_0 := f$.
- If b_α is in the domain of f_α then set $f'_{\alpha+1} = f_\alpha$. If not, then $f'_{\alpha+1}$ is the p.e.m. which extends f_α by sending b_α to a realisation of $f_\alpha(\text{tp}(b_\alpha/\text{Dom}(f_\alpha)))$. As $|\text{Dom}(f_\alpha)| < \kappa$, κ -saturation of \mathcal{M} , together with Lemmas 7.32 and 7.33, make this possible.
- If b_α is in the range of $f'_{\alpha+1}$, then we set $f_{\alpha+1} = f'_{\alpha+1}$. If not, then $f_{\alpha+1}$ is the p.e.m. which extends $f'_{\alpha+1}$ by sending a realisation of $f'^{-1}_{\alpha+1}(\text{tp}(b_\alpha/\text{Range}(f'_{\alpha+1})))$ to b_α . As $|\text{Range}(f'_{\alpha+1})| < \kappa$, κ -saturation of \mathcal{M} , together with Lemmas 7.32 and 7.33, make this possible.
- At limits we take unions. I leave it to you to check that the union of a chain of p.e.m.s is a p.e.m.

The construction ensures that b_α is in both the domain and range of $f_{\alpha+1}$, for all $\alpha < \kappa$. So, if we set $g := \bigcup_{\alpha < \kappa} f_\alpha$ then we get an automorphism of \mathcal{M} extending f , as desired. \square

This is very nice, except that without assuming GCH we do not know that there always exists a saturated elementary extension. So the best we can do is:

THEOREM 7.38. *Suppose κ is an infinite cardinal. Every L -structure has a κ -saturated and strongly κ -homogeneous elementary extension.*

PROOF. Given \mathcal{M} we actually build a κ^+ -saturated and strongly κ^+ -homogeneous elementary extension. We build an elementary chain of extensions $\{\mathcal{M}_\alpha : \alpha < \kappa^+\}$ recursively as follows:

- $\mathcal{M}_0 := \mathcal{M}$
- $\mathcal{M}_{\alpha+1}$ is an $|M_\alpha|^+$ -saturated elementary extension of \mathcal{M}_α , which exists by Proposition 7.28.
- At limits we take unions.

Let $\mathcal{N} := \bigcup_{\alpha < \kappa^+} \mathcal{M}_\alpha$. So $\mathcal{M} \leq \mathcal{N}$. Moreover \mathcal{N} is κ^+ -saturated: any type in \mathcal{N} over $\leq \kappa$ -many parameters is by regularity of κ^+ a type in some $\mathcal{M}_{\alpha+1}$ with parameters in \mathcal{M}_α , and hence realised in $\mathcal{M}_{\alpha+1}$ by $|M_\alpha|^+$ -saturation of the latter.

It remains to show that \mathcal{N} is strongly κ^+ -homogeneous. Suppose $f : A \rightarrow N$ is a p.e.m. with $|A| \leq \kappa$. We need to extend f to an automorphism. First note that, by regularity of κ^+ again, both A and its image under f lie in M_α for some $\alpha < \kappa^+$. Viewing f as a p.e.m. in $\mathcal{M}_{\alpha+1}$, and using the $|M_\alpha|^+$ -saturation, we can, exactly as in the proof of Proposition 7.37, extend f to a p.e.m. f_α in $\mathcal{M}_{\alpha+1}$, whose domain and range contain M_α . Now, working in $\mathcal{M}_{\alpha+1}$ we can extend f_α to $f_{\alpha+1}$ with domain and range containing $M_{\alpha+1}$. We can iterate this recursively, taking unions at limits, to produce a chain of p.e.m.s, f_β in $\mathcal{M}_{\beta+1}$ for all $\alpha \leq \beta < \kappa^+$, whose domain and range both contain M_β . Letting $g := \bigcup_{\alpha \leq \beta < \kappa^+} f_\alpha$, we get an automorphism of \mathcal{N} that extends f , as desired. \square

We collect together here the principal uses of saturation and homogeneity.

THEOREM 7.39. *Suppose \mathcal{M} is κ -saturated and strongly κ -homogeneous.*

- (a) *If $\mathcal{N} \equiv \mathcal{M}$ and $|N| \leq \kappa$, then there is an elementary embedding of \mathcal{N} into \mathcal{M} .*
- (b) *Suppose $b, b' \in M^n$ and $A \subset M$ with $|A| < \kappa$. Then $\text{tp}(b/A) = \text{tp}(b'/A)$ if and only if there exists $f \in \text{Aut}_A(\mathcal{M})$ with $f(b) = b'$. Here $\text{Aut}_A(\mathcal{M})$ denotes the group of automorphisms f of \mathcal{M} such that $f \upharpoonright A = \text{id}$.*
- (c) *Suppose $X \subseteq M^n$ is a definable set and $A \subset M$ with $|A| < \kappa$. Then X is A -definable if and only if whenever $f \in \text{Aut}_A(\mathcal{M})$, $f(X) = X$.*

PROOF. For part (a) we just iterate Lemma 7.33. Enumerate $N = \{a_\alpha : \alpha < \kappa\}$ and for each $\alpha < \kappa$ let $A_\alpha := \{a_\beta : \beta < \alpha\}$. Note that each $|A_\alpha| < \kappa$. We build a chain of p.e.m.s, $(f_\alpha : A_\alpha \rightarrow M : \alpha < \kappa)$ recursively as follows:

- f_0 is the empty map, which is a p.e.m. as $\mathcal{N} \equiv \mathcal{M}$.
- $f_{\alpha+1}$ is the p.e.m. which extends f_α by sending a_α to a realisation of $f_\alpha(\text{tp}(a_\alpha/A_\alpha))$. It is by Lemmas 7.32 and 7.33, as well as the κ -saturation of \mathcal{M} , that this is possible.
- At limits we take unions.

Now set $f := \bigcup_{\alpha < \kappa} f_\alpha$. Then f is a p.e.m. with domain N , hence it is an elementary embedding of \mathcal{N} into \mathcal{M} .

Part (b) right-to-left is an easy exercise. For the left-to-right direction assume $\text{tp}(b/A) = \text{tp}(b'/A)$. It follows that the map $g : A \cup \{b_1, \dots, b_n\} \rightarrow M$ given by $g \upharpoonright A = \text{id}$ and $g(b_i) = b'_i$ for $i = 1, \dots, n$, is a p.e.m. Hence by strong κ -homogeneity, g extends to an automorphism $f \in \text{Aut}_A(\mathcal{M})$ with $f(b) = b'$.

Finally, let's consider part (c). This time the left-to-right direction is an easy exercise, and we prove the converse. Let X be defined by $\phi(x, b)$ where $x = (x_1, \dots, x_n)$ and $b = (b_1, \dots, b_m)$. Consider the following set of $L_{A \cup \{b_1, \dots, b_m\}}$ -formulas:

$$\Phi(x, y) := \{\psi(x) \leftrightarrow \psi(y) : \psi \text{ an } L_A \text{- formula}\} \cup \{\phi(x, b), \neg\phi(y, b)\}$$

where $y = (y_1, \dots, y_n)$ is another n -tuple of variables. We claim that Φ is *not* a $2n$ -type in \mathcal{M} . Indeed, if it were, then by κ -saturation it would be realised and we would have a pair of n -tuples d and d' which have the same trype over A but such that $d \in X$ while $d' \notin X$. But by part (a) it would follow that $f(d) = d'$ for some $f \in \text{Aut}_A(\mathcal{M})$, and hence $f(X) \neq X$. This contradiction proves that Φ cannot be a type in \mathcal{M} . Hence it is not finitely satisfiable in \mathcal{M} . That is, there exist L_A -formulas ψ_1, \dots, ψ_ℓ such that

$$\mathcal{M} \models \forall x \forall y \left(\bigwedge_{i=1}^{\ell} (\psi_i(x) \leftrightarrow \psi_i(y)) \longrightarrow (\phi(x, b) \leftrightarrow \phi(y, b)) \right)$$

That is, if d and d' satisfy the same ψ_i s then they are either both in X or both not in X . So, if for each $\tau : \{1, \dots, \ell\} \rightarrow \{0, 1\}$ we define

$$\theta_\tau(x) := \bigwedge_{\tau(i)=1} \psi_i \wedge \bigwedge_{\tau(j)=0} \neg\psi_j$$

then the θ_τ s form a disjoint cover of M^n that partitions X . Hence X is a union of sets defined by $\theta_\tau(x)$, for various τ . That is, there are τ_1, \dots, τ_N such that X is defined by $\bigvee_{i=1}^N \theta_{\tau_i}(x)$. As all the θ_{τ_i} are L_A -formulas, X must be A -definable. \square

As a kind of illustration of how the above theorem can be used, but also as an important pair of notions in their own right, we discuss algebraic and definable closure.

DEFINITION 7.40 (Algebraic/definable closure). Suppose \mathcal{M} is an L -structure, $b \in M$, and $A \subseteq M$. We say that b is *algebraic over* A , denoted $b \in \text{acl}_{\mathcal{M}}(A)$, if b is contained in a finite A -definable set. That is, if there is an L_A -formula $\phi(x)$ such that $\mathcal{M} \models \phi(b)$ and ϕ has only finitely many realisations in \mathcal{M} . We also call $\text{acl}_{\mathcal{M}}(A)$ the *algebraic closure* of A in \mathcal{M} .

Similarly, b is *definable over* A , denoted $b \in \text{dcl}_{\mathcal{M}}(A)$, if $\{b\}$ is an A -definable set. That is, if there is an L_A -formula $\phi(x)$ such that $\mathcal{M} \models \phi(b)$ and b is the only realisation of ϕ in \mathcal{M} . We also call $\text{dcl}_{\mathcal{M}}(A)$ the *definable closure* of A in \mathcal{M} .

EXERCISE 7.41. Suppose $\mathcal{M} \preceq \mathcal{N}$, $A \subseteq M$, and $b \in M$. Then $b \in \text{acl}_{\mathcal{M}}(A)$ if and only if $b \in \text{acl}_{\mathcal{N}}(A)$. Similarly for definable closure.

Often when the ambient structure is understood we write $\text{acl}(A)$ and $\text{dcl}(A)$ rather than $\text{acl}_{\mathcal{M}}(A)$ and $\text{dcl}_{\mathcal{M}}(A)$.

PROPOSITION 7.42. *Suppose \mathcal{M} is a κ -saturated and strongly κ -homogeneous L -structure, $A \subset M$ with $|A| < \kappa$, and $b \in M$.*

- (a) *The following are equivalent*
 - (i) $b \in \text{acl}(A)$
 - (ii) b is one of finitely many realisations of $\text{tp}(b/A)$ in \mathcal{M}
 - (iii) the orbit of b under $\text{Aut}_A(\mathcal{M})$, $\{f(b) : f \in \text{Aut}_A(\mathcal{M})\}$, is finite.
- (b) *The following are equivalent*
 - (i) $b \in \text{dcl}(A)$
 - (ii) b is the only realisation of $\text{tp}(b/A)$ in \mathcal{M}
 - (iii) $f(b) = b$ for all $f \in \text{Aut}_A(\mathcal{M})$.

PROOF. We prove part (a) and leave part (b) as an exercise. Suppose $b \in \text{acl}(A)$. Then there is a finite A -definable set X with $b \in X$. The formula defining X must be in $\text{tp}(b/A)$, and hence the realisations of this type form a subset of X , which is therefore finite. Now assume $\text{tp}(b/A)$ has only finitely many realisations in \mathcal{M} . Theorem 7.39 (b) says that $\text{Aut}_A(\mathcal{M})$ acts transitively on the set of realisations of $\text{tp}(b/A)$ – since $|A| < \kappa$. Since this set contains b , it must be the orbit of b , which is therefore finite. Finally, suppose $X := \{f(b) : f \in \text{Aut}_A(\mathcal{M})\}$ is finite. As X is finite it is a definable set. Moreover, it is clear that X is preserved setwise by every automorphism of \mathcal{M} that fixes A pointwise. Hence by Theorem 7.39 (c), X is A -definable. As it contains b and is finite, $b \in \text{acl}(A)$ as desired. \square

EXERCISE 7.43. Suppose $K \models \text{ACF}$ and $A \subseteq K$. Then $\text{acl}(A) = \mathbb{F}(A)^{\text{alg}}$. That is, the model-theoretic algebraic closure of a set is equal to the field-theoretic algebraic closure of the field generated by that set. On the other hand

$$\text{dcl}(A) = \begin{cases} \mathbb{Q}(A) & \text{if } \text{char}(K) = 0 \\ \mathbb{F}_p(A)^{\text{per}} & \text{if } \text{char}(K) = p > 0 \end{cases}$$

where $k^{\text{per}} := \{a \in k^{\text{alg}} : a^{p^n} \in k, \text{ some } n\}$ is the perfect closure of a field k in positive characteristic. (*Hint: Pass to a $|A|$ -saturated and strongly $|A|$ -homogeneous elementary extension, and then use Proposition 7.42.*)

CHAPTER 8

Totally Transcendental Theories

8.1. Morley Rank

We wish to define a certain ordinal- (or ∞ -) valued rank for formulas. The main idea is rather simple: we will want a definable set to have rank $\geq \alpha + 1$ if it contains infinitely many pairwise disjoint definable subsets of rank $\geq \alpha$. But a few subtleties appear, mostly to do with the fact that we don't want the computation of rank to change as we pass to elementary extensions. This forces the following slightly convoluted definition, which becomes clarified when we pass to an \aleph_0 -saturated elementary extension.

DEFINITION 8.1 (Morley rank). Suppose \mathcal{M} is an L -structure and $\phi(x)$ is an $L_{\mathcal{M}}$ -formula, where $x = (x_1, \dots, x_n)$. We defined inductively on ordinals α what it means for ϕ have *Morley rank* $\geq \alpha$, denoted $\text{RM}_{\mathcal{M}}(\phi) \geq \alpha$:

- $\text{RM}_{\mathcal{M}}(\phi) \geq 0$ if $\phi^{\mathcal{M}} \neq \emptyset$
- $\text{RM}_{\mathcal{M}}(\phi) \geq \alpha + 1$ if there exist an elementary extension, $\mathcal{M} \preceq \mathcal{N}$, and $L_{\mathcal{N}}$ -formulas, $\psi_i(x)$ for $i < \omega$, such that: $\text{RM}_{\mathcal{N}}(\psi_i) \geq \alpha$ and $\mathcal{N} \models \forall x(\psi_i(x) \rightarrow \phi(x))$ for all $i < \omega$, and $\mathcal{N} \models \neg \exists x(\psi_i(x) \wedge \psi_j(x))$ for all $i \neq j$.
- for α a limit ordinal, $\text{RM}_{\mathcal{M}}(\phi) \geq \alpha$ if $\text{RM}_{\mathcal{M}}(\phi) \geq \beta$ for all $\beta < \alpha$.

We define $\text{RM}_{\mathcal{M}}(\phi) = \alpha$ if $\text{RM}_{\mathcal{M}}(\phi) \geq \alpha$ but $\text{RM}_{\mathcal{M}}(\phi) \not\geq \alpha + 1$. If, on the other hand, $\text{RM}_{\mathcal{M}}(\phi) \geq \alpha$ for all α , then we write $\text{RM}_{\mathcal{M}}(\phi) = \infty$. Finally, if ϕ has no realisations in \mathcal{M} then we write $\text{RM}_{\mathcal{M}}(\phi) = -1$.

If $X \subseteq M^n$ is a definable set then the *Morley rank of X in \mathcal{M}* , denoted by $\text{RM}_{\mathcal{M}}(X)$, is just the Morley rank of $\phi(x)$ where $X = \phi^{\mathcal{M}}$. (Check that this is well-defined: it does not depend on the choice of formula defining X .)

EXERCISE 8.2. If $\mathcal{M} \preceq \mathcal{N}$ and ϕ is an $L_{\mathcal{M}}$ -formula, then $\text{RM}_{\mathcal{M}}(\phi) = \text{RM}_{\mathcal{N}}(\phi)$. (*Hint: This is a straightforward induction on α , the only complication is showing that if $\text{RM}_{\mathcal{M}}(\phi) \geq \alpha + 1$ then the same holds in \mathcal{N} . For this recall the following amalgamation fact: if \mathcal{N} and \mathcal{N}' are elementary extensions of \mathcal{M} then there exists an elementary extension \mathcal{R} of \mathcal{N} and an elementary embedding of \mathcal{N}' into \mathcal{R} that is the identity on \mathcal{M} . This latter fact can be proved using compactness and Proposition 4.45 and also follows from Exercise 5(c) of Assignment 5, PMath 433, Fall 2008.*)

Partly justified by the above exercise, we will usually write RM instead of $\text{RM}_{\mathcal{M}}$ when this does not cause confusion.

PROPOSITION 8.3. Suppose \mathcal{M} is an L -structure, $\phi(x, y)$ is an L -formula with $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_m)$, and $b, b' \in M^m$. If $\text{tp}(b) = \text{tp}(b')$ then $\text{RM}(\phi(x, b)) = \text{RM}(\phi(x, b'))$.

PROOF. We show by induction on α that if $\text{RM}(\phi(x, b)) \geq \alpha$ then $\text{RM}(\phi(x, b')) \geq \alpha$ – by symmetry this will suffice. For $\alpha = 0$ this follows because if $\mathcal{M} \models \exists x \phi(x, b)$ then $\exists x \phi(x, y) \in \text{tp}(b) = \text{tp}(b')$ and so $\mathcal{M} \models \exists x \phi(x, b')$. At limit stages the result follows easily by the inductive hypothesis.

Hence we are left to consider the case when $\text{RM}(\phi(x, b)) \geq \alpha + 1$. Let $\mathcal{N} \succeq \mathcal{M}$ and L_N -formulas $\psi_0(x, d_0), \psi_1(x, d_1), \dots$ witness this. Here the ψ_i are L -formulas and the d_i are tuples from N . Let $\mathcal{N}' \succeq \mathcal{N}$ be an \aleph_0 -saturated and strongly ω -homogeneous elementary extension. Note that $\text{RM}_{\mathcal{N}'}(\psi_i(x, d_i)) \geq \alpha$ by Exercise 8.2, and still they define pairwise disjoint subsets of $\phi(x, b)^{\mathcal{N}'}$. Also, $\text{tp}_{\mathcal{N}'}(b) = \text{tp}_{\mathcal{N}'}(b')$ as $\mathcal{M} \preceq \mathcal{N}'$. By Theorem 7.39 (b) there exists $f \in \text{Aut}(\mathcal{N}')$ with $f(b) = b'$. By the inductive hypothesis, $\text{RM}_{\mathcal{N}'}(\psi_i(x, f(d_i))) \geq \alpha$ as $\text{tp}(d_i) = \text{tp}(f(d_i))$, for each $i < \omega$. Moreover, the $\psi_i(x, f(d_i))$ define pairwise disjoint subsets of $\phi(x, f(b))^{\mathcal{N}'} = \phi(x, b')^{\mathcal{N}'}$. Hence, \mathcal{N}' together with $\psi_i(x, f(d_i))$ for $i < \omega$ witness that $\text{RM}(\phi(x, b')) \geq \alpha + 1$, as desired. \square

COROLLARY 8.4. *Suppose \mathcal{M} is \aleph_0 -saturated and $\phi(x)$ is an L_M -formula. Then for every ordinal α , $\text{RM}(\phi) \geq \alpha + 1$ if and only if there are in \mathcal{M} infinitely many pairwise disjoint definable subsets of $\phi^{\mathcal{M}}$ of Morley rank $\geq \alpha$.*

PROOF. The right-to-left direction is easy and does not use ω -saturation. So we assume that $\text{RM}(\phi) \geq \alpha + 1$. Let $\mathcal{N} \succeq \mathcal{M}$ and L_N -formulas $\psi_0(x, d_0), \psi_1(x, d_1), \dots$ witness this. Here the ψ_i are L -formulas and the d_i are tuples from N . Let $A \subseteq M$ be a finite set containing the parameters of ϕ . Let $\{d'_0, d'_1, \dots\}$ be tuples in \mathcal{M} such that for each ℓ , $\text{tp}(d'_0, \dots, d'_\ell/A) = \text{tp}(d_0, \dots, d_\ell/A)$. This is possible by ω -saturation: we inductively define this sequence as follows. First use ω -saturation of \mathcal{M} to find d'_0 in M realising $\text{tp}(d_0/A)$. Having defined d'_0, \dots, d'_ℓ , let $f : A \cup \{d_0, \dots, d_\ell\} \rightarrow N$ be the partial elementary map that is the identity on A and sends d_i to d'_i . Now use ω -saturation of \mathcal{M} to find $d'_{\ell+1}$ in M realising $f(\text{tp}(d_{\ell+1}/A \cup \{d_0, \dots, d_\ell\}))$. Then $\text{tp}(d_0, \dots, d_{\ell+1}/A) = \text{tp}(d'_0, \dots, d'_{\ell+1}/A)$, as desired.

Now we have

$$\begin{aligned} \text{RM}_{\mathcal{M}}(\psi(x, d'_i)) &= \text{RM}_{\mathcal{N}}(\psi(x, d'_i)) \quad \text{by Exercise 8.2} \\ &= \text{RM}_{\mathcal{N}}(\psi(x, d_i)) \quad \text{by Proposition 8.3 as } \text{tp}_{\mathcal{N}}(d_i) = \text{tp}_{\mathcal{N}}(d'_i) \\ &\geq \alpha \end{aligned}$$

Moreover, as $\forall x(\psi_i(x, z_i) \rightarrow \phi(x)) \in \text{tp}_{\mathcal{N}}(d_i/A) = \text{tp}(d'_i/A)$, we have that

$$\mathcal{M} \models \forall x(\psi_i(x, d'_i) \rightarrow \phi(x)),$$

for all $i < \omega$. Similarly, for $i \neq j$,

$$\mathcal{M} \models \neg \exists x(\psi_i(x, d'_i) \wedge \psi_j(x, d'_j))$$

as $\neg \exists x(\psi_i(x, z_i) \wedge \psi_j(x, z_j)) \in \text{tp}(d_i, d_j) = \text{tp}(d'_i, d'_j)$. So we have found pairwise disjoint definable subsets of $\phi^{\mathcal{M}}$, namely $\psi_i(x, d'_i)^{\mathcal{M}}$ for $i < \omega$, each of Morley rank $\geq \alpha$. \square

Here are some first properties of Morley rank.

LEMMA 8.5. *Suppose \mathcal{M} is an L -structure and $X, Y \subseteq M^n$ are definable sets.*

- (a) *If $X \subseteq Y$ then $\text{RM}(X) \leq \text{RM}(Y)$.*
- (b) *$\text{RM}(X \cup Y) = \max\{\text{RM}(X), \text{RM}(Y)\}$*
- (c) *$\text{RM}(X) = 0$ if and only if X is a nonempty finite set.*

PROOF. By Exercise 8.2 we can compute the Morley ranks in an elementary extension. Hence, without loss of generality, we may assume that \mathcal{M} is \aleph_0 -saturated. Now using the characterisation of Morley rank given by Corollary 8.4, one shows by induction on α that if $X \subseteq Y$ and $\text{RM}(X) \geq \alpha$ then $\text{RM}(Y) \geq \alpha$. Indeed, the case of $\alpha = 0$ is trivial, and the limit stage follows immediately from the inductive hypothesis. At the successor stage one merely observes that any infinite collection of pairwise disjoint definable subsets of X is also an infinite collection of pairwise disjoint definable subsets of Y . This gives (a).

For (b) we show by induction that if $\text{RM}(X \cup Y) \geq \alpha$ then $\max\{\text{RM}(X), \text{RM}(Y)\} \geq \alpha$ – the converse follows immediately from part (a). When $\alpha = 0$ this is clear. Suppose α is a limit ordinal. Then

$$\begin{aligned}
\text{RM}(X \cup Y) \geq \alpha &\implies \text{RM}(X \cup Y) \geq \beta \text{ for all } \beta < \alpha \\
&\implies \max\{\text{RM}(X), \text{RM}(Y)\} \geq \beta \text{ for all } \beta < \alpha, \text{ by IH} \\
&\implies \text{either } \text{RM}(X) \geq \beta \text{ for all } \beta < \alpha \text{ or } \text{RM}(Y) \geq \beta \text{ for all } \beta < \alpha \\
&\implies \text{either } \text{RM}(X) \geq \alpha \text{ or } \text{RM}(Y) \geq \alpha \\
&\implies \max\{\text{RM}(X), \text{RM}(Y)\} \geq \alpha.
\end{aligned}$$

Finally, if $\text{RM}(X \cup Y) \geq \alpha + 1$ then $X \cup Y$ contains an infinite collection of pairwise disjoint definable subsets, say Z_0, Z_1, \dots , with $\text{RM}(Z_i) \geq \alpha$ for each $i < \omega$. By the inductive hypothesis $\max\{\text{RM}(Z_i \cap X), \text{RM}(Z_i \cap Y)\} \geq \alpha$. Hence either infinitely many of the $Z_i \cap X$ – which are all pairwise disjoint definable subsets of X – are of Morley rank $\geq \alpha$, or infinitely many of the $Z_i \cap Y$ are. In the former case we have $\text{RM}(X) \geq \alpha + 1$, and in the latter $\text{RM}(Y) \geq \alpha + 1$. So $\max\{\text{RM}(X), \text{RM}(Y)\} \geq \alpha + 1$, as desired.

For part (c) note that, as we are in an \aleph_0 -saturated structure, $\text{RM}(X) \geq 1$ if and only if it contains an infinite collection of pairwise disjoint definable nonempty sets. That is, $\text{RM}(X) \geq 1$ if and only if it is infinite. Hence, $\text{RM}(X) = 0$ if and only if X is a nonempty finite subset. \square

If a definable set X is of Morley rank α then it does not contain infinitely many pairwise disjoint definable subset of Morley rank α . But can it contain arbitrarily large finite numbers of pairwise disjoint definable subsets of Morley rank α ? The answer is no, and this gives rise to a notion of *degree*.

LEMMA 8.6. *Suppose \mathcal{M} is an L -structure and X is definable set with $\text{RM}(X) = \alpha$, an ordinal. Then there exists a positive integer d such that if Y_1, \dots, Y_ℓ are pairwise disjoint definable subsets of X of Morley rank α , then $\ell \leq d$.*

PROOF. We build a binary tree $S \subseteq 2^{<\omega}$ and definable subsets $(X_\sigma \subseteq X : \sigma \in S)$ of Morley rank α , as follows:

- $\emptyset \in S$ and $X_\emptyset = X$.
- If $\sigma \in S$ then there are two possibilities: either (a) there exists an M -definable set $Y \subseteq X_\sigma$ with $\text{RM}(Y) = \text{RM}(X_\sigma \setminus Y) = \alpha$, in which case we put $(\sigma, 0), (\sigma, 1) \in S$ and set $X_{(\sigma,0)} = Y$ and $X_{(\sigma,1)} = X_\sigma \setminus Y$; or (b), no such Y exists in which case σ is a terminal node of the tree S .

If S is infinite then by König's Lemma S has an infinite branch $f : \omega \rightarrow 2$, and the sets $Z_m := X_{f \upharpoonright m} \setminus X_{f \upharpoonright m+1}$ form an infinite pairwise disjoint collection of definable subsets of X each of Morley rank α . This would contradict $\text{RM}(X) = \alpha$. Hence S must be finite. Let

$\sigma_1, \dots, \sigma_d$ be the terminal nodes of S , and $Z_i := X_{\sigma_i}$, for $i = 1, \dots, d$ the corresponding definable subsets. Then by construction each Z_i has Morley rank α and X is the pairwise disjoint union of Z_1, \dots, Z_d . Moreover, also by construction, as $\sigma_1, \dots, \sigma_d$ were terminal nodes of S , no Z_i can contain two disjoint definable subset of Morley rank α .

Suppose Y_1, \dots, Y_ℓ are pairwise disjoint definable subsets of X of Morley rank α . Then for each $i \leq d$ there is at most one $j \leq \ell$ such that $Z_i \cap Y_j$ has Morley rank α . So if $\ell > d$, then there must be some $j \leq \ell$ such that $Z_i \cap Y_j$ has Morley rank $< \alpha$ for all $i \leq d$. Hence

$Y_j = \bigcup_{i=1}^d (Z_i \cap Y_j)$ must have Morley rank $< \alpha$ by Lemma 8.5, contradicting our assumption that each Y_j is of Morley rank α . So $\ell \leq d$, as desired. \square

PROPOSITION 8.7. *Suppose \mathcal{M} is an L -structure and $\phi(x)$ is an L_M -formula with $\text{RM}(\phi) = \alpha$ an ordinal. Then there exists a positive integer d such that if $\mathcal{N} \succeq \mathcal{M}$ and $\psi_1(x), \dots, \psi_\ell(x)$ are L_N -formulas of Morley rank α with $\mathcal{N} \models \forall x(\psi_i(x) \rightarrow \phi(x))$ for all $i = 1, \dots, \ell$, and $\mathcal{N} \models \neg \exists x(\psi_i(x) \wedge \psi_j(x))$ for all $i \neq j$, then $\ell \leq d$.*

PROOF. Let $\mathcal{M}' \succeq \mathcal{M}$ be an \aleph_0 -saturated elementary extension of \mathcal{M} . Let d be a positive integer satisfying the conclusion of Lemma 8.6 applied to \mathcal{M}' . We will prove that this d works. Now suppose $\mathcal{N} \succeq \mathcal{M}$ and $\phi^{\mathcal{N}}$ contains ℓ -many pairwise disjoint definable sets, given by L_N -formulas $\psi_1(x, b_1), \dots, \psi_\ell(x, b_\ell)$, of rank α in \mathcal{N} . Let \mathcal{R} be a common elementary extension of \mathcal{N} and \mathcal{M}' over \mathcal{M} (see the hint for Exercise 8.2). Then also in \mathcal{R} , $\psi_1(x, b_1), \dots, \psi_\ell(x, b_\ell)$ define pairwise disjoint subsets of $\phi^{\mathcal{R}}$ of Morley rank α . By ω -saturation, let (b'_1, \dots, b'_ℓ) in \mathcal{M}' realise $\text{tp}(b_1, \dots, b_\ell/A)$, where $A \subset M$ is a finite set containing the parameters for ϕ . Then in \mathcal{M}' , $\psi_1(x, b'_1), \dots, \psi_\ell(x, b'_\ell)$ define pairwise disjoint subsets of $\phi^{\mathcal{M}'}$. Moreover, by Proposition 8.3, these are each of Morley rank α (in \mathcal{R} and hence in \mathcal{M}'). Hence $\ell \leq d$, as desired. \square

DEFINITION 8.8 (Morley degree). We call the least positive integer d satisfying Proposition 8.7 the *Morley degree of ϕ in \mathcal{M}* , and denoted it by $\text{dM}_{\mathcal{M}}(\phi)$. Again, if X is a definable set then by $\text{dM}_{\mathcal{M}}(X)$ we mean the Morley degree of a formula defining X (and this is independent of which formula defining X we choose).

PROPOSITION 8.9. *Suppose \mathcal{M} is an L -structure and $\phi(x)$ is an L_M -formula with $\text{RM}(\phi) = \alpha$ an ordinal.*

- (a) *If $\mathcal{M} \preceq \mathcal{N}$ then $\text{dM}_{\mathcal{M}}(\phi) = \text{dM}_{\mathcal{N}}(\phi)$.*
- (b) *Suppose \mathcal{M} is \aleph_0 -saturated. Then $\text{dM}(\phi)$ is the greatest positive integer d such that $\phi^{\mathcal{M}}$ contains d -many pairwise disjoint definable subsets of Morley rank α .*
- (c) *Suppose \mathcal{M} is \aleph_0 -saturated. Then $\text{dM}(\phi)$ is the greatest positive integer d such that $\phi^{\mathcal{M}}$ is a pairwise disjoint union of d -many definable subsets of Morley rank α .*
- (d) *Suppose $\phi(x) = \psi(x, b)$ where ψ is an L -formula and b is a tuple from M . If $\text{tp}(b') = \text{tp}(b)$ then $\text{dM}(\psi(x, b)) = \text{dM}(\psi(x, b'))$.*

PROOF. Toward a proof of (a), let \mathcal{N}' be an \aleph_0 -saturated elementary extension of \mathcal{N} . Let $d > 0$ be the least positive integer satisfying Lemma 8.6 applied to $\phi^{\mathcal{N}'}$. Since \mathcal{N}' is \aleph_0 -saturated the proof of Proposition 8.7 shows that $\text{dM}_{\mathcal{N}'}(\phi) \leq d$. On the other hand, by the minimal choice of d , $\phi^{\mathcal{N}'}$ contains d -many pairwise disjoint definable subsets of Morley

rank α . As $\mathcal{N}' \succeq \mathcal{N}$, $d \leq \text{dM}_{\mathcal{N}}(\phi)$ by definition. So $d = \text{dM}_{\mathcal{N}}(\phi)$. But $\mathcal{N}' \succeq \mathcal{M}$ also, so the same argument shows that $d = \text{dM}_{\mathcal{M}}(\phi)$. Hence $\text{dM}_{\mathcal{M}}(\phi) = \text{dM}_{\mathcal{N}}(\phi)$.

To prove part (b), we need only prove, assuming that \mathcal{M} is \aleph_0 -saturated, that $\phi^{\mathcal{M}}$ does contain $\text{dM}(\phi)$ -many pairwise disjoint definable subsets of Morley rank α . By definition, \mathcal{M} has an elementary extension \mathcal{N} where this is true, witnessed say by $L_{\mathcal{N}}$ -formulas $\psi_1(x, b_1), \dots, \psi_d(x, b_d)$, where $d = \text{dM}(\phi)$. Let (b'_1, \dots, b'_d) in M realise $\text{tp}(b_1, \dots, b_d/A)$, where $A \subset M$ is a finite set containing the parameters of ϕ – possible as \mathcal{M} is \aleph_0 -saturated. Then the L_M -formulas $\psi_1(x, b'_1), \dots, \psi_d(x, b'_d)$ define pairwise disjoint definable subsets of $\phi^{\mathcal{M}}$ of Morley rank α , as desired.

Now (c) follows from (b) once we note that if Y_1, \dots, Y_d are pairwise disjoint definable subsets of $\phi^{\mathcal{M}}$, where $d = \text{dM}(\phi)$ of Morley rank α , then $Z := \phi^{\mathcal{M}} \setminus (\bigcup_{i=1}^d Y_i)$ must have Morley rank $< \alpha$. Hence $(Y_1 \cup Z)$ has Morley rank α , and so $\{(Y_1 \cup Z), Y_2, \dots, Y_d\}$ is a collection of d -many pairwise disjoint definable subsets of $\phi^{\mathcal{M}}$ of Morley rank α , whose union is $\phi^{\mathcal{M}}$.

Finally we prove (d). Let \mathcal{N} be an \aleph_0 -saturated and strongly ω -homogeneous elementary extension of \mathcal{M} . Then $\text{tp}(b') = \text{tp}(b)$ is witnessed by an automorphism f of \mathcal{N} such that $f(b) = b'$. It is easy to see, from say part (c), that Morley degree is an automorphism invariant. That is $\text{dM}_{\mathcal{N}}(\psi(x, b')) = \text{dM}_{\mathcal{N}}(\psi(x, b))$. Hence, by part (a), $\text{dM}_{\mathcal{M}}(\psi(x, b')) = \text{dM}_{\mathcal{M}}(\psi(x, b))$. \square

EXERCISE 8.10. Suppose X and Y are disjoint definable sets with ordinal-valued Morley rank. Then

$$\text{dM}(X \cup Y) = \begin{cases} \text{dM}(X) + \text{dM}(Y) & \text{if } \text{RM}(X) = \text{RM}(Y) \\ \text{dM}(X) & \text{if } \text{RM}(X) > \text{RM}(Y) \\ \text{dM}(Y) & \text{if } \text{RM}(Y) > \text{RM}(X) \end{cases}$$

We end this section by extending Morley rank and degrees to types.

DEFINITION 8.11 (Morley rank/degree for types). Suppose \mathcal{M} is an L -structure, $A \subseteq M$, and $p(x) \in S_n(A)$. Then the *Morley rank of p in \mathcal{M}* is

$$\text{RM}_{\mathcal{M}}(p) := \inf\{\text{RM}_{\mathcal{M}}(\phi) : \phi \in p\}.$$

If $\text{RM}_{\mathcal{M}}(p) = \alpha$ an ordinal, then we define *the Morley degree of p in \mathcal{M}* to be

$$\text{dM}_{\mathcal{M}}(p) := \inf\{\text{dM}_{\mathcal{M}}(\phi) : \phi \in p \text{ with } \text{RM}_{\mathcal{M}}(\phi) = \alpha\}.$$

If $p = \text{tp}(a/A)$ then we define $\text{RM}_{\mathcal{M}}(a/A) := \text{RM}_{\mathcal{M}}(p)$ and $\text{dM}_{\mathcal{M}}(a/A) := \text{dM}_{\mathcal{M}}(p)$.

REMARK 8.12. The Morley rank and degree of a complete type is preserved when one passes to an elementary extension. When there is no risk of confusion we simply write $\text{RM}(p)$ and $\text{dM}(p)$ rather than $\text{RM}_{\mathcal{M}}(p)$ and $\text{dM}_{\mathcal{M}}(p)$.

For ease of notation, if ϕ has ordinal-valued Morley rank, we will often write $(\text{RM}, \text{dM})(\phi)$ for the pair $(\text{RM}(\phi), \text{dM}(\phi))$. Also, we implicitly put the lexicographic order on such pairs; that is, if ψ also has ordinal-valued Morley rank, then $(\text{RM}, \text{dM})(\phi) < (\text{RM}, \text{dM})(\psi)$ means that $\text{RM}(\phi) < \text{RM}(\psi)$ or $\text{RM}(\phi) = \text{RM}(\psi)$ and $\text{dM}(\phi) < \text{dM}(\psi)$.

The following technical lemma will be very useful.

LEMMA 8.13. *Suppose \mathcal{M} is an L -structure, $A \subseteq M$, and $p \in S_n(A)$ with $\text{RM}(p) = \alpha$ an ordinal and $\text{dM}(p) = d$. Let $\phi \in p$ be such that $(\text{RM}, \text{dM})(\phi) = (\alpha, d)$. Then for any L_A -formula $\psi(x)$, $\psi \in p$ if and only if $\text{RM}(\phi \wedge \psi) = \alpha$. In particular, ϕ isolates p from among the types in $S_n(A)$ of Morley rank $\geq \alpha$. That is, if $q \in S_n(A)$ with $\text{RM}(q) \geq \alpha$ and $\phi \in q$, then $p = q$.*

PROOF. For left-to-right notice that $\psi \in p$ implies $\phi \wedge \psi \in p$ and so $\text{RM}(\phi \wedge \psi) \geq \text{RM}(p) = \alpha$, while clearly $\text{RM}(\phi \wedge \psi) \leq \text{RM}(\phi) = \alpha$. To prove the converse let us assume that $\text{RM}(\phi \wedge \psi) = \alpha$ but $\psi \notin p$, and seek a contradiction. Then $\neg\psi \in p$ and hence $\text{RM}(\phi \wedge \neg\psi) = \alpha$ also. So

$$\begin{aligned} \text{dM}(\phi) &= \text{dM}((\phi \wedge \psi) \vee (\phi \wedge \neg\psi)) \\ &= \text{dM}(\phi \wedge \psi) + \text{dM}(\phi \wedge \neg\psi) \quad \text{by Exercise 8.10} \\ &\geq d + 1 \end{aligned}$$

which is a contradiction.

For the ‘‘in particular’’ clause, suppose $\text{RM}(q) \geq \alpha$ and $\phi \in q$. So if $\psi \in q$, then $\phi \wedge \psi \in q$, and hence $\text{RM}(\phi \wedge \psi) = \alpha$. By the main part of the lemma, we get that $\psi \in p$. So $q \subseteq p$, and as they are both complete, $q = p$ as desired. \square

PROPOSITION 8.14. *Suppose \mathcal{M} is \aleph_0 -saturated, $A \subseteq M$, $x = (x_1, \dots, x_n)$ and $\phi(x)$ is an L_A -formula.*

- (a) $\text{RM}(\phi) = \sup\{\text{RM}(p) : p \in S_n(A), \phi \in p\}$, and this supremum is attained.
- (b) If ϕ has ordinal-valued Morley rank then

$$\text{dM}(\phi) := |\{p(x) \in S_n(M) : \text{RM}(p) = \text{RM}(\phi), \phi \in p\}|.$$

PROOF. By definition, if $\phi \in p \in S_n(A)$, then $\text{RM}(\phi) \geq \text{RM}(p)$. Next we find $p \in S_n(A)$ such that $\text{RM}(p) = \text{RM}(\phi)$. Consider

$$\Phi(x) := \{\phi(x)\} \cup \{\neg\psi(x) : \psi \text{ an } L_A\text{-formula with } \text{RM}(\phi \wedge \psi) < \text{RM}(\phi)\}$$

This collection of L_A -formulas is finitely satisfiable in \mathcal{M} since $\phi^{\mathcal{M}}$ cannot be a finite union of definable subsets of strictly lesser Morley rank. Hence $\Phi(x)$ is an n -type, and can therefore be extended to a complete n -type $p \in S_n(A)$. Now $\text{RM}(p) \leq \text{RM}(\phi)$ as $\phi \in p$. If $\psi \in p$ then $\phi \wedge \psi \in p$ and by definition $\text{RM}(\psi \wedge \phi) \geq \text{RM}(\phi)$. So $\text{RM}(p) \geq \text{RM}(\phi)$, as desired. Notice that we did not use ω -saturation for part (a).

For part (b), let $X := \phi^{\mathcal{M}}$, $(\text{RM}, \text{dM})(X) = (\alpha, d)$, and write $X = Y_1 \cup \dots \cup Y_d$ where Y_1, \dots, Y_d are disjoint definable subsets of X of Morley rank α . Applying part (a) to each Y_i (as M -definable sets), we get d -many distinct complete n -types over M extending ϕ , each of Morley rank α . On the other hand, suppose we had $(d + 1)$ -many such types, p_1, \dots, p_{d+1} . For each $i = 1, \dots, d + 1$, let $\psi_i \in p_i$ be such that $(\text{RM}, \text{dM})(\psi_i) = (\text{RM}, \text{dM})(p_i)$. Taking conjunctions with ϕ if necessary, we may assume that each $\psi_i^{\mathcal{M}} \subseteq X$. By Lemma 8.13, as the types are distinct, $\psi_i \notin p_j$ for any $i \neq j$. Hence $\theta_j := \psi_j \wedge \bigwedge_{i \neq j} \neg\psi_i \in p_j$, and so $\text{RM}(\theta_j) = \alpha$ also. But $\{\theta_1^{\mathcal{M}}, \dots, \theta_{d+1}^{\mathcal{M}}\}$ are pairwise disjoint. This contradicts $\text{dM}(X) = d$. \square

REMARK 8.15. Suppose \mathcal{M} is κ -saturated, $A \subseteq M$ is of cardinality $< \kappa$, and X is an A definable set. Then Proposition 8.14(a) says that $\text{RM}(X) = \sup\{\text{RM}(a/A) : a \in X\}$, and that this supremum is attained.

We end with the following useful lemma about Morley rank.

LEMMA 8.16. *Suppose $\mathcal{M} \models T$, $A \subseteq M$, $a \in M^n$, and $b \in M$. If $b \in \text{acl}(Aa)$ then $\text{RM}(ab/A) = \text{RM}(a/A)$.*

PROOF. We may assume that \mathcal{M} is \aleph_0 -saturated. I leave it as an exercise for you to prove that $\text{RM}(ab/A) \geq \text{RM}(a/A)$. In particular, if $\text{RM}(a/A) = \infty$ then so is $\text{RM}(ab/A)$. So for the converse, it suffices to prove by induction on α that if $\text{RM}(a/A) = \alpha$ and $b \in \text{acl}(Aa)$ then $\text{RM}(ab/A) \leq \alpha$. The case of $\alpha = 0$ is clear.

Assume $\text{RM}(a/A) = \alpha > 0$ and let $x = (x_1, \dots, x_n)$ and let $\phi(x, y) \in \text{tp}(ab/A)$ with $\text{RM}(\phi) = \text{RM}(ab/A)$. There is an L_A formula in $\text{tp}(a/A)$ of Morley rank α . This formula is of course also in $\text{tp}(ab/A)$. So by replacing $\phi(x, y)$ by its conjunction with such a formula, we may assume that

$$(1) \quad \text{RM}(\exists y \phi(x, y)) = \alpha.$$

As $b \in \text{acl}(Aa)$ there is an L_A formula $\psi(x, y)$ such that $\psi(a, y)^{\mathcal{M}}$ is a finite set, say of size ℓ . Then the formula $\psi(x, y) \wedge (\exists^{\leq \ell} y \psi(x, y))$ is in $\text{tp}(ab/A)$, where $\exists^{\leq \ell} y$ is shorthand for the formula which says that there are at most ℓ -many y 's with $\psi(x, y)$. Replacing $\phi(x, y)$ by its conjunction with the above formula we may assume that

$$(2) \quad \text{for all } a' \text{ in any elementary extension, } \phi(a', y) \text{ defines a finite set.}$$

We now prove that $\text{RM}(\phi) = \alpha$. Suppose, toward a contradiction, that $\text{RM}(\phi) \geq \alpha + 1$. Then there exists an infinite collection of pairwise disjoint definable subsets of $\phi^{\mathcal{M}}$ each of Morley rank $\geq \alpha$. By Proposition 8.14(a) we can find a complete type over M in each of them, of Morley rank $\geq \alpha$. So we have an infinite set of types $P \subseteq S_{n+1}(M)$ such that for each $p(x, y) \in P$, $\phi \in p$ and $\text{RM}(p) \geq \alpha$. Given $p(x, y) \in P$, let $\hat{p}(x) \in S_n(M)$ be the restriction of p to the variables x , and let $R := \{\hat{p} : p \in P\} \subseteq S_n(M)$. Facts (1) and (2) ensure that for all (a', b') realising $p(x, y)$ in an elementary extension, $\text{RM}(a'/M) \leq \alpha$ and $b' \in \text{acl}(Ma')$. Hence, by the induction hypothesis, $\text{RM}(r) = \alpha$ for all $r \in R$. Fact (2) also implies that for any fixed $r \in S_n(M)$, $r(x) \cup \{\phi(x, y)\}$ has only finitely many completions (exercise). It follows that each element of R is contained in only finitely many elements of P . As P is infinite, R must be infinite. We have found infinitely many complete types over M extending $\exists y \phi(x, y)$, of the same rank as $\exists y \phi(x, y)$. This contradicts Proposition 8.14(b) (which we can use because $\exists y \phi(x, y)$, being in $\text{tp}(a/A)$, does have ordinal-valued Morley rank). \square

8.2. Total transcendental and ω -stability

We now address the question of when Morley rank is ordinal valued.

DEFINITION 8.17 (Totally transcendental). A consistent theory T is *totally transcendental* (respectively *of finite Morley rank*) if every definable set in every model of T has bounded (respectively finite) Morley rank.

DEFINITION 8.18 (κ -Stable). Suppose L is a countable language and κ is an infinite cardinal. A consistent L -theory T is κ -*stable* if for every $\mathcal{M} \models T$ and all $A \subseteq M$ of cardinality at most κ , $|S_n^{\mathcal{M}}(A)| \leq \kappa$. We say ω -*stable* instead of \aleph_0 -stable.

THEOREM 8.19. *Suppose L is a countable language and T is a consistent L -theory. Then the following are equivalent:*

- (i) T is ω -stable.
- (ii) T is totally transcendental.
- (iii) T is κ -stable for all infinite κ .

PROOF. We prove the contrapositive of (i) implies (ii). Suppose that there exist $\mathcal{M} \models T$ and an L_M -formula $\phi(x)$, where $x = (x_1, \dots, x_n)$, such that $\text{RM}(\phi) = \infty$. Passing to an elementary extension we may assume that \mathcal{M} is \aleph_0 -saturated. We construct an infinite binary tree of L_M -formulas $(\phi_\sigma(x) : \sigma \in 2^{<\omega})$ such that

- $\phi_\emptyset = \phi$
- For each $\sigma \in 2^{<\omega}$, $\phi_{\sigma,0} = \phi_\sigma \wedge \psi_\sigma$ and $\phi_{\sigma,1} = \phi_\sigma \wedge \neg\psi_\sigma$, for some L_M -formula ψ_σ .
- For each $\sigma \in 2^\omega$, $\text{RM}(\phi_\sigma) = \infty$.

In order to construct such a tree we need only observe that if X is a definable set of Morley rank ∞ then it contains a definable subset that is of Morley rank ∞ and whose complement in X is also of Morley rank ∞ . Indeed, as there are only set-many L_M -formulas, there is an ordinal γ such that a definable set is of Morley rank ∞ if and only if it is of Morley rank $\geq \gamma$. Now X , being of Morley rank $\geq \gamma + 1$, contains infinitely many pairwise disjoint definable subsets of Morley rank $\geq \gamma$. Letting Y be one of these we have that $\text{RM}(Y) \geq \gamma$ and $\text{RM}(X \setminus Y) \geq \gamma$, and hence both Y and its complement in X are of Morley rank ∞ .

Having constructed the binary tree, let $A \subset M$ be a countable set containing all the parameters appearing in the (countably many) formulas of the tree. For each $f : \omega \rightarrow 2$, let

$$\Phi_f := \{\phi_{f \upharpoonright \ell} : \ell < \omega\}.$$

Note that Φ_f is an n -type in \mathcal{M} over A : since the formulas become stronger it suffices to observe that every formula in Φ_f is realisable, which follows from the fact that they all have Morley rank ≥ 0 (indeed, they all have Morley rank ∞). Let $p_f \in S_n(A)$ be a complete n -type extending Φ_f . By construction it should be clear that if $f \neq g$ then $p_f \neq p_g$: indeed, if $f(\ell) \neq g(\ell)$ then $\phi_{f \upharpoonright \ell+1}$ and $\phi_{g \upharpoonright \ell+1}$ are inconsistent. We have thus constructed 2^{\aleph_0} -many distinct n -types over the countable set A . So T is not ω -stable. We have shown that (i) implies (ii).

Next we show that (ii) implies (iii). Assume T is totally transcendental, κ is an infinite cardinal, $\mathcal{M} \models T$, and $A \subseteq M$ with $|A| \leq \kappa$. For each $p(x) \in S_n(A)$, let $\phi_p(x) \in p(x)$ be such that $(\text{RM}, \text{dM})(\phi_p) = (\text{RM}, \text{dM})(p)$. By Lemma 8.13, if $\phi_p = \phi_q$ then $p = q$. So we have an injective map from $S_n(A)$ to the set of L_A -formulas in free variables x . As there are at most κ -many such formulas (L is countable and κ is infinite), $|S_n(A)| \leq \kappa$. Hence, T is κ -stable.

That (iii) implies (ii) is clear. \square

EXERCISE 8.20. Show that a theory T is totally transcendental if and only if the formula $(x = x)$ has bounded Morley rank.

So, for example, every definable set in a model of ACF (which is ω -stable) has ordinal-valued Morley rank, while this is not the case for DLO (which is not ω -stable).

The following is useful for showing that a given theory is ω -stable, and is a good example of using automorphism arguments.

PROPOSITION 8.21. *A consistent theory T in a countable language is ω -stable if and only if for every $\mathcal{M} \models T$ and all $A \subseteq M$ countable, $S_1^{\mathcal{M}}(A)$ is countable.*

PROOF. Left-to-right is clear. For the converse we show by induction on n that $S_n(A)$ is countable for all countable subsets A of models of T . The case of $n = 1$ is given. Suppose, toward a contradiction, that $n > 1$, $A \subseteq M$ is countable, but $|S_n(A)| > \aleph_0$. Passing to an elementary extension if necessary, we may assume that \mathcal{M} is \aleph_1 -saturated and strongly \aleph_1 -homogeneous. Now, by the induction hypothesis, $|S_{n-1}(A)| \leq \aleph_0$. Hence, there must be uncountably many complete n -types over A whose restrictions to their first $n - 1$ variables agree. That is, we have $q(x_1, \dots, x_{n-1}) \in S_{n-1}(A)$ and $\{p_\alpha(x_1, \dots, x_n) \in S_n(A) : \alpha < \aleph_1\}$ such that $p_\alpha \neq p_\beta$ for all $\alpha \neq \beta$, and $\{\phi(x_1, \dots, x_{n-1}) : \phi \in p_\alpha\} = q$ for all α . Let $a = (a_1, \dots, a_{n-1})$ realise q , and for each α , let $b_\alpha = (b_{1,\alpha}, \dots, b_{n,\alpha})$ realise p_α . Then we have, for each $\alpha < \aleph_1$, $f_\alpha \in \text{Aut}_A(\mathcal{M})$ such that $f_\alpha(b_{1,\alpha}, \dots, b_{n-1,\alpha}) = a$. Now, for any α, β , if $\text{tp}(f_\alpha(b_{n,\alpha})/Aa) = \text{tp}(f_\beta(b_{n,\beta})/Aa)$ and this is witness by some $g \in \text{Aut}_{Aa}(\mathcal{M})$ with $g(f_\alpha(b_{n,\alpha})) = f_\beta(b_{n,\beta})$, then $f_\beta^{-1} \circ g \circ f_\alpha$ is an automorphism fixing A pointwise and taking b_α to b_β . But that would imply that they have the same type over A , and so $p_\beta = p_\alpha$. We have therefore shown that $\{\text{tp}(f_\alpha(b_{n,\alpha})/Aa) : \alpha < \aleph_1\}$ is an uncountable subset of $S_1(Aa)$ – which contradicts the assumption that there are only countable many complete 1-types over a countable set. \square

8.3. The strongly minimal case

In this section we consider the case of theories whose universe is of rank and degree one.

DEFINITION 8.22 (Strongly minimal theory). A consistent theory T is *strongly minimal* if for all $\mathcal{M} \models T$, M is infinite and every definable subset of M is either finite or cofinite. A structure is called *strongly minimal* if its theory is strongly minimal.

One way to understand this definition is to notice that a theory is strongly minimal if and only if every definable subset of the universe is definable in the empty language. In any case, we have the following more or less obvious characterisation in terms of Morley rank and degree.

PROPOSITION 8.23. *A consistent theory T is strongly minimal if and only if the universe of every model of T is of Morley rank 1 and Morley degree 1.*

PROOF. Suppose T is strongly minimal and $\mathcal{M} \models T$. Passing to an elementary extension if necessary we may assume that \mathcal{M} is \aleph_0 -saturated. Then, as M is infinite, it is of rank ≥ 1 . Since M does not contain even two disjoint infinite definable subsets, it must be of Morley rank 1 and Morley degree 1. For the converse, suppose $\mathcal{M} \models T$, $(\text{RM}, \text{dM})(M) = 1$, and $X \subseteq M$ is a definable subset. If X were infinite and co-infinite then $\{X, (M \setminus X)\}$ would contradict the fact that $\text{dM}(M) = 1$. Hence every definable subset is either finite or cofinite. \square

By Exercise 8.20, strongly minimal theories are ω -stable.

It is important to note here that the definition and the above proposition only refer to definable subsets of the universe itself; we say nothing about definable subsets of cartesian

powers of the universe. However, as we will see, strong minimality has strong consequences regarding the structure of definable subsets of cartesian powers as well. To start with, we get a very concrete characterisation of Morley rank for all definable sets. To describe this characterisation it is useful to first discuss some notions from combinatorial geometry.

DEFINITION 8.24 (Matroid). A *matroid*¹ is a set X together with a *closure* operation $\text{cl} : \mathcal{P}(X) \rightarrow \mathcal{P}(X)$ which satisfies:

- (1) $A \subseteq \text{cl}(A)$ and $\text{cl}(\text{cl}(A)) = \text{cl}(A)$,
- (2) if $A \subseteq B$ then $\text{cl}(A) \subseteq \text{cl}(B)$,
- (3) if $a \in \text{cl}(A)$ then there exists a finite subset $A_0 \subseteq A$ with $a \in \text{cl}(A_0)$, and
- (4) (exchange) if $a \in \text{cl}(A \cup \{b\}) \setminus \text{cl}(A)$ then $b \in \text{cl}(A \cup \{a\})$.

A subset $Y \subseteq X$ of a matroid is said to be *cl-independent* if $a \notin \text{cl}(Y \setminus \{a\})$ for all $a \in Y$. It is a *cl-basis* for X if it is independent and $X = \text{cl}(Y)$.

An example of a matroid is a vector space with the closure operation given by linear span. Indeed, the above definition abstracts exactly those properties of linear span in vector spaces needed to prove the basic properties of linear independence and bases. One thus obtains the following fact, whose proof I leave to you to either find in the literature or prove for yourself.

FACT 8.25. *Suppose (X, cl) is a matroid.*

- (a) *A subset $Y \subseteq X$ is a cl-basis for X if and only if Y is a maximally cl-independent subset of X , if and only if Y is minimal among subsets of X whose closure is X .*
- (b) *Every matroid has a cl-basis, and any two cl-bases have the same cardinality. This cardinality is called the cl-dimension of X .*

We get from this a relative notion of cl-independence and dimension:

DEFINITION 8.26. Suppose X is a matroid and $A \subseteq X$. A subset $Y \subseteq X$ is said to be *cl-independent over A* if $a \notin \text{cl}(A \cup (Y \setminus \{a\}))$ for all $a \in Y$. By the *cl-dimension of Y over A* , denoted $\text{cl-dim}(Y/A)$, we mean the cardinality of any maximally cl-independent subset of Y over A .

Note that $\text{cl-dim}(Y/A)$ is well-defined because of Fact 8.25. Indeed, $\text{cl-dim}(Y/A)$ is just the dimension of the induced matroid on $\text{cl}(Y)$ “localised” at A , that is, where the closure of a set $B \subseteq \text{cl}(Y)$ is taken to be $\text{cl}(A \cup B) \cap \text{cl}(Y)$.

As a matter of notation, if $a_1, \dots, a_n \in X$ then we say that the n -tuple (a_1, \dots, a_n) is *cl-independent over A* to mean that the finite set $\{a_1, \dots, a_n\}$ is cl-independent over A . More generally, by $\text{cl-dim}((a_1, \dots, a_n)/A)$ we actually mean the cl-dimension of the finite set $\{a_1, \dots, a_n\}$ over A . In particular, the cl-dimension of an n -tuple is always $\leq n$.

EXERCISE 8.27. A finite set $\{a_1, \dots, a_n\}$ is cl-independent over A if and only if for all $i = 1, \dots, n$, $a_i \notin \text{cl}(A \cup \{a_1, \dots, a_{i-1}\})$.

What does any of this have to do with strongly minimal theories? First of all, notice that if \mathcal{M} is any structure then (M, acl) satisfies the first three defining properties of being a matroid. The following proposition says that if \mathcal{M} is strongly minimal then the exchange property is also satisfied.

¹Model-theorists often use the term “pregeometry” instead of matroid.

PROPOSITION 8.28. *If \mathcal{M} is strongly minimal then (M, acl) is a matroid.*

PROOF. As we have mentioned, conditions (1) through (3) are easy to check and do not require strong minimality. We need to prove the exchange property. Suppose $A \subseteq M$, $a, b \in M$, and $a \in \text{acl}(Ab) \setminus \text{acl}(A)$. We need to prove that $b \in \text{acl}(Aa)$.

Let $\phi(x, y)$ be an L_A -formula such that $\phi(x, b)$ has only finitely many realisations, one of which is a . For any $b' \in M$, let $X_{b'} := \phi(x, b')^{\mathcal{M}}$. So $a \in X_b$ and X_b is finite, say $|X_b| = \ell$. Let $B := \{b' \in M : |X_{b'}| = \ell\}$. Note that B is an A -definable subset of M containing b . Now consider $B(a) := \{b' \in B : a \in X_{b'}\}$. This is an Aa -definable subset of M containing b . We claim that $B(a)$ is finite, in which case $b \in \text{acl}(Aa)$ and we are done.

Assume toward a contradiction, that $B(a)$ is not finite. Hence it is cofinite, say $|M \setminus B(a)| = s$. Consider $E := \{a' \in M : |M \setminus B(a')| = s\}$. Here by $B(a')$ we mean of course the Aa' -definable set $\{b' \in B : a' \in X_{b'}\}$. Then E is an A -definable set containing a , and hence cannot be finite. Choose distinct $a_1, \dots, a_{\ell+1} \in E$. As each $B(a_i)$ is cofinite (indeed its complement is of size s), $\bigcap_{i=1}^{\ell+1} B(a_i)$ is cofinite, and hence nonempty. Let $\hat{b} \in \bigcap_{i=1}^{\ell+1} B(a_i)$.

Then by definition $X_{\hat{b}}$ contains $a_1, \dots, a_{\ell+1}$. But this contradicts the fact that $\hat{b} \in B$. \square

LEMMA 8.29. *Suppose \mathcal{M} is strongly minimal, $A \subset M$, and $a, b \in M^n$ are acl-independent tuples over A . Then $\text{tp}(a/A) = \text{tp}(b/A)$. This unique type of an acl-independent tuple is called the acl-generic type of M^n over A .*

PROOF. This is proved by induction on n . For $n = 1$, saying that a is acl-independent over A just means that $a \notin \text{acl}(A)$. Let $\phi(x)$ be an L_A -formula. By strong minimality either $\phi^{\mathcal{M}}$ or $(\neg\phi)^{\mathcal{M}}$ is finite. Hence $\phi \in \text{tp}(a/A)$ if and only if $(\neg\phi)^{\mathcal{M}}$ is finite. Similarly, as $b \notin \text{acl}(A)$, $\phi \in \text{tp}(b/A)$ if and only if $(\neg\phi)^{\mathcal{M}}$ is finite. Hence $\text{tp}(a/A) = \text{tp}(b/A)$ as desired.

Now suppose $n > 1$. By the induction hypothesis $\text{tp}(a_1, \dots, a_{n-1}/A) = \text{tp}(b_1, \dots, b_{n-1}/A)$. Passing to an elementary extension if necessary we may assume that \mathcal{M} is $|A|^+$ -saturated and strongly $|A|^+$ -homogeneous. So we have $f \in \text{Aut}_A(\mathcal{M})$ such that $f(a_1, \dots, a_{n-1}) = (b_1, \dots, b_{n-1})$. Since a is acl-independent over A , $a_n \notin \text{acl}(A \cup \{a_1, \dots, a_{n-1}\})$. Applying f we have that $f(a_n) \notin \text{acl}(A \cup \{b_1, \dots, b_{n-1}\})$. As $b_n \notin \text{acl}(A \cup \{b_1, \dots, b_{n-1}\})$ also, $\text{tp}(f(a_n)/Ab_1, \dots, b_{n-1}) = \text{tp}(b_n/Ab_1, \dots, b_{n-1})$. So we have $g \in \text{Aut}_{Ab_1, \dots, b_{n-1}}(\mathcal{M})$ such that $g(f(a_n)) = b_n$. Hence $g(f(a)) = b$ and $g \circ f \in \text{Aut}_A(\mathcal{M})$. This implies that $\text{tp}(a/A) = \text{tp}(b/A)$, as desired. \square

PROPOSITION 8.30. *Suppose \mathcal{M} is strongly minimal, $A \subset M$, and $a \in M^n$. Then a is acl-independent over A if and only if $\text{RM}(a/A) = n$.*

PROOF. As both acl and RM are invariant under elementary extensions, we may assume that \mathcal{M} is $|A|^+$ -saturated and strongly $|A|^+$ -homogeneous.

We prove this by induction on n . For $n = 1$,

$$\begin{aligned}
a \text{ is acl-independent over } A &\iff a \notin \text{acl}(A) \\
&\iff \text{every } A\text{-definable set containing } a \text{ is infinite} \\
&\iff \text{every } A\text{-definable set containing } a \text{ has Morley rank } \geq 1 \\
&\iff \text{RM}(a/A) \geq 1 \\
&\iff \text{RM}(a/A) = 1
\end{aligned}$$

where the last equivalence is because in a strongly minimal theory $\text{RM}(x = x) = 1$.

Now suppose $n > 1$. We first show that $\text{RM}(a/A) = n$ implies that a is acl-independent over A . Indeed, after rearrangement, we may assume that $\{a_1, \dots, a_\ell\}$ is an acl-basis for $\{a_1, \dots, a_n\}$ over A , so that $a_j \in \text{acl}(Aa_1 \dots a_\ell)$ for all $j > \ell$. By Lemma 8.16, $\text{RM}(a/A) = \text{RM}(a_1, \dots, a_\ell/A)$. If a is not acl-independent over A then $\ell < n$, and by induction $\text{RM}(a_1, \dots, a_\ell/A) = \ell$.

Next we assume that a is acl-independent over A and prove that $\text{RM}(a/A) \geq n$. Letting $\phi \in \text{tp}(a/A)$, we show that $\text{RM}(\phi) \geq n$. Fix distinct elements $b_1, b_2, \dots \in M \setminus \text{acl}(A)$, and set $\psi_i := \phi(x) \wedge (x_1 = b_i)$. The ψ_i s define pairwise disjoint definable subsets of ϕ^M . It remains to show therefore that each ψ has Morley rank $\geq n - 1$. Fix i and choose c_2, \dots, c_n with $c_j \notin \text{acl}(Ab_i c_2 \dots c_{j-1})$ for $j = 2, \dots, n$. Then (b_i, c_2, \dots, c_n) is an acl-independent tuple over A , hence has the same type as a over A by Lemma 8.29, and hence realises ϕ . So (b_i, c_2, \dots, c_n) realises ψ_i . Moreover, $\text{RM}(b_i, c_2, \dots, c_n/Ab_i) \geq \text{RM}(c_2, \dots, c_n/Ab_i) = n - 1$ by the induction hypothesis. Hence $\text{RM}(\psi_i) \geq n - 1$, as desired.

Finally we show that $\text{RM}(a/A) \leq n$. In fact we show that $\text{RM}(M^n) \leq n$. If not, there would exist at least two disjoint definable subsets of M^n of Morley rank n . Say these sets are defined over some finite B , and use Proposition 8.14(a) to find distinct complete types $p, q \in S_n(B)$ of Morley rank n . Let c and d realise p and q respectively. As we have already shown that Morley rank n for n -tuples implies acl-independence, c and d are both acl-independent tuples over B . But then c and d must have the same type over B , contradicting $p \neq q$. \square

THEOREM 8.31. *Suppose \mathcal{M} is strongly minimal, $A \subseteq M$, and $a \in M^n$. Then $\text{RM}(a/A) = \text{acl-dim}(a/A)$.*

PROOF. Let $\ell = \text{acl-dim}(a/A)$. Assume, without loss of generality, that $\{a_1, \dots, a_\ell\}$ is an acl-basis for $\{a_1, \dots, a_n\}$ over A . Then by Lemma 8.16, $\text{RM}(a/A) = \text{RM}(a_1, \dots, a_\ell/A)$. As (a_1, \dots, a_ℓ) is acl-independent, $\text{RM}(a_1, \dots, a_\ell/A) = \ell$ by Proposition 8.30. \square

COROLLARY 8.32. *Suppose T is a strongly minimal theory, \mathcal{M} is a κ -saturated and strongly κ -homogeneous model of T , $A \subset M$ with $|A| < \kappa$, and $X \subseteq M^n$ is an A -definable set. Then $\text{RM}(X) = \max\{\text{acl-dim}(a/A) : a \in X\}$. In particular, strongly minimal theories are of finite Morley rank.*

PROOF. The main clause follows from Theorem 8.31 by Remark 8.15. The ‘‘in particular’’ clause follows since the acl-dim of an n -tuple is bounded by n . \square

Let us consider some basic examples. I leave the details as exercises.

Suppose L is the empty language and T is the theory of infinite sets, and $\mathcal{M} \models T$. Then T is strongly minimal. For any set $A \subseteq M$, $\text{acl}(A) = A$. Hence, for any $a = (a_1, \dots, a_n) \in M^n$, $\text{RM}(a/A) = |\{a_i : a_i \notin A, i = 1, \dots, n\}|$. The generic n -type over A is just the type of an n -tuple whose co-ordinates are distinct and not in A .

Suppose F is a field and L is the language of F -vector spaces, and T is the theory of infinite F -vector spaces, and $V \models T$. Then T is strongly minimal. For any set $A \subseteq V$, $\text{acl}(A)$ is the F -linear span of A . Hence, for any $a \in M^n$, $\text{RM}(a/A) = |B' \setminus B|$ where B is a linear basis for $\text{span}(A)$ and B' is an extension of B to a linear basis for $\text{span}(Aa)$. In particular, if $\dim_F \text{span}(A) < \omega$ then $\text{RM}(a/A) = \dim_F \text{span}(Aa) - \dim_F \text{span}(A)$. Here,

\dim_F denotes the linear dimension as an F -vector space. The generic n -type over A is the type of an n -tuple that is linearly independent over $\text{span}(A)$.

Suppose L is the language of rings, and T is the theory of algebraically closed fields, and $K \models T$. Then T is strongly minimal. For any set $A \subseteq V$, $\text{acl}(A) = \mathbb{F}(A)^{\text{alg}}$, the field-theoretic algebraic closure of the field generated by A . Hence, for any $a \in M^n$, $\text{RM}(a/A) = \text{trdeg}(\mathbb{F}(Aa)^{\text{alg}}/\mathbb{F}(A))$. The generic n -type over A is the type of an n -tuple that is algebraically independent over $\mathbb{F}(A)$.

Before we end this section and return to the more general ω -stable case, let us use our characterisation of Morley rank in strongly minimal theories to prove that Morley rank is “a definable property of the parameters”. First a lemma.

LEMMA 8.33. *Suppose \mathcal{M} is strongly minimal, $X \subseteq M^{n+1}$ is A -definable, and for $b \in X^n$ let $X_b := \{a \in M : (a, b) \in X\}$. Set $F := \{b \in M^n : X_b \text{ is nonempty and finite}\}$ and $G := \{X_b \text{ is infinite}\}$. Then F and G are A -definable and $\text{RM}(X) = \max\{\text{RM}(F), \text{RM}(G) + 1\}$.*

PROOF. By strong minimality there exists an $N < \omega$ such that for all $b \in M^n$, either $|X_b| \leq N$ or $|M \setminus X_b| \leq N$. We leave as an exercise the proof that such an N exists. In any case, it follows that $F = \{b : 0 < |X_b| \leq N\}$ and $G = \{b : |X_b| > N\}$ are A -definable.

Let $\pi : M^{n+1} \rightarrow M^n$ be the projection onto the last n co-ordinates. Then $X = (\pi^{-1}(F) \cap X) \cup (\pi^{-1}(G) \cap X)$. Note that if $(a, b) \in \pi^{-1}(F) \cap X$ then $a \in \text{acl}(Ab)$, and hence $\text{acl-dim}((a, b)/A) = \text{acl-dim}(b/A)$. It follows by Corollary 8.32 that $\text{RM}(\pi^{-1}(F) \cap X) = \text{RM}(F)$. Now let us consider $\pi^{-1}(G) \cap X$. By Theorem 8.31, for every $(a, b) \in \pi^{-1}(G) \cap X$, $\text{RM}((a, b)/A) = \text{acl-dim}((a, b)/A) \leq \text{acl-dim}(b/A) + 1 = \text{RM}(b/A) + 1 \leq \text{RM}(G) + 1$. On the other hand, choose $b \in G$ such that $\text{RM}(b/A) = \text{RM}(G)$. As X_b is infinite, by saturation X_b must contain some element $a \notin \text{acl}(Ab)$ (exercise). Hence $(a, b) \in \pi^{-1}(G) \cap X$ and $\text{acl-dim}((a, b)/A) = \text{acl-dim}(b/A) + 1 = \text{RM}(G) + 1$. By Corollary 8.32, $\text{RM}(\pi^{-1}(G) \cap X) = \text{RM}(G) + 1$. \square

PROPOSITION 8.34. *Suppose \mathcal{M} is strongly minimal, $A \subseteq M$, and $X \subseteq M^{m+n}$ is an A -definable set. For each $b \in M^m$ we set $X_b := \{a \in M^n : (a, b) \in X\}$. Then for each $k \leq n$, $B_k := \{b \in M^m : \text{RM}(X_b) \geq k\}$ is A -definable.*

PROOF. We proceed by induction on n . The case of $n = 1$ is done by Lemma 8.33 since Morley rank 0 means finite nonempty and Morley rank 1 means infinite. Suppose $n = r + 1$ and let $Y = \pi(X) \subseteq M^{m+r}$ where π is the projection onto the last $m + r$ co-ordinates. Let

$$F := \{c \in Y : \text{fibre of } X \rightarrow Y \text{ over } c \text{ is nonempty and finite}\}$$

and

$$G := \{c \in Y : \text{fibre of } X \rightarrow Y \text{ over } c \text{ is infinite}\}.$$

By Lemma 8.33 these subsets of M^{m+r} are A -definable. For any $b \in M^m$, let $Y_b := \{a \in M^r : (a, b) \in Y\}$, $F_b := \{a \in M^r : (a, b) \in F\}$, and $G_b := \{a \in M^r : (a, b) \in G\}$. Then by Lemma 8.33 again, $\text{RM}(X_b) = \max\{\text{RM}(F_b), \text{RM}(G_b) + 1\}$. Hence $\text{RM}(X_b) \geq k$ if and only if $\text{RM}(Y_b) \geq k$ or $\text{RM}(G_b) \geq k - 1$. So by the induction hypothesis, the set of such b is A -definable. \square

Maybe it is worth restating the above proposition in terms of formulas. It says that in a strongly minimal theory, if $\phi(x, y)$ is an L_A -formula where $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_m)$, then for every $k = 0, \dots, n$, $\{b \in M^m : \text{RM}(\phi(x, b)) \geq k\}$ is A -definable.

CHAPTER 9

Nonforking in Totally Transcendental Theories

For this chapter we fix a complete totally transcendental theory T and a sufficiently saturated model $\mathcal{U} \models T$. By a “small” set we will mean a set whose cardinality is less than the saturation and strong homogeneity of \mathcal{U} . All tuples and parameter sets mentioned, unless otherwise specified, will be assumed to come from \mathcal{U} and be small. All models of T mentioned, unless otherwise specified, will be assumed to be small elementary substructures of \mathcal{U} (cf. Theorem 7.39(a) for a justification of this). By $\models \sigma$ we will mean $\mathcal{U} \models \sigma$. We say that \mathcal{U} is a *universal domain*.

9.1. First properties

In the strongly minimal case we used the matroid geometry to get a notion of acl-independence which in turn gave rise to a dimension function (on complete types, say), namely acl-dim, that turned out to coincide with Morley rank. In our current setting (U, acl) is no longer a matroid and hence acl-independence is not well-defined. On the other hand, we do have bounded (i.e., ordinal-valued) Morley rank. Working somewhat in the opposite direction we can use this rank to define a notion of independence.

DEFINITION 9.1 (Nonforking and independence). Suppose $A \subseteq B$, $p \in S_n(A)$, $q \in S_n(B)$, and $p \subseteq q$. We say that q is a *nonforking extension* of p , or that q *does not fork over* A , if $\text{RM}(q) = \text{RM}(p)$. Otherwise, that is if $\text{RM}(q) < \text{RM}(p)$, then we say that q is a *forking extension*. If a is an n -tuple then we say that a is *independent of* B over A , denoted by $a \downarrow_A B$, if $\text{tp}(a/B)$ is a nonforking extension of $\text{tp}(a/A)$; that is, if $\text{RM}(a/B) = \text{RM}(a/A)$.

We think of the nonforking extensions as “free” extensions; in some sense q does not impose to many more restrictions than p already imposed. At least if we interpret “restriction” and decreasing Morley rank. The intuition behind a being independent of B over A is similar: over B there are essentially no more relations among the co-ordinates of a than already existed over A . Again, here “essentially no more relations” means the Morley rank does not go down. As an example, consider the case when $n = 1$ and T is strongly minimal: then $a \downarrow_A B$ means that if $a \notin \text{acl}(A)$ then $a \notin \text{acl}(B)$.

NOTATION 9.2. Suppose a is an n -tuple, A is a set, and B is a set – but it is not necessarily the case that $A \subseteq B$. Then by $a \downarrow_A B$ we really mean $a \downarrow_A (A \cup B)$.

REMARK 9.3. The following properties of nonforking extensions for complete types is immediate from the definitions:

- (a) Transitivity: if q is a nonforking extension of p and r is a nonforking extension of q , then r is a nonforking extension of p .
- (b) Monotonicity: if q is a nonforking extension of p and $p \subseteq r \subseteq q$, then r is a nonforking extension of p and q is a nonforking extension of r .

These properties can be written more succinctly in terms of independence as follows: Suppose $A \subseteq B \subseteq C$. Then $a \underset{A}{\perp} C$ if and only if $a \underset{A}{\perp} B$ and $a \underset{B}{\perp} C$.

The following property of nonforking is not immediate, but follows rather easily from thing we have already done.

PROPOSITION 9.4 (Finite character). *Suppose $p \in S_n(A)$, $A \subseteq B$, and $q \in S_n(B)$. If $q \upharpoonright B_0$ is a nonforking extension of p for all $A \subseteq B_0 \subseteq B$ with $B_0 \setminus A$ finite, then q is a nonforking extension of p .*

PROOF. Suppose q is a forking extension of p . Then $\text{RM}(q) < \text{RM}(p)$. Let $\phi \in q$ witness the Morley rank of q . Let B_0 be A together with the finite set of parameters appearing in ϕ . Then $\phi \in q \upharpoonright B_0$. Hence $\text{RM}(q \upharpoonright B_0) \leq \text{RM}(\phi) < \text{RM}(p)$. \square

In terms of independence, finite character says that $a \underset{A}{\not\perp} B$ if and only if there exists a finite tuple b from B such that $a \underset{A}{\not\perp} b$.

PROPOSITION 9.5. *For every $p(x) \in S_n(B)$ there exists a finite subset $A \subseteq B$ such that p does not fork over A .*

PROOF. Let $\phi(x) \in p(x)$ be such that $\text{RM}(\phi) = \text{RM}(p)$, and let A be the finite set of parameters actually appearing in ϕ . Then $\phi \in p \upharpoonright A$ and hence $\text{RM}(p) = \text{RM}(\phi) \geq \text{RM}(p \upharpoonright A) \geq \text{RM}(p)$. So $\text{RM}(p) = \text{RM}(p \upharpoonright A)$, as desired. \square

PROPOSITION 9.6. $a \underset{A}{\perp} \text{acl}(A)$

PROOF. Suppose $\text{RM}(a/\text{acl}(A)) = \alpha$. Suppose $\phi(x, b) \in \text{tp}(a/\text{acl}(A))$ is of rank α , where $\phi(x, y)$ is an L -formula and all the co-ordinates of b are in $\text{acl}(A)$. Let $\{b_1, \dots, b_r\}$ be the orbit of b under $\text{Aut}_A(\mathcal{U})$. Then $\bigvee_{i=1}^r \phi(x, b_i)$ defines a set that is $\text{Aut}_A(\mathcal{U})$ -invariant, has Morley rank α , and contains a . Hence there exists an L_A -formula $\psi(x) \in \text{tp}(a/A)$ with $\models \forall x(\psi(x) \leftrightarrow (\bigvee_{i=1}^r \phi(x, b_i)))$. So $\text{RM}(a/A) \leq \text{RM}(\psi) = \alpha$, and so $\text{tp}(a/\text{acl}(A))$ does not fork over A . \square

Our next property is important, but also rather easy.

THEOREM 9.7 (Existence of nonforking extensions). *Suppose $A \subseteq B$, $p \in S_n(A)$. Then there exists at least one and at most $\text{dM}(p)$ -many nonforking extension of p to B . Moreover, if B is the universe of an \aleph_0 -saturated model then there exists exactly $\text{dM}(p)$ -many nonforking extensions of p to B .*

PROOF. We first prove the “moreover” clause of the theorem: assume that B is the universe of an \aleph_0 -saturated model $\mathcal{B} \preceq \mathcal{U}$. Let $\phi \in p$ be such that $(\text{RM}, \text{dM})(\phi) = (\text{RM}, \text{dM})(p) = (\alpha, d)$. By Proposition 8.14(b), there are exactly d -many types $\{q_1, \dots, q_d\} \subseteq S_n(B)$ such that $\phi \in q_i$ and $\text{RM}(q_i) = \alpha$. Note that for each i , if ψ is an L_A -formula in q_i then $\text{RM}(\phi \wedge \psi) = \alpha$, and hence, by Lemma 8.13, $\psi \in p$. That is, each q_i extends p ; it is a nonforking extension of p . Conversely, every nonforking extension of p to B is of Morley rank α and contains ϕ , and hence is one of the q_i s. So there exists exactly d -many nonforking extensions of p to B .

Now consider the general case. There exists a small \aleph_0 -saturated elementary substructure, $\mathcal{N} \preceq \mathcal{U}$, containing B .¹ By the “moreover clause”, p has d -many nonforking extensions to N . By monotonicity, the restriction to B of any such is also a nonforking extension of p . Hence p has at least one nonforking extension to B . Also by the “moreover clause” any nonforking extension of p to B has at least one nonforking extension to N , which by transitivity will be a nonforking extension of p to N – hence there are at most d -many nonforking extensions of p to B . \square

9.2. Definability of types and stable formulas

To prove further properties of nonforking independence, such as symmetry, we will need the notion of a “definable type” over a model. What should this mean? To know a type $p(x) \in S_n(M)$ is to know what L_M -formulas are in p . Hence for p to be “definable” will mean that for an arbitrary L_M -formula, $\phi(x, b)$, whether or not this formula is in p is a definable property of the parameter b .

DEFINITION 9.8 (Definability of types). Suppose $p(x) \in S_n(M)$ where $x = (x_1, \dots, x_n)$, $\mathcal{M} \preceq \mathcal{U}$, and $A \subseteq M$. We say that p is *definable over A* (or *A -definable*) if for every L -formula $\phi(x, y)$, where $y = (y_1, \dots, y_\ell)$, there is an L_A -formula $d_p\phi(y)$ such that

$$\text{for all } b \in M^\ell, \phi(x, b) \in p(x) \text{ if and only if } \models d_p\phi(b).$$

That is, in \mathcal{M} , the set of tuples b such that $\phi(x, b)$ is in p is defined by $d_p\phi$. We say that $d_p\phi(y)$ is a ϕ -*definition* of $p(x)$.

We wish to show that in a totally transcendental theory, every complete type over a model is definable. This requires a little detour into the more general terrain of *stable formulas*.

DEFINITION 9.9 (Stable formula). Suppose $\delta(x, y)$ is an L_U -formula, $x = (x_1, \dots, x_n)$, and $y = (y_1, \dots, y_\ell)$. We say that $\delta(x, y)$ is *stable* if there do not exist sequences $(a_i : i < \omega)$ and $(b_i : i < \omega)$ such that $\models \delta(a_i, b_j)$ if and only if $i < j$.

Note that the notion of a stable formula depends not only on the formula but also on its presentation in terms of two sets of free variables x and y .

Here is a motivating example of an unstable formula: Suppose $(X, <)$ is an infinite linear ordering definable in \mathcal{U} ; that is, $X \subseteq U^n$ is an infinite definable set and there is a formula

¹Indeed, by downward Löwenheim-Skolem (Theorem 4.27) we can find a small elementary substructure of $\mathcal{N}_0 \preceq \mathcal{U}$ that contains B . Let \mathcal{N}_1 be a small \aleph_0 -saturated elementary extension of \mathcal{N}_0 . By saturation we can elementarily embed \mathcal{N}_1 into \mathcal{U} over \mathcal{N}_0 .

$\phi(x, y)$ which defines a subset of X^2 that forms a linear ordering on X . Let $(a_i : i < \omega)$ be any infinite increasing sequence in X , and set $b_i = a_i$. Then $\models \phi(a_i, b_j)$ if and only if $a_i < a_j$ if and only if $i < j$. So $\phi(x, y)$ is unstable. However, as the following proposition establishes, this can never happen in a totally transcendental theory. In particular no infinite linear ordering is definable in a totally transcendental theory.

PROPOSITION 9.10. *In a totally transcendental theory every formula is stable.*

PROOF. Suppose $\phi(x, y)$ is an unstable formula, and seek a contradiction. Let L' be the language obtained from L by adding new tuples of constant symbols $c_q = (c_{q,1}, \dots, c_{q,n})$ and $d_q = (d_{q,1}, \dots, d_{q,\ell})$, for each rational number $q \in \mathbb{Q}$. Consider the L' -theory $T' := T \cup \{\phi(c_q, d_r) : q < r \in \mathbb{Q}\}$. Since ϕ is unstable, every finite subset of T' is consistent and hence by compactness T' is consistent. It follows by saturation that there exists $\{a_q : q \in \mathbb{Q}\} \subseteq U^n$ and $\{b_q : q \in \mathbb{Q}\} \subseteq U^\ell$ such that $\models \phi(a_q, b_r)$ if and only if $q < r$.

Note that for any fixed $r \in \mathbb{Q}$, the set of rational numbers q such that $\models \phi(a_q, b_r)$ is an infinite convex subset of \mathbb{Q} – namely $(-\infty, r)$. Hence there exists an formula $\psi(x)$ of minimal (RM, dM) with the property that $C = C(\psi) := \{q : \models \psi(a_q)\}$ is an infinite convex subset of \mathbb{Q} . Let r be an interior point of C . Then both $C(\psi(x) \wedge \phi(x, b_r))$ and $C(\psi(x) \wedge \neg\phi(x, b_r))$ are infinite convex subsets. But as $\psi^{\mathcal{U}}$ is the disjoint union of the sets defined by $\psi(x) \wedge \phi(x, b_r)$ and $\psi(x) \wedge \neg\phi(x, b_r)$, one of them must have strictly lower (RM, dM) – which is a contradiction. \square

LEMMA 9.11 (T arbitrary). *Suppose \mathcal{M} is a model, $\delta(x, y)$ is a stable L_M -formula, and $a \in U^n$. Then there exists an L_M -formula $\phi(y)$ such that $\phi(y)^{\mathcal{M}} = \{b \in M^\ell : \models \delta(a, b)\}$.*

PROOF. Since $\delta(x, y)$ is stable, by saturation there exists $N < \omega$ such that there do not exist sequences $(a_i : i \leq N)$ and $(b_i : i \leq N)$ with $\models \delta(a_i, b_j)$ if and only if $i < j$.

CLAIM 9.12. *Suppose $\psi(x) \in \text{tp}(a/M)$. There exists $k \leq N$ and $a_1, \dots, a_k \in M^n$ satisfying $\psi(x)$, such that for all $b \in M^\ell$, if $\models \bigwedge_{i=1}^k \delta(a_i, b)$ then $\models \delta(a, b)$.*

PROOF OF CLAIM 9.12. Suppose not and seek a contradiction to the stability of $\delta(x, y)$. Indeed, we will build (inductively) sequences $(a_i : i \leq N)$ and $(b_i : i \leq N)$ in \mathcal{M} such that $\models \neg\delta(a, b_i)$ for all i and $\models \delta(a_i, b_j)$ if and only if $i < j$. First let us find a_1, b_1 . If $\delta(a, y)^{\mathcal{U}} \cap M^\ell = M^\ell$ then setting $\phi(y)$ to be $(y = y)$ proves the lemma. So we may assume that there is $b_1 \in M^\ell$ such that $\models \neg\delta(a, b_1)$. Hence, as $\mathcal{M} \preceq \mathcal{U}$, there is also $a_1 \in M^n$ with $\models \neg\delta(a_1, b_1) \wedge \psi(a_1)$.

Now suppose we have already constructed $a_1, \dots, a_k, b_1, \dots, b_k$. Since the claim does not hold for (a_1, \dots, a_k) , there is a $b_{k+1} \in M^\ell$ such that $\models \bigwedge_{i=1}^k \delta(a_i, b_{k+1})$ but $\models \neg\delta(a, b_{k+1})$.

So $\bigwedge_{i=1}^{k+1} \neg\delta(x, b_i) \wedge \psi(x)$ is an L_M -formula realised by a , and hence realised by some $a_{k+1} \in M^n$. Then $a_1, \dots, a_{k+1}, b_1, \dots, b_{k+1}$ have the desired properties and we can continue the construction. \square

Now let $\xi(\bar{x}, y)$ be $\bigwedge_{i=1}^k \delta(x^i, y)$, where k is as given by the claim. Then $\xi(\bar{x}, y)$ is stable (exercise). Then there exists $N' < \omega$ such that there do not exist sequences $(\bar{a}_i : i \leq N')$ and $(b_i : i \leq N')$ with $\models \xi(\bar{a}_i, b_j)$ if and only if $i > j$ (exercise).

CLAIM 9.13. *There exists $r \leq N'$ and $\bar{a}_1, \dots, \bar{a}_r \in M^{nk}$ such that for all $b \in M^\ell$, b realises $\bigvee_{i=1}^r \xi(\bar{a}_i, y)$ if and only if b realises $\delta(a, y)$.*

PROOF OF CLAIM 9.13. Suppose not and seek a contradiction to the stability of $\theta(\bar{x}, y)$. Indeed, we build (inductively) sequences $(\bar{a}_i : i \leq N')$ and $(b_i : i \leq N')$ in \mathcal{M} with $\models \theta(\bar{a}_i, b_j)$ if and only if $i > j$ and such that the following additional properties hold:

- (i) $\bar{a}_i = (a_i^1, \dots, a_i^k)$ satisfies the statement of Claim 9.12 (with ψ being just $x = x$),
- (ii) $\models \delta(a, b_i)$ for all i .

To start the construction let \bar{a}_1 be the (a_1, \dots, a_k) given by Claim 9.12. So if $b \in M^\ell$ realises $\xi(\bar{a}_1, y)$ then b realises $\delta(a, y)$. If the converse holds the claim is proved with $r = 1$, so we may assume there is $b_1 \in M^\ell$ such that $\models \neg \xi(\bar{a}_1, b_1)$ while $\models \delta(a, b_1)$.

Now suppose we have already constructed $\bar{a}_1, \dots, \bar{a}_{r-1}$ and b_1, \dots, b_{r-1} . Apply Claim 9.12 with $\psi(x) = \bigwedge_{i=1}^{r-1} \delta(x, b_i)$ to get $\bar{a}_r = (a_r^1, \dots, a_r^k)$ such that $\models \xi(\bar{a}_r, b_i)$ for all $i = 1, \dots, r-1$

and such that for any $b \in M^\ell$, if $\models \xi(\bar{a}_r, b)$ then $\delta(a, b)$. As $\bigvee_{i=1}^r \xi(\bar{a}_i, y)$ does not satisfy the claim there must exist $b_r \in M^\ell$ such that $\models \delta(a, b_r)$ but $\models \neg \xi(\bar{a}_i, b_r)$ for all $i = 1, \dots, r$. Then the sequences $\bar{a}_1, \dots, \bar{a}_r$ and b_1, \dots, b_r satisfying the desired properties and we can continue the construction. \square

Claim 9.13 proves the lemma as it says that $\phi(y) := \bigvee_{i=1}^r \xi(\bar{a}_i, y)$ works. \square

THEOREM 9.14 (*T* totally transcendental). *All complete types over models are definable.*

PROOF. Suppose $p(x) \in S_n(M)$ and let $\phi(x, y)$ be an L -formula with $y = (y_1, \dots, y_\ell)$. Let $a \in U^n$ realise $p(x)$. Then $\{b \in M^\ell : \phi(x, b) \in p\} = \{b \in M^\ell : \models \phi(a, b)\}$. As $\phi(x, y)$ is stable, there is, by Lemma 9.11, an L_M -formula defining this set in \mathcal{M} . \square

We can use definability of types to prove that given a fixed definable set X and a definable family of sets Y_b , the property of intersecting X in a set of maximal Morley rank is a definable property of the parameters b .

PROPOSITION 9.15. *Suppose $\theta(x)$ is an L_A -formula, and $\phi(x, y)$ is an L -formula with $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_\ell)$. Then*

- (a) $\{b \in U^\ell : \text{RM}(\theta(x) \wedge \phi(x, b)) = \text{RM}(\theta(x))\}$ is an A -definable set,
- (b) $\{b \in U^\ell : (\text{RM}, \text{dM})(\theta(x) \wedge \phi(x, b)) = (\text{RM}, \text{dM})(\theta(x))\}$ is an A -definable set.

PROOF. For part (a) let $X := \{b \in U^\ell : \text{RM}(\theta(x) \wedge \phi(x, b)) = \text{RM}(\theta(x))\}$. It suffices to prove that X is definable, A -definability will follow by automorphisms. Let $\text{RM}(\theta) = \alpha$.

We may assume that $\text{dM}(\theta) = 1$. Indeed, otherwise we find $\theta_1(x), \dots, \theta_d(x)$, each of Morley rank α and degree 1, such that $\models \forall x(\theta(x) \leftrightarrow \bigvee_{i=1}^d \theta_i(x))$. In that case $\text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha$ if and only if $\text{RM}(\theta_i(x) \wedge \phi(x, b)) = \alpha$ for some $i = 1, \dots, d$.

We may assume that A is finite. Let \mathcal{M} be an \aleph_0 -saturated model containing A and let $p(x) \in S_n(M)$ extend $\theta(x)$ with $\text{RM}(p) = \alpha$. Note that $\text{dM}(p) = 1$ also. Hence by Lemma 8.13, for any $b \in M^\ell$, $\phi(x, b) \in p(x)$ if and only if $\text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha$. As p is definable, we have that for any $b \in M^\ell$, $\models d_p \phi(b)$ if and only if $\text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha$. We show that by \aleph_0 -saturation of \mathcal{M} the same is true for all $b \in U^\ell$. Let B be a finite set containing A which includes the parameters of $d_p \phi$. Now suppose that $b \in U^\ell$ and let $b' \in M^\ell$ realise $\text{tp}(b/B)$ in \mathcal{M} . Then

$$\begin{aligned} \models d_p \phi(b) &\iff \models d_p \phi(b') \\ &\iff \text{RM}(\theta(x) \wedge \phi(x, b')) = \alpha \\ &\iff \text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha \text{ as Morley rank is automorphism invariant.} \end{aligned}$$

So X is defined by $d_p \phi$.

For part (b) let $\theta_1(x), \dots, \theta_d(x)$ be as above. That is, each θ_i has Morley rank α and degree 1 and $\models \forall x(\theta(x) \leftrightarrow \bigvee_{i=1}^d \theta_i(x))$. Then $(\text{RM}, \text{dM})(\theta(x) \wedge \phi(x, b)) = (\alpha, d)$ if and only if $\text{RM}(\theta_i(x) \wedge \phi(x, b)) = \alpha$ for all $i = 1, \dots, d$. Hence, by part (a), the set of such bs is definable. \square

The above application of definability of types itself implies a refinement of the definability of types:

COROLLARY 9.16. *Every type over a model is definable over a finite set.*

PROOF. Suppose $p(x) \in S_n(M)$. Note that definability of p does not immediately imply that p is definable over a finite set; for each $\phi(x, y)$ we have a ϕ -definition of $p(x)$ which uses only finitely many parameters, but that finite set could *a priori* vary with ϕ . The corollary is claiming that it does not.

Let $\theta \in p$ be such that $(\text{RM}, \text{dM})(\theta) = (\text{RM}, \text{dM})(p) = (\alpha, d)$. Let A be the finite set of parameters appearing in θ . We show that p is A -definable. For any L -formula $\phi(x, y)$ and $b \in M^\ell$

$$\begin{aligned} \phi(x, b) \in p &\iff \text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha \\ &\iff \models \psi(b) \end{aligned}$$

where ψ is the L_A -formula given by Proposition 9.15(a). So ψ is a ϕ -definition of p . Hence p is definable over A . \square

Here is another application of the definability of types:

PROPOSITION 9.17. *Every type over a model has Morley rank 1.*

PROOF. Suppose $p(x) \in S_n(M)$. Note, as an aside, that if \mathcal{M} were \aleph_0 -saturated then the proposition would follow from Theorem 9.7. But we do not here assume that \mathcal{M} is \aleph_0 -saturated.

Let $\theta(x) \in p(x)$ be such that $(\text{RM}, \text{dM})(\theta) = (\text{RM}, \text{dM})(p) = (\alpha, d)$. For each $\phi(x, y)$ let $\psi_1(y)$ witness Proposition 9.15(a) and let ψ_2 witness Proposition 9.15(b). That is,

$$\psi_1^{\mathcal{U}} = \{b \in U^\ell : \text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha\}$$

and

$$\psi_2^{\mathcal{U}} = \{b \in U^\ell : (\text{RM}, \text{dM})(\theta(x) \wedge \phi(x, b)) = (\alpha, d)\}.$$

Both ψ_1 and ψ_2 can be taken to be L_M -formulas (indeed, L_A -formulas where θ is over A). But note that for $b \in M^\ell$, both $\text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha$ and $(\text{RM}, \text{dM})(\theta(x) \wedge \phi(x, b)) = (\alpha, d)$ are equivalent to $\phi(x, b) \in p$. Hence $\psi_1^{\mathcal{M}} = \psi_2^{\mathcal{M}}$, and so in fact $\psi_1^{\mathcal{U}} = \psi_2^{\mathcal{U}}$. That is, for all L_U -formulas $\phi(x, b)$, if $\text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha$ then $\text{dM}(\theta(x) \wedge \phi(x, b)) = d$. But if $d > 1$ then there exists a definable subset of $\theta(x)^{\mathcal{U}}$ that is of rank α and whose complement is of rank α . If such a set is defined by $\phi(x, b)$ then we have $\text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha$ while $\text{dM}(\theta(x) \wedge \phi(x, b)) < d$ – a contradiction. So $d = 1$, as desired. \square

COROLLARY 9.18. *If $p(x) \in S_n(M)$ is a type over a model and $B \supseteq M$, then $p(x)$ has a unique nonforking extension to B .*

PROOF. This is just Proposition 9.17 together with Theorem 9.7. \square

9.3. Further properties of nonforking

DEFINITION 9.19. A type $p(x) \in S_n(A)$ is *stationary* if for all $B \supseteq A$, $p(x)$ has a unique nonforking extension to B . In that case we denote by $p(x) \upharpoonright B$ this unique nonforking extension.

Corollary 9.18 tells us that types over models are stationary.

Suppose $p(x) \in S_n(M)$ is a complete type over a model, $\phi(x, y)$ is an L -formula and $d_p\phi(y)$ is a ϕ -definition of p (and hence an L_M -formula). Then $\{b \in M^\ell : \models d_p\phi(b)\}$ tells us which $\phi(x, b)$ s are in p . What about realisations b' of $d_p\phi(y)$ outside of \mathcal{M} ? The next proposition tells us that if b' is such then $\phi(x, b')$ is in the nonforking extension of p .

THEOREM 9.20. *Suppose \mathcal{M} is a model, $p(x) \in S_n(M)$, and for each L -formula $\phi(x, y)$, let $d_p\phi(y)$ be a ϕ -definition of $p(x)$ that is over M . Suppose $B \supset M$. Then*

$$p(x) \upharpoonright B = \{\phi(x, b) : \phi(x, y) \text{ an } L\text{-formula}, b \in B^\ell, \models d_p\phi(b)\}.$$

In particular, if $M \subseteq N$ are models and $r(x) \in S_n(N)$, then $r(x)$ does not fork over M if and only if $r(x)$ is definable over M . In that case, a ϕ -definition of $r \upharpoonright M$ is also a ϕ -definition of r , for all $\phi(x, y)$.

PROOF. Let us first show how the “in particular” clause follows from the main statement. Let $p(x) := r \upharpoonright M$. If $r(x)$ does not fork over M then $r(x) = p(x) \upharpoonright N$. By the main statement of the theorem, $p(x) \upharpoonright N$ is definable over M , and the $d_p\phi(y)$ are ϕ -definitions for it. Conversely, suppose r is definable over M , and let $d_r\phi(y)$ be a ϕ -definition for r that is over M . Then $d_r\phi(y)$ is also a ϕ -definition for $r \upharpoonright M = p$, so let us take $d_p\phi := d_r\phi$. We

have $\phi(x, b) \in r(x)$ if and only if $\models d_p\phi(b)$. By the main statement of the theorem, we see that $r = p(x) \upharpoonright N$, and it therefore does not fork over M .

We now prove the main statement of the theorem. Let

$$q(x) := \{\phi(x, b) : \phi(x, y) \text{ an } L\text{-formula}, b \in B^\ell, \models d_p\phi(b)\}.$$

We first show that q is a type. Suppose $\phi_1(x, b_1), \dots, \phi_k(x, b_k) \subseteq q(x)$. If this collection is not finitely satisfiable then

$$\models \neg \exists x \left(\bigwedge_{i=1}^k \phi_i(x, b_i) \right) \wedge \bigwedge_{i=1}^k d_p\phi_i(b_i).$$

As $\mathcal{M} \preceq \mathcal{U}$, we have b'_1, \dots, b'_k from \mathcal{M} such that

$$\models \neg \exists x \left(\bigwedge_{i=1}^k \phi_i(x, b'_i) \right) \wedge \bigwedge_{i=1}^k d_p\phi_i(b'_i).$$

But as each b'_i realises $d_p\phi_i(y)$ in \mathcal{M} , $\phi_i(x, b_i) \in p$ for all $i = 1, \dots, k$; and hence $\left(\bigwedge_{i=1}^k \phi_i(x, b'_i) \right)$ is realised by any realisation of p . This contradiction proves that $q(x)$ is finitely satisfiable and hence $q(x)$ is a type.

Next we show it is complete. For any L -formula $\phi(x, y)$, note that for all $b' \in M^\ell$, $\phi(x, b') \in p(x)$ if and only if $\neg\phi(x, b') \notin p(x)$ by completeness of $p(x)$. So, $d_p\phi(y)^{\mathcal{M}} = (-d_p\neg\phi(y))^{\mathcal{M}}$. Hence, $d_p\phi(y)^{\mathcal{U}} = (-d_p\neg\phi(y))^{\mathcal{U}}$, and so for all $b \in B^\ell$, $\phi(x, b) \in q(x)$ if and only if $\neg\phi(x, b) \notin q(x)$. Now suppose $\psi(x)$ is an L_B -formula. Write $\psi(x) = \phi(x, b)$ for some L -formula $\phi(x, y)$ and some $b \in B^\ell$. Then by what we have just proved, either $\psi \in q$ or $\neg\psi \in q$. That is, q is complete.

Next we show that q is a nonforking extension of p . It is clearly an extension of p . Let $\theta(x) \in p$ with $\text{RM}(\theta) = \alpha$. Toward a contradiction, suppose there is $\phi(x, b) \in q(x)$ with $\text{RM}(\phi(x, b)) < \alpha$. Then also $\text{RM}(\theta(x) \wedge \phi(x, b)) < \alpha$. Let $\psi(y)$ be the L_M -formula given by Proposition 9.15(a) – that is, the formula defining the set of $b' \in U^\ell$ with $\text{RM}(\theta(x) \wedge \phi(x, b')) = \alpha$. So $\models d_p\phi(b) \wedge \neg\psi(b)$. Hence there is $b' \in M^\ell$ with $\models d_p\phi(b') \wedge \neg\psi(b')$. This means that $\phi(x, b') \in p$ but $\text{RM}(\theta(x) \wedge \phi(x, b')) < \alpha$, contradicting the minimality of α .

Finally, to see that it is the unique nonforking extension to B , use the fact that $\text{dM}(p) = 1$ by Proposition 9.17, and then apply Theorem 9.7. \square

LEMMA 9.21. *Suppose \mathcal{M} is model, $p(x) \in S_n(M)$, and $q(y) \in S_\ell(M)$. Let $\phi(x, y)$ be an L -formula, and let $\psi(y, x) := \phi(x, y)$. Let $d_p\phi(y)$ be a ϕ -definition of $p(x)$, and let $d_q\psi(x)$ be a ψ -definition of $q(y)$. So*

$$d_p\phi(y)^{\mathcal{M}} = \{b \in M^\ell : \phi(x, b) \in p\} \quad \text{and} \quad d_q\psi(x)^{\mathcal{M}} = \{a \in M^n : \phi(a, y) \in q\}.$$

Then, $d_p\phi(y) \in q(y)$ if and only if $d_q\psi(x) \in p(x)$.

PROOF. We first observe that we may assume \mathcal{M} is \aleph_0 -saturated. Indeed, let \mathcal{N} be an \aleph_0 -saturated model containing \mathcal{M} . Then by Theorem 9.20, $d_p\phi$ is the ϕ -definition of $p \upharpoonright N$ and $d_q\psi$ is the ψ -definition of $q \upharpoonright N$. So if we know the result for $p \upharpoonright N$ and $q \upharpoonright N$, then we know the result for p and q .

Suppose, toward a contradiction, that $d_p\phi(y) \in q(y)$ but $\neg d_q\psi(x) \in p(x)$. Let $A \subset M$ be a finite set over which both $d_p\phi$ and $d_q\psi$ are defined. Let $a_1 \in M^n$ realise $p \upharpoonright A$ and $b_1 \in M^\ell$ realise $p \upharpoonright (Aa_1)$. Iterate this to build sequences $(a_i : i < \omega)$ in M^n and $(b_i : i < \omega)$ in M^ℓ such that a_{r+1} realises $p \upharpoonright (Ab_1 \dots b_r)$ and b_{r+1} realises $q \upharpoonright (Aa_1 \dots a_{r+1})$.

If $\models \phi(a_i, b_j)$ and $i \leq j$ then $\phi(a_i, y) \in q$ and so $\models d_q\psi(a_i)$, which contradicts the fact that $\neg d_q\psi(x) \in p(x) \upharpoonright A$. Conversely, if $i > j$ and $\models \neg\phi(a_i, b_j)$, then $\neg\phi(x, b_j) \in p(x)$ and so $\models \neg d_p\phi(b_j)$ which contradicts the fact that $d_p\phi(y) \in q(y) \upharpoonright A$. Hence, $\models \phi(a_i, b_j)$ if and only if $i > j$. This contradicts the stability of $\phi(x, y)$. \square

PROPOSITION 9.22 (Symmetry). *If $a \downarrow_A b$ then $b \downarrow_A a$.*

PROOF. Let us first deal with the case when $A = M$ is a model. Let $p(x) = \text{tp}(a/M)$ and $q(y) = \text{tp}(b/M)$. Note that

$$\begin{aligned} a \downarrow_M b &\iff \text{tp}(a/Mb) \text{ is a nonforking extension of } \text{tp}(a/M) \\ &\iff a \text{ realises } p(x) \upharpoonright (Mb) \end{aligned}$$

For each L_M -formula $\phi(x, y)$ let $d_p\phi(y)$ be a ϕ -definition for $p(x)$, working in the language L_M and the theory $\text{Th}(\mathcal{U}_M)$. (Note that $\text{Th}(\mathcal{U}_M)$ is still totally transcendental and \mathcal{U}_M is still a sufficiently saturated model.) Then by the characterisation of $p(x) \upharpoonright (Mb)$ given in Theorem 9.20,

$$\begin{aligned} a \downarrow_M b &\iff a \text{ realises } p(x) \upharpoonright (Mb) \\ &\iff \models \phi(a, b) \text{ for all } L_M\text{-formulas } \phi(x, y) \text{ such that } \models d_p\phi(b) \\ &\iff \models \phi(a, b) \text{ for all } L_M\text{-formulas } \phi(x, y) \text{ such that } d_p\phi(y) \in q(y) \end{aligned}$$

Now repeating the argument with the places of $a, p(x)$ and $b, q(y)$ reversed we have

$$b \downarrow_M a \iff \models \phi(a, b) \text{ for all } L_M\text{-formulas } \phi(x, y) \text{ such that } d_q\psi(x) \in p(x)$$

So by Lemma 9.21, $a \downarrow_M b$ if and only if $b \downarrow_M a$.

Finally, let us consider the case when A is arbitrary. We assume $a \downarrow_A b$. Let \mathcal{M} be a model containing A and let b' realise a nonforking extension of $\text{tp}(b/A)$ to \mathcal{M} . So there is an $f \in \text{Aut}_A(\mathcal{U})$ be such that $f(b) = b'$. Let a' realise a nonforking extension of $\text{tp}(f(a)/Ab')$ to Mb' . So $a' \downarrow_{Ab'} Mb'$.

We claim that $a' \downarrow_M b'$. Indeed, as $a \downarrow_A b$, we get $f(a) \downarrow_A b'$ by automorphisms, and hence $a' \downarrow_A b'$ since a' and $f(a)$ have the same type over Ab' . This together with $a' \downarrow_{Ab'} Mb'$ implies, by transitivity, that $a' \downarrow_A Mb'$. So by monotonicity, $a' \downarrow_M b'$.

It follows by symmetry over models that $b' \downarrow_M a'$. On the other hand, $b' \downarrow_A M$ by choice, and so by transitivity $b' \downarrow_A Ma'$. Monotonicity yields $b' \downarrow_A a'$. By automorphisms we first get $b' \downarrow_A f(a)$ and then $b \downarrow_A a$, as desired. \square

An Introduction to Totally Transcendental Groups

In this final chapter we will provide fewer details, leaving more up to the reader.

10.1. Definable groups

By a *definable group* in a structure \mathcal{M} we mean a definable set $G \subseteq M^m$ together with a definable function $\cdot : G \times G \rightarrow G$. Note that the identity $e \in G$ and the inverse function $^{-1} : G \rightarrow G$ are then also definable.

The motivating idea is that of an *algebraic group*; an algebraic variety (say over an algebraically closed field K) together with a group operation that is a morphism of algebraic varieties over K . So an algebraic group is a group-object in the category of algebraic varieties over K . For example, an elliptic curve is an algebraic group. Another class of examples is the class of linear algebraic groups. We can view $\mathrm{GL}_n(K)$ as a Zariski closed subset K^{2n+1} by identifying an $n \times n$ matrix with an n^2 -tuple from K and then viewing

$$\mathrm{GL}_n(K) = \{(A, w) \in K^{n^2+1} : w(\det A) = 1\}$$

with the group operation given by $(A, w) \cdot (B, v) = (AB, vw)$. A *linear algebraic group* is then a subgroup of $\mathrm{GL}_n(K)$, for some n , that is Zariski closed. It is a (nontrivial) fact, requiring both explanation and proof, that the algebraic groups are exactly the groups definable in algebraically closed fields.

In analogy with algebraic groups, definable groups in \mathcal{U} are group-objects in the category $\mathrm{Def}(\mathcal{M})$. However the connection with algebraic groups is even stronger: a still open conjecture (of Cherlin–Zil’ber) says that if $\mathrm{RM}(G) < \omega$, and G is an infinite simple group, then there is an algebraically closed field K definable in \mathcal{M} and G is definably isomorphic to an algebraic group over K .

Suppose (G, \cdot) is a group definable in \mathcal{M} . Note that by a definable set $X \subseteq G^\ell$ we mean that X is definable in \mathcal{M} ; not that X is definable in (G, \cdot) viewed in its own right as a structure in the language of groups. Of course, if X is definable in the group structure (G, \cdot) then it is also definable in \mathcal{U} . However, there may be many more definable subsets of G^ℓ , depending on what the signature of \mathcal{M} is. (It may happen that $\mathcal{M} = (G, \cdot)$ and then there are no other definable sets, but that is a very special case.)

What about parameters? Well, G and \cdot only involve some fixed finite set of parameters, and so we will often just name these to the language and assume that the group is 0-definable. The definable subsets of G (and its cartesian powers) will of course involve additional parameters, which *a priori* may come from outside G . However, in a totally transcendental theory you can always take the parameters as coming from inside G . To see this we first point out that types over arbitrary sets are definable.

LEMMA 10.1 (Compare with Theorem 9.14). *Suppose T is totally transcendental, $\mathcal{M} \models T$, $A \subset M$, and $p(x) \in S_n(A)$. Then for each L -formula $\phi(x, y)$, p has a ϕ -definition. That is, there exists an L_A -formula, $d_p\phi(y)$, such that*

$$\mathcal{M} \models d_p\phi(b) \iff \phi(x, b) \in p(x)$$

for all tuples b from A .

PROOF. Work in a sufficiently saturated model $\mathcal{U} \preceq \mathcal{M}$. Let $\theta(x) \in p(x)$ be such that $\text{RM}(\theta) = \text{RM}(p) =: \alpha$. Let $d_p\phi(y)$ be the L_A -formula given by Proposition 9.15(a). That is,

$$d_p\phi(y)^{\mathcal{U}} = \{b \in U^\ell : \text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha\}.$$

But for b from A , $\phi(x, b) \in p$ if and only if $\text{RM}(\theta(x) \wedge \phi(x, b)) = \alpha$. \square

PROPOSITION 10.2. *Suppose T is totally transcendental, $\mathcal{M} \models T$, $A \subseteq M$, and X is an A -definable set in \mathcal{M} . Then every definable subset of X is definable over (a finite subset of) $A \cup X$.¹*

PROOF. Suppose $X \subseteq M^\ell$ and $Y \subseteq X$ is defined by $\phi(a, y)$ where $a \in M^n$ and $y = (y_1, \dots, y_\ell)$. Let $p(x) = \text{tp}(a/X)$ and consider $d_p\phi(y)$ be an L_X -formula given by Lemma 10.1. Then for $b \in M^\ell$,

$$\begin{aligned} b \in Y &\iff \mathcal{M} \models \phi(a, b) \\ &\iff b \in X \text{ and } \phi(x, b) \in p(y) \\ &\iff a \in X \text{ and } \mathcal{M} \models d_p\phi(b) \end{aligned}$$

So $Y = X \cap (d_p\phi(y)^{\mathcal{M}})$. \square

In particular, if (G, \cdot) is a 0-definable group in \mathcal{M} then every definable subset of G (or its cartesian powers) is defined using parameters from G . So in studying the structure on (G, \cdot) induced from \mathcal{M} , we do not need to worry about parameters from outside G .

10.2. Descending chain condition and connected components

For the rest of this chapter we renew the conventions of the last: T is a totally transcendental theory and \mathcal{U} is a sufficiently saturated model that serves as a universal domain. Moreover, we fix an infinite 0-definable group (G, \cdot) in \mathcal{U} .

PROPOSITION 10.3 (DCC). *There is no infinite decreasing chain of definable subgroups.*

PROOF. Suppose $G = G_0 \geq G_1 \geq \dots$ is a decreasing chain of definable subgroups. Note that G_i is a union of cosets of G_{i+1} , and each coset being definably isomorphic to G_{i+1} is of the same Morley rank as G_{i+1} . Since distinct cosets are disjoint, if $G_{i+1} \neq G_i$, then $(\text{RM}, \text{dM})(G_i) > (\text{RM}, \text{dM})(G_{i+1})$. As this lexicographic ordering is a well-ordering, the chain must eventually stabilise. \square

COROLLARY 10.4. (a) *Every injective definable group homomorphism, $f : G \rightarrow G$, is surjective.*

¹By $A \cup X$ here I mean the set made up of A together with all the co-ordinates of elements of X .

- (b) If $\{H_i : i \in I\}$ is a set of definable subgroups of G then $\bigcap_{i \in I} H_i = \bigcap_{j \in J} H_j$ where $J \subseteq I$ is a finite subset. In particular, $\bigcap_{i \in I} H_i$ is a definable subgroup of G .
- (c) For any $A \subseteq G$, the centraliser of A , $C(A) := \{g \in G : ga = ag \text{ for all } a \in A\}$, is a definable subgroup of G .

PROOF. For part (a) let $G_i = f^i(G)$. This forms a decreasing chain of definable subgroups. By the DCC, $G_{i+1} = G_i$ for some i . Applying $(f^{-1})^i$ to both sides we get $G = f(G)$, as desired.

For part (b), enumerate $I = \{i_\alpha : \alpha < |I|\}$, set $G_0 := G$, and for each $\alpha > 0$, set $G_\alpha := \bigcap_{\beta < \alpha} H_{i_\beta}$. Now apply DCC to this decreasing chain of definable subgroups.

For part (c) just note that $C(A) = \bigcap_{a \in A} C(\{a\})$ and apply part (b). □

DEFINITION 10.5 (Connected component). Let G^0 be the intersection of all the definable subgroups of G of finite index in G . We call G^0 the *connected component* of G and say that G is *connected* if $G = G^0$.

LEMMA 10.6. Suppose G is connected and $h : G \rightarrow G$ is a definable group homomorphism with finite kernel. Then h is surjective.

PROOF. Let h be A -definable. First note that for every $a \in G$, $a \in \text{acl}(Ah(a))$ and $h(a) \in \text{acl}(Aa)$. Hence, $\text{RM}(a/A) = \text{RM}(h(a)/A)$ by Lemma 8.16. It follows that $\text{RM}(h(G)) = \text{RM}(G)$, and so $h(G)$ is of finite index in G . By connectedness, $h(G) = G$. □

PROPOSITION 10.7. G^0 is a 0-definable normal subgroup of G of finite index in G .

PROOF. That G^0 is definable and of finite index is just Corollary 10.4(b). Since being a definable subgroup of finite index in G is preserved by all automorphisms of \mathcal{U} , G^0 being the intersection of all such, is preserved by all automorphism of \mathcal{U} . Hence G^0 is 0-definable.

We show that G^0 is normal. Fix $h \in G$ and consider the group homomorphism $f : G \rightarrow G$ given by $f(x) = h x h^{-1}$. Note that f is in fact an automorphism and it is definable. Hence it preserves the set of all definable subgroups of finite index. It therefore preserves the intersection of this set, namely G^0 . We have shown that for all $h \in G$, $h G^0 h^{-1} = G^0$. That is, G^0 is normal. □

10.3. Generics

We begin with a definition we could have (and maybe should have) made in the last, or even prior, chapter.

DEFINITION 10.8 (Generic type). Suppose $X \subseteq U^n$ is an A -definable set and $p(x) \in S_n(A)$ is a complete type extending X . Then p is a *generic type of X over A* if $\text{RM}(p) = \text{RM}(X)$. If a realises p then we say that a is *generic in X over A* .

Note that for any $A \subseteq B$, a is generic in X over B if and only if a is generic over A and $a \downarrow_A B$. That generic types always exist is Proposition 8.14(a). This definition extends the notion of an acl-generic type in strongly minimal theories, introduced in Lemma 8.29.

LEMMA 10.9. *If $\text{dM}(X) = 1$ then X has a unique generic type over A .*

PROOF. Toward a contradiction suppose $p_1(x) \neq p_2(x)$ are distinct generic types in X over A . Then there exists $\theta_i \in p_i$ such that $\theta_1^{\mathcal{U}} \cap \theta_2^{\mathcal{U}} = \emptyset$. Taking conjunctions with other formulas in $p_i(x)$ we may assume that each $\theta_i^{\mathcal{U}} \subseteq X$ and that $\text{RM}(\theta_i) = \text{RM}(p_i) = \text{RM}(X)$. So X contains two disjoint definable subsets of rank $\text{RM}(X)$, contradicting $\text{dM}(X) = 1$. \square

PROPOSITION 10.10. *Suppose \mathcal{M} is a model and X is M -definable. Then there exist exactly $\text{dM}(X)$ -many distinct generic types in X over M .*

PROOF. It suffices to prove this in the case that \mathcal{M} is \aleph_0 -saturated (because of Corollary 9.18, exercise). Write X as a pairwise disjoint union of M -definable sets Y_1, \dots, Y_d where $d = \text{dM}(X)$, and each $(\text{RM}, \text{dM})(Y_i) = (\text{RM}(X), 1)$. Let p_1, \dots, p_d be the generic types of Y_1, \dots, Y_d respectively. Then these are the generic types of X over M (exercise). \square

We now study the generics of our 0-definable group (G, \cdot) .

LEMMA 10.11. *Suppose $a \in G$. The following are equivalent*

- (i) *a is generic in G over A*
- (ii) *For all $g \in G$, if $a \downarrow g$ then ga is generic in G over Ag .*
- (iii) *For all $g \in G$, if $a \downarrow_A g$ then $ga \downarrow_A g$.*

Similar equivalences hold for right multiplication.

PROOF. First of all, note that if $\phi(x)$ is a formula with $\phi^{\mathcal{U}} \subseteq G$ and $g \in G$, then $g \cdot (\phi^{\mathcal{U}})$ is defined by the formula $\phi(g^{-1} \cdot x)$.² Since left multiplication by g is a definable bijection of G onto itself, it follows that $\text{RM}(\phi(x)) = \text{RM}(\phi(g^{-1} \cdot x))$. On the other hand, notice that for any $a \in G$, $\text{tp}(ga/Ag) = \{\phi(g^{-1} \cdot x) : \phi(x) \in \text{tp}(a/Ag)\}$. Hence, $\text{RM}(a/Ag) = \text{RM}(ga/Ag)$.

Now suppose a is generic in G over A and $a \downarrow_A g$. Then $\text{RM}(ga/Ag) = \text{RM}(a/Ag) = \text{RM}(a/A) = \text{RM}(G)$, as desired. Similarly (i) implies (ii) with right multiplication.

That (ii) implies (iii) is clear for both right and left multiplication.

Assuming (iii) we now show (i). Let $b \in G$ be generic in G over Aa . Then $b \downarrow_A a$ and so by (i) implies (ii) for right multiplication, we get that ba is generic over Aa , and hence over A . On the other hand, $a \downarrow_A b$ by symmetry, and so by (iii), $ba \downarrow_A b$. It follows that $\text{RM}(a/Ab) = \text{RM}(ba/Ab) = \text{RM}(ba/A) = \text{RM}(G)$. So a is generic over Ab , and hence over A . Similarly (iii) implies (i) for right multiplication. \square

LEMMA 10.12. *Suppose \mathcal{M} is a model, $p(x) \in S_m(M)$ is a generic type of G over M , and $g \in G$. Then there exists a generic type $q(x)$ of G over M such that for any a realising $p(x)$ with $a \downarrow_M g$, ga realises $q(x)$. We write this as $g \cdot p = q$.*

²This is just shorthand for an actual formula which exists because both $^{-1}$ and \cdot are 0-definable.

PROOF. By Lemma 10.11 we know that ga is generic over M ; what we need to show is that $\text{tp}(ga/M)$ depends only on g and p , and not on the choice of a . Toward this, suppose a' is another realisation of $p(x)$ with $a' \downarrow_M g$. Since $p(x)$ is stationary we have that $\text{tp}(a/Mg) = \text{tp}(a'/Mg)$. Let $f \in \text{Aut}_{Mg}(\mathcal{U})$ with $f(a) = a'$. Then $f(ga) = ga'$. Hence, in particular, $\text{tp}(ga/M) = \text{tp}(ga'/M)$. \square

Lemma 10.12 defines an action of G on the set of generic types of G over M . We leave it as an exercise for you to show that this is a group action.

PROPOSITION 10.13. *Suppose \mathcal{M} is a model and S is the set of generic types of G over M . Then*

- (a) G acts transitively on S , and
- (b) for each $p(x), q(x) \in S$, $\{g \in G : g \cdot p = q\}$ is M -definable.

PROOF. Suppose $p(x), q(x) \in S$, and let a realise p , b realise q , with $b \downarrow_M a$. By Proposition 9.6, $b \downarrow_M a^{-1}$. By Lemma 10.11, $b \cdot a^{-1} \downarrow_M a^{-1}$, and hence $b \cdot a^{-1} \downarrow_M a$. Setting $g = b \cdot a^{-1}$ we have $g \downarrow_M a$ and $ga = b$. So $g \cdot p = q$. This proves the transitivity of the action.

Toward a proof of (b), for each $p(x) \in S$ and let $\theta_p(x) \in p$ be such that $\theta_p^{\mathcal{U}} \subseteq G$ and $(\text{RM}, \text{dM})(\theta_p) = (\text{RM}, \text{dM})(p) =: (\alpha, 1)$. Recall that θ_p isolates p from among all types over M of rank $\geq \alpha$ (cf. Lemma 8.13). In particular, θ_p isolates p from among the types in S .

For each $p \in S$ we let $\mathbf{p}(x) \in S_m(\mathcal{U})$ denote the nonforking extension of $p(x)$ to the whole universe. I use the boldface font here because such a type contradicts our conventions; it is not over a small set of parameters. We call such a type a *global type*.

CLAIM 10.14. *Suppose $p, q \in S$ and $g \in G$. Then $g \cdot p = q$ if and only if $\theta_p(g^{-1} \cdot x) \in \mathbf{q}(x)$.*

PROOF OF CLAIM 10.14. Suppose $\theta_p(g^{-1} \cdot x) \in \mathbf{q}(x)$. Then $\theta_p(g^{-1} \cdot x) \in \mathbf{q}(x) \upharpoonright Mg = q \upharpoonright Mg$. Hence there exists b realising $q(x)$ with $b \downarrow_M g$, such that $\models \theta_p(g^{-1}b)$. But by Lemma 10.11, $a := g^{-1}b$ is generic in G over Mg . So $\text{tp}(a/M) \in S$, but it contains θ_p , and hence $\text{tp}(a/M) = p$. We have that $ga = b$ and $a \downarrow_M g$. So $g \cdot p = q$ by definition. I leave the the converse as an exercise; noting that each of the steps in the above argument are reversible. \square

Now, by Theorem 9.20 (applied in a sufficiently saturated elementary extension of \mathcal{U} !), each $\mathbf{q}(x)$ is definable over M . Hence, the set of all $g \in G$ such that $\theta_p(g^{-1} \cdot x) \in \mathbf{q}(x)$ is M -definable – indeed, it is defined by the $\phi_p(y)$ -definition of \mathbf{q} where $\phi_p(x, y) := (y \in G) \wedge \theta_p(y^{-1} \cdot x)$. This, together with Claim 10.14, completes the proof of part (b). \square

THEOREM 10.15. *The index of G^0 in G is $\text{dM}(G)$. In particular, G is connected if and only if it has a unique generic type over the empty set. Moreover, over any model M , the generic types of G over M correspond to the generic types of the cosets of G^0 over M .*

PROOF. Let $r = [G : G^0]$ and $d = \text{dM}(G)$. Then G is a pairwise disjoint union of r -many cosets of G^0 , hence $\text{RM}(G) = \text{RM}(G^0)$ and $d = r \text{dM}(G^0) \geq r$. For the converse, let \mathcal{M} be a model and let $S = \{p_1(x), \dots, p_d(x)\}$ be the set of generic types in G over M . We have

a transitive group action of G on S . Moreover, $G_1 := \text{Stab}(p_1) = \{g \in G : g \cdot p_1 = p_1\}$ is an M -definable subgroup of G . Now G/G_1 is in bijection with the orbit of p_1 which is S , so $[G : G_1] = d$. It follows that $G^0 \leq G_1$ and hence $r \geq d$. So $r = d$.

Fix a model \mathcal{M} and let $g_1, \dots, g_d \in G \cap M^m$ be such that $G = \bigcup_{i=1}^d g_i \cdot G^0$. It follows that $\text{dM}(g_i \cdot G^0) = 1$ for each i , and that the generic types of G over M are exactly the generic types of the $g_i \cdot G^0$ over M (cf. the proof of Proposition 10.10). \square

10.4. Totally transcendental fields

By a definable field $(K, +, \times)$ in \mathcal{U} we mean a definable set K together with definable functions $+$: $K^2 \rightarrow K$ and \times : $K^2 \rightarrow K$ satisfying the axioms of a field.

THEOREM 10.16 (Macintyre's Theorem). *Every infinite field definable in a totally transcendental theory is algebraically closed. In particular, if $(K, 0, 1, +, -, \times)$ is an infinite totally transcendental field then K is algebraically closed.*

PROOF. We proceed by a series of four claims.

CLAIM 10.17. *Both the additive and multiplicative groups of an infinite field definable in a totally transcendental theory are connected.*

PROOF OF CLAIM 10.17. Suppose $(K, +, \times)$ is an infinite definable in \mathcal{U} (a sufficiently saturated model of a totally transcendental theory). Consider the additive group $(G, \cdot) := (K, +)$, which is 0-definable in \mathcal{U} . Let $a \in K \setminus \{0\}$, and consider the group isomorphism $x \mapsto ax$, which is a definable map from G to itself. Since this map fixes the set of all finite index definable subgroups of G , it fixes the connected component G^0 setwise. What this means is that G^0 is an ideal of K , and hence is equal to K . So $(K, +)$ is connected. It follows by Theorem 10.15 that $\text{dM}(K) = 1$, and hence $(K \setminus \{0\}, \times)$ is also a connected group. \square

CLAIM 10.18. *Every infinite field definable in a totally transcendental theory is perfect.*

PROOF OF CLAIM 10.18. For each $n > 1$, consider the 0-definable multiplicative homomorphism $x \mapsto x^n$. As it has a finite kernel, and the multiplicative group is connected, this map is surjective by Lemma 10.6. That is, every element of our definable field K has an n th root in K , for all $n > 1$. In particular, taking $n = p = \text{char}(K)$, we see that K is perfect. \square

CLAIM 10.19. *Suppose $(K, +, \times)$ is an infinite field definable in a totally transcendental theory, $1 < n < \omega$, and K contains all m th roots of unity for $m \leq n$. Then K has no Galois extensions of degree n .*

PROOF OF CLAIM 10.19. Assume this claim is false and seek a contradiction. Let $n > 1$ be least such that there exists an infinite field $(K, +, \times)$ definable in a totally transcendental theory, with the property that K has all m th roots of unity for $m \leq n$, but K has a Galois extension L of degree n .

We wish to argue that n is prime. If q is a prime dividing n , then by Galois theory there is an intermediate field $K \subseteq F \subseteq L$ such that L/F is a Galois extension of degree q . Now

F is a finite simple extension of K (again by Galois theory), so let's write it as $F = K(\alpha)$. Then $\{1, \alpha, \dots, \alpha^{r-1}\}$ is a K -basis for F , where $r = \frac{n}{q}$ is the degree of the extension F/K . Let us define $\otimes : (K^r)^2 \rightarrow K^r$ by $(a_0, \dots, a_{r-1}) \otimes (b_0, \dots, b_{r-1}) = (c_0, \dots, c_{r-1})$ if

$$\left(\sum_{i=0}^{r-1} a_i \alpha^i\right) \left(\sum_{i=0}^{r-1} b_i \alpha^i\right) = \left(\sum_{i=0}^{r-1} c_i \alpha^i\right).$$

I leave it as an exercise for you to check that \otimes is a definable function. Define $\oplus : (K^r)^2 \rightarrow K^r$ by co-ordinate wise addition, which is clearly definable. Then (K^r, \oplus, \otimes) is a definable field in (the totally transcendental \mathcal{U}) which is isomorphic to $(F, +, \times)$. Since K contains all m th roots of unity for $m \leq n$, the same is true of F and hence of the definable field (K^r, \oplus, \otimes) . In particular, (K^r, \oplus, \otimes) contains all m th roots of unity for $m \leq q$. On the other hand, F has a Galois extension of degree q (namely L), and hence so does (K^r, \oplus, \otimes) . By our minimal choice of n it must be the case that $n = q$.

So L/K is a Galois extension of degree a prime q . Now there are two cases. The first case is that $\text{char}(K)$ is not q (so it is either 0 or it is some other prime). In that case, Galois theory tells us that the minimal polynomial of L/K is of the form $X^q - a$ for some $a \in K$. But we have seen (in the proof of Claim 10.18) that $x \mapsto x^q$ is a surjective map from K to itself. Hence $X^q - a$ has a root in K , which implies that $X^q - a$ is reducible over K , which is a contradiction.

The other case is that $\text{char}(K) = q$. In that case Galois theory tells us that the minimal polynomial of L/K is of the form $X^q + X - a$ for some $a \in K$. But the additive homomorphism $x \mapsto (x^q + x)$ also has a finite kernel, and so by the connectedness of $(K, +)$, this map is also surjective. So $X^q + X - a$ must have a root in K and hence is reducible over K , which is a contradiction. This completes the proof of Claim 10.19. \square

CLAIM 10.20. *Every infinite field definable in a totally transcendental theory contains all roots of unity.*

PROOF OF CLAIM 10.20. Suppose $(K, +, \times)$ is a counter-example and let d be least such that K does not contain all the d th roots of unity. Let ξ be a primitive d th root of unity. Then $K(\xi)$ is a proper Galois extension of K and its degree, say n , is at most $d - 1$ (as K does contain some d th root of unity, namely 1). By Claim 10.19, for some $m \leq n$, K does not contain all m th roots of unity. But this contradicts the minimal choice of d . \square

We now prove the theorem. Suppose $(K, +, \times)$ is an infinite field definable in the totally transcendental \mathcal{U} . Claims 10.20 and 10.19 imply that K has no proper Galois extensions. But by Claim 10.18 K is perfect, so every polynomial over K is separable. This means that K has no algebraic extensions whatsoever. That is, K is algebraically closed, as desired. \square