

Tame and wild theories

An introduction to the logic of dependence and independence

José Iovino

DEPARTMENT OF MATHEMATICS
THE UNIVERSITY OF TEXAS AT SAN ANTONIO

© 2026 José Iovino. All rights reserved.

These notes are a work in progress. Please do not post or redistribute copies;
link instead to the current version at [utsa-quantum-math.github.io](https://github.com/utsa-quantum-math).

Comments are welcome at jose.iovino@utsa.edu.

To Xavier Caicedo

Introduction

Some theories are tame and some are wild. The distinction looks like a matter of taste, and the purpose of this book is to show that it is not: it can be drawn precisely, and once drawn it decides what can be proved. A formula in two variables defines a relation between two sets. Shelah measured the wildness of such a relation by the configurations it admits. A formula is *stable* if it cannot order arbitrarily long sequences. It has the *negation of the independence property* if it cannot shatter arbitrarily large finite sets. A theory is *treeless* if none of its formulas grows a certain kind of tree. These are the dividing lines between the tame and the wild. On the tame side of each, a theory admits a structure theory. On the wild side, examples show that none is to be had.

The names deserve a word, for names in mathematics are not idle: they record what their authors thought they had found. *Treeless* and *negation of the independence property* are Shelah's original names.¹ Each says exactly what the class is: in a treeless theory no formula grows the tree, and the second name does not pretend to be more than a negation. The literature has since preferred *simple* for the first and the abbreviation NIP for the second. In this introduction I use the original names. In the chapters that follow I use the established ones, so that the reader can open the literature.

A definition by negation states what a thing is not, and the old *via negativa* held that some things can be described in no other way. Each dividing line is drawn in this manner: a configuration is named and then forbidden. The stable theories cannot order, the theories with the negation of the independence property cannot shatter, and the treeless theories cannot grow trees. The negative definition has nevertheless a positive content. From the absence of a configuration follows the presence of a structure. In a stable theory the structure is the definability of types. In a treeless theory it is a calculus of independence. The proofs pass from what a theory is not to what it is.

The title should not mislead. The lines are three, not one, and a theory may be tame by one measure and wild by another. The random graph is treeless and has the independence property (see Example 4.1.4 and Theorem 5.4.6). The rational order has the negation of the independence property and is not treeless (see Example 4.1.3 and Proposition 5.3.2). In Theorem 5.3.4 we show that the stable theories are tame by both measures, and that they are the only such theories.

Mathematics is full of independence: of vectors, of transcendental numbers, of events in a probability space. In each case, to say that a is independent from B is to say that B tells us nothing new about a . The notion obeys the same laws in all three cases. What is independent from a set is independent from its parts. Independence is symmetric. Every element can be moved, without change to what is known of it, into a position of independence. Whether these laws follow from a single principle is one of the subjects of this book. In Chapter 6

¹The word *treeless* has recently been given a different meaning by Kaplan, Ramsey and Simon [KRS24], for a class of theories that contains the stable theories and the binary theories; in this book it is used only in Shelah's sense.

we show that in a stable theory, independence obeys all the laws of the classical examples, including uniqueness of independent extensions. In Chapter 7 we show that in a treeless theory that is not stable, uniqueness fails and amalgamation survives (see Theorems 7.5.2 and 7.4.7), and that once the tree property appears, no relation with the laws of the treeless case exists (see Theorem 7.6.5). The dividing lines therefore separate the theories in which independence makes sense from those in which it does not. Shelah also called the theories without the independence property *dependent*, so that even the terminology takes sides.

Each line is presented twice: from within a single formula, by the configuration that the formula forbids, and from without, by the relation of independence that the theory carries. The second presentation rests on the first. That a stable formula has definable types is a theorem about one formula (see Theorem 1.1.6). That a stable theory carries a notion of independent extension with uniqueness is a theorem about all formulas at once (see Theorem 6.2.2). In the same way, the tree property of a single formula and the independence theorem of Kim and Pillay stand at the two ends of the treatment of treeless theories. The arguments between the ends are given in full.

A second theme runs alongside the first. A formula is a function with values 0 and 1. The definition of stability does not require the values to be so restricted. In Section 1.2 we show that, when the values are allowed to be real, stability becomes the statement that two iterated limits may be exchanged. The subject then meets Grothendieck's double limit criterion, Rosenthal's dichotomy and the reflexive Banach spaces. This theme is followed where it leads to a proof, and no further. Several chapters close with their real-valued counterparts, and the analytic background is collected in an appendix. The thesis of the second theme may be put briefly: the types of the logician and the limits of the analyst are the same objects under two modes of presentation, and what is proved with effort under one is sometimes proved with ease under the other (compare Corollary 1.2.13 with Chapter 2).

The reader is assumed to know the compactness theorem and not to be afraid of ultrafilters; nothing else is presupposed. Three rules have been kept throughout: to separate what belongs to a single formula from what belongs to the theory as a whole; never to give a definition without first saying what it is meant to accomplish; and never to confuse a type with the set of its realizations. The historical remarks at the end of each chapter record where the ideas came from, and the mathematical line is kept free of them. The Leitfaden that follows the table of contents shows how the chapters depend on one another, for the reader who prefers not to read them in order, which is to say, for most readers.


Contents

Introduction	5
Leitfaden	9
Chapter 1. Stable formulas, both 2-valued and real-valued	11
1.1. The 2-valued case: φ -types and type definability	12
1.2. The real-valued case	14
1.3. Exercises	21
1.4. Historical remarks	26
Chapter 2. Stability, trees, and definability of types	29
2.1. The equivalence proved in this chapter	29
2.2. From the absence of the order property to type definability	31
2.3. From the order property to tall trees	32
2.4. Tall trees produce non-definable types	33
2.5. Types over countable models: a perfect set dichotomy	35
2.6. Exercises	36
2.7. Historical remarks	41
Chapter 3. More stability, trees, and definability	43
3.1. Trees give the order property: a quantitative proof by Ramsey's theorem	43
3.2. The classical proof of definability, via ranks	46
3.3. Exercises	48
3.4. Historical remarks	53
Chapter 4. From stability to NIP	55
4.1. The independence property	56
4.2. The real-valued case	58
4.3. Exercises	60
4.4. Historical remarks	63
Chapter 5. From stability to simplicity	67
5.1. Conventions	68
5.2. Dividing and the tree property	68
5.3. Simplicity between stability and NIP	72
5.4. The Rado graph	75
5.5. Fraïssé limits	82
5.6. The real-valued case	88
5.7. Exercises	91
5.8. Historical remarks	94

Chapter 6. Stability and notions of independence	97
6.1. Definitions and heirs	98
6.2. Independence relations	102
6.3. Local character, uniqueness, and forking	104
6.4. Exercises	105
6.5. Historical remarks	110
Chapter 7. Simplicity and notions of independence	113
7.1. Dividing over models	114
7.2. Kim’s lemma	117
7.3. Forking, symmetry and transitivity	121
7.4. The independence theorem	124
7.5. Stability through forking	128
7.6. The Kim–Pillay theorem	130
7.7. Exercises	133
7.8. Historical remarks	135
Appendix A. Grothendieck’s double limit criterion	137
A.1. The criterion	137
A.2. Historical remarks	140
Bibliography	141

Leitfaden

The chapters do not form a single line. Chapter 1 is the trunk: in it we introduce stable formulas, in the 2-valued and in the real-valued form, and state the definability of types, which we prove in Chapter 2. From there the book branches three ways: toward trees and the perfect set dichotomy (Chapter 3), toward NIP (Chapter 4), and toward simplicity and independence (Chapters 5–7). In the first five chapters we study the dividing lines from inside the formulas: definability and trees for stability, shattering for NIP, dividing and the tree property for simplicity, which meet in the theorem that the stable theories are the simple NIP theories (see Theorem 5.3.4). In the last two we describe stability and simplicity globally, by the independence relations they carry: uniqueness of independent extensions for stability, amalgamation for simplicity. In the diagram, an arrow from one chapter to another means that the second uses results of the first, a heavy arrow marks heavy use, and arrows that follow from others are omitted. A reader interested in simple theories may read Chapters 1, 2, 5, 6 and 7; a reader interested in NIP may read Chapters 1 and 4, with Appendix A.



In the appendix (dashed) we collect the analytic background shared by Chapters 1 and 4. In Chapter 6 we use from Chapter 5 only the conventions on monster models (see Section 5.1) and the axioms for independence (see Definition 5.4.7); the real source of Chapter 6 is the definability of types of Chapters 1 and 2.

Chapter 1

Stable formulas, both 2-valued and real-valued

Stability is a property of a single formula $\varphi(x, y)$, and in this chapter we look at stability from two sides. On the logical side, a stable formula is one that cannot order arbitrarily long sequences of elements, and the central fact about such formulas is that all their types are definable: in Section 1.1 we set up this side and state the definability theorem, Theorem 1.1.6, which we prove in Chapter 2. On the analytic side, a formula is simply a $\{0, 1\}$ -valued function of two variables, and nothing prevents us from replacing such functions by bounded real-valued ones. Once we do, stability becomes the statement that two iterated limits can be exchanged (see Subsection 1.2.1), and in Section 1.2 we show how closely this property is tied to reflexive Banach spaces.

Roadmap

The question with which this chapter begins is an old one in a new dress: the question when one thing depends on another in a manner that can be written down. A formula $\varphi(x, y)$ makes each element depend on the parameters, since an element a may fall under $\varphi(x, b)$ for some b and not for others, and the φ -type of a is the record of these dependences. Such a record may be lawless, or it may be governed by a rule; when it is governed by a rule, we call the type definable. The answer given in Section 1.1 is that every record is governed by a rule exactly when φ cannot order long sequences. This is the definability theorem, and everything that follows depends on it.

In Section 1.2 we put the same question to bounded real-valued functions and show that its four answers are one. A function is stable when its two iterated limits may be exchanged (see Subsection 1.2.1), when it extends to a separately continuous function on pairs of types (see Subsection 1.2.2), when it is the pairing between the unit balls of a reflexive Banach space and of its dual (see Subsection 1.2.3), and when its types are uniform limits of polynomials in its instances (see Subsection 1.2.4). The last answer is the real-valued form of definability, and if the chapter has a moral, it is this: a definition is a limit that has been made explicit.

Chapters 2 and 3 do not depend on Section 1.2. A reader mainly interested in the classical theory may go directly to Chapter 2, where Theorem 1.1.6 is proved, and return to Section 1.2 later; the analytic proof of $(1) \Rightarrow (3)$ in Theorem 1.1.6 given in Section 1.2 (see Corollary 1.2.13) is perhaps best appreciated after the combinatorial one. The two sides meet again in Section 4.2, where the independence property is characterized in terms of convergent subsequences by means of Rosenthal's dichotomy, and in Chapter 6, where Harrington's lemma, the fact behind the symmetry of independence, turns out to be the exchange of iterated ultrafilter limits (see Subsection 6.1.1).

1.1 THE 2-VALUED CASE: φ -TYPES AND TYPE DEFINABILITY

Throughout this section and Chapters 2 and 3, we fix a first-order theory T in a language L and an L -formula $\varphi(x, y)$ consistent with T , where x and y are finite tuples of variables. Tuples of elements are written $a, b, \dots$ and are always of the appropriate length. We write $\varphi^1 = \varphi$ and $\varphi^0 = \neg\varphi$.

Definition 1.1.1. Let M be a model of T .

- (a) A φ -formula over M is a formula $\varphi(x, b)^u$ with $b \in M$ and $u \in \{0, 1\}$.
- (b) A complete φ -type over M (in the variables x) is a set p of φ -formulas over M such that
 - (i) the set p is consistent with the elementary diagram of M , and
 - (ii) for each $b \in M$, exactly one of $\varphi(x, b)$ and $\neg\varphi(x, b)$ belongs to p .

The set of complete φ -types over M will be denoted $S_\varphi(M)$.

If $N \succeq M$ and $c \in N$, the set

$$\text{tp}_\varphi(c/M) = \{\varphi(x, b)^u : b \in M, u \in \{0, 1\}, N \models \varphi(c, b)^u\}$$

is a complete φ -type over M . Conversely, every complete φ -type p over M is consistent with the elementary diagram of M , so it is realized by some element c of some $N \succeq M$, and then we have $p = \text{tp}_\varphi(c/M)$.

For $p \in S_\varphi(M)$ we write $D_p = \{b \in M : \varphi(x, b) \in p\}$. By condition (ii), we have $p = \{\varphi(x, b) : b \in D_p\} \cup \{\neg\varphi(x, b) : b \in M \setminus D_p\}$, so the type p is determined by D_p . The set D_p records, for each parameter b , whether $\varphi(x, b)$ holds of the elements that realize p . Each formula $\varphi(x, b)$ is, as Frege would say, unsaturated: it has an empty place, and it becomes true or false only when an object is put there. A type says, for every b , whether the objects that realize it fall under $\varphi(x, b)$.

A type over M need not be realized in M , but each finite part of it is realized there.

Lemma 1.1.2 (Finite satisfiability). *Let $p \in S_\varphi(M)$ be given. For every finite subset $\{\varphi(x, b_l)^{u_l} : l < k\}$ of p , there is $a \in M$ with $M \models \bigwedge_{l < k} \varphi(a, b_l)^{u_l}$.*

Proof. The $L(M)$ -sentence $\exists x \bigwedge_{l < k} \varphi(x, b_l)^{u_l}$ is consistent with the elementary diagram of M . That diagram is a complete $L(M)$ -theory, so the sentence belongs to the diagram, that is, holds in M . ☺

Lemma 1.1.2 gives φ -types a natural topological home. The sets $\varphi(M, b) = \{a \in M^{|x|} : M \models \varphi(a, b)\}$, for $b \in M^{|y|}$, generate a Boolean algebra $\mathcal{B}_\varphi(M)$ of subsets of $M^{|x|}$. For $p \in S_\varphi(M)$, the sets $\varphi(M, b)$ with $\varphi(x, b) \in p$ and the complements $M^{|x|} \setminus \varphi(M, b)$ with $\neg\varphi(x, b) \in p$ have the finite intersection property by Lemma 1.1.2, and they generate an ultrafilter on $\mathcal{B}_\varphi(M)$. Conversely, an ultrafilter on $\mathcal{B}_\varphi(M)$ determines a set of φ -formulas over M that is finitely satisfiable in M , hence consistent with the elementary diagram of M , and so a complete φ -type. A complete φ -type over M is therefore the same thing as an ultrafilter on $\mathcal{B}_\varphi(M)$, and we regard $S_\varphi(M)$ as the Stone space of $\mathcal{B}_\varphi(M)$: a compact, Hausdorff, zero-dimensional space whose basic clopen sets consist of the types containing a given finite set of φ -formulas over M .

Definition 1.1.3. Let $p \in S_\varphi(M)$ be given.

- (a) The type p is *definable* if there is an $L(M)$ -formula $\psi(y)$, a φ -definition of p , such that for all $b \in M$, we have

$$\varphi(x, b) \in p \iff M \models \psi(b).$$

Equivalently, the set $D_p = \psi(M)$ is a definable subset of M .

(b) A *positive Boolean combination of instances of φ over M* is a formula of the form

$$\bigvee_{i < n} \bigwedge_{a \in W_i} \varphi(a, y),$$

where $n < \omega$ and $W_0, \dots, W_{n-1}$ are finite subsets of M . By convention the empty disjunction is $\perp$ and the empty conjunction is $\top$.

In general the set D_p is just a subset of M , with no reason to be describable inside M : when M is infinite there are far more subsets of M than formulas over M , and a definable type is one whose set D_p is the extension of a single formula. Realized types are the basic example, since for $c \in M$ the type $\text{tp}_\varphi(c/M)$ is defined by the formula $\varphi(c, y)$. A definable type thus behaves like the type of an element of M , even when the type is realized only in an elementary extension.

Definability also has a practical side. A formula $\psi(y)$ makes sense in every elementary extension $N \succeq M$, so a definition of p tells us how to extend p to parameters in N : put $\varphi(x, b)$ in the extension exactly when $N \models \psi(b)$. This is the source of the canonical extensions used to define independence in Chapter 6.

Let us notice that definability typically fails because of an order: in $(\mathbb{Q}, <)$, let $\varphi(x, y)$ be $x < y$ and let p be the type of $\sqrt{2}$ over $\mathbb{Q}$. The set $D_p = \{b \in \mathbb{Q} : b > \sqrt{2}\}$ is a cut with an irrational edge, whereas the subsets of $\mathbb{Q}$ definable in $(\mathbb{Q}, <)$ with parameters are finite unions of intervals with rational or infinite endpoints; hence the type p is not definable. The type of an element just above a rational q , by contrast, is defined by $y > q$.

In $(\mathbb{Q}, <)$ the cut at $\sqrt{2}$ escapes every formula because the formula $x < y$ orders the rationals. The next definition isolates this behavior in finite pieces.

Definition 1.1.4. Let M be a model of T , and let $n \geq 1$ be given. An *n -ladder for φ in M* consists of tuples $a_1, \dots, a_n$ and $b_1, \dots, b_n$ of elements of M such that, for all $1 \leq i, j \leq n$, we have

$$M \models \varphi(a_i, b_j) \iff i \leq j.$$

We say that φ has the *order property* if for every n some model of T contains an n -ladder for φ , and that φ is *stable* otherwise.




FIGURE 1.1. A 4-ladder. Left: the elements a_i and b_j are joined when $\varphi(a_i, b_j)$ holds, that is, when $i \leq j$. The pairs (a_i, b_i) are the rungs (thick, pink), and each a_i also reaches every b_j with $j > i$ (purple). Right: the truth table of $\varphi(a_i, b_j)$, with pink cells where φ holds. The table is a staircase, which is why an n -ladder is also called a *half-graph*.

Figure 1.1 explains the name. Put the a_i on one rail and the b_j on the other, both in increasing order, and join a_i to b_j whenever $\varphi(a_i, b_j)$ holds; the “horizontal” pairs (a_i, b_i) are then the rungs of a ladder. What matters is the asymmetry between the two directions: each a_i is related to every b_j with $j \geq i$ and to none with $j < i$, so the formula φ linearly orders the rungs, in the sense that $i \leq j$ if and only if $\varphi(a_i, b_j)$.



Remark 1.1.5. An $(n+1)$ -ladder contains an n -ladder. Hence, by compactness, the formula φ has the order property if and only if some model of T contains an infinite sequence $(a_i, b_i)_{i < \omega}$ with $\varphi(a_i, b_j) \iff i \leq j$. The choice of “ $\leq$ ” is immaterial. Suppose that $a_1, \dots, a_n$ and $b_1, \dots, b_n$ are tuples in a model of T such that $\varphi(a_i, b_j) \iff i > j$ for $1 \leq i, j \leq n$. Reversing the indices by putting $a'_i = a_{n+1-i}$ and $b'_j = b_{n+1-j}$ gives $\varphi(a'_i, b'_j) \iff i < j$. Shifting by putting $b''_j = b'_{j+1}$ then gives $\varphi(a'_i, b''_j) \iff i \leq j$ for $1 \leq i, j \leq n-1$. Thus a “ $>$ ”-pattern of length n yields an $(n-1)$ -ladder.

The main theorem of this chapter confirms that the order property is the only obstruction to definability, and that when the obstruction is absent the definitions can be chosen of the simplest possible kind: positive Boolean combinations of instances of φ (see Definition 1.1.3(b)). These are, up to equivalence, all the formulas built from instances $\varphi(a, y)$ with $a \in M$ using only $\wedge$ and $\vee$, by distributivity; if φ is quantifier-free, so is every such combination.

Theorem 1.1.6 (Definability of types). *The following are equivalent:*

- (1) *The formula φ does not have the order property.*
- (2) *For every model M of T , every $p \in S_\varphi(M)$ has a φ -definition that is a positive Boolean combination of instances $\varphi(a, y)$ with $a \in M$.*
- (3) *For every model M of T , every $p \in S_\varphi(M)$ is definable.*

In short: *all φ -types are definable if and only if φ does not have the order property.* The implication (2) $\Rightarrow$ (3) is trivial, and the implication (1) $\Rightarrow$ (2) is proved in Section 2.2. The remaining implication (3) $\Rightarrow$ (1) takes a detour through binary trees: in Sections 2.3–2.4 we show that if φ has the order property, then the formula φ encodes tall trees, and that tall trees produce a φ -type that is not definable.

1.2 THE REAL-VALUED CASE

Once a formula is identified with a $\{0, 1\}$ -valued function, stability can be read as a property of bounded real-valued functions, namely the exchangeability of iterated limits. We follow Iovino [Iov99]. Seen this way, stability meets the reflexivity of Banach spaces, through a characterization of reflexivity due to James [Jam64] that looks very much like the order property.

1.2.1 Stability and iterated limit exchange

We first state the order property for arbitrary functions of two variables, and then show that its absence is exactly the exchangeability of two iterated limits.

Let M be an L -structure and $\varphi(x, y)$ an L -formula. We identify φ with the function $\varphi: M^{|x|} \times M^{|y|} \rightarrow \{0, 1\}$ that takes the value 1 at (a, b) if and only if $M \models \varphi(a, b)$. More generally, a function $\varphi: A \times B \rightarrow \{0, 1\}$ on arbitrary sets A, B has the *order property* if there

are sequences $(a_m)_{m < \omega}$ in A and $(b_n)_{n < \omega}$ in B such that

$$\varphi(a_m, b_n) = 1 \iff m \leq n.$$

Definition 1.2.1. A function $\varphi: A \times B \rightarrow \{0, 1\}$ is *unstable* if either φ or $\neg\varphi = 1 - \varphi$ has the order property, and *stable* otherwise. A formula $\varphi(x, y)$ is *stable on M* if the corresponding function on $M^{|x|} \times M^{|y|}$ is stable.



Remark 1.2.2. A formula $\varphi(x, y)$ is stable in the sense of Definition 1.1.4 if and only if φ is stable on every model of T . Indeed, if φ has the order property, then by compactness (see Remark 1.1.5) some model contains an infinite ladder, so the formula φ is unstable on that model. Conversely, if φ or $\neg\varphi$ has the order property on some model M of T , then the model M contains n -ladders for φ for every n : directly in the first case, and by Remark 1.1.5 in the second, since $\neg\varphi(a_m, b_n) \iff m \leq n$ means $\varphi(a_m, b_n) \iff m > n$.

To state the analytic form of stability, we use limits along ultrafilters. For a family $(a_i)_{i \in I}$ in a topological space X , a point $a \in X$, and an ultrafilter $\mathcal{U}$ on I , we write

$$(1.1) \quad \lim_{i, \mathcal{U}} a_i = a$$

if a is a limit of (a_i) along $\mathcal{U}$. We use freely the standard facts about such limits: limits along ultrafilters are unique in Hausdorff spaces, always exist in compact spaces, describe closures, and are preserved by continuous maps.

Along an infinite ladder, the two iterated limits of φ along a nonprincipal ultrafilter on ω already differ: for fixed m the values $\varphi(a_m, b_n)$ are eventually 1, and for fixed n the values $\varphi(a_m, b_n)$ are eventually 0, so one iterated limit is 1 and the other is 0. This is the phenomenon that the analytic form of stability forbids. Ultrafilters on arbitrary index sets are convenient, but ladders are indexed by ω . The next proposition, which we use several times, bridges the gap: from iterated limits taken along arbitrary families it extracts sequences on which the two iterated limits are already visible, one above the diagonal and the other below.

Proposition 1.2.3. Let $\varphi: A \times B \rightarrow [0, 1]$ be a function, let $(a_i)_{i \in I}$ be a family in A and $(b_j)_{j \in J}$ a family in B , and let $\mathcal{U}, \mathcal{V}$ be ultrafilters on I, J with

$$\lim_{i, \mathcal{U}} \lim_{j, \mathcal{V}} \varphi(a_i, b_j) = \alpha, \quad \lim_{j, \mathcal{V}} \lim_{i, \mathcal{U}} \varphi(a_i, b_j) = \beta.$$

Then there are $i_0, i_1, \dots \in I$ and $j_0, j_1, \dots \in J$ such that

$$|\varphi(a_{i_m}, b_{j_n}) - \alpha| \leq 2^{1-m} \text{ if } m < n, \quad |\varphi(a_{i_m}, b_{j_n}) - \beta| \leq 2^{1-n} \text{ if } n < m.$$

Proof. By definition of the limits, for every $\varepsilon > 0$ there are $U_\varepsilon \in \mathcal{U}$ and $V_\varepsilon \in \mathcal{V}$ such that

$$i \in U_\varepsilon \implies \left| \lim_{j, \mathcal{V}} \varphi(a_i, b_j) - \alpha \right| \leq \varepsilon,$$

$$j \in V_\varepsilon \implies \left| \lim_{i, \mathcal{U}} \varphi(a_i, b_j) - \beta \right| \leq \varepsilon.$$

Moreover, for all $\hat{i} \in I, \hat{j} \in J$ and $\varepsilon > 0$ there are $V_\varepsilon^{\hat{i}} \in \mathcal{V}$ and $U_\varepsilon^{\hat{j}} \in \mathcal{U}$ such that

$$j \in V_\varepsilon^{\hat{i}} \implies \left| \varphi(a_{\hat{i}}, b_j) - \lim_{j, \mathcal{V}} \varphi(a_{\hat{i}}, b_j) \right| \leq \varepsilon,$$

$$i \in U_\varepsilon^{\hat{j}} \implies \left| \varphi(a_i, b_{\hat{j}}) - \lim_{i, \mathcal{U}} \varphi(a_i, b_{\hat{j}}) \right| \leq \varepsilon.$$

Choose recursively

$$i_N \in U_{2^{-N}} \cap \bigcap_{k < N} U_{2^{-k}}^{j_k}, \quad j_N \in V_{2^{-N}} \cap \bigcap_{k \leq N} V_{2^{-k}}^{i_k}.$$

These sets are finite intersections of members of an ultrafilter, hence nonempty. If $m < n$, then we have $j_n \in V_{2^{-n}}^{i_m}$ and $i_m \in U_{2^{-m}}$, so we have $|\varphi(a_{i_m}, b_{j_n}) - \alpha| \leq 2^{-n} + 2^{-m} \leq 2^{1-m}$. If $n < m$, then we have $i_m \in U_{2^{-m}}^{j_n}$ and $j_n \in V_{2^{-n}}$, so we have $|\varphi(a_{i_m}, b_{j_n}) - \beta| \leq 2^{-m} + 2^{-n} \leq 2^{1-n}$. ☺

Figure 1.2 shows the array produced by Proposition 1.2.3.




FIGURE 1.2. The array $\varphi(a_{i_m}, b_{j_n})$ of Proposition 1.2.3. Above the diagonal the values approach α , below it they approach β , and the deepening shades show the tolerances 2^{1-m} and 2^{1-n} shrinking. Nothing is claimed on the diagonal. When $\alpha \neq \beta$ the array is a ladder in the limit: compare the staircase of Figure 1.1.

With Proposition 1.2.3 in hand, we can show that for $\{0, 1\}$ -valued functions the order property is exactly the failure of limit exchange.

Proposition 1.2.4. *Let $\varphi: A \times B \rightarrow \{0, 1\}$ be a function. The following are equivalent:*

- (1) *The function φ is stable.*
- (2) *For all sequences (a_m) in A and (b_n) in B and all ultrafilters $\mathcal{U}, \mathcal{V}$ on $\mathbb{N}$, we have*

$$\lim_{m, \mathcal{U}} \lim_{n, \mathcal{V}} \varphi(a_m, b_n) = \lim_{n, \mathcal{V}} \lim_{m, \mathcal{U}} \varphi(a_m, b_n).$$

Proof. (2) $\Rightarrow$ (1). Suppose that φ is unstable. If φ has the order property, there are sequences (a_m) in A and (b_n) in B with $\varphi(a_m, b_n) = 1 \iff m \leq n$. Let $\mathcal{U} = \mathcal{V}$ be a nonprincipal ultrafilter on $\mathbb{N}$. For fixed m we have $\varphi(a_m, b_n) = 1$ for all $n \geq m$, so the left-hand side of (2) is 1. For fixed n we have $\varphi(a_m, b_n) = 0$ for all $m > n$, so the right-hand side is 0. If instead $\neg\varphi$ has the order property, the two sides are 0 and 1.

(1) $\Rightarrow$ (2). Suppose that (2) fails: for some sequences (a_m) , (b_n) and ultrafilters $\mathcal{U}, \mathcal{V}$, the two iterated limits are $\alpha \neq \beta$. Since φ is $\{0, 1\}$ -valued, both limits lie in $\{0, 1\}$. By Proposition 1.2.3 there are sequences (a'_m) in A and (b'_n) in B with $|\varphi(a'_m, b'_n) - \alpha| \leq 2^{1-m}$ for $m < n$ and $|\varphi(a'_m, b'_n) - \beta| \leq 2^{1-n}$ for $n < m$. Since φ is $\{0, 1\}$ -valued, we have $\varphi(a'_m, b'_n) = \alpha$ for $2 \leq m < n$ and $\varphi(a'_m, b'_n) = \beta$ for $2 \leq n < m$. Put $c_m = a'_{2m+2}$ and $d_n = b'_{2n+3}$. Since $2m+2 < 2n+3 \iff m \leq n$, we get $\varphi(c_m, d_n) = \alpha$ if $m \leq n$ and $\varphi(c_m, d_n) = \beta$ if $m > n$. Thus the function φ has the order property if $\alpha = 1$, and the function $\neg\varphi$ does if $\alpha = 0$. ☺

Condition (2) of Proposition 1.2.4 makes sense for every bounded real-valued function, and we take this condition as the definition of stability in the real-valued setting.

Definition 1.2.5. A bounded function $\varphi: A \times B \rightarrow \mathbb{R}$ is *stable* if, for all sequences (a_m) in A and (b_n) in B and all ultrafilters $\mathcal{U}, \mathcal{V}$ on $\mathbb{N}$, we have

$$\lim_{m, \mathcal{U}} \lim_{n, \mathcal{V}} \varphi(a_m, b_n) = \lim_{n, \mathcal{V}} \lim_{m, \mathcal{U}} \varphi(a_m, b_n).$$

By Proposition 1.2.4, the new definition agrees with Definition 1.2.1 for $\{0, 1\}$ -valued functions, and by Proposition 1.2.3, the sequences in Definition 1.2.5 may be replaced by arbitrary families, as the next corollary records. For a bounded set of continuous functions on a compact space, the same condition is the hypothesis of Grothendieck's double limit criterion (see Appendix A).

Corollary 1.2.6. *If $\varphi: A \times B \rightarrow [0, 1]$ is stable, then we have $\lim_{i, \mathcal{U}} \lim_{j, \mathcal{V}} \varphi(a_i, b_j) = \lim_{j, \mathcal{V}} \lim_{i, \mathcal{U}} \varphi(a_i, b_j)$ for all families $(a_i)_{i \in I}$, $(b_j)_{j \in J}$ and all ultrafilters $\mathcal{U}$ on I , $\mathcal{V}$ on J .*

Proof. If the two limits are $\alpha \neq \beta$, take (i_m) , (j_n) as in Proposition 1.2.3 and a nonprincipal ultrafilter $\mathcal{W}$ on $\mathbb{N}$. For fixed m , the limit $\lim_{n, \mathcal{W}} \varphi(a_{i_m}, b_{j_n})$ is within 2^{1-m} of α , so we have $\lim_{m, \mathcal{W}} \lim_{n, \mathcal{W}} \varphi(a_{i_m}, b_{j_n}) = \alpha$. Symmetrically, the other iterated limit is β , which contradicts the stability of φ . ☹

The proof of Corollary 1.2.6 used a single nonprincipal ultrafilter $\mathcal{W}$ on $\mathbb{N}$, in both coordinates, and this is no accident. For ultrafilters $\mathcal{U}$ on I and $\mathcal{V}$ on J , let $\mathcal{U} \otimes \mathcal{V}$ be the ultrafilter on $I \times J$ consisting of the sets X with $\{i : \{j : (i, j) \in X\} \in \mathcal{V}\} \in \mathcal{U}$. For every bounded $f: I \times J \rightarrow \mathbb{R}$ we have $\lim_{(i, j), \mathcal{U} \otimes \mathcal{V}} f(i, j) = \lim_{i, \mathcal{U}} \lim_{j, \mathcal{V}} f(i, j)$, so the two iterated limits of Definition 1.2.5 are the limits of $\varphi(a_m, b_n)$ along $\mathcal{U} \otimes \mathcal{V}$ and along the transpose of $\mathcal{V} \otimes \mathcal{U}$. These two ultrafilters can differ, because the tensor product of ultrafilters is not commutative, and the simplest witness already lives on $\mathbb{N} \times \mathbb{N}$: if $\mathcal{W}$ is a nonprincipal ultrafilter on $\mathbb{N}$, then the upper triangle $\{(m, n) : m < n\}$ belongs to $\mathcal{W} \otimes \mathcal{W}$, because each of its rows is cofinite, while the lower triangle $\{(m, n) : m > n\}$ does not, because each of its rows is finite. A ladder is exactly a pair of sequences along which φ is the indicator function of the closed upper triangle $\{(m, n) : m \leq n\}$. Thus the order property says that φ can tell $\mathcal{W} \otimes \mathcal{W}$ from its own transpose, and the proofs of Proposition 1.2.4 and Corollary 1.2.6 show the following: a bounded function φ is stable if and only if, for one (equivalently, for every) nonprincipal ultrafilter $\mathcal{W}$ on $\mathbb{N}$ and all sequences (a_m) and (b_n) , we have

$$\lim_{m, \mathcal{W}} \lim_{n, \mathcal{W}} \varphi(a_m, b_n) = \lim_{n, \mathcal{W}} \lim_{m, \mathcal{W}} \varphi(a_m, b_n).$$

Stable functions are blind to the noncommutativity of the tensor square of a single ultrafilter.

1.2.2 Stability and continuous extendability

Exchanging iterated limits is a statement about continuity in disguise. In Proposition 1.2.8 below we show that a function $\varphi: A \times B \rightarrow [0, 1]$ is stable exactly when the function extends to a separately continuous function on pairs of types. The types in question live in compact spaces built from A and B .

Definition 1.2.7. Let $\varphi: A \times B \rightarrow [0, 1]$ be a function. For $a \in A$, the *left φ -type* of a is the function $\text{ltp}_\varphi(a): y \mapsto \varphi(a, y)$ on B . For $b \in B$, the *right φ -type* of b is the function $\text{rtp}_\varphi(b): x \mapsto \varphi(x, b)$ on A . The *space of left φ -types* $\text{LS}(\varphi)$ is the closure of $\{\text{ltp}_\varphi(a) : a \in A\}$ in $[0, 1]^B$ with the product topology, and the *space of right φ -types* $\text{RS}(\varphi)$ is the closure of $\{\text{rtp}_\varphi(b) : b \in B\}$ in $[0, 1]^A$.

By Tychonoff's theorem, the spaces $[0, 1]^A$ and $[0, 1]^B$ are compact, and so are their closed subspaces $\text{RS}(\varphi)$ and $\text{LS}(\varphi)$. Both spaces are quotients of Stone–Čech compactifications. The map $a \mapsto \text{ltp}_\varphi(a)$ from the discrete space A into the compact space $[0, 1]^B$ extends to a continuous map $\beta A \rightarrow [0, 1]^B$, namely $\mathcal{U} \mapsto \lim_{a, \mathcal{U}} \text{ltp}_\varphi(a)$, whose image is compact and contains the realized types, and is therefore $\text{LS}(\varphi)$. Thus every left φ -type is the limit of the realized types along an ultrafilter on A , the principal ultrafilters give the realized types, and

the space $\text{LS}(\varphi)$ is a continuous image of βA ; likewise the space $\text{RS}(\varphi)$ is a continuous image of βB .

Proposition 1.2.8. *For $\varphi: A \times B \rightarrow [0, 1]$, the following are equivalent:*

- (1) *The function φ is stable.*
- (2) *There is a separately continuous function $\hat{\varphi}: \text{LS}(\varphi) \times \text{RS}(\varphi) \rightarrow [0, 1]$ such that*

$$\hat{\varphi}(\text{ltp}_\varphi(a), \text{rtp}_\varphi(b)) = \varphi(a, b) \quad \text{for all } (a, b) \in A \times B.$$

Proof. (1) $\Rightarrow$ (2). Let $p \in \text{LS}(\varphi)$ and $q \in \text{RS}(\varphi)$ be given. There are families $(a_i)_{i \in I}$ in A , $(b_j)_{j \in J}$ in B and ultrafilters $\mathcal{U}$ on I , $\mathcal{V}$ on J with

$$p = \lim_{i, \mathcal{U}} \text{ltp}_\varphi(a_i), \quad q = \lim_{j, \mathcal{V}} \text{rtp}_\varphi(b_j).$$

By Corollary 1.2.6, we have $\lim_{i, \mathcal{U}} \lim_{j, \mathcal{V}} \varphi(a_i, b_j) = \lim_{j, \mathcal{V}} \lim_{i, \mathcal{U}} \varphi(a_i, b_j)$. Define $\hat{\varphi}(p, q)$ to be this common value. Since limits in the product topology are computed coordinatewise, this common value satisfies

$$(1.2) \quad \hat{\varphi}(p, q) = \lim_{i, \mathcal{U}} q(a_i) = \lim_{j, \mathcal{V}} p(b_j).$$

The map $\hat{\varphi}$ is well defined. Suppose also that $p = \lim_{i, \mathcal{U}'} \text{ltp}_\varphi(a'_i)$ and $q = \lim_{j, \mathcal{V}'} \text{rtp}_\varphi(b'_j)$. Then we have $\lim_{j, \mathcal{V}} \varphi(x, b_j) = q(x) = \lim_{j, \mathcal{V}'} \varphi(x, b'_j)$ for every $x \in A$, and $\lim_{i, \mathcal{U}} \varphi(a_i, y) = p(y) = \lim_{i, \mathcal{U}'} \varphi(a'_i, y)$ for every $y \in B$. Hence, using Corollary 1.2.6 for the middle step, we have

$$\lim_{i, \mathcal{U}} \lim_{j, \mathcal{V}} \varphi(a_i, b_j) = \lim_{i, \mathcal{U}} \lim_{j, \mathcal{V}'} \varphi(a_i, b'_j) = \lim_{j, \mathcal{V}'} \lim_{i, \mathcal{U}} \varphi(a_i, b'_j) = \lim_{j, \mathcal{V}'} \lim_{i, \mathcal{U}'} \varphi(a'_i, b'_j).$$

The map $\hat{\varphi}$ extends φ . Take constant families.

The map $\hat{\varphi}$ is separately continuous. Fix p , and let $q = \lim_{k, \mathcal{W}} q_k$ in $\text{RS}(\varphi)$, where $\mathcal{W}$ is an ultrafilter on a set K . Write $q_k = \lim_{j, \mathcal{V}_k} \text{rtp}_\varphi(b_{k,j})$ with j ranging over a set J_k . On $J' = \{(k, j) : k \in K, j \in J_k\}$ let $\mathcal{V}'$ be the ultrafilter of all $X \subseteq J'$ with $\{k : \{j : (k, j) \in X\} \in \mathcal{V}_k\} \in \mathcal{W}$. For every family $(z_{k,j})$ in a compact Hausdorff space we have $\lim_{(k,j), \mathcal{V}'} z_{k,j} = \lim_{k, \mathcal{W}} \lim_{j, \mathcal{V}_k} z_{k,j}$. It follows that $q = \lim_{(k,j), \mathcal{V}'} \text{rtp}_\varphi(b_{k,j})$, and by (1.2), we have

$$\hat{\varphi}(p, q) = \lim_{(k,j), \mathcal{V}'} p(b_{k,j}) = \lim_{k, \mathcal{W}} \lim_{j, \mathcal{V}_k} p(b_{k,j}) = \lim_{k, \mathcal{W}} \hat{\varphi}(p, q_k).$$

Thus the function $\hat{\varphi}(p, \cdot)$ is continuous. Continuity of $\hat{\varphi}(\cdot, q)$ is proved symmetrically.

(2) $\Rightarrow$ (1). Let (a_m) , (b_n) be sequences in A , B , let $\mathcal{U}, \mathcal{V}$ be ultrafilters on $\mathbb{N}$, and put $p = \lim_{m, \mathcal{U}} \text{ltp}_\varphi(a_m)$ and $q = \lim_{n, \mathcal{V}} \text{rtp}_\varphi(b_n)$. Then the pair (p, q) lies in $\text{LS}(\varphi) \times \text{RS}(\varphi)$, and by separate continuity

$$\lim_{m, \mathcal{U}} \lim_{n, \mathcal{V}} \varphi(a_m, b_n) = \lim_{m, \mathcal{U}} \hat{\varphi}(\text{ltp}_\varphi(a_m), q) = \hat{\varphi}(p, q) = \lim_{n, \mathcal{V}} \lim_{m, \mathcal{U}} \varphi(a_m, b_n). \quad \odot$$

1.2.3 Stability and reflexivity

The pairing between a reflexive Banach space and its dual, restricted to the product of the two unit balls with their weak topologies, is a separately continuous function on a product of compact spaces. In Theorem 1.2.9 below we state Iovino's theorem that this example is universal: the stable functions are exactly the restrictions of the pairing between the unit balls of a reflexive Banach space E and of E^* .

For a Banach space E with dual E^* , we write B_E for the closed unit ball of E and $\langle x, y \rangle$ for $y(x)$, where $x \in E$ and $y \in E^*$. Recall that E is *reflexive* if the canonical embedding of E into E^{**} is surjective.

Theorem 1.2.9 (Iovino). *For $\varphi: A \times B \rightarrow [0, 1]$, the following are equivalent:*

- (1) *The function φ is stable.*
- (2) *There are a reflexive Banach space E and maps $u: A \rightarrow B_E$ and $v: B \rightarrow B_{E^*}$ such that*

$$\varphi(x, y) = \langle u(x), v(y) \rangle \quad \text{for all } (x, y) \in A \times B.$$

Proof. See [Iov99, Theorem 3.3]. For (1) $\Rightarrow$ (2), stability is used to show that the operator $T: \ell^\infty(A)^* \rightarrow \ell^\infty(B)$ given by $T\mu(y) = \int \varphi(x, y) d\mu(x)$ is weakly compact; this step goes back to Krivine and Maurey [KM81] and Raynaud [Ray83]. The theorem of Davis, Figiel, Johnson and Pełczyński [DFJP74] then factors T through a reflexive space. The implication (2) $\Rightarrow$ (1) follows from the weak compactness of B_E and the weak* compactness of B_{E^*} . $\odot$

Read back into model theory, the theorem says the following: by Remark 1.2.2 and Proposition 1.2.4, a formula $\varphi(x, y)$ is stable in the sense of Definition 1.1.4 if and only if, in every model M of T , the function φ on M is the restriction of the pairing map $B_E \times B_{E^*} \rightarrow \mathbb{R}$ for some reflexive Banach space E (depending on M), composed with maps $M^{|x|} \rightarrow B_E$ and $M^{|y|} \rightarrow B_{E^*}$.

1.2.4 Stability and polynomial approximation

By Theorem 1.1.6, if φ does not have the order property, every φ -type over a model of T is defined by a Boolean combination of instances $\varphi(a, y)$. On $\{0, 1\}$ the Boolean connectives are polynomials, namely $\neg t = 1 - t$ and $t \wedge s = ts$. The real-valued analogue of a definition is therefore a *uniform limit of polynomials in instances of φ* , and in Theorem 1.2.11 below we use the Stone–Weierstrass theorem to show that stability is exactly what is needed for such approximations to exist. Continuous logic treats definability in the same way [BYU10], and Ben Yaacov [BY14] used polynomial approximation to derive definability of types from Grothendieck’s double limit criterion (see the historical remarks at the end of this chapter, and Appendix A). We use the Stone–Weierstrass theorem in the following form.

Fact 1.2.10 (Stone–Weierstrass). *Let K be a compact Hausdorff space and let $\mathcal{F} \subseteq C(K)$ be a family of continuous real-valued functions that separates the points of K . Then every $f \in C(K)$ is a uniform limit on K of functions of the form $P(g_1, \dots, g_k)$, where $g_1, \dots, g_k \in \mathcal{F}$ and P is a real polynomial in k variables.*

To apply the Stone–Weierstrass theorem we need continuous functions on the type spaces that separate points, and the instances of φ supply them. Fix $\varphi: A \times B \rightarrow [0, 1]$. For $a \in A$, let $e_a: \text{RS}(\varphi) \rightarrow [0, 1]$ be the evaluation map $e_a(q) = q(a)$, and for $b \in B$, let $e_b: \text{LS}(\varphi) \rightarrow [0, 1]$ be $e_b(p) = p(b)$. These maps are continuous, because $\text{RS}(\varphi)$ and $\text{LS}(\varphi)$ carry the product topology, and on realized types they are instances of φ , since $e_a(\text{rtp}_\varphi(b)) = \varphi(a, b) = e_b(\text{ltp}_\varphi(a))$. The family $\{e_a : a \in A\}$ separates the points of $\text{RS}(\varphi)$, since two distinct functions $q \neq q'$ in $[0, 1]^A$ differ at some $a \in A$, and likewise the family $\{e_b : b \in B\}$ separates the points of $\text{LS}(\varphi)$.

Theorem 1.2.11. *For $\varphi: A \times B \rightarrow [0, 1]$, the following are equivalent:*

- (1) *The function φ is stable.*
- (2) *There is a function $\hat{\varphi}: \text{LS}(\varphi) \times \text{RS}(\varphi) \rightarrow [0, 1]$ with $\hat{\varphi}(\text{ltp}_\varphi(a), \text{rtp}_\varphi(b)) = \varphi(a, b)$ for all $(a, b) \in A \times B$, such that for every $\varepsilon > 0$:*
 - (a) *for every $p \in \text{LS}(\varphi)$ there are $a_1, \dots, a_k \in A$ and a real polynomial P with*

$$|\hat{\varphi}(p, q) - P(q(a_1), \dots, q(a_k))| \leq \varepsilon \quad \text{for all } q \in \text{RS}(\varphi);$$

(b) for every $q \in \text{RS}(\varphi)$ there are $b_1, \dots, b_l \in B$ and a real polynomial Q with

$$|\hat{\varphi}(p, q) - Q(p(b_1), \dots, p(b_l))| \leq \varepsilon \quad \text{for all } p \in \text{LS}(\varphi).$$

Proof. (1) $\Rightarrow$ (2). Let $\hat{\varphi}$ be the separately continuous function given by Proposition 1.2.8. For fixed p , the function $\hat{\varphi}(p, \cdot)$ is continuous on the compact Hausdorff space $\text{RS}(\varphi)$. Applying Fact 1.2.10 with $\mathcal{F} = \{e_a : a \in A\}$ gives (a). Clause (b) follows in the same way, using the continuity of $\hat{\varphi}(\cdot, q)$ on $\text{LS}(\varphi)$ and $\mathcal{F} = \{e_b : b \in B\}$.

(2) $\Rightarrow$ (1). Each function $q \mapsto P(q(a_1), \dots, q(a_k))$ is continuous on $\text{RS}(\varphi)$. By (a), the function $\hat{\varphi}(p, \cdot)$ is a uniform limit of such functions, hence continuous. By (b), the function $\hat{\varphi}(\cdot, q)$ is continuous for the same reason. Hence the function $\hat{\varphi}$ is separately continuous, and the function φ is stable by Proposition 1.2.8. $\odot$

Evaluated at realized right types, clause (a) says that left types are approximated by polynomials in instances. Suppose that φ is stable, and let $\hat{\varphi}$ be the separately continuous extension given by Proposition 1.2.8, as in the proof of (1) $\Rightarrow$ (2). Then we have $\hat{\varphi}(p, \text{rtp}_\varphi(b)) = p(b)$ for $p \in \text{LS}(\varphi)$ and $b \in B$: if $p = \lim_{i, \mathcal{U}} \text{ltp}_\varphi(a_i)$, continuity in the first variable gives $\hat{\varphi}(p, \text{rtp}_\varphi(b)) = \lim_{i, \mathcal{U}} \varphi(a_i, b) = p(b)$. Taking $q = \text{rtp}_\varphi(b)$ in (a) therefore gives

$$|p(b) - P(\varphi(a_1, b), \dots, \varphi(a_k, b))| \leq \varepsilon \quad \text{for all } b \in B.$$

Thus every left φ -type is a uniform limit on B of polynomials in instances $\varphi(a, y)$ with $a \in A$, and symmetrically for right types.

The approximation in Theorem 1.2.11 is required on all of $\text{RS}(\varphi)$ (respectively $\text{LS}(\varphi)$), not only at realized types, and this requirement cannot be weakened: the next example shows that approximation on B alone does not imply stability, since there every type is trivially approximable on B and yet the function has the order property.

Example 1.2.12. Let $A = B = \mathbb{N}$, and define $\varphi(m, n) = 1$ if $m \leq n$ and $\varphi(m, n) = 0$ otherwise. Then the function φ has the order property and is therefore unstable. Here the left type $\text{ltp}_\varphi(m)$ is the indicator function of $[m, \infty)$, and the space $\text{RS}(\varphi)$ consists of the indicator functions of the intervals $[0, n]$ together with the constant function **1**. Similarly, the space $\text{LS}(\varphi)$ consists of the functions $\text{ltp}_\varphi(m)$ together with the constant function **0**. Thus every $p \in \text{LS}(\varphi)$, viewed as a function on B , is either an instance $\varphi(m, \cdot)$ or the constant 0, which is trivially a polynomial in instances. Symmetrically, every $q \in \text{RS}(\varphi)$ is an instance or the constant 1. The failure appears at the unrealized pair $(\mathbf{0}, \mathbf{1})$. If $\hat{\varphi}$ were separately continuous and extended φ as in Theorem 1.2.11, then we would have $\hat{\varphi}(\mathbf{0}, \text{rtp}_\varphi(n)) = \mathbf{0}(n) = 0$ for all n , and since $\text{rtp}_\varphi(n) \rightarrow \mathbf{1}$, we would get $\hat{\varphi}(\mathbf{0}, \mathbf{1}) = 0$. On the other hand, we would also have $\hat{\varphi}(\text{ltp}_\varphi(m), \mathbf{1}) = \mathbf{1}(m) = 1$ for all m , and since $\text{ltp}_\varphi(m) \rightarrow \mathbf{0}$, we would get $\hat{\varphi}(\mathbf{0}, \mathbf{1}) = 1$.

We now return to formulas, whose left types are their φ -types in another guise. Let M be a model of T , and identify each $p \in S_\varphi(M)$ with the characteristic function of D_p on $M^{|y|}$. With this identification we have $\text{LS}(\varphi) = S_\varphi(M)$. Indeed, for $a \in M$ the left type $\text{ltp}_\varphi(a)$ is $\text{tp}_\varphi(a/M)$, and by Lemma 1.1.2, every $p \in S_\varphi(M)$ agrees on each finite set of parameters with some realized type, so the type p is a pointwise limit of realized types. Conversely, a pointwise limit of realized types determines a set of φ -formulas over M that is finitely satisfiable in M , hence consistent with the elementary diagram of M , and so a complete φ -type.

In this two-valued setting, the quotient map $\beta M^{|x|} \rightarrow S_\varphi(M)$ is transparent. Recall from Section 1.1 that $S_\varphi(M)$ is the Stone space of the Boolean algebra $\mathcal{B}_\varphi(M)$ generated by the

sets $\varphi(M, b)$ with $b \in M^{|y|}$. An ultrafilter $\mathcal{U}$ on $M^{|x|}$ goes to its trace on $\mathcal{B}_\varphi(M)$, that is, to the type containing $\varphi(x, b)$ exactly when $\varphi(M, b) \in \mathcal{U}$. The realized types are the traces of principal ultrafilters, and Lemma 1.1.2 says that every type is the trace of some ultrafilter.

For $\{0, 1\}$ -valued types, an approximation within $\frac{1}{3}$ can be rounded, and polynomials in instances become Boolean combinations of instances.

Corollary 1.2.13. *Let M be an L -structure and let $\varphi(x, y)$ be a formula that is stable on M . Then every $p \in \text{LS}(\varphi)$ is defined by a Boolean combination of instances: there are $a_1, \dots, a_k \in M^{|x|}$ and a Boolean combination $\psi(y)$ of $\varphi(a_1, y), \dots, \varphi(a_k, y)$ such that, for all $b \in M^{|y|}$, we have*

$$p(b) = 1 \iff M \models \psi(b).$$

Proof. Apply Theorem 1.2.11 with $\varepsilon = \frac{1}{3}$ to get $a_1, \dots, a_k$ and P . Both $p(b)$ and $\varphi(a_i, b)$ lie in $\{0, 1\}$, so we have $p(b) = 1$ if and only if the tuple $t = (\varphi(a_1, b), \dots, \varphi(a_k, b))$ belongs to $S = \{t \in \{0, 1\}^k : P(t) > \frac{1}{2}\}$. Take $\psi(y) = \bigvee_{t \in S} \bigwedge_{i \leq k} \varphi(a_i, y)^{t_i}$ (recall that $\varphi^1 = \varphi$ and $\varphi^0 = \neg\varphi$). ☺

Since a formula without the order property is stable on every model of T (see Remark 1.2.2), Corollary 1.2.13 already gives (1) $\Rightarrow$ (3) of Theorem 1.1.6, by analytic means. The price is positivity: the definitions obtained this way are Boolean combinations of instances rather than positive ones, and their size depends on p .



Remark 1.2.14(1) *Lattice version.* The lattice form of the Stone–Weierstrass theorem replaces polynomials by the operations max, min and affine maps. Continuous logic uses this form to show that a small set of connectives generates all continuous connectives up to uniform approximation [BYU10, Proposition 1.4].

- (2) *Convex combinations.* Stability gives more than a polynomial approximation. In Proposition A.1.7 below we show, using Grothendieck’s criterion, that stability of φ implies that $\{e_a : a \in A\}$ is relatively weakly compact in $C(\text{RS}(\varphi))$. In that case the equality $p = \lim_{i, \mathcal{U}} \text{ltp}_\varphi(a_i)$ implies that $\hat{\varphi}(p, \cdot)$ is the weak limit of the functions e_{a_i} along $\mathcal{U}$. Mazur’s lemma [Maz33] then shows that $\hat{\varphi}(p, \cdot)$ is a uniform limit of convex combinations $\sum_j \lambda_j e_{a_{i_j}}$. Thus the left type p is a uniform limit on B of averages of instances $\varphi(a, y)$. Ben Yaacov [BY14] uses this convex approximation to recover the familiar forms of the defining formulas.

1.3 EXERCISES

The exercises follow the sections of the chapter. Throughout, the letter θ denotes the *damped ladder*, the function $\theta: \mathbb{N} \times \mathbb{N} \rightarrow [0, 1]$ given by $\theta(m, n) = 2^{-m-1}$ if $m \leq n$ and $\theta(m, n) = 0$ if $m > n$. The mark (*) indicates the exercises that ask for more than a direct application of the text.

- (Section 1.1.) Let M be an infinite model of a theory T , and let $\varphi(x, y)$ be the formula $x = y$, where x and y are single variables. Show that, for every $p \in S_\varphi(M)$, the set D_p is either a singleton $\{a\}$ with $a \in M$ or empty, and that each of these sets is D_p for exactly one $p \in S_\varphi(M)$. For each type, find a φ -definition that is a positive Boolean combination of instances of φ (see Definition 1.1.3). Let p_∞ be the type with $D_{p_\infty} = \emptyset$. Show that p_∞ is the only point of $S_\varphi(M)$ that is not isolated, so that the space $S_\varphi(M)$ is the one-point compactification of the discrete space M . What changes when M is finite?

Next, let T_E be the theory of an equivalence relation E with infinitely many classes, all of them infinite, let M be a model of T_E , and let $\varphi(x, y)$ be the formula xEy . Show that, for every $p \in S_\varphi(M)$, the set D_p is either an E -class of M or empty, that every E -class of M and the empty set occur in this way, and that the space $S_\varphi(M)$ is the one-point compactification of the discrete set of E -classes of M . (Hint: if b and b' belong to D_p , then the formula $xEb \wedge xEb'$ is consistent with the elementary diagram of M , and that diagram contains transitivity. A set of φ -formulas over M that is finitely satisfiable in M is consistent with the elementary diagram of M , by compactness.)

2. (Section 1.1.) Let K be a finite field, and let T be the theory of infinite vector spaces over K , in the language with $+$, $-$, 0 and a unary function symbol for multiplication by each $\lambda \in K$. Let $n \geq 1$ be given, and let $\varphi(x_1, \dots, x_n; y)$ be the formula

$$\bigvee_{(\lambda_1, \dots, \lambda_n) \in K^n} y = \lambda_1 x_1 + \dots + \lambda_n x_n,$$

which says that y lies in the span of $x_1, \dots, x_n$. The object variables are $x = (x_1, \dots, x_n)$, and the parameter variable is y . Let M be a model of T . Show that, for every $p \in S_\varphi(M)$, the set D_p is a subspace of M of dimension at most n . Show that every subspace of M of dimension at most n is D_p for some type p realized in M , and that every $p \in S_\varphi(M)$ is defined by a single instance $\varphi(d, y)$ with $d \in M^n$. Conclude that every complete φ -type over M is realized in M . Let $N \succeq M$ be given, and let $c \in N^n$ be a tuple whose span meets M only in 0 . Show that the tuples c and $(0, \dots, 0)$ have the same φ -type over M , and that they have different complete types over M when $c \neq (0, \dots, 0)$. (Hint: if $c \in N^n$ realizes p , then we have $D_p = M \cap \text{span}(c)$.)

3. (Section 1.1.) Work in $(\mathbb{Q}, <)$, with $\varphi(x, y)$ the formula $x < y$.
- (a) Show that, for every $p \in S_\varphi(\mathbb{Q})$, the set D_p is a final segment of $\mathbb{Q}$, that is, if $b \in D_p$ and $b < b'$, then we have $b' \in D_p$. Show that every final segment of $\mathbb{Q}$ is D_p for exactly one $p \in S_\varphi(\mathbb{Q})$.
 - (b) Show that the final segments of $\mathbb{Q}$ are $\emptyset$, $\mathbb{Q}$, the sets $\{b : b \geq q\}$ and $\{b : b > q\}$ with $q \in \mathbb{Q}$, and the sets $\{b : b > r\}$ with r irrational. Conclude that the set $S_\varphi(\mathbb{Q})$ has cardinality $2^{\aleph_0}$.
 - (c) Let $A \subseteq \mathbb{Q}$ be finite. Show that every subset of $\mathbb{Q}$ definable in $(\mathbb{Q}, <)$ with parameters from A is a union of singletons $\{a\}$ with $a \in A$ and of open intervals whose endpoints are consecutive elements of $A \cup \{-\infty, +\infty\}$. (Hint: such a set is invariant under the automorphisms of $(\mathbb{Q}, <)$ that fix A pointwise. If c and d lie in the same such interval, a piecewise linear automorphism fixing A pointwise maps c to d .)
 - (d) Deduce that the types p with D_p equal to $\emptyset$, $\mathbb{Q}$, $\{b : b \geq q\}$ or $\{b : b > q\}$ are definable, and give their definitions. Deduce also that the types with $D_p = \{b : b > r\}$ and r irrational are not definable.
 - (e) Let $q \in \mathbb{Q}$ be given. Show that the type p with $D_p = \{b : b \geq q\}$ has no φ -definition that is a positive Boolean combination of instances $\varphi(a, y)$ with $a \in \mathbb{Q}$. (Hint: compute the set defined by $\bigvee_{i < n} \bigwedge_{a \in W_i} a < y$.)
 - (f) Show directly that φ has the order property, with $a_i = i$ and $b_j = j + \frac{1}{2}$. Explain why (d), (e) and (f) agree with Theorem 1.1.6.

4. (*) (Section 1.1.) Show that no structure contains a 2-ladder for the formula $x = y$, and that no model of T_E (see Exercise 1) contains a 2-ladder for the formula xEy . Compare with Exercise 1 and Theorem 1.1.6. Next, in the setting of Exercise 2, show that every model of T contains an $(n+1)$ -ladder for φ and that no model of T contains an $(n+2)$ -ladder for φ . (Hint: given an m -ladder $a_1, \dots, a_m, b_1, \dots, b_m$, let V_i be the span of the entries of a_i , and put $W_j = V_1 \cap \dots \cap V_j$. Show that W_{j+1} is a proper subspace of W_j for $j < m$. For the $(n+1)$ -ladder, let $e_1, \dots, e_n$ be linearly independent, let $a_i = (e_i, \dots, e_n, 0, \dots, 0)$ for $i \leq n$, and let $a_{n+1} = (0, \dots, 0)$, $b_j = e_j$ for $j \leq n$ and $b_{n+1} = 0$.)
5. (Section 1.1.) Let φ^{op} be the formula $\varphi(x, y)$ read with y as the tuple of object variables and x as the tuple of parameter variables. Thus an n -ladder for φ^{op} consists of tuples $c_1, \dots, c_n$ of the length of y and $d_1, \dots, d_n$ of the length of x such that $\varphi(d_j, c_i)$ holds if and only if $i \leq j$. Show that the following are equivalent: the formula φ has the order property; the formula $\neg\varphi$ has the order property; the formula φ^{op} has the order property. (Hint: for $\neg\varphi$, use Remark 1.1.5. For φ^{op} , reverse the indices of an n -ladder for φ .) Conclude from Exercise 3 that, in $(\mathbb{Q}, <)$, the formula $x \leq y$ has the order property.
6. (Section 1.1.) Assume that φ is stable, and let M be a model of T . Show that $|S_\varphi(M)| \leq |M| + \aleph_0$. (Hint: by Theorem 1.1.6, each $p \in S_\varphi(M)$ has a definition that is a positive Boolean combination of instances of φ over M . Such a combination is determined by a finite set of finite subsets of $M^{|x|}$, and the set D_p determines p .) Deduce from Exercise 3 a second proof that the formula $x < y$ has the order property in the theory of $(\mathbb{Q}, <)$. Compare with Exercise 1, where the number of φ -types over M is one more than the number of E -classes of M .
7. (Subsection 1.2.1.) Let $\mathcal{U}$ be a nonprincipal ultrafilter on $\mathbb{N}$.
- Let $\varphi: \mathbb{N} \times \mathbb{N} \rightarrow \{0, 1\}$ be given by $\varphi(m, n) = 1$ if $m < n$ and $\varphi(m, n) = 0$ otherwise. Compute $\lim_{m, \mathcal{U}} \lim_{n, \mathcal{U}} \varphi(m, n)$ and $\lim_{n, \mathcal{U}} \lim_{m, \mathcal{U}} \varphi(m, n)$. Find sequences (a_m) and (b_n) in $\mathbb{N}$ with $\varphi(a_m, b_n) = 1$ if and only if $m \leq n$.
 - Let $\varphi: A \times B \rightarrow \mathbb{R}$ be a bounded function, let (a_m) and (b_n) be sequences in A and B , and let $\mathcal{V}$ and $\mathcal{W}$ be ultrafilters on $\mathbb{N}$. Show that, if $\mathcal{V}$ or $\mathcal{W}$ is principal, then the two iterated limits $\lim_{m, \mathcal{V}} \lim_{n, \mathcal{W}} \varphi(a_m, b_n)$ and $\lim_{n, \mathcal{W}} \lim_{m, \mathcal{V}} \varphi(a_m, b_n)$ are equal. Thus only pairs of nonprincipal ultrafilters can witness instability.
 - Show that the function $(x, y) \mapsto xy$ on $[0, 1] \times [0, 1]$ is stable. (Hint: given sequences (a_m) and (b_n) in $[0, 1]$ and ultrafilters $\mathcal{V}$ and $\mathcal{W}$ on $\mathbb{N}$, compute both iterated limits in terms of $\lim_{m, \mathcal{V}} a_m$ and $\lim_{n, \mathcal{W}} b_n$.)
 - Let $\delta: [0, 1] \times [0, 1] \rightarrow \{0, 1\}$ be given by $\delta(x, y) = 1$ if $x = y$ and $\delta(x, y) = 0$ otherwise. Show that δ is stable, although no function $\delta(\cdot, y)$ is continuous on $[0, 1]$. (Hint: regard $[0, 1]$ as a model of the theory of infinite sets in the empty language, and use Exercise 4 and Remark 1.2.2.)
8. (Subsection 1.2.1.)
- Let $\varphi, \psi: A \times B \rightarrow \mathbb{R}$ be bounded stable functions, and let $c \in \mathbb{R}$ be given. Show that the functions $\varphi + \psi$, $c\varphi$ and $\varphi\psi$ are stable, and that every constant function is stable. (Hint: limits along an ultrafilter commute with sums and products of bounded real sequences.)

- (b) Let $(\varphi_k)_{k < \omega}$ be a sequence of stable functions that converges uniformly on $A \times B$ to a function φ . Show that φ is stable.
- (c) Let φ be a stable function with values in an interval $[c, d]$, and let $g: [c, d] \rightarrow \mathbb{R}$ be continuous. Show that $g \circ \varphi$ is stable. (Hint: apply Fact 1.2.10 with $K = [c, d]$ and $\mathcal{F}$ consisting of the identity.)
- (d) Let $\varphi(x, y)$ and $\psi(x, y)$ be stable formulas. Show that $\neg\varphi$, $\varphi \wedge \psi$ and $\varphi \vee \psi$ are stable, and hence that every Boolean combination of φ and ψ is stable. (Hint: on a model of T , the function of $\varphi \wedge \psi$ is the product of the functions of φ and ψ , and the function of $\neg\varphi$ is $1 - \varphi$. Use Proposition 1.2.4 and Remark 1.2.2.) Explain how to treat a formula $\psi(x, z)$ whose parameter variables differ from those of φ .
9. (*) (Subsection 1.2.1.) Let $\varphi: A \times B \rightarrow \mathbb{R}$ be a bounded function. Show that φ is unstable if and only if there are sequences (a_m) in A and (b_n) in B and real numbers $r \neq s$ such that, with $\delta = |r - s|/3$, we have $|\varphi(a_m, b_n) - r| \leq \delta$ whenever $m \leq n$ and $|\varphi(a_m, b_n) - s| \leq \delta$ whenever $m > n$. (Hint: for the forward direction, rescale φ to take values in $[0, 1]$, apply Proposition 1.2.3, and pass to subsequences as in the proof of Proposition 1.2.4.) Check that, for $\{0, 1\}$ -valued functions, the condition says that φ or $\neg\varphi$ has the order property. Next, show that the damped ladder satisfies $\theta(m, n) > 0$ if and only if $m \leq n$, and that both iterated limits of $\theta(m, n)$ along a nonprincipal ultrafilter on $\mathbb{N}$ are 0. Thus, in the real-valued case, a ladder witnesses instability only when its values are separated by a margin. (The function θ is in fact stable; see Exercises 10 and 13.)
10. (Subsection 1.2.2.) Let K and L be compact Hausdorff spaces, and let $f: K \times L \rightarrow \mathbb{R}$ be a bounded function that is separately continuous. Show that f is stable, as a function on the sets K and L . (Hint: given sequences (a_m) in K and (b_n) in L and ultrafilters $\mathcal{U}$ and $\mathcal{V}$ on $\mathbb{N}$, put $a = \lim_{m, \mathcal{U}} a_m$ and $b = \lim_{n, \mathcal{V}} b_n$, and show that both iterated limits equal $f(a, b)$.) Apply this to the following functions.
- (a) The function f on $[0, 1] \times [0, 1]$ given by $f(x, y) = 2xy/(x^2 + y^2)$ for $(x, y) \neq (0, 0)$ and $f(0, 0) = 0$. Show that f is stable and takes values in $[0, 1]$, but is not continuous at $(0, 0)$.
- (b) The damped ladder θ . Let $\mathbb{N} \cup \{\infty\}$ be the one-point compactification of $\mathbb{N}$, and put $\theta(m, \infty) = 2^{-m-1}$, $\theta(\infty, n) = 0$ and $\theta(\infty, \infty) = 0$. Show that the extended function is separately continuous, and conclude that θ is stable.
- The function δ of Exercise 7(d) is stable, but it is not separately continuous on $[0, 1] \times [0, 1]$ with the usual topology. Explain why this does not contradict Proposition 1.2.8.
11. (Subsection 1.2.2.) Let A be an infinite set, and let $\varphi: A \times A \rightarrow \{0, 1\}$ be given by $\varphi(a, b) = 1$ if $a = b$ and $\varphi(a, b) = 0$ otherwise. This function is stable by the argument of Exercise 7(d). For $a \in A$, let $\mathbf{1}_a$ be the indicator function of $\{a\}$.
- (a) Show that $\text{LS}(\varphi)$ consists of the functions $\mathbf{1}_a$ with $a \in A$ together with the constant function $\mathbf{0}$, that each point $\mathbf{1}_a$ is isolated in $\text{LS}(\varphi)$, and that every neighborhood of $\mathbf{0}$ contains all but finitely many of the points $\mathbf{1}_a$. Show that the same holds for $\text{RS}(\varphi)$. Compare $\text{LS}(\varphi)$ with the space $S_\varphi(M)$ of Exercise 1.
- (b) Let $\hat{\varphi}$ be the function given by Proposition 1.2.8. Using (1.2), show that $\hat{\varphi}(p, q) = 1$ if $p = q = \mathbf{1}_a$ for some $a \in A$, and $\hat{\varphi}(p, q) = 0$ otherwise.

- (c) Show that the function $\hat{\varphi}$ is not continuous at $(\mathbf{0}, \mathbf{0})$ as a function on $\text{LS}(\varphi) \times \text{RS}(\varphi)$. (Hint: take the limit of the pairs $(\mathbf{1}_a, \mathbf{1}_a)$ along a nonprincipal ultrafilter on A .) Thus the extension in Proposition 1.2.8 is separately continuous, and in general no more.
- 12.** (*Subsection 1.2.3.*) The spaces in this exercise are real, and e_k denotes the k -th unit vector.
- (a) Show that the function $(x, y) \mapsto \langle x, y \rangle$ on $B_{\ell^2} \times B_{\ell^2}$ is stable. (Hint: the ball B_{ℓ^2} is compact and Hausdorff in the weak topology, and the inner product is separately continuous for it. Use Exercise 10.)
- (b) Put $y_n = (n+1)^{-1/2} \sum_{k \leq n} e_k$. Show that $y_n \in B_{\ell^2}$, that $\langle e_m, y_n \rangle > 0$ if and only if $m \leq n$, and that both iterated limits of $\langle e_m, y_n \rangle$ along a nonprincipal ultrafilter on $\mathbb{N}$ are 0. Compare with Exercise 9.
- (c) In ℓ^1 , whose dual is ℓ^∞ , put $f_n = \sum_{k \leq n} e_k \in B_{\ell^\infty}$. Let $A = \{e_m : m < \omega\}$ and $B = \{f_n : n < \omega\}$. Show that the pairing restricted to $A \times B$ takes values in $\{0, 1\}$ and has the order property. Deduce from Theorem 1.2.9 that ℓ^1 is not reflexive.
- (d) In c_0 , whose dual is ℓ^1 , put $s_m = \sum_{k \leq m} e_k \in B_{c_0}$. Compute both iterated limits of $\langle s_m, e_n \rangle$ along a nonprincipal ultrafilter on $\mathbb{N}$, and deduce that c_0 is not reflexive.
- 13.** (*Subsection 1.2.3.*) For each of the following functions $\varphi: A \times B \rightarrow [0, 1]$, find a reflexive Banach space X and maps $u: A \rightarrow B_X$ and $v: B \rightarrow B_{X^*}$ with $\varphi(x, y) = \langle u(x), v(y) \rangle$, as in Theorem 1.2.9 (see condition (2) there, where the space is called E).
- (a) The function of Exercise 11 on an infinite set A . (Hint: take $X = \ell^2(A)$.)
- (b) The function of the formula xEy on a model M of T_E (see Exercise 1). (Hint: let C be the set of E -classes of M , take $X = \ell^2(C)$, and send each element of M to the unit vector of its class.)
- (c) The damped ladder θ . (Hint: take $X = \ell^2$, $u(m) = e_m$ and $v(n) = \sum_{k \leq n} 2^{-k-1} e_k$, and check that $\|v(n)\| \leq 1$.) This gives a second proof that θ is stable.
- Let $\lambda: \mathbb{N} \times \mathbb{N} \rightarrow \{0, 1\}$ be given by $\lambda(m, n) = 1$ if $m \leq n$ and $\lambda(m, n) = 0$ otherwise. Show that the choice $u(m) = e_m$ and $v(n) = \sum_{k \leq n} e_k$ in ℓ^2 gives $\lambda(m, n) = \langle u(m), v(n) \rangle$, but that $v(n)$ leaves the unit ball. Explain, using Theorem 1.2.9, why there are no reflexive space X and maps u and v into B_X and B_{X^*} (or, after rescaling, into any bounded subsets of X and X^*) with $\lambda(m, n) = \langle u(m), v(n) \rangle$, and compare with Exercise 12(c).
- 14.** (*Subsection 1.2.4.*) This exercise computes the types of the damped ladder θ , which is stable by Exercise 10.
- (a) Show that $\text{LS}(\theta)$ consists of the functions $\text{lt}_\theta(m)$ with $m \in \mathbb{N}$ together with the constant function $\mathbf{0}$. Show that $\text{RS}(\theta)$ consists of the functions $\text{rt}_\theta(n)$ with $n \in \mathbb{N}$ together with the function q_∞ given by $q_\infty(m) = 2^{-m-1}$, and that q_∞ is not a realized type.
- (b) Let $\hat{\theta}$ be the function given by Proposition 1.2.8. Using (1.2), show that $\hat{\theta}(\text{lt}_\theta(m), q) = q(m)$, $\hat{\theta}(\mathbf{0}, q) = 0$ and $\hat{\theta}(p, \text{rt}_\theta(n)) = p(n)$ for all m, n, p and q .
- (c) Let $k \in \mathbb{N}$ be given. Show that $|\hat{\theta}(p, q_\infty) - p(k)| \leq 2^{-k-2}$ for all $p \in \text{LS}(\theta)$. This is clause (b) of Theorem 1.2.11 for q_∞ , with $l = 1$, $b_1 = k$ and $Q(t) = t$. Evaluated at realized types, it says that $|q_\infty(m) - \theta(m, k)| \leq 2^{-k-2}$ for all m : the type q_∞ is the


uniform limit of the instances $\theta(x, k)$ as $k \rightarrow \infty$. Compare q_∞ with the norm limit of the vectors $v(n)$ of Exercise 13(c).

- (d) Show that no exact definition exists: there are no $b_1, \dots, b_l \in \mathbb{N}$ and no polynomial Q with $\hat{\theta}(p, q_\infty) = Q(p(b_1), \dots, p(b_l))$ for all $p \in \text{LS}(\theta)$. (Hint: evaluate at $p = \text{ltp}_\theta(m)$ with $m > \max_i b_i$.) Contrast this with Corollary 1.2.13, where rounding turns approximate definitions of $\{0, 1\}$ -valued types into exact ones.

1.4 HISTORICAL REMARKS

✂ **Section 1.1.** Counting types begins with Morley’s proof of Loś’s conjecture [Mor65]. Morley introduced the totally transcendental theories, defined by a condition on the topology of their type spaces. For countable theories he showed this condition to be equivalent to a condition on the number of types: over a countable set there are only countably many. This is the condition now called ω -stability. The general notion followed. Setting out to extend Morley’s categoricity theorem to uncountable theories, Rowbottom generalized ω -stability to λ -stability: a theory is λ -stable if there are at most λ types over sets of size λ [Row64]. Shelah then called a theory *stable* when λ -stable for some λ , and began the study of stable theories in general in [She69]. In [She71] he moved from theories to single formulas, defining when a formula $\varphi(x; y)$ is *stable in λ* : the formula φ has at most λ φ -types over sets of size λ [She71, Definition 2.1]. He also defined the *order property* [She71, Definition 2.2] and proved that stability is equivalent to the absence of the order property. The ladder formulation of Definition 1.1.4 is a finitary restatement of Shelah’s condition. Definable types were introduced by Gaifman in his study of end extensions of models of Peano arithmetic [Gai76], and Shelah arrived at them independently. The theorem that φ -types are definable when φ is stable is [She71, Theorem 3.1]; see also [She90, Chapter II]. It is often called the fundamental theorem of stability theory. Ben Yaacov [BY14] observed that the same result follows from Grothendieck’s 1952 work on weak compactness in function spaces [Gro52] (see Appendix A below). This is the viewpoint of Section 1.2.

✂ **Section 1.2.1.** Interchanging iterated limits is an old theme. For double sequences of reals, the classical sufficient condition is uniform convergence in one variable (the Moore–Osgood theorem). The condition of Definition 1.2.5 requires the iterated limits to agree along *all* pairs of sequences. This condition comes from functional analysis, and reached model theory through Banach space geometry.



1940–1947

Eberlein–Šmulian. Šmulian (1940) and Eberlein [Ebe47] (1947) prove that in a Banach space weak compactness and weak sequential compactness coincide, so that weak compactness can be tested along sequences.

- 1950–1952** • **Grothendieck.** Grothendieck proves the *double limit criterion*, announced in [Gro50] and proved in [Gro52]. For compact K , a bounded set $H \subseteq C(K)$ is relatively weakly compact if and only if $\lim_m \lim_n f_n(x_m) = \lim_n \lim_m f_n(x_m)$ for all sequences (f_n) in H and (x_m) in K for which both iterated limits exist. For a countably compact space, the same condition decides when pointwise limits of functions in H are continuous. See Appendix A.
- 1961** • **Namioka.** Namioka proves a combinatorial “substitute for Lebesgue’s bounded convergence theorem” [Nam61]. From this result he obtains, without measure theory, Krein’s theorem: in a Banach space, the closed convex hull of a weakly compact set is weakly compact.
- 1963–1964** • **Pták.** Pták [Ptá63a] finds a purely combinatorial lemma about averages, from which both the double limit criterion and Krein’s theorem follow. In [Ptá63b] he uses an interchange of iterated limits to decide which bounded functions on a product $X \times Y$, continuous in each variable separately, extend to bilinear forms with the same kind of continuity.
- 1968** • **Simons.** Simons [Sim68] introduces the *Fubini characterization* of the interchange of double limits. A bounded function on a product satisfies the iterated limit condition if and only if the two iterated integrals of the function agree, for every pair of measures on the factors. This characterization plays an important role in the proof of the factorization theorem of Davis, Figiel, Johnson and Pełczyński [DFJP74].
- 1971** • **Young; Pryce.** N. J. Young [You71] gives topological proofs of Pták’s double-limit theorems, based on Grothendieck’s work rather than combinatorics, and generalizes Pták’s bilinear extension theorem. Pryce [Pry71] uses a trick of Whitley’s to simplify proofs about pointwise compact sets of continuous functions.
- 1972** • **Simons.** In three consecutive papers, Simons gives an elementary proof of Pták’s lemma [Sim72c]. He also proves an inequality about suprema and upper limits, now known as *Simons’ inequality* [Sim72a], and relates minimax theorems, which say when $\max \min = \min \max$, to James’s characterization of reflexivity [Sim72b].
- 1974** • **Davis–Figiel–Johnson–Pełczyński.** Davis, Figiel, Johnson and Pełczyński [DFJP74] prove that every weakly compact operator factors through a reflexive Banach space. This is the key step in the proof of Theorem 1.2.9.
- 1974** • **Namioka.** Namioka [Nam74] studies when separately continuous functions on products are jointly continuous at many points. This is the question behind Proposition 1.2.8.
- 1981** • **Krivine–Maurey.** Krivine and Maurey [KM81] apply the interchange of limits to the norm: a Banach space is *stable* if $\lim_m \lim_n \|x_m + y_n\| = \lim_n \lim_m \|x_m + y_n\|$ for all bounded sequences for which both limits exist. They prove that every stable Banach space contains, for some p , copies of

ℓ_p that are as close to isometric as desired, which generalizes a theorem of Aldous about subspaces of L_1 .

- 1987 • **König.** König [Kön87] studies the double limit relation in angelic spaces, where compactness can be tested by sequences.
- 1994 • **Iovino.** Iovino's Ph.D. thesis [Iov94] introduces stability for real-valued logic.
- 1999 • **Iovino.** Iovino [Iov99] observes that Shelah's order property is the 2-valued case of the same condition, and proves Theorem 1.2.9. Propositions 1.2.4 and 1.2.3 are taken from his paper.
- 2014 • **Ben Yaacov.** Ben Yaacov [BY14] derives the definability of types in stable theories directly from Grothendieck's criterion.
- 2017 • **Pillay.** Pillay [Pil17] comments on the model-theoretic reading of Grothendieck's criterion, and derives from the criterion the positive form of the definitions.

⌘ **Section 1.2.2.** Spaces of left and right φ -types (see Definition 1.2.7) are the analytic counterpart of Shelah's spaces of φ -types. The use of *types* in Banach space theory, as limits of functions $y \mapsto \|x_n + y\|$, is due to Krivine and Maurey [KM81]. Proposition 1.2.8, which characterizes stability by the existence of a separately continuous extension to pairs of types, is Proposition 2.6 of Iovino [Iov99]. Separately continuous functions on products of compact spaces, and their relation to iterated limits, are the subject of Grothendieck's paper [Gro52].

⌘ **Section 1.2.3.** James [Jam64] characterized reflexivity by the absence of sequences (a_m) in B_E and (f_n) in B_{E^*} whose pairing takes one value for $m < n$ and another for $m \geq n$. This is a Banach-space form of the order property. In the early 1980s, Krivine and Maurey [KM81] and Raynaud [Ray83] showed that stable functions give rise to weakly compact operators. The factorization theorem of Davis, Figiel, Johnson and Pełczyński [DFJP74] (1974) says that every weakly compact operator factors through a reflexive space. Iovino [Iov99] combined these results to prove Theorem 1.2.9: stable functions are exactly the pairings between the unit balls of a reflexive space E and of E^* .

⌘ **Section 1.2.4.** Weierstrass proved in 1885 that every continuous function on a compact interval is a uniform limit of polynomials [Wei85]. Stone replaced the interval by an arbitrary compact Hausdorff space and the coordinate function by any family of functions that separates points [Sto37, Sto48]. Ben Yaacov and Usvyatsov [BYU10] used Stone's theorem in the foundations of continuous first-order logic, where it shows that finitely many connectives suffice and that the definable predicates are exactly the continuous functions on type spaces. Ben Yaacov [BY14] combined Grothendieck's criterion with Mazur's lemma [Maz33] (a weakly convergent sequence has convex combinations converging in norm), and so obtained the defining formulas of stable types as uniform limits of averages of instances; see Remark 1.2.14.

Chapter 2

Stability, trees, and definability of types

In this chapter we prove the definability theorem, Theorem 1.1.6, which we stated without proof in Chapter 1. The argument passes through binary trees, and along the way we show that a formula is stable exactly when it cannot encode tall trees. As in Section 1.1, we fix a first-order theory T in a language L and an L -formula $\varphi(x, y)$ consistent with T .

Roadmap

This chapter has a single purpose, to prove the definability theorem, and a single instrument, the binary tree. A theorem about formulas that cannot order is proved here by means of trees because orders and definitions are hard to compare directly, whereas each is easily compared with trees: a long ladder can be folded into a tall tree (see Section 2.3), and a tall tree has more branches than there are definitions to describe them (see Section 2.4). In Section 2.1 we state the four equivalent conditions of Theorem 2.1.3 and explain the order in which they are proved. Section 2.2 contains the substantial argument: in the absence of the order property, the definition of a type is found among positive Boolean combinations of instances of φ , and the whole difficulty lies in choosing finitely many parameters well. In Sections 2.3 and 2.4 we close the circle, by turning ladders into trees and trees into types that no formula defines. Section 2.5 is a topological epilogue, in which we show that over a countable model, a stable formula has countably many types and an unstable one has a perfect set of them, so that the logician's count and the topologist's shape agree. In the language of the subtitle: a definable type depends on its parameters through a single formula, and a tree is the shape that dependence takes when no single formula will do.

2.1 THE EQUIVALENCE PROVED IN THIS CHAPTER

In Theorem 2.1.3 below we show that four conditions on φ coincide: instability, the order property, the presence of tall binary trees, and the existence of a non-definable φ -type. The notions of φ -type, definable type, ladder, order property and stability come from Section 1.1; the newcomer, a binary tree for φ , needs a little notation for finite 0–1 sequences.

Notation 2.1.1. Recall that $\varphi^1 = \varphi$ and $\varphi^0 = \neg\varphi$. For an ordinal μ , we write $2^{<\mu}$ for the set of 0–1 sequences of length $<\mu$, and we write $2^{\leq n} = 2^{<n} \cup 2^n$. We write $\sigma \sqsubseteq \tau$ if σ is an initial segment of τ , $\sigma \sqsubset \tau$ if moreover $\sigma \neq \tau$, $|\sigma|$ for the length of σ , $\sigma^\frown \langle u \rangle$ for the extension of σ by one entry u , and $\eta|k$ for the initial segment of length k of a sequence η of length $>k$.

Definition 2.1.2. A *tree of height n for φ in a model M of T* consists of families $(b_\sigma : \sigma \in 2^{<n})$ and $(a_\eta : \eta \in 2^n)$ of tuples of M such that

$$\sigma^\frown \langle u \rangle \sqsubseteq \eta \implies M \models \varphi(a_\eta, b_\sigma)^u.$$

In other words, the leaf a_η satisfies $\varphi(x, b_\sigma)$ exactly when the branch η turns right ($u = 1$) at the node σ . We say that φ *encodes binary trees of arbitrarily large finite height* if for every n some model of T contains a tree of height n for φ .

Figure 2.1 shows a tree of height 3. Each node b_σ splits the leaves above it into two halves: the leaves on the right satisfy $\varphi(x, b_\sigma)$, and the leaves on the left satisfy $\neg\varphi(x, b_\sigma)$. The branch leading to a leaf a_η is thus a record, node by node, of which of the formulas $\varphi(x, b_\sigma)$ the leaf satisfies, and a tree of height n realizes all 2^n such records.




FIGURE 2.1. A tree of height 3 for φ . Each round node is labeled by the formula $\varphi(x, b_\sigma)$ given by its parameter b_σ , and the squares a_η are the leaves. An edge going up and to the right from the node $\varphi(x, b_\sigma)$ is labeled φ (purple), and an edge going up and to the left is labeled $\neg\varphi$ (pink): every leaf above a right edge out of that node satisfies $\varphi(x, b_\sigma)$, and every leaf above a left edge satisfies $\neg\varphi(x, b_\sigma)$. The edges into the leaves are not labeled, but their colors follow the same rule. For example, the thick branch leading up to a_{101} goes right, left, right, so that $M \models \varphi(a_{101}, b_\emptyset) \wedge \neg\varphi(a_{101}, b_1) \wedge \varphi(a_{101}, b_{10})$. The formula φ is not required to say anything about a_η and b_σ when σ is not an initial segment of η , such as a_{101} and b_0 .

Like ladders, trees can be shortened. A tree of height $n + 1$ restricts to one of height n : keep the nodes b_σ with $\sigma \in 2^{<n}$, and for each $\eta \in 2^n$ keep the leaf $a_{\eta \smallfrown \langle 0 \rangle}$.

Theorem 2.1.3. *Let $\varphi(x, y)$ be an L -formula. The following are equivalent:*

- (1) *The formula φ is unstable (see Definition 1.1.4).*
- (2) *The formula φ has the order property (see Definition 1.1.4).*
- (3) *The formula φ encodes binary trees of arbitrarily large finite height.*
- (4) *Some φ -type over some model of T is not definable (see Definitions 1.1.1 and 1.1.3).*

The proof goes around a cycle of implications. The equivalence (1) $\Leftrightarrow$ (2) is the definition of stability (see Definition 1.1.4). In Section 2.3 we turn long ladders into tall trees, which gives (2) $\Rightarrow$ (3) (see Proposition 2.3.1). In Section 2.4 we use tall trees to produce more φ -types than there are formulas to define them, which gives (3) $\Rightarrow$ (4) (see Proposition 2.4.2). The cycle closes with the definability theorem, which we prove first, in Section 2.2, and whose contrapositive is (4) $\Rightarrow$ (2) (see Theorem 2.2.1 and Corollary 2.2.2).

In Chapter 3 we give alternative proofs of two of these implications, each of which yields information that the proofs in this chapter do not provide. Here the way from tall trees back to the order property passes through non-definable types, and so requires counting types over models of large cardinality. In Section 3.1 we prove (3) $\Rightarrow$ (2) directly instead, by finite combinatorics based on Ramsey's theorem (see Proposition 3.1.5 and Corollary 3.1.6). This direct proof stays inside a single model and is quantitative: a tree of height $2^{\binom{2n}{n}} - 1$ in a model M already yields an n -ladder in M . In Section 3.2 the classical rank argument proves (4) $\Rightarrow$ (3), again in contrapositive form: Proposition 3.2.3 shows that the absence of tall trees alone makes every φ -type definable, without going through the order property, and gives the definition explicitly in terms of a rank. Unlike the definitions of Section 2.2, however, these definitions are not positive Boolean combinations of instances of φ .

2.2 FROM THE ABSENCE OF THE ORDER PROPERTY TO TYPE DEFINABILITY

We begin with the implication of the definability theorem that carries the weight, (1) $\Rightarrow$ (2) of Theorem 1.1.6: if φ does not have the order property, then every φ -type over a model has a definition built from instances of φ using only $\wedge$ and $\vee$. The candidate definitions are easy to write down, and the work lies in choosing the finitely many parameters that they use.

Let M be a model of T , let $p \in S_\varphi(M)$ be given, and recall that $D_p = \{b \in M : \varphi(x, b) \in p\}$. For a finite set $F \subseteq M$, we shall write

$$\mathcal{W}_F = \left\{ \{a \in F : M \models \varphi(a, b)\} : b \in D_p \right\} \quad \text{and} \quad d_F \varphi(y) = \bigvee_{W \in \mathcal{W}_F} \bigwedge_{a \in W} \varphi(a, y).$$

The family $\mathcal{W}_F$ consists of subsets of the finite set F , so the family is finite and the formula $d_F \varphi$ is a positive Boolean combination of instances of φ over M . In words, the formula $d_F \varphi(y)$ says that, for some $b' \in D_p$, every $a \in F$ with $M \models \varphi(a, b')$ also satisfies $\varphi(a, y)$. Every parameter $b \in D_p$ satisfies $d_F \varphi$, whatever F is (take $b' = b$). The task is to choose F so that no parameter outside D_p satisfies $d_F \varphi$, and this is where the absence of the order property is used.

Theorem 2.2.1. *Assume that $\varphi(x, y)$ does not have the order property, let M be a model of T , and let $p \in S_\varphi(M)$ be given. Then there is a finite set $F \subseteq M$ such that $d_F \varphi$ is a φ -definition of p ; that is, for all $b \in M$ we have $\varphi(x, b) \in p$ if and only if $M \models d_F \varphi(b)$.*

Proof. For $b \in M$ we shall write $S_b = \{a \in M : M \models \varphi(a, b)\}$. The only property of p we use is that p is finitely satisfiable in M (see Lemma 1.1.2).

© **Claim.** There is a finite $F \subseteq M$ such that $S_b \cap F \not\subseteq S_{b'}$ for all $b \in D_p$ and $b' \in M \setminus D_p$.

Granting the claim, we check that $d_F \varphi$ defines p . Let $b \in M$ be given.

- (a) If $b \in D_p$, then the set $W := S_b \cap F$ belongs to $\mathcal{W}_F$ and satisfies $W \subseteq S_b$, so we have $M \models \bigwedge_{a \in W} \varphi(a, b)$. Hence we have $M \models d_F \varphi(b)$.
- (b) If $b \notin D_p$ and $W \in \mathcal{W}_F$, say $W = S_{b_0} \cap F$ with $b_0 \in D_p$, then we have $W \not\subseteq S_b$ by the claim. Thus some $a \in W$ has $M \models \neg \varphi(a, b)$, and hence we have $M \models \neg \bigwedge_{a \in W} \varphi(a, b)$. As this holds for every $W \in \mathcal{W}_F$, we conclude that $M \models \neg d_F \varphi(b)$.

© **Proof of the claim.** Suppose the claim fails. Then for every finite $F \subseteq M$ there are $b \in D_p$ and $b' \in M \setminus D_p$ with $S_b \cap F \subseteq S_{b'}$, that is, with $M \models \varphi(a, b) \rightarrow \varphi(a, b')$ for all $a \in F$. Letting the failure of the claim and the finite satisfiability of p take turns, we build a

configuration in which Ramsey's theorem will find a ladder. We choose recursively, for $i < \omega$, elements $b_i \in D_p$, $b'_i \in M \setminus D_p$ and $a_i \in M$ as follows.

- (a) Apply the failure of the claim to $F = \{a_j : j < i\}$ to get $b_i \in D_p$ and $b'_i \in M \setminus D_p$ with

$$(2.1) \quad M \models \varphi(a_j, b_i) \rightarrow \varphi(a_j, b'_i) \quad \text{for all } j < i.$$

- (b) Since $\varphi(x, b_k) \in p$ and $\neg\varphi(x, b'_k) \in p$ for $k \leq i$, Lemma 1.1.2 gives $a_i \in M$ with

$$(2.2) \quad M \models \varphi(a_i, b_k) \wedge \neg\varphi(a_i, b'_k) \quad \text{for all } k \leq i.$$

Color each pair $j < i$ by the pair of truth values $(\varphi(a_j, b_i), \varphi(a_j, b'_i)) \in \{0, 1\}^2$. By (2.1), the color $(1, 0)$ never occurs. By Ramsey's theorem there is an infinite $H = \{h_0 < h_1 < \dots\} \subseteq \omega$ on which all pairs have the same color c . Let $s, t < \omega$ be given.

⊙ **Case $c \in \{(0, 1), (1, 1)\}$.** If $s < t$, then we have $M \models \varphi(a_{h_s}, b'_{h_t})$ by homogeneity. If $s \geq t$, then we have $M \models \neg\varphi(a_{h_s}, b'_{h_t})$ by (2.2). Thus the formula $\varphi(a_{h_s}, b'_{h_t})$ holds if and only if $s < t$, and with $\beta_t = b'_{h_{t+1}}$ we get $M \models \varphi(a_{h_s}, \beta_t) \iff s \leq t$.

⊙ **Case $c = (0, 0)$.** If $s < t$, then we have $M \models \neg\varphi(a_{h_s}, b_{h_t})$ by homogeneity. If $s \geq t$, then we have $M \models \varphi(a_{h_s}, b_{h_t})$ by (2.2). Thus the formula $\varphi(a_{h_s}, b_{h_t})$ holds if and only if $s \geq t$. For a given n , reversing the first n indices by putting $\alpha_s = a_{h_{n-1-s}}$ and $\beta_t = b_{h_{n-1-t}}$ gives $M \models \varphi(\alpha_s, \beta_t) \iff s \leq t$ for $s, t < n$.

In both cases the model M contains an n -ladder for every n , so the formula φ has the order property, a contradiction. This proves the claim and the theorem. ☺

Corollary 2.2.2. *If φ does not have the order property, then for every model M of T , every $p \in S_\varphi(M)$ is definable by a positive Boolean combination of instances $\varphi(a, y)$ with $a \in M$. In particular, if φ is quantifier-free, then the type p has a quantifier-free φ -definition.*

The proof of Theorem 2.2.1 applies Ramsey's theorem for pairs to an infinite sequence, inside an argument by contradiction, and so says nothing about sizes. A finite Ramsey-type argument, adapted to binary trees rather than to sequences, reappears in Section 3.1, where it extracts ladders of prescribed length directly from trees of explicit height.

2.3 FROM THE ORDER PROPERTY TO TALL TREES

To prove (3) $\Rightarrow$ (1) of Theorem 1.1.6, we need a supply of many φ -types over a small model, and binary trees (see Definition 2.1.2) are the natural source. The first step, taken in this section, is to grow tall trees out of long ladders, which is (2) $\Rightarrow$ (3) of Theorem 2.1.3. The idea is to place the nodes and the leaves of a binary tree on a line, in such a way that the leaves above the right edge out of each node lie to the left of that node and the leaves above the left edge lie to the right of the node. In Proposition 2.3.1 below we show that a ladder read along this line then has exactly the pattern of a tree.

Proposition 2.3.1. *If a model M of T contains an N -ladder for φ , where $N = 2^{n+1} - 1$, then it contains a tree of height n . Hence, if φ has the order property, then it encodes binary trees of arbitrarily large finite height.*

Proof. For $\sigma \in 2^{<n}$ put

$$d(\sigma) = \sum_{k < |\sigma|} (1 - \sigma(k)) 2^{-(k+1)} + 2^{-(|\sigma|+1)} \in (0, 1).$$

The binary expansion of $d(\sigma)$ is $(1 - \sigma(0), \dots, 1 - \sigma(|\sigma| - 1), 1)$, which ends in 1. Thus the map d is injective on $2^{\leq n}$, and the image V of d has exactly $|2^{\leq n}| = 2^{n+1} - 1 = N$ elements.

⊙ **Claim.** If $\sigma \frown \langle u \rangle \subseteq \eta$ with $\eta \in 2^n$, then we have $d(\eta) < d(\sigma) \iff u = 1$.

Indeed, let $k = |\sigma|$. Since $\eta \upharpoonright k = \sigma$ and $\eta(k) = u$, we have

$$d(\eta) - d(\sigma) = -u 2^{-(k+1)} + R, \quad R = \sum_{k < j < n} (1 - \eta(j)) 2^{-(j+1)} + 2^{-(n+1)}.$$

Here we have $0 < R \leq (2^{-(k+1)} - 2^{-n}) + 2^{-(n+1)} < 2^{-(k+1)}$. Hence the difference $d(\eta) - d(\sigma)$ is negative if $u = 1$ and positive if $u = 0$, which proves the claim.

Let $a_1, \dots, a_N, b_1, \dots, b_N \in M$ satisfy $\varphi(a_i, b_j) \iff i \leq j$, and let $e: V \rightarrow \{1, \dots, N\}$ be the order isomorphism. Put $b_\sigma := b_{e(d(\sigma))}$ for $\sigma \in 2^{< n}$ and $a_\eta := a_{e(d(\eta))}$ for $\eta \in 2^n$. If $\sigma \frown \langle u \rangle \subseteq \eta$, then we have $d(\eta) \neq d(\sigma)$, and so we have

$$M \models \varphi(a_\eta, b_\sigma) \iff e(d(\eta)) \leq e(d(\sigma)) \iff d(\eta) < d(\sigma) \iff u = 1.$$

Thus the families (b_σ) and (a_η) form a tree of height n in M . ⊙

Figure 2.2 shows the construction for $n = 2$.




FIGURE 2.2. The proof of Proposition 2.3.1 for $n = 2$. The seven sequences of length at most 2 are placed at the positions $e(d(\sigma)) = 1, \dots, 7$ of a ladder of length 7, and the tree grows upward. In a ladder, the formula $\varphi(a_i, b_j)$ holds when a_i lies to the left of b_j (that is, $i \leq j$), so here the φ -edges (purple) go up and to the left and the $\neg\varphi$ -edges (pink) go up and to the right. The bound $R < 2^{-(k+1)}$ in the claim is what keeps every leaf on the correct side of each node below it: the leaf a_{10} , for instance, lies to the left of b_0 and to the right of b_1 .

2.4 TALL TREES PRODUCE NON-DEFINABLE TYPES

Tall trees produce more φ -types over a small model than there are formulas to define them, and this proves (3) $\Rightarrow$ (4) of Theorem 2.1.3. The count is done in Proposition 2.4.2 below. To prepare for it, in Lemma 2.4.1 below we use compactness to turn finite trees into a tree of infinite height, whose branches will give the types. When L is countable, the count is easiest to see over a countable model, and this case is the picture to keep in mind: a tree of height ω has $2^{\aleph_0}$ branches, each branch gives a φ -type over a countable model, and only countably many formulas are available to define them. In Section 2.5 the same tree of height ω reappears as a copy of the Cantor space inside $S_\varphi(M)$.

Lemma 2.4.1 (Compactness). *Suppose that φ encodes binary trees of arbitrarily large finite height, and let μ be an infinite cardinal. Then there are a model M' of T and tuples $(b_s)_{s \in 2^{<\mu}}$ in M' such that for every $\eta \in 2^\mu$ the set*

$$\Sigma_\eta(x) = \{\varphi(x, b_{\eta|_\alpha})^{\eta(\alpha)} : \alpha < \mu\}$$

is finitely satisfiable in M' .

Proof. Add new tuples of constants c_s ($s \in 2^{<\mu}$) and let Γ be T together with the sentences

$$\gamma_{\eta, F} := \exists x \bigwedge_{\alpha \in F} \varphi(x, c_{\eta|_\alpha})^{\eta(\alpha)} \quad (\eta \in 2^\mu, F \subseteq \mu \text{ finite}).$$

By compactness, we need only show that every finite $\Gamma_0 \subseteq \Gamma$ has a model.

Let γ_{η_i, F_i} ($i < m$) be the sentences of this form in Γ_0 (if $m = 0$ there is nothing to do). Only finitely many nodes of $2^{<\mu}$ are mentioned in these sentences, and the plan is to embed them into a finite tree, preserving the direction in which each branch leaves each node. Let S be the finite set of all $\eta_i|_\alpha$ and $\eta_i|_{(\alpha+1)}$ with $i < m$ and $\alpha \in F_i$. Note that $\alpha + 1 < \mu$ because μ is a limit ordinal. Let $S' = \{s \wedge s' : s, s' \in S\}$, where $s \wedge s'$ is the longest common initial segment. The set S' is finite and closed under $\wedge$, because in a tree the meet $s \wedge s' \wedge s''$ equals one of the pairwise meets. Thus the set S' has a least element r , and every $s \in S' \setminus \{r\}$ has a longest predecessor $s^- \in S'$ with $s^- \sqsubset s$.

Define $f: S' \rightarrow 2^{<N}$, where $N = |S'| + 1$, by recursion on length, setting $f(r) = \emptyset$ and $f(s) = f(s^-) \frown \langle s(|s^-|) \rangle$. Then we have $|f(s)| < N$.

© **Claim.** If $s, s' \in S'$ and $s \sqsubset s'$, then we have $f(s) \frown \langle s'(|s|) \rangle \sqsubseteq f(s')$.

The proof is by induction on $|s'|$. If $s'^- = s$, this is the definition of f . Otherwise we have $s \sqsubset s'^- \sqsubset s'$ by maximality of s'^- . By induction, we have $f(s) \frown \langle s'^- (|s|) \rangle \sqsubseteq f(s'^-) \sqsubseteq f(s')$, and finally we have $s'^- (|s|) = s' (|s|)$.

By hypothesis, take a model M_0 of T containing a tree $(b_\sigma)_{\sigma \in 2^{<N}}$, $(a_\nu)_{\nu \in 2^N}$ of height N . Interpret c_s as $b_{f(s)}$ for $s \in S'$, and arbitrarily otherwise. Fix $i < m$, let $\alpha^* = \max F_i$ and $s^* = \eta_i|_{(\alpha^*+1)} \in S'$. Pick $\nu \in 2^N$ extending $f(s^*)$. For $\alpha \in F_i$ we have $\eta_i|_\alpha \sqsubset s^*$ and $s^*(\alpha) = \eta_i(\alpha)$. By the Claim, we have $f(\eta_i|_\alpha) \frown \langle \eta_i(\alpha) \rangle \sqsubseteq \nu$, which gives $M_0 \models \varphi(a_\nu, b_{f(\eta_i|_\alpha)})^{\eta_i(\alpha)}$. Hence the tuple a_ν witnesses γ_{η_i, F_i} , and the structure M_0 with this interpretation is a model of Γ_0 .

By compactness, the theory Γ has a model. The L -reduct M' of that model, with $b_s = c_s^{M'}$, is as required. 

In the next proposition we write $\kappa = |L| + \aleph_0$, which is the number of L -formulas.

Proposition 2.4.2. *If φ encodes binary trees of arbitrarily large finite height, then for every cardinal $\lambda \geq \kappa$ there is a model M of T with $|M| \leq \lambda$ and $|S_\varphi(M)| > \lambda$. In particular, there are a model M of T with $|M| \leq \kappa$ and a type $p \in S_\varphi(M)$ such that p is not definable.*

Proof. Let μ be the least cardinal with $2^\mu > \lambda$. Then the cardinal μ is infinite and satisfies $\mu \leq \lambda$. Since $2^\nu \leq \lambda$ for every $\nu < \mu$, we have $|2^{<\mu}| = \sum_{\nu < \mu} 2^\nu \leq \mu \cdot \lambda = \lambda$.

Take M' and (b_s) as in Lemma 2.4.1. By the downward Löwenheim–Skolem theorem, there is $M \preceq M'$ with $|M| \leq \lambda$ containing all b_s . Each finite part of Σ_η is satisfiable in M' , hence in M , since the formula $\exists x \bigwedge \dots$ is an $L(M)$ -sentence. Thus the set Σ_η is consistent with the elementary diagram of M and extends to some $p_\eta \in S_\varphi(M)$.

If $\eta \neq \eta'$, let α be least with $\eta(\alpha) \neq \eta'(\alpha)$. Put $s = \eta|_\alpha = \eta'|_\alpha$; then the types p_η and $p_{\eta'}$ disagree on $\varphi(x, b_s)$. Hence we have $|S_\varphi(M)| \geq 2^\mu > \lambda$.

For the last statement take $\lambda = \kappa$. A definable type is determined by any φ -definition of the type, and there are only $|L(M)| \leq \kappa$ formulas over M . Thus at most κ types in $S_\varphi(M)$ are definable, and some $p \in S_\varphi(M)$ is not. ☺

Proof of (3) $\Rightarrow$ (1) in Theorem 1.1.6. If φ has the order property, then it encodes binary trees of arbitrarily large finite height by Proposition 2.3.1, and so some φ -type over some model of T is not definable by Proposition 2.4.2. ☺

This completes the proof of Theorem 1.1.6, and with it that of Theorem 2.1.3. The equivalence of (2) and (3) in Theorem 2.1.3 is the form in which trees are used later (see Section 2.5), so we record it separately.

Theorem 2.4.3. *The formula φ has the order property if and only if it encodes binary trees of arbitrarily large finite height.*

Proof. The forward direction is Proposition 2.3.1. Conversely, if φ encodes tall trees, then some φ -type is not definable (see Proposition 2.4.2), so the formula φ has the order property by (1) $\Rightarrow$ (3) of Theorem 1.1.6. ☺

2.5 TYPES OVER COUNTABLE MODELS: A PERFECT SET DICHOTOMY

In Section 2.4 we measured instability by counting types; topology measures instability by shape, and over countable models the two measurements agree. Throughout this section the language L is countable. For a model M of T , the space $S_\varphi(M)$ is the Stone space of the Boolean algebra $\mathcal{B}_\varphi(M)$ (see Section 1.1). If M is countable, the algebra $\mathcal{B}_\varphi(M)$ is countable, so its Stone space $S_\varphi(M)$ is a compact metrizable zero-dimensional space. By the Cantor–Bendixson theorem, such a space is either countable, in which case the space is scattered, or contains a homeomorphic copy of the Cantor space 2^ω . In Theorem 2.5.1 below we show that stability decides which alternative occurs.

Theorem 2.5.1. *Assume that L is countable. The following are equivalent:*

- (1) *The formula φ is stable.*
- (2) *For every countable model M of T , the space $S_\varphi(M)$ is countable.*
- (3) *For every countable model M of T , the space $S_\varphi(M)$ is scattered.*

Moreover, if φ is unstable, then some countable model M of T has a closed subset of $S_\varphi(M)$ homeomorphic to 2^ω .

Proof. (1) $\Rightarrow$ (2). By Theorem 1.1.6, every $p \in S_\varphi(M)$ has a φ -definition, which is an $L(M)$ -formula, and the definition determines p . Since there are only countably many $L(M)$ -formulas, the space $S_\varphi(M)$ is countable.

(2) $\Leftrightarrow$ (3). This is the Cantor–Bendixson theorem for compact metrizable spaces.

(2) $\Rightarrow$ (1) and the last statement. Suppose that φ is unstable. By Theorem 2.4.3, the formula φ encodes binary trees of arbitrarily large finite height, so Lemma 2.4.1 with $\mu = \omega$ gives a model M' of T and tuples $(b_\sigma)_{\sigma \in 2^{<\omega}}$ in M' such that every set $\Sigma_\eta(x)$, for $\eta \in 2^\omega$, is finitely satisfiable in M' . By the downward Löwenheim–Skolem theorem, there is a countable model $M \preceq M'$ containing every b_σ , and by elementarity every $\Sigma_\eta(x)$ is finitely satisfiable in M . For $\sigma \in 2^{<\omega}$, let C_σ be the clopen set of types $p \in S_\varphi(M)$ that contain $\varphi(x, b_{\sigma|_\alpha})^{\sigma(\alpha)}$ for every $\alpha < |\sigma|$. Each C_σ is nonempty, because the formulas in question form a finitely satisfiable subset of Σ_η for any η extending σ , and the set C_σ is the disjoint union of $C_{\sigma \frown \langle 0 \rangle}$

and $C_{\sigma \smallfrown \langle 1 \rangle}$, which disagree on $\varphi(x, b_\sigma)$. Thus the sets C_σ form a Cantor scheme. The closed set $D = \bigcap_n \bigcup_{|\sigma|=n} C_\sigma$ maps continuously onto 2^ω , by sending p to the unique η with $p \in C_{\eta|n}$ for every n ; the map is onto because each $\bigcap_n C_{\eta|n}$ is a decreasing intersection of nonempty closed sets in a compact space. Hence the space $S_\varphi(M)$ is uncountable, so (2) fails, and by the Cantor–Bendixson theorem the uncountable compact metrizable space $S_\varphi(M)$ contains a closed copy of 2^ω . ☺




FIGURE 2.3. The Cantor scheme in the proof of Theorem 2.5.1, growing upward. Each bar is a nonempty clopen set C_σ of φ -types, and C_σ splits into $C_{\sigma \smallfrown \langle 0 \rangle}$ (left) and $C_{\sigma \smallfrown \langle 1 \rangle}$ (right), which disagree on $\varphi(x, b_\sigma)$. The pink bars $C_1 \supseteq C_{11} \supseteq C_{110} \supseteq \dots$ follow one branch η , and their intersection contains a type p with $p \in C_{\eta|n}$ for every n . Distinct branches give distinct types, so the types along the $2^{\aleph_0}$ branches form the closed set D , which maps onto 2^ω .

A tree of infinite height is thus nothing but a Cantor scheme of clopen sets in a type space (see Figure 2.3), and stability says that type spaces over countable models are scattered. Morley [Mor65] used the same topology to define the totally transcendental theories, whose type spaces over countable sets are scattered, and whose Morley rank is the Cantor–Bendixson rank on type spaces over $\aleph_0$ -saturated models. The rank R introduced in Section 3.2 below is a binary splitting rank on the clopen sets $[\theta]$, a finite combinatorial cousin of the Cantor–Bendixson rank: Lemma 3.2.2(iii) says that splitting a clopen set into two pieces of rank at least m raises the rank.

2.6 EXERCISES

The exercises follow the sections of the chapter. Throughout, we write T_E for the theory of an equivalence relation E with infinitely many classes, all of them infinite, and θ for the damped ladder, as in the exercises of Section 1.3. For every finite 0–1 sequence σ we put $d(\sigma) = \sum_{k < |\sigma|} (1 - \sigma(k)) 2^{-(k+1)} + 2^{-(|\sigma|+1)}$, as in the proof of Proposition 2.3.1; the number $d(\sigma)$ does not depend on n . The mark (*) indicates the exercises that ask for more than a direct application of the text.

1. (Section 2.1.) Work in $(\mathbb{Q}, <)$, with $\varphi(x, y)$ the formula $x < y$.

- (a) Put $b_0 = 4$, $b_1 = 2$, $b_2 = 6$, $a_{11} = 1$, $a_{10} = 3$, $a_{01} = 5$ and $a_{00} = 7$. Check that these elements form a tree of height 2 for φ , and draw it in the manner of Figure 2.1.

- (b) Find a tree of height 3 for φ whose fifteen elements are the integers $1, \dots, 15$. (Hint: place the root at 8, and place the leaves above the right edge out of each node to the left of that node.)
- (c) Let $(I, <)$ be a linear order, and let $(b_\sigma), (a_\eta)$ be a tree of height n for $x < y$ in $(I, <)$. Let $\sigma \in 2^{<n}$ be given, and let $\eta, \eta' \in 2^n$ satisfy $\sigma \cap \langle 1 \rangle \subseteq \eta$ and $\sigma \cap \langle 0 \rangle \subseteq \eta'$. Show that $a_\eta < b_\sigma \leq a_{\eta'}$. Deduce that the leaves are pairwise distinct, and that $a_{\nu'} < a_\nu$ whenever $\nu, \nu' \in 2^n$ and ν precedes ν' in the lexicographic order. Thus the arrangement of part (a), in which the leaves turning right lie to the left, is forced.

2. (Section 2.1.)

- (a) Let $(b_\sigma), (a_\eta)$ be a tree of height n for φ in a model M of T . For a sequence σ , let $\bar{\sigma}$ be the sequence of the same length with $\bar{\sigma}(k) = 1 - \sigma(k)$. Show that the families $(b_{\bar{\sigma}})$ and $(a_{\bar{\eta}})$ form a tree of height n for $\neg\varphi$ in M . Conclude that φ and $\neg\varphi$ admit trees of the same heights in M .
- (b) Let M be an infinite set, regarded as a structure in the empty language, and let $\varphi(x, y)$ be the formula $x = y$. Show that M contains a tree of height 1 for φ and no tree of height 2. (Hint: in a tree of height 2, the leaf a_{11} equals both $b_\emptyset$ and b_1 .)
- (c) Let M be a model of T_E . Show that the largest height of a tree for xEy in M is 1.
- (d) In the setting of part (b), show that the largest m such that M contains an m -ladder is 1 for the formula $x = y$ and 2 for the formula $x \neq y$. Thus the length of the longest ladder may change when φ is replaced by $\neg\varphi$, whereas by part (a) the height of the tallest tree does not.

3. (Section 2.1.) Let K be a finite field, let T be the theory of infinite vector spaces over K , and let $n \geq 1$ be given. As in the exercises of Section 1.3, let $\varphi(x_1, \dots, x_n; y)$ be the formula which says that y lies in the span of $x_1, \dots, x_n$. Let M be a model of T . Show that the largest height of a tree for φ in M is n . (Hint: in a tree of height h , for $\sigma \in 2^{<h}$, let U_σ be the span of the nodes b_τ with $\tau \cap \langle 1 \rangle \subseteq \sigma$. Show that $b_\sigma \notin U_\sigma$ by looking at a leaf above $\sigma \cap \langle 0 \rangle$, and follow the branch that always turns right. For the construction, take the nodes linearly independent.) Compare with the exercises of Section 1.3, where the longest ladder for φ in M has length $n + 1$.

4. (Section 2.2.) Let M be a model of T , let $p \in S_\varphi(M)$ be given, and let $\mathcal{W}_F$ and $d_F\varphi$ be as in Section 2.2.

- (a) Let $F \subseteq F'$ be finite. Show that every $b \in M$ with $M \models d_{F'}\varphi(b)$ also satisfies $M \models d_F\varphi(b)$. Deduce that, if $d_F\varphi$ defines p , then so does $d_{F'}\varphi$.
- (b) Show that $d_F\varphi$ defines p if and only if F has the property stated in the claim in the proof of Theorem 2.2.1.
- (c) Let M be an infinite set in the empty language, and let $\varphi(x, y)$ be the formula $x = y$. Let $c \in M$ be given, and let p be the type with $D_p = \{c\}$. Compute $\mathcal{W}_F$ and $d_F\varphi$ for every finite $F \subseteq M$, and show that $d_F\varphi$ defines p if and only if $c \in F$. Do the same for the type with $D_p = \emptyset$.

- (d) Let M be a model of T_E , let $\varphi(x, y)$ be the formula xEy , and let p be a type with D_p an E -class C . Show that $\mathcal{W}_F = \{F \cap C\}$, and that $d_F\varphi$ defines p if and only if F meets C . Determine the least size of a set F that works.
- (e) In the setting of Exercise 3, let $p \in S_\varphi(M)$ be given. Show that there is a single n -tuple $c \in M^n$ such that $d_{\{c\}}\varphi$ is the formula $\varphi(c, y)$ and defines p . (Hint: the set D_p is a subspace of M of dimension at most n (see the exercises of Section 1.3).)
5. (Section 2.2.) Work in $(\mathbb{Q}, <)$, with $\varphi(x, y)$ the formula $x < y$. Let r be an irrational real number, and let $p \in S_\varphi(\mathbb{Q})$ be the type with $D_p = \{b \in \mathbb{Q} : b > r\}$.
- (a) Let $F \subseteq \mathbb{Q}$ be finite, and put $F_r = \{a \in F : a < r\}$. Show that every member of $\mathcal{W}_F$ contains F_r , and that $F_r \in \mathcal{W}_F$. Deduce that the formula $d_F\varphi(y)$ is equivalent in $(\mathbb{Q}, <)$ to $\bigwedge_{a \in F_r} a < y$, which defines the set $\{b : b > \max F_r\}$ when F_r is nonempty and the set $\mathbb{Q}$ otherwise.
- (b) Conclude that no finite $F \subseteq \mathbb{Q}$ makes $d_F\varphi$ a definition of p .
- (c) Thus the claim in the proof of Theorem 2.2.1 fails for p . Let $(c_i)_{i < \omega}$ be a strictly increasing sequence of rationals below r , and let c be a rational above r . Show that the choices $b_i = c$, $b'_i = c_i$ and $a_i = c_i$ satisfy (2.1) and (2.2). Determine the color of each pair, and the ladder that the proof extracts.
- Compare with the exercises of Section 1.3 on $(\mathbb{Q}, <)$, where the same type is shown not to be definable at all.
6. (Section 2.3.) Work in $(\mathbb{Q}, <)$, with $\varphi(x, y)$ the formula $x < y$, and let $a_i = i$ and $b_j = j + \frac{1}{2}$ for $1 \leq i, j \leq 7$.
- (a) Show that these elements form a 7-ladder for φ .
- (b) Carry out the proof of Proposition 2.3.1 with $n = 2$ and this ladder. Compute $d(\sigma)$ for $\sigma \in 2^{\leq 2}$, the index $e(d(\sigma))$, and the elements b_σ and a_η that the proof assigns. Check the tree conditions directly, and compare with Figure 2.2.
- (c) Let n be given, and let V be the image of d on $2^{\leq n}$. Show that $V = \{k 2^{-(n+1)} : 0 < k < 2^{n+1}\}$, so that $e(d(\sigma)) = 2^{n+1}d(\sigma)$. Show that this index is odd if σ is a leaf and even if σ is a node, and that the root receives the index 2^n . Identify the elements of the ladder that the proof leaves unused.
7. (*) (Section 2.3.) Proposition 2.3.1 turns a ladder of length $2^{n+1} - 1$ into a tree of height n . This exercise shows that length 2^n suffices, and that length $2^n - 2$ does not suffice in general.
- (a) Let $a_1, \dots, a_{2^n}, b_1, \dots, b_{2^n}$ be a 2^n -ladder for φ in a model M of T . For $\sigma \in 2^{<n}$, put $k_\sigma = 2^{n+1}d(\sigma)$. Show that the elements $b_\sigma := b_{k_\sigma/2}$ for $\sigma \in 2^{<n}$ and $a_\eta := a_{(k_\eta+1)/2}$ for $\eta \in 2^n$ form a tree of height n in M . (Hint: use Exercise 6(c), and note that $k_\eta < k_\sigma$ if and only if $k_\eta + 1 \leq k_\sigma$.)
- (b) Let $k \in \mathbb{N}$ be given, let R_k be the binary relation on $\mathbb{N}$ given by $R_k(m, m')$ if and only if $m \leq m'$ and $m \leq k$, and let T_k be the theory of $(\mathbb{N}, R_k)$, with $\varphi(x, y)$ the formula $R_k(x, y)$. Show that T_k implies the following two sentences: exactly $k + 1$ elements x satisfy $R_k(x, x)$, and every x with $\neg R_k(x, x)$ satisfies $\neg R_k(x, y)$ for all y .

- (c) Show that every model of T_k contains a $(k+1)$ -ladder and no $(k+2)$ -ladder for φ . Show that a model of T_k contains a tree of height h for φ if and only if $2^h \leq k+2$. (Hint: two distinct leaves of a tree disagree on some formula $\varphi(x, b_\sigma)$, and by part (b) a leaf a with $\neg R_k(a, a)$ satisfies none of these formulas, so at most $k+2$ leaves are available. For the construction, work in $(\mathbb{N}, R_k)$, put $a_\eta = (k_\eta - 1)/2$ and $b_\sigma = k_\sigma/2 - 1$ with $n = h$, and note that the existence of a tree of height h is expressed by a sentence.)
- (d) Let $N \geq 1$ be given. Deduce that the structure $(\mathbb{N}, R_{N-1})$ contains an N -ladder for φ and no tree of height n when $N < 2^n - 1$. Thus the length of the ladder needed for a tree of height n lies between $2^n - 1$ and 2^n .
- (e) Show that $R_k(m, m')$ holds if and only if $\theta(m, m') \geq 2^{-k-1}$, and that $\theta(m, m') > 0$ holds if and only if $m \leq m'$. Conclude that the formula R_k , which defines the level set $\{\theta \geq 2^{-k-1}\}$ of the damped ladder, is stable in T_k , that the heights of its trees tend to infinity with k , and that the formula $\theta > 0$ has the order property. Explain why this is consistent with the stability of θ (see the exercises of Section 1.3).
8. (Section 2.4.) Let M be a model of T . We say that M contains a *tree of height ω* for φ if there are tuples $(b_\sigma)_{\sigma \in 2^{<\omega}}$ in M such that, for every $\eta \in 2^\omega$, the set $\Sigma_\eta(x) = \{\varphi(x, b_{\eta|k})^{\eta(k)} : k < \omega\}$ is finitely satisfiable in M , as in Lemma 2.4.1 with $\mu = \omega$.
- (a) Show that, if M contains a tree of height ω , then the model M contains a tree of height n for every n , and the formula φ is unstable. (Hint: for $\eta \in 2^n$, choose $\eta' \in 2^\omega$ extending η , and let $a_\eta \in M$ realize the finite part $\{\varphi(x, b_{\eta|k})^{\eta(k)} : k < n\}$ of $\Sigma_{\eta'}$. For instability, use Theorem 2.4.3.)
- (b) Assume that L is countable. Show that φ is unstable if and only if some countable model of T contains a tree of height ω for φ . (Hint: Lemma 2.4.1 and the downward Löwenheim–Skolem theorem.)
- (c) Let M be countable and contain a tree of height ω . Show that $|S_\varphi(M)| = 2^{\aleph_0}$. Conclude that a countable model with only countably many φ -types contains no tree of height ω (for the converse, see Exercise 15).
9. (Section 2.4.) Work in $(\mathbb{Q}, <)$, with $\varphi(x, y)$ the formula $x < y$, and put $b_\sigma = d(\sigma)$ for $\sigma \in 2^{<\omega}$.
- (a) Show that the set $\{b_\sigma : \sigma \in 2^{<\omega}\}$ is the set of dyadic rationals in the interval $(0, 1)$, which is countable and dense in $(0, 1)$. (Hint: Exercise 6(c).)
- (b) Show that the elements b_σ form a tree of height ω in $\mathbb{Q}$ (see Exercise 8). (Hint: for $\eta \in 2^\omega$ and $n < \omega$, the claim in the proof of Proposition 2.3.1 shows that $d(\eta|n)$ satisfies $\varphi(x, b_{\eta|k})^{\eta(k)}$ for all $k < n$.) No appeal to compactness is needed here.
- (c) Deduce that $|S_\varphi(\mathbb{Q})| = 2^{\aleph_0}$, and that already the instances $\varphi(x, b)$ with b a dyadic rational in $(0, 1)$ separate $2^{\aleph_0}$ types. Conclude that some φ -type over $\mathbb{Q}$ is not definable, without using the description of definable sets of $(\mathbb{Q}, <)$.
- (d) For $\eta \in 2^\omega$, put $r_\eta = \sum_{k < \omega} (1 - \eta(k)) 2^{-(k+1)}$. Assume that η is not eventually constant. Show that $r_\eta < b_{\eta|k}$ if $\eta(k) = 1$ and $r_\eta > b_{\eta|k}$ if $\eta(k) = 0$, and that every $p \in S_\varphi(\mathbb{Q})$ containing Σ_η contains $x < b$ for each rational $b > r_\eta$ and $\neg(x < b)$ for each rational $b < r_\eta$. Determine the types that contain Σ_η when r_η is irrational.

- 10. (*) (Section 2.4.)** This exercise shows that the passage to a new model in Lemma 2.4.1 cannot be avoided. For $N \geq 1$, let C_N be a chain of N elements. Let M be the disjoint union of the chains C_N , ordered by letting $c < c'$ hold exactly when c and c' lie in the same chain and c precedes c' there. Let $\varphi(x, y)$ be the formula $x < y$, and let T be the theory of M .
- (a) Show that M contains a tree of height n for φ for every n . (Hint: Exercise 1.)
 - (b) Let $(b_\sigma)_{\sigma \in 2^{<\omega}}$ be a tree of height ω in M (see Exercise 8). Show that every node b_σ with $\sigma \neq \emptyset$ and $\sigma(0) = 1$ lies in the chain of $b_\emptyset$, and that, for every n , the 2^n leaves a_η with $\eta(0) = 1$ of the tree of height $n + 1$ given by Exercise 8(a) are pairwise distinct and lie in the chain of $b_\emptyset$. Conclude that M contains no tree of height ω .
- Since M contains trees of every finite height, Lemma 2.4.1, applied to the elementary diagram of M in place of T , gives an elementary extension M' of M containing a tree of height ω . Show that, in such an M' , the set of elements comparable with $b_\emptyset$ is infinite, although in M every such set is finite.
- 11. (Section 2.4.)** This exercise concerns the choice of μ in the proof of Proposition 2.4.2. Assume that L is countable, so that $\kappa = \aleph_0$.
- (a) Show that, for $\lambda = \aleph_0$, the proof takes $\mu = \aleph_0$ and produces a countable model M with $|S_\varphi(M)| = 2^{\aleph_0}$.
 - (b) Let $\lambda = \aleph_1$, and assume that $2^{\aleph_0} > \aleph_1$. Show that the proof takes $\mu = \aleph_0$, and that the set $2^{<\omega_1}$ has more than $\aleph_1$ elements. Thus, with $\mu = \lambda$, the bound $|2^{<\mu}| \leq \lambda$ used in the proof would fail.
 - (c) Let $\lambda = \aleph_1$, and assume that $2^{\aleph_0} = \aleph_1$. Show that the proof takes $\mu = \aleph_1$, and that $|2^{<\omega_1}| = \aleph_1$.
 - (d) Deduce from the proposition that, if φ encodes binary trees of arbitrarily large finite height, then for every cardinal $\lambda \geq \kappa$ some model M of T with $|M| \leq \lambda$ carries a φ -type that is not definable.
- 12. (Section 2.5.)** For a topological space X , let X' be the set of points of X that are not isolated in X . Put $X^{(0)} = X$, $X^{(\alpha+1)} = (X^{(\alpha)})'$, and $X^{(\delta)} = \bigcap_{\alpha < \delta} X^{(\alpha)}$ for limit δ . The *Cantor–Bendixson rank* of a point $p \in X$ is the ordinal α with $p \in X^{(\alpha)} \setminus X^{(\alpha+1)}$, when such an ordinal exists. Recall that X is scattered if every nonempty closed subset of X has a point that is isolated in that subset.
- (a) Let M be a countably infinite set in the empty language, and let $\varphi(x, y)$ be the formula $x = y$. Using the description of $S_\varphi(M)$ as a one-point compactification in the exercises of Section 1.3, show that the realized types have rank 0, that the remaining type has rank 1, and that $S_\varphi(M)^{(2)} = \emptyset$.
 - (b) Do the same for a countable model of T_E and the formula xEy .
 - (c) Let M be an uncountable set in the empty language, with φ the formula $x = y$. Show that $S_\varphi(M)$ is uncountable and scattered. Explain why this does not contradict Theorem 2.5.1.
- 13. (*) (Section 2.5.)** Work in the setting of Exercise 3, and let M be a model of T . By the exercises of Section 1.3, the map $p \mapsto D_p$ is a bijection from $S_\varphi(M)$ onto the set of

subspaces of M of dimension at most n . Show by induction on $j \leq n$ that $S_\varphi(M)^{(j)}$ consists of the types p with $\dim D_p \leq n - j$, and conclude that the Cantor–Bendixson rank of the type with $D_p = \{0\}$ is n and that $S_\varphi(M)^{(n+1)} = \emptyset$. (Hint: let p be a type with $\dim D_p = k$. The conjunction of the formulas $\varphi(x, c)$ with $c \in D_p$ defines a neighborhood of p in which every other type q satisfies $D_q \supsetneq D_p$. When $k < n - j$, a basic neighborhood of p is determined by finitely many formulas, among them finitely many $\neg\varphi(x, b)$ with $b \notin D_p$; find infinitely many subspaces $D_p + Ke$ of dimension $k + 1$ that avoid these b , using that K and D_p are finite and M is infinite.) Compare with Exercise 3: for this formula, the largest Cantor–Bendixson rank of a point of $S_\varphi(M)$ equals the largest height of a tree for φ .

14. (Section 2.5.) Work in $(\mathbb{Q}, <)$, with $\varphi(x, y)$ the formula $x < y$.
- (a) For $\eta \in 2^\omega$, put $s_\eta = \sum_{k < \omega} (1 + \eta(k)) 10^{-(k+1)!}$. Show that the map $\eta \mapsto s_\eta$ is continuous and injective, and that every s_η is irrational. (Hint: the decimal expansion of s_η has its nonzero digits exactly at the positions $(k + 1)!$.)
 - (b) Let $p_\eta \in S_\varphi(\mathbb{Q})$ be the type with $D_{p_\eta} = \{b \in \mathbb{Q} : b > s_\eta\}$. Show that the map $\eta \mapsto p_\eta$ is a homeomorphism of 2^ω onto a closed subset of $S_\varphi(\mathbb{Q})$.
 - (c) In the proof of Theorem 2.5.1, take $M = \mathbb{Q}$ and use the tree of Exercise 9 in place of the one obtained there. Show that the set D of that proof is the whole space $S_\varphi(\mathbb{Q})$, and that the map from D onto 2^ω sends the two types with $D_p = \{b : b \geq \frac{1}{3}\}$ and $D_p = \{b : b > \frac{1}{3}\}$ to the same sequence. (Hint: $\frac{1}{3}$ has binary expansion $0.0101\dots$) Thus the set D need not itself be a copy of 2^ω , which is why the proof appeals to the Cantor–Bendixson theorem.
15. (*) (Section 2.5.) Assume that L is countable, and let M be a countable model of T . Show that $S_\varphi(M)$ is uncountable if and only if M contains a tree of height ω for φ (see Exercise 8). (Hint: by the Cantor–Bendixson theorem, an uncountable compact metrizable space contains a nonempty closed subset P without isolated points. Choose the nodes b_σ by recursion on $|\sigma|$, so that the clopen set of types containing $\varphi(x, b_{\sigma|k})^{\sigma(k)}$ for all $k < |\sigma|$ meets P ; two distinct types in P differ on some instance $\varphi(x, b)$.) Deduce from Exercise 8(a) a second proof of the implication (1) $\Rightarrow$ (2) of Theorem 2.5.1: if some countable model has uncountably many φ -types, then the formula φ encodes binary trees of arbitrarily large finite height, and hence is unstable. If Corollary 3.1.6 is used in place of Theorem 2.4.3, this proof avoids type definability altogether.

2.7 HISTORICAL REMARKS

✂ **Section 2.2.** Definability of φ -types over models for formulas without the order property is due to Shelah [She71, She90]. The sharper form proved here, in which the definition is a positive Boolean combination of instances of φ itself, goes back to Hrushovski and Pillay [HP94], as Pillay points out in [Pil96, Chapter 1]. Pillay [Pil17, Remark 2.4] derives the positive form directly from Grothendieck’s double limit criterion (see Theorem A.1.3 below) by a compactness argument. Positivity matters in applications: a quantifier-free formula receives quantifier-free definitions.

- ⌘ **Section 2.3.** Binary trees enter stability theory through Shelah's local rank $R(-, \varphi, 2)$. In [She71, Lemma 2.7] the rank of a type is computed from the consistency of families of formulas indexed by the branches of binary trees of height n . The formulas of infinite 2-rank are thus exactly those that encode binary trees of arbitrarily large finite height (see Definition 2.1.2), and Shelah showed that these are exactly the formulas with the order property [She71, She90]. Quantitatively, the two directions differ. A ladder of length exponential in n yields a tree of height n (see Proposition 2.3.1). In the other direction, a tree yields only a ladder of length logarithmic in the height of the tree; see Section 3.1.
- ⌘ **Section 2.4.** The construction of this section is the local form of Shelah's counting argument [She71]. A tree of infinite height gives $2^{\aleph_0}$ φ -types over a countable set of parameters, which is incompatible with definability of all φ -types. This is also the origin of the *stability spectrum*: a theory with the order property has more than λ types over some set of size λ , for every $\lambda \geq |T|$, so it is λ -stable for no λ [She90].

Chapter 3

More stability, trees, and definability

In Chapter 2 we established the equivalences of Theorem 2.1.3 by the most direct route to the definability theorem. In that route, the way from tall trees back to the order property goes through non-definable types and a count of types over large models. In this chapter we develop the quantitative side of the theory, which that route leaves out. In Section 3.1 we show by a direct combinatorial argument, based on Ramsey’s theorem and carried out inside a single model, how tall a tree must be to contain a ladder of a given length. In Section 3.2 we present the classical proof of definability of types, which measures types by the height of the trees that can be grown above them and produces the definitions explicitly from this rank, without passing through the order property. The conventions are those of Chapter 2.

Roadmap

In Chapter 2 we settled *whether*; in this chapter we ask *how much* and *how*. A proof by counting shows that a non-definable type exists without pointing to it, and a proof through the order property shows that a type is definable without saying what its definition is. Neither defect touches the truth of the theorems, but both matter to anyone who wishes to use them. In Section 3.1 we remove the first defect: there we show, by Ramsey’s theorem, that a tree of height $2^{\binom{2n}{n}} - 1$ in a model already contains a ladder of length n in the same model, so that trees and ladders are compared without leaving the model and with an explicit bound. In Section 3.2 we remove the second. Types are measured by the height of the trees that can be grown above them, and the definition of a type is read off from this rank. Nothing later in the book depends on this chapter. The reader whose interest is in the truth of the theorems may omit it; the reader whose interest is in their mechanism will not wish to.

3.1 TREES GIVE THE ORDER PROPERTY: A QUANTITATIVE PROOF BY RAMSEY’S THEOREM

Theorem 2.4.3 was derived from the definability theorem; here we give a direct and quantitative proof of the implication “tall trees $\Rightarrow$ order property”. The proof is pure finite combinatorics, the Boolean case of a Ramsey-type argument due to Hodges, and we follow the presentation of Conant and Terry [CT26]. Finding a ladder inside a tree requires passing to smaller trees embedded in it, and the first definition says what such an embedding is.

Definition 3.1.1. Let $s, t \geq 1$ be given.

- (a) A *node embedding* is a map $\iota: 2^{<s} \rightarrow 2^{<t}$ such that $\rho \frown \langle u \rangle \sqsubseteq \tau$ implies $\iota(\rho) \frown \langle u \rangle \sqsubseteq \iota(\tau)$.
- (b) A *full embedding* is a map $\iota: 2^{\leq s} \rightarrow 2^{\leq t}$ sending $2^{<s}$ into $2^{<t}$ and 2^s into 2^t , with the same property for all $\rho, \tau \in 2^{\leq s}$.

Lemma 3.1.2. *Every node embedding $\iota: 2^{<s} \rightarrow 2^{<t}$ extends to a full embedding. If (b_σ) , (a_η) is a tree of height t in M and ι is a full embedding $2^{<s} \rightarrow 2^{<t}$, then the families $(b_{\iota(\sigma)} : \sigma \in 2^{<s})$ and $(a_{\iota(\eta)} : \eta \in 2^s)$ form a tree of height s in M .*

Proof. For $\eta \in 2^s$, let $\iota(\eta)$ be any element of 2^t extending $\iota(\eta|(s-1)) \frown \langle \eta(s-1) \rangle$. Suppose that $\rho \frown \langle u \rangle \sqsubseteq \eta$. If $\rho = \eta|(s-1)$, the requirement holds by construction. Otherwise we have $\rho \frown \langle u \rangle \sqsubseteq \eta|(s-1)$, and therefore $\iota(\rho) \frown \langle u \rangle \sqsubseteq \iota(\eta|(s-1)) \sqsubseteq \iota(\eta)$. The second statement is immediate, since whether $\varphi(a_\eta, b_\sigma)$ holds depends only on the direction u with $\sigma \frown \langle u \rangle \sqsubseteq \eta$, and full embeddings preserve this direction. $\odot$

To find a ladder, the proof of Proposition 3.1.5 below splits the nodes of a tree into two classes and needs a tall tree embedded in one of them. The next lemma, a Ramsey theorem for node embeddings, provides such an embedding; Figure 3.1 illustrates it for $t_0 = t_1 = 2$. In its proof and in the rest of this chapter, for $u \in \{0, 1\}$ and a finite 0–1 sequence σ , we write $\langle u \rangle \frown \sigma$ for the sequence obtained by putting u in front of σ .




FIGURE 3.1. The nodes $2^{<3}$, growing upward, coloured with two colours C_0 (pink) and C_1 (purple). The dashed arrows show a node embedding $\iota: 2^{<2} \rightarrow 2^{<3}$ into C_1 , with $\iota(\emptyset) = \emptyset$, $\iota(0) = 00$ and $\iota(1) = 1$: the image of the left child lies above the image of the root in the left direction, and the image of the right child in the right direction, but the map ι need not send children to children. This is the conclusion of Lemma 3.1.3 for $t_0 = t_1 = 2$ and $t = 3$. The colour C_0 admits no such embedding here.

Lemma 3.1.3 (Ramsey lemma for trees). *Let $t_0, t_1 \geq 1$ be given, let $t = t_0 + t_1 - 1$, and suppose that $2^{<t} = C_0 \cup C_1$. Then for some $u \in \{0, 1\}$ there is a node embedding $\iota: 2^{<t_u} \rightarrow 2^{<t}$ with image contained in C_u .*

Proof. Induction on t . If $t = 1$, then we have $t_0 = t_1 = 1$, and the map $\emptyset \mapsto \emptyset$ works for the u with $\emptyset \in C_u$. Now suppose that $t > 1$, and choose w with $\emptyset \in C_w$. If $t_w = 1$, the map $\emptyset \mapsto \emptyset$ works, so assume that $t_w \geq 2$. For $v \in \{0, 1\}$ and $i \in \{0, 1\}$ let $C_i^v = \{\sigma \in 2^{<t-1} : \langle v \rangle \frown \sigma \in C_i\}$. Then the union $C_0^v \cup C_1^v$ covers $2^{<t-1}$, and we have $(t_w - 1) + t_{1-w} - 1 = t - 1$. By induction, for each v either

- (i) there is a node embedding $\chi: 2^{<t_{1-w}} \rightarrow C_{1-w}^v$, in which case the map $\sigma \mapsto \langle v \rangle \frown \chi(\sigma)$ is a node embedding $2^{<t_{1-w}} \rightarrow 2^{<t}$ with image in C_{1-w} , and we are done; or
- (ii) there is a node embedding $\chi_v: 2^{<t_w-1} \rightarrow C_w^v$.

If (ii) holds for both v , define $\iota: 2^{<t_w} \rightarrow 2^{<t}$ by $\iota(\emptyset) = \emptyset$ and $\iota(\langle v \rangle \frown \sigma) = \langle v \rangle \frown \chi_v(\sigma)$. The image of ι lies in C_w . Suppose that $\rho \frown \langle u \rangle \sqsubseteq \tau$. If $\rho = \emptyset$, then we have $\tau = \langle u \rangle \frown \sigma$ and $\emptyset \frown \langle u \rangle \sqsubseteq \iota(\tau)$. If $\rho = \langle v \rangle \frown \rho'$, then we have $\tau = \langle v \rangle \frown \tau'$ with $\rho' \frown \langle u \rangle \sqsubseteq \tau'$, and therefore $\chi_v(\rho') \frown \langle u \rangle \sqsubseteq \chi_v(\tau')$. In both cases we have $\iota(\rho) \frown \langle u \rangle \sqsubseteq \iota(\tau)$, so the map ι is a node embedding. $\odot$

The ladders found inside a tree come in two orientations, according to the side of the split on which the Ramsey lemma lands, and the induction keeps track of both.

Definition 3.1.4. Let $(b_\sigma)_{\sigma \in 2^{<t}}$, $(a_\eta)_{\eta \in 2^t}$ be a tree in M . An *up-ladder of length k* in this tree consists of nodes $\sigma_1, \dots, \sigma_k$ and leaves $\eta_1, \dots, \eta_k$ such that $M \models \varphi(a_{\eta_p}, b_{\sigma_q}) \iff p \leq q$. A *down-ladder of length ℓ* is defined in the same way with “ $p > q$ ” in place of “ $p \leq q$ ”.

The height needed to force a ladder is governed by a familiar recursion. Define $B(k, \ell)$ for $k, \ell \geq 1$ by $B(k, 1) = B(1, \ell) = 1$ and $B(k, \ell) = B(k-1, \ell) + B(k, \ell-1) + 1$ for $k, \ell \geq 2$. By induction on $k + \ell$, using Pascal’s rule $\binom{k+\ell-3}{k-2} + \binom{k+\ell-3}{k-1} = \binom{k+\ell-2}{k-1}$, we have $B(k, \ell) = 2\binom{k+\ell-2}{k-1} - 1$. The recursion is the Erdős–Szekeres recursion for the Ramsey numbers $R(k, \ell)$, and the bound mirrors $R(k, \ell) \leq \binom{k+\ell-2}{k-1}$. The proof below follows the Erdős–Szekeres pattern as well: we fix a leaf η_* and split the nodes according to whether a_{η_*} satisfies the corresponding instance of φ . Lemma 3.1.3 finds a large subtree whose nodes all lie on one side of the split, induction finds a ladder inside that subtree, and the pair formed by η_* and the root of the subtree extends the ladder by one rung.

Proposition 3.1.5. *For all $k, \ell \geq 1$, every tree of height $B(k, \ell)$ for φ in a model of T contains an up-ladder of length k or a down-ladder of length ℓ .*

Proof. Induction on $k + \ell$. In a tree of height 1, the root $\emptyset$ with the leaf $\langle 1 \rangle$ is an up-ladder of length 1, and the root with the leaf $\langle 0 \rangle$ is a down-ladder of length 1. This settles the case $k = 1$ or $\ell = 1$.

Let $k, \ell \geq 2$ be given, let $t = B(k, \ell)$, and let $(b_\sigma), (a_\eta)$ be a tree of height t in M . Fix a leaf $\eta_* \in 2^t$ and put

$$P^+ = \{\sigma \in 2^{<t} : M \models \varphi(a_{\eta_*}, b_\sigma)\}, \quad P^- = 2^{<t} \setminus P^+.$$

Let $t^+ = B(k-1, \ell) + 1$ and $t^- = B(k, \ell-1) + 1$, so that $t^+ + t^- - 1 = t$. By Lemma 3.1.3, one of the two cases below occurs.

Case +: there is a node embedding $\iota: 2^{<t^+} \rightarrow 2^{<t}$ with image in P^+ . Extend ι to a full embedding $\hat{\iota}$ (see Lemma 3.1.2) and set $\iota'(\tau) = \hat{\iota}(\langle 0 \rangle \cap \tau)$ for $\tau \in 2^{\leq t^+ - 1}$. Then the map ι' is a full embedding of height $B(k-1, \ell)$. Put $\sigma_0 = \iota(\emptyset)$. Since $\emptyset \cap \langle 0 \rangle \subseteq \langle 0 \rangle \cap \tau$, we have $\sigma_0 \cap \langle 0 \rangle \subseteq \iota'(\tau)$ for every τ . By Lemma 3.1.2, the embedding ι' induces a tree of height $B(k-1, \ell)$ in M . By induction, the induced tree contains a down-ladder of length ℓ , which is then a down-ladder in the original tree, or an up-ladder $(\sigma'_p, \eta'_p)_{p=1}^{k-1}$. In the latter case, let

$$(\sigma_1, \eta_1) = (\sigma_0, \eta_*), \quad (\sigma_{p+1}, \eta_{p+1}) = (\iota'(\sigma'_p), \iota'(\eta'_p)) \quad (1 \leq p < k).$$

We check that $\varphi(a_{\eta_p}, b_{\sigma_q}) \iff p \leq q$.

- (a) For $p = 1$, every σ_q is in the image of ι , which lies in P^+ , so the formula $\varphi(a_{\eta_*}, b_{\sigma_q})$ holds, and indeed $1 \leq q$.
- (b) For $p \geq 2$ and $q = 1$, we have $\sigma_0 \cap \langle 0 \rangle \subseteq \eta_p$, so the formula $\neg \varphi(a_{\eta_p}, b_{\sigma_0})$ holds, and indeed $p > 1$.
- (c) For $p, q \geq 2$: this is the up-ladder property of the smaller tree.

This is an up-ladder of length k .

Case -: there is a node embedding $\iota: 2^{<t^-} \rightarrow 2^{<t}$ with image in P^- . Proceed symmetrically, with $\iota'(\tau) = \hat{\iota}(\langle 1 \rangle \cap \tau)$, a full embedding of height $B(k, \ell-1)$ whose leaves all extend $\sigma_0 \cap \langle 1 \rangle$. By induction the induced tree contains an up-ladder of length k , and we are done, or a down-ladder $(\sigma'_p, \eta'_p)_{p=1}^{\ell-1}$. Prepend (σ_0, η_*) as before.

- (a) For $p = 1$, the formula $\neg\varphi(a_{\eta_*}, b_{\sigma_q})$ holds for all q , because the nodes lie in P^- , and indeed $1 \not\preceq q$.
- (b) For $p \geq 2$ and $q = 1$, the formula $\varphi(a_{\eta_p}, b_{\sigma_0})$ holds, because $\sigma_0 \frown \langle 1 \rangle \sqsubseteq \eta_p$, and indeed $p > 1$.

This is a down-ladder of length ℓ . 😊

Corollary 3.1.6. *If a model M of T contains a tree of height $B(n+1, n+1)$ for φ , then it contains an n -ladder for φ .*

Proof. By Proposition 3.1.5, the model M contains an up-ladder or a down-ladder of length $n+1$. An up-ladder of length $n+1$ is an $(n+1)$ -ladder, and so contains an n -ladder. A down-ladder of length $n+1$ yields an n -ladder by Remark 1.1.5. 😊

3.2 THE CLASSICAL PROOF OF DEFINABILITY, VIA RANKS

In Proposition 3.2.3 below we give the classical argument that, in the absence of tall trees, every φ -type is definable. Together with Proposition 2.3.1, the argument gives another route to (1) $\Rightarrow$ (3) in Theorem 1.1.6, at the cost of positivity: the definitions obtained in this section are not positive Boolean combinations of instances of φ , since they contain quantifiers and negations that do not occur in φ (see Lemma 3.2.2(iv)). The idea is to measure each finite piece of a type by the height of the tallest tree that can be grown above the piece, and to define the type by a piece of minimal height.

Until the end of the proof of Proposition 3.2.3, we assume that there is n such that no model of T contains a tree of height n , and we fix a model M of T . By a φ -conjunction over M we mean a finite conjunction $\theta(x)$ of formulas $\varphi(x, b)^u$ with $b \in M$.

Definition 3.2.1. For a φ -conjunction θ over M and $m < \omega$, we write $R(\theta) \geq m$ if there are $b_\sigma \in M$ ($\sigma \in 2^{<m}$) such that for every $\eta \in 2^m$, we have

$$M \models \exists x \left(\theta(x) \wedge \bigwedge_{k < m} \varphi(x, b_{\eta|k})^{\eta(k)} \right).$$

Thus we have $R(\theta) \geq 0$ if and only if $M \models \exists x \theta(x)$. When θ is satisfiable in M , we write $R(\theta)$ for the largest m with $R(\theta) \geq m$.

In Lemma 3.2.2(i) below we show that this largest m exists.

Lemma 3.2.2.

- (i) *The inequality $R(\theta) < n$ holds for every θ ; in particular, the rank $R(\theta)$ is well defined whenever θ is satisfiable in M .*
- (ii) *If $\theta' = \theta \wedge \chi$, then we have $R(\theta') \leq R(\theta)$.*
- (iii) *If $b \in M$ satisfies $R(\theta \wedge \varphi(x, b)^u) \geq m$ for both $u \in \{0, 1\}$, then we have $R(\theta) \geq m + 1$.*
- (iv) *For each θ and m there is an $L(M)$ -formula $\psi_{\theta, m}(y)$ such that for all $b \in M$ we have $R(\theta \wedge \varphi(x, b)) \geq m \iff M \models \psi_{\theta, m}(b)$.*

Proof. (i) If $R(\theta) \geq n$, witnessed by (b_σ) , then choosing a witness $a_\eta \in M$ for each $\eta \in 2^n$ gives a tree of height n in M , which is impossible.

(ii) A witness for $R(\theta') \geq m$ is a witness for $R(\theta) \geq m$.

(iii) Let $(b_\sigma^u)_{\sigma \in 2^{<m}}$ witness $R(\theta \wedge \varphi(x, b)^u) \geq m$. Put $b_\emptyset = b$ and $b_{\langle u \rangle \frown \sigma} = b_\sigma^u$. A branch of length $m+1$ has the form $\langle u \rangle \frown \eta$ with $\eta \in 2^m$, and the corresponding formula is $\theta \wedge \varphi(x, b)^u \wedge \bigwedge_{k < m} \varphi(x, b_{\eta|k}^u)^{\eta(k)}$, which is satisfiable in M .

(iv) Take

$$\psi_{\theta, m}(y) := \exists (z_\sigma)_{\sigma \in 2^{<m}} \bigwedge_{\eta \in 2^m} \exists x \left(\theta(x) \wedge \varphi(x, y) \wedge \bigwedge_{k < m} \varphi(x, z_{\eta|k})^{\eta(k)} \right),$$

whose parameters are those of θ . ☺

Proposition 3.2.3. *Let M be a model of T . If for some n no model of T contains a tree of height n , then every $p \in S_\varphi(M)$ is definable.*

Proof. Let Φ_p be the set of finite conjunctions of members of p . Each $\theta \in \Phi_p$ is satisfiable in M , because the sentence $\exists x \theta$ is an $L(M)$ -sentence consistent with the elementary diagram of M , hence true in M . By Lemma 3.2.2(i), choose $\theta \in \Phi_p$ with $m = R(\theta)$ minimal. We show that $\psi_{\theta, m}$ is a φ -definition of p , that is, that for $b \in M$, we have

$$\varphi(x, b) \in p \iff R(\theta \wedge \varphi(x, b)) \geq m.$$

If $\varphi(x, b) \in p$, then the conjunction $\theta \wedge \varphi(x, b)$ belongs to Φ_p , so we have $R(\theta \wedge \varphi(x, b)) \geq m$ by minimality. Conversely, suppose that $\varphi(x, b) \notin p$. Then the formula $\neg \varphi(x, b)$ belongs to p , so we have $R(\theta \wedge \neg \varphi(x, b)) \geq m$ by minimality. If also $R(\theta \wedge \varphi(x, b)) \geq m$, then Lemma 3.2.2(iii) would give $R(\theta) \geq m+1$, a contradiction. The equivalence together with Lemma 3.2.2(iv) shows that p is definable. ☺



Remark 3.2.4.

- (a) The passages between ladders and trees are effective. An up- or down-ladder of length $n+1$ appears in any tree of height $B(n+1, n+1) = 2^{\binom{2n}{n}} - 1$ (see Proposition 3.1.5). Conversely, a ladder of length $2^{n+1} - 1$ yields a tree of height n (see Proposition 2.3.1).
- (b) Theorem 2.2.1 uses only that p is finitely satisfiable in M (see Lemma 1.1.2). Consequently, the theorem applies to any φ -type over a set A that is finitely satisfiable in A (see Definition 3.2.5(a) below). The definition given by the theorem has its parameters in the finite set F , whereas the rank argument of Proposition 3.2.3 yields a definition whose parameters are those of the formula θ of minimal rank.
- (c) Without finite satisfiability, definitions over the parameter set cannot be expected in general, and the restriction to types over models is essential. Definition 3.2.5 and Example 3.2.6 below make this precise.

Definition 3.2.5. Let N be a model of T , and let $A \subseteq N$ and $p \in S_\varphi(N)$ be given.

- (a) A *complete φ -type over A* is defined as in Definition 1.1.1, with b ranging over A instead of M ; such a type is *finitely satisfiable in A* if every finite subset of the type is realized by an element of A .
- (b) The type p is *definable over A* if p has a φ -definition that is an $L(A)$ -formula.
- (c) The type p is *definable almost over A* if p has a φ -definition of the form $\psi(y, c)$, with $\psi(y, z)$ an L -formula and c a tuple of N , such that only finitely many distinct subsets of N are of the form $\psi(N, c')$ with $c' \in N$ realizing the same complete type over A as c . In other words, the definition of p has only finitely many conjugates over A .

Definable over A implies definable almost over A : take c to enumerate the parameters from A , so that $c' = c$ is the only possibility. The converse fails, as the next example shows with two equivalence classes.

Example 3.2.6. Let T be the theory of an equivalence relation E with exactly two classes, both infinite, let $\varphi(x, y)$ be $E(x, y)$, and let $A = \emptyset$. Let N be a model of T with classes C_1 and C_2 , and let $p \in S_\varphi(N)$ be the φ -type saying that x lies in C_1 , that is, the type that contains $E(x, b)$ if and only if $b \in C_1$. For any $c \in C_1$, the formula $E(y, c)$ is a φ -definition of p . All elements of N have the same complete type over $\emptyset$, so the sets $E(N, c')$ are exactly C_1 and C_2 . Hence the type p is definable almost over $\emptyset$. The type p is not definable over $\emptyset$: since N has an automorphism swapping C_1 and C_2 , every L -formula $\psi(y)$ without parameters defines $\emptyset$ or N , never C_1 . Note that p is not finitely satisfiable in $\emptyset$.

3.3 EXERCISES

The exercises follow the sections of the chapter, and we write T_E for the theory of an equivalence relation E with infinitely many classes, all of them infinite, as in the exercises of Section 2.6. The mark (*) indicates the exercises that ask for more than a direct application of the text.

1. (Section 3.1.) Let $s, t \geq 1$ be given, and let $\iota: 2^{<s} \rightarrow 2^{<t}$ be a node embedding.
 - (a) Show that $|\iota(\sigma)| \geq |\sigma|$ for every $\sigma \in 2^{<s}$. Show that the map ι is injective and sends incomparable sequences to incomparable sequences, and that, for all $\rho, \tau \in 2^{<s}$ and $u \in \{0, 1\}$, we have $\rho \frown \langle u \rangle \sqsubseteq \tau$ if and only if $\iota(\rho) \frown \langle u \rangle \sqsubseteq \iota(\tau)$. (Hint: if ρ and τ are incomparable, let π be their longest common initial segment, and apply the definition to π . For the converse implication, treat separately the cases $\tau \sqsubseteq \rho$, $\rho \frown \langle 1 - u \rangle \sqsubseteq \tau$, and ρ, τ incomparable.)
 - (b) Show that there is no node embedding $2^{<s} \rightarrow 2^{<t}$ when $t < s$, and that the identity is the only node embedding $2^{<s} \rightarrow 2^{<s}$.
 - (c) Show that the map sending σ to $\langle 0 \rangle \frown \sigma$ is a node embedding $2^{<s} \rightarrow 2^{<s+1}$, and that so is the map δ which repeats every entry, given by $\delta(\sigma) = \langle \sigma(0), \sigma(0), \sigma(1), \sigma(1), \dots \rangle$ with $|\delta(\sigma)| = 2|\sigma|$, as a map $2^{<s} \rightarrow 2^{<2s-1}$. Show that, for $s \geq 2$, the map sending σ to $\sigma \frown \langle 0 \rangle$ is not a node embedding $2^{<s} \rightarrow 2^{<s+1}$.
 - (d) For $s, t \geq 1$, let $N(s, t)$ be the number of node embeddings $2^{<s} \rightarrow 2^{<t}$, and put $N(s, 0) = 0$. Show that $N(1, t) = 2^t - 1$ and that

$$N(s, t) = \sum_{d < t} 2^d N(s - 1, t - d - 1)^2 \quad (s \geq 2).$$

(Hint: classify the node embeddings by $d = |\iota(\emptyset)|$, and write $\iota(\langle v \rangle \frown \sigma) = \iota(\emptyset) \frown \langle v \rangle \frown \chi_v(\sigma)$, as in the proof of Lemma 3.1.3; check that χ_0 and χ_1 are arbitrary node embeddings $2^{<s-1} \rightarrow 2^{<t-d-1}$.) Deduce that $N(2, 3) = 11$, $N(2, 4) = 71$ and $N(3, 4) = 123$, and list the eleven node embeddings $2^{<2} \rightarrow 2^{<3}$.

2. (Section 3.1.)

- (a) Let $s, t \geq 1$ be given, let $\iota: 2^{<s} \rightarrow 2^{<t}$ be a node embedding, and let $\hat{\iota}: 2^{\leq s} \rightarrow 2^{\leq t}$ be a map that extends ι and sends 2^s into 2^t . Show that the map $\hat{\iota}$ is a full embedding if and only if $\iota(\eta|(s-1)) \cap \langle \eta|(s-1) \rangle \subseteq \hat{\iota}(\eta)$ for every $\eta \in 2^s$. Deduce that the number of full embeddings extending ι is $\prod_{\eta \in 2^s} 2^{t-1-|\iota(\eta|(s-1))|}$.
- (b) Let $\iota: 2^{<2} \rightarrow 2^{<3}$ be the node embedding of Figure 3.1, given by $\iota(\emptyset) = \emptyset$, $\iota(0) = 00$ and $\iota(1) = 1$. Show that ι has exactly four extensions to a full embedding $2^{\leq 2} \rightarrow 2^{\leq 3}$, and list them.
- (c) Work in $(\mathbb{Q}, <)$, with $\varphi(x, y)$ the formula $x < y$. For $\sigma \in 2^{\leq 3}$ put $b_\sigma = 16d(\sigma)$ if $|\sigma| < 3$ and $a_\sigma = 16d(\sigma)$ if $|\sigma| = 3$, with d as in the proof of Proposition 2.3.1. Explicitly, the nodes are $b_\emptyset = 8$, $b_0 = 12$, $b_1 = 4$, $b_{00} = 14$, $b_{01} = 10$, $b_{10} = 6$, $b_{11} = 2$, and the leaves are $a_{000} = 15$, $a_{001} = 13$, $a_{010} = 11$, $a_{011} = 9$, $a_{100} = 7$, $a_{101} = 5$, $a_{110} = 3$, $a_{111} = 1$. Show, by the claim in the proof of Proposition 2.3.1 or directly, that these elements form a tree of height 3 for φ . Let $\hat{\iota}$ be the extension of the node embedding of part (b) with $\hat{\iota}(00) = 000$, $\hat{\iota}(01) = 001$, $\hat{\iota}(10) = 100$ and $\hat{\iota}(11) = 110$. Compute the tree of height 2 given by Lemma 3.1.2, and check the tree conditions directly. Show that, if the value $\hat{\iota}(00)$ is replaced by 010, the resulting families are no longer a tree.
3. (Section 3.1.) For a finite 0–1 sequence σ and $v \in \{0, 1\}$, let $\#_v(\sigma)$ be the number of $k < |\sigma|$ with $\sigma(k) = v$, and for $k < \omega$ let 0^k and 1^k be the constant sequences of length k .
- (a) In the colouring of Figure 3.1, where $C_1 = \{\emptyset, 1, 00\}$ and $C_0 = \{0, 01, 10, 11\}$, show that the embedding drawn there is the only node embedding $2^{<2} \rightarrow 2^{<3}$ with image in C_1 , and that no node embedding $2^{<2} \rightarrow 2^{<3}$ has image in C_0 .
- (b) Let $\iota: 2^{<s} \rightarrow 2^{<t}$ be a node embedding. Show that $\#_v(\iota(\sigma)) \geq \#_v(\sigma)$ for every $\sigma \in 2^{<s}$ and $v \in \{0, 1\}$.
- (c) Let $t_0, t_1 \geq 1$ be given with $t_0 + t_1 \geq 3$, put $t = t_0 + t_1 - 2$, and let $C_0 = \{\sigma \in 2^{<t} : \#_0(\sigma) \leq t_0 - 2\}$ and $C_1 = 2^{<t} \setminus C_0$. Show that no node embedding $2^{<t_0} \rightarrow 2^{<t}$ has image in C_0 , and that no node embedding $2^{<t_1} \rightarrow 2^{<t}$ has image in C_1 . (Hint: consider the images of 0^{t_0-1} and of 1^{t_1-1} , and note that every $\sigma \in C_1$ satisfies $\#_1(\sigma) \leq t_1 - 2$.) Conclude that the number $t_0 + t_1 - 1$ in Lemma 3.1.3 cannot be replaced by a smaller one.
4. (Section 3.1.) Let $B(k, \ell)$ be defined by the recursion of Section 3.1.
- (a) Compute $B(k, \ell)$ for $1 \leq k, \ell \leq 5$ from the recursion, and check the values against the closed form $2^{\binom{k+\ell-2}{k-1}} - 1$. Compute the bound $B(n+1, n+1) = 2^{\binom{2n}{n}} - 1$ of Corollary 3.1.6 for $1 \leq n \leq 5$; the values are 3, 11, 39, 139 and 503.
- (b) Show that the numbers $B(k, \ell) + 1$ satisfy Pascal's rule $B(k, \ell) + 1 = (B(k-1, \ell) + 1) + (B(k, \ell-1) + 1)$ for $k, \ell \geq 2$, with the value 2 when $k = 1$ or $\ell = 1$, and derive the closed form from this without guessing it. Recall that the Ramsey numbers satisfy $R(k, 1) = R(1, \ell) = 1$ and the Erdős–Szekeres inequality $R(k, \ell) \leq R(k-1, \ell) + R(k, \ell-1)$. Show that the numbers $\binom{k+\ell-2}{k-1}$ satisfy the corresponding equality with the same boundary values, so that $B(k, \ell) + 1$ is exactly twice the Erdős–Szekeres bound for $R(k, \ell)$. Locate the summand 1 of the recursion for B in the proof of Proposition 3.1.5.

- (c) Show that every tree of height $B(n, n+1)$ for φ in a model M of T yields an n -ladder for φ in M , and that $B(n, n+1) = \binom{2n}{n} - 1$. Thus the height in Corollary 3.1.6 may be lowered to less than half of its value; for $1 \leq n \leq 5$ the new values are 1, 5, 19, 69 and 251.
- (d) Show that $4^n/(2n+1) \leq \binom{2n}{n} < 4^n$ for $n \geq 1$. Deduce that a tree of height $h \geq 4$ for φ in M yields an n -ladder in M with $n = \lfloor \log_4 h \rfloor$, and that the bound $B(n+1, n+1)$ grows exponentially in n . Compare with Proposition 2.3.1, by which a ladder of length L yields a tree of height $\lfloor \log_2(L+1) \rfloor - 1$: in both directions the passage loses a logarithm.
5. (Section 3.1.) Work with the tree of height 3 for $x < y$ in $(\mathbb{Q}, <)$ of Exercise 2(c), and carry out the proof of Proposition 3.1.5 with $k = \ell = 2$, for which $B(2, 2) = 3$ and $t^+ = t^- = 2$. In parts (a) and (b), whenever the proof needs a node embedding $2^{<2} \rightarrow 2^{<3}$, take the inclusion map, and extend it to the full embedding $\hat{i}$ given by $\hat{i}(\eta) = \eta \frown \langle 0 \rangle$ for $\eta \in 2^2$.
- (a) Take $\eta_* = 111$. Show that $P^+ = 2^{<3}$, so that the proof is in Case +. Compute the map ι' and the induced tree of height $B(1, 2) = 1$, whose node is $b_0 = 12$ and whose leaves are $a_{000} = 15$ and $a_{010} = 11$. Find the up-ladder of length 1 given by the base case of the induction, and show that the proof returns the up-ladder with $\sigma_1 = \emptyset$, $\eta_1 = 111$, $\sigma_2 = 0$ and $\eta_2 = 010$. Check directly that the elements $a_{111} = 1$, $a_{010} = 11$ and $b_\emptyset = 8$, $b_0 = 12$ form a 2-ladder.
- (b) Take $\eta_* = 000$. Show that $P^+ = \emptyset$, so that the proof is in Case -, and that it returns the down-ladder with $\sigma_1 = \emptyset$, $\eta_1 = 000$, $\sigma_2 = 1$ and $\eta_2 = 100$, with elements 15, 7 and 8, 4. Check this down-ladder directly, and determine the ladder that Remark 1.1.5 extracts from it.
- (c) For each of the eight leaves η_* , determine the set P^+ , and decide whether a node embedding $2^{<2} \rightarrow 2^{<3}$ with image in P^+ , or one with image in P^- , exists. Show that the first exists exactly for $\eta_* \in \{011, 100, 101, 110, 111\}$ and the second exactly for $\eta_* \in \{000, 001, 010, 011, 100\}$, in accordance with Lemma 3.1.3.
6. (*) (Section 3.1.) For $m \geq 2$, let $H(m)$ be the least h with the following property: for every theory T , every formula $\varphi(x, y)$ and every model M of T , if M contains a tree of height h for φ , then the model M contains an m -ladder for φ . By Exercise 4(c), we have $H(m) \leq \binom{2m}{m} - 1$.
- (a) Let $n \geq 1$ be given, let T be the theory of infinite sets in the empty language, and let $\varphi(x_1, \dots, x_n; y)$ be the formula $\bigvee_{i=1}^n y = x_i$, which says that y is an entry of the tuple x . Show that every model of T contains an n -ladder for φ and no $(n+1)$ -ladder. (Hint: in an $(n+1)$ -ladder, the elements $b_1, \dots, b_{n+1}$ are pairwise distinct entries of a_1 .)
- (b) In the setting of part (a), show that every model of T contains a tree of height n for φ and no tree of height $n+1$. (Hint: in a tree of height h , let $j < k < h$ be given, and look at a leaf above $1^k \frown \langle 0 \rangle$ to see that $b_{1j} \neq b_{1k}$; all the nodes b_{1k} are entries of a_{1h} . For the construction, take the nodes pairwise distinct, and let the leaf a_η list the nodes b_σ with $\sigma \frown \langle 1 \rangle \subseteq \eta$, repeated as needed, or an element distinct from all the nodes if there are none.) For $n = 1$, compare with Exercise 2(b) of Section 2.6.

- (c) Deduce that $H(m) \geq m$. Let $(b_\sigma), (a_\eta)$ be a tree of height 2 for an arbitrary formula φ in a model M . Show that the elements a_{11}, a_{10} and $b_1, b_\emptyset$, in this order, form a 2-ladder. Conclude that $H(2) = 2$, whereas Exercise 4(c) gives the bound 5 and Corollary 3.1.6 the bound 11. (A finite search by computer shows that $H(3) = 3$ and $H(4) \geq 5$.)
- (d) In the setting of Exercise 3 of Section 2.6, the span formula over a finite field has trees of height n , and its longest ladders have length $n + 1$ (see the exercises of Section 1.3). Show that this example yields the inequality $H(n + 2) \geq n + 1$, which is weaker than the one of part (c). The structures $(\mathbb{N}, R_k)$ of Exercise 7 of Section 2.6 contain $(k + 1)$ -ladders, no $(k + 2)$ -ladder, and no tree of height $\lfloor \log_2(k + 2) \rfloor + 1$. Show that they yield only the inequality $H(m) \geq \lfloor \log_2 m \rfloor + 1$, which for $m \geq 3$ is weaker than the one of part (c), and explain why: their ladders are long compared with their trees. Show that they bear instead on the converse passage: for $n \geq 2$, a ladder of length $2^n - 2$ need not yield a tree of height n , so that Proposition 2.3.1 needs ladders of length at least $2^n - 1$ in general (see Exercise 7(d) of Section 2.6).
7. (Section 3.2.) Let M be a model of T . The condition $R(\theta) \geq m$ makes sense, by the same definition, for every $L(M)$ -formula $\theta(x)$, and without the standing assumption of Section 3.2. In this and the following exercises we use it in this generality, we write $R(\theta) = -1$ if θ is not satisfiable in M and $R(\theta) = \infty$ if $R(\theta) \geq m$ for every m , and we write R^M for R when the model must be indicated.
- (a) Show that $R(\theta) \geq m$ if and only if M contains a tree of height m for φ whose leaves all satisfy θ . In particular, the rank $R(x = x)$ is the largest height of a tree for φ in M , or ∞ if there is no largest height.
- (b) Show that, for each m , the condition $R(\theta) \geq m$ is expressed in M by an $L(M)$ -sentence whose parameters are those of θ . Deduce that $R^M(\theta) = R^N(\theta)$ whenever $M \preceq N$. Show that $R(\theta) \leq R(\theta')$ whenever $\theta(M) \subseteq \theta'(M)$; in particular, the rank $R(\theta)$ depends only on the set $\theta(M)$, and Lemma 3.2.2(ii) holds for every $L(M)$ -formula χ .
- (c) Show the converse of Lemma 3.2.2(iii): if $R(\theta) \geq m + 1$, then some $b \in M$ satisfies $R(\theta \wedge \varphi(x, b)^u) \geq m$ for both $u \in \{0, 1\}$. Deduce that, if $R(x = x)$ is finite, then every θ satisfiable in M satisfies

$$R(\theta) = 1 + \max_{b \in M} \min_{u \in \{0, 1\}} R(\theta \wedge \varphi(x, b)^u).$$

8. (Section 3.2.) Let M be an infinite set in the empty language, and let $\varphi(x, y)$ be the formula $x = y$. By Exercise 2(b) of Section 2.6, the standing assumption of Section 3.2 holds with $n = 2$. Use the conventions of Exercise 7.
- (a) Show that $R(x = x) = 1$, and that a φ -conjunction θ over M that is satisfiable in M has $R(\theta) = 0$ if it has a conjunct of the form $x = c$ and $R(\theta) = 1$ otherwise.
- (b) Let $a, a' \in M$ be distinct, and let θ be the formula $x = a \vee x = a'$. Show that $R(\theta) = 1$, although $R(x = a) = R(x = a') = 0$. Show that, for $b = a$, both formulas $\theta \wedge \varphi(x, b)^u$ have rank 0, and check Lemma 3.2.2(iii) in this case. Thus a single instance of φ can lower the rank on both sides.

- (c) For $c \in M$, let p_c be the type with $D_{p_c} = \{c\}$, and let p_∞ be the type with $D_{p_\infty} = \emptyset$ (see Exercise 1 of Section 1.3). Show that the least rank of a member of Φ_{p_c} is 0 and that of a member of Φ_{p_∞} is 1. For θ the formula $x = c$, write out the formula $\psi_{\theta,0}(y)$ of Lemma 3.2.2(iv) and show that it is equivalent to $y = c$. For θ the formula $x \neq c$, which belongs to p_∞ , show that

$$\psi_{\theta,1}(y) = \exists z (\exists x (x \neq c \wedge x = y \wedge x \neq z) \wedge \exists x (x \neq c \wedge x = y \wedge x = z)),$$

and that this formula is equivalent in M to $y \neq y$. Compare with the definition $d_F \varphi$ of Exercise 4(c) of Section 2.6, which for p_∞ is the empty disjunction, and compare the least ranks with the Cantor–Bendixson ranks of Exercise 12(a) of Section 2.6.

- (d) Do the same for a model M of T_E and the formula xEy : show that $R(x = x) = 1$, that a satisfiable φ -conjunction has rank 0 if it has a conjunct of the form xEc and rank 1 otherwise, that for θ the formula xEc the formula $\psi_{\theta,0}(y)$ is $\exists x (xEc \wedge xEy)$, which is equivalent to yEc in M , and that for θ the formula $\neg xEc$ the formula $\psi_{\theta,1}(y)$ is equivalent to $y \neq y$. Compare with Exercises 4(d) and 12(b) of Section 2.6.

In both settings, the definitions produced by Proposition 3.2.3 contain a quantifier and, for the type whose least rank is 1, the negation of an instance of φ , although each of them is equivalent to a positive Boolean combination of instances of φ .

9. (*) (Section 3.2.) Work in the setting of Exercise 3 of Section 2.6: let K be a finite field, let T be the theory of infinite vector spaces over K , let $n \geq 1$ be given, let $\varphi(x_1, \dots, x_n; y)$ say that y lies in the span of $x_1, \dots, x_n$, and let M be a model of T . Use the conventions of Exercise 7.

- (a) Let θ be the φ -conjunction $\bigwedge_{i < r} \varphi(x, c_i) \wedge \bigwedge_{j < s} \neg \varphi(x, d_j)$, and let U be the span of $c_0, \dots, c_{r-1}$. Show that θ is satisfiable in M if and only if $\dim U \leq n$ and no d_j lies in U , and that in this case $R(\theta) = n - \dim U$. (Hint: for the upper bound, let (b_σ) , (a_η) be a tree of height h whose leaves satisfy θ , and show, by looking at a leaf above $1^{k \smallfrown} \langle 0 \rangle$, that b_{1^k} does not lie in the span of U and the nodes b_{1^j} with $j < k$; then count the dimension of the span of a_{1^h} . For the lower bound, note that M has infinite dimension, and take the nodes linearly independent over the span of U and the d_j .)
- (b) Deduce that $R(x = x) = n$, the largest height of a tree found in Exercise 3 of Section 2.6. Thus the standing assumption of Section 3.2 holds with $n + 1$ in place of n , and the bound of Lemma 3.2.2(i) is attained.
- (c) Let $p \in S_\varphi(M)$ be given, so that D_p is a subspace of M of dimension at most n (see the exercises of Section 1.3). Show that the least rank of a member of Φ_p is $n - \dim D_p$, and that it is attained by $\theta_p = \varphi(x, 0) \wedge \bigwedge_{c \in C} \varphi(x, c)$ for any basis C of D_p . Show that this least rank equals the Cantor–Bendixson rank of p found in Exercise 13 of Section 2.6.
- (d) Show directly from part (a) that $\psi_{\theta_p, m}(y)$, with $m = n - \dim D_p$, defines D_p in M . For $n = 1$ and $D_p = \{0\}$, write out $\psi_{\theta_p, 1}(y)$ and show that it is equivalent to $y = 0$. Compare with the definition $\varphi(c, y)$ by a single instance in Exercise 4(e) of Section 2.6.

10. (Section 3.2.) Work in $(\mathbb{Q}, <)$, with $\varphi(x, y)$ the formula $x < y$, so that the standing assumption of Section 3.2 fails. Use the conventions of Exercise 7.

- (a) Show that every φ -conjunction θ over $\mathbb{Q}$ that is satisfiable in $\mathbb{Q}$ defines a set containing an open interval, and that $R(\theta) = \infty$. (Hint: the numbers $2^{m+1}d(\sigma)$ with $\sigma \in 2^{\leq m}$,

for d as in the proof of Proposition 2.3.1, form a tree of height m , and a tree for $x < y$ is carried to a tree by every order-preserving map.)

- (b) Let r be an irrational real number, and let p be the type with $D_p = \{b \in \mathbb{Q} : b > r\}$. Show that, for every $\theta \in \Phi_p$ and every m , there is a rational $b < r$ with $R(\theta \wedge \varphi(x, b)) \geq m$, so that the formula $\psi_{\theta, m}$ does not define p . Identify the step of the proof of Proposition 3.2.3 that fails. Recall from Exercise 3 of Section 1.3 that p has no definition at all.

11. (*) (Section 3.2.) Use the conventions of Exercise 7.

- (a) Show that the following are equivalent: for some n , no model of T contains a tree of height n for φ ; for every model M of T , the rank $R^M(x = x)$ is finite. (Hint: for the nontrivial implication, use Lemma 2.4.1 with $\mu = \omega$.) Show that, if T is complete, these conditions are also equivalent to the finiteness of $R^M(x = x)$ for a single model M of T . Show that completeness cannot be dropped, by taking for T the theory of linear orders and for φ the formula $x < y$.
- (b) Let M be the disjoint union of finite chains of Exercise 10 of Section 2.6, with $\varphi(x, y)$ the formula $x < y$. Show that $R(x = x) = \infty$ in M , although M contains no tree of height ω . Thus infinite rank is witnessed by finite trees.
- (c) Let $k \in \mathbb{N}$ be given, and let T_k be as in Exercise 7 of Section 2.6. Show that $R(x = x) = \lfloor \log_2(k+2) \rfloor$ in every model of T_k , so that the standing assumption of Section 3.2 holds with $n = \lfloor \log_2(k+2) \rfloor + 1$, although every model of T_k contains a $(k+1)$ -ladder.

3.4 HISTORICAL REMARKS

- ✂ **Section 3.1.** The Ramsey-type argument that extracts ladders from trees is due to Hodges, and appears in his book [Hod93]. It proves directly, without definability of types, that tall trees force the order property. Conant and Terry [CT26] gave the argument explicit bounds and extended it from $\{0, 1\}$ -valued formulas to real-valued functions; we follow their version. Bounds of this kind are what the applications of stability to finite combinatorics require, the regularity lemma for stable graphs among them.
- ✂ **Section 3.2.** The rank $R(\theta)$ is Shelah's local rank $R(\theta, \varphi, 2)$ [She71, She90]. The proof that a type of minimal rank is definable is the classical route to definability of types. It is older than the combinatorial and analytic proofs given elsewhere in this book, and it is the one found in most textbooks, for instance Pillay [Pil96] and Tent–Ziegler [TZ12]. Its definitions involve quantifiers; those of Section 2.2 do not.

Chapter 4

From stability to NIP

In a stable theory no formula can encode arbitrarily long ladders. Many of the best understood structures in mathematics fail this test for a simple reason: they carry a definable linear order. Dense linear orders, the ordered field of real numbers and, more generally, o-minimal structures are all unstable, and yet their definable sets are remarkably well behaved [PS86]. To account for this kind of tameness we introduce the *independence property*, a configuration much wilder than a ladder, and study the formulas and theories that avoid the independence property, called *NIP*. Since a ladder is one of the patterns that the independence property demands, every stable formula is NIP (see Proposition 4.1.2); the orders just mentioned are NIP as well (see Example 4.1.3 and [PS86]), so that NIP generalizes both stability and the tameness of order. In Section 4.1 we treat formulas, with the conventions of Section 1.1. In Section 4.2 we treat bounded functions, with the conventions of Section 1.2; there NIP becomes the existence of pointwise convergent subsequences, and the independence property a copy of $\beta\mathbb{N}$.

Roadmap

The definability theorem shows that stability is a condition of great power, and the examples show that it is also a condition of great severity. The real line with its order is as well understood as any structure in mathematics, and yet it is not stable, because an order is the very configuration that stability forbids. In this chapter we propose a weaker condition that admits the orders while still excluding what is truly wild. In Section 4.1 we define the independence property, under which a formula imposes every pattern on arbitrarily large finite sets of parameters, and its absence, NIP. A ladder is one such pattern, so every stable formula is NIP (see Proposition 4.1.2); the order of the rationals is NIP and unstable (see Example 4.1.3); and the edge relation of the random graph has the independence property (see Example 4.1.4). In Section 4.2 we pass to bounded functions and, by Rosenthal’s dichotomy, find that NIP is the existence of pointwise convergent subsequences (see Theorem 4.2.4), while the independence property produces a copy of $\beta\mathbb{N}$ (see Proposition 4.2.5).

The word “independence” is used in this book in two senses, and the equivocation is historical. A family of parameters is independent, in the sense of this chapter, when the formula can impose every pattern on it; the independence property is thus a superabundance of choices, and it is not a notion of independence of the kind studied in Chapters 6 and 7. The two meanings are nevertheless related, and in Section 7.5 we show how: where independent extensions of a type fail to be unique, an independent family appears.

4.1 THE INDEPENDENCE PROPERTY

A ladder (see Definition 1.1.4) is a special pattern: the element a_i satisfies $\varphi(x, b_j)$ exactly for $j \geq i$. The independence property asks instead that *every* pattern be realized.

Definition 4.1.1. Let M be a model of T , and let $n \geq 1$ be given.

- (a) Tuples $b_1, \dots, b_n$ of M are *shattered by φ in M* if for every $S \subseteq \{1, \dots, n\}$ there is a tuple a_S of M such that, for all $1 \leq i \leq n$, we have

$$M \models \varphi(a_S, b_i) \iff i \in S.$$

- (b) The formula φ has the *independence property* if for every n some model of T contains n tuples shattered by φ . Otherwise the formula φ is *NIP*; the letters abbreviate Shelah's name, the *negation of the independence property* (see the Introduction). NIP formulas are also called *dependent*.
- (c) The *VC dimension* of φ is the largest n such that some model of T contains n tuples shattered by φ , and ∞ if there is no largest such n .

A theory is *NIP* if every formula is NIP.

Thus the formula φ is NIP if and only if the VC dimension of φ is finite. The terminology comes from learning theory, where shattering is a property of a family of sets. For $a \in M$ we write $\varphi(a, M) = \{b \in M : M \models \varphi(a, b)\}$. If $b_1, \dots, b_n$ are shattered by φ in M , then every subset of $\{b_1, \dots, b_n\}$ is cut out by one of the sets $\varphi(a, M)$, so that the family $\{\varphi(a, M) : a \in M\}$ shatters $\{b_1, \dots, b_n\}$ in the sense of learning theory.

As for ladders, an $(n+1)$ -shattered family contains an n -shattered one, so by compactness the formula φ has the independence property if and only if some model of T contains an infinite sequence $(b_i)_{i < \omega}$ such that for every $S \subseteq \omega$ there is a_S with $\varphi(a_S, b_i)$ if and only if $i \in S$.

Proposition 4.1.2. *If φ does not have the order property, then the formula φ is NIP.*

Proof. We prove the contrapositive, by finding a ladder among the patterns that shattering provides. Suppose that $b_1, \dots, b_n$ are shattered by φ in a model M of T , witnessed by the tuples a_S . For $1 \leq i \leq n$ put $S_i = \{j : i \leq j \leq n\}$ and $a_i = a_{S_i}$. We then have $M \models \varphi(a_i, b_j)$ if and only if $j \in S_i$, that is, if and only if $i \leq j$. Thus the suffixes S_i cut out a staircase pattern, and the tuples $a_1, \dots, a_n$ and $b_1, \dots, b_n$ form an n -ladder. Since n is arbitrary, the formula φ has the order property whenever it has the independence property. ☺

Figure 4.1 shows the proof for $n = 3$.

The converse of Proposition 4.1.2 fails, and the first example of the chapter, the order of the rationals, shows why NIP is the right notion for ordered structures.

Example 4.1.3. Let T be the theory of $(\mathbb{Q}, <)$ and let $\varphi(x, y)$ be $x < y$. Then the elements $a_i = i$ and $b_j = j + \frac{1}{2}$ satisfy $\varphi(a_i, b_j)$ if and only if $i \leq j$, so the formula φ has the order property and is unstable. However, the formula φ is NIP, with VC dimension 1. Indeed, let b_1, b_2 be elements of a model, say $b_1 \leq b_2$. Shattering requires some a with $a < b_1$ and $a \not< b_2$, that is, some a with $b_2 \leq a < b_1$, which is impossible. Hence no two elements are shattered. A single element b is shattered, by any $a < b$ and any $a \geq b$.

The random graph lies at the other extreme: its edge relation realizes every pattern on every finite set of vertices.




FIGURE 4.1. The proof of Proposition 4.1.2 for $n = 3$. Left: shattered tuples b_1, b_2, b_3 , with one row a_S for each $S \subseteq \{1, 2, 3\}$ and a pink cell where $\varphi(a_S, b_j)$ holds. Right: the three rows indexed by the suffixes $S_1 = \{1, 2, 3\}$, $S_2 = \{2, 3\}$ and $S_3 = \{3\}$ (framed) form the staircase of a 3-ladder (compare Figure 1.1). A ladder uses only n of the 2^n rows, which is why the order property does not give the independence property (see Example 4.1.3).

Example 4.1.4. Let T be the theory of the random graph, whose models are the graphs in which, for all disjoint finite sets U, V of vertices, some vertex outside $U \cup V$ is adjacent to every vertex of U and to no vertex of V . Let $\varphi(x, y)$ be the edge relation $E(x, y)$. Any n distinct vertices $b_1, \dots, b_n$ are shattered: for $S \subseteq \{1, \dots, n\}$ apply the extension property to $U = \{b_i : i \in S\}$ and $V = \{b_i : i \notin S\}$. Thus the relation E has the independence property.

The independence property of $\varphi(x, y)$ refers to shattering tuples on the y -side. In Proposition 4.1.5 below we show that the choice of side does not matter, at the price of an exponential loss in the number of shattered tuples. We write $\varphi^{\text{op}}(y, x)$ for $\varphi(x, y)$ with the roles of the variables exchanged.

Proposition 4.1.5. *The formula φ is NIP if and only if φ^{op} is NIP. More precisely, if some model contains 2^n tuples shattered by φ , then that model contains n tuples shattered by φ^{op} .*

Proof. The idea is to index the shattered tuples by sets, so that membership itself becomes the pattern, and then to read the pattern from the other side. Let M be a model of T containing 2^n tuples shattered by φ , index them as b_I with $I \subseteq \{1, \dots, n\}$, and let the tuples a_S with $S \subseteq \mathcal{P}(\{1, \dots, n\})$ witness that they are shattered by φ . For $1 \leq i \leq n$ put $S_i = \{I : i \in I\}$ and $a_i = a_{S_i}$. We then have $M \models \varphi(a_i, b_I)$ if and only if $I \in S_i$, that is, if and only if $i \in I$. Thus for every $I \subseteq \{1, \dots, n\}$ the tuple b_I satisfies $\varphi^{\text{op}}(b_I, a_i)$ exactly for $i \in I$, and the tuples $a_1, \dots, a_n$ are shattered by φ^{op} . Applying this in both directions gives the equivalence. ☺



Remark 4.1.6. Examples 4.1.3 and 4.1.4 are unstable for different reasons, and a theorem of Shelah says that there are no others. A formula $\psi(x, y)$ has the *strict order property* if for every n some model M of T contains tuples $b_1, \dots, b_n$ with $\psi(M, b_1) \subsetneq \psi(M, b_2) \subsetneq \dots \subsetneq \psi(M, b_n)$, where $\psi(M, b) = \{a \in M : M \models \psi(a, b)\}$. The formula $x < y$ of Example 4.1.3 is the basic case. Shelah's theorem states that a theory is unstable if and only if some formula has the independence property or some formula has the strict order property [She90, Chapter II]. Thus an unstable theory either encodes arbitrary subsets, as the random graph does, or encodes an infinite order; NIP theories are those in which only the second reason can occur.

4.2 THE REAL-VALUED CASE

As in Section 1.2, formulas give way to bounded functions on arbitrary sets, and the goal is an analytic form of NIP: in Theorem 4.2.4 below we show that a function is NIP exactly when every sequence of its evaluation maps has a pointwise convergent subsequence. The definition must first be adapted. On arbitrary sets compactness is not available, so shattering is required only for finite patterns, with a uniform gap between the values on the two sides.

Definition 4.2.1. Let A, B be sets and $\varphi: A \times B \rightarrow \mathbb{R}$ a bounded function. A sequence $(b_n)_{n < \omega}$ in B is *independent for φ* if there are reals $r < s$ such that for all disjoint finite sets $P, Q \subseteq \omega$ there is $a \in A$ with

$$\varphi(a, b_n) \geq s \text{ for } n \in P \quad \text{and} \quad \varphi(a, b_n) \leq r \text{ for } n \in Q.$$

The function φ has the *independence property* if some sequence in B is independent for φ , and is *NIP* otherwise.

For a $\{0, 1\}$ -valued function, any $0 < r < s < 1$ may be used, and a sequence is independent exactly when each finite part of the sequence is shattered in the sense of Definition 4.1.1. By compactness, a formula $\varphi(x, y)$ is NIP in the sense of Definition 4.1.1 if and only if the corresponding function on $M^{|x|} \times M^{|y|}$ is NIP for every model M of T .

As in the $\{0, 1\}$ -valued case (see Proposition 4.1.2), stability implies NIP. The proof there turned a shattered set into a ladder, a staircase being one of the patterns that shattering provides. In the proof of Proposition 4.2.2 below, the ladder appears among the types: ultrafilter limits of witnesses give types p_m that form a ladder against the b_n , and the separate continuity of the extension $\hat{\varphi}$ of Proposition 1.2.8 rules the ladder out.

Proposition 4.2.2. *Every stable function $\varphi: A \times B \rightarrow [0, 1]$ is NIP.*

Proof. Let $\hat{\varphi}: \text{LS}(\varphi) \times \text{RS}(\varphi) \rightarrow [0, 1]$ be the separately continuous function of Proposition 1.2.8. Taking a constant family in (1.2) gives $\hat{\varphi}(p, \text{rtp}_\varphi(b)) = p(b)$ for all $p \in \text{LS}(\varphi)$ and $b \in B$.

Suppose, towards a contradiction, that (b_n) is independent for φ , with reals $r < s$. For $m < N$ choose $a_{m,N} \in A$ with

$$\varphi(a_{m,N}, b_n) \geq s \text{ for } m \leq n < N \quad \text{and} \quad \varphi(a_{m,N}, b_n) \leq r \text{ for } n < m,$$

and let $a_{m,N}$ be arbitrary for $N \leq m$. Fix a nonprincipal ultrafilter $\mathcal{U}$ on ω and put

$$p_m = \lim_{N, \mathcal{U}} \text{ltp}_\varphi(a_{m,N}) \in \text{LS}(\varphi).$$

For fixed n the displayed inequalities hold for all $N > \max(m, n)$, so we have

$$(4.1) \quad p_m(b_n) \geq s \text{ if } n \geq m, \quad p_m(b_n) \leq r \text{ if } n < m.$$

Now put $q = \lim_{n, \mathcal{U}} \text{rtp}_\varphi(b_n) \in \text{RS}(\varphi)$ and $p = \lim_{m, \mathcal{U}} p_m \in \text{LS}(\varphi)$. We compute $\hat{\varphi}(p, q)$ in two ways.

First, by continuity of $\hat{\varphi}(p_m, \cdot)$ and (4.1), we have $\hat{\varphi}(p_m, q) = \lim_{n, \mathcal{U}} p_m(b_n) \geq s$ for every m , since $p_m(b_n) \geq s$ for all $n \geq m$. By continuity of $\hat{\varphi}(\cdot, q)$, we obtain $\hat{\varphi}(p, q) = \lim_{m, \mathcal{U}} \hat{\varphi}(p_m, q) \geq s$.

Second, limits in $\text{LS}(\varphi)$ are computed coordinatewise. Since $p_m(b_n) \leq r$ for all $m > n$, we have

$$p(b_n) = \lim_{m, \mathcal{U}} p_m(b_n) \leq r \text{ for every } n.$$

By continuity of $\hat{\varphi}(p, \cdot)$, we obtain

$$\hat{\varphi}(p, q) = \lim_{n, \mathcal{U}} \hat{\varphi}(p, \text{rtp}_\varphi(b_n)) = \lim_{n, \mathcal{U}} p(b_n) \leq r.$$

Hence we have $s \leq \hat{\varphi}(p, q) \leq r$, contradicting $r < s$. $\odot$

Stability was characterized in Section 1.2 by the exchange of iterated limits. NIP, the weaker condition by Proposition 4.2.2, is characterized by the evaluation maps $e_b: \text{LS}(\varphi) \rightarrow [0, 1]$, given by $e_b(p) = p(b)$ for $b \in B$ as in Subsection 1.2.4: every sequence of them has a pointwise convergent subsequence. The analytic input is Rosenthal's dichotomy: a bounded sequence of functions has either a pointwise convergent subsequence or a subsequence that behaves like an independent sequence.

Fact 4.2.3 (Rosenthal). *Let X be a set and (f_n) a uniformly bounded sequence of real-valued functions on X . Then the sequence (f_n) has a subsequence that converges pointwise on X , or a subsequence (f_{n_k}) for which there are reals $r < s$ such that for all disjoint finite sets $P, Q \subseteq \omega$ there is $x \in X$ with $f_{n_k}(x) \geq s$ for $k \in P$ and $f_{n_k}(x) \leq r$ for $k \in Q$.*

Theorem 4.2.4. *For $\varphi: A \times B \rightarrow [0, 1]$, the following are equivalent:*

- (1) *The function φ is NIP.*
- (2) *For every sequence (b_n) in B , the sequence of functions (e_{b_n}) has a subsequence that converges pointwise on $\text{LS}(\varphi)$.*

Proof. (2) $\Rightarrow$ (1). Suppose that (b_n) is independent for φ , with reals $r < s$, and let (b_{n_k}) be a subsequence. For each $K < \omega$ choose $a_K \in A$ with $\varphi(a_K, b_{n_k}) \geq s$ for even $k < K$ and $\varphi(a_K, b_{n_k}) \leq r$ for odd $k < K$. Let $\mathcal{U}$ be a nonprincipal ultrafilter on ω , and put $p = \lim_{K, \mathcal{U}} \text{ltp}_\varphi(a_K)$, a point of $\text{LS}(\varphi)$. For each fixed k the inequality for k holds for all $K > k$, so the value $e_{b_{n_k}}(p) = p(b_{n_k})$ is at least s for even k and at most r for odd k . Hence the sequence $(e_{b_{n_k}}(p))_k$ does not converge. Since the subsequence was arbitrary, condition (2) fails.

(1) $\Rightarrow$ (2). Suppose that condition (2) fails for (b_n) . By Fact 4.2.3 applied to $X = \text{LS}(\varphi)$ and $f_n = e_{b_n}$, there are a subsequence (b_{n_k}) and reals $r < s$ such that for all disjoint finite $P, Q \subseteq \omega$ there is $p \in \text{LS}(\varphi)$ with $p(b_{n_k}) \geq s$ for $k \in P$ and $p(b_{n_k}) \leq r$ for $k \in Q$. Fix r' and s' with $r < r' < s' < s$. Since p is a pointwise limit of left types $\text{ltp}_\varphi(a)$ and only finitely many coordinates are involved, some $a \in A$ satisfies $\varphi(a, b_{n_k}) > s'$ for $k \in P$ and $\varphi(a, b_{n_k}) < r'$ for $k \in Q$. Thus the sequence (b_{n_k}) is independent for φ , with $r' < s'$, and the function φ is not NIP. $\odot$

In Proposition 4.2.5 below we show that the failure of condition (2) has a striking topological form: an independent sequence produces a copy of $\beta\mathbb{N}$.

Proposition 4.2.5. *If (b_n) is independent for φ , then the map $\mathcal{U} \mapsto \lim_{n, \mathcal{U}} e_{b_n}$ is a homeomorphism from $\beta\mathbb{N}$ onto a closed subset of the pointwise closure of $\{e_b : b \in B\}$ in $[0, 1]^{\text{LS}(\varphi)}$.*

Proof. The map $n \mapsto e_{b_n}$ from $\mathbb{N}$ into the compact space $[0, 1]^{\text{LS}(\varphi)}$ extends to the continuous map $\mathcal{U} \mapsto \lim_{n, \mathcal{U}} e_{b_n}$ on $\beta\mathbb{N}$, whose image lies in the pointwise closure of $\{e_b : b \in B\}$. Since $\beta\mathbb{N}$ is compact and $[0, 1]^{\text{LS}(\varphi)}$ is Hausdorff, it remains to be shown that the map is injective. Let $r < s$ witness the independence of (b_n) , let $\mathcal{U} \neq \mathcal{V}$ be ultrafilters on $\mathbb{N}$, and choose $X \in \mathcal{U}$ with $\mathbb{N} \setminus X \in \mathcal{V}$. For each finite $F \subseteq \mathbb{N}$ there is $a_F \in A$ with $\varphi(a_F, b_n) \geq s$ for $n \in F \cap X$ and $\varphi(a_F, b_n) \leq r$ for $n \in F \setminus X$. Let $\mathcal{D}$ be an ultrafilter on the finite subsets of $\mathbb{N}$ that contains $\{G : G \supseteq F\}$ for every finite F , and put $p = \lim_{F, \mathcal{D}} \text{ltp}_\varphi(a_F) \in \text{LS}(\varphi)$. Then we

have $p(b_n) \geq s$ for $n \in X$ and $p(b_n) \leq r$ for $n \notin X$. Hence the limit of $e_{b_n}(p)$ along $\mathcal{U}$ is at least s , the limit along $\mathcal{V}$ is at most r , and the two limits of (e_{b_n}) differ at p . ☺




FIGURE 4.2. The proof of Proposition 4.2.5. Independence of (b_n) lets a single point p of $\text{LS}(\varphi)$ cut out any prescribed set $X \subseteq \mathbb{N}$ (pink bars): the values $p(b_n)$ lie above s on X and below r off X . Two distinct ultrafilters $\mathcal{U} \ni X$ and $\mathcal{V} \ni \mathbb{N} \setminus X$ therefore give limits of (e_{b_n}) that differ at p , and $\beta\mathbb{N}$ embeds in the closure of the instances.

Thus the instances of φ , viewed as continuous functions on $\text{LS}(\varphi)$, face a dichotomy. If φ is NIP, every sequence of instances has a pointwise convergent subsequence (see Theorem 4.2.4); if φ has the independence property, the closure of the instances contains $\beta\mathbb{N}$, a compact space in which no sequence of distinct points converges. The dichotomy is the topological face of Rosenthal’s ℓ_1 dichotomy [Ros74]. In the theory of Bourgain, Fremlin and Talagrand [BFT78], the same dichotomy separates the *Rosenthal compacta*, the compact spaces of Baire class 1 functions on Polish spaces, from compact spaces of functions that contain $\beta\mathbb{N}$, and Simon [Sim15b] develops the NIP side of this picture for spaces of types. Stable functions sit at the rigid end of the spectrum: their type spaces are Eberlein compacta (see Corollary A.1.8 below).



Remark 4.2.6. Proposition 4.2.2 can also be derived, with heavier machinery, from Theorem 4.2.4. By Grothendieck’s double limit criterion, in the form of Proposition A.1.7, stability of φ makes $\{e_b : b \in B\}$ relatively weakly compact in $C(\text{LS}(\varphi))$. By Corollary A.1.6, which rests on the Eberlein–Šmulian theorem (see Fact A.1.5 below), every sequence (e_{b_n}) then has a pointwise convergent subsequence. Thus condition (2) of Theorem 4.2.4 holds, and the function φ is NIP by the implication (2) $\Rightarrow$ (1), which does not use Rosenthal’s dichotomy. The direct proof of Proposition 4.2.2 avoids both Grothendieck’s criterion and the Eberlein–Šmulian theorem.

4.3 EXERCISES

The exercises follow the sections of the chapter. Throughout, we write T_{RG} for the theory of the random graph (see Example 4.1.4), and a family $\mathcal{F}$ of subsets of a set X *shatters* a set $Z \subseteq X$ if every subset of Z is of the form $Y \cap Z$ with $Y \in \mathcal{F}$. The mark $(*)$ indicates the exercises that ask for more than a direct application of the text.

1. (Section 4.1.) Let T be the theory of $(\mathbb{Q}, <)$, and let $\varphi(x_1x_2, y)$ be the formula $x_1 < y \wedge y < x_2$. For a pair $a = (a_1, a_2)$ in a model M , the set $\varphi(a, M)$ is the open interval

with endpoints a_1 and a_2 , and it is empty when $a_2 \leq a_1$. Show that the VC dimension of φ is 2. Show that the VC dimension of φ^{op} is also 2. (Hint: let three pairs $a_i = (l_i, r_i)$ be shattered by φ^{op} , choose i with l_i least and j with r_j greatest, and, when $i \neq j$, use the pattern $\{i, j\}$; show that for some index k the interval $\varphi(a_k, M)$ is contained in the union of the other two.) Show finally that the formula φ has the order property (see Definition 1.1.4).

2. (Section 4.1.) Let T be the theory of an equivalence relation E with infinitely many classes, all of them infinite. Show that the formula $E(x, y)$ has VC dimension 1 and that no model of T contains a 2-ladder for it. For $k \geq 1$, let $\psi_k(x_1 \dots x_k, y)$ be the formula $\bigvee_{l=1}^k E(x_l, y)$. Show that the VC dimension of ψ_k is k , and that no model of T contains a $(k+1)$ -ladder for ψ_k . Thus for every k there is a stable formula of VC dimension exactly k . (Hint: in a ladder $a_1, \dots, a_n$ and $b_1, \dots, b_n$ for ψ_k , the elements $b_1, \dots, b_n$ lie in pairwise distinct classes, and each of these classes contains a coordinate of a_1 .)
3. (Section 4.1.) Let T be the theory of $(\mathbb{Q}, <)$. For $R \subseteq \{<, =, >\}$, let $\varphi_R(x, y)$ be the disjunction of the formulas $x < y$, $x = y$ and $y < x$ that correspond to the members of R , the empty disjunction being $\perp$. Show that every quantifier-free formula $\varphi(x, y)$, with x and y single variables, is equivalent modulo T to some φ_R . Show that every φ_R has VC dimension at most 1. (Hint: for $b_1 < b_2$ there are only five positions of an element a relative to b_1 and b_2 ; list the pairs of truth values of $\varphi_R(a, b_1)$ and $\varphi_R(a, b_2)$.) Show that φ_R is stable if and only if R contains both or neither of the symbols $<$ and $>$. Thus four of the eight formulas are unstable and NIP, as in Example 4.1.3.
4. (Section 4.1.) Work in T_{RG} , and let $\varphi(x, y)$ be a quantifier-free formula in the language of graphs, with x and y single variables. Show that φ is equivalent modulo T_{RG} to a formula $(x = y \wedge \alpha) \vee (x \neq y \wedge \beta)$, where α is $\top$ or $\perp$ and β is one of $\top$, $\perp$, $E(x, y)$ and $\neg E(x, y)$. Show that the formula φ is NIP if and only if β is $\top$ or $\perp$, that is, if and only if φ is equivalent modulo T_{RG} to a Boolean combination of the formula $x = y$. (Hint: when β is $\neg E(x, y)$, argue as in Example 4.1.4 with the sets U and V exchanged.)
5. (Section 4.1.) Let T be the theory of $(\mathbb{N}, |)$, where $x \mid y$ means that x divides y , and let $\varphi(x, y)$ be the formula $x \mid y$. Let $p_1, \dots, p_n$ be distinct primes. Show that the elements $b_i = \prod_{j \neq i} p_j$, for $1 \leq i \leq n$, are shattered by φ in $\mathbb{N}$, so that the formula φ has the independence property. Show directly, without Proposition 4.1.5, that the elements $p_1, \dots, p_n$ are shattered by φ^{op} in $\mathbb{N}$. Show also that the formula φ has the strict order property of Remark 4.1.6. Thus the structure $(\mathbb{N}, |)$ is unstable for both of the reasons named in that remark.
6. (Section 4.1.) Let $m \geq 1$ be given, put $X = \{1, \dots, 2^m\}$, and let M be the finite structure whose universe is the disjoint union of X and $\mathcal{P}(X)$, with one binary relation R such that $R(c, d)$ holds if and only if $c \in X$, $d \in \mathcal{P}(X)$ and $c \in d$. Let T be the complete theory of M , and let $\varphi(x, y)$ be the formula $R(x, y)$. Every model of T is isomorphic to M . Show that the VC dimension of φ^{op} is 2^m and that the VC dimension of φ is m . (Hint: for the lower bound, label the elements of X by the subsets of $\{1, \dots, m\}$; for the upper bound, count the witnesses a_S .) Conclude that the bound of Proposition 4.1.5 cannot be improved: a model may contain 2^m tuples shattered by one formula and no $m+1$ tuples shattered by its opposite.

7. (*) (Section 4.1.) The following bound is the Sauer–Shelah lemma mentioned in the historical remarks.

(a) Let X be a set with m elements, let $d \geq 0$ be given, and let $\mathcal{F}$ be a family of subsets of X that shatters no subset of X with $d + 1$ elements. Show that $|\mathcal{F}| \leq \sum_{i=0}^d \binom{m}{i}$. (Hint: induction on m . Fix $x \in X$, let $\mathcal{F}'$ be the family of the sets $Y \setminus \{x\}$ with $Y \in \mathcal{F}$, and let $\mathcal{F}''$ be the family of the sets $Y \subseteq X \setminus \{x\}$ with $Y \in \mathcal{F}$ and $Y \cup \{x\} \in \mathcal{F}$. Show that $|\mathcal{F}| = |\mathcal{F}'| + |\mathcal{F}''|$ and that the family $\mathcal{F}''$ shatters no subset with d elements. For $d = 0$ the last condition means that $\mathcal{F}''$ is empty.)

(b) Suppose that the formula $\varphi(x, y)$ has finite VC dimension d , let M be a model of T , and let $b_1, \dots, b_m$ be tuples of M . Show that, as a ranges over M , the sets $\{i : M \models \varphi(a, b_i)\}$ take at most $\sum_{i=0}^d \binom{m}{i} \leq (m + 1)^d$ values.

(c) Suppose that the formulas $\varphi(x, y)$ and $\psi(x, y)$ are NIP. Show that the formulas $\neg\varphi$, $\varphi \wedge \psi$ and $\varphi \vee \psi$ are NIP. (Hint: for the conjunction, let d and e be bounds for the VC dimensions of φ and ψ , and compare $(m + 1)^{d+e}$ with 2^m .)

8. (Section 4.2.) For $n \geq 1$, let P_n be the set of pairs (n, i) with $1 \leq i \leq n$, and let Q_n be the set of pairs (n, Y) with $Y \subseteq \{1, \dots, n\}$. Let M be the structure whose universe is the disjoint union of the sets P_n and Q_n , with one binary relation R that holds of $((n, Y), (n, i))$ if and only if $i \in Y$, and of no other pair. Let T be the complete theory of M , and let $\varphi(x, y)$ be the formula $R(x, y)$. Show that the formula φ has the independence property in the sense of Definition 4.1.1. Show that the corresponding function $\varphi: M \times M \rightarrow \{0, 1\}$ is NIP in the sense of Definition 4.2.1. (Hint: all terms of an independent sequence lie in a single set P_n , and they are pairwise distinct.) Show finally, by compactness, that some model of T contains a sequence independent for φ . Thus, in the equivalence stated after Definition 4.2.1, the words “for every model” cannot be replaced by “for the given model”.

9. (Section 4.2.) Let $\varphi: A \times \mathbb{N} \rightarrow [0, 1]$ be a function such that, for every $a \in A$, the sequence $(\varphi(a, n))_n$ is monotone, and in the same direction for all a .

(a) Show from Definition 4.2.1 that the function φ is NIP. (Hint: two terms of an independent sequence are distinct, and one pattern of size two already fails.)

(b) Verify condition (2) of Theorem 4.2.4 directly. (Hint: every $p \in \text{LS}(\varphi)$ is monotone on $\mathbb{N}$, and every sequence in $\mathbb{N}$ has a subsequence that is constant or strictly increasing.)

(c) Apply (a) to the function $\varphi(x, n) = x^n$ on $[0, 1] \times \mathbb{N}$ of Example A.1.4 and to the function of Example 1.2.12. Show that both functions are unstable in the sense of Definition 1.2.5.

(d) Let A be a topological space, let B be a set, and let $\varphi: A \times B \rightarrow [0, 1]$ be an NIP function such that every function $\varphi(\cdot, b)$ is continuous on A . Show that for every sequence (b_n) in B some subsequence of the functions $\varphi(\cdot, b_n)$ converges pointwise on A , so that its limit is of the first Baire class. Show that for x^n on $[0, 1]$ the limit is not continuous.

10. (Section 4.2.) Let $\varphi: [0, 1] \times \mathbb{N} \rightarrow [0, 1]$ be given by $\varphi(t, k) = \frac{1}{2}(1 + \sin(2\pi 8^k t))$. Show that the sequence $(k)_{k < \omega}$ is independent for φ , with $r = \frac{1}{2} - \frac{\sqrt{2}}{4}$ and $s = \frac{1}{2} + \frac{\sqrt{2}}{4}$. (Hint: write $t = \sum_{j \geq 1} d_j 8^{-j}$ with digits $d_j \in \{0, 2, 6\}$; the value $\sin(2\pi 8^k t)$ is at least $\frac{\sqrt{2}}{2}$ when $d_{k+1} =$

2, and at most $-\frac{\sqrt{2}}{2}$ when $d_{k+1} = 6$.) Show also that no subsequence of the continuous functions $t \mapsto \varphi(t, k)$ converges pointwise on $[0, 1]$. (Hint: for a subsequence (k_l) , choose t with $d_{k_l+1} = 2$ for even l , with $d_{k_l+1} = 6$ for odd l , and with all other digits 0; independence alone would not suffice, by Exercise 12(a).) Deduce from Proposition 4.2.5 that the pointwise closure of $\{e_k : k \in \mathbb{N}\}$ in $[0, 1]^{\text{LS}(\varphi)}$ contains a copy of $\beta\mathbb{N}$.


11. (*) (Section 4.2.) Let $\varphi: [0, 2\pi] \times \mathbb{N} \rightarrow [0, 1]$ be given by $\varphi(x, n) = \frac{1}{2}(1 + \sin nx)$. Show that $\int_0^{2\pi} (\sin nx - \sin n'x)^2 dx = 2\pi$ for distinct $n, n' \geq 1$. Deduce from the dominated convergence theorem that no subsequence of the functions $\varphi(\cdot, n)$ converges pointwise on $[0, 2\pi]$. Apply Fact 4.2.3 with $X = [0, 2\pi]$, not $X = \text{LS}(\varphi)$, to conclude that some subsequence of $(n)_{n < \omega}$ is independent for φ . Deduce that the function $(x, y) \mapsto \frac{1}{2}(1 + \sin xy)$ on $\mathbb{R} \times \mathbb{R}$ has the independence property. Compare Exercise 10, which gives an explicit independent sequence.
12. (Section 4.2.) Let A be the set of finite subsets of $\mathbb{N}$, and let $\varphi: A \times \mathbb{N} \rightarrow \{0, 1\}$ be given by $\varphi(F, n) = 1$ if $n \in F$ and $\varphi(F, n) = 0$ otherwise.
 - (a) Show that the sequence $(n)_{n < \omega}$ is independent for φ . Show that for every sequence (n_k) of distinct natural numbers the functions $\varphi(\cdot, n_k)$ converge pointwise to 0 on A . Thus condition (2) of Theorem 4.2.4 cannot be weakened to convergence at the realized types $\text{lt}_\varphi(F)$.
 - (b) Show that $\text{LS}(\varphi)$ is the set of all functions from $\mathbb{N}$ to $\{0, 1\}$, which we identify with $\mathcal{P}(\mathbb{N})$, and that $e_n(Z) = 1$ if and only if $n \in Z$.
 - (c) Let $\mathcal{U}$ be an ultrafilter on $\mathbb{N}$. Show that the function $\lim_{n, \mathcal{U}} e_n$ on $\mathcal{P}(\mathbb{N})$ takes the value 1 at Z if and only if $Z \in \mathcal{U}$. Conclude that the map of Proposition 4.2.5 sends each ultrafilter to its own characteristic function, and that it is a homeomorphism from $\beta\mathbb{N}$ onto the pointwise closure of $\{e_n : n \in \mathbb{N}\}$.
 - (d) Show that the function φ is unstable in the sense of Definition 1.2.1. (Hint: consider $\neg\varphi$.)
13. (*) (Section 4.2.) Let A and φ be as in Exercise 12, and let $\psi: \mathbb{N} \times A \rightarrow \{0, 1\}$ be given by $\psi(n, F) = \varphi(F, n)$. Show that the function ψ is NIP. (Hint: in an independent sequence (F_k) , the finite set F_0 would have to contain 2^N distinct elements for every N .) Show that the function ψ has the order property. Thus, for functions on arbitrary sets, the independence property of a function does not pass to its opposite, in contrast with Proposition 4.1.5. Show, however, that for every k some sets $F_1, \dots, F_k \in A$ have all 2^k patterns realized by elements of $\mathbb{N}$, so that the failure comes from the requirement that a single infinite sequence carry all finite patterns.

4.4 HISTORICAL REMARKS

✂ **Section 4.1.** The independence property was introduced by Shelah [She71], in the same paper as the order property. His theorem that an unstable theory has the independence property or the strict order property is in [She90, Chapter II]. The combinatorial notion of shattering, and the VC dimension, were introduced independently in probability and learning theory by Vapnik and Chervonenkis [VC71]; the polynomial bound on the number of traces

(the subsets cut out on an n -element set) of a family of finite VC dimension was found independently by Sauer [Sau72] and by Shelah [She72]. Laskowski [Las92] observed that a formula is NIP exactly when the family of sets it defines has finite VC dimension. The observation joined model theory to learning theory.

✂ **Section 4.2.** Rosenthal’s dichotomy comes from his characterization of Banach spaces containing ℓ_1 [Ros74], and the systematic study of pointwise compact sets of Baire class 1 functions is due to Rosenthal [Ros77] and to Bourgain, Fremlin and Talagrand [BFT78]. Their work entered model theory later. The development is summarized in the following timeline.



- 1971** • **Shelah; Vapnik–Chervonenkis.** Shelah [She71] introduces the independence property. Independently, Vapnik and Chervonenkis [VC71] introduce shattering and what is now called VC dimension, in the study of uniform convergence of empirical frequencies.
- 1972** • **Sauer; Shelah.** Sauer [Sau72] and Shelah [She72] independently prove the polynomial bound on the number of traces of a set system of finite VC dimension.
- 1974–1978** • **Rosenthal; Bourgain–Fremlin–Talagrand.** Rosenthal proves the ℓ_1 dichotomy [Ros74] and studies pointwise compact sets of Baire class 1 functions [Ros77]. Bourgain, Fremlin and Talagrand [BFT78] develop the theory of these “Rosenthal compacta”.
- 1981** • **Poizat.** Poizat [Poi81] characterizes countable NIP theories by a bound on the number of types, and develops the theory of unstable theories without the independence property.
- 1986** • **Pillay–Steinhorn.** Pillay and Steinhorn [PS86] develop the theory of o-minimal structures, following van den Dries. These structures are NIP and are the main source of unstable NIP theories.
- 1992** • **Laskowski.** Laskowski [Las92] identifies NIP formulas with formulas defining families of finite VC dimension.
- 2008–2011** • **Hrushovski–Peterzil–Pillay; Hrushovski–Pillay.** The study of invariant and generically stable measures in NIP theories [HPP08, HP11], motivated by Pillay’s conjecture on definable groups in o-minimal structures, turns NIP into a central topic of model theory.
- 2009** • **Ben Yaacov.** Ben Yaacov [BY09] extends VC classes to $[0, 1]$ -valued functions, the real-valued analogue of NIP used in Section 4.2.

- 2015** • **Simon.** Simon [Sim15b] applies the theory of Rosenthal compacta to spaces of φ -types of NIP formulas, and publishes the monograph [Sim15a] on NIP theories.

Chapter 5

From stability to simplicity

A stable formula has definable types, and a definition extends its type canonically to every larger set of parameters (see Theorem 1.1.6 and the remarks following Definition 1.1.3). A stable theory therefore has, over a model, a preferred way of extending a type without adding information, and in Chapter 6 we show that this way is unique. The random graph of Example 4.1.4 is not stable, since its edge relation has the independence property, and yet it carries an evident notion of independence. Over a set C , a vertex a is independent from a set B when a is not one of the vertices of B outside C ; once this holds, the extension axioms leave a free to be joined to the vertices of B in any way whatever. Uniqueness is lost precisely here: a new vertex may be joined to a given vertex b or not, and neither choice tells us more about b than the other. What survives is amalgamation. Suppose that c_1 and c_2 have the same type over C , that c_1 is independent from b_1 and c_2 from b_2 , and that b_1 and b_2 are independent from each other. In Theorem 5.4.9(c) we show that then a single vertex realizes the type of c_1 over $C \cup \{b_1\}$ and the type of c_2 over $C \cup \{b_2\}$ at once. This is the *independence theorem*. In Chapter 7 we show that it holds in the *simple* theories of Shelah, and that, together with the other properties of independence, it characterizes them (see Theorems 7.4.7 and 7.6.5). The theme of this chapter fits in one line: from stability to simplicity, uniqueness gives way to amalgamation.

Simplicity has a combinatorial definition in the spirit of the order property and of the independence property: a theory is simple if no formula has the *tree property* (see Definition 5.2.7), a configuration assembled from *dividing*, the way in which a formula can say something essentially new about a set.

Roadmap

In this chapter we proceed from definitions to an example, and from the example to the property that made the example work. In Section 5.1 we fix the conventions on monster models, and in Section 5.2 we define dividing and the tree property. In Section 5.3 we place simplicity in the classification: every stable theory is simple, a theory is stable if and only if it is simple and NIP (see Theorem 5.3.4), and in a simple theory every type is nondividing over a small part of its domain. In Section 5.4 we compute dividing in the random graph from the extension axioms alone, prove that the random graph is simple, and prove the independence theorem there. In Section 5.5 we isolate what these proofs used: a single property of the finite graphs, *disjoint 3-amalgamation*. Every Fraïssé limit with this property is simple (see Theorem 5.5.9), and two classical Fraïssé limits without it, the rational order and the Henson graph, are not (see Proposition 5.3.2 and Example 5.5.10). In Section 5.6 we pass to real-valued functions, as we did in Chapters 1 and 4, and show by counting approximate definitions that stable functions have no tree property (see Theorem 5.6.4).

The order of exposition is deliberate. A general theorem is easier to believe after one has seen it hold in a particular case, and easier to prove after one has seen which features of the particular case the proof actually used. The random graph is studied here in detail so that in Chapter 7 we may set it aside.

5.1 CONVENTIONS

In this section we fix the setting for the rest of the book. From now on T is a *complete* theory in the language L , and $\kappa = |T| = |L| + \aleph_0$. We work inside a monster model $\mathbb{M}$ of T : a $\bar{\kappa}$ -saturated and strongly $\bar{\kappa}$ -homogeneous model for some very large cardinal $\bar{\kappa}$ (larger than 2^{2^κ} , say). A set is *small* if of cardinality $< \bar{\kappa}$. Letters A, B, C denote small subsets of $\mathbb{M}$; tuples may be infinite, but are always small. We freely identify a set with an enumerating tuple. A *model* is a small elementary substructure $M \preceq \mathbb{M}$. Every model of T of size $< \bar{\kappa}$ embeds elementarily in $\mathbb{M}$, and every type over a small set is realized in $\mathbb{M}$.

We write $\text{tp}(a/A)$ for the complete type of a over A , $S_x(A)$ for the set of complete types over A in the variables x , and $a \equiv_A a'$ for $\text{tp}(a/A) = \text{tp}(a'/A)$. We write MB for $M \cup B$. By strong homogeneity, we have $a \equiv_A a'$ if and only if $\sigma(a) = a'$ for some $\sigma \in \text{Aut}(\mathbb{M}/A)$.

Until now, stability has been a property of single formulas. We record the corresponding notion for the theory T , together with a sufficient condition by counting types, which we use in Section 5.5 to recognize stable Fraïssé limits.

Definition 5.1.1.

- (a) The theory T is *stable* if no L -formula $\varphi(x, y)$ has the order property (see Definition 1.1.4).
- (b) For an infinite cardinal λ , the theory T is λ -*stable* if $|S_x(A)| \leq \lambda$ whenever $|A| \leq \lambda$ and x is a finite tuple of variables.

Lemma 5.1.2. *If T is λ -stable for some $\lambda \geq \kappa$, then the theory T is stable.*

Proof. Suppose that $\varphi(x, y)$ has the order property. By Proposition 2.3.1, the formula φ encodes binary trees of arbitrarily large finite height. By Proposition 2.4.2 there is a model N of T with $|N| \leq \lambda$ and $|S_\varphi(N)| > \lambda$. Since T is complete, we may assume $N \preceq \mathbb{M}$. Every φ -type over N is consistent with the elementary diagram of N , so each φ -type over N extends to a complete type in $S_x(N)$, and distinct φ -types have distinct extensions. Hence we have $|S_x(N)| > \lambda \geq |N|$, and T is not λ -stable. ☺

5.2 DIVIDING AND THE TREE PROPERTY

The goal of this section is to define dividing and the tree property, and with them the simple theories. A small set of formulas with parameters in $\mathbb{M}$ is *consistent* if it is realized in $\mathbb{M}$. We often identify a tuple with the set of its coordinates; thus, for tuples a, b and a set C , the inclusion $a \cap b \subseteq C$ says that every element that is a coordinate of both a and b belongs to C .

A formula $\varphi(x, b)$ may say about x something that a set C already knows, or something new. Two examples show the difference. In a model of the theory of $(\mathbb{Q}, <)$, the formula $b_1 < x < b_2$ is one of infinitely many formulas $b_1^i < x < b_2^i$, with $b_2^i < b_1^{i+1}$, which look alike over $\emptyset$ and are pairwise contradictory. In the random graph, the formula $E(x, b)$, where b is

a vertex, looks like $E(x, b')$ for every other vertex b' , but any finite number of these formulas hold together, because the extension axioms provide a common neighbour. The first formula *divides* over $\emptyset$, and the second does not (see Example 5.2.6 and Theorem 5.4.5). Indiscernible sequences make “look alike” precise.

Definition 5.2.1. Let C be a small set. A sequence $(b_i)_{i < \omega}$ of tuples of the same length is *indiscernible over C* if for every n and all $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$ in ω , we have

$$(b_{i_1}, \dots, b_{i_n}) \equiv_C (b_{j_1}, \dots, b_{j_n}).$$

Every infinite subsequence of a sequence indiscernible over C is indiscernible over C ; a sequence indiscernible over C is indiscernible over every subset of C ; and, taking $n = 1$, all its terms have the same type over C . Leibniz held that things which cannot be told apart are one and the same. The terms of an indiscernible sequence may be distinct, and yet no formula over C tells one increasing tuple of them from another.

Indiscernible sequences are plentiful. Every sequence leaves a trace of the formulas that hold of all its increasing tuples, and in Lemma 5.2.2 below we show that Ramsey’s theorem, together with compactness, produces an indiscernible sequence with the same trace. For a sequence $(b_i)_{i < \omega}$ and a small set C , the *Ehrenfeucht–Mostowski type* of (b_i) over C is the set of $L(C)$ -formulas $\psi(y_1, \dots, y_n)$ such that $\mathbb{M} \models \psi(b_{i_1}, \dots, b_{i_n})$ for all $i_1 < \dots < i_n$ in ω .

Lemma 5.2.2 (Extraction of indiscernible sequences). *Let C be a small set, and let $(b_i)_{i < \omega}$ be tuples of the same length. Then there is a sequence $(b'_i)_{i < \omega}$, indiscernible over C , that realizes the Ehrenfeucht–Mostowski type of (b_i) over C : for every formula $\psi(y_1, \dots, y_n)$ of that type and all $j_1 < \dots < j_n$, we have $\mathbb{M} \models \psi(b'_{j_1}, \dots, b'_{j_n})$. In particular, if all the b_i have the same type over C , then we have $b'_j \equiv_C b_0$ for every j .*

Proof. Let Σ be the set of $L(C)$ -formulas in variables y_i ($i < \omega$) consisting of

- (a) $\psi(y_{j_1}, \dots, y_{j_n})$, for every ψ in the Ehrenfeucht–Mostowski type of (b_i) over C and all $j_1 < \dots < j_n$, and
- (b) $\theta(y_{i_1}, \dots, y_{i_n}) \leftrightarrow \theta(y_{j_1}, \dots, y_{j_n})$, for every $L(C)$ -formula θ and all $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$.

A realization of Σ is a sequence as required: (b) says that it is indiscernible over C , and (a) that it realizes the Ehrenfeucht–Mostowski type. Since Σ has at most $|C| + \aleph_0$ parameters, by saturation it suffices to show that every finite $\Sigma_0 \subseteq \Sigma$ is satisfiable.

Let $\theta_1, \dots, \theta_m$ be the formulas that occur in the formulas of type (b) in Σ_0 , and let n be the largest number of variables among them (if there are none, take $H = \omega$ below). Colour each set $\{i_1 < \dots < i_n\} \subseteq \omega$ by the sequence of truth values of $\theta_l(b_{i_1}, \dots, b_{i_{n_l}})$, $l = 1, \dots, m$, where $n_l \leq n$ is the number of variables of θ_l . There are at most 2^m colours, so Ramsey’s theorem gives an infinite set $H \subseteq \omega$ all of whose n -element subsets have the same colour. Let $h_0 < h_1 < \dots$ enumerate H , and assign b_{h_i} to y_i . Every increasing n_l -tuple from H is the beginning of an increasing n -tuple from H , because H is infinite, so the truth value of θ_l on increasing n_l -tuples from H does not depend on the tuple; hence the formulas of type (b) in Σ_0 hold. The formulas of type (a) hold because they hold of all increasing tuples of the original sequence. For the last assertion, every formula of $\text{tp}(b_0/C)$, in one variable y_1 , belongs to the Ehrenfeucht–Mostowski type. ☺

In the Rado graph and in Fraïssé limits, quantifier elimination gives a sharper conclusion over finite sets: an infinite subsequence of the given sequence is itself indiscernible (see Lemmas 5.4.4 and 5.5.7).

Definition 5.2.3. Let C be a small set, let $\varphi(x, y)$ be an L -formula, and let b be a tuple of the length of y .

- (a) The formula $\varphi(x, b)$ *divides over C* if there is a sequence $(b_i)_{i < \omega}$, indiscernible over C , with $b_0 = b$, such that the set $\{\varphi(x, b_i) : i < \omega\}$ is inconsistent.
- (b) For $k \geq 2$, a set of formulas is *k -inconsistent* if each of its subsets with k elements is inconsistent, and $\varphi(x, b)$ *k -divides over C* if the sequence in (a) can be chosen so that $\{\varphi(x, b_i) : i < \omega\}$ is k -inconsistent.
- (c) Let $B \supseteq C$ be small. A type $p \in S_x(B)$ *divides over C* if some formula of p divides over C .

Every $L(B)$ -formula has the form $\varphi(x, b)$ with $\varphi \in L$ and b a finite tuple from B , so clause (c) makes sense. The next lemma shows that dividing is k -dividing for a suitable k , so that the two notions of (a) and (b) agree.

Lemma 5.2.4. *Let $(b_i)_{i < \omega}$ be indiscernible over C , and suppose that $\{\varphi(x, b_i) : i < \omega\}$ is inconsistent. Then there is $k \geq 2$ such that $\{\varphi(x, b_i) : i < \omega\}$ is k -inconsistent. In particular, a formula divides over C if and only if it k -divides over C for some $k \geq 2$.*

Proof. By compactness and saturation, some finite subset $\{\varphi(x, b_{i_1}), \dots, \varphi(x, b_{i_k})\}$, with $i_1 < \dots < i_k$, is inconsistent; adding formulas keeps it inconsistent, so we may assume $k \geq 2$. Thus we have $\mathbb{M} \models \neg \exists x \bigwedge_{l=1}^k \varphi(x, b_{i_l})$. For $j_1 < \dots < j_k$, indiscernibility gives $(b_{j_1}, \dots, b_{j_k}) \equiv_C (b_{i_1}, \dots, b_{i_k})$, and hence we have $\mathbb{M} \models \neg \exists x \bigwedge_{l=1}^k \varphi(x, b_{j_l})$. ☺



Remark 5.2.5. If $C \subseteq C'$ and $\varphi(x, b)$ k -divides over C' , then the formula $\varphi(x, b)$ k -divides over C , because a sequence indiscernible over C' is indiscernible over C . Consequently, if $C \subseteq C' \subseteq B$ and $p \in S_x(B)$ divides over C' , then the type p divides over C . The terms of the sequence in Definition 5.2.3 all have the type of b over C , so dividing is a strengthening of Shelah's original condition, which asks only for infinitely many realizations of $\text{tp}(b/C)$ whose formulas are k -inconsistent; in fact, the two conditions are equivalent (see [TZ12, Chapter 7]), but we do not need this.

Example 5.2.6. Let T be the theory of $(\mathbb{Q}, <)$, let $b_1 < b_2$ be elements of $\mathbb{M}$, and let $\theta(x; y_1, y_2)$ be the formula $y_1 < x < y_2$. Choose pairs $b^i = (b_1^i, b_2^i)$ with $b^0 = (b_1, b_2)$, with $b_1^i < b_2^i$, and with $b_2^i < b_1^{i+1}$ for every i . By quantifier elimination for dense linear orders without endpoints (see Fact 5.5.2(b) and Example 5.5.3(2) below, or [Hod93, Theorem 7.4.1]), the type over $\emptyset$ of a tuple is determined by the order between its coordinates, and any two increasing n -tuples of the sequence (b^i) are ordered alike. Hence the sequence (b^i) is indiscernible over $\emptyset$. The intervals (b_1^i, b_2^i) are pairwise disjoint, so the set $\{\theta(x, b^i) : i < \omega\}$ is 2-inconsistent, and $\theta(x, b_1, b_2)$ 2-divides over $\emptyset$.

In Example 5.2.6 we can go on: inside the interval (b_1, b_2) there are infinitely many disjoint subintervals, inside each of these infinitely many more, and so on. Choosing one interval at each stage, we obtain a consistent descending sequence of formulas; at each stage, the available choices are pairwise inconsistent. The tree property abstracts this picture. We use the notation of Chapter 2 for sequences, with ω in place of 2. Thus the set $\omega^{<\omega}$ consists of the finite sequences of natural numbers, ω^ω is the set of infinite ones, $\sigma \frown \langle i \rangle$ is the extension of σ by the entry i , and $\eta|n$ is the initial segment of η of length n .

Definition 5.2.7. Let $\varphi(x, y)$ be an L -formula, and let $k \geq 2$ be given. The formula φ has the k -tree property if there are tuples b_σ , for $\sigma \in \omega^{<\omega}$ with $\sigma \neq \emptyset$, such that the following hold:

- (a) for every $\eta \in \omega^\omega$, the set $\{\varphi(x, b_{\eta|n}) : 1 \leq n < \omega\}$ is consistent;
- (b) for every $\sigma \in \omega^{<\omega}$, the set $\{\varphi(x, b_{\sigma \smallfrown \langle i \rangle}) : i < \omega\}$ is k -inconsistent.

The formula φ has the *tree property* if it has the k -tree property for some $k \geq 2$. The theory T is *simple* if no L -formula has the tree property.

Along each branch the formulas are consistent; among the immediate successors of a node, called *siblings*, they are k -inconsistent. Each step up a branch therefore chooses one of infinitely many incompatible options, and a simple theory is one that cannot accommodate an ω -sequence of such choices made by a single formula. Shelah's original name for the simple theories, *treeless*, said precisely this (see the Introduction). The trees of Chapter 2 were binary, with parameters at the nodes and elements at the leaves; the trees here branch infinitely, carry parameters at every node, and have no leaves. Figure 5.1 shows the configuration.




FIGURE 5.1. The k -tree property (see Definition 5.2.7). The tree grows upward; the node σ carries the parameter b_σ . Along every branch (pink) the formulas $\varphi(x, b_{\eta|n})$ have a common solution, while the formulas attached to the siblings of any node (framed for the successors of the root) are k -inconsistent: no k of them hold together.

In the next remark we show that, by compactness, trees of finite height suffice.



Remark 5.2.8. Let $\varphi(x, y)$ be an L -formula, and let $k \geq 2$ be given. Suppose that for every $n \geq 1$ there are tuples b_σ , for $\sigma \in \omega^{<n}$ with $\sigma \neq \emptyset$, such that condition (a) of Definition 5.2.7 holds for the branches of length n (the set $\{\varphi(x, b_{\eta|m}) : 1 \leq m \leq n\}$ is consistent for every $\eta \in \omega^n$) and condition (b) holds for every $\sigma \in \omega^{<n}$. Then the formula φ has the k -tree property. Indeed, by compactness a branch $\{\varphi(x, b_{\eta|m}) : 1 \leq m < \omega\}$ is consistent if and only if each of its finite initial segments is. Hence the k -tree property for a family of variables (y_σ) is expressed by the set of formulas

$$\exists x \bigwedge_{m=1}^n \varphi(x, y_{\eta|m}) \quad (\eta \in \omega^\omega, n \geq 1), \quad \neg \exists x \bigwedge_{l=1}^k \varphi(x, y_{\sigma \smallfrown \langle i_l \rangle}) \quad (\sigma \in \omega^{<\omega}, i_1 < \dots < i_k).$$

A finite subset of this set mentions only nodes of length at most some n , and the given family of height n satisfies it. By compactness and saturation, the whole set is realized in $\mathbb{M}$.

5.3 SIMPLICITY BETWEEN STABILITY AND NIP

The goal of this section is to place simplicity in the classification. We show that every stable theory is simple, that a theory with the strict order property is not, and that the simple NIP theories are exactly the stable ones. We close with the first structural consequence of simplicity: every type is nondividing over a subset of its domain of size at most κ .

Proposition 5.3.1. *If T is stable, then the theory T is simple.*

Proof. The idea is to count branches. A tree with λ successors at every node has $\lambda^{\aleph_0}$ branches, whereas a single φ -type can follow at most $2^{\aleph_0}$ of them, because at each node it accepts fewer than k successors; a stable formula has too few types to follow all the branches.

Suppose that $\varphi(x, y)$ has the k -tree property. Put $\lambda_0 = 2^\kappa$, $\lambda_{n+1} = 2^{\lambda_n}$, and $\lambda = \sup_n \lambda_n$. Then we have $\lambda \geq 2^\kappa$, the set $\lambda^{<\omega}$ has cardinality λ , and λ has cofinality ω , so that $\lambda^{\aleph_0} > \lambda$ by König's theorem.

© **Step 1: a wide tree.** There are tuples b_σ , for $\sigma \in \lambda^{<\omega}$ with $\sigma \neq \emptyset$, satisfying the two conditions of Definition 5.2.7 with λ in place of ω . Indeed, the conditions are expressed by formulas as in Remark 5.2.8, and a finite set of them mentions finitely many nodes. Let $J \subseteq \lambda$ be the finite set of entries of these nodes, and let $e: J \rightarrow \omega$ be injective. The map sending σ to the sequence $e \circ \sigma$ preserves initial segments and the relation of being siblings, so the tuples $b_{e \circ \sigma}$ of the given tree satisfy the finite set. By compactness and saturation, the wide tree exists in $\mathbb{M}$. By the downward Löwenheim–Skolem theorem, there is a model $M \preceq \mathbb{M}$ with $|M| = \lambda$ that contains all the tuples b_σ .

© **Step 2: few types.** By Theorem 1.1.6, every $p \in S_\varphi(M)$ has a φ -definition that is a positive Boolean combination of instances $\varphi(a, y)$ with $a \in M$. The type p is determined by the set D_p , hence by its definition, and there are at most $|M| + \aleph_0 = \lambda$ such combinations. Therefore $|S_\varphi(M)| \leq \lambda$.

© **Step 3: many branches.** For each $\eta \in \lambda^\omega$, choose a_η realizing $\{\varphi(x, b_{\eta|n}) : n \geq 1\}$, and put $p_\eta = \text{tp}_\varphi(a_\eta/M)$. Fix $p \in S_\varphi(M)$. For each node σ , the set $\text{Acc}_\sigma = \{\alpha < \lambda : \varphi(x, b_{\sigma \smallfrown \langle \alpha \rangle}) \in p\}$ of successors accepted by p has fewer than k elements: otherwise p would contain k siblings, whose conjunction is inconsistent in $\mathbb{M}$ and hence in M , whereas p is consistent with the elementary diagram of M . If $p_\eta = p$, then we have $\eta(n) \in \text{Acc}_{\eta|n}$ for every n . Hence the map sending such an η to the sequence of positions of $\eta(n)$ in $\text{Acc}_{\eta|n}$ is injective into $(k-1)^\omega$, and at most $2^{\aleph_0}$ branches η satisfy $p_\eta = p$. It follows that

$$\lambda^{\aleph_0} = |\lambda^\omega| \leq |S_\varphi(M)| \cdot 2^{\aleph_0} \leq \lambda,$$

because $\lambda \geq 2^\kappa \geq 2^{\aleph_0}$; this contradicts $\lambda^{\aleph_0} > \lambda$. ☹

Simplicity is thus weaker than stability. In Proposition 5.3.2 below we show that a definable infinite order is enough to destroy it: the strict order property (see Remark 4.1.6) produces the picture that follows Example 5.2.6, nested intervals with infinitely many disjoint subintervals inside each interval.

Proposition 5.3.2. *If some formula has the strict order property, then the theory T is not simple.*

Proof. Let $\psi(x, y)$ have the strict order property, and write $\psi(\mathbb{M}, c)$ for the set of solutions of $\psi(x, c)$ in $\mathbb{M}$. By compactness and saturation, there are tuples c_q , for $q \in \mathbb{Q}$, with $\psi(\mathbb{M}, c_q) \subsetneq \psi(\mathbb{M}, c_r)$ whenever $q < r$: a finite set of these conditions mentions finitely many rationals $q_1 < \dots < q_n$, and is satisfied by a strictly increasing chain of length n , which exists in some model of T and hence in $\mathbb{M}$. Let $\theta(x; y_1, y_2)$ be the formula $\psi(x, y_2) \wedge \neg\psi(x, y_1)$, and for rationals $q < r$ put

$$\Delta(q, r) = \psi(\mathbb{M}, c_r) \setminus \psi(\mathbb{M}, c_q),$$

the set of solutions of $\theta(x; c_q, c_r)$; it is nonempty. Two properties of these sets follow from monotonicity of the chain:

- (i) if $q \leq q' < r' \leq r$, then we have $\Delta(q', r') \subseteq \Delta(q, r)$;
- (ii) if $r \leq q' < r'$, then we have $\Delta(q, r) \cap \Delta(q', r') = \emptyset$, since $\Delta(q, r) \subseteq \psi(\mathbb{M}, c_r) \subseteq \psi(\mathbb{M}, c_{q'})$ and $\Delta(q', r')$ is disjoint from $\psi(\mathbb{M}, c_{q'})$.

Assign rationals $q_\sigma < r_\sigma$ to the nodes $\sigma \in \omega^{<\omega}$ by recursion, starting from $(q_\emptyset, r_\emptyset) = (0, 1)$ and putting, with $d = r_\sigma - q_\sigma$,

$$q_{\sigma \smallfrown \langle i \rangle} = q_\sigma + d(1 - 2^{-i}), \quad r_{\sigma \smallfrown \langle i \rangle} = q_\sigma + d(1 - 2^{-i-1}).$$

Then we have $q_\sigma \leq q_{\sigma \smallfrown \langle i \rangle} < r_{\sigma \smallfrown \langle i \rangle} \leq r_\sigma$, and $r_{\sigma \smallfrown \langle i \rangle} \leq q_{\sigma \smallfrown \langle j \rangle}$ for $i < j$. Put $b_\sigma = (c_{q_\sigma}, c_{r_\sigma})$ for $\sigma \neq \emptyset$. For a branch η and $n \geq 1$, property (i) gives $\Delta(q_{\eta|n}, r_{\eta|n}) \subseteq \Delta(q_{\eta|m}, r_{\eta|m})$ for $m \leq n$, so any element of the nonempty set $\Delta(q_{\eta|n}, r_{\eta|n})$ satisfies $\theta(x, b_{\eta|m})$ for all $m \leq n$. By property (ii), the sets attached to two siblings are disjoint. Thus the formula θ has the 2-tree property. ☺



FIGURE 5.2. The proof of Proposition 5.3.2. The node σ carries the rational interval $(q_\sigma, r_\sigma]$, drawn at the height of the length of σ ; the successors of a node carry disjoint subintervals of its interval. The formula $\theta(x, b_\sigma)$ defines $\Delta(q_\sigma, r_\sigma) = \psi(\mathbb{M}, c_{r_\sigma}) \setminus \psi(\mathbb{M}, c_{q_\sigma})$, a band of the chain $(\psi(\mathbb{M}, c_q))_{q \in \mathbb{Q}}$. Nested intervals give nested bands, so a branch (pink) is consistent; disjoint intervals give disjoint bands, so siblings are 2-inconsistent.

The two propositions combine with Shelah's theorem, quoted in Remark 4.1.6, to characterize the stable theories as the simple NIP ones.

Fact 5.3.3 (Shelah [She90, Chapter II]). *If T is unstable, then some formula has the independence property or some formula has the strict order property.*

Theorem 5.3.4. *A complete theory is stable if and only if it is simple and NIP.*

Proof. If T is stable, then the theory T is simple by Proposition 5.3.1, and every formula is NIP by Proposition 4.1.2. Conversely, if T is unstable, then by Fact 5.3.3 some formula has the independence property, so that T is not NIP, or some formula has the strict order property, so that T is not simple by Proposition 5.3.2. ☺

Thus simplicity and NIP are the two halves of stability: an unstable theory encodes either arbitrary subsets or an infinite order, simplicity forbids the order, and NIP forbids the subsets. Neither half contains the other: the random graph is simple and not NIP (see Theorem 5.4.6 and Example 4.1.4), and the rational order is NIP and not simple (see Example 4.1.3 and Proposition 5.3.2). Figure 5.3 places the Fraïssé limits studied in Section 5.5 in the four regions that this division determines.




FIGURE 5.3. Theorem 5.3.4: the stable theories are the simple NIP theories. The pure infinite set is stable (see Example 5.5.3). The Rado graph is simple (see Theorem 5.4.6) and has the independence property (see Example 4.1.4); the random tournament and the random hypergraphs are simple (see Theorem 5.5.9), and their extension properties shatter every finite set, as in Example 4.1.4. The rational order is NIP, being o-minimal [PS86], and is not simple (see Proposition 5.3.2). The Henson graph is neither (see Example 5.5.10).

The first consequence of simplicity concerns dividing. In a stable theory a type over a model is determined by its restriction to a submodel of size at most κ , as we show in Chapter 6; in Theorem 5.3.6 below we show that in a simple theory a type is nondividing over a subset of size at most κ . The proof runs the construction of the tree property backwards: a long sequence of dividing formulas is turned into a tree. The tree is built one level at a time, using automorphisms to copy a sequence that witnesses dividing above each leaf.

Lemma 5.3.5. *Let $\varphi(x, y)$ be an L -formula, and let $k \geq 2$ be given.*

- (1) *Let $n \geq 1$, and let $a_0, \dots, a_{n-1}$ be tuples such that $\{\varphi(x, a_m) : m < n\}$ is consistent and, for every $m < n$, the formula $\varphi(x, a_m)$ k -divides over $A_m = \{a_l : l < m\}$. Then there are tuples b_σ , for $\sigma \in \omega^{\leq n}$ with $\sigma \neq \emptyset$, such that*
 - (i) *for every $\eta \in \omega^n$, we have $(b_{\eta|1}, \dots, b_{\eta|n}) \equiv (a_0, \dots, a_{n-1})$, and*
 - (ii) *for every $\sigma \in \omega^{<n}$, the set $\{\varphi(x, b_{\sigma \smallfrown \langle i \rangle}) : i < \omega\}$ is k -inconsistent.**In particular, these tuples form a tree of height n as in Remark 5.2.8.*
- (2) *If for every $n \geq 1$ there are tuples $a_0, \dots, a_{n-1}$ as in (1), possibly depending on n , then the formula φ has the k -tree property. This holds in particular if $(a_n)_{n < \omega}$ is a sequence such that $\{\varphi(x, a_n) : n < \omega\}$ is consistent and, for every n , the formula $\varphi(x, a_n)$ k -divides over $A_n = \{a_m : m < n\}$.*

Proof. (1) For each $m < n$, fix a sequence $(a_m^i)_{i < \omega}$, indiscernible over A_m , with $a_m^0 = a_m$ and $\{\varphi(x, a_m^i) : i < \omega\}$ k -inconsistent. We build the tuples level by level: for $p = 1, \dots, n$ we find b_σ for $\sigma \in \omega^{\leq p}$, $\sigma \neq \emptyset$, satisfying (i) and (ii) with p in place of n and $(a_0, \dots, a_{p-1})$ in place of $(a_0, \dots, a_{n-1})$. Condition (i) gives condition (a) of Definition 5.2.7 for the branches of length p , because $\{\varphi(x, a_m) : m < p\}$ is consistent. For $p = 1$, put $b_{\langle i \rangle} = a_0^i$; all terms of an indiscernible sequence have the same type, so (i) holds, and (ii) holds by the choice of (a_0^i) .

Suppose that the tuples have been built for some $p < n$. For each $\eta \in \omega^p$, by (i) and strong homogeneity there is $f_\eta \in \text{Aut}(\mathbb{M})$ with $f_\eta(a_0, \dots, a_{p-1}) = (b_{\eta|1}, \dots, b_{\eta|p})$. Keep the tuples already built, and put $b_{\eta \frown \langle i \rangle} = f_\eta(a_p^i)$ for $i < \omega$. Since $a_p^i \equiv_{A_p} a_p$, we have

$$(b_{\eta|1}, \dots, b_{\eta|p}, b_{\eta \frown \langle i \rangle}) = f_\eta(a_0, \dots, a_{p-1}, a_p^i) \equiv (a_0, \dots, a_{p-1}, a_p^i) \equiv (a_0, \dots, a_p),$$

which is (i) for $p+1$. The set $\{\varphi(x, b_{\eta \frown \langle i \rangle}) : i < \omega\}$ is the image under f_η of a k -inconsistent set, and hence is k -inconsistent; this is (ii) at the new level, and (ii) at the old levels is unchanged.

(2) By (1), for every $n \geq 1$ there is a tree of height n whose branches are consistent and whose siblings are k -inconsistent, and Remark 5.2.8 gives the k -tree property. For the last assertion, apply (1) to the initial segments $a_0, \dots, a_{n-1}$ of the sequence. ☺

Part (2) gives the local character of dividing at once. The finite form (1) is used in Chapter 7, where the tuples $a_0, \dots, a_{n-1}$ are produced afresh for each n .

Theorem 5.3.6 (Local character of dividing). *Assume that T is simple, and let a be a finite tuple and B a small set. Then there is $A \subseteq B$ with $|A| \leq \kappa$ such that $\text{tp}(a/B)$ does not divide over A .*

Proof. Suppose not. We choose, by recursion on $\alpha < \kappa^+$, finite tuples b_α from B , formulas $\varphi_\alpha(x, y) \in L$ and integers $k_\alpha \geq 2$. Put $B_\alpha = \bigcup_{\beta < \alpha} b_\beta$, a subset of B of size at most κ . By assumption $\text{tp}(a/B)$ divides over B_α , so some formula of $\text{tp}(a/B)$, say $\varphi_\alpha(x, b_\alpha)$ with $\varphi_\alpha \in L$ and b_α from B , divides over B_α , and by Lemma 5.2.4 it k_α -divides over B_α for some $k_\alpha \geq 2$.

There are only κ pairs (φ, k) with $\varphi \in L$ and $k < \omega$, and κ^+ is regular, so some pair (φ, k) equals $(\varphi_\alpha, k_\alpha)$ for unboundedly many α . Choose such $\alpha_0 < \alpha_1 < \dots$, and put $a_n = b_{\alpha_n}$. The set $A_n = \{a_m : m < n\}$ is contained in B_{α_n} , so the formula $\varphi(x, a_n)$ k -divides over A_n by Remark 5.2.5. The formulas $\varphi(x, a_n)$ all belong to $\text{tp}(a/B)$, and so they form a consistent set. By Lemma 5.3.5(2), the formula φ has the k -tree property, contradicting simplicity. ☺



Remark 5.3.7. The converse of Theorem 5.3.6 holds as well: if every type over a model is nondividing over a submodel of size at most κ , then the theory T is simple (see Corollary 7.6.4); the version over arbitrary sets is due to Shelah [She80]. Kim and Pillay [KP97], building on Kim [Kim98], characterized the simple theories by an abstract relation of independence, as we do for the stable theories in Chapter 6; their theorem is the subject of Chapter 7. In Sections 5.4 and 5.5 we prove the corresponding statements directly, from the axioms, for the random graph and for a class of Fraïssé limits.

5.4 THE RADO GRAPH


The goal of this section is to compute dividing in the random graph, from the extension axioms alone, and to derive from the computation that the random graph is simple and that its independence relation satisfies the independence theorem but not uniqueness. Let T_{RG} be the theory of Example 4.1.4: the axioms say that E is irreflexive and symmetric and, for all

m, n , that for all disjoint finite sets U, V of vertices with $|U| = m$ and $|V| = n$ some vertex outside $U \cup V$ is adjacent to every vertex of U and to no vertex of V . These are the *extension axioms*. A countable model is easy to write down.

Example 5.4.1. Let the vertices be the natural numbers, and for $i < j$ join i and j if the i -th binary digit of j is 1 (digits counted from the right, starting at 0). This graph satisfies the extension axioms. Let U, V be disjoint finite sets of natural numbers, choose N larger than every element of $U \cup V$, and put

$$z = 2^N + \sum_{u \in U} 2^u.$$

Then we have $z \geq 2^N > N$, so we have $z \notin U \cup V$ and z is larger than every element of $U \cup V$. The binary digits of z at the positions $u \in U$ are 1 and at the positions $v \in V$ are 0, so the element z is adjacent to every vertex of U and to no vertex of V . The graph is called the *Rado graph* $\mathcal{R}$ (see Figure 5.4).



j	binary	neighbours $i < j$
1	001	$\{0\}$
2	010	$\{1\}$
3	011	$\{0, 1\}$
4	100	$\{2\}$
5	101	$\{0, 2\}$
6	110	$\{1, 2\}$
7	111	$\{0, 1, 2\}$

FIGURE 5.4. The vertices $0, \dots, 7$ of the Rado graph of Example 5.4.1. For $i < j$, the vertex j is joined to i exactly when the binary expansion of j has digit 1 in position i ; the table lists the expansions. The new neighbour of U that avoids V is obtained by writing 1 in the positions of U , 0 in the positions of V , and a 1 far to the left.

Fact 5.4.2. *The theory T_{RG} is complete, $\aleph_0$ -categorical, and admits elimination of quantifiers; its countable model is the Rado graph (see [Hod93, Theorem 7.4.1], or Fact 5.5.2 below).*

From now until the end of this section, the theory T is T_{RG} . By quantifier elimination, the type of a tuple over a set is determined by the equalities and edges that it satisfies with parameters from the set. We use this in the following form: for small tuples a, a' and a small set C ,

$$(5.1) \quad a \equiv_C a' \iff \begin{array}{l} \text{the map that fixes } C \text{ pointwise and sends each} \\ \text{coordinate of } a \text{ to the corresponding coordinate of } a' \text{ is} \\ \text{a well defined bijection from } C \cup a \text{ onto } C \cup a' \text{ that} \\ \text{preserves edges and non-edges.} \end{array}$$

An injective map between graphs that preserves edges and non-edges is an *embedding*. The computation of dividing rests on two tools. The first is the extension axioms, which let us embed into $\mathbb{M}$ any graph that agrees with $\mathbb{M}$ on a given small set.

Lemma 5.4.3. *Let C be a small set of vertices, and let X be a graph whose vertex set is the disjoint union of C and a small set X_0 , and whose restriction to C is the graph induced by $\mathbb{M}$ on C . Then there is an embedding $g: X \rightarrow \mathbb{M}$ with $g(c) = c$ for every $c \in C$.*

Proof. Introduce a variable x_v for each $v \in X_0$, and let Σ be the set of formulas $x_v \neq x_w$, $x_v \neq c$, $E(x_v, x_w)^u$ and $E(x_v, c)^u$ (for $v \neq w$ in X_0 and $c \in C$) that describe X , where $u \in \{0, 1\}$ records whether the edge is present in X . A realization of Σ is an embedding of the required kind, and Σ is a small set of formulas over C , so by saturation it suffices to show that every finite subset of Σ is satisfiable. A finite subset mentions finite sets $C_0 \subseteq C$ and $\{v_1, \dots, v_m\} \subseteq X_0$. Choose images $w_1, \dots, w_m \in \mathbb{M}$ one at a time: if $w_1, \dots, w_{l-1}$ have been chosen, apply the extension axiom to U , the set of elements of $C_0 \cup \{w_1, \dots, w_{l-1}\}$ that must be adjacent to w_l , and V , the set of the others. The vertex w_l so obtained lies outside $C_0 \cup \{w_1, \dots, w_{l-1}\}$, so the new vertices are distinct from each other and from C_0 . ☺

The second tool is Ramsey's theorem, in the sharper form announced after Lemma 5.2.2: over a finite set, an infinite subsequence of the given sequence is itself indiscernible.

Lemma 5.4.4. *Let D be a finite set of vertices, and let $(b_i)_{i < \omega}$ be tuples of the same finite length. Then there is an infinite set $H \subseteq \omega$ such that $(b_i)_{i \in H}$ is indiscernible over D .*

Proof. Call the *diagram* over D of a tuple the set of equalities and edges, and of their negations, that hold among its coordinates and elements of D . Since D is finite and the tuples have fixed length, there are finitely many possible diagrams over D of one tuple b_i , and of a pair (b_i, b_j) . By the pigeonhole principle, all b_i with i in some infinite set have the same diagram over D ; colouring each pair $i < j$ of this set by the diagram of (b_i, b_j) over D , Ramsey's theorem gives an infinite homogeneous set H . Each equality or edge involves at most two vertices, hence at most two of the tuples. Hence, for $i_1 < \dots < i_n$ in H , the diagram of $(b_{i_1}, \dots, b_{i_n})$ over D is determined by the diagrams of the b_{i_l} and of the pairs (b_{i_l}, b_{i_m}) with $l < m$, which do not depend on the choice of $i_1 < \dots < i_n$ in H . By (5.1), the sequence $(b_i)_{i \in H}$ is indiscernible over D . ☺

We can now compute dividing. The answer is the one suggested by the introduction to this chapter: a type over $C \cup B$ divides over C exactly when it puts one of the variables equal to an element of B outside C .

Theorem 5.4.5 (Dividing in the Rado graph). *Let a be a finite tuple, and let C and B be small sets. Then the type $\text{tp}(a/C \cup B)$ divides over C if and only if $a \cap B \not\subseteq C$.*

Proof. Suppose first that some coordinate a_j of a lies in $B \setminus C$, and write $d = a_j$. The formula $x_j = d$ belongs to $\text{tp}(a/C \cup B)$. By Lemma 5.4.3, applied to the graph on $C \cup \{d'_i : i < \omega\}$ in which each new vertex d'_i is adjacent to $c \in C$ exactly when d is, and no two new vertices are adjacent, there are distinct vertices $d''_i \notin C$ with these adjacencies. By (5.1), any two increasing n -tuples of (d''_i) have the same type over C , and $d''_0 \equiv_C d$. Choose $f \in \text{Aut}(\mathbb{M}/C)$ with $f(d''_0) = d$. The sequence $(f(d''_i))_{i < \omega}$ is indiscernible over C , begins with d , and consists of distinct vertices, so the set $\{x_j = f(d''_i) : i < \omega\}$ is 2-inconsistent. Hence the type $\text{tp}(a/C \cup B)$ divides over C .

Conversely, suppose that $a \cap B \subseteq C$. Let $\varphi(x, b)$ be a formula of $\text{tp}(a/C \cup B)$, with $\varphi \in L$ and b a finite tuple from $C \cup B$, and let $(b_i)_{i < \omega}$ be indiscernible over C with $b_0 = b$. We find a single tuple a'' with $a''b_i \equiv_C ab$ for all i ; then the tuple a'' satisfies every $\varphi(x, b_i)$, so the formula $\varphi(x, b)$ does not divide over C . The idea is to join a fresh copy of a to each b_i exactly

as a is joined to b . For this to be well defined, coordinates shared by different terms of the sequence must play the same role in b . Write $e[p]$ for the p -th coordinate of a tuple e .

© **Claim.** If $i \neq j$ and $b_i[p] = b_j[q]$, then we have $b[p] = b[q]$, and all terms of the sequence have the same p -th coordinate and the same q -th coordinate.

Exchanging the roles of (i, p) and (j, q) if necessary, we may assume $i < j$. Choose $l > j$. Since $(b_i, b_j) \equiv_C (b_i, b_l)$, we have $b_i[p] = b_l[q]$, and since $(b_i, b_j) \equiv_C (b_j, b_l)$, we have $b_j[p] = b_l[q]$. Hence we have $b_j[p] = b_i[p] = b_j[q]$. As $b_j \equiv_C b$, this gives $b[p] = b[q]$. Moreover, since $b_i[p] = b_j[p]$ with $i < j$, indiscernibility gives $b_{i'}[p] = b_{j'}[p]$ for all $i' < j'$; thus the p -th coordinate is constant. Finally, $b_{i'} \equiv_C b$ and $b[p] = b[q]$ give $b_{i'}[p] = b_{i'}[q]$ for every i' , so the q -th coordinate is constant as well. This proves the claim (see Figure 5.5).

	[1]	[2]	[3]	[4]	
b_0	u_0	e	v_0	e'	e, e' : constant columns, the only elements shared by two different terms
b_1	u_1	e	v_1	e'	
b_2	u_2	e	v_2	e'	
b_3	u_3	e	v_3	e'	
b_4	u_4	e	v_4	e'	
					u_i, v_i : new in each term

FIGURE 5.5. The claim in the proof of Theorem 5.4.5. The rows are the terms of a sequence indiscernible over C , and the columns are their coordinates. An element outside C that occurs in two different terms occupies a whole column (pink): it is the same coordinate of every term. The fresh copy of a can therefore be joined to e and e' without ambiguity, exactly as a is joined to the corresponding coordinates of b .

With the claim in hand, we build the fresh copy of a . Let I be the set of indices v with $a_v \notin C$. For $v \in I$ we have $a_v \notin b$, because $b \subseteq C \cup B$ and $a \cap B \subseteq C$. Let $V = \bigcup_i b_i \setminus C$, and let X be the graph on $C \cup V$ together with fresh vertices a'_v for $v \in I$, where $a'_v = a'_w$ exactly when $a_v = a_w$, with the following edges:

- (i) on $C \cup V$, the edges of $\mathbb{M}$;
- (ii) the vertex a'_v is adjacent to $c \in C$ if and only if a_v is adjacent to c , and a'_v is adjacent to a'_w if and only if a_v is adjacent to a_w ;
- (iii) if $e = b_i[p] \in V$, then the element a'_v is adjacent to e if and only if a_v is adjacent to $b[p]$.

Clause (iii) does not depend on the choice of i and p with $e = b_i[p]$: if also $e = b_j[q]$, then we have $b[p] = b[q]$, by the claim when $i \neq j$, and because $b_i \equiv_C b$ when $i = j$. By Lemma 5.4.3, there is an embedding g of X into $\mathbb{M}$ fixing $C \cup V$. Let a'' be the tuple with $a''_v = g(a'_v)$ for $v \in I$ and $a''_v = a_v$ for $v \notin I$. Fix i . The map that fixes C , sends a to a'' and b to b_i coordinatewise is a well defined bijection from $C \cup a \cup b$ onto $C \cup a'' \cup b_i$: on $C \cup b$ it is the bijection given by $b \equiv_C b_i$, the coordinates a_v with $v \in I$ lie outside $C \cup b$, and their images lie outside $C \cup V$. By (i)–(iii) this map preserves edges and non-edges, so we have $a''b_i \equiv_C ab$ by (5.1). ☺

A dividing step therefore uses up a coordinate of a , and a tuple of length n can divide at most n times. This is the heart of the proof that the random graph is simple.

Theorem 5.4.6. *Let a be a tuple of length n , and let $D_0 \subseteq D_1 \subseteq \dots \subseteq D_m$ be small sets such that $\text{tp}(a/D_{j+1})$ divides over D_j for every $j < m$. Then we have $m \leq n$. The theory T_{RG} is simple.*

Proof. By Theorem 5.4.5, applied with $C = D_j$ and $B = D_{j+1}$, the tuple a meets D_{j+1} in a coordinate outside D_j . Thus the sets $a \cap D_0 \subsetneq a \cap D_1 \subsetneq \dots \subsetneq a \cap D_m$ increase strictly, and a has at least m distinct coordinates, so we have $m \leq n$.

Suppose now that $\varphi(x, y)$ has the k -tree property, witnessed by tuples b_σ , and let n be the length of x . The idea is to find, along a branch, $n + 1$ dividing steps, which the first part forbids. For this we prune the tree so that siblings become indiscernible over the nodes below them. By recursion, define a set $P \subseteq \omega^{\leq n+1}$ of nodes: the empty sequence belongs to P , and if $\sigma \in P$ has length at most n , let $D_\sigma = \bigcup_{1 \leq l \leq |\sigma|} b_{\sigma|l}$, a finite set (empty when $\sigma = \emptyset$), choose by Lemma 5.4.4 an infinite set $H_\sigma \subseteq \omega$ such that $(b_{\sigma \smallfrown \langle i \rangle})_{i \in H_\sigma}$ is indiscernible over D_σ , and put $\sigma \smallfrown \langle i \rangle \in P$ for $i \in H_\sigma$. Choose $\eta \in \omega^{n+1}$ with $\eta|l \in P$ for every l , let a realize $\{\varphi(x, b_{\eta|l}) : 1 \leq l \leq n + 1\}$, and put $D_l = D_{\eta|l}$ for $l \leq n + 1$, where $D_{\eta|n+1}$ is defined by the same formula.

Let $l \leq n$ and $\sigma = \eta|l$. The terms $b_{\sigma \smallfrown \langle i \rangle}$ with $i \in H_\sigma$ and $i \geq \eta(l)$ form a sequence indiscernible over D_l that begins with $b_{\eta|l+1}$, and the corresponding formulas are k -inconsistent. Hence the formula $\varphi(x, b_{\eta|l+1})$ divides over D_l , and, since this formula belongs to $\text{tp}(a/D_{l+1})$, the type $\text{tp}(a/D_{l+1})$ divides over D_l . The chain $D_0 \subseteq \dots \subseteq D_{n+1}$ has $n + 1$ dividing steps, contradicting the first part of the theorem (see Figure 5.6). ☺




FIGURE 5.6. The proof of Theorem 5.4.6 for a pair $x = (x_1, x_2)$, so we have $n = 2$. The tree is pruned so that the siblings above each node are indiscernible over the parameters below it (framed). Each step up the branch (pink) is a dividing step, and by Theorem 5.4.5 it catches a new coordinate of the realization a in the parameters D_l collected so far. After two steps both coordinates are caught, and a third dividing step is impossible; a tree of height 3 cannot exist.

Simplicity settled, we turn to the second aim of this section, the independence relation of the random graph. An independence relation over a model M is a relation $A \perp_M B$, read “ A is independent from B over M ”, and the following axioms are the yardstick for such relations. In Chapter 6 we show that a theory is stable exactly when it carries a relation satisfying all of them (see Theorem 6.2.2); in Theorem 5.4.9 below we show that the random graph carries a relation satisfying all of them but one.

Before the axioms are stated, we explain what they are meant to accomplish. A relation deserves the name of independence only if it expresses the thought that B tells us nothing about A beyond what M already tells us, and several requirements follow from this thought at once. Whether A is independent from B over M should depend only on how A , B and M sit in the monster model, and not on how they are named (INV). What is independent from a set should be independent from its parts (MON), and whether it is should be decided by finite parts alone (FC). If B tells us nothing new about A , then the set A should tell us nothing new about B (SYM). Every type over M should be realizable in a position of independence from any given set (EX). Enlarging the base within B should not destroy independence (BMON), and independence over a chain of bases should compose (TR). A finite tuple should be independent from a model over some small submodel (LOC), so that independence can always be witnessed by a small base. Finally, if B tells us nothing new about A , there is nothing to choose between two ways of placing A independently from B , and so the independent extension should be unique (UN). The last requirement is the strongest, and it is the one that will fail in the random graph (see Theorem 5.4.9(b)). The axioms below are these requirements, made exact.

Definition 5.4.7. An *independence relation (over models)* is a relation $A \downarrow_M B$ between small sets A, B and models M satisfying the following axioms, for all small A, A', B, B' and all models M, N .

- (INV) *Invariance.* If $A \downarrow_M B$ and $\sigma \in \text{Aut}(\mathbb{M})$, then we have $\sigma(A) \downarrow_{\sigma(M)} \sigma(B)$.
- (MON) *Monotonicity.* If $A \downarrow_M B$, $A' \subseteq A$ and $B' \subseteq B$, then we have $A' \downarrow_M B'$.
- (BMON) *Base monotonicity.* If $A \downarrow_M B$ and N is a model with $M \subseteq N \subseteq B$, then we have $A \downarrow_N B$.
- (TR) *Transitivity.* If $M \subseteq N$, $A \downarrow_M N$ and $A \downarrow_N B$, then we have $A \downarrow_M NB$.
- (FC) *Finite character.* The relation $A \downarrow_M B$ holds if and only if $a \downarrow_M b$ for all finite tuples a from A and b from B .
- (SYM) *Symmetry.* The relation $A \downarrow_M B$ holds if and only if $B \downarrow_M A$.
- (EX) *Existence (extension).* For all A, M, B there is $A' \equiv_M A$ with $A' \downarrow_M B$.
- (UN) *Uniqueness (stationarity).* If $A \equiv_M A'$, $A \downarrow_M B$ and $A' \downarrow_M B$, then we have $A \equiv_{MB} A'$.
- (LOC) *Local character.* For every finite tuple a and every model N there is a model $M \subseteq N$ with $|M| \leq \kappa$ and $a \downarrow_M N$.

The relation is suggested by the computation of Theorem 5.4.5: a tuple should be independent from a set when it shares with it nothing outside the base.

Definition 5.4.8. For small sets A, B, C , write $A \downarrow_C^t B$ if $A \cap B \subseteq C$. We call $\downarrow^t$ *trivial independence*; the superscript distinguishes it from the abstract relations $\downarrow$ of Definition 5.4.7.

For a finite tuple a , Theorem 5.4.5 says that $a \downarrow_C^t B$ if and only if $\text{tp}(a/C \cup B)$ does not divide over C . We compare the relation $\downarrow^t$ with the axioms of Definition 5.4.7, read with arbitrary models M, N of T_{RG} as bases, and with one further axiom, *anti-reflexivity* (AR): if $a \downarrow_M a$ for a finite tuple a , then we have $a \subseteq M$. It reappears for stable theories in Theorem 6.2.2.

Theorem 5.4.9 (Independence in the Rado graph). *The following hold in T_{RG} :*

- (a) *The relation $\perp^t$ satisfies all the axioms of Definition 5.4.7 except (UN), and it satisfies (AR).*
- (b) *The relation $\perp^t$ does not satisfy (UN).*
- (c) (Independence theorem.) *Let C be a small set and b_1, b_2, c_1, c_2 finite tuples. If $b_1 \perp_C^t b_2$, $c_1 \perp_C^t b_1$, $c_2 \perp_C^t b_2$ and $c_1 \equiv_C c_2$, then there is c with $c \equiv_{Cb_1} c_1$, $c \equiv_{Cb_2} c_2$ and $c \perp_C^t b_1 b_2$.*

Proof. (a) Invariance, monotonicity, finite character, symmetry and anti-reflexivity are immediate from the definition. *Base monotonicity:* if $A \cap B \subseteq M$ and $M \subseteq N \subseteq B$, then we have $A \cap B \subseteq N$. *Transitivity:* if $M \subseteq N$, $A \cap N \subseteq M$ and $A \cap B \subseteq N$, then we have $A \cap B \subseteq A \cap N \subseteq M$, so we have $A \cap (N \cup B) \subseteq M$. *Local character:* given a finite tuple a and a model N , the downward Löwenheim–Skolem theorem gives a model $M \preceq N$ with $a \cap N \subseteq M$ and $|M| \leq \kappa$, and then we have $a \perp_M^t N$. *Existence:* given A , M and B , let X be the graph on $M \cup B$ together with fresh vertices a'_v , one for each coordinate a_v of A outside M (equal coordinates receiving equal vertices), where the edges on $M \cup B$ are those of $\mathbb{M}$, the vertices a'_v are joined to each other and to M as the a_v are, and no a'_v is joined to a vertex of $B \setminus M$. Lemma 5.4.3 embeds X into $\mathbb{M}$ over $M \cup B$, and the resulting tuple A' satisfies $A' \equiv_M A$ by (5.1) and $A' \cap B \subseteq M$.

(b) Let M be a model, and let b be a vertex not in M . By Lemma 5.4.3 there are vertices $a, a' \notin M \cup \{b\}$, adjacent to no vertex of M , with a adjacent to b and a' not. Then we have $a \equiv_M a'$ by (5.1), and $a \perp_M^t b$ and $a' \perp_M^t b$, but $a \not\equiv_{Mb} a'$.

(c) Let I be the set of indices v with $(c_1)_v \notin C$; since $c_1 \equiv_C c_2$, it is also the set of indices with $(c_2)_v \notin C$, and $(c_1)_v = (c_2)_v$ for $v \notin I$. Let X be the graph on $C \cup b_1 \cup b_2$, with the edges of $\mathbb{M}$, together with fresh vertices c'_v for $v \in I$ (equal coordinates receiving equal vertices), where the new vertices are joined to each other and to $C \cup b_1$ as the coordinates of c_1 are, and to $b_2 \setminus C$ as the coordinates of c_2 are. The two prescriptions overlap only on C , because $b_1 \cap b_2 \subseteq C$, and on C (and among the new vertices) they agree, because $c_1 \equiv_C c_2$. By Lemma 5.4.3, there is an embedding g of X into $\mathbb{M}$ fixing $C \cup b_1 \cup b_2$. Let c be the tuple with $c_v = g(c'_v)$ for $v \in I$ and $c_v = (c_1)_v$ for $v \notin I$. The coordinates $(c_1)_v$ with $v \in I$ lie outside $C \cup b_1$, since $c_1 \cap b_1 \subseteq C$, and the coordinates c_v with $v \in I$ lie outside $C \cup b_1 \cup b_2$. Hence the map fixing $C \cup b_1$ and sending c_1 to c is a bijection preserving edges and non-edges, and $c \equiv_{Cb_1} c_1$ by (5.1). In the same way $c \equiv_{Cb_2} c_2$, and $c \cap (b_1 \cup b_2) \subseteq C$. ☺

Figure 5.7 shows the construction of (c) in the smallest interesting case.

Stability and simplicity, side by side

In a stable theory, as we show in Chapter 6, every type over a model M has exactly one independent extension to $M \cup B$: the heir, given by the definition scheme of the type. In the Rado graph, a type over M has many independent extensions to $M \cup \{b\}$, one for each way of joining the new vertices to b , and uniqueness fails (see Theorem 5.4.9(b)). In its place stands the independence theorem (see Theorem 5.4.9(c)): independent extensions over two independent sets can always be realized together. Dividing is measured by coordinates (see Theorem 5.4.5), each dividing step uses up one coordinate, and so no formula can have the tree property (see Theorem 5.4.6).




FIGURE 5.7. The independence theorem in the Rado graph, with $C = \{e\}$ (see Theorem 5.4.9(c)). Solid lines are edges and dashed lines non-edges. The vertices c_1 and c_2 have the same type over C (both are adjacent to e); the vertex c_1 is adjacent to b_1 , and c_2 is not adjacent to b_2 . The vertex c inherits its edges to $C \cup \{b_1\}$ from c_1 and to $C \cup \{b_2\}$ from c_2 . The two prescriptions can only conflict on C , where they agree. A second vertex, joined to b_1 and b_2 in the same way but differing elsewhere, would be just as good: the amalgam exists, but it is not unique.

5.5 FRAÏSSÉ LIMITS

The goal of this section is to find the property of finite graphs on which Section 5.4 depends, and to prove that every Fraïssé limit with this property is simple. In the proofs of Theorems 5.4.5 and 5.4.9, new vertices were joined to old ones as prescribed by several types at once, and Lemma 5.4.3 embedded the resulting graph into $\mathbb{M}$. Every graph is allowed in the random graph, so the prescriptions only had to be consistent with one another. In a Fraïssé limit only the structures of a given class are allowed, and the question becomes whether the structure assembled from the prescriptions belongs to the class. In the proof of the independence theorem, three structures are given pairwise, on $C \cup c \cup b_1$, on $C \cup c \cup b_2$ and on $C \cup b_1 \cup b_2$, and they must be completed to a structure on the union. This is a 3-amalgamation problem.

Throughout this section, the language L is finite and relational. For an L -structure N , the *age* of N is the class of finite L -structures that embed into N .

Definition 5.5.1. A *Fraïssé class* is a nonempty class $\mathcal{K}$ of finite L -structures, closed under isomorphism, with the following properties:

- (a) *hereditary property*: every substructure of a structure in $\mathcal{K}$ belongs to $\mathcal{K}$;
- (b) *joint embedding property*: any two structures in $\mathcal{K}$ embed into a common structure in $\mathcal{K}$;
- (c) *amalgamation property*: if $A, B_1, B_2 \in \mathcal{K}$ and $f_l: A \rightarrow B_l$ are embeddings, then there are $B \in \mathcal{K}$ and embeddings $g_l: B_l \rightarrow B$ with $g_1 \circ f_1 = g_2 \circ f_2$.

An L -structure N is *ultrahomogeneous* if every isomorphism between finite substructures of N extends to an automorphism of N . Given a Fraïssé class $\mathcal{K}$, a structure N has the *extension property* for $\mathcal{K}$ if, whenever $A \subseteq B$ belong to $\mathcal{K}$, with $|B \setminus A| = 1$ and A possibly empty, every embedding of A into N extends to an embedding of B into N .

Fact 5.5.2 (Fraïssé [Fra54]; see [Hod93, Theorems 7.1.2 and 7.4.1]). *Let $\mathcal{K}$ be a Fraïssé class.*

- (a) *There is a countable ultrahomogeneous L -structure $M_{\mathcal{K}}$ whose age is $\mathcal{K}$, and it is unique up to isomorphism.*

- (b) The theory $T_{\mathcal{K}}$ of $M_{\mathcal{K}}$ is $\aleph_0$ -categorical and admits elimination of quantifiers.
- (c) An L -structure N is a model of $T_{\mathcal{K}}$ if and only if every finite substructure of N belongs to $\mathcal{K}$ and N has the extension property for $\mathcal{K}$.

The structure $M_{\mathcal{K}}$ is the *Fraïssé limit* of $\mathcal{K}$.

Example 5.5.3. The following are Fraïssé classes; the amalgamation property is checked by taking unions (orienting the new pairs, for tournaments) or, for linear orders, by interleaving.

- (1) The finite sets, in the language with equality only. The limit is a countably infinite set. Its theory is stable: by quantifier elimination, the type of an n -tuple over a set A is determined by the equalities among its coordinates and with elements of A , so there are at most $|A| + \aleph_0$ such types; the theory is λ -stable for every infinite λ , and stable by Lemma 5.1.2.
- (2) The finite linear orders. The limit is $(\mathbb{Q}, <)$.
- (3) The finite graphs. The limit is the Rado graph, and $T_{\mathcal{K}}$ is T_{RG} .
- (4) The finite tournaments (complete graphs with an orientation of each edge). The limit is the *random tournament*.
- (5) For $r \geq 2$, the finite r -uniform hypergraphs. The limit is the *random r -hypergraph*.
- (6) The finite triangle-free graphs. The limit is the *Henson graph* $\mathcal{H}$.

In the proofs of Section 5.4, the new vertices were always fresh: no new vertex was identified with an old one. In Definition 5.5.4 below, we ask the same of amalgamation in $\mathcal{K}$. For Fraïssé classes, we treat structures as sets with relations and say that a structure D *extends* a structure B if B is a substructure of D .

Definition 5.5.4. Let $\mathcal{K}$ be a Fraïssé class.

- (a) The class $\mathcal{K}$ has *disjoint 2-amalgamation* if whenever $B_1, B_2 \in \mathcal{K}$ have a common substructure A with $B_1 \cap B_2 = A$, some structure in $\mathcal{K}$ on $B_1 \cup B_2$ extends B_1 and B_2 .
- (b) The class $\mathcal{K}$ has *disjoint 3-amalgamation* if the following holds. Let B_1, B_2, B_3 be structures with a common substructure $A \in \mathcal{K}$ and with $B_l \cap B_m = A$ for $l \neq m$, and let $B_{12}, B_{13}, B_{23} \in \mathcal{K}$ be structures on $B_1 \cup B_2, B_1 \cup B_3, B_2 \cup B_3$ such that B_{lm} extends B_l and B_m . Then some structure in $\mathcal{K}$ on $B_1 \cup B_2 \cup B_3$ extends B_{12}, B_{13} and B_{23} .

In a 3-amalgamation problem every tuple contained in the union of two of the sets B_l already has its relations prescribed, by one of the structures B_{lm} ; what remains free is the relations on tuples that meet all three of the sets $B_l \setminus A$. Figure 5.8 shows a 3-amalgamation problem in general, and two problems without solution.

Example 5.5.5. All classes of Example 5.5.3 have disjoint 2-amalgamation. For graphs, triangle-free graphs and hypergraphs, take the union of the two structures, with no relations between $B_1 \setminus A$ and $B_2 \setminus A$; a triangle in the union that is contained neither in B_1 nor in B_2 would need two adjacent vertices, one in each of $B_1 \setminus A$ and $B_2 \setminus A$. For tournaments, orient the new pairs from B_1 to B_2 . For linear orders, put the elements of $B_1 \setminus A$ before those of $B_2 \setminus A$ within each interval determined by A .

The classes (1), (3), (4) and (5) have disjoint 3-amalgamation. For sets, graphs and tournaments, every relation involves at most two elements, so in a 3-amalgamation problem every relation is already prescribed and the union of B_{12} , B_{13} and B_{23} is a solution. For r -uniform hypergraphs, declare that no hyperedge meets all three of the sets $B_i \setminus A$.

The classes (2) and (6) do not have disjoint 3-amalgamation. Take $A = \emptyset$ and one-element sets $B_1 = \{x\}$, $B_2 = \{y\}$, $B_3 = \{z\}$. For linear orders, let B_{12} , B_{23} , B_{13} say that $x < y$, $y < z$ and $z < x$: no linear order extends them. For triangle-free graphs, let each B_{lm} be an edge: the only graph on $\{x, y, z\}$ extending them is a triangle.




FIGURE 5.8. Left: a disjoint 3-amalgamation problem (see Definition 5.5.4(b)). The sets B_1, B_2, B_3 meet pairwise in A ; the structures B_{12}, B_{13}, B_{23} are given on the three pairwise unions, and a solution is a structure in $\mathcal{K}$ on the whole union extending all three. Right: two problems without solution, with $A = \emptyset$ and one point in each B_i (see Example 5.5.5). Each pair is consistent, and only the triple is forbidden.

For the rest of this section, we fix a Fraïssé class $\mathcal{K}$, put $T = T_{\mathcal{K}}$, and let $\mathbb{M}$ be a monster model of T . By quantifier elimination, (5.1) holds with “preserves edges and non-edges” replaced by “is an isomorphism of substructures”; we refer to this version as (5.1) as well. The two tools of Section 5.4, the embedding of graphs over a set (see Lemma 5.4.3) and the extraction of indiscernible subsequences by Ramsey’s theorem (see Lemma 5.4.4), generalize directly.

Lemma 5.5.6. *Let C be a small subset of $\mathbb{M}$, and let X be an L -structure whose universe is the disjoint union of C and a small set X_0 , whose restriction to C is the substructure of $\mathbb{M}$ on C , and whose finite substructures all belong to $\mathcal{K}$. Then there is an embedding $g: X \rightarrow \mathbb{M}$ with $g(c) = c$ for every $c \in C$.*

Proof. Argue as in the proof of Lemma 5.4.3, with the atomic formulas of L in place of edges. Given finite sets $C_0 \subseteq C$ and $\{v_1, \dots, v_m\} \subseteq X_0$, the substructures of X on $C_0 \cup \{v_1, \dots, v_l\}$, for $l \leq m$, belong to $\mathcal{K}$, and each extends the previous one by one element. Applying the extension property of Fact 5.5.2(c) m times, starting from the inclusion of C_0 into $\mathbb{M}$, gives an embedding of the substructure on $C_0 \cup \{v_1, \dots, v_m\}$ into $\mathbb{M}$ over C_0 . ☺

Lemma 5.5.7. *Let D be a finite subset of $\mathbb{M}$, and let $(b_i)_{i < \omega}$ be tuples of the same finite length. Then there is an infinite set $H \subseteq \omega$ such that $(b_i)_{i \in H}$ is indiscernible over D .*

Proof. Let r be the larger of 2 and the maximal arity of a symbol of L . Argue as in the proof of Lemma 5.4.4, colouring each r -element set $i_1 < \dots < i_r$ by the diagram over D of $(b_{i_1}, \dots, b_{i_r})$, that is, by the set of atomic formulas with parameters from D , and of their negations, satisfied by it. Let H be infinite and homogeneous. Each atomic formula involves at most r of the tuples. If it involves the tuples $b_{j_1}, \dots, b_{j_s}$, with $j_1 < \dots < j_s$ in H and $s \leq r$, then its truth value is read off from the diagram of any increasing r -tuple from H that begins with $j_1, \dots, j_s$, and so depends only on the position of the tuples involved. Hence all increasing n -tuples from H have the same diagram over D , and the sequence is indiscernible over D by (5.1). $\odot$

The first half of the proof of Theorem 5.4.5 needed an indiscernible sequence of distinct conjugates of a vertex, and in the Rado graph it was written down by hand. In general we obtain it from disjoint 2-amalgamation and the extraction of indiscernible sequences (see Lemma 5.2.2).

Lemma 5.5.8. *Assume that $\mathcal{K}$ has disjoint 2-amalgamation. Let C be a small set, and let d be an element of $\mathbb{M}$ not in C . Then there is a sequence $(d_i)_{i < \omega}$ of distinct elements, indiscernible over C , with $d_0 = d$.*

Proof. First, there are infinitely many distinct realizations of $\text{tp}(d/C)$. By compactness and saturation, it suffices to find, for every finite $C_0 \subseteq C$ and every n , distinct elements with the same diagram over C_0 as d . Let Y be the substructure of $\mathbb{M}$ on $C_0 \cup \{d\}$. By $n-1$ applications of disjoint 2-amalgamation, there is a structure in $\mathcal{K}$ on $C_0 \cup \{d^1, \dots, d^n\}$, with new elements d^l , in which each $C_0 \cup \{d^l\}$ is a copy of Y over C_0 . By Lemma 5.5.6, it embeds into $\mathbb{M}$ over C_0 . Let $(d_i^l)_{i < \omega}$ be distinct realizations of $\text{tp}(d/C)$.

Next, apply Lemma 5.2.2 over C to the sequence (d_i^l) . Its Ehrenfeucht–Mostowski type over C contains every formula of $\text{tp}(d/C)$, in one variable, and the formula $y_1 \neq y_2$, so the lemma gives a sequence (d_i'') , indiscernible over C , of distinct realizations of $\text{tp}(d/C)$. An automorphism of $\mathbb{M}$ over C sending d_0'' to d carries it to the required sequence. $\odot$

In Theorem 5.5.9 below we show that, with these lemmas, the arguments of Section 5.4 go through for Fraïssé limits: disjoint 2-amalgamation supplies the conjugates that make a type divide, and disjoint 3-amalgamation supplies the common realizations that keep it from dividing.

Theorem 5.5.9. *Let $\mathcal{K}$ be a Fraïssé class in a finite relational language with disjoint 2-amalgamation and disjoint 3-amalgamation, and let $T = T_{\mathcal{K}}$.*

- (1) *For a finite tuple a and small sets C, B , the type $\text{tp}(a/C \cup B)$ divides over C if and only if $a \cap B \not\subseteq C$.*
- (2) *If a is a tuple of length n and $D_0 \subseteq \dots \subseteq D_m$ are small sets such that $\text{tp}(a/D_{j+1})$ divides over D_j for every $j < m$, then we have $m \leq n$. The theory T is simple.*
- (3) *The relation $\perp^t$ of Definition 5.4.8 satisfies all the axioms of Definition 5.4.7 except possibly (UN), together with (AR) and the independence theorem stated in Theorem 5.4.9(c).*

Proof. (1) If $a \cap B \not\subseteq C$, argue as in the first part of the proof of Theorem 5.4.5, with Lemma 5.5.8 in place of the explicit sequence (d_i'') .

Suppose that $a \cap B \subseteq C$. Let $\varphi(x, b)$ belong to $\text{tp}(a/C \cup B)$, with b a finite tuple from $C \cup B$, and let $(b_i)_{i < \omega}$ be indiscernible over C with $b_0 = b$. Let $\Sigma(x)$ be the set of formulas

$\psi(x, b_i)$, for $i < \omega$ and $\psi(x, y)$ an $L(C)$ -formula with $\psi(x, b) \in \text{tp}(a/C \cup b)$. A realization a'' of Σ satisfies $a''b_i \equiv_C ab$ for every i , and in particular every $\varphi(x, b_i)$, so it suffices to show, by saturation, that every finite subset Σ_0 of Σ is realized. Here, unlike in the Rado graph, the structure to be embedded is not determined by the prescriptions, and 3-amalgamation supplies it.

The set Σ_0 involves finitely many indices $i_1 < \dots < i_h$ and finitely many formulas, which by quantifier elimination we may take to be quantifier-free with parameters from a finite set $C_0 \subseteq C$; enlarge C_0 to contain $a \cap C$ and $b \cap C$. Let F be the set of elements not in C that belong to two distinct terms of the sequence. By the claim in the proof of Theorem 5.4.5, which used only indiscernibility, each element of F is a coordinate that is constant along the sequence, so we have $F \subseteq b$ and $F \subseteq b_i$ for every i . Since $b_i \cap C = b \cap C$ for every i , we obtain $b_i \cap b_j \subseteq C_0 \cup F$ for $i \neq j$. Put $A = C_0 \cup F$.

Let I be the set of indices v with $a_v \notin C$; as before, $a_v \notin b$ for $v \in I$. Take fresh elements a'_v for $v \in I$, equal exactly when the a_v are, and let a' be the tuple with these coordinates and with $a'_v = a_v$ for $v \notin I$. For $l \leq h$, let π_l be the map from $A \cup b_{i_l} \cup a'$ onto $A \cup b \cup a$ that fixes A , sends b_{i_l} to b coordinatewise, and sends a' to a . It is a well defined bijection: a coordinate $b_{i_l}[p]$ lies in A if and only if $b[p]$ does, and then the two are equal (for coordinates in C_0 because $b_{i_l} \equiv_C b$, and for coordinates in F because these are constant); the map is injective on b_{i_l} because $b_{i_l} \equiv_C b$; and the coordinates a_v with $v \in I$ lie outside $A \cup b$, because $a \cap b \subseteq C$ and $F \subseteq b$. Let Z_l be the structure on $A \cup b_{i_l} \cup a'$ that makes π_l an isomorphism onto the substructure of $\mathbb{M}$ on $A \cup b \cup a$. Then the structure Z_l belongs to $\mathcal{K}$, its restriction to $A \cup b_{i_l}$ is the substructure of $\mathbb{M}$ (because π_l there is the restriction of an automorphism of $\mathbb{M}$ over C sending b_{i_l} to b , and fixing F), and its restriction to $A \cup a'$ does not depend on l .

By induction on $l \leq h$, we build $Y_l \in \mathcal{K}$ on $A \cup a' \cup b_{i_1} \cup \dots \cup b_{i_l}$ extending $Z_1, \dots, Z_l$ and the substructure of $\mathbb{M}$ on $A \cup b_{i_1} \cup \dots \cup b_{i_l}$. Put $Y_1 = Z_1$. For $l \geq 2$, apply disjoint 3-amalgamation over A to

$$B_1 = A \cup a', \quad B_2 = A \cup b_{i_1} \cup \dots \cup b_{i_{l-1}}, \quad B_3 = A \cup b_{i_l},$$

with $B_{12} = Y_{l-1}$, $B_{13} = Z_l$, and B_{23} the substructure of $\mathbb{M}$ on $B_2 \cup B_3$. The sets B_1, B_2, B_3 meet pairwise in A : the elements a'_v with $v \in I$ are fresh, and $b_{i_l} \cap b_{i_j} \subseteq A$. The three structures agree on B_1, B_2 and B_3 , by the properties of the Z_l . A solution is the required Y_l (see Figure 5.9).

By Lemma 5.5.6, applied to Y_h over the finite set $A \cup b_{i_1} \cup \dots \cup b_{i_h}$, there is an embedding g of Y_h into $\mathbb{M}$ fixing this set. Put $a'' = g(a')$. For each l , the map $g \circ \pi_l^{-1}$ is an isomorphism from the substructure on $A \cup b \cup a$ onto the substructure on $A \cup b_{i_l} \cup a''$ that fixes C_0 , sends b to b_{i_l} , and sends a to a'' . Hence the tuples $a''b_{i_l}$ and ab satisfy the same quantifier-free formulas over C_0 , and a'' realizes Σ_0 .

(2) The first statement follows from (1) as in the proof of Theorem 5.4.6, and the second follows from the first as there, with Lemma 5.5.7 in place of Lemma 5.4.4.

(3) The axioms other than existence are checked as in the proof of Theorem 5.4.9(a). Existence and the independence theorem are proved as in Theorem 5.4.9(a) and (c), with compactness reducing each to finitely many formulas over a finite part C_0 of the base, and with the structure to be embedded supplied by amalgamation. For existence, given A, M and B , let A_0 and B_0 be finite parts of A and B , and let $C_0 \subseteq M$ be finite and contain $A_0 \cap M$ and $B_0 \cap M$. Apply disjoint 2-amalgamation over the substructure on C_0 to the substructure on $C_0 \cup B_0$ and to a copy over C_0 of the substructure on $C_0 \cup A_0$ whose new elements are disjoint from B_0 . For the independence theorem, apply disjoint 3-amalgamation




FIGURE 5.9. The inductive step in the proof of Theorem 5.5.9(1). The fresh copy a' of a has already been joined to the first $l - 1$ terms (the structure Y_{l-1}); it must now be joined to b_{i_l} as a is joined to b (the structure Z_l), while the terms of the sequence keep the structure they have in $\mathbb{M}$. Disjoint 3-amalgamation over A glues the three prescriptions into Y_l . In the Rado graph nothing has to be glued, because every relation is binary and every graph is allowed.

over the substructure on C_0 , where C_0 contains $(b_1 \cup b_2 \cup c_1) \cap C$, to

$$B_1 = C_0 \cup c', \quad B_2 = C_0 \cup b_1, \quad B_3 = C_0 \cup b_2,$$

where c' is a fresh copy of the coordinates of c_1 outside C , B_{12} makes $C_0 \cup c' \cup b_1$ a copy of $C_0 \cup c_1 \cup b_1$, B_{13} makes $C_0 \cup c' \cup b_2$ a copy of $C_0 \cup c_2 \cup b_2$, and B_{23} is the substructure of $\mathbb{M}$. The sets meet pairwise in C_0 because $b_1 \cap b_2 \subseteq C$, and B_{12} and B_{13} agree on B_1 because $c_1 \equiv_C c_2$. Embed the solution over $C_0 \cup b_1 \cup b_2$ by Lemma 5.5.6, and conclude by saturation. ☺

By Example 5.5.5, Theorem 5.5.9 applies to the Rado graph, the random tournament and the random r -hypergraphs, and also to the pure infinite set, which is stable. The rational order has disjoint 2-amalgamation but not disjoint 3-amalgamation, and it is not simple by Proposition 5.3.2. The Henson graph, which we treat in the next example, shows that disjoint 3-amalgamation cannot be dropped even when the forbidden configuration is a triangle rather than an order.

Example 5.5.10. Let $\mathcal{K}$ be the class of finite triangle-free graphs, so that $T_{\mathcal{K}}$ is the theory of the Henson graph $\mathcal{H}$. We show that the formula $\varphi(x; y_1, y_2)$, defined as $E(x, y_1) \wedge E(x, y_2)$, has the 2-tree property.

Let G be the graph with vertices u_σ, v_σ , for $\sigma \in \omega^{<\omega}$ with $\sigma \neq \emptyset$, whose edges are the pairs $\{u_{\sigma \smallfrown \langle i \rangle}, v_{\sigma \smallfrown \langle j \rangle}\}$ with $\sigma \in \omega^{<\omega}$ and $i < j$. Every edge joins a vertex u_τ to a vertex $v_{\tau'}$, so the graph G is bipartite and in particular triangle-free; by Lemma 5.5.6 with $C = \emptyset$, we may assume that G is an induced subgraph of $\mathbb{M}$. Put $b_\sigma = (u_\sigma, v_\sigma)$.

Siblings are 2-inconsistent. Let $i < j$. A solution of $\varphi(x, b_{\sigma \smallfrown \langle i \rangle}) \wedge \varphi(x, b_{\sigma \smallfrown \langle j \rangle})$ would be adjacent to $u_{\sigma \smallfrown \langle i \rangle}$ and to $v_{\sigma \smallfrown \langle j \rangle}$, which are adjacent to each other, and would form a triangle with them; every finite substructure of a model of $T_{\mathcal{K}}$ is triangle-free.

Branches are consistent. Two distinct nodes of a branch have different lengths and so are never siblings. Hence, for $\eta \in \omega^\omega$ and $n \geq 1$, the $2n$ vertices $u_{\eta \smallfrown m}, v_{\eta \smallfrown m}$ ($1 \leq m \leq n$) are pairwise nonadjacent. The graph obtained from them by adding one vertex adjacent to all of them is triangle-free, and Lemma 5.5.6 embeds it over these vertices; the new vertex

satisfies $\varphi(x, b_{\eta|m})$ for every $m \leq n$. By Remark 5.2.8, condition (a) of Definition 5.2.7 follows.

Hence the theory T_K is not simple. It also has the independence property: for pairwise nonadjacent vertices $b_1, \dots, b_n$ and a set $S \subseteq \{1, \dots, n\}$, adding a vertex adjacent exactly to the b_l with $l \in S$ creates no triangle, so the edge relation shatters $\{b_1, \dots, b_n\}$. Figure 5.10 shows the siblings of one node.




FIGURE 5.10. The siblings of one node in Example 5.5.10, written $b_i = (u_i, v_i)$. The vertex u_i is joined to v_j exactly when $i < j$: the bipartite graph is a ladder (compare Figure 1.1). A common solution x of $\varphi(x, b_1)$ and $\varphi(x, b_3)$ would close the triangle x, u_1, v_3 (pink), so siblings are 2-inconsistent. Along a branch, the pairs are pairwise nonadjacent, and a common neighbour exists.



Remark 5.5.11. The proof of Theorem 5.5.9(2) shows that a type of an n -tuple admits at most n successive dividing steps. In the language of simplicity theory, this says that the theory T is *supersimple*, and that every type of a single element has *SU-rank* at most 1. A theory is supersimple if the type of every finite tuple over every set B does not divide over some finite subset of B . The *SU-rank* of a type p over A is at least $\alpha + 1$ if p has an extension of *SU-rank* at least α that divides over A , and it is at least a limit ordinal λ if it is at least every $\alpha < \lambda$ (see [Wag00] or [Cas11]). Kruckman [Kru19] studied Fraïssé classes with disjoint n -amalgamation for every n , and showed that their limits have pseudofinite theories. The class of finite triangle-free graphs is the basic example of a *free amalgamation* class, in which the union of two structures over a common substructure, with no new relations, is always an amalgam. Shelah showed that the Henson graph has the property SOP_3 of his hierarchy of strict order properties, but not SOP_4 [She96, Claim 2.8], and Conant [Con17] developed the model theory of free amalgamation theories in general, showing in particular that they are NSOP_4 .

5.6 THE REAL-VALUED CASE

At the end of Chapters 1 and 4 we replaced formulas with bounded functions on arbitrary sets, and the same can be done for the tree property. Two things are lost in the passage. There is no compactness, so a tree must carry its own witnesses, branch by branch; and there is no definability in the sense of Chapter 1, so the proof that stable theories are simple, which counted definitions (see Proposition 5.3.1), has nothing to count. What takes the place of a definition is an approximation by polynomials in instances (see Theorem 1.2.11), and in

Theorem 5.6.4 below we show that approximations can be counted just as well, provided that we count them only up to a tolerance. The tolerance is supplied by a gap between the values that the tree requires on its branches and the values that it forbids among siblings. We treat the case $k = 2$ of Definition 5.2.7, in which siblings are pairwise incompatible.

Definition 5.6.1. Let A, B be sets, let $\varphi: A \times B \rightarrow [0, 1]$ be a function, and let $r < s$ be reals. The function φ has the *2-tree property with gap* (r, s) if there are $b_\sigma \in B$, for $\sigma \in \omega^{<\omega}$ with $\sigma \neq \emptyset$, such that the following hold:

- (a) for every $\sigma \in \omega^{<\omega}$ with $\sigma \neq \emptyset$, there is $a \in A$ with $\varphi(a, b_{\sigma|_m}) \geq s$ for $1 \leq m \leq |\sigma|$;
- (b) for every $\tau \in \omega^{<\omega}$ and all $i \neq j$, there is no $a \in A$ with $\varphi(a, b_{\tau \smallfrown \langle i \rangle}) > r$ and $\varphi(a, b_{\tau \smallfrown \langle j \rangle}) > r$.

The function φ has the *2-tree property* if it has the 2-tree property with gap (r, s) for some reals $r < s$.

Condition (a) asks for witnesses to the finite parts of the branches only; on arbitrary sets there is nothing that would turn them into witnesses for whole branches. Only gaps with $0 \leq r < s \leq 1$ matter for a $[0, 1]$ -valued function: a gap with $s > 1$ violates condition (a), and a gap with $r < 0$ violates condition (b), since condition (a) provides elements of A and every value exceeds r . For a formula the two definitions agree. Fix reals $0 \leq r < s \leq 1$; for a $\{0, 1\}$ -valued function, a value is at least s , or larger than r , exactly when it is 1. A formula $\varphi(x, y)$ has the 2-tree property of Definition 5.2.7 if and only if, for some model M , the function φ on $M^{|x|} \times M^{|y|}$ has the 2-tree property with gap (r, s) . Indeed, if the tuples b_σ witness the 2-tree property in $\mathbb{M}$, a model $M \preceq \mathbb{M}$ containing all of them (by the downward Löwenheim–Skolem theorem) contains witnesses for the finite parts of the branches, by elementarity, and pairs of siblings remain inconsistent in M . Conversely, every model of T embeds elementarily in $\mathbb{M}$, and a tree as in Definition 5.6.1 in a model $M \preceq \mathbb{M}$ is a tree in $\mathbb{M}$ whose finite branches are consistent and whose siblings are 2-inconsistent; Remark 5.2.8 gives the 2-tree property.

The proof that a stable function has no such tree counts branches, and the branches are counted in binary trees of infinite height, the real-valued form of the trees of Definition 2.1.2. In Lemma 5.6.3 below, we show that every 2-tree contains one.

Definition 5.6.2. Let $\varphi: A \times B \rightarrow [0, 1]$ be a function, and let $r < s$ be reals. A *binary tree of infinite height for φ with gap* (r, s) consists of families $(b_\sigma : \sigma \in 2^{<\omega})$ in B and $(a_\tau : \tau \in 2^{<\omega})$ in A such that, for all $\sigma, \tau \in 2^{<\omega}$,

$$\sigma \smallfrown \langle 1 \rangle \sqsubseteq \tau \implies \varphi(a_\tau, b_\sigma) \geq s, \quad \sigma \smallfrown \langle 0 \rangle \sqsubseteq \tau \implies \varphi(a_\tau, b_\sigma) \leq r.$$

Lemma 5.6.3. *If a function $\varphi: A \times B \rightarrow [0, 1]$ has the 2-tree property with gap (r, s) , then it has a binary tree of infinite height with gap (r, s) .*

Proof. Let (b_σ) witness the 2-tree property. Define $f: 2^{<\omega} \rightarrow \omega^{<\omega}$ by recursion: $f(\emptyset) = \emptyset$, $f(\sigma \smallfrown \langle 1 \rangle) = f(\sigma) \smallfrown \langle 0 \rangle$ and $f(\sigma \smallfrown \langle 0 \rangle) = f(\sigma) \smallfrown \langle 1 \rangle$. Then the map f preserves initial segments, and $f(\sigma \smallfrown \langle 0 \rangle)$ and $f(\sigma \smallfrown \langle 1 \rangle)$ are siblings. For $\sigma \in 2^{<\omega}$ put $b'_\sigma = b_{f(\sigma \smallfrown \langle 1 \rangle)}$, and for $\tau \in 2^{<\omega}$ let $a'_\tau \in A$ satisfy $\varphi(a'_\tau, b_{f(\tau)|_m}) \geq s$ for $1 \leq m \leq |f(\tau)|$, as condition (a) allows (for $\tau = \emptyset$ any element of A will do).

Let $\sigma, \tau \in 2^{<\omega}$ be given. If $\sigma \smallfrown \langle 1 \rangle \sqsubseteq \tau$, then we have $f(\sigma \smallfrown \langle 1 \rangle) \sqsubseteq f(\tau)$, so the element b'_σ lies on the branch below $f(\tau)$ and $\varphi(a'_\tau, b'_\sigma) \geq s$. If $\sigma \smallfrown \langle 0 \rangle \sqsubseteq \tau$, then we have $f(\sigma \smallfrown \langle 0 \rangle) \sqsubseteq f(\tau)$, so we have $\varphi(a'_\tau, b_{f(\sigma \smallfrown \langle 0 \rangle)}) \geq s > r$. The nodes $f(\sigma \smallfrown \langle 0 \rangle)$ and $f(\sigma \smallfrown \langle 1 \rangle)$ are siblings, and condition (b) gives $\varphi(a'_\tau, b'_\sigma) \leq r$. Thus the families (b'_σ) and (a'_τ) form a binary tree of infinite height with gap (r, s) . ☺

Figure 5.11 shows the embedding f . At the node $f(\sigma)$, the binary tree keeps two children of the ω -branching tree. The child 0 carries the parameter b'_σ of the binary node σ , and a

leaf that turns right at σ passes through it. A leaf that turns left passes through the child 1 instead, takes values at least s there, and so, by the incompatibility of siblings, takes values at most r at the child 0.




FIGURE 5.11. Lemma 5.6.3. The binary tree lives inside the ω -branching tree on the children 0 and 1 of the nodes it uses (pink); the other children (pale) are discarded. The parameter of the binary node σ is the child 0 of $f(\sigma)$. A leaf that turns right at σ passes through this child; a leaf that turns left passes through its sibling, the child 1, and the incompatibility of siblings forces the value at the child 0 down to at most r .

Theorem 5.6.4. *A stable function $\varphi: A \times B \rightarrow [0, 1]$ has no binary tree of infinite height with a gap. In particular, a stable function does not have the 2-tree property.*

Proof. The idea is to count. A binary tree of infinite height has $2^{\aleph_0}$ branches, and the types along distinct branches differ by at least $s - r$ at some node. A stable function, on the other hand, has only countably many approximate definitions over a countable set, and two branches with nearly the same definition cannot differ by $s - r$ anywhere.

Suppose that (b_σ) and (a_τ) form a binary tree of infinite height for φ with gap (r, s) , and put $\varepsilon = (s - r)/5$. Let $A_0 = \{a_\tau : \tau \in 2^{<\omega}\}$ and $B_0 = \{b_\sigma : \sigma \in 2^{<\omega}\}$, and let φ_0 be the restriction of φ to $A_0 \times B_0$. Every instance of Definition 1.2.5 for φ_0 is an instance for φ , so the function φ_0 is stable.

© **Step 1: types along branches.** Fix a nonprincipal ultrafilter $\mathcal{U}$ on ω , and for $\eta \in 2^\omega$ put $p_\eta = \lim_{n, \mathcal{U}} \text{ltp}_{\varphi_0}(a_{\eta|n}) \in \text{LS}(\varphi_0)$. Let $\sigma \sqsubseteq \eta$ be given. For every $n > |\sigma|$ we have $\sigma \smallfrown \langle \eta(|\sigma|) \rangle \sqsubseteq \eta|n$, so the value $\varphi(a_{\eta|n}, b_\sigma)$ is at least s if $\eta(|\sigma|) = 1$ and at most r if $\eta(|\sigma|) = 0$, and the same bound holds for $p_\eta(b_\sigma)$. If $\eta \neq \eta'$ and σ is the longest common initial segment of η and η' , then we have $\eta(|\sigma|) \neq \eta'(|\sigma|)$, and therefore

$$(5.2) \quad |p_\eta(b_\sigma) - p_{\eta'}(b_\sigma)| \geq s - r.$$

© **Step 2: approximate definitions.** Let $\hat{\varphi}_0$ be the separately continuous function of Proposition 1.2.8 for φ_0 ; as in the proof of Proposition 4.2.2, we have $\hat{\varphi}_0(p, \text{rtp}_{\varphi_0}(b)) = p(b)$ for all $p \in \text{LS}(\varphi_0)$ and $b \in B_0$. Let $\eta \in 2^\omega$ be given. By Theorem 1.2.11, in the form given by its proof, there are $a_1, \dots, a_k \in A_0$ and a real polynomial P with $|\hat{\varphi}_0(p_\eta, q) - P(q(a_1), \dots, q(a_k))| \leq \varepsilon$ for all $q \in \text{RS}(\varphi_0)$. Taking $q = \text{rtp}_{\varphi_0}(b)$ gives

$$|p_\eta(b) - P(\varphi(a_1, b), \dots, \varphi(a_k, b))| \leq \varepsilon \quad \text{for all } b \in B_0.$$

Each monomial is bounded by 1 on $[0, 1]^k$, so changing the coefficients of P by small enough amounts gives a polynomial P_η with rational coefficients with $|P - P_\eta| \leq \varepsilon$ on $[0, 1]^k$. Put $\bar{a}_\eta = (a_1, \dots, a_k)$. Then we have $|p_\eta(b) - P_\eta(\varphi(a_1, b), \dots, \varphi(a_k, b))| \leq 2\varepsilon$ for all $b \in B_0$.

⊙ **Step 3: counting.** The set A_0 is countable and there are countably many polynomials with rational coefficients, so there are only countably many pairs $(\bar{a}_\eta, P_\eta)$, while 2^ω is uncountable. Hence there are $\eta \neq \eta'$ with $(\bar{a}_\eta, P_\eta) = (\bar{a}_{\eta'}, P_{\eta'})$, and then we have $|p_\eta(b) - p_{\eta'}(b)| \leq 4\varepsilon < s - r$ for every $b \in B_0$. This contradicts (5.2). The last statement follows from Lemma 5.6.3. ☺

For formulas, Theorem 5.6.4 gives a second proof of the case $k = 2$ of Proposition 5.3.1. If $\varphi(x, y)$ is stable, then it is stable on every model (see Remark 1.2.2), so the corresponding $\{0, 1\}$ -valued function on every model is stable in the sense of Definition 1.2.5 (see Proposition 1.2.4). By Theorem 5.6.4 this function has no 2-tree property with any gap, and by the paragraph following Definition 5.6.1 the formula φ does not have the 2-tree property. The two proofs count different things: definitions over a model of size λ in Proposition 5.3.1, approximate definitions over a countable set here. The second count is the real-valued form of the counting in the proof of (1) $\Rightarrow$ (2) in Theorem 2.5.1: a binary tree of infinite height is a Cantor scheme of types, and stability leaves room only for countably many types up to a tolerance.

The converse of the last statement of Theorem 5.6.4 fails. Let M be a model of T_{RG} . By the extension axioms (see Example 4.1.4), every sequence of distinct vertices of M is independent for the edge relation, in the sense of Definition 4.2.1, so the edge relation is not NIP and hence, by Proposition 4.2.2, not stable. On the other hand, the random graph is simple (see Theorem 5.4.6), so by the paragraph following Definition 5.6.1 the edge relation does not have the 2-tree property on any model.

5.7 EXERCISES

The exercises follow the sections of the chapter. Throughout, we write N_E for the structure $(\omega \times \omega, E)$ in which $(m, n) E (m', n')$ holds if and only if $m = m'$, and T_E for its theory; thus the relation E is an equivalence relation with infinitely many classes, all of them infinite. The mark (*) indicates the exercises that ask for more than a direct application of the text.

1. (Section 5.1.) Assume that T_E admits elimination of quantifiers (see Exercise 10). Let A be a small set and let $n \geq 1$ be given. Show that the type over A of an n -tuple a is determined by the following data: which coordinates of a are equal and which are E -equivalent; which coordinates are equal to an element of A , and to which one; and which coordinates are E -equivalent to an element of A , and to which E -class of elements of A . Deduce that $|S_x(A)| \leq |A| + \aleph_0$, that T_E is λ -stable for every infinite λ , and that T_E is stable and simple (see Lemma 5.1.2 and Proposition 5.3.1). Show, in contrast, that the theory of $(\mathbb{Q}, <)$ has $2^{\aleph_0}$ complete 1-types over a countable model, and deduce from Lemma 5.1.2 that it is λ -stable for no infinite λ . (Hint: for the last assertion, find n -ladders for $x < y$; see Definition 1.1.4.)
2. (Section 5.2.) Let C be a small set and b a tuple. Show that the formula $x = b$ divides over C if and only if $\text{tp}(b/C)$ has infinitely many realizations, and that it then 2-divides over C . (Hint: if $(b_i)_{i < \omega}$ is indiscernible over C and $b_i = b_j$ for some $i < j$, then the sequence is constant. For the converse, apply Lemma 5.2.2 to distinct realizations of

$\text{tp}(b/C)$, and note that the formula $y_1 \neq y_2$ belongs to their Ehrenfeucht–Mostowski type.) Let T be the theory of a structure (N, E) in which E is an equivalence relation with infinitely many classes, each with exactly two elements. Let c and b be distinct E -equivalent elements. Show that $x = b$ does not divide over $\{c\}$, although $b \notin \{c\}$.

3. (Section 5.2.) Let C be a small set, let $\varphi(x, y)$ and $\psi(x, z)$ be L -formulas, and let b and c be tuples such that $\mathbb{M} \models \forall x (\psi(x, c) \rightarrow \varphi(x, b))$. Show that if $\varphi(x, b)$ k -divides over C , then the formula $\psi(x, c)$ k -divides over C . (Hint: let $(b_i)_{i < \omega}$ witness k -dividing, choose c_i with $b_i c_i \equiv_C bc$, and apply Lemma 5.2.2 over C to the sequence $(b_i c_i)_{i < \omega}$; the formulas $\forall x (\psi(x, z_1) \rightarrow \varphi(x, y_1))$ and $\neg \exists x \bigwedge_{l=1}^k \varphi(x, y_l)$ belong to its Ehrenfeucht–Mostowski type.) Deduce that if $\varphi(x, b)$ divides over C , then so does $\varphi(x, b) \wedge \chi(x, d)$ for every L -formula $\chi(x, w)$ and every tuple d .
4. (Section 5.2.) Work in T_E , and assume that T_E admits elimination of quantifiers (see Exercise 10). Let C be a small set, and let b be an element.
 - (a) Show that the formula xEb divides over C if and only if b is E -equivalent to no element of C , and that it then 2-divides over C . In particular, the formula xEb divides over $\emptyset$. (Hint: for one direction, choose $b_1, b_2, \dots$ in distinct E -classes that meet neither C nor the class of b , and show that $(b, b_1, b_2, \dots)$ is indiscernible over C .)
 - (b) Show that the formula $x = b$ divides over C if and only if $b \notin C$ (see Exercise 2).
 - (c) Let a be an element and B a small set. Show that $\text{tp}(a/C \cup B)$ divides over C if and only if $a \in B \setminus C$, or a is E -equivalent to some element of B and to no element of C . (Hint: in the remaining cases, given a formula of $\text{tp}(a/C \cup B)$ with parameters from a finite tuple d and a sequence $(d_i)_{i < \omega}$ indiscernible over C with $d_0 = d$, find a single element a'' with $a'' d_i \equiv_C ad$ for every i . Distinguish the cases $a \in C$, $a \notin C \cup B$ with a equivalent to an element of C , and $a \notin C \cup B$ with a equivalent to no element of $C \cup B$.)

Compare (c) with Theorem 5.4.5: in T_E a type may divide over C although $a \cap B \subseteq C$.

5. (Section 5.2.) Work in the theory of $(\mathbb{Q}, <)$, let C be a small set, and let $b_1 < b_2$ be elements. Show that the formula $b_1 < x < b_2$ divides over C if and only if no element c of C satisfies $b_1 \leq c \leq b_2$, and that it then 2-divides over C . Show also that the formulas $x < b_1$ and $b_1 < x$ divide over no small set. (Hint: for the first assertion, follow Example 5.2.6, choosing the pairs b^i in the cut of C that contains b_1 and b_2 . If $c = b_1$ for some $c \in C$, then every sequence indiscernible over C that begins with (b_1, b_2) has first coordinate constantly equal to c .)
6. (Section 5.3.) Work in the theory of $(\mathbb{Q}, <)$, and let $\theta(x; y_1, y_2)$ be the formula $y_1 < x < y_2$. Find pairs $a_n = (a_n^1, a_n^2)$, for $n < \omega$, such that $\{\theta(x, a_n) : n < \omega\}$ is consistent and $\theta(x, a_n)$ 2-divides over the set A_n of coordinates of the pairs a_m with $m < n$ (see Exercise 5). Deduce from Lemma 5.3.5 that θ has the 2-tree property, and compare the discussion following Example 5.2.6. Deduce also that there are an element a and small sets $D_0 \subseteq D_1 \subseteq \dots$ such that $\text{tp}(a/D_{n+1})$ divides over D_n for every n ; compare Theorem 5.4.6. Finally, show that the formula $x < y$ has the k -tree property for no $k \geq 2$, although it has the order property. Thus, in a theory that is not simple, a formula with the order property need not have the tree property.


7. (Section 5.3.) Let T be the theory of the structure $(\omega^{<\omega}, \sqsubseteq)$, where $\sqsubseteq$ is the relation of being an initial segment. Show that the formula $y \sqsubseteq x$ has the 2-tree property, with $b_\sigma = \sigma$ for every nonempty $\sigma \in \omega^{<\omega}$. (Hint: the structure $(\omega^{<\omega}, \sqsubseteq)$ embeds elementarily in $\mathbb{M}$; two siblings have no common extension, and this is expressed by a sentence with parameters.) Show that the formula $x \sqsubseteq y$ has the k -tree property for no $k \geq 2$. (Hint: the theory T says that some element is an initial segment of every element.) Deduce that T is neither simple nor stable.
8. (Section 5.4.) Work in T_{RG} . Let C be a small set, let b be a finite tuple, and let $\varphi(x, b)$ be an L -formula in which x is a single variable. Show that $\varphi(x, b)$ divides over C if and only if every solution of $\varphi(x, b)$ in $\mathbb{M}$ is a coordinate of b that does not belong to C . (Hint: if some solution a is not such a coordinate, then we have $a \cap b \subseteq C$; use Theorem 5.4.5. For the converse, use Lemma 5.4.3 and (5.1) to build a sequence indiscernible over C that begins with b and whose terms are pairwise disjoint outside C , as in the first part of the proof of Theorem 5.4.5.) Deduce that the formulas $E(x, b_1)$ and $\neg E(x, b_1)$, for a single vertex b_1 , divide over no small set.
9. (Section 5.4.) Work in T_{RG} . Let a be a finite tuple and B a small set. Show that $\text{tp}(a/B)$ does not divide over the finite set $a \cap B$. Thus, in Theorem 5.3.6 for T_{RG} , the set A can be chosen finite (see Remark 5.5.11). Show also that the bound of Theorem 5.4.6 is attained: for every $n \geq 1$, give a tuple a of length n and sets $D_0 \subseteq \cdots \subseteq D_n$ such that $\text{tp}(a/D_{j+1})$ divides over D_j for every $j < n$.
10. (*) (Section 5.5.) Let $\mathcal{K}_E$ be the class of finite structures (X, E) in which E is an equivalence relation.
- (a) Show that $\mathcal{K}_E$ is a Fraïssé class with disjoint 2-amalgamation. (Hint: on $B_1 \cup B_2$, make an element of $B_1 \setminus A$ equivalent to an element of $B_2 \setminus A$ exactly when both are equivalent to a common element of A .)
 - (b) Show that N_E is ultrahomogeneous and has age $\mathcal{K}_E$. Deduce from Fact 5.5.2 that $T_E = T_{\mathcal{K}_E}$ and that T_E admits elimination of quantifiers.
 - (c) Show that $\mathcal{K}_E$ does not have disjoint 3-amalgamation. (Hint: take $A = \emptyset$ and one-element sets B_1, B_2, B_3 .)
 - (d) Conclude, by Exercise 1, that the converse of Theorem 5.5.9 fails, and, by Exercise 4(c), that conclusion (1) of Theorem 5.5.9 fails for $\mathcal{K}_E$.
 - (e) Show that the class of finite structures (X, E_1, E_2) in which E_1 and E_2 are equivalence relations is a Fraïssé class with disjoint 2-amalgamation and without disjoint 3-amalgamation, and that the theory of its limit is stable. (Hint: count types as in Exercise 1, using Fact 5.5.2(b).)
11. (Section 5.5.) Let L consist of one unary relation symbol P , and let $\mathcal{K}$ be the class of finite L -structures in which at most one element satisfies P . Show that $\mathcal{K}$ is a Fraïssé class whose limit is a countably infinite set with exactly one element in P . Show that $\mathcal{K}$ has disjoint 3-amalgamation and does not have disjoint 2-amalgamation. Let d be the element of P in $\mathbb{M}$. Show that no sequence of distinct elements beginning with d is indiscernible over $\emptyset$, and that $\text{tp}(d/\{d\})$ does not divide over $\emptyset$ (every sequence indiscernible over $\emptyset$ that begins with d is constant). Thus the conclusions of Lemma 5.5.8 and of Theorem 5.5.9(1) fail without disjoint 2-amalgamation.

- 12. (*) (Section 5.5.)** Let L consist of one ternary relation symbol R , and let $\mathcal{K}$ be the class of finite 3-uniform hypergraphs (see Example 5.5.3(5)), that is, L -structures in which R holds only of triples of distinct elements and is invariant under permutations of its arguments, with no *tetrahedron*, that is, with no four distinct elements of which every three form a hyperedge.
- (a) Show that $\mathcal{K}$ is a Fraïssé class with disjoint 2-amalgamation. (Hint: take the union, with no new hyperedges, as in Example 5.5.5.)
 - (b) Show that $\mathcal{K}$ has disjoint 3-amalgamation. (Hint: declare that no hyperedge meets all three of the sets $B_i \setminus A$. A tetrahedron that meets all three sets contains a triple that meets all three.)
 - (c) Deduce from Theorem 5.5.9 that $T_{\mathcal{K}}$ is simple. Show that the formula $R(x; y_1, y_2)$ has the independence property (see Definition 4.1.1), and deduce that $T_{\mathcal{K}}$ is not stable.
- 13. (*) (Section 5.5.)** Let $\mathcal{K}$ be the class of finite graphs with no four pairwise adjacent vertices.
- (a) Show that $\mathcal{K}$ is a Fraïssé class with disjoint 2-amalgamation, and that it does not have disjoint 3-amalgamation. (Hint: take $A = \emptyset$, $B_1 = \{x\}$, $B_2 = \{y\}$ and $B_3 = \{z, w\}$ with z adjacent to w .)
 - (b) Let $\psi(x; y_1, y_2, y_3)$ be the formula $E(x, y_1) \wedge E(x, y_2) \wedge E(x, y_3)$. Show that ψ has the 2-tree property in $T_{\mathcal{K}}$, and conclude that $T_{\mathcal{K}}$ is not simple. (Hint: adapt Example 5.5.10. Take vertices $u_\sigma, v_\sigma, w_\sigma$ for nonempty $\sigma \in \omega^{<\omega}$, join v_σ to w_σ , and join $u_{\sigma \smallfrown \langle i \rangle}$ to $v_{\sigma \smallfrown \langle j \rangle}$ and to $w_{\sigma \smallfrown \langle j \rangle}$ whenever $i < j$. Show that this graph has no four pairwise adjacent vertices, and put $b_\sigma = (u_\sigma, v_\sigma, w_\sigma)$.)
- 14. (Section 5.6.)** Let $A = \omega^{<\omega}$ and $B = \omega^{<\omega} \setminus \{\emptyset\}$, and define $\varphi: A \times B \rightarrow [0, 1]$ by $\varphi(a, b) = 2/3$ if $b \sqsubseteq a$, by $\varphi(a, b) = 0$ if a is a proper initial segment of b , and by $\varphi(a, b) = 1/3$ if a and b are incomparable.
- (a) Show that, for reals $r < s$, the function φ has the 2-tree property with gap (r, s) (see Definition 5.6.1) if and only if $1/3 \leq r < s \leq 2/3$. (Hint: for $r < 1/3$, use condition (b) with $\tau = \emptyset$ and an element a of length 1.)
 - (b) Fix a gap (r, s) as in (a), and apply the construction of Lemma 5.6.3 to the tree $b_\sigma = \sigma$, with $a'_\tau = f(\tau)$. Compute $b'_\emptyset, b'_{\langle 0 \rangle}, b'_{\langle 1 \rangle}, a'_{\langle 0 \rangle}$ and $a'_{\langle 1 \rangle}$, and check the two inequalities of Definition 5.6.2 for $\sigma = \emptyset$.
 - (c) Show directly from Definition 1.2.5 that φ is not stable, as Theorem 5.6.4 requires. (Hint: let a_m and b_n be the sequences of zeros of lengths m and $n + 1$, and let $\mathcal{U}$ and $\mathcal{V}$ be nonprincipal.)

5.8 HISTORICAL REMARKS

⌘ Section 5.1. Working inside a single large saturated and strongly homogeneous model goes back to Shelah [She90], who wrote $\mathfrak{C}$ for this model. The name *monster model* became standard later. The notion of ω -stability comes from Morley's categoricity theorem [Mor65]. The study of λ -stability for arbitrary λ , including the fact that stability is equivalent to λ -stability for all λ with $\lambda^{|T|} = \lambda$, is due to Shelah [She71, She90].

- ✂ **Section 5.2.** Dividing, forking and the tree property were introduced by Shelah in *Classification Theory* ([She90, Chapter III]; first edition 1978), and the simple theories, those in which no formula has the tree property, were named and studied in his paper [She80]. Indiscernible sequences go back to Ehrenfeucht and Mostowski [EM56], and were made into a standard tool of classification theory by Morley and Shelah [She90]. For the next fifteen years the simple unstable theories received little attention. Two examples brought them back: the pseudofinite fields, whose theories Hrushovski showed to be simple and unstable (in work circulated as a manuscript from 1991 and published in [Hru02]), and the difference fields of Chatzidakis and Hrushovski [CH99].
- ✂ **Section 5.3.** That stable theories are simple, and that the strict order property implies the tree property, go back to Shelah [She90, Chapter III], [She80]; the dichotomy of Fact 5.3.3 is in [She90, Chapter II]. Kim [Kim98] proved that forking and dividing agree in simple theories and that forking independence is symmetric; Kim and Pillay [KP97] proved the independence theorem, abstracting it from Hrushovski’s work on pseudofinite fields, and characterized simple theories by their independence relation. The title of this chapter is that of their survey [KP98]. The monographs of Wagner [Wag00] and Casanovas [Cas11] develop the theory in full.
- ✂ **Section 5.4.** The graph of Example 5.4.1 is implicit in Ackermann’s coding of the hereditarily finite sets by natural numbers [Ack37]: the relation “the i -th binary digit of j is 1” encodes membership, and symmetrizing it gives the Rado graph. Erdős and Rényi [ER63] proved that a random countable graph, with each edge chosen independently with probability $\frac{1}{2}$, is isomorphic with probability 1 to a fixed graph, and Rado [Rad64] gave the explicit construction of Example 5.4.1 and proved that the graph is universal for countable graphs. The theory T_{RG} is also the almost sure theory of finite graphs: by the zero–one law of Glebskiĭ, Kogan, Liogon’kiĭ and Talanov [GKLT69] and of Fagin [Fag76], a first-order sentence holds in almost all finite graphs if and only if it holds in the Rado graph. Cameron’s survey [Cam97] collects its many descriptions.
- ✂ **Section 5.5.** Fraïssé [Fra54] introduced ages and proved Fact 5.5.2(a), generalizing Cantor’s characterization of $(\mathbb{Q}, <)$ as the unique countable dense linear order without endpoints. Henson [Hen71] constructed the homogeneous K_n -free graphs, and Lachlan and Woodrow [LW80] proved that the countably infinite ultrahomogeneous graphs are, up to complementation, the disjoint unions of complete graphs of equal size, the Rado graph, and the Henson graphs. Kruckman [Kru19] made disjoint n -amalgamation a systematic tool for the theories of Fraïssé limits.
- ✂ **Section 5.6.** Simplicity for structures with real-valued relations was developed by Ben Yaacov [BY03] in the framework of compact abstract theories, which includes continuous logic. The argument of Theorem 5.6.4 is elementary by comparison: it combines the polynomial approximation of Theorem 1.2.11 with the counting of branches in the proof of Theorem 2.5.1.



- 1937 • **Ackermann.** Ackermann [Ack37] codes the hereditarily finite sets by natural numbers through their binary digits; symmetrized, the coding gives the Rado graph.
- 1954 • **Fraïssé.** Fraïssé [Fra54] introduces ages and proves that every Fraïssé class has a unique countable ultrahomogeneous limit.
- 1963 • **Erdős–Rényi.** Erdős and Rényi [ER63] show that the random countable graph is unique up to isomorphism with probability 1.
- 1964 • **Rado.** Rado [Rad64] gives the explicit construction of Example 5.4.1 and proves its universality.
- 1969–1976 • **zero–one law.** Glebskiĭ, Kogan, Liogon’kiĭ and Talanov [GKLT69], and independently Fagin [Fag76], prove the zero–one law for first-order sentences about finite graphs; the almost sure theory is T_{RG} .
- 1971 • **Henson.** Henson [Hen71] constructs the homogeneous K_n -free graphs.
- 1978 • **Shelah.** Shelah’s *Classification Theory* [She90] introduces forking, dividing and the tree property.
- 1980 • **Shelah; Lachlan–Woodrow.** Shelah [She80] names and studies the simple theories, those without the tree property. Lachlan and Woodrow [LW80] classify the countable ultrahomogeneous graphs.
- 1991 • **Hrushovski.** Hrushovski’s manuscript on pseudofinite fields, published in [Hru02], shows that their theories are simple and proves an independence theorem for them.
- 1996 • **Shelah.** Shelah [She96] introduces the hierarchy of properties SOP_n ; the Henson graph has SOP_3 .
- 1997–1998 • **Kim; Kim–Pillay.** Kim [Kim98] proves that forking equals dividing in simple theories and that forking is symmetric. Kim and Pillay [KP97] prove the independence theorem; their survey [KP98] is titled *From stability to simplicity*.
- 1999 • **Chatzidakis–Hrushovski.** Chatzidakis and Hrushovski [CH99] develop the model theory of difference fields, a central example of a simple unstable theory.
- 2000–2011 • **Wagner; Casanovas.** The monographs of Wagner [Wag00] and Casanovas [Cas11].
- 2003 • **Ben Yaacov.** Ben Yaacov [BY03] extends simplicity to compact abstract theories, a framework that includes continuous logic.
- 2017–2019 • **Conant; Kruckman.** Conant [Con17] develops free amalgamation theories; Kruckman [Kru19] studies disjoint n -amalgamation and pseudofinite countably categorical theories.

Chapter 6

Stability and notions of independence

Ladders, trees and definitions all describe stability from inside a single formula. Independence describes stability globally: in this chapter we show that the stable theories are exactly those carrying a notion of independence, written $A \downarrow_M B$ and read “ A is independent from B over M ”, that obeys a short list of natural axioms. The most important of these are finite character, symmetry, and the existence and uniqueness of independent extensions. The idea is that, given a type p over a model M and a set B , an independent extension of p to $M \cup B$ adds no information beyond p . Stability is precisely the condition that such an extension always exists and is unique (see Theorem 6.2.2).

Both directions rest on Chapter 2. To obtain independence from stability we use definability of types: the independent extension is the *heir*, the extension of p given by the same definition scheme as p (see Section 6.1). The converse is a counting argument of the kind used there (see Section 6.2): uniqueness and local character together bound the number of types.

We keep the conventions of Section 5.1, and we measure independence against the axioms of Definition 5.4.7. In Chapter 5 the random graph carried a relation with every axiom but uniqueness; here uniqueness is the protagonist.

Roadmap

In the preceding chapters we studied dependence one formula at a time. In this chapter we take a global view and explain what it means for a set A to be independent from a set B over a model M , when all formulas are considered at once. The explanation is not an explicit definition but an implicit one, a list of axioms (see Definition 5.4.7), and the theorem of the chapter is that a theory is stable exactly when some relation satisfies all of them, uniqueness included (see Theorem 6.2.2). In Section 6.1 we construct the relation. The independent extension of a type is its heir, the extension given by the same definitions, and the symmetry of the relation rests on Harrington’s lemma, which in Subsection 6.1.1 we show to be an exchange of iterated limits, the analyst’s form of stability met in Chapter 1. In Section 6.2 we state the axioms and prove the theorem in both directions, the converse by a count of types in the manner of Chapter 2. In Section 6.3 we record how little of the list is really needed, and how the relation extends, as nonforking, to arbitrary base sets. If the chapter had to be summarized in a sentence, the sentence would be this: in a stable theory there is exactly one way of learning nothing new.

6.1 DEFINITIONS AND HEIRS

In a stable theory every type over a model has a definition scheme, one definition for each formula (see Lemma 6.1.1). A definition can be evaluated at parameters outside the model, so the scheme extends the type to larger sets of parameters, and the result is the heir (see Definition 6.1.2). The section closes with the fact that will make independence symmetric: definitions over a model commute (see Lemma 6.1.4). In Subsection 6.1.1 we show that this commutation is the exchange of iterated ultrafilter limits that defines stability in Section 1.2.

Lemma 6.1.1. *Assume that T is stable, and let $p \in S_x(M)$ be given, where x may be infinite. For every $L(M)$ -formula $\varphi(x', y)$, with x' a finite subtuple of x and y finite, there is an $L(M)$ -formula $d_p\varphi(y)$ such that for all $b \in M$, we have*

$$\varphi(x', b) \in p \iff M \models d_p\varphi(b).$$

Any two formulas with this property are equivalent in $\mathbb{M}$.

Proof. Write $\varphi(x', y) = \varphi_0(x', y, m)$ with $\varphi_0 \in L$ and $m \in M$. The set of formulas $\varphi_0(x', c)^u$ ($c \in M$) that belong to p is a complete φ_0 -type over M in the sense of Definition 1.1.1. Since φ_0 does not have the order property, Theorem 1.1.6 shows that this φ_0 -type has a definition $\psi(z)$ over M , which is a positive Boolean combination of instances of φ_0 . Here the tuple z corresponds to the parameter variables y followed by those replaced by m . Then the formula $d_p\varphi(y) := \psi(y, m)$ works. If $\psi_1(y)$ and $\psi_2(y)$ both work, then we have $M \models \forall y (\psi_1 \leftrightarrow \psi_2)$. This is an $L(M)$ -sentence, so this sentence also holds in $\mathbb{M}$. ☺

The family $(d_p\varphi)_\varphi$ given by Lemma 6.1.1, determined up to equivalence in $\mathbb{M}$, is called the *definition scheme* of p .

Definition 6.1.2. Assume that T is stable. Let $p \in S_x(M)$ be given, and let B be small. Define

$$p|B := \{\varphi(x', b) : \varphi(x', y) \in L(M), b \in B, \mathbb{M} \models d_p\varphi(b)\}.$$

The set $p|B$ is called the *heir* of p over MB .

Every $L(MB)$ -formula has the form $\varphi(x', b)$ with $\varphi \in L(M)$ and $b \in B$, so the heir $p|B$ is a set of $L(MB)$ -formulas. Figure 6.1 shows the construction. The next lemma shows that the heir is in fact a complete type extending p , and that heirs are compatible with enlarging the parameter set and the base model.

Lemma 6.1.3. *Assume that T is stable, let $p \in S_x(M)$ be given, and let B and D be small sets. Then the following hold:*

- (a) *The set $p|B$ belongs to $S_x(MB)$ and extends p .*
- (b) *If $B \subseteq D$, then the type $p|B$ is the restriction of $p|D$ to MB .*
- (c) *If N is a model with $M \subseteq N$, then the type $p|N$ has the same definition scheme as p , and we have $(p|N)|B = p|NB$.*

Proof. (a) *Completeness.* Let $\varphi \in L(M)$ be given. Since p is complete, for every $c \in M$ exactly one of $\varphi(x', c)$, $\neg\varphi(x', c)$ lies in p . Hence we have $M \models \forall y (d_p(\neg\varphi)(y) \leftrightarrow \neg d_p\varphi(y))$, and the same holds in $\mathbb{M}$. Thus for $b \in B$ exactly one of $\varphi(x', b)$, $\neg\varphi(x', b)$ belongs to $p|B$.

Consistency. Let $\varphi_i(x_i, b_i) \in p|B$ be given for $i < n$, and put $\chi(\bar{x}; y_0, \dots, y_{n-1}) = \bigwedge_{i < n} \varphi_i(x_i, y_i)$, where $\bar{x}$ lists the variables in the x_i . For $\bar{c} \in M$, the formula $\chi(\bar{x}, \bar{c})$ is in




FIGURE 6.1. Heirs (see Definition 6.1.2). The definition $d_p\varphi$ is a formula with parameters in M , so it defines a set of parameters in all of $\mathbb{M}$ (pink). Inside M (shaded) it records which instances $\varphi(x, c)$ belong to p ; outside M the same set decides which instances $\varphi(x, b)$ belong to the heir $p|B$. The heir is the type p evaluated, by its own definitions, at parameters that M has never seen.

p if and only if each $\varphi_i(x_i, c_i)$ is, and in that case the formula $\chi(\bar{x}, \bar{c})$ is satisfiable in M . Therefore we have

$$M \models \forall \bar{y} \left(\bigwedge_{i < n} d_p \varphi_i(y_i) \rightarrow d_p \chi(\bar{y}) \right) \quad \text{and} \quad M \models \forall \bar{y} (d_p \chi(\bar{y}) \rightarrow \exists \bar{x} \chi(\bar{x}, \bar{y})).$$

Both sentences transfer to $\mathbb{M}$. Applied to $\bar{b} = (b_0, \dots, b_{n-1})$, they give $\mathbb{M} \models \exists \bar{x} \bigwedge_i \varphi_i(x_i, b_i)$. By compactness, the set $p|B$ is consistent.

Extension. For $c \in M$ we have $\varphi(x', c) \in p \iff M \models d_p \varphi(c)$, which holds if and only if $\varphi(x', c) \in p|B$.

(b) This is immediate from the definition.

(c) By (a), the type $p|N$ extends p . For $n \in N$ we have $\varphi(x', n) \in p|N$ if and only if $\mathbb{M} \models d_p \varphi(n)$, so the formula $d_p \varphi$ is a definition of $p|N$ for $\varphi \in L(M)$. For an $L(N)$ -formula $\varphi(x', y) = \varphi_0(x', y, n)$ with $\varphi_0 \in L(M)$, the formula $d_p \varphi_0(y, n)$ defines $p|N$. By the uniqueness clause of Lemma 6.1.1, the definition scheme of $p|N$ is thus given, up to equivalence, by that of p . Now for $b \in B$ and $n \in N$, we have

$$\varphi_0(x', b, n) \in (p|N)|B \iff \mathbb{M} \models d_p \varphi_0(b, n) \iff \varphi_0(x', b, n) \in p|NB. \quad \odot$$

The fulcrum of the chapter is the following lemma, which says that definitions of types over a model commute; this lemma is the only place in this chapter where the absence of the order property is used directly. To state the lemma, let $p \in S_x(M)$ and $q \in S_y(M)$ be given, with x and y finite, and let $\varphi(x, y) \in L(M)$ be given. As before, we write $d_p \varphi(y)$ for the definition of p for $\varphi(x, y)$, and we write $d_q \varphi(x)$ for the definition of q for φ regarded as a formula in the variables y , with parameter variables x .

Lemma 6.1.4 (Harrington's lemma). *Assume that T is stable, and let $p \in S_x(M)$, $q \in S_y(M)$ and $\varphi(x, y) \in L(M)$ be given, where x and y are finite. Then we have*

$$d_p \varphi(y) \in q \iff d_q \varphi(x) \in p.$$

Proof. The idea is to realize p and q alternately, so that each pair (a_i, b_j) consults one definition below the diagonal and the other on and above the diagonal; if the two definitions disagreed, the result would be a ladder. Choose recursively $a_i, b_i \in \mathbb{M}$ ($i < \omega$) such that a_i

$M^{|x|} \times M^{|y|}$. For $b' \in M^{|y|}$ we have $\mathbb{M} \models d_p \varphi(b')$ if and only if $\varphi(x, b') \in p$, that is, if and only if $\varphi(M, b') \in \mathcal{U}$, that is, if and only if $\lim_{a', \mathcal{U}} \varphi(a', b') = 1$. Therefore we have

$$d_p \varphi(y) \in q \iff \{b' \in M^{|y|} : \mathbb{M} \models d_p \varphi(b')\} \in \mathcal{V} \iff \lim_{b', \mathcal{V}} \lim_{a', \mathcal{U}} \varphi(a', b') = 1,$$

and in the same way we have

$$d_q \varphi(x) \in p \iff \lim_{a', \mathcal{U}} \lim_{b', \mathcal{V}} \varphi(a', b') = 1.$$

Thus the two sides of Harrington's lemma are the statements that the two iterated limits of φ along $\mathcal{U}$ and $\mathcal{V}$ equal 1, and the lemma says exactly that these iterated limits agree. Note that no heirs, saturation or realizations are involved: the computation takes place entirely in M .

Let us notice that, when T is stable, the limits do agree: write $\varphi(x, y) = \varphi_0(x, y, m)$ with $\varphi_0 \in L$ and $m \in M$. The formula $\varphi_0(x, y, z)$ does not have the order property, so by Remark 1.2.2 it is stable on M , and therefore so is its restriction φ to the parameters of the form (b', m) , since a ladder for φ or $\neg \varphi$ would be one for φ_0 or $\neg \varphi_0$. By Proposition 1.2.4, the function φ is then stable in the sense of Definition 1.2.5, and by Corollary 1.2.6, applied to the families $(a')_{a' \in M^{|x|}}$ and $(b')_{b' \in M^{|y|}}$, the two iterated limits are equal. This proves Lemma 6.1.4 again.

The two proofs are the same argument at different levels of abstraction. The alternating realizations (a_i, b_j) in the first proof consult one definition below the diagonal and the other on and above it, exactly as the sequences of Proposition 1.2.3 display one iterated limit above the diagonal and the other below it; and in both cases a disagreement is converted into a ladder.

The computation also explains what the two orders of realization mean. For a small set B , the *coheir* of $\mathcal{U}$ over MB is the type

$$q_{\mathcal{U}}|B = \{\varphi(x, b) : \varphi(x, y) \in L(M), b \text{ a finite tuple from } B, \{a \in M^{|x|} : \mathbb{M} \models \varphi(a, b)\} \in \mathcal{U}\}.$$

We first collect its basic properties, which are used again in Chapter 7, and then return to the two orders of realization.

Lemma 6.1.5 (Coheirs). *Let M be a model, let $\mathcal{U}$ be an ultrafilter on $M^{|x|}$, and let B and B' be small sets.*

- (a) *The set $q_{\mathcal{U}}|B$ is a complete type over MB , and it is finitely satisfiable in M .*
- (b) *Every type over MB that is finitely satisfiable in M equals $q_{\mathcal{U}'}|B$ for some ultrafilter $\mathcal{U}'$ on $M^{|x|}$.*
- (c) *If $B \subseteq B'$, then we have $q_{\mathcal{U}}|B \subseteq q_{\mathcal{U}}|B'$.*
- (d) *The coheir is invariant over M : if $\sigma \in \text{Aut}(\mathbb{M}/M)$, then we have $\sigma(q_{\mathcal{U}}|B) = q_{\mathcal{U}}|\sigma(B)$.*

Proof. For an $L(MB)$ -formula $\varphi(x, b)$, write $\varphi(M, b) = \{a \in M^{|x|} : \mathbb{M} \models \varphi(a, b)\}$. These sets satisfy $(\neg \varphi)(M, b) = M^{|x|} \setminus \varphi(M, b)$ and $(\varphi \wedge \psi)(M, b) = \varphi(M, b) \cap \psi(M, b)$.

(a) Since $\mathcal{U}$ is an ultrafilter, exactly one of $\varphi(M, b)$ and its complement belongs to $\mathcal{U}$, so exactly one of $\varphi(x, b)$ and $\neg \varphi(x, b)$ belongs to $q_{\mathcal{U}}|B$. If $\varphi_1(x, b_1), \dots, \varphi_n(x, b_n)$ belong to $q_{\mathcal{U}}|B$, the intersection of the sets $\varphi_i(M, b_i)$ belongs to $\mathcal{U}$ and so is non-empty, and any of its elements is a tuple from M satisfying all n formulas. Hence the set $q_{\mathcal{U}}|B$ is finitely satisfiable in M , in particular consistent, and complete.

(b) Let p be a type over MB finitely satisfiable in M . The sets $\varphi(M, b)$, for $\varphi(x, b) \in p$, have the finite intersection property, so they are contained in an ultrafilter $\mathcal{U}'$ on $M^{|x|}$. Then we have $p \subseteq q_{\mathcal{U}'}|B$, and both are complete, so they are equal.

(c) This is immediate from the definition.

(d) Let a be a tuple from M . Since σ fixes a and the parameters of φ from M , we have $\mathbb{M} \models \varphi(a, b)$ if and only if $\mathbb{M} \models \varphi(a, \sigma(b))$. Hence we have $\varphi(M, b) = \varphi(M, \sigma(b))$, and it follows that $\varphi(x, b) \in q_{\mathcal{U}}|B$ if and only if $\varphi(x, \sigma(b)) \in q_{\mathcal{U}}|B$. $\odot$

We return to the two orders of realization. If b realizes $q_{\mathcal{V}}|\emptyset$ and a realizes $q_{\mathcal{U}}|\{b\}$, then unwinding the definitions gives

$$\mathbb{M} \models \varphi(a, b) \iff \lim_{a', \mathcal{U}} \lim_{b', \mathcal{V}} \varphi(a', b') = 1,$$

and realizing the two coheirs in the other order replaces this iterated limit by the one in the other order. Thus the two orders of realizing coheirs give the same type exactly when the iterated limits commute. In a stable theory, moreover, the coheir $q_{\mathcal{U}}|B$ is the heir of $q_{\mathcal{U}}|\emptyset$, because both are nonforking extensions and nonforking extensions of types over models are unique (see Remark 6.3.2 and [Pil96, TZ12]). Read in this way, the symmetry of independence proved in Section 6.2 is the commutation of ultrafilter limits: the logical counterpart of exchanging the order of two limits in analysis.

6.2 INDEPENDENCE RELATIONS

The axioms of Definition 5.4.7 record how the heir relation behaves. When T is stable the heir relation satisfies all of them, and conversely a relation satisfying just two of them, uniqueness and local character, forces T to be stable (see Theorem 6.2.2). In Chapter 5 we met relations satisfying all the axioms except uniqueness (see Theorems 5.4.9 and 5.5.9); the theorem below explains why uniqueness is exactly what they lack.

In terms of types, the axioms (EX) and (UN) say that a type $p \in S_x(M)$ has exactly one *independent extension* to MB , namely the type over MB of any $A \models p$ with $A \perp_M B$. All the other axioms are statements about which extensions of p count as independent. In a stable theory the natural candidate for this extension is the heir of Section 6.1, and we take it as the definition.

Definition 6.2.1. Assume that T is stable. For small A, B and a model M we define

$$A \perp_M^h B \iff \text{tp}(A/MB) = \text{tp}(A/M)|B.$$

In words, the type $\text{tp}(A/MB)$ is the *heir* of $\text{tp}(A/M)$, the extension defined by the same definition scheme. We call $\perp_M^h$ *heir independence*.

For finite tuples the definition becomes a test, formula by formula, which is the form used in the proofs below. If a, b are finite tuples and $p = \text{tp}(a/M)$, then every $L(Mb)$ -formula in x has the form $\varphi(x, b)$ with $\varphi(x, y) \in L(M)$. Hence we have

$$(6.1) \quad a \perp_M^h b \iff \text{for all } \varphi(x, y) \in L(M) : \quad \mathbb{M} \models \varphi(a, b) \leftrightarrow d_p \varphi(b).$$

Theorem 6.2.2. *For a complete theory T , the following are equivalent:*

- (1) *The theory T is stable.*
- (2) *There is an independence relation satisfying all the axioms of Definition 5.4.7.*
- (3) *There is a relation $A \perp_M B$ satisfying (UN) and (LOC).*

Moreover, if T is stable, the relation of Definition 6.2.1 satisfies all the axioms. This relation also satisfies

- (AR) Anti-reflexivity. *If $a \perp_M^h a$ for a finite tuple a , then we have $a \subseteq M$.*

6.2.1 Proof of (1) $\Rightarrow$ (2)

Assume that T is stable, and let $\perp^h$ be the relation of Definition 6.2.1. Throughout, we write $p = \text{tp}(A/M)$. Most of the axioms follow by unwinding the definitions with the help of Lemma 6.1.3; symmetry is the one place where Harrington's lemma is needed.

(INV). Let $\sigma \in \text{Aut}(\mathbb{M})$ be given. Applying σ to the parameters of the definition scheme of p gives a definition scheme for $\sigma(p) = \text{tp}(\sigma A/\sigma M)$. Hence we have $\sigma(p)|\sigma(B) = \sigma(p|B)$. If $\text{tp}(A/MB) = p|B$, then we have $\text{tp}(\sigma A/\sigma M \sigma B) = \sigma(p|B) = \sigma(p)|\sigma B$.

(MON). The definition scheme of $\text{tp}(A'/M)$ is the part of that of p that involves only variables for A' . Hence the heir $\text{tp}(A'/M)|B'$ is the restriction of $p|B$ to these variables and to MB' (see Lemma 6.1.3(b)), which is $\text{tp}(A'/MB')$.

(FC). The condition $\text{tp}(A/MB) = p|B$ says that for every $L(M)$ -formula $\varphi(x', y)$ and every $b \in B$ we have $\mathbb{M} \models \varphi(a', b) \leftrightarrow d_p \varphi(b)$, where a' is the finite subtuple of A corresponding to x' . Each instance involves only finite $a' \subseteq A$ and $b \subseteq B$, and is exactly an instance of $a' \perp_M^h b$.

(BMON). Let N be a model with $M \subseteq N \subseteq B$, and suppose that $\text{tp}(A/MB) = p|B$. Then we have $\text{tp}(A/N) = p|N$ by Lemma 6.1.3(b). By Lemma 6.1.3(c), we have $\text{tp}(A/N)|B = (p|N)|B = p|NB = p|B = \text{tp}(A/NB)$, which means that $A \perp_N^h B$.

(TR). Let N be a model with $M \subseteq N$, and suppose that $\text{tp}(A/N) = p|N$ and $\text{tp}(A/NB) = \text{tp}(A/N)|B$. By Lemma 6.1.3(c), we have $\text{tp}(A/NB) = (p|N)|B = p|NB$, which means that $A \perp_M^h NB$.

(SYM). By (FC) we need only treat finite tuples a, b . Put $p = \text{tp}(a/M)$ and $q = \text{tp}(b/M)$, and assume that $a \perp_M^h b$. For $\varphi(x, y) \in L(M)$, using (6.1) and Lemma 6.1.4, we have

$$\mathbb{M} \models \varphi(a, b) \iff \mathbb{M} \models d_p \varphi(b) \iff d_p \varphi(y) \in q \iff d_q \varphi(x) \in p \iff \mathbb{M} \models d_q \varphi(a).$$

By (6.1) with the roles of a and b exchanged, this says $b \perp_M^h a$.

(EX). By Lemma 6.1.3(a) and saturation, the type $p|B$ has a realization A' . Then we have $A' \equiv_M A$ and $A' \perp_M^h B$.

(UN). If $A \equiv_M A'$, then both have type p over M . If both are independent from B over M , then we have $\text{tp}(A/MB) = p|B = \text{tp}(A'/MB)$.

(LOC). Let a be finite, let N be a model, and put $p = \text{tp}(a/N)$. For each L -formula $\varphi(x, y)$ fix a definition $d_p \varphi(y) \in L(N)$, and let $D \subseteq N$ be the set of all parameters used, so that we have $|D| \leq \kappa$. By downward Löwenheim-Skolem there is a model $M \preceq N$ with $D \subseteq M$ and $|M| \leq \kappa$. Put $q = \text{tp}(a/M)$. Each $d_p \varphi$ is an $L(M)$ -formula, and on M it agrees with the definition of p , so the formula $d_p \varphi$ defines q for φ . These formulas therefore form a definition scheme for q (for $L(M)$ -formulas $\varphi_0(x, y, m)$ use $d_p \varphi_0(y, m)$), and so the heir $q|N$

is defined by the same formulas as p . Since $q|N$ and p are both complete types over N , we have $q|N = p = \text{tp}(a/N)$, which means that $a \downarrow_M^h N$.

(AR). Suppose that $a \downarrow_M^h a$ with $p = \text{tp}(a/M)$, and suppose that some coordinate a_0 of a is not in M . Let $\varphi(x, y)$ be the formula $x_0 = y_0$, where x_0 and y_0 are the corresponding coordinates of x and y . For $c \in M$ we have $\varphi(x, c) \notin p$, because $a_0 \neq c_0$. Hence we have $M \models \neg \exists y d_p \varphi(y)$, and the same holds in $\mathbb{M}$. By (6.1), we have $\mathbb{M} \models \varphi(a, a) \leftrightarrow d_p \varphi(a)$. However, the formula $\varphi(a, a)$ holds and the formula $d_p \varphi(a)$ fails, a contradiction.

6.2.2 Proof of (2) $\Rightarrow$ (3) $\Rightarrow$ (1)

The implication (2) $\Rightarrow$ (3) is immediate, because a relation satisfying all the axioms satisfies (UN) and (LOC) in particular. For (3) $\Rightarrow$ (1), the idea is that the axiom (LOC) lets us describe each type over a model N by a type over a small submodel, and the axiom (UN) guarantees that this description loses no information; counting the descriptions then bounds the number of types. Let $\downarrow$ satisfy (UN) and (LOC), and put $\lambda = 2^\kappa$, so that $\lambda^\kappa = \lambda$. By Lemma 5.1.2, it suffices to show that the theory T is λ -stable.

Let A be a set with $|A| \leq \lambda$, and let x be a finite tuple of variables. By downward Löwenheim–Skolem there is a model $N \supseteq A$ with $|N| \leq \lambda$, and restriction maps $S_x(N)$ onto $S_x(A)$. Thus we need only show that $|S_x(N)| \leq \lambda$.

For each $p \in S_x(N)$ choose a realization a_p . By (LOC) choose a model $M_p \subseteq N$ with $|M_p| \leq \kappa$ and $a_p \downarrow_{M_p} N$. Consider the map

$$p \mapsto (M_p, \text{tp}(a_p/M_p)).$$

This map is injective. Indeed, suppose that $M_p = M_{p'} =: M$ and $a_p \equiv_M a_{p'}$. Since $a_p \downarrow_M N$ and $a_{p'} \downarrow_M N$, the axiom (UN) gives $a_p \equiv_N a_{p'}$, which means that $p = p'$.

The number of subsets of N of size $\leq \kappa$ is at most $\lambda^\kappa = \lambda$. For each such M we have $|S_x(M)| \leq 2^\kappa = \lambda$. Hence we have $|S_x(N)| \leq \lambda \cdot \lambda = \lambda$. ☺

6.3 LOCAL CHARACTER, UNIQUENESS, AND FORKING



Remark 6.3.1. The bound $|M| \leq \kappa$ in (LOC) plays no special role in the proof of (3) $\Rightarrow$ (1): it may be weakened to $|M| < \kappa_0$ for any fixed cardinal κ_0 , by taking $\lambda = 2^{\kappa_0 + \kappa}$. Thus stability is equivalent to the existence of a relation for which types over models have unique independent extensions and every type is independent over a submodel of bounded size; all the other axioms then hold automatically, by (1) $\Rightarrow$ (2).



Remark 6.3.2. The relation $\downarrow^h$ has a counterpart over arbitrary base sets, Shelah’s nonforking independence, which we develop in Chapter 7. A type over $B \supseteq C$ *forks* over C if it implies a finite disjunction of formulas each of which divides over C (see Definition 5.2.3); a type that divides over C forks over C , and Theorem 7.3.3 shows that over models the two notions agree in simple theories. In a stable theory the relation $\downarrow^h$ of Definition 6.2.1 is *nonforking independence* restricted to models as bases. For $M \subseteq B$, the following are equivalent: the type $\text{tp}(a/B)$ does not fork over M ; this type is the heir of $\text{tp}(a/M)$; this type is definable over M ; see Pillay [Pil96] or Tent–Ziegler [TZ12]. Nonforking makes sense over arbitrary base sets C . Over such sets, nonforking satisfies all the axioms above except uniqueness, which holds only over algebraically closed sets in T^{eq} (Harnik–Harrington [HH84]). Example 3.2.6 shows the failure over $\emptyset$. In that theory there is only one 1-type over $\emptyset$, and it has at least two nonforking extensions to N : the type of a new element of C_1 and the type of a new

element of C_2 . Their φ -parts are defined by the conjugate formulas $E(y, c)$ with $c \in C_1$ and with $c \in C_2$, so both types are definable almost over $\emptyset$, and neither is definable over $\emptyset$. Kim and Pillay [KP97] proved that in a stable (indeed in a simple; see Chapters 5 and 7) theory, nonforking is the *only* relation on arbitrary small sets satisfying invariance, monotonicity, transitivity, finite character, symmetry, extension, local character and the independence theorem. Developing forking over arbitrary sets requires considerably more machinery; for this reason we have taken models as base sets in this chapter.

6.4 EXERCISES

The exercises follow the sections of the chapter. Throughout, M is a model, and we write T_E for the theory of an equivalence relation E with infinitely many classes, all of them infinite. We use that T_E admits elimination of quantifiers (see Exercise 10 of Section 5.7), the description of its types in Exercise 1 there, the fact that the theory of $(\mathbb{Q}, <)$ admits elimination of quantifiers (see Fact 5.5.2(b) and Example 5.5.3(2)), and the fact that it is not stable (see Exercise 1 there). The mark (*) indicates the exercises that ask for more than a direct application of the text.

1. (Section 6.1.) Assume that T is stable, let x be a finite tuple of variables, let $p \in S_x(M)$ be given, let B be a small set, and let $r \in S_x(MB)$ be a type extending p .
 - (a) Show that $r = p|B$ if and only if r has the following property: for every $L(M)$ -formula $\varphi(x, y)$ and every tuple b from B with $\varphi(x, b) \in r$, there is a tuple c from M with $\varphi(x, c) \in p$. (Hint: for one direction, note that $\exists y d_p \varphi(y)$ is an $L(M)$ -sentence. For the other, apply the property to the formula $\varphi(x, y) \wedge \neg d_p \varphi(y)$.) This property is the definition of an heir given by Lascar and Poizat.
 - (b) Let a be a finite tuple with $a \downarrow_M^h B$, and suppose that some formula $\psi(x) \in \text{tp}(a/MB)$ has only finitely many realizations in M . Show that a is a tuple from M . (Hint: write $\psi(x)$ as $\psi_0(x, b)$ with $\psi_0 \in L(M)$ and b from B , and apply (a) to the conjunction of $\psi_0(x, y)$ with a formula saying that $\psi_0(x', y)$ has at most k solutions x' . A finite set defined by a formula with parameters from M is contained in M .) Deduce from (MON) that $a \downarrow_M^h B$ implies $a \cap B \subseteq M$, and obtain (AR) again.
2. (Section 6.1.) Let T_∞ be the theory of an infinite set in the empty language, and assume that T_∞ admits elimination of quantifiers.
 - (a) Let A be a small set. Show that the type over A of an n -tuple is determined by which of its coordinates are equal to each other, and by which of its coordinates are equal to an element of A , and to which one. Deduce that $|S_x(A)| \leq |A| + \aleph_0$ for finite x , and that T_∞ is stable (see Lemma 5.1.2).
 - (b) Let $c \in M$ be given, let p_c be the type of c over M , and let $p_\infty \in S_1(M)$ be the type of an element outside M ; by (a), this type does not depend on the element. Show that, for the formula $x = y$, the formula $y = c$ is a definition of p_c and the formula $y \neq y$ is a definition of p_∞ . Let B be a small set. Deduce that the heir $p_\infty|B$ is the type over MB of an element outside MB , and that the heir $p_c|B$ is the type of c over MB .

- (c) Let a be a finite tuple and B a small set. Show that $a \perp_M^h B$ if and only if $a \cap B \subseteq M$. (Hint: use Exercise 1(b) for one direction; for the other, compare a realization of $\text{tp}(a/M)|B$ with a by means of (a).) Deduce by (FC) that, in T_∞ , heir independence is trivial independence (see Definition 5.4.8) with models as bases.
3. (Section 6.1.) Work in T_E . Let $m \in M$ be given, let a be an element outside M with aEm , and let a' be an element that is E -equivalent to no element of M . Put $p = \text{tp}(a/M)$ and $p' = \text{tp}(a'/M)$, and let B be a small set.
- (a) Show that, for the formula $x = y$, the formula $y \neq y$ is a definition of p and of p' . Show that, for the formula xEy , the formula yEm is a definition of p and the formula $y \neq y$ is a definition of p' .
- (b) Deduce that the heir $p|B$ is the type over MB of an element outside MB that is E -equivalent to m , and that the heir $p'|B$ is the type over MB of an element that is E -equivalent to no element of MB . (Hint: by elimination of quantifiers, a complete type over MB in one variable x is determined by the formulas $x = d$ and xEd , with $d \in MB$, that it contains.) Thus the heir of the type of a new element of an old class lies in that class, and the heir of the type of an element of a new class lies in a class that is new over MB .
- (c) Let e be an element. Show that $e \perp_M^h B$ holds if and only if $e \notin B \setminus M$ and, whenever e is E -equivalent to an element of B , the element e is E -equivalent to an element of M . Compare with Exercise 4(c) of Section 5.7, and conclude that $e \perp_M^h B$ holds if and only if $\text{tp}(e/MB)$ does not divide over M .
4. (*) (Section 6.1.) Let K be a finite field, and let T be the theory of infinite vector spaces over K , in the language of Exercise 2 of Section 1.3. Assume that T admits elimination of quantifiers. For a set D , write $\langle D \rangle$ for the subspace of $\mathbb{M}$ spanned by D .
- (a) Let A be a small set and $a = (a_1, \dots, a_n)$ a tuple. Show that $\text{tp}(a/A)$ is determined by the set of pairs (λ, w) with $\lambda \in K^n$, $w \in \langle A \rangle$ and $\sum_i \lambda_i a_i = w$. Deduce that $|S_x(A)| \leq |A| + \aleph_0$ for finite x , and that T is stable.
- (b) Show that all elements outside M have the same type p_g over M , the *generic type*. Let $\varphi(x, y_1, \dots, y_k)$ be the formula $x = \mu_1 y_1 + \dots + \mu_k y_k + m$, with $k \geq 1$, $\mu_j \in K$ and $m \in M$. Show that the formula $y_1 \neq y_1$ is a definition of p_g for φ . Let B be a small set. Deduce that the heir $p_g|B$ is the type over MB of an element outside $\langle MB \rangle$.
- (c) Let a be a finite tuple and B a small set. Show that $a \perp_M^h B$ if and only if $\langle Ma \rangle \cap \langle MB \rangle = M$. (Hint: for one direction, apply Exercise 1(a) to a formula $\sum_i \lambda_i x_i = \sum_j \mu_j b_j + m$; for the other, use (a).)
- (d) Let a and b be elements. Show that $a \perp_M^h b$ holds if and only if $a \in M$, or $b \in M$, or a and b are linearly independent over M . Observe that the conditions in (c) and (d) are visibly symmetric, and derive (AR) from (c).
5. (Subsection 6.1.1.) Work in T_E , and let $\varphi(x, y)$ be the formula xEy . By Exercise 1 of Section 5.7, the types in $S_1(M)$ are the following: for each $c \in M$, the type r_c of c over M ; for each E -class C of M , the type s_C of an element outside M that is E -equivalent to the elements of C ; and the type t of an element that is E -equivalent to no element of

- M . Say that r_c determines the class of c , that s_C determines C , and that t determines no class.
- (a) For each type in $S_1(M)$, find a definition for φ . Since φ is symmetric, the same formulas, with x in place of y , are definitions for φ regarded as a formula in the variable y .
- (b) Let $p, q \in S_1(M)$ be given. Verify Harrington's lemma (see Lemma 6.1.4) for p, q and φ directly: show that $d_p\varphi(y) \in q$ holds if and only if some E -class of M is determined both by p and by q , and that the same holds for $d_q\varphi(x) \in p$.
- (c) Let $\mathcal{U}$ and $\mathcal{V}$ be ultrafilters on M with traces p and q . Show that an E -class C of M belongs to $\mathcal{U}$ if and only if p determines C . Show that $\lim_{a', \mathcal{U}} \lim_{b', \mathcal{V}} \varphi(a', b') = 1$ holds if and only if some E -class of M belongs to both $\mathcal{U}$ and $\mathcal{V}$, and that the same holds for the other iterated limit. Thus here both iterated limits depend only on the traces p and q ; compare Exercise 7.
6. (Subsection 6.1.1.) Work in the theory of $(\mathbb{Q}, <)$, let $\varphi(x, y)$ be the formula $x < y$, and let $p \in S_1(M)$ be the type of an element greater than every element of M ; by elimination of quantifiers, this type does not depend on the element.
- (a) For every $L(M)$ -formula $\chi(x, y)$, let $d_p\chi(y)$ be the formula $\exists z \forall t (z < t \rightarrow \chi(t, y))$. Show that, for every tuple b from M , we have $\chi(x, b) \in p$ if and only if $M \models d_p\chi(b)$. (Hint: by elimination of quantifiers, the truth value of $\chi(t, b)$ is the same for all t greater than the finitely many parameters of $\chi(x, b)$.) Thus the type p has a definition scheme, although the theory is not stable. Show that the proof of Lemma 6.1.3(a) uses stability only through the existence of a definition scheme, so that, for every small set B , the set $p|B$ given by the formula of Definition 6.1.2 is a complete type over MB extending p .
- (b) Let b be an element greater than every element of M . Show that $p|\{b\}$ contains the formula $b < x$, and that every type over Mb that extends p and is finitely satisfiable in M contains the formula $x < b$. Let $\mathcal{U}$ be an ultrafilter on M that contains the set $\{c \in M : m < c\}$ for every $m \in M$. Show that $q_{\mathcal{U}}|\emptyset = p$ and that $q_{\mathcal{U}}|\{b\} \neq p|\{b\}$. Thus the heir and the coheir of p differ; compare Exercise 8.
- (c) Let $q \in S_1(M)$ be the type p written in the variable y , and let $d_q\varphi(x)$ be the definition of q given by (a) for φ regarded as a formula in the variable y . Show that the formula $d_p\varphi(y)$ is equivalent to $y \neq y$ and that the formula $d_q\varphi(x)$ is equivalent to $x = x$. Conclude that $d_p\varphi(y) \notin q$ and $d_q\varphi(x) \in p$, so that the conclusion of Harrington's lemma fails, and compute the two iterated limits of φ along $\mathcal{U}$ in both variables.
- (d) Let b realize q and let a realize $p|\{b\}$; let a' realize p and let b' realize $q|\{a'\}$. Show that $b < a$ and $a' < b'$. Thus the two orders of realization give pairs with different types over M , although in both pairs each coordinate realizes p . Deduce that $\text{tp}(a/Mb) = \text{tp}(a/M)|\{b\}$ but $\text{tp}(b/Ma) \neq \text{tp}(b/M)|\{a\}$, so that the relation “ $\text{tp}(a/Mb) = \text{tp}(a/M)|\{b\}$ ”, for elements whose types over M have definition schemes, is not symmetric here.
7. (Subsection 6.1.1.) Work in the theory of $(\mathbb{Q}, <)$, regard $(\mathbb{Q}, <)$ as a model, and let $\varphi(x, y)$ be the formula $x < y$. Let r be an irrational real number, and let $p_r \in S_1(\mathbb{Q})$ be the type that contains $c < x$ for every rational $c < r$ and $x < c$ for every rational $c > r$.

- (a) Show that there are ultrafilters $\mathcal{U}$ and $\mathcal{V}$ on $\mathbb{Q}$ such that $\mathcal{U}$ contains the set $\{c \in \mathbb{Q} : s < c < r\}$ for every rational $s < r$, and $\mathcal{V}$ contains the set $\{c \in \mathbb{Q} : r < c < s\}$ for every rational $s > r$. Show that $q_{\mathcal{U}}|\emptyset = q_{\mathcal{V}}|\emptyset = p_r$.
- (b) Show that $\lim_{a', \mathcal{U}} \lim_{b', \mathcal{U}} \varphi(a', b') = 1$ and $\lim_{b', \mathcal{U}} \lim_{a', \mathcal{U}} \varphi(a', b') = 0$, while $\lim_{a', \mathcal{U}} \lim_{b', \mathcal{V}} \varphi(a', b') = \lim_{b', \mathcal{V}} \lim_{a', \mathcal{U}} \varphi(a', b') = 1$. Thus whether the iterated limits agree depends on the ultrafilters, and not only on their traces; compare Exercise 5(c).
- (c) Let b realize p_r and let a realize $q_{\mathcal{U}}|\{b\}$; let a' realize p_r and let b' realize $q_{\mathcal{U}}|\{a'\}$. Show that $a < b$ and $b' < a'$, so that the two orders of realization give pairs with different types over $\mathbb{Q}$. Show also that $q_{\mathcal{U}}|\{b\} \neq q_{\mathcal{V}}|\{b\}$, so that p_r has two different extensions to $\mathbb{Q} \cup \{b\}$ that are finitely satisfiable in $\mathbb{Q}$; compare Exercise 8(b).
8. (*) (Subsection 6.1.1.) Assume that T is stable, let x be a finite tuple of variables, let $p \in S_x(M)$ be given, and let B be a small set.
- (a) Let $\mathcal{U}$ be an ultrafilter on $M^{|x|}$ with $q_{\mathcal{U}}|\emptyset = p$. Show that $q_{\mathcal{U}}|B = p|B$. (Hint: let $\varphi(x, y) \in L(M)$ and a tuple b from B be given, and put $q = \text{tp}(b/M)$. Show that $\{a \in M^{|x|} : \mathbb{M} \models \varphi(a, b)\}$ is the set of tuples of M that satisfy $d_q\varphi(x)$, and apply Lemma 6.1.4.)
- (b) Deduce that $p|B$ is the only type over MB that extends p and is finitely satisfiable in M . (Hint: Lemma 6.1.5(b).)
- (c) Let a and b be finite tuples. Deduce that $a \downarrow_M^h B$ holds if and only if $\text{tp}(a/MB)$ is finitely satisfiable in M , and that $\text{tp}(a/Mb)$ is finitely satisfiable in M if and only if $\text{tp}(b/Ma)$ is. Thus, in a stable theory, coheir independence is symmetric; compare Exercise 2 of Section 7.7, and Exercises 6 and 7.
9. (Section 6.2.) Work in T_E .
- (a) Let A and B be small sets. Show that $A \downarrow_M^h B$ holds if and only if the following two conditions hold: the intersection $A \cap B$ is contained in M , and every element of A that is E -equivalent to an element of B is E -equivalent to an element of M . (Hint: by (FC), it suffices to treat finite tuples a and b . For one direction, apply Exercise 1(a) to the formulas $x_i = y_j$ and $x_i E y_j$. For the other, let a' realize $\text{tp}(a/M)|b$, and compare the data of Exercise 1 of Section 5.7 for a and a' over Mb .)
- (b) Using only the description in (a), verify the axioms (SYM), (BMON) and (TR) and anti-reflexivity (AR) for $\downarrow^h$ in T_E .
- (c) Verify (UN) directly: let $a \equiv_M a'$ with $a \downarrow_M^h B$ and $a' \downarrow_M^h B$ be given, and show by Exercise 1 of Section 5.7 that $a \equiv_{MB} a'$.
- (d) Let a be an n -tuple and N a model. Show that $a \downarrow_M^h N$ holds for every model $M \subseteq N$ that contains the coordinates of a lying in N and, for every coordinate of a that is E -equivalent to an element of N , some element of N that is E -equivalent to it. Deduce (LOC), and show that there is a set $F \subseteq N$ with at most n elements such that every countable model M with $F \subseteq M \subseteq N$ satisfies $a \downarrow_M^h N$.
10. (Section 6.2.) Work in the theory of $(\mathbb{Q}, <)$, and consider trivial independence $\downarrow^t$ (see Definition 5.4.8) with models as bases.
- (a) Show that $\downarrow^t$ satisfies (LOC).

- (b) Let p be as in Exercise 6, let b be an element greater than every element of M , and let a and a' be elements with $M < a < b < a'$. Show that $a \equiv_M a'$, $a \downarrow_M^t b$ and $a' \downarrow_M^t b$, but $a \not\equiv_{Mb} a'$. Thus the relation $\downarrow^t$ does not satisfy (UN); by Exercise 6, the type $\text{tp}(a/Mb)$ is the coheir of p , and the type $\text{tp}(a'/Mb)$ is the set $p|\{b\}$.
- (c) Let N be a model, let $M \subseteq N$ be a countable model with $M \neq N$, and let $n \in N \setminus M$ be given. Find elements a and a' outside N with $a < n < a'$ and $a \equiv_M a'$. Show that $a \downarrow_M^t N$ and $a' \downarrow_M^t N$, but $a \not\equiv_N a'$. Thus, if the proof of (3) $\Rightarrow$ (1) in Theorem 6.2.2 is run with $\downarrow^t$, the map $p \mapsto (M_p, \text{tp}(a_p/M_p))$ need not be injective.
- (d) Deduce from Theorem 6.2.2 that no relation satisfies both (UN) and (LOC) in this theory.
- 11.** (Section 6.2.) Work in T_E with heir independence, and let N be a model. Fix a countable model $M_0 \subseteq N$ and, for every $n \in N$, a countable model M_n with $M_0 \cup \{n\} \subseteq M_n \subseteq N$.
- (a) Let $p \in S_1(N)$ be given, and let a_p realize p . Show that $a_p \downarrow_M^h N$ holds for some $M \in \{M_0\} \cup \{M_n : n \in N\}$. (Hint: Exercise 9(d).)
- (b) Run the proof of (3) $\Rightarrow$ (1) in Theorem 6.2.2 with this choice of the models M_p . Show that the map $p \mapsto (M_p, \text{tp}(a_p/M_p))$ is injective and that its image has at most $|N|$ elements, so that $|S_1(N)| \leq |N|$. (Hint: each M_n is countable, so we have $|S_1(M_n)| \leq \aleph_0$ by Exercise 1 of Section 5.7.)
- (c) Show directly that $|S_1(N)| = |N|$. Compare with the bound $|S_1(N)| \leq 2^{\aleph_0}$ that the proof gives, with $\lambda = 2^{\aleph_0}$, for models N with $|N| \leq 2^{\aleph_0}$.
- 12.** (Section 6.3.)
- (a) Let κ_0 be an infinite cardinal, and let $\downarrow$ be a relation satisfying (UN) and the weakening of (LOC) in which the condition $|M| \leq \kappa$ is replaced by $|M| < \kappa_0$ (see Remark 6.3.1). Put $\lambda = 2^{\kappa_0 + \kappa}$. Show that $\lambda^{\kappa_0} = \lambda$, that a set of size at most λ has at most λ subsets of size less than κ_0 , and that $|S_x(M)| \leq \lambda$ for every finite x and every model M with $|M| < \kappa_0$. Complete the proof that T is stable, following the proof of (3) $\Rightarrow$ (1) in Theorem 6.2.2.
- (b) Let $A \downarrow_M B$ mean that $A \subseteq M$ or $B \subseteq M$. Show that this relation satisfies (UN) in every complete theory, and that it does not satisfy (LOC) in the theory of $(\mathbb{Q}, <)$. (Hint: consider a model N with $|N| > \aleph_0$ and an element outside N .) Together with Exercise 10, conclude that neither (UN) nor (LOC) alone implies stability, so that both are needed in condition (3) of Theorem 6.2.2.
- 13.** (*) (Section 6.3.) Let T_2 be the theory of an equivalence relation E with exactly two classes, both infinite (see Example 3.2.6), and assume that T_2 admits elimination of quantifiers. Let N be a countable model with classes C_1 and C_2 , and, for $i = 1, 2$, let $p_i \in S_1(N)$ be the type of an element outside N that is E -equivalent to the elements of C_i .
- (a) Show that T_2 is stable, that $S_1(\emptyset)$ has exactly one element, and that $S_1(N)$ consists of the types realized in N together with p_1 and p_2 .
- (b) Show that every type realized in N divides over $\emptyset$ (see Exercise 2 of Section 5.7), and that p_1 and p_2 do not divide over $\emptyset$ (see Definition 5.2.3). (Hint: let $(d_k)_{k < \omega}$ be a

sequence of tuples indiscernible over $\emptyset$. Since there are only two classes, corresponding coordinates of the tuples d_k are E -equivalent.) Thus the unique type over $\emptyset$ has exactly two extensions to N that do not divide over $\emptyset$; compare Remark 6.3.2.

- (c) Let $M \subseteq N$ be a model. Show that M meets both classes, that the restrictions of p_1 and p_2 to M are different, and that p_i is the heir of its restriction to M . Thus over a model the two extensions are already distinguished by their restrictions.
- (d) Let $\perp$ be a relation between small sets over arbitrary small base sets C that satisfies invariance, in the sense that $A \perp_C B$ implies $\sigma(A) \perp_{\sigma(C)} \sigma(B)$ for every $\sigma \in \text{Aut}(\mathbb{M})$, and existence over $\emptyset$, in the sense that for all A and B there is $A' \equiv_{\emptyset} A$ with $A' \perp_{\emptyset} B$. Show that uniqueness over $\emptyset$ fails: there are elements a' and a'' with $a' \equiv_{\emptyset} a''$, $a' \perp_{\emptyset} N$ and $a'' \not\perp_{\emptyset} N$, but $a' \not\equiv_N a''$. (Hint: an automorphism of N that exchanges C_1 and C_2 extends, by strong homogeneity, to an automorphism σ of $\mathbb{M}$ with $\sigma(N) = N$.)

6.5 HISTORICAL REMARKS

✂ **Abstract independence.** The axiomatic study of independence is older than model theory. Steinitz, in his foundational paper on the algebraic theory of fields [Ste10], introduced transcendence bases and showed that any two transcendence bases have the same cardinality. The argument is an exchange argument of the kind now called the Steinitz exchange lemma, the same one that gives the dimension of a vector space. Van der Waerden, in *Moderne Algebra* [vdW30], observed that linear dependence and algebraic dependence obey the same few rules, and derived the existence of bases and the invariance of their cardinality from an abstract notion of dependence whose central axiom is exchange. Whitney [Whi35] introduced *matroids*, structures that isolate the abstract properties of linear dependence, and axiomatized them in several equivalent ways: by independent sets, bases, circuits and rank. Nakasawa developed the same theory independently in the mid-1930s, and Birkhoff [Bir35] and Mac Lane [ML36] related matroid theory to lattices and to projective geometry.

After a period in which matroids were studied mainly by Rado and Tutte, Rota made them central to combinatorics. In his monograph with Crapo [CR70] they become *combinatorial geometries*, and a closure operator with the exchange property is called a *pregeometry*, the term still used in model theory. Rota's view was that dependence is a geometric notion with invariants of its own, such as the Möbius function and the characteristic polynomial of the associated lattice.

Model theory met these ideas in the study of uncountably categorical theories. Baldwin and Lachlan [BL71] showed that in a strongly minimal set algebraic closure has the exchange property, so algebraic closure is a pregeometry and every model has a well-defined dimension. They used this to give a new proof of Morley's categoricity theorem [Mor65]. Shelah's forking [She90] carries independence beyond pregeometries. Nonforking independence is not given by a closure operator; yet in a stable theory it has the symmetry, existence and uniqueness on which arguments about dimension depend. The work of Zilber and Hrushovski on the geometry of forking made of it a subject of its own, geometric stability theory [Pil96]. The point of view of this chapter, which characterizes independence by axioms, as van der Waerden and Whitney did for linear dependence, is that of Harnik and Harrington, discussed below.

- ⌘ **Section 6.1.** Heirs and coheirs (*héritiers* and *cohéritiers*) were introduced by Lascar and Poizat [LP79]. They gave an alternative presentation of Shelah’s forking in stable theories, based only on definability of types and finite satisfiability. Lemma 6.1.4, which says that the definitions of two types over a model commute, is due to Harrington; it appears in Harnik–Harrington [HH84]. It is the key step in the proof that independence in stable theories is symmetric (see the proof of Theorem 6.2.2).
- ⌘ **Section 6.2.** Nonforking independence was introduced by Shelah [She90], and its reformulation through heirs is the one of Lascar and Poizat recalled above. The characterization of stability through abstract independence notions is due to Harnik and Harrington [HH84]. Kim and Pillay [KP97] extended it to simple theories, with the independence theorem in place of uniqueness (see Chapter 7). Theorem 6.2.2 is the stable case over models, with uniqueness (stationarity over models) as the key axiom. The anchor symbol $\downarrow$ for independence was introduced by Makkai in a lecture in Jerusalem, during the Model Theory Year at the Institute for Advanced Studies (1980–81), and was adopted at once. The symbol appears in print in Makkai’s survey [Mak84].

Chapter 7

Simplicity and notions of independence

This chapter is in preparation: Sections 7.1–7.6, the exercises and the historical remarks are written; the roadmap below indicates what remains.

In Chapter 6 we characterized the stable theories as those carrying a relation of independence over models with all the properties of Definition 5.4.7, uniqueness included. In Chapter 5 the random graph, and every Fraïssé limit with disjoint 3-amalgamation, carried a relation with all these properties except uniqueness, and with the independence theorem in its place. The purpose of this chapter is to show that this is no accident of the examples. By a theorem of Kim and Pillay [KP97], a theory is simple exactly when it carries such a relation, and nondividing independence is then the weakest such relation (Kim and Pillay show that it is the only one); we prove this in Theorem 7.6.5. As Chapter 6 rested on Harrington’s lemma, this chapter rests on three results of Kim [Kim98] and of Kim and Pillay: dividing is witnessed by any Morley sequence, independence is symmetric, and independent extensions over independent sets amalgamate (see Theorems 7.2.2, 7.3.5 and 7.4.7).

Roadmap

In Chapter 6 we found, in a stable theory, exactly one way of learning nothing new. In a simple theory there are many, and in this chapter we find the order among them. The route is the one Kim found, and it has three steps, each resting on the one before. First, dividing is witnessed by a single sequence and need not be tested along every indiscernible sequence: in Section 7.1 we set up nondividing independence and Morley sequences in an arbitrary theory, and in Section 7.2 we show that in a simple theory a single Morley sequence suffices (see Theorem 7.2.2). Second, independence is symmetric: in Section 7.3 we show this, after showing that forking equals dividing over models, and we prove transitivity as well. Third, independent extensions amalgamate: in Section 7.4 we show this; this is the independence theorem (see Theorem 7.4.7), proved after Casanovas [Cas11].

The next two sections draw the consequences. In Section 7.5 we show that uniqueness, added to simplicity, is exactly stability, because two different independent extensions of one type produce the independence property. In Section 7.6 we show that simplicity is exactly the existence of a relation with the remaining axioms and the independence theorem, that every such relation implies nondividing independence, and, on the way, that local character of dividing alone characterizes simplicity. A final section, still to be written, will return to the Rado graph and the Fraïssé limits of Chapter 5, with pictures of Morley sequences, of the reversal argument, and of the amalgamation in the independence theorem. The moral of the chapter is the counterpart of the moral of the last: where there are many ways of learning nothing new, simplicity is the condition under which they can always be reconciled.

7.1 DIVIDING OVER MODELS

A type over $M \cup B$ that does not divide over M says nothing about B that M could not have foreseen: whatever sequence of copies of B the indiscernibility of M allows, the type can be realized simultaneously against all of them. In this section we turn that sentence into a lemma, and draw three consequences that hold in every theory: types finitely satisfiable in M do not divide over M , so nondividing extensions always exist; such types generate indiscernible sequences of a special kind, the Morley sequences; and nondividing is transitive on the left. Simplicity enters only in the next section.

Definition 7.1.1. For small sets A, B, C , write $A \downarrow_C^d B$ if for every finite tuple a from A the type $\text{tp}(a/CB)$ does not divide over C . We call $\downarrow_C^d$ *nondividing independence*.

A formula of $\text{tp}(a/CB)$ mentions only finitely many elements of B , and a formula that divides over C belongs to the type of a over C together with those elements. Hence we have $A \downarrow_C^d B$ if and only if $a \downarrow_C^d b$ for all finite tuples a from A and b from B ; this is finite character, built into the definition. For a finite tuple b , the notation $\text{tp}(a/Cb)$ is short for $\text{tp}(a/C \cup b)$.

Parameters from C cause no difficulty. Let $\psi(x, y) \in L(C)$ be given, and write $\psi(x, y) = \psi_0(x, y, c)$ with $\psi_0 \in L$ and c a tuple from C ; the formula $\psi(x, b)$ divides over C when $\psi_0(x, b, c)$ does, in the sense of Definition 5.2.3. A sequence indiscernible over C that begins with (b, c) has the form $((b_i, c))_{i < \omega}$, because each of its terms has the type of (b, c) over C , and conversely $((b_i, c))_{i < \omega}$ is indiscernible over C whenever $(b_i)_{i < \omega}$ is. Hence the formula $\psi(x, b)$ divides over C if and only if $\{\psi(x, b_i) : i < \omega\}$ is inconsistent for some sequence $(b_i)_{i < \omega}$, indiscernible over C , with $b_0 = b$.

7.1.1 Nondividing and indiscernible sequences

The definition of dividing quantifies over indiscernible sequences, one formula at a time. The first lemma quantifies over sequences once and for all: a type does not divide exactly when it can be moved, without changing it, to a position from which a given indiscernible sequence still looks indiscernible (see Figure 7.1).

Lemma 7.1.2. *Let C be a small set, let a be a small tuple, and let b be a finite tuple. The following are equivalent:*

- (1) *the type $\text{tp}(a/Cb)$ does not divide over C ;*
- (2) *for every sequence $(b_i)_{i < \omega}$, indiscernible over C , with $b_0 = b$, there is a' with $a' \equiv_{Cb} a$ such that $(b_i)_{i < \omega}$ is indiscernible over $C \cup a'$.*

Proof. Write $p(x, b) = \text{tp}(a/Cb)$, where $p(x, y)$ is the set of $L(C)$ -formulas $\psi(x, y)$ with $\mathbb{M} \models \psi(a, b)$, and write $p(x, b_i)$ for the set of formulas $\psi(x, b_i)$ with $\psi \in p$.

(2) $\Rightarrow$ (1). Let $\psi(x, b)$ be a formula of $\text{tp}(a/Cb)$, with $\psi(x, y) \in L(C)$, and let $(b_i)_{i < \omega}$ be indiscernible over C with $b_0 = b$. By (2) there is $a' \equiv_{Cb} a$ such that the sequence is indiscernible over $C \cup a'$. Since $\mathbb{M} \models \psi(a', b_0)$, indiscernibility over $C \cup a'$ gives $\mathbb{M} \models \psi(a', b_i)$ for every i . Hence the set $\{\psi(x, b_i) : i < \omega\}$ is consistent. By the observation following Definition 7.1.1, the formula $\psi(x, b)$ does not divide over C .

(1) $\Rightarrow$ (2). Let $(b_i)_{i < \omega}$ be indiscernible over C with $b_0 = b$. The proof has two steps: first realize all the copies $p(x, b_i)$ at once, then make the sequence indiscernible over the realization by extracting a new sequence and moving it back.

⊙ **Step 1.** The set $\Sigma(x) = \bigcup_{i < \omega} p(x, b_i)$ is consistent. Otherwise, by compactness, some finite subset is inconsistent. It involves finitely many formulas $\psi_1, \dots, \psi_r$ of p and finitely many indices $i_1, \dots, i_n$. Let ψ be the conjunction of $\psi_1, \dots, \psi_r$; then we have $\psi(x, b) \in p(x, b)$, and $\{\psi(x, b_{i_l}) : 1 \leq l \leq n\}$ is inconsistent. Hence the set $\{\psi(x, b_i) : i < \omega\}$ is inconsistent, so the formula $\psi(x, b)$ divides over C , contradicting (1). Let a'' realize Σ .

⊙ **Step 2.** Apply Lemma 5.2.2 over $C \cup a''$ to the sequence $(b_i)_{i < \omega}$. Its Ehrenfeucht–Mostowski type over $C \cup a''$ contains, for every n , the type of $(b_0, \dots, b_{n-1})$ over C , because the sequence is indiscernible over C ; it also contains every formula $\psi(a'', y_1)$ with $\psi \in p$, because a'' realizes every $p(x, b_i)$. The lemma therefore gives a sequence $(b'_i)_{i < \omega}$, indiscernible over $C \cup a''$, such that $(b'_0, \dots, b'_{n-1}) \equiv_C (b_0, \dots, b_{n-1})$ for every n and $\mathbb{M} \models \psi(a'', b'_0)$ for every $\psi \in p$. The first property gives $(b'_i)_{i < \omega} \equiv_C (b_i)_{i < \omega}$, and the second gives $a''b'_0 \equiv_C ab$. By strong homogeneity there is $\sigma \in \text{Aut}(\mathbb{M}/C)$ with $\sigma(b'_i) = b_i$ for every i . Put $a' = \sigma(a'')$. Then we have $a'b = \sigma(a''b'_0) \equiv_C ab$, that is, $a' \equiv_{Cb} a$, and the sequence $(b_i) = (\sigma(b'_i))$ is indiscernible over $\sigma(C \cup a'') = C \cup a'$. ☺




FIGURE 7.1. Lemma 7.1.2. A type $\text{tp}(a/Cb)$ that does not divide over C can be moved, over $C \cup b$, to a tuple a' that stands in the same relation to every term of a given indiscernible sequence starting with b , so that the sequence does not notice a' .

7.1.2 Coheirs do not divide

Dividing is a statement about many copies of a parameter, and a type that is finitely satisfiable in M is witnessed by single elements of M , which every copy over M must treat alike. That is the whole proof of the next lemma.

Lemma 7.1.3. *Let M be a model, let B be a small set, and let a be a small tuple such that $\text{tp}(a/MB)$ is finitely satisfiable in M . Then we have $a \downarrow_M^d B$.*

Proof. Let $\psi(x, b)$ be a formula of $\text{tp}(a/MB)$, with $\psi(x, y) \in L(M)$ and b a finite tuple from B , and let $(b_i)_{i < \omega}$ be indiscernible over M with $b_0 = b$. By finite satisfiability there is a tuple m from M with $\mathbb{M} \models \psi(m, b_0)$. The tuple m and the parameters of ψ lie in M , so indiscernibility over M gives $\mathbb{M} \models \psi(m, b_i)$ for every i . Hence the set $\{\psi(x, b_i) : i < \omega\}$ is consistent, and $\psi(x, b)$ does not divide over M . ☺

Every type over a model is finitely satisfiable in it (see Subsection 6.1.1), and such a type extends to a type over any larger set MB that is still finitely satisfiable in M (see Proposition 7.1.4 below). The following proposition, together with Lemma 7.1.3, gives the existence axiom (EX) for $\downarrow^d$ in every theory.

Proposition 7.1.4. *Let M be a model, and let A and B be small sets. Then there is $A' \equiv_M A$ such that $\text{tp}(A'/MB)$ is finitely satisfiable in M . In particular, we have $A' \downarrow_M^d B$.*

Proof. Let $\pi = \text{tp}(A/M)$, in variables X , and let F be the set of $L(MB)$ -formulas $\chi(X)$ that are satisfied by no tuple from M . We claim that $\pi \cup \{\neg\chi : \chi \in F\}$ is consistent. A finite subset is contained in $\{\theta, \neg\chi_1, \dots, \neg\chi_n\}$ for some $\theta \in \pi$ and $\chi_1, \dots, \chi_n \in F$. Since θ is an $L(M)$ -formula realized in $\mathbb{M}$ and $M \preceq \mathbb{M}$, some tuple m from M satisfies θ , and m satisfies each $\neg\chi_i$ because $\chi_i \in F$. By saturation, the set is realized by some A' . Then we have $A' \equiv_M A$, and every formula χ of $\text{tp}(A'/MB)$ is satisfied by a tuple from M , since otherwise $\chi \in F$ and $\neg\chi \in \text{tp}(A'/MB)$. Hence the type $\text{tp}(A'/MB)$ is finitely satisfiable in M , and Lemma 7.1.3 gives $A' \downarrow_M^d B$. $\odot$

For a finite tuple a , the types produced in this way are exactly the coheirs $q_{\mathcal{U}}|B$ of Lemma 6.1.5: a type over MB that is finitely satisfiable in M has the form $q_{\mathcal{U}}|B$ for some ultrafilter $\mathcal{U}$ on $M^{|a|}$ with $q_{\mathcal{U}}|\emptyset = \text{tp}(a/M)$.

7.1.3 Morley sequences

Realizing a coheir over and over, each time over everything realized so far, produces an indiscernible sequence in which every term is independent from the earlier ones (see Lemma 7.1.6 below). Such sequences are the test sequences of simplicity theory.

Definition 7.1.5. Let M be a model. A *Morley sequence over M* is a sequence $(b_i)_{i < \omega}$, indiscernible over M , such that $b_i \downarrow_M^d \{b_j : j < i\}$ for every i .

Lemma 7.1.6. Let M be a model, let b be a finite tuple, and let $\mathcal{U}$ be an ultrafilter on $M^{|b|}$ with $q_{\mathcal{U}}|\emptyset = \text{tp}(b/M)$. Put $b_0 = b$, and for $i \geq 1$ let b_i realize $q_{\mathcal{U}}|\{b_j : j < i\}$. Then the sequence $(b_i)_{i < \omega}$ is a Morley sequence over M with $b_0 = b$.

The ultrafilter exists by Lemma 6.1.5(b), since $\text{tp}(b/M)$ is finitely satisfiable in M , and the tuples b_i exist by saturation. We write $b_{<i}$ for $(b_j)_{j < i}$.

Proof. For $i \geq 1$ and $i_1 < \dots < i_n < i$, Lemma 6.1.5(c) gives $q_{\mathcal{U}}|\{b_{i_1}, \dots, b_{i_n}\} \subseteq q_{\mathcal{U}}|b_{<i} = \text{tp}(b_i/Mb_{<i})$; since the smaller set is a complete type over $M \cup \{b_{i_1}, \dots, b_{i_n}\}$ by Lemma 6.1.5(a), it equals the type of b_i over that set. In particular, every b_i realizes $q_{\mathcal{U}}|\emptyset = \text{tp}(b/M)$, for $i = 0$ by hypothesis.

We show by induction on n that all increasing n -tuples of the sequence have the same type over M . For $n = 1$ this was just observed. Let $i_1 < \dots < i_{n+1}$ and $j_1 < \dots < j_{n+1}$ be given. By the induction hypothesis and strong homogeneity there is $\sigma \in \text{Aut}(\mathbb{M}/M)$ with $\sigma(b_{i_l}) = b_{j_l}$ for $l \leq n$. By Lemma 6.1.5(d), the automorphism σ carries $q_{\mathcal{U}}|\{b_{i_1}, \dots, b_{i_n}\}$, the type of $b_{i_{n+1}}$ over $M \cup \{b_{i_1}, \dots, b_{i_n}\}$, to $q_{\mathcal{U}}|\{b_{j_1}, \dots, b_{j_n}\}$, the type of $b_{j_{n+1}}$ over $M \cup \{b_{j_1}, \dots, b_{j_n}\}$. Hence we have $\sigma(b_{i_{n+1}}) \equiv_{Mb_{j_1} \dots b_{j_n}} b_{j_{n+1}}$, and therefore $(b_{i_1}, \dots, b_{i_{n+1}}) \equiv_M (b_{j_1}, \dots, b_{j_{n+1}})$. Thus the sequence is indiscernible over M .

For every i , the type $\text{tp}(b_i/Mb_{<i})$ is a coheir. By Lemma 6.1.5(a), this type is finitely satisfiable in M , and by Lemma 7.1.3 we have $b_i \downarrow_M^d b_{<i}$. $\odot$

7.1.4 Transitivity on the left

Nondividing independence can be built up one coordinate block at a time on the left: if A is independent from B over C , and D is independent from B over $C \cup A$, then the set $A \cup D$ is independent from B over C . The proof uses the implication (1) $\Rightarrow$ (2) of Lemma 7.1.2 twice, first for A and then for D , concludes with the implication (2) $\Rightarrow$ (1), and uses no hypothesis on T .

Lemma 7.1.7. Let A, B, C, D be small sets. If $A \downarrow_C^d B$ and $D \downarrow_{CA}^d B$, then we have $AD \downarrow_C^d B$.

Proof. Regard A as a tuple, and let finite tuples e from D and b from B be given. By finite character it suffices to show that $\text{tp}(Ae/Cb)$ does not divide over C ; the types of the finite tuples from $A \cup e$ are restrictions of this type. Every formula of $\text{tp}(A/Cb)$ mentions a finite tuple from A , so the hypothesis $A \downarrow_C^d B$ says that $\text{tp}(A/Cb)$ does not divide over C ; likewise $D \downarrow_{CA}^d B$ says that $\text{tp}(e/CAb)$ does not divide over $C \cup A$. We verify condition (2) of Lemma 7.1.2 for the tuple Ae .

Let $(b_i)_{i < \omega}$ be indiscernible over C with $b_0 = b$. By Lemma 7.1.2, applied to A , there is $A' \equiv_{Cb} A$ such that (b_i) is indiscernible over $C \cup A'$. Choose $\sigma \in \text{Aut}(\mathbb{M}/Cb)$ with $\sigma(A') = A$. Then the sequence $(\sigma(b_i))_{i < \omega}$ is indiscernible over $C \cup A$ and begins with $\sigma(b_0) = b$. By Lemma 7.1.2, applied to e over the base $C \cup A$, there is $e' \equiv_{CAb} e$ such that $(\sigma(b_i))$ is indiscernible over $C \cup Ae'$. Applying σ^{-1} , which fixes $C \cup b$, we see that (b_i) is indiscernible over $C \cup A'\sigma^{-1}(e')$, and

$$A'\sigma^{-1}(e') = \sigma^{-1}(Ae') \equiv_{Cb} Ae' \equiv_{Cb} Ae.$$

By Lemma 7.1.2, the type $\text{tp}(Ae/Cb)$ does not divide over C . ☺

7.1.5 The axioms that hold in every theory

We collect what we have shown in this section, measured against Definition 5.4.7.

Proposition 7.1.8. *In every complete theory, the relation $\downarrow^d$, with models as bases, satisfies (INV), (MON), (BMON), (FC) and (EX). If T is simple, it also satisfies (LOC).*

Proof. *Invariance:* an automorphism σ of $\mathbb{M}$ carries sequences indiscernible over C to sequences indiscernible over $\sigma(C)$ and preserves consistency, so the formula $\varphi(x, b)$ divides over C if and only if $\varphi(x, \sigma(b))$ divides over $\sigma(C)$. *Monotonicity* is immediate from the definition, and finite character was observed after Definition 7.1.1. *Base monotonicity:* let $M \subseteq N \subseteq B$, and let a be a finite tuple from A . Since $NB = B = MB$, a formula of $\text{tp}(a/NB)$ that divided over N would divide over M by Remark 5.2.5. *Existence* is Proposition 7.1.4. *Local character:* let a be a finite tuple and N a model. If T is simple, Theorem 5.3.6 gives $A_0 \subseteq N$ with $|A_0| \leq \kappa$ such that $\text{tp}(a/N)$ does not divide over A_0 . By the downward Löwenheim–Skolem theorem there is a model $M \preceq N$ with $A_0 \subseteq M$ and $|M| \leq \kappa$, and $\text{tp}(a/N)$ does not divide over M , by Remark 5.2.5 once more. Thus we have $a \downarrow_M^d N$. ☺

Symmetry, transitivity and the independence theorem remain. They can fail outside simple theories; indeed Kim [Kim01] showed that symmetry of forking independence characterizes the simple theories. We prove all three from simplicity, through Kim's lemma, which is the subject of the next section.

7.2 KIM'S LEMMA

A formula divides when some indiscernible sequence of copies of its parameter makes it inconsistent. The definition asks for one such sequence, and gives no hint where to look for it. In a simple theory there is no need to look: any Morley sequence will do, and so will any sequence of realizations of the type of the parameter in which each term is independent from the earlier ones. This is Kim's lemma [Kim98], and the rest of the chapter rests on it. With the preparations of Section 7.1 the proof is short. In the proof of Theorem 7.2.2 below we show that a Morley sequence, read backwards, becomes a chain of formulas each of which divides over the earlier ones, and Lemma 5.3.5(1) turns such a chain into a tree. A simple

theory has no such trees, and so the instances of the formula along the sequence cannot be consistent. Throughout the section, M is a model.

7.2.1 Enlarging the base

The proof of Kim's lemma needs a formula that divides over M to divide also over M together with later terms of the sequence. Enlarging the base can only make dividing harder (see Remark 5.2.5); the next lemma shows that it does not make it harder when the new parameters are independent from the parameter of the formula. Since dividing will be counted with a fixed k , we note first that the parameter convention following Definition 7.1.1 applies to k -dividing as well, with the same proof: for $\psi(x, y) \in L(C)$, the formula $\psi(x, b)$ k -divides over C if and only if $\{\psi(x, b_i) : i < \omega\}$ is k -inconsistent for some sequence $(b_i)_{i < \omega}$, indiscernible over C , with $b_0 = b$.

Lemma 7.2.1. *Let C be a small set, let b be a finite tuple, and let e be a small tuple such that $\text{tp}(e/Cb)$ does not divide over C . Let $\varphi(x, y) \in L(C)$, and let $k \geq 2$ be given. If $\varphi(x, b)$ k -divides over C , then the formula $\varphi(x, b)$ k -divides over $C \cup e$.*

Proof. Let $(b_i)_{i < \omega}$ be indiscernible over C , with $b_0 = b$ and $\{\varphi(x, b_i) : i < \omega\}$ k -inconsistent. By Lemma 7.1.2 there is $e' \equiv_{Cb} e$ such that (b_i) is indiscernible over $C \cup e'$. Choose $\sigma \in \text{Aut}(\mathbb{M}/Cb)$ with $\sigma(e') = e$. The sequence $(\sigma(b_i))_{i < \omega}$ is indiscernible over $\sigma(C \cup e') = C \cup e$, and it begins with $\sigma(b) = b$. Since σ fixes the parameters of φ , which lie in C , the set $\{\varphi(x, \sigma(b_i)) : i < \omega\}$ is the image under σ of a k -inconsistent set, and so it is k -inconsistent. Hence the formula $\varphi(x, b)$ k -divides over $C \cup e$. ☺

7.2.2 The lemma

We write $b_{<i}$ for $(b_j)_{j < i}$, as in Section 7.1.

Theorem 7.2.2 (Kim's lemma). *Assume that T is simple. Let M be a model, let b be a finite tuple, and let $\varphi(x, y) \in L(M)$.*

- (a) *Suppose that $\varphi(x, b)$ divides over M , and let $(b_i)_{i < \omega}$ be a sequence such that $b_i \equiv_M b$ and $b_i \perp_M^d b_{<i}$ for every i . Then the set $\{\varphi(x, b_i) : i < \omega\}$ is inconsistent.*
- (b) *The following are equivalent:*
 - (1) *the formula $\varphi(x, b)$ divides over M ;*
 - (2) *for every Morley sequence $(b_i)_{i < \omega}$ over M with $b_0 = b$, the set $\{\varphi(x, b_i) : i < \omega\}$ is inconsistent;*
 - (3) *for some Morley sequence $(b_i)_{i < \omega}$ over M with $b_0 = b$, the set $\{\varphi(x, b_i) : i < \omega\}$ is inconsistent.*

Part (a) does not require the sequence to be indiscernible: it asks only that every term have the type of b over M and be independent from the earlier terms. This generality will be used in the proof of symmetry in the next section. The proof of (a) reads the sequence backwards, because the Morley condition points the wrong way: it makes each term independent from the *earlier* terms, while the tree construction of Lemma 5.3.5 needs each formula to divide over the parameters of the earlier formulas of the chain. Reversing a finite piece of the sequence makes the earlier terms of the chain the later terms of the sequence; transitivity on the left (see Lemma 7.1.7) shows that these later terms, taken together, are independent from the term at hand, and Lemma 7.2.1 then moves the dividing over to them.

An infinite sequence cannot be reversed into an ω -sequence, which is why the finite form of Lemma 5.3.5 is needed (see Figure 7.2).

Proof. (a) Suppose, for a contradiction, that $\{\varphi(x, b_i) : i < \omega\}$ is consistent. By Lemma 5.2.4, applied as in the parameter convention, there is $k \geq 2$ such that $\varphi(x, b)$ k -divides over M . Write $\varphi(x, y) = \varphi_0(x, y, m)$ with $\varphi_0 \in L$ and m a finite tuple from M , and regard φ_0 as a formula $\varphi_0(x; y, z)$ with parameter variables y, z . We show that φ_0 has the k -tree property, which contradicts the simplicity of T . By Lemma 5.3.5(2), it suffices to find, for every $n \geq 1$, tuples $a_0, \dots, a_{n-1}$ as in Lemma 5.3.5(1) for the formula φ_0 . Let $n \geq 1$ be given.

© **Step 1.** For $i < l < n$, we have $b_{i+1} \dots b_l \downarrow_M^d b_i$. We argue by induction on l . For $l = i + 1$, the hypothesis $b_{i+1} \downarrow_M^d b_{<i+1}$ gives $b_{i+1} \downarrow_M^d b_i$ by monotonicity. Suppose that the assertion holds for some l with $i < l < n - 1$, and put $A = \{b_{i+1}, \dots, b_l\}$. The hypothesis on b_{l+1} says that $\text{tp}(b_{l+1}/Mb_{<l+1})$ does not divide over M . Its restriction to $M \cup A \cup b_i$ does not divide over M either, and so, by Remark 5.2.5 applied to $M \subseteq M \cup A \subseteq M \cup A \cup b_i$, it does not divide over $M \cup A$. Thus we have $b_{l+1} \downarrow_{MA}^d b_i$. Together with $A \downarrow_M^d b_i$, Lemma 7.1.7 gives $Ab_{l+1} \downarrow_M^d b_i$, which is the assertion for $l + 1$.

© **Step 2.** For $i < n$, the formula $\varphi(x, b_i)$ k -divides over $M \cup \{b_{i+1}, \dots, b_{n-1}\}$. Since $b_i \equiv_M b$, there is $\sigma \in \text{Aut}(\mathbb{M}/M)$ with $\sigma(b) = b_i$; it fixes the parameters of φ , and it carries a sequence witnessing that $\varphi(x, b)$ k -divides over M to one witnessing that $\varphi(x, b_i)$ k -divides over M . For $i = n - 1$ there is nothing more to show. For $i < n - 1$, Step 1 with $l = n - 1$ says that $\text{tp}(b_{i+1} \dots b_{n-1}/Mb_i)$ does not divide over M , and Lemma 7.2.1, with $e = b_{i+1} \dots b_{n-1}$, gives the assertion.

© **Step 3.** We reverse the order. For $j < n$, put $a_j = (b_{n-1-j}, m)$. Then we have $\{\varphi_0(x, a_j) : j < n\} = \{\varphi(x, b_i) : i < n\}$, which is consistent. Let $j < n$ be given, and let $A_j = \{a_l : l < j\}$, the set of coordinates of m and of $b_{n-j}, \dots, b_{n-1}$ (for $j = 0$, the empty set). By Step 2, with $i = n - 1 - j$, the formula $\varphi(x, b_{n-1-j})$ k -divides over $C_j = M \cup \{b_{n-j}, \dots, b_{n-1}\}$; that is, $\varphi_0(x, a_j)$ k -divides over C_j . Since $A_j \subseteq C_j$, Remark 5.2.5 shows that $\varphi_0(x, a_j)$ k -divides over A_j . Thus the elements $a_0, \dots, a_{n-1}$ satisfy the hypotheses of Lemma 5.3.5(1) for the formula φ_0 , as required.

(b) (1) $\Rightarrow$ (2). A Morley sequence over M with $b_0 = b$ is indiscernible over M , so we have $b_i \equiv_M b$ for every i , and it satisfies $b_i \downarrow_M^d b_{<i}$ by definition. Part (a) applies.

(2) $\Rightarrow$ (3). There is a Morley sequence over M beginning with b , by Lemma 7.1.6.

(3) $\Rightarrow$ (1). A Morley sequence over M is indiscernible over M , so this is the definition of dividing, with the parameter convention. ☺

7.2.3 Without simplicity

In a theory that is not simple, dividing may be visible along one Morley sequence and invisible along another. The theory of $(\mathbb{Q}, <)$ shows how. It is not simple, by Proposition 5.3.2, since the formula $x < y$ has the strict order property. We use quantifier elimination as in Example 5.2.6, now with parameters: applied to a tuple together with finite tuples from a set D , it shows that the type of the tuple over D is determined by the order between its coordinates and the elements of D . For an element c , write $M < c$ if c is greater than every element of M .

Example 7.2.3. Let T be the theory of $(\mathbb{Q}, <)$, let M be a model, and let $\varphi(x; y_1, y_2)$ be the formula $y_1 < x < y_2$. Let $b = (d_1, d_2)$ with $M < d_1 < d_2$.




FIGURE 7.2. The proof of Kim's lemma, for $n = 5$. The Morley condition makes each term independent from the terms on its left (arrows above). Transitivity on the left makes the terms on the right of b_1 , taken together, independent from b_1 (shaded), and Lemma 7.2.1 then makes $\varphi(x, b_1)$ divide over M together with them. Read from right to left, the sequence is a chain in which every formula divides over the earlier ones, and Lemma 5.3.5(1) builds a tree of height 5 from it.

The formula $\varphi(x, b)$ divides over M . Choose pairs $b^i = (b_1^i, b_2^i)$ with $b^0 = b$, with $M < b_1^i < b_2^i$, and with $b_2^i < b_1^{i+1}$ for every i . Every increasing tuple of the sequence (b^i) is ordered in the same way, and all its coordinates are greater than every element of M . By quantifier elimination, the sequence is indiscernible over M . The intervals (b_1^i, b_2^i) are pairwise disjoint, so the formulas $\varphi(x, b^i)$ are 2-inconsistent. A Morley sequence along which the formula is consistent. Choose pairs $b_i = (b_{i,1}, b_{i,2})$ with $b_0 = b$, and with

$$M < b_{i+1,1} < b_{i,1} \quad \text{and} \quad b_{i,2} < b_{i+1,2}$$

for every i , so that the intervals $(b_{i,1}, b_{i,2})$ are nested and grow outward (see Figure 7.3). As before, the sequence (b_i) is indiscernible over M ; and every element of the interval (d_1, d_2) satisfies all the formulas $\varphi(x, b_i)$, so these are consistent. It remains to show that $b_i \downarrow_M^d b_{<i}$ for $i \geq 1$ (for $i = 0$ there is nothing to prove), and we verify condition (2) of Lemma 7.1.2 for the finite tuple $c = b_{<i}$.

Let $(c_j)_{j < \omega}$ be indiscernible over M with $c_0 = c$. Every coordinate of c is greater than every element of M , and each c_j has the type of c over M , so the same holds for every coordinate of every c_j . By saturation there are a'_1, a'_2 with $M < a'_1$, with a'_1 smaller than every coordinate of every c_j , and with a'_2 greater than all these coordinates. Put $a' = (a'_1, a'_2)$. For $j_1 < \dots < j_r$, the order between the coordinates of a' , the coordinates of $c_{j_1}, \dots, c_{j_r}$, and the elements of M is determined by the order between the coordinates of $c_{j_1}, \dots, c_{j_r}$ and the elements of M , which does not depend on $j_1 < \dots < j_r$. By quantifier elimination, (c_j) is indiscernible over $M \cup a'$. Finally, $b_{i,1}$ lies above M and below every coordinate of c , and $b_{i,2}$ lies above every coordinate of c , exactly as a'_1 and a'_2 do; so we have $a' \equiv_{M \cup c} b_i$. By Lemma 7.1.2, we have $b_i \downarrow_M^d b_{<i}$, and (b_i) is a Morley sequence over M .

The coheir sequence sees the dividing. Let $\mathcal{U}$ be an ultrafilter on M^2 with $q_{\mathcal{U}}|_{\emptyset} = \text{tp}(b/M)$, and let $(b'_i)_{i < \omega}$ be the Morley sequence of Lemma 7.1.6, with $b'_0 = b$ and $b'_i =$

$(b'_{i,1}, b'_{i,2})$. Let $j < i$ be given. Since $b'_j \equiv_M b$, we have $M < b'_{j,1}$, so every pair (m_1, m_2) from M satisfies $m_2 < b'_{j,1}$, and the formula $y_2 < b'_{j,1}$ belongs to $q_U|_{b'_{<i}} = \text{tp}(b'_i/Mb'_{<i})$. Hence we have $b'_{i,2} < b'_{j,1}$: each interval lies entirely below the earlier ones, and the formulas $\varphi(x, b'_i)$ are 2-inconsistent, as Kim's lemma would require.




FIGURE 7.3. Example 7.2.3, in the theory of $(\mathbb{Q}, <)$. The model M is shaded, and every interval lies above it. Both sequences of intervals are Morley sequences over M beginning with the same interval $b_0 = b$. Along the coheir sequence (top) the intervals move down toward M and are pairwise disjoint; along the nested sequence (bottom) they grow outward and all contain b_0 . Kim's lemma says that in a simple theory this cannot happen.

Kim's lemma is the form in which simplicity enters the rest of the chapter. In the next section we use it to show that a formula forks over a model exactly when it divides, so that nondividing types have nondividing extensions to any larger set, and that nondividing independence is symmetric and transitive.

7.3 FORKING, SYMMETRY AND TRANSITIVITY

Dividing is easy to test and hard to extend; forking is easy to extend and hard to test. A formula divides when one sequence of its conjugates is inconsistent, which can be checked; but a type that does not divide over C need not have an extension to a larger set that does not divide over C , because a disjunction of formulas that divide need not divide (see Remark 7.3.2). Forking closes dividing under disjunctions and implications, and so makes extension automatic, at the price of an extra quantifier. In Theorem 7.3.3 below we show that in a simple theory, over a model, the two notions coincide, and one has the advantages of both. From this and Kim's lemma, symmetry follows in a few lines (see Theorem 7.3.5), and transitivity follows from symmetry and Lemma 7.1.7 (see Theorem 7.3.6). Throughout the section, M and N are models.

7.3.1 Forking

Definition 7.3.1. Let C be a small set, and let $\varphi(x, y) \in L(C)$.

- (a) The formula $\varphi(x, b)$ *forks over* C if there are $r \geq 1$ and formulas $\psi_1(x, c_1), \dots, \psi_r(x, c_r)$, with $\psi_l(x, z_l) \in L(C)$, each of which divides over C (in the sense of the parameter convention following Definition 7.1.1), such that

$$\mathbb{M} \models \forall x \left(\varphi(x, b) \rightarrow \bigvee_{l=1}^r \psi_l(x, c_l) \right).$$

- (b) Let $B \supseteq C$ be small. A type p over B *forks over* C if $p \vdash \bigvee_{l=1}^r \psi_l(x, c_l)$ for some formulas $\psi_l(x, c_l)$ as in (a).

This agrees with Remark 6.3.2. A formula that divides over C forks over C (take $r = 1$ and $\psi_1 = \varphi$), and so does every formula that implies a formula forking over C , and every finite disjunction of formulas forking over C . By compactness, a type that is closed under finite conjunctions, in particular a complete type, forks over C if and only if one of its formulas forks over C . Hence a complete type over B that forks over C divides over C if forking and dividing agree for formulas.



Remark 7.3.2. Forking and dividing differ in general. The standard example is the theory of the dense circular order, with a ternary relation $R(x, y, z)$, read “going around the circle in the positive direction from x , one meets y before z ” (see [TZ12, Chapter 7]). Choose three distinct elements c_1, c_2, c_3 in this order around the circle. Every element is equal to some c_l or lies on one of the three open arcs between them, so the formula $x = x$ implies the disjunction of the six formulas $x = c_l$ and $R(c_l, x, c_{l+1})$ (indices modulo 3). Each of the six divides over $\emptyset$. Indeed, by quantifier elimination for this theory, all pairs of distinct points have the same type over $\emptyset$, and a sequence of points $e_0, e_1, \dots$ in positive order around the circle, or of pairs (e_i, f_i) with $e_0, f_0, e_1, f_1, \dots$ in positive order, is indiscernible over $\emptyset$. An automorphism moves e_0 to c_l , or (e_0, f_0) to (c_l, c_{l+1}) , and the formulas $x = e_i$, like the formulas $R(e_i, x, f_i)$ of the pairwise disjoint arcs, are 2-inconsistent. Hence the formula $x = x$ forks over $\emptyset$. It does not divide over $\emptyset$, since it has no parameters and is consistent. Over models of a simple theory this cannot happen, as we now show.

7.3.2 Forking equals dividing

Theorem 7.3.3. Assume that T is simple. Let b be a finite tuple, and let $\varphi(x, y) \in L(M)$. Then the formula $\varphi(x, b)$ forks over M if and only if $\varphi(x, b)$ divides over M . Consequently, for a small tuple a and a small set B , the type $\text{tp}(a/MB)$ forks over M if and only if it divides over M .

Proof. A formula that divides over M forks over M . For the converse, suppose that $\varphi(x, b)$ does not divide over M , and let formulas $\psi_1(x, c_1), \dots, \psi_r(x, c_r)$ be given, with $\psi_l(x, z_l) \in L(M)$, such that $\varphi(x, b)$ implies their disjunction. We show that some $\psi_l(x, c_l)$ does not divide over M .

Let $c = (b, c_1, \dots, c_r)$, and let $(c^i)_{i < \omega}$ be a Morley sequence over M with $c^0 = c$, which exists by Lemma 7.1.6. Write $c^i = (b^i, c_1^i, \dots, c_r^i)$, with the blocks of the lengths of $b, c_1, \dots, c_r$. The sequence $(b^i)_{i < \omega}$ is indiscernible over M , and $b^i \downarrow_M^d b^{<i}$ for every i , by monotonicity on both sides; so it is a Morley sequence over M with $b^0 = b$. The formula $\varphi(x, b)$ does not divide over M , so the set $\{\varphi(x, b^i) : i < \omega\}$ is consistent; let a realize it.

The $L(M)$ -formula $\forall x (\varphi(x, y) \rightarrow \bigvee_l \psi_l(x, z_l))$ is satisfied by c , and $c^i \equiv_M c$; so the formula $\varphi(x, b^i)$ implies $\bigvee_l \psi_l(x, c_l^i)$ for every i . Hence for every i there is l with $\mathbb{M} \models \psi_l(a, c_l^i)$, and by the pigeonhole principle there are a single l and an infinite set $I \subseteq \omega$ such that $\mathbb{M} \models \psi_l(a, c_l^i)$ for every $i \in I$. Let $i_0 < i_1 < \dots$ enumerate I , and put $d_k = c_l^{i_k}$. As before,

$(d_k)_{k<\omega}$ is a Morley sequence over M , and a realizes $\{\psi_l(x, d_k) : k < \omega\}$. By Kim's lemma (see Theorem 7.2.2(b)), the formula $\psi_l(x, d_0)$ does not divide over M . Since $d_0 \equiv_M c_l$, an automorphism of $\mathbb{M}$ over M carries d_0 to c_l , and so the formula $\psi_l(x, c_l)$ does not divide over M either.

The statement about types follows from the observation after Definition 7.3.1. $\odot$

The first use of forking is the one it was invented for.

Corollary 7.3.4 (Extension). *Assume that T is simple. Let a be a small tuple, and let $B \subseteq B'$ be small sets such that $a \perp_M^d B$. Then there is $a' \equiv_{MB} a$ with $a' \perp_M^d B'$.*

Proof. Let $\Sigma(x)$ be the union of $\text{tp}(a/MB)$ and the set of formulas $\neg\psi(x, c)$ such that $\psi(x, z) \in L(M)$, the tuple c is from B' , and $\psi(x, c)$ divides over M . If Σ were inconsistent, then by compactness, and since $\text{tp}(a/MB)$ is closed under conjunctions, some formula $\varphi(x, b)$ of $\text{tp}(a/MB)$ would imply a finite disjunction of formulas dividing over M ; so the formula $\varphi(x, b)$ would fork over M , and by Theorem 7.3.3 it would divide over M , contradicting $a \perp_M^d B$. By saturation, the set Σ is realized by some a' . Then we have $a' \equiv_{MB} a$, and no formula of $\text{tp}(a'/MB')$ divides over M , since the negation of each formula of this kind over MB' belongs to Σ . Hence we have $a' \perp_M^d B'$. $\odot$

7.3.3 Symmetry

Symmetry says that if a tells us nothing new about b , then the tuple b tells us nothing new about a . The proof turns a failure of symmetry into a sequence of copies of a , each independent from the earlier ones and all in the same relation to b , and Kim's lemma forbids such a sequence (see Figure 7.4).

Theorem 7.3.5 (Symmetry). *Assume that T is simple, and let A and B be small sets. If $A \perp_M^d B$, then we have $B \perp_M^d A$.*

Proof. By finite character, we may assume that A and B are finite tuples a and b . Suppose, for a contradiction, that $a \perp_M^d b$ and that $\text{tp}(b/Ma)$ divides over M . Let $\psi(y, a)$ be a formula of $\text{tp}(b/Ma)$, with $\psi(y, z) \in L(M)$, that divides over M .

We choose a_i for $i < \omega$ such that $a_i \equiv_{Mb} a$ and $a_i \perp_M^d ba_{<i}$. Put $a_0 = a$, which satisfies both conditions because $a \perp_M^d b$. Given $a_{<i}$ with $i \geq 1$, apply Corollary 7.3.4 to a , with $\{b\}$ and $\{b\} \cup a_{<i}$ in place of B and B' ; it gives $a_i \equiv_{Mb} a$ with $a_i \perp_M^d ba_{<i}$.

Every a_i has the type of a over M , and $a_i \perp_M^d a_{<i}$ by monotonicity. Since $\psi(y, a)$ divides over M , Kim's lemma (see Theorem 7.2.2(a)) says that $\{\psi(y, a_i) : i < \omega\}$ is inconsistent. On the other hand, the tuple b satisfies $\psi(y, a)$, and $a_i \equiv_{Mb} a$, so it satisfies $\psi(y, a_i)$ for every i and realizes the set. This is a contradiction. $\odot$

7.3.4 Transitivity, and the axioms that hold in simple theories

With symmetry in hand, transitivity reduces to transitivity on the left (see Lemma 7.1.7): one turns the relations around, applies that lemma, and turns the result around again. The theorem below collects the axioms that nondividing independence satisfies in a simple theory.

Theorem 7.3.6. *Assume that T is simple. Then the relation $\perp^d$, with models as bases, satisfies (INV), (MON), (BMON), (TR), (FC), (SYM), (EX) and (LOC).*

Proof. All but transitivity and symmetry are in Proposition 7.1.8, and symmetry is Theorem 7.3.5. *Transitivity:* let $M \subseteq N$, and suppose that $A \perp_M^d N$ and $A \perp_N^d B$. By symmetry, we have $N \perp_M^d A$ and $B \perp_N^d A$; since $M \cup N = N$, the second relation says




FIGURE 7.4. The proof of symmetry (see Theorem 7.3.5). If $a \downarrow_M^d b$ but a formula $\psi(y, a)$ of $\text{tp}(b/Ma)$ divides over M , then extension produces copies a_i of a over $M \cup b$, each independent from the earlier ones. The tuple b realizes every $\psi(y, a_i)$, while Kim's lemma says that no tuple realizes all of them.

that $B \downarrow_{MN}^d A$. Lemma 7.1.7, with M, N, A and B in the roles of C, A, B and D , gives $NB \downarrow_M^d A$, and symmetry gives $A \downarrow_M^d NB$. ☺

By Theorem 7.3.3, over models, nondividing independence is nonforking independence; we keep the notation $\downarrow^d$. Of the properties of Definition 5.4.7, only uniqueness is missing, and it must fail in general: uniqueness together with local character implies stability (see Theorem 6.2.2), while the random graph is simple and unstable (see Theorem 5.4.6). In its place, simple theories have the independence theorem, which we prove in the next section.

7.4 THE INDEPENDENCE THEOREM

Uniqueness fails in an unstable simple theory (see Theorem 7.5.2): a type over a model may have many independent extensions to a larger set, and nothing singles one of them out. What survives is a weaker and, as we show in Theorem 7.6.5, more characteristic property. Two independent extensions of the same type, one to $M \cup a$ and one to $M \cup b$, can always be realized together, independently, provided only that a and b are themselves independent over M . This is the independence theorem of Kim and Pillay [KP97]; in the random graph it is the amalgamation of Theorem 5.4.9(c). In a stable theory it is a consequence of uniqueness; in a simple theory it has to be earned.

The proof moves a parameter step by step. Suppose that a formula $\varphi(x, a) \wedge \psi(x, b')$ does not divide over M , where b' has the type of b over M ; we want the same of $\varphi(x, a) \wedge \psi(x, b)$. If b' and b begin an indiscernible sequence, the parameter can be moved from b' to b (see Proposition 7.4.5); in general a third tuple, a coheir over everything in sight, serves as a stepping stone between them. The step itself rests on a fact about conjunctions along indiscernible sequences (see Proposition 7.4.4), and here a second idea enters: an arbitrary indiscernible sequence over M can be replaced by a Morley sequence over a larger model, obtained as a coheir of a reversed copy of the sequence, and to Morley sequences Kim's lemma applies. Throughout the section, T is simple and M is a model.

7.4.1 Two tools

The first tool extends the corollary on extension (see Corollary 7.3.4) from complete types to partial ones. A set of formulas is *closed under conjunctions* if it contains the conjunction of any two of its members.

Lemma 7.4.1. *Let $B \subseteq B'$ be small sets, and let $\pi(x)$ be a set of $L(MB)$ -formulas such that no finite conjunction of members of π divides over M . Then some a realizes π with $a \downarrow_M^d B'$.*

Proof. Arguing as in the proof of Corollary 7.3.4, let $\Sigma(x)$ be the union of π and the set of formulas $\neg\chi(x, c)$ with $\chi(x, z) \in L(M)$, the tuple c from B' , and $\chi(x, c)$ dividing over M . If Σ were inconsistent, some finite conjunction of members of π would imply a finite disjunction of formulas dividing over M (or be inconsistent, and then divide over M); so it would fork over M , and by Theorem 7.3.3 it would divide over M . Hence the set Σ is consistent, and any realization a of Σ satisfies $a \downarrow_M^d B'$. ☺

We also use a consequence of Theorem 7.3.3 that does not depend on the lemma: a formula implied by a formula that does not divide over M does not divide over M either, since forking is closed under implication (see the remarks after Definition 7.3.1) and forking equals dividing over models (see Theorem 7.3.3). The same holds over any model in place of M .

The second tool turns indiscernible sequences around and attaches companions to their terms. We say that tuples d_0, d_1 *begin an indiscernible sequence over M* if there are d_i for $2 \leq i < \omega$ such that $(d_i)_{i < \omega}$ is indiscernible over M .

Lemma 7.4.2. *If d_0, d_1 begin an indiscernible sequence over M , then so do d_1, d_0 .*

Proof. Let $(d_i)_{i < \omega}$ be indiscernible over M . By compactness and saturation there is a sequence $(d'_i)_{i \in \mathbb{Z}}$ such that $(d'_{i_1}, \dots, d'_{i_n}) \equiv_M (d_0, \dots, d_{n-1})$ for every n and all $i_1 < \dots < i_n$ in $\mathbb{Z}$: a finite part of the required type mentions finitely many indices, and an increasing map from them into ω realizes it. Put $e_i = d'_{1-i}$ for $i < \omega$. An increasing tuple of (e_i) is a decreasing tuple of (d'_i) , and every decreasing n -tuple of (d'_i) has the type of $(d_{n-1}, \dots, d_0)$ over M ; hence the sequence $(e_i)_{i < \omega}$ is indiscernible over M , and $(e_0, e_1) \equiv_M (d_1, d_0)$. An automorphism of $\mathbb{M}$ over M carrying (e_0, e_1) to (d_1, d_0) carries (e_i) to the required sequence. ☺

Lemma 7.4.3. *Suppose that finite tuples b_0, b_1 begin an indiscernible sequence over M , and let a be a tuple with $a \downarrow_{Mb_0}^d b_1$. Then there is a^* such that the pairs (b_0, a) and (b_1, a^*) begin an indiscernible sequence over M .*

Proof. Let $(b_i)_{i < \omega}$ be indiscernible over M . The sequence $(b_i)_{1 \leq i < \omega}$ is indiscernible over $C = M \cup b_0$ and begins with b_1 , and $\text{tp}(a/Cb_1)$ does not divide over C . By Lemma 7.1.2 there is $a' \equiv_{Cb_1} a$ such that $(b_i)_{i \geq 1}$ is indiscernible over $C \cup a'$. Applying an automorphism of $\mathbb{M}$ over $C \cup b_1$ that carries a' to a , and replacing b_i by its image for $i \geq 2$, we may assume that $(b_i)_{i \geq 1}$ is indiscernible over $M \cup b_0 \cup a$; the whole sequence remains indiscernible over M .

For $n < \omega$ the sequences $(b_i)_{i < \omega}$ and $(b_{n+i})_{i < \omega}$ have the same type over M , so there is $\sigma_n \in \text{Aut}(\mathbb{M}/M)$ with $\sigma_n(b_i) = b_{n+i}$ for every i . Put $a_n = \sigma_n(a)$. For $n < m$ we have

$$(b_n, a_n, b_m) = \sigma_n(b_0, a, b_{m-n}) \equiv_M (b_0, a, b_{m-n}) \equiv_M (b_0, a, b_1),$$

the last step because $(b_i)_{i \geq 1}$ is indiscernible over $M \cup b_0 \cup a$. Apply Lemma 5.2.2 over M to the sequence of pairs $((b_n, a_n))_{n < \omega}$. Its Ehrenfeucht–Mostowski type over M contains $\text{tp}(b_0, a, b_1/M)$, read as a type of the first coordinate of one pair, the second coordinate of

that pair, and the first coordinate of a later pair. The lemma gives an indiscernible sequence $((b_n^*, a_n^*))_{n < \omega}$ over M with $(b_0^*, a_0^*, b_1^*) \equiv_M (b_0, a, b_1)$. An automorphism of $\mathbb{M}$ over M that carries (b_0^*, a_0^*, b_1^*) to (b_0, a, b_1) carries a_1^* to the required a^* . $\odot$

7.4.2 Conjunctions along indiscernible sequences

If $\psi(x, b)$ does not divide over M and (b_i) is an indiscernible sequence over M beginning with b , the formulas $\psi(x, b_i)$ are consistent, by the definition of dividing. The next proposition says more: finitely many of them together still do not divide. For a Morley sequence this is an application of Kim's lemma to blocks of the sequence (see Step 3 of the proof below). For an arbitrary indiscernible sequence we first replace it by a Morley sequence over a larger model, with the same type over M .

Proposition 7.4.4. *Let $\psi(x, y) \in L(M)$, let b be a finite tuple such that $\psi(x, b)$ does not divide over M , and let $(b_i)_{i < \omega}$ be indiscernible over M with $b_0 = b$. Then for every $n \geq 1$ the formula $\bigwedge_{i < n} \psi(x, b_i)$ does not divide over M .*

Proof. $\odot$ **Step 1: a reversed copy.** By compactness and saturation, as in the proof of Lemma 7.4.2, there is a sequence $(h_j)_{j < \omega}$ such that $(h_{j_0}, \dots, h_{j_{n-1}}) \equiv_M (b_0, \dots, b_{n-1})$ for every n and all $j_0 > j_1 > \dots > j_{n-1}$: the decreasing tuples of (h_j) have the types of the increasing tuples of (b_i) . Let N be a small model containing $M \cup \{h_j : j < \omega\}$.

$\odot$ **Step 2: a Morley sequence over N .** Let $\mathcal{D}$ be a nonprincipal ultrafilter on ω , and let $\mathcal{V}$ be the ultrafilter on $N^{|y|}$ consisting of the sets X with $\{j : h_j \in X\} \in \mathcal{D}$. Let t_0 realize $q_{\mathcal{V}}|\emptyset$, the type over N determined by $\mathcal{V}$, and for $k \geq 1$ let t_k realize $q_{\mathcal{V}}|\{t_l : l < k\}$. By Lemma 7.1.6, applied over the model N , the sequence $(t_k)_{k < \omega}$ is a Morley sequence over N .

We claim that $(t_0, \dots, t_{n-1}) \equiv_M (b_0, \dots, b_{n-1})$ for every n . By induction on n , the definition of the coheir gives, for every $\theta(y_0, \dots, y_{n-1}) \in L(N)$,

$$\mathbb{M} \models \theta(t_0, \dots, t_{n-1}) \iff \left\{ j_{n-1} : \left\{ j_{n-2} : \dots \left\{ j_0 : \mathbb{M} \models \theta(h_{j_0}, \dots, h_{j_{n-1}}) \right\} \in \mathcal{D} \dots \right\} \in \mathcal{D} \right\} \in \mathcal{D}.$$

Indeed, t_{n-1} satisfies $\theta(t_0, \dots, t_{n-2}, y)$ exactly when the set of j with $\mathbb{M} \models \theta(t_0, \dots, t_{n-2}, h_j)$ belongs to $\mathcal{D}$, and for each j the formula $\theta(y_0, \dots, y_{n-2}, h_j)$ has parameters in N . Let $\theta \in L(M)$ be given, and suppose that $\mathbb{M} \models \theta(b_0, \dots, b_{n-1})$; then the formula θ holds at every tuple $(h_{j_0}, \dots, h_{j_{n-1}})$ with $j_0 > j_1 > \dots > j_{n-1}$. The set $\{j : j > j'\}$ belongs to $\mathcal{D}$ for every j' , because $\mathcal{D}$ is nonprincipal. Hence, by induction from the innermost set outward, for all $j_{r+1} > \dots > j_{n-1}$ the set of j_r at the level of nesting r contains $\{j_r : j_r > j_{r+1}\}$, and so belongs to $\mathcal{D}$; at the outermost level we obtain $\mathbb{M} \models \theta(t_0, \dots, t_{n-1})$. Applying this to $\neg\theta$ gives the converse. This proves the claim. In particular $\psi(x, t_0)$ does not divide over M .

$\odot$ **Step 3: over N .** By Lemma 7.4.1 there is a_0 with $\mathbb{M} \models \psi(a_0, t_0)$ and $a_0 \downarrow_M^d N \cup \{t_0\}$. Let $p(x, y)$ be the set of $L(N)$ -formulas $\theta(x, y)$ with $\mathbb{M} \models \theta(a_0, t_0)$, so that $p(x, t_0) = \text{tp}(a_0/Nt_0)$, and write $p(x, t_k)$ for the set of formulas $\theta(x, t_k)$ with $\theta \in p$. By Remark 5.2.5, the type $p(x, t_0)$ does not divide over N . Let $\theta(x, t_0) \in p$ be given. Since $\theta(x, t_0)$ does not divide over N and (t_k) is indiscernible over N with first term t_0 , the set $\{\theta(x, t_k) : k < \omega\}$ is consistent. The blocks $e_m = (t_{mn}, \dots, t_{mn+n-1})$, for $m < \omega$, form an indiscernible sequence over N , and $e_m \downarrow_N^d e_{<m}$: indeed, for $r < n - 1$, if $(t_{mn}, \dots, t_{mn+r}) \downarrow_N^d t_{<mn}$, then $t_{mn+r+1} \downarrow_{N \cup \{t_{mn}, \dots, t_{mn+r}\}}^d t_{<mn}$ by monotonicity and Remark 5.2.5, and Lemma 7.1.7 gives $(t_{mn}, \dots, t_{mn+r+1}) \downarrow_N^d t_{<mn}$; so the sequence (e_m) is a Morley sequence over N . The

formulas $\bigwedge_{k < n} \theta(x, t_{mn+k})$, for $m < \omega$, are consistent together. Hence Kim's lemma, in the form of Theorem 7.2.2(b) applied over N , shows that $\bigwedge_{k < n} \theta(x, t_k)$ does not divide over N .

Therefore no finite conjunction of members of $\bigcup_{k < n} p(x, t_k)$ divides over N : each such conjunction is implied by a formula $\bigwedge_{k < n} \theta(x, t_k)$ with $\theta \in p$, because p is closed under conjunctions. By Lemma 7.4.1, applied over N , some a realizes this union with $a \perp_N^d \{t_k : k < n\}$.

⊙ **Step 4: back to M .** The type $p(x, t_0)$ contains $\text{tp}(a_0/N)$, so we have $a \equiv_N a_0$, and hence we have $a \perp_M^d N$. Transitivity (see Theorem 7.3.6) gives $a \perp_M^d N \cup \{t_k : k < n\}$. The tuple a satisfies $\psi(x, t_k)$ for every $k < n$, because $\psi(x, t_0) \in p$; so the formula $\bigwedge_{k < n} \psi(x, t_k)$ belongs to a type that does not divide over M , and does not divide over M . By the claim of Step 2, the same holds for $\bigwedge_{k < n} \psi(x, b_k)$. ⊙

7.4.3 Moving a parameter

The tools are now in place for the step announced at the start of the section: a parameter is moved from the first term of an indiscernible sequence to the second, by giving the second term a companion (see Lemma 7.4.3) and applying Proposition 7.4.4 to the resulting sequence of pairs.

Proposition 7.4.5. *Let $\varphi(x, z), \psi(x, y) \in L(M)$, and let a, b_0, b_1 be finite tuples such that b_0, b_1 begin an indiscernible sequence over M and $a \perp_{Mb_0}^d b_1$. If $\varphi(x, a) \wedge \psi(x, b_0)$ does not divide over M , then the formula $\varphi(x, a) \wedge \psi(x, b_1)$ does not divide over M .*

Proof. By Lemma 7.4.3 there is a^* such that (b_0, a) and (b_1, a^*) begin an indiscernible sequence over M . Apply Proposition 7.4.4 to the formula $\chi(x; y, z) = \psi(x, y) \wedge \varphi(x, z)$ and this sequence of pairs: the formula $\chi(x; b_0, a) \wedge \chi(x; b_1, a^*)$ does not divide over M . It implies $\varphi(x, a) \wedge \psi(x, b_1)$, which therefore does not divide over M . ⊙

Two tuples with the same type over a model need not begin an indiscernible sequence; a coheir over both of them joins them (see Figure 7.5).

Proposition 7.4.6. *Let $\varphi(x, z), \psi(x, y) \in L(M)$, and let a, b, b' be finite tuples with $b \equiv_M b'$ and $a \perp_M^d bb'$. If $\varphi(x, a) \wedge \psi(x, b')$ does not divide over M , then the formula $\varphi(x, a) \wedge \psi(x, b)$ does not divide over M .*

Proof. Let $\mathcal{U}$ be an ultrafilter on $M^{|b|}$ with $q_{\mathcal{U}}|\emptyset = \text{tp}(b/M) = \text{tp}(b'/M)$ (see Lemma 6.1.5(b)), and let e realize $q_{\mathcal{U}}| \{a, b, b'\}$. Then the tuple e realizes $q_{\mathcal{U}}| \{b'\}$ and $q_{\mathcal{U}}| \{b\}$, so by Lemma 7.1.6 each of the pairs b', e and b, e begins an indiscernible sequence over M , and by Lemma 7.4.2 so does the pair e, b . The type $\text{tp}(e/Mabb')$ is finitely satisfiable in M , so we have $e \perp_M^d abb'$ by Lemma 7.1.3. By symmetry we have $bb' \perp_M^d a$, and $e \perp_{Mbb'}^d a$ by monotonicity and Remark 5.2.5; Lemma 7.1.7 and symmetry give $a \perp_M^d bb'e$. Consequently $a \perp_{Mb'}^d e$ and $a \perp_{Me}^d b$.

Proposition 7.4.5, applied to the pair b', e , shows that $\varphi(x, a) \wedge \psi(x, e)$ does not divide over M ; applied to the pair e, b , it shows that $\varphi(x, a) \wedge \psi(x, b)$ does not divide over M . ⊙

7.4.4 The theorem

Proposition 7.4.6 carries the weight of the theorem. What remains is to find, for given formulas over A and over B , a copy of the parameter from B for which the conjunction is known not to divide, and such that the parameter from A is independent over M from the original and the copy together, as the proposition requires.




FIGURE 7.5. Proposition 7.4.6. A coheir e over $M \cup \{a, b, b'\}$ joins b' to b : the pair b', e begins an indiscernible sequence, and so, after reversal, does the pair e, b . Along each of the two steps Proposition 7.4.5 carries non-dividing of $\varphi(x, a) \wedge \psi(x, \cdot)$ from one end to the other.

Theorem 7.4.7 (Independence theorem). *Assume that T is simple. Let M be a model, let A and B be small sets with $A \downarrow_M^d B$, and let c_1, c_2 be tuples with $c_1 \equiv_M c_2$, $c_1 \downarrow_M^d A$ and $c_2 \downarrow_M^d B$. Then there is c with $c \equiv_{MA} c_1$, $c \equiv_{MB} c_2$ and $c \downarrow_M^d AB$.*

Proof. Let $\Sigma(x)$ be the closure under conjunctions of the union of the types $\text{tp}(c_1/MA)$ and $\text{tp}(c_2/MB)$. By Lemma 7.4.1 it suffices to show that no formula of Σ divides over M . Such a formula is implied by a formula $\varphi(x, a) \wedge \psi(x, b)$ with $\varphi(x, a) \in \text{tp}(c_1/Ma)$, $\psi(x, b) \in \text{tp}(c_2/Mb)$, $\varphi, \psi \in L(M)$, and a and b finite tuples from A and B . Fix such formulas and tuples; by monotonicity, $a \downarrow_M^d b$, $c_1 \downarrow_M^d a$ and $c_2 \downarrow_M^d b$.

Let $\sigma \in \text{Aut}(M/M)$ carry c_2 to c_1 , and put $b_0 = \sigma(b)$. Then we have $c_1 b_0 \equiv_M c_2 b$, and $b_0 \downarrow_M^d c_1$ because $b \downarrow_M^d c_2$ by symmetry. By Corollary 7.3.4 there is $b' \equiv_{Mc_1} b_0$ with $b' \downarrow_M^d c_1 a b$. Thus we have $c_1 b' \equiv_M c_2 b$, and in particular $b' \equiv_M b$.

We have $a \downarrow_M^d c_1$ by symmetry, and $b' \downarrow_{Ma}^d c_1$ by monotonicity and Remark 5.2.5, so Lemma 7.1.7 and symmetry give $c_1 \downarrow_M^d a b'$. Hence the formula $\varphi(x, a) \wedge \psi(x, b')$, which belongs to $\text{tp}(c_1/Mab')$ because $c_1 b' \equiv_M c_2 b$, does not divide over M . Similarly, $b \downarrow_M^d a$ and $b' \downarrow_{Mb}^d a$ give $b b' \downarrow_M^d a$, that is, $a \downarrow_M^d b b'$. By Proposition 7.4.6, the formula $\varphi(x, a) \wedge \psi(x, b)$ does not divide over M . ☺

In the Rado graph, Theorem 5.4.9(c) is the analogue for trivial independence, over arbitrary small sets, and Theorem 5.5.9 proves the corresponding statement for Fraïssé limits with disjoint 3-amalgamation. With Theorem 7.4.7, nondividing independence over models has every property of Definition 5.4.7 except uniqueness, together with the independence theorem. In Section 7.5 we show that uniqueness, added to these properties, gives exactly the stable theories, and in Section 7.6 that these properties, suitably strengthened, characterize the simple theories.

7.5 STABILITY THROUGH FORKING

A simple theory has, over a model, a notion of independence with every property of Definition 5.4.7 except uniqueness. In Theorem 7.5.2 below we show that uniqueness adds exactly

stability. One direction is the theorem of Chapter 6, for uniqueness together with local character already forces stability. The other direction is a short argument with the independence theorem (see Lemma 7.5.1). If a type over a model had two different nondividing extensions, the independence theorem would let us realize them in any pattern we like along an independent sequence, and a formula that distinguishes the two extensions would then shatter the sequence. Uniqueness fails, in other words, only in the presence of the independence property. Throughout the section, M is a model.

Lemma 7.5.1. *Assume that T is simple. Let a, a' be tuples with $a \equiv_M a'$, and let B be a small set with $a \downarrow_M^d B$ and $a' \downarrow_M^d B$. If $a \not\equiv_{MB} a'$, then some formula has the independence property.*

Proof. Choose $\varphi(x, y) \in L(M)$, with x a finite tuple of variables, and a finite tuple b from B with $M \models \varphi(a, b) \wedge \neg\varphi(a', b)$. By monotonicity, $a \downarrow_M^d b$ and $a' \downarrow_M^d b$. Let $(b_i)_{i < \omega}$ be a Morley sequence over M with $b_0 = b$ (see Lemma 7.1.6); then we have $b_i \equiv_M b$ and $b_i \downarrow_M^d b_{<i}$ for every i .

Let $n \geq 1$ and $S \subseteq \{0, \dots, n-1\}$ be given. We show, by induction on $k \leq n$, that there is c_k with $c_k \equiv_M a$, $c_k \downarrow_M^d b_{<k}$, and $M \models \varphi(c_k, b_i)$ if and only if $i \in S$, for every $i < k$. For $k = 0$ take $c_0 = a$. Suppose that c_k has been found, with $k < n$. Let $\sigma \in \text{Aut}(M/M)$ carry b to b_k , and put $d = \sigma(a)$ if $k \in S$ and $d = \sigma(a')$ if $k \notin S$. Then we have $d \equiv_M a \equiv_M c_k$, and $d \downarrow_M^d b_k$ by invariance; moreover, $\varphi(d, b_k)$ holds if and only if $k \in S$. By symmetry, $b_{<k} \downarrow_M^d b_k$. The independence theorem (see Theorem 7.4.7), applied to $A = b_{<k}$, $B = b_k$, $c_1 = c_k$ and $c_2 = d$, gives c_{k+1} with $c_{k+1} \equiv_{Mb_{<k}} c_k$, $c_{k+1} \equiv_{Mb_k} d$ and $c_{k+1} \downarrow_M^d b_{<k+1}$, as required.

The tuple c_n satisfies $\varphi(x, b_i)$ exactly for $i \in S$. Write $\varphi(x, y) = \varphi_0(x, y, m)$ with $\varphi_0 \in L$ and m a finite tuple from M . For every n , the tuples $(b_0, m), \dots, (b_{n-1}, m)$ are shattered by $\varphi_0(x; y, z)$, so the formula φ_0 has the independence property. 

Theorem 7.5.2 (Stability through forking). *For a complete theory T , the following are equivalent:*

- (1) *The theory T is stable.*
- (2) *The theory T is simple, and the relation $\downarrow^d$, with models as bases, satisfies (UN).*
- (3) *The relation $\downarrow^d$, with models as bases, satisfies (UN) and (LOC).*

Moreover, if T is stable, then the relation $\downarrow^d$ coincides, over models, with the heir independence $\downarrow^h$ of Definition 6.2.1.

Proof. (1) $\Rightarrow$ (2). By Proposition 5.3.1, the theory T is simple. If (UN) failed for $\downarrow^d$, Lemma 7.5.1 would give a formula with the independence property, and by Proposition 4.1.2 that formula would be unstable.

(2) $\Rightarrow$ (3). In a simple theory, $\downarrow^d$ satisfies (LOC) (see Proposition 7.1.8).

(3) $\Rightarrow$ (1). This is the implication (3) $\Rightarrow$ (1) of Theorem 6.2.2.

For the last statement, assume that T is stable, and let A and B be small sets. Heirs do not divide: if $\varphi(x, b)$ belongs to the heir $\text{tp}(A/M)|B$, with $\varphi(x, y) \in L(M)$ and x a finite subtuple of the variables of A , and $(b_i)_{i < \omega}$ is indiscernible over M with $b_0 = b$, then we have $M \models d_p \varphi(b_i)$ for every i , where $p = \text{tp}(A/M)$, because $d_p \varphi$ has parameters in M ; hence every $\varphi(x, b_i)$ belongs to the heir $p| \{b_i : i < \omega\}$, which is consistent by Lemma 6.1.3. Thus the

condition $A \downarrow_M^h B$ implies $A \downarrow_M^d B$. Conversely, suppose that $A \downarrow_M^d B$, and let A' realize $p|B$. Then we have $A' \equiv_M A$ and $A' \downarrow_M^h B$, so we have $A' \downarrow_M^d B$; by (UN) for $\downarrow^d$, which holds by (2), we have $A' \equiv_{MB} A$, and so we have $A \downarrow_M^h B$. $\odot$

The theorem also gives a second proof of the harder half of Theorem 5.3.4, one that does not pass through Shelah's dichotomy (see Fact 5.3.3). If T is simple and every formula is NIP, then Lemma 7.5.1 shows that $\downarrow^d$ satisfies (UN), and the theory T is stable by the implication (2) $\Rightarrow$ (1). The two proofs divide the work differently. Shelah's dichotomy finds, in an unstable theory, either an independent configuration or an infinite order; here the strict order is excluded by simplicity at the outset, through the independence theorem, and the independent configuration is built by hand from two incompatible extensions of one type. In the random graph the construction is visible directly: two new vertices with the same type over M , one joined to a vertex $b \notin M$ and one not, are both independent from b , and the independence theorem along a sequence of vertices produces the shattering of Example 4.1.4.

7.6 THE KIM-PILLAY THEOREM

In Chapter 6 we characterized the stable theories by the existence of an independence relation with uniqueness. Kim and Pillay [KP97] did the same for the simple theories, with the independence theorem in place of uniqueness. Their theorem says that simplicity is not one property among several that a good notion of independence might have: it is what the existence of such a notion amounts to. Existence must be taken in the stronger form of extension, since uniqueness, which made the weak form sufficient in Definition 5.4.7, is no longer available; and the independence theorem takes the place of uniqueness.

Definition 7.6.1. Let $\downarrow$ be a relation $A \downarrow_M B$ between small sets A, B and models M .

(EXT) *Extension.* If $A \downarrow_M B$ and $B \subseteq B'$, then there is $A' \equiv_{MB} A$ with $A' \downarrow_M B'$.

(IT) *Independence theorem.* If $A \downarrow_M B$, $c_1 \equiv_M c_2$, $c_1 \downarrow_M A$ and $c_2 \downarrow_M B$, then there is c with $c \equiv_{MA} c_1$, $c \equiv_{MB} c_2$ and $c \downarrow_M AB$.

A *Kim–Pillay relation* is a relation satisfying (INV), (MON), (BMON), (TR), (FC), (SYM), (EX) and (LOC) of Definition 5.4.7, together with (EXT) and (IT).

In a simple theory, nondividing independence is a Kim–Pillay relation: Theorem 7.3.6 gives the axioms of Definition 5.4.7 other than uniqueness, Corollary 7.3.4 gives (EXT), and Theorem 7.4.7 gives (IT). The converse needs two tools. The first, Proposition 7.6.2, produces from an indiscernible sequence an independent one with the same type; the second, Proposition 7.6.3, turns the tree property into a type that divides over every small submodel of a model.

7.6.1 Independent copies of indiscernible sequences

Local character, applied along a long sequence, yields independence over a model; Fodor's lemma on regressive functions supplies the model. To keep the type of the sequence, we build it, as in the proof of Proposition 7.4.4, as a coheir of a reversed copy.

Proposition 7.6.2. Let $\downarrow$ be a relation over models satisfying (MON), (BMON) and (LOC). Let M be a model, and let $(b_i)_{i < \omega}$ be a sequence of finite tuples, indiscernible over M ; let y be a tuple of variables of the length of b_0 . Then there are a model $N \supseteq M$ and a sequence $(e_i)_{i < \omega}$,

indiscernible over N , such that $(e_0, \dots, e_{n-1}) \equiv_M (b_0, \dots, b_{n-1})$ for every n and $e_i \perp_N e_{<i}$ for every i .

Proof. **Step 1: a long coheir sequence.** As in Steps 1 and 2 of the proof of Proposition 7.4.4, let $(h_j)_{j<\omega}$ be a sequence whose decreasing tuples have the types over M of the increasing tuples of (b_i) , let N_0 be a small model containing $M \cup \{h_j : j < \omega\}$, let $\mathcal{D}$ be a nonprincipal ultrafilter on ω , and let $\mathcal{V}$ be the ultrafilter on $N_0^{|y|}$ induced by $j \mapsto h_j$. Let $\lambda = (|N_0| + \kappa^+)^+$, a regular cardinal, and choose by recursion on $\alpha < \lambda$ small models N_α and tuples t_α : the models form an increasing chain, continuous at limits, with $N_\alpha \supseteq N_0 \cup \{t_\gamma : \gamma < \alpha\}$ and $|N_\alpha| \leq |N_0| + |\alpha| + \kappa$, and t_α realizes $q_{\mathcal{V}}|N_\alpha$.

For $\beta < \lambda$, all $\beta \leq \alpha_0 < \dots < \alpha_{n-1}$ and every $\theta \in L(N_\beta)$, the computation of Step 2 of the proof of Proposition 7.4.4 gives

$$\mathbb{M} \models \theta(t_{\alpha_0}, \dots, t_{\alpha_{n-1}}) \iff \left\{ j_{n-1} : \dots \left\{ j_0 : \mathbb{M} \models \theta(h_{j_0}, \dots, h_{j_{n-1}}) \right\} \in \mathcal{D} \dots \right\} \in \mathcal{D},$$

since t_{α_r} realizes the coheir of $\mathcal{V}$ over a model containing N_β and $t_{\alpha_0}, \dots, t_{\alpha_{r-1}}$. The right side does not depend on $\alpha_0 < \dots < \alpha_{n-1}$. Hence the sequence $(t_\alpha)_{\beta \leq \alpha < \lambda}$ is indiscernible over N_β , and, as there, $(t_{\alpha_0}, \dots, t_{\alpha_{n-1}}) \equiv_M (b_0, \dots, b_{n-1})$ for all $\alpha_0 < \dots < \alpha_{n-1}$.

Step 2: local character and Fodor's lemma. For each $\alpha < \lambda$, (LOC) gives a model $M_\alpha \subseteq N_\alpha$ with $|M_\alpha| \leq \kappa$ and $t_\alpha \perp_{M_\alpha} N_\alpha$. Let E be the set of $\alpha < \lambda$ of cofinality κ^+ , a stationary subset of λ . For $\alpha \in E$ we have $N_\alpha = \bigcup_{\gamma < \alpha} N_\gamma$, and $|M_\alpha| < \text{cf } \alpha$, so we have $M_\alpha \subseteq N_{f(\alpha)}$ for some $f(\alpha) < \alpha$. By Fodor's lemma, the function f is constant, say with value β , on a stationary set $S \subseteq E$. For $\alpha \in S$ we have $M_\alpha \subseteq N_\beta \subseteq N_\alpha$, so (BMON) gives $t_\alpha \perp_{N_\beta} N_\alpha$, and (MON) gives $t_\alpha \perp_{N_\beta} \{t_\gamma : \gamma \in S, \gamma < \alpha\}$.

Step 3. Let $\alpha_0 < \alpha_1 < \dots$ be the first ω elements of S above β , and put $N = N_\beta$ and $e_i = t_{\alpha_i}$. By Step 1 the sequence (e_i) is indiscernible over N and has the required type over M , and by Step 2 we have $e_i \perp_N e_{<i}$. ☺

7.6.2 Trees give dividing over every small submodel

The second tool is the converse of the local character of dividing (see Theorem 5.3.6). It grows a branch of length κ^+ through the tree property, one step at a time, keeping above the branch a whole tree with which the branch is consistent, and makes each new formula divide over everything chosen before it.

Proposition 7.6.3. *Suppose that some formula $\varphi(x, y)$ has the k -tree property. Then there are a model N and a finite tuple a such that $\text{tp}(a/N)$ divides over every model $M \subseteq N$ with $|M| \leq \kappa$.*

Proof. By recursion on $\alpha \leq \kappa^+$ we choose models N_α with $|N_\alpha| \leq \kappa^+$, forming an increasing chain continuous at limits, and tuples $b_\gamma \in N_{\gamma+1}$ for $\gamma < \alpha$, such that, with $\Sigma_\alpha = \{\varphi(x, b_\gamma) : \gamma < \alpha\}$:

- (i) for $\gamma < \alpha$, the formula $\varphi(x, b_\gamma)$ k -divides over N_γ ;
- (ii) there are tuples c_σ , for $\sigma \in \omega^{<\omega}$ with $\sigma \neq \emptyset$, satisfying condition (b) of Definition 5.2.7, such that for every $\eta \in \omega^\omega$ the set $\Sigma_\alpha \cup \{\varphi(x, c_{\eta|n}) : n \geq 1\}$ is consistent.

Condition (ii) implies that Σ_α is consistent. For $\alpha = 0$, let N_0 be any model of size κ ; condition (ii) is the k -tree property.

Successor step. Let (c_σ) witness (ii) for α . For $i < \omega$ let s_i be the tuple $(c_{\langle i \rangle \smallfrown \tau})_{\tau \in \omega^{<\omega}}$, listing the subtree above the node $\langle i \rangle$, with $c_{\langle i \rangle}$ as its first entry. Apply Lemma 5.2.2 over

N_α to the sequence $(s_i)_{i < \omega}$; it gives a sequence $(s'_i)_{i < \omega}$, indiscernible over N_α , realizing the Ehrenfeucht–Mostowski type of (s_i) over N_α . That type contains the following formulas, because they hold of all the tuples, or all the increasing k -tuples, of (s_i) :

- $\neg \exists x \bigwedge_{l < k} \varphi(x, z_l)$, where z_l is the first entry of the l -th tuple;
- for each finite $\Sigma_0 \subseteq \Sigma_\alpha$, each n and each path ρ of length n in the subtree, the formula saying that Σ_0 , $\varphi(x, z)$ and the formulas at the nodes of ρ are consistent, where z is the first entry;
- the formulas saying that siblings within one subtree are k -inconsistent.

Let b_α be the first entry of s'_0 , and let $N_{\alpha+1}$ be a model of size $|N_\alpha|$ containing $N_\alpha \cup b_\alpha$. The first entries of the s'_i form a sequence indiscernible over N_α , beginning with b_α , along which φ is k -inconsistent, so the formula $\varphi(x, b_\alpha)$ k -divides over N_α . The remaining entries of s'_0 , indexed by $\tau \neq \emptyset$, satisfy condition (b) of Definition 5.2.7, and by compactness every branch through them is consistent with $\Sigma_\alpha \cup \{\varphi(x, b_\alpha)\}$. This is (ii) for $\alpha + 1$.

Limit step. Let $\delta \leq \kappa^+$ be a limit, and put $N_\delta = \bigcup_{\alpha < \delta} N_\alpha$. Condition (ii) for δ is expressed by a set of formulas in variables z_σ , with parameters among the b_γ ; a finite part of it mentions only Σ_α for some $\alpha < \delta$, and is satisfied by the tree given by (ii) for α . By compactness and saturation, (ii) holds for δ , and (i) holds by the induction hypothesis.

Put $N = N_{\kappa^+}$, and let a realize Σ_{κ^+} . Let $M \subseteq N$ be a model with $|M| \leq \kappa$. Since κ^+ is regular, we have $M \subseteq N_\alpha$ for some $\alpha < \kappa^+$. The formula $\varphi(x, b_\alpha)$ belongs to $\text{tp}(a/N)$ and k -divides over N_α , hence over M (see Remark 5.2.5). ☺

Proposition 7.6.3 gives the converse of Theorem 5.3.6 promised in Remark 5.3.7, in the form with models.

Corollary 7.6.4. *A complete theory T is simple if and only if for every finite tuple a and every model N there is a model $M \subseteq N$ with $|M| \leq \kappa$ such that $\text{tp}(a/N)$ does not divide over M .*

Proof. If T is simple, this is (LOC) in Proposition 7.1.8. If T is not simple, some formula has the k -tree property for some k , and Proposition 7.6.3 gives a and N for which no such M exists. ☹

7.6.3 The theorem

Theorem 7.6.5 (Kim–Pillay). *A complete theory T is simple if and only if it carries a Kim–Pillay relation. Moreover, nondividing independence is then a Kim–Pillay relation, and every Kim–Pillay relation $\perp$ satisfies: if $A \perp_M B$, then we have $A \perp_M^d B$.*

Proof. If T is simple, nondividing independence is a Kim–Pillay relation, as observed after Definition 7.6.1. Conversely, let $\perp$ be a Kim–Pillay relation.

☺ **Step 1: the relation $\perp$ implies nondividing.** Let $a \perp_M b$ with a and b finite, let $p(x, y)$ be the set of $L(M)$ -formulas $\psi(x, y)$ with $\mathbb{M} \models \psi(a, b)$, and let $(b_i)_{i < \omega}$ be indiscernible over M with $b_0 = b$. We show that $\bigcup_{i < \omega} p(x, b_i)$ is consistent; then no formula of $\text{tp}(a/Mb)$ divides over M , and by (FC) this proves the claim for arbitrary A and B . By compactness it suffices to realize $\bigcup_{i < n} p(x, b_i)$ for each n .

Let N and (e_i) be as in Proposition 7.6.2. Since $e_0 \equiv_M b$, there is $\sigma \in \text{Aut}(\mathbb{M}/M)$ with $\sigma(b) = e_0$, and $\sigma(a) \perp_M e_0$ by (INV). By (EXT) there is $a^* \equiv_{Me_0} \sigma(a)$ with $a^* \perp_M N \cup e_0$, and (BMON) and (MON) give $a^* \perp_N e_0$. Let $q(x, e_0) = \text{tp}(a^*/Ne_0)$; it contains $p(x, e_0)$. The sequence (e_i) is indiscernible over N , so for each i an automorphism of $\mathbb{M}$ over N carries e_0 to e_i and a^* to some d_i with $d_i \equiv_N a^*$, $d_i \perp_N e_i$, and d_i realizing $q(x, e_i)$.

We show by induction on $m \leq n$ that some c_m realizes $\bigcup_{i < m} q(x, e_i)$ with $c_m \equiv_N a^*$ and $c_m \perp_N e_{<m}$. For $m = 0$ take $c_0 = a^*$. Given c_m , we have $e_{<m} \perp_N e_m$ by (SYM), and (IT), applied over N to $A = e_{<m}$, $B = e_m$, $c_1 = c_m$ and $c_2 = d_m$, gives c_{m+1} as required. Thus the type $\bigcup_{i < n} p(x, e_i)$ is realized, and since $(e_0, \dots, e_{n-1}) \equiv_M (b_0, \dots, b_{n-1})$, so is $\bigcup_{i < n} p(x, b_i)$.

© **Step 2: the theory T is simple.** Let a be a finite tuple and N a model. By (LOC) there is a model $M \subseteq N$ with $|M| \leq \kappa$ and $a \perp_M N$, and by Step 1 the type $\text{tp}(a/N)$ does not divide over M . By Corollary 7.6.4, the theory T is simple. ☺

The theorem identifies nondividing independence as the weakest Kim–Pillay relation. It is in fact the only one: Kim and Pillay [KP97] show, for the analogous relations over arbitrary base sets, that such a relation coincides with nonforking independence. In a stable theory this follows at once from what we have proved. If $\perp$ is a Kim–Pillay relation and $A \perp_M^d B$, let $A' \equiv_M A$ satisfy $A' \perp_M B$, by (EX). Then we have $A' \perp_M^d B$ by Theorem 7.6.5, and uniqueness for $\perp^d$ (see Theorem 7.5.2) gives $A' \equiv_{MB} A$, so that $A \perp_M B$ by (INV).

7.7 EXERCISES

The exercises follow the sections of the chapter. Throughout, M is a model, and we write T_{RG} for the theory of the random graph (see Example 5.4.1). The mark (*) indicates the exercises that ask for more than a direct application of the text.

1. (Section 7.1.) Let b be an element with $b \notin M$. Show that the formula $x = b$ divides over M . Deduce that, in every complete theory, the relation $\perp^d$ satisfies anti-reflexivity: if $a \perp_M^d a$ for a finite tuple a , then we have $a \subseteq M$. Deduce also that $a \perp_M^d B$ implies $a \cap B \subseteq M$. (Hint: a model is algebraically closed, so the type $\text{tp}(b/M)$ has infinitely many realizations; use Lemma 5.2.2.)
2. (Section 7.1.) Work in T_{RG} . Let a and b be vertices outside M such that a is adjacent to every vertex of M and to b , while b is adjacent to no vertex of M . Show that such vertices exist, that $\text{tp}(b/Ma)$ is finitely satisfiable in M , and that $\text{tp}(a/Mb)$ is not. Conclude that *coheir independence*, the relation “ $\text{tp}(A/MB)$ is finitely satisfiable in M ”, is not symmetric, even in a simple theory, and that $\text{tp}(a/Mb)$ is a type that does not divide over M (see Theorem 5.4.5) but is not finitely satisfiable in M . Thus the converse of Lemma 7.1.3 fails.
3. (Section 7.2.) Assume that T is simple, let b be a finite tuple, and let $\varphi(x, y), \psi(x, y) \in L(M)$ be given. Show that $\varphi(x, b) \vee \psi(x, b)$ divides over M if and only if both $\varphi(x, b)$ and $\psi(x, b)$ divide over M . (Hint: for the nontrivial direction, fix one Morley sequence $(b_i)_{i < \omega}$ over M with $b_0 = b$. By Kim’s lemma (see Theorem 7.2.2), the instances of φ and of ψ along it are inconsistent, hence k -inconsistent and l -inconsistent for some k and l (see Lemma 5.2.4); show that the instances of the disjunction are $(k+l-1)$ -inconsistent.) Identify where the argument uses simplicity.
4. (Section 7.2.) In Example 7.2.3, the formula $\varphi(x, b)$ divides over M along some Morley sequence over M and not along another. Show that the formula $x = d_1$, with d_1 as in that example, divides along every indiscernible sequence over M that begins with d_1 and is not constant. Show that no Morley sequence over M beginning with d_1 is constant

(see Exercise 1), so that, in contrast with $\varphi(x, b)$, the formula $x = d_1$ divides along every Morley sequence over M beginning with d_1 .


5. (Section 7.3.) Let $\varphi(x, b)$ be a formula over $M \cup b$ that is satisfied by some tuple from M . Show that $\varphi(x, b)$ does not fork over M , in every complete theory. Deduce that in the dense circular order of Remark 7.3.2 the formula $x = x$ forks over $\emptyset$ but over no model.
6. (Section 7.3.) Symmetry does not follow from invariance, monotonicity, finite character and existence alone. Using Exercise 2, show that the relation of coheir independence over models satisfies (INV), (MON), (FC) and (EX) in every complete theory (see Proposition 7.1.4) and fails (SYM) in T_{RG} .
7. (Section 7.4.) Work in the theory of $(\mathbb{Q}, <)$. Let $a < b$ be elements greater than every element of M , and let c_1 and c_2 be elements with $M < c_1 < a$ and $b < c_2$. Show that $a \perp_M^d b$, $c_1 \equiv_M c_2$, $c_1 \perp_M^d a$ and $c_2 \perp_M^d b$, and that no c satisfies both $c \equiv_{Ma} c_1$ and $c \equiv_{Mb} c_2$. Thus the independence theorem fails for $\perp^d$ in this theory. Compare Theorem 7.6.5: since the theory is not simple, the relation $\perp^d$ is not a Kim–Pillay relation, but the theorem does not say which axiom fails; local character fails as well (see Exercise 13). (Hint: a finite conjunction of formulas from $\text{tp}(a/Mb)$ remains consistent when b is replaced by finitely many terms of an indiscernible sequence over M .)
8. (Section 7.4.) Assume that T is stable. Show that the independence theorem holds for $\perp^d$ over models without the hypothesis $A \perp_M^d B$: if $c_1 \equiv_M c_2$, $c_1 \perp_M^d A$ and $c_2 \perp_M^d B$, there is c with $c \equiv_{MA} c_1$, $c \equiv_{MB} c_2$ and $c \perp_M^d AB$. (Hint: realize a nondividing extension of $\text{tp}(c_1/M)$ to $M \cup A \cup B$, and use (MON) and (UN); see Theorem 7.5.2.) Show by an example in T_{RG} that the hypothesis $A \perp_M^d B$ cannot be dropped in a simple theory.
9. (*) (Section 7.5.) Assume that T is stable, and let C be a small set. Show that every sequence $(b_i)_{i < \omega}$ indiscernible over C is an indiscernible set over C : for all n , all distinct $i_1, \dots, i_n$ and all distinct $j_1, \dots, j_n$, the tuples $(b_{i_1}, \dots, b_{i_n})$ and $(b_{j_1}, \dots, b_{j_n})$ have the same type over C . (Hint: by compactness, index the sequence by $\mathbb{Q}$. If a transposition of two adjacent positions changes the truth value of a formula, fix the other terms and obtain a formula that orders an infinite subsequence; see Definition 1.1.4.) Deduce that every Morley sequence over a model M of a stable theory is an indiscernible set over M .
10. (Section 7.5.) Work in T_{RG} . Find vertices a, a' and b , not in M , with $a \equiv_M a'$, $a \perp_M^d b$, $a' \perp_M^d b$ and $a \not\equiv_{Mb} a'$. Follow the proof of Lemma 7.5.1 for the formula xEy , and compare the tuples c_n that it produces with those given directly by the extension axioms.
11. (*) (Section 7.6.) Work in T_{RG} , with trivial independence $\perp^t$ (see Definition 5.4.8) and models as bases. Show that $\perp^t$ satisfies (EXT), and that Theorem 5.4.9(c), together with compactness, gives (IT) for small sets. Conclude from Theorem 7.6.5 a second proof that T_{RG} is simple (see Theorem 5.4.6), and a second proof of one direction of Theorem 5.4.5 over models.
12. (Section 7.6.) Show that every Kim–Pillay relation $\perp$ satisfies anti-reflexivity: if $a \perp_M a$ for a finite tuple a , then we have $a \subseteq M$. (Hint: Theorem 7.6.5 and Exercise 1.)

13. (*) (Section 7.6.) Work in the theory of $(\mathbb{Q}, <)$, so that $\kappa = |T| = \aleph_0$ in Corollary 7.6.4. Let N be a model, and let a realize a cut (L, R) of N such that no subset of L of size at most κ is cofinal in L and no subset of R of size at most κ is coinitial in R . Show that such N and a exist, and that $\text{tp}(a/N)$ divides over every model $M \subseteq N$ with $|M| \leq \kappa$. (Hint: find $l \in L$ and $r \in R$ with no element of M between them, and argue as in Example 7.2.3.) Conclude from Corollary 7.6.4 that the theory is not simple.

7.8 HISTORICAL REMARKS

- ✂ **Section 7.1.** Dividing and forking were introduced by Shelah in the first edition of *Classification Theory* [She90], where they serve stable theories. Dividing is the combinatorial notion; forking is its closure under finite disjunctions, designed so that nonforking extensions always exist. Indiscernible sequences go back to Ehrenfeucht and Mostowski [EM56], who extracted them with Ramsey's theorem, as in Lemma 5.2.2. Morley [Mor65] made them a central tool of the subject, and sequences whose terms are independent from the earlier ones bear his name. That finitely satisfiable types do not divide, and the use of coheirs to produce such sequences, come from the work of Lascar and Poizat [LP79] on forking in stable theories. The characterization of nondividing by indiscernible sequences (see Lemma 7.1.2) is standard; we follow the presentation of Grossberg, Iovino and Lessmann [GIL02].
- ✂ **Section 7.2.** Simple theories were introduced by Shelah [She80], who showed that local character holds in them. Kim [Kim98] proved that in a simple theory a formula divides exactly when it divides along some, equivalently every, Morley sequence, and deduced that forking and dividing agree and that independence is symmetric and transitive. In a later paper [Kim01] he showed that each of symmetry and transitivity of forking characterizes the simple theories.
- ✂ **Section 7.3.** Forking and dividing differ in general, and the dense circular order is the usual example (see [TZ12, Chapter 7]). The derivation of symmetry from Kim's lemma, by a sequence of copies of a that a single tuple b cannot follow, goes back to Kim [Kim98].
- ✂ **Section 7.4.** The independence theorem is due to Kim and Pillay [KP97], who abstracted it from Hrushovski's work on pseudofinite fields [Hru02]. They proved it for Lascar strong types over arbitrary sets; over a model these are types. In that form it appears in the monographs of Wagner [Wag00] and Casanovas [Cas11], in Casanovas's notes [Cas07], and in the primer of Grossberg, Iovino and Lessmann [GIL02]. Our proof follows the route of Casanovas: a parameter is moved along indiscernible sequences (see Propositions 7.4.5 and 7.4.6), and the moving rests on the fact that conjunctions along an indiscernible sequence do not divide (see Proposition 7.4.4). Casanovas proves that fact over a set that is not a model, using symmetry and transitivity of forking over arbitrary sets. To stay with models, we have replaced that step by an argument specific to this book: the indiscernible sequence is traded for a Morley sequence over a larger model, obtained as a coheir of a reversed copy of the sequence, to which Kim's lemma applies directly.
- ✂ **Section 7.5.** That a theory is stable exactly when nonforking extensions of types over models are unique goes back to Shelah [She90]. The derivation of the independence property from two distinct nonforking extensions by means of the independence theorem, and the resulting proof that simple NIP theories are stable, belong to the folklore of simple theories (see, for instance, Casanovas [Cas11]).

✂ **Section 7.6.** The characterization of simple theories by an abstract independence relation is due to Kim and Pillay [KP97], who also proved that such a relation must be nonforking independence. Corollary 7.6.4, the converse of local character, is Shelah’s [She80]; the construction of a long chain of dividing formulas from the tree property, keeping a tree above the chain, is the classical one. Our proof of Step 1 of Theorem 7.6.5 uses the reversed copy of Section 7.4 together with Fodor’s lemma, so as to stay with models as bases.



Appendix A

Grothendieck's double limit criterion

Grothendieck's double limit criterion is used at several points in this book: in Subsection 1.2.4 and in Remarks 1.2.14 and 4.2.6. In this appendix we state the criterion in the form we need and explain how the criterion relates to the stability of functions (see Definition 1.2.5). Throughout, we fix a compact Hausdorff space K , and $C(K)$ denotes the Banach space of continuous real-valued functions on K with the supremum norm. The *weak topology* on $C(K)$ is the coarsest topology making every bounded linear functional continuous, and the *pointwise topology* is the one induced by the product topology of $\mathbb{R}^K$. Point evaluations $f \mapsto f(x)$ are bounded linear functionals, so the weak topology is finer than the pointwise topology.

Roadmap

Grothendieck's criterion is the bridge by which the exchange of limits, a condition on numbers, becomes compactness, a condition on sets of functions. In this appendix we cross the bridge once, in the direction this book requires. In Appendix A.1 we state the double limit condition, which is precisely the stability of the evaluation function, and show that the ultrafilter form used in this book agrees with Grothendieck's form with ordinary limits (see Lemma A.1.2). We then prove the criterion: a bounded set of continuous functions on a compact space is relatively weakly compact exactly when it satisfies the double limit condition (see Theorem A.1.3). With the theorem of Eberlein and Šmulian, the criterion yields convergent subsequences (see Corollary A.1.6), and, for a stable function, it shows that the spaces of types are Eberlein compacta, so that every type is the limit of a sequence of realized types (see Corollary A.1.8). The last statement deserves to be read twice. It says that, when a function is stable, every ideal element is the limit of a sequence of actual ones.

A.1 THE CRITERION

Definition A.1.1. A bounded set $H \subseteq C(K)$ satisfies the *double limit condition* if for all sequences (x_m) in K and (f_n) in H and all ultrafilters $\mathcal{U}, \mathcal{V}$ on $\mathbb{N}$, we have

$$\lim_{m, \mathcal{U}} \lim_{n, \mathcal{V}} f_n(x_m) = \lim_{n, \mathcal{V}} \lim_{m, \mathcal{U}} f_n(x_m).$$

In the language of Definition 1.2.5, the double limit condition says that the evaluation function $K \times H \rightarrow \mathbb{R}$, $(x, f) \mapsto f(x)$, is stable. Grothendieck stated the condition with ordinary limits; the next lemma shows that the two formulations are equivalent.

Lemma A.1.2. *A bounded set $H \subseteq C(K)$ satisfies the double limit condition if and only if*

$$\lim_m \lim_n f_n(x_m) = \lim_n \lim_m f_n(x_m)$$

for all sequences (x_m) in K and (f_n) in H for which both iterated limits exist.

Proof. If both ordinary iterated limits exist, these limits are equal to the iterated limits along any nonprincipal ultrafilters, so the double limit condition implies the second condition.

Conversely, suppose the double limit condition fails for (x_m) , (f_n) , $\mathcal{U}$ and $\mathcal{V}$, with iterated limits $\alpha \neq \beta$. After an affine change of scale we may assume that all values lie in $[0, 1]$. Proposition 1.2.3, applied to the function $(x, f) \mapsto f(x)$, gives subsequences, which we still denote (x_m) and (f_n) , with $|f_n(x_m) - \alpha| \leq 2^{1-m}$ for $m < n$ and $|f_n(x_m) - \beta| \leq 2^{1-n}$ for $n < m$. By a diagonal argument we may pass to further subsequences along which $\lim_n f_n(x_m)$ exists for every m and $\lim_m f_n(x_m)$ exists for every n ; the inequalities are preserved, with the indices renumbered. Then the limit $\lim_n f_n(x_m)$ is within 2^{1-m} of α , so we have $\lim_m \lim_n f_n(x_m) = \alpha$, and symmetrically $\lim_n \lim_m f_n(x_m) = \beta$. Hence the second condition fails. ☹

Theorem A.1.3 (Grothendieck). *Let $H \subseteq C(K)$ be bounded. The following are equivalent:*

- (1) *The set H is relatively weakly compact in $C(K)$.*
- (2) *Every function in the pointwise closure of H in $\mathbb{R}^K$ is continuous.*
- (3) *The set H satisfies the double limit condition.*

The equivalence of (2) and (3) is an exchange of ultrafilter limits, together with Proposition 1.2.3; the measure-theoretic input and the angelicity of pointwise compact sets enter only in (2) $\Rightarrow$ (1).

Proof. (2) $\Rightarrow$ (3). Let (x_m) , (f_n) , $\mathcal{U}$, $\mathcal{V}$ be given, and put $x = \lim_{m, \mathcal{U}} x_m \in K$ and $f = \lim_{n, \mathcal{V}} f_n \in \mathbb{R}^K$. By (2), the function f is continuous. Using the continuity of f and of each f_n , we obtain

$$\lim_{m, \mathcal{U}} \lim_{n, \mathcal{V}} f_n(x_m) = \lim_{m, \mathcal{U}} f(x_m) = f(x) = \lim_{n, \mathcal{V}} f_n(x) = \lim_{n, \mathcal{V}} \lim_{m, \mathcal{U}} f_n(x_m).$$

(3) $\Rightarrow$ (2). Suppose that f is in the pointwise closure of H but is not continuous at some $x \in K$. Then we have $f = \lim_{i, \mathcal{U}} f_i$ for a family $(f_i)_{i \in I}$ in H and an ultrafilter $\mathcal{U}$ on I , and there are a family $(x_j)_{j \in J}$ in K and an ultrafilter $\mathcal{V}$ on J with $\lim_{j, \mathcal{V}} x_j = x$ and $\lim_{j, \mathcal{V}} f(x_j) = \beta \neq f(x)$. (Fix $\varepsilon > 0$ such that every neighborhood O of x contains a point x_O with $|f(x_O) - f(x)| \geq \varepsilon$, and let $\mathcal{V}$ be an ultrafilter on the set of neighborhoods of x containing, for each O , the set of neighborhoods contained in O .) By continuity of each f_i , we have

$$\lim_{i, \mathcal{U}} \lim_{j, \mathcal{V}} f_i(x_j) = \lim_{i, \mathcal{U}} f_i(x) = f(x), \quad \lim_{j, \mathcal{V}} \lim_{i, \mathcal{U}} f_i(x_j) = \lim_{j, \mathcal{V}} f(x_j) = \beta.$$

After an affine change of scale, Proposition 1.2.3 and the proof of Corollary 1.2.6 turn these families into sequences whose iterated limits along a nonprincipal ultrafilter are β and $f(x)$. This contradicts (3).

(1) $\Rightarrow$ (2). Let W be the weak closure of H , which is weakly compact by (1). The identity map from $C(K)$ with the weak topology to $C(K)$ with the pointwise topology is continuous, so the set W is pointwise compact, hence closed in $\mathbb{R}^K$. Hence the pointwise closure of H is contained in $W \subseteq C(K)$.

(2) $\Rightarrow$ (1). Let $\bar{H}$ be the pointwise closure of H . This set is compact in $\mathbb{R}^K$, because H is bounded, and is contained in $C(K)$ by (2). We need only show that the weak and pointwise topologies agree on $\bar{H}$, since then the set $\bar{H}$ is weakly compact. By the Riesz representation theorem, the weak topology is generated by the maps $h \mapsto \int h d\mu$ with μ a Radon measure on K . On the uniformly bounded set $\bar{H}$ each such map is sequentially continuous for the

pointwise topology, by Lebesgue's dominated convergence theorem. Grothendieck proved that a pointwise compact subset of $C(K)$ is *angelic*: every point in the closure of a subset is the limit of a sequence from that subset [Gro52]; see also Pryce [Pry71]. On an angelic space sequentially continuous functions are continuous, which completes the proof. ☺

The simplest family of continuous functions with a discontinuous pointwise limit already shows the two iterated limits parting company.

Example A.1.4. Let $K = [0, 1]$ and $H = \{f_n : n \geq 1\}$ with $f_n(x) = x^n$. The pointwise limit of (f_n) is 0 on $[0, 1)$ and 1 at 1, which is not continuous, so (2) of Theorem A.1.3 fails. Condition (3) fails along $x_m = 1 - 1/m$: for each m we have $\lim_n f_n(x_m) = 0$, while for each n we have $\lim_m f_n(x_m) = 1$, so the two iterated limits are 0 and 1 (see Figure A.1). Hence the set H is not relatively weakly compact in $C[0, 1]$.




FIGURE A.1. Example A.1.4: the functions $f_n(x) = x^n$ (purple, darker for larger n) converge pointwise to a function (pink) that is 0 on $[0, 1)$ and jumps to 1 at $x = 1$. Along the points $x_m = 1 - 1/m$, letting $n \rightarrow \infty$ first gives 0, and letting $m \rightarrow \infty$ first gives 1.

Weak compactness becomes a statement about sequences through a classical theorem of Eberlein and Šmulian, and the criterion then yields convergent subsequences.

Fact A.1.5 (Eberlein–Šmulian). *In a Banach space, every sequence in a relatively weakly compact set has a weakly convergent subsequence [Š40, Ebe47].*

Corollary A.1.6. *If a bounded set $H \subseteq C(K)$ satisfies the double limit condition, then every sequence in H has a subsequence that converges pointwise on K , to a continuous function.*

Proof. By Theorem A.1.3, the set H is relatively weakly compact. By Fact A.1.5, every sequence in H has a subsequence converging weakly to some $f \in C(K)$. Point evaluations are bounded linear functionals, so the subsequence converges pointwise to f . ☺

The criterion meets stability through the spaces of types of Section 1.2. For $\varphi: A \times B \rightarrow [0, 1]$, recall the evaluation maps $e_b: \text{LS}(\varphi) \rightarrow [0, 1]$, given by $e_b(p) = p(b)$, and $e_a: \text{RS}(\varphi) \rightarrow [0, 1]$, given by $e_a(q) = q(a)$, of Subsection 1.2.4.

Proposition A.1.7. *Let $\varphi: A \times B \rightarrow [0, 1]$ be stable. Then the family $\{e_b : b \in B\}$ is relatively weakly compact in $C(\text{LS}(\varphi))$, and $\{e_a : a \in A\}$ is relatively weakly compact in $C(\text{RS}(\varphi))$. More precisely, the pointwise closure of $\{e_b : b \in B\}$ consists of the functions $\hat{\varphi}(\cdot, q)$ with $q \in \text{RS}(\varphi)$, where $\hat{\varphi}$ is the function of Proposition 1.2.8.*

Proof. By (1.2) with a constant family, we have $e_b = \hat{\varphi}(\cdot, \text{rtp}_\varphi(b))$. If $q = \lim_{j, \gamma} \text{rtp}_\varphi(b_j)$, then by continuity of $\hat{\varphi}(p, \cdot)$ we get $\lim_{j, \gamma} e_{b_j}(p) = \hat{\varphi}(p, q)$ for every p . Thus every pointwise limit of functions e_b is of the form $\hat{\varphi}(\cdot, q)$, and conversely every such function is a pointwise limit of this kind. These functions are continuous on $\text{LS}(\varphi)$ by separate continuity. Hence condition (2) of Theorem A.1.3 holds, and $\{e_b : b \in B\}$ is relatively weakly compact. The statement for $\{e_a : a \in A\}$ is proved symmetrically. ☺

Proposition A.1.7 places the type spaces of a stable function among the most rigid compact spaces of topology. Recall that a compact space is an *Eberlein compactum* if the space is homeomorphic to a weakly compact subset of a Banach space with the weak topology.

Corollary A.1.8. *Let $\varphi : A \times B \rightarrow [0, 1]$ be stable. Then the spaces $\text{RS}(\varphi)$ and $\text{LS}(\varphi)$ are Eberlein compacta. In particular, every right φ -type is the limit of a sequence of realized right types, and likewise for left types.*

Proof. Let H be the pointwise closure of $\{e_b : b \in B\}$ in $\mathbb{R}^{\text{LS}(\varphi)}$. By Proposition A.1.7, the weak closure W of $\{e_b : b \in B\}$ in $C(\text{LS}(\varphi))$ is weakly compact. The identity map from W with the weak topology to W with the coarser pointwise topology is a continuous bijection from a compact space to a Hausdorff space, so the two topologies agree on W . In particular, the set W is pointwise compact, hence pointwise closed, so we have $H \subseteq W$; and $W \subseteq H$ because the weak topology is finer than the pointwise topology. Thus we have $H = W$, and on H the weak and pointwise topologies agree. The map $q \mapsto \hat{\varphi}(\cdot, q)$ from $\text{RS}(\varphi)$ onto H is continuous for the pointwise topology, by the separate continuity of $\hat{\varphi}$, and injective, because $\hat{\varphi}(\text{ltp}_\varphi(a), q) = q(a)$ for every $a \in A$. A continuous bijection from a compact space to a Hausdorff space is a homeomorphism, so the space $\text{RS}(\varphi)$ is homeomorphic to the weakly compact set H . The same argument, with the roles of A and B exchanged, applies to $\text{LS}(\varphi)$. Finally, the set H is a pointwise compact subset of $C(\text{LS}(\varphi))$, hence angelic (see the discussion before Fact A.1.5), so each point of H is the limit of a sequence from $\{e_b : b \in B\}$; through the homeomorphism, each $q \in \text{RS}(\varphi)$ is the limit of a sequence of realized types $\text{rtp}_\varphi(b_n)$. ☺

Sequences are a genuine restriction. In the unstable structure $(\omega_1, <)$ with $\varphi(x, y)$ the formula $x < y$, the left type p given by $p(\beta) = 0$ for every $\beta < \omega_1$ (an element above every countable ordinal) is a limit of realized types along an ultrafilter, but not along any sequence: if $\alpha_n < \omega_1$ for every n , then we have $\alpha = \sup_n \alpha_n < \omega_1$, and the realized types $\text{ltp}_\varphi(\alpha_n)$ all take the value 1 at $\alpha + 1$, so no subsequence of them converges to p .

A.2 HISTORICAL REMARKS

✂ **Appendix A.1.** Grothendieck proved the criterion of Theorem A.1.3 in a paper of 1952 [Gro52], announced two years earlier [Gro50]; the same paper proves that pointwise compact subsets of $C(K)$ are angelic. The Eberlein–Šmulian theorem is due to Šmulian [Š40] and Eberlein [Ebe47]. Later proofs and extensions of the criterion, by Namioka, Pták, Young, Simons, Pryce and König, are listed in the timeline in the historical remarks on Subsection 1.2.1. The reading of the criterion as a statement about stable functions, which is the point of view of Definition A.1.1 and Proposition A.1.7, is due to Ben Yaacov [BY14] and Pillay [Pil17] for formulas, following Iovino [Iov99] and Krivine–Maurey [KM81].

Bibliography

- [Ack37] W. Ackermann, *Die Widerspruchsfreiheit der allgemeinen Mengenlehre*, Math. Ann. **114** (1937), 305–315.
- [BFT78] J. Bourgain, D. H. Fremlin, and M. Talagrand, *Pointwise compact sets of Baire-measurable functions*, Amer. J. Math. **100** (1978), no. 4, 845–886. MR 509077
- [Bir35] G. Birkhoff, *Abstract linear dependence and lattices*, Amer. J. Math. **57** (1935), 800–804.
- [BL71] J. T. Baldwin and A. H. Lachlan, *On strongly categorical theories*, J. Symbolic Logic **36** (1971), 79–96.
- [BY03] I. Ben Yaacov, *Simplicity in compact abstract theories*, J. Math. Log. **3** (2003), 163–191.
- [BY09] ———, *Continuous and random Vapnik-Chervonenkis classes*, Israel J. Math. **173** (2009), 309–333.
- [BY14] Itai Ben Yaacov, *Model theoretic stability and definability of types, after A. Grothendieck*, Bull. Symb. Log. **20** (2014), no. 4, 491–496. MR 3294276
- [BYU10] Itai Ben Yaacov and Alexander Usvyatsov, *Continuous first order logic and local stability*, Trans. Amer. Math. Soc. **362** (2010), no. 10, 5213–5259. MR 2657678
- [Cam97] P. J. Cameron, *The random graph*, The Mathematics of Paul Erdős II, Algorithms Combin., vol. 14, Springer, Berlin, 1997, pp. 333–351.
- [Cas07] E. Casanovas, *Simplicity simplified*, Rev. Colombiana Mat. **41** (2007), 263–277.
- [Cas11] ———, *Simple Theories and Hyperimaginaries*, Lecture Notes in Logic, vol. 39, Cambridge University Press, 2011.
- [CH99] Z. Chatzidakis and E. Hrushovski, *Model theory of difference fields*, Trans. Amer. Math. Soc. **351** (1999), 2997–3071.
- [Con17] G. Conant, *An axiomatic approach to free amalgamation*, J. Symbolic Logic **82** (2017), 648–671.
- [CR70] H. H. Crapo and G.-C. Rota, *On the Foundations of Combinatorial Theory: Combinatorial Geometries*, preliminary ed., MIT Press, Cambridge, MA, 1970.
- [CT26] G. Conant and C. Terry, *Encoding orders and trees in real-valued functions*, arXiv:2607.21761, 2026.
- [DFJP74] W. J. Davis, T. Figiel, W. B. Johnson, and A. Pelczynski, *Factoring weakly compact operators*, J. Functional Analysis **17** (1974), 311–327.
- [Ebe47] W. F. Eberlein, *Weak compactness in Banach spaces. I*, Proc. Nat. Acad. Sci. U.S.A. **33** (1947), 51–53.
- [EM56] A. Ehrenfeucht and A. Mostowski, *Models of axiomatic theories admitting automorphisms*, Fund. Math. **43** (1956), 50–68. MR 0084456 (18,863a)
- [ER63] P. Erdős and A. Rényi, *Asymmetric graphs*, Acta Math. Acad. Sci. Hungar. **14** (1963), 295–315.
- [Fag76] R. Fagin, *Probabilities on finite models*, J. Symbolic Logic **41** (1976), 50–58.
- [Fra54] R. Fraïssé, *Sur l’extension aux relations de quelques propriétés des ordres*, Ann. Sci. École Norm. Sup. **71** (1954), 363–388.
- [Gai76] Haim Gaifman, *Models and types of Peano’s arithmetic*, Ann. Math. Logic **9** (1976), no. 3, 223–306. MR 0406791
- [GIL02] Rami Grossberg, José N. Iovino, and Olivier Lessmann, *A primer of simple theories*, Arch. Math. Logic **41** (2002), no. 6, 541–580. MR MR1923196 (2003f:03043)
- [GKLT69] Y. V. Glebskiĭ, D. I. Kogan, M. I. Liogon’kiĭ and V. A. Talanov, *Range and degree of realizability of formulas in the restricted predicate calculus*, Kibernetika (Kiev) (1969), no. 2, 17–28.
- [Gro50] A. Grothendieck, *Sur la complétion du dual d’un espace vectoriel localement convexe*, C. R. Acad. Sci. Paris **230** (1950), 605–606.
- [Gro52] ———, *Critères de compacité dans les espaces fonctionnels généraux*, Amer. J. Math. **74** (1952), 168–186. MR 47313

- [Hen71] C. W. Henson, *A family of countable homogeneous graphs*, Pacific J. Math. **38** (1971), 69–83.
- [HH84] V. Harnik. and L. Harrington, *Fundamentals of forking*, Ann. Pure Appl. Logic **26** (1984), 245–286.
- [Hod93] Wilfrid Hodges, *Model theory*, Encyclopedia of Mathematics and its Applications, vol. 42, Cambridge University Press, Cambridge, 1993. MR 94e:03002
- [HP94] E. Hrushovski and A. Pillay, *Groups definable in local fields and pseudofinite fields*, Israel J. Math. **85** (1994), 203–262.
- [HP11] E. Hrushovski and A. Pillay, *On NIP and invariant measures*, J. Eur. Math. Soc. **13** (2011), 1005–1061.
- [HPP08] Ehud Hrushovski, Ya'acov Peterzil, and Anand Pillay, *Groups, measures, and the NIP*, J. Amer. Math. Soc. **21** (2008), no. 2, 563–596. MR 2373360 (2008k:03078)
- [Hru02] E. Hrushovski, *Pseudo-finite fields and related structures*, Model Theory and Applications, Quad. Mat., vol. 11, Aracne, Rome, 2002, pp. 151–212.
- [Iov94] José N. Iovino, *Stable theories in functional analysis*, Ph.D. thesis, University of Illinois at Urbana-Champaign, 1994.
- [Iov99] ———, *Stable models and reflexive Banach spaces*, J. Symbolic Logic **64** (1999), no. 4, 1595–1600. MR 2001e:03068
- [Jam64] R. C. James, *Uniformly non-square Banach spaces*, Ann. of Math. **80** (1964), 542–550.
- [Kim98] B. Kim, *Forking in simple unstable theories*, J. London Math. Soc. (2) **57** (1998), 257–267.
- [Kim01] ———, *Simplicity, and stability in there*, J. Symbolic Logic **66** (2001), 822–836.
- [KM81] Jean-Louis Krivine and Bernard Maurey, *Espaces de Banach stables*, Israel J. Math. **39** (1981), no. 4, 273–295. MR 83a:46030
- [Kön87] H. König, *Angelic spaces and the double limit relation*, J. London Math. Soc. (2) **35** (1987), 454–470.
- [KP97] Byunghan Kim and Anand Pillay, *Simple theories*, Ann. Pure Appl. Logic **88** (1997), no. 2-3, 149–164, Joint AILA-KGS Model Theory Meeting (Florence, 1995). MR MR1600895 (99b:03049)
- [KP98] B. Kim and A. Pillay, *From stability to simplicity*, Bull. Symbolic Logic **4** (1998), 17–36.
- [KRS24] Itay Kaplan, Nicholas Ramsey, and Pierre Simon, *Generic stability independence and treeless theories*, Forum of Mathematics, Sigma **12** (2024), e49.
- [Kru19] A. Kruckman, *Disjoint n -amalgamation and pseudofinite countably categorical theories*, Notre Dame J. Formal Logic **60** (2019), 139–160.
- [Las92] M. C. Laskowski, *Vapnik–Chervonenkis classes of definable sets*, J. London Math. Soc. (2) **45** (1992), 377–384.
- [LP79] D. Lascar and B. Poizat, *An introduction to forking*, J. Symbolic Logic **44** (1979), 330–350.
- [LW80] A. H. Lachlan and R. E. Woodrow, *Countable ultrahomogeneous undirected graphs*, Trans. Amer. Math. Soc. **262** (1980), 51–94.
- [Mak84] M. Makkai, *A survey of basic stability theory, with particular emphasis on orthogonality and regular types*, Israel J. Math. **49** (1984), 181–238.
- [Maz33] S. Mazur, *Über konvexe Mengen in linearen normierten Räumen*, Studia Math. **4** (1933), 70–84.
- [ML36] S. Mac Lane, *Some interpretations of abstract linear dependence in terms of projective geometry*, Amer. J. Math. **58** (1936), 236–240.
- [Mor65] M. Morley, *Categoricity in power*, Trans. Amer. Math. Soc. **114** (1965), 514–538.
- [Nam61] I. Namioka, *A substitute for Lebesgue's bounded convergence theorem*, Proc. Amer. Math. Soc. **12** (1961), 713–716.
- [Nam74] Isaac Namioka, *Separate continuity and joint continuity*, Pacific J. Math. **51** (1974), 515–531.
- [Pil96] A. Pillay, *Geometric stability theory*, The Clarendon Press Oxford University Press, New York, 1996, Oxford Science Publications.
- [Pil17] ———, *Stability and Grothendieck*, arXiv:1703.04761, 2017.
- [Poi81] B. Poizat, *Théories instables*, J. Symbolic Logic **46** (1981), 513–522.
- [Pry71] J. D. Pryce, *A device of R. J. Whitley's applied to pointwise compactness in spaces of continuous functions*, Proc. London Math. Soc. (3) **23** (1971), 532–546.
- [PS86] A. Pillay and C. Steinhorn, *Definable sets in ordered structures. I*, Trans. Amer. Math. Soc. **295** (1986), 565–592.
- [Ptá63a] V. Pták, *A combinatorial lemma on the existence of convex means and its application to weak compactness*, Convexity, Proc. Sympos. Pure Math., vol. 7, Amer. Math. Soc., 1963, pp. 437–450.
- [Ptá63b] Vlastimil Pták, *An extension theorem for separately continuous functions and its application to functional analysis*, Comment. Math. Univ. Carolinae **4** (1963), 109–116. MR 0166599
- [Rad64] R. Rado, *Universal graphs and universal functions*, Acta Arith. **9** (1964), 331–340.

- [Ray83] Y. Raynaud, *Espaces de Banach superstables, distances stables et homéomorphismes uniformes*, Israel J. Math. **44** (1983), no. 1, 33–52. MR 84j:46024
- [Ros74] H. P. Rosenthal, *A characterization of Banach spaces containing l^1* , Proc. Nat. Acad. Sci. U.S.A. **71** (1974), 2411–2413. MR 358307
- [Ros77] ———, *Pointwise compact subsets of the first Baire class*, American Journal of Mathematics **99** (1977), no. 2, 362–378.
- [Row64] F. Rowbottom, *The ι_0 conjecture for uncountable theories*, Notices of the Amer. Math. Soc. **11** (1964), 248–249.
- [Sau72] N. Sauer, *On the density of families of sets*, J. Combinatorial Theory Ser. A **13** (1972), 145–147. MR 307902
- [She69] S. Shelah, *Stable theories*, Israel J. Math. **7** (1969), 187–202.
- [She71] Saharon Shelah, *Stability, the f.c.p., and superstability; model theoretic properties of formulas in first order theory*, Ann. Math. Logic **3** (1971), no. 3, 271–362. MR 0317926
- [She72] ———, *A combinatorial problem; stability and order for models and theories in infinitary languages*, Pacific J. Math. **41** (1972), 247–261. MR 307903
- [She80] S. Shelah, *Simple unstable theories*, Annals of Math. Logic **19** (1980), 177–203.
- [She90] S. Shelah, *Classification theory and the number of nonisomorphic models*, second ed., Studies in Logic and the Foundations of Mathematics, vol. 92, North-Holland Publishing Co., Amsterdam, 1990. MR MR1083551 (91k:03085)
- [She96] ———, *Toward classifying unstable theories*, Ann. Pure Appl. Logic **80** (1996), 229–255.
- [Sim68] S. Simons, *The iterated limit condition, a Fubini theorem and weak compactness*, Math. Ann. **176** (1968), 87–95.
- [Sim72a] ———, *A convergence theorem with boundary*, Pacific J. Math. **40** (1972), 703–708.
- [Sim72b] ———, *Maximinimax, minimax, and antiminimax theorems and a result of R. C. James*, Pacific J. Math. **40** (1972), 709–718.
- [Sim72c] ———, *On Pták’s combinatorial lemma*, Pacific J. Math. **40** (1972), 719–721.
- [Sim15a] Pierre Simon, *A guide to NIP theories*, Lecture Notes in Logic, vol. 44, Association for Symbolic Logic, Chicago, IL; Cambridge Scientific Publishers, Cambridge, 2015. MR 3560428
- [Sim15b] ———, *Rosenthal compacta and NIP formulas*, Fund. Math. **231** (2015), no. 1, 81–92. MR 3361236
- [Ste10] E. Steinitz, *Algebraische Theorie der Körper*, J. Reine Angew. Math. **137** (1910), 167–309.
- [Sto37] M. H. Stone, *Applications of the theory of Boolean rings to general topology*, Trans. Amer. Math. Soc. **41** (1937), 375–481.
- [Sto48] ———, *The generalized Weierstrass approximation theorem*, Math. Mag. **21** (1948), 167–184 and 237–254.
- [TZ12] K. Tent and M. Ziegler, *A Course in Model Theory*, Lecture Notes in Logic, vol. 40, Cambridge University Press, 2012.
- [VC71] Vladimir N Vapnik and Alexey Ya Chervonenkis, *On the uniform convergence of relative frequencies of events to their probabilities*, Theory of Probability & Its Applications **16** (1971), no. 2, 264–280.
- [vdW30] B. L. van der Waerden, *Moderne Algebra*, Springer, Berlin, 1930.
- [Š40] V. Šmulian, *Über lineare topologische Räume*, Mat. Sbornik N.S. **7(49)** (1940), 425–448.
- [Wag00] Frank O. Wagner, *Simple theories*, Mathematics and its Applications, vol. 503, Kluwer Academic Publishers, Dordrecht, 2000. MR MR1747713 (2001b:03035)
- [Wei85] K. Weierstrass, *Über die analytische Darstellbarkeit sogenannter willkürlicher Functionen einer reellen Veränderlichen*, Sitzungsber. Königl. Preuß. Akad. Wiss. Berlin (1885), 633–639 and 789–805.
- [Whi35] H. Whitney, *On the abstract properties of linear dependence*, Amer. J. Math. **57** (1935), 509–533.
- [You71] N. J. Young, *On Pták’s double-limit theorems*, Proc. Edinburgh Math. Soc. (2) **17** (1970/71), 194–200. MR 0310587