

GRADUATE STUDIES
IN MATHEMATICS **220**

Ultrafilters Throughout Mathematics

Isaac Goldbring



Ultrafilters Throughout Mathematics

GRADUATE STUDIES **220**
IN MATHEMATICS

Ultrafilters Throughout Mathematics

Isaac Goldbring

EDITORIAL COMMITTEE

Marco Gualtieri
Bjorn Poonen
Gigliola Staffilani (Chair)
Jeff A. Viaclovsky
Rachel Ward

2020 *Mathematics Subject Classification*. Primary 03C20, 54D80, 03H05, 03E55, 03C50.

During the writing of this book, the author was partially supported by NSF grants DMS-1349399 (CAREER) and DMS-2054477.

For additional information and updates on this book, visit
www.ams.org/bookpages/gsm-220

Library of Congress Cataloging-in-Publication Data

Names: Goldbring, Isaac, author.

Title: Ultrafilters throughout mathematics / Isaac Goldbring.

Description: Providence, Rhode Island : American Mathematical Society, [2022] | Series: Graduate studies in mathematics, 1065-7339 ; 220 | Includes bibliographical references and index.

Identifiers: LCCN 2021055552 | ISBN 9781470469009 (hardback) | ISBN 9781470469610 (paperback) | ISBN 9781470469603 (ebook)

Subjects: LCSH: Ultrafilters (Mathematics) | AMS: Mathematical logic and foundations – Model theory – Ultraproducts and related constructions. | General topology – Fairly general properties of topological spaces – Special constructions of topological spaces (spaces of ultrafilters, etc.). | Mathematical logic and foundations – Nonstandard models – Nonstandard models in mathematics. | Mathematical logic and foundations – Set theory – Large cardinals. | Mathematical logic and foundations – Model theory – Models with special properties (saturated, rigid, etc.).

Classification: LCC QA248 .G574 2022 | DDC 511.3–dc23/eng/20220112

LC record available at <https://lcn.loc.gov/2021055552>

Copying and reprinting. Individual readers of this publication, and nonprofit libraries acting for them, are permitted to make fair use of the material, such as to copy select pages for use in teaching or research. Permission is granted to quote brief passages from this publication in reviews, provided the customary acknowledgment of the source is given.

Republication, systematic copying, or multiple reproduction of any material in this publication is permitted only under license from the American Mathematical Society. Requests for permission to reuse portions of AMS publication content are handled by the Copyright Clearance Center. For more information, please visit www.ams.org/publications/pubpermissions.

Send requests for translation rights and licensed reprints to reprint-permission@ams.org.

© 2022 by the American Mathematical Society. All rights reserved.

The American Mathematical Society retains all rights
except those granted to the United States Government.

Printed in the United States of America.

⊗ The paper used in this book is acid-free and falls within the guidelines
established to ensure permanence and durability.

Visit the AMS home page at <https://www.ams.org/>

10 9 8 7 6 5 4 3 2 1 27 26 25 24 23 22

To Karina, Kaylee, and Daniella

Contents

Preface	xiii
Part 1. Ultrafilters and their applications	
Chapter 1. Ultrafilter basics	3
§1.1. Basic definitions	3
§1.2. The ultrafilter quantifier	6
§1.3. The category of ultrafilters	7
§1.4. The number of ultrafilters	11
§1.5. The ultrafilter number \mathfrak{u}	13
§1.6. The Rudin-Keisler order	14
§1.7. Notes and references	17
Chapter 2. Arrow's theorem on fair voting	19
§2.1. Statement of the theorem	19
§2.2. The connection with ultrafilters	21
§2.3. Block voting	23
§2.4. Finishing the proof	24
§2.5. Notes and references	26
Chapter 3. Ultrafilters in topology	27
§3.1. Ultralimits	27
§3.2. The Stone-Čech compactification: the discrete case	31
§3.3. z -ultrafilters and the Stone-Čech compactifications in general	33
§3.4. The Stone representation theorem	38

§3.5.	Notes and References	44
Chapter 4.	Ramsey theory and combinatorial number theory	45
§4.1.	Ramsey's theorem	45
§4.2.	Idempotent ultrafilters and Hindman's theorem	46
§4.3.	Banach density, means, and measures	50
§4.4.	Furstenberg's correspondence principle	53
§4.5.	Jin's sunset theorem	55
§4.6.	Notes and references	57
Chapter 5.	Foundational concerns	59
§5.1.	The ultrafilter theorem and the axiom of choice: Part I	59
§5.2.	Can there exist a "definable" ultrafilter on \mathbb{N} ?	62
§5.3.	The ultrafilter game	66
§5.4.	Selective ultrafilters and P-points	70
§5.5.	Notes and references	77
Part 2. Classical ultraproducts		
Chapter 6.	Classical ultraproducts	81
§6.1.	Motivating the definition of ultraproducts	82
§6.2.	Ultraproducts of sets	83
§6.3.	Ultraproducts of structures	85
§6.4.	Łoś's theorem	86
§6.5.	The ultrafilter theorem and the axiom of choice: Part II	89
§6.6.	Countably incomplete ultrafilters	92
§6.7.	Revisiting the Rudin-Keisler order	94
§6.8.	Cardinalities of ultraproducts	96
§6.9.	Iterated ultrapowers	98
§6.10.	A category-theoretic perspective on ultraproducts	101
§6.11.	The Feferman-Vaught theorem	104
§6.12.	Notes and references	108
Chapter 7.	Applications to geometry, commutative algebra, and number theory	109
§7.1.	Ax's theorem on polynomial functions	109
§7.2.	Bounds in the theory of polynomial rings	111
§7.3.	The Ax-Kochen theorem and Artin's conjecture	116

§7.4. Notes and references	122
Chapter 8. Ultraproducts and saturation	123
§8.1. Saturation	123
§8.2. First saturation properties of ultraproducts	127
§8.3. Regular ultrafilters	128
§8.4. Good ultrafilters: Part 1	134
§8.5. Good ultrafilters: Part 2	141
§8.6. Keisler's order	145
§8.7. Notes and references	155
Chapter 9. Nonstandard analysis	157
§9.1. Naïve axioms for nonstandard analysis	157
§9.2. Nonstandard numbers big and small	159
§9.3. Some nonstandard calculus	161
§9.4. Ultrapowers as a model of nonstandard analysis	163
§9.5. Complete extensions and limit ultrapowers	164
§9.6. Many-sorted structures and internal sets	168
§9.7. Nonstandard generators of ultrafilters	173
§9.8. Hausdorff ultrafilters	177
§9.9. Notes and references	178
Chapter 10. Limit groups	181
§10.1. Introducing the class of limit groups	181
§10.2. First examples and properties of limit groups	183
§10.3. Connection with fully residual freeness	186
§10.4. Explaining the terminology: the space of marked groups	189
§10.5. Notes and references	191
 Part 3. Metric ultraproducts and their applications	
Chapter 11. Metric ultraproducts	195
§11.1. Definition of the metric ultraproduct	195
§11.2. Metric ultraproducts and nonstandard hulls of metric spaces	198
§11.3. Completeness properties of the metric ultraproduct	199
§11.4. Continuous logic	201
§11.5. Reduced products of metric structures	208
§11.6. Notes and references	209

Chapter 12. Asymptotic cones and Gromov's theorem	211
§12.1. Some group-theoretic preliminaries	212
§12.2. Growth rates of groups	213
§12.3. Gromov's theorem on polynomial growth	216
§12.4. Definition of asymptotic cones	221
§12.5. General properties of asymptotic cones	223
§12.6. Growth functions and properness of the asymptotic cones	226
§12.7. Properness of asymptotic cones revisited	229
§12.8. Nonhomeomorphic asymptotic cones	231
§12.9. Notes and references	232
Chapter 13. Sofic groups	233
§13.1. Ultraproducts of bi-invariant metric groups	233
§13.2. Definition of sofic groups	235
§13.3. Examples of sofic groups	240
§13.4. An application of sofic groups	245
§13.5. Notes and references	248
Chapter 14. Functional analysis	249
§14.1. Banach space ultraproducts	249
§14.2. Applications to local geometry of Banach spaces	254
§14.3. Commutative C^* -algebras and ultracoproducts of compact spaces	259
§14.4. The tracial ultraproduct construction	264
§14.5. The Connes embedding problem	273
§14.6. Notes and references	277
Part 4. Advanced topics	
Chapter 15. Does an ultrapower depend on the ultrafilter?	281
§15.1. Statement of results	281
§15.2. The case when \mathcal{M} is unstable	284
§15.3. The case when \mathcal{M} is stable	292
§15.4. Notes and references	295
Chapter 16. The Keisler-Shelah theorem	297
§16.1. The Keisler-Shelah theorem	297
§16.2. Application: Elementary classes	305

§16.3. Application: Robinson's joint consistency theorem	306
§16.4. Application: Elementary equivalence of matrix rings	307
§16.5. Notes and references	307
Chapter 17. Large cardinals	309
§17.1. Worldly cardinals	309
§17.2. Inaccessible cardinals	311
§17.3. Measurable cardinals	314
§17.4. Strongly and weakly compact cardinals	320
§17.5. Ramsey cardinals	325
§17.6. Measurable cardinals as critical points of elementary embeddings	327
§17.7. An application of large cardinals	333
§17.8. Notes and references	338
Part 5. Appendices	
Appendix A. Logic	341
§A.1. Languages and structures	341
§A.2. Syntax and semantics	342
§A.3. Embeddings	344
§A.4. References	346
Appendix B. Set theory	347
§B.1. The axioms of ZFC	347
§B.2. Ordinals	350
§B.3. Cardinals	351
§B.4. V and L	353
§B.5. Relative consistency statements	353
§B.6. Relativization and absoluteness	355
§B.7. References	356
Appendix C. Category theory	357
§C.1. Categories	357
§C.2. Functors, natural transformations, and equivalences of categories	359
§C.3. Limits	360
§C.4. References	362

Appendix D. Hints and solutions to selected exercises	363
Bibliography	385
Index	395

Preface

What is this book about?

This book is about **ultrafilters**. So what is an ultrafilter? Given a set X , an ultrafilter on X is simply a “sensible” division of all of the subsets of X into two categories: small and large. For this division to be sensible, one should impose some axioms:

- X should be a large subset of X , while \emptyset should be a small subset of X .
- If Y is a large subset of X and $Y \subseteq Z \subseteq X$, then Z should also be large; that is, a set containing a large set should also be large.
- If Y and Z are two large subsets of X , then so is $Y \cap Z$.

The last axiom is perhaps not entirely intuitive, but becomes more intuitive when stated in terms of small sets: the union of two small sets is once again small. The axioms also imply that a set is large precisely when its complement is small.

Why write a book about such a seemingly simple notion? It turns out that this notion is very useful for describing limits of various objects. For example, much to the chagrin of many calculus students, one knows that there are many sequences $(a_n)_{n \in \mathbb{N}}$ from $[0, 1]$ that have no limit. However, limit in the usual sense is very restrictive in that it requires a_n to be close to the limit for a large number of n , where large here means for all but finitely many n . Note that this restrictive notion of largeness does not lead to an ultrafilter on \mathbb{N} as there are certainly sets that are infinite and which have infinite complement. However, if one works with a notion of largeness as given by an ultrafilter, then all of a sudden *every sequence in $[0, 1]$ has a*

limit! This fact can be used as a powerful tool in analytic and topological endeavors.

The notion of ultrafilter also allows one to consider limits of families of structures like groups, rings, graphs, or Banach spaces. The limiting structures alluded to here are called **ultraproducts** and will become a central part of this book. These limiting objects can be very useful in solving problems, for often various desirable properties are approximately true in the individual structures of the family, while in the limit they become exactly true.

Who should read this book?

The short answer is: everyone! More precisely, the thesis of this book is that, while ultrafilters and ultraproducts are often relegated to graduate-level courses in logic, we believe that this practice is entirely misguided. Indeed, the notion of ultrafilter and ultraproduct are entirely accessible to an advanced undergraduate or beginning graduate student in mathematics (the target audience of this book). Moreover, as we will see throughout the course of this book, ultrafilters and ultraproducts have had numerous applications to nearly every area of mathematics. Thus, no matter what area of mathematics the reader is interested in, it is quite likely that ultrafilters and ultraproducts have made an impact in that area. An attempt has been made to present as diverse a sample of such applications as possible.

That being said, this book is being written by a logician, and ultrafilters present numerous fascinating foundational concerns, many of which are discussed in this book. If the reader is purely interested in mathematical applications, they may safely skip the portions of this book discussing these metamathematical issues.

What is in this book?

Let us briefly summarize the contents of this book. Part 1 is entirely devoted to ultrafilters. Chapter 1 introduces the basic facts about ultrafilters, including what it means for them to be isomorphic and how many of them there are. Chapter 2 provides one with a first application of ultrafilters, namely to a proof of Arrow's theorem on fair voting. This application is nice in the sense that it requires little to no mathematical background and yet exemplifies a perfect use of ultrafilters. Chapter 3 introduces the use of ultrafilters in topology, including the aforementioned facts about generalized limits. This chapter also shows how ultrafilters can be used to describe the important Stone-Ćech compactification construction. Chapter 4 is a brief introduction to how ultrafilters can be used in certain parts of combinatorics; a much more detailed investigation of that line of research can be found in

the book [42], written by the author with Mauro Di Nasso and Martino Lupini. Chapter 5, the last chapter in Part 1 of the book, discusses many of the interesting foundational issues presented by the existence of ultrafilters.

Part 2 of the book is concerned with the classical ultraproduct construction. As alluded to above, this construction allows one to take the limit of families of objects such as groups, rings, graphs, etc., . . . The lengthy Chapter 6 introduces this construction and proves the Fundamental Theorem of Ultraproducts (otherwise known as Łoś's theorem), which states that the truth of a first-order sentence in an ultraproduct is determined by whether or not the sentence is true in a large (as measured by the ultrafilter) number of the individual structures. This chapter includes many other important facts about ultraproducts, including cardinalities of ultraproducts and a discussion of what happens when one tries to iterate the ultraproduct construction.

Chapter 7 gives one a first look at how ultraproducts can be used "in practice." The applications in this chapter are all algebraic in nature, and include Ax's theorem on polynomial functions and the Ax-Kochen theorem relating the rings \mathbb{Z}_p of p -adic integers with the power series rings $\mathbb{F}_p[[T]]$. One important feature of ultraproducts is that they are often very "rich" in the precise sense of being saturated. Chapter 8 gives a detailed discussion of exactly how saturated ultraproducts can be. Chapter 9 gives a brief introduction to nonstandard analysis. While nonstandard analysis is a subject of its own, it is often presented using ultraproducts and we discuss this approach here. This chapter is far from a complete story on nonstandard analysis and we refer the interested reader to [42] for a more thorough discussion. Chapter 10 discusses the class of subgroups of nonstandard (in the sense of Chapter 9) free groups; the finitely generated such subgroups are called limit groups and have become a widely studied class of groups in geometric group theory.

The ultraproduct construction referred to above is suitable for discrete spaces such as those arising in algebra and combinatorics, but is not very useful for structures appearing in analysis. Part 3 of the book is concerned with a modification of the ultraproduct construction for structures based on metric spaces. Chapter 11 introduces this metric ultraproduct and discusses some of its basic properties. That chapter also includes a discussion of a relatively new logic, aptly called continuous logic, which is the logic naturally connected to this metric ultraproduct construction.

The remainder of Part 3 details several applications of the metric ultraproduct construction. Chapter 12 describes a fascinating theorem of Gromov from geometric group theory, where the key ingredient to the proof is a particular metric ultraproduct called an asymptotic cone. Chapter 13

discusses the class of sofic groups, which can be defined in terms of metric ultraproducts of symmetric groups. Chapter 14, the final chapter of Part 3, discusses some applications of metric ultraproducts to functional analysis. One might argue that functional analysis is an area of mathematics where ultraproducts have played an increasingly more important role. Unfortunately, the mathematical background needed by the reader is much larger in this area of mathematics and thus this section cannot quite do justice to the importance of ultraproducts in functional analysis.

Part 4, the last part of this book, is devoted to three advanced topics. Chapter 15 discusses a question that often arises to many people seeing ultraproducts for the first time: does the ultraproduct depend on the ultrafilter being used? The answer to this question is surprisingly subtle and a more or less complete answer to a specific case of this question is discussed. Chapter 16 discusses the fantastic Keisler-Shelah theorem, which shows how elementary equivalence, a notion from logic, can be reformulated in terms of isomorphic ultrapowers, a purely algebraic notion. This chapter also includes a few applications of the Keisler-Shelah theorem. Chapter 17, the final chapter of the book, shows how the study of large cardinals in set theory can be recast in terms of ultrafilters satisfying certain extra properties. This part of the book might require a bit more maturity and/or background from the reader.

What are the prerequisites for reading this book?

We have no illusions that any one student has all of the prerequisites necessary to read the entire book. However, this fact is by design! As discussed above, we are trying to convey to the reader that ultrafilters and ultraproducts are applicable in most areas of mathematics and thus we have tried to describe a wide variety of applications.

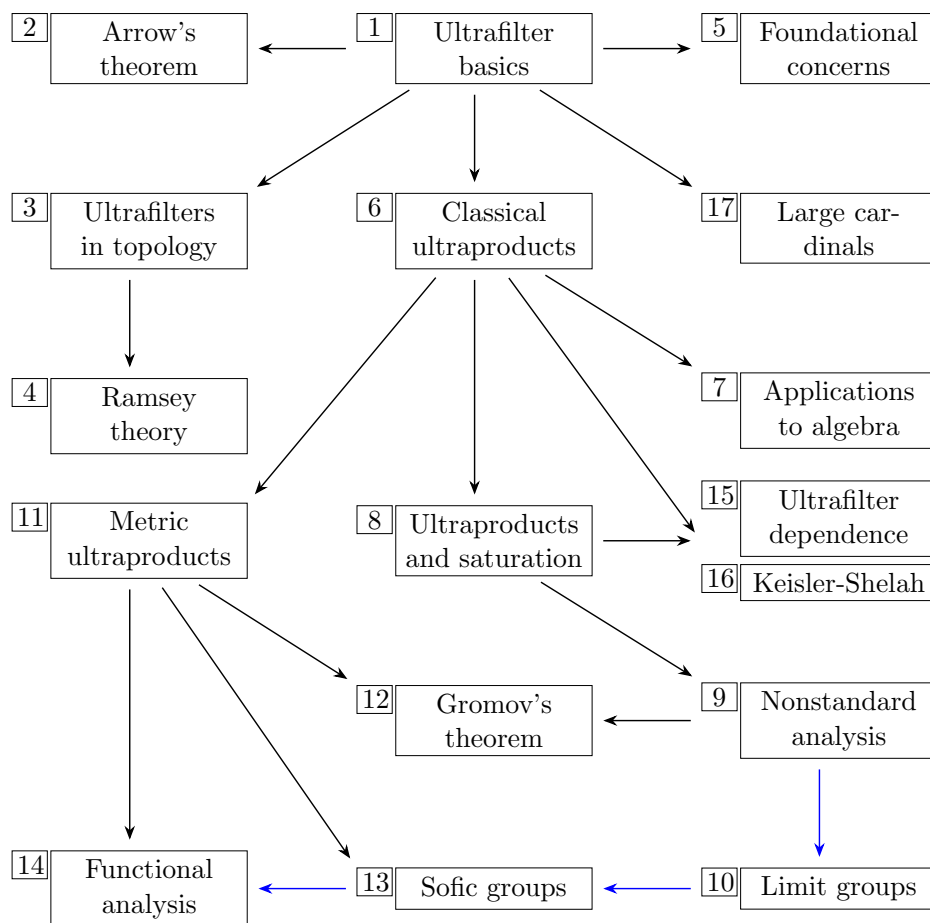
That being said, we have assumed that the reader is familiar with some basic facts from real analysis, topology, and algebra. Any facts that we believe are not part of the usual curricula from those disciplines are often described in full detail here. Sometimes certain topics are outside of the scope of this book and we provide references to the reader for places in the literature where they can learn more. It is also our hope that a reader interested in, for example, algebra sees the chapter on, say, functional analysis, and finds the general idea interesting enough that they decide to learn more about this area. In today's mathematical world, breadth is everything and an aspiring mathematician should keep their eyes open to all areas of mathematics.

In discussing ultrafilters, one cannot hide the fact that logic and set theory play an important role. Moreover, there is a high probability that

the average reader might not have the requisite knowledge in these areas to follow the main parts of this book. For the reader's convenience, appendices on these subjects are included in this book. Also, occasionally in the text, very basic parts of category theory are needed and the necessary facts from category theory are collected in the final appendix.

How to read this book

Some later chapters rely somewhat heavily on earlier chapters. The following flowchart lists some of these dependencies. The blue arrows indicate dependencies that are not strictly necessary but possibly helpful.



Exercises

Rather than ending each section or chapter with a list of exercises, we have instead sprinkled them throughout the text itself. Some of the exercises are simply checks for understanding, but others are more involved. Often,

the exercises themselves will be used in the proofs of later results. We recommend that the reader stop reading when they encounter an exercise and attempt a solution at that moment. Solutions to a handful of exercises appear in Appendix D but we urge the reader not to consult these solutions unless the situation becomes dire!

Acknowledgements

We thank Andreas Blass, Bradd Hart, Alessandro Sisto, Eleftherios Tachtsis, and Martin Zeman for helpful discussions during the writing of this book.

We also extend our deepest gratitude to Ryan Burkhart, Alec Fox, Michael Hehmann, Jennifer Pi, Brian Ransom, and Jessica Schirle for a very thorough reading of an earlier draft of this book and for offering many, many corrections and suggestions for improvement.

Isaac Goldbring
Irvine, CA

Part 1

Ultrafilters and their applications

Ultrafilter basics

In this chapter, we present the basic theory of ultrafilters. Section 1.1 contains the basic definitions of filters and ultrafilters and proves the existence of a nonprincipal ultrafilter on an infinite set. Section 1.2 is a short section devoted to explaining how one can view an ultrafilter on a set as a kind of quantifier. Section 1.3 gives a category-theoretic perspective on ultrafilters. In Section 1.4, we compute the cardinality of the set of ultrafilters on a given set. In Section 1.5, we introduce the cardinal characteristic \mathfrak{u} , which, roughly speaking, is the smallest number of sets needed to specify a nonprincipal ultrafilter on \mathbb{N} , while in Section 1.6, we introduce the Rudin-Keisler ordering on the collection of all ultrafilters, which is a relative measure of complexity for ultrafilters.

1.1. Basic definitions

Throughout this section, we let S denote a set.

Definition 1.1.1. A (proper) **filter** on S is a set \mathcal{F} of subsets of S (that is, $\mathcal{F} \subseteq \mathcal{P}(S)$) such that:

- (1) $\emptyset \notin \mathcal{F}$ but $S \in \mathcal{F}$;
- (2) if $A, B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$;
- (3) if $A \in \mathcal{F}$ and $A \subseteq B$, then $B \in \mathcal{F}$.

We think of elements of \mathcal{F} as “big” sets (because that is what filters do, they catch the big objects). The first and third axioms are (hopefully) intuitive properties of big sets. Perhaps the second axiom is not as intuitive, but if one thinks of the complement of a big set as a “small” set, then the

second axiom asserts that the union of two small sets is small (which is hopefully more intuitive).

Exercise 1.1.2. Suppose that S is infinite. Set $\mathcal{F} := \{A \subseteq S \mid S \setminus A \text{ is finite}\}$. Prove that \mathcal{F} is a filter on S , called the **Fréchet** or **cofinite** filter on S .

One often describes a filter by specifying a base:

Definition 1.1.3. Suppose that \mathcal{F} is a filter on S . Then a **base** for \mathcal{F} is a collection \mathcal{B} of subsets of S such that $\mathcal{F} = \{A \subseteq S \mid B \subseteq A \text{ for some } B \in \mathcal{B}\}$.

Exercise 1.1.4. Suppose that \mathcal{B} is a collection of nonempty subsets of S . Prove that \mathcal{B} is a base for a (necessarily unique) filter on S if and only if, for any $A, B \in \mathcal{B}$, there is $C \in \mathcal{B}$ such that $C \subseteq A \cap B$.

In practice, one has a collection \mathcal{D} of subsets of S which they would like to belong to some filter \mathcal{F} on S but which does not satisfy the criterion in the previous exercise for being a base for a filter. One can of course try to force \mathcal{D} to satisfy the criterion by closing \mathcal{D} under finite intersections. However, a base for a filter is required to consist of nonempty sets and, in closing \mathcal{D} under finite intersections, one may accidentally stumble upon the emptyset. Thus, the following definition becomes crucial:

Definition 1.1.5. Suppose that \mathcal{D} is a collection of subsets of S . We say that \mathcal{D} has the **finite intersection property** (or **FIP** for short) if, whenever $D_1, \dots, D_n \in \mathcal{D}$, we have $D_1 \cap \dots \cap D_n \neq \emptyset$.

The previous discussion thus establishes:

Theorem 1.1.6. *Suppose that \mathcal{D} is a collection of subsets of S with the finite intersection property. Then $\{D_1 \cap \dots \cap D_n \mid D_1, \dots, D_n \in \mathcal{D}\}$ is a base for a filter on S , called the **filter generated by \mathcal{D}** , denoted $\langle \mathcal{D} \rangle$.*

To be explicit, we have

$$\langle \mathcal{D} \rangle = \{E \subseteq S \mid D_1 \cap \dots \cap D_n \subseteq E \text{ for some } D_1, \dots, D_n \in \mathcal{D}\}.$$

Exercise 1.1.7. Suppose that \mathcal{F} is a filter on S and $A \subseteq S$. Prove that $\mathcal{F} \cup \{A\}$ has the FIP if and only if $(S \setminus A) \notin \mathcal{F}$.

If \mathcal{F} is a filter on S , then a subset of S cannot be simultaneously big and small (that is, both it and its complement belong to \mathcal{F}), but there is no requirement that one of the two be big. It will be desirable (for reasons that will become clear in a moment) to add this as an additional property:

Definition 1.1.8. If \mathcal{F} is a filter on S , then \mathcal{F} is an **ultrafilter** if, for any $A \subseteq S$, either $A \in \mathcal{F}$ or $S \setminus A \in \mathcal{F}$ (but not both!).

Ultrafilters are usually denoted by \mathcal{U} or \mathcal{V} , and the set of ultrafilters on S is usually denoted βS (for a topological reason that we will discuss in Chapter 3). Observe that the Fréchet filter on S is not an ultrafilter since there are sets $A \subseteq S$ such that A and $S \setminus A$ are both infinite.

Exercise 1.1.9 (For those who are familiar with measure theory). Given a set S and a collection \mathcal{U} of subsets of S , we have that \mathcal{U} is an ultrafilter on S if and only if there is a finitely additive probability measure μ on S that only takes the values 0 or 1 such that, for all $A \subseteq S$, we have $A \in \mathcal{U}$ if and only if $\mu(A) = 1$. In this case, μ is unique, whence we may denote it by $\mu_{\mathcal{U}}$.

Exercise 1.1.10. Suppose that \mathcal{U} is an ultrafilter on S and A_1, \dots, A_n are subsets of S such that $A_1 \cup \dots \cup A_n \in \mathcal{U}$. Prove that there is $i \in \{1, \dots, n\}$ such that $A_i \in \mathcal{U}$. Moreover, if the A_1, \dots, A_n are pairwise disjoint, prove that there is a unique such i .

There is actually a strong converse to the previous exercise that we will use in Chapter 2:

Exercise 1.1.11. Suppose that \mathcal{U} is a collection of nonempty subsets of S with the following property: Whenever A_1, A_2 , and A_3 are pairwise disjoint subsets of S with $S = A_1 \cup A_2 \cup A_3$ (with perhaps one or more of the $A_i = \emptyset$), then there is exactly one $i \in \{1, 2, 3\}$ with $A_i \in \mathcal{U}$. Prove that \mathcal{U} is an ultrafilter on S . (*Hint.* This is a fun exercise with Venn diagrams.)

Exercise 1.1.12. Suppose that \mathcal{U} is an ultrafilter on S and that $A \in \mathcal{U}$. Prove that

$$A \cap \mathcal{U} := \{A \cap B : B \in \mathcal{U}\}$$

is an ultrafilter on A , called the **ultrafilter on A induced by \mathcal{U}** .

We have yet to see an example of an ultrafilter. Here is a “trivial” source of ultrafilters:

Definition 1.1.13. Given $s \in S$, set $\mathcal{U}_s := \{A \subseteq S \mid s \in A\}$.

Exercise 1.1.14. For $s \in S$, prove that \mathcal{U}_s is an ultrafilter on S , called the **principal ultrafilter generated by s** .

We say that an ultrafilter \mathcal{U} on S is **principal** if $\mathcal{U} = \mathcal{U}_s$ for some $s \in S$; otherwise, we say that \mathcal{U} is **nonprincipal**. Although principal ultrafilters settle the question of the existence of ultrafilters, they will turn out to be useless for most purposes, as we will see later on. From a philosophical viewpoint, principal ultrafilters fail to capture the idea that sets belonging to the ultrafilter are large, for $\{s\}$ belongs to the ultrafilter \mathcal{U}_s and yet hardly anyone would dare say that the set $\{s\}$ is large!

Exercise 1.1.15. Prove that an ultrafilter \mathcal{U} on S is principal if and only if there is a *finite* set $A \subseteq S$ such that $A \in \mathcal{U}$. In particular, every ultrafilter on a finite set is principal.

Exercise 1.1.16. Suppose that \mathcal{F} is a filter on S . Then \mathcal{F} is an ultrafilter on S if and only if it is a maximal filter, that is, if and only if, whenever \mathcal{F}' is a filter on S such that $\mathcal{F} \subseteq \mathcal{F}'$, we have $\mathcal{F} = \mathcal{F}'$.

Fix a filter \mathcal{F} . Since it is readily verified that the union of an increasing chain of filters on S containing \mathcal{F} is once again a filter on S containing \mathcal{F} , the previous exercise and Zorn's lemma (see Appendix B) yield the following:

Corollary 1.1.17. *Given any filter \mathcal{F} on S , there is an ultrafilter \mathcal{U} on S such that $\mathcal{F} \subseteq \mathcal{U}$.*

We refer to the previous statement as the **ultrafilter theorem for S** . By the ultrafilter theorem, we mean the statement that the ultrafilter theorem for S holds for every set S . We will have a lot more to say about the ultrafilter theorem from a foundational perspective in Chapter 5. For now, we note that, by applying the ultrafilter theorem for S to the Fréchet filter on S (when S is infinite), we obtain the following:

Corollary 1.1.18. *If S is infinite, then there is a nonprincipal ultrafilter on S .*

Exercise 1.1.19. Suppose that S is an infinite set and \mathcal{D} is a collection of subsets of S such that $D_1 \cap \cdots \cap D_n$ is infinite for any finitely many $D_1, \dots, D_n \in \mathcal{D}$. Prove that there is a nonprincipal ultrafilter \mathcal{U} on S such that $\mathcal{D} \subseteq \mathcal{U}$.

Exercise 1.1.20. Prove that there is an ultrafilter \mathcal{U} on \mathbb{N} such that, for every $A \in \mathcal{U}$, we have that $\sum_{n \in A} \frac{1}{n}$ diverges.

Definition 1.1.21. An ultrafilter \mathcal{U} on I is called **uniform** if $|A| = |I|$ for every $A \in \mathcal{U}$.

Exercise 1.1.22. Suppose that \mathcal{U} is an ultrafilter on I and $J \in \mathcal{U}$ has minimal cardinality (amongst sets in \mathcal{U}). Prove that $\mathcal{U} \cap J$ is a uniform ultrafilter on J .

1.2. The ultrafilter quantifier

In this section, given $A \subseteq S$, we view A both as a subset of S and a relation on S , and thus the expressions “ $s \in A$ ” and “ $A(s)$ ” are synonymous.

Definition 1.2.1. Given a set S , a subset A of S , and an ultrafilter \mathcal{U} on S , we write $(\mathcal{U}s)A(s)$ if $A \in \mathcal{U}$, and we say that “ \mathcal{U} -almost all s in S satisfy $A(s)$ ”.

Remark 1.2.2. Using the notation from Exercise 1.1.9, we have that $(\mathcal{U}s)A(s)$ holds if and only if $A(s)$ holds $\mu_{\mathcal{U}}$ -almost everywhere.

We think of the formation $\mathcal{U}s$ as a quantifier of sorts. We can translate many of the basic properties of ultrafilters into properties of the ultrafilter quantifier.

Exercise 1.2.3. Prove the following properties of the ultrafilter quantifier:

- (1) $(\forall s A(s)) \Rightarrow (\mathcal{U}s)A(s)$.
- (2) $\neg((\mathcal{U}s)A(s)) \Leftrightarrow (\mathcal{U}s)(\neg A(s))$.
- (3) $(\mathcal{U}s)(A(s) \wedge B(s)) \Leftrightarrow ((\mathcal{U}s)A(s) \wedge (\mathcal{U}s)B(s))$.
- (4) $(\mathcal{U}s)(A(s) \vee B(s)) \Leftrightarrow ((\mathcal{U}s)A(s) \vee (\mathcal{U}s)B(s))$.

One must take care in manipulations with the ultrafilter quantifier as it does not always behave like its more familiar counterparts \forall and \exists :

Exercise 1.2.4. Fix an ultrafilter \mathcal{U} on \mathbb{N} .

- (1) If \mathcal{U} is nonprincipal, prove that the quantifiers $\mathcal{U}s$ and $\mathcal{U}t$ do not commute, that is, there is $A \subseteq \mathbb{N}^2$ such that $(\mathcal{U}s)(\mathcal{U}t)A(s, t)$ but $\neg(\mathcal{U}t)(\mathcal{U}s)A(s, t)$.
- (2) For any $n \in \mathbb{N}$ and any $B \subseteq \mathbb{N}^2$, prove that $(\mathcal{U}s)(\mathcal{U}_n t)B(s, t)$ holds if and only if $(\mathcal{U}_n t)(\mathcal{U}s)B(s, t)$ holds. In other words, the ultrafilter quantifier corresponding to a principal ultrafilter commutes with any other ultrafilter quantifier.

At first glance, it might seem that what we have defined is a sort of universal quantifier. Temporarily, let us rewrite our quantifier as $\forall^{\mathcal{U}}$, that is, $(\forall^{\mathcal{U}}s)A(s)$ holds precisely when $\{s \in S : A(s) \text{ holds}\} \in \mathcal{U}$. In analogy with the usual quantifiers, one might be tempted to then define the corresponding existential quantifier $\exists^{\mathcal{U}}$ by declaring $(\exists^{\mathcal{U}}s)A(s)$ holds if and only if it is not the case that $(\forall^{\mathcal{U}}s)A(s)$ fails, or symbolically, $\exists^{\mathcal{U}}s = \neg\forall^{\mathcal{U}}s\neg$.

Exercise 1.2.5. Prove that the quantifier $\exists^{\mathcal{U}}$ coincides with the quantifier $\forall^{\mathcal{U}}$, that is, for any set S , any subset $A \subseteq S$, and any ultrafilter \mathcal{U} on S , prove that $(\exists^{\mathcal{U}}s)A(s)$ holds if and only if $(\forall^{\mathcal{U}}s)A(s)$ holds.

For this reason, we only consider the single ultrafilter quantifier introduced above.

1.3. The category of ultrafilters

In this section, we define what it means for two ultrafilters to be isomorphic. Naïvely speaking, one might expect ultrafilters \mathcal{U} and \mathcal{V} on index sets S and T to be isomorphic if there is a bijection $f : S \rightarrow T$ such that, for $A \subseteq S$,

we have $A \in \mathcal{U}$ if and only if $f(A) \in \mathcal{V}$. However, this turns out to be a bit too restrictive and does not quite capture the spirit of things needing only to occur on “large” sets.

To aid ourselves in coming up with the correct notion of isomorphic ultrafilters, it behooves us to take a categorical perspective on the matter. (The reader unfamiliar with basic category theory may consult Appendix C.)

Definition 1.3.1.

- (1) If $f : S \rightarrow T$ is a function and \mathcal{U} is an ultrafilter on S , then the **pushforward of \mathcal{U} along f** is the ultrafilter $f(\mathcal{U})$ on T defined by setting, for $A \subseteq T$:

$$A \in f(\mathcal{U}) \Leftrightarrow f^{-1}(A) \in \mathcal{U}.$$

- (2) Given two functions $f, f' : S \rightarrow T$, we say that f and f' are **equal modulo \mathcal{U}** , written $f =_{\mathcal{U}} f'$, if $(\mathcal{U}s)(f(s) = f'(s))$.

Remark 1.3.2. Using the notation from Exercise 1.1.9, the pushforward ultrafilter $f(\mathcal{U})$ is the ultrafilter of measure 1 sets corresponding to the pushforward measure $f^*\mu_{\mathcal{U}}$ on T .

Exercise 1.3.3. Suppose that $f : S \rightarrow T$ is a function and $s \in S$. Prove that $f(\mathcal{U}_s) = \mathcal{U}_{f(s)}$.

Exercise 1.3.4.

- (1) Prove that $=_{\mathcal{U}}$ is an equivalence relation on functions from S to T . We denote the equivalence class of $f : S \rightarrow T$ by $[f]_{\mathcal{U}}$.
- (2) Prove that if $f =_{\mathcal{U}} f'$, then $f(\mathcal{U}) = f'(\mathcal{U})$.
- (3) If $f : S \rightarrow T$ and $g : T \rightarrow U$ are functions and \mathcal{U} is an ultrafilter on S , prove that $(g \circ f)(\mathcal{U}) = g(f(\mathcal{U}))$.
- (4) If $f, f' : S \rightarrow T$ and $g, g' : T \rightarrow U$ are functions, \mathcal{U} is an ultrafilter on S , and $f =_{\mathcal{U}} f'$ and $g =_{f(\mathcal{U})} g'$, prove that $g \circ f =_{\mathcal{U}} g' \circ f'$.

Definition 1.3.5. Given two ultrafilters \mathcal{U} and \mathcal{V} on sets S and T , respectively, a **morphism** between \mathcal{U} and \mathcal{V} is an equivalence class $[f]_{\mathcal{U}}$ such that $f(\mathcal{U}) = \mathcal{V}$. As is customary in category theory, we write $[f]_{\mathcal{U}} : \mathcal{U} \rightarrow \mathcal{V}$ if $[f]_{\mathcal{U}}$ is a morphism.

By Exercise 1.3.4(2), this notion is well-defined, that is, independent of representative. Moreover, Exercise 1.3.4(3) allows us to unambiguously define the composition of two morphisms $[f]_{\mathcal{U}}$ and $[g]_{\mathcal{V}}$ to be $[g]_{\mathcal{V}} \circ [f]_{\mathcal{U}} := [g \circ f]_{\mathcal{U}}$ (defined when $f(\mathcal{U}) = \mathcal{V}$). It is easy to see that this notion of composition is associative. Moreover, denoting the identity function on S

by id_S , we see that $[\text{id}_S]_{\mathcal{U}}$ is an identity morphism for \mathcal{U} . Summarizing, we thus have:

Theorem 1.3.6. *The collection of all ultrafilters equipped with the above notion of morphism forms a category.*

Now that we have a category, we obtain a natural notion of isomorphism between ultrafilters:

Definition 1.3.7. If \mathcal{U} and \mathcal{V} are ultrafilters on index sets S and T , we say that \mathcal{U} and \mathcal{V} are **isomorphic**, denoted $\mathcal{U} \cong \mathcal{V}$, if they are isomorphic in the category-theoretic sense, that is, if there are morphisms $[f]_{\mathcal{U}} : \mathcal{U} \rightarrow \mathcal{V}$ and $[g]_{\mathcal{V}} : \mathcal{V} \rightarrow \mathcal{U}$ such that $[g]_{\mathcal{V}} \circ [f]_{\mathcal{U}} = [\text{id}_S]_{\mathcal{U}}$ and $[f]_{\mathcal{U}} \circ [g]_{\mathcal{V}} = [\text{id}_T]_{\mathcal{V}}$.

Exercise 1.3.8. Prove that any two principal ultrafilters are isomorphic but that a principal ultrafilter is never isomorphic to a nonprincipal ultrafilter.

Remark 1.3.9. A particular consequence of the preceding exercise is that ultrafilters on index sets of different cardinalities can still be isomorphic.

It is desirable to have a more concrete description of isomorphic ultrafilters that avoids category-theoretic language. An essential tool in this endeavor is the following:

Theorem 1.3.10. *If \mathcal{U} is a nonprincipal ultrafilter on S , then the only morphism from \mathcal{U} to itself is $[\text{id}_S]_{\mathcal{U}}$. In other words, if $f : S \rightarrow S$ is such that $f(\mathcal{U}) = \mathcal{U}$, then $f =_{\mathcal{U}} \text{id}_S$.*

To prove Theorem 1.3.10, we need to prove a combinatorial fact. First, we establish a piece of notation that will be used many times throughout this book:

Notation 1.3.11. Given any set S , we let $\mathcal{P}_f(S)$ denote the set of *finite* subsets of S .

Lemma 1.3.12. *Suppose that S is infinite and $g : S \rightarrow S$ is fixed-point free, that is, $g(s) \neq s$ for all $s \in S$. Then there is a partition of $S := S_1 \cup S_2 \cup S_3$ such that, for all $n = 1, 2, 3$, if $s \in S_n$, then $g(s) \notin S_n$.*

Proof. We first establish the following:

Claim. *For every finite subset $F \subseteq S$, there is a partition $F = S_{1,F} \cup S_{2,F} \cup S_{3,F}$ such that, for all $s \in F$ and $n = 1, 2, 3$, if $s \in S_{n,F}$ and $g(s) \in F$, then $g(s) \notin S_{n,F}$.*

Proof of Claim. We prove the Claim by induction on $|F|$. The Claim is obvious when $|F| = 1$. Now suppose that $|F| > 1$ and the Claim has been proven for all finite sets of smaller size. By the pigeonhole principle, there

is $s \in F$ such that there is at most one $t \in F$ such that $g(t) = s$. Set $G := F \setminus \{s\}$. By induction, we may find a partition $G = S_{1,G} \cup S_{2,G} \cup S_{3,G}$ as in the statement of the Claim. We now take $n \in \{1, 2, 3\}$ such that $g(s), g(t) \notin S_n$, where t is the unique element of G such that $g(t) = s$, should it exist. Set $S_{n,F} = S_{n,G} \cup \{s\}$ for this n and set $S_{m,F} := S_{m,G}$ for $m \in \{1, 2, 3\} \setminus \{n\}$. It is clear that this partition of F is as desired. Thus, the Claim is proven.

For each $s \in S$, let $A_s := \{F \in \mathcal{P}_f(S) : s \in F\}$. Since the family $(A_s)_{s \in S}$ has the finite intersection property (as $\{s_1, \dots, s_n\} \subseteq \bigcap_{i=1}^n A_{s_i}$ for any finitely many $s_1, \dots, s_n \in S$), there is an ultrafilter \mathcal{U} on $\mathcal{P}_f(S)$ such that $A_s \in \mathcal{U}$ for all $s \in S$. We use \mathcal{U} to define a partition $S = S_1 \cup S_2 \cup S_3$ as follows. Given $s \in S$, put $s \in S_n$ if and only if n is the unique number in $\{1, 2, 3\}$ such that $s \in S_{n,F}$ for \mathcal{U} -almost all $F \in A_s$, where $S_{n,F}$ is as in the Claim. Note that this partition is as desired: if $g(s) \in S_n$ as well, then for \mathcal{U} -almost all $F \in \mathcal{P}_f(S)$, we have $s, g(s) \in F$ and $s, g(s) \in S_{n,F}$, contradicting the choice of partition $S_{n,F}$. \square

Proof of Theorem 1.3.10. Suppose that $f(\mathcal{U}) = \mathcal{U}$ and yet, toward a contradiction, that $f \neq_{\mathcal{U}} \text{id}_S$. Thus, there is $A \in \mathcal{U}$ such that $f(s) \neq s$ for all $s \in A$. Let $g : S \rightarrow S$ be fixed-point free and such that $g(s) = f(s)$ for all $s \in A$, whence $g =_{\mathcal{U}} f$. Since \mathcal{U} is nonprincipal (by assumption), S is infinite. Let $S = S_1 \cup S_2 \cup S_3$ be the partition of S as guaranteed by Lemma 1.3.12. By Exercise 1.1.10, there is a unique n such that $S_n \in \mathcal{U}$. Since $f(\mathcal{U}) = \mathcal{U}$, we have that $f^{-1}(S_n) \in \mathcal{U}$, whence $A \cap S_n \cap f^{-1}(S_n) \in \mathcal{U}$. In particular, there is $s \in A \cap S_n$ such that $f(s) \in S_n$, contradicting the defining property of the partition and the fact that $f(s) = g(s)$ for this particular s . \square

Corollary 1.3.13. *Suppose that \mathcal{U} and \mathcal{V} are ultrafilters on the sets I and J , respectively, and that there are morphisms $[f]_{\mathcal{U}} : \mathcal{U} \rightarrow \mathcal{V}$ and $[g]_{\mathcal{V}} : \mathcal{V} \rightarrow \mathcal{U}$.*

- (1) $[f]_{\mathcal{U}}$ and $[g]_{\mathcal{V}}$ are inverse isomorphisms.
- (2) $[f]_{\mathcal{U}}$ is the only morphism from \mathcal{U} to \mathcal{V} , and $[g]_{\mathcal{V}}$ is the only morphism from \mathcal{V} to \mathcal{U} .

Proof. Item (1) is clear when \mathcal{U} and \mathcal{V} are principal, and it follows immediately from Theorem 1.3.10 when they are nonprincipal. For (2), notice that if $[f']_{\mathcal{U}} : \mathcal{U} \rightarrow \mathcal{V}$ were also a morphism, then by item (1) (applied to $[f']_{\mathcal{U}}$ and $[g]_{\mathcal{V}}$), we would have that $[f']_{\mathcal{U}}$ is the inverse of $[g]_{\mathcal{V}}$; by uniqueness of inverses, we have that $[f]_{\mathcal{U}} = [f']_{\mathcal{U}}$. The analogous statement for $[g]_{\mathcal{V}}$ follows by the same argument. \square

Corollary 1.3.14. *If \mathcal{U} is an ultrafilter on S and $f : S \rightarrow T$ is a function, then $[f]_{\mathcal{U}}$ is an isomorphism if and only if there is $A \in \mathcal{U}$ such that $f \upharpoonright A$ is injective.*

Proof. First suppose that $[f]_{\mathcal{U}}$ is an isomorphism with inverse $[g]_{\mathcal{V}}$, where $\mathcal{V} := f(\mathcal{U})$. Let

$$A := \{s \in S : (g \circ f)(s) = s\}.$$

Since $g \circ f =_{\mathcal{U}} \text{id}_S$, we have that $A \in \mathcal{U}$. It remains to note that $f \upharpoonright A$ is injective.

Conversely, suppose that $A \in \mathcal{U}$ is such that $f \upharpoonright A$ is injective. Let $g : T \rightarrow S$ be such that $g \upharpoonright f(A)$ is the inverse of the bijection $f \upharpoonright A : A \rightarrow f(A)$. We leave it to the reader to check that $[g]_{\mathcal{V}}$ is the inverse of $[f]_{\mathcal{U}}$. \square

Exercise 1.3.15. Finish the proof of the previous corollary by checking that $[g]_{\mathcal{V}}$ is indeed the inverse of $[f]_{\mathcal{U}}$.

We can now provide a more concrete description of isomorphic ultrafilters:

Corollary 1.3.16. *If \mathcal{U} and \mathcal{V} are ultrafilters on sets S and T , respectively, then $\mathcal{U} \cong \mathcal{V}$ if and only if there is a function $f : S \rightarrow T$ such that $f(\mathcal{U}) = \mathcal{V}$ and for which there is $A \in \mathcal{U}$ such that $f \upharpoonright A$ is injective.*

Note that, in the notation from the previous corollary, we have that $\mathcal{U} \cong \mathcal{V}$ if and only if $f \upharpoonright A$ witnesses that $\mathcal{U} \cap A$ and $\mathcal{V} \cap f(A)$ are isomorphic in the naïve sense introduced in the beginning of this section (where $\mathcal{U} \cap A := \{B \cap A : B \in \mathcal{U}\}$ is the ultrafilter on A induced by \mathcal{U} as introduced in Exercise 1.1.12 and likewise for $\mathcal{V} \cap f(A)$).

1.4. The number of ultrafilters

Fix an infinite cardinal κ . Since the set of ultrafilters on κ is a subset of $\mathcal{P}(\mathcal{P}(\kappa))$, a naïve upper bound for the cardinality of the set of ultrafilters on κ is 2^{2^κ} . In this section, we show that this upper bound is actually achieved:

Theorem 1.4.1. *For any infinite cardinal κ , there are 2^{2^κ} many ultrafilters on κ .*

The plan of the proof is as follows. For each $C \subseteq 2^\kappa$, we would like to construct an ultrafilter $\mathcal{U}(C)$ on κ such that $C_1 \neq C_2$ implies $\mathcal{U}(C_1) \neq \mathcal{U}(C_2)$, whence there will be at least (and hence exactly) 2^{2^κ} many ultrafilters on κ . We start by momentarily fixing a particular subset $X \subseteq 2^\kappa$ (assumptions on which will be forthcoming) and defining, for $C \subseteq X$,

$$\mathcal{B}(C) := \{f^{-1}(0) : f \in C\} \cup \{f^{-1}(1) : f \in X \setminus C\}.$$

In this display, we are viewing elements of 2^κ not as subsets of κ but rather as functions $\kappa \rightarrow 2 = \{0, 1\}$.

If $\mathcal{B}(C)$ were to have the finite intersection property, then we can extend $\mathcal{B}(C)$ to an ultrafilter $\mathcal{U}(C)$ on κ . It then remains to show that $C_1 \neq C_2$ implies $\mathcal{U}(C_1) \neq \mathcal{U}(C_2)$. Without loss of generality, we may take $f \in C_1 \setminus C_2$. Then $f^{-1}(0) \in \mathcal{B}(C_1) \subseteq \mathcal{U}(C_1)$ and $f^{-1}(1) \in \mathcal{B}(C_2) \subseteq \mathcal{U}(C_2)$; since $f^{-1}(0) \cap f^{-1}(1) = \emptyset$, this shows that $\mathcal{U}(C_1) \neq \mathcal{U}(C_2)$.

Consequently, the above proof hinges on the sets $\mathcal{B}(C)$ having the finite intersection property. For this to happen, we would need to be able to take arbitrary $f_1, \dots, f_m \in C$ and $g_1, \dots, g_n \in X \setminus C$ and find $x \in \kappa$ such that $f_i(x) = 0$ and $g_j(x) = 1$ for each i, j , whence it would follow that $x \in \bigcap_{i=1}^m f_i^{-1}(0) \cap \bigcap_{j=1}^n f_j^{-1}(1)$. Unfortunately, this statement is not true for a general subset X of 2^κ . (Exercise!) Thankfully, we can show that there is a set $X \subseteq 2^\kappa$ such that $|X| = 2^\kappa$ and such that the sets $\mathcal{B}(C)$ do have the finite intersection property whenever $C \subseteq X$, whence the above proof can be rescued.

Definition 1.4.2. If A is a set, then a set of functions $X \subseteq 2^A$ is **independent** if, for any finitely many distinct functions $f_1, \dots, f_n \in X$ and finitely many elements $y_1, \dots, y_n \in \{0, 1\}$, there is $x \in A$ such that $f_i(x) = y_i$ for $i = 1, \dots, n$.

Remark 1.4.3. The terminology in the previous definition is motivated by the fact that, given $f_1, \dots, f_n \in X$, each of the 2^n possible intersections $\bigcap_{i=1}^n f_i^{-1}(y_i)$, as (y_1, \dots, y_n) ranges over 2^n , is nonempty, whence these intersections are independent in the sense of Venn diagrams.

From the above discussion, Theorem 1.4.1 will follow from the following theorem:

Theorem 1.4.4. *For any infinite cardinal κ , there is an independent set $X \subseteq 2^\kappa$ with $|X| = 2^\kappa$.*

Proof. By set-theoretic trickery, it will suffice to find a set B with $|B| = \kappa$ and an independent set $X \subseteq 2^B$ with $|X| = 2^\kappa$. Here is the B that will work:

$$B := \{(F, G, s) : F \subseteq \kappa \text{ is finite, } G \subseteq \mathcal{P}(F), s \in 2^G\}.$$

It is an easy exercise to see that $|B| = \kappa$. For $A \subseteq \kappa$, consider $f_A \in 2^B$ defined by

$$f_A(F, G, s) = \begin{cases} s(A \cap F) & \text{if } A \cap F \in G \\ 0 & \text{otherwise.} \end{cases}$$

We note that the function $A \mapsto f_A$ is injective: if $A_1 \neq A_2$, we may take, without loss of generality, $x \in A_1 \setminus A_2$. Let $F := \{x\}$, $G = \{F\}$, and $s(F) = 1$. Then $f_{A_1}(F, G, s) = 1$ while $f_{A_2}(F, G, s) = 0$, so $f_{A_1} \neq f_{A_2}$.

Thus $\{f_A : A \subseteq \kappa\}$ is a subset of 2^B of cardinality 2^κ . It remains to see that it is an independent set. Toward this end, fix finitely many distinct subsets A_1, \dots, A_n of κ and $y_1, \dots, y_n \in \{0, 1\}$. For each $1 \leq l < m \leq n$, take $a_{l,m} \in A_l \triangle A_m$. Let F be the set of $a_{l,m}$'s thus obtained, and let $G = \{A_m \cap F : 1 \leq m \leq n\}$. Note that $A_l \cap F \neq A_m \cap F$ for $l \neq m$ as $a_{l,m}$ belongs to one but not the other. We are thus allowed to unambiguously define $s : G \rightarrow \{0, 1\}$ by $s(A_m \cap F) := y_m$. By definition, $f_{A_m}(F, G, s) = y_m$, as desired. \square

Since there are exactly κ many principal ultrafilters on κ , we immediately obtain:

Corollary 1.4.5. *For any infinite cardinal κ , there are 2^{2^κ} many nonprincipal ultrafilters on κ .*

Exercise 1.4.6. Prove that there are 2^{2^κ} many nonisomorphic nonprincipal ultrafilters on κ .

1.5. The ultrafilter number \mathfrak{u}

Recall the definition for a base for a filter given in Definition 1.1.3.

Definition 1.5.1. The **ultrafilter number** \mathfrak{u} is the minimum of all cardinals κ for which there is a nonprincipal ultrafilter on \mathbb{N} with a base of cardinality κ .

Exercise 1.5.2. Prove that $\aleph_1 \leq \mathfrak{u} \leq \mathfrak{c}$.

By the preceding exercise, if the Continuum Hypothesis (CH) holds, then $\mathfrak{u} = \aleph_1 = \mathfrak{c}$. Thus, it is only interesting to consider \mathfrak{u} in the case that CH fails. Under the negation of CH, anything can happen, namely there are models of the negation of CH where $\mathfrak{u} = \aleph_1$ (see, for example, [5]), where $\mathfrak{u} = \mathfrak{c}$ (e.g., any model of Martin's axiom, see [89, Section 23]), and where $\aleph_1 < \mathfrak{u} < \mathfrak{c}$ [15].

The ultrafilter number \mathfrak{u} is an example of a so-called **cardinal characteristic of the continuum**, which, roughly speaking, is an example of a cardinality reflecting some combinatorial property that holds for \mathfrak{c} but not for \aleph_0 . In general, one aims to see what comparisons hold between these cardinal characteristics in Zermelo-Fraenkel set theory with choice (ZFC) and which comparisons depend on further axioms. To give a feel for this area, we give one such comparison as an exercise. First, a definition:

Definition 1.5.3. The **unbounding number** \mathfrak{b} is the minimal cardinality of a family $X \subseteq \omega^\omega$ for which there does *not* exist $g \in \omega^\omega$ with the property that, for all $f \in \omega^\omega$, there is $N \in \omega$ such that $f(n) \leq g(n)$ for all $n \geq N$.

Exercise 1.5.4. Show that $\mathfrak{b} \leq \mathfrak{u}$. (*Hint.* For each X in a base \mathcal{B} for a nonprincipal ultrafilter on \mathbb{N} , define $f_X \in \omega^\omega$ by $f_X(n) :=$ the least $m \geq n$ that belongs to X . Prove that $(f_X)_{X \in \mathcal{B}}$ is unbounded.)

1.6. The Rudin-Keisler order

In this section, we introduce the Rudin-Keisler order \leq_{RK} on the collection of all ultrafilters. Roughly speaking, $\mathcal{V} \leq_{RK} \mathcal{U}$ indicates that \mathcal{V} is “no more complicated” than \mathcal{U} . It turns out that we have already met this notion in Section 1.3:

Definition 1.6.1. If \mathcal{U} and \mathcal{V} are ultrafilters on the sets S and T , respectively, then we say that \mathcal{V} is below \mathcal{U} in the **Rudin-Keisler order**, denoted $\mathcal{V} \leq_{RK} \mathcal{U}$, if there is a morphism from \mathcal{U} to \mathcal{V} in the category of ultrafilters, that is, if there is a function $f : S \rightarrow T$ such that $f(\mathcal{U}) = \mathcal{V}$. We also write $\mathcal{U} <_{RK} \mathcal{V}$ to mean $\mathcal{U} \leq_{RK} \mathcal{V}$ but $\mathcal{V} \not\leq_{RK} \mathcal{U}$.

At first glance, it is not clear how this definition matches up with our rough description at the beginning of this section. Let us take a moment to elaborate further.

Exercise 1.6.2. Suppose that $f : S \rightarrow T$ and $g_1, g_2 : T \rightarrow U$ are functions and \mathcal{U} and \mathcal{V} are ultrafilters on S and T , respectively, such that $f(\mathcal{U}) = \mathcal{V}$ and $[g_1]_{\mathcal{V}} \circ [f]_{\mathcal{U}} = [g_2]_{\mathcal{V}} \circ [f]_{\mathcal{U}}$. Prove that $[g_1]_{\mathcal{V}} = [g_2]_{\mathcal{V}}$.

In category-theoretic terminology (see Appendix C), the above exercise says that every morphism in the category of ultrafilters is an epimorphism. In the category of sets, the epimorphisms are exactly the surjections (Exercise!), whence the existence of an epimorphism from a set X to a set Y is an indication that the set Y is “no larger than” or “no more complicated than” X . It is for this reason that the existence of a morphism from \mathcal{U} to \mathcal{V} (which is automatically an epimorphism) is an indication that \mathcal{V} is no more complicated than \mathcal{U} .

Another heuristic behind the definition of \leq_{RK} is that the function f takes queries about whether or not a subset A of T belongs to \mathcal{V} and converts it into the question of whether or not $f^{-1}(A)$ belongs to \mathcal{U} . Thus, with total knowledge about \mathcal{U} , one can answer all queries about \mathcal{V} , whence \mathcal{V} is no more complicated than \mathcal{U} .

Exercise 1.6.3. Suppose that \mathcal{V} is a principal ultrafilter. Prove that $\mathcal{V} \leq_{RK} \mathcal{U}$ for any ultrafilter \mathcal{U} .

It is clear that \leq_{RK} is reflexive ($\mathcal{U} \leq_{RK} \mathcal{U}$) and transitive ($\mathcal{U} \leq_{RK} \mathcal{V}$ and $\mathcal{V} \leq_{RK} \mathcal{W}$ implies $\mathcal{U} \leq_{RK} \mathcal{W}$), whence \leq_{RK} is a preorder on the set of ultrafilters. It is not a partial order as $\mathcal{U} \leq_{RK} \mathcal{V}$, and $\mathcal{V} \leq_{RK} \mathcal{U}$ does not

imply that $\mathcal{U} = \mathcal{V}$. Following usual nomenclature with preorders, we write $\mathcal{U} \equiv_{RK} \mathcal{V}$ to mean that $\mathcal{U} \leq_{RK} \mathcal{V}$ and $\mathcal{V} \leq_{RK} \mathcal{U}$. However, this notion is not new to us, for Corollary 1.3.13(1) implies the following:

Corollary 1.6.4. *For any two ultrafilters \mathcal{U} and \mathcal{V} , we have that $\mathcal{U} \equiv_{RK} \mathcal{V}$ if and only if $\mathcal{U} \cong \mathcal{V}$.*

Consequently, \leq_{RK} induces a partial order on the set of isomorphism classes of ultrafilters. In what follows, we often blur this distinction and speak of \leq_{RK} both as the preorder on ultrafilters and the partial order on the set of isomorphism classes of ultrafilters.

By Exercises 1.3.8 and 1.6.3, there is a unique isomorphism class that is a minimum under the ordering \leq_{RK} , namely the isomorphism class of principal ultrafilters. In what follows, we discard this (uninteresting) class and consider only the partial ordering on classes of nonprincipal ultrafilters.

The following questions naturally arise:

Question 1.6.5.

- (1) Is \leq_{RK} linear?
- (2) Is there a \leq_{RK} -maximal element?
- (3) Is there a \leq_{RK} -minimal element?

The answers are: No, no, and maybe!

Fact 1.6.6 (Rudin and Shelah [162]). There are 2^{2^κ} many \leq_{RK} -incomparable elements in $\beta\kappa$.

The combinatorics involved in this result are quite intricate. Instead, we prove here the following easier result. Note that we make an extra (simplifying) set-theoretic assumption in the statement of the next result, while the previous fact is indeed a theorem of ZFC.

Proposition 1.6.7. *Assume that $\mathfrak{u} = \mathfrak{c}$. Then there are nonprincipal $\mathcal{U}, \mathcal{V} \in \beta\mathbb{N}$ that are \leq_{RK} -incomparable.*

Proof. Let $(f_\alpha : \alpha < \mathfrak{c})$ enumerate all elements of ω^ω . We construct filters \mathcal{F}_α and \mathcal{G}_α on \mathbb{N} with the following properties:

- (1) $\mathcal{F}_0 = \mathcal{G}_0 =$ the Fréchet filter on \mathbb{N} ;
- (2) If $\beta < \alpha$, then $\mathcal{F}_\beta \subseteq \mathcal{F}_\alpha$ and $\mathcal{G}_\beta \subseteq \mathcal{G}_\alpha$;
- (3) \mathcal{F}_α has cardinality $< \mathfrak{c}$;
- (4) If \mathcal{U} and \mathcal{V} are any ultrafilters containing \mathcal{F}_α and \mathcal{G}_α , respectively, then $f_\alpha(\mathcal{U}) \neq \mathcal{V}$ and $f_\alpha(\mathcal{V}) \neq \mathcal{U}$.

If we let \mathcal{U} and \mathcal{V} be any ultrafilters extending $\bigcup_{\alpha} \mathcal{F}_{\alpha}$ and $\bigcup_{\alpha} \mathcal{G}_{\alpha}$, respectively, then it follows that \mathcal{U} and \mathcal{V} are \leq_{RK} -incomparable.

Suppose that \mathcal{F}_{β} and \mathcal{G}_{β} have been defined for all $\beta < \alpha$; we show how to define \mathcal{F}_{α} and \mathcal{G}_{α} . If $\alpha = \beta + 1$, set $\mathcal{F} := \mathcal{F}_{\beta}$. If α is a limit ordinal, set $\mathcal{F} := \bigcup_{\beta < \alpha} \mathcal{F}_{\beta}$. Make the analogous definition for \mathcal{G} . Since $|\mathcal{F}| < \mathfrak{c}$, the assumption that $\mathfrak{u} = \mathfrak{c}$ implies that \mathcal{F} is not an ultrafilter. Consequently, there is $A \subseteq \mathbb{N}$ such that $A, \mathbb{N} \setminus A \notin \mathcal{F}$. Note that, by Exercise 1.1.7, both $\mathcal{F} \cup \{A\}$ and $\mathcal{F} \cup \{\mathbb{N} \setminus A\}$ have the FIP. We consider three cases:

- If $f_{\alpha}^{-1}(A) \in \mathcal{G}$, then we set \mathcal{F}' to be the filter generated by $\mathcal{F} \cup \{\mathbb{N} \setminus A\}$ and we set $\mathcal{G}' := \mathcal{G}$.
- If $f_{\alpha}^{-1}(\mathbb{N} \setminus A) \in \mathcal{G}$, then we set \mathcal{F}' to be the filter generated by $\mathcal{F} \cup \{A\}$ and we set $\mathcal{G}' := \mathcal{G}$.
- If neither of the above two cases hold, then we may set \mathcal{F}' to be the filter generated by $\mathcal{F} \cup \{\mathbb{N} \setminus A\}$ and \mathcal{G}' to be the filter generated by $\mathcal{G} \cup \{f_{\alpha}^{-1}(A)\}$.

At this point, we have guaranteed that $f_{\alpha}(\mathcal{V}) \neq \mathcal{U}$ whenever \mathcal{U} is an ultrafilter extending \mathcal{F}' and \mathcal{V} is an ultrafilter extending \mathcal{G}' . Since $|\mathcal{G}'| < \mathfrak{c}$, we can reverse the procedure to find $\mathcal{F}_{\alpha} \supseteq \mathcal{F}'$ and $\mathcal{G}_{\alpha} \supseteq \mathcal{G}'$ as desired. \square

We now move on to the second question. First, some notation: given sets S and T , a subset $Y \subseteq S \times T$, and $t \in T$, we set $Y_t := \{s \in S : (s, t) \in Y\}$.

Definition 1.6.8. Suppose that \mathcal{U}, \mathcal{V} are ultrafilters on index sets S and T . We define the **product** of \mathcal{U} and \mathcal{V} to be

$$\mathcal{U} \times \mathcal{V} := \{Y \subseteq S \times T : \{t \in T : Y_t \in \mathcal{U}\} \in \mathcal{V}\}.$$

In other words, $Y \in \mathcal{U} \times \mathcal{V} \Leftrightarrow (\forall t)(\mathcal{U}s)Y(s, t)$.

Exercise 1.6.9. Prove that $\mathcal{U} \times \mathcal{V}$ is an ultrafilter on $S \times T$.

A particular consequence of the next proposition is that there is no \leq_{RK} -maximal ultrafilter.

Proposition 1.6.10. For all ultrafilters \mathcal{U} and \mathcal{V} , we have $\mathcal{U} <_{RK} \mathcal{U} \times \mathcal{V}$ and $\mathcal{V} <_{RK} \mathcal{U} \times \mathcal{V}$.

Proof. We only prove that $\mathcal{U} <_{RK} \mathcal{U} \times \mathcal{V}$, the other assertion being analogous. Let $\pi : S \times T \rightarrow S$ be the function $\pi(s, t) = s$. Note then that, for $A \subseteq S$, we have that $\pi^{-1}(A) = A \times T$ and $A \times T \in \mathcal{U} \times \mathcal{V}$ if and only if $A \in \mathcal{U}$. It follows that $\pi(\mathcal{U} \times \mathcal{V}) = \mathcal{U}$, so $\mathcal{U} \leq_{RK} \mathcal{U} \times \mathcal{V}$. If $\mathcal{U} \times \mathcal{V} \leq_{RK} \mathcal{U}$ were to hold, then by Corollary 1.3.13(1), $[\pi]_{\mathcal{U} \times \mathcal{V}}$ would be an isomorphism, contradicting Corollary 1.3.16 and the fact that $\pi \upharpoonright Y$ is never injective for any $Y \in \mathcal{U} \times \mathcal{V}$. \square

With more work, one can actually prove that every element of $\beta\kappa$ has 2^{2^κ} many successors in $\beta\kappa$; see [83, Theorem 11.9].

We address the last question in Section 5.4.

1.7. Notes and references

The notion of ultrafilter was introduced by H. Cartan [23, 24] in 1937 to study convergence in topological spaces, a topic we will study in Chapter 3. The existence of a nonprincipal ultrafilter was first proven by Tarski in [173]. Our approach to the ultrafilter quantifier is motivated by Todorčević's book [177]. Our treatment of the category of ultrafilters follows Blass's thesis [16]. The fact that there are the maximal possible number of ultrafilters on a given set is due to Pospíšil [144]. A nice survey on cardinal characteristics of the continuum, including more information about the ultrafilter number, is Blass's survey [14]. The Rudin-Keisler order was independently introduced by M. Rudin in [150] and by Keisler in lectures given at UCLA.

Arrow's theorem on fair voting

In this chapter, we give our first application of ultrafilters by proving Arrow's theorem, a classical result in voting theory. In Section 2.1, we introduce the statement of the result, while in Section 2.2 we explain the approach to proving Arrow's theorem via ultrafilters. In Section 2.3, we show how the notion of block voting reduces the proof of Arrow's theorem to its version for at most three voters, and we prove this latter statement in Section 2.4.

2.1. Statement of the theorem

Throughout this chapter, V denotes the set of **voters** in an election. The voters are ranking their preferences amongst a finite set of, say, n **candidates**, which we label, for the sake of simplicity, as $1, \dots, n$. They express their preference using a permutation σ of the set $\{1, \dots, n\}$. We will refer to the pair (V, n) as an **election**.

Example 2.1.1. If $n = 4$ and a voter has preference σ such that $\sigma(1) = 2$, $\sigma(2) = 4$, $\sigma(3) = 3$, and $\sigma(4) = 1$, then this means that they prefer candidate 4 the most, followed by candidate 1, then candidate 3, and finally they prefer candidate 2 the least.

Remark 2.1.2. While permutation notation for voters' preferences is mathematically natural, it leads to statements that appear counterintuitive. For example, in the previous example, the fact that $\sigma(1) < \sigma(2)$ actually signifies that voters prefer candidate 1 over candidate 2, the appearance of the $<$ symbol might indicate that the opposite preference held.

A **state of the election** is a function $\pi : V \rightarrow S_n$. (Recall that S_n is the set of permutations of the set $\{1, \dots, n\}$.) In other words, a state of the election is simply a record of the preferences of each voter involved: for $v \in V$, $\pi(v)$ is voter v 's preferences.

The following is the central question in this chapter: given a state of the election, how does one get a final ranking of the candidates that takes into account the individual voter rankings? Or even more to the point, one would like to determine, in advance, a method of turning any state of the election into a final ranking of the candidates. The following definition formalizes this idea:

Definition 2.1.3. An **election procedure** is a function $f : S_n^V \rightarrow S_n$, that is, for any state of the election π , $f(\pi)$ is the final ranking of the candidates. An election procedure f is called **fair** if it satisfies:

- (U) **unanimity**: if $\sigma \in S_n$ is such that $\pi(v) = \sigma$ for all $v \in V$, then $f(\pi) = \sigma$.
- (IA) **irrelevant alternatives**: if π and π' are states of the election and $i, j \in \{1, \dots, n\}$ are such that, for all $v \in V$, $\pi(v)(i) > \pi(v)(j)$ if and only if $\pi'(v)(i) > \pi'(v)(j)$, then $f(\pi)(i) > f(\pi)(j)$ if and only if $f(\pi')(i) > f(\pi')(j)$.

In English, unanimity expresses the fact that if all voters have the same preferences, then the outcome of the election is that common preference, while irrelevant alternatives says that the final ranking of any two candidates should only depend on how the voters feel about those two candidates.

Exercise 2.1.4. If f is a fair election procedure, prove that f further satisfies **local unanimity** (LU): if $\pi(v)(i) > \pi(v)(j)$ for all $v \in V$, then $f(\pi)(i) > f(\pi)(j)$.

At the (seemingly) opposite extreme of a fair election procedure is an election procedure f that possesses a **dictator**, which is a voter $v \in V$ such that, for every state of the election π , we have $f(\pi) = \pi(v)$; in other words, the outcome of the election is always v 's ranking of the candidates. It would seem that a fair election procedure would preclude the existence of a dictator. Such a sentiment is precisely why the following theorem of the economist Kenneth Arrow is so intriguing:

Theorem 2.1.5 (Arrow's theorem). *Suppose that V is a finite set of voters, $n \geq 3$, and $f : S_n^V \rightarrow S_n$ is a fair election procedure. Then there is a dictator for f .*

It is the goal of this chapter to prove Theorem 2.1.5.

Remark 2.1.6. Note that we must assume $n \geq 3$. Indeed, if $n = 2$ and (for simplicity) $|V|$ is odd, we can let $f : S_n^V \rightarrow S_n$ be the election procedure that just picks the candidate with the most votes. It is clear that this election procedure has no dictator.

2.2. The connection with ultrafilters

Arrow's original proof of Theorem 2.1.5 did not use ultrafilters, but the most transparent explanation of the story does. Indeed, let us first take up the question: how can one define an election procedure? One naïve idea is to just take the permutation that “appears most often” in the state of the election. If V is infinite, then what does one mean by the permutation that appears most often? Well, if we fix an ultrafilter \mathcal{U} on V , then there is a unique permutation $\sigma \in S_n$ such that $\{v \in V : \pi(v) = \sigma\} \in \mathcal{U}$ (we are using here that S_n is finite!). We can then define an election procedure $f_{\mathcal{U}}$ by setting $f_{\mathcal{U}}(\pi) :=$ this unique σ .

Proposition 2.2.1. *Given an ultrafilter \mathcal{U} on V , the election procedure $f_{\mathcal{U}}$ defined above is fair. Moreover, $v \in V$ is a dictator for $f_{\mathcal{U}}$ if and only if \mathcal{U} is the principal ultrafilter generated by v .*

Proof. (U) follows from the fact that $V \in \mathcal{U}$. To prove (IA), fix states of the election π and π' and candidates $i, j \in \{1, \dots, n\}$ such that, for all $v \in V$, $\pi(v)(i) > \pi(v)(j)$ if and only if $\pi'(v)(i) > \pi'(v)(j)$. Suppose, without loss of generality, that $f_{\mathcal{U}}(\pi)(i) > f_{\mathcal{U}}(\pi)(j)$. Then $\{v \in V : \pi(v)(i) > \pi(v)(j)\} \in \mathcal{U}$. Since this set is precisely the same as $\{v \in V : \pi'(v)(i) > \pi'(v)(j)\}$, it follows that $f_{\mathcal{U}}(\pi')(i) > f_{\mathcal{U}}(\pi')(j)$.

The second statement of the proposition is obvious from the definition of $f_{\mathcal{U}}$. \square

We thus have a function $\mathcal{U} \mapsto f_{\mathcal{U}}$ mapping the set of ultrafilters on V to the set of fair election procedures which maps principal ultrafilters to those election procedures possessing dictators. The key to the ultrafilter proof of Arrow's theorem is the assertion that this function is a bijection when $n \geq 3$:

Theorem 2.2.2. *If $n \geq 3$, the map $\mathcal{U} \mapsto f_{\mathcal{U}}$ from the set of ultrafilters on V to the set of fair election procedures for the election (V, n) is a bijection.*

Arrow's theorem follows immediately from Theorem 2.2.2:

Proof of Theorem 2.1.5. Suppose that V is a finite set of voters and $n \geq 3$. Suppose further that $f : S_n^V \rightarrow S_n$ is a fair election procedure. By Theorem 2.2.2, there is an ultrafilter \mathcal{U} on V such that $f = f_{\mathcal{U}}$. Since V is finite, by Exercise 1.1.15, \mathcal{U} is principal, whence $f = f_{\mathcal{U}}$ has a dictator. \square

Since Theorem 2.2.2 holds when V is infinite as well, we can conclude:

Corollary 2.2.3. *Suppose that V is an infinite set of voters. Then for any n , there is a fair election procedure for (V, n) that does not have a dictator.*

Proof. By Corollary 1.1.18, there is a nonprincipal ultrafilter \mathcal{U} on V . The corresponding election procedure $f_{\mathcal{U}}$ is fair and does not have a dictator. \square

Remark 2.2.4. The proof of the preceding corollary implicitly used the axiom of choice. In Chapter 5, we will see that there are models of set theory in which no nonprincipal ultrafilters exist, whence some form of the axiom of choice is needed in the previous result. Although not entirely accurate, the following phrase provides a humorous summary: In a universe without choice, there will always be a dictator!

We now work toward the proof of Theorem 2.2.2. We will proceed by defining an inverse to the function $\mathcal{U} \mapsto f_{\mathcal{U}}$. Until further notice, we fix an election (V, n) with $n \geq 3$ and a fair election procedure $f : S_n^V \rightarrow S_n$. Here is the key notion:

Definition 2.2.5. We call $F \subseteq V$ a **decisive set of voters for f** (or simply decisive for f) if, whenever there is a state of the election π and $\sigma \in S_n$ such that $\pi(v) = \sigma$ for all $v \in F$, we have $f(\pi) = \sigma$.

In other words, F is decisive for f if, whenever every member of F votes the same way, the outcome of the election procedure f is that common preference. Note that (U) states that V is decisive for f . Note also that, for $v \in V$, we have that $\{v\}$ is decisive for f precisely when v is a dictator for f .

Exercise 2.2.6. If \mathcal{U} is an ultrafilter on V and $f_{\mathcal{U}}$ is the associated election procedure, show that the decisive sets for $f_{\mathcal{U}}$ are precisely the elements of \mathcal{U} .

Let $\mathcal{U}_f := \{F \subseteq V : F \text{ is a decisive set of voters for } f\}$. The previous exercise can then be formulated as $\mathcal{U}_{f_{\mathcal{U}}} = \mathcal{U}$ and gives us a hint as to how to prove Theorem 2.2.2, namely we should prove the following theorem:

Theorem 2.2.7. \mathcal{U}_f is an ultrafilter on V .

Exercise 2.2.8. Suppose that Theorem 2.2.7 has been proven. Prove that $f_{\mathcal{U}_f} = f$.

Thus, if we can prove Theorem 2.2.7, then the map $f \mapsto \mathcal{U}_f$ is the desired inverse to $\mathcal{U} \mapsto f_{\mathcal{U}}$, finishing the proof of Theorem 2.2.2.

2.3. Block voting

We now work toward proving Theorem 2.2.7. Using Exercise 1.1.11, it suffices to prove the following statement:

(*) Whenever $n \geq 3$, f is a fair election procedure for (V, n) , and $V = V_1 \cup V_2 \cup V_3$ with V_1, V_2 , and V_3 pairwise disjoint (and perhaps some $V_i = \emptyset$), then exactly one V_i is decisive for f .

Set $\mathcal{P} := \{V_1, V_2, V_3\}$, a partition of V into at most three pieces. The key idea is to consider the new election (\mathcal{P}, n) (which has at most three voters). We can then associate to any state of the election π for (\mathcal{P}, n) a state of the election $\pi_{\mathcal{P}}$ for (V, n) by having all members of V_i vote in the same manner according to π , that is, for each $i = 1, 2, 3$ and each $v \in V_i$, $\pi_{\mathcal{P}}(v) = \pi(V_i)$; we refer to this situation as **block voting**. Note that any election procedure $f : S_n^V \rightarrow S_n$ gives rise to an election procedure $f_{\mathcal{P}} : S_n^{\mathcal{P}} \rightarrow S_n$ by defining $f_{\mathcal{P}}(\pi) := f(\pi_{\mathcal{P}})$.

The election (\mathcal{P}, n) is significantly simpler than the original election (V, n) as it has at most three voters; it turns out that one can prove this case directly by hand:

Theorem 2.3.1 (Arrow's theorem for few voters). *If $|V| \leq 3$ and $n \geq 3$, then every fair election procedure for (V, n) has a dictator.*

We will prove Theorem 2.3.1 in the next section. In the remainder of this section, we see how it implies statement (*) above.

Fix a set V of voters, $n \geq 3$, and a fair election procedure f for (V, n) . By Theorem 2.3.1, $f_{\mathcal{P}}$ has a dictator, say, without loss of generality, it is V_1 . We will show that V_1 is decisive for f . At first glance, this is not obvious. However, V_1 does have the seemingly weaker property of being *block decisive* for f :

Definition 2.3.2. $F \subseteq V$ is **block decisive** for f if, whenever π is a state of the election for (V, n) such that π is constantly σ on F and constantly σ' on $V \setminus F$, then $f(\pi) = \sigma$.

Exercise 2.3.3. Verify that V_1 as above is block decisive for f .

Now, a minor miracle occurs:

Proposition 2.3.4. *If $F \subseteq V$ is block decisive, then it is decisive.*

Proof. Suppose, toward a contradiction, that F is block decisive but not decisive. Since F is not decisive, there is a state of the election $\pi : V \rightarrow S_n$ and a permutation $\sigma \in S_n$ such that $\pi(v) = \sigma$ for all $v \in F$ and yet $f(\pi) \neq \sigma$. Take distinct $i, j \in \{1, \dots, n\}$ such that $\sigma(i) > \sigma(j)$ and yet $f(\pi)(i) < f(\pi)(j)$.

Since $n \geq 3$, we may consider $k \in \{1, \dots, n\} \setminus \{i, j\}$. Let $\sigma' \in S_n$ be such that $\sigma'(i) > \sigma'(k) > \sigma'(j)$. Let $\pi' : V \rightarrow S_n$ be a new state of the election such that:

- if $v \in F$, then $\pi'(v) = \sigma'$;
- if $v \notin F$, then
 - $\pi'(v)(k) < \pi'(v)(i)$,
 - $\pi'(v)(k) < \pi'(v)(j)$, and
 - $\pi'(v)(i) < \pi'(v)(j)$ if and only if $\pi(v)(i) < \pi(v)(j)$.

By (IA), we have that $f(\pi')(i) < f(\pi')(j)$. By (LU), we have $f(\pi')(i) > f(\pi')(k)$, whence $f(\pi')(j) > f(\pi')(k)$.

Finally, let $\sigma'' \in S_n$ be such that $\sigma''(j) > \sigma''(k)$, and define a new state of the election $\pi'' : V \rightarrow S_n$ such that $\pi''(v) = \sigma'$ for all $v \in F$ and $\pi''(v) = \sigma''$ for all $v \in V \setminus F$. By (IA) again, $f(\pi'')(j) > f(\pi'')(k)$. However, since F is block decisive, we have that $f(\pi'') = \sigma'$, so $f(\pi'')(k) > f(\pi'')(j)$, yielding the desired contradiction. \square

Combining Exercise 2.3.3 and Proposition 2.3.4, we have that V_1 is decisive. Thus, we have succeeded in proving statement (*) using Theorem 2.3.1 and, as mentioned above, statement (*) and Exercise 1.1.11 imply Theorem 2.2.7.

2.4. Finishing the proof

In this section, we finish the proof of Arrow's theorem by proving Theorem 2.3.1. It will behoove us to first prove the case of two voters. We do this in a series of lemmas. First, one final definition:

Definition 2.4.1. Given $v \in V$ and distinct $i, j \in \{1, \dots, n\}$, we call v a **decisive voter for f with respect to (i, j)** if, for every state of the election π for which $\pi(v)(i) < \pi(v)(j)$, we have $f(\pi)(i) < f(\pi)(j)$.

We also introduce some useful notation: when $|V| = 2$, we write $V = \{v, w\}$. In this case, if i, j, k are candidates and f is some fixed election procedure for (V, n) , we write $(ijk, ikj) \rightsquigarrow kij$ as an abbreviation for the statement: given a state π of the election (V, n) for which $\pi(v)(i) < \pi(v)(j) < \pi(v)(k)$ and $\pi(w)(i) < \pi(w)(k) < \pi(w)(j)$, we have $f(\pi)(k) < f(\pi)(i) < f(\pi)(j)$. Note that, by (IA), this notation is well-defined and independent of the choice of π . We use the same notation for all other possible permutations and also in the case of preferences for two candidates.

Lemma 2.4.2. *Suppose that $|V| = 2$, i, j, k are distinct candidates, and f is a fair election procedure for the election (V, n) . If v is decisive for f with respect to (i, j) , then v is decisive for f with respect to both (i, k) and (j, k) .*

Proof. By our assumption, we have $(ij, ji) \rightsquigarrow ij$. By (LU), we then have $(ijk, jki) \rightsquigarrow (ijk)$. By (IA), we then have $(ik, ki) \rightsquigarrow ik$, whence v is decisive for f with respect to (i, k) . Using (LU) again, we have that $(jik, kji) \rightsquigarrow jik$. By (IA) again, we then have $(jk, kj) \rightsquigarrow jk$, whence v is decisive for f with respect to (j, k) . \square

Lemma 2.4.3. *Suppose that $|V| = 2$, i and j are distinct candidates, and f is a fair election procedure for the election (V, n) . If v is decisive for f with respect to (i, j) , then v is decisive for f with respect to (j, i) .*

Proof. Since $n \geq 3$, we may consider $k \in \{1, \dots, n\} \setminus \{i, j\}$. By Lemma 2.4.2, we have that v is decisive for f with respect to (j, k) . By Lemma 2.4.2 again (applied to j, k, i), we obtain that v is decisive for f with respect to (j, i) . \square

Theorem 2.4.4 (Arrow's theorem for two voters). *Suppose that $|V| = 2$, $n \geq 3$, and f is a fair election procedure for (V, n) . Then there is a dictator for f .*

Proof. Fix distinct candidates i and j . By (LU) and (IA), we see that either v or w is decisive for f with respect to (i, j) . Without loss of generality, suppose it is v . We claim that v is a dictator for f . Indeed, this follows immediately from Lemmas 2.4.2 and 2.4.3. \square

Exercise 2.4.5. Verify the last sentence in the previous proof.

We are now ready to prove Theorem 2.3.1. In the proof, we extend our above notation so that when $|V| = 3$, we write $V = \{u, v, w\}$ and we extend our \rightsquigarrow notation in the obvious way.

Proof of Theorem 2.3.1. Suppose that $|V| = 3$, $n \geq 3$, and f is a fair election procedure for (V, n) . We define $V_u := \{u, vw\}$, where now vw is treated as one single voter (so $|V_u| = 2$) and a fair election procedure f_u for (V_u, n) as in the previous section on block voting (so v and w are always voting in the same manner). We define the sets V_v and V_w and the fair election procedures f_v and f_w in the analogous way. By Theorem 2.4.4, each of the election procedures f_u , f_v , and f_w , has a dictator.

Claim. *There is $x \in V$ such that x is the dictator for f_x .*

Proof of Claim. Suppose that the Claim is false. Fix distinct candidates i and j . By assumption, uw is the dictator for f_v , whence we have:

$$(1) \quad (ij, ji, ij) \rightsquigarrow ij.$$

Once again, by assumption, uv is the dictator for f_w , whence we have:

$$(2) \quad (ij, ij, ji) \rightsquigarrow ij.$$

Consider now the set $V' := \{v, w\}$. Fix $\sigma \in S_n$ such that $\sigma(i) < \sigma(j)$ and define the fair election procedure f' for (V', n) by setting $f'(\sigma_v, \sigma_w) := f(\sigma, \sigma_v, \sigma_w)$. By Theorem 2.4.4 again, there is a dictator for f' . If v is the dictator for f' , then we get a contradiction to (1), while if w is the dictator for f' , we get a contradiction to (2). Thus, the Claim is proven.

By the Claim, we may assume, without loss of generality, that u is the dictator for f_u . But then $\{u\}$ is block decisive for f , whence it is decisive for f by Proposition 2.3.4, and thus u is the dictator for f . \square

2.5. Notes and references

Arrow's original proof can be found in [2]. Our treatment of the ultrafilter proof of Arrow's theorem has borrowed heavily from Galvin's notes [64] and the article of Komj ath and Tatik [108].

Ultrafilters in topology

In this chapter, we investigate some uses of ultrafilters in topology. In Section 3.1, we define the notion of an ultralimit of a sequence in a topological space and show how this notion can be used to give nice characterizations of familiar topological notions, such as the closure of a set, the compactness of a space, and the continuity of a function. In Section 3.2, we show how the set of ultrafilters on a discrete space can be topologized so that it becomes the familiar Stone-Čech compactification of the discrete space; this analysis is extended to the nondiscrete setting in Section 3.3. In Section 3.4, we introduce the notion of an ultrafilter on a Boolean algebra so as to be able to prove the Stone duality theorem, which expresses the dual equivalence of the categories of Boolean algebras and compact, Hausdorff, totally disconnected spaces (otherwise known as Stone spaces).

3.1. Ultralimits

Definition 3.1.1. For X a topological space, $(x_i)_{i \in I}$ a sequence from X , and \mathcal{U} an ultrafilter on I , a \mathcal{U} -**ultralimit** of $(x_i)_{i \in I}$ is a point $x \in X$ such that, for all open neighborhoods U of x , we have $\{i \in I : x_i \in U\} \in \mathcal{U}$.

Exercise 3.1.2. Suppose that $(x_i)_{i \in I}$ is a sequence from X and \mathcal{U} is the principal ultrafilter on I generated by j . Show that x_j is a \mathcal{U} -ultralimit of $(x_i)_{i \in I}$. If, in addition, X is a T_1 space, prove that x_j is the only \mathcal{U} -ultralimit of $(x_i)_{i \in I}$.

Exercise 3.1.3. Suppose that X is Hausdorff. Show that, for any sequence $(x_i)_{i \in I}$ from X and any ultrafilter \mathcal{U} on I , there can be at most one \mathcal{U} -ultralimit of $(x_i)_{i \in I}$.

By the previous exercise, in the case of Hausdorff spaces, we write $\lim_{\mathcal{U}} x_i$ (or $\lim_{i, \mathcal{U}} x_i$) for *the* \mathcal{U} -ultralimit of $(x_i)_{i \in I}$ when it exists.

Exercise 3.1.4. Suppose that $(x_n)_{n \in \mathbb{N}}$ is a sequence in a metric space X . Suppose also that $\lim_{n \rightarrow \infty} x_n = x$. Show that for any nonprincipal ultrafilter \mathcal{U} on \mathbb{N} , we have $\lim_{\mathcal{U}} x_n = x$.

Exercise 3.1.5. Suppose that $(x_n)_{n \in \mathbb{N}}$ and $(y_n)_{n \in \mathbb{N}}$ are two sequences from \mathbb{R} and \mathcal{U} is an ultrafilter on \mathbb{N} such that $\lim_{\mathcal{U}} x_n$ and $\lim_{\mathcal{U}} y_n$ both exist. Prove the following:

- (1) $\lim_{\mathcal{U}}(x_n \pm y_n) = (\lim_{\mathcal{U}} x_n) \pm (\lim_{\mathcal{U}} y_n)$.
- (2) $\lim_{\mathcal{U}}(x_n \cdot y_n) = (\lim_{\mathcal{U}} x_n) \cdot (\lim_{\mathcal{U}} y_n)$.
- (3) If $\lim_{\mathcal{U}} y_n \neq 0$, then $y_n \neq 0$ for \mathcal{U} -almost all n and $\lim_{\mathcal{U}} \frac{x_n}{y_n} = \frac{\lim_{\mathcal{U}} x_n}{\lim_{\mathcal{U}} y_n}$. (Part of the exercise is to explain what the left hand side of the equation even means!)

Ultralimits become a convenient tool for speaking about limit points:

Theorem 3.1.6. *Suppose that X is a topological space, A is a subset of X , and $x \in X$. Then $x \in \overline{A}$ if and only if there is a sequence $(x_i)_{i \in I}$ from A and an ultrafilter \mathcal{U} on I such that x is a \mathcal{U} -ultralimit of $(x_i)_{i \in I}$.*

Proof. First suppose that $x \in \overline{A}$. If $x \in A$, then x is the ultralimit of the constant sequence (x, x, x, \dots) with respect to any ultrafilter \mathcal{U} on \mathbb{N} . Thus we may suppose that x is a limit point of A . For every open neighborhood U of x , take $x_U \in (A \cap U) \setminus \{x\}$. For each open neighborhood U of x , let $F_U := \{V \subseteq X : V \text{ is open and } V \subseteq U\}$. Set $\mathcal{D} := \{F_U : U \text{ an open neighborhood of } x\}$. Note that \mathcal{D} has the finite intersection property as $F_{U \cap U'} \subseteq F_U \cap F_{U'}$. Consequently, there is an ultrafilter \mathcal{U} on the set of open neighborhoods of x such that $\mathcal{D} \subseteq \mathcal{U}$. It is then easy to verify that x is a \mathcal{U} -ultralimit of (x_U) .

Conversely, suppose that x is a \mathcal{U} -ultralimit of the sequence $(x_i)_{i \in I}$ from A . Let U be an open neighborhood of x . Then there are \mathcal{U} -many i (so in particular one i) such that $x_i \in U$. It follows that $x \in \overline{A}$. \square

Corollary 3.1.7. *Suppose that $(x_n)_{n \in \mathbb{N}}$ is a sequence from a metric space X and $x \in X$. Then:*

- (1) x is a subsequential limit of $(x_n)_{n \in \mathbb{N}}$ if and only if there is a non-principal ultrafilter \mathcal{U} on \mathbb{N} such that $x = \lim_{\mathcal{U}} x_n$.
- (2) $(x_n)_{n \in \mathbb{N}}$ converges to x if and only if $x = \lim_{\mathcal{U}} x_n$ for every non-principal ultrafilter \mathcal{U} on \mathbb{N} .

Exercise 3.1.8. Prove the previous corollary.

Ultralimits provide a very convenient characterization of compactness:

Theorem 3.1.9. *For a topological space X , the following are equivalent:*

- (1) X is compact.
- (2) Given any sequence $(x_i)_{i \in I}$ from X and any ultrafilter \mathcal{U} on I , $(x_i)_{i \in I}$ has a \mathcal{U} -ultralimit in X .

Proof. (1) implies (2): Suppose that $(x_i)_{i \in I}$ is a sequence from X and \mathcal{U} is an ultrafilter on I such that $(x_i)_{i \in I}$ has no \mathcal{U} -ultralimit. Consequently, for each $x \in X$, there is an open neighborhood U_x of x such that $\{i \in I : x_i \in U_x\} \notin \mathcal{U}$. If there were a finite subcover $X = \bigcup_{j=1}^n U_{x_j}$, then we would have

$$\emptyset = \left\{ i \in I : x_i \notin \bigcup_{j=1}^n U_{x_j} \right\} = \bigcap_{j=1}^n \{i \in I : x_i \notin U_{x_j}\} \in \mathcal{U},$$

which is a contradiction. Thus X is not compact.

(2) implies (1): Suppose that X is not compact. Take an open cover $(U_i)_{i \in I}$ with no finite subcover. Without loss of generality, each $U_i \neq \emptyset$. Let Y be the set of nonempty, finite subsets of I . For each $i \in I$, let $A_i := \{J \in Y : i \in J\}$. Observe that $(A_i)_{i \in I}$ has the finite intersection property, whence we may take an ultrafilter \mathcal{U} on Y for which $A_i \in \mathcal{U}$ for each $i \in I$. Since $(U_i)_{i \in I}$ has no finite subcover, for each $J \in Y$, we may fix some $x_J \in X \setminus \bigcup_{j \in J} U_j$. We claim that $(x_J)_{J \in Y}$ has no \mathcal{U} -ultralimit in X . Indeed, suppose that x was a \mathcal{U} -ultralimit of $(x_J)_{J \in Y}$. Take $i \in I$ such that $x \in U_i$. Then there is $A \in \mathcal{U}$ such that $x_J \in U_i$ for all $J \in A$. Since $A_i \in \mathcal{U}$, we have $x_J \in U_i$ for $J \in A \cap A_i$; but $J \in A_i$ implies $i \in J$, whence $x_J \notin U_i$, yielding a contradiction. \square

Theorem 3.1.10. *A topological space X is compact and Hausdorff if and only if, given any sequence $(x_i)_{i \in I}$ from X and any ultrafilter \mathcal{U} on I , $(x_i)_{i \in I}$ has a unique \mathcal{U} -ultralimit.*

Proof. The forward direction follows from Theorem 3.1.9 and Exercise 3.1.3. We now prove the backward direction. Suppose that X is not Hausdorff; we find some sequence from X and some ultrafilter on the index set such that the sequence does not have a unique ultralimit with respect to that ultrafilter. Since X is not Hausdorff, there are distinct $x, y \in X$ such that every open neighborhood of x intersects every open neighborhood of y . We let \mathcal{O}_x denote the set of open neighborhoods of x and similarly for \mathcal{O}_y , and we let $\mathcal{F} := \mathcal{O}_x \cup \mathcal{O}_y$. By our assumption that every open neighborhood of x intersects every open neighborhood of y , we have that \mathcal{F} is a filter on X . Let \mathcal{U} be any ultrafilter on X extending \mathcal{F} . We claim that x and y are

both \mathcal{U} -ultralimits of the sequence $(z)_{z \in X}$. Indeed, if U is an open neighborhood of x , then $U \in \mathcal{O}_x \subseteq \mathcal{F} \subseteq \mathcal{U}$, whence $z \in U$ for \mathcal{U} -almost all $z \in X$. Consequently, x is a \mathcal{U} -ultralimit of the sequence. A symmetric argument shows that y is also a \mathcal{U} -ultralimit of the sequence. \square

The following special case of the previous theorem is worth recording:

Corollary 3.1.11. *Suppose that $(a_n)_{n \in \mathbb{N}}$ is a bounded sequence in \mathbb{R} and \mathcal{U} is an ultrafilter on \mathbb{N} . Then $\lim_{\mathcal{U}} a_n$ exists.*

Thus, while regular limits (even of bounded sequences) need not exist, ultralimits of bounded sequences always exist!

Exercise 3.1.12. Use ultralimits to show that a closed subspace of a compact space is compact.

Theorem 3.1.9 also yields a very simple proof of Tychonoff's theorem:

Theorem 3.1.13 (Tychonoff's theorem). *Given a family $(X_j)_{j \in J}$ of compact spaces, the product space $\prod_{j \in J} X_j$ is also compact.*

Proof. We use Theorem 3.1.9 to prove that $\prod_{j \in J} X_j$ is compact. Let $(x_i)_{i \in I}$ be a sequence from $\prod_{j \in J} X_j$ and write $x_i = (x_i(j))_{j \in J}$. Let \mathcal{U} be an ultrafilter on I . Since each X_j is compact, we may consider a \mathcal{U} -ultralimit $x(j)$ of the sequence $(x_i(j))_{i \in I}$. We show that $x = (x(j))_{j \in J}$ is a \mathcal{U} -ultralimit of $(x_i)_{i \in I}$. Take $j_1, \dots, j_n \in J$ and open sets $U_k \subseteq X_{j_k}$ containing $x(j_k)$ for $k = 1, \dots, n$. For each $k = 1, \dots, n$, there is $A_k \in \mathcal{U}$ such that $x_i(j_k) \in U_k$ for $i \in A_k$. Let $A = A_1 \cap \dots \cap A_n$. Then if $i \in A$, we have $x_i(j_k) \in U_k$ for all $k = 1, \dots, n$, whence x_i belongs to the basic open set determined by U_1, \dots, U_n . \square

We next present the ultralimit characterization of continuity:

Theorem 3.1.14. *Suppose that $f : X \rightarrow Y$ is a function between topological spaces and $x \in X$. Then f is continuous at x if and only if, for any sequence $(x_i)_{i \in I}$ from X and any ultrafilter \mathcal{U} on I for which x is a \mathcal{U} -ultralimit of $(x_i)_{i \in I}$, we have that $f(x)$ is a \mathcal{U} -ultralimit of $(f(x_i))_{i \in I}$.*

Proof. First suppose that f is continuous at x and x is a \mathcal{U} -ultralimit of $(x_i)_{i \in I}$. Fix an open neighborhood U of $f(x)$ and take an open neighborhood V of x such that $f(V) \subseteq U$. Since $\{i \in I : x_i \in V\} \in \mathcal{U}$, we have $\{i \in I : f(x_i) \in U\} \in \mathcal{U}$; since U was an arbitrary neighborhood of $f(x)$, we have that $f(x)$ is a \mathcal{U} -ultralimit of $(f(x_i))_{i \in I}$.

Conversely, suppose that f is not continuous at x . Take an open neighborhood U of $f(x)$ such that, for every open neighborhood V of x , $f(V) \not\subseteq U$. Take $x_V \in V$ such that $f(x_V) \notin U$. Let \mathcal{U} be any ultrafilter on the set of

open neighborhoods of x such that, for any given open neighborhood V_0 of x , we have $\{V : V \subseteq V_0\} \in \mathcal{U}$. Consequently, x is a \mathcal{U} -ultralimit of the sequence (x_V) , whence, by assumption, $f(x)$ is a \mathcal{U} -ultralimit of the sequence $(f(x_V))$. Since $Y \setminus U$ is closed and $f(x_V) \in Y \setminus U$ for all V , we have $f(x) \in Y \setminus U$, yielding a contradiction. \square

We close this section with an application of Theorem 3.1.9 and Theorem 3.1.14 that will be useful in Chapter 15:

Proposition 3.1.15. *Suppose that $f : X \rightarrow Y$ is a surjective continuous map between topological spaces. Further suppose that X is compact and Hausdorff. Then there exists a closed subset K of X such that $f(K) = Y$ but $f(K') \neq Y$ for any proper closed subset K' of K .*

Proof. Let P denote the set of all closed subsets K of X such that $f(K) = Y$. Note that P is nonempty since f is surjective. Put a partial order \leq on P by declaring $K_1 \leq K_2$ if and only if $K_1 \supseteq K_2$. Consequently, we seek a maximal element in the partial order (P, \leq) . To prove that such a maximal element exists, we apply Zorn's lemma. Toward that end, suppose that $(K_i)_{i \in I}$ is a chain in P ; we show that $K := \bigcap_{i \in I} K_i$ is an upper bound for the chain in P . It is clear that K is a closed subset of X . To see that $f(K) = Y$, fix $y \in Y$. For each $i \in I$, take $x_i \in K_i$ such that $f(x_i) = y$; this is possible since $f(K_i) = Y$ for each $i \in I$. For each $i \in I$, let $D_i := \{j \in I : K_i \supseteq K_j\}$. Note that $(D_i)_{i \in I}$ has the FIP since $(K_i)_{i \in I}$ is a chain. Consequently, there is an ultrafilter \mathcal{U} on I such that $D_i \in \mathcal{U}$ for all $i \in I$. Set $x := \lim_{\mathcal{U}} x_i$. Since f is continuous, Theorem 3.1.14 implies that $f(x) = \lim_{\mathcal{U}} f(x_i) = y$. It remains to verify that $x \in K$. To see this, suppose, toward a contradiction, that $x \notin K_i$ for some $i \in I$. Since $X \setminus K_i$ is open, we have that $x_j \in X \setminus K_i$ for \mathcal{U} -almost all $j \in I$, contradicting that $D_i \in \mathcal{U}$ and the fact that $K_j \subseteq K_i$ for $j \in D_i$. \square

3.2. The Stone-Čech compactification: the discrete case

Until further notice, we fix an infinite set X , which we also think of as a topological space equipped with the discrete topology. For $A \subseteq X$, we set $U_A := \{\mathcal{U} \in \beta X : A \in \mathcal{U}\}$.

Theorem 3.2.1. *The sets U_A form a neighborhood base for a topology on βX . When equipped with this topology, we have the following:*

- (1) βX is Hausdorff.
- (2) βX is **zero-dimensional**, that is, has a base of clopen subsets.
- (3) βX is compact.

Setting $\iota : X \rightarrow \beta X$ to be the function defined by $\iota(x) := U_x$, we also have:

- (4) For every $x \in X$, $\iota(x)$ is an isolated point of βX .
 (5) ι is a homeomorphism between X and $\iota(X)$.
 (6) $\iota(X)$ is dense in βX .

Proof. The fact that $U_{A \cap B} \subseteq U_A \cap U_B$ implies that the sets U_A form a base for a topology on βX . We now prove the remaining items.

For (1), note that, if \mathcal{U}, \mathcal{V} are distinct elements of βX , then taking $A \subseteq X$ for which $A \in \mathcal{U}$ and $X \setminus A \in \mathcal{V}$, we have that U_A and $U_{X \setminus A}$ are disjoint open sets containing \mathcal{U} and \mathcal{V} , respectively, whence the topology is Hausdorff.

(2) follows from the observation that $U_{X \setminus A} := \beta X \setminus U_A$.

To prove (3), it suffices to show, by Theorem 3.1.9, that given any sequence $(\mathcal{U}_i)_{i \in I}$ from βX (here, \mathcal{U}_i is not a principal ultrafilter but rather some arbitrary element of βX) and ultrafilter \mathcal{V} on I , that $(\mathcal{U}_i)_{i \in I}$ has a \mathcal{V} -ultralimit. Let \mathcal{U} denote those $A \subseteq X$ such that $A \in \mathcal{U}_i$ for \mathcal{V} almost all $i \in I$. We leave it to the reader to check that \mathcal{U} is a \mathcal{V} -ultralimit of $(\mathcal{U}_i)_{i \in I}$.

(4) follows from the observation that $\iota(x) = U_{\{x\}}$.

(5) ι is injective as, whenever $x, y \in X$ are distinct, then $\{x\} \in \mathcal{U}_x \setminus \mathcal{U}_y$. ι is continuous as X is discrete. To see that ι is a homeomorphism onto $\iota(X)$, it suffices to observe that, for any $A \subseteq X$, we have $\iota(A) = U_A \cap \iota(X)$.

(6) follows from the observation that, if $A \subseteq X$ is nonempty and $x \in A$, then $\mathcal{U}_x \in U_A$. \square

Exercise 3.2.2. Verify the claims in the previous theorem left to the reader.

In what follows, we identify X with its image $\iota(X)$.

Exercise 3.2.3. For $A \subseteq X$, the closure of A in βX is U_A .

Exercise 3.2.4. For $\mathcal{U} \in \beta X$, we have $\lim_{x, \mathcal{U}} x = \mathcal{U}$.

We recall the following:

Definition 3.2.5. If Y is a topological space and K is a compact space, then K is a **compactification** of Y if Y is a dense subspace of K .

Theorem 3.2.1 thus implies that βX is a compactification of X whenever X is an infinite discrete space. The reader may have encountered the one-point compactification of a locally compact space, which is a very “small” compactification for it merely adds one point. The compactification βX is a much, much larger compactification of X .

Definition 3.2.6. Let Y be a topological space. A Hausdorff compactification K of Y is called a **Stone-Ćech compactification** of Y if, for every compact Hausdorff space Z and every continuous function $f : Y \rightarrow Z$, there is a unique continuous function $\tilde{f} : K \rightarrow Z$ extending f .

Exercise 3.2.7. Prove that a topological space Y can have at most one Stone-Čech compactification in the following strong form: if K_1 and K_2 are both Stone-Čech compactifications of Y , then there is a unique homeomorphism $\Phi : K_1 \rightarrow K_2$ such that $\Phi(y) = y$ for all $y \in Y$.

By the previous exercise, we may unambiguously denote the Stone-Čech compactification of a topological space Y , when it exists, by βY . The reader may note that this is also the notation we used in Chapter 1 to denote the set of all ultrafilters on a set. The reason that this double-use of notation is not problematic is actually the main result of this section:

Theorem 3.2.8. *For an infinite discrete set X , βX is the Stone-Čech compactification of X .*

Proof. We already know that βX is a Hausdorff compactification of X . We now need to verify its “universal” property. Suppose that $f : X \rightarrow Y$ is a continuous function into a compact Hausdorff space. Given $\mathcal{U} \in \beta X$, define $\tilde{f}(\mathcal{U}) := \lim_{\mathcal{U}} f(x)$, which exists by Theorem 3.1.10 and the fact that Y is compact and Hausdorff. Note that \tilde{f} extends f by Exercise 3.1.2. We must show that \tilde{f} is continuous. Since every point of X is isolated in βX , it suffices to show continuity at \mathcal{U} for nonprincipal \mathcal{U} . Toward this end, let U be an open neighborhood of $\tilde{f}(\mathcal{U})$ in Y . Let $V \subseteq U$ be an open neighborhood of $\tilde{f}(\mathcal{U})$ in Y such that $\overline{V} \subseteq U$. Take $A \in \mathcal{U}$ such that $f(x) \in V$ for $x \in A$. Suppose $\mathcal{V} \in U_A$, so $A \in \mathcal{V}$; then $\lim_{\mathcal{V}} f(x) \in \overline{V} \subseteq U$, so $\tilde{f}(U_A) \subseteq U$ and \tilde{f} is continuous at \mathcal{U} .

For the uniqueness of \tilde{f} , suppose that $g : \tilde{X} \rightarrow Y$ is a continuous function that extends f . By Theorem 3.1.14 and Exercise 3.2.4, $g(\mathcal{U}) = g(\lim_{\mathcal{U}} x) = \lim_{\mathcal{U}} g(x) = \lim_{\mathcal{U}} f(x) = \tilde{f}(\mathcal{U})$, whence $g = \tilde{f}$. \square

Exercise 3.2.9. Suppose that $f : I \rightarrow J$ is a function. Then f extends to a continuous function $\beta f : \beta I \rightarrow \beta J$. Show that, for all $\mathcal{U} \in \beta I$, we have $(\beta f)(\mathcal{U}) = \tilde{f}(\mathcal{U})$, the pushforward ultrafilter.

3.3. z -ultrafilters and the Stone-Čech compactifications in general

In this section, we turn to the problem of constructing the Stone-Čech compactification in general, that is, for not necessarily discrete spaces. We first take up the question: which spaces have Hausdorff compactifications? Note that a space has a Hausdorff compactification if and only if it can be embedded into a compact Hausdorff space.

How can one go about trying to embed a space X into a compact Hausdorff space? Set $C := C(X)$ to be the set of continuous, real-valued functions on the set X . Consider the mapping $e : X \rightarrow [0, 1]^C$ given by $e(x)(f) = f(x)$.

Since $[0, 1]^C$ is a compact Hausdorff space, in order to show that X has a compactification, it suffices to show that, under certain assumptions on X , that e is a homeomorphism of X onto its image.

Continuity of the map e holds for any space X . Indeed, it suffices to show, for a fixed $f \in C$ and open $U \subseteq [0, 1]$, that the preimage of the subbasic open set $\{F \in [0, 1]^C : F(f) \in U\}$ under e is open in X . However, this preimage is simply $f^{-1}(U)$, which is open by the continuity of f .

In order to check that e is injective, it suffices to know that, for any distinct $x, y \in X$, there is $f \in C$ such that $f(x) \neq f(y)$. When this is indeed the case for all $x, y \in X$, we say that $C(X)$ **separates points in X** .

Finally, we want to know that e is a homeomorphism onto its image, that is, $e : X \rightarrow e(X)$ is an open map. To verify this, take $U \subseteq X$ open; we must show that $e(U)$ is open in $e(X)$. Take $x \in U$; we need $e(x)$ to be in the interior of $e(U)$ as calculated in $e(X)$. If there is $f \in C$ such that $f = 0$ for all $y \in X \setminus U$ while $f(x) = 1$, then, setting V to be the subbasic open subset of $[0, 1]^C$ determined by the condition $f > \frac{1}{2}$, we have that $V \cap e(X) \subseteq e(U)$, as desired.

The condition appearing in the previous paragraph has a name:

Definition 3.3.1. A topological space X is **completely regular** if, for every closed $C \subseteq X$ and every $x \in X \setminus C$, there is $f \in C(X)$ such that $f(y) = 0$ for all $y \in C$ while $f(x) = 1$.

While it appears that complete regularity of a space X ensures that $C(X)$ separates points, this is only the case if points in X are closed, that is, if X is T_1 .

Definition 3.3.2. A topological space X is called a **Tychonoff space** if it is a completely regular T_1 space.

The previous discussion shows:

Theorem 3.3.3. *If X is a Tychonoff space, then X is homeomorphic to a subspace of a compact Hausdorff space, whence has a Hausdorff compactification.*

Thankfully, the converse is also true and follows from the following two exercises:

Exercise 3.3.4. A subspace of a Tychonoff space is completely regular.

Exercise 3.3.5. Compact Hausdorff spaces are Tychonoff spaces.

To summarize:

Theorem 3.3.6 (Tychonoff). *A topological space X has a Hausdorff compactification if and only if X is a Tychonoff space.*

The main theorem of this section is that for the spaces that do possess Hausdorff compactifications, namely the Tychonoff spaces, the Stone-Čech compactification also exists. The naïve idea might be to mimic the construction of the Stone-Čech compactification of a discrete space except to deal only with maximal filters of closed sets rather than all sets. It turns out that this does not work in general (but does work if the space is *perfectly normal* as defined below). Instead, we need to work with special kinds of closed sets:

Definition 3.3.7.

- (1) For $f \in C(X)$, we define the **zeroset of f** to be $Z(f) := \{x \in X : f(x) = 0\}$.
- (2) We call $Z \subseteq X$ a **zeroset of X** if there is $f \in C(X)$ such that $Z = Z(f)$.
- (3) We let $\mathcal{Z}(X)$ denote the set of zerosets in X .

Exercise 3.3.8. Show that $\mathcal{Z}(X)$ is closed under finite unions and intersections.

Exercise 3.3.9. Suppose that X is completely regular, $x \in X$, and $C \subseteq X$ is closed with $x \notin C$. Show that there is $Z \in \mathcal{Z}(X)$ such that $x \in Z$ and $Z \cap C = \emptyset$.

Clearly, every zeroset is closed. We say that X is **perfectly normal** if, conversely, every closed set is a zeroset. We will not need to know too much about perfectly normal spaces, but rather only the following:

Exercise 3.3.10. Show that $[0, 1]$ is perfectly normal.

The importance of zerosets in completely regular spaces is explained by the next theorem:

Theorem 3.3.11. *X is completely regular if and only if the zerosets of X form a base for the closed sets of X .*

Proof. First assume that X is completely regular. By the proof of Theorem 3.3.3, we may assume that X is a subspace of $[0, 1]^I$ for some index set I . Let $C \subseteq X$ be closed. Then C is an intersection of sets of the form $C_J := \{\vec{x} \in X : x_i \in C_i \text{ for } i \in J\}$, where $J \subseteq I$ is finite and each $C_i \subseteq [0, 1]$ is closed. Since $[0, 1]$ is perfectly normal, we may assume that $C_i := Z(f_i)$ for some $f_i \in C([0, 1])$. Then $C_J = Z(\sum_{j \in J} g_j^2)$, where $g_j(\vec{x}) = f_j(x_j)$.

For the converse, suppose that $C \subseteq X$ is closed and $x \in X \setminus C$. By assumption, there is $f \in C(X)$ such that $f(y) = 0$ for all $y \in C$ while $f(x) \neq 0$. By multiplying f by a suitable multiple, we may assume that $f(x) = 1$, whence X is completely regular. \square

In the rest of this section, we assume that X is a Tychonoff space.

By the previous theorem, we see that zerosets play a fundamental role in the theory of completely regular spaces. As we will now see, by proceeding as in the previous section, but only working with zerosets, we can construct the Stone-Ćech compactification of X . Toward this end, we make the following definition.

Definition 3.3.12. We call $\mathcal{F} \subseteq \mathcal{Z}(X) \setminus \{\emptyset\}$ a **z -filter on X** if:

- (1) $Z_1, Z_2 \in \mathcal{F} \Rightarrow Z_1 \cap Z_2 \in \mathcal{F}$;
- (2) $Z_1 \in \mathcal{F}, Z_2 \in \mathcal{Z}(X)$ and $Z_1 \subseteq Z_2 \Rightarrow Z_2 \in \mathcal{F}$.

A maximal z -filter on X is called a **z -ultrafilter** on X . We let ζX denote the set of z -ultrafilters on X .

Exercise 3.3.13. Suppose that \mathcal{U} is a z -filter on X . Prove that \mathcal{U} is a z -ultrafilter on X if and only if, whenever $C \in \mathcal{Z}(X)$ is such that $C \cap Z \neq \emptyset$ for all $Z \in \mathcal{U}$, then we have $C \in \mathcal{U}$.

Exercise 3.3.14. Prove that any family $(Z_i)_{i \in I}$ of elements of $\mathcal{Z}(X)$ with the finite intersection property is contained in a z -ultrafilter on X .

We will soon see that ζX “is” βX . We first need to put a topology on ζX . We take our cue from the discrete case: given $Z \in \mathcal{Z}(X)$, we set $C_Z := \{\mathcal{U} \in \zeta X : Z \in \mathcal{U}\}$. We then give ζX the topology where the C_Z ’s form a base for the closed sets. Unlike the discrete case, the basic closed sets are not also open. In fact, given an open set $O \subseteq X$, we set

$$U_O := \{\mathcal{U} \in \zeta X : Z \subseteq O \text{ for some } Z \in \mathcal{U}\}.$$

Exercise 3.3.15. Prove that $\zeta X \setminus C_Z = U_{X \setminus Z}$.

Consequently, U_O , for O a complement of a zeroset, is a basic open set in ζX .

Theorem 3.3.16. ζX is a compact Hausdorff space.

Proof. We first show that ζX is Hausdorff. Toward this end, fix distinct $\mathcal{U}, \mathcal{V} \in \zeta X$. By Exercise 3.3.13, there are $Z_1, Z_2 \in \mathcal{Z}(X)$ such that $Z_1 \in \mathcal{U}$, $Z_2 \in \mathcal{V}$ and $Z_1 \cap Z_2 = \emptyset$. Let $f \in C(X)$ be such that $Z_1 = Z(f)$. Since Z_2 is closed, there is some $\epsilon > 0$ such that $f(x) \geq \epsilon$ for all $x \in Z_2$. Define $g_1, g_2 \in C(X)$ by setting $g_1(x) = \max(\frac{\epsilon}{2} - f(x), 0)$ and $g_2(x) = \max(f(x) - \frac{\epsilon}{2}, 0)$. Set $W_1 := Z(g_1)$ and $W_2 := Z(g_2)$. Then, for $i = 1, 2$, we have $Z_i \subseteq (X \setminus W_i)$ and $(X \setminus W_1) \cap (X \setminus W_2) = \emptyset$. We thus have that $U_{(X \setminus W_1)}$ and $U_{(X \setminus W_2)}$ are disjoint open neighborhood of \mathcal{U} and \mathcal{V} , whence ζX is Hausdorff.

To see that ζX is compact, it suffices to show that any family $(C_{Z_i})_{i \in I}$ of basic closed subsets of ζX with the finite intersection property has a nonempty intersection. Since $(C_{Z_i})_{i \in I}$ has the finite intersection property, it follows that $(Z_i)_{i \in I}$ has the finite intersection property: given $i_1, \dots, i_n \in I$, take $\mathcal{U} \in \bigcap_{j=1}^n C_{Z_{i_j}}$ and note that $\bigcap_{j=1}^n Z_{i_j} \in \mathcal{U}$, whence $\bigcap_{j=1}^n Z_{i_j} \neq \emptyset$. By Exercise 3.3.14, there is a z -ultrafilter \mathcal{U} containing the family $(Z_i)_{i \in I}$, whence $\mathcal{U} \in \bigcap_{i \in I} C_{Z_i}$. \square

We now show that ζX is a compactification of X :

Theorem 3.3.17. *For $x \in X$, set $\iota(x) := \{Z \in \mathcal{Z}(X) : x \in Z\}$. Then:*

- (1) $\iota(x) \in \zeta X$;
- (2) the mapping $\iota : X \rightarrow \zeta X$ is a homeomorphism of X onto its image;
- (3) $\iota(X)$ is dense in ζX .

Consequently, identifying X with $\iota(X)$, we have that ζX is a compactification of X .

Proof. (1) follows immediately from Exercises 3.3.9 and 3.3.13.

For (2), first notice that ι is injective: if $x, y \in X$ are distinct, then by complete regularity, there are $Z_1, Z_2 \in \mathcal{Z}(X)$ such that $x \in Z_1, y \in Z_2$, and $Z_1 \cap Z_2 = \emptyset$. It follows that $\iota(x) \neq \iota(y)$. To see that ι is continuous, it suffices to notice that $\iota^{-1}(C_Z) = Z$. Finally, to check that $\iota : X \rightarrow \iota(X)$ is open, it is enough to show that $\iota(Z)$ is closed in $\iota(X)$ for every $Z \in \mathcal{Z}(X)$, which follows from the fact that $\iota(Z) = C_Z \cap \iota(X)$.

For (3), take a nonempty basic open set U_O and take $\mathcal{U} \in U_O$. Take $Z \in \mathcal{U}$ such that $Z \subseteq O$ and let $x \in Z$. Then $Z \in \iota(x)$, so $\iota(x) \in U_O$. \square

It remains to see that ζX is the Stone-Čech compactification of X . To prove this, it will be convenient to use a different characterization of βX due to Čech. First, we say that $A, B \subseteq X$ are **completely separated** if there is $f \in C(X)$ such that $f(x) = 0$ for all $x \in A$ while $f(x) = 1$ for all $x \in B$.

Theorem 3.3.18. *βX is the unique compactification of X such that completely separated subsets of X have disjoint closures in βX .*

Proof. We first show that βX has the stated property. Indeed, suppose that $A, B \subseteq X$ are completely separated by $f \in C(X)$. Without loss of generality, we may assume that $f(X) \subseteq [0, 1]$. Note then that if $x, y \in \beta X$ are in the closures of A and B , respectively, then $\beta f(x) = 0$ while $\beta f(y) = 1$, whence $x \neq y$. (Here, βf denotes the unique continuous extension of f to βX .)

Now suppose that K is a compactification of X with the stated property. Let $g : \beta X \rightarrow K$ be a continuous function that is the identity on X . Since X is dense in K , it follows that g is onto. If we can show that g is also injective, then since βX is compact and K is Hausdorff, it will follow that g is a homeomorphism. Let $p, q \in \beta X$ be distinct. Let $f : \beta X \rightarrow [0, 1]$ be such that $f(p) = 0$ and $f(q) = 1$. (This is possible since compact spaces are completely regular.) Set $A := \{x \in X : f(x) \leq \frac{1}{3}\}$ and $B = \{x \in X : f(x) \geq \frac{2}{3}\}$. Note that A and B are completely separated. (Exercise.) Since p is in the closure of A in βX , we have that $g(p)$ is in the closure of A in K . Similarly, $g(q)$ is in the closure of B in K . By assumption, these closures are disjoint, whence $g(p) \neq g(q)$, as desired. \square

In order to use the previous theorem, we need to understand what closures of zerosets in X inside ζX look like:

Lemma 3.3.19. *Given $Z \in \mathcal{Z}(X)$, the closure of Z in ζX is C_Z .*

Proof. Temporarily, set \overline{Z} to be the closure of Z in ζX . We clearly have $\overline{Z} \subseteq C_Z$. To obtain the reverse direction, let C_W be a basic closed set containing Z . Then $W = C_W \cap X \supseteq Z$, whence $C_Z \subseteq C_W$. It follows that $C_Z \subseteq \overline{Z}$. \square

Finally, we are ready to prove the main result of this section:

Theorem 3.3.20. *ζX is the Stone-Čech compactification of X .*

Proof. Let $A, B \subseteq X$ be completely separated. We show that A and B have disjoint closures in ζX . Since A and B are completely separated, there are $Z_1, Z_2 \in \mathcal{Z}(X)$ such that $A \subseteq Z_1$, $B \subseteq Z_2$, and $Z_1 \cap Z_2 = \emptyset$. It thus suffices to show that Z_1 and Z_2 have disjoint closures in ζX . By the previous lemma, we must show that $C_{Z_1} \cap C_{Z_2} = \emptyset$. However, it is easy to see that $C_{Z_1} \cap C_{Z_2} = C_{Z_1 \cap Z_2} = \emptyset$, whence we are finished. \square

3.4. The Stone representation theorem

When X is discrete, Theorem 3.2.1 shows that βX is a Stone space in the sense of the following definition:

Definition 3.4.1. A topological space is a **Stone space** if it is compact, Hausdorff, and totally disconnected.

We remind the reader that a topological space is totally disconnected if singletons are the only connected sets. In order to see that βX is indeed a Stone space, we need:

Exercise 3.4.2. A zero-dimensional space is totally disconnected.

In this section, we show that all Stone spaces can be obtained in this manner provided we are willing to work with ultrafilters on arbitrary Boolean algebras, as defined here:

Definition 3.4.3. A **Boolean algebra** is a structure $\mathbb{B} = (\mathbb{B}, 0, 1, \wedge, \vee, \neg)$, where 0 and 1 are elements of \mathbb{B} , \wedge and \vee are binary operations on \mathbb{B} , and \neg is a unary operation on \mathbb{B} for which we have, for all $a, b, c \in \mathbb{B}$:

- (1) $a \wedge a = a \vee a = a$.
- (2) $a \wedge b = b \wedge a$ and $a \vee b = b \vee a$.
- (3) $a \wedge (b \wedge c) = (a \wedge b) \wedge c$ and $a \vee (b \vee c) = (a \vee b) \vee c$.
- (4) $a \wedge (a \vee b) = a \vee (a \wedge b) = a$.
- (5) $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$ and $a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c)$.
- (6) $0 \wedge a = 0$ and $0 \vee a = a$.
- (7) $1 \wedge a = a$ and $1 \vee a = 1$.
- (8) $a \wedge \neg a = 0$ and $a \vee \neg a = 1$.

Exercise 3.4.4. Given any set X , prove that $(\mathcal{P}(X), \emptyset, X, \cap, \cup, X \setminus \cdot)$ is a Boolean algebra.

We refer to a Boolean algebra as in the previous exercise as a **powerset algebra**. A Boolean subalgebra (defined in the obvious way) of a powerset algebra will be referred to as a **concrete Boolean algebra**. The Stone representation theorem, to be proven below, will show that every Boolean algebra is isomorphic to a concrete Boolean algebra.

Since Boolean algebras can be viewed as abstract generalizations of power sets, we can further abstract the notion of an ultrafilter. First, since the usual notion of an ultrafilter refers to the subset relation, we need to identify the abstracted version of this relation. We take our cue from the obvious fact that $A \subseteq B$ if and only if $A \cap B = A$.

Definition 3.4.5. If \mathbb{B} is a Boolean algebra, we define the binary relation \leq on \mathbb{B} by declaring $a \leq b$ if and only if $a \wedge b = a$.

Definition 3.4.6. If \mathbb{B} is a Boolean algebra, a **filter** on \mathbb{B} is a subset $\mathcal{F} \subseteq \mathbb{B}$ satisfying the following properties:

- (1) $0 \notin \mathcal{F}$, $1 \in \mathcal{F}$.
- (2) If $a, b \in \mathbb{B}$ are such that $a \in \mathcal{F}$ and $a \leq b$, then $b \in \mathcal{F}$.
- (3) If $a, b \in \mathcal{F}$, then $a \wedge b \in \mathcal{F}$.

Furthermore, \mathcal{F} is called an **ultrafilter** on \mathbb{B} if it also satisfies:

- (4) For all $a \in \mathbb{B}$, either $a \in \mathcal{F}$ or $\neg a \in \mathcal{F}$.

We let $S(\mathbb{B})$ denote the set of all ultrafilters on \mathbb{B} and refer to it as the **Stone space of \mathbb{B}** .

As before, we often denote ultrafilters on Boolean algebras by \mathcal{U} and \mathcal{V} .

Exercise 3.4.7. Verify that, for a powerset algebra $\mathcal{P}(X)$, the Boolean algebra notion of ultrafilter agrees with our earlier notion of ultrafilter, whence $S(\mathcal{P}(X)) = \beta X$.

Exercise 3.4.8. Prove the **ultrafilter theorem for Boolean algebras**: for any Boolean algebra \mathbb{B} and any filter \mathcal{F} on \mathbb{B} , there is an ultrafilter \mathcal{U} on \mathbb{B} extending \mathcal{F} .

We would like to define a topology on $S(\mathbb{B})$ for \mathbb{B} an arbitrary Boolean algebra so that the resulting space is a Stone space (and agrees with the earlier topology on βX in the case of a powerset algebra). Taking our cue from the case of powersets, we topologize $S(\mathbb{B})$ by declaring the basic open sets to be of the form U_a for $a \in \mathbb{B}$. Note that, in the case of powerset algebras, this agrees with the usual topology on βX .

In order to prove that $S(\mathbb{B})$, when endowed with the above topology, is a Stone space, it helps to introduce homomorphisms between Boolean algebras:

Definition 3.4.9. If \mathbb{B} and \mathbb{B}' are Boolean algebras and $h : \mathbb{B} \rightarrow \mathbb{B}'$ is a function, we say that h is a **homomorphism** if, $h(0) = 0'$, $h(1) = 1'$, and for all $a, b \in \mathbb{B}$, we have:

$$(1) \quad h(a \wedge b) = h(a) \wedge h(b).$$

$$(2) \quad h(a \vee b) = h(a) \vee h(b).$$

$$(3) \quad h(\neg a) = \neg h(a).$$

Exercise 3.4.10. Suppose $f : X \rightarrow Y$ is a function. Show that $h_f : \mathcal{P}(Y) \rightarrow \mathcal{P}(X)$ given by $h_f(A) := f^{-1}(A)$ is a Boolean algebra homomorphism.

Definition 3.4.11. $\mathbf{2}$ is the unique Boolean algebra with two elements $\{0, 1\}$.

Exercise 3.4.12. If \mathbb{B} is a Boolean algebra, then $\mathcal{U} \subseteq \mathbb{B}$ is an ultrafilter on \mathbb{B} if and only if there is a Boolean algebra homomorphism $h : \mathbb{B} \rightarrow \mathbf{2}$ such that $\mathcal{U} = \{a \in \mathbb{B} : h(a) = 1\}$. In this case, the homomorphism h is unique.

In light of the above exercise, we may view $S(\mathbb{B})$ as a subspace of $2^{\mathbb{B}}$. Notice that the topology on $S(\mathbb{B})$ is the subspace topology it inherits from $2^{\mathbb{B}}$.

Exercise 3.4.13. $S(\mathbb{B})$ is a closed subspace of $2^{\mathbb{B}}$.

Exercise 3.4.14. For any set X , 2^X is a Stone space.

Exercise 3.4.15. Any closed subspace of a Stone space is a Stone space.

Combining the previous exercises, we have proven:

Theorem 3.4.16. For any Boolean algebra \mathbb{B} , $S(\mathbb{B})$ is a Stone space.

In what follows, we consider the categories of Boolean algebras (where morphisms are Boolean algebra homomorphisms) and Stone spaces (where the morphisms are the continuous functions, yielding a full subcategory of the category of topological spaces). The association $\mathbb{B} \mapsto S(\mathbb{B})$ is then a contravariant functor. To see this, note that, given a Boolean algebra homomorphism $h : \mathbb{B} \rightarrow \mathbb{B}'$, we get a continuous function $S(h) : S(\mathbb{B}') \rightarrow S(\mathbb{B})$ given by $S(h)(\mathcal{U}) := h^{-1}(\mathcal{U})$.

Exercise 3.4.17. Verify the validity of the last sentence in the previous paragraph, namely that $h^{-1}(\mathcal{U})$ is indeed an ultrafilter on \mathbb{B} and that $S(h)$ is a continuous function.

Exercise 3.4.18. Using the notation from Exercise 3.4.10, show that $S(h_f)(\mathcal{U}) = f(\mathcal{U})$, the pushforward of \mathcal{U} along f as defined in Definition 1.3.1.

Exercise 3.4.19. Verify that S is indeed a contravariant functor from the category of Boolean algebras (with Boolean algebra homomorphisms) to the category of Stone spaces (with continuous functions) by verifying:

- (1) For all Boolean algebras \mathbb{B} , $S(\text{id}_{\mathbb{B}}) = \text{id}_{S(\mathbb{B})}$.
- (2) For all Boolean algebras \mathbb{B} , \mathbb{B}' , and \mathbb{B}'' and Boolean algebra homomorphisms $h : \mathbb{B} \rightarrow \mathbb{B}'$ and $h' : \mathbb{B}' \rightarrow \mathbb{B}''$, we have $S(h' \circ h) = S(h) \circ S(h')$.

The main result of this section is that S witnesses that the categories of Boolean algebras and Stone spaces are *dually equivalent*. In order to establish this result, we need to define an “inverse” functor from the category of Stone spaces to the category of Boolean algebras.

Definition 3.4.20. For any topological space X , we let $\text{Cl}(X)$ denote the set of clopen subsets of X .

Exercise 3.4.21. For any topological space X , $\text{Cl}(X)$ is a Boolean subalgebra of $\mathcal{P}(X)$. Moreover, X is discrete if and only if $\text{Cl}(X) = \mathcal{P}(X)$.

Just as in the case of S , Cl is a contravariant functor from the category of topological spaces to the category of Boolean algebras: If $f : X \rightarrow Y$ is a continuous function between topological spaces, we set $\text{Cl}(f) : \text{Cl}(Y) \rightarrow \text{Cl}(X)$ to be given by $\text{Cl}(f)(A) := f^{-1}(A)$. (By continuity of f , $f^{-1}(A)$ is indeed clopen when A is clopen.)

Exercise 3.4.22. Verify that $\text{Cl}(f)$ as above is indeed a Boolean algebra homomorphism.

Exercise 3.4.23. Verify that Cl is indeed a contravariant functor from the category of topological spaces to the category of Boolean algebras.

For some spaces, $\text{Cl}(X)$ is very small:

Exercise 3.4.24. Prove that a topological space X is connected if and only if $\text{Cl}(X) \cong \mathbf{2}$.

In order to show that Cl , when restricted to the full subcategory of Stone spaces, is indeed an inverse to S , we need to show that $\text{Cl}(X)$ is quite large when X is a Stone space.

Definition 3.4.25. Let X be a topological space and $x \in X$. We set $\mathcal{U}_x := \{A \in \text{Cl}(X) : x \in A\}$.

Exercise 3.4.26. \mathcal{U}_x is an ultrafilter on the Boolean algebra $\text{Cl}(X)$.

Note that, in the case that X is discrete, we have that this usage of the terminology \mathcal{U}_x agrees with the previous usage, namely the principal ultrafilter generated by x .

Lemma 3.4.27. *If X is a Stone space, then for every $x \in X$, $\bigcap \mathcal{U}_x = \{x\}$.*

Proof. Suppose, toward a contradiction, that $\bigcap \mathcal{U}_x$ contained at least two elements. Since X is totally disconnected, there are disjoint open sets $U_1, U_2 \subseteq X$ such that $\bigcap \mathcal{U}_x \cap U_1, \bigcap \mathcal{U}_x \cap U_2 \neq \emptyset$ and $\bigcap \mathcal{U}_x = (\bigcap \mathcal{U}_x \cap U_1) \cup (\bigcap \mathcal{U}_x \cap U_2)$. Since $\bigcap \mathcal{U}_x \cap (U_1 \cup U_2)^c = \emptyset$, by compactness there are finitely many $V_1, \dots, V_n \in \mathcal{U}_x$ such that $V_1 \cap \dots \cap V_n \cap (U_1 \cup U_2)^c = \emptyset$. Set $V := V_1 \cap \dots \cap V_n$. Suppose, without loss of generality, that $x \in U_1$. Note then that $(V \cap U_1)^c = (V \cap U_2) \cup V^c$, whence $V \cap U_1$ is a clopen set containing x , and thus $\bigcap \mathcal{U}_x \subseteq V \cap U_1 \subseteq U_1$, contradicting the fact that $\bigcap \mathcal{U}_x \cap U_2 \neq \emptyset$. \square

Lemma 3.4.28. *If X is a Stone space, then $\text{Cl}(X)$ forms a base for the topology on X (whence X is zero dimensional).*

Proof. Let U be an open subset of X and take $x \in U$. By the previous lemma, $\bigcap \mathcal{U}_x \cap U^c = \emptyset$, whence, by compactness, there are $U_1, \dots, U_n \in \mathcal{U}_x$ such that $U_1 \cap \dots \cap U_n \cap U^c = \emptyset$. It follows that $U_1 \cap \dots \cap U_n$ is a clopen neighborhood of x contained in U . \square

We now investigate what happens to objects in these categories when they are evaluated on the compositions of the above functors.

Lemma 3.4.29. *For any Boolean algebra \mathbb{B} , we have $\text{Cl}(S(\mathbb{B})) = \{U_a : a \in \mathbb{B}\}$. Moreover, $U_a = U_b$ if and only if $a = b$.*

Proof. Let U be a clopen subset of $S(\mathbb{B})$. Since U is open, we may write $U = \bigcup_a U_a$. Since U is closed (and thus compact), there are a_1, \dots, a_n such that $U = U_{a_1} \cup \dots \cup U_{a_n} = U_{a_1 \vee \dots \vee a_n}$.

For the moreover part, if $a \neq b$, then we have that $a \wedge \neg b \neq 0$ or $b \wedge \neg a \neq 0$. Without loss of generality, assume that it is the former. Then, by the ultrafilter theorem for Boolean algebras (Exercise 3.4.8 above), there is an ultrafilter \mathcal{U} on \mathbb{B} with $a \wedge \neg b \in \mathcal{U}$. It follows that $\mathcal{U} \in U_a \setminus U_b$. \square

Lemma 3.4.30. *For any Stone space X , we have $S(\text{Cl}(X)) = \{\mathcal{U}_x : x \in X\}$. Moreover, $\mathcal{U}_x = \mathcal{U}_y$ if and only if $x = y$.*

Proof. Fix $\mathcal{U} \in S(\text{Cl}(X))$. Since \mathcal{U} has the finite intersection property, we have that $\bigcap \mathcal{U} \neq \emptyset$ by compactness of X . Fixing $x \in \bigcap \mathcal{U}$, we have that $\mathcal{U} \subseteq \mathcal{U}_x$, whence $\mathcal{U} = \mathcal{U}_x$.

For the moreover part, assume that $x \neq y$. Then by Lemma 3.4.28, there is a clopen set $U \subseteq X$ such that $x \in U$ and $y \in X \setminus U$. It follows that $U \in \mathcal{U}_x \setminus \mathcal{U}_y$. \square

Before proving our main theorem, the following calculations will be relevant. Fix a Boolean algebra homomorphism $h : \mathbb{B} \rightarrow \mathbb{B}'$ and a continuous function between Stone spaces $f : X \rightarrow X'$. We then have that $(\text{Cl} \circ S)(h) : \text{Cl}(S(\mathbb{B})) \rightarrow \text{Cl}(S(\mathbb{B}'))$ is given by

$$(\text{Cl} \circ S)(h)(U_a) = \text{Cl}(S(h))(U_a) = S(h)^{-1}(U_a) = U_{h(a)}$$

and $(S \circ \text{Cl})(f) : S(\text{Cl}(X)) \rightarrow S(\text{Cl}(X'))$ is given by

$$(S \circ \text{Cl})(f)(\mathcal{U}_x) = S(\text{Cl}(f))(\mathcal{U}_x) = \text{Cl}(f)^{-1}(\mathcal{U}_x) = \mathcal{U}_{f(x)}.$$

We are now ready to prove our main theorem, stating that the functors S and Cl are “inverses” of one another. Note that this statement cannot literally be true as $(\text{Cl} \circ S)(\mathbb{B})$ is not literally \mathbb{B} and similarly $(S \circ \text{Cl})(X)$ is not literally X . The next theorem states that $(\text{Cl} \circ S)(\mathbb{B})$ is isomorphic to \mathbb{B} and $(S \circ \text{Cl})(X)$ is homeomorphic to X , and that these isomorphisms and homeomorphisms are “natural” in a precise sense.

Theorem 3.4.31 (Stone duality theorem). *S and Cl witness that the category of Boolean algebras and the category of Stone spaces are dually equivalent. More precisely, for every Boolean algebra \mathbb{B} and every Stone space X , we have an isomorphism $\epsilon_{\mathbb{B}} : (\text{Cl} \circ S)(\mathbb{B}) \rightarrow \mathbb{B}$ and a homeomorphism $\eta_X : (S \circ \text{Cl})(X) \rightarrow X$ satisfying, for every Boolean algebra homomorphism $h : \mathbb{B} \rightarrow \mathbb{B}'$ and every continuous function $f : X \rightarrow X'$, that $\epsilon_{\mathbb{B}'} \circ (\text{Cl} \circ S)(h) = h \circ \epsilon_{\mathbb{B}}$ and $\eta_{X'} \circ (S \circ \text{Cl})(f) = f \circ \eta_X$.*

Proof. By Lemma 3.4.29, we may define a bijection $\epsilon_{\mathbb{B}} : (\text{Cl} \circ \text{S})(\mathbb{B}) \rightarrow \mathbb{B}$ by $\epsilon_{\mathbb{B}}(U_a) := a$. It is also clear that this is a Boolean algebra homomorphism, whence an isomorphism. Also, the above calculation shows that $(\epsilon_{\mathbb{B}'} \circ (\text{Cl} \circ \text{S})(h))(U_a) = \epsilon_{\mathbb{B}'}(U_{h(a)}) = h(a) = (h \circ \epsilon_{\mathbb{B}})(U_a)$.

By Lemma 3.4.30, we can define $\eta_X : (\text{S} \circ \text{Cl})(X) \rightarrow X$ by $\eta_X(\mathcal{U}_x) = x$. It is clear that η_X is a bijection. η_X is continuous as $\eta_X^{-1}(U) = \bigcup_{a \in U} \mathcal{U}_a$, whence it is open. Since the domain and range of η_X are compact and Hausdorff, it follows that η_X is a homeomorphism. Finally, from the above calculation, we have

$$(\eta_{X'} \circ (\text{S} \circ \text{Cl})(f))(\mathcal{U}) = \eta_{X'}(\mathcal{U}_{f(x)}) = f(x) = (f \circ \eta_X)(\mathcal{U}_x). \quad \square$$

The following two results were promised earlier in this section:

Corollary 3.4.32. *Every Stone space X is the Stone space of a unique (up to isomorphism) Boolean algebra, namely $\text{Cl}(X)$.*

Corollary 3.4.33 (Stone representation theorem). *Every abstract Boolean algebra \mathbb{B} is isomorphic to a concrete Boolean algebra, namely $\text{Cl}(\text{S}(\mathbb{B}))$.*

Exercise 3.4.34. Prove, without using the axiom of choice, that the ultrafilter theorem is equivalent to the ultrafilter theorem for Boolean algebras.

3.5. Notes and References

As mentioned in Chapter 1, the notion of ultrafilter was introduced by H. Cartan [23, 24] in 1937 to study convergence in topological spaces. A more thorough treatment of topology using filters and ultrafilters can be found in Bourbaki [19]. The Stone-Čech compactification of a space was introduced by Stone [170] and Čech [25] in 1937. Our approach to the nondiscrete case follows Gillman and Jerison's book [67]. Stone's representation theorem and Stone duality were proven in his paper [169].

Ramsey theory and combinatorial number theory

In this chapter, we give a taste of some combinatorial applications of ultrafilters. In Section 4.1, we give an ultrafilter proof of the infinite version of Ramsey's theorem. In Section 4.2, we introduce a binary operation \oplus on $\beta\mathbb{Z}$ and prove that sets that belong to elements of $\beta\mathbb{Z}$ that are idempotent with respect to this operation have interesting combinatorial structure, ultimately leading to a proof of a celebrated theorem of Hindman. In Section 4.3, we introduce a measure of relative size of a set of integers known as Banach density and relate this notion of density to probability measures on $\beta\mathbb{Z}$; this analysis is then used in Section 4.4 to give a proof of the Furstenberg correspondence principle, which is a technique used to translate combinatorial questions about sets of integers into ergodic-theoretic questions in an associated dynamical system. We illustrate this technique with a couple of examples. In Section 4.5, we present a theorem of Jin known as the Sunset theorem. Instead of giving Jin's original proof, which was phrased in the language of nonstandard analysis (see Chapter 9), we give Beiglböck's proof, which uses the connection between Banach density and measures on $\beta\mathbb{Z}$ presented in Section 4.3.

4.1. Ramsey's theorem

In this section, we show how to use ultrafilters to give a very short proof of the infinite version of Ramsey's theorem. First, we need some notation:

Given a set X and $k \in \mathbb{N}$, we let $X^{[k]}$ denote the set of k -element subsets of X . When $X = \mathbb{N}$, we identify this with the set of tuples (n_1, \dots, n_k) with $n_1 < \dots < n_k$.

Theorem 4.1.1 (Ramsey's theorem, Infinite version). *For every $k \in \mathbb{N}$ and every partition $\mathbb{N}^{[k]} = C_1 \sqcup C_2$, there is $i \in \{1, 2\}$ and infinite $X \subseteq \mathbb{N}$ such that $X^{[k]} \subseteq C_i$.*

This theorem is often stated in terms of *colorings*: if we color the k -element subsets of \mathbb{N} in two colors, then there will be some infinite subset X of \mathbb{N} that is *homogeneous for the coloring*, that is, all k -element subsets from X receive the same color. Note that there is no real need to restrict to two colors in the above theorem; indeed, by induction, one could prove the same result using any finite number of colors. (Exercise.)

Proof of Theorem 4.1.1. For notational simplicity, let us assume that $k = 3$ and fix $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. For $(x, y) \in \mathbb{N}^{[2]}$, let $A_{(x,y)} := \{z \in \mathbb{N} : (x, y, z) \in C_1\}$. Without loss of generality, we may assume that $(\mathcal{U}x)(\mathcal{U}y)A_{(x,y)} \in \mathcal{U}$. (Otherwise we may switch the roles of C_1 and C_2 .) For $x \in \mathbb{N}$, let $B_x := \{y \in \mathbb{N} : A_{(x,y)} \in \mathcal{U}\}$ and let $C := \{x \in \mathbb{N} : B_x \in \mathcal{U}\}$. By assumption, we have that $C \in \mathcal{U}$. Fix $x_1 \in C$ arbitrary and take $x_2 \in B_{x_1} \cap C$ with $x_2 > x_1$; note that this is possible as $B_{x_1} \cap C \in \mathcal{U}$ and \mathcal{U} is nonprincipal. Now take $x_3 \in A_{(x_1, x_2)} \cap B_{x_1} \cap B_{x_2} \cap C$ with $x_3 > x_2$. Now choose x_4 belonging to

$$A_{(x_1, x_2)} \cap A_{(x_1, x_3)} \cap A_{(x_2, x_3)} \cap B_{x_1} \cap B_{x_2} \cap B_{x_3} \cap C$$

with $x_4 > x_3$. In this way, we construct an infinite set $X := \{x_1, x_2, \dots\}$ such that $X^{[k]} \subseteq C_1$, as desired. \square

Exercise 4.1.2. Turn the “vague” conclusion of the previous proof into something rigorous by formulating a precise inductive construction of the sequence and proving that this inductive construction can be continued indefinitely.

Exercise 4.1.3. Prove Ramsey's theorem for $k > 3$. (The proof is in the same spirit as ours above, the inductive construction being slightly more complicated to describe.)

In Section 6.4, we will show how to derive the finite version of Ramsey's theorem from the above infinite version.

4.2. Idempotent ultrafilters and Hindman's theorem

In this section, we use a special kind of ultrafilter on \mathbb{Z} to prove an important theorem in Ramsey theory known as Hindman's theorem. To state the theorem, we need some definitions.

Definition 4.2.1.

- (1) Given a nonempty, finite subset $F \subseteq \mathbb{N}$ and $c = (c_n)_{n \in \mathbb{Z}}$ a sequence of distinct elements from \mathbb{Z} , define $c_F := \sum_{n \in F} c_n$.
- (2) Given a sequence $c = (c_n)_{n \in \mathbb{N}}$ of distinct elements from \mathbb{Z} , set

$$\text{FS}(c) := \{c_F : F \subseteq \mathbb{N} \text{ nonempty, finite}\}.$$
- (3) We say that $A \subseteq \mathbb{Z}$ is an **FS-set** if there is a sequence $c = (c_n)_{n \in \mathbb{N}}$ of distinct elements from \mathbb{Z} such that $\text{FS}(c) \subseteq A$.

We are now ready to state the main combinatorial result of this section:

Theorem 4.2.2 (Hindman's theorem). *For every partition $\mathbb{Z} = C_1 \sqcup \dots \sqcup C_n$ of \mathbb{Z} , there is $i \in \{1, \dots, n\}$ such that C_i is an FS-set.*

The key to proving the above theorem is the realization that, in the case of the discrete space \mathbb{Z} , $\beta\mathbb{Z}$ has some algebraic structure in addition to the topological structure considered in the previous chapter. Indeed:

Definition 4.2.3. For $A \subseteq \mathbb{Z}$ and $\mathcal{U} \in \beta\mathbb{Z}$, set $A - \mathcal{U} := \{k \in \mathbb{Z} : A - k \in \mathcal{U}\}$.

In other words, $k \in A - \mathcal{U}$ if and only if $(\mathcal{U})(k + l \in A)$, that is, \mathcal{U} -many shifts of k are in A .

Definition 4.2.4. For $\mathcal{U}, \mathcal{V} \in \beta\mathbb{Z}$, we define $\mathcal{U} \oplus \mathcal{V} := \{A \subseteq \mathbb{Z} : A - \mathcal{V} \in \mathcal{U}\}$.

Exercise 4.2.5.

- (1) If $A \subseteq \mathbb{Z}$, then $A \in \mathcal{U} \oplus \mathcal{V}$ if and only if $(\mathcal{U}k)(\mathcal{V}l)(k + l \in A)$, that is, \mathcal{U} -many shifts of A are \mathcal{V} -large.
- (2) For $\mathcal{U}, \mathcal{V} \in \beta\mathbb{Z}$, we have $\mathcal{U} \oplus \mathcal{V} \in \beta\mathbb{Z}$.
- (3) \oplus is associative: for $\mathcal{U}, \mathcal{V}, \mathcal{W} \in \beta\mathbb{Z}$, $(\mathcal{U} \oplus \mathcal{V}) \oplus \mathcal{W} = \mathcal{U} \oplus (\mathcal{V} \oplus \mathcal{W})$.
- (4) For $\mathcal{U}, \mathcal{V} \in \beta\mathbb{Z}$, $\mathcal{U} \oplus \mathcal{V}$ is principal if and only if both \mathcal{U} and \mathcal{V} are principal. In this case, if \mathcal{U} and \mathcal{V} are the principal ultrafilters generated by k and l , respectively, then $\mathcal{U} \oplus \mathcal{V}$ is the principal ultrafilter generated by $k + l$, that is, $\mathcal{U}_k \oplus \mathcal{U}_l = \mathcal{U}_{k+l}$.
- (5) $\mathcal{U} \oplus \mathcal{U}_0 = \mathcal{U}$ for all $\mathcal{U} \in \beta\mathbb{Z}$.

The algebraic properties of \oplus are quite complicated and interesting; see the encyclopedia on the subject [83]. For the proof of Hindman's theorem, we need to know that \oplus behaves well with respect to the topology on $\beta\mathbb{Z}$. It would be nice if we could just say that \oplus is a continuous function, but unfortunately that is not true. Thankfully, something weaker will be good enough:

Proposition 4.2.6. \oplus is left semicontinuous, meaning that for all $\mathcal{V} \in \beta\mathbb{Z}$, the map $\mathcal{U} \mapsto \mathcal{U} \oplus \mathcal{V} : \beta\mathbb{Z} \rightarrow \beta\mathbb{Z}$ is continuous.

Proof. Fix $\mathcal{V} \in \beta\mathbb{Z}$ and a basic open set U_A in $\beta\mathbb{Z}$. Its inverse image under $\mathcal{U} \mapsto \mathcal{U} \oplus \mathcal{V}$ is $U_{A-\mathcal{V}}$, which is open. \square

A crucial role in the use of ultrafilters in combinatorics is played by the following kind of ultrafilter:

Definition 4.2.7. $\mathcal{U} \in \beta\mathbb{Z}$ is **idempotent** if $\mathcal{U} \oplus \mathcal{U} = \mathcal{U}$.

Exercise 4.2.8. Suppose that $\mathcal{U} \in \beta\mathbb{Z}$ is idempotent and principal. Prove that $\mathcal{U} = \mathcal{U}_0$.

Theorem 4.2.9. *Nonprincipal idempotent ultrafilters exist.*

Proof. Let $\mathcal{Y} := \{Y \subseteq \beta\mathbb{Z} \setminus \mathbb{Z} : Y \text{ is nonempty, closed, and } Y \oplus Y \subseteq Y\}$. Here, $Y \oplus Y := \{\mathcal{U} \oplus \mathcal{V} : \mathcal{U}, \mathcal{V} \in Y\}$. By Theorem 3.2.1(4) and Exercise 4.2.5(4), $\beta\mathbb{Z} \setminus \mathbb{Z} \in \mathcal{Y}$. By compactness of $\beta\mathbb{Z} \setminus \mathbb{Z}$, any descending chain in \mathcal{Y} has nonempty intersection. It follows from Zorn's lemma that \mathcal{Y} has a minimal element Y_0 . We show that every element of Y_0 is idempotent (and nonprincipal).

Fix $\mathcal{U} \in Y_0$. We first claim that $Y_0 \oplus \mathcal{U} \in \mathcal{Y}$. Indeed, it is clearly nonempty. By left semicontinuity, it is closed. Finally, using associativity of \oplus and the fact that Y_0 is closed under \oplus , we have

$$(Y_0 \oplus \mathcal{U}) \oplus (Y_0 \oplus \mathcal{U}) \subseteq (Y_0 \oplus Y_0) \oplus (Y_0 \oplus \mathcal{U}) \subseteq Y_0 \oplus \mathcal{U}.$$

Since $Y_0 \oplus \mathcal{U} \subseteq Y_0$, minimality of Y_0 implies that $Y_0 \oplus \mathcal{U} = Y_0$. Set

$$Y_1 := \{\mathcal{V} \in Y_0 : \mathcal{V} \oplus \mathcal{U} = \mathcal{U}\}.$$

We just showed that $Y_1 \neq \emptyset$. Note also that Y_1 is closed as it is the preimage of the closed set $\{\mathcal{U}\}$ under the continuous map $\mathcal{V} \mapsto \mathcal{V} \oplus \mathcal{U}$. Finally, if $\mathcal{V}_1, \mathcal{V}_2 \in Y_1$, then

$$(\mathcal{V}_1 \oplus \mathcal{V}_2) \oplus \mathcal{U} = \mathcal{V}_1 \oplus (\mathcal{V}_2 \oplus \mathcal{U}) = \mathcal{V}_1 \oplus \mathcal{U} = \mathcal{U},$$

so $Y_1 \oplus Y_1 \subseteq Y_1$. It follows that $Y_1 \in \mathcal{Y}$; by minimality of Y_0 again, we have that $Y_1 = Y_0$, whence $\mathcal{U} \in Y_1$ and $\mathcal{U} \oplus \mathcal{U} = \mathcal{U}$, as desired. \square

Here is the connection between idempotent ultrafilters and FS-sets:

Theorem 4.2.10. *If \mathcal{U} is a nonprincipal idempotent ultrafilter, then every $A \in \mathcal{U}$ is an FS-set.*

Proof. We define a sequence $c = (c_n)_{n \in \mathbb{N}}$ in strictly increasing order so that $\text{FS}(c) \subseteq A$.

Suppose inductively that $c_1 < \dots < c_n$ have been constructed so that, for all nonempty $F \subseteq \{1, \dots, n\}$, we have $c_F \in A$ and $A - c_F \in \mathcal{U}$. We show how to continue this construction. First note that since \mathcal{U} is idempotent, we can conclude that $A - c_F - \mathcal{U} \in \mathcal{U}$. Using the fact that \mathcal{U} is nonprincipal

(and idempotent again), we may take $c_{n+1} > c_n$ so that $c_{n+1} \in A \cap (A - \mathcal{U}) \cap (A - c_F) \cap (A - c_F - \mathcal{U})$ for every F as above. Note then that c_1, \dots, c_{n+1} is a valid continuation of the construction. \square

The proof of Hindman's theorem is now immediate:

Proof of Theorem 4.2.2. Fix a partition $\mathbb{Z} = C_1 \sqcup \dots \sqcup C_n$ and a non-principal idempotent ultrafilter \mathcal{U} . Since $\mathbb{Z} \in \mathcal{U}$, there is some $i \in \{1, \dots, n\}$ such that $C_i \in \mathcal{U}$, whence C_i is an FS-set. \square

If we work a bit harder, we can improve the previous statement. First, a lemma:

Lemma 4.2.11. *Suppose that $c = (c_n)_{n \in \mathbb{N}}$ is a sequence of distinct elements from \mathbb{Z} . Then there is a nonprincipal idempotent ultrafilter \mathcal{U} on \mathbb{Z} such that $\text{FS}(c) \in \mathcal{U}$.*

Proof. For each $m \in \mathbb{N}$, let $c^m := (c_n)_{n \geq m}$ and set $U_m := \{\mathcal{U} \in \beta\mathbb{Z} \setminus \mathbb{Z} : \text{FS}(c^m) \in \mathcal{U}\}$, a nonempty closed subset of $\beta\mathbb{Z}$. Since $U_{m+1} \subseteq U_m$, by compactness we have that $S := \bigcap_{m \in \mathbb{N}} U_m$ is a nonempty closed subset of $\beta\mathbb{Z}$.

Claim. $S \oplus S \subseteq S$.

Proof of Claim. Suppose that $\mathcal{U}, \mathcal{V} \in S$; we show that $\mathcal{U} \oplus \mathcal{V} \in S$. To see this, fix $m \in \mathbb{N}$. We must show that $\text{FS}(c^m) \in \mathcal{U} \oplus \mathcal{V}$. Fix $a \in \text{FS}(c^m)$ and write $a = c_{n_1} + \dots + c_{n_l}$ with $m \leq n_1 < n_2 < \dots < n_l$. Note then that $\text{FS}(c^{n_i+1}) \subseteq \text{FS}(c^m) - a$. Since $\text{FS}(c^{n_i+1}) \in \mathcal{V}$, it follows that $\text{FS}(c^m) - a \in \mathcal{V}$. Since $a \in \text{FS}(c^m)$ was arbitrary and $\text{FS}(c^m) \in \mathcal{U}$, it follows that $\text{FS}(c^m) - \mathcal{V} \in \mathcal{U}$, whence $\text{FS}(c^m) \in \mathcal{U} \oplus \mathcal{V}$, as desired. This proves the claim.

We may thus repeat the proof of Theorem 4.2.9, but this time only using ultrafilters from S ; it follows that the nonprincipal idempotent ultrafilter thus constructed belongs to S , which, in particular, implies that $\text{FS}(c) \in \mathcal{U}$. \square

Corollary 4.2.12 (Strong Hindman's theorem). *Suppose that C is an FS-set and C is partitioned into finitely many pieces $C = C_1 \sqcup \dots \sqcup C_n$. Then there is $i \in \{1, \dots, n\}$ such that C_i is an FS-set.*

Proof. Take $c = (c_n)_{n \in \mathbb{N}}$ such that $\text{FS}(c_n) \subseteq C$. Take a nonprincipal idempotent ultrafilter \mathcal{U} such that $\text{FS}(c) \in \mathcal{U}$. Then $C \in \mathcal{U}$ as well, whence $C_i \in \mathcal{U}$ for a unique $i = 1, \dots, n$, and this C_i is itself an FS-set. \square

In combinatorial terms, the above corollary shows that being an FS-set is a **partition regular** property.

Corollary 4.2.13. $\mathcal{U} \in \beta\mathbb{Z}$ is in the closure of the set of nonprincipal idempotent ultrafilters if and only if every element of \mathcal{U} is an FS-set.

Proof. First suppose that \mathcal{U} is in the closure of the set of nonprincipal idempotent ultrafilters and suppose that $A \in \mathcal{U}$. Then U_A is a neighborhood of \mathcal{U} , whence there is $\mathcal{V} \in U_A$ that is idempotent. Thus, $A \in \mathcal{V}$, whence A is an FS-set by Theorem 4.2.10.

Conversely, suppose that every element of \mathcal{U} is an FS-set. Let U_A be a basic neighborhood of \mathcal{U} . By assumption, A is an FS-set, whence $A \in \mathcal{V}$ for some nonprincipal idempotent ultrafilter \mathcal{V} by Lemma 4.2.11. It follows that $\mathcal{V} \in U_A$. Since U_A is an arbitrary basic open neighborhood of \mathcal{U} , it follows that \mathcal{U} is in the closure of the set of nonprincipal idempotent ultrafilters. \square

4.3. Banach density, means, and measures

A central notion in combinatorial number theory is the attempt to describe “what proportion” of the integers lie in a given set $A \subseteq \mathbb{Z}$. Since the set A is presumably infinite, one cannot simply just consider the fraction $\frac{|A|}{|\mathbb{Z}|}$. Instead, one must first consider what proportion of A lies in various finite intervals and then use some limiting process. The resulting quantity is referred to as a *density* of A . While there are many notions of densities for subsets of \mathbb{Z} , we will only concern ourselves with the following:

Definition 4.3.1. Given $A \subseteq \mathbb{Z}$, the **Banach density** of A is the quantity

$$\text{BD}(A) := \lim_{n \rightarrow \infty} \max \left\{ \frac{|A \cap I|}{n} : I \subseteq \mathbb{Z} \text{ an interval of length } n \right\}.$$

A few words are in order concerning the previous definition. Note that, given an interval $I \subseteq \mathbb{Z}$ of length n , $\frac{|A \cap I|}{n}$ belongs to the finite set $\{0, \frac{1}{n}, \frac{2}{n}, \dots, 1\}$. Thus, as I ranges over all such intervals, the above fraction takes a maximal value, which we temporarily denote by $\delta(A, n)$.

Exercise 4.3.2. For $A \subseteq \mathbb{Z}$ and $m, n \in \mathbb{N}$, prove that $\delta(A, m+n) \leq \delta(A, m) + \delta(A, n)$.

An elementary real analysis fact known as *Fekete’s lemma* allows one to conclude that $\lim_{n \rightarrow \infty} \delta(A, n)$ actually exists. This limit is then the Banach density of A .

The following is a notion of largeness for subsets of \mathbb{Z} :

Definition 4.3.3. $A \subseteq \mathbb{Z}$ is called **thick** if it contains arbitrarily long intervals, that is, for each $n \in \mathbb{N}$, there is an interval $I \subseteq \mathbb{Z}$ with $|I| \geq n$ such that $I \subseteq A$.

Exercise 4.3.4. For $A \subseteq \mathbb{Z}$, prove that $\text{BD}(A) = 1$ if and only if A is thick.

Definition 4.3.5. Given $A \subseteq \mathbb{Z}$, we say that a sequence of intervals (I_n) in \mathbb{Z} witness the Banach density of A if $\lim_{n \rightarrow \infty} |I_n| = \infty$ and $\lim_{n \rightarrow \infty} \frac{|A \cap I_n|}{|I_n|} = \text{BD}(A)$.

While the notion of Banach density satisfies many natural properties similar to those satisfied by measures, it differs from measures in many unsatisfying ways. For example:

Exercise 4.3.6. Show that Banach density is subadditive, that is, for $A, B \subseteq \mathbb{Z}$, we have $\text{BD}(A \cup B) \leq \text{BD}(A) + \text{BD}(B)$. Show that, even when A and B are disjoint, $\text{BD}(A \cup B)$ need not equal $\text{BD}(A) + \text{BD}(B)$.

That is, while measures are additive (even countably additive), densities are merely subadditive. In the remainder of this section, we will show how ultrafilters can be used to bring actual measures into the picture when discussing density calculations.

Measures will be introduced using the *Riesz representation theorem*. To explain this, first suppose that X is a compact Hausdorff space and μ is a finite Borel measure on X . Let $C(X)$ denote the \mathbb{R} -vector space of continuous functions $X \rightarrow \mathbb{R}$. Then we can consider the integration functional $\ell_\mu : C(X) \rightarrow \mathbb{R}$ given by $\ell_\mu(f) := \int f d\mu$. The function ℓ_μ is an \mathbb{R} -linear map which is moreover *positive*, meaning that if $f \geq 0$, then $\ell_\mu(f) \geq 0$. Also notice that $\ell_\mu(1) = \mu(X)$, where 1 denotes the function on X that is constantly 1. In particular, μ is a probability measure if and only if $\ell(1) = 1$.

The Riesz representation theorem gives a converse to the previous paragraph: given any positive linear functional $\ell : C(X) \rightarrow \mathbb{R}$, there is a finite Borel measure μ on X such that $\ell = \ell_\mu$. (μ is actually unique if one assumes that it is a so-called regular measure.)

The above discussion motivates the need to produce positive linear functionals on $C(X)$ for some compact space X . While \mathbb{Z} is not compact, $\beta\mathbb{Z}$ is compact, and it is this compact space for which we will apply the Riesz representation theorem.

Given any set X , let $B(X)$ denote the \mathbb{R} -vector space of bounded functions $X \rightarrow \mathbb{R}$.

Exercise 4.3.7. Prove that the map $f \mapsto \beta f : B(X) \rightarrow C(\beta X)$ is an isomorphism of \mathbb{R} -vector spaces. Moreover, prove that:

- (1) for $f \in B(X)$, we have $f \geq 0$ if and only if $\beta f \geq 0$, and
- (2) $\beta 1 = 1$.

Motivated by the previous discussion, we consider the following:

Definition 4.3.8. A **mean** on X is a positive linear functional $\ell : B(X) \rightarrow \mathbb{R}$ such that $\ell(1) = 1$.

We summarize this discussion as follows:

Theorem 4.3.9. *For every mean ℓ on X , there is a Borel probability measure μ on βX such that $\ell(f) = \int_{\beta X} \beta f d\mu$ for every $f \in B(X)$.*

Thus, in order to produce measures on $\beta\mathbb{Z}$, we need a way of constructing means on \mathbb{Z} . The following is our main method of accomplishing this task:

Proposition 4.3.10. *Suppose that $\mathcal{I} := (I_n)_{n \in \mathbb{N}}$ is a sequence of nonempty, finite subsets of \mathbb{Z} and \mathcal{U} is an ultrafilter on \mathbb{N} . For $f \in B(\mathbb{Z})$, define $\ell(f) := \ell_{\mathcal{I}, \mathcal{U}}(f) := \lim_{\mathcal{U}} \frac{1}{|I_n|} \sum_{x \in I_n} f(x)$. Then ℓ is a mean on \mathbb{Z} .*

Exercise 4.3.11. Prove the previous proposition.

In Proposition 4.3.10, we think of $\frac{1}{|I_n|} \sum_{x \in I_n} f(x)$ as the average of $f(x)$ on the set I_n and then use the ultrafilter to see what those averages converge to.

We will need to consider means on \mathbb{Z} satisfying an extra property:

Definition 4.3.12. If ℓ is a mean on \mathbb{Z} , we say that ℓ is an **invariant mean** if $\ell(k.f) = \ell(f)$ for all $k \in \mathbb{Z}$ and $f \in B(\mathbb{Z})$, where $(k.f)(x) := f(x - k)$.

Exercise 4.3.13. Suppose, using the notation from Proposition 4.3.10, that each I_n is an interval in \mathbb{Z} and $\lim_{n \rightarrow \infty} |I_n| = \infty$. Further suppose that \mathcal{U} is nonprincipal. Prove that $\ell_{\mathcal{I}, \mathcal{U}}$ is an invariant mean on \mathbb{Z} .

The invariant means $\ell_{\mathcal{I}, \mathcal{U}}$ have a connection to Banach density:

Exercise 4.3.14. Suppose that \mathcal{I} and \mathcal{U} are as in Exercise 4.3.13 and that $A \subseteq \mathbb{Z}$.

- (1) Show that $\text{BD}(A) \geq \ell_{\mathcal{I}, \mathcal{U}}(1_A)$, where 1_A denotes the characteristic function of A .
- (2) Suppose that \mathcal{I} witnesses the Banach density of A . Prove that $\text{BD}(A) = \ell_{\mathcal{I}, \mathcal{U}}(1_A)$.

To round out the discussion, we mention the following:

Lemma 4.3.15. *For any subset $A \subseteq \mathbb{Z}$, we have $\beta(1_A) = 1_{U_A}$.*

Exercise 4.3.16. Prove the previous lemma.

Corollary 4.3.17. *For any mean ℓ on \mathbb{Z} , there is a Borel probability measure μ on $\beta\mathbb{Z}$ such that, for any $A \subseteq \mathbb{Z}$, we have $\ell(1_A) = \mu(U_A)$.*

We will use the fact that Banach densities can be mirrored by genuine measures in the next two sections.

4.4. Furstenberg's correspondence principle

In this section, we explain an extremely powerful extension of the ideas used in the previous section that has proven to be a valuable tool in the application of ideas from *ergodic theory* to combinatorial number theory. First, some definitions regarding measure-preserving transformations.

Definition 4.4.1. Given a probability space (Ω, μ) , a measurable map $T : \Omega \rightarrow \Omega$ is called a **measure-preserving transformation** if, for each measurable $E \subseteq \Omega$, we have $\mu(E) = \mu(T^{-1}(E))$.

To avoid multiple parentheses in the next result, for $n \in \mathbb{Z}$, we will write $T^{-n}E$ instead of the more cumbersome $T^{-n}(E)$. Note that, for $n > 0$, this means those $x \in \Omega$ for which $T^n(x) \in E$, where T^n denotes the map $T \circ T \circ \cdots \circ T$ (n times), while for $n < 0$, this means the image of E under the map T^{-n} .

Theorem 4.4.2 (Furstenberg's correspondence principle). *Given $A \subseteq \mathbb{Z}$, there is a probability space (Ω, μ) , a measure-preserving transformation $T : \Omega \rightarrow \Omega$, and a measurable set E such that:*

- (1) $\mu(E) = \text{BD}(A)$ and
- (2) for all $l_1, \dots, l_k \in \mathbb{Z}$, we have

$$\text{BD}(A \cap (A - l_1) \cap \cdots \cap (A - l_k)) \geq \mu(E \cap T^{-l_1}E \cap \cdots \cap T^{-l_k}E).$$

Proof. Set $\Omega := \beta\mathbb{Z}$. Let \mathcal{I} witness the Banach density of A , and let \mathcal{U} be a nonprincipal ultrafilter on \mathbb{N} . Set $\ell := \ell_{\mathcal{I}, \mathcal{U}}$ and let μ be the probability measure on $\beta\mathbb{Z}$ as in Corollary 4.3.17 corresponding to ℓ , except that we restrict to the σ -algebra generated by the basic open sets in $\beta\mathbb{Z}$. Let $T : \Omega \rightarrow \Omega$ be given by $T(\mathcal{U}) := \mathcal{U} \oplus \mathcal{U}_1$ and let $E := U_A$. We show that these are as desired.

First, observe that $T^{-1}(U_B) := U_{B-1}$ for any $B \subseteq \mathbb{Z}$. Indeed, $\mathcal{U} \in T^{-1}(U_B)$ if and only if $\mathcal{U} \oplus \mathcal{U}_1 \in U_B$ if and only if $B \in \mathcal{U} \oplus \mathcal{U}_1$ if and only if $B - 1 \in \mathcal{U}$ if and only if $\mathcal{U} \in U_{B-1}$. In particular, T is a measurable map. Also, it follows by induction that $T^{-n}(U_B) = U_{B-n}$ for any $n > 0$. Consequently, $T^{-l}(U_B) = U_{B-l}$ for any $l \in \mathbb{Z}$.

Next observe that $\mu(T^{-1}U_B) = \mu(U_{B-1}) = \ell(1_{B-1}) = \ell(1_B) = \mu(U_B)$, where the equality $\ell(1_{B-1}) = \ell(1_B)$ follows from invariance of ℓ . Since we have restricted attention to the σ -algebra generated by the basic open sets, it follows that T is measure-preserving.

Item (1) now follows from the choice of μ . To see (2), fix $l_1, \dots, l_k \in \mathbb{Z}$. Since

$$E \cap T^{-l_1}E \cap \cdots \cap T^{-l_k}E = U_{A \cap (A - l_1) \cap \cdots \cap (A - l_k)},$$

the desired inequality follows from Exercise 4.3.14(1). \square

Furstenberg originally introduced his correspondence principle to give an ergodic theoretic proof of Szemerédi's famous theorem on arithmetic progressions. Recall that an arithmetic progression in \mathbb{Z} of length k is a sequence of the form $a, a + n, \dots, a + (k - 1)n$, where $a \in \mathbb{Z}$ and $n > 0$. Szemerédi proved the following amazing theorem:

Theorem 4.4.3 (Szemerédi's theorem). *If $A \subseteq \mathbb{Z}$ is such that $\text{BD}(A) > 0$, then for any $k \in \mathbb{N}$, there is an arithmetic progression of length k contained in A .*

Having positive Banach density is a very mild condition assuring that a set is “not too sparse”. For example, a set A having Banach density $\frac{1}{100}$ means that there are longer and longer intervals in \mathbb{Z} in which the proportion of the interval that lies in A hovers around $\frac{1}{100}$. While such a set can seem quite small, Szemerédi's theorem ensures that it must contain arbitrarily long arithmetic progressions.

Szemerédi's original proof was combinatorial and very (very!) complicated. Furstenberg's ergodic theoretic proof is conceptually much simpler, combining the above correspondence principle with the following (difficult) theorem in ergodic theory:

Theorem 4.4.4 (Furstenberg's recurrence theorem). *For any probability space (Ω, μ) , any measure-preserving transformation $T : \Omega \rightarrow \Omega$, any measurable $E \subseteq \Omega$ with $\mu(E) > 0$, and any $k \in \mathbb{N}$, there is $n > 0$ such that $\mu(E \cap T^{-n}E \cap T^{-2n}E \cap \dots \cap T^{-kn}E) > 0$.*

Szemerédi's theorem is now an immediate consequence of the correspondence principle and the recurrence theorem. Indeed, let Ω, μ, E , and T be as in the Furstenberg correspondence principle for A , and let $n > 0$ be as in the Furstenberg recurrence theorem for this data and $k - 1$. Then

$$\text{BD}(A \cap (A - n) \cap (A - 2n) \cap \dots \cap (A - (k - 1)n)) > 0.$$

In particular, there is $x \in A \cap (A - n) \cap (A - 2n) \cap \dots \cap (A - (k - 1)n)$, that is, $x, x + n, x + 2n, \dots, x + (k - 1)n$, all belong to A , whence we have found an arithmetic progression of length k in A .

We conclude this section with another application of Furstenberg's correspondence principle. First, we need another structural notion of largeness for subsets of \mathbb{Z} :

Definition 4.4.5. $B \subseteq \mathbb{Z}$ is called **syndetic** if there is $m \in \mathbb{N}$ such that, whenever $I \subseteq \mathbb{Z}$ is an interval with $B \cap I = \emptyset$, then $|I| \leq m$.

In other words, B is syndetic if there is a uniform bound on the size of gaps of B . We use the Furstenberg correspondence principle to prove the following:

Theorem 4.4.6. *Suppose that $A \subseteq \mathbb{Z}$ is such that $\text{BD}(A) > 0$. Then the difference set*

$$A - A := \{x - y : x, y \in A\}$$

is syndetic.

We accomplish this as follows.

Definition 4.4.7. A Δ_r -set is a set of the form $\{n_j - n_i : 1 \leq i < j \leq r\}$, where $n_1 < \dots < n_r$ are nonnegative integers. $A \subseteq \mathbb{Z}$ is called a Δ_r^* -set if $A \cap B \neq \emptyset$ for every Δ_r -set $B \subseteq \mathbb{Z}$.

Exercise 4.4.8. Prove that Δ_r^* -sets are syndetic.

Theorem 4.4.9. *Suppose that (Ω, μ) is a probability space and $T : \Omega \rightarrow \Omega$ is measure-preserving. Then for any measurable set $E \subseteq \Omega$ with $\mu(E) > 0$, the return set*

$$R_E := \{n \in \mathbb{Z} : \mu(E \cap T^{-n}E) > 0\}$$

is a Δ_r^ -set for any $r > \frac{1}{\mu(A)}$. In particular, R_E is piecewise syndetic.*

Proof. Fix nonnegative integers $n_1 < \dots < n_r$. Suppose that

$$R_E \cap \{n_j - n_i : 1 \leq i < j \leq r\} = \emptyset.$$

Then, for all $1 \leq i < j \leq r$, we have that $\mu(T^{-n_i}E \cap T^{-n_j}E) = \mu(E \cap T^{-(n_j - n_i)}E) = 0$, so $1 \geq \mu(\bigcup_{i=1}^r \mu(T^{-n_i}E)) = r \cdot \mu(E)$, whence $r \leq \frac{1}{\mu(E)}$. \square

Proof of Theorem 4.4.6. Let Ω , μ , E , and T be as in the Furstenberg correspondence principle for A . If $n \in R_E$, then $\text{BD}(A \cap (A - n)) > 0$, so $n \in A - A$. Thus $A - A$ contains the syndetic set R_E , so is itself syndetic. \square

4.5. Jin's sumset theorem

In this section, we prove the following theorem of Jin, which is in a similar spirit as Theorem 4.4.6. First, we need yet another structural notion of largeness for subsets of \mathbb{Z} :

Definition 4.5.1. $C \subseteq \mathbb{Z}$ is called **piecewise syndetic** if there is $m \in \mathbb{N}$ and intervals I_1, I_2, \dots in \mathbb{Z} satisfying:

- (1) $\lim_{n \rightarrow \infty} |I_n| = \infty$, and
- (2) for any n and any interval $J \subseteq I_n$, if $C \cap J = \emptyset$, then $|J| \leq m$.

We thus see that piecewise syndeticity is a weakening of the notion of syndeticity in that we only require that the gaps of C have bounded size on longer and longer intervals in \mathbb{Z} . Although at first it might seem like a strange notion, it is of extreme importance in combinatorial number theory

and, in some sense, is a more natural notion. For example, one can prove that piecewise syndeticity is a partition regular notion (as defined in Section 4.2), while neither thickness nor syndeticity are partition regular.

In what follows, given $X, Y \subseteq \mathbb{Z}$, we define their sumset to be

$$X + Y := \{x + y : x \in X, y \in Y\}.$$

Exercise 4.5.2. Prove that $C \subseteq \mathbb{Z}$ is piecewise syndetic if and only if there is a finite set $F \subseteq \mathbb{Z}$ such that, for all $n \in \mathbb{N}$, there is an interval $I \subseteq F + C$ with $|I| \geq n$.

Here is the main result of this section:

Theorem 4.5.3 (Jin's sumset theorem). *Suppose that $A, B \subseteq \mathbb{Z}$ are such that $\text{BD}(A), \text{BD}(B) > 0$. Then $A + B$ is piecewise syndetic.*

Jin's original theorem used nonstandard analysis, a topic we discuss in Chapter 9. In fact, Jin's theorem was one of the early successes of non-standard analysis applied to combinatorial number theory and this area of research is currently extremely active; see [42] for an entire monograph on the subject.

Instead of following Jin's original proof, we will use an ultrafilter proof due to Beiglböck which makes substantial use of analysis on $\beta\mathbb{Z}$ and the conversion from densities to means described in Section 4.3. The following lemma is the key to Beiglböck's proof of Theorem 4.5.3.

Lemma 4.5.4. *For any $A, B \subseteq \mathbb{Z}$, there is $\mathcal{U} \in \beta\mathbb{Z}$ such that*

$$\text{BD}(A \cap (B - \mathcal{U})) \geq \text{BD}(A) \cdot \text{BD}(B).$$

Proof. Fix an invariant mean ℓ on \mathbb{Z} such that $\ell(1_B) = \text{BD}(B)$, and let μ be the associated Borel probability measure on $\beta\mathbb{Z}$. Let $(I_n)_{n \in \mathbb{N}}$ be a sequence of intervals witnessing the Banach density of A . Define $f_n : \beta\mathbb{Z} \rightarrow [0, 1]$ by

$$f_n(\mathcal{U}) := \frac{1}{|I_n|} \sum_{k \in A \cap I_n} 1_{U_{B-k}}(\mathcal{U}).$$

Note that each f_n is measurable. Set $f(\mathcal{U}) := \limsup_n f_n(\mathcal{U})$ (which is also measurable) and note that $f(\mathcal{U}) \leq \text{BD}(A \cap (B - \mathcal{U}))$ for all $\mathcal{U} \in \beta\mathbb{Z}$. Fatou's lemma implies

$$\int_{\beta\mathbb{Z}} f d\mu \geq \limsup_n \int_{\beta\mathbb{Z}} \frac{1}{|I_n|} \sum_{k \in A \cap I_n} 1_{U_{B-k}} d\mu = \limsup_n \frac{1}{|I_n|} \sum_{k \in I_n \cap A} \ell(1_{B-k}).$$

Since ℓ is invariant, the latter term is equal to $\limsup_n \frac{|A \cap I_n|}{|I_n|} \cdot \ell(1_B) = \text{BD}(A) \cdot \text{BD}(B)$. Thus, we have shown $\int_{\beta\mathbb{Z}} f d\mu \geq \text{BD}(A) \cdot \text{BD}(B)$. In particular, there is some $\mathcal{U} \in \beta\mathbb{Z}$ such that $f(\mathcal{U}) \geq \text{BD}(A) \cdot \text{BD}(B)$, as desired. \square

Notice that, in the notation of the above proof, $\mu(\mathbb{Z}) = 0$, whence we can take \mathcal{U} as in the conclusion of the lemma to be nonprincipal.

Beiglböck’s proof of Theorem 4.5.3. Assume that $\text{BD}(A), \text{BD}(B) > 0$. Apply Lemma 4.5.4 with A replaced by $-A$ (which has the same Banach density), obtaining $\mathcal{U} \in \beta\mathbb{Z}$ such that $C := (-A) \cap (B - \mathcal{U})$ has positive Banach density. By Theorem 4.4.6, $C - C$ is syndetic; since $C - C \subseteq A + (B - \mathcal{U})$, we have that $A + (B - \mathcal{U})$ is also syndetic.

Suppose $s \in A + (B - \mathcal{U})$. Then for some $a \in A$, $B - (s - a) \in \mathcal{U}$, whence $a + B - s \in \mathcal{U}$ and hence $A + B - s \in \mathcal{U}$. Thus, for any finite set $s_1, \dots, s_n \in A + (B - \mathcal{U})$, we have $\bigcap_{i=1}^n (A + B - s_i) \in \mathcal{U}$, and, in particular, is nonempty, meaning there is $t \in \mathbb{Z}$ such that $t + \{s_1, \dots, s_n\} \subseteq A + B$. We claim that this implies that $A + B$ is piecewise syndetic. Indeed, take $F \subseteq \mathbb{Z}$ such that $F + A + (B - \mathcal{U}) = \mathbb{Z}$. By Exercise 4.5.2, it suffices to check that $F + A + B$ contains arbitrarily long intervals. To see this, fix $n \in \mathbb{N}$ and, for $i = 1, \dots, n$, take $s_i \in A + (B - \mathcal{U})$ such that $i \in F + s_i$. Take $t \in \mathbb{Z}$ such that $t + \{s_1, \dots, s_n\} \subseteq A + B$. Then $t + [1, n] \subseteq t + F + \{s_1, \dots, s_n\} \subseteq F + (A + B)$, completing the proof. \square

4.6. Notes and references

An entire book devoted to the subject matter of this chapter is [42]. Ramsey’s theorem was proven in his paper [146]. Hindman’s theorem was proven in his paper [82]. However, Hindman’s original proof was purely combinatorial and very difficult to follow. (Hindman himself even once suggested that one could torture a graduate student by asking them to read the original proof.) That being said, Baumgartner gave a short combinatorial proof in [4]. It had been previously observed by Galvin (see also [81]) that the existence of an idempotent ultrafilter (which was unknown at the time) yields the conclusion of Hindman’s theorem. The existence of idempotent ultrafilters was later established by Glazer; see [33]. The idempotent ultrafilter proof of Hindman’s theorem paved the way for many results in Ramsey theory; see [177] and [42].

The connection between measures on $\beta\mathbb{Z}$ and Banach density can be found, for example, in Bergelson’s article [11]. Furstenberg’s correspondence principle was originally proven in [62], where he gave his proof of Szemerédi’s theorem, which itself was originally proven in [171]. Our proof of Theorem 4.4.6 is based on the article [12]. Jin’s sumset theorem was originally proven in [92]; our proof is based on that of Beiglböck given in [8]. Many generalizations of Jin’s sumset theorem have been proven over the years; see [42].

Foundational concerns

In this chapter, we consider some foundational concerns related to the existence of nonprincipal ultrafilters. In Section 5.1, we give a detailed account of the various ultrafilter existence axioms and how they compare in strength with each other and with the axiom of choice. In Section 5.2, we show that there cannot exist a definable, in the sense of descriptive set theory, nonprincipal ultrafilter on \mathbb{N} , while in Section 5.3 we consider the connection between the existence of nonprincipal ultrafilters on \mathbb{N} and various forms of the axiom of determinacy. In Section 5.4, we consider special kinds of nonprincipal ultrafilters called selective ultrafilters and P-points, which are ultrafilters whose existence is independent of ZFC.

5.1. The ultrafilter theorem and the axiom of choice: Part I

In Section 1.1, we proved the ultrafilter theorem (UT), namely that every filter on every set can be extended to an ultrafilter on that set. The proof used Zorn's lemma, one of the avatars of the axiom of choice (AC). In this section, we describe, in more detail, the connection between UT and AC. In particular, we will describe several variations of the UT and discuss their relative strengths. Since all of the proofs of the results discussed involve set theory way out of the scope of this book, we simply point the interested reader to references. Our discussion here will involve the notion of consistency strengths of axioms, and we refer the reader to Appendix B for more information.

We first recall from Appendix B that AC is independent of the axioms ZF of set theory, that is, ZF cannot prove nor disprove AC. In particular, assuming the consistency of ZF, there is a model of ZF where AC is true

(Gödel's constructible universe L) and there is a model of ZF where AC is false (the so-called basic Cohen model).

We consider the following statements:

- (1) $\text{WUT}(X)$ is the statement: there is a nonprincipal ultrafilter on X . We refer to this statement as the **weak ultrafilter theorem for X** .
- (2) WUT is the statement: there is an infinite X such that $\text{WUT}(X)$ holds. We refer to this statement as the **weak ultrafilter theorem**.
- (3) IUT is the statement: for every infinite X , $\text{WUT}(X)$ holds. We refer to this statement as the **intermediate ultrafilter theorem**.
- (4) $\text{UT}(X)$ is the statement: for every filter \mathcal{F} on X , there is an ultrafilter \mathcal{U} on X extending \mathcal{F} .
- (5) UT is the statement: for every infinite set X , $\text{UT}(X)$ holds.

Exercise 5.1.1. In ZF, show that $\text{UT}(X)$ implies $\text{WUT}(X)$. Consequently, in ZF, UT implies IUT , which in turn implies WUT .

We thus see that WUT is the weakest possible ultrafilter existence axiom one might hope to consider. We already have models of ZF where this axiom is not true, as proven by Blass [13].

Theorem 5.1.2. *There is a model of ZF where the WUT fails.*

Of course, $\text{WUT}(\mathbb{N})$ implies WUT . However, the converse is false:

Theorem 5.1.3. *There is a model of ZF where the WUT is true but $\text{WUT}(\mathbb{N})$ fails, whence WUT is true but IUT is false.*

Proof. For example, see the model constructed in [90, Chapter 5, Problem 24]. □

At this point, we can continue in two different directions.

Theorem 5.1.4. *There is a model of ZF where $\text{WUT}(\mathbb{N})$ is true but IUT is false.*

Proof. This is [84, FM model N51]. □

The following is [90, Chapter 8, Problem 5].

Theorem 5.1.5. *There is a model of ZF such that IUT is true and yet UT fails.*

Going back to $\text{WUT}(\mathbb{N})$, we note the following result proven in [77]:

Theorem 5.1.6. *There is a model of ZF where $\text{WUT}(\mathbb{N})$ is true and yet $\text{UT}(\mathbb{N})$ fails.*

Surprisingly, it is not known if there is any implication or lack thereof between the statements WUT and $\text{UT}(\mathbb{N})$.

Finally, we have:

Theorem 5.1.7. *There is a model of ZF where UT is true and yet AC fails.*

Proof. The basic Cohen model is in fact a model of UT . □

Thus, we really have seen that most of the various existence axioms for nonprincipal ultrafilters are independent of one another.

We now turn to the notion of idempotent ultrafilters first defined in Section 4.2. The proof we gave that idempotent ultrafilters on \mathbb{N} exist used Zorn's lemma in a seemingly essential way. However, with a more careful analysis, the following was proven by Di Nasso and Tachtsis [43]:

Theorem 5.1.8. *ZF plus $\text{UT}(\mathbb{R})$ proves that there exist idempotent ultrafilters on \mathbb{N} .*

For this to really be an improvement of the original proof, one needs to know the following (see [79]):

Theorem 5.1.9. *There is a model of ZF where $\text{UT}(\mathbb{R})$ is true and yet UT fails (whence AC fails).*

It is seemingly strange to need to use $\text{UT}(\mathbb{R})$ to prove a result about certain kinds of ultrafilters on \mathbb{N} . However, what was really used in the Di Nasso and Tachtsis proof was the existence of a choice function which, to each filter \mathcal{F} on \mathbb{N} , associated an ultrafilter \mathcal{U} on \mathbb{N} extending \mathcal{F} . It is shown in [43, Proposition 3.4] that $\text{UT}(\mathbb{R})$ implies the existence of such a choice function.

The exact strength of the existence of idempotent ultrafilters on \mathbb{N} is unknown. In particular, it is not known if one can replace the assumption of $\text{UT}(\mathbb{R})$ with the weaker hypothesis of $\text{UT}(\mathbb{N})$, which is indeed a weaker hypothesis (see, for example, [104]):

Theorem 5.1.10. *There is a model of ZF where $\text{UT}(\mathbb{N})$ is true and yet $\text{UT}(\mathbb{R})$ fails.*

It is well known (see, for example, [90, Chapter 2, Problem 8]) that in ZF, AC is equivalent to Tychonoff's theorem (Theorem 3.1.13). We end this section by mentioning the following interesting fact:

Theorem 5.1.11. *In ZF, UT is equivalent to Tychonoff's theorem for compact Hausdorff spaces.*

Proof. The proof we gave of Theorem 3.1.13 used AC in two places: (1) to take an ultrafilter extending a given filter, and (2) to choose, for each sequence in a particular family, an ultralimit of that sequence. However, the Hausdorffness assumption implies that ultralimits are unique, whence the second use of AC is not needed and the proof goes through under the weaker assumption of UT.

We leave the converse as an exercise for the reader. □

Exercise 5.1.12. Finish the proof of the previous theorem by showing that, in ZF, Tychonoff’s theorem for compact Hausdorff spaces implies UT. (*Hint.* Filters on a set X are elements of 2^{2^X} .)

We remark that some mathematicians are often hesitant about using ultrafilters as they are “nonconstructive” in the sense that one cannot prove (without some form of choice) that they exist and that they cannot name a “concrete” nonprincipal ultrafilter. However, most mathematicians do not even hesitate to use AC; since UT is strictly weaker than AC, proofs that use UT should be viewed with *less* suspicion than those that use the full strength of AC. Moreover, the fact that UT is equivalent to the statement of Tychonoff’s theorem for compact Hausdorff spaces further solidifies (in this author’s biased opinion) that arguments using the existence of nonprincipal ultrafilters should not be viewed with any sort of prejudice.

5.2. Can there exist a “definable” ultrafilter on \mathbb{N} ?

Roughly speaking, *descriptive set theory* is the study of “definable” subsets of \mathbb{R} (and, more generally, *Polish spaces*, as defined below.) The word definable in the previous sentence is somewhat imprecise and we will consider some specific formalizations of it throughout this section. The motivation, however, comes from the fact that the axiom of choice can be used to construct “pathological” subsets of \mathbb{R} , such as, for example, sets that are not Lebesgue measurable. The hope is that if one restricts one’s attention to “nice” sets, then the pathologies that arise from the axiom of choice should disappear; e.g., all “nice” sets should be Lebesgue measurable.

Since we saw in the last section that the existence of nonprincipal ultrafilters on \mathbb{N} is intimately tied up with the axiom of choice (although not quite equivalent to it), it is natural to expect that there cannot be any nonprincipal ultrafilters that are definable when viewed as subsets of $2^{\mathbb{N}}$ (which is indeed a Polish space). We will see in this section that this intuition is indeed correct for the most part, although that the situation becomes a little murkier for certain notions of definability (and in certain models of set theory).

We first define the spaces that are the subject of descriptive set theory.

Definition 5.2.1. A **Polish space** is a separable topological space that is moreover completely metrizable, meaning that there is some complete metric on the space that induces the topology.

Examples 5.2.2. The following spaces are Polish spaces:

- (1) \mathbb{R}^n for any $n \geq 1$.
- (2) **Baire space** $\mathbb{N}^{\mathbb{N}}$.
- (3) **Cantor space** $2^{\mathbb{N}}$. This is a compact subspace of Baire space and is homeomorphic to the usual Cantor set in $[0, 1]$.

Exercise 5.2.3. Prove that the above examples are indeed Polish spaces.

While there are many (many!) interesting examples of Polish spaces, we will restrict our attention to those given above. We next define the most basic class of definable sets:

Definition 5.2.4. Given a topological space X , the set of all **Borel subsets of X** is the σ -algebra generated by the open subsets of X , that is, the smallest collection of subsets of X containing the open sets and closed under complementation and countable unions.

We will soon see that no nonprincipal ultrafilter on \mathbb{N} can be a Borel subset of $2^{\mathbb{N}}$. First, we want to extend our definition of definable set to allow for sets that are “almost Borel”. There are actually two ways of accomplishing this, depending on whether we are speaking about *measure* or *category*.

We first treat the case of measure. Although this can be made much more general, we stick with the concrete case of Lebesgue measure λ on \mathbb{R} . As noted above, Cantor space $2^{\mathbb{N}}$ is homeomorphic to the usual two-thirds Cantor set in $[0, 1]$. Moreover, one can show that Lebesgue measure on $[0, 1]$ restricts to a measure on the two-thirds Cantor set which in fact agrees with the usual product measure on $2^{\mathbb{N}}$ (after identifying the two spaces).

Definition 5.2.5.

- (1) $N \subseteq \mathbb{R}$ is a **null set** if there is a Borel set $B \subseteq \mathbb{R}$ with $\lambda(B) = 0$ such that $N \subseteq B$.
- (2) The set of **Lebesgue measurable** subsets of \mathbb{R} is the σ -algebra generated by the Borel sets and the null sets.

Exercise 5.2.6. $A \subseteq \mathbb{R}$ is Lebesgue measurable if and only if there is a Borel set B such that $A \triangle B$ is a null set.

We now treat the case of category.

Definition 5.2.7. Let X be a topological space.

- (1) We say that $A \subseteq X$ is **nowhere dense** if the interior of the closure of A is empty: $\text{int}(\overline{A}) = \emptyset$.
- (2) We say that $A \subseteq X$ is **meager** if it is a countable union of nowhere dense sets. The complement of a meager set is called **comeager**.
- (3) The set of **Baire measurable** subsets of X is the σ -algebra generated by the Borel sets and the meager sets.

Exercise 5.2.8. $A \subseteq X$ is Baire measurable if and only if there is a Borel set B such that $A \Delta B$ is a meager set.

The collections of null sets and meager sets both encompass notions of “smallness,” whence we can view a Lebesgue measurable set or a Baire measurable set as a set that is “almost” Borel in that it differs from a Borel set from a very small set. The precise definition of small is captured by the following notion dual to that of a filter:

Definition 5.2.9. If X is a set, then $\mathcal{I} \subseteq \mathcal{P}(X)$ is called an **ideal** on X if $\{X \setminus A : A \in \mathcal{I}\}$ is a filter on X .

Note that ideals are closed under finite unions. If we demand closure under countable unions, we arrive at:

Definition 5.2.10. An ideal is called a σ -**ideal** if it is also closed under countable unions.

Exercise 5.2.11. Prove that the collections of null sets and meager sets are (possibly improper) σ -ideals.

We will need the following fact, which is a special case of the Baire category theorem:

Fact 5.2.12. If X is a Polish space, then X is not a meager subset of itself. In other words, no subset of X can be both meager and comeager.

We now move toward the proof that no nonprincipal ultrafilter on \mathbb{N} can be almost Borel. We first need a few more facts.

Definition 5.2.13. $A \subseteq 2^{\mathbb{N}}$ is a **tail set** if, whenever $(x_n)_{n \in \mathbb{N}} \in A$ and $(y_n)_{n \in \mathbb{N}} \in 2^{\mathbb{N}}$ are such that $x_n = y_n$ eventually, then $(y_n)_{n \in \mathbb{N}} \in A$.

Exercise 5.2.14. If \mathcal{U} is a nonprincipal ultrafilter on \mathbb{N} , then \mathcal{U} (viewed as a subset of $2^{\mathbb{N}}$) is a tail set.

Finally, we need the following important facts (see [96]):

Fact 5.2.15 (0-1 laws). Suppose that $A \subseteq 2^{\mathbb{N}}$ is a tail set.

- (1) If A is Lebesgue measurable, then $\lambda(A) = 0$ or $\lambda(A) = 1$.
- (2) If A is Baire measurable, then A is either meager or comeager.

We can now prove our first nondefinability result for nonprincipal ultrafilter on \mathbb{N} :

Theorem 5.2.16. *If \mathcal{U} is a nonprincipal ultrafilter on \mathbb{N} , then \mathcal{U} , viewed as a subset of $2^{\mathbb{N}}$, is neither Lebesgue measurable nor Baire measurable.*

Proof. We prove only the assertion for Baire measurability, the assertion for Lebesgue measurability being nearly identical. Suppose, toward a contradiction, that \mathcal{U} is Baire measurable. By Exercise 5.2.14 and Fact 5.2.15, we have that \mathcal{U} is either meager or comeager. Let $f : 2^{\mathbb{N}} \rightarrow 2^{\mathbb{N}}$ be the map defined by $f(x_n) := 1 - x_n$. In other words, f flips the digits of all the coordinates. Since \mathcal{U} is an ultrafilter, we have that $f(\mathcal{U}) = 2^{\mathbb{N}} \setminus \mathcal{U}$. As f is clearly a homeomorphism of $2^{\mathbb{N}}$, we have that \mathcal{U} is both meager and comeager, contradicting Fact 5.2.12. \square

There is a wider class of definable sets beyond the class of Borel sets, namely the class of **projective sets**.

Definition 5.2.17. We define, by recursion on n , the classes of Σ_n^1 and Π_n^1 subsets of Polish spaces as follows. Throughout the definition, X is a Polish space.

- (1) $A \subseteq X$ is Σ_1^1 if there is a Polish space Y and a Borel subset B of $X \times Y$ such that $A = \{x \in X : \text{there is } y \in Y \text{ such that } (x, y) \in B\}$.
- (2) $A \subseteq X$ is Π_n^1 if $X \setminus A$ is Σ_n^1 .
- (3) $A \subseteq X$ is Σ_{n+1}^1 if there is a Polish space Y and a Π_n^1 subset B of $X \times Y$ such that $A = \{x \in X : \text{there is } y \in Y \text{ such that } (x, y) \in B\}$.

For a subset A of a Polish space X , we say that A is **projective** if it is Σ_n^1 or Π_n^1 for some n . We also define the class Δ_n^1 to consist of those sets which belong to both Σ_n^1 and Π_n^1 .

In some sense, the class of projective sets is what one gets from the class of Borel sets if one is allowed to quantify over Polish spaces. In order to get a picture for this class of sets, we mention the following:

Fact 5.2.18.

- (1) The class Δ_1^1 coincides with the class of Borel sets.
- (2) For every n , we have $\Sigma_n^1 \cup \Pi_n^1 \subseteq \Delta_{n+1}^1$.

- (3) If X is a perfect Polish space, then for every n , we can find a subset of X that is Σ_n^1 but not Π_n^1 .

The projective sets provide a natural class of definable sets properly extending the class of Borel sets. It is thus natural to wonder if a nonprincipal ultrafilter on \mathbb{N} could be a projective subset of $2^{\mathbb{N}}$.

The following is one of the central early results in descriptive set theory (see [96]):

Fact 5.2.19. If X is a Polish space, then every Σ_1^1 subset and every Π_1^1 subset of X is Baire measurable.

Using the previous fact and Theorem 5.2.16, we have:

Corollary 5.2.20. *If \mathcal{U} is a nonprincipal ultrafilter on \mathbb{N} , then \mathcal{U} is neither a Σ_1^1 nor a Π_1^1 subset of $2^{\mathbb{N}}$.*

Can we say more? Can we go higher in the projective hierarchy? Unfortunately not in ZFC:

Fact 5.2.21. The statement “there is a nonprincipal ultrafilter \mathcal{U} that is a Σ_1^2 subset of $2^{\mathbb{N}}$ ” is independent of ZFC.

While the proof of the above fact is beyond the scope of this book, let us at least roughly indicate why it is true. First, in L (see Appendix B), there is a Σ_2^1 well-ordering of $\mathcal{P}(\mathbb{N})$; this ordering can be used to construct a nonprincipal ultrafilter \mathcal{U} on \mathbb{N} that is Σ_2^1 . On the other hand, in [129], Martin and Solovay proved that the statement “every Σ_2^1 subset of $2^{\mathbb{N}}$ is Baire measurable” is consistent with ZFC. It follows from this statement and Theorem 5.2.16 that ZFC cannot prove that there is a nonprincipal ultrafilter that is a Σ_2^1 subset of $2^{\mathbb{N}}$.

5.3. The ultrafilter game

Fix a set X and a subset $D \subseteq X^{\mathbb{N}}$. We consider a two-player game $\mathcal{G}(D)$ defined as follows. Players I and II take turns playing elements of X : Player I plays $a_0 \in X$, then player II responds with $a_1 \in X$, then player I responds with $a_2 \in X$, etc... They play this game for countably many rounds, producing a **play of the game** $a = (a_0, a_1, a_2, \dots) \in X^{\mathbb{N}}$. We say that player I wins this play of the game if the play of the game a belongs to D ; otherwise, player II wins.

A **strategy** for player I is, informally speaking, a rule that tells player I what play to make at a given stage given what moves have been made thus far in the game. A strategy for player I is called **winning** if, regardless of how player II plays, player I follows the strategy, then they are guaranteed to

win. The notions of strategy and winning strategy for player II are defined in the analogous way.

We say that $D \subseteq X^{\mathbb{N}}$ is **determined** if one of the two players has a winning strategy for the game $\mathcal{G}(D)$. (Note that not both players can have a winning strategy.)

Using the axiom of choice, one can show that not all sets are determined:

Exercise 5.3.1. Show that there is $D \subseteq \mathbb{N}^{\mathbb{N}}$ that is not determined. (*Hint.* “Diagonalize” over all possible strategies.)

Once again, we continue our theme that pathologies occurring due to the axiom of choice continue to occur if instead we merely consider the existence of a nonprincipal ultrafilter. First, given an increasing sequence $a = (a_0, a_1, \dots) \in \mathbb{N}^{\mathbb{N}}$, we define a set $A_a \subseteq \mathbb{N}$ by

$$A_a := \bigcup_{n=0}^{\infty} [a_{2n-1}, a_{2n}) = [0, a_0) \cup [a_1, a_2) \cup [a_3, a_4) \cup \dots,$$

where we set $a_{-1} := 0$ for convenience. Given an ultrafilter \mathcal{U} on \mathbb{N} , we set $D_{\mathcal{U}} \subseteq \mathbb{N}^{\mathbb{N}}$ to consist of all sequences $a \in \mathbb{N}^{\mathbb{N}}$ such that either:

- a is not strictly increasing and the minimal n such that $a_n \leq a_{n-1}$ is odd, or
- a is strictly increasing and A_a belongs to \mathcal{U} .

Theorem 5.3.2. *If \mathcal{U} is a nonprincipal ultrafilter on \mathbb{N} , then $D_{\mathcal{U}}$ is not determined.*

Proof. We will only show that player II cannot have a winning strategy in $\mathcal{G}(D_{\mathcal{U}})$ and leave it as an exercise to show that player I also cannot have a winning strategy.

Suppose, toward a contradiction, that player II has a winning strategy. Note, in particular, that this strategy forces player II to always play a natural number strictly larger than player I’s previous move. The argument now proceeds by so-called “strategy stealing”. We consider two runs of the game being played simultaneously as follows: first, player I plays $a_0 \in \mathbb{N}$ and then II responds with $a_1 \in \mathbb{N}$ according to the winning strategy. Now the players start a second run of the game and player I’s first move a'_0 is actually player II’s first move from the first game, that is, $a'_0 = a_1$. Now player II responds with $a'_1 \in \mathbb{N}$ according to their winning strategy. The players now return to the first game and player I’s next move in this game is player II’s first move from the second game, that is, $a_2 := a'_1$, with which player II responds with $a_3 \in \mathbb{N}$ according to their winning strategy.

The players continue playing both games in this manner, each time player I playing player II’s previous move from the other game and then

player II always responding according to their winning strategy. By assumption, player II wins both runs of the game, that is, $\mathbb{N} \setminus A_a$ and $\mathbb{N} \setminus A_{a'}$ both belong to \mathcal{U} .

The two plays of the game (a_n) and (a'_n) satisfy the equations $a_{2n} = a'_{2n-1}$ and $a'_{2n} = a_{2n+1}$. In particular, for any n , we have $a'_n = a_{n+1}$. It follows that

$$\begin{aligned} A_a &= \bigcup_{n=0}^{\infty} [a'_{2n-1}, a'_{2n}] = [0, a_0] \cup \bigcup_{n=1}^{\infty} [a_{2n-1}, a_{2n}] \\ &= [0, a_0] \cup \bigcup_{n=0}^{\infty} [a'_{2n}, a'_{2n+1}] = [0, a_0] \cup (\mathbb{N} \setminus A_{a'}). \end{aligned}$$

Since $[0, a_0]$ is finite and \mathcal{U} is nonprincipal, we see that $A_{a'} \in \mathcal{U}$ if and only if $\mathbb{N} \setminus A_a \in \mathcal{U}$, contradicting the fact that player II won both games. \square

Exercise 5.3.3. Complete the proof of Theorem 5.3.2 by showing that player I cannot have a winning strategy in $\mathcal{G}(D_{\mathcal{U}})$ when \mathcal{U} is nonprincipal.

Definition 5.3.4. The **axiom of determinacy (AD)** is the axiom that states that every subset of $\mathbb{N}^{\mathbb{N}}$ is determined.

As we have just seen, AD is incompatible with both AC and the existence of nonprincipal ultrafilters on \mathbb{N} . At first glance, it might seem strange to consider an axiom that lies in such drastic contradistinction with AC. However, by considering axioms of *definable determinacy*, that is, versions of the axiom of determinacy that only ask for definable sets to be determined, a lot of insight into definable subsets of the reals can be drawn.

By a *pointclass*, we mean some class of subsets of Polish spaces of a certain kind, e.g., the pointclass of Borel sets, the pointclass of Σ^1_1 sets, etc....

Definition 5.3.5. If Γ is a pointclass, then the **axiom of determinacy for Γ** , denoted $\text{AD}(\Gamma)$, states that every subset of $\mathbb{N}^{\mathbb{N}}$ that belongs to Γ is determined.

We say that a pointclass Γ is *preserved under continuous substitution* if, whenever $f : X \rightarrow Y$ is a continuous function between Polish spaces and $A \subseteq Y$ belongs to Γ , then so does $f^{-1}(A) \subseteq X$.

Exercise 5.3.6. Show that the class of Borel sets is closed under continuous substitution as are the pointclasses Σ^1_n and Π^1_n for each $n \geq 1$.

Exercise 5.3.7. Suppose that Γ is a pointclass closed under continuous substitution. Show that if the nonprincipal ultrafilter \mathcal{U} , viewed as a subset of $2^{\mathbb{N}}$, belongs to Γ , then $D_{\mathcal{U}}$, as a subset of $\mathbb{N}^{\mathbb{N}}$, also belongs to Γ .

An immediate corollary of the previous exercise and Theorem 5.3.2 is the following:

Corollary 5.3.8. *Suppose that Γ is a pointclass closed under continuous substitution. Then $ZFC+AD(\Gamma)$ implies that there is no nonprincipal ultrafilter on \mathbb{N} that belongs to Γ .*

The previous corollary has earlier predecessors:

Fact 5.3.9. Assume that Γ is an *adequate* pointclass and that $AD(\Gamma)$ holds. Then:

- (1) (Banach, Mazur, Oxtoby) Every subset of \mathbb{R} in Γ is Baire measurable.
- (2) (Mycielski, Swierczkowski) Every subset of \mathbb{R} in Γ is Lebesgue measurable.

In fact, these results remain true even for sets in the pointclass $\exists^{\mathbb{R}}\Gamma$, which is the class of sets $A \subseteq X$, X a Polish space, for which there exists $Y \subseteq X \times \mathbb{R}$ such that Y belongs to Γ and A is the projection of Y onto X . (See [96] for details.)

It is of course natural to wonder for which pointclasses Γ is $AD(\Gamma)$ a sensible axiom. It turns out, for the most basic class of definable sets, namely the Borel sets, definable determinacy is not an axiom at all, but rather a *theorem* of ZFC:

Fact 5.3.10 (Martin). $AD(\Delta_1^1)$ is a theorem of ZFC. (See [96, Theorem 20.5] for a proof.)

Combining the previous theorem with Fact 5.3.9 implies that Σ_1^1 subsets of \mathbb{R} are Baire and Lebesgue measurable, a fact we referred to in the previous section.

Martin's theorem cannot be extended to the next level of the projective hierarchy. Indeed, again using Fact 5.3.9, if $AD(\Pi_1^1)$ were an axiom of ZFC, then every Σ_1^2 subset of \mathbb{R} would be both Baire and Lebesgue measurable; however, we mentioned in the previous section that this latter fact is independent of ZFC. That being said, if one assumes the existence of a *measurable cardinal* (a certain kind of large cardinal that we discuss in greater detail in Chapter 17), then one can in fact prove $AD(\Pi_1^1)$. In fact, from even stronger large cardinal assumptions, one can prove that, in a certain model called $L(\mathbb{R})$, *all* sets are determined (whence there are no nonprincipal ultrafilters at all!).

5.4. Selective ultrafilters and P-points

In the previous sections, we concerned ourselves with the connection between the axiom of choice and the existence of nonprincipal ultrafilters on \mathbb{N} . In this section, we will consider the existence of nonprincipal ultrafilters on \mathbb{N} satisfying some extra natural properties. Surprisingly, the existence of such ultrafilters will be independent of ZFC.

We begin by introducing three kinds of ultrafilters that appear frequently throughout the literature.

Definition 5.4.1. Let \mathcal{U} be a nonprincipal ultrafilter on \mathbb{N} . We say that \mathcal{U} is **selective** if, whenever $\mathbb{N} = \bigsqcup_{i=0}^{\infty} A_i$ is a partition with each $A_i \notin \mathcal{U}$, then there is $B \in \mathcal{U}$ such that $|B \cap A_i| \leq 1$ for each $i \in \mathbb{N}$.

Note that, if each A_i is nonempty in the above definition, then by adding one element of A_i to B in case $B \cap A_i = \emptyset$, we can ensure that $|B \cap A_i| = 1$ for each i ; in this way, B *selects* one element of each A_i , whence the name. Here is a useful reformulation of the notion of selective ultrafilter:

Exercise 5.4.2. Suppose that \mathcal{U} is a nonprincipal ultrafilter on \mathbb{N} . Then \mathcal{U} is selective if and only if, whenever $f : \mathbb{N} \rightarrow \mathbb{N}$ is a function, then f is either constant on a set in \mathcal{U} or injective on a set in \mathcal{U} .

Our next kind of ultrafilter asks us to be able to witness the truth of Ramsey's theorem (see Section 4.1) using an ultrafilter:

Definition 5.4.3. Let \mathcal{U} be a nonprincipal ultrafilter on \mathbb{N} . We say that \mathcal{U} is **Ramsey** if, for each $k \in \mathbb{N}$ and each 2-coloring of $\mathbb{N}^{[k]}$, there is $X \in \mathcal{U}$ such that X is homogeneous for the coloring.

We have encountered our final kind of ultrafilter back in Section 1.6:

Definition 5.4.4. Let \mathcal{U} be a nonprincipal ultrafilter on \mathbb{N} . We say that \mathcal{U} is **minimal** if there is no nonprincipal ultrafilter \mathcal{V} on \mathbb{N} such that $\mathcal{V} <_{RK} \mathcal{U}$.

It turns out that the three kinds of ultrafilters above are actually the same. We prove this equivalence in the next theorem, along with two other equivalent formulations:

Theorem 5.4.5. *For \mathcal{U} a nonprincipal ultrafilter on \mathbb{N} , the following are equivalent:*

- (1) \mathcal{U} is Ramsey.
- (2) If R is a binary relation on \mathbb{N} satisfying, for each $m \in \mathbb{N}$, the property that $(\mathcal{U}n)R(m, n)$, then there is $A \in \mathcal{U}$ such that, enumerating $A = (a_n)_{n \in \mathbb{N}}$ in increasing order, we have $R(a_n, a_{n+1})$ for each n .

- (3) \mathcal{U} is selective.
- (4) \mathcal{U} is minimal.
- (5) \mathcal{U} is **quasi-normal**: for every family $(A_n)_{n \in \mathbb{N}}$ of members of \mathcal{U} , there is a set $A \in \mathcal{U}$ such that, for all $m, n \in A$, if $m > n$, then $m \in A_n$.

Proof. (1) implies (2): Fix R as in (2) and define a coloring $c : \mathbb{N}^{[2]} \rightarrow \{0, 1\}$ such that $c(\{m, n\}) = 1$ if and only if $R(m, n)$ (and $m < n$). Since \mathcal{U} is Ramsey, there is $A \in \mathcal{U}$ homogeneous for c ; by assumption, A must be homogeneous of color 1, whence this A is as desired.

(2) implies (3): Take a partition $\mathbb{N} = \bigsqcup_{i=0}^{\infty} A_i$ with each $A_i \notin \mathcal{U}$ and define a binary relation R on \mathbb{N} by $R(m, n)$ if and only if there is $k < l$ with $m \in A_k$ and $n \in A_l$. Setting $B_l := \bigcap_{k \leq l} A_k^c$, we see that $B_l \in \mathcal{U}$ and every element of A_l is R -related to every element of B_l (as the A_i 's form a partition). Thus, by (2), there is $B \in \mathcal{U}$ such that, when enumerated in increasing order as $(b_n)_{n \in \mathbb{N}}$, we have $R(b_n, b_{n+1})$ for every $i \in \mathbb{N}$. It is clear that B is as desired in the conclusion of selective ultrafilter.

(3) implies (4): Suppose that \mathcal{U} is selective and let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a function that is not constant on a \mathcal{U} -large set (so $f(\mathcal{U})$ is not principal). We must show that $f(\mathcal{U}) \equiv_{RK} \mathcal{U}$. However, Exercise 5.4.2 implies that there is $B \in \mathcal{U}$ such that $f \upharpoonright B$ is injective, whence $f(\mathcal{U}) \equiv_{RK} \mathcal{U}$ by Corollary 1.3.16.

(4) implies (5): Suppose that \mathcal{U} is minimal and suppose that $(A_n)_{n \in \mathbb{N}}$ is a family of elements of \mathcal{U} . We are looking for $A \in \mathcal{U}$ such that, for all $x, y \in A$, if $x < y$, then $y \in A_x$. Without loss of generality, we may assume that $\bigcap_{x \in \mathbb{N}} A_x = \emptyset$. Thus, we may define $f : \mathbb{N} \rightarrow \mathbb{N}$ by $f(y)$ equals the least x such that $y \notin A_x$. Since each $A_x \in \mathcal{U}$, it follows that f cannot be constant on any set in \mathcal{U} , whence, by minimality, there is $B \in \mathcal{U}$ such that f is injective on B . Let $g : \mathbb{N} \rightarrow \mathbb{N}$ be defined by $g(x) = \max(\max\{y \in B : f(y) \leq x\}, x+1)$, which is a legitimate definition since f is injective on B . By construction, g is increasing, $g(x) > x$, and, for $y \in B$, if $y > g(x)$, then $f(y) > x$, whence $y \in A_x$. Recursively define a sequence $(\alpha_n)_{n \in \mathbb{N}}$ by setting $\alpha_0 = 0$ and $\alpha_{n+1} = g(\alpha_n)$. Define $h : \mathbb{N} \rightarrow \mathbb{N}$ by $h(y)$ equals the least n such that $y \leq \alpha_n$. Note that h is not constant on any set in \mathcal{U} (as h constant on a set implies that the set is bounded), whence, by minimality, h is injective on some $C \in \mathcal{U}$. To create some space, we take $A \in \mathcal{U}$ such that $A \subseteq B \cap C$ and $h(A)$ contains no two consecutive integers. We claim that this A is as desired. Suppose that $x, y \in A$ and $x < y$. Since h injective on C , we have that $h(x) < h(y)$, whence $h(x) + 1 < h(y)$ (since $x, y \in A$). Then

$$x \leq \alpha_{h(x)} \Rightarrow g(x) \leq g(\alpha_{h(x)}) = \alpha_{h(x)+1}.$$

Since $h(x) + 1 < h(y)$, we have that $\alpha_{h(x)+1} < y$, whence $g(x) < y$; since $y \in B$, it follows from our earlier observation that $y \in A_x$.

(5) implies (1): We proceed by induction on k , the case $k = 1$ being immediate from the definition of ultrafilter. So suppose that $\mathbb{N}^{[k+1]}$ is partitioned into P_1 and P_2 . For each $x \in \mathbb{N}$ and $(y_1, \dots, y_k) \in \mathbb{N}^{[k]}$, say that $P_1(x)(y_1, \dots, y_k)$ holds precisely when $x < y_1$ and $P_1(x, y_1, \dots, y_k)$; otherwise, $P_2(x)(y_1, \dots, y_k)$ holds. By induction, for each $x \in \mathbb{N}$, there is $A_x \in \mathcal{U}$ and $i(x) \in \{1, 2\}$ such that, for all $(y_1, \dots, y_k) \in A_x^{[k]}$, we have that $(y_1, \dots, y_k) \in P_{i(x)}(x)$. Without loss of generality, we may assume that $y > x$ for all $y \in A_x$; making this assumption has the nice benefit of yielding that $(y_1, \dots, y_k) \in [A_x]^k$ implies that $(x, y_1, \dots, y_k) \in P_{i(x)}$. Let $A \in \mathcal{U}$ witness quasi-normality of \mathcal{U} for the family $(A_x)_{x \in \mathbb{N}}$ and let $i \in \{1, 2\}$ and $B \in \mathcal{U}$ be such that $i(x)$ is constantly i on B . It follows that, for $(x, y_1, \dots, y_k) \in (A \cap B)^{[k+1]}$, we have that $(y_1, \dots, y_k) \in A_x^{[k]}$ (by quasi-normality), whence $(x, y_1, \dots, y_k) \in P_{i(x)} = P_i$ (since $x \in B$), as desired. \square

From now on, we refer to the ultrafilters in the above theorem as selective (for the sake of definiteness). We now turn to the question of the existence of selective ultrafilters. We first need a bit of terminology and an exercise.

Given $A, B \subseteq \mathbb{N}$, we write $A \subseteq^* B$ if $A \setminus B$ is finite. (Thus, A is “almost contained” in B .)

Exercise 5.4.6. Prove that \subseteq^* is a transitive relation on subsets of \mathbb{N} : if $C \subseteq^* B \subseteq^* A$, then $C \subseteq^* A$.

Given a family $(B_i)_{i \in I}$ of subsets of \mathbb{N} , a **pseudo-intersection** of the family is a set A such that $A \subseteq^* B_i$ for all $i \in I$.

Exercise 5.4.7. Suppose that α is a countable ordinal and $(B_\beta)_{\beta < \alpha}$ is a family of infinite subsets of \mathbb{N} such that $B_\beta \subseteq^* B_\gamma$ for all $\gamma < \beta < \alpha$. Prove that the family $(B_\beta)_{\beta < \alpha}$ has an infinite pseudo-intersection.

Theorem 5.4.8. *If CH holds, then there exists a selective ultrafilter.*

Proof. By CH, we may enumerate all countable partitions of \mathbb{N} (as in the definition of selective ultrafilter) by $(\mathcal{A}_\alpha)_{\alpha < \omega_1}$. We now construct a sequence $(X_\alpha)_{\alpha < \omega_1}$ of infinite subsets of \mathbb{N} so that $X_\beta \subseteq^* X_\alpha$ for all $\alpha < \beta < \omega_1$ as follows. Set $X_0 := \mathbb{N}$. Suppose that X_α has been constructed. If there is $A \in \mathcal{A}_\alpha$ for which $X_\alpha \cap A$ is infinite, then set $X_{\alpha+1} := X_\alpha \cap A$ for some such A . Otherwise, take infinite $X_{\alpha+1} \subseteq X_\alpha$ such that $|X_{\alpha+1} \cap A| \leq 1$ for all $A \in \mathcal{A}_\alpha$. Note that $X_{\alpha+1} \subseteq^* X_\beta$ for all $\beta \leq \alpha$ by Exercise 5.4.6. Assume now that α is a limit ordinal and that X_β has been defined for all $\beta < \alpha$ in

such a way that $X_\gamma \subseteq^* X_\beta$ for all $\beta < \gamma < \alpha$. We now take X_α to be an infinite pseudo-intersection of $(X_\beta)_{\beta < \alpha}$; this is possible by Exercise 5.4.7.

Now that $(X_\alpha)_{\alpha < \omega_1}$ has been defined, we note that $(X_\alpha)_{\alpha < \omega_1}$ has the FIP and in fact each finite intersection has infinitely many elements, whence, by Exercise 1.1.19, there is a nonprincipal ultrafilter \mathcal{U} generated by the family $(X_\alpha)_{\alpha < \omega_1}$.

We claim that \mathcal{U} is selective. Indeed, first note that, since each X_α is infinite, we have that \mathcal{U} is nonprincipal. To see the main defining property of being selective, consider a partition $\mathbb{N} = \bigsqcup_{i=0}^{\infty} A_i$ into countably many sets A_i with $A_i \notin \mathcal{U}$ for each i . Take $\alpha < \omega_1$ such that this partition is \mathcal{A}_α . If $X_\alpha \cap A$ were infinite for some $A \in \mathcal{A}_\alpha$, then $X_{\alpha+1} = X_\alpha \cap A$ for such an A ; since $X_{\alpha+1} \in \mathcal{U}$, it would follow that $A \in \mathcal{U}$, yielding a contradiction. Thus, $X_\alpha \cap A$ is finite for each $A \in \mathcal{A}_{\alpha+1}$, whence, by construction, $|X_{\alpha+1} \cap A| \leq 1$ for all $A \in \mathcal{A}_{\alpha+1}$. Since $X_{\alpha+1} \in \mathcal{U}$, we have that \mathcal{U} is selective. \square

As a result, the statement “there is a selective ultrafilter on \mathbb{N} ” is consistent with ZFC. We will soon see that it is also consistent that selective ultrafilters do not exist. In fact, we will see that it is consistent that there are no ultrafilters satisfying the following weakening of the notion of selectivity:

Definition 5.4.9. A nonprincipal ultrafilter \mathcal{U} on \mathbb{N} is **weakly selective** if, whenever $\mathbb{N} = \bigsqcup_{i=0}^{\infty} A_i$ is a partition with each $A_i \notin \mathcal{U}$, then there is $B \in \mathcal{U}$ such that $B \cap A_i$ is finite for each $i \in \mathbb{N}$.

Clearly, a selective ultrafilter is weakly selective, whence CH implies that weakly selective ultrafilters exist.

Remark 5.4.10. In general, being weakly selective is a genuine weakening of being selective. For example, in [132], Mathias proved that, under CH, there is a weakly selective ultrafilter that is not selective. In [135], Miller proved that, in certain models of set theory, there are weakly selective ultrafilters but no selective ultrafilters. On the other hand, it can happen that, in certain models of set theory, the two notions coincide. In fact, in [161, Section XVIII.4], Shelah constructed a model where there exists exactly one weakly selective ultrafilter (up to isomorphism), which is, in fact, actually selective.

The following definition is standard in topology:

Definition 5.4.11. If X is a topological space, then $x \in X$ is a **P-point** if, for every countable family $(U_n)_{n \in \mathbb{N}}$ of neighborhoods of x , we have that $\bigcap_{n \in \mathbb{N}} U_n$ is also a neighborhood of x .

Theorem 5.4.12. For a nonprincipal ultrafilter \mathcal{U} on \mathbb{N} , the following are equivalent:

- (1) \mathcal{U} is weakly selective.
- (2) \mathcal{U} is a P-point of $\beta\mathbb{N} \setminus \mathbb{N}$.

Proof. In this proof, we slightly abuse notation and write U_A for those elements of $\beta\mathbb{N} \setminus \mathbb{N}$ that contain A . (In other words, we are identifying U_A with $U_A \cap (\beta\mathbb{N} \setminus \mathbb{N})$.)

First assume that \mathcal{U} is weakly selective. We show that \mathcal{U} is a P-point. Notice that it suffices to assume that, whenever each U_n is a basic open neighborhood of \mathcal{U} , we have that $\bigcap_{n \in \mathbb{N}} U_n$ is also a neighborhood of \mathcal{U} . Take $A_n \subseteq \mathbb{N}$ such that $U_n = U_{A_n}$. Since $\mathcal{U} \in U_n$ for each n , we have that $A_n^c \notin \mathcal{U}$. Since \mathcal{U} is weakly selective, there is $B \in \mathcal{U}$ such that $B \cap A_n^c$ is finite for all n . Set $U := U_B$. Then $\mathcal{U} \in U$. We claim that $U \subseteq U_n$ for each n . Indeed, if $\mathcal{V} \in U$, then $B \in \mathcal{V}$; since $B \cap A_n^c$ is finite and \mathcal{V} is nonprincipal, we have that $A_n^c \notin \mathcal{V}$, so $\mathcal{V} \in U_n$, as desired.

Now suppose that \mathcal{U} is a P-point and that $A_n \notin \mathcal{U}$ for each n . Let $U_n := U_{A_n^c}$ be an open neighborhood of \mathcal{U} . Take $B \in \mathcal{U}$ such that $U_B \subseteq \bigcap_n U_n$. If $B \cap A_n$ were infinite, then there would be a nonprincipal ultrafilter \mathcal{V} containing $B \cap A_n$, whence $\mathcal{V} \in U_B \setminus U_n$, a contradiction. \square

From now on, we refer to nonprincipal ultrafilters on \mathbb{N} that are P-points of $\beta\mathbb{N} \setminus \mathbb{N}$ as simply P-points.

Exercise 5.4.13. If \mathcal{U} is a P-point and \mathcal{V} is a nonprincipal ultrafilter on \mathbb{N} such that $\mathcal{V} \leq_{RK} \mathcal{U}$, then \mathcal{V} is also a P-point.

The original interest in P-point ultrafilters was that they were used to settle a question about the topological space $\beta\mathbb{N} \setminus \mathbb{N}$. Before explaining this, we need an exercise:

Exercise 5.4.14. Suppose that X is a Hausdorff topological space such that every point is a P-point. Prove that:

- (1) Every countable union of closed sets is closed.
- (2) Every countable set is discrete.

Conclude that if X is also compact, then it is finite.

The following is immediate from the previous exercise applied to $\beta\mathbb{N} \setminus \mathbb{N}$:

Corollary 5.4.15. *Not every nonprincipal ultrafilter on \mathbb{N} is a P-point.*

One can actually describe a nonprincipal ultrafilter on \mathbb{N} that is not a P-point:

Exercise 5.4.16. Let $(X_n)_{n \in \mathbb{N}}$ be a partition of \mathbb{N} into infinitely many infinite sets. Set \mathcal{F} to be those $Z \subseteq \mathbb{N}$ such that, for all but finitely many $n \in \mathbb{N}$, we have that $X_n \subseteq^* Z$. Prove that:

- (1) \mathcal{F} is a filter on \mathbb{N} containing the Fréchet filter.
- (2) No ultrafilter extending \mathcal{F} is a P-point.

Corollary 5.4.17 (Kunen [112]). *Assuming CH, $\beta\mathbb{N} \setminus \mathbb{N}$ is not **point-homogeneous**, that is, there are $\mathcal{U}, \mathcal{V} \in \beta\mathbb{N} \setminus \mathbb{N}$ for which there does not exist a self-homeomorphism σ of $\beta\mathbb{N} \setminus \mathbb{N}$ with $\sigma(\mathcal{U}) = \mathcal{V}$.*

Proof. Let $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$ be a P-point. Observe that, for any homeomorphism σ of $\beta\mathbb{N} \setminus \mathbb{N}$, $\sigma(\mathcal{U})$ is also a P-point. Consequently, by Corollary 5.4.15, there is some $\mathcal{V} \in \beta\mathbb{N} \setminus \mathbb{N}$ that is not of the form $\sigma(\mathcal{U})$ for some homeomorphism σ of $\beta\mathbb{N} \setminus \mathbb{N}$. \square

In [61], Frolik obtained the same conclusion as in the previous corollary without assuming CH.

As shown above, it is consistent with ZFC that weakly selective ultrafilters exist. On the other hand, we have the following difficult theorem of Shelah [159, VI, §4]:

Theorem 5.4.18 (Shelah). *The existence of a weakly selective ultrafilter cannot be proven in ZFC.*

We conclude this section by describing an interesting connection between weakly selective ultrafilters and the idempotent ultrafilters introduced in Section 4.2. Recall that every element of an idempotent ultrafilter on \mathbb{N} is an FS-set, meaning that it contains $\text{FS}(c)$ for some sequence $c = (c_n)_{n \in \mathbb{N}}$ of distinct elements of \mathbb{N} . (Technically speaking, we worked with ultrafilters on \mathbb{Z} , but the exact same analysis works for ultrafilters on \mathbb{N} .) Let us call an ultrafilter on \mathbb{N} for which every member is an FS-set a **weakly summable** ultrafilter. (By Corollary 4.2.13, these are precisely the ultrafilters that are in the closure of the set of idempotent ultrafilters.) It is natural to consider the following variation:

Definition 5.4.19. An ultrafilter \mathcal{U} on \mathbb{N} is called **strongly summable** if for every $A \in \mathcal{U}$, there is a sequence $c = (c_n)_{n \in \mathbb{N}}$ of distinct elements of \mathbb{N} such that $\text{FS}(c) \subseteq A$ and $\text{FS}(c) \in \mathcal{U}$.

Surprisingly, this seemingly harmless improvement leads us to a kind of ultrafilter that cannot be proven to exist in ZFC:

Theorem 5.4.20. *The existence of a strongly summable ultrafilter implies the existence of a weakly selective ultrafilter. Consequently, the existence of strongly summable ultrafilters cannot be proven in ZFC.*

Remark 5.4.21. It is, on the other hand, consistent with ZFC that strongly summable ultrafilters exist; see [83, Theorem 12.29].

We sketch a proof of Theorem 5.4.20 now. The proof goes via a different kind of ultrafilter. First, if $\mathcal{F} = (F_n)_{n \in \mathbb{N}}$ is a sequence of subsets of \mathbb{N} , we write

$$\text{FU}(\mathcal{F}) := \left\{ \bigcup_{n \in G} F_n : G \subseteq \mathbb{N} \text{ finite} \right\}.$$

Definition 5.4.22. \mathcal{U} is a **union ultrafilter** if it is a nonprincipal ultrafilter on $\mathcal{P}_f(\mathbb{N})$ and, for each $A \in \mathcal{U}$, there is a sequence $\mathcal{F} = (F_n)_{n \in \mathbb{N}}$ of pairwise disjoint elements of $\mathcal{P}_f(\mathbb{N})$ such that $\text{FU}(\mathcal{F}) \subseteq A$ and $\text{FU}(\mathcal{F}) \in \mathcal{U}$.

We will need the following fact [83, Theorems 12.31 and 12.35]:

Theorem 5.4.23. *A strongly summable ultrafilter exists if and only if a union ultrafilter exists.*

The proof of the previous theorem is not difficult and involves some arithmetic trickery. Thus, we are left to show:

Theorem 5.4.24. *If a union ultrafilter exists, then a weakly selective ultrafilter exists.*

Proof. Let \mathcal{U} be a union ultrafilter. Let $\max : \mathcal{P}_f(\mathbb{N}) \rightarrow \mathbb{N}$ be the usual maximum function and consider its unique continuous extension $\beta \max : \beta \mathcal{P}_f(\mathbb{N}) \rightarrow \beta \mathbb{N}$. We show that $\mathcal{V} := (\beta \max)(\mathcal{U})$ is a P-point.

We first observe that \mathcal{V} is nonprincipal: if $\mathcal{V} = \mathcal{U}_n$ for some $n \in \mathbb{N}$, then $\max^{-1}(\{n\}) \in \mathcal{U}$. Since $\max^{-1}(\{n\})$ is a finite set (there are only finitely many finite sets whose maximum is n), this would imply that \mathcal{U} is principal, leading to a contradiction.

In the remainder of the proof, we abuse notation by writing, for any $A \subseteq \mathbb{N}$, U_A instead of $U_A \cap (\beta \mathbb{N} \setminus \mathbb{N})$.

We now check the defining property of being a P-point. Fix a sequence $(A_n)_{n \in \mathbb{N}}$ from \mathcal{V} ; we seek $B \in \mathcal{V}$ such that $U_B \subseteq \bigcap_{n=0}^{\infty} U_{A_n}$. Without loss of generality, we may suppose that $A_0 = \mathbb{N}$, $A_{n+1} \subseteq A_n$, and $n \notin A_{n+1}$ for all n . These conditions ensure that $\bigcap_{n \in \mathbb{N}} A_n = \emptyset$, which allows us to consider $f : \mathbb{N} \rightarrow \mathbb{N}$ given by

$$f(x) := \max\{n \in \mathbb{N} : x \in A_n\}.$$

Set $\mathcal{B} := \{F \in \mathcal{P}_f(\mathbb{N}) : f(\max F) \leq \min F\}$.

Claim. $\mathcal{B} \notin \mathcal{U}$.

Proof of Claim. Suppose, toward a contradiction, that $\mathcal{B} \in \mathcal{U}$. Since \mathcal{U} is a union ultrafilter, there is $\mathcal{F} = (F_n)_{n \in \mathbb{N}}$, a collection of pairwise disjoint elements of $\mathcal{P}_f(\mathbb{N})$, such that $\text{FU}(\mathcal{F}) \subseteq \mathcal{B}$ and $\text{FU}(\mathcal{F}) \in \mathcal{U}$. Without loss of generality, we may suppose that $\min F_n < \min F_{n+1}$ for all n . Note that

$\{\max F_n : n \in \mathbb{N}\} = \{\max G : G \in \text{FU}(\mathcal{F})\} \in \mathcal{V}$ as $\text{FU}(\mathcal{F}) \in \mathcal{U}$. Now pick $k \in \mathbb{N}$ such that, for all $n \geq k$, we have $\max F_n > \max F_0$. It then follows that $\{\max F_n : n \geq k\} \in \mathcal{V}$. Set $l := \min F_0$. Then $A_{l+1} \in \mathcal{V}$, whence there is $n \geq k$ such that $\max F_n \in A_{l+1}$. But then

$$l + 1 \leq f(\max F_n) = f(\max(F_n \cup F_0)) \leq \min(F_n \cup F_0) = \min F_0 = l,$$

a contradiction.

By the Claim, we have that $\mathcal{B}^c \in \mathcal{U}$. Once again, take $\mathcal{F} = (F_n)_{n \in \mathbb{N}}$, a collection of pairwise disjoint elements of $\mathcal{P}_f(\mathbb{N})$ with $\min F_n < \min F_{n+1}$ for all n , such that $\text{FU}(\mathcal{F}) \subseteq \mathcal{B}^c$ and $\text{FU}(\mathcal{F}) \in \mathcal{U}$. Set $B := \{\max F_n : n \in \mathbb{N}\} \in \mathcal{V}$; we claim that this B is as desired. Indeed, suppose that n and k are such that $\max F_n \in B \setminus A_k$. We then have that $n \leq \min F_n < f(\max F_n)$ (this latter fact following from the fact that $F_n \in \mathcal{B}^c$). However, $f(\max F_n) < k$ since $\max F_n \notin A_k$. We thus see that if $\max F_n \in B \setminus A_k$, then $n < k$. Consequently, $|B \setminus A_k| < k$ whence $B \subseteq^* A_k$ for each k and thus $U_B \subseteq U_{A_k}$, as desired. \square

5.5. Notes and references

A great reference for the various weakenings of the axiom of choice, the connections with ultrafilters, and an introduction to forcing with a perspective on such statements is Jech's book [90]. The fact that there can be no measurable ultrafilter on \mathbb{N} is essentially due to Sierpiński [163]. A more thorough treatment of descriptive set theory can be found in Kechris's book [96]. We stumbled upon the ultrafilter game in a mathstackexchange of Noah Schweber

<https://mathoverflow.net/questions/109739/determinacy-and-definable-ultrafilters>.

A nice article about large cardinals and determinacy (and their philosophical implications) can be found at

<https://plato.stanford.edu/entries/large-cardinals-determinacy/>.

Our treatment of Ramsey ultrafilters and P-points is a mixture of Blass's thesis [16] and Booth's article [18], the latter of which contains other interesting reformulations of these notions. Our proof of Theorem 5.4.8 is taken from [89, Theorem 7.8].

Part 2

Classical ultraproducts

Classical ultraproducts

In this chapter, we introduce the fundamental construction of taking the ultraproduct of a family of structures in a given first-order language. In Section 6.1, we motivate the idea behind the construction by considering the Stone duality theorem from Section 3.4 applied to the Lindenbaum algebra of a first-order language. In Section 6.2, we undertake the construction of an ultraproduct of a family of sets and then extend the construction to a family of structures in a given language in Section 6.3. In Section 6.4 we prove the Fundamental Theorem of Ultraproducts, otherwise known as Łoś's theorem, which states that truth in an ultraproduct is given by almost-everywhere truth in the individual structures. In Section 6.5, we revisit the discussion from Section 5.1 connecting the ultrafilter theorem and the axiom of choice, now using Łoś's theorem and the compactness theorem as two new players in the story. In Section 6.6, we consider the question of when an ultrapower of a set is the same as the set itself, leading to the notion of (in)complete ultrafilters. Section 6.7 re-examines the Rudin-Keisler order introduced in Section 1.6 through the lens of embeddings of ultrapowers of structures. Section 6.8 presents some results concerning the cardinalities of ultraproducts, while Section 6.9 considers the possibility of iterating the ultrapower construction. In Section 6.10, we consider a category-theoretic take on the ultraproduct construction, allowing us to generalize beyond the case of first-order structures as well as to consider a dual version of the construction known as the ultracoproduct construction. Finally, in Section 6.11, we present the proof of the Feferman-Vaught theorem, which is a result in the spirit of Łoś's theorem connecting truth in a reduced product with truth in the constituent structures.

6.1. Motivating the definition of ultraproducts

Fix a first-order language \mathcal{L} . The collection of all \mathcal{L} -sentences very closely resembles a Boolean algebra, with the Boolean operations being interpreted by conjunction, disjunction, and negation. However, it is not entirely clear which sentences should play the roles of 0 and 1. For example, since $\forall x(x = x)$ is true in all \mathcal{L} -structures, perhaps it should play the role of 1. But then again, the sentence $\exists x(x = x)$ possesses the same feature. This leads us to consider not the set of \mathcal{L} -sentences as a Boolean algebra, but rather equivalence classes $[\sigma]$ of \mathcal{L} -sentences modulo the notion of *logical equivalence*, where \mathcal{L} -sentences σ and τ are logically equivalent if $\mathcal{M} \models (\sigma \leftrightarrow \tau)$ for all \mathcal{L} -structures \mathcal{M} . The operations of conjunction, disjunction, and negation yield well-defined operations on the equivalence classes and it is readily verified that the resulting structure is indeed a Boolean algebra, called the **Lindenbaum algebra** for \mathcal{L} , denoted $\mathbb{B}_{\mathcal{L}}$.

Exercise 6.1.1. Check all of the assertions made in the previous paragraph.

The Stone space $S(\mathbb{B}_{\mathcal{L}})$ of $\mathbb{B}_{\mathcal{L}}$ is essentially the same as the set of complete \mathcal{L} -theories. More precisely, given an ultrafilter \mathcal{U} on $\mathbb{B}_{\mathcal{L}}$, the set $T_{\mathcal{U}} := \{\sigma : [\sigma] \in \mathcal{U}\}$ is readily verified to be a complete \mathcal{L} -theory. Conversely, given any complete \mathcal{L} -theory T , the set $\{[\sigma] : \sigma \in T\}$ is an ultrafilter on $\mathbb{B}_{\mathcal{L}}$. In what follows, we will blur the distinction between these two sets and simply consider elements of $S(\mathbb{B}_{\mathcal{L}})$ as complete \mathcal{L} -theories. In this way, one obtains a compact topology on the set of complete \mathcal{L} -theories whose basic open sets are those of the form $U_{\sigma} := \{T \in S(\mathbb{B}_{\mathcal{L}}) : \sigma \in T\}$.

Since $S(\mathbb{B}_{\mathcal{L}})$ is a compact Hausdorff space, given any family $(T_i)_{i \in I}$ from $S(\mathbb{B}_{\mathcal{L}})$ and ultrafilter \mathcal{U} on I , we can consider the complete \mathcal{L} -theory $\lim_{\mathcal{U}} T_i$, which is characterized by the property that, given any \mathcal{L} -sentence σ , we have $\sigma \in \lim_{\mathcal{U}} T_i$ if and only if $\sigma \in T_i$ for \mathcal{U} -almost all i . Said differently, given models $\mathcal{M}_i \models T_i$ and an \mathcal{L} -structure \mathcal{M} , we have that $\mathcal{M} \models \lim_{\mathcal{U}} T_i$ if and only if, for any \mathcal{L} -sentence σ , if $\mathcal{M}_i \models \sigma$ for \mathcal{U} -almost all i , then $\mathcal{M} \models \sigma$. Any such \mathcal{M} can be considered a \mathcal{U} -ultralimit of the \mathcal{L} -structures \mathcal{M}_i .

We now ask the question: given a family $(\mathcal{M}_i)_{i \in I}$ of \mathcal{L} -structures and an ultrafilter \mathcal{U} on I , how can we construct *some* model $\lim_{\mathcal{U}} \mathcal{M}_i$ of $\lim_{\mathcal{U}} \text{Th}(\mathcal{M}_i)$, where $\text{Th}(\mathcal{M}_i)$ denotes the complete theory of \mathcal{M}_i as defined in Definition A.2.9? (The notation $\lim_{\mathcal{U}} \mathcal{M}_i$ is not any sort of official notation, but is merely a notation that we are using for the current discussion.) The idea in constructing $\lim_{\mathcal{U}} \mathcal{M}_i$ that eventually will work comes about by asking us to strengthen the connection between the model $\lim_{\mathcal{U}} \mathcal{M}_i$ of the limiting theory and the original family $(\mathcal{M}_i)_{i \in I}$ of models by considering formulae with parameters rather than just sentences. For example, given an \mathcal{L} -formula $\varphi(x)$ and elements $a(i) \in M_i$, if $\mathcal{M}_i \models \varphi(a(i))$ for \mathcal{U} -almost all i ,

then the desired structure $\lim_{\mathcal{U}} \mathcal{M}_i$ should behave as if this is true as well. But what does that even mean?!? Such an assertion somehow implies that the structure $\lim_{\mathcal{U}} \mathcal{M}_i$ has access to the element $a \in \prod_{i \in I} M_i$.

Based on the last sentence of the previous paragraph, one might wonder if some direct product construction might yield the desired model. However, this approach can be quickly dismissed:

Example 6.1.2. Suppose that $\mathcal{L} = \{\cdot\}$, where \cdot is a binary function symbol. Suppose that, for each $n \in \mathbb{N}$, G_n is a group and that G_n is abelian for $n > 0$ while G_0 is not abelian. Then $\prod_{n \in \mathbb{N}} G_n \models \exists x \exists y (xy \neq yx)$. However, if \mathcal{U} is any nonprincipal ultrafilter on \mathbb{N} , then $G_n \models \neg \exists x \exists y (xy \neq yx)$ for \mathcal{U} -almost all n .

The issue with the direct product construction is that it pays too much attention to what happens in particular structures. It turns out, however, that a modification of the direct product construction that only keeps track of what happens on a \mathcal{U} -large set of coordinates will indeed yield the desired result.

Returning to our earlier idea, consider the formula $\varphi(x, y)$ that is simply $x = y$ and two elements $a, b \in \prod_{i \in I} M_i$. If $\mathcal{M}_i \models a(i) = b(i)$ for \mathcal{U} -almost all i , then we want our structure $\lim_{\mathcal{U}} \mathcal{M}_i$ to also believe this is the case. In other words, the (possibly) distinct elements $a, b \in \prod_{i \in I} M_i$ should be identified in the model we wish to construct. This *quotient* of the cartesian product is known as the ultraproduct of the family $(M_i)_{i \in I}$ with respect to the ultrafilter \mathcal{U} and is the subject of the next section. In Section 6.3, we show that this ultraproduct of sets is naturally the universe of an \mathcal{L} -structure, called the ultraproduct of the family $(\mathcal{M}_i)_{i \in I}$ with respect to \mathcal{U} , which will indeed be the structure $\lim_{\mathcal{U}} \mathcal{M}_i$ we have been searching for.

6.2. Ultraproducts of sets

In this section, we carry out the first part of the plan laid out at the end of the previous section. Fix a family $(M_i)_{i \in I}$ of sets and a filter \mathcal{F} on I . We define a relation $\sim_{\mathcal{F}}$ on the cartesian product $\prod_{i \in I} M_i$ by declaring

$$a \sim_{\mathcal{F}} b \Leftrightarrow \{i \in I : a(i) = b(i)\} \in \mathcal{F}.$$

In other words, $a \sim_{\mathcal{F}} b$ holds when a and b agree on an \mathcal{F} -large set of coordinates.

Example 6.2.1. Suppose that $\mathcal{F} = \{I\}$. Then $a \sim_{\mathcal{F}} b$ if and only if $a = b$.

Example 6.2.2. Suppose that $\mathcal{F} = \mathcal{U}_j$, the principal ultrafilter generated by j . Then $a \sim_{\mathcal{F}} b$ if and only if $a(j) = b(j)$.

The next exercise is integral in what is to follow; it simply uses the axioms for being a filter.

Exercise 6.2.3. Prove that $\sim_{\mathcal{F}}$ is an equivalence relation on $\prod_{i \in I} M_i$.

We denote the $\sim_{\mathcal{F}}$ -equivalence class of f by $[a]_{\mathcal{F}}$ (or sometimes simply $[a]$ if \mathcal{F} is clear from context).

Definition 6.2.4. The **reduced product** of the family $(M_i)_{i \in I}$ with respect to the filter \mathcal{F} is the set of equivalence classes $[a]_{\mathcal{F}}$ and is denoted by $\prod_{\mathcal{F}} M_i$. When \mathcal{F} is an ultrafilter, we refer to the reduced product as an **ultraproduct**. When each $M_i = M$ for some fixed set M , we refer to the reduced product (resp., ultraproduct) as the **reduced power** (resp., **ultrapower**) of the set M with respect to \mathcal{F} , denoted $M^{\mathcal{F}}$ (resp., $M^{\mathcal{U}}$).

Example 6.2.5. By Example 6.2.1, when $\mathcal{F} = \{I\}$, we have that $[a]_{\mathcal{F}} = \{a\}$ for each $a \in \prod_{i \in I} M_i$, whence $\prod_{\mathcal{F}} M_i = \prod_{i \in I} M_i$ (after identifying $[a]_{\mathcal{F}}$ with a itself).

Example 6.2.6. When $\mathcal{F} = \mathcal{U}_j$, then the map $[a]_{\mathcal{U}_j} \mapsto a(j)$ is a bijection between $\prod_{\mathcal{U}_j} M_i$ and M_j .

Remark 6.2.7. Recall from Exercise 1.1.9 that ultrafilters on I are the same thing as $\{0, 1\}$ -valued finitely additive probability measures on I . Thus, the ultrapower $M^{\mathcal{U}}$ is the result of considering the set of functions $I \rightarrow M$ and identifying two such functions if they agree on a set of measure 1 (in the sense of $\mu_{\mathcal{U}}$). This procedure is very common in measure theory, e.g., in the study of L^p -spaces.

For reasons that will become clear in Section 6.4, ultraproducts are a much more useful tool than arbitrary reduced products. In the rest of this section (and essentially the rest of this book with the exception of Section 6.11), we restrict our attention to ultraproducts and ultrapowers.

Given a set M , an element $x \in M$, and an index set I , we let $a_x : I \rightarrow M$ be the function that is constantly equal to x . If we are also given an ultrafilter \mathcal{U} on I , we then have a function $d : M \rightarrow M^{\mathcal{U}}$ given by $d(x) := [a_x]_{\mathcal{U}}$.

Exercise 6.2.8. Prove that $d : M \rightarrow M^{\mathcal{U}}$ is injective.

Definition 6.2.9. The function d above is referred to as the **diagonal embedding** of M into its ultrapower $M^{\mathcal{U}}$.

In the remainder of this book, we often identify M with its image in $M^{\mathcal{U}}$ and view M as literally contained in its ultrapowers. For applications, we usually want ultrapowers to be bigger (and usually much bigger) than the original sets themselves. In other words, we do not want d to be onto. We will characterize when this happens in Section 6.6.

6.3. Ultraproducts of structures

Suppose now that we have a family $(\mathcal{M}_i)_{i \in I}$ of \mathcal{L} -structures and a filter \mathcal{F} on I . We set $N := \prod_{\mathcal{F}} M_i$, the reduced product of the underlying universes with respect to \mathcal{F} . We would like to make N the underlying universe of an \mathcal{L} -structure \mathcal{N} in a natural way.

When $\mathcal{F} := \{I\}$, N is simply the direct product of the family $(M_i)_{i \in I}$, and in many algebraic settings, the natural structure to put on the direct product is simply that of pointwise operations. We would like to take our cue from this particular instance of the reduced product and define the interpretation of the symbols from \mathcal{L} in N to be those induced by the pointwise operations. For this to work, we need to know that this is independent of representatives:

Exercise 6.3.1. Suppose that R is an n -ary relation symbol from \mathcal{L} and that F is an n -ary function symbol from \mathcal{L} . Suppose that $a_1, \dots, a_n, b_1, \dots, b_n \in \prod_{i \in I} M_i$ are such that $a_i \sim_{\mathcal{F}} b_i$ for all $i = 1, \dots, n$. Prove that:

- (1) $\{i \in I : (a_1(i), \dots, a_n(i)) \in R^{\mathcal{M}_i}\} \in \mathcal{F}$ if and only if $\{i \in I : (b_1(i), \dots, b_n(i)) \in R^{\mathcal{M}_i}\} \in \mathcal{F}$.
- (2) $\{i \in I : F^{\mathcal{M}_i}(a_1(i), \dots, a_n(i)) = F^{\mathcal{M}_i}(b_1(i), \dots, b_n(i))\} \in \mathcal{F}$.

By the previous exercise, we are entitled to consider the structure \mathcal{N} whose underlying universe is N and which interprets symbols as follows: Suppose that $[a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}} \in N$.

- If R is an n -ary relation symbol in \mathcal{L} , then $([a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}}) \in R^{\mathcal{N}}$ if and only if

$$\{i \in I : R^{\mathcal{M}_i}(a_1(i), \dots, a_n(i))\} \in \mathcal{F}.$$

- If F is an n -ary function symbol in \mathcal{L} , then $F^{\mathcal{N}}([a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}}) := [b]_{\mathcal{F}}$, where $b \in N$ satisfies $b(i) := F^{\mathcal{M}_i}(a_1(i), \dots, a_n(i))$.

Definition 6.3.2. The structure \mathcal{N} defined above is referred to as the **reduced product** of the family $(\mathcal{M}_i)_{i \in I}$ of structures and is denoted $\prod_{\mathcal{F}} \mathcal{M}_i$. One defines **ultraproducts**, **reduced powers**, and **ultrapowers** of structures analogously as in Definition 6.2.4.

Example 6.3.3. If $\mathcal{F} = \{I\}$, then the reduced product of the family $(\mathcal{M}_i)_{i \in I}$ is referred to as the **direct product** of the structures.

Example 6.3.4. If $\mathcal{F} = \mathcal{U}_j$, then the bijection $[a]_{\mathcal{U}_j} \mapsto a(j)$ yields an isomorphism between the ultraproduct $\prod_{\mathcal{U}_j} \mathcal{M}_i$ and the structure \mathcal{M}_j .

Exercise 6.3.5. Show that the diagonal embedding $d : M \rightarrow M^{\mathcal{U}}$ yields an embedding of \mathcal{L} -structures $d : \mathcal{M} \rightarrow \mathcal{M}^{\mathcal{U}}$.

In some algebraic cases, the reduced product can be rephrased in algebraic terms:

Exercise 6.3.6. Suppose that $(G_i)_{i \in I}$ is a family of groups and \mathcal{F} is a filter on I . Let

$$N_{\mathcal{F}} := \{a \in \prod_{i \in I} G_i : \{i \in I : a(i) = e_{G_i}\} \in \mathcal{F}\}.$$

Show that $N_{\mathcal{F}}$ is a normal subgroup of $\prod_{i \in I} G_i$ and $\prod_{i \in I} G_i / N_{\mathcal{F}} \cong \prod_{\mathcal{F}} G_i$ (as structures in the language of groups). Conclude that $\prod_{\mathcal{F}} G_i$ is a group.

For the next exercise, we recall the notion of induced ultrafilter from Exercise 1.1.12.

Exercise 6.3.7. Suppose that \mathcal{U} is an ultrafilter on I and $J \in \mathcal{U}$. Show that the map

$$[a]_{\mathcal{U}} \mapsto [a \upharpoonright J]_{\mathcal{U} \cap J} : \prod_{\mathcal{U}} \mathcal{M}_i \rightarrow \prod_{\mathcal{U} \cap J} \mathcal{M}_i$$

is an isomorphism.

By the previous exercise and Exercise 1.1.22, in connection with ultraproducts, one may essentially always assume that one is working with uniform ultrafilters.

The following simple exercise can be quite useful in various situations (see, for example, Example 8.6.5); it simply says that taking ultraproducts commutes with taking reducts.

Exercise 6.3.8. Suppose that $(\mathcal{M}_i)_{i \in I}$ is a family of \mathcal{L} -structures, $\mathcal{L}_0 \subseteq \mathcal{L}$ a sublanguage, and \mathcal{U} an ultrafilter on I . Prove that the identity map $\prod_{\mathcal{U}} \mathcal{M}_i \rightarrow \prod_{\mathcal{U}} \mathcal{M}_i$ yields an isomorphism $(\prod_{\mathcal{U}} \mathcal{M}_i) \upharpoonright \mathcal{L}_0 \cong \prod_{\mathcal{U}} (\mathcal{M}_i \upharpoonright \mathcal{L}_0)$.

6.4. Łoś's theorem

In this section, we prove the principal result explaining the connection between truth in ultraproducts and truth in the individual models. This theorem is called **Łoś's theorem** or sometimes the **Fundamental Theorem of Ultraproducts** and is the connection alluded to in the discussion in Section 1.

To motivate this theorem and why it is specific to ultraproducts (rather than arbitrary reduced products), let us consider an example.

Suppose that $(K_i)_{i \in I}$ is a family of fields. In a first course in algebra, one encounters the sad fact that the direct product $\prod_{i \in I} K_i$ is no longer a field (although it is still a commutative ring with unity). Indeed, if $i_0 \in I$ is a fixed index and one considers the element $a \in \prod_{i \in I} K_i$ for which $a(i) = 1$ for all $i \neq i_0$ and $a(i_0) = 0$, then a has no multiplicative inverse in $\prod_{i \in I} K_i$.

In this example, a “bad” phenomenon in one coordinate ruined the possibility that the corresponding element of the direct product could be invertible. However, from the point of view of an ultraproduct, such an isolated occurrence would have no effect on the corresponding element of the ultraproduct. In fact, the ultraproduct $\prod_{\mathcal{U}} K_i$ is a field! Let us check only the fact that every nonzero entry has a multiplicative inverse. Suppose that $[a]_{\mathcal{U}} \in \prod_{\mathcal{U}} K_i$ is nonzero. This means that $a(i) \in K_i \setminus \{0\}$ for \mathcal{U} -almost all $i \in I$. For these i , set $b(i) := a(i)^{-1}$; for the \mathcal{U} -small set of i for which $a(i) = 0$, define $b(i) \in K_i$ arbitrarily. It follows that $a(i)b(i) = 1$ for \mathcal{U} -almost all i , whence $[a]_{\mathcal{U}} \cdot [b]_{\mathcal{U}} = [1]_{\mathcal{U}}$ and $[b]_{\mathcal{U}}$ is the multiplicative inverse of $[a]_{\mathcal{U}}$.

Łoś's theorem provides an explanation for the previous example and many other applications of ultraproducts throughout mathematics: if the structures involved all satisfy some particular first-order property, then so will the ultraproduct.

Theorem 6.4.1 (Łoś's theorem). *Suppose that $(\mathcal{M}_i)_{i \in I}$ is a family of \mathcal{L} -structures and \mathcal{U} is an ultrafilter on I . Further suppose that $\varphi(x_1, \dots, x_m)$ is an \mathcal{L} -formula and $[a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}} \in \prod_{\mathcal{U}} M_i$. Then*

$$\prod_{\mathcal{U}} \mathcal{M}_i \models \varphi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}}) \Leftrightarrow \{i \in I : \mathcal{M}_i \models \varphi(a_1(i), \dots, a_m(i))\} \in \mathcal{U}.$$

Proof. We proceed by induction on the complexity of φ . Observe that the statement of Łoś's theorem when φ is atomic follows immediately from the definition of interpretations in ultraproducts. (Exercise.)

We now assume that $\varphi = \neg\psi$ and that the theorem holds for ψ . We then have that the following statements are equivalent:

- (1) $\prod_{\mathcal{U}} \mathcal{M}_i \models \varphi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}})$
- (2) $\prod_{\mathcal{U}} \mathcal{M}_i \not\models \psi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}})$
- (3) $\{i \in I : \mathcal{M}_i \models \psi(a_1(i), \dots, a_m(i))\} \notin \mathcal{U}$
- (4) $\{i \in I : \mathcal{M}_i \models \varphi(a_1(i), \dots, a_m(i))\} \in \mathcal{U}$

The equivalence between (2) and (3) follows from the induction hypothesis applied to ψ while the equivalence between (3) and (4) follows from the fact that \mathcal{U} is an ultrafilter (rather than just a filter).

We now assume that $\varphi = \psi \wedge \theta$ and that the theorem holds for ψ and θ . First assume that $\prod_{\mathcal{U}} \mathcal{M}_i \models \varphi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}})$, so $\prod_{\mathcal{U}} \mathcal{M}_i \models \psi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}})$ and $\prod_{\mathcal{U}} \mathcal{M}_i \models \theta([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}})$. By induction, each of the sets $X_{\psi} := \{i \in I : \mathcal{M}_i \models \psi(a_1(i), \dots, a_m(i))\}$ and $X_{\theta} := \{i \in I : \mathcal{M}_i \models \theta(a_1(i), \dots, a_m(i))\}$ belong to \mathcal{U} . Since $\{i \in I : \mathcal{M}_i \models \varphi(a_1(i), \dots, a_m(i))\}$ is precisely the intersection $X_{\psi} \cap X_{\theta}$, we are finished with this direction. Conversely, assume that $\{i \in I : \mathcal{M}_i \models \varphi(a_1(i), \dots, a_m(i))\} \in \mathcal{U}$

belongs to \mathcal{U} . Since this set is contained in X_ψ and X_φ , we can conclude that each of those sets belong to \mathcal{U} . It follows by induction that $\prod_{\mathcal{U}} \mathcal{M}_i \models \psi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}})$ and $\prod_{\mathcal{U}} \mathcal{M}_i \models \theta([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}})$, whence $\prod_{\mathcal{U}} \mathcal{M}_i \models \varphi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}})$.

We now finish by dealing with the case that φ is $\exists y\psi(\vec{x}, y)$ and that the theorem holds for ψ . First assume that $\prod_{\mathcal{U}} \mathcal{M}_i \models \varphi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}})$. It follows that there is $[b]_{\mathcal{U}} \in \prod_{\mathcal{U}} M_i$ such that

$$\prod_{\mathcal{U}} \mathcal{M}_i \models \psi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}}, [b]_{\mathcal{U}}),$$

whence by induction, $\{i \in I : \mathcal{M}_i \models \psi(a_1(i), \dots, a_m(i), b(i))\}$ belongs to \mathcal{U} . Since the latter set is included in $\{i \in I : \mathcal{M}_i \models \varphi(a_1(i), \dots, a_m(i))\}$, we are finished with this direction. Conversely, suppose that $\{i \in I : \mathcal{M}_i \models \varphi(a_1(i), \dots, a_m(i))\} \in \mathcal{U}$. For each i in that set, choose $b_i \in M_i$ such that $\mathcal{M}_i \models \psi(a_1(i), \dots, a_m(i), b_i)$. For the other (\mathcal{U} -small set of) i , let $b(i) \in M_i$ be arbitrary. Then $\{i \in I : \mathcal{M}_i \models \psi(a_1(i), \dots, a_m(i), b(i))\} \in \mathcal{U}$, so by induction $\prod_{\mathcal{U}} \mathcal{M}_i \models \psi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}}, [b]_{\mathcal{U}})$, which of course implies $\prod_{\mathcal{U}} \mathcal{M}_i \models \varphi([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}})$. \square

Remark 6.4.2. We note that in the proof of the existential case above, we used the axiom of choice to pick witnesses to the existential statements in each of the models for which that existential statement was true. In Section 6.5, we will go further into the connection between the axiom of choice and Łoś's theorem and show that this use of choice is unavoidable.

Exercise 6.4.3. Show that the diagonal embedding $d : \mathcal{M} \rightarrow \mathcal{M}^{\mathcal{U}}$ is an elementary embedding.

It is worth singling out the special case of Łoś's theorem for sentences:

Corollary 6.4.4. *Suppose that $(\mathcal{M}_i)_{i \in I}$ is a family of \mathcal{L} -structures and \mathcal{U} is an ultrafilter on I . Further suppose that σ is an \mathcal{L} -sentence. Then*

$$\prod_{\mathcal{U}} \mathcal{M}_i \models \sigma \Leftrightarrow \{i \in I : \mathcal{M}_i \models \sigma\} \in \mathcal{U}.$$

The previous corollary in particular implies that

$$\text{Th}\left(\prod_{\mathcal{U}} \mathcal{M}_i\right) = \lim_{\mathcal{U}} \text{Th}(\mathcal{M}_i)$$

in the Stone space $S(\mathbb{B}_{\mathcal{L}})$ as discussed in Section 1.

Exercise 6.4.5. Suppose that, for each $i \in I$, $f_i : \mathcal{M}_i \rightarrow \mathcal{N}_i$ is an elementary embedding. Further suppose that \mathcal{U} is an ultrafilter on I . Prove that the **ultraproduct embedding** $\prod_{\mathcal{U}} f_i : \prod_{\mathcal{U}} \mathcal{M}_i \rightarrow \prod_{\mathcal{U}} \mathcal{N}_i$ defined by $(\prod_{\mathcal{U}} f_i)([a]_{\mathcal{U}}) = [b]_{\mathcal{U}}$, where $b(i) := f_i(a(i))$, is an elementary embedding. In particular, if $\mathcal{M}_i \preceq \mathcal{N}_i$ for all $i \in I$, then $\prod_{\mathcal{U}} \mathcal{M}_i \preceq \prod_{\mathcal{U}} \mathcal{N}_i$.

Using Loś's theorem, we can fulfill a promise made in Section 4.1:

Theorem 6.4.6 (Ramsey's theorem, finite version). *Given $n, k, l \in \mathbb{N}$, there is $m \in \mathbb{N}$ such that whenever X is a set with $|X| \geq m$ and $X^{[k]} = \bigsqcup_{i=1}^l X_i$ is a partition of $X^{[k]}$ into l pieces, there is $Y \subseteq X$ with $|Y| \geq n$ and $i \in \{1, \dots, l\}$ such that $Y^{[k]} \subseteq X_i$.*

Proof. For notational simplicity, we assume that $k = l = 2$. (The proof of the general case is no more complicated, just notationally messier.) Suppose, toward a contradiction, that no such m exists. Then for each $m \in \mathbb{N}$, we may find a finite set $X(m)$ with $|X(m)| \geq m$ and a partition $X(m)^{[2]} = X_1(m) \sqcup X_2(m)$ such that there is no $Y \subseteq X(m)$ with $|Y| \geq n$ homogeneous for the partition. Let \mathcal{U} be any nonprincipal ultrafilter on \mathbb{N} and let $Z := \prod_{\mathcal{U}} X(m)$. Note that Z is infinite. We define a partition $Z^{[2]} = Z_1 \sqcup Z_2$ as follows. If $\{[a]_{\mathcal{U}}, [b]_{\mathcal{U}}\} \in Z^{[2]}$, then we have that $\{a(m), b(m)\} \in X(m)^{[2]}$ for \mathcal{U} -almost all m . There is then a unique $i \in \{1, 2\}$ such that $\{a(m), b(m)\} \in X_i(m)$ for \mathcal{U} -almost all m , and then we declare $\{[a]_{\mathcal{U}}, [b]_{\mathcal{U}}\} \in Z_i$ for this i .

By the infinite version of Ramsey's theorem (Theorem 4.1.1), there is an infinite set $Y \subseteq Z$ homogeneous for this coloring, say, without loss of generality, that $Y^{[2]} \subseteq Z_1$. Fix distinct elements $[a_1]_{\mathcal{U}}, \dots, [a_n]_{\mathcal{U}}$ from Y . It follows that there is a \mathcal{U} -large set of m such that $a_1(m), \dots, a_n(m)$ are distinct elements of $X(m)$ such that $\{a_i(m), a_j(m)\} \in X_1(m)$ for all $1 \leq i < j \leq n$, yielding the desired contradiction. \square

Exercise 6.4.7. Explain exactly how Loś's theorem was used in the previous proof.

Another nice application of Loś's theorem is an ultraproduct proof of the **compactness theorem**:

Theorem 6.4.8 (Compactness theorem). *If T is a finitely satisfiable set of \mathcal{L} -sentences, then T is satisfiable.*

Exercise 6.4.9. Prove the compactness theorem directly from Loś's theorem by finding a suitable ultrafilter \mathcal{U} on $I := \mathcal{P}_f(T)$ such that, letting $\mathcal{M}_{\Delta} \models \Delta$ for each $\Delta \in I$, we have $\prod_{\mathcal{U}} \mathcal{M}_{\Delta} \models T$.

This proof of the compactness theorem used AC via Loś's theorem. In the next section, we offer a different proof of the compactness theorem that reduces the use of AC to UT.

6.5. The ultrafilter theorem and the axiom of choice: Part II

We return to our study of the connection between AC and UT from Section 5.1 and show how the two new characters, Loś's theorem (Loś) and the

compactness theorem (CT), are connected. Our first main result is the following:

Theorem 6.5.1. *In ZF, the following are equivalent:*

- (1) *For every Boolean algebra \mathbb{B} , $S(\mathbb{B})$ is compact.*
- (2) CT.
- (3) UT.

Before proving this result, we need a quick detour.

Definition 6.5.2. Suppose that T is an \mathcal{L} -theory. We say that T is:

- (1) **Maximal finitely satisfiable** if T is finitely satisfiable and, for every \mathcal{L} -sentence σ , either $\sigma \in T$ or $\neg\sigma \in T$;
- (2) **Henkin** if, for every \mathcal{L} -formula $\varphi(x)$, there is a constant symbol c such that the \mathcal{L} -sentence $\exists x\varphi(x) \rightarrow \varphi(c)$ belongs to T .

Fact 6.5.3 (ZF). Suppose that T is a maximal finitely satisfiable Henkin theory. Then T is satisfiable.

Fact 6.5.4 (ZF). For any finitely satisfiable \mathcal{L} -theory T , there is a language \mathcal{L}' containing \mathcal{L} and a finitely satisfiable Henkin \mathcal{L}' -theory T' containing T .

Proof of Theorem 6.5.1. First suppose that (1) holds and T is a finitely satisfiable set of \mathcal{L} -sentences. We wish to show that T is satisfiable. By Fact 6.5.4, we may as well assume that T is also a Henkin theory. We may also assume that T is closed under logical implication, that is, if $\sigma \in T$ and $\sigma \rightarrow \tau$ is logically valid, then $\tau \in T$. By finite satisfiability of T , the collection $(U_\sigma)_{\sigma \in T}$ of basic closed sets in $S(\mathbb{B}_{\mathcal{L}})$ has the FIP, whence, by compactness of $S(\mathbb{B}_{\mathcal{L}})$, there is $T' \in \bigcap_{\sigma \in T} U_\sigma$. Note then that T' is a maximal finitely satisfiable set of \mathcal{L} -sentences containing T which is still Henkin. Thus, by Fact 6.5.3, T' is satisfiable, whence so is T .

Now suppose that (2) holds and let \mathcal{F} be a filter on a set I . Let \mathcal{L} be the language with constant symbols c_A for all $A \subseteq I$ and a single unary predicate symbol P . Let T be the following set of \mathcal{L} -sentences:

- $P(c_A)$ for all $A \in \mathcal{F}$;
- $P(c_A) \wedge P(c_B) \rightarrow P(c_{A \cap B})$;
- $P(c_A) \rightarrow P(c_B)$ whenever $A \subseteq B$;
- $P(c_A) \vee P(c_{I \setminus A})$ for all $A \subseteq I$.

We leave it as an exercise to check that T is finitely satisfiable, whence it is satisfiable by CT. Fix $\mathcal{M} \models T$ and define \mathcal{U} by setting $A \in \mathcal{U}$ if and only if $\mathcal{M} \models P(c_A)$. It is immediate from the definition of T that \mathcal{U} is an ultrafilter on I extending \mathcal{F} .

Finally, suppose that (3) holds and fix a Boolean algebra \mathbb{B} . We mentioned in Section 5.1 that Tychonoff's theorem for compact Hausdorff spaces is equivalent to UT, whence $2^{\mathbb{B}}$ is compact. Since the proof that $S(\mathbb{B})$ is closed in $2^{\mathbb{B}}$ does not make any use of AC, it follows that $S(\mathbb{B})$ is also compact. \square

Exercise 6.5.5. Prove that T in the (2)-implies-(3) direction of the above proof is finitely satisfiable.

We now turn to the question about the connection between AC and Loś. One might wonder if there is not a more clever proof of Loś that uses a weaker version of AC, perhaps UT? It turns out that this is not the case:

Theorem 6.5.6 (ZF). *UT + Los implies AC, whence AC is equivalent to UT + Los.*

Proof. Work in ZF + UT + Los. Suppose, toward a contradiction, that X is a set of nonempty sets without a choice function. Without loss of generality, we may assume that the sets in X are pairwise disjoint and no element of X is itself an element of an element of X . Let $\mathcal{L} = \{R\}$, where R is a single binary relation symbol. We consider the \mathcal{L} -structure \mathcal{M} which has universe $X \cup \bigcup X$, where $\bigcup X$ denotes the set of elements of elements of X , and such that $(t, y) \in R^{\mathcal{M}}$ if and only if either (i) $y \in X$ and $t \in y$ or (2) $t = y \in \bigcup X$.

Set

$$\mathcal{F} := \{z \subseteq X : X \setminus z \text{ has a choice function}\}.$$

Clearly, $X \in \mathcal{F}$ (as the emptyset has a choice function) and by assumption $\emptyset \notin \mathcal{F}$. If $z \in \mathcal{F}$ and $z \subseteq w$, then $X \setminus w \subseteq X \setminus z$, whence $X \setminus w$ has a choice function (the restriction of the choice function on $X \setminus z$), whence $w \in \mathcal{F}$. Finally, if $w, z \in \mathcal{F}$, then $(X \setminus w) \cup (X \setminus z)$ has a choice function (exercise), whence $w \cap z \in \mathcal{F}$.

It follows that \mathcal{F} is a filter on X , whence we may extend it to an ultrafilter \mathcal{U} on X by UT. Since $\mathcal{M} \models \forall y \exists t R(t, y)$, we conclude from Loś that $\mathcal{M}^{\mathcal{U}} \models \forall y \exists t R(t, y)$. We apply this in the case that $y = [\text{id}_X]_{\mathcal{U}}$, which makes sense as an element of $\mathcal{M}^{\mathcal{U}}$ since the index set for the filter is X and X is a subset of the universe of \mathcal{M} . We thus have $f : X \rightarrow \mathcal{M}$ such that $\mathcal{M}^{\mathcal{U}} \models R([f]_{\mathcal{U}}, [\text{id}_X]_{\mathcal{U}})$. By the definition of $R^{\mathcal{M}}$, this means that $\{y \in X : f(y) \in y\} \in \mathcal{U}$. But $\{y \in X : f(y) \in y\}$ has a choice function (tautologically!), whence its complement belongs to \mathcal{F} and thus \mathcal{U} , yielding a contradiction. \square

Corollary 6.5.7. *There is a model of ZF where Loś is true but AC is false.*

Proof. In a model of ZF where WUT fails, Loś holds vacuously (as all ultraproducts are principal) but AC fails (else WUT would hold). \square

Corollary 6.5.8. *There is a model of $ZF+UT$ in which $Loś$ fails.*

Proof. If $ZF+UT$ proved $Loś$, then $ZF+UT$ proves AC by Theorem 6.5.6, which we know it does not. \square

6.6. Countably incomplete ultrafilters

In this section, we return to the question discussed at the end of Section 6.2: When is the diagonal embedding $d : M \rightarrow M^{\mathcal{U}}$ onto?

In the case of a finite set, we have the following result:

Exercise 6.6.1. Suppose that M is finite. Then for any index set I and any ultrafilter \mathcal{U} on I , prove that $d : M \rightarrow M^{\mathcal{U}}$ is a bijection.

In the case of a countable index set, it is easy to verify when d is onto:

Exercise 6.6.2. Suppose that M is an infinite set, I is a countable set, and \mathcal{U} is an ultrafilter on I . Prove that $d : M \rightarrow M^{\mathcal{U}}$ is a bijection if and only if \mathcal{U} is a principal ultrafilter.

To answer the above question for an arbitrary index set, we need some new definitions:

Definition 6.6.3. Suppose that \mathcal{U} is an ultrafilter and κ is a cardinal. We say that \mathcal{U} is κ -**complete** if whenever $Y \subseteq \mathcal{U}$ is such that $|Y| < \kappa$, then $\bigcap Y \in \mathcal{U}$. (In other words, \mathcal{U} is κ -complete if and only if \mathcal{U} is closed under intersections of families of size $< \kappa$.) \aleph_1 -complete ultrafilters are often referred to as **countably complete** and an ultrafilter that is not countably complete is called **countably incomplete**.

Some exercises to get us acquainted:

Exercise 6.6.4. Prove that every ultrafilter is \aleph_0 -complete.

Exercise 6.6.5. Prove that an ultrafilter \mathcal{U} is κ -complete for all κ if and only if \mathcal{U} is principal.

Exercise 6.6.6. Suppose that \mathcal{U} is κ -complete and $\lambda < \kappa$. Prove that \mathcal{U} is also λ -complete.

Exercise 6.6.7. Prove that an ultrafilter \mathcal{U} on the index set I is countably incomplete if and only if there is a sequence $(E_n)_{n \in \mathbb{N}}$ from \mathcal{U} such that $I = E_0 \supseteq E_1 \supseteq E_2 \supseteq \dots$ and $\bigcap_{n \in \mathbb{N}} E_n = \emptyset$.

Exercise 6.6.8. Suppose that \mathcal{U} is κ -complete and $\mathcal{V} \leq_{RK} \mathcal{U}$. Prove that \mathcal{V} is κ -complete.

The next lemma shows us that a nonprincipal ultrafilter on a set of size κ is never κ^+ -complete.

Lemma 6.6.9. *Suppose that \mathcal{U} is an ultrafilter on a set I with $|I| = \kappa$. If \mathcal{U} is κ^+ -complete, then \mathcal{U} is principal. In particular, an ultrafilter on a countable set is countably incomplete if and only if it is nonprincipal.*

Proof. Suppose that \mathcal{U} is nonprincipal. Then for each $i \in I$, we have that $I \setminus \{i\} \in \mathcal{U}$. Since $\bigcap_{i \in I} (I \setminus \{i\}) = \emptyset$ and $|I| = \kappa$, we see that \mathcal{U} is not κ^+ -complete. \square

Remark 6.6.10. By the previous lemma, the most complete that a nonprincipal ultrafilter on a set of cardinality κ can be is κ -complete. A cardinal κ that possesses a nonprincipal κ -complete ultrafilter is called **measurable**. The existence of an uncountable measurable cardinal cannot be proven in ZFC. In fact, if there exists an uncountable cardinal that possesses a countably complete ultrafilter, then there exists a measurable cardinal. (All of these facts will be discussed and proven in Chapter 17.) For this reason, when one restricts one's attention to countably incomplete ultrafilters, this is really no loss of generality if one wants to stay within the confines of ZFC.

Exercise 6.6.11. Prove that an ultrafilter \mathcal{U} is countably incomplete if and only if there is a nonprincipal ultrafilter \mathcal{V} on \mathbb{N} such that $\mathcal{V} \leq_{RK} \mathcal{U}$. (*Hint.* For the forward direction, let $(E_n)_{n \in \mathbb{N}}$ be as in Exercise 6.6.7 and define $f : I \rightarrow \mathbb{N}$ by $f(x)$ equals the maximal n such that $x \in E_n$. Show that $f(\mathcal{U})$ is nonprincipal.)

It will be useful to reformulate the notion of κ -completeness in terms of partitions of the index set:

Lemma 6.6.12. *Suppose that \mathcal{U} is an ultrafilter on the index set I . Then \mathcal{U} is κ -complete if and only if, for every partition of I into fewer than κ many pieces, exactly one of the pieces belongs to \mathcal{U} .*

Proof. First suppose that \mathcal{U} is κ -complete and $I = \bigcup_{\alpha < \lambda} X_\alpha$, where $\lambda < \kappa$. Note then that $\bigcap_{\alpha < \lambda} (I \setminus X_\alpha) = \emptyset$. Since \mathcal{U} is κ -complete, it follows that $I \setminus X_\alpha \notin \mathcal{U}$ for some $\alpha < \lambda$, whence $X_\alpha \in \mathcal{U}$. The uniqueness of X_α follows from the fact that the X_α 's are pairwise disjoint.

Suppose now that for every partition of I into fewer than κ many pieces, exactly one of the pieces belongs to \mathcal{U} . We show that \mathcal{U} is κ -complete. Suppose that $Y \subseteq I$ is such that $|Y| < \kappa$. Enumerate $Y = \{Y_\alpha : \alpha < \lambda\}$, where $\lambda < \kappa$. We define a partition $(X_\alpha)_{\alpha \in \lambda \cup \{\lambda\}}$ of I as follows. First, set $X_\lambda := \bigcap Y$. Next, if $i \notin \bigcap Y$, we put $i \in X_\alpha$ if $\alpha < \lambda$ is least with $i \notin Y_\alpha$. By our assumption, $X_\alpha \in \mathcal{U}$ for a unique $\alpha \leq \lambda$. However, for $\alpha < \lambda$, since $X_\alpha \cap Y_\alpha = \emptyset$ and $Y_\alpha \in \mathcal{U}$, we see that $X_\alpha \notin \mathcal{U}$. Consequently, $X_\lambda = \bigcap Y \in \mathcal{U}$, as desired. \square

We are now ready to see the connection between complete ultrafilters and the surjectivity of the diagonal embedding.

Proposition 6.6.13. *Suppose that \mathcal{U} is an ultrafilter on I .*

- (1) *If \mathcal{U} is countably complete, then for any countable set M , $d : M \rightarrow M^{\mathcal{U}}$ is onto.*
- (2) *If there is an infinite M such that $d : M \rightarrow M^{\mathcal{U}}$ is onto, then \mathcal{U} is countably complete.*

Proof. (1) Suppose that \mathcal{U} is countably complete and M is countable. Fix $a : I \rightarrow M$; we show that $[a]_{\mathcal{U}}$ is in the image of d . Enumerate $M = \{x_j : j \in \mathbb{N}\}$ and set $X_j := \{i \in I : a(i) = x_j\}$. Then X_j forms a countable partition of I , whence, by Lemma 6.6.12, there is a unique j for which $X_j \in \mathcal{U}$, and hence $[a]_{\mathcal{U}} = d(x_j)$.

(2) Suppose that M is an infinite set such that d is onto and that $(X_j)_{j \in \mathbb{N}}$ is a countable partition of I . Let $(x_j)_{j \in \mathbb{N}}$ be a collection of pairwise distinct elements of M and define $a : I \rightarrow M$ by setting $a(i) = x_j$ if and only if $i \in X_j$. Take j such that $[a]_{\mathcal{U}} = d(x_j)$; it follows that $X_j \in \mathcal{U}$, whence \mathcal{U} is countably complete by Lemma 6.6.12. \square

Exercise 6.6.14. Adapt the proof of the previous proposition to show the following: If M is a set with $|M| = \kappa$ and \mathcal{U} is an ultrafilter on a set I , then the diagonal embedding $d : M \rightarrow M^{\mathcal{U}}$ is onto if and only if \mathcal{U} is κ^+ -complete.

6.7. Revisiting the Rudin-Keisler order

In this section, we show that there is a connection between the Rudin-Keisler order introduced in Section 1.6 and the embeddability relation between ultrapowers. More precisely, we have:

Theorem 6.7.1. *Suppose that \mathcal{U} and \mathcal{V} are ultrafilters on I and J , respectively. We then have:*

- (1) *$\mathcal{U} \leq_{RK} \mathcal{V}$ if and only if, for every structure \mathcal{M} (in any language), $\mathcal{M}^{\mathcal{U}}$ elementarily embeds into $\mathcal{M}^{\mathcal{V}}$.*
- (2) *$\mathcal{U} \equiv_{RK} \mathcal{V}$ if and only if, for every structure \mathcal{M} (again, in any language), $\mathcal{M}^{\mathcal{U}} \cong \mathcal{M}^{\mathcal{V}}$.*

Proof. First suppose that $\mathcal{U} \leq_{RK} \mathcal{V}$ and take $f : J \rightarrow I$ such that $\mathcal{U} = f(\mathcal{V})$. Fix a structure \mathcal{M} . We check that the map $[a]_{\mathcal{U}} \mapsto [a \circ f]_{\mathcal{V}} : \mathcal{M}^{\mathcal{U}} \rightarrow \mathcal{M}^{\mathcal{V}}$ is an elementary embedding, proving the forward direction of (1). Indeed, suppose that $\mathcal{M}^{\mathcal{U}} \models \varphi([a_1]_{\mathcal{U}}, \dots, [a_n]_{\mathcal{U}})$. By Łoś's theorem, we have that $\{i \in I \mid \mathcal{M} \models \varphi(a_1(i), \dots, a_n(i))\} \in \mathcal{U}$. By the choice of f , we have that the

preimage of this latter set under f belongs to \mathcal{V} , that is,

$$\{j \in J : \mathcal{M} \models \varphi(a_1(f(j)), \dots, a_n(f(j)))\} \in \mathcal{V},$$

whence $\mathcal{M}^\mathcal{V} \models \varphi([a_1 \circ f]_\mathcal{V}, \dots, [a_n \circ f]_\mathcal{V})$, as desired.

Suppose, in addition, that $\mathcal{V} \leq_{RK} \mathcal{U}$. We claim that the above embedding of $\mathcal{M}^\mathcal{U}$ into $\mathcal{M}^\mathcal{V}$ is actually surjective, whence an isomorphism between $\mathcal{M}^\mathcal{U}$ and $\mathcal{M}^\mathcal{V}$, proving the forward direction of (2). By Corollary 1.3.16, we may assume that f can be chosen so that there is $X \in \mathcal{V}$ such that $f \upharpoonright X$ is injective. Fix $[b]_\mathcal{V} \in M^\mathcal{V}$; we seek $[a]_\mathcal{U} \in M^\mathcal{U}$ such that $[a \circ f]_\mathcal{V} = [b]_\mathcal{V}$. For $i \in f(X)$, define $a(i) := b(f^{-1}(i))$, which is well defined as f is injective on X . For $i \in I \setminus f(X)$, define $a(i)$ arbitrarily. For $j \in X$, we then have $a(f(j)) = b(j)$, whence $[a \circ f]_\mathcal{V} = [b]_\mathcal{V}$, as desired.

We now suppose that $\mathcal{M}^\mathcal{U}$ elementarily embeds into $\mathcal{M}^\mathcal{V}$ for every structure \mathcal{M} , and we show that $\mathcal{U} \leq_{RK} \mathcal{V}$. Let \mathcal{L}_I be the language with a unary predicate P_A for each $A \subseteq I$. Let \mathcal{I} be the structure with universe I and for which $P_A^\mathcal{I} = A$. Let $e : \mathcal{I}^\mathcal{U} \rightarrow \mathcal{I}^\mathcal{V}$ be an elementary embedding, which exists by our assumption. Let id denote the identity function on I and let $f : J \rightarrow I$ be such that $e([\text{id}]_\mathcal{U}) = [f]_\mathcal{V}$. We then note that, for $A \subseteq I$, the following are equivalent:

- $A \in \mathcal{U}$;
- $\{i \in I : \mathcal{I} \models P_A(i)\} \in \mathcal{U}$;
- $\mathcal{I}^\mathcal{U} \models P_A([\text{id}]_\mathcal{U})$;
- $\mathcal{I}^\mathcal{V} \models P_A([f]_\mathcal{V})$;
- $\{j \in J : \mathcal{I} \models P_A(f(j))\} \in \mathcal{V}$;
- $f^{-1}(A) \in \mathcal{V}$.

It follows that $\mathcal{U} = f(\mathcal{V})$, so $\mathcal{U} \leq_{RK} \mathcal{V}$, proving the backward direction of (1).

Finally, if $\mathcal{M}^\mathcal{U} \cong \mathcal{M}^\mathcal{V}$ for every \mathcal{M} , then by the previous paragraph, $\mathcal{U} \leq_{RK} \mathcal{V}$ and $\mathcal{V} \leq_{RK} \mathcal{U}$, whence $\mathcal{U} \equiv_{RK} \mathcal{V}$, proving the backward direction of (2). \square

We can use the above interpretation to give a nice characterization of minimal ultrafilters in terms of substructures of the corresponding ultrapowers. First, some preparation.

We fix an infinite structure \mathcal{M} , a nonprincipal ultrafilter \mathcal{U} on I , and we set $\mathcal{N} := \mathcal{M}^\mathcal{U}$. For each $f : I \rightarrow M$, we set $\mathcal{N}[f] := \{[g \circ f]_\mathcal{U} : g : M \rightarrow M\}$. This is a substitute for the substructure of \mathcal{N} generated $[f]_\mathcal{U}$. In fact:

Exercise 6.7.2. If each function $M \rightarrow M$ is the interpretation of a function symbol in the language, prove that $\mathcal{N}[f]$ is the substructure of \mathcal{N} generated by $[f]_\mathcal{U}$.

The following exercise highlights some properties of the above construction:

Exercise 6.7.3.

- (1) Suppose that $f, h : I \rightarrow M$ and $h(\mathcal{U}) \leq_{RK} f(\mathcal{U})$. Then $[h]_{\mathcal{U}} \in \mathcal{N}[f]$.
- (2) If f is a constant function, then $\mathcal{N}[f] = \mathcal{M}$ (viewed as a substructure of \mathcal{N} via the diagonal embedding).
- (3) $\mathcal{N}[f]$ is an elementary substructure of \mathcal{N} and $\mathcal{N}[f] \cong \mathcal{M}^{f(\mathcal{U})}$ via the isomorphism $[g \circ f]_{\mathcal{U}} \mapsto [g \upharpoonright f(I)]_{f(\mathcal{U})}$.

The following theorem highlights the connection between the Rudin-Keisler ordering and the substructure $\mathcal{N}[f]$:

Theorem 6.7.4. *If $f(\mathcal{U}) \equiv_{RK} \mathcal{U}$, then $\mathcal{N}[f] = \mathcal{N}$. If $|I| \leq |M|$, then the converse holds.*

Proof. First suppose that $f(\mathcal{U}) \equiv_{RK} \mathcal{U}$. By Corollary 1.3.16, we may suppose that f is chosen so that there is $X \in \mathcal{U}$ such that $f \upharpoonright X$ is injective. As argued in the proof of Theorem 6.7.1, given any $h : I \rightarrow M$, there is $g : M \rightarrow M$ such that $[g \circ f]_{\mathcal{U}} = [h]_{\mathcal{U}}$. It follows that $\mathcal{N}[f] = \mathcal{N}$.

For the converse, assume that $|I| \leq |M|$ and that $\mathcal{N}[f] = \mathcal{N}$. Fix an injective function $h : I \rightarrow M$ and take $g : M \rightarrow M$ such that $[h]_{\mathcal{U}} = [g \circ f]_{\mathcal{U}}$. By Exercise 1.3.4, we have that $h(\mathcal{U}) = (g \circ f)(\mathcal{U}) = g(f(\mathcal{U}))$, whence $h(\mathcal{U}) \leq_{RK} f(\mathcal{U})$. On the other hand, since h is injective, we have that $\mathcal{U} \equiv_{RK} h(\mathcal{U})$. Clearly, $f(\mathcal{U}) \leq_{RK} \mathcal{U}$. Altogether, we have $f(\mathcal{U}) \equiv_{RK} \mathcal{U}$, as desired. \square

Here is the promised characterization of minimal ultrafilters in terms of the corresponding ultrapowers:

Corollary 6.7.5. *Suppose that \mathcal{M} is an infinite structure, \mathcal{U} a nonprincipal ultrafilter over \mathbb{N} , and set $\mathcal{N} := \mathcal{M}^{\mathcal{U}}$. Then \mathcal{U} is minimal if and only if, for every $f : \mathbb{N} \rightarrow M$, either $\mathcal{N}[f] = \mathcal{M}$ or $\mathcal{N}[f] = \mathcal{N}$. In particular, if \mathcal{M} has function symbols for every function $M \rightarrow M$, then \mathcal{U} is minimal if and only if the only substructures of \mathcal{N} are \mathcal{M} and \mathcal{N} .*

Exercise 6.7.6. Prove Corollary 6.7.5.

6.8. Cardinalities of ultraproducts

In this section, we discuss some results concerning the cardinalities of ultraproducts. First, some easy facts:

Exercise 6.8.1. Let \mathcal{U} be an ultrafilter on the index set I .

- (1) If $|M_i| = |N_i|$ for all $i \in I$, then $|\prod_{\mathcal{U}} M_i| = |\prod_{\mathcal{U}} N_i|$.

- (2) If $|M_i| \leq |N_i|$ for all $i \in I$, then $|\prod_{\mathcal{U}} M_i| \leq |\prod_{\mathcal{U}} N_i|$.
 (3) $|\prod_{\mathcal{U}} M_i| \leq |\prod_{i \in I} M_i|$.
 (4) $|M| \leq |M^{\mathcal{U}}| \leq |M|^{|I|}$.

Exercise 6.8.2. Prove that $\prod_{\mathcal{U}} M_i$ is finite if and only if there is $n \in \mathbb{N}$ such that $|M_i| \leq n$ for \mathcal{U} -almost all $i \in I$.

The next theorem shows that ultraproducts are usually always large in size:

Theorem 6.8.3. *Suppose that \mathcal{U} is a countably incomplete ultrafilter over an index set I and $(M_i)_{i \in I}$ is a family of sets. Then $\prod_{\mathcal{U}} M_i$ is either finite or has size at least \mathfrak{c} .*

Proof. Suppose that $\prod_{\mathcal{U}} M_i$ is infinite. We separate into two cases:

Case 1. M_i is finite for \mathcal{U} -almost all i . Without loss of generality, we may assume that M_i is finite for all $i \in I$. For each $i \in I$, set $n_i := |M_i|$. By Exercise 6.8.2, we know that $\lim_{\mathcal{U}} n_i = \infty$. By Exercise 6.8.1, we may assume that $M_i = [n_i] := \{0, 1, \dots, n_i - 1\}$. Define a function $f: \prod_{\mathcal{U}} M_i \rightarrow [0, 1]$ by defining $f([a]_{\mathcal{U}}) := \lim_{\mathcal{U}} \frac{a(i)}{n_i}$. It suffices to show that $[0, 1)$ is contained in the range of f . Toward this end, given $x \in [0, 1)$, define $a_x \in \prod_{i \in I} M_i$ by $a_x(i) = k$, where $k \in [n_i]$ is such that $\frac{k}{n_i} \leq x < \frac{k+1}{n_i}$. Note that $|\frac{a_x(i)}{n_i} - x| < \frac{1}{n_i}$ for all $i \in I$. We claim that $f([a_x]_{\mathcal{U}}) = x$. To see this, fix $N \in \mathbb{N}$ and note that $n_i > N$ for \mathcal{U} -almost all $i \in I$. For these $i \in I$, we have that $|\frac{a_x(i)}{n_i} - x| < \frac{1}{n_i} < \frac{1}{N}$, whence $|f([a_x]_{\mathcal{U}}) - x| \leq \frac{1}{N}$. Since $N \in \mathbb{N}$ was arbitrary, we see that $f([a_x]_{\mathcal{U}}) = x$, as desired.

Case 2. M_i is infinite for \mathcal{U} -almost all i . In this case, by Exercise 6.8.1, it is enough to show that $|\mathbb{N}^{\mathcal{U}}| \geq \mathfrak{c}$. By Exercise 6.6.11 and Theorem 6.7.1, it is enough to assume that \mathcal{U} is a nonprincipal ultrafilter on \mathbb{N} . (It is here that we have used that \mathcal{U} is countably incomplete.) However, by Case 1, $\prod_{\mathcal{U}} [n]$ has cardinality $\geq \mathfrak{c}$, whence, by Exercise 6.8.1 again, so does $\mathbb{N}^{\mathcal{U}}$. \square

The following corollary of the previous theorem is worth singling out:

Corollary 6.8.4. *Suppose that \mathcal{U} is a nonprincipal ultrafilter on \mathbb{N} and $(M_n)_{n \in \mathbb{N}}$ is a family of sets with $|M_n| \leq \mathfrak{c}$ for all $n \in \mathbb{N}$. Then $\prod_{\mathcal{U}} M_n$ is either finite or has size exactly \mathfrak{c} .*

Proof. The result follows immediately from Theorem 6.8.3 and the fact that $\prod_{\mathcal{U}} M_n$ has size at most \mathfrak{c} under the current assumptions as its cardinality is bounded by the cardinality of the direct product $\prod_{n \in \mathbb{N}} M_n$, which has size at most $\mathfrak{c}^{\aleph_0} = \mathfrak{c}$. \square

The following theorem can often be used to show that a particular structure is *not* an ultrapower.

Theorem 6.8.5. *Suppose that \mathcal{U} is a countably incomplete ultrafilter on I and M is an infinite set such that $M^{\mathcal{U}}$ has cardinality κ . Then $\kappa^\omega = \kappa$. In particular, $\text{cof}(\kappa) > \omega$.*

Proof. Set $N := M^{<\omega}$. Since $|N| = |M|$, it suffices to prove that $\kappa^\omega \leq |N^{\mathcal{U}}|$. Since \mathcal{U} is countably incomplete, by Exercise 6.6.7, we may fix $I = X_0 \supseteq X_1 \supseteq X_2 \supseteq \dots$ with each $X_n \in \mathcal{U}$ and $\bigcap_{n \in \mathbb{N}} X_n = \emptyset$. This allows us to define, for each $i \in I$, the number $n(i) :=$ the maximal n such that $i \in X_n$.

We now define $\sigma : (M^I)^\omega \rightarrow N^I$ by setting

$$\sigma(g)(i) := (g(1)(i), \dots, g(n(i))(i)).$$

We would like to define a sort of “inverse” to σ , but at the level of ultrapowers. More specifically, set $N_0 \subseteq N^{\mathcal{U}}$ to consist of those elements of the form $[\sigma(g)]_{\mathcal{U}}$ for some $g \in (M^I)^\omega$. We would like to define $\tau : N_0 \rightarrow (M^{\mathcal{U}})^\omega$ by setting $\tau([\sigma(g)]_{\mathcal{U}}) := ([g(0)]_{\mathcal{U}}, [g(1)]_{\mathcal{U}}, \dots)$. If this is possible, then since τ is clearly surjective, we achieve the desired result. In order for τ to be well defined, we need to know that $\sigma(g) \equiv_{\mathcal{U}} \sigma(h)$ implies that $g(n) \equiv_{\mathcal{U}} h(n)$ for all n . So suppose that $\sigma(g) \equiv_{\mathcal{U}} \sigma(h)$ and set $X := \{i \in I : \sigma(g)(i) = \sigma(h)(i)\} \in \mathcal{U}$. Fix n ; we aim to show that $g(n) \equiv_{\mathcal{U}} h(n)$. For $i \in X \cap X_n$, we have that $n \leq n(i)$, whence $g(n)(i) = h(n)(i)$. Since $X \cap X_n \in \mathcal{U}$, the desired result follows. \square

We will present one more result on cardinalities of ultrapowers in Section 8.3.

6.9. Iterated ultrapowers

It is natural to wonder what happens if you take an ultrapower of an ultrapower. It turns out that the resulting structure is itself an ultrapower, as the next theorem indicates. (We ask the reader to recall the notion of product ultrafilter from Exercise 1.6.9.)

Theorem 6.9.1. *Suppose that \mathcal{M} is a structure and \mathcal{U} and \mathcal{V} are ultrafilters on sets I and J , respectively. Then $(\mathcal{M}^{\mathcal{U}})^{\mathcal{V}} \cong \mathcal{M}^{\mathcal{U} \times \mathcal{V}}$.*

Exercise 6.9.2. Prove Theorem 6.9.1. (*Hint.* For $a : I \times J \rightarrow M$, set $a_j : I \rightarrow M$ to be the function $a_j(i) := a(i, j)$. Show that the map $[a]_{\mathcal{U} \times \mathcal{V}} \mapsto [a^*]_{\mathcal{V}}$, where $a^*(j) := [a_j]_{\mathcal{U}}$, is an isomorphism.)

Remark 6.9.3. In Section 8.4, we will see examples of ultrafilters \mathcal{U} and \mathcal{V} such that $\mathcal{U} \times \mathcal{V}$ is not Rudin-Keisler equivalent to $\mathcal{V} \times \mathcal{U}$. Consequently, by Theorem 6.7.1, there will be a structure \mathcal{M} such that $\mathcal{M}^{\mathcal{U} \times \mathcal{V}} \not\cong \mathcal{M}^{\mathcal{V} \times \mathcal{U}}$, whence $(\mathcal{M}^{\mathcal{U}})^{\mathcal{V}} \not\cong (\mathcal{M}^{\mathcal{V}})^{\mathcal{U}}$.

Exercise 6.9.4. Prove that the product operation on ultrafilters is associative. More precisely, if \mathcal{U} , \mathcal{V} , and \mathcal{W} are ultrafilters on sets I , J , and K , prove that $(\mathcal{U} \times \mathcal{V}) \times \mathcal{W}$ and $\mathcal{U} \times (\mathcal{V} \times \mathcal{W})$ are the same ultrafilter on $I \times J \times K$.

By Theorem 6.9.1 and Exercise 6.9.4, given any ultrafilters $\mathcal{U}_1, \dots, \mathcal{U}_n$ and any structure \mathcal{M} , the structure one gets by iterating the ultrapower construction relative to the various \mathcal{U}_i 's is isomorphic to the ultrapower $\mathcal{M}^{\mathcal{U}_1 \times \dots \times \mathcal{U}_n}$. In other words, finite iterated ultrapowers do not yield us genuinely new structures in the sense that they are just isomorphic to ordinary ultrapowers. However, it turns out that there is a construction of iterated ultrapowers for an arbitrary linearly ordered set of ultrafilters which can produce structures that are not obtainable as ordinary ultrapowers. We now explain this procedure.

Fix a linearly ordered set $(X, <)$ and, for each $x \in X$, suppose that \mathcal{U}_x is an ultrafilter on some set I_x . For each nonempty finite $Y \subseteq X$, we let $\mathcal{U}_Y := \mathcal{U}_{y_1} \times \dots \times \mathcal{U}_{y_n}$, where $y_1 < \dots < y_n$ is an increasing enumeration of Y . When $Y = \emptyset$, we define $\mathcal{M}^{\mathcal{U}_Y} := \mathcal{M}$.

Viewing $\mathcal{P}_f(X)$ as a directed set under inclusion, given any structure \mathcal{M} , we notice that there are natural embeddings $\mathcal{M}^{\mathcal{U}_Y} \rightarrow \mathcal{M}^{\mathcal{U}_Z}$ whenever $Y \subseteq Z$ belong to $\mathcal{P}_f(X)$. Indeed, for simplicity, suppose that $y_1 < \dots < y_n$ is an increasing enumeration of Y and that $Z = Y \cup \{z\}$. (The general case is no more difficult, just notationally messier.). Suppose that $i \in \{1, \dots, n\}$ is such that $y_i < z < y_{i+1}$. (If $i = 0$, this just means that $z < y_0$ while if $i = n$, this just means that $y_n < z$.) Set $Y_1 := \{y_1, \dots, y_i\}$, $Y_2 := \{y_{i+1}, \dots, y_n\}$, and $Z_1 = Y_1 \cup \{z\}$. Then $\mathcal{M}^{\mathcal{U}_{Z_1}} \cong (\mathcal{M}^{\mathcal{U}_{Y_1}})^{\mathcal{U}_{Z_2}}$. Let $i : \mathcal{M}^{\mathcal{U}_{Y_1}} \rightarrow \mathcal{M}^{\mathcal{U}_{Z_1}}$ be the elementary embedding obtained by composing the diagonal embedding of $\mathcal{M}^{\mathcal{U}_{Y_1}}$ into $(\mathcal{M}^{\mathcal{U}_{Y_1}})^{\mathcal{U}_{Z_2}}$ with the isomorphism $(\mathcal{M}^{\mathcal{U}_{Y_1}})^{\mathcal{U}_{Z_2}} \rightarrow \mathcal{M}^{\mathcal{U}_{Z_1}}$ given in the proof of Theorem 6.9.1. By Exercise 6.4.5, we have that the ultrapower embedding $i^{\mathcal{U}_{Y_2}} : (\mathcal{M}^{\mathcal{U}_{Y_1}})^{\mathcal{U}_{Y_2}} \rightarrow (\mathcal{M}^{\mathcal{U}_{Z_1}})^{\mathcal{U}_{Y_2}}$ is an elementary embedding. By using the isomorphisms from Theorem 6.9.1 again, this yields an elementary embedding $\mathcal{M}^{\mathcal{U}_Y} \rightarrow \mathcal{M}^{\mathcal{U}_Z}$, as desired.

We define the direct limit of the directed system from the previous paragraph to be the **iterated ultrapower** of the structure \mathcal{M} relative to the family $(\mathcal{U}_x)_{x \in X}$ and denoted it by $\mathcal{M}^{\mathcal{U}_X}$. We warn the reader that this is simply notation and that there is not actually an ultrafilter \mathcal{U}_X being defined. We also note that the natural embeddings of each $\mathcal{M}^{\mathcal{U}_Y}$ into $\mathcal{M}^{\mathcal{U}_X}$ are elementary. In particular, the natural embedding $\mathcal{M} \rightarrow \mathcal{M}^{\mathcal{U}_X}$ is elementary.

There is a special case of the above construction that is especially appealing. First, a couple of definitions:

Definition 6.9.5.

- (1) Suppose that \mathcal{M} is a substructure of a structure \mathcal{N} . We say that \mathcal{N} is an **ultrapower extension** of \mathcal{M} if the diagonal embedding $d : \mathcal{M} \rightarrow \mathcal{M}^{\mathcal{U}}$ can be extended to an isomorphism $d' : \mathcal{N} \rightarrow \mathcal{M}^{\mathcal{U}}$.
- (2) An **ultrapower chain over \mathcal{M}** is a chain of structures $\mathcal{M} \subseteq \mathcal{M}_1 \subseteq \mathcal{M}_2 \subseteq \dots$ such that each \mathcal{M}_{n+1} is an ultrapower extension of \mathcal{M}_n . We will refer to the limit of this chain as \mathcal{M}_∞ .

It is clear that an ultrapower extension is an elementary extension and thus an ultrapower chain is an elementary chain, whence the limit \mathcal{M}_∞ is an elementary extension of \mathcal{M} . Consequently, if \mathcal{M} and \mathcal{N} are structures for which there are ultrapower chains over \mathcal{M} and \mathcal{N} , respectively, whose limits \mathcal{M}_∞ and \mathcal{N}_∞ are isomorphic, then \mathcal{M} and \mathcal{N} are elementarily equivalent. It turns out that the converse is also true, which we will prove in Section 8.3:

Fact 6.9.6. Structures \mathcal{M} and \mathcal{N} are elementarily equivalent if and only if there are ultrapower chains over \mathcal{M} and \mathcal{N} , respectively, whose limits \mathcal{M}_∞ and \mathcal{N}_∞ are isomorphic.

Returning to iterated ultrapowers, we now suppose that $(X, <) = (\mathbb{N}, <)$. Fix a structure \mathcal{M} . For each $n \in \mathbb{N}$, fix ultrafilters \mathcal{U}_n on index sets I_n and set $\mathcal{M}_n := \mathcal{M}^{\mathcal{U}_0 \times \dots \times \mathcal{U}_n}$. Note then that $\mathcal{M} \subseteq \mathcal{M}_0 \subseteq \mathcal{M}_1 \subseteq \dots$ is an ultrapower chain over \mathcal{M} .

Exercise 6.9.7. In the notation from the paragraph preceding Fact 6.9.6, prove that the iterated ultrapower $\mathcal{M}^{\mathcal{U}_X}$ is isomorphic to the limit \mathcal{M}_∞ of the ultrapower chain from above.

Combining Fact 6.9.6 and Exercise 6.9.7, we arrive at:

Corollary 6.9.8. *Structures \mathcal{M} and \mathcal{N} are elementarily equivalent if and only if they have isomorphic iterated ultrapowers.*

The appeal of the previous corollary is that it provides a reformulation of elementary equivalence that does not mention first-order logic and only mentions the “algebraic” notion of iterated ultrapower. In Chapter 16, we will improve upon this latter fact by proving the **Keisler-Shelah theorem**, which states that structures \mathcal{M} and \mathcal{N} are elementarily equivalent if and only if they have isomorphic ultrapowers. While the statement of the Keisler-Shelah theorem is obviously more aesthetically pleasing than Corollary 6.9.8, it is much more difficult to prove.

Finally, we use Exercise 6.9.7 to show that not every iterated ultrapower is obtainable as an ordinary ultrapower, although we will need to borrow a notion and a result from Section 8.3:

Proposition 6.9.9. *For any countable structure \mathcal{M} , there is an iterated ultrapower of \mathcal{M} that is not isomorphic to any ultrapower of \mathcal{M} .*

Proof. First, we define a sequence of cardinals $(\kappa_n)_{n \in \mathbb{N}}$ by setting $\kappa_0 := \aleph_0$ and $\kappa_{n+1} := \kappa_n^{\text{cof}(\kappa_n)}$. By Proposition B.3.21, $(\kappa_n)_{n \in \mathbb{N}}$ is a strictly increasing sequence of cardinals. For each $n \in \omega$, let \mathcal{U}_n be a *regular* ultrafilter on $\text{cof}(\kappa_n)$. (We will define the notion of regular ultrafilter in Section 8.3). Recursively define a sequence $(\mathcal{M}_n)_{n \in \mathbb{N}}$ of structures by setting $\mathcal{M}_0 := \mathcal{M}$ and $\mathcal{M}_{n+1} := \mathcal{M}_n^{\mathcal{U}_n}$. By Theorem 8.3.9, $|M_{n+1}| = |M_n|^{\text{cof}(\kappa_n)}$, whence, by induction, we see that $|M_n| = \kappa_n$ for all $n \in \omega$. Note that $(\mathcal{M}_n)_{n \in \mathbb{N}}$ is an ultrapower chain over \mathcal{M} whose limit \mathcal{M}_∞ is thus an iterated ultrapower of \mathcal{M} whose cardinality has cofinality ω by construction. Theorem 6.8.5 shows that \mathcal{M}_∞ cannot be isomorphic to an ultrapower of \mathcal{M} with respect to a countably incomplete ultrafilter. An ultrapower of \mathcal{M} with respect to a countably complete ultrafilter is isomorphic to \mathcal{M} , and thus countable whence it is not isomorphic to \mathcal{M}_∞ . Thus \mathcal{M}_∞ is an iterated ultrapower of \mathcal{M} not isomorphic to any ultrapower of \mathcal{M} . \square

6.10. A category-theoretic perspective on ultraproducts

In this section, we consider a category-theoretic perspective on ultraproducts that allows us to take ultraproducts of families of objects besides first-order structures. Moreover, this perspective will allow us to consider a dual notion of the ultraproduct, naturally called the ultracoproduct, which leads to other interesting examples. We will freely use the language of category theory as discussed in Appendix C.

Let us begin by considering a simple example, namely the ultraproduct construction for groups. Consider a family of groups $(G_i)_{i \in I}$ and an ultrafilter \mathcal{U} on I . For each $J \in \mathcal{U}$, consider the group $G_J := \prod_{i \in J} G_i$, the direct product of groups. We consider the set \mathcal{U} as a directed set under reverse inclusion, that is, for $J, K \in \mathcal{U}$, we set $J \leq K$ if and only if $J \supseteq K$. Notice that \mathcal{U} is indeed a directed set, for given $J, K \in \mathcal{U}$, we have that $J \leq J \cap K$ and $K \leq J \cap K$. Given $J \leq K$ (that is, $J \supseteq K$), we have a homomorphism $f_{JK} : G_J \rightarrow G_K$ given by restriction, that is, $f_{JK}(a) := a \upharpoonright K$. It is clear that (G_J, f_{JK}) form a directed system, that is, each f_{JJ} is the identity on G_J and $f_{KL} \circ f_{JK} = f_{JL}$ whenever $J \leq K \leq L$.

We can thus consider the direct limit $G := \varinjlim G_J$ of the directed system. It will not be necessary to recall the exact construction of the direct limit (although it is given in Appendix C) but rather that it satisfies the following universal properties:

- (1) For each $J \in \mathcal{U}$, there are homomorphisms $g_J : G_J \rightarrow G$ such that $g_J = g_K \circ f_{JK}$ whenever $J \leq K$.

- (2) Whenever H is a group equipped with homomorphisms $h_J : G_J \rightarrow H$ satisfying $h_J = h_K \circ f_{JK}$, then there is a unique homomorphism $\phi : G \rightarrow H$ such that $\phi \circ g_J = h_J$ for all $J \in \mathcal{U}$.

The direct limit of groups is unique up to unique isomorphism and this isomorphism commutes with the morphisms above.

Theorem 6.10.1. *In the notation of the discussion above, we have that $\varinjlim G_J = \prod_{\mathcal{U}} G_i$.*

Proof. We need to verify that $\prod_{\mathcal{U}} G_i$ satisfies the universal property of the direct limit given above. To verify property (1), given $J \in \mathcal{U}$, we define $g_J : G_J \rightarrow \prod_{\mathcal{U}} G_i$ by setting $g_J(a) := [b]_{\mathcal{U}}$, where $b \in G_I$ is such that $f_{IJ}(b) = a$. (In other words, b is an arbitrary extension of a that is defined on all of I rather than only on the subset J .) We leave it as an exercise to check that, if $J \leq K$, then $g_J = g_K \circ f_{JK}$.

To verify property (2), suppose that H is a group equipped with morphisms $h_J : G_J \rightarrow H$ satisfying $h_J = h_K \circ f_{JK}$. Let $\phi : \prod_{\mathcal{U}} G_i \rightarrow H$ be given by $\phi([a]_{\mathcal{U}}) := h_I(a)$. To see that this map is well defined, notice that if $[a]_{\mathcal{U}} = [a']_{\mathcal{U}}$, then there is $J \in \mathcal{U}$ such that $f_{IJ}(a) = f_{IJ}(a')$, whence $h_I(a) = h_J(f_{IJ}(a)) = h_J(f_{IJ}(a')) = h_I(a')$, as desired. It is clear that ϕ is a group homomorphism. Moreover, for $a \in G_J$, we have

$$(\phi \circ g_J)(a) = \phi([a]_{\mathcal{U}}) = h_I(a) = h_J(f_{IJ}(a)) = h_J(a),$$

as desired. If ϕ' is another such function, then, in particular, $\phi' \circ g_I = h_I$, that is, $\phi'([a]_{\mathcal{U}}) = h_I(a) = \phi([a]_{\mathcal{U}})$, completing the proof. \square

The preceding discussion motivates us to generalize the notion of ultraproduct to certain categories. Suppose that \mathcal{C} is a category that has arbitrary products and direct limits. Fix a family $(A_i)_{i \in I}$ of objects from \mathcal{C} and an ultrafilter \mathcal{U} on I . For $J \in \mathcal{U}$, we set A_J to be the product (in the category-theoretic sense) of the family $(A_i)_{i \in J}$. By the universal property of product, if $J \leq K$, there is a canonical morphism $f_{JK} : A_J \rightarrow A_K$. As above, the family (A_J, f_{JK}) forms a directed system.

Definition 6.10.2. Suppose that \mathcal{C} is a category that has products and direct limits. Given a family $(A_i)_{i \in I}$ of objects from \mathcal{C} and an ultrafilter \mathcal{U} on I , we define the **ultraproduct of the family with respect to \mathcal{U}** , denoted $\prod_{\mathcal{U}} A_i$, to be the direct limit of the directed system (A_J, f_{JK}) above.

Let us say a few words as to why this categorical ultraproduct really is a generalization of our earlier model-theoretic ultraproduct. Fix a first-order language \mathcal{L} . Given two \mathcal{L} -structures \mathcal{M} and \mathcal{N} , a **homomorphism** f from \mathcal{M} to \mathcal{N} is defined exactly like an embedding from \mathcal{M} to \mathcal{N} except

that one only requires $(a_1, \dots, a_n) \in R^{\mathcal{M}}$ implies $(f(a_1), \dots, f(a_n)) \in R^{\mathcal{N}}$. In particular, homomorphisms need not be injective. We let $\mathcal{C}_{\mathcal{L}}$ denote the category of \mathcal{L} -structures, where morphisms are homomorphisms in the above sense.

Exercise 6.10.3. Suppose that $(\mathcal{M}_i)_{i \in I}$ is a directed family of L -structures. Let M be the direct limit of the sets $(M_i)_{i \in I}$. Explain how to view M as the universe of an \mathcal{L} -structure \mathcal{M} such that \mathcal{M} is the direct limit of the family $(\mathcal{M}_i)_{i \in I}$ in the category $\mathcal{C}_{\mathcal{L}}$.

Exercise 6.10.4. Suppose that $(\mathcal{M}_i)_{i \in I}$ is a family of \mathcal{L} -structures and \mathcal{U} is an ultrafilter on I . Show that the model-theoretic ultraproduct $\prod_{\mathcal{U}} \mathcal{M}_i$ coincides with the category-theoretic ultraproduct of the family $(\mathcal{M}_i)_{i \in I}$ with respect to \mathcal{U} .

Now that we have abstracted things to the category-theoretic level, there is nothing stopping us from considering the dual situation. Suppose that \mathcal{C} is a category that has coproducts and inverse limits. Fix a family $(A_i)_{i \in I}$ of objects from \mathcal{C} and an ultrafilter \mathcal{U} on I . For $J \in \mathcal{U}$, we set $A_J := \coprod_{i \in J} A_i$ to be the coproduct of the family $(A_i)_{i \in J}$. By the universal property of coproducts, if $J \leq K$, there is a morphism $f_{JK} : A_K \rightarrow A_J$. Now the family (A_J, f_{JK}) forms an inverse system.

Definition 6.10.5. Suppose that \mathcal{C} is a category that has coproducts and inverse limits. Given a family $(A_i)_{i \in I}$ of objects from \mathcal{C} and an ultrafilter \mathcal{U} on I , we define the **ultracoproduct of the family with respect to \mathcal{U}** , denoted $\prod_{\mathcal{U}} A_i$, to be the inverse limit of the inverse system (A_J, f_{JK}) above.

Exercise 6.10.6. Suppose that $F : \mathcal{C} \rightarrow \mathcal{D}$ yields an equivalence of categories. Then whenever an ultraproduct $\prod_{\mathcal{U}} A_i$ (resp., an ultracoproduct $\prod_{\mathcal{U}} A_i$) of objects of \mathcal{C} exists, we have that $F(\prod_{\mathcal{U}} A_i) = \prod_{\mathcal{U}} F(A_i)$ (resp., $F(\prod_{\mathcal{U}} A_i) = \prod_{\mathcal{U}} F(A_i)$). Similarly, if $F : \mathcal{C} \rightarrow \mathcal{D}$ yields a dual equivalence of categories, then whenever an ultraproduct $\prod_{\mathcal{U}} A_i$ (resp., an ultracoproduct $\prod_{\mathcal{U}} A_i$) of objects of \mathcal{C} exists, we have that $F(\prod_{\mathcal{U}} A_i) = \prod_{\mathcal{U}} F(A_i)$ (resp., $F(\prod_{\mathcal{U}} A_i) = \prod_{\mathcal{U}} F(A_i)$).

Example 6.10.7. Recall from Section 3.4 that the functors S and Cl yield a dual equivalence of categories between the category of Boolean algebras and the category of Stone spaces. By Exercises 6.10.4 and 6.10.6, it follows that ultracoproducts exist in the category of Stone spaces: if $(X_i)_{i \in I}$ is a family of Stone spaces and \mathcal{U} is an ultrafilter on I , then $\prod_{\mathcal{U}} X_i$ is once again a Stone space. Moreover, $S(\prod_{\mathcal{U}} Cl(X_i)) \cong \prod_{\mathcal{U}} X_i$ and $Cl(\prod_{\mathcal{U}} X_i) \cong \prod_{\mathcal{U}} Cl(X_i)$.

The category of Stone spaces is a full subcategory of the category of compact Hausdorff spaces. It is natural to wonder if this larger category has an ultracoproduct construction. This is indeed the case:

Example 6.10.8. The category of compact Hausdorff spaces is closed under coproducts and inverse limits. To see the former, recall that if $(X_i)_{i \in I}$ is a family of topological spaces, then one can consider the direct sum $\bigoplus_{i \in I} X_i$, which as a set is the disjoint union $\bigsqcup_{i \in I} X_i$, and whose topology is given by declaring $U \subseteq \bigsqcup_{i \in I} X_i$ to be open if and only if $U \cap X_i$ is an open subset of X_i for each $i \in I$. Unfortunately, if each X_i is compact, $\bigoplus_{i \in I} X_i$ need not be compact. Nevertheless, $\bigoplus_{i \in I} X_i$ is a Tychonoff space, whence one can consider its Stone-Čech compactification $\beta(\bigoplus_{i \in I} X_i)$. We leave it to the reader to check that $\prod_{i \in I} X_i = \beta(\bigoplus_{i \in I} X_i)$.

It is a standard fact that the category of compact spaces admits inverse limits. We merely outline here what the inverse limit construction is. Suppose that $(X_i)_{i \in I}$ is an inverse limit of compact Hausdorff spaces. Let X denote the inverse limit of the sets $(X_i)_{i \in I}$. One can then endow X with the smallest topology so that all projection maps $X \rightarrow X_i$ are continuous. It can then be verified that X is once again a compact Hausdorff space.

It follows that the category of compact Hausdorff spaces has ultracoproducts.

Exercise 6.10.9. Fill in the details in the previous example.

It is natural to wonder if Stone duality, the dual equivalence of categories between Stone spaces and Boolean algebras, “extends” to a dual equivalence of categories between all compact Hausdorff spaces and some other category of “algebraic” objects. This is indeed the case, and this dual equivalence of categories is given by Gelfand duality, the algebraic objects being so-called C^* -algebras. The category of C^* -algebras has a natural ultraproduct construction and the Gelfand functor takes ultraproducts of C^* -algebras to the ultracoproducts of the corresponding compact Hausdorff spaces. We will explore this in more detail in Section 14.3.

6.11. The Feferman-Vaught theorem

Although the remainder of this book is about ultraproducts, we would be remiss if we did not mention one of the more important results about arbitrary reduced products, namely the Feferman-Vaught theorem, which is the analogue of the Łoś theorem for arbitrary reduced products in that it connects truth in the reduced product with truth in the individual structures.

There are two main complications in generalizing the Łoś theorem to arbitrary reduced products. First, the truth of a formula φ in an ultraproduct depends on the truth of φ itself in the individual structures. In the case

of reduced products, we will have to relate the truth of φ in the reduced product to the truth of some *related* formulas ψ_1, \dots, ψ_m in the individual structures. It is worth noting that the ψ_i 's depend only on φ and not on the actual reduced product or parameters involved. (In some sense, one can “effectively” calculate ψ_1, \dots, ψ_m from φ .)

Secondly, while φ is true in an ultraproduct if and only if it is true in almost every structure, the analogous statement is not true in the reduced product structure, even when replacing φ by the aforementioned formulae ψ_1, \dots, ψ_m . This is essentially because the quotient Boolean algebra $\mathcal{P}(I)/\mathcal{U}$ (to be defined below) has only two elements, which are determined by whether or not a set belongs to \mathcal{U} . In the case of an arbitrary filter \mathcal{F} on I , the quotient Boolean algebra $\mathcal{P}(I)/\mathcal{F}$ is much more complicated and thus, in general, it will be some complicated Boolean algebra statement about the truths of the various ψ_i 's that will determine whether or not φ is true in the reduced product.

We now make the above discussion precise and prove the Feferman-Vaught theorem. First, suppose that we have a set I and a filter \mathcal{F} on I . We define a relation $\sim_{\mathcal{F}}$ on $\mathcal{P}(I)$ by declaring $X \sim_{\mathcal{F}} Y$ if and only if $\chi_X \sim_{\mathcal{F}} \chi_Y$, where $\chi_X : I \rightarrow \{0, 1\}$ is the characteristic function of X and similarly for Y .

Exercise 6.11.1. Prove that $X \sim_{\mathcal{F}} Y$ if and only if $I \setminus (X \Delta Y) \in \mathcal{F}$.

In other words, X and Y are “almost equal” since their symmetric difference is “small”, where “small” here means has “large complement” where “large” means belongs to \mathcal{F} .

Exercise 6.11.2. Prove that $\sim_{\mathcal{F}}$ is an equivalence relation on $\mathcal{P}(I)$.

Exercise 6.11.3. Suppose that $X_i, Y_i \in \mathcal{P}(I)$ for $i = 1, 2$ are such that $X_i \sim_{\mathcal{F}} Y_i$. Show that $X_1 \cup X_2 \sim_{\mathcal{F}} Y_1 \cup Y_2$, $X_1 \cap X_2 \sim_{\mathcal{F}} Y_1 \cap Y_2$, and $I \setminus X_1 \sim_{\mathcal{F}} I \setminus Y_1$.

We let $\mathcal{P}(I)/\mathcal{F}$ denote the set of equivalence classes and set $[X]_{\mathcal{F}}$ for the equivalence class of X . By the previous exercise, we may define $[X]_{\mathcal{F}} \wedge [Y]_{\mathcal{F}} := [X \cap Y]_{\mathcal{F}}$, and similarly for \vee and \neg .

Exercise 6.11.4. Prove that $(\mathcal{P}(I)/\mathcal{F}, \wedge, \vee, \neg, [I]_{\mathcal{F}}, [\emptyset]_{\mathcal{F}})$ is a Boolean algebra that is naturally isomorphic to $\mathbf{2}^{\mathcal{F}}$, the reduced power of the Boolean algebra $\mathbf{2}$. In particular, if \mathcal{F} is an ultrafilter, prove that $\mathcal{P}(I)/\mathcal{F}$ is isomorphic to $\mathbf{2}$.

We also let $\mathcal{L}_{\text{BA}} := \{0, 1, \wedge, \vee, \neg\}$ denote the natural first-order language for studying Boolean algebras. The axioms introduced in Section 3.4 are obviously first order; we let T_{BA} denote the \mathcal{L}_{BA} -theory axiomatizing Boolean algebras.

Definition 6.11.5. An \mathcal{L}_{BA} -formula $\gamma(y_1, \dots, y_m)$ is **monotonic** if

$$T_{\text{BA}} \models \forall y_1, \dots, \forall y_m \forall z_1 \cdots \forall z_m ((\gamma(y_1, \dots, y_m) \wedge \bigwedge_{i=1}^m y_i \leq z_i) \rightarrow \gamma(z_1, \dots, z_m)).$$

In other words, thinking of concrete Boolean algebras, if $\gamma(A_1, \dots, A_m)$ is true, where A_1, \dots, A_m are subsets of some set, and one enlarges each A_i to some superset B_i , then $\gamma(B_1, \dots, B_m)$ still holds true.

We now fix an arbitrary language \mathcal{L} . The next definition makes precise the intuition from the introduction of this section.

Definition 6.11.6. Given \mathcal{L} -formulae $\varphi(x_1, \dots, x_n)$, $\psi_1(x_1, \dots, x_n)$, \dots , $\psi_m(x_1, \dots, x_n)$ and a monotonic \mathcal{L}_{BA} -formula $\gamma(y_1, \dots, y_m)$, we say that φ is **determined by** $(\gamma; \psi_1, \dots, \psi_m)$ if, given any filter \mathcal{F} on any set I , any family $(\mathcal{M}_i)_{i \in I}$ of \mathcal{L} -structures and any $a_1, \dots, a_n \in \prod_{i \in I} M_i$, setting $X_j := \{i \in I : \mathcal{M}_i \models \psi_j(a_1(i), \dots, a_n(i))\}$, we have

$$\prod_{\mathcal{F}} \mathcal{M}_i \models \varphi([a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}}) \Leftrightarrow \mathcal{P}(I)/\mathcal{F} \models \gamma([X_1]_{\mathcal{F}}, \dots, [X_m]_{\mathcal{F}}).$$

We say that φ is **determined** if there are ψ_1, \dots, ψ_m and γ such that φ is determined by $(\gamma; \psi_1, \dots, \psi_m)$.

We are now ready for the main result of this section:

Theorem 6.11.7 (Feferman-Vaught). *Every \mathcal{L} -formula is determined.*

Proof. We proceed by induction on the complexity of the formula φ . We leave it to the reader to check that, when φ is atomic, φ is determined by $(y = 1; \varphi)$.

Now suppose that φ is determined by $(\gamma; \psi_1, \dots, \psi_m)$. Set

$$\delta(y_1, \dots, y_m) := \neg\gamma(\neg y_1, \dots, \neg y_m).$$

We show that $\neg\varphi$ is determined by $(\delta; \neg\psi_1, \dots, \neg\psi_m)$. Set

$$X_j := \{i \in I : \mathcal{M}_i \models \neg\psi_j(a_1(i), \dots, a_n(i))\}.$$

It remains to notice that the following statements are equivalent:

- $\prod_{\mathcal{F}} \mathcal{M}_i \models \neg\varphi([a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}})$,
- $\prod_{\mathcal{F}} \mathcal{M}_i \not\models \varphi([a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}})$,
- $\mathcal{P}(I)/\mathcal{F} \not\models \gamma([\neg X_1]_{\mathcal{F}}, \dots, [\neg X_m]_{\mathcal{F}})$,
- $\mathcal{P}(I)/\mathcal{F} \models \neg\gamma([\neg X_1]_{\mathcal{F}}, \dots, [\neg X_m]_{\mathcal{F}})$,
- $\mathcal{P}(I)/\mathcal{F} \models \delta([X_1]_{\mathcal{F}}, \dots, [X_m]_{\mathcal{F}})$.

Note that the equivalence between the second and third line is exactly the inductive assumption that φ is determined by $(\gamma; \psi_1, \dots, \psi_m)$.

We leave the proof of the conjunction case to the reader. We come to the final, and hardest, part of the proof, the existential case. Suppose that $\varphi(w, x_1, \dots, x_n)$ is determined by $(\gamma; \psi_1, \dots, \psi_m)$. We show that $\exists w \varphi(w, x_1, \dots, x_m)$ is determined.

Let s_1, \dots, s_{2^m} enumerate $\mathcal{P}(\{1, \dots, m\})$ with $s_i = \{i\}$ for each $i = 1, \dots, m$. For each $k \in \{1, \dots, 2^m\}$, define the formula

$$\theta_k(x_1, \dots, x_n) := \exists w \bigwedge_{j \in s_k} \psi_j(w, x_1, \dots, x_n).$$

We now set δ to be the formula

$$\exists z_1 \cdots \exists z_{2^m} \left(\bigwedge_{1 \leq k \leq 2^m} (z_k \leq y_k) \wedge \bigwedge_{s_i \cup s_j = s_k} (z_i \wedge z_j = z_k) \wedge \gamma(z_1, \dots, z_m) \right).$$

Note that δ is clearly a monotonic formula. The following claim finishes the proof of the theorem.

Claim. $\exists w \varphi$ is determined by $(\delta; \theta_1, \dots, \theta_{2^m})$.

Proof of Claim. Fix \mathcal{L} -structures $(\mathcal{M}_i)_{i \in I}$ and a filter \mathcal{F} on I . For $k = 1, \dots, 2^m$, set $Y_k := \{i \in I : \mathcal{M}_i \models \theta_k(a_1(i), \dots, a_n(i))\}$.

First suppose that $\prod_{\mathcal{F}} \mathcal{M}_i \models \exists w \varphi(w, [a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}})$. Take $[b]_{\mathcal{F}} \in \prod_{\mathcal{F}} \mathcal{M}_i$ such that $\prod_{\mathcal{F}} \mathcal{M}_i \models \varphi([b]_{\mathcal{F}}, [a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}})$. Set $Z_k := \{i \in I : \mathcal{M}_i \models \bigwedge_{j \in s_k} \psi_j(b(i), a_1(i), \dots, a_n(i))\}$. Note, in particular, that for $k = 1, \dots, m$, that $Z_k = \{i \in I : \mathcal{M}_i \models \psi_k(b(i), a_1(i), \dots, a_n(i))\}$. It is clear that $Z_k \subseteq Y_k$ for all $k = 1, \dots, 2^m$ and $Z_i \cap Z_j = Z_k$ whenever $s_i \cup s_j = s_k$. Also, since $(\gamma; \psi_1, \dots, \psi_m)$ determines φ , we have that $\mathcal{P}(I)/\mathcal{F} \models \gamma([Z_1]_{\mathcal{F}}, \dots, [Z_m]_{\mathcal{F}})$. It follows that $\mathcal{P}(I)/\mathcal{F} \models \delta([Y_1]_{\mathcal{F}}, \dots, [Y_m]_{\mathcal{F}})$.

We now suppose that $\mathcal{P}(I)/\mathcal{F} \models \delta([Y_1]_{\mathcal{F}}, \dots, [Y_m]_{\mathcal{F}})$ as witnessed by $[Z_k]_{\mathcal{F}}$ for $k = 1, \dots, 2^m$. Note that, for example, $[Z_1]_{\mathcal{F}} \leq [Y_1]_{\mathcal{F}}$ does not imply that $Z_1 \subseteq Y_1$ but merely that there is a set $X \in \mathcal{F}$ such that $Z_1 \cap X \subseteq Y_1$. However, since there are only finitely many such conditions, we can find $X \in \mathcal{F}$ such that:

- (1) $Z_k \cap X \subseteq Y_k$ for $k = 1, \dots, 2^m$, and
- (2) $Z_i \cap Z_j \cap X = Z_k \cap X$ whenever $s_i \cup s_j = s_k$.

Fix $i \in X$. Let t_i consist of those $j = 1, \dots, m$ for which $i \in Z_j$. Let l be such that $t_i = s_l$. Since $t_i = \bigcup_{j \in t_i} s_j$, (2) above tells us that $i \in Z_l$, whence by (1), $i \in Y_l$, that is, $\mathcal{M}_i \models \exists w \bigwedge_{j \in t_i} \psi_j(w, a_1(i), \dots, a_n(i))$. Fix $b(i) \in M_i$ such that $\mathcal{M}_i \models \bigwedge_{j \in t_i} \psi_j(b(i), a_1(i), \dots, a_n(i))$. For $i \notin X$, define $b(i) \in M_i$

arbitrarily. Now, for $k = 1, \dots, m$, set

$$W_k := \{i \in I : \mathcal{M}_i \models \psi_k(b(i), a_1(i), \dots, a_m(i))\}.$$

Note that $Z_k \cap X \subseteq W_k$ for each $k = 1, \dots, m$, whence $[Z_k]_{\mathcal{F}} \leq [W_k]_{\mathcal{F}}$. Since $\mathcal{P}(I)/\mathcal{F} \models \gamma([Z_1]_{\mathcal{F}}, \dots, [Z_m]_{\mathcal{F}})$ and γ is monotonic, we have $\mathcal{P}(I)/\mathcal{F} \models \gamma([W_1]_{\mathcal{F}}, \dots, [W_m]_{\mathcal{F}})$. Since $(\gamma; \psi_1, \dots, \psi_m)$ determined φ , we have that $\prod_{\mathcal{F}} \mathcal{M}_i \models \varphi([b]_{\mathcal{F}}, [a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}})$, and hence

$$\prod_{\mathcal{F}} \mathcal{M}_i \models \exists w \varphi(w, [a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}}),$$

as desired. □

Exercise 6.11.8. Verify the atomic and conjunction cases in the previous proof.

Here is a sample application of the Feferman-Vaught theorem:

Exercise 6.11.9. Suppose that $\mathcal{M}_i \equiv \mathcal{N}_i$ for all $i \in I$. Then for any filter \mathcal{F} on I , prove that $\prod_{\mathcal{F}} \mathcal{M}_i \equiv \prod_{\mathcal{F}} \mathcal{N}_i$.

6.12. Notes and references

The idea behind the ultraproduct construction goes back to Skolem's work [164] from 1934 on nonstandard models of arithmetic. In 1948, Hewitt [80] studies ultraproducts of fields. The ultraproduct construction for general first-order structures is due to Łoś [113], where he also proved what is now known as Łoś's theorem. The proof of the compactness theorem using ultraproducts is from [59]. Theorem 6.5.1 is based on a similar discussion in [90, Theorem 2.2]. Theorem 6.5.6 is from the article [85]. Much of Section 6.6 comes from the book [28]. Theorem 6.7.1 comes from Blass's thesis [16]. The discussion around the model $\mathcal{N}[f]$ comes from Keisler's article [102]. Most of Section 6.8 comes from the book [28], although the proof of Theorem 6.8.3 we believe to be our own. Our presentation of iterated ultrapowers borrows substantially from [28], although we simplify things in many respects. The category-theoretic perspective on ultraproducts seems to be well known but we struggled to find a precise reference. The notion of ultracoproducts of compact spaces seems to be discussed for the first time in [76]. Feferman and Vaught proved their theorem in [56], although our treatment follows that of [28] very closely.

Applications to geometry, commutative algebra, and number theory

In this chapter, we present three applications of the ultraproduct construction of an algebraic nature. In Section 7.1, we present Ax's theorem on polynomial functions, which is an ingenious use of ultraproducts that transfers a problem about the field of complex numbers to a problem about finite fields. Section 7.2 presents several results about bounds in the theory of polynomial rings, whose proofs are obtained by contradiction using a commutative-algebraic analysis of an ultraproduct of counterexamples. In Section 7.3, we list some number-theoretic applications of ultraproducts by discussing simple instances of a powerful result in model-theoretic algebra known as the Ax-Kochen theorem. The material in this last section is considerably more advanced and so we merely content ourselves with a presentation of the results and refer the interested reader elsewhere for full proofs.

7.1. Ax's theorem on polynomial functions

We start this chapter with a seemingly silly exercise:

Exercise 7.1.1. Suppose that X is a finite set and $f : X \rightarrow X$ is a function. Prove that f is injective if and only if f is surjective.

Why would we start the chapter this way? The answer is that, surprisingly, it is one of the key steps in the proof of the following, significantly deeper, result:

Theorem 7.1.2 (Ax's theorem). *Suppose that $f : \mathbb{C}^n \rightarrow \mathbb{C}^n$ is a polynomial function. If f is injective, then f is surjective.*

Here, when we say that f is a polynomial function, we mean that there are polynomials $P_1(\vec{X}), \dots, P_n(\vec{X}) \in \mathbb{C}[X_1, \dots, X_n]$ such that, for all $\vec{x} \in \mathbb{C}^n$, we have $f(\vec{x}) = (P_1(\vec{x}), \dots, P_n(\vec{x}))$.

How do we connect these seemingly disparate situations? Well, ultraproducts, of course! First, some intermediate steps.

Fact 7.1.3. If K_1 and K_2 are two algebraically closed fields of the same characteristic and the same *uncountable cardinality*, then $K_1 \cong K_2$.

Exercise 7.1.4. If $(K_i)_{i \in I}$ is a family of algebraically closed fields and \mathcal{U} is an ultrafilter on I , prove that $\prod_{\mathcal{U}} K_i$ is also an algebraically closed field.

For each prime p , let $\overline{\mathbb{F}_p}$ denote the algebraic closure of the finite field \mathbb{F}_p . Putting this all together, we have:

Theorem 7.1.5. *For any nonprincipal ultrafilter \mathcal{U} on the set of prime numbers, we have $\prod_{\mathcal{U}} \overline{\mathbb{F}_p} \cong \mathbb{C}$.*

Proof. By Exercise 7.1.4, we have that $\prod_{\mathcal{U}} \overline{\mathbb{F}_p}$ is an algebraically closed field. Furthermore, it must have characteristic 0, for, given any prime p , the sentence $p \cdot 1 \neq 0$ is true in all but one of the fields involved, whence by Łoś's theorem it is true in the ultraproduct. Finally, by Corollary 6.8.4, $\prod_{\mathcal{U}} \overline{\mathbb{F}_p}$ has cardinality \mathfrak{c} . Thus, by Fact 7.1.3, $\prod_{\mathcal{U}} \overline{\mathbb{F}_p} \cong \mathbb{C}$. \square

We are thus left proving Ax's theorem for the field $\prod_{\mathcal{U}} \overline{\mathbb{F}_p}$.

Exercise 7.1.6. For all $m, n \in \mathbb{N}$, prove that there is a sentence $\sigma_{m,n}$ in the language of rings so that, for any field K , we have $K \models \sigma_{m,n}$ if and only if, for any polynomial function $f : K^n \rightarrow K^n$ with degrees bounded by m , if f is injective, then f is surjective.

Recall now that $\overline{\mathbb{F}_p} = \bigcup_t \mathbb{F}_{p^t}$.

Lemma 7.1.7. *For any prime number p and any polynomial function $f : \overline{\mathbb{F}_p}^n \rightarrow \overline{\mathbb{F}_p}^n$, if f is injective, then f is surjective.*

Proof. Fix a prime p and consider an injective polynomial function $f : \overline{\mathbb{F}_p}^n \rightarrow \overline{\mathbb{F}_p}^n$. Let t be such that all of the coefficients of f belong to \mathbb{F}_{p^t} . It follows that for any $s \geq t$, we can consider the function $f_s : \mathbb{F}_{p^s}^n \rightarrow \mathbb{F}_{p^s}^n$ which is simply the restriction of f . Since f is injective, so is each f_s , whence each

f_s is also surjective by Exercise 7.1.1. However, if each f_s is surjective, it follows immediately that f is also surjective. \square

The previous lemma can be restated by saying that $\overline{\mathbb{F}_p} \models \sigma_{m,n}$ for all primes p and all m, n . By Łoś's theorem, it follows that $\prod_{\mathcal{U}} \overline{\mathbb{F}_p} \models \sigma_{m,n}$ for each m, n , whence $\mathbb{C} \models \sigma_{m,n}$ for all m, n , and thus Ax's theorem is proved.

With minimal effort, we can extend Ax's theorem to a broader class of functions. First, we say that $X \subseteq \mathbb{C}^n$ is **definable** if there is a formula $\varphi(\vec{x}, \vec{y})$ and elements $\vec{b} \in \mathbb{C}$ such that $X = \{\vec{a} \in \mathbb{C}^n : \mathbb{C} \models \varphi(\vec{a}, \vec{b})\}$. Also, given $X \subseteq \mathbb{C}^n$, we say that a function $f : X \rightarrow \mathbb{C}^n$ is **definable** if the graph of f , which is the set $\{(x, f(x)) : x \in X\}$, is a definable subset of \mathbb{C}^{2n} .

Exercise 7.1.8. If $f : X \rightarrow \mathbb{C}^n$ is a definable function, then X is a definable set.

Theorem 7.1.9 (Strong form of Ax's theorem). *Suppose that $X \subseteq \mathbb{C}^n$ is a definable set and $f : X \rightarrow X$ is an injective definable function. Then f is surjective.*

Exercise 7.1.10. Prove the strong form of Ax's theorem.

The definable sets in \mathbb{C}^n have clear geometric meaning. Indeed, define $X \subseteq \mathbb{C}^n$ to be **Zariski closed** if there are polynomials $P_1, \dots, P_m \in \mathbb{C}[\vec{X}]$ such that

$$X = \{\vec{x} \in \mathbb{C}^n : P_1(\vec{x}) = \dots = P_m(\vec{x}) = 0\}.$$

We then say that X is **constructible** if X can be obtained from Zariski closed sets by taking (finite) unions, intersections, and complements. Clearly, constructible sets are definable. It is a fact (known as the Chevalley-Tarski theorem) that, conversely, every definable set is constructible. (This holds, more generally, for any algebraically closed field.) Thus, defining a function to be constructible if its graph is constructible, we can restate the strong form of Ax's theorem in the following geometric form: if $f : X \rightarrow X$ is an injective constructible function, then f is surjective.

7.2. Bounds in the theory of polynomial rings

In this section, we will be considering the following situation: R is a commutative ring with unity, A is a $k \times l$ matrix with entries from R , and we are looking at the linear homogeneous system

$$(\star) \quad A\vec{y} = \vec{0}.$$

Associated to the system (\star) is the **solution submodule**

$$\mathcal{S}_A := \mathcal{S}_A(R) := \{\vec{r} \in R^l : A \cdot \vec{r} = \vec{0}\}.$$

We say that $\vec{r}_1, \dots, \vec{r}_p \in \mathcal{S}_A$ **generate** \mathcal{S}_A if, for every $\vec{r} \in \mathcal{S}_A$, there are $s_1, \dots, s_p \in R$ such that $\vec{r} = s_1\vec{r}_1 + \dots + s_p\vec{r}_p$. If there are finitely many $\vec{r}_1, \dots, \vec{r}_p \in \mathcal{S}_A$ that generate \mathcal{S}_A , we say that \mathcal{S}_A is **finitely generated**.

Fact 7.2.1. If R is Noetherian, then \mathcal{S}_A is finitely generated.

In this chapter, we will be interested in the case $R = K[X_1, \dots, X_n]$, where K is a field. By the Hilbert basis theorem, R is Noetherian, whence \mathcal{S}_A is finitely generated. Thus, there is clearly a bound α on the degrees of the polynomials in the generating set for \mathcal{S}_A . In theory, this bound α depends on many of the characters involved: a bound d on the degrees of the polynomials appearing in A , the dimensions k and l of the matrix A , the number n of indeterminates in $K[X_1, \dots, X_n]$, the field K itself, and the coefficients of the polynomials in the matrix A . The main theorem in this section is that, in fact, the bound α depends only on n , d , and k :

Theorem 7.2.2. *Given $n, d, k \in \mathbb{N}$, there is $\alpha = \alpha(n, d, k) \in \mathbb{N}$ such that the following holds: Suppose that K is a field and A is a $k \times l$ -matrix over $K[X_1, \dots, X_n]$ such that each polynomial in A has degree at most d . Then the solution submodule \mathcal{S}_A is generated by polynomials of degree at most α .*

Remark 7.2.3. It is quite clear that any system (\star) as above where all polynomials involved have degree at most d is equivalent to one where l equals the number of monomials $X_1^{e_1} \cdots X_n^{e_n}$ with $e_1 + \dots + e_n \leq d$. We refer to this quantity as $\lambda(n, d)$.

The key to the proof of Theorem 7.2.2 is the following algebraic concept.

Definition 7.2.4. Suppose that $R \subseteq S$ are rings. We say that S is **flat over** R if, for any system (\star) (where the entries in A come from R) and any solution $\vec{s} \in \mathcal{S}_A(S)$, there are $\vec{r}_1, \dots, \vec{r}_n \in \mathcal{S}_A(R)$ and $b_1, \dots, b_n \in S$ such that $\vec{s} = b_1\vec{r}_1 + \dots + b_n\vec{r}_n$.

In other words, S is flat over R if any solution in S to a linear system with coefficients in R is equal to an S -linear combination of solutions in R . We will need the following two standard facts about flatness. While these results are not difficult, their proofs would take us too far afield. We refer the reader to [133].

Facts 7.2.5. Suppose that $R \subseteq S$ are rings.

- (1) S is flat over R if and only if the criteria in the definition holds for systems with $k = 1$ (that is, for a single linear homogeneous equation).
- (2) If S is flat over R , then $S[X]$ is flat over $R[X]$.

For us, here are the rings R and S that are going to be relevant to the proof of Theorem 7.2.2. Fix a family $(K_t)_{t \in \mathbb{N}}$ of fields and $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. Set $K := \prod_{\mathcal{U}} K_t$ and $R := K[X_1, \dots, X_n]$. Thus, R is an ordinary polynomial ring over a field K , where the field K happens to be an ultraproduct of a family of fields. What is the ring S ? Well, we can instead consider the ordinary polynomial rings $K_t[X_1, \dots, X_n]$ over the fields K_t and then take the ultraproduct of these rings, yielding the ring $S := \prod_{\mathcal{U}} (K_t[X_1, \dots, X_n])$. There is an obvious way of viewing R as a subring of S , namely by the identification

$$\sum_{i_1, \dots, i_n} [c_{i_1, \dots, i_n}(t)]_{\mathcal{U}} X_1^{i_1} \cdots X_n^{i_n} \leftrightarrow \left[\sum_{i_1, \dots, i_n} c_{i_1, \dots, i_n}(t) X_1^{i_1} \cdots X_n^{i_n} \right]_{\mathcal{U}}.$$

Note that R is indeed a proper subring of S . For example, the polynomial $[X_1^t]_{\mathcal{U}}$ is an element of S that is not an element of R . We can view this element as a “nonstandard polynomial” (see Chapter 9) whose “degree” is the element $[\text{id}]_{\mathcal{U}} \in \mathbb{N}^{\mathcal{U}}$. Since $[\text{id}]_{\mathcal{U}} > n$ for all $n \in \mathbb{N}$, this polynomial can be thought of as having “infinite degree”.

Exercise 7.2.6. Prove that the elements of R are precisely the elements of S of finite degree.

Here is the main algebraic fact underlying the proof of Theorem 7.2.2:

Theorem 7.2.7. *Let $(K_t)_{t \in \mathbb{N}}$ be a family of fields and let $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. Then $\prod_{\mathcal{U}} (K_t[X_1, \dots, X_n])$ is flat over $(\prod_{\mathcal{U}} K_t)[X_1, \dots, X_n]$.*

First, we will need a change of variable trick. In the lemma that follows, K is an arbitrary field. We view an element of $K[X_1, \dots, X_n]$ as an element of $K[X_1, \dots, X_{n-1}][X_n]$, that is, as a polynomial in the variable X_n whose coefficients come from the ring $K[X_1, \dots, X_{n-1}]$. It thus makes sense to speak of the leading coefficient of such a polynomial. We also use the multi-index notation for polynomials in $K[X_1, \dots, X_n]$, namely a term of such a polynomial may be written as $a_j X^j$ instead of the more cumbersome notation $a_{j_1, \dots, j_n} X_1^{j_1} \cdots X_n^{j_n}$.

Lemma 7.2.8. *Given $f \in K[X_1, \dots, X_n] = K[X_1, \dots, X_{n-1}][X_n]$, there are numbers $d_1, \dots, d_{n-1} > 0$ such that, setting $Z_i := X_i - X_n^{d_i}$ for $i = 1, \dots, n-1$, $Z_n = X_n$, and*

$$f^{\#}(Z_1, \dots, Z_n) := f(Z_1 + Z_n^{d_1}, \dots, Z_{n-1} + Z_n^{d_{n-1}}, Z_n),$$

we have that $f^{\#}(Z_1, \dots, Z_n)$ has an element of K as its leading coefficient (as opposed to an element of $K[Z_1, \dots, Z_{n-1}]$).

Proof. Write $f = \sum_j a_j X^j$ with each $a_j \neq 0$. Using the change of variable appearing in the statement of the lemma, we have

$$f^\#(Z_1, \dots, Z_n) = \sum_j a_j (Z_1 + Z_n^{d_1})^{j_1} \cdots (Z_{n-1} + Z_n^{d_{n-1}})^{j_{n-1}} \cdot Z_n^{j_n}.$$

Rewriting this, we have

$$f^\#(Z_1, \dots, Z_n) = \left(\sum_j a_j Z_n^{d_1 j_1 + \cdots + d_{n-1} j_{n-1} + j_n} \right) + g(Z_1, \dots, Z_n),$$

where g has no occurrence of a monomial of the form cZ_n^k ($c \in K$). If we choose $d > 0$ large enough and set $d_i := d^{n-i}$, we leave it to the reader to verify that none of the exponents $d_1 j_1 + \cdots + d_{n-1} j_{n-1} + j_n$ are equal. Thus, the leading coefficient of $f^\#$ is a_j for some j , proving the lemma. \square

Proof of Theorem 7.2.7. We proceed by induction on n . When $n = 0$, there is nothing to prove. Now assume that $n > 0$ and the theorem is true for $n - 1$; we show that it is true for n . For simplicity, set $K := \prod_{\mathcal{U}} K_t$, $R := K[X_1, \dots, X_n]$, and $S := \prod_{\mathcal{U}} (K_t[X_1, \dots, X_n])$. By Facts 7.2.5(1), it is enough to prove the following: given $f_1, \dots, f_l \in R$ and a solution $\vec{g} = (g_1, \dots, g_l) \in S^l$ of

$$(\dagger) \quad f_1 Y_1 + \cdots + f_l Y_l = 0,$$

then \vec{g} is an S -linear combination of solutions to (\dagger) in R^l . Without loss of generality, X_n appears in f_1 . By Lemma 7.2.8, we may make a change of coordinates and thus assume, without loss of generality, that f_1 is monic in X_n , that is, when f_1 is viewed as a polynomial in $(\prod_{\mathcal{U}} K_t)[X_1, \dots, X_{n-1}][X_n]$, its leading coefficient is 1. (We leave it to the reader to verify that this change of coordinates is harmless for the task at hand.)

Set $d := \deg_{X_n} f_1$. For $i = 2, 3, \dots, l$, let $\hat{f}_i \in R^l$ be the vector with $-f_i$ in the first component, f_1 in the i th component, and zeroes elsewhere, e.g., $\hat{f}_2 = (-f_2, f_1, 0, 0, \dots, 0)$. Note that each \hat{f}_i is a solution to the equation (\dagger) in R^l . By choosing $h_2, \dots, h_l \in S$ appropriately, we see that $\vec{g}' := \vec{g} - h_2 \hat{f}_2 - \cdots - h_l \hat{f}_l$ is a solution to the equation (\dagger) with $\deg_{X_n}(g'_2), \dots, \deg_{X_n}(g'_n) < d$. Since $f_1 g'_1 + \cdots + f_l g'_l = 0$, it follows that $\deg_{X_n}(g'_1)$ is also finite. In other words, each component of \vec{g}' is an ordinary polynomial in the variable X_n with coefficients in the ring $\prod_{\mathcal{U}} (K_t[X_1, \dots, X_{n-1}])$. By the inductive hypothesis, $\prod_{\mathcal{U}} (K_t[X_1, \dots, X_{n-1}])$ is flat over $K[X_1, \dots, X_{n-1}]$, whence, by Facts 7.2.5(2),

$$\left(\prod_{\mathcal{U}} (K_t[X_1, \dots, X_{n-1}]) \right) [X_n]$$

is flat over

$$K[X_1, \dots, X_{n-1}][X_n] = K[X_1, \dots, X_n].$$

It follows that \vec{g}' is a $\prod_{\mathcal{U}}(K_t[X_1, \dots, X_{n-1}])[X_n]$ -linear combination of solutions to equation (\dagger) in R ; in particular, \vec{g}' is an S -linear combination of solutions to equation (\dagger) in R^l . Since $\vec{g} = \vec{g}' + h_2\hat{f}_2 + \dots + h_l\hat{f}_l$, we have that \vec{g} is also an S -linear combination of solutions to equation (\dagger) in R^l , as desired. \square

Proof of Theorem 7.2.2. By Remark 7.2.3, we may assume that, in equation (\star) , we have that $l = \lambda(n, d)$. Suppose that Theorem 7.2.2 is false for a given n, d, k . For each t , let K_t be a field, and let $A(t)$ be a $k \times l$ -matrix with entries from $K_t[X_1, \dots, X_n]$ with degrees bounded by d such that (\star) has a solution $\vec{s}(t)$ from $K_t[X_1, \dots, X_n]$ that is not a linear combination of solutions of degree bounded by t . Let K, R , and S be as defined in the proof of Theorem 7.2.7. Let A be the matrix over S whose entries are the ultraproducts of the entries from $A(t)$, that is, $A_{ij} = [A(t)_{ij}]_{\mathcal{U}}$. Since all entries of each $A(t)$ have degree bounded by d , it follows that A is actually a matrix over R . Now $[\vec{s}(t)]_{\mathcal{U}}$ is a solution of $A \cdot \vec{y} = \vec{0}$ in S , whence, by Theorem 7.2.7, there are solutions $\vec{r}_1, \dots, \vec{r}_m \in R$ of $A \cdot \vec{y} = \vec{0}$ and $[s_1]_{\mathcal{U}}, \dots, [s_m]_{\mathcal{U}} \in S$ such that $[\vec{s}]_{\mathcal{U}} = [s_1]_{\mathcal{U}}\vec{r}_1 + \dots + [s_m]_{\mathcal{U}}\vec{r}_m$. Let $q \in \mathbb{N}$ be an upper bound for the degrees of $\vec{r}_1, \dots, \vec{r}_m$. For each t and $i = 1, \dots, m$, let $\vec{r}_i(t) \in K_t[X_1, \dots, X_n]$ be such that $\vec{r}_i = [\vec{r}_i(t)]_{\mathcal{U}}$. Since \mathcal{U} is nonprincipal, there is $t > q$ such that $\vec{r}_1(t), \dots, \vec{r}_m(t)$ are solutions of $A(t) \cdot \vec{y} = \vec{0}$ and $\vec{s}(t) = s_1(t)\vec{r}_1(t) + \dots + s_m(t)\vec{r}_m(t)$, contradicting the choice of $\vec{s}(t)$. \square

There is also something to be said about nonhomogeneous equations. This time, we consider the equation

$$(\dagger\dagger) \quad A \cdot \vec{y} = \vec{f},$$

where each entry from A and each entry from \vec{f} come from the ring R . The relevant algebraic notion is the following:

Definition 7.2.9. If $R \subseteq S$ are rings, we say that S is **faithfully flat** over R if S is flat over R and every system $(\dagger\dagger)$ with a solution in S^l also has a solution in R^l .

The proof of the next theorem is more difficult than the proof of Theorem 7.2.7, and we refer the reader to [180] for a proof:

Theorem 7.2.10. *Let $(K_t)_{t \in \mathbb{N}}$ be a family of fields and let $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. Then $\prod_{\mathcal{U}}(K_t[X_1, \dots, X_n])$ is faithfully flat over $(\prod_{\mathcal{U}} K_t)[X_1, \dots, X_n]$.*

Corollary 7.2.11. *Given $n, d, k \in \mathbb{N}$, there is $\beta = \beta(n, d, k) \in \mathbb{N}$ such that the following holds: whenever K is a field and $(\dagger\dagger)$ is such that all entries from A and \vec{f} come from $K[X_1, \dots, X_n]$ and have degree at most d , then if $(\dagger\dagger)$ has a solution, it has a solution with all entries of degree at most β .*

Exercise 7.2.12. Prove Corollary 7.2.11 from Theorem 7.2.10.

When $k = 1$, we can reformulate the conclusion of the previous corollary in terms of ideals.

Corollary 7.2.13. *Given $n, d \in \mathbb{N}$, there is $\gamma = \gamma(n, d)$ such that the following holds: whenever K is a field and $f, f_1, \dots, f_l \in K[X_1, \dots, X_n]$ all have degree at most d , then $f \in (f_1, \dots, f_l)$ if and only if there are $h_1, \dots, h_l \in K[X_1, \dots, X_n]$ of degree at most γ with $f = h_1 f_1 + \dots + h_l f_l$.*

Exercise 7.2.14. Suppose that K is a field. Then, for any $f_0, f_1, \dots, f_m \in \mathbb{Z}[C, X]$, we have $\{c \in K^n : f_0(c, X) \in (f_1(c, X), \dots, f_m(c, X))\}$ is definable. In particular, if K is algebraically closed, then this set is constructible.

There are many other results of the kind described in this section. However, the algebraic arguments needed are beyond the scope of this book and for that we reason we have chosen to only describe the aforementioned results in detail. Still, one of the results from [180] is compelling enough to state here, without proof.

Theorem 7.2.15. *Let $(K_t)_{t \in \mathbb{N}}$ be a family of fields and let $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. Fix also $f_1, \dots, f_m \in (\prod_{\mathcal{U}} K_t)[X_1, \dots, X_n]$. Then f_1, \dots, f_m generate a prime ideal of $(\prod_{\mathcal{U}} K_t)[X_1, \dots, X_n]$ if and only if they generate a prime ideal of $\prod_{\mathcal{U}} (K_t[X_1, \dots, X_n])$.*

Exercise 7.2.16. Use the previous theorem to prove the following “bounds” result: given $n, d \in \mathbb{N}$, there is $\delta = \delta(n, d) \in \mathbb{N}$ such that the following holds: if K is a field and $f_1, \dots, f_m \in K[X_1, \dots, X_n]$ all have degree at most d and $gh \in (f_1, \dots, f_m)$ implies $g \in (f_1, \dots, f_m)$ or $h \in (f_1, \dots, f_m)$ for all g, h of degree at most δ , then (f_1, \dots, f_m) is a prime ideal or all of $K[X_1, \dots, X_n]$.

7.3. The Ax-Kochen theorem and Artin’s conjecture

Let us begin this section by comparing the rings $\mathbb{Z}/(p^2\mathbb{Z})$ and $\mathbb{F}_p[T]/(T^2)$. In the former ring, we may think of elements as being of the form $a_0 + a_1 p$, where $a_0, a_1 \in \{0, 1, \dots, p-1\}$ (as the map which sends such an element to its coset modulo $p^2\mathbb{Z}$ is a bijection). Likewise, elements of the latter ring may be viewed as polynomials $b_0 + b_1 T \in \mathbb{F}_p[T]$ (so $b_0, b_1 \in \mathbb{F}_p$). Consequently, both rings have p^2 elements. There are other ways in which these rings are similar (they are both henselian local rings with maximal ideal generated by an element whose square is 0; these terms will be defined shortly). On the other hand, these rings are different in other ways, as, for example, the latter ring contains a subfield (namely \mathbb{F}_p) while the former ring does not!

Nevertheless, one can capture the intuition that the rings $\mathbb{Z}/(p^2\mathbb{Z})$ and $\mathbb{F}_p[T]/(T^2)$ are similar, and that this similarity actually “improves” as p gets larger. Moreover, this result holds for any n instead of just for $n = 2$:

Theorem 7.3.1. *For any $n \geq 1$ and any nonprincipal ultrafilter \mathcal{U} on the set of primes, we have*

$$\prod_{\mathcal{U}} \mathbb{Z}/(p^n \mathbb{Z}) \cong \prod_{\mathcal{U}} \mathbb{F}_p[T]/(T^n).$$

Corollary 7.3.2. *For any sentence σ in the language of rings, we have, for all but finitely many primes p , that*

$$\mathbb{Z}/(p^n \mathbb{Z}) \models \sigma \Leftrightarrow \mathbb{F}_p[T]/(T^n) \models \sigma.$$

Proof. Suppose that the corollary was false. Then there would be an infinite set A of primes such that, without loss of generality, for all $p \in A$, we have $\mathbb{Z}/(p^n \mathbb{Z}) \models \sigma$ while $\mathbb{F}_p[T]/(T^n) \models \neg\sigma$. Let \mathcal{U} be a nonprincipal ultrafilter on the set of primes such that $A \in \mathcal{U}$. We then have $\prod_{\mathcal{U}} \mathbb{Z}/(p^n \mathbb{Z}) \models \sigma$ while $\prod_{\mathcal{U}} \mathbb{F}_p[T]/(T^n) \models \neg\sigma$, contradicting Theorem 7.3.1. \square

We will discuss some further ramifications of Theorem 7.3.1 in a moment, but let us pause briefly to discuss some of the ingredients behind its proof. The key idea is that of a local ring. For simplicity, in this section, when we use the word “ring” we always mean “commutative ring with unity.”

Definition 7.3.3. A ring R is a **local ring** if it has a unique maximal ideal.

If R is a local ring, we usually denote its unique maximal ideal by \mathfrak{m} and we let $\mathfrak{k} := R/\mathfrak{m}$; \mathfrak{k} is referred to as the **residue field** of the local ring R .

Exercise 7.3.4.

- (1) Prove that $\mathbb{Z}/(p^n \mathbb{Z})$ is a local ring whose unique maximal ideal is the ideal generated by the coset of p and whose residue field is isomorphic to \mathbb{F}_p .
- (2) Prove that $\mathbb{F}_p[T]/(T^n)$ is a local ring whose unique maximal ideal is the ideal generated by the coset of T and whose residue field is isomorphic to \mathbb{F}_p .

Contrast the previous exercise with the fact that neither \mathbb{Z} nor $F[T]$ (F any field) are local rings.

Exercise 7.3.5. Suppose that R is a ring.

- (1) Prove that R is a local ring if and only if the set of noninvertible elements forms an ideal of R .
- (2) Prove that there is a sentence σ_{local} in the language of rings such that $R \models \sigma_{\text{local}}$ if and only if R is a local ring.
- (3) Prove that there is a formula $\varphi_{\text{max}}(x)$ in the language of rings such that, for all $R \models \sigma_{\text{local}}$, $\varphi_{\text{max}}(R)$ is the maximal ideal of R .

- (4) Suppose that $(R_i)_{i \in I}$ is a family of local rings with maximal ideals \mathfrak{m}_i and residue fields \mathbf{k}_i . Fix an ultrafilter \mathcal{U} on I and set $R := \prod_{\mathcal{U}} R_i$. Show that R is a local ring. Moreover, denoting the maximal ideal and residue field for R by \mathfrak{m} and \mathbf{k} , show that $\mathfrak{m} \cong \prod_{\mathcal{U}} \mathfrak{m}_i$ and $\mathbf{k} \cong \prod_{\mathcal{U}} \mathbf{k}_i$.

The following result is the key to proving Theorem 7.3.1 above. Its proof would take us too far afield so we refer the reader to the excellent treatment in [179].

Theorem 7.3.6. *Suppose that R is a local ring and that $t \in R$ and $n \in \mathbb{N}$ are such that:*

- (1) $\text{char}(\mathbf{k}) = 0$;
- (2) $\mathfrak{m} = tR$;
- (3) $t^n \neq 0$ but $t^{n+1} = 0$.

Then $R \cong \mathbf{k}[T]/(T^{n+1})$ by an isomorphism that sends t to the coset of T .

Proof of Theorem 7.3.1. Let

$$R := \prod_{\mathcal{U}} \mathbb{Z}/(p^n \mathbb{Z}) \quad \text{and} \quad S := \prod_{\mathcal{U}} \mathbb{F}_p[T]/(T^n).$$

By Exercise 7.3.5, R and S are local rings. For each p , let $t_R(p)$ denote the generator of the maximal ideal of $\mathbb{Z}/(p^n \mathbb{Z})$. Likewise, let $t_S(p)$ denote the generator of the maximal ideal of $\mathbb{F}_p[T]/(T^n)$. Then by Exercise 7.3.5, letting \mathfrak{m}_R and \mathfrak{m}_S denote the maximal ideals of R and S , respectively, we have $\mathfrak{m}_R = [t_R]_{\mathcal{U}} R$ and $\mathfrak{m}_S = [t_S]_{\mathcal{U}} S$. Note also that Loś's theorem implies that $[t_R]_{\mathcal{U}}^n, [t_S]_{\mathcal{U}}^n \neq 0$ but $[t_R]_{\mathcal{U}}^{n+1} = [t_S]_{\mathcal{U}}^{n+1} = 0$. Finally, Exercise 7.3.5 implies that $\mathbf{k}_R \cong \mathbf{k}_S \cong \prod_{\mathcal{U}} \mathbb{F}_p$, a field \mathbf{k} of characteristic 0. Thus, by Theorem 7.3.6, $R \cong \mathbf{k}[T]/(T^{n+1}) \cong S$, as desired. \square

Remark 7.3.7. Note that the proof of the previous theorem shows that the isomorphism between $\prod_{\mathcal{U}} \mathbb{Z}/(p^n \mathbb{Z})$ and $\prod_{\mathcal{U}} \mathbb{F}_p[T]/(T^n)$ maps $[t_R]_{\mathcal{U}}$ to $[t_S]_{\mathcal{U}}$. Consequently, in Corollary 7.3.2, we could even allow sentences in an extension of the language of rings by adding a constant, and where this constant is interpreted in both $\mathbb{Z}/(p^n \mathbb{Z})$ and $\mathbb{F}_p[T]/(T^n)$ as the generator of the maximal ideal.

We now consider what happens if we let $n \rightarrow \infty$ in the above discussion. What exactly does that mean? Well, in the case of $\mathbb{F}_p[T]/(T^n)$, letting $n \rightarrow \infty$ means we should consider the power series ring $\mathbb{F}_p[[T]]$ defined as the collection of all formal power series $\sum_{i=0}^{\infty} a_i T^i$ with each $a_i \in \mathbb{F}_p$. One adds and multiplies such formal power series in the same way as with ordinary polynomials. Of course this construction makes sense over any field, not just \mathbb{F}_p .

Exercise 7.3.8. For any field \mathbf{k} , show that $R := \mathbf{k}[[T]]$ is a local ring with unique maximal ideal $\mathfrak{m} = T\mathbf{k}[[T]]$ and residue field \mathbf{k} .

What does it mean to let $n \rightarrow \infty$ in the case of $\mathbb{Z}/(p^n\mathbb{Z})$? Naïvely speaking, we should consider formal power series $\sum_{i=0}^{\infty} a_i p^i$, with each $a_i \in \{0, 1, \dots, p-1\}$. To make this rigorous, one should think of such an infinite sum as the limit as $n \rightarrow \infty$ of the partial sums $\sum_{i=0}^n a_i p^i$. But in what topology is this limit taken? Well, we should think of tails $\sum_{i=n+1}^m a_i p^i$ as being negligible for large $m > n$. Thus, large powers of p should be small, leading us to consider the norm $|\cdot|_p$ on \mathbb{Z} given by $|a|_p := p^{-v_p(a)}$, where $v_p(a) :=$ the largest n such that $p^n \mid a$. We then define a metric d_p on \mathbb{Z} given by $d_p(a, b) := |a - b|_p$. One then lets \mathbb{Z}_p denote the completion of \mathbb{Z} as a metric space. It follows that the sequence $(\sum_{i=0}^n a_i p^i)_{n \in \mathbb{N}}$ converges to an element of \mathbb{Z}_p , which we write as $\sum_{i=0}^{\infty} a_i p^i$. One can check that the ring operations on \mathbb{Z} extend to ring operations on \mathbb{Z}_p . \mathbb{Z}_p is referred to as the **ring of p -adic integers**.

Fact 7.3.9. \mathbb{Z}_p is a local ring with unique maximal ideal $p\mathbb{Z}_p$ and residue field \mathbb{F}_p .

We need to introduce one further notion:

Definition 7.3.10. Suppose that R is a local ring. We say that R is **henselian** if, whenever $f(X) \in R[X]$ and $a \in R$ is such that $f(a) \in \mathfrak{m}$ but $f'(a) \notin \mathfrak{m}$, then there is $x \in R$ with $f(x) = 0$ and $x - a \in \mathfrak{m}$.

Although we will not get too much into the details, here is how one should think of the henselian property. One thinks of elements of \mathfrak{m} as “small” or “infinitesimal” in some sense. In this case, $f(a) \in \mathfrak{m}$ means that $f(a)$ is close to 0. The definition says that as long as the slope of the tangent line is not infinitesimal, then a is really close to an actual root of f .

One can also think of the henselian property in terms of the residue field. Letting \bar{a} denote the image of a in \mathbf{k} and $\bar{f}(X) \in \mathbf{k}[X]$ as the result of applying the quotient map to all coefficients of f , the hypothesis is that $\bar{f}(\bar{a}) = 0$ and $\bar{f}'(\bar{a}) \neq 0$. In other words, one has a *nonsingular zero* of \bar{f} in $\mathbf{k}[X]$ and the henselian property says that the resulting root of \bar{f} can be *lifted* to an actual root of f .

Exercise 7.3.11. Show that there is a theory T_{hens} in the language of rings such that, for any local ring R , we have $R \models T_{\text{hens}}$ if and only if R is henselian.

Example 7.3.12.

- (1) The rings $\mathbb{Z}/(p^n\mathbb{Z})$ and $\mathbb{F}_p/(T^n)$ are henselian.

- (2) Any local ring satisfying the assumptions of Theorem 7.3.6 above is henselian.
- (3) \mathbb{Z}_p is henselian.
- (4) For any field \mathbf{k} , $\mathbf{k}[[T_1, \dots, T_n]]$ is a henselian local ring.

A crucial fact used in the proof of Theorem 7.3.6 above is the following:

Theorem 7.3.13. *Suppose that R is a henselian local ring with $\text{char } \mathbf{k} = 0$. Then there is a subfield E of R (that is, a subring of R that happens to be a field) such that the residue map maps E isomorphically onto \mathbf{k} .*

Consequently, we have:

Corollary 7.3.14. *Suppose that R is a henselian local ring with $\text{char } \mathbf{k} = 0$ and $f_1, \dots, f_k \in \mathbb{Z}[X]$, where $X = (X_1, \dots, X_n)$. Then any solution of $f_1(X) = \dots = f_k(X) = 0$ in \mathbf{k} can be lifted to a solution in R .*

A subfield E as in the previous theorem is called a *lift* of \mathbf{k} . Putting together all of the previous information, we can now prove the following theorem, which answered a question of Lang from the 1950s:

Theorem 7.3.15 (Greenleaf-Ax-Kochen). *Fix $f_1, \dots, f_k \in \mathbb{Z}[X]$, where $X = (X_1, \dots, X_n)$. Then for all but finitely many primes p , we have that every solution of $f_1(X) = \dots = f_k(X) = 0$ in \mathbb{F}_p can be lifted to a solution in \mathbb{Z}_p .*

Proof. Suppose, toward a contradiction, that the theorem is false. Then there is an infinite set A of primes such that, for each $p \in A$, there is a solution $a(p) \in \mathbb{F}_p$ of the system that cannot be lifted to a solution in \mathbb{Z}_p . Let \mathcal{U} be a nonprincipal ultrafilter on the set of primes such that $A \in \mathcal{U}$. Consider $R := \prod_{\mathcal{U}} \mathbb{Z}_p$, which is a henselian local ring whose residue field \mathbf{k} is isomorphic to $\prod_{\mathcal{U}} \mathbb{F}_p$, which has characteristic 0. Then $[a]_{\mathcal{U}}$ is a solution in \mathbf{k} to the system, whence, by Corollary 7.3.14, there is $[b]_{\mathcal{U}} \in R$ that is a solution to the system as well. By Łoś's theorem, we have that, for \mathcal{U} -many primes p , the residue of $b(p)$ is $a(p)$ and $b(p)$ is a solution to the system. In particular, there is some $p \in A$ for which this is true, contradicting the definition of A . \square

One can use the Greenleaf-Ax-Kochen theorem to prove one further nice result. But first, we need:

Fact 7.3.16 (Chevalley-Warning). *Suppose that q is a power of a prime p and $f_1, \dots, f_k \in \mathbb{F}_q[X] \setminus \{0\}$, where $X = (X_1, \dots, X_n)$, are such that $\sum \deg f_i < n$. Then p divides $|\{x \in \mathbb{F}_q^n : f_1(x) = \dots = f_k(x) = 0\}|$. In particular, if each f_i has constant term 0, then there is at least one*

such solution (namely the zero solution), whence there are at least p many solutions, and hence one nonzero solution.

Corollary 7.3.17. *Suppose that $f_1, \dots, f_k \in \mathbb{Z}[X] \setminus \{0\}$ all have constant term 0. Then for all but finitely many primes p , there is a nonzero solution in \mathbb{Z}_p^n to the system.*

It turns out that Theorem 7.3.1 above also holds true “in the limit”:

Theorem 7.3.18. *For any nonprincipal ultrafilter \mathcal{U} on the set of primes, we have*

$$\prod_{\mathcal{U}} \mathbb{Z}_p \cong \prod_{\mathcal{U}} \mathbb{F}_p[[T]].$$

Unlike Theorem 7.3.1, the proof of the previous theorem is much more difficult. The previous theorem is actually a special case of a more general theorem known as the **Ax-Kochen-Ershov** theorem, a fundamental result in model-theoretic algebra.

As above, we have the following:

Corollary 7.3.19. *For any sentence σ in the language of rings, for all but finitely many primes p , we have*

$$\mathbb{Z}_p \models \sigma \Leftrightarrow \mathbb{F}_p[[T]] \models \sigma.$$

The previous corollary has a spectacular application. First, we set \mathbb{Q}_p to be the fraction field of \mathbb{Z}_p .

Theorem 7.3.20. *If $f(X_1, \dots, X_5) \in \mathbb{Q}_p[X_1, \dots, X_5]$ is homogeneous of degree 2, then f has a zero in \mathbb{Q}_p .*

Motivated by the previous result, Artin conjectured the following generalization:

Conjecture 7.3.21 (Artin). *If $f(X_1, \dots, X_n) \in \mathbb{Z}_p[X_1, \dots, X_n]$ is homogeneous of degree d with $n = d^2 + 1$, then f has a zero in \mathbb{Z}_p .*

The previous conjecture was later proven to be true for $d = 3$ [37] and [114], but false in general for $d = 4$ [175]. However, the conjecture does hold if we replace \mathbb{Z}_p by $\mathbb{F}_p[[T]]$:

Theorem 7.3.22 (Lang). *If $f(X_1, \dots, X_n) \in (\mathbb{F}_p[[T]])[X_1, \dots, X_n]$ is homogeneous of degree d with $n = d^2 + 1$, then f has a zero in $\mathbb{F}_p[[T]]$.*

By Corollary 7.3.19, we have a positive solution to an asymptotic version of Artin's conjecture:

Corollary 7.3.23. *Fix $d \in \mathbb{N}$ and set $n := d^2 + 1$. Then for all but finitely many primes p , we have the following: for any $f(X_1, \dots, X_n) \in \mathbb{Z}_p[X_1, \dots, X_n]$ that is homogeneous of degree d , f has a zero in \mathbb{Z}_p .*

7.4. Notes and references

Ax's theorem was originally proven in [3]. Grothendieck [75] independently proved this theorem using techniques from algebraic geometry, whence the theorem is sometimes known as the Ax-Grothendieck theorem. The material presented in Section 7.2 is taken directly from [180] except that we use the language of ultraproducts rather than the language of nonstandard analysis (see Chapter 9). Our treatment of the Ax-Kochen theorem and Artin's conjecture is essentially an excerpt from the incredible lecture notes [179], which include full proofs and further historical context.

A book devoted entirely to the use of ultraproducts in commutative algebra is Schouten's book [153].

Ultraproducts and saturation

One of the main benefits of the ultraproduct construction is that it often yields structures that are very “rich” in a sense that is made precise using the model-theoretic notion of saturation. In Section 8.1, we introduce this notion and the related notion of universality. In Section 8.2, we show that ultraproducts with respect to countably incomplete ultrafilters are always countably saturated, which is the weakest nontrivial level of saturation. Section 8.3 describes the class of regular ultrafilters; ultraproducts with respect to regular ultrafilters are always universal. In order to obtain ultraproducts that are fairly saturated, one must consider the class of good ultrafilters, which is the subject of Sections 8.4 and 8.5. Finally, Section 8.6 briefly discusses Keisler’s order, which is a measure of relative complexity of first-order theories based on the relative level of saturation of ultrapowers of their models.

8.1. Saturation

In order to motivate the notion of saturation, we first mention the following consequence of the compactness theorem.

Proposition 8.1.1. *Suppose that \mathcal{M} is an \mathcal{L} -structure and $A \subseteq M$ is a set of parameters. Suppose that $\Sigma(x)$ is a set of \mathcal{L}_A -formulas in the free variables x (x being some finite tuple of variables) that is finitely satisfiable in \mathcal{M}_A , that is, for every finite subset $\Delta(x)$ of $\Sigma(x)$, there is $b \in M$ such that $\mathcal{M}_A \models \Delta(b)$. Then there is an elementary extension \mathcal{N} of \mathcal{M} such that $\Sigma(x)$ is satisfied in \mathcal{N}_A , meaning that there is $b \in N$ such that $\mathcal{N}_A \models \Sigma(b)$.*

Proof. Let $\mathcal{L}' := \mathcal{L}_M \cup \{c\}$, where c is a tuple of new constant symbols of the same length as x . Let $\Sigma(c)$ be the set of \mathcal{L}' -sentences obtained by replacing each occurrence of the tuple x by the new tuple c . We consider the set Σ' of \mathcal{L}' -sentences given by

$$\Sigma' := \Sigma(c) \cup \{\varphi : \varphi \text{ is an } \mathcal{L}_M\text{-sentence such that } \mathcal{M}_M \models \varphi\}.$$

Since Σ is finitely satisfiable in \mathcal{M}_A , we see that any finite subset of Σ' is modeled by an appropriate expansion of \mathcal{M}_M to an \mathcal{L}' -structure. By the compactness theorem, Σ' has a model \mathcal{N}' . Setting \mathcal{N} to be the \mathcal{L} -reduct of \mathcal{N}' yields the desired elementary extension of \mathcal{M} . \square

In Proposition 8.1.1, we think of $\Sigma(x)$ being finitely satisfiable in \mathcal{M}_A as saying that Σ is describing some properties of an element that *could* exist in \mathcal{M} . The compactness argument allows us to conclude that the element does indeed exist, albeit in an elementary extension. Often, passing to an elementary extension is necessary:

Example 8.1.2. Suppose that $\mathcal{M} = (\mathbb{N}, <)$ and let $\Sigma := \{n < x : n \in \mathbb{N}\}$. Then Σ is finitely satisfiable in \mathcal{M} but obviously not satisfiable in \mathcal{M} .

The point of saturated structures is that, as long as the number of parameters used in the description is small, then any element that could exist in \mathcal{M} does in fact exist in \mathcal{M} :

Definition 8.1.3. Suppose that \mathcal{M} is an \mathcal{L} -structure and κ is an infinite cardinal. We say that \mathcal{M} is κ -**saturated** if, whenever $A \subseteq M$ is a set of parameters with $|A| < \kappa$ and $\Sigma(x)$ is a set of \mathcal{L}_A -formulae that is finitely satisfiable in \mathcal{M}_A , then $\Sigma(x)$ is satisfiable in \mathcal{M}_A .

It is clear that any κ -saturated structure is also λ -saturated for $\lambda < \kappa$.

Exercise 8.1.4. Suppose that \mathcal{M} is κ -saturated. Prove that \mathcal{M}_A is also κ -saturated for any $A \subseteq M$ with $|A| < \kappa$.

Exercise 8.1.5. Prove that \mathcal{M} is κ -saturated if and only if it is κ -saturated for sets consisting of formulae with just one free variable.

Exercise 8.1.6. Suppose that \mathcal{M} is κ -saturated. Prove that $|M| \geq \kappa$.

Consequently, the most saturated that a structure \mathcal{M} can be is $|M|$ -saturated. In this case, we simply say that \mathcal{M} is **saturated**.

Example 8.1.7. Any uncountable algebraically closed field is saturated. To see this, suppose that K is an uncountable algebraically closed field and that Σ is a set of \mathcal{L}_A -formulae in the free variable x , where $A \subseteq K$ is such that $|A| < |K|$ (and \mathcal{L} is the language of rings). By the Chevalley-Tarski theorem referred to in Section 7.1, we may as well assume that Σ consists of

either equations $f(x) = 0$ or inequations $f(x) \neq 0$, where $f(x) \in k[x]$ and k is the subfield of K generated by A . If there is any equation in Σ , then since K is algebraically closed, all of the roots of that equation belong to K and, by finite satisfiability, one of those roots satisfies Σ . If all of the elements of Σ are inequations, then each inequation asks that x not be one of the finitely many roots of the polynomial in question. Since there are at most $|k| = \max(|A|, \aleph_0) < |K|$ many such polynomials, we can find an element of K satisfying all of the inequations, whence satisfying Σ .

Note, however, that not every countable algebraically closed fields is saturated. For example, if K is the algebraic closure of the prime field, then K is not saturated as it does not satisfy the finitely satisfiable set of sentences asking that x be transcendental. (The set is finitely satisfiable as algebraically closed fields are infinite.) In fact, of all of the countable algebraically closed fields of a given characteristic, only the one of transcendence degree \aleph_0 over the prime field is saturated.

Exercise 8.1.8. Prove the last statement in Example 8.1.7.

The following important theorem will be used several times throughout this book.

Theorem 8.1.9. *Suppose that \mathcal{M} and \mathcal{N} are saturated structures, $\mathcal{M} \equiv \mathcal{N}$, and $|M| = |N|$. Then $\mathcal{M} \cong \mathcal{N}$.*

Proof. Set $\kappa := |M| = |N|$. Enumerate $M = (a_\alpha)_{\alpha < \kappa}$ and $N := (b_\alpha)_{\alpha < \kappa}$. For $\alpha < \kappa$, we build sets $A_\alpha \subseteq M$ and $B_\alpha \subseteq N$ and bijections $f_\alpha : A_\alpha \rightarrow B_\alpha$ satisfying:

- (1) $f_\alpha \subseteq f_{\alpha+1}$;
- (2) $|A_\alpha| < \kappa$ and $|B_\alpha| < \kappa$;
- (3) $a_\alpha \in A_{\alpha+1}$ and $b_\alpha \in B_{\alpha+1}$;
- (4) f_α is *partial elementary*: for any formula $\varphi(x)$ and any tuple $a \in A_\alpha$, we have $\mathcal{M} \models \varphi(a)$ if and only if $\mathcal{N} \models \varphi(f_\alpha(a))$.

If we then let $f := \bigcup_{\alpha < \kappa} f_\alpha$, we get that $f : M \rightarrow N$ is an isomorphism. To begin, when $\alpha = 0$, we let f_0 to be the empty function. Note that, in this case, (4) follows from the fact that \mathcal{M} and \mathcal{N} are elementarily equivalent. When α is a limit ordinal, we just take $f_\alpha := \bigcup_{\beta < \alpha} f_\beta$, which clearly satisfies (2) and (4).

Supposing now that f_α has been constructed, we now show how to construct $f_{\alpha+1}$. We actually build $f_{\alpha+1}$ in two steps, first by ensuring that a_α is in the domain of $f_{\alpha+1}$ and then ensuring that b_α is in the range of $f_{\alpha+1}$. To accomplish the first step, we are searching for $c \in N$ so that, roughly speaking, the relationship between a_α and A_α is the same as the relationship

between c and B_α . More formally, we are looking to find $c \in N$ that belongs to $\bigcap_{\varphi \in X_\alpha} B_\varphi$, where,

$$X_\alpha := \{\varphi(x, e) : \varphi(x, e) \text{ is a } \mathcal{L}_{A_\alpha}\text{-formula such that } \mathcal{M} \models \varphi(a_\alpha, e)\}$$

and

$$B_\varphi := \{c \in N : \mathcal{N} \models \varphi(x, f_\alpha(e))\}.$$

Note that if $c \in \bigcap_{\varphi} B_\varphi$, then extending f_α to $g : A_\alpha \cup \{a_\alpha\} \rightarrow N$ by defining $g(a_\alpha) := c$ is a partial elementary map extending f_α and which a_α is in its domain. Since $|A_\alpha| < \kappa$ and \mathcal{N} is κ -saturated, it suffices to show that, for any finite $X \subseteq X_\varphi$, $\bigcap_{\varphi \in X} B_\varphi$ is satisfiable in \mathcal{N} . To see this, note that $\mathcal{M} \models \exists x \bigwedge_{\varphi \in X} \varphi(x, e)$, whence $\mathcal{N} \models \exists x \bigwedge_{\varphi \in X} \varphi(x, f_\alpha(e))$, which follows from the fact that f_α is partial elementary.

One proceeds in a similar fashion to extend g to a partial elementary map $f_{\alpha+1}$ which has b_α in its range. We leave the details to the reader. \square

We will also need to consider the following related notion:

Definition 8.1.10. Suppose that κ is an infinite cardinal. An \mathcal{L} -structure \mathcal{M} is κ -**universal** if, for every \mathcal{L} -structure \mathcal{N} with $|N| < \kappa$ and $\mathcal{M} \equiv \mathcal{N}$, we have that \mathcal{N} embeds elementarily into \mathcal{M} .

Once again, κ -universal clearly implies λ -universal for $\lambda < \kappa$. It is also clear that if \mathcal{M} is κ -universal, then $\kappa \leq |M|^+$. Consequently, we refer to a $|M|^+$ -universal structure simply as **universal**.

Universality is a weakening of saturation:

Proposition 8.1.11. *If \mathcal{M} is κ -saturated, then \mathcal{M} is κ^+ -universal.*

Exercise 8.1.12. Prove Proposition 8.1.11. (*Hint.* Proceed as in the proof of Theorem 8.1.9 but without needing the “two-step” procedure.)

There is a partial converse to Proposition 8.1.11:

Proposition 8.1.13. *Suppose that $|\mathcal{L}| \leq \kappa$. Then the \mathcal{L} -structure \mathcal{M} is κ -saturated if and only if, for every subset $A \subseteq M$ with $|A| < \kappa$, \mathcal{M}_A is κ^+ -universal.*

Exercise 8.1.14. Prove Proposition 8.1.13.

The following reformulation of universality will prove useful in the next section. Given a language \mathcal{L} and a cardinal κ , let $\mathcal{L}(\kappa)$ denote the extension of \mathcal{L} obtained by adding κ new constant symbols c_α , $\alpha < \kappa$. For $a \in M^\kappa$, we let $(\mathcal{M}; a)$ denote the expansion of \mathcal{M} to an $\mathcal{L}(\kappa)$ -structure obtained by interpreting c_α as $a(\alpha)$.

Proposition 8.1.15. *Suppose that $|\mathcal{L}| \leq \kappa$. For an \mathcal{L} -structure \mathcal{M} , the following are equivalent:*

- (1) \mathcal{M} is κ^+ -universal
- (2) For any set Σ of $\mathcal{L}(\kappa)$ -sentences, if $\text{Th}(\mathcal{M}) \cup \Sigma$ is satisfiable, then Σ is satisfiable in $(\mathcal{M}; a)$ for some $a \in M^\kappa$.

Proof. (1) implies (2): First suppose that \mathcal{M} is κ^+ -universal and Σ is as in (2). Let $(\mathcal{N}; b) \models \text{Th}(\mathcal{M}) \cup \Sigma$. By the Downward Löwenheim-Skolem theorem, we may assume that $|N| \leq \kappa$. Since $\mathcal{N} \equiv \mathcal{M}$ and \mathcal{M} is κ^+ -universal, there is an elementary embedding $j : \mathcal{N} \rightarrow \mathcal{M}$. It follows that $(\mathcal{M}; j(b)) \models \Sigma$.

Now suppose that (2) holds and take $\mathcal{N} \equiv \mathcal{M}$ with $|N| \leq \kappa$. Let $\Sigma := \text{Th}(\mathcal{N}; b)$, where $b \in N^\kappa$ is an enumeration of N . Then Σ is as in (2), whence there is a such that $(\mathcal{M}; a) \models \Sigma$. The map which sends b to a yields the desired elementary embedding of \mathcal{N} into \mathcal{M} . \square

8.2. First saturation properties of ultraproducts

In this short section, we show how ultraproducts (usually) yield somewhat saturated structures. More specifically:

Theorem 8.2.1. *Suppose that \mathcal{U} is a countably incomplete ultrafilter on a set I , \mathcal{L} is a countable language, and $(\mathcal{M}_i)_{i \in I}$ is a family of \mathcal{L} -structures. Then $\prod_{\mathcal{U}} \mathcal{M}_i$ is \aleph_1 -saturated.*

Proof. Set $\mathcal{N} := \prod_{\mathcal{U}} \mathcal{M}_i$ and suppose that $A \subseteq N$ is countable. Suppose also that $\Sigma(x)$ is a set of \mathcal{L}_A -formulae in the variable x that is finitely satisfiable in \mathcal{N}_A . We must show that Σ is satisfiable in \mathcal{N}_A .

Without loss of generality, we may assume that $\Sigma = \{\varphi_n(x, [a_n]_{\mathcal{U}}) : n \in \mathbb{N}\}$ and that $\mathcal{N} \models \forall x(\varphi_{n+1}(x, [a_{n+1}]_{\mathcal{U}}) \rightarrow \varphi_n(x, [a_n]_{\mathcal{U}}))$. Since \mathcal{U} is countably incomplete, we may fix sets $I_n \in \mathcal{U}$ such that $I_n \supseteq I_{n+1}$ and $\bigcap_{n \in \mathbb{N}} I_n = \emptyset$. Set $J_n := \{i \in I_n : \mathcal{M}_i \models \exists x \varphi_n(x, a_n(i))\}$; since Σ is finitely satisfiable in \mathcal{N}_A , we have that each $J_n \in \mathcal{U}$. By assumption, $J_n \supseteq J_{n+1}$ and clearly $\bigcap_{n \in \mathbb{N}} J_n = \emptyset$.

For each $i \in I$, let $n(i)$ be the maximal n such that $i \in J_n$, and let $b(i) \in M_i$ be such that $\mathcal{M}_i \models \varphi_n(b(i), a_n(i))$. We claim that $[b]_{\mathcal{U}}$ satisfies Σ . To see this, fix $n \in \mathbb{N}$; we show that $\mathcal{N} \models \varphi_n([b]_{\mathcal{U}}, [a_n]_{\mathcal{U}})$. To see this, take $i \in J_n$ and note that $n \leq n(i)$. Consequently, $\mathcal{M}_i \models \varphi_{n(i)}(b(i), a_{n(i)}(i))$ and thus $\mathcal{M}_i \models \varphi_n(b(i), a_n(i))$. Since $J_n \in \mathcal{U}$, we have that $\mathcal{N} \models \varphi_n([b]_{\mathcal{U}}, [a_n]_{\mathcal{U}})$, as desired. \square

Remark 8.2.2. One cannot drop the countably incomplete assumption in the previous theorem, for if \mathcal{U} is a countably complete ultrafilter, then for

any countable structure \mathcal{M} , $\mathcal{M}^{\mathcal{U}}$ is isomorphic to \mathcal{M} , whence it need not be \aleph_1 -saturated.

What if one desires ultraproducts that are more than just \aleph_1 -saturated? In general, given $\kappa > \aleph_1$, one cannot hope to prove, in ZFC, that all ultraproducts with respect to countably incomplete ultrafilters are κ -saturated. Indeed, if \mathcal{U} is a nonprincipal ultrafilter on \mathbb{N} and \mathcal{M} is a countable structure, then $|M^{\mathcal{U}}| = \mathfrak{c}$ by Corollary 6.8.4. In a model of ZFC plus CH, $|M^{\mathcal{U}}| = \aleph_1$. By Exercise 8.1.6, one cannot have that $\mathcal{M}^{\mathcal{U}}$ is κ -saturated for any $\kappa > \aleph_1$.

In Section 8.4, we will introduce so-called “good” ultrafilters that guarantee that ultraproducts with respect to them have higher levels of saturation; at the same time, we will also consider the case of uncountable languages.

Before moving on, we mention one nice application of the \aleph_1 -saturation of ultraproducts. First, we need a definition:

Definition 8.2.3. Suppose that \mathcal{M} and \mathcal{N} are \mathcal{L} -structures with $\mathcal{M} \subseteq \mathcal{N}$. We say that \mathcal{M} is **existentially closed (e.c.)** in \mathcal{N} if, for any quantifier-free \mathcal{L}_M -formula $\varphi(x)$, we have

$$\mathcal{M}_M \models \exists x \varphi(x) \Leftrightarrow \mathcal{N}_M \models \exists x \varphi(x).$$

One thinks of existential closedness as a model-theoretic generalization of (relative) algebraic closedness. Ultrapowers provide a nice semantic reformulation of this notion:

Exercise 8.2.4. Suppose that \mathcal{M} and \mathcal{N} are countable structures in the countable language \mathcal{L} with $\mathcal{M} \subseteq \mathcal{N}$. Prove that \mathcal{M} is e.c. in \mathcal{N} if and only if, for some (equivalently any) nonprincipal ultrafilter \mathcal{U} on \mathbb{N} , the diagonal embedding $d : \mathcal{M} \rightarrow \mathcal{M}^{\mathcal{U}}$ extends to an embedding $j : \mathcal{N} \rightarrow \mathcal{M}^{\mathcal{U}}$.

We can drop the countability assumptions in the previous exercise once we encounter the aforementioned good ultrafilters that yield more saturation.

8.3. Regular ultrafilters

Before moving on to the class of good ultrafilters that yield ultraproducts with higher levels of saturation, it behooves us to first consider the class of regular ultrafilters. While the definition of good ultrafilters is somewhat complicated and their existence is quite cumbersome to prove, the class of regular ultrafilters are quite easy to define and their existence is a simple exercise. We will see that ultrapowers (not arbitrary ultraproducts) with

respect to regular ultrafilters are highly universal (in the sense of the previous section). Ultrapowers with respect to regular ultrafilters possess other desirable properties that will also be mentioned.

We begin this section with an exercise that motivates what is to come:

Exercise 8.3.1. Suppose that \mathcal{U} is an ultrafilter on an index set I . Prove that \mathcal{U} is countably incomplete if and only if there is $E \subseteq \mathcal{U}$, $|E| = \aleph_0$, such that each element of I belongs to only finitely many elements of E .

We strengthen the notion of countably incomplete ultrafilter by demanding that there exists a set E as in Exercise 8.3.1 that has cardinality larger than \aleph_0 :

Definition 8.3.2. If κ is a cardinal, then an ultrafilter \mathcal{U} on I is κ -**regular** if there is $E \subseteq \mathcal{U}$, $|E| = \kappa$, such that each element of I belongs to only finitely many elements of E . We call such a set E a κ -**regularizing set for** \mathcal{U} .

Thus, the previous exercise can be rephrased as:

Lemma 8.3.3. *An ultrafilter is countably incomplete if and only if it is \aleph_0 -regular.*

Exercise 8.3.4. Suppose that \mathcal{U} is a κ -regular ultrafilter on an index set I . Prove the following:

- (1) \mathcal{U} is nonprincipal.
- (2) If $\lambda < \kappa$, then \mathcal{U} is also λ -regular.
- (3) $\kappa \leq |I|$.

Definition 8.3.5. An ultrafilter \mathcal{U} on an index set I is called **regular** if it is $|I|$ -regular.

By the previous exercise, regular ultrafilters are “maximally” regular. The following reformulation of κ -regularity is often useful:

Exercise 8.3.6. Suppose that \mathcal{U} is an ultrafilter on I and κ is an infinite cardinal. Prove that \mathcal{U} is κ -regular if and only if there is $f : I \rightarrow \mathcal{P}_f(\kappa)$ such that, for each $\alpha < \kappa$, we have $\{i \in I : \alpha \in f(i)\} \in \mathcal{U}$.

Exercise 8.3.7. If \mathcal{U} and \mathcal{V} are ultrafilters with $\mathcal{U} \leq_{RK} \mathcal{V}$ and \mathcal{U} is κ -regular, prove that \mathcal{V} is also κ -regular. In particular, if \mathcal{U} and \mathcal{V} are ultrafilters and either \mathcal{U} or \mathcal{V} are κ -regular (maybe both), then so is $\mathcal{U} \times \mathcal{V}$.

We now show that regular ultrafilters exist:

Proposition 8.3.8. *For each cardinal κ , there is a regular ultrafilter \mathcal{U} on κ .*

Proof. It suffices to find a regular ultrafilter on some set of size κ ; we will use $I := \mathcal{P}_f(\kappa)$. For $\alpha < \kappa$, let $\hat{\alpha} := \{u \in I : \alpha \in u\}$. Let $E := \{\hat{\alpha} : \alpha < \kappa\}$. Clearly, $|E| = \kappa$. Notice that E has the finite intersection property: $\{\alpha_1, \dots, \alpha_n\} \in \bigcap_{i=1}^n \hat{\alpha}_i$. Let \mathcal{U} be an ultrafilter on I containing E . Then E is a κ -regularizing set for \mathcal{U} : $u \in \hat{\alpha}$ if and only if $\alpha \in u$. \square

Ultrapowers with respect to regular ultrafilters have predictable cardinalities:

Theorem 8.3.9. *Suppose that \mathcal{U} is a regular ultrafilter on κ and M is infinite. Then $|M^{\mathcal{U}}| = |M|^\kappa$.*

Proof. It suffices to prove that $|M|^\kappa \leq |M^{\mathcal{U}}|$ (the other inequality always being true). Let $E \subseteq \mathcal{U}$ be a κ -regularizing set for \mathcal{U} . Let $N := M^{<\omega}$. It suffices to find an injection $\rho : M^E \rightarrow N^{\mathcal{U}}$. For each $g \in M^E$, define $g' : I \rightarrow N$ by defining $g'(i) := (g(A_1), \dots, g(A_n))$, where A_1, \dots, A_n enumerate all of the elements of E containing i (with respect to some fixed ordering of E). We now define $\rho(g) := [g']_{\mathcal{U}}$. It remains to show that ρ is injective. Suppose that $g \neq h$ and take A such that $g(A) \neq h(A)$. Suppose that $i \in A$. Then it is clear that $g'(i) \neq h'(i)$. Since $A \in \mathcal{U}$, we have that $g' \not\equiv_{\mathcal{U}} h'$, that is, $\rho(g) \neq \rho(h)$. \square

We now state the connection between regularity and universality:

Theorem 8.3.10. *Suppose that \mathcal{U} is an ultrafilter on an index set I . Then \mathcal{U} is κ -regular if and only if, whenever $|\mathcal{L}| \leq \kappa$ and \mathcal{M} is an \mathcal{L} -structure, we have that $\mathcal{M}^{\mathcal{U}}$ is κ^+ -universal.*

Proof. First suppose that \mathcal{U} is κ -regular, $|\mathcal{L}| \leq \kappa$, and \mathcal{M} is an \mathcal{L} -structure. Set $\mathcal{N} := \mathcal{M}^{\mathcal{U}}$. We use the criteria established in Proposition 8.1.15 to show that \mathcal{N} is κ^+ -universal. Let Σ be a set of $\mathcal{L}(\kappa)$ -sentences of size at most κ such that $\text{Th}(\mathcal{N}) \cup \Sigma$ is satisfiable. Let $E \subseteq \mathcal{U}$ be a κ -regularizing set for \mathcal{U} . Let $h : \Sigma \rightarrow E$ be any injection. Let $\Sigma(i) := \{\sigma \in \Sigma : i \in h(\sigma)\}$. By assumption, $\Sigma(i)$ is finite for each $i \in I$. Since $\Sigma(i)$ is finite and consistent with $\text{Th}(\mathcal{M})$, there is $a(i) \in \mathcal{M}^\kappa$ such that $\Sigma(i)$ is realized in $(\mathcal{M}; a(i))$. (Note that only finitely many constants are mentioned, so most of the choice of $a(i)$ is irrelevant.) We set $a \in N^\kappa$ to be given by $a(\alpha) := [i \mapsto a(i)(\alpha)]_{\mathcal{U}}$. We claim that Σ is realized in $(\mathcal{N}; a)$. Fix $\sigma \in \Sigma$. Then $h(\sigma) \in \mathcal{U}$ and for $i \in h(\sigma)$, we have that $\sigma \in \Sigma(i)$, so $(\mathcal{M}; a(i)) \models \sigma$, as desired.

We now prove the converse. Consider the language \mathcal{L} with one binary relation R and constant symbols d_i for $i < \kappa$. Let \mathcal{M} be the \mathcal{L} -structure whose universe is $\mathcal{P}_f(\kappa)$, that interprets R as \subseteq , and interprets d_i as $\{i\}$. By assumption, $\mathcal{M}^{\mathcal{U}}$ is κ^+ -universal. Let $\Sigma := \{R(d_i, c_0) : i < \kappa\}$, a set of $\mathcal{L}(\kappa)$ -sentences (that happens to only mention one of the new constants, namely

c_0). It is clear that $\text{Th}(\mathcal{M}) \cup \Sigma$ is finitely satisfiable in a suitable expansion of \mathcal{M} , whence the compactness theorem implies that $\text{Th}(\mathcal{M}) \cup \Sigma$ is satisfiable. Since $\mathcal{M}^{\mathcal{U}}$ is κ^+ -universal, there is $b \in \mathcal{M}^{\mathcal{U}}$ such that Σ is satisfied in $(\mathcal{M}^{\mathcal{U}}; [b]_{\mathcal{U}})$. (Here, we slightly change notation and let $[b]_{\mathcal{U}}$ denote a single element of $M^{\mathcal{U}}$, meant to be the interpretation of c_0 ; the interpretations of the other constants are irrelevant.) For each $i < \kappa$, consider the set $A_i := \{j \in I : \mathcal{M} \models R(\{i\}, b(j))\}$, an element of \mathcal{U} , and let $E := \{A_i : i < \kappa\}$. Note that $|E| = \kappa$. It is clear that a given $j \in I$ can only be contained in finitely many A_i , so E is a κ -regularizing set for \mathcal{U} . \square

The following corollary is immediate:

Corollary 8.3.11 (Frayne). *\mathcal{L} -structures \mathcal{M} and \mathcal{N} are elementarily equivalent if and only if \mathcal{M} is elementarily embeddable in an ultrapower of \mathcal{N} .*

We can use the preceding corollary to establish a fact mentioned in Section 6.9; we refer the reader to the discussion there for the definitions of the relevant terminology:

Corollary 8.3.12. *\mathcal{L} -structures \mathcal{M} and \mathcal{N} are elementarily equivalent if and only if there are ultrapower chains $(\mathcal{M}_n)_{n \in \mathbb{N}}$ and $(\mathcal{N}_n)_{n \in \mathbb{N}}$ over \mathcal{M} and \mathcal{N} , respectively, whose limits \mathcal{M}_{∞} and \mathcal{N}_{∞} are isomorphic.*

Proof sketch. The backward direction is immediate from the fact that \mathcal{M} and \mathcal{N} are elementary substructures of \mathcal{M}_{∞} and \mathcal{N}_{∞} , respectively. We now sketch a proof of the forward direction. Suppose that \mathcal{M} and \mathcal{N} are elementarily equivalent. By Corollary 8.3.11, there is an ultrapower \mathcal{N}'_1 of \mathcal{N} such that \mathcal{M} is elementarily embeddable in \mathcal{N}'_1 . Let $f_1 : \mathcal{M} \rightarrow \mathcal{N}'_1$ be such an elementary embedding. Let \mathcal{L}' denote the extension of \mathcal{L} by adding constants for all elements of M and let \mathcal{M}' and \mathcal{N}'_1 denote the expansions of \mathcal{M} and \mathcal{N}'_1 to \mathcal{L}' -structures where, for $a \in M$, $c_a^{\mathcal{M}'} = a$ and $c_a^{\mathcal{N}'_1} = f_1(a)$. Since f_1 is elementary, we have that $\mathcal{M}' \equiv \mathcal{N}'_1$. Using Corollary 8.3.11 again, there is an ultrapower \mathcal{M}'_1 of \mathcal{M}' such that \mathcal{N}'_1 elementarily embeds into \mathcal{M}'_1 ; let $g_1 : \mathcal{N}'_1 \rightarrow \mathcal{M}'_1$ denote such an elementary embedding. We let \mathcal{M}_1 denote the reduct of \mathcal{M}'_1 to \mathcal{L} , which is clearly an ultrapower extension of \mathcal{M} . By design, $g_1 \circ f_1 : M \rightarrow M_1$ is the inclusion of M into M_1 . By the same argument, we can construct an ultrapower extension \mathcal{N}_2 of \mathcal{N}_1 and an elementary embedding $f_2 : \mathcal{M}_1 \rightarrow \mathcal{N}_2$ such that $f_2 \circ g_1 : N_1 \rightarrow N_2$ is just the inclusion of N_1 into N_2 . This also implies that f_2 extends f_1 . Continuing back and forth in this manner, we get the desired ultrapower chains and maps $f_n : \mathcal{M}_n \rightarrow \mathcal{N}_{n+1}$ and $g_n : \mathcal{N}_n \rightarrow \mathcal{M}_n$ that continue to extend each other. Thus, setting $f := \bigcup_{n \in \mathbb{N}} f_n$ and $g := \bigcup_{n \in \mathbb{N}} g_n$, we have that f and g are inverse isomorphisms between \mathcal{M}_{∞} and \mathcal{N}_{∞} . \square

Exercise 8.3.13. Verify all of the details in the previous proof.

Here are two more applications of Theorem 8.3.10. First, given an \mathcal{L} -structure \mathcal{M} , define the **universal theory** of \mathcal{M} , denoted $\text{Th}_{\forall}(\mathcal{M})$, to be the set of universal \mathcal{L} -sentences σ such that $\mathcal{M} \models \sigma$.

Exercise 8.3.14. Given \mathcal{L} -structures \mathcal{M} and \mathcal{N} , prove that the following are equivalent:

- (1) $\mathcal{M} \models \text{Th}_{\forall}(\mathcal{N})$.
- (2) For every κ -regular ultrafilter \mathcal{U} , where $\kappa = \max(|M|, |\mathcal{L}|)$, one has that \mathcal{M} embeds into $\mathcal{N}^{\mathcal{U}}$.
- (3) \mathcal{M} embeds into some ultrapower $\mathcal{N}^{\mathcal{U}}$ of \mathcal{N} .

Exercise 8.3.15. Remove the cardinality restrictions in Exercise 8.2.4. More precisely, suppose that \mathcal{L} is any language and \mathcal{M} and \mathcal{N} are \mathcal{L} -structures with $\mathcal{M} \subseteq \mathcal{N}$. Prove that \mathcal{M} is e.c. in \mathcal{N} if and only if, for some ultrafilter \mathcal{U} , the diagonal embedding $d : \mathcal{M} \rightarrow \mathcal{M}^{\mathcal{U}}$ extends to an embedding $j : \mathcal{N} \rightarrow \mathcal{M}^{\mathcal{U}}$.

The following is one of the fundamental facts about regular ultrafilters. We will make use of it once in the next section, but it will be the driving force behind the discussion of Keisler's order in Section 8.6.

Theorem 8.3.16 (Keisler). *Suppose that \mathcal{U} is a κ -regular ultrafilter on the set I . Suppose that $|\mathcal{L}| \leq \kappa$ and \mathcal{M} and \mathcal{N} are \mathcal{L} -structures with $\mathcal{M} \equiv \mathcal{N}$. Then $\mathcal{M}^{\mathcal{U}}$ is κ^+ -saturated if and only if $\mathcal{N}^{\mathcal{U}}$ is κ^+ -saturated.*

Proof. Suppose that $\mathcal{M}^{\mathcal{U}}$ is κ^+ -saturated and let $E \subseteq \mathcal{U}$ be a κ -regularizing subset for \mathcal{U} . Suppose that $\Sigma(x)$ is a set of $\mathcal{L}(\kappa)$ -formulae that is finitely satisfiable in $(\mathcal{N}^{\mathcal{U}}; b)$. We show that $\Sigma(x)$ is satisfiable in $(\mathcal{N}^{\mathcal{U}}; b)$. As in the proof of Theorem 8.3.10, let $h : \Sigma \rightarrow E$ be an injection and set $\Sigma(i) := \{\sigma \in \Sigma : i \in h(\sigma)\}$. Once again, $\Sigma(i)$ is finite for each $i \in I$. Let $\Gamma(i)$ be the finite set of sentences of the form $\exists x \bigwedge_{\sigma \in u} \sigma$, where u is a nonempty subset of $\Sigma(i)$. Since $\mathcal{M} \equiv \mathcal{N}$, there is $a(i) \in M^{\kappa}$ such that

$$(\ddagger) \quad \text{Th}(\mathcal{M}, a(i)) \cap \Gamma(i) = \text{Th}(\mathcal{N}, b(i)) \cap \Gamma(i).$$

We now define $a \in (M^{\mathcal{U}})^{\kappa}$ by $a(\alpha) := [i \mapsto a(i)(\alpha)]_{\mathcal{U}}$. We claim that $\Sigma(x)$ is finitely satisfiable in $(\mathcal{M}^{\mathcal{U}}; a)$. Fix $\sigma_1, \dots, \sigma_n \in \Sigma$ and let $\tau := \exists x \bigwedge_{j=1}^n \sigma_j$. If $i \in \bigcap_{j=1}^n h(\sigma_j)$, then $\tau \in \Gamma(i)$. Since Σ is finitely satisfiable in $(\mathcal{N}^{\mathcal{U}}; b)$, for \mathcal{U} -almost all i , we have that $(\mathcal{N}; b(i)) \models \tau$. It follows that for \mathcal{U} -almost all i , $(\mathcal{M}; a(i)) \models \tau$, whence $(\mathcal{M}^{\mathcal{U}}; a) \models \tau$, as desired.

Since $\mathcal{M}^{\mathcal{U}}$ is κ^+ -saturated, Σ is satisfied in $(\mathcal{M}^{\mathcal{U}}; a)$, say by an element $[c]_{\mathcal{U}}$. For each $i \in I$, let $\Psi(i) := \{\sigma \in \Sigma(i) : (\mathcal{M}; a(i)) \models \sigma(c(i))\}$. It follows by (\ddagger) that there is $d(i) \in N$ such that $(\mathcal{N}; b(i)) \models \sigma(d(i))$ for each $\sigma \in \Psi(i)$.

We conclude by proving that that $[d]_{\mathcal{U}}$ realizes Σ in $(\mathcal{N}^{\mathcal{U}}; b)$. Fix $\sigma \in \Sigma$. For $i \in h(\sigma)$, we have that $\sigma \in \Sigma(i)$. Since $(\mathcal{M}^{\mathcal{U}}; a) \models \sigma(c)$, we have that $(\mathcal{M}; a(i)) \models \sigma(c(i))$ for \mathcal{U} -almost all i . It follows that, for \mathcal{U} -almost all i , we have $\sigma \in \Psi(i)$. For these i , $(\mathcal{N}; b(i)) \models \sigma(d(i))$. It follows that $(\mathcal{N}^{\mathcal{U}}; b) \models \sigma([d]_{\mathcal{U}})$, as desired. \square

Let us return to the statement of Theorem 8.3.10 for another moment. Suppose that \mathcal{U} is a regular ultrafilter on κ , $|\mathcal{L}| \leq \kappa$, and \mathcal{M} is an \mathcal{L} -structure with $|M| \leq \kappa$. By Theorem 8.3.9, $|M^{\mathcal{U}}| = |M|^{\kappa} \geq \kappa^+$. Consequently, in theory, $\mathcal{M}^{\mathcal{U}}$ could even be κ^{++} -universal, not just merely κ^+ -universal as guaranteed by Theorem 8.3.10. (Note that κ^{++} -universality is the most that one could hope to prove in ZFC in this situation, for if the GCH at κ holds, we have that $|M^{\mathcal{U}}| = \kappa^+$.) The κ^+ -good ultrafilters that we will encounter in the next section yield κ^+ -saturated, and hence κ^{++} -universal, ultrapowers. However, we will soon see that there are κ -regular ultrafilters that are not κ^+ -good.

Let $\text{Univ}(\kappa, \mathcal{U})$ denote the following statement: \mathcal{U} is a regular ultrafilter on κ and for every language \mathcal{L} with $|\mathcal{L}| \leq \kappa$ and every \mathcal{L} -structure \mathcal{M} , $\mathcal{M}^{\mathcal{U}}$ is κ^{++} -universal.

The question of whether or not $\text{Univ}(\kappa, \mathcal{U})$ holds for all κ and \mathcal{U} was posed in [28], and only recently have some results been proven along these lines. We summarize some of them here:

- Assuming GCH, $\text{Univ}(\kappa, \mathcal{U})$ holds for every regular cardinal κ and every regular ultrafilter \mathcal{U} on κ .
- In ZFC, it is shown that, for every regular cardinal κ , there are regular ultrafilters \mathcal{U} on κ for which $\text{Univ}(\kappa, \mathcal{U})$ holds and yet \mathcal{U} is not good.
- Assuming large cardinals, there is a singular strong limit cardinal κ and a regular ultrafilter \mathcal{U} on κ such that $\text{Univ}(\kappa, \mathcal{U})$ fails. (It is known that a large cardinal assumption is necessary in this case.) Consequently, one cannot prove in ZFC that $\text{Univ}(\kappa, \mathcal{U})$ holds for all κ and \mathcal{U} .

All of the above results are from the article [103]. In the same article, it is conjectured that it is consistent that there is a regular cardinal κ and a regular ultrafilter \mathcal{U} on κ for which $\text{Univ}(\kappa, \mathcal{U})$ fails.

We end this section with one more set-theoretical remark. In [45], Donder proved that it is consistent with ZFC that every uniform ultrafilter on an infinite set is regular. On the other hand, if \mathcal{U} is a countably complete, uniform ultrafilter on some set I and \mathcal{V} is a nonprincipal ultrafilter on \mathbb{N} ,

then $\mathcal{U} \times \mathcal{V}$ is a uniform ultrafilter that is not regular. Thus, whether or not all uniform ultrafilters on infinite sets are regular is independent of ZFC.

8.4. Good ultrafilters: Part 1

We are now ready to consider the ultrafilters for which all ultraproducts are κ^+ -saturated (given that the language is of size at most κ). Rather than just jump right in to the definition, we prefer to let the notion arise organically.

Until further notice, we fix the following data:

- an index set I ;
- an ultrafilter \mathcal{U} on I ;
- a language \mathcal{L} with $|\mathcal{L}| \leq \kappa$;
- a family of \mathcal{L} -structures $(\mathcal{M}_i)_{i \in I}$ whose ultraproduct we denote by $\mathcal{N} := \prod_{\mathcal{U}} \mathcal{M}_i$;
- a set Σ of \mathcal{L}_A -formulae, where $A \subseteq N$ is such that $|A| \leq \kappa$, for which Σ is finitely satisfiable in \mathcal{N}_A .

We want to see what it would take to have that Σ is satisfiable in \mathcal{N} .

In what follows, we simplify notation by hiding any mention of parameters in the formulae from Σ . Thus, when referring to a formula $\varphi(x)$ from Σ , really $\varphi(x)$ is $\varphi(x, a)$ for some parameters a from \mathcal{N} . Likewise, when we write $\mathcal{M}_i \models \exists x \varphi(x)$, we really mean $\mathcal{M}_i \models \exists x \varphi(x, a(i))$.

We begin naïvely by considering a *concern* function $C : \Sigma \rightarrow \mathcal{P}(I)$, where we view the fact that $i \in C(\varphi)$ as telling us that \mathcal{M}_i should be “concerned” about satisfying $\exists x \varphi(x)$. The goal is to find $b \in \mathcal{N}$ that realizes Σ by asking that $b(i)$ realizes $\varphi(x)$ whenever $i \in C(\varphi)$. Of course, for this strategy to work, we must have:

- for each $\varphi \in \Sigma$, $C(\varphi) \in \mathcal{U}$.

If \mathcal{M}_i cannot satisfy $\varphi(x)$, then \mathcal{M}_i should not be concerned about φ , so we also require:

- $C(\varphi) \subseteq \{i \in I : \mathcal{M}_i \models \exists x \varphi(x)\}$.

Now suppose that $i \in C(\varphi) \cap C(\psi)$. There is no reason a priori why $\mathcal{M}_i \models \exists x(\varphi(x) \wedge \psi(x))$. But it would be great if it did! So let us try to add this as a requirement. However, $\varphi \wedge \psi$ may not be an element of Σ . There is an easy fix for this: we simply extend the concern function now to $C : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{P}(I)$ and view $i \in C(u)$ as saying that \mathcal{M}_i should be concerned with satisfying $\bigwedge_{\varphi \in u} \varphi(x)$. Our earlier requirements now become:

- for each $u \in \mathcal{P}_f(\Sigma)$, $C(u) \in \mathcal{U}$;
- for each $u \in \mathcal{P}_f(\Sigma)$, $C(u) \subseteq \{i \in I : \mathcal{M}_i \models \exists x \bigwedge_{\varphi \in u} \varphi(x)\}$.

There is a nice reformulation of the second requirement if we make two important definitions. First, given $f, g : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$, write $f \leq g$ if $f(u) \subseteq g(u)$ for all $u \in \mathcal{P}_f(\Sigma)$; in this case, we say that f **refines** g .

Second, define the **Łoś map** to be the map $\text{Łoś} : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$ given by

$$\text{Łoś}(u) := \{i \in I : \mathcal{M}_i \models \exists x \bigwedge_{\varphi \in u} \varphi(x)\}.$$

We see now that the second condition merely asks that our concern function C refines the Łoś map.

Next note that $u \subseteq v$ implies $C(v) \subseteq C(u)$; we refer to this property as **antimonotonicity**. Note that this implies that $C(u \cup v) \subseteq C(u) \cap C(v)$. Our earlier desire about conjunctions (now extended to the setting of finite sets of formulae) translates to wanting the previous inclusion to become an equality:

- $C(u \cup v) = C(u) \cap C(v)$.

We refer to this last property as **multiplicativity**.

Suppose now that we want to find $[b]_{\mathcal{U}} \in \mathcal{N}$ that realizes Σ . How should we define $b(i) \in M_i$? Well, we should simply collect all the $\varphi \in \Sigma$ that \mathcal{M}_i is concerned about and then take $b(i) \in M_i$ such that $\mathcal{M}_i \models \varphi(b(i))$ for each such φ . But how do we know that there are only finitely many such φ ? It seems that we should add that to our list of requirements:

- For each $i \in I$, there are only finitely many $u \in \mathcal{P}_f(\Sigma)$ such that $i \in C(u)$.

We call this property **local finiteness**.

Exercise 8.4.1. For a multiplicative concern function $C : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$, prove that local finiteness is equivalent to, for each $i \in I$, there is a finite set $C_i \in \mathcal{P}_f(\Sigma)$ such that, for all $u \in \mathcal{P}_f(\Sigma)$, we have $i \in C(u)$ if and only if $u \subseteq C_i$.

Given the previous exercise, we can thus define $b(i) \in M_i$ so that $\mathcal{M}_i \models \bigwedge_{\varphi \in C_i} \varphi(b(i))$, and we see that the corresponding $[b]_{\mathcal{U}} \in \mathcal{N}$ is as desired: given $\varphi \in \Sigma$ and $i \in C(\{\varphi\})$, we have $\{\varphi\} \subseteq C_i$, so $\mathcal{M}_i \models \varphi(b(i))$.

We have just argued:

Proposition 8.4.2. *Suppose that there is a multiplicative, locally finite concern function $C : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$ that refines the Łoś map. Then Σ is realized in \mathcal{N} .*

So, the question becomes: how hard is it to come by such a concern function? It turns out that, for regular ultrafilters, it is fairly easy to come by a concern function that is almost as desired:

Proposition 8.4.3. *If \mathcal{U} is κ -regular, then there is an antimonotonic, locally finite concern function $C : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$ that refines the Łoś map.*

Proof. Let $(\varphi_i)_{i < \kappa}$ enumerate Σ and let $E = \{E_i : i < \kappa\}$ be a κ -regularizing set for \mathcal{U} . We then define $C : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$ by setting

$$C(\{\varphi_{i_1}, \dots, \varphi_{i_k}\}) := \text{Łoś}(\{\varphi_{i_1}, \dots, \varphi_{i_k}\}) \cap E_{i_1} \cap \dots \cap E_{i_k}.$$

It is readily verified that C is as desired. \square

The following terminology is more standard:

Definition 8.4.4. A **distribution** for Σ is a function $C : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$ that is antimonotonic, refines the Łoś map, and is locally finite.

Proposition 8.4.3 above shows that if \mathcal{U} is κ -regular, then there is a distribution for Σ . Note that a refinement of a distribution is once again a distribution. Consequently, we have:

Corollary 8.4.5. *If some distribution for Σ has a multiplicative refinement, then Σ is realized in \mathcal{N} .*

We now want a criteria on \mathcal{U} that implies all distributions have multiplicative refinements, yet that is intrinsic in that it only mentions \mathcal{U} itself. Here is the key definition in this section:

Definition 8.4.6. We say that \mathcal{U} is κ^+ -**good** if every antimonotonic function $\mathcal{P}_f(\kappa) \rightarrow \mathcal{U}$ has a multiplicative refinement.

We first note:

Proposition 8.4.7. *If \mathcal{U} is countably incomplete and κ^+ -good, then \mathcal{U} is κ -regular.*

Proof. Suppose that $(I_n)_{n \in \mathbb{N}}$ is a sequence from \mathcal{U} such that $I_n \supseteq I_{n+1}$ and $\bigcap_{n \in \mathbb{N}} I_n = \emptyset$. Define $f : \mathcal{P}_f(\kappa) \rightarrow \mathcal{U}$ by $f(w) = I_{|w|}$. Clearly, f is antimonotonic. Let $g : \mathcal{P}_f(\kappa) \rightarrow \mathcal{U}$ be a multiplicative refinement of f . We claim that, setting $E_\alpha := g(\{\alpha\})$, we have that $E := \{E_\alpha : \alpha \in \kappa\}$ is a κ -regularizing set for \mathcal{U} . To see this, suppose, toward a contradiction, that $w \subseteq \kappa$ is infinite and $x \in \bigcap_{\alpha \in w} E_\alpha$. For each n , let $w(n) \subseteq w$ be finite of size n . Then $x \in \bigcap_{\alpha \in w(n)} E_\alpha = g(w(n)) \subseteq I_n$ for each n , a contradiction. We leave it to the reader to verify that $|E| = \kappa$. \square

Corollary 8.4.8. *If $|I| = \kappa$ and \mathcal{U} is a countably incomplete ultrafilter on I , then \mathcal{U} is not κ^{++} -good.*

By the above corollary, the most good (horrible English!) an ultrafilter on I can be is $|I|^+$ -good; for this reason, we simply call such ultrafilters **good**.

Proposition 8.4.3, Corollary 8.4.5, and Proposition 8.4.7 imply:

Proposition 8.4.9. *If \mathcal{U} is a countably incomplete κ^+ -good ultrafilter on I , then for any language \mathcal{L} with $|\mathcal{L}| \leq \kappa$ and any family $(\mathcal{M}_i)_{i \in I}$ of \mathcal{L} -structures, $\prod_{\mathcal{U}} \mathcal{M}_i$ is κ^+ -saturated.*

In order for the previous proposition to be interesting, we need to know that κ^+ -good ultrafilters exist. We will take care of this in the next section.

Remark 8.4.10. Suppose that \mathcal{M} is an \mathcal{L} -structure with $|\mathcal{L}|, |\mathcal{M}| \leq \kappa$. Suppose that \mathcal{U} is a κ^+ -good ultrafilter on κ . Then $\mathcal{M}^{\mathcal{U}}$ is κ^+ -saturated. Since κ^+ -good implies κ -regular, we see, assuming GCH, that $|\mathcal{M}^{\mathcal{U}}| = \kappa^+$. Thus, in general, one cannot hope to prove in ZFC that κ^+ -good ultrafilters yield ultrapowers which are κ^{++} -saturated

We now work toward proving that the converse of Proposition 8.4.9 holds. Such an argument presumably would proceed by assuming that \mathcal{U} is not κ^+ -good and finding a structure \mathcal{M} in a language of size at most κ for which $\mathcal{N} := \mathcal{M}^{\mathcal{U}}$ is not κ^+ -saturated, that is, by finding some finitely satisfiable set Σ of \mathcal{L}_A sentences ($A \subseteq \mathcal{N}$, $|A| \leq \kappa$) that is not realized in \mathcal{N} . One might guess that if Σ is realized in \mathcal{N} , then every distribution for Σ has a multiplicative refinement, whence if we found some distribution for Σ that did not have a multiplicative refinement, then we could conclude that Σ could not be realized in \mathcal{N} . It turns out that it is not quite true that if Σ is realized in \mathcal{N} , then every distribution for Σ has a multiplicative refinement. Let us try to prove this, see where we get stuck, and make a key definition.

Suppose that Σ is realized by $[b]_{\mathcal{U}} \in N$ and that C is a distribution for Σ . Let us define a new function $C' : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$ by

$$C'(u) = \left\{ i \in \kappa : \mathcal{M}_i \models \bigwedge_{\varphi \in u} \varphi(b(i)) \right\} \cap C(u).$$

Note that C' really does take values in \mathcal{U} as $[b]_{\mathcal{U}}$ realizes Σ . Clearly, C' refines C , whence C' is also a distribution. We would like to have that C' is multiplicative. Toward that end, suppose that $i \in C'(u) \cap C'(v)$; we need $i \in C'(u \cup v)$. It is clear that i belongs to the first set involved in the definition of $C'(u \cup v)$; the question becomes whether or not $i \in C(u \cup v)$. While this seems to be begging the question (namely in trying to prove that C has a multiplicative refinement, we would need to assume that C itself is multiplicative), we actually have some more information, namely that \mathcal{M}_i does actually satisfy all elements of $u \cup v$. This leads to the following:

Definition 8.4.11. We call a distribution $C : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$ **accurate** if, for every $u \in \mathcal{P}_f(\Sigma)$, if $i \in C(\{\varphi\})$ for all $\varphi \in u$, then $[i \in C(u)$ if and only if $\mathcal{M}_i \models \exists x \bigwedge_{\varphi \in u} \varphi(x)]$.

We have thus proven:

Proposition 8.4.12. *If Σ is realized in \mathcal{N} , then every accurate distribution has a multiplicative refinement.*

Accurate distributions exist in any regular ultraproduct:

Lemma 8.4.13. *Suppose that \mathcal{U} is κ -regular. Then Σ has an accurate distribution.*

Exercise 8.4.14. Prove the previous lemma. (*Hint.* The distribution constructed in Proposition 8.4.3 is accurate.)

Summarizing thus far:

Theorem 8.4.15. *Suppose that \mathcal{U} is a κ -regular ultrafilter. Then the following are equivalent:*

- (1) *Some distribution of Σ has a multiplicative refinement.*
- (2) *Σ is realized in \mathcal{N} .*
- (3) *Every accurate distribution of Σ has a multiplicative refinement.*

Proof. (1) implies (2) is Corollary 8.4.5. (2) implies (3) is Proposition 8.4.12. (3) implies (1) follows from Lemma 8.4.13. \square

We are now ready to carry out the key argument for establishing the converse of Proposition 8.4.9.

Theorem 8.4.16. *Suppose that \mathcal{U} is a κ -regular ultrafilter that is not κ^+ -good. Let \mathcal{M} be the structure whose universe is $\mathcal{P}_f(\omega)$ and which possesses a single binary relation \subseteq to be interpreted as subset. Then $\mathcal{M}^{\mathcal{U}}$ is not κ^+ -saturated.*

Proof. By Theorem 8.3.16, we may replace \mathcal{M} with a κ^+ -saturated elementarily equivalent structure \mathcal{M}' . Since \mathcal{U} is not κ^+ -good, there is an antimonotonic $f : \mathcal{P}_f(\kappa) \rightarrow \mathcal{U}$ that has no multiplicative refinement. We will use this f to produce a set Σ of formulae with parameters from $(\mathcal{M}')^{\mathcal{U}}$ of size at most κ and an *accurate* distribution C for Σ that has no multiplicative refinement; by Theorem 8.4.15, this establishes that Σ cannot be realized in $(\mathcal{M}')^{\mathcal{U}}$, and hence $(\mathcal{M}')^{\mathcal{U}}$ is not κ^+ -saturated.

Fix $i \in I$ and temporarily call $j \in \kappa$ *i -relevant* if $i \in f(\{j\})$. Consider the set $\Gamma_i(y_j)$ of formulae in the free variables y_j for those j which are

i -relevant, defined by

$$\Gamma_i(y_j) := \left\{ \exists x \bigwedge_{l \leq k} x \subseteq y_{j_l} : i \in f(\{j_1, \dots, j_k\}) \right\} \\ \cup \left\{ \neg \exists x \bigwedge_{l \leq k} x \subseteq y_{j_l} : i \notin f(\{j_1, \dots, j_k\}) \right\}.$$

We leave it to the reader to check that $\Gamma_i(y_j)$ is finitely satisfiable in \mathcal{M}' , hence satisfiable in \mathcal{M}' by the assumption that \mathcal{M}' is κ^+ -saturated. For $j < \kappa$ that is i -relevant, we let $a_j(i) \in \mathcal{M}'$ realize Γ_i ; otherwise, take $a_j(i) \in \mathcal{M}'$ for j not i -relevant in an arbitrary fashion. (That set is outside of \mathcal{U} , whence it is negligible.)

We have now defined elements $[a_j]_{\mathcal{U}} \in (\mathcal{M}')^{\mathcal{U}}$ for $j < \kappa$. Set $\varphi_j(x)$ to be the formula $x \subseteq [a_j]_{\mathcal{U}}$ and $\Sigma := \{\varphi_j(x) : j \in \kappa\}$. Let $\{E_j : j < \kappa\}$ be a regularizing subset of \mathcal{U} and define $C : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$ by first setting $C(\{\varphi_j\}) := f(\{j\}) \cap E_j$ and then extending it to all of $\mathcal{P}_f(\Sigma)$ by setting

$$C(\{\varphi_{j_1}, \dots, \varphi_{j_k}\}) := \bigcap_{l \leq k} C(\{\varphi_{j_l}\}) \cap \text{LoS}(\{\varphi_{j_1}, \dots, \varphi_{j_k}\}) \cap f(\{j_1, \dots, j_k\}).$$

We leave it to the reader to check that C is an accurate distribution for Σ . Since C refines f , C has no multiplicative refinement taking values in \mathcal{U} by assumption. Consequently, Σ is not realized in $(\mathcal{M}')^{\mathcal{U}}$ and thus, $(\mathcal{M}')^{\mathcal{U}}$ is not κ^+ -saturated, as desired. \square

We have now seen that the κ^+ -good ultrafilters are exactly the ultrafilters that yield κ^+ -saturated ultraproducts (when the language has size at most κ). We can also see that other variations hold:

Theorem 8.4.17. *Let \mathcal{U} be a countably incomplete ultrafilter on I . The following are equivalent:*

- (1) \mathcal{U} is κ^+ -good.
- (2) For any language \mathcal{L} with $|\mathcal{L}| \leq \kappa$, all \mathcal{U} -ultraproducts of \mathcal{L} -structures are κ^+ -saturated.
- (3) For any language \mathcal{L} with $|\mathcal{L}| \leq \kappa$, all \mathcal{U} -ultraproducts of \mathcal{L} -structures are κ^+ -universal.
- (4) For any language \mathcal{L} with $|\mathcal{L}| \leq \kappa$, all \mathcal{U} -ultrapowers of \mathcal{L} -structures are κ^+ -saturated.

Proof. (1) implies (2) is Proposition 8.4.9 above and (2) implies (3) follows from Proposition 8.1.11.

(3) implies (4): Suppose that $|\mathcal{L}| \leq \kappa$ and \mathcal{M} is an \mathcal{L} -structure; we wish to show that $\mathcal{M}^{\mathcal{U}}$ is κ^+ -saturated. By Proposition 8.1.13, it suffices to show that $(\mathcal{M}^{\mathcal{U}}; a)$ is κ^+ -universal for every $a \in (M^{\mathcal{U}})^{\kappa}$. However, this latter statement follows from (3), the observation that $(\mathcal{M}^{\mathcal{U}}; a) = \prod_{\mathcal{U}}(\mathcal{M}; a(i))$, and the fact that $|\mathcal{L}(\kappa)| \leq \kappa$.

(4) implies (1): Suppose that (4) holds and \mathcal{U} is not κ^+ -good. Since κ^+ -saturation implies κ^+ -universal, we see from Theorem 8.3.10 that \mathcal{U} is κ -regular. We thus obtain a contradiction using Theorem 8.4.16. \square

Corollary 8.4.18. *Every nonprincipal ultrafilter on ω is good.*

Proof. This follows immediately from Theorems 8.2.1 and 8.4.17. \square

Exercise 8.4.19. Prove Corollary 8.4.18 directly from the definition of \aleph_1 -good.

Exercise 8.4.20. Suppose that \mathcal{V} is κ -regular. Prove that $\mathcal{U} \times \mathcal{V}$ is κ^+ -good if and only if \mathcal{V} is κ^+ -good. (*Hint.* Use Theorems 6.9.1 and 8.4.17.)

The previous exercise yields the following:

Corollary 8.4.21. *For each uncountable κ , there is a countably incomplete ultrafilter on κ that is regular but not good.*

Proof. Let \mathcal{U} be a regular ultrafilter on κ and let \mathcal{V} be a nonprincipal ultrafilter on ω . Then $\mathcal{U} \times \mathcal{V}$ is a κ -regular ultrafilter on $\kappa \times \omega$ that is not \aleph_2 -good, so not good. \square

Exercise 8.4.22. For each uncountable κ , give an example of ultrafilters \mathcal{U} and \mathcal{V} on κ such that $\mathcal{U} \times \mathcal{V} \not\equiv_{RK} \mathcal{V} \times \mathcal{U}$.

We end this section with an interesting application of Theorem 8.4.17.

Corollary 8.4.23. *Suppose that $|\mathcal{L}| \leq \kappa$ and that \mathcal{M} and \mathcal{N} are \mathcal{L} -structures with $|M|, |N| \leq \kappa^+$. Suppose that $2^\kappa = \kappa^+$. Suppose that \mathcal{U} is a good ultrafilter on κ . Then $\mathcal{M} \equiv \mathcal{N}$ if and only if $\mathcal{M}^{\mathcal{U}} \cong \mathcal{N}^{\mathcal{U}}$.*

Proof. The only direction that needs proving is the forward direction. If $\mathcal{M} \equiv \mathcal{N}$, then $\mathcal{M}^{\mathcal{U}}$ and $\mathcal{N}^{\mathcal{U}}$ are elementarily equivalent κ^+ -saturated models of size 2^κ by Theorems 8.3.9 and 8.4.17. Since we are assuming that $2^\kappa = \kappa^+$, we conclude that $\mathcal{M}^{\mathcal{U}} \cong \mathcal{N}^{\mathcal{U}}$ by Theorem 8.1.9. \square

Corollary 8.4.24. *Assume GCH. Then for any language \mathcal{L} and any \mathcal{L} -structures \mathcal{M} and \mathcal{N} , we have that $\mathcal{M} \equiv \mathcal{N}$ if and only if \mathcal{M} and \mathcal{N} have isomorphic ultrapowers.*

The previous corollary is interesting in that it gives an “algebraic” or “logic-free” characterization of elementary equivalence. The issue is that it has a heavy set-theoretic hypothesis, namely the GCH. Thankfully, one can remove the GCH assumption (at the expense of a much more complicated proof); the resulting theorem is called the **Keisler-Shelah theorem** and will be discussed in Chapter 16.

8.5. Good ultrafilters: Part 2

In this section, we prove the existence of good ultrafilters. The rough idea is to construct an increasing sequence of filters that eventually become an ultrafilter (we take care of this at “even” steps of the construction) and which satisfies the goodness property (which we take care of at “odd” steps of the construction). As we go along, in order to ensure that we can continue the construction, we need to insist that there are “many” possible extensions of our filter; the notion of many possible extensions will be made precise in the notion of consistent pair to be defined below. Lemmas 8.5.7 and 8.5.8 will take care of the aforementioned even and odd steps, while Lemma 8.5.6 gets us started with many possible extensions in the first place.

We recommend that the reader read the following definition and the statements of the lemmas and then proceed to the proof of the existence of good ultrafilters. They can then return to read the proofs of the lemmas (if they desire).

The key ingredient in establishing the existence of good ultrafilters is a strengthening of the notion of independent set of functions as introduced in Definition 1.4.2. Recall that $X \subseteq 2^\kappa$ is called independent if, for any $f_1, \dots, f_n \in X$ and any $y_1, \dots, y_n \in 2 = \{0, 1\}$, there is $\alpha < \kappa$ such that $f_i(\alpha) = y_i$ for all $i = 1, \dots, n$. We will generalize this notion in two ways.

First, we will consider subsets of κ^κ rather than 2^κ . Consequently, we will demand that, for any $f_1, \dots, f_n \in X$ and any $\alpha_1, \dots, \alpha_n \in \kappa$, there is $\alpha < \kappa$ such that $f_i(\alpha) = \alpha_i$ for all $i = 1, \dots, n$.

Before describing the second generalization, we first ask the reader to verify the following:

Exercise 8.5.1. Suppose that $X \subseteq 2^\kappa$ is independent. Prove that, for any $f_1, \dots, f_n \in X$ and any $y_1, \dots, y_n \in 2$, there are *infinitely* many $\alpha < \kappa$ such that $f_i(\alpha) = y_i$ for all $i = 1, \dots, n$.

Said in filter terms, the previous exercise guarantees that, for any $f_1, \dots, f_n \in X$ and any $y_1, \dots, y_n \in 2$, the filter generated by the Fréchet filter on κ and the set $\bigcap_{i=1}^n f_i^{-1}(\{y_i\})$ is proper. Our second generalization involves replacing the Fréchet filter by an arbitrary filter on κ .

Here is the precise definition:

Definition 8.5.2. Suppose that κ is a cardinal, $F \subseteq \kappa^\kappa$ is a set of functions, and \mathcal{F} is a filter on κ . We say that F is of **large oscillation modulo \mathcal{F}** if, given any set $X \in \mathcal{F}$, any distinct functions $f_1, \dots, f_n \in F$, and any $\alpha_1, \dots, \alpha_n < \kappa$, we have that

$$X \cap \{\alpha < \kappa : f_i(\alpha) = \alpha_i \text{ for } i = 1, \dots, n\} \neq \emptyset.$$

In this section, we will take a different (but equivalent), perspective on this notion:

Definition 8.5.3. Fix a cardinal κ . Let Π be a nonempty set of partitions of κ such that each partition in Π has κ many (nonempty) pieces. Let \mathcal{F} be a filter on κ . We say that (Π, \mathcal{F}) is **consistent** if, given any $X \in \mathcal{F}$, any distinct $P_1, \dots, P_n \in \Pi$, and any $X_1 \in P_1, \dots, X_n \in P_n$, we have that $X \cap \bigcap_{i=1}^n X_i \neq \emptyset$.

The previous two definitions really are equivalent:

Exercise 8.5.4. Suppose that κ is a cardinal and \mathcal{F} is a filter on κ . Given a function $f \in \kappa^\kappa$, let π_f be the partition of κ defined by $\pi_f = \{f^{-1}(\gamma) : \gamma < \kappa\}$. (Note that some elements of this partition might be empty.) Given $F \subseteq \kappa^\kappa$, set $\Pi_F := \{\pi_f : f \in F\}$. Prove that $F \subseteq \kappa^\kappa$ is of large oscillation modulo \mathcal{F} if and only if (Π_F, \mathcal{F}) is a consistent pair. Moreover, prove that, for any $\lambda \leq 2^\kappa$, this assignment provides a bijection between the set of subsets of κ^κ of size λ that are of large oscillation modulo \mathcal{F} and the set of sets Π of partitions of κ of size λ such that (Π, \mathcal{F}) is a consistent pair.

We now work toward proving the existence of good ultrafilters. Our first lemma will actually need its own set-theoretic lemma.

Lemma 8.5.5. *Suppose that κ is an infinite cardinal and, for each $\alpha < \kappa$, Y_α is a set of cardinality κ . Then, for each $\alpha < \kappa$, there is $Z_\alpha \subseteq Y_\alpha$ of cardinality κ such that $Z_\alpha \cap Z_\beta = \emptyset$ for each $\alpha \neq \beta$.*

Proof. For each $\alpha \leq \kappa$, we let $T_\alpha := \{(\beta, \gamma) : \beta \leq \gamma < \alpha\}$. In other words, if one thinks of the “discrete grid” $\alpha \times \alpha$, then T_α is simply the set of points on or above the diagonal. For each $\beta < \alpha < \kappa$, we let $T_{\alpha, \beta} := T_\alpha \cap (\{\beta\} \times \alpha)$, that is, the vertical slice of T_α above β . We define, by induction on $\alpha < \kappa$, an increasing chain of injective functions f_α with domain T_α such that $f(T_{\alpha, \beta}) \subseteq Y_\beta$ for each $\beta < \alpha$. Indeed, suppose that f_α has been constructed as such. Note that $T_{\alpha+1}$ simply adds one more point to each vertical slice of T_α ; since $|T_\alpha| < \kappa$ and each $|Y_\beta| = \kappa$, it is clear how to extend f_α to an injective function $f_{\alpha+1}$ on $T_{\alpha+1}$. Also, taking unions at limits preserves the requirements of the recursion.

We now set $f := \bigcup_{\alpha < \kappa}$. Note that f is injective and $f(T_{\kappa,\beta}) \subseteq Y_\beta$ for each $\beta < \kappa$. Since $|T_{\kappa,\beta}| = \kappa$ and f is injective, setting $Z_\beta := f(T_{\kappa,\beta})$ yields the desired sets. \square

The following generalizes Theorem 1.4.4:

Lemma 8.5.6. *Suppose that \mathcal{F} is a uniform filter on κ generated by $E \subseteq \mathcal{F}$, $|E| \leq \kappa$. Then there is a set Π of partitions of κ with $|\Pi| = 2^\kappa$ and such that (Π, \mathcal{F}) is consistent.*

Proof. Without loss of generality, E is closed under finite intersections. Let $(J_\alpha : \alpha < \kappa)$ enumerate E ; since \mathcal{F} is uniform, each J_α has cardinality κ . By Lemma 8.5.5, there are $I_\alpha \subseteq J_\alpha$ such that each I_α has cardinality κ and $I_\alpha \cap I_\beta = \emptyset$ for $\alpha \neq \beta$.

Let $(f_\gamma)_{\gamma < \kappa}$ be an enumeration of all functions $f : \mathcal{P}(u) \rightarrow \kappa$, where $u \subseteq \kappa$ is finite. We let u_γ denote the domain of f_γ . Without loss of generality, we assume that, for every $\alpha < \kappa$, $\{f_\gamma : \gamma \in I_\alpha\}$ itself is an enumeration of the set of such functions. (Recall that each I_α has cardinality κ .)

We are now ready to define our partitions. First, for each $X \subseteq \kappa$, define $f_X : \kappa \rightarrow \kappa$ by $f_X(\gamma) := f_\gamma(X \cap u_\gamma)$ if $\gamma \in I_\alpha$ for some $\alpha < \kappa$; otherwise, set $f_X(\gamma) = \emptyset$. We then let $P_X := \{f_X^{-1}(\{\delta\}) : \delta < \kappa\}$. It is clear that each P_X is a partition of κ into precisely κ many pieces.

We first show that $P_X \neq P_Y$ for distinct subsets X and Y of κ (whence there are 2^κ many such partitions). Indeed, suppose, without loss of generality, that $\eta \in X \setminus Y$. Let $f : \mathcal{P}(\{\eta\}) \rightarrow \kappa$ be defined by $f(\{\eta\}) = 0$ while $f(\emptyset) = 1$. Let $\alpha < \kappa$ be arbitrary and take $\gamma \in I_\alpha$ such that $f = f_\gamma$. Then $f_X(\gamma) = f_\gamma(X \cap u_\gamma) = f(\{\eta\}) = 0$ while $f_Y(\gamma) = f_\gamma(Y \cap u_\gamma) = f(\emptyset) = 1$. It follows that P_X and P_Y are distinct partitions of κ .

We finish the proof by showing that (Π, \mathcal{F}) is consistent. To see this, it suffices to show that $J_\alpha \cap f_{X_1}^{-1}(\{\delta_1\}) \cap \cdots \cap f_{X_n}^{-1}(\{\delta_n\}) \neq \emptyset$ for any α , $\delta_1, \dots, \delta_n < \kappa$ and distinct subsets X_1, \dots, X_n of κ . Toward this end, take $u \subseteq \kappa$ finite such that $X_i \cap u \neq X_j \cap u$ for distinct $i, j = 1, \dots, n$. Define $f : \mathcal{P}(u) \rightarrow \kappa$ so that $f(X_i \cap u) = \delta_i$ (and defined arbitrarily for other subsets of u). Take $\gamma \in I_\alpha$ such that $f = f_\gamma$. Then, for each $i = 1, \dots, n$, we have $f_{X_i}(\gamma) = f_\gamma(X_i \cap u_\gamma) = \delta_i$, as desired. \square

Lemma 8.5.7. *Suppose that (Π, \mathcal{F}) is consistent and $J \subseteq \kappa$. Then either $(\Pi, \langle \mathcal{F}, \{J\} \rangle)$ is consistent or else there is a cofinite $\Pi' \subseteq \Pi$ for which $(\Pi', \langle \mathcal{F}, \{\kappa \setminus J\} \rangle)$ is consistent.*

Proof. Suppose that $(\Pi, \langle \mathcal{F}, \{J\} \rangle)$ is not consistent. By definition, this means that there is $X \in \mathcal{F}$, distinct $P_1, \dots, P_n \in \Pi$, and $X_i \in P_i$, $i = 1, \dots, n$, for which $X \cap J \cap \bigcap_{i=1}^n X_i = \emptyset$. Let $\Pi' := \Pi \setminus \{P_1, \dots, P_n\}$, so Π'

is a cofinite subset of Π . We claim that $(\Pi', \langle \mathcal{F}, \{\kappa \setminus J\} \rangle)$ is consistent. To see this, fix $Y \in \mathcal{F}$, distinct $Q_1, \dots, Q_m \in \Pi'$, and $X'_j \in Q_j$, $j = 1, \dots, m$. Since (Π, \mathcal{F}) is consistent, there is $x \in X \cap Y \cap \bigcap_{i=1}^n X_i \cap \bigcap_{j=1}^m X'_j$. By the above, we see that $x \notin J$, whence $(\kappa \setminus J) \cap Y \cap \bigcap_{j=1}^m X'_j \neq \emptyset$, as desired. \square

Lemma 8.5.8. *Suppose that (Π, \mathcal{F}) is consistent. Let $p : \mathcal{P}_f(\kappa) \rightarrow \mathcal{F}$ be an antimonotonic map and let $P \in \Pi$. Then there is a filter \mathcal{F}' extending \mathcal{F} such that $(\Pi \setminus \{P\}, \mathcal{F}')$ is consistent and a multiplicative refinement $q : \mathcal{P}_f(\kappa) \rightarrow \mathcal{F}'$ of p .*

Proof. Let $(X_\delta)_{\delta < \kappa}$ be an injective enumeration of P . Let $(t_\delta)_{\delta < \kappa}$ be an enumeration of $\mathcal{P}_f(\kappa)$. For $\delta < \kappa$, define $q_\delta : \mathcal{P}_f(\kappa) \rightarrow \mathcal{P}(\kappa)$ by $q_\delta(u) = p(t_\delta) \cap X_\delta$ if $u \subseteq t_\delta$; otherwise, $q_\delta(u) = \emptyset$. Some observations:

- $q_\delta(u) \subseteq p(t_\delta)$;
- $q_\delta(u) \neq \emptyset$ if $u \subseteq t_\delta$ (by consistency of (Π, \mathcal{F}));
- q_δ is multiplicative.

To see the last point, take $u, v \in \mathcal{P}_f(\kappa)$. If either $u \not\subseteq t_\delta$ or $v \not\subseteq t_\delta$, then $q_\delta(u \cup v) = q_\delta(u) \cap q_\delta(v) = \emptyset$. Otherwise, $u, v \subseteq t_\delta$, whence $u \cup v \subseteq t_\delta$, and thus $q_\delta(u \cup v) = q_\delta(u) = q_\delta(v) = p(t_\delta) \cap X_\delta$.

Define $q : \mathcal{P}_f(\kappa) \rightarrow \mathcal{P}(\kappa)$ by $q(u) := \bigcup_{\delta < \kappa} q_\delta(u)$. First note that q refines p . To see this, fix $u \in \mathcal{P}_f(\kappa)$ and note that, if $u \subseteq t_\delta$, then $q_\delta(u) = p(t_\delta) \cap X_\delta \subseteq p(u) \cap X_\delta$ by antimonotonicity of p ; it follows that $q(u) \subseteq p(u)$, as desired.

We next note that q is multiplicative. This follows immediately from the fact that each q_δ is multiplicative and the X_δ 's are disjoint.

Let J denote the range of q . Let $\mathcal{F}' := \langle \mathcal{F}, J \rangle$. We conclude by showing that $(\Pi \setminus \{P\}, \mathcal{F}')$ is consistent. To see this, take $X \in \mathcal{F}$, $u \in \mathcal{P}_f(\kappa)$, distinct $P_1, \dots, P_n \in \Pi \setminus \{P\}$, and $X_i \in P_i$, $i = 1, \dots, n$; we show that $X \cap q(u) \cap \bigcap_{i=1}^n X_i \neq \emptyset$. (Note that, since q is multiplicative, an arbitrary element of \mathcal{F}' is of the form $X \cap q(u)$.) Take δ such that $u = t_\delta$. Then $q(u) \supseteq q_\delta(u) = p(t_\delta) \cap X_\delta$. Since (Π, \mathcal{F}) is consistent and $p(t_\delta) \in \mathcal{F}$, we have that $X \cap p(t_\delta) \cap X_\delta \cap \bigcap_{i=1}^n X_i \neq \emptyset$, concluding the proof. \square

We now have all of the necessary ingredients to prove:

Theorem 8.5.9. *For each κ , there is a countably incomplete κ^+ -good ultrafilter on κ .*

Proof. Let $(X_\alpha)_{\alpha < 2^\kappa}$ enumerate $\mathcal{P}(\kappa)$ and let $(f_\alpha)_{\alpha < 2^\kappa}$ enumerate all antimonotonic functions from $\mathcal{P}_f(\kappa)$ to $\mathcal{P}(\kappa)$, with each such function enumerated 2^κ many times. We define, by transfinite recursion, a sequence of

consistent pairs $(\Pi_\alpha, \mathcal{F}_\alpha)$ for $\alpha < 2^\kappa$ satisfying the following conditions:

- (1) Each Π_α has cardinality 2^κ .
- (2) If $\alpha \leq \beta < \kappa$, then $\Pi_\alpha \supseteq \Pi_\beta$ and $\mathcal{F}_\alpha \subseteq \mathcal{F}_\beta$.
- (3) If $\alpha < 2^\kappa$ is a limit ordinal, then $\Pi_\alpha := \bigcap_{\beta < \alpha} \Pi_\beta$ and $\mathcal{F}_\alpha := \bigcup_{\beta < \alpha} \mathcal{F}_\beta$.
- (4) Each $\Pi_\alpha \setminus \Pi_{\alpha+1}$ is finite.
- (5) Either $X_\alpha \in \mathcal{F}_{\alpha+1}$ or $\kappa \setminus X_\alpha \in \mathcal{F}_{\alpha+1}$.
- (6) If the range of f_α is contained in \mathcal{F}_α , then there is a multiplicative refinement $g : \mathcal{P}_f(\kappa) \rightarrow \mathcal{F}_{\alpha+1}$ of f_α .

Once this has been completed, we set $\mathcal{U} := \bigcup_{\alpha < 2^\kappa} \mathcal{F}_\alpha$. By item (5), \mathcal{U} is an ultrafilter on κ . By item (6), \mathcal{U} is κ^+ -good. Indeed, suppose that $f : \mathcal{P}_f(\kappa) \rightarrow \mathcal{U}$ is antimonic. Since $\text{cof}(2^\kappa) > \kappa$, there is some $\alpha < 2^\kappa$ such that the range of f is contained in \mathcal{F}_α . Since each such f is enumerated 2^κ times, there is some $\beta > \alpha$ such that $f = f_\beta$. It follows that f has a multiplicative refinement that takes values in $\mathcal{F}_{\beta+1}$, and hence in \mathcal{U} , as desired.

We now show how to carry out the recursion. Item (3) tells us what to do at limit ordinals α ; note that item (4) ensures that the resulting Π_α still has cardinality 2^κ .

We now carry out the induction step. Suppose that $(\Pi_\beta, \mathcal{F}_\beta)$ have been constructed for all $\beta \leq \alpha$. By Lemma 8.5.7, there is $\Pi' \subseteq \Pi_\beta$ with $\Pi_\beta \setminus \Pi'$ finite and $\mathcal{F}' \supseteq \mathcal{F}$ such that (Π', \mathcal{F}') is consistent and either $X_\alpha \in \mathcal{F}'$ or $\kappa \setminus X_\alpha \in \mathcal{F}'$. If the range of f_α is not contained in \mathcal{F}' , then set $\Pi_{\alpha+1} := \Pi'$ and $\mathcal{F}_{\alpha+1} := \mathcal{F}'$. Otherwise, by Lemma 8.5.8, there is $\Pi_{\alpha+1} \subseteq \Pi'$ with $\Pi' \setminus \Pi_{\alpha+1}$ finite (a singleton even!) and $\mathcal{F}_{\alpha+1} \supseteq \mathcal{F}'$ such that $(\Pi_{\alpha+1}, \mathcal{F}_{\alpha+1})$ is consistent for which f_α has a multiplicative refinement with range included in $\mathcal{F}_{\alpha+1}$. This concludes the proof. \square

Exercise 8.5.10. Suppose that $E \subseteq \mathcal{P}(\kappa)$ is such that $|E| \leq \kappa$, each element of E has cardinality κ , and E is closed under finite intersections. Prove that there is a κ^+ -good ultrafilter on κ containing E .

Exercise 8.5.11. Prove that there are 2^{2^κ} many κ^+ -good ultrafilters on κ .

8.6. Keisler's order

Model theorists are fascinated by studying the “complexities” of theories in a myriad of different ways. In this section, we show how the ideas from earlier in this chapter lend themselves to one natural comparison of the complexity of theories due to Keisler.

We recall the statement of Theorem 8.3.16 above: Suppose that \mathcal{U} is a regular ultrafilter over a set of cardinality κ . Suppose that $|\mathcal{L}| \leq \kappa$ and \mathcal{M} and \mathcal{N} are \mathcal{L} -structures such that $\mathcal{M} \equiv \mathcal{N}$. Then $\mathcal{M}^{\mathcal{U}}$ is κ^+ -saturated if and only if $\mathcal{N}^{\mathcal{U}}$ is κ^+ -saturated.

This theorem allows us to make the following definition:

Definition 8.6.1. For a complete theory T in a countable language and a regular ultrafilter \mathcal{U} on κ , we say that \mathcal{U} **saturates** T if $\mathcal{M}^{\mathcal{U}}$ is κ^+ -saturated for some (equivalently, any) $\mathcal{M} \models T$.

The preceding definition allows us to compare theories:

Definition 8.6.2. For complete theories T_1 and T_2 in countable languages, we write $T_1 \trianglelefteq T_2$ if every regular ultrafilter that saturates T_2 also saturates T_1 .

When $T_1 \trianglelefteq T_2$, we view T_1 as no more complicated than T_2 (in some sense) as any ultrafilter that is “smart enough” to encode the intricacies of models of T_2 by saturating ultrapowers of its models can also do the same thing for models of T_1 . It is clear that \trianglelefteq is a preorder on complete theories, called **Keisler’s order**. As with any preorder, one obtains an equivalence relation by declaring T_1 and T_2 equivalent if $T_1 \trianglelefteq T_2$ and $T_2 \trianglelefteq T_1$. The preorder \trianglelefteq induces a partial order (also denoted \trianglelefteq) on the equivalence classes. We write $T_1 \triangleleft T_2$ if $T_1 \trianglelefteq T_2$ but $T_2 \not\trianglelefteq T_1$.

Example 8.6.3. Recall that in Example 8.1.7 we showed that every uncountable algebraically closed field is saturated. In particular, if K is an algebraically closed field and \mathcal{U} is a regular ultrafilter on κ , then $|K^{\mathcal{U}}| = |K|^{\kappa} \geq \kappa^+$, whence $K^{\mathcal{U}}$ is κ^+ -saturated. Since \mathcal{U} was arbitrary, this shows that ACF_p (the theory of algebraically closed fields of characteristic p) is a minimum element in Keisler’s order.

The previous example generalizes: call a complete theory T in a countable language **uncountably categorical** if for some uncountable cardinal κ , T has a unique model of cardinality κ up to isomorphism. A famous theorem of Morley states that a theory is uncountably categorical if and only if, for *every* uncountable cardinal κ , T has a unique model of size κ . Part of the proof of Morley’s theorem shows that all uncountable models of T are saturated. Thus, any uncountably categorical theory is also a minimum element in Keisler’s order. In particular, all such theories are equivalent (in the sense of Keisler’s order).

Proposition 8.6.4. *The complete theory T is a minimum in Keisler’s order if and only if every regular ultrafilter saturates T .*

Proof. The backward direction is exactly as in Example 8.6.3. For the forward direction, it just suffices to note that, given any regular ultrafilter \mathcal{U} , some theory is saturated by \mathcal{U} (e.g., ACF_p). \square

There are examples of theories that are a minimum in Keisler's order that are not uncountably categorical:

Example 8.6.5. Let $\mathcal{L} = \{E\}$, where E is a single binary relation symbol. Let T be the \mathcal{L} -theory which states that E is an equivalence relation with exactly two classes, both of which are infinite. Then T is not uncountably categorical. For example, T has two nonisomorphic models of cardinality \aleph_1 .

On the other hand, suppose that \mathcal{M} is the unique countable model of T and \mathcal{U} is an ultrafilter on an index set I . Set $\mathcal{N} := \mathcal{M}^{\mathcal{U}}$. We claim that both equivalence classes of \mathcal{N} have the same cardinality as \mathcal{N} itself. Indeed, let $\mathcal{L}' := \mathcal{L} \cup \{f\}$, where f is a new unary function symbol. Let \mathcal{M}' denote the expansion of \mathcal{M} which interprets f as a bijection between the two equivalence classes, and let \mathcal{N}' denote the expansion of $\mathcal{M}^{\mathcal{U}}$ which interprets f as the ultrapower of $f^{\mathcal{M}'}$. It follows from Łoś's theorem and Exercise 6.3.8 that $f^{\mathcal{N}'}$ is a bijection between the two equivalence classes of \mathcal{N} , as desired.

As a consequence, all ultrapowers of \mathcal{M} with respect to a regular ultrafilter on a cardinal κ are isomorphic and thus T is a minimum in Keisler's order.

In [101], Keisler introduced a syntactic notion, called the finite cover property, that he believed would characterize the theories in the minimum class. He was successful in showing that the finite cover property was a necessary condition for theories in the minimum class, but it was not until the work of Shelah [158] that the finite cover property was shown to be sufficient as well, thus confirming Keisler's suspicion.

Definition 8.6.6. A formula $\varphi(x, y)$ (where x and y are two finite tuples of variables) has the **finite cover property** (fcp) in T if there is $\mathcal{M} \models T$ such that, for arbitrarily large $n \in \mathbb{N}$, there are tuples $a_1, \dots, a_n \in M$ such that:

- $\mathcal{M} \models \neg \exists x \bigwedge_{1 \leq i \leq n} \varphi(x, a_i)$, but
- for every $m < n$, $\mathcal{M} \models \exists x \bigwedge_{1 \leq i \leq n, i \neq m} \varphi(x, a_i)$.

We refer to the set $\{\varphi(x, a_1), \dots, \varphi(x, a_n)\}$ as a (φ, n) -cover of \mathcal{M} . We say that T has the **finite cover property** if some formula has the fcp in T . If $\varphi(x, y)$ does not have the fcp in T , then we abuse grammar and say that φ has the **nfc**p in T . Similarly, if T does not have the fcp, we also say that T has the **nfc**p.

The “cover” terminology comes from the fact that one views the formulae $\neg\varphi(x, a_i)$, for $i = 1, \dots, n$, as covering the universe M (as every element of M makes one of the formulae true); the second item in the definition then states that this cover of M has no proper subcover. Consequently, the formula φ having the fcp in T is a kind of finitary nonsaturation property.

Exercise 8.6.7. Suppose that φ has the fcp in T as witnessed by $\mathcal{M} \models T$. Prove that all $\mathcal{N} \models T$ witness that T has the fcp.

Exercise 8.6.8. Prove that $\varphi(x, y)$ has the nfcps in T if and only if there is $k \in \mathbb{N}$ such that the following holds: for every $\mathcal{M} \models T$ and every $A \subseteq M$, setting $\Sigma := \{\varphi(x, a) : a \in A\}$, if every subset of Σ of size at most k is satisfiable in \mathcal{M} , then Σ is finitely satisfiable in \mathcal{M} .

Here is the prototypical example of a theory with the fcp:

Example 8.6.9. Let T be the theory of an equivalence relation with a unique equivalence class of size n for every $n \geq 1$. Then the formula $E(x, y) \wedge x \neq z$ has the fcp in T . To see this, let \mathcal{M} be the unique model of T with only finite equivalence classes. Fix $n \geq 1$ and let a_1, \dots, a_n enumerate the unique class of size n . Then $\mathcal{M} \models \neg\exists x \bigwedge_{1 \leq i \leq n} (E(x, a_i) \wedge x \neq a_i)$ while, for any $m \leq n$, $\mathcal{M} \models \exists x \bigwedge_{1 \leq i \leq n, i \neq m} (E(x, a_i) \wedge x \neq a_i)$.

As noted above, the fcp is some finitary failure of saturation; Keisler’s observation was that this finitary failure of saturation propagates somewhat to ultrapowers:

Proposition 8.6.10. *Suppose that T has the fcp and \mathcal{U} is a countably incomplete ultrafilter on κ . Then for any $\mathcal{M} \models T$, we have that $\mathcal{M}^{\mathcal{U}}$ is not $(2^\kappa)^+$ -saturated.*

Proof. Suppose that $\varphi(x, y)$ have the fcp in T . Consequently, for arbitrarily large $n \in \mathbb{N}$, there is a (φ, n) -cover in \mathcal{M} as witnessed by the tuples $a_1(n), \dots, a_n(n) \in M$. Take $(I_n)_{n \in \mathbb{N}}$ from \mathcal{U} such that $I_n \supseteq I_{n+1}$ for all $n \in \mathbb{N}$ and such that $\bigcap_{n \in \mathbb{N}} I_n = \emptyset$. For each $i \in I$, let $m(i)$ be the least $m \in \mathbb{N}$ such that $i \notin I_m$ and take $n(i) \geq m(i)$ such that there is a $(\varphi, n(i))$ -cover in \mathcal{M} . Take $S \subseteq \prod_{i \in I} [n(i)]$ such that every element of $\prod_{\mathcal{U}} [n(i)]$ is equivalent modulo \mathcal{U} to a unique element of S . Note that $|S| \leq 2^\kappa$. For each $s \in S$, let $f_s \in M^I$ be defined by $f_s(i) := a_{s(i)}(n(i))$. We show that $\mathcal{M}^{\mathcal{U}}$ is not $(2^\kappa)^+$ -saturated by showing that $\Sigma := \{\varphi(x, [f_s]_{\mathcal{U}}) : s \in S\}$ is finitely satisfiable in $\mathcal{M}^{\mathcal{U}}$ but not satisfiable in $\mathcal{M}^{\mathcal{U}}$.

To see that Σ is not satisfiable in \mathcal{U} , fix any $g \in M^I$; we show that $[g]_{\mathcal{U}}$ does not satisfy Σ . For each $i \in I$, take $k(i) \leq n(i)$ such that $\mathcal{M} \models \neg\varphi(g(i), a_{k(i)}(n(i)))$. Let $s \in S$ be such that $s(i) = k(i)$ for \mathcal{U} -almost all $i \in I$. Then $\mathcal{M}^{\mathcal{U}} \models \neg\varphi([g]_{\mathcal{U}}, [f_s]_{\mathcal{U}})$, whence $[g]_{\mathcal{U}}$ does not satisfy Σ .

We conclude by showing that Σ is finitely satisfiable in $\mathcal{M}^{\mathcal{U}}$. In fact, we prove something much stronger, namely, for every $s \in S$, $\Sigma \setminus \{\varphi(x, [f_s]_{\mathcal{U}})\}$ is satisfiable in $\mathcal{M}^{\mathcal{U}}$. Fix $s \in S$ and take $g \in M^I$ such that

$$\mathcal{M} \models \bigwedge_{1 \leq j \leq n(i), j \neq s(i)} \varphi(g(i), a_j(n(i))).$$

It follows that $[g]_{\mathcal{U}}$ satisfies $\Sigma \setminus \{\varphi(x, [f_s]_{\mathcal{U}})\}$. \square

Corollary 8.6.11. *Suppose that T is in the minimum class in Keisler's order. Then T has the nfcp.*

Proof. We prove the contrapositive. Suppose that T has the fcp. Let \mathcal{U}_1 be any regular ultrafilter on 2^κ and let \mathcal{U}_2 be any countably incomplete ultrafilter on κ . Set $\mathcal{U} := \mathcal{U}_1 \times \mathcal{U}_2$. By Exercise 8.3.7, \mathcal{U} is also a regular ultrafilter on 2^κ . Fix $\mathcal{M} \models T$. By Proposition 8.6.10, $\mathcal{M}^{\mathcal{U}} \cong (\mathcal{M}^{\mathcal{U}_1})^{\mathcal{U}_2}$ is not $(2^\kappa)^+$ -saturated. It follows that \mathcal{U} does not saturate T , whence T is not in the minimum class in Keisler's order. \square

Shelah was able to prove the converse of the previous corollary. This converse follows from the following more general result:

Theorem 8.6.12. *Suppose that T has the nfcp and \mathcal{U} is a countably incomplete ultrafilter. Set $\lambda := |\mathbb{N}^{\mathcal{U}}|$. Then for all $\mathcal{M} \models T$, \mathcal{M} is λ -saturated.*

Before discussing the proof of Theorem 8.6.12, we first see how it yields the promised syntactic characterization of the minimum class in Keisler's order:

Corollary 8.6.13. *A theory T is in the minimum class in Keisler's order if and only if T has the nfcp.*

Proof. The forward direction is Corollary 8.6.11. To prove the backward direction, suppose that T has the nfcp and take any regular ultrafilter \mathcal{U} on κ . By Theorem 8.3.9, $|\mathbb{N}^{\mathcal{U}}| = 2^\kappa \geq \kappa^+$. By Theorem 8.6.12, \mathcal{U} saturates T . Since \mathcal{U} was an arbitrary regular ultrafilter, it follows that T is in the minimum class in Keisler's order. \square

We now discuss the proof of Theorem 8.6.12, which makes heavy use of ideas from **stability theory**. The class of **stable theories** is an important class of theories in model theory where an in-depth analysis of the class of models is possible. We defer the definition of a stable theory to Chapter 15, where the notion will play a central role. For our purposes, it suffices to know the following fact, whose proof we relegate to an exercise in Chapter 15 (see Exercise 15.1.5):

Proposition 8.6.14. *If T has the nfcp, then T is stable.*

The key result from stability theory that we will need is Theorem 8.6.16 below. In order to state it, we need the following definitions.

Definition 8.6.15. Suppose that \mathcal{M} is an \mathcal{L} -structure and $A = \{a_i : i < \alpha\}$ is a subset of M .

- (1) For any finite set Δ of \mathcal{L} -formulae and any $n \in \mathbb{N}$, we say that A is a **Δ - n -indiscernible set in \mathcal{M}** if, for all $\varphi \in \Delta$ and all pairs (i_1, \dots, i_n) and (j_1, \dots, j_n) of n -tuples of ordinals below α , we have $\mathcal{M} \models \varphi(a_{i_1}, \dots, a_{i_n}) \leftrightarrow \varphi(a_{j_1}, \dots, a_{j_n})$.
- (2) We say that A is an **indiscernible set in \mathcal{M}** if it is a Δ - n -indiscernible set in \mathcal{M} for all finite sets Δ of \mathcal{L} -formulae and all $n \in \mathbb{N}$.

Theorem 8.6.16. *Suppose that T is a stable theory and $\mathcal{N} \models T$. For any uncountable cardinal κ , we have that \mathcal{N} is κ -saturated if and only if:*

- (1) \mathcal{N} is \aleph_1 -saturated.
- (2) *Every countably infinite indiscernible set in \mathcal{N} can be extended to an indiscernible set in \mathcal{N} of cardinality κ .*

Exercise 8.6.17. Prove that the forward direction of Theorem 8.6.16 holds in any model of any theory.

The main use of the assumption that T has the nfcp in Theorem 8.6.12 (besides the fact that it implies that T is stable) is the following fact:

Theorem 8.6.18. *Suppose that T has the nfcp. Then for every finite set Δ of \mathcal{L} -formulae and every $n \in \mathbb{N}$, there is $m = m(\Delta, n) \in \mathbb{N}$ with the following property: for every $\mathcal{M} \models T$, every Δ - n -indiscernible set in \mathcal{M} of length at least m can be extended to an infinite Δ - n -indiscernible set in \mathcal{M} .*

We are now ready to prove Theorem 8.6.12:

Proof of Theorem 8.6.12. Fix $\mathcal{M} \models T$. By Theorem 8.2.1, $\mathcal{M}^{\mathcal{U}}$ is \aleph_1 -saturated. Thus, by Theorem 8.6.16, it suffices to show that any infinite indiscernible set $\{c_n : n \in \mathbb{N}\}$ in $\mathcal{M}^{\mathcal{U}}$ can be extended to one of size λ . Fix a family S of subsets of M with the following three properties:

- $|S| = |M|$;
- $\mathcal{P}_f(M) \subseteq S$;
- for every finite set Δ of \mathcal{L} -formulae and every $n \in \mathbb{N}$, if $w \in S$ is a Δ - n -indiscernible set with the property that there is an infinite Δ - n -indiscernible set in $\mathcal{M}^{\mathcal{U}}$ containing w , then such an extension of w also exists in S .

Fix enumerations $\{a_\alpha : \alpha < |M|\}$ and $\{w_\alpha : \alpha < |M|\}$ of M and S , respectively. Consider the language $\mathcal{L}' := \mathcal{L} \cup \{E, P\}$, where E is a new binary relation symbol and P is a new unary relation symbol. We expand \mathcal{M} to an \mathcal{L}' -structure \mathcal{M}' by declaring, for $\alpha, \beta < |M|$, that:

- $(a_\alpha, a_\beta) \in E^{\mathcal{M}'}$ if and only if $a_\alpha \in w_\beta$, and
- $a_\alpha \in P^{\mathcal{M}'}$ if and only if w_α is infinite.

Next note that, for each finite set Δ of \mathcal{L} -formulae and $n \in \mathbb{N}$, there is an \mathcal{L}' -formula $\varphi_{\Delta, n}(x)$ such that, for all $\alpha < |M|$, we have $\mathcal{M} \models \varphi_{\Delta, n}(a_\alpha)$ if and only if w_α is a Δ - n -indiscernible set. Fix $m(\Delta, n)$ as in Theorem 8.6.18. By the definition of S , for each $\alpha < |M|$, we have that $\mathcal{M}' \models (|a_\alpha| \geq m(\Delta, n) \wedge \varphi_{\Delta, n}(a_\alpha)) \rightarrow \exists x(P(x) \wedge a_\alpha \subseteq x \wedge \varphi_{\Delta, n}(x))$. Here, we are abusing notation and writing $|a_\alpha| \geq m(\Delta, n)$ as shorthand for the statement $\exists x_1 \cdots \exists x_{m(\Delta, n)} (\bigwedge_{1 \leq i < j < m(\Delta, n)} x_i \neq x_j \wedge \bigwedge_{1 \leq i \leq m(\Delta, n)} E(x_i, a_\alpha))$; similar abuse of notation is being used in writing $a_\alpha \subseteq x$. Let

$$\Sigma := \{P(x)\} \cup \{E(c_n, x) : n \in \mathbb{N}\} \cup \{\varphi_{\Delta, n}(x) : \Delta \in \mathcal{P}_f(T), n \in \mathbb{N}\}.$$

By the above discussion (and the fact that S contains all finite subsets of M), we have that Σ is finitely satisfiable in $(\mathcal{M}')^{\mathcal{U}}$. Since $(\mathcal{M}')^{\mathcal{U}}$ is \aleph_1 -saturated, there is $[b]_{\mathcal{U}} \in M^{\mathcal{U}}$ that satisfies Σ . Set $X := \{[f]_{\mathcal{U}} \in M^{\mathcal{U}} : (\mathcal{M}')^{\mathcal{U}} \models E([f]_{\mathcal{U}}, [b]_{\mathcal{U}})\}$. It is clear that X is an indiscernible set in $\mathcal{M}^{\mathcal{U}}$ containing $\{c_n : n \in \mathbb{N}\}$. Moreover, setting $X_i := \{a \in M : \mathcal{M}' \models E(a, b(i))\}$, we have that X_i is infinite for \mathcal{U} -almost all i . Since $X = \prod_{\mathcal{U}} X_i$, it follows that $|X| \geq \lambda$, as desired. \square

What is interesting about the proof of Theorem 8.6.12 is that it also gives information about ultrapowers of models of stable theories *with* the fcp. In order to state this result, we first introduce a key definition:

Definition 8.6.19. Given an ultrafilter \mathcal{U} on an index set I , we set

$$\text{pfc}(\mathcal{U}) := \min \left\{ \left| \prod_{\mathcal{U}} X_i \right| : X_i \text{ a finite set for all } i \in I, \lim_{\mathcal{U}} |X_i| = \infty \right\}.$$

Note that, by Theorem 6.8.3, $\text{pfc}(\mathcal{U}) \geq \mathfrak{c}$ whenever \mathcal{U} is a countably incomplete ultrafilter.

Theorem 8.6.20. *Suppose that T is a stable theory with the fcp and that \mathcal{U} is a countably incomplete ultrafilter. For any $\mathcal{M} \models T$, we have that $\mathcal{M}^{\mathcal{U}}$ is $\text{pfc}(\mathcal{U})$ -saturated.*

Exercise 8.6.21. Prove Theorem 8.6.20.

It might appear that the conclusion of Theorem 8.6.20 is quite weak and that with a more clever proof, the conclusion can be strengthened. However, this is not the case:

Example 8.6.22. Recall the theory of an equivalence class with a unique class of size n for every $n \geq 1$ introduced in Example 8.6.9. As mentioned there, this theory has the fcp. It is in fact also a stable theory. Let \mathcal{M} be the model with only finite equivalence classes. For each $n \geq 1$, let $a_n \in M$ be an element of the equivalence class of size n . Fix an ultrafilter \mathcal{U} on λ and a sequence $(n_i)_{i < \lambda}$ of natural numbers such that $\text{pfc}(\mathcal{U}) = |\prod_{\mathcal{U}} [n_i]|$. Setting $a := [a_{n_i}]_{\mathcal{U}}$, we see that $E(x, a)$ defines the set $\prod_{\mathcal{U}} E(x, a_{n_i})$ in $\mathcal{M}^{\mathcal{U}}$, which has cardinality $\text{pfc}(\mathcal{U})$. It follows from Exercise 8.6.23 below that $\mathcal{M}^{\mathcal{U}}$ is not $\text{pfc}(\mathcal{U})^+$ -saturated. Consequently, $\text{pfc}(\mathcal{U})$ -saturation is the optimal level of saturation for $\mathcal{M}^{\mathcal{U}}$.

Exercise 8.6.23. Suppose that \mathcal{M} is a κ -saturated structure and $X \subseteq M^n$ is an infinite definable set. Prove that $|X| \geq \kappa$.

It turns out that the conclusion of Example 8.6.22 is always the case in a stable theory with the fcp:

Theorem 8.6.24. *Suppose that T is a stable theory with the fcp. For any $\mathcal{M} \models T$ and any ultrafilter \mathcal{U} , we have that $\mathcal{M}^{\mathcal{U}}$ is not $\text{pfc}(\mathcal{U})^+$ -saturated.*

Proof sketch. The proof relies on Shelah's "fcp theorem", which yields a formula $\varphi(x, y, z)$, where x, y , and z are finite tuples of variables, x and y having equal length, with the following two properties:

- for all $c \in M$, $\varphi(x, y, c)$ defines an equivalence relation;
- for all $n \in \mathbb{N}$, there is a tuple $c_n \in M$ such that $\varphi(x, y, c_n)$ has at least n , but only finitely many, equivalence classes.

Taking this fact for granted, the proof of the theorem is fairly straightforward. Indeed, suppose that $(n_i)_{i < \lambda}$ is such that $\text{pfc}(\mathcal{U}) = |\prod_{\mathcal{U}} [n_i]|$ and define $c \in M^I$ by setting $c(i) := c_{m_i}$, where m_i is the maximal $l \in \mathbb{N}$ such that $\varphi(x, y, c_l)$ has at most n_i classes. (If no such l exists, set $c(i) = 1$; note that this only happens for a \mathcal{U} -small set of $i \in I$.) It follows that $\varphi(x, y, [c]_{\mathcal{U}})$ defines an equivalence relation in $\mathcal{M}^{\mathcal{U}}$ with infinitely many, but at most $\text{pfc}(\mathcal{U})$ many classes. It follows that $\mathcal{M}^{\mathcal{U}}$ is not $\text{pfc}(\mathcal{U})^+$ -saturated. \square

Since there are stable theories with the fcp, Theorem 8.6.20 tells us that there must be some κ and some regular ultrafilter \mathcal{U} on κ such that $\text{pfc}(\mathcal{U}) \leq \kappa$. Shelah considered the question: which cardinals can arise as $\text{pfc}(\mathcal{U})$ for \mathcal{U} a regular ultrafilter on κ ? Clearly, $\text{pfc}(\mathcal{U}) \leq 2^\kappa$ and an argument analogous to the proof of Theorem 6.8.5 shows that $\text{pfc}(\mathcal{U})^{\aleph_0} = \text{pfc}(\mathcal{U})$. Surprisingly, these are the only two restrictions:

Theorem 8.6.25. *Suppose that κ is an infinite cardinal and λ is an infinite cardinal such that $\lambda \leq 2^\kappa$ and $\lambda^{\aleph_0} = \lambda$. Then there is a regular ultrafilter \mathcal{U} on κ such that $\text{pfc}(\mathcal{U}) = \lambda$.*

Returning to the study of Keisler's order, Theorems 8.6.20 and 8.6.24 imply the following:

Corollary 8.6.26. *Any two stable theories with the fcp are equivalent in Keisler's order.*

To see that every stable theory is below every unstable theory in Keisler's order, we need to know the following fact for unstable theories, again due to Shelah:

Theorem 8.6.27. *Suppose that T is unstable and \mathcal{U} is a countably incomplete ultrafilter with $\text{pcf}(\mathcal{U}) < 2^\lambda$. For any $\mathcal{M} \models T$, we have that $\mathcal{M}^{\mathcal{U}}$ is not λ^+ -saturated.*

Proof sketch: Throughout the proof, we fix $\mathcal{M} \models T$ and a sequence $(m_i)_{i \in I}$ of natural numbers such that $\text{pcf}(\mathcal{U}) = |\prod_{\mathcal{U}} [m_i]|$. The proof of the theorem relies on a fundamental fact from stability theory: if T is an unstable theory, then T has either the **strict order property** or T has the **independence property**.

First suppose that T has the strict order property. By definition, this means that there is a formula $\varphi(x, y)$, where x and y are finite tuples of the same length, which has the following two properties:

- $\varphi(x, y)$ defines a (strict) partial order on M , and
- for each $n \in \mathbb{N}$, there are $a_1(n), \dots, a_n(n) \in M$, for all $i = 1, \dots, n-1$, such that $\varphi(a_i(n), a_{i+1}(n))$.

Set $A_i := \{a_1(m(i)), \dots, a_{m(i)}(m(i))\}$ and set $A := \prod_{\mathcal{U}} A_i$. By the definition of $\text{pcf}(\mathcal{U})$, we have that $|A| < 2^\lambda$.

Suppose, toward a contradiction, that $\mathcal{M}^{\mathcal{U}}$ is λ^+ -saturated. This assumption allows us to construct, by induction on the length of η , $(a_\eta)_{\eta \in 2^{\leq \lambda}}, (b_\eta)_{\eta \in 2^{\leq \lambda}} \in A$, satisfying the following:

- For every $\eta \in 2^{\leq \lambda}$, we have $\mathcal{M}^{\mathcal{U}} \models \varphi(a_\eta, b_\eta)$.
- For every $n \in \mathbb{N}$, there are $c_0, \dots, c_n \in A$ such that $c_0 = a_\eta, c_n = b_\eta$, and $\mathcal{M}^{\mathcal{U}} \models \bigwedge_{i < n} \varphi(c_i, c_{i+1})$.
- If η is an initial segment of ν , then $\mathcal{M}^{\mathcal{U}} \models \varphi(a_\eta, a_\nu) \wedge \varphi(b_\nu, b_\eta) \wedge \varphi(b_{\eta \frown (0)}, a_{\eta \frown (1)})$.

It is clear from the properties of the construction that the a_η 's are distinct, whence $|A| \geq 2^\lambda$, yielding the desired contradiction.

Now suppose instead that T has the independence property. By definition, there is a formula $\varphi(x, y)$ such that, for every $n \in \mathbb{N}$, there are

$a_1(n), \dots, a_n(n) \in M$ for which, for all $J \subseteq [n]$, we have

$$\mathcal{M} \models \exists x \left(\bigwedge_{i \in J} \varphi(x, a_i(n)) \wedge \bigwedge_{i \notin J} \neg \varphi(x, a_i(n)) \right).$$

For each $J \subseteq [n]$, let $b_J(n) \in M$ witness the above existential statement. For each $i \in I$, set $k(i) := \lfloor \log_2(m_i) \rfloor$ if $m_i \neq 1$; if $m_i = 1$ (which happens for a \mathcal{U} -small set of $i \in I$), set $k(i) := 1$. Set $A_i := \{a_1(k(i)), \dots, a_{k(i)}(k(i))\}$ and set $A := \prod_{\mathcal{U}} A_i$. Similarly, set $B_i := \{b_J(k(i)) : J \subseteq [k(i)]\}$ and set $B = \prod_{\mathcal{U}} B_i$. Note that $\aleph_0 \leq |A| \leq |B| \leq \mu$, whence, by the definition of μ , we have $|A| = |B| = \mu$. Let $\eta := \min(\text{pfc}(\mathcal{U}), \lambda)$ and let $A' \subseteq A$ be such that $|A'| = \eta$. For each $C \subseteq A'$, let $\Sigma_C := \{\varphi(x, a) : a \in C\} \cup \{\neg \varphi(x, a) : a \notin C\}$. It is readily verified that each Σ_C is finitely satisfiable in $\mathcal{M}^{\mathcal{U}}$. Since each Σ_C mentions at most η parameters, in order to show that $\mathcal{M}^{\mathcal{U}}$ is not λ^+ -saturated, it suffices to show that some Σ_C is not realized in $\mathcal{M}^{\mathcal{U}}$. To see this, it suffices to note that if Σ_C is realized in $\mathcal{M}^{\mathcal{U}}$, then it is in fact realized by some element of B . Indeed, if this is the case, then at most $\text{pfc}(\mathcal{U})$ many of the Σ_C 's can be realized in $\mathcal{M}^{\mathcal{U}}$; since there are $2^\eta = 2^{\min(\text{pfc}(\mathcal{U}), \lambda)} > \mu$ many Σ_C 's, it follows that at least one of them is not realized in $\mathcal{M}^{\mathcal{U}}$, as desired. Thus, to finish, suppose that Σ_C is realized by $[c]_{\mathcal{U}} \in M^{\mathcal{U}}$. For each $i \in I$, set $w_i := \{e \in M : \mathcal{M} \models \varphi(c(i), e)\} \cap A_i$. It then follows that, for $[f]_{\mathcal{U}} \in A'$, we have $[f]_{\mathcal{U}} \in C$ if and only if $[f]_{\mathcal{U}} \in \prod_{\mathcal{U}} w_i$. Consequently, if we consider $[d]_{\mathcal{U}} \in B$ defined by $d(i) = b_{w_i}(k(i))$ for each $i \in I$, we have that $[d]_{\mathcal{U}}$ realizes Σ_C . \square

Exercise 8.6.28. Verify all of the details in the above proof sketch.

As noted earlier, if \mathcal{U} is a regular ultrafilter on λ , then the maximal value for $\text{pfc}(\mathcal{U})$ is 2^λ . Theorem 8.6.27 implies that as soon as $\text{pfc}(\mathcal{U})$ is less than maximal, then \mathcal{U} does not saturate any unstable theory. We should also note that the converse to Theorem 8.6.27 is not true, as will become evident from our discussion in Chapter 15.

Corollary 8.6.29. *Suppose that T_1 is stable and T_2 is unstable. Then $T_1 \trianglelefteq T_2$.*

Proof. Without loss of generality, assume that T_1 has the fcp. Suppose that \mathcal{U} is a regular ultrafilter on κ that does not saturate T_1 . By Theorem 8.6.20, $\text{pfc}(\mathcal{U}) \leq \kappa$; since $\kappa < 2^\kappa$, Theorem 8.6.27 implies that \mathcal{U} does not saturate T_2 either. \square

It is in fact true that, in the statement of the previous corollary, we have that $T_1 \triangleleft T_2$, whence the stable theories comprise the two smallest classes in Keisler's order. We will prove this fact in Corollary 15.2.13.

We now consider the opposite end of Keisler's order:

Example 8.6.30. $\text{Th}(\mathcal{P}_f(\omega), \subseteq)$ is a maximum element in Keisler's order. Indeed, let T be any complete theory. Suppose that \mathcal{U} is a regular ultrafilter that saturates $\text{Th}(\mathcal{P}_f(\omega), \subseteq)$. By Theorem 8.4.16, \mathcal{U} is κ^+ -good, whence it saturates T by Proposition 8.4.9. Consequently, $T \leq \text{Th}(\mathcal{P}_f(\omega), \subseteq)$.

Proposition 8.6.31. *The complete theory T is a maximum in Keisler's order if and only if, for every regular ultrafilter \mathcal{U} : \mathcal{U} saturates T if and only if \mathcal{U} is good.*

Proof. The backward direction is essentially the argument used in the previous example. For the forward direction, any maximum element is equivalent to $\text{Th}(\mathcal{P}_f(\omega), \subseteq)$ and the corresponding property of this theory is the content of Theorem 8.4.16. \square

Remark 8.6.32. Unlike the case of the minimum class, an exact syntactic description of the theories belonging to the maximum class is unknown. At the time of writing, the best known result is due to Malliaris and Shelah, who showed in [121] that theories with the so-called SOP_2 are in the maximum class. (They actually conjecture that SOP_2 characterizes the theories in the maximum class in the same way that the nfc p characterizes the theories in the minimum class.) Incidentally, the techniques used in [121] also settled the oldest problem in the study of cardinal characteristics of the continuum (see Section 1.5), namely by showing that $\mathfrak{p} = \mathfrak{t}$.

8.7. Notes and references

A good reference for further facts about saturation and universality is Marker's introductory textbook on model theory [126]. Theorem 8.2.1 is due to Keisler [100]. Part of our treatment on regular ultrafilters is based on the analogous discussion in [28]. Theorem 8.3.9 is due to Frayne, Morel, and Scott [59]. Theorem 8.3.10 and Theorem 8.3.16 are due to Keisler [101]. Good ultrafilters were introduced in [97] and [99], where their existence was proven using GCH. The use of GCH was removed by Kunen in [111]. Our presentation of the existence of good ultrafilters follows [28] very closely. The definition of good ultrafilter can appear quite strange at first sight, but the article [120] does a great job making the definition appear more intuitive; our presentation is inspired heavily by this latter article, although we try to clarify matters even further. Corollary 8.4.24 is due to Keisler [97].

After Keisler's original paper [101], besides the results of Shelah mentioned above, the study of Keisler's order remained somewhat dormant. Interest in the subject was renewed starting with Malliaris's thesis [119]. She then collaborated with Shelah to prove many more interesting results on the structure of Keisler's order (e.g., the result about theories with SOP_2

belonging to the maximum class mentioned above). They also proved that there are infinitely many classes in Keisler's order, and in fact the partial order induced on the set of classes has an infinite strictly descending chain [123].

Malliaris-Shelah [125] and Ulrich [178] (independently) proved, from a large cardinal assumption, that Keisler's order is not a linear order; later, Malliaris and Shelah were able to remove the large cardinal assumption [122]. Even more recently, Malliaris and Shelah were able to prove that there are the maximal number of equivalence classes of theories with respect to Keisler's order, that is, they found continuum-many theories that are pairwise inequivalent with respect to Keisler's order [124], solving a problem that had remained open since Keisler's original paper [101].

Nonstandard analysis

In this chapter, we develop some basic nonstandard analysis and highlight the connection with the ultraproduct construction. Section 9.1 introduces the basic idea behind nonstandard analysis and enumerates a naïve axiomatization for what properties a nonstandard universe should satisfy. Section 9.2 examines the kind of new numbers that appear in the nonstandard extension of the real field while Section 9.3 uses these new numbers to develop some basic nonstandard calculus. Section 9.4 makes explicit the connection with ultrapowers by giving the ultrapower model for nonstandard extensions. The main result of Section 9.5 is that the ultrapower model of nonstandard extensions is almost general in the sense that every nonstandard extension is locally isomorphic to an ultrapower, or, more precisely, is isomorphic to a limit ultrapower. Section 9.6 describes an extension of the previous setup suitable for studying the nonstandard extension of more complex objects such as topological spaces or measure spaces. Section 9.7 highlights the way in which nonstandard elements can be used to generate ultrafilters, a method which has been very successful recently in connection with combinatorial applications, some of which are sketched in this section. Finally, in Section 9.8, we examine the Hausdorff condition on ultrafilters, a condition which arises naturally in the context of the discussion from Section 9.7, and we ponder the existence of such ultrafilters.

9.1. Naïve axioms for nonstandard analysis

The initial approach to calculus made free use of “infinitesimal” and “infinite” elements. For example, if $f : \mathbb{R} \rightarrow \mathbb{R}$ is a function, then one might say $\lim_{x \rightarrow 0} f(x) = 0$ if, whenever a is an infinitesimal number, then $f(a)$ is also infinitesimal. As the ontological status of these infinitesimal elements

was called into question, calculus was not deemed to be on a firm foundation until the advent of the “ ϵ - δ ” style approach that we are accustomed to today.

However, in the 1960s, the model theorist Abraham Robinson realized that techniques from model theory could be used to rescue the use of these infinitesimal and infinite elements. Indeed, these elements, while not elements of the usual field \mathbb{R} , do belong in some field \mathbb{R}^* extending \mathbb{R} , often termed a field of **hyperreal numbers**. Moreover, this field \mathbb{R}^* behaves “logically” like the usual field \mathbb{R} of real numbers. This latter fact, often referred to as the *transfer principle*, is crucial in nonstandard arguments. (As we will soon see, the transfer principle, in some sense, is merely Łoś’s theorem in disguise.) After its inception, this technique, now known by the unfortunate name of nonstandard analysis, has been used to prove significant theorems in a wide variety of areas of mathematics. There have been a number of excellent books written about nonstandard analysis, so we make no attempt in this chapter to provide a complete introduction. We merely take the opportunity to point out some of the more important topics.

To begin our introduction to nonstandard analysis, we will work in a nonstandard universe \mathbb{R}^* that satisfies the following properties:

- (1) $(\mathbb{R}; +, \cdot, 0, 1)$ is an ordered subfield of $(\mathbb{R}^*; +, \cdot, 0, 1)$.
- (2) \mathbb{R}^* has a positive **infinitesimal** element, that is, there is $\epsilon \in \mathbb{R}^*$ such that $\epsilon > 0$ but $\epsilon < r$ for every $r \in \mathbb{R}^{>0}$.
- (3) For every $n \in \mathbb{N}$ and every function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, there is a “natural extension” $f : (\mathbb{R}^*)^n \rightarrow \mathbb{R}^*$. The natural extensions of the operations $+, \cdot : \mathbb{R}^2 \rightarrow \mathbb{R}$ coincide with the field operations in \mathbb{R}^* . Similarly, for every $A \subseteq \mathbb{R}^n$, there is a subset $A^* \subseteq (\mathbb{R}^*)^n$ such that $A^* \cap \mathbb{R}^n = A$.
- (4) \mathbb{R}^* , equipped with the above assignment of extensions of functions and subsets, “behaves logically” like \mathbb{R} .

We have chosen not to extend partial functions. This can be taken care of in two equivalent ways using the above setup:

Exercise 9.1.1. Suppose that $A \subseteq \mathbb{R}^n$ and let $f : A \rightarrow \mathbb{R}$. Consider the graph of f

$$\Gamma_f := \{(a, f(a)) : a \in A\} \subseteq \mathbb{R}^{n+1}.$$

Prove that:

- (1) Γ_f^* is the graph of a function from A^* to \mathbb{R}^* extending f .
- (2) For any (total) function $F : \mathbb{R}^n \rightarrow \mathbb{R}$, if $F \upharpoonright A = f$, then function from part (1) is $F \upharpoonright A^*$.

The unique function from the preceding exercise will be the nonstandard extension of our partial function f .

In particular, if $(s_n)_{n \in \mathbb{N}}$ is a sequence from \mathbb{R} , then viewing it as a function $s : \mathbb{N} \rightarrow \mathbb{R}$, we get the nonstandard extension $s : \mathbb{N}^* \rightarrow \mathbb{R}^*$. For $N \in \mathbb{N}^*$, we write $s(N)$ as s_N to remind us of the original sequence perspective.

9.2. Nonstandard numbers big and small

Since \mathbb{R}^* is an ordered field, we can start performing the field operations to our positive infinitesimal ϵ . For example, ϵ has an additive inverse $-\epsilon$, which is then a negative infinitesimal. Also, we can consider $\pi \cdot \epsilon$; it is reasonably easy to see that $\pi \cdot \epsilon$ is also a positive infinitesimal. (This will follow from a more general principle that we will shortly see.)

Since $\epsilon \neq 0$, it has a multiplicative inverse ϵ^{-1} . For a given $r \in \mathbb{R}^{>0}$, since $\epsilon < \frac{1}{r}$, we see that $\epsilon^{-1} > r$. Since r was an arbitrary positive real number, we see that ϵ^{-1} is a positive infinite element. And, of course, $-\epsilon^{-1}$ is a negative infinite element. But now we can continue playing, considering numbers like $\sqrt{2} \cdot \epsilon^{-1}$ and so on. . .

Besides algebraic manipulations, we also have transcendental matters to consider. Indeed, we have the nonstandard extension of the function $\sin : \mathbb{R}^* \rightarrow \mathbb{R}^*$; what is $\sin(\epsilon)$? All in due time. . . First, let us make precise some of the words we have been thus far freely tossing around.

Definition 9.2.1.

- (1) The set of **finite hyperreals** is

$$\mathbb{R}_{\text{fin}} := \{x \in \mathbb{R}^* : |x| \leq n \text{ for some } n \in \mathbb{N}\}.$$

- (2) The set of **infinite hyperreals** is $\mathbb{R}_{\text{inf}} := \mathbb{R}^* \setminus \mathbb{R}_{\text{fin}}$.

- (3) The set of **infinitesimal hyperreals** is

$$\mu := \{x \in \mathbb{R}^* : |x| \leq \frac{1}{n} \text{ for all } n \in \mathbb{N}^{>0}\}.$$

Observe that $\mu \subseteq \mathbb{R}_{\text{fin}}$, $\mathbb{R} \subseteq \mathbb{R}_{\text{fin}}$, and $\mu \cap \mathbb{R} = \{0\}$. Also note that if $\delta \in \mu \setminus \{0\}$, then $\delta^{-1} \in \mathbb{R}_{\text{inf}}$.

Exercise 9.2.2. Prove that \mathbb{R}_{fin} is a subring of \mathbb{R}^* and μ is an ideal in \mathbb{R}_{fin} .

Definition 9.2.3. For $x, y \in \mathbb{R}^*$, we say that x and y are **infinitely close**, denoted $x \approx y$, if $x - y \in \mu$.

Exercise 9.2.4.

- (1) Show that \approx is an equivalence relation on \mathbb{R}^* .

- (2) Show that \approx is a congruence relation on \mathbb{R}_{fin} ; that is, it is an equivalence relation on \mathbb{R}_{fin} such that, for all $x, y, u, v \in \mathbb{R}_{\text{fin}}$, if $x \approx u$ and $y \approx v$, then $x \pm y \approx u \pm v$ and $xy \approx uv$.

Exercise 9.2.5. Show that \mathbb{R} and μ are both nonempty subsets of \mathbb{R}^* that are bounded above but yet have no least upper bound. Consequently, the completeness property is not true for the ordered field \mathbb{R}^* .

Theorem 9.2.6 (The existence of standard parts). *If $r \in \mathbb{R}_{\text{fin}}$, then there is a unique $s \in \mathbb{R}$ such that $r \approx s$.*

Proof. Uniqueness is immediate: if $r \approx s_1$ and $r \approx s_2$ with $s_1, s_2 \in \mathbb{R}$, then $s_1 \approx s_2$, so $s_1 - s_2 \in \mu \cap \mathbb{R} = \{0\}$, whence $s_1 = s_2$.

We now prove existence. Without loss of generality, we may assume that $r > 0$. (Why?) We then set $A := \{x \in \mathbb{R} : x < r\}$. Since $r \in \mathbb{R}_{\text{fin}}$, A is bounded above. Also, $0 \in A$, so $A \neq \emptyset$. Thus, by the completeness property of \mathbb{R} , $\sup(A)$ exists. Set $s := \sup(A)$. We claim that this is the desired s . Toward this end, fix $\delta \in \mathbb{R}^{>0}$; we show that $|r - s| \leq \delta$. Since s is an upper bound for A , $s + \delta \notin A$, that is, $r \leq s + \delta$. If $r \leq s - \delta$, then $s - \delta$ would also be an upper bound for A , contradicting the fact that s was the least upper bound. Thus, $r > s - \delta$, and thus $|r - s| \leq \delta$, as desired. \square

Definition 9.2.7. Given $r \in \mathbb{R}_{\text{fin}}$, the unique $s \in \mathbb{R}$ such that $r \approx s$ is called the **standard part** of r and is denoted $\text{st}(r)$.

Exercise 9.2.8. Fix $x, y \in \mathbb{R}_{\text{fin}}$.

- (1) Prove that $x \approx y$ if and only if $\text{st}(x) = \text{st}(y)$.
- (2) If $x \leq y$, then $\text{st}(x) \leq \text{st}(y)$. Give a counterexample to the converse statement.
- (3) If $x \in \mathbb{R}$, then $\text{st}(x) = x$.

Exercise 9.2.9. Prove that $\text{st} : \mathbb{R}_{\text{fin}} \rightarrow \mathbb{R}$ is a surjective ring homomorphism with kernel μ . Conclude that $\mathbb{R}_{\text{fin}}/\mu \cong \mathbb{R}$.

We conclude this section by taking a closer look at \mathbb{N}^* .

Proposition 9.2.10.

- (1) $\mathbb{N}^* \setminus \mathbb{N} \neq \emptyset$.
- (2) If $N \in \mathbb{N}^* \setminus \mathbb{N}$, then N is a positive infinite element.

Proof. For (1), fix $y \in \mathbb{R}^*$ a positive infinite element. Since the statement “for all $x \in \mathbb{R}$, if $x > 0$, then there is $n \in \mathbb{N}$ such that $x \leq n$ ” is true in \mathbb{R} , the statement “for all $x \in \mathbb{R}^*$, if $x > 0$, then there is $n \in \mathbb{N}^*$ such that $x \leq n$ ” is true in \mathbb{R}^* by the transfer principle. Thus, there is $N \in \mathbb{N}^*$ such

that $y \leq N$. If $N \in \mathbb{N}$, then $y \in \mathbb{R}_{\text{fin}}$, a contradiction. Thus, $N \in \mathbb{N}^* \setminus \mathbb{N}$, as desired.

For (2), we note that $N \geq 0$ by the transfer principle. We next note that if $N \in \mathbb{R}_{\text{fin}}$, then there is $n \in \mathbb{N}$ such that $n \leq N \leq n + 1$. However, the statement “for all $m \in \mathbb{N}$, if $n \leq m \leq n + 1$, then $m = n$ or $m = n + 1$ ” is true in \mathbb{R} , whence, by the transfer principle, we conclude that $N = n$ or $N = n + 1$, a contradiction. Consequently, $N \notin \mathbb{R}_{\text{fin}}$, that is, N is a positive infinite element. \square

We next examine how many nonstandard natural numbers there are. Toward this end, for $N \in \mathbb{N}^*$ (potentially standard), we set $\gamma(N) := \{N \pm m : m \in \mathbb{N}\}$ and call this the **galaxy** or **archimedean class** of N . Clearly, $N \in \mathbb{N}$ if and only if $\gamma(N) = \mathbb{Z}$; this is called the **finite galaxy**, while all other galaxies are referred to as **infinite galaxies**.

Exercise 9.2.11.

- (1) If $N \in \mathbb{N}^* \setminus \mathbb{N}$, then $\gamma(N) \subseteq \mathbb{N}^*$.
- (2) For $M, N \in \mathbb{N}^*$, $\gamma(M) = \gamma(N)$, if and only if $|M - N| \in \mathbb{N}$.
- (3) If $\gamma(M) = \gamma(M')$ and $\gamma(N) = \gamma(N')$ and $\gamma(M) \neq \gamma(N)$, then $M < N$ if and only if $M' < N'$.

The last item allows us to define an ordering on galaxies: if $\gamma(M) \neq \gamma(N)$, then declare $\gamma(M) < \gamma(N)$ if and only if $M < N$; in this case, we often say that M is **infinitely less** than N .

Exercise 9.2.12. Prove that there is no largest nor smallest infinite galaxy and that between any two infinite galaxies lies a third infinite galaxy.

9.3. Some nonstandard calculus

Theorem 9.3.1. *Suppose that $a, L \in \mathbb{R}$ and $f : \mathbb{R} \rightarrow \mathbb{R}$ is a function. Then $\lim_{x \rightarrow a} f(x) = L$ if and only if, for all $y \in \mathbb{R}^*$, if $y \approx a$ but $y \neq a$, then $f(y) \approx L$.*

Proof. First suppose that $\lim_{x \rightarrow a} f(x) = L$ and take $y \in \mathbb{R}^*$ with $y \approx a$ and $y \neq a$. Fix $\epsilon > 0$ and take $\delta \in \mathbb{R}^{>0}$ such that, for all $b \in \mathbb{R}$, if $0 < |x - b| < \delta$, then $|f(x) - L| < \epsilon$. Since $y \approx a$ but $y \neq a$, we have that $0 < |y - a| < \delta$, whence, by transfer, we have that $|f(y) - L| < \epsilon$. Since $\epsilon > 0$ was arbitrary, we have that $f(y) \approx L$, as desired.

Now assume that the converse assumption holds and fix $\epsilon > 0$. In the nonstandard extension, the statement “there is $\delta > 0$ such that, for all $y \in \mathbb{R}^*$, if $0 < |x - y| < \delta$, then $|f(x) - L| < \epsilon$ ” is true, as one can let δ be any positive infinitesimal. Thus, the corresponding statement is true in \mathbb{R} , as desired. \square

Exercise 9.3.2. Give nonstandard proofs of all of the usual limit laws from calculus. For example, if $\lim_{x \rightarrow a} f(x) = L$ and $\lim_{x \rightarrow a} g(x) = M$, give a nonstandard proof of the fact that $\lim_{x \rightarrow a} (f \pm g)(x) = L \pm M$.

Corollary 9.3.3. *Given $f : \mathbb{R} \rightarrow \mathbb{R}$ and $a \in \mathbb{R}$, then f is continuous at a if and only if, whenever $x \approx a$, then $f(x) \approx f(a)$.*

Theorem 9.3.4 (Intermediate value theorem). *Suppose that f is continuous on $[a, b]$. Then for every d in between $f(a)$ and $f(b)$, there is $c \in (a, b)$ such that $f(c) = d$.*

Proof. Without loss of generality, suppose $f(a) < f(b)$, so $f(a) < d < f(b)$. Define a sequence $(s_n)_{n \in \mathbb{N}}$ as follows. For $n > 0$, let $\{p_0, \dots, p_n\}$ denote the partition of $[a, b]$ into n equal pieces of width $\frac{b-a}{n}$. Since $f(p_0) < d$, we can define the number $s_n := \max\{p_k : f(p_k) < d\}$, that is, p_k is the last partition point for which $f(p_k) < d$. Note that $s_n < b$.

Now fix $N \in \mathbb{N}^* \setminus \mathbb{N}$. We claim that $c := \text{st}(s_N)$ is as desired, that is, that $f(c) = d$. (Note that $s_N \in [a, b]^*$ by transfer whence $\text{st}(s_N)$ is defined.) Indeed, since $s_N < b$ by transfer, we have that $s_N + \frac{b-a}{N} \leq b$, whence, by transfer again, $d \leq f(s_N + \frac{b-a}{N})$. However, since $s_N + \frac{b-a}{N} \approx s_N \approx c$, we have that

$$f(c) \approx f(s_N) < d \leq f\left(s_N + \frac{b-a}{N}\right) \approx f(c),$$

whence $f(c) \approx d$. Since $f(c), d \in \mathbb{R}$, we get that $f(c) = d$. □

Theorem 9.3.5 (Extreme value theorem). *Suppose that $f : [a, b] \rightarrow \mathbb{R}$ is continuous. Then there are $c, d \in [a, b]$ such that $f(c) \leq f(x) \leq f(d)$ for all $x \in [a, b]$.*

Proof. We only prove the existence of the maximum. Let $\{p_0, \dots, p_n\}$ be as above. This time define s_n to be some partition point p_k such that $f(p_j) \leq f(p_k)$ for all $j = 0, \dots, n$. Fix $N \in \mathbb{N}^* \setminus \mathbb{N}$ and set $d := \text{st}(s_N)$. We claim that $f(d)$ is a maximum. (The intuition here is that we are breaking $[a, b]$ up into hyperfinitely many pieces of infinitesimal width, considering the maximum value f takes on these hyperfinitely many elements, and then noting that that value is infinitely close to the maximum value of the original function.)

Take $x \in [a, b]$. By transfer, there is $k \in \mathbb{N}^*$ with $0 \leq k < N$ such that $x \in [a + \frac{k(b-a)}{N}, a + \frac{(k+1)(b-a)}{N}]$. By continuity of f , we have that $f(a + \frac{k(b-a)}{N}) \approx f(x)$ while $f(a + \frac{k(b-a)}{N}) \leq f(s_N)$. Since $f(s_N) \approx f(d)$, it follows that $f(x) \leq f(d)$. □

The nonstandard perspective is particularly well suited for explaining the difference between continuity and uniform continuity:

Theorem 9.3.6. *Suppose that I is an interval. Then $f : I \rightarrow \mathbb{R}$ is uniformly continuous if and only if, for all $x, y \in I^*$, if $x \approx y$, then $f(x) \approx f(y)$.*

Exercise 9.3.7. Prove Theorem 9.3.6.

Thus, the nonstandard explanation for the difference between continuity and uniform continuity is that continuity requires one of the elements in the pair of infinitely close numbers to be standard.

Example 9.3.8. $f : (0, 1) \rightarrow \mathbb{R}$ given by $f(x) = \frac{1}{x}$ is continuous but not uniformly continuous. Indeed, for $N \in \mathbb{N}^* \setminus \mathbb{N}$, $\frac{1}{N} \approx \frac{1}{N+1}$, but $f(\frac{1}{N}) = N \not\approx N + 1 = f(\frac{1}{N+1})$.

Theorem 9.3.9. *If $f : [a, b] \rightarrow \mathbb{R}$ is continuous, then it is uniformly continuous.*

Proof. Fix $x, y \in [a, b]^*$ with $x \approx y$. Let $c := \text{st}(x) = \text{st}(y) \in [a, b]$. Then f being continuous at c implies $f(x) \approx f(c) \approx f(y)$, as desired. \square

Exercise 9.3.10. Suppose that $f : [0, 1] \rightarrow \mathbb{R}$ is Riemann integrable. Then for every infinite N , one has

$$\int_0^1 f(x) dx \approx \sum_{i=0}^{N-1} f\left(\frac{i}{N}\right) \frac{1}{N}.$$

(*Hint.* Part of the exercise is making sense of the right-hand side.)

9.4. Ultrapowers as a model of nonstandard analysis

In this section, we show how ultrapowers can be used to put the discussion in the previous sections on firm footing. Fix $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. We show that $\mathbb{R}^* := \mathbb{R}^{\mathcal{U}}$ serves as a suitable model for our nonstandard extension.

We first note that, by Łoś's theorem, $\mathbb{R}^{\mathcal{U}}$ is an ordered field, and the diagonal embedding allows us to view \mathbb{R} as a subfield of $\mathbb{R}^{\mathcal{U}}$.

In order to get a positive infinitesimal element, set $\epsilon := [(1, \frac{1}{2}, \frac{1}{3}, \dots)]_{\mathcal{U}}$. Since $\frac{1}{n} > 0$ for all n , $\epsilon > 0$. In order to see that ϵ is infinitesimal, fix $r \in \mathbb{R}^{>0}$; we then have that $\frac{1}{n} < r$ for all but finitely many n . Since cofinite sets belong to \mathcal{U} (as \mathcal{U} is nonprincipal), we have that $\frac{1}{n} < r$ for \mathcal{U} -almost all n , whence $\epsilon < r$.

We define nonstandard extensions of functions and sets as before: given $f : \mathbb{R}^n \rightarrow \mathbb{R}$, we define $f : (\mathbb{R}^n)^{\mathcal{U}} \rightarrow \mathbb{R}^{\mathcal{U}}$ by $f([a_1]_{\mathcal{U}}, \dots, [a_n]_{\mathcal{U}}) := [i \mapsto f(a_1(i), \dots, a_n(i))]_{\mathcal{U}}$ and $A^* := A^{\mathcal{U}}$.

Finally, when we say that $\mathbb{R}^{\mathcal{U}}$ “behaves logically like” \mathbb{R} , we are really referring to the fact that Łoś's theorem holds. Although to be truthful, we have not really specified a first-order language, so this is a bit vague.

We will pin down precisely the first-order formalism for doing this in the next section, but, in theory, one could avoid all discussions of first-order logic if one is willing to constantly reprove Loś's theorem during each argument (although we do not recommend it). We illustrate this by giving the nonstandard characterization of limit from Theorem 9.3.1 above in the ultrapower language.

First suppose that $\lim_{x \rightarrow a} f(x) = L$ and $x \approx a$, $x \neq a$. We show that $f(x) \approx L$. Fix $\epsilon > 0$ and take $\delta > 0$ such that, for all $y \in \mathbb{R}$, if $0 < |y - a| < \delta$, then $|f(y) - L| < \epsilon$. Since $x \approx a$ but $x \neq a$, this means that $0 < |x(n) - a| < \delta$ for \mathcal{U} -almost all n . Consequently, for \mathcal{U} -almost all n , we have $|f(x(n)) - L| < \epsilon$, and thus $|f(x) - L| < \epsilon$.

We prove the converse by contrapositive: suppose that $\lim_{x \rightarrow a} f(x) \neq L$. Then there is $\epsilon > 0$ such that, for all $n \in \mathbb{N}$, there is $x(n)$ with $0 < |x(n) - a| < \frac{1}{n}$ and yet $|f(x(n)) - L| \geq \epsilon$. Then $[x]_{\mathcal{U}} \approx a$, $[x]_{\mathcal{U}} \neq a$, and yet $|f([x]_{\mathcal{U}}) - L| \geq \epsilon$.

9.5. Complete extensions and limit ultrapowers

In order to apply nonstandard analysis to more complicated situations, one needs to be able to work with nonstandard extensions of objects besides \mathbb{R} . In the next two sections, we explain the model-theoretic approach to this endeavor.

Definition 9.5.1. Suppose that \mathcal{M} is an \mathcal{L} -structure. Let $\mathcal{L}_M^\#$ denote the language obtained from \mathcal{L} by adding the following symbols:

- For each $a \in M$, a constant symbol c_a .
- For each n -ary function $f : M^n \rightarrow M$, an n -ary function symbol F_f .
- For each set $A \subseteq M^n$, an n -ary relation symbol R_A .

One then considers the expansion $\mathcal{M}^\#$ of \mathcal{M} to an $\mathcal{L}_M^\#$ -structure by interpreting c_a as a , F_f as f , and R_A as A . This is called the **complete expansion of \mathcal{M}** . If M is just a set, then we write $M^\#$ for the complete expansion of the structure \mathcal{M} in the empty language, and we call $M^\#$ the **complete structure on the set M** .

Definition 9.5.2. Suppose that \mathcal{M} and \mathcal{N} are \mathcal{L} -structures and $i : \mathcal{M} \rightarrow \mathcal{N}$ is an embedding. We say that i is a **complete embedding** if there is an expansion $\mathcal{N}^\#$ of \mathcal{N} to an $\mathcal{L}_M^\#$ -structure such that i is an embedding $i : \mathcal{M}^\# \rightarrow \mathcal{N}^\#$. When i is the inclusion map (so \mathcal{M} is a substructure of \mathcal{N}), we say that \mathcal{N} is a **complete extension of \mathcal{M}** .

Exercise 9.5.3. Suppose that $i : \mathcal{M} \rightarrow \mathcal{N}$ is a complete embedding. Prove that $i : \mathcal{M}^\# \rightarrow \mathcal{N}^\#$ is an elementary embedding.

Exercise 9.5.4. Given a structure \mathcal{M} and an ultrafilter \mathcal{U} , prove that the diagonal embedding $d : \mathcal{M} \rightarrow \mathcal{M}^{\mathcal{U}}$ is a complete embedding.

Exercise 9.5.5. Suppose that \mathbb{R}^* is a proper complete extension of \mathbb{R} . Prove that \mathbb{R}^* satisfies all of the “naïve axioms” introduced in Section 9.1.

Exercise 9.5.5 tells us that one can do nonstandard analysis by working in a proper complete extension of \mathbb{R} , while Exercise 9.5.4 tells us that this proper complete extension can be achieved using an ultrapower of \mathbb{R} . However, there is nothing special about \mathbb{R} here; if one wants to study some mathematical object, say a group G , then one merely needs to pass to a proper complete extension of that object, which can be obtained by using an ultrapower.

It is natural to wonder if, conversely, every complete extension of a structure \mathcal{M} is isomorphic to an ultrapower of \mathcal{M} . This is unfortunately not the case:

Exercise 9.5.6. Suppose that \mathcal{N} is an iterated ultrapower of \mathcal{M} (see Section 6.9). Prove that \mathcal{N} is a complete extension of \mathcal{M} .

Recall that Proposition 6.9.9 shows that, given any countable structure \mathcal{M} , there is an iterated ultrapower of \mathcal{M} not isomorphic to any ordinary ultrapower of \mathcal{M} , whence there are complete extensions of \mathcal{M} that are not obtainable using an ultrapower.

It is natural to wonder if every complete extension of \mathcal{M} is isomorphic to an iterated ultrapower of \mathcal{M} . At the time of the writing of this book, the answer to this question is unknown. However, a small modification of the notion of iterated ultrapower will lead us to all complete extensions. First, a preliminary notion:

Definition 9.5.7. Suppose that \mathcal{M} is an \mathcal{L} -structure and \mathcal{U} and \mathcal{V} are ultrafilters on sets I and J , respectively. Given a function $\sigma : J \rightarrow I$ such that $\sigma(\mathcal{V}) = \mathcal{U}$, we get an **induced embedding** $\tilde{\sigma} : \mathcal{M}^{\mathcal{U}} \rightarrow \mathcal{M}^{\mathcal{V}}$ given by $\tilde{\sigma}([f]_{\mathcal{U}}) = [f \circ \sigma]_{\mathcal{V}}$. An embedding $\mathcal{M}^{\mathcal{U}} \rightarrow \mathcal{M}^{\mathcal{V}}$ is called **induced** if it is of the form $\tilde{\sigma}$ for some function $\sigma : J \rightarrow I$ with $\sigma(\mathcal{V}) = \mathcal{U}$.

Exercise 9.5.8. Suppose that $\tilde{\sigma} : \mathcal{M}^{\mathcal{U}} \rightarrow \mathcal{M}^{\mathcal{V}}$ is an induced embedding. Prove that $\tilde{\sigma}$ is an $\mathcal{L}_{\mathcal{M}}^{\#}$ -elementary embedding (where each ultrapower is equipped with its “canonical” expansion to an $\mathcal{L}_{\mathcal{M}}^{\#}$ -structure).

Exercise 9.5.9. Show that the composition of two induced embeddings is induced.

Exercise 9.5.10. Show that the map $\mathcal{M}^{\mathcal{U}} \rightarrow \mathcal{M}^{\mathcal{U} \times \mathcal{V}}$ from Theorem 6.9.1 is induced.

Definition 9.5.11. Suppose that \mathcal{M} is a structure. An **ultrapower system over \mathcal{M}** is a directed family $(\mathcal{M}_i)_{i \in I}$ of ultrapower extensions of \mathcal{M} such that all maps in the system are induced. A **limit ultrapower of \mathcal{M}** is the direct limit of an ultrapower system over \mathcal{M} .

Exercise 9.5.12. Show that every iterated ultrapower of \mathcal{M} is a limit ultrapower of \mathcal{M} .

Exercise 9.5.13. Show that limit ultrapowers of \mathcal{M} are complete extensions of \mathcal{M} .

Exercise 9.5.14. Show that $\mathcal{M} \equiv \mathcal{N}$ if and only if they have isomorphic limit ultrapowers.

We will soon see that the converse of Exercise 9.5.13 is true. First, we need a different presentation of limit ultrapowers that will be convenient for proving the converse.

Given a function $g : I \rightarrow M$, we set

$$\text{eq}(g) := \{(i, j) \in I \times I : g(i) = g(j)\}.$$

Definition 9.5.15. Suppose that M is a set, \mathcal{U} is an ultrafilter on I , and \mathcal{F} is a filter on $I \times I$. We define the set $M^{\mathcal{U}|\mathcal{F}}$ to be

$$M^{\mathcal{U}|\mathcal{F}} := \{[g]_{\mathcal{U}} : g \in M^I, \text{eq}(g) \in \mathcal{F}\}.$$

Exercise 9.5.16. For any structure \mathcal{M} , any ultrafilter \mathcal{U} on a set I , and any filter \mathcal{F} on $I \times I$, prove that $M^{\mathcal{U}|\mathcal{F}}$ is the universe of a substructure of $\mathcal{M}^{\mathcal{U}}$.

Definition 9.5.17. With the notation of the previous exercise, we let $\mathcal{M}^{\mathcal{U}|\mathcal{F}}$ denote the substructure of $\mathcal{M}^{\mathcal{U}}$ with universe $M^{\mathcal{U}|\mathcal{F}}$.

Exercise 9.5.18. Prove that $\mathcal{M}^{\mathcal{U}|\mathcal{F}}$ is a complete extension of \mathcal{M} .

Here are two “degenerate” examples:

Example 9.5.19. Suppose $\mathcal{F} = \mathcal{P}(I \times I)$. Then $\mathcal{M}^{\mathcal{U}|\mathcal{F}} = \mathcal{M}^{\mathcal{U}}$.

Example 9.5.20. Suppose $\mathcal{F} = \{I \times I\}$. Then $\mathcal{M}^{\mathcal{U}|\mathcal{F}}$ equals the diagonal image of \mathcal{M} in $\mathcal{M}^{\mathcal{U}}$.

Exercise 9.5.21. Suppose that \mathcal{M} is a structure, \mathcal{U} is an ultrafilter on I , and \mathcal{F} and \mathcal{G} are filters on $I \times I$ with $\mathcal{F} \subseteq \mathcal{G}$. Prove that $\mathcal{M}^{\mathcal{U}|\mathcal{F}}$ is an $\mathcal{L}_M^{\#}$ -elementary substructure of $\mathcal{M}^{\mathcal{U}|\mathcal{G}}$.

The following example will be crucial for proving Theorem 9.5.25:

Example 9.5.22. Suppose that \mathcal{F} is a principal filter on $I \times I$, that is, suppose that there is $X \subseteq I \times I$ such that, for all $Y \subseteq I \times I$, $Y \in \mathcal{F}$ if and only if $X \subseteq Y$. We claim that, in this case, $\mathcal{M}^{\mathcal{U}|\mathcal{F}}$ is isomorphic to an ordinary ultrapower of \mathcal{M} . To see this, we first note that, for $g : I \rightarrow M$, we have that $\text{eq}(g) \in \mathcal{F}$ if and only if $g(i) = g(j)$ whenever $(i, j) \in X$. Let E_X be the equivalence relation on I generated by the relation X , that is, the smallest equivalence relation E on I for which $E(a, b)$ holds whenever $(a, b) \in X$. It follows that if $g : I \rightarrow M$ is such that $\text{eq}(g) \in \mathcal{F}$, then g is constant on E_X -equivalence classes. We let $I_X := I/E_X$, the set of E_X -equivalence classes, and let $\pi_X : I \rightarrow I_X$ be the canonical quotient map. We now define $\mathcal{U}_X := \pi_X(\mathcal{U})$, an ultrafilter on I_X . For $g : I \rightarrow M$ for which $\text{eq}(g) \in \mathcal{F}$, we define a new map $g_X : I_X \rightarrow M$ by declaring $g_X([i]_{E_X}) :=$ the constant value g takes on the equivalence class of i . It follows that the map which sends, for $g : I \rightarrow M$ with $\text{eq}(g) \in \mathcal{F}$, $[g]_{\mathcal{U}}$ to $[g_X]_{\mathcal{U}_X}$, is an isomorphism $\mathcal{M}^{\mathcal{U}|\mathcal{F}}$ to $\mathcal{M}^{\mathcal{U}_X}$.

Exercise 9.5.23. Verify all of the details of Example 9.5.22.

Exercise 9.5.24. Fix $X, Y \subseteq I \times I$ with $X \subseteq Y$. In the setting of the previous example, prove that there is an induced embedding $\mathcal{M}^{\mathcal{U}_Y} \rightarrow \mathcal{M}^{\mathcal{U}_X}$.

We are now ready to prove the main result of this section:

Theorem 9.5.25. *Suppose that \mathcal{M} and \mathcal{N} are \mathcal{L} -structures with $\mathcal{M} \subseteq \mathcal{N}$. The following are equivalent:*

- (1) \mathcal{N} is a complete extension of \mathcal{M} .
- (2) There is an ultrafilter \mathcal{U} on a set I and a filter \mathcal{F} on $I \times I$ such that \mathcal{N} is isomorphic to $\mathcal{M}^{\mathcal{U}|\mathcal{F}}$ over \mathcal{M} .
- (3) There is a limit ultrapower \mathcal{M}' of \mathcal{M} such that \mathcal{N} is isomorphic to \mathcal{M}' over \mathcal{M} .

Proof. (1) (\Rightarrow) (2): Let $\mathcal{N}^\#$ be an $\mathcal{L}_M^\#$ -expansion of \mathcal{N} such that $\mathcal{M}^\# \subseteq \mathcal{N}^\#$. By Exercise 9.5.3, $\mathcal{M}^\# \equiv \mathcal{N}^\#$. By Theorem 8.3.10, there is an ultrafilter \mathcal{U} on an index set I such that there is an elementary embedding $\pi : \mathcal{N}^\# \rightarrow (\mathcal{M}^\#)^\mathcal{U}$. Note that, for $a \in M$, we have that $\pi(a) = d(a)$ as elements of \mathcal{M} are named in $\mathcal{L}^\#$.

Set C to be the range of π . Let \mathcal{F} be the filter over $I \times I$ generated by the set $\{\text{eq}(f) : f : I \rightarrow M, [f]_{\mathcal{U}} \in C\}$. It suffices to show that $C = M^{\mathcal{U}|\mathcal{F}}$. Since it is clear that $C \subseteq M^{\mathcal{U}|\mathcal{F}}$, we prove the other inclusion. Fix $[g]_{\mathcal{U}} \in M^{\mathcal{U}|\mathcal{F}}$. By assumption, there is $f : I \rightarrow M$ such that $f \equiv_{\mathcal{U}} g$ and $\text{eq}(f) \in \mathcal{F}$. Take $h_1, \dots, h_n : I \rightarrow M$ such that $[h_i]_{\mathcal{U}} \in C$ for each $i = 1, \dots, n$ and such that $\bigcap_{i=1}^n \text{eq}(h_i) \subseteq \text{eq}(f)$. Fix a function $G : M^n \rightarrow M$ such that, for all $i \in I$, $G(h_1(i), \dots, h_n(i)) = f(i)$ for all $i \in I$. It follows that $G^{(\mathcal{M}^\#)^\mathcal{U}}([h_1]_{\mathcal{U}}, \dots, [h_n]_{\mathcal{U}}) = [f]_{\mathcal{U}}$, whence $[g]_{\mathcal{U}} = [f]_{\mathcal{U}} \in C$, as desired.

(2) (\Rightarrow) (3): It suffices to show that $\mathcal{M}^{\mathcal{U}|\mathcal{F}}$ is a limit ultrapower. For each $X \in \mathcal{F}$, let \mathcal{F}_X be the principal filter on $I \times I$ generated by X . Then $\mathcal{M}^{\mathcal{U}|\mathcal{F}_X}$ is isomorphic to the ordinary ultrapower $\mathcal{M}^{\mathcal{U}_X}$ by Example 9.5.22. Moreover, if $Y \in \mathcal{F}$, then there are induced embeddings $\mathcal{M}^{\mathcal{U}_X} \rightarrow \mathcal{M}^{\mathcal{U}_{X \cap Y}}$ and $\mathcal{M}^{\mathcal{U}_Y} \rightarrow \mathcal{M}^{\mathcal{U}_{X \cap Y}}$ by Exercise 9.5.24. It follows that the family $(\mathcal{M}^{\mathcal{U}_X})_{X \in \mathcal{F}}$ is an ultrapower system over \mathcal{M} . It remains to see that the direct limit of this system is isomorphic to $\mathcal{M}^{\mathcal{U}|\mathcal{F}}$. To see this, given $[b]_{\mathcal{U}} \in \mathcal{M}^{\mathcal{U}|\mathcal{F}}$, set $X := \{(i, j) \in I \times I : b(i) = b(j)\}$; we then have that $X \in \mathcal{F}$ and $[b]_{\mathcal{U}} \in \mathcal{M}^{\mathcal{U}|\mathcal{F}_X}$, as desired.

(3) (\Rightarrow) (1) is the statement of Exercise 9.5.13. □

As mentioned earlier in this section, the following question remains open:

Question 9.5.26. Is every limit ultrapower isomorphic to an iterated ultrapower?

9.6. Many-sorted structures and internal sets

In many areas of mathematics, we study many different sets at a time as well as functions between these various sets.

Example 9.6.1. A vector space is a set V together with two functions: vector addition, which is a function $+$: $V \times V \rightarrow V$, and scalar multiplication, which is a function \cdot : $\mathbb{F} \times V \rightarrow V$, where \mathbb{F} is some field.

Example 9.6.2. A metric space is a set X together with a metric, which is a function d : $X \times X \rightarrow \mathbb{R}$ satisfying certain first-order axioms.

Example 9.6.3. A measure space is a triple (X, \mathcal{B}, μ) , where X is a set, \mathcal{B} is a σ -algebra of subsets of X (so $B \subseteq \mathcal{P}(X)$), and μ : $\mathcal{B} \rightarrow \mathbb{R}$ is a measure.

We now develop a nonstandard framework suitable for studying such situations. Before, we were working with a structure consisting of just a single “sort”, namely a sort for M . Now, we will work in a structure \mathcal{M} with a (nonempty) collection of sorts S . For each $s \in S$, we have a set M_s , the universe of the sort s in \mathcal{M} . So, for example, in the linear algebra situation, we might have $S = \{s, t\}$, with $M_s = V$ and $M_t = \mathbb{F}$. Often we will write a many-sorted structure as $M = (M_s : s \in S)$. Thus, we might write the linear algebra example as (V, \mathbb{F}) , suppressing mention of the names of the sorts. For any finite sequence $\vec{s} = (s_1, \dots, s_n)$ of sorts, we have the product set $M_{\vec{s}} := M_{s_1} \times \dots \times M_{s_n}$.

We now consider a language which was just as expressive as before. Namely, we have:

- For every finite sequence \vec{s} of sorts and every $A \subseteq M_{\vec{s}}$, we have a predicate symbol P_A .

- For every sort s and every $a \in M_s$, we have a constant symbol c_a .
- For every finite sequence \vec{s} of sorts, every sort t , and every function $f : M_{\vec{s}} \rightarrow M_t$, we have a function symbol F_f .

This is the many-sorted analogue of taking the complete language for the standard structure as done in the one-sorted case in the previous section.

One now builds terms and formulae just as in ordinary logic, with the understanding that each sort comes equipped with its own collection of variables. If we need to be clear, we might decorate a variable with the name of the sort it is intended to range over, e.g., x^s .

Returning to the vector space example, let us see how we might write the distributive law $c \cdot (x + y) = c \cdot x + c \cdot y$. Recall that $S = \{s, t\}$, with $M_s = V$ and $M_t = \mathbb{F}$. Let $f : M_s \times M_s \rightarrow M_s$ denote vector addition and $g : M_t \times M_s \rightarrow M_s$ denote scalar multiplication. Then the axiom for the distributive law would be written as:

$$\forall x^s \forall y^s \forall z^t g(z, f(x, y)) = f(g(z, x), g(z, y)).$$

Of course, for the purpose of sanity, in practice we will continue to write things as they might naturally be written in ordinary mathematics; however, one must be aware of the formal way that such sentences would be written.

As in the one-sorted case of the previous section, one can take a complete extension of a many-sorted structure and this complete extension will serve as our nonstandard universe. As before, taking an ultrapower of the many-sorted structure can serve as such a complete extension.

As we saw with the measure theory example, it will often be convenient to have a sort for $\mathcal{P}(X)$ whenever X is itself a sort. For simplicity of the coming discussion, let us consider the many-sorted structure $(X, \mathcal{P}(X))$; what we say now is easily adapted to the more general situation that X and $\mathcal{P}(X)$ are sorts in a many-sorted structure containing other sorts.

We have the nonstandard extension $(X, \mathcal{P}(X)) \subseteq (X^*, \mathcal{P}(X)^*)$. We must be careful not to confuse $\mathcal{P}(X)^*$ with $\mathcal{P}(X^*)$, the latter retaining its usual meaning as the set of subsets of X^* . At the moment, $\mathcal{P}(X)^*$ is some abstract set, perhaps having no affiliation with an actual powerset. We now discuss how to relate $\mathcal{P}(X)^*$ and $\mathcal{P}(X^*)$. Set $E = \{(x, A) \in X \times \mathcal{P}(X) \mid x \in A\}$, the symbol for the membership relation.

Lemma 9.6.4. *We may assume that our nonstandard extension is such that $\mathcal{P}(X)^* \subseteq \mathcal{P}(X^*)$ and E^* is the membership relation restricted to $X^* \times \mathcal{P}(X)^*$.*

Proof. For $A \in \mathcal{P}(X)^*$, set $\Phi(A) := \{x \in X^* : (x, A) \in E^*\}$. We claim that $\Phi : \mathcal{P}(X)^* \rightarrow \mathcal{P}(X^*)$ is injective. Indeed, suppose that $A_1, A_2 \in \mathcal{P}(X)^*$ are such that $A_1 \neq A_2$. By the transfer principle, there is, without loss

of generality, $x \in X^*$ such that $(x, A_1) \in E^*$ but $(x, A_2) \notin E^*$. Then $x \in \Phi(A_1) \setminus \Phi(A_2)$, whence $\Phi(A_1) \neq \Phi(A_2)$.

Now one makes $(X, \Phi(\mathcal{P}(X)^*))$ into a structure so that the map $(x, A) \mapsto (x, \Phi(A))$ is an isomorphism. \square

Remark 9.6.5. In the preceding sections of this chapter, one often blurred the distinction between an element of the standard universe and its image in the nonstandard extension. This is tantamount to identifying a structure with its image in an ultrapower via the diagonal embedding. However, this identification is no longer possible for higher-order objects such as $A \in \mathcal{P}(X)$, for there is a drastic difference between a subset A of X and its nonstandard extension A^* .

There is some potential confusion that we should clear up now. Suppose that $A \subseteq X$. Then we have $A^* \subseteq X^*$ from the interpretation of the symbol P_A . However, $A \in \mathcal{P}(X)$, so it is mapped by the embedding to an element of $\mathcal{P}(X)^*$, which we temporarily denote by $(A)^*$. Fortunately, all is well:

Lemma 9.6.6. $A^* = (A)^*$.

Proof. By the transfer principle, we have that, for $x \in X^*$, $x \in A^*$ if and only if $(X^*, \mathcal{P}(X)^*) \models P_A(x)$. On the other hand, we have that $x \in (A)^*$ if and only if $(x, (A)^*) \in E^*$. However, $(X, \mathcal{P}(X)) \models \forall x(P_A(x) \leftrightarrow P_E(x, A))$, so the desired result follows from transfer. \square

Definition 9.6.7. A subset A of X^* is called **internal** if $A \in \mathcal{P}(X)^*$; otherwise, A is called **external**.

Thus, the transfer principle applies to the internal subsets of X^* :

Example 9.6.8. Let us consider $(\mathbb{N}, \mathcal{P}(\mathbb{N}))$ and its nonstandard extension $(\mathbb{N}^*, \mathcal{P}(\mathbb{N})^*)$. We claim that \mathbb{N} is an external subset of \mathbb{N}^* . To see this, note that the following sentence is true in $(\mathbb{N}, \mathcal{P}(\mathbb{N}))$:

$$\begin{aligned} \forall X \in \mathcal{P}(\mathbb{N})((\exists x \in \mathbb{N})(P_E(x, A)) \wedge \exists y \in \mathbb{N} \forall z \in \mathbb{N}(P_E(z, A) \rightarrow z \leq y)) \\ \rightarrow \exists y \in \mathbb{N}(P_E(y, A) \wedge \forall z \in \mathbb{N}(P_E(z, A) \rightarrow z \leq y)). \end{aligned}$$

This sentence says that if $A \subseteq \mathbb{N}$ is bounded above, then A has a maximum element. By transfer, the same holds true for any $A \in \mathcal{P}(\mathbb{N})^*$, that is, for any internal subset of \mathbb{N}^* . If \mathbb{N} were internal, then since it is bounded above (by an infinite element), it would have a maximum, which is clearly not true.

Example 9.6.9. We continue to work with the setup of the previous example. Since

$$(\mathbb{N}, \mathcal{P}(\mathbb{N})) \models (\forall n \in \mathbb{N})(\exists A \in \mathcal{P}(\mathbb{N}))(\forall m \in \mathbb{N})(P_E(m, A) \leftrightarrow m \leq n),$$

by transfer we have

$$(\mathbb{N}^*, \mathcal{P}(\mathbb{N})^*) \models (\forall n \in \mathbb{N}^*)(\exists A \in \mathcal{P}(\mathbb{N})^*)(\forall m \in \mathbb{N}^*)(P_E(m, A) \leftrightarrow m \leq n).$$

Fixing $N \in \mathbb{N}^*$, we suggestively let $\{0, 1, \dots, N\}$ denote the internal subset of \mathbb{N}^* consisting of all the elements of \mathbb{N}^* that are no greater than N . This is a prototypical example of a *hyperfinite* set, to be defined below.

Here are some exercises to get us acquainted with internal sets:

Exercise 9.6.10. Suppose that the nonstandard extension is obtained via an ultrapower model, say using the ultrafilter \mathcal{U} on the index set I . Prove that $A \subseteq X^*$ is internal if and only if there is a family $(A_i)_{i \in I}$ of subsets of X such that $A = \prod_{\mathcal{U}} A_i$.

Exercise 9.6.11. Prove that μ and \mathbb{R}_{fin} are external subsets of \mathbb{R}^* .

Exercise 9.6.12. Suppose that $X \subseteq \mathbb{N}^*$ is internal.

- (1) (Overflow) Suppose that $X \cap \mathbb{N}$ is unbounded in \mathbb{N} . Prove that there is $N \in \mathbb{N}^* \setminus \mathbb{N}$ such that $N \in X$.
- (2) (Underflow) Suppose that, for every $N \in \mathbb{N}^* \setminus \mathbb{N}$, there is $K \in \mathbb{N}^*$ such that $K < N$ and $K \in X$. Prove that $X \cap \mathbb{N} \neq \emptyset$.

The following principle is useful in practice; it says that sets defined (in the first-order logic sense) from internal parameters are internal.

Theorem 9.6.13 (Internal definition principle). *Let*

$$\varphi(x, x_1, \dots, x_m, y_1, \dots, y_n)$$

be a formula, where x, x_1, \dots, x_m range over the sort for X and y_1, \dots, y_n range over the sort for $\mathcal{P}(X)$. Suppose that $a_1, \dots, a_m \in X^$ and $A_1, \dots, A_n \in \mathcal{P}(X)^*$ are given. Set*

$$B := \{b \in X^* : (X^*, \mathcal{P}(X)^*) \models \varphi(b, a_1, \dots, a_m, A_1, \dots, A_n)\}.$$

Then B is internal.

Proof. The following sentence is true in $(X, \mathcal{P}(X))$:

$$\forall x_1 \cdots \forall x_m \forall y_1 \cdots \forall y_n \exists z \forall x (P_E(x, z) \leftrightarrow \varphi(x, x_1, \dots, x_m, y_1, \dots, y_n)).$$

By transfer, this remains true in $(X^*, \mathcal{P}(X)^*)$. Plugging in a_i for x_i and A_j for y_j , we see that

$$(X^*, \mathcal{P}(X)^*) \models \exists z \forall x (P_E(x, z) \leftrightarrow \varphi(x, a_1, \dots, a_m, A_1, \dots, A_n)).$$

The set asserted to exist is B , which then belongs to $\mathcal{P}(X)^*$, that is, B is internal. \square

Example 9.6.14. For any finite collection $a_1, \dots, a_m \in X^*$, the set $\{a_1, \dots, a_m\}$ is internal. Indeed, let $\phi(x, x_1, \dots, x_m)$ be the formula $x = x_1 \vee \dots \vee x_m$. Then

$$\{a_1, \dots, a_m\} = \{b \in X^* : (X^*, \mathcal{P}(X)^*) \models \phi(b, a_1, \dots, a_m)\}.$$

Exercise 9.6.15. Prove that, for any $A \subseteq X^*$, that A is internal if and only if A is a definable set (in the sense of the complete extension of the standard universe).

Exercise 9.6.16. Suppose that the nonstandard universe is κ -saturated. Prove that, for any $\alpha < \kappa$ and any family $(A_\beta)_{\beta < \alpha}$ of internal subsets of X^* , we have that $\bigcap_{\beta < \alpha} A_\beta \neq \emptyset$.

Remark 9.6.17. The conclusion of the previous exercise is often taken to be the definition of κ -saturation for nonstandard universes. While this definition is convenient when presenting nonstandard universes in an informal way, it is a little unfair as it is a weaker notion than actual κ -saturation of the complete extension of the standard universe.

It will also prove useful to have a notion of internal function. To do this, we need to expand our setup a bit. We now consider the many-sorted structure $(X, \mathcal{P}(X), \mathcal{P}(X \times X))$ with an embedding into a nonstandard extension $(X^*, \mathcal{P}(X)^*, \mathcal{P}(X \times X)^*)$. We set:

- $E_1 := \{(x, A) \in X \times \mathcal{P}(X) : x \in A\}$ and
- $E_2 := \{(x, y, A) \in X \times X \times \mathcal{P}(X \times X) : (x, y) \in A\}$.

The proof of the following lemma is exactly like the proof of Lemma 9.6.4.

Lemma 9.6.18. *We may assume that our nonstandard extension satisfies the additional three conditions:*

- $X \subseteq X^*$ and $x = x^*$ for all $x \in X$;
- $\mathcal{P}(X)^* \subseteq \mathcal{P}(X^*)$ and E_1^* is the membership relation restricted to $X^* \times \mathcal{P}(X)^*$;
- $\mathcal{P}(X \times X)^* \subseteq \mathcal{P}(X^* \times X^*)$ and E_2^* is the membership relation restricted to $X^* \times X^* \times \mathcal{P}(X \times X)^*$.

Definition 9.6.19. $B \subseteq X^* \times X^*$ is **internal** if $B \in \mathcal{P}(X \times X)^*$; otherwise, it is **external**. If $A, B \subseteq X^*$ and $f : A \rightarrow B$ is a function, then we say that f is **internal** if the graph of f , $\Gamma(f) := \{(x, y) \in X^* \times X^* : f(x) = y\} \subseteq X^* \times X^*$ is an internal set.

Exercise 9.6.20. Suppose that $f : A \rightarrow B$ is an internal function.

- (1) Prove that A and $\text{range}(f)$ are internal sets.

- (2) Suppose that $B = \mathbb{N}^*$ and $\text{range}(f) \subseteq \mathbb{N}$. Prove that there is $n \in \mathbb{N}$ such that $f(a) \leq n$ for all $a \in A$.

At this point, the reader should verify that they would know how to escape the friendly confines of considering many-sorted structures of the form $(X, \mathcal{P}(X))$ or $(X, \mathcal{P}(X), \mathcal{P}(X \times X))$ and instead be able to consider much wilder many-sorted structures that might contain many sets and their powersets.

Definition 9.6.21. We say that $B \subseteq X^*$ is **hyperfinite** if $B \in \mathcal{P}_f(X)^*$.

Exercise 9.6.22. Assume that $(\mathbb{N}, \mathcal{P}(\mathbb{N}), \mathcal{P}(\mathbb{N} \times \mathbb{N}))$ is part of our structure.

- (1) Prove that hyperfinite sets are internal.
- (2) Prove that an internal subset of a hyperfinite set is hyperfinite.
- (3) Prove that $B \subseteq X^*$ is hyperfinite if and only if there is an internal function $f : B \rightarrow \mathbb{N}^*$ such that f is a bijection between B and $\{0, 1, \dots, N\}$ for some $N \in \mathbb{N}^*$; we then refer to $N + 1$ as the **internal cardinality** of B .
- (4) Prove that finite subsets of X^* are hyperfinite and that their internal cardinality agrees with their usual cardinality.

9.7. Nonstandard generators of ultrafilters

We have seen how, from ultrafilters, we can construct nonstandard extensions. In this section, we explore the other direction: from nonstandard extensions, we can construct ultrafilters. The interplay between these two perspectives has been extremely useful in combinatorial applications as we will see with a couple of sample applications.

For the moment, we fix a set S and consider some nonstandard extension S^* of S .

Definition 9.7.1. Given $\alpha \in S^*$, set $\mathcal{U}_\alpha := \{A \subseteq S : \alpha \in A^*\}$.

Exercise 9.7.2. Prove that \mathcal{U}_α is an ultrafilter on S which agrees with our earlier use of the notation \mathcal{U}_α in case $\alpha \in S$.

So from nonstandard extensions, we get ultrafilters. Under enough saturation, we get all ultrafilters this way:

Exercise 9.7.3. Suppose that the nonstandard extension is $(2^{|S|})^+$ -saturated. Given $\mathcal{U} \in \beta S$, prove that there is some $\alpha \in \bigcap_{A \in \mathcal{U}} A^*$ and that, for such an α , we have $\mathcal{U} = \mathcal{U}_\alpha$.

In other words, every ultrafilter is “principal” if we allow nonstandard generators.

From now on, we assume sufficient saturation so that the conclusion of the previous exercise holds. We let $\pi : S^* \rightarrow \beta S$ be the canonical surjection given by $\pi(\alpha) := \mathcal{U}_\alpha$. While we just saw that π is surjective, it may not be injective, that is, there may be many nonprincipal generators for a given ultrafilter. (For a discussion of when this map is injective, see the next section.) We define an equivalence relation \sim on S^* by setting $\alpha \sim \beta$ if $\mathcal{U}_\alpha = \mathcal{U}_\beta$. In other words, $\alpha \sim \beta$ if and only if, for every $A \subseteq \mathcal{U}$, we have $\alpha \in A^* \Leftrightarrow \beta \in A^*$. Thus π descends to a bijection $\bar{\pi} : S^*/\sim \rightarrow \beta S$.

The **S-topology** (S for “standard”) on S^* has as a basis of clopen sets the sets A^* for $A \subseteq S$. Note that the S-topology on S^* is compact but not necessarily Hausdorff (again, see Section 9.8) and, in fact, the map $\bar{\pi}$ witnesses that βS is homeomorphic to the Hausdorff separation of S^* .

We now restrict our attention to $S = \mathbb{N}$. The naïve expectation would be that $\pi : \mathbb{N}^* \rightarrow \beta\mathbb{N}$ is a semigroup homomorphism, that is, $\mathcal{U}_{\alpha \cdot \beta} = \mathcal{U}_\alpha \cdot \mathcal{U}_\beta$. (We recall the definition of \oplus on $\beta\mathbb{N}$ from Section 4.2.). This is unfortunately not the case; see [42, Example 3.8] for a concrete counterexample. However, there is still a viable formula along these lines whose validity allows the nonstandard method to be applicable to the algebra of ultrafilters.

Fix $\alpha, \beta \in \mathbb{N}^*$ and $A \subseteq \mathbb{N}$. Note then that

$$\begin{aligned} A - \mathcal{U}_\beta &= \{s \in \mathbb{N} : A - s \in \mathcal{U}_\beta\} \\ &= \{s \in \mathbb{N} : \beta \in (A - s)^*\} = \{s \in \mathbb{N} : s + \beta \in A^*\}. \end{aligned}$$

By the definition of the semigroup operation on $\beta\mathbb{N}$, we have that

$$A \in \mathcal{U}_\alpha \oplus \mathcal{U}_\beta \Leftrightarrow A - \mathcal{U}_\beta \in \mathcal{U}_\alpha \Leftrightarrow \alpha \in (A - \mathcal{U}_\beta)^*.$$

Working naïvely (and motivated by some kind of transfer principle), the latter should in turn be equivalent to $\alpha + \beta^* \in A^{**}$. Of course, for this to make any sense, one needs to give meaning to the objects β^* and A^{**} .

One can indeed give concrete meaning to objects like β^* and A^{**} . This idea was first pursued by Mauro Di Nasso in [40], where he used this technique to give an ultrafilter generalization of Rado’s classical theorem on partition regularity of linear equations. One works in a framework for nonstandard analysis where one can *iterate* the $*$ operation, whence β^* above is an element of \mathbb{N}^{**} and A^{**} is a subset of \mathbb{N}^{**} . There is an obvious transfer principle between one level of the tower of iterated nonstandard extensions and the next level. For complete details, see [40] or [42], the latter of which contains many applications of this technique to Ramsey theory. Admittedly this approach takes some getting used to (e.g., unlike the usual convention that $s^* = s$ for $s \in \mathbb{N}$, we now have that $\alpha^* \neq \alpha$ for $\alpha \in \mathbb{N}^* \setminus \mathbb{N}$); however, once one is familiarized with this framework, it proves to be extremely convenient.

Now given $\alpha \in \mathbb{N}^{**}$, we can define an ultrafilter \mathcal{U}_α on \mathbb{N} just as before, namely $A \in \mathcal{U}_\alpha$ if and only if $\alpha \in A^{**}$. For $\alpha, \beta \in \mathbb{N}^* \cup \mathbb{N}^{**}$, we once again set $\alpha \sim \beta$ if and only if $\mathcal{U}_\alpha = \mathcal{U}_\beta$.

Returning to the earlier context, for $\alpha, \beta \in \mathbb{N}^*$ and $A \subseteq \mathbb{N}$, we now have

$$A \in \mathcal{U}_\alpha \oplus \mathcal{U}_\beta \Leftrightarrow \alpha + \beta^* \in A^{**} \Leftrightarrow A \in \mathcal{U}_{\alpha+\beta^*}.$$

In other words, $\mathcal{U}_\alpha \oplus \mathcal{U}_\beta = \mathcal{U}_{\alpha+\beta^*}$.

We now use this framework to reprove Hindman's theorem first proven in Section 4.2 using ultrafilter arithmetic.

Definition 9.7.4. We say that $\alpha \in \mathbb{N}^*$ is **idempotent** if $\alpha + \alpha^* \sim \alpha$.

Exercise 9.7.5. For $\alpha \in \mathbb{N}^*$, prove that α is idempotent if and only if \mathcal{U}_α is idempotent.

The nonstandard version of a subsemigroup of βS is the following:

Definition 9.7.6. $T \subseteq \mathbb{N}^*$ is a **S-subsemigroup** if, for every $\alpha, \beta \in T$, there is $\gamma \in T$ such that $\gamma \sim \alpha + \beta^*$.

Exercise 9.7.7. Prove that $T \subseteq \mathbb{N}^*$ is a S-subsemigroup if and only if $\pi(T)$ is a subsemigroup of $\beta\mathbb{N}$.

Recalling Theorem 4.2.9, Exercise 9.7.7 yields:

Corollary 9.7.8. *Every nonempty closed S-subsemigroup of \mathbb{N}^* contains an idempotent element.*

We now wish to show that if α is idempotent and $\alpha \in A^*$, then A is an FS-set. The following definitions will become useful:

Definition 9.7.9. For $A \subseteq \mathbb{N}$ and $\alpha \in \mathbb{N}^*$, we set

- $A_\alpha := \{s \in \mathbb{N} : s + \alpha \in A^*\}$ and
- ${}_\alpha A := A \cap A_\alpha$.

Exercise 9.7.10. Given $\alpha \in \mathbb{N}^*$, prove that α is idempotent if and only if, for every $A \subseteq \mathbb{N}$, if $\alpha \in A^*$, then $\alpha \in {}_\alpha A^*$. Moreover, in this case, prove that if $s \in A_\alpha$ (resp., $s \in {}_\alpha A$), then $s + \alpha \in A_\alpha^*$ (resp., $s + \alpha \in {}_\alpha A^*$).

We can now prove:

Proposition 9.7.11. *Suppose that $\alpha \in \mathbb{N}^*$ is idempotent and $\alpha \in A^*$. Then A is an FS-set.*

Proof. We recursively construct a sequence $(x_n)_{n=1}^\infty$ such that, for all $m \in \mathbb{N}$, we have $\text{FS}((x_n)_{n=1}^m) \subseteq {}_\alpha A$. Since $\alpha \in {}_\alpha A^*$, there is $x_1 \in {}_\alpha A$. Suppose now that $(x_n)_{n=1}^m$ has been defined with $\text{FS}((x_n)_{n=1}^m) \subseteq {}_\alpha A$. By the previous lemma, we have $\text{FS}((x_n)_{n=1}^m) + \alpha \subseteq {}_\alpha A^*$. By transfer, there is $x_{m+1} \in {}_\alpha A$ with $\text{FS}((x_n)_{n=1}^m) + x_{m+1} \subseteq {}_\alpha A$, whence x_{m+1} is as desired. \square

Hindman’s theorem follows immediately from Proposition 9.7.11: if $\mathbb{N} = C_1 \cup \dots \cup C_r$ and $\alpha \in \mathbb{N}^*$ is an idempotent, then there is $i \in \{1, \dots, r\}$ such that $\alpha \in C_i^*$, whence, by the previous proposition, C_i is an FS-set.

We end this section by using the iterated nonstandard extension framework to give a different proof of Ramsey’s theorem:

Theorem 9.7.12. *Suppose that $\mathbb{N}^{[2]} = C_1 \cup \dots \cup C_r$. Then there is an infinite set $X \subseteq \mathbb{N}$ and $i \in \{1, \dots, r\}$ such that $X^{[2]} \subseteq C_i$.*

Proof. Fix $\alpha \in \mathbb{N}^* \setminus \mathbb{N}$ and take $i \in \{1, \dots, r\}$ such that $\{\alpha, \alpha^*\} \in C_i^{**}$. For simplicity, set $C := C_i$. We now recursively construct a sequence $(x_n)_{n \in \mathbb{N}}$ such that, for all $i < j \leq n$, we have

- $x_i < x_j$,
- $(x_i, x_j) \in C$ (whence the sequence is the desired infinite monochromatic set), and
- $(x_i, \alpha^*) \in C_i^*$.

Suppose that the sequence has been constructed through x_{n-1} ; we show how to construct x_n . Since α witnesses the truth of

$$(\exists y \in \mathbb{N}^*)(x_{n-1} < y \wedge \bigwedge_{i < n} (x_i, y) \in C^* \wedge (y, \alpha^*) \in C^{**}),$$

by transfer we have that

$$(\exists y \in \mathbb{N})(x_{n-1} < y \wedge \bigwedge_{i < n} (x_i, y) \in C \wedge (y, \alpha) \in C^*).$$

Letting x_n witness the truth of this statement, we see that x_n is as desired. □

Remark 9.7.13. One can adapt the above proof to consider partitions of $\mathbb{N}^{[m]}$ for $m > 2$ using m iterated nonstandard extensions instead.

From ultrafilters we can get ultrapowers (and thus nonstandard extensions), whilst from nonstandard extensions we can construct ultrafilters. The last couple of results in this section are motivated by the question of what happens when we compose these two constructions.

For example, suppose that $\mathcal{U} \in \beta I$ and $[\alpha]_{\mathcal{U}} \in I^{\mathcal{U}}$ (which we think of as a nonstandard extension of I). We can use this nonstandard element to generate an ultrafilter on I which we call \mathcal{U}_{α} . (This looks like bad notation as perhaps the choice of representative of $[\alpha]_{\mathcal{U}}$ might change the ultrafilter; we will soon see that it does not.) It is natural to wonder how this ultrafilter \mathcal{U}_{α} compares to the original ultrafilter \mathcal{U} . There is in fact a direct connection:

Proposition 9.7.14. *Suppose that $\mathcal{U} \in \beta I$ and $\alpha \in I^{\mathcal{U}}$. Then $\mathcal{U}_{\alpha} = \alpha(\mathcal{U})$.*

Proof. Given $A \subseteq I$, we note that $A \in \mathcal{U}_\alpha$ if and only if $[\alpha]_{\mathcal{U}} \in A^{\mathcal{U}}$ if and only if $\alpha(i) \in A$ for \mathcal{U} -almost all $i \in I$ if and only if $A \in \alpha(\mathcal{U})$. \square

We now consider the other direction:

Proposition 9.7.15. *Suppose that $\alpha \in S^*$. Then $S^{\mathcal{U}_\alpha}$ embeds in S^* .*

Proof. The desired embedding is given by $[f]_{\mathcal{U}_\alpha} \mapsto f(\alpha)$. This is well defined and an injection since $[f]_{\mathcal{U}_\alpha} = [g]_{\mathcal{U}_\alpha}$ if and only if $\alpha \in \{s \in S : f(s) = g(s)\}^*$ if and only if $f(\alpha) = g(\alpha)$. \square

9.8. Hausdorff ultrafilters

In Section 1.3, we proved that if \mathcal{U} is an ultrafilter and $f, g : I \rightarrow I$ are functions for which $f \equiv_{\mathcal{U}} g$, then $f(\mathcal{U}) = g(\mathcal{U})$. We give a name for ultrafilters satisfying the converse implication:

Definition 9.8.1. Let \mathcal{U} be a nonprincipal ultrafilter on the set I . We say that \mathcal{U} is **Hausdorff** if, given any two functions $f, g : I \rightarrow I$, if $f(\mathcal{U}) = g(\mathcal{U})$, then $f \equiv_{\mathcal{U}} g$.

Exercise 9.8.2. Suppose that \mathcal{U} is a Hausdorff ultrafilter and $\mathcal{V} \leq_{RK} \mathcal{U}$. Show that \mathcal{V} is also a Hausdorff ultrafilter.

We now explain the terminology:

Theorem 9.8.3. *Suppose that \mathcal{U} is a nonprincipal ultrafilter on I . Then \mathcal{U} is Hausdorff if and only if the S -topology on $I^{\mathcal{U}}$ is Hausdorff.*

Proof. We first note that, given $f : I \rightarrow I$ and $A \subseteq I$, we have that $A \in f(\mathcal{U})$ if and only if $f^{-1}(A) \in \mathcal{U}$ if and only if $[f]_{\mathcal{U}} \in A^{\mathcal{U}}$.

Now suppose that \mathcal{U} is Hausdorff and that $[f]_{\mathcal{U}} \neq [g]_{\mathcal{U}}$. By the definition of being Hausdorff, $f(\mathcal{U}) \neq g(\mathcal{U})$, whence there is some $A \in f(\mathcal{U})$ with $A \notin g(\mathcal{U})$. By the above calculation, this means that $[f]_{\mathcal{U}} \in A^{\mathcal{U}}$ while $[g]_{\mathcal{U}} \notin A^{\mathcal{U}}$. Thus, the basic open neighborhoods $A^{\mathcal{U}}$ and $(I \setminus A)^{\mathcal{U}}$ separate $[f]_{\mathcal{U}}$ and $[g]_{\mathcal{U}}$, whence $I^{\mathcal{U}}$ is Hausdorff.

Conversely, suppose that $I^{\mathcal{U}}$ is Hausdorff. Suppose that $f, g : I \rightarrow I$ are such that $f(\mathcal{U}) = g(\mathcal{U})$. By the above calculation, this shows, for any $A \subseteq I$, we have that $[f]_{\mathcal{U}} \in A^{\mathcal{U}}$ if and only if $[g]_{\mathcal{U}} \in A^{\mathcal{U}}$; since $I^{\mathcal{U}}$ is Hausdorff, this implies that $[f]_{\mathcal{U}} = [g]_{\mathcal{U}}$, as desired. \square

In other words: if I^* is the nonstandard extension we get from \mathcal{U} , then \mathcal{U} is Hausdorff if and only if the map $\alpha \mapsto \mathcal{U}_\alpha : I^* \rightarrow \beta I$ is injective.

We now turn to the existence of Hausdorff ultrafilters. Surprisingly, the following question is still open:

Question 9.8.4. Can one prove in ZFC that Hausdorff ultrafilters exist?

However, under stronger set-theoretic axioms, we can prove that Hausdorff ultrafilters exist:

Theorem 9.8.5. *Selective ultrafilters are Hausdorff. In particular, it is consistent with ZFC that Hausdorff ultrafilters exist.*

Proof. Suppose that \mathcal{U} is selective and that $f, g : I \rightarrow I$ are such that $f(\mathcal{U}) = g(\mathcal{U})$. If $f(\mathcal{U}) = \mathcal{U}_n$, then f is constantly n , which forces g to be constantly n and thus f and g are the same function. If $f(\mathcal{U})$ is nonprincipal, then by minimality of \mathcal{U} , $f(\mathcal{U}) = \mathcal{U}$, whence there is $A \in \mathcal{U}$ such that $f(n) = n$ for all $n \in A$. Likewise, there is $B \in \mathcal{U}$ such that $g(n) = n$ for all $n \in B$. It follows that $f(n) = g(n) = n$ for all $n \in A \cap B$, whence $[f]_{\mathcal{U}} = [g]_{\mathcal{U}}$ and \mathcal{U} is Hausdorff. \square

The final result in this section shows that if we are to search for regular Hausdorff ultrafilters, then they cannot exist on sets of “large” size. We recall that \mathfrak{u} denotes the ultrafilter number, as defined in Section 1.5.

Theorem 9.8.6. *Suppose that κ is a cardinal with $\kappa \geq \mathfrak{u}$. If \mathcal{U} is a regular ultrafilter on κ , then \mathcal{U} is not Hausdorff.*

Proof. Since \mathcal{U} is regular and $\kappa \geq \mathfrak{u}$, there are sets $(A_\alpha)_{\alpha < \mathfrak{u}}$ in \mathcal{U} such that, for every $i < \kappa$, $F_i := \{\alpha < \mathfrak{u} : i \in A_\alpha\}$ is finite. Without loss of generality, we may assume that each $F_i \neq \emptyset$. Let $(B_\alpha)_{\alpha < \mathfrak{u}}$ be a base for a nonprincipal ultrafilter \mathcal{V} on ω . Choose functions $f, g : \kappa \rightarrow \omega$ such that, for each $i \in I$, we have $f(i), g(i) \in \bigcap_{\alpha \in F_i} B_\alpha$ while $f(i) \neq g(i)$; this is clearly possible since $\bigcap_{\alpha \in F_i} B_\alpha$ is infinite. Note that $A_\alpha \subseteq f^{-1}(B_\alpha) \cap g^{-1}(B_\alpha)$, whence every B_α belongs to both $f(\mathcal{U})$ and $g(\mathcal{U})$, whence $f(\mathcal{U}) = g(\mathcal{U}) = \mathcal{V}$. However, $f \not\equiv_{\mathcal{U}} g$. It follows that \mathcal{U} is not Hausdorff. \square

In particular, in a model of ZFC for which $\mathfrak{u} = \aleph_1$, a regular Hausdorff ultrafilter, should it exist, must be an ultrafilter on a countably infinite set.

9.9. Notes and references

There are many introductory texts on nonstandard analysis. My favorite, and the one for which most of the first four sections is based on, is Goldblatt’s book [68]. Of course, Robinson’s original treatise [148] is still a great read. The discussion on complete extensions and limit ultrapowers is taken from [28], although we have reorganized things dramatically in what we hope makes for a cleaner perspective. The introduction of limit ultrapowers and their equivalence with complete extensions is due to Keisler [98]. Our

treatment of many-sorted structures is heavily influenced by Henson's article [78]. Although essentially just the notion of realization of a type, the perspective of hyperfinite generator of an ultrafilter has proven very useful recently in applications of nonstandard methods to combinatorics. A nice introduction to this topic, containing more details than those presented in Section 9.7, is DiNasso's article [39], which also has further information on Hausdorff ultrafilters. Theorem 9.8.6 was taken from [41].

Limit groups

Section 10.1 introduces the class of limit groups, which is the class of finitely generated subgroups of ultrapowers of free groups. Section 10.2 gives some examples of limit groups and discusses some algebraic properties true of all limit groups. Section 10.3 offers a purely algebraic description of limit groups as those finitely generated groups that are fully residually free, and connects this characterization with the algebraic properties discussed in Section 10.2. Finally, in Section 10.4, we explain the nomenclature by introducing the topological space of marked groups and prove that limit groups are exactly those marked groups that are in the closure of the set of marked free groups.

10.1. Introducing the class of limit groups

In this chapter, we explore the class of *limit groups*, which is the class of finitely generated subgroups of *nonstandard free groups*. It is the purpose of this introductory section to elucidate the precise meaning of the last sentence.

We first mention a special case of Exercise 8.3.14:

Proposition 10.1.1. *Suppose that G and H are groups. The following are equivalent:*

- (1) $G \models \text{Th}_\forall(H)$.
- (2) For every finitely generated subgroup G_0 of G and every nonprincipal ultrafilter \mathcal{U} on \mathbb{N} , G_0 embeds into $H^\mathcal{U}$.
- (3) For every κ -regular ultrafilter \mathcal{U} with $\kappa := \max(|G|, \aleph_0)$, G embeds into $H^\mathcal{U}$.
- (4) For some ultrafilter \mathcal{U} , G embeds into $H^\mathcal{U}$.

In the literature, one often finds the previous fact stated as “ G embeds into some nonstandard extension H^* of H ”, with the type of nonstandard extension and level of saturation left intentionally vague. We will follow this practice, but remind the reader of the precise reading given in the previous proposition.

We assume that the reader is familiar with the basic definitions and facts concerning free groups. For every cardinal κ , \mathbb{F}_κ will denote the free group on κ generators. (In Chapter 7, we used \mathbb{F}_p , for p an integer, to denote the finite field with p elements. In the rest of this book, we stick with the current convention.) This includes the case $\mathbb{F}_0 = \{e\}$ and $\mathbb{F}_1 = \mathbb{Z}$. For $\kappa \geq 2$, \mathbb{F}_κ is nonabelian. For $\kappa \leq \lambda$, note that \mathbb{F}_κ is a subgroup of \mathbb{F}_λ . We also recall that, for every $\kappa \leq \aleph_0$, \mathbb{F}_κ can be embedded into \mathbb{F}_2 . By Proposition 10.1.1, it follows that \mathbb{F}_κ can be embedded into \mathbb{F}_2^* for any cardinal κ . (As a reminder, the saturation of the nonstandard extension will have to increase as κ increases.) Consequently, all the nonabelian free groups have the same universal theory. For this reason, we will often speak about groups embedding into \mathbb{F}_2^* for the sake of simplicity.

Remark 10.1.2. A remarkable generalization of the conclusion of the previous fact holds, namely all nonabelian free groups are elementarily equivalent. This is a very difficult result due to Sela, answering a question of Tarski.

Definition 10.1.3. A group G is called a **universally free group** if it can be embedded into \mathbb{F}_2^* . A finitely generated universally free group is called a **limit group**.

In Section 10.4, we will see the reason behind the nomenclature “limit group”.

Exercise 10.1.4. Prove that universally free groups are torsion-free.

For abelian groups, being torsion-free characterizes being universally free:

Theorem 10.1.5. *Suppose that G is an abelian group. The following are equivalent:*

- (1) G is torsion-free.
- (2) G embeds into \mathbb{Z}^* .
- (3) G is universally free.

Proof. The only direction that needs proof is (1) \Rightarrow (2). By Proposition 10.1.1 above, it suffices to assume that G is finitely generated. Thus, by the *fundamental theorem of finitely generated abelian groups*, G is isomorphic to \mathbb{Z}^k for some k . For ease of exposition, we assume that $k = 2$. Thus, we are

trying to show that $\mathbb{Z} \times \mathbb{Z}$ embeds into \mathbb{Z}^* . By the compactness theorem, it suffices to show that the following is a consistent set of formulae:

$$\text{Th}(\mathbb{Z}) \cup \{mx \neq ny : m, n \in \mathbb{Z} \setminus \{0\}\},$$

for if $a, b \in \mathbb{Z}^*$ realize the set of formulae, then the subgroup of \mathbb{Z}^* generated by a and b is isomorphic to $\mathbb{Z} \times \mathbb{Z}$, as desired. However, this set is clearly finitely satisfied in \mathbb{Z} by taking $x = k$ and $y = l$ for k and l sufficiently far apart. \square

An algebraic characterization of the class of all universally free groups is possible and will be discussed later in the chapter.

10.2. First examples and properties of limit groups

Clearly, free groups are universally free. Here is a nonfree example:

Example 10.2.1. Let a and b be the two generators of \mathbb{F}_2 and let A be the cyclic subgroup of \mathbb{F}_2 generated by a . We claim that the *amalgamated free product* $G := \mathbb{F}_2 *_A (A \times \mathbb{Z})$ is a limit group that is not a free group. (For the reader unfamiliar with the amalgamated free product construction, just think of G as being the “freest” combination of the two groups \mathbb{F}_2 and $A \times \mathbb{Z}$ subject to the requirement that their two copies of A must be identified.) It is clear that G is not free. To see that G is universally free, fix $N \in \mathbb{N}^* \setminus \mathbb{N}$ and define a group homomorphism $\phi : G \rightarrow \mathbb{F}_2^*$ by declaring that ϕ is the identity on \mathbb{F}_2 and such that $\phi(e, 1) = a^N$. (By the universal property of amalgamated free products, this description does in fact yield a well-defined homomorphism.) It remains to see that ϕ is injective. To see this, it is useful to recall that elements of G can be put in a *normal form* $g_1 \cdots g_n$, where the g_i ’s alternate between being elements of \mathbb{F}_2 and $A \times \mathbb{Z}$, no g_i belongs to A if $i > 1$ and $g_1 \neq e$. For sake of concreteness, consider the element $g_1(a^{k_1}, l_1)g_2(a^{k_2}, l_2) \cdots g_n(a^{k_n}, l_n)$ of G , where each $g_i \in \mathbb{F}_2 \setminus \{e\}$, and $l_i \neq 0$ for all $i = 1, \dots, n$. This element gets mapped by ϕ to $g_1 a^{k_1 + Nl_1} g_2 a^{k_2 + Nl_2} \cdots g_n a^{k_n + Nl_n}$. We leave it to the reader to verify that this is a nontrivial element of \mathbb{F}_2^* .

One can generalize the method of proof in the previous example to prove the following:

Theorem 10.2.2. *Suppose that A is a maximal cyclic subgroup of \mathbb{F}_2 and B is a finitely generated free abelian group. Then the group $\mathbb{F}_2 *_A (A \times B)$ is a limit group.*

The previous theorem can then be used to show that many so-called *surface groups* from geometric group theory are limit groups.

The above technique can be pushed even further to produce a way of creating new limit groups from old ones:

Theorem 10.2.3. *Suppose that L is a limit group, A is a maximal abelian subgroup of L , and B is a finitely generated free abelian group. Then the group $L *_A (A \times B)$ is a limit group.*

We now move on to discussing algebraic properties of universally free groups. As discussed in the previous section, universally free groups are torsion-free. Here is another algebraic property they possess:

Definition 10.2.4. A group G is **commutative transitive** if commutativity is a transitive relation on $G \setminus \{e\}$, that is, for all $a, b, c \in G \setminus \{e\}$, if $ab = ba$ and $bc = cb$, then $ac = ca$.

It is clear that any abelian group is commutative transitive. The following exercise gives some more interesting examples:

Exercise 10.2.5.

- (1) Show that free groups are commutative transitive.
- (2) Show that universally free groups are commutative transitive.

Exercise 10.2.6.

- (1) Suppose that G is a nonabelian commutative transitive group. Prove that the center of G is trivial.
- (2) Suppose that G is a nonabelian group and H is a nontrivial abelian group. Prove that $G \times H$ is not commutative transitive. Conclude that none of the following classes of groups is closed under direct products: commutative transitive groups, universally free groups, limit groups.

Here are some other characterizations of being commutative transitive:

Exercise 10.2.7. For a given group G , prove that the following are equivalent:

- (1) G is commutative transitive.
- (2) The centralizer of any nontrivial element is abelian.
- (3) If two abelian subgroups intersect nontrivially, then their union generates an abelian subgroup.

In order to motivate our next algebraic property of universally free groups, we consider the following:

Exercise 10.2.8. Suppose that G is a commutative transitive group and H is a maximal abelian subgroup of G . Prove that for any $g \in G$, one has that $H \cap gHg^{-1}$ is either H or $\{e\}$.

The previous exercise motivates:

Definition 10.2.9. A group G is **CSA** if, whenever H is a maximal abelian subgroup of G , then for all $g \in G \setminus H$, $H \cap gHg^{-1} = \{e\}$.

A subgroup H of a group G is called **malnormal** if $gHg^{-1} = \{e\}$ for all $g \in G \setminus H$. Thus, being malnormal is the diametrically opposite property of being normal. Another name for being malnormal is conjugately separated. CSA is an abbreviation for **conjugately separated abelian**: a group is CSA if and only if all of its maximal abelian subgroups are conjugately separated.

Exercise 10.2.10. Prove that CSA groups are commutative transitive.

On the other hand, it is possible to find examples (both finite and infinite) of commutative transitive groups that are not CSA groups [57].

Exercise 10.2.11. Prove that free groups are CSA.

We would eventually like to use the previous exercise to show that universally free groups are CSA. Unfortunately, the definition of CSA is not first order, for it quantifies over all subgroups of G . Thankfully, there is a first-order reformulation:

Exercise 10.2.12. Prove that a group G is CSA if and only if it is commutative transitive and satisfies:

- for all $g, h \in G \setminus \{e\}$, if h and ghg^{-1} commute, then g and h commute.

Since the above characterization of CSA is in terms of two first-order universal statements, we have:

Corollary 10.2.13. *Universally free groups are CSA.*

We have just seen that universally free groups are torsion-free CSA groups. It is not the case that the nonabelian torsion-free CSA groups are exactly the nonabelian universally free groups, as the following example shows. Since the proof involves material outside the scope of this book, we content ourselves with a sketch.

Example 10.2.14. Let G be the group generated by two elements a and b subject to the single relation $abab^2 \cdots ab^n = e$. Since the relation is not a proper power, a theorem of Karrass, Magnus, and Solitar [95] shows that G is torsion free. If n is sufficiently large, then one can show that G has the so-called *metric small cancellation property*, which, in turn, implies that it is a so-called *word-hyperbolic group*, and then a theorem of Gromov implies that the group is CSA. Note that G is also nonabelian. However, in the

next section, we show that \mathbb{F}_2 is the only nonabelian limit group generated by two elements. It follows that G is not a limit group.

The previous example notwithstanding, there is quite a close connection between torsion-free CSA groups and universally free groups as we will see in the next section.

10.3. Connection with fully residual freeness

In this section, we give several purely algebraic reformulations of being universally free and connect it with the algebraic notions described at the end of the previous section. Here is the crucial definition:

Definition 10.3.1. Let \mathcal{C} denote a class of groups (such as the class of finite groups, the class of abelian groups, the class of free groups, etc. . . .). We say that a group G is **residually \mathcal{C}** if, for any $g \in G \setminus \{e\}$, there is a group $H \in \mathcal{C}$ and a group homomorphism $\phi : G \rightarrow H$ such that $\phi(g) \neq e$.

Clearly, free groups are residually free. Here is a nonfree counterexample:

Exercise 10.3.2. Show that the direct product of two free groups is residually free. In particular, show that $\mathbb{F}_2 \times \mathbb{Z}$ is residually free.

Exercise 10.3.3. Prove that residually free groups are torsion free.

By Exercises 10.2.5 and 10.2.6(2), $\mathbb{F}_2 \times \mathbb{Z}$ is not a limit group; combined with Exercise 10.3.2, we see that residual freeness does not imply being a limit group (for a finitely generated group). However, a seemingly small tweak in the definition does in fact imply being a limit group, and, in fact, characterizes being a limit group:

Definition 10.3.4. Let \mathcal{C} be a class of groups. We say that a group G is **fully residually \mathcal{C}** if, for any finite $F \subseteq G \setminus \{e\}$, there is a group $H \in \mathcal{C}$ and a group homomorphism $\phi : G \rightarrow H$ such that $\phi(g) \neq e$ for all $g \in F$.

For some classes of groups, there is no difference in the two notions:

Exercise 10.3.5. Suppose that the class \mathcal{C} is closed under finite direct products. Show that G is residually \mathcal{C} if and only if it is fully residually \mathcal{C} .

Thus, for example, there is no difference between being residually finite and fully residually finite. However, the class of free groups is not closed under direct products, whence there could in fact be a difference in the two notions. Keeping in mind the case of $\mathbb{F}_2 \times \mathbb{Z}$, we see that being fully residually free is in fact stronger than being residually free:

Proposition 10.3.6. *If G is a fully residually \mathcal{C} group, then G embeds in an ultraproduct of members of \mathcal{C} .*

Proof. For each finite $F \subseteq G \setminus \{e\}$, take $H_F \in \mathcal{C}$ and a homomorphism $\phi_F : G \rightarrow H_F$ such that $\phi_F(g) \neq e$ for all $g \in F$. Set $I := \mathcal{P}_f(G)$. For each $g \in G$, let $A_g := \{F \in I : g \in F\}$. Then the family $(A_g)_{g \in G}$ has the FIP, whence there is an ultrafilter \mathcal{U} on I such that $A_g \in \mathcal{U}$ for all $g \in G$. We claim that G embeds in $\prod_{\mathcal{U}} H_F$. To see this, for each $g \in G$, let $a_g \in \prod_{F \in I} H_F$ be such that $a_g(F) = \phi_F(g)$ when $g \in F$ and $a_g(F) = e$ when $g \notin F$. Define $\phi : G \rightarrow \prod_{\mathcal{U}} H_F$ to be given by $\phi(g) := [a_g]_{\mathcal{U}}$. We leave it to the reader to verify that ϕ is an injective homomorphism. \square

Corollary 10.3.7. *If G is a fully residually free group, then G is a universally free group.*

We now show that being fully residually free characterizes being a limit group (amongst finitely generated groups). First, we need:

Lemma 10.3.8. *Suppose that R is a finitely generated subring of \mathbb{Z}^* . Then for any $a_1, \dots, a_n \in R \setminus \{0\}$, there is a ring homomorphism $\phi : R \rightarrow \mathbb{Z}$ such that $\phi(a_i) \neq 0$ for $i = 1, \dots, n$.*

Proof. Take variables X_1, \dots, X_k and a surjective ring homomorphism $\theta : \mathbb{Z}[X_1, \dots, X_k] \rightarrow R$. (This is possible since R is finitely generated.) Let J be the kernel of θ , which is finitely generated since the polynomial ring is Noetherian. Let f_1, \dots, f_m generate J . Take $g_i \in \mathbb{Z}[X_1, \dots, X_k]$ such that $\theta(g_i) = a_i$. Note then that

$$\mathbb{Z}^* \models \exists z_1 \cdots \exists z_k \left(\bigwedge_{i=1}^m f_i(\vec{z}) = 0 \wedge \bigwedge_{i=1}^n g_i(\vec{z}) \neq 0 \right),$$

as witnessed by $z_i := \theta(X_i)$. By elementarity, \mathbb{Z} believes the same statement, say as witnessed by u_1, \dots, u_k . Then the map $\phi : R \rightarrow \mathbb{Z}$ given by setting $\phi(\theta(X_i)) = u_i$ is as desired. \square

In the proof of Theorem 10.3.10 below, we will also need the following:

Fact 10.3.9. Fix an odd prime p and consider the group homomorphism $\mathrm{SL}_2(\mathbb{Z}) \rightarrow \mathrm{SL}_2(\mathbb{Z}/p\mathbb{Z})$. Then the kernel K is a nonabelian free group.

Proof. This is a standard consequence of the *ping-pong lemma* from geometric group theory. For example, see [31, Chapter 5] for a very friendly account. \square

Theorem 10.3.10. *A finitely generated group is a limit group if and only if it is fully residually free.*

Proof. Suppose that G is a limit group. By Fact 10.3.9, we may assume that G is a subgroup of K^* . Let g_1, \dots, g_n generate G . Without loss of generality, we may assume that this generating set is closed under inverse.

Write $g_j = \begin{pmatrix} 1 + pa_j & pb_j \\ pc_j & 1 + pd_j \end{pmatrix}$ with each $a_j, b_j, c_j, d_j \in \mathbb{Z}^*$. Let R be the subring of \mathbb{Z}^* generated by these elements. Then $G \subseteq \mathrm{SL}_2(R)$.

To show that G is fully residually free, fix $h_1, \dots, h_k \in G \setminus \{1\}$. Take a ring homomorphism $\phi : R \rightarrow \mathbb{Z}$ such that, for any entry a of some h_j , if $a \neq 0$, then $\phi(a) \neq 0$ and if $a \neq 1$, then $\phi(a) \neq 1$. This induces a group homomorphism $\psi : \mathrm{SL}_2(R) \rightarrow \mathrm{SL}_2(\mathbb{Z})$. Note that $\psi(G)$ is contained in the free group K and $\psi(h_j) \neq 1$ for $j = 1, \dots, k$, as desired. \square

One cannot drop the finitely generated assumption. For example, \mathbb{F}_2^* contains a copy of \mathbb{Q} and any homomorphism from \mathbb{Q} onto a free group is trivial. Nevertheless, we can say the following:

Corollary 10.3.11. *A group is universally free if and only if every finitely generated subgroup is fully residually free.*

Using the above algebraic reformulation of being a limit group, we can fulfill a promise made in the last section.

Proposition 10.3.12. *Suppose that G is a residually free group and $a, b \in G$. Then the group generated by a and b is isomorphic to \mathbb{Z} , $\mathbb{Z} \times \mathbb{Z}$, or \mathbb{F}_2 .*

Proof. Since G is torsion free, if a and b commute, then they generate \mathbb{Z} or $\mathbb{Z} \times \mathbb{Z}$. Suppose that a and b do not commute, that is, $aba^{-1}b^{-1} \neq e$. Since G is residually free, there is a homomorphism $\phi : G \rightarrow \mathbb{F}_2$ such that $\phi(aba^{-1}b^{-1}) \neq e$, that is, $\phi(a)$ and $\phi(b)$ do not commute. It follows that the subgroup of \mathbb{F}_2 generated by $\phi(a)$ and $\phi(b)$ is freely generated by them, whence a and b generate a free group. \square

Corollary 10.3.13. *\mathbb{F}_2 is the only nonabelian limit group generated by two elements.*

There is a connection between the current discussion and the algebraic properties discussed in Section 10.2. The proof, however, is beyond the scope of this book and is a combination of the results in [147], [63], and [6].

Theorem 10.3.14. *For any group G , the following are equivalent:*

- (1) G is a universally free group.
- (2) G is residually free and commutative transitive.
- (3) G is residually free and CSA.
- (4) G is residually free and does not contain a subgroup isomorphic to $\mathbb{F}_2 \times \mathbb{Z}$.

Remark 10.3.15. In the case of abelian groups, the previous theorem simply says that being residually free and fully residually free coincide with being torsion free.

Exercise 10.3.16. Prove that the class of universally free groups is closed under free products.

10.4. Explaining the terminology: the space of marked groups

In this section, we explain the terminology “limit groups” by showing that a finitely generated group is a limit group if and only if it is a limit of finitely generated free groups in the space of *marked groups*.

Definition 10.4.1. A **marked group** is a pair (G, S) , where G is a group and $S = (s_1, \dots, s_n)$ is a finite ordered tuple from G such that the set $\{s_1, \dots, s_n\}$ generates G .

Example 10.4.2. The same group can have many markings. For example, $(\mathbb{Z}, (1))$, $(\mathbb{Z}, (1, 2))$, $(\mathbb{Z}, (2, 1))$, and $(\mathbb{Z}, (1, 1, 3))$ are all different markings of the group \mathbb{Z} .

We let \mathcal{G}_n denote the set of marked groups (G, S) such that S is an n -tuple from G . We now explain the appropriate topology on \mathcal{G}_n .

Definition 10.4.3. Given $(G, S), (G', S') \in \mathcal{G}_n$, let $\nu((G, S), (G', S'))$ denote the maximal $k \in \mathbb{N}$ such that, for any word $w(x_1, \dots, x_n)$ of length at most k , $w(s_1, \dots, s_n) = e$ in G if and only if $w(s'_1, \dots, s'_n) = e$ in G' . We then define $d((G, S), (G', S')) = 2^{-\nu((G, S), (G', S'))}$.

Note that it is possible that $\nu((G, S), (G', S')) = \infty$, from which $d((G, S), (G', S')) = 0$. This happens if and only if the marked groups are isomorphic in the sense of the following definition:

Definition 10.4.4. If (G, S) and (G', S') belong to \mathcal{G}_n , then (G, S) and (G', S') are **isomorphic** if the map $s_i \mapsto s'_i$ for $i = 1, \dots, n$ yields an isomorphism between G and G' .

It thus behooves us to redefine \mathcal{G}_n to consist of isomorphism classes of marked groups with an n -tuple of generators (rather than the marked groups themselves).

Exercise 10.4.5. Prove that the function d on \mathcal{G}_n is a metric on \mathcal{G}_n .

The **metric space of n -generated marked groups** is the set \mathcal{G}_n equipped with the above metric. We can now work toward the aforementioned description of limit groups as limits (in the sense of the above metric) of free groups. We first describe an appropriate ultraproduct construction for marked groups.

Definition 10.4.6. Suppose that $(G_i, S_i)_{i \in I}$ is a family from \mathcal{G}_n and \mathcal{U} is an ultrafilter on I . Write $S_i = (s_1(i), \dots, s_n(i))$. Then the **marked group ultraproduct of the family $(G_i, S_i)_{i \in I}$ with respect to \mathcal{U}** , denoted, $\prod_{\mathcal{U}}^m (G_i, S_i)$, is the subgroup of the usual ultraproduct $\prod_{\mathcal{U}} G_i$ generated by $[s_1]_{\mathcal{U}}, \dots, [s_n]_{\mathcal{U}}$ equipped with the marking $([s_1]_{\mathcal{U}}, \dots, [s_n]_{\mathcal{U}})$.

Theorem 10.4.7. *Given any family $(G_i, S_i)_{i \in I}$ and ultrafilter \mathcal{U} on I , we have that*

$$\prod_{\mathcal{U}}^m (G_i, S_i) = \lim_{\mathcal{U}} (G_i, S_i),$$

where the ultralimit is calculated in the topology on \mathcal{G}_n given above.

Proof. Fix $k \in \mathbb{N}$. Take $J \in \mathcal{U}$ such that, for all $j \in J$ and all words $w(x_1, \dots, x_n)$ of length at most k , we have $w([s_1]_{\mathcal{U}}, \dots, [s_n]_{\mathcal{U}}) = e$ in $\prod_{\mathcal{U}} G_i$ if and only if $w(s_1(j), \dots, s_n(j)) = e$ in G_j . It follows that

$$d\left(\prod_{\mathcal{U}}^m (G_i, S_i), (G_j, S_j)\right) \leq 2^{-k}$$

for all $j \in J$. The result now follows. \square

Corollary 10.4.8. *For each n , \mathcal{G}_n is compact.*

Corollary 10.4.9. *Suppose that $(G, S) \in \mathcal{G}_n$ is in the closure of the set of marked free groups in \mathcal{G}_n . Then G is a limit group.*

We now prove the converse of the previous corollary. First:

Proposition 10.4.10. *Suppose that G is a finitely generated subgroup of an ultraproduct $\prod_{\mathcal{U}} G_i$ and fix a marking (G, S) of G . Write $S = ([s_1]_{\mathcal{U}}, \dots, [s_n]_{\mathcal{U}})$. For each $i \in I$, let H_i be the subgroup of G_i generated by $S_i := (s_1(i), \dots, s_n(i))$. Then $\lim_{\mathcal{U}} (H_i, S_i) = (G, S)$.*

Exercise 10.4.11. Prove Proposition 10.4.10.

Corollary 10.4.12. *If G is a limit group, then for any marking $(G, S) \in \mathcal{G}_n$ of G , (G, S) is a limit of marked free groups.*

Summarizing the above discussion, we have:

Theorem 10.4.13. *Given a finitely generated group G , the following are equivalent:*

- (1) G is a limit group.
- (2) For any marking (G, S) in \mathcal{G}_n of G , (G, S) is a limit of marked free groups.
- (3) There is a marking (G, S) of G in \mathcal{G}_n such that (G, S) is a limit of marked free groups.

Exercise 10.4.14. Use Theorem 10.4.13 to give a different proof of Corollary 10.3.13.

Some intuition can be gained from this topological perspective; see [26].

10.5. Notes and references

The class of limit groups was introduced by Sela in [155] in his first paper presenting the solution of the Tarski problem, namely that all nonabelian free groups are elementarily equivalent. An excellent introduction to the class of limit groups is the survey article by Champetier and Guirardel [26]. Theorem 10.2.2 is due to Baumslag [6] while its generalization can be found in [7]. Much of the discussion on CT and CSA groups is from [57]. The equivalence with fully residually groups is due to Remeslennikov [147] but our treatment follows Chiswell's book [30]. The space of marked groups was introduced by Grigorchuk in [70]. The main result of Section 10.4 is from [26].

Part 3

Metric ultraproducts and their applications

Metric ultraproducts

In this chapter, we study the metric ultraproduct, the construction that will occupy our attention in this part of the book. Section 11.1 introduces the definition of the ultraproduct, while Section 11.2 explains the connection between the metric ultraproduct and the nonstandard hull construction from nonstandard analysis. A nice feature of the metric ultraproduct is that it is often automatically a complete metric space, as detailed in Section 11.3. In particular, the metric ultraproduct construction yields a nice construction of the completion of a metric space. In Section 11.4, we give an outline of relatively recent continuous logic suitable for studying structures based on metric spaces and for which an analogue of Łoś's theorem holds. Unlike the classical ultraproduct, our treatment of the metric ultraproduct dives straight into the ultraproduct construction without first defining the reduced product with respect to an arbitrary filter. Such reduced products do play a role in some areas of analysis and thus we remedy this gap in Section 11.5 by introducing the reduced product of metric structures.

11.1. Definition of the metric ultraproduct

Suppose that $(M_i, d_i)_{i \in I}$ is a family of metric spaces and that \mathcal{U} is an ultrafilter on I . We would like to form the metric ultraproduct $\prod_{\mathcal{U}} M_i$ which should be, once again, a metric space. We take our cue from the discrete setting described in metric language: if each M_i were instead just a set, we could consider them as metric spaces as equipped with the discrete metric, that is, distinct points are considered at distance 1 from each other. Then in forming the ultraproduct, we consider the usual cartesian product $\prod_{i \in I} M_i$ and decide whether or not to identify two elements $a, b \in \prod_{i \in I} M_i$ by calculating $d(a, b) := \lim_{\mathcal{U}} d_i(a(i), b(i))$, where d_i denotes the discrete metric

on M_i . Indeed, $d(a, b) = 0$ if and only if $d_i(a(i), b(i)) = 0$ for \mathcal{U} -almost all $i \in I$, that is, if and only if $a_i = b_i$ for \mathcal{U} -almost all $i \in I$, in which case we identify a and b in the ultraproduct.

Proceeding naïvely along these lines, given a family $(M_i, d_i)_{i \in I}$ of metric spaces and an ultrafilter \mathcal{U} on I , we define d on $\prod_{i \in I} M_i$ by setting

$$d(a, b) := \lim_{\mathcal{U}} d_i(a(i), b(i)).$$

In analogy with the discrete case, it seems that we should quotient out by the elements of distance 0 to each other. While we will eventually perform such a quotient, we must deal with the following issue:

- If the sequence of real numbers $(d_i(a(i), b(i)))_{i \in I}$ is bounded, say by K , then the above ultralimit exists by the compactness of the interval $[0, K]$. However, that sequence may not be bounded, whence the above procedure would assign the two sequences a distance of ∞ .

There are two ways to take care of this issue. One way is to simply only consider an ultraproduct of metric spaces where there is a uniform bound on the diameters of all the spaces involved, that is, there is $K > 0$ such that $d_i(x, y) \leq K$ for all $i \in I$ and all $x, y \in M_i$. We choose to take a different route (which subsumes the previous suggestion):

Definition 11.1.1. A **pointed metric space** is a triple (M, d, o) , where (M, d) is a metric space and $o \in M$ is a distinguished point.

If each metric space (M_i, d_i) above comes equipped with a distinguished point $o(i) \in M_i$, then we can consider the set of tuples $a \in \prod_{i \in I} M_i$ such that $\sup_{i \in I} d_i(a(i), o(i)) < \infty$. The issue above now disappears if we only define d on such sequences, for, by the triangle inequality, it follows that

$$\sup_{i \in I} d_i(a(i), b(i)) \leq \sup_{i \in I} d_i(a(i), o(i)) + \sup_{i \in I} d_i(b(i), o(i)) < \infty.$$

Assuming that we only consider such sequences, the function d above is now easily verified to be a pseudo-metric, that is, a function that satisfies all of the definitions of a metric except that distinct points may be assigned distance 0 from one another. For example, if $I = \mathbb{N}$, all of the M_i above are $[0, 1]$, and $a(i) = \frac{1}{i}$ and $b(i) = 0$ for all $i \in \mathbb{N}$, then $d(a, b) = 0$ even though $a \neq b$.

One now proceeds by doing what one does when one encounters a pseudo-metric space and wants a metric space: just mod out by the relation $d = 0$. That is, look at equivalence classes of sequences a under the equivalence relation $d(a, b) = 0$; the pseudo-metric now induces a metric on the set equivalence classes. This is tantamount to, in the discrete case, modding out by \mathcal{U} -almost everywhere agreement. In fact, as discussed above, in the

case that each M_i is equipped with the discrete metric, these two operations coincide.

We summarize the above discussion:

Definition 11.1.2. Let $(M_i, d_i, o(i))_{i \in I}$ be a family of pointed metric spaces and \mathcal{U} an ultrafilter on I . We set

$$\ell^\infty(M_i) := \ell^\infty(M_i, d_i, o(i)) := \left\{ a \in \prod_{i \in I} M_i : \sup_{i \in I} d(a(i), o(i)) < \infty \right\}.$$

We define the pseudo-metric d on $\ell^\infty(M_i)$ by setting

$$d(a, b) := \lim_{\mathcal{U}} d_i(a(i), b(i)).$$

The **ultraproduct of the family** $(M_i, d_i, o(i))_{i \in I}$ **with respect to** \mathcal{U} is the set of equivalence classes of $\ell^\infty(M_i)$ modulo the equivalence relation $d = 0$ with the induced metric. We denote the ultraproduct by $\prod_{\mathcal{U}}(M_i, d_i, o(i))$ or sometimes $\prod_{\mathcal{U}}(M_i, o(i))$ or even $\prod_{\mathcal{U}} M_i$ if there is no source of confusion. Given $a \in \ell^\infty(M_i)$, we let $[a]_{\mathcal{U}}$ denote its equivalence class in $\prod_{\mathcal{U}} M_i$.

When writing the ultraproduct in the last way, one has to make sure to point out that one is thinking of a metric ultraproduct and not the classical ultraproduct discussed earlier in the book.

Remark 11.1.3. If the family $(M_i, d_i)_{i \in I}$ has uniformly bounded diameter, then $\ell^\infty(M_i) = \prod_{i \in I} M_i$ regardless of the choice of basepoint. In this case, we drop any mention of the basepoint and simply write $\prod_{\mathcal{U}}(M_i, d_i)$ or just $\prod_{\mathcal{U}} M_i$.

Exercise 11.1.4. Consider a metric space (M, d) and two basepoints $o, o' \in M$. Show that, for any index set I and any ultrafilter \mathcal{U} on I , the ultraproducts $\prod_{\mathcal{U}}(M, d, o)$ and $\prod_{\mathcal{U}}(M, d, o')$ are the same.

The preceding exercise allows us to make the following definition:

Definition 11.1.5. Given a metric space (M, d) and an ultrafilter \mathcal{U} over some index set, the **ultrapower of** (M, d) **with respect to** \mathcal{U} , denoted $(M, d)^{\mathcal{U}}$ or simply $M^{\mathcal{U}}$, is the ultraproduct $\prod_{\mathcal{U}}(M, d, o)$, where $o \in M$ is any basepoint.

Given any $a \in M$, we can once again consider the element of M^I that is constantly equal to a . This element belongs to $\ell^\infty(M)$, whence we may consider its image $[a]_{\mathcal{U}}$ in $M^{\mathcal{U}}$. We define the **diagonal embedding** to be the map $d : M \rightarrow M^{\mathcal{U}}$ given by $d(a) := [a]_{\mathcal{U}}$. (It is a bit unfortunate that we are using the letter d for both the induced metric on the ultrapower and the diagonal embedding, but hopefully each use of d will be clear from context.)

Exercise 11.1.6. Prove that $d : M \rightarrow M^{\mathcal{U}}$ is an isometric embedding.

In the discrete case, the diagonal embedding is onto for every \mathcal{U} precisely when the structure involved is finite. What is the corresponding property of M in the metric case? Well, note that $[a]_{\mathcal{U}} = d(b)$ if and only if $\lim_{\mathcal{U}} a(i) = b$. So, for the diagonal embedding to always be onto is the same as saying that every sequence in $\ell^\infty(M)$ has an ultralimit with respect to every ultrafilter. This almost sounds like our characterization of compactness from Section 3.1 except that it only asks for bounded sequences to have an ultralimit with respect to every ultrafilter. In other words, it asks for every bounded set to be compact. This notion has a name:

Definition 11.1.7. A metric space is **proper** (or **Heine-Borel**) if every bounded set is compact.

We can thus conclude:

Theorem 11.1.8. *If (M, d) is a metric space, then $d : M \rightarrow M^{\mathcal{U}}$ is surjective for every ultrafilter \mathcal{U} if and only if (M, d) is proper.*

11.2. Metric ultraproducts and nonstandard hulls of metric spaces

In this short section, we relate the metric ultraproduct construction with a known construction in nonstandard analysis, namely the nonstandard hull construction. This perspective will be especially useful in the next chapter on Gromov's theorem.

To begin, we work in some nonstandard extension of our universe and suppose that (M, d) is an **internal metric space**, that is, M is an internal set (see Definition 9.6.7) and $d : M \times M \rightarrow \mathbb{R}^*$ is an internal function that satisfies the usual axioms of a metric (except now that it takes values in \mathbb{R}^* instead of \mathbb{R}). Fix also a basepoint $o \in M$. Now consider

$$M_{\text{fin}} := M_{\text{fin}, o} := \{x \in M : d(a, o) \in \mathbb{R}_{\text{fin}}\}.$$

For $a, b \in M_{\text{fin}}$, one can define $\hat{d}(a, b) := \text{st}(d(a, b))$. It follows that \hat{d} is a pseudo-metric on M_{fin} and the resulting quotient is a metric space, called the **nonstandard hull** of the internal metric space (M, d) .

Suppose now that the nonstandard universe is constructed by taking the ultrapower of the standard universe with respect to some ultrafilter \mathcal{U} . In this case, for (M, d) to be an internal metric space is precisely the same as saying that $M = \prod_{\mathcal{U}} M_i$ (as a discrete ultraproduct) for some family $(M_i, d_i)_{i \in I}$ of metric spaces and that $d(a, b) := [d_i(a(i), b(i))] \in \mathbb{R}^{\mathcal{U}} = \mathbb{R}^*$ (again, discrete ultraproduct). Moreover, fixing a basepoint $o \in M$, we have that $M_{\text{fin}, o}$ consists of those elements $[a]_{\mathcal{U}} \in \prod_{\mathcal{U}} M_i$ (discrete ultraproduct) such that $a \in \ell^\infty(M_i, d_i, o(i))$. Since taking standard part in this context is

the same as taking an ultralimit, we see that the nonstandard hull of (M, d) is nothing more than the metric ultraproduct $\prod_{\mathcal{U}}(M_i, d_i, o(i))$.

11.3. Completeness properties of the metric ultraproduct

In Chapter 8, we saw that discrete ultraproducts are suitably rich (in the precise sense of saturation). Once an appropriate logical formalism for the metric ultraproduct is given (see the next section on continuous logic), the analogous fact for metric structures also holds. In this section, we will see a special case of this fact, namely that ultraproducts of metric spaces are often complete. Here is a first version of this phenomenon:

Theorem 11.3.1. *Suppose that $(M_i, o(i))_{i \in I}$ is a family of complete pointed metric spaces and that \mathcal{U} is an ultrafilter on I . Then $\prod_{\mathcal{U}}(M_i, o(i))$ is complete.*

Proof. Suppose that $([x_n]_{\mathcal{U}})_{n \in \mathbb{N}}$ is a Cauchy sequence from $\prod_{\mathcal{U}}(M_i, o(i))$; we wish to show that $([x_n]_{\mathcal{U}})_{n \in \mathbb{N}}$ converges. Without loss of generality, we may assume that $d([x_n]_{\mathcal{U}}, [x_{n+1}]_{\mathcal{U}}) < 2^{-n}$ for each n . (This is a basic result from real analysis, namely that every Cauchy sequence has such a “fast” subsequence.) For each $m \in \mathbb{N}$, let

$$A_m := \{i \in I : d(x_n(i), x_{n+1}(i)) < 2^{-n} \text{ for } n = 0, \dots, m\}.$$

By assumption, each $A_m \in \mathcal{U}$. We now define an element $y \in \ell^\infty(M_i)$ such that $\lim_{n \rightarrow \infty} [x_n]_{\mathcal{U}} = [y]_{\mathcal{U}}$. For $i \in A_m \setminus A_{m+1}$, set $y(i) := x_{m+1}(i)$. Note that in this case that $d(x_m(i), y(i)) < 2^{-m}$. If $i \in \bigcap_{m \in \mathbb{N}} A_m$, then $(x_n(i))_{n \in \mathbb{N}}$ is a Cauchy sequence in M_i , whence converges to some $y(i) \in M_i$ by completeness of M_i . Note in this case that $d(x_m(i), y(i)) < 2^{-m+1}$. (Think geometric series.) Finally, if $i \notin A_0$, set $y(i) := o(i)$.

We leave it to the reader to check that y really is an element of $\ell^\infty(M_i)$. To see that $\lim_{n \rightarrow \infty} [x_n]_{\mathcal{U}} = [y]_{\mathcal{U}}$, note that, given $m \in \mathbb{N}$, if $i \in A_m$, then $d(x_m(i), y(i))$ is either $d(x_m(i), x_{m+1}(i)) < 2^{-m}$ or $d(x_m(i), y(i)) < 2^{-m+1}$. In either event, for a \mathcal{U} -large set of i , we have that $d(x_m(i), y(i)) < 2^{-m+1}$, whence $d([x_m]_{\mathcal{U}}, [y]_{\mathcal{U}}) < 2^{-m+1}$. The desired result now follows. \square

Notice that the previous proof only used completeness of the factor spaces to deal with the case that some i belonged to $\bigcap_{m \in \mathbb{N}} A_m$. Thus, essentially the same proof yields:

Theorem 11.3.2. *Suppose that $(M_i, o(i))_{i \in I}$ is any family of pointed metric spaces and that \mathcal{U} is a countably incomplete ultrafilter on I (see Definition 6.6.3). Then $\prod_{\mathcal{U}}(M_i, o(i))$ is complete.*

Proof. Let $(B_m)_{m \in \mathbb{N}}$ be a family from \mathcal{U} such that $\bigcap_{m \in \mathbb{N}} B_m = \emptyset$. For $m \in \mathbb{N}$, define $A_m := \{i \in B_m : d(x_n(i), x_{n+1}(i)) < 2^{-n} \text{ for } n = 0, \dots, m\}$.

Now one proceeds exactly as in the previous proof, never running into the case that $i \in \bigcap_{m \in \mathbb{N}} A_m$ and thus never needing to quote the completeness of the space M_i . \square

Note that if \mathcal{U} is countably complete, then the preceding theorem may fail as then $M^{\mathcal{U}} \cong M$ for such \mathcal{U} regardless of the space M . However, as we remarked earlier, the existence of nonprincipal countably complete ultrafilters enters the territory of set theory, so the above assumption in the previous theorem can be viewed as a mild one.

Exercise 11.3.3. Suppose that $(M_i, d_i)_{i \in I}$ is a family of metric spaces and $X_i \subseteq M_i$ is dense for all $i \in I$. If \mathcal{U} is a countably incomplete ultrafilter on I , prove that the natural inclusion $\prod_{\mathcal{U}} X_i \subseteq \prod_{\mathcal{U}} M_i$ is actually an equality.

Corollary 11.3.4 (Existence of completion). *Let M be a metric space and $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. Then the closure \overline{M} of the diagonal image of M in $M^{\mathcal{U}}$ is the completion of M , that is, a complete metric space in which M embeds densely. This can be described concretely as*

$$\overline{M} := \{[x]_{\mathcal{U}} \in M^{\mathcal{U}} : \text{for all } \epsilon > 0, \text{ there is } y \in M \\ \text{such that } d(x(n), y) < \epsilon \text{ for } \mathcal{U}\text{-almost all } n\}.$$

In general, $M^{\mathcal{U}}$ is much bigger than \overline{M} . For example, by Exercise 11.3.3 above, $\mathbb{Q}^{\mathcal{U}} = \mathbb{R}^{\mathcal{U}}$, but $\overline{\mathbb{Q}} = \mathbb{R}$.

We end this section with the metric analogue of Theorem 6.8.3 on “sizes” of metric ultraproducts. To keep things simple, we consider the case of a family of metric spaces of uniformly bounded diameter. Recall that the *density character* of a metric space (M, d) is the smallest cardinality of a dense subset of M ; this is the appropriate measure of the size of a metric space. Also, given $\epsilon > 0$, recall that a subset F of M is an ϵ -net if, given any $x \in M$, there is $y \in F$ such that $d(x, y) < \epsilon$.

Theorem 11.3.5. *Suppose that $(M_i, d_i)_{i \in I}$ is a family of metric spaces of uniformly bounded diameter and \mathcal{U} is a countably incomplete ultrafilter on I . Then either $\prod_{\mathcal{U}} M_i$ is compact or else has density character at least \mathfrak{c} .*

Proof. For $\epsilon > 0$, let $n_{\epsilon}(i)$ be the minimum size of a finite ϵ -net in M_i , if such a finite net exists; otherwise, set $n_{\epsilon}(i) := \infty$.

First suppose that $\lim_{\mathcal{U}} n_{\epsilon}(i) < \infty$ for all ϵ . We claim that in this case, $M := \prod_{\mathcal{U}} M_i$ is compact. Indeed, for each i , let $A_{\epsilon, i} \subseteq M_i$ be a an ϵ -net of size $n_{\epsilon}(i)$. Let $A_{\epsilon} := \prod_{\mathcal{U}} A_{\epsilon, i}$. Then A_{ϵ} is finite (of size $\lim_{\mathcal{U}} n_{\epsilon}(i)$) and is a 2ϵ -net in M . Thus, M is totally bounded. Since M is also complete, it follows that M is compact.

Now suppose that there is $\epsilon > 0$ such that $\lim_{\mathcal{U}} n_{\epsilon}(i) = \infty$. This means that there are sets $A_i \subseteq M_i$ with all elements of A_i at least ϵ -apart and such

that $\lim_{\mathcal{U}} |A_i| = \infty$. Note then that in this case, the discrete and metric ultraproducts of the sequence $(A_i)_{i \in I}$ agree and that this ultraproduct has size at least \mathfrak{c} by Theorem 6.8.3. Since all elements of this ultraproduct are at least ϵ -apart, we see that M has density character at least \mathfrak{c} . \square

11.4. Continuous logic

As we will see in the coming chapters, we will need to consider metric spaces with extra structure just as we encountered plain sets with extra structure in Part 2 of this book. We now briefly discuss a logic suitable for handling such **metric structures**. While many predecessors to this logic have been around for quite some time, the current incarnation of **continuous logic** has only been around for 15 years and has achieved much of its success due to its striking similarities with classical logic. (In fact, we will soon see that continuous logic is a direct generalization of classical logic in a precise sense.)

We motivate the introduction of continuous logic by desiring that the metric ultraproduct described above be the correct ultraproduct construction for our logic, meaning that we would like the ultraproduct of a family of structures to once again be a structure. For simplicity, let us suppose that we have a family $(M_i, d_i, f_i)_{i \in I}$ of metric spaces of diameter at most 1, each equipped with a distinguished function $f_i : M_i \rightarrow M_i$. (Of course, this is analysis, so in all likelihood these functions will at least be continuous, but let us see how the ultraproduct construction guides us to the correct assumption on the f_i 's.)

The naïve idea would be to define a function $f : M \rightarrow M$, where $M := \prod_{\mathcal{U}} M_i$ (metric ultraproduct), by setting $f([a]_{\mathcal{U}}) := [f \circ a]_{\mathcal{U}}$, in analogy with the discrete setting. In order for this to be well-defined this time, we would need to know that if $a, b \in \prod_i M_i$ are such that $\lim_{\mathcal{U}} d_i(a(i), b(i)) = 0$, then $\lim_{\mathcal{U}} d_i(f(a(i)), f(b(i))) = 0$. If this were not the case, then there would be $\epsilon > 0$ such that, no matter how small a $\delta > 0$ we consider, we would have, for \mathcal{U} -almost all i , that $d(a(i), b(i)) < \delta$ and yet $d(f_i(a(i)), f_i(b(i))) \geq \epsilon$.

Turning this argument on its head, we see that if we want to ensure that the function f defined on the ultraproduct is well-defined, then we should assume that each f_i is uniformly continuous and, moreover, that the “witness” to uniform continuity is the same for each f_i in a sense to be made precise by the following definition. We warn the reader that a seemingly strange asymmetry in our use of weak and strong inequalities will appear in this definition, but we promise a satisfactory explanation for this shortly.

Definition 11.4.1.

- (1) A **modulus of uniform continuity** is simply a function $\Delta : (0, 1] \rightarrow (0, 1]$.
- (2) If $f : M \rightarrow N$ is a function between metric spaces and Δ is a modulus of uniform continuity, then we say that Δ is a **modulus of uniform continuity for f** if, for every $a, b \in M$ and every $\epsilon \in (0, 1]$, if $d(a, b) < \Delta(\epsilon)$, then $d(f(a), f(b)) \leq \epsilon$.

Notice that if f has a modulus of uniform continuity, then f is uniformly continuous. Now notice that, in the above argument, if there was a modulus of uniform continuity Δ such that Δ was a modulus of uniform continuity for each f_i , then f is a well-defined function as we had hoped for. However, something even better is true, namely that Δ is once again a modulus of uniform continuity for f : if $d([a]_{\mathcal{U}}, [b]_{\mathcal{U}}) < \Delta(\epsilon)$, then $d(a(i), b(i)) < \Delta(\epsilon)$ for \mathcal{U} -almost all i , whence $d(f(a(i)), f(b(i))) \leq \epsilon$ for \mathcal{U} -almost all i , which implies that $d(f([a]_{\mathcal{U}}), f([b]_{\mathcal{U}})) \leq \epsilon$. Notice that if we had modified the definition of modulus of uniform continuity for a function by replacing “ $< \Delta(\epsilon)$ ” with “ $\leq \Delta(\epsilon)$ ” and/or “ $\leq \epsilon$ ” with “ $< \epsilon$ ”, then the preceding proof would break down and Δ need not be a modulus of uniform continuity for f anymore.

Of course we will want to consider functions $M^n \rightarrow M$ rather than just $M \rightarrow M$ and the same ideas persist. In order for the above discussion to be applicable, we will need to endow M^n with a particular metric. While there are many ways of doing this, we always assume that M^n is equipped with its so-called max metric, that is, $d(\vec{a}, \vec{b}) = \max_{1 \leq i \leq n} d(a_i, b_i)$.

We now move on to considering distinguished predicates. Here, there is a substantial shift in perspective. In classical logic, $=$ is a distinguished (logical) predicate. For metric spaces, considering $=$ as a distinguished predicate is not a good idea, especially if one is hoping for some form of the Łoś theorem to hold. Indeed, it is often the case that, in a metric ultraproduct, we have that $[a]_{\mathcal{U}} = [b]_{\mathcal{U}}$ and yet $a(i) \neq b(i)$ for all i . Of course, by its very definition, the metric in the ultraproduct and the metric in the various factors are related and so we should consider the metric symbol d as a distinguished (logical) predicate. But the metric does not take values in the set $\{0, 1\}$, but rather the interval $[0, 1]$. (Recall our simplifying assumption that our metric spaces have diameter bounded by 1.) Thus a paradigm shift: predicates should now be functions into $[0, 1]$. If we now consider a family $(M_i, d_i, P_i)_{i \in I}$, where each $P_i : M \rightarrow [0, 1]$, and attempt to define P on M by $P([a]_{\mathcal{U}}) := [P(a(i))]_{\mathcal{U}}$, then we do not quite get another structure of the same form as P now takes values in $[0, 1]^{\mathcal{U}}$ rather than $[0, 1]$. However, $[0, 1]$ is compact, whence $\lim_{\mathcal{U}}$ provides an isomorphism between $[0, 1]^{\mathcal{U}}$ and

$[0, 1]$. Thus, we *redefine* $P([a]_{\mathcal{U}}) := \lim_{\mathcal{U}} P_i(a(i))$ and now we get a function $P : M \rightarrow [0, 1]$, a function of the same kind. An argument as before shows that to have a well-defined function P , then we should insist that there be a function Δ that is a modulus of uniform continuity for each P_i and, after insisting upon this, we then have that Δ is a modulus of uniform continuity for P as well. Once again, all of this holds for predicates $M^n \rightarrow [0, 1]$ as well.

Now that we know what we want our structures to be, it is clear what information a language for continuous logic should prescribe:

Definition 11.4.2. A **continuous language** \mathcal{L} consists of function and relation symbols along with the following information:

- For every function symbol F , \mathcal{L} should provide an arity n_F and a modulus of uniform continuity Δ_F .
- For every predicate symbol P , \mathcal{L} should provide an arity n_P and a modulus of uniform continuity Δ_P .

As usual, constant symbols can be treated as function symbols of arity 0 (and thus there is no modulus of uniform continuity requirement; they simply name elements of the structure).

Definition 11.4.3. Suppose that \mathcal{L} is a continuous language. An **\mathcal{L} -structure** \mathcal{M} is a *complete* metric space (M, d) , together with:

- For each function symbol F , we have a distinguished function $F^{\mathcal{M}} : M^{n_F} \rightarrow M$ which has Δ_F as a modulus of uniform continuity.
- For each predicate symbol P , we have a distinguished function $P^{\mathcal{M}} : M^{n_P} \rightarrow [0, 1]$ which has Δ_P as a modulus of uniform continuity.

Remark 11.4.4. It might seem a bit artificial to ask that the the metric spaces underlying our structures be complete. However, there is no real loss of generality to assume this. Indeed, one can show that the completion of a metric structure in the above sense whose underlying universe need not be complete is naturally a metric structure in the obvious way and that these two structures are “indistinguishable” in a precise model-theoretic sense. In fact, one could even start with a structure whose underlying space is only a *pseudo-metric* space, mod out by the relation of having distance 0, and then complete, and still one would arrive at a structure which is logically the same as the original structure. See [9] for more details.

Exercise 11.4.5. Suppose that \mathcal{L} is a continuous language, $(\mathcal{M}_i)_{i \in I}$ a family of \mathcal{L} -structures, and \mathcal{U} an ultrafilter on I . Guided by our earlier discussion, explain how to make the metric ultraproduct $\prod_{\mathcal{U}} M_i$ into the universe of

an \mathcal{L} -structure $\prod_{\mathcal{U}} \mathcal{M}_i$, called the **ultraproduct of the family** $(\mathcal{M}_i)_{i \in I}$ **with respect to** \mathcal{U} .

Now that we know what languages and structures are, we should describe formulae and their interpretations. Fix a continuous language \mathcal{L} . We define \mathcal{L} -terms just as usual: they are function symbols applied to constant symbols and variables. We define the interpretation $t^{\mathcal{M}}$ of an \mathcal{L} -term t in an L -structure \mathcal{M} just as in the classical case.

Exercise 11.4.6. For each \mathcal{L} -term $t(x_1, \dots, x_n)$, prove that there is a modulus of uniform continuity Δ_t such that, for every \mathcal{L} -structure \mathcal{M} , the function $t^{\mathcal{M}} : M^n \rightarrow M$ has Δ_t as a modulus of uniform continuity. (*Hint.* Proceed by induction on the complexity of t .)

We now describe \mathcal{L} -formulae. In analogy with classical logic, and with our earlier decision to replace $=$ with d , \mathcal{L} -atomic formulae are expressions of the form $P(t_1, \dots, t_n)$ and $d(t_1, t_2)$, where t_1, \dots, t_n are \mathcal{L} -terms and P is an n -ary predicate symbol. It is clear how we should interpret these formulae: if each t_i has free variables amongst x_1, \dots, x_m and $\varphi(x_1, \dots, x_m) = P(t_1(x_1, \dots, x_m), \dots, t_n(x_1, \dots, x_m))$, then $\varphi^{\mathcal{M}} : \mathcal{M}^m \rightarrow [0, 1]$ is given by $\varphi^{\mathcal{M}}(a_1, \dots, a_m) = P^{\mathcal{M}}(t_1^{\mathcal{M}}(a_1, \dots, a_m), \dots, t_n^{\mathcal{M}}(a_1, \dots, a_m))$. (Thinking of the metric symbol as a distinguished binary predicate symbol, this definition covers all atomic formulae.). Once again, we note that, our “truth values” for formulae no longer land in the set $\{0, 1\}$, but rather the interval $[0, 1]$. (For this reason, it sometimes helps to abandon the idea that formulae are true or false; they simply, upon interpretation, return a number in $[0, 1]$.)

We now come to the discussion of connectives. Given pre-existing formulae $\varphi_1, \dots, \varphi_n$, we are going to want to combine them in some way to create a new formula. Since the formulae, once interpreted, take values in $[0, 1]$, we should plug them into some function $u : [0, 1]^n \rightarrow [0, 1]$ to create the new formula $u(\varphi_1, \dots, \varphi_n)$. And since we are doing continuous logic, the function u should probably be continuous. But which continuous functions should we allow as connectives?

In the classical case, depending on presentation, one usually restricts formulae to those built with the connectives $\neg : \{0, 1\} \rightarrow \{0, 1\}$ and $\vee : \{0, 1\}^2 \rightarrow \{0, 1\}$. (Some authors choose to also include some or all of \wedge , \rightarrow , and \leftrightarrow as well. Also, one can simply use the single connective $|$, the Sheffer stroke.) However, a basic result of propositional logic says that, with the above connectives, one can actually generate every function $\{0, 1\}^n \rightarrow \{0, 1\}$ for all n . Thus, if one chose to do classical logic allowing all functions $\{0, 1\}^n \rightarrow \{0, 1\}$ as connectives, then nothing would change.

With the preceding paragraph in mind, we choose to allow *all* continuous functions as connectives: if $\varphi_1, \dots, \varphi_n$ are \mathcal{L} -formulae and $u : [0, 1]^n \rightarrow [0, 1]$

is any continuous function, then $\varphi := u(\varphi_1, \dots, \varphi_n)$ is also an \mathcal{L} -formula. Presuming we already know how to interpret $\varphi_1^{\mathcal{M}}, \dots, \varphi_n^{\mathcal{M}}$, then we interpret $\varphi^{\mathcal{M}}$ as $u(\varphi_1^{\mathcal{M}}, \dots, \varphi_n^{\mathcal{M}})$.

We should make two comments about the above choice. First, even if the language \mathcal{L} is countable, the number of \mathcal{L} -formulae is uncountable. Of course, in analysis, cardinality is not necessarily the right notion but rather density character, as was discussed in the previous section. In this regards, for each n , there is a *countable* set of continuous functions $[0, 1]^n \rightarrow [0, 1]$ that is dense in the set of all such continuous functions and the set of \mathcal{L} -formulae constructed from those is dense (in a certain precise sense) in the set of all \mathcal{L} -formulae. In many respects, working with these formulae is often enough. In the case that the language is countable, one then has a countable dense set of formulae.

On the other hand, our choice of connectives is actually not large enough for some of the more advanced model-theoretic considerations. Indeed, one often needs to consider *infinitary formulae* in the sense that if $\varphi_1, \varphi_2, \dots$ are countably many formulae and $u : [0, 1]^{\mathbb{N}} \rightarrow [0, 1]$ is a continuous function, then $u(\varphi_1, \varphi_2, \dots)$ is also a formula. For example, $\sum_{n=1}^{\infty} 2^{-n} \varphi_n$ would also be a formula. Such generalized formulae are called **definable predicates**, but we will not say anything further about them.

We now come to the choice of quantifiers. We aim to motivate this by analogy with the discrete world, but in doing so we ask the reader to abandon a traditional identification (one we made above, in fact): usually the truth value T is identified with the number 1 and the truth value F identified with the number 0. In continuous logic, we often make the reverse identification. One rationale for doing so is that, in a metric space, the statement $a = b$ is true if and only if the quantity $d(a, b)$ equals 0. Given this reversal in viewpoint, we now see that the statement “for all x , $\varphi(x) = 0$ ” is equivalent to the statement $\sup_x \varphi(x) = 0$. (Recall our convention that formulae take values in $[0, 1]$.) So, in this regards, the “quantifier” \sup_x acts like the universal quantifier $\forall x$. One can make a similar case for the use of the quantifier \inf_x except, unfortunately, it only behaves like an approximate existential quantifier: $\inf_x \varphi(x) = 0$ if and only if, for every ϵ , there is x such that $\varphi(x) < \epsilon$. (If ϵ is small, we think of this as saying that such an x almost makes φ true.) In ω -saturated structures (which are defined in a way analogous to the discrete setting), the quantifier \inf_x does behave exactly like an existential quantifier, that is, $\inf_x \phi(x) = 0$ if and only if there is x such that $\phi(x) = 0$.

The preceding paragraph in mind, we now declare that if φ is an \mathcal{L} -formula and x is a single variable, then the expressions $\sup_x \varphi$ and $\inf_x \varphi$ are once again \mathcal{L} -formulae. The interpretations are given by $(\sup_x \varphi)^{\mathcal{M}}(b) :=$

$\sup\{\varphi^{\mathcal{M}}(a, b) : a \in M\}$ and likewise for inf. These are well-defined as the values of φ lie in $[0, 1]$.

This completes the recursive definition of \mathcal{L} -formulae.

Exercise 11.4.7. For every \mathcal{L} -formula $\varphi(x_1, \dots, x_n)$, prove that there is a modulus of uniform continuity Δ_φ such that, for all \mathcal{L} -structures \mathcal{M} , we have that Δ_φ is a modulus of uniform continuity for the function $\varphi^{\mathcal{M}} : M^n \rightarrow [0, 1]$.

One can now state the continuous logic version of Łoś's theorem:

Theorem 11.4.8 (Łoś's theorem for continuous logic). *Suppose that \mathcal{L} is a continuous language, $(\mathcal{M}_i)_{i \in I}$ is a family of \mathcal{L} -structures, and \mathcal{U} is an ultrafilter on I . Set $\mathcal{M} := \prod_{\mathcal{U}} \mathcal{M}_i$. Further suppose that $\varphi(x_1, \dots, x_m)$ is an \mathcal{L} -formula and $[a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}} \in \prod_{\mathcal{U}} M_i$. Then*

$$\varphi^{\mathcal{M}}([a_1]_{\mathcal{U}}, \dots, [a_m]_{\mathcal{U}}) = \lim_{\mathcal{U}} \varphi^{\mathcal{M}_i}(a_1(i), \dots, a_m(i)).$$

Exercise 11.4.9. Prove the previous theorem.

We now give some of the basic model-theoretic definitions in this setting:

Definition 11.4.10. Suppose that \mathcal{L} is a continuous language and \mathcal{M} and \mathcal{N} are \mathcal{L} -structures.

- (1) A **embedding** $i : \mathcal{M} \rightarrow \mathcal{N}$ is a function $i : M \rightarrow N$ such that:
- for every n -ary function symbol F and every $a_1, \dots, a_n \in M$, we have

$$i(F^{\mathcal{M}}(a_1, \dots, a_n)) = F^{\mathcal{N}}(i(a_1), \dots, i(a_n));$$

- for every n -ary predicate symbol P and every $a_1, \dots, a_n \in M$, we have

$$P^{\mathcal{M}}(a_1, \dots, a_n) = P^{\mathcal{N}}(i(a_1), \dots, i(a_n)).$$

- (2) An **elementary embedding** is an embedding $i : \mathcal{M} \rightarrow \mathcal{N}$ for which, given any \mathcal{L} -formula $\varphi(x_1, \dots, x_n)$ and any $a_1, \dots, a_n \in M$, we have

$$\varphi^{\mathcal{M}}(a_1, \dots, a_n) = \varphi^{\mathcal{N}}(i(a_1), \dots, i(a_n)).$$

- (3) \mathcal{M} is a **substructure** of \mathcal{N} if $M \subseteq N$ and the inclusion map $i : M \rightarrow N$ is an embedding $i : \mathcal{M} \rightarrow \mathcal{N}$. If i is moreover an elementary embedding, then we say that \mathcal{M} is an **elementary substructure** of \mathcal{N} , denoted $\mathcal{M} \preceq \mathcal{N}$.

- (4) \mathcal{M} is **elementarily equivalent** to \mathcal{N} , denoted $\mathcal{M} \equiv \mathcal{N}$, if, for every \mathcal{L} -sentence σ , $\sigma^{\mathcal{M}} = \sigma^{\mathcal{N}}$.

Note that an embedding is automatically an isometric embedding (by considering $P = d$). Note also that, by Łoś's theorem for continuous logic, we have that the diagonal embedding is an elementary embedding.

An interesting feature of continuous logic is that it is a *positive* logic in that one is unable to talk about negations. Still, there are often substitutes for related things, such as implications, as we will see in the following example. First, define the function $\dot{-} : [0, 1]^2 \rightarrow [0, 1]$ to be given by $x \dot{-} y := \max(x - y, 0)$.

Example 11.4.11. Suppose that $\varphi(x)$ and $\psi(x)$ are \mathcal{L} -formulae, \mathcal{M} is an \mathcal{L} -structure, and $r, s \in [0, 1]$. Suppose that we want to say: for all $a \in M$, if $\varphi(a) < r$, then $\psi(a) \leq s$. This can be accomplished by asserting $(\sup_x \min(r \dot{-} \varphi(x), \psi(x) \dot{-} s))^{\mathcal{M}} = 0$. Indeed, suppose that $a \in M$ and $\varphi(a) < r$. Then $r \dot{-} \varphi(a) > 0$. Thus, for the minimum to be 0, one must have $\psi(a) \dot{-} s = 0$, which is the same as saying that $\psi(a) \leq s$. Note that the function $z \mapsto r \dot{-} z : [0, 1] \rightarrow [0, 1]$ is a continuous function, whence $r \dot{-} \varphi(x)$ is a formula once again. Similarly, $\psi(x) \dot{-} s$ is also a formula. Since $\min : [0, 1]^2 \rightarrow [0, 1]$ is also a continuous function, the expression written above is indeed a formula in continuous logic. Note however that, in general, we could not express $\varphi(a) \leq r$ implying $\psi(a) \leq s$; the strong inequality was important in the argument above.

Earlier, we said that continuous logic is a generalization of ordinary first-order logic. When we say this, we mean that, if we view every classical structure as a metric structure by equipping the universe of the structure with the discrete metric, then many of the notions and results in continuous logic are simply generalizations of their discrete counterpart. For example:

Fact 11.4.12. Suppose that \mathcal{L} is a classical language and \mathcal{M} and \mathcal{N} are classical \mathcal{L} -structures. Then \mathcal{M} and \mathcal{N} are elementarily equivalent as discrete structures if and only if they are elementarily equivalent as metric structures when equipped with the discrete metric.

The proof of the previous fact is complicated by the fact that continuous logic allows more connectives, whence there are more sentences that one has to check when testing elementary equivalence. We refer the reader to [9] for more details.

Many of the notions and results from classical model theory have their continuous counterparts, and in many cases the continuous translation of a notion or result from classical model theory is routine to identify. That being said, there are certainly subtleties that arise in continuous logic that do not arise in the classical case, such as the notion of definable set in continuous logic as well as the Omitting Types Theorem in continuous logic. Since our

treatment here is meant to cover only the very basics, we refer the reader to [9] once again for more details.

While we have restricted our attention to languages where all predicates take values in $[0, 1]$, there is nothing fundamentally different about languages where the predicates are allowed to take values in arbitrary compact intervals $[a, b]$ instead (which may even vary as the predicates vary), although the details are a bit more messy. More care must be taken to treat structures where the predicates are allowed to take any real values, but we will say nothing about this case here.

Continuous logic has considerably expanded the class of mathematical objects that are now amenable to model-theoretic tools. In later chapters, we will meet objects like metric groups, Banach spaces, Hilbert spaces, and operator algebras which can now be studied through the model-theoretic lens. This is currently a very active and fruitful area of research.

11.5. Reduced products of metric structures

The reader may have noticed that, unlike in the case of discrete structures, we did not introduce the notion of reduced product of metric structures and instead skipped straight to the definition of ultraproduct. There is indeed a natural notion of reduced product of metric structures which we introduce in this section. For simplicity, we consider the case of $[0, 1]$ -valued structures.

To motivate the definition to come, consider the idea of how the discrete metric in a reduced product of discrete metric spaces is related to the discrete metric in the factors. More precisely, given a family $(M_i)_{i \in I}$ of sets equipped with the discrete metric, a filter \mathcal{F} on I , and $[a]_{\mathcal{F}}, [b]_{\mathcal{F}} \in \prod_{\mathcal{F}} M_i$, we have

$$\begin{aligned} d([a]_{\mathcal{F}}, [b]_{\mathcal{F}}) = 0 &\Leftrightarrow [a]_{\mathcal{F}} = [b]_{\mathcal{F}} \\ &\Leftrightarrow (\exists J \in \mathcal{F})(\forall i \in J)(a(i) = b(i)) \Leftrightarrow \inf_{J \in \mathcal{F}} \sup_{i \in J} d(a(i), b(i)) = 0. \end{aligned}$$

Definition 11.5.1. Given $f : I \rightarrow [0, 1]$ and a filter \mathcal{F} on I , we set

$$\limsup_{\mathcal{F}} f := \inf_{J \in \mathcal{F}} \sup_{i \in J} f(i).$$

Exercise 11.5.2. Suppose that $(M_i, d_i)_{i \in I}$ is a family of metric spaces with diameter bounded by 1. Define d on $\prod_{i \in I} M_i$ by setting $d(a, b) := \limsup_{\mathcal{F}} d_i(a(i), b(i))$. Prove that d is a pseudo-metric.

Definition 11.5.3. Given a family $(M_i, d_i)_{i \in I}$ of metric spaces with diameter bounded by 1 and a filter \mathcal{F} on I , we define the **reduced product**, denoted $\prod_{\mathcal{F}} M_i$, to be the metric space obtained from considering the pseudo-metric d from the previous exercise quotiented by the relation $d = 0$.

In order to see that this agrees with the notion of ultraproduct in the case that \mathcal{F} is an ultrafilter, the following alternative formula for $\limsup_{\mathcal{F}}$ is useful:

Lemma 11.5.4. *For any function $f : I \rightarrow [0, 1]$ and any filter \mathcal{F} on I , we have*

$$\limsup_{\mathcal{F}} f = \sup \left\{ \lim_{\mathcal{U}} f : \mathcal{F} \subseteq \mathcal{U} \in \beta I \right\}.$$

In particular, if \mathcal{F} is an ultrafilter, then $\limsup_{\mathcal{F}} f = \lim_{\mathcal{F}} f$.

Proof. Set $r := \sup\{\lim_{\mathcal{U}} f : \mathcal{F} \subseteq \mathcal{U} \in \beta I\}$. We first show that $r \leq \limsup_{\mathcal{F}} f$. Toward this end, fix $\mathcal{U} \in \beta I$ with $\mathcal{F} \subseteq \mathcal{U}$ and $J \in \mathcal{F}$; it suffices to show that $\lim_{\mathcal{U}} f \leq \sup_{i \in J} f(i)$. But this is clear: if $f(i) \leq s$ for all $i \in J$, then $\lim_{\mathcal{U}} f \leq s$.

For the other direction, fix $\epsilon > 0$ and for each $J \in \mathcal{F}$, take $i_J \in J$ such that $f(i_J) \geq \limsup_{\mathcal{F}} f - \epsilon$. Let $X := \{i \in I : f(i) \geq \limsup_{\mathcal{F}} f - \epsilon\}$. Then $\mathcal{F} \cup \{X\}$ generates a proper filter on I (as $i_J \in J \cap X$ for each $J \in \mathcal{F}$), whence there is $\mathcal{U} \in \beta I$ for which $\mathcal{F} \subseteq \mathcal{U}$ and $X \in \mathcal{U}$. Since $X \in \mathcal{U}$, we have $\lim_{\mathcal{U}} f \geq \limsup_{\mathcal{F}} f - \epsilon$. It follows that $r \geq \limsup_{\mathcal{F}} f - \epsilon$. Since $\epsilon > 0$ was arbitrary, we get $r \geq \limsup_{\mathcal{F}} f$, as desired. \square

Exercise 11.5.5. Suppose that \mathcal{L} is a continuous language, $(\mathcal{M}_i)_{i \in I}$ is a family of \mathcal{L} -structures, and \mathcal{F} is a filter on I . Set $M := \prod_{\mathcal{F}} M_i$. For each n -ary function symbol F in \mathcal{L} and $[a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}} \in M$, set

$$F^{\mathcal{M}}([a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}}) := [i \mapsto F^{\mathcal{M}_i}(a_1(i), \dots, a_n(i))]_{\mathcal{F}}$$

and for each n -ary predicate symbol P in \mathcal{L} and $[a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}} \in M$, set

$$P^{\mathcal{M}}([a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}}) := \limsup_{\mathcal{F}} P^{\mathcal{M}_i}([a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}}).$$

Verify that these definitions are independent of representatives and that the resulting structure \mathcal{M} is indeed an \mathcal{L} -structure.

The structure introduced in the previous exercise is called the **reduced product of the family $(\mathcal{M}_i)_{i \in I}$ with respect to the filter \mathcal{F}** . By Lemma 11.5.4, in the case of an ultrafilter, the reduced product structure coincides with the ultraproduct structure.

There is also a continuous version of the Feferman-Vaught theorem, whose classical version was proven in Section 6.11; see [66] for the continuous version.

11.6. Notes and references

The history of the metric ultraproduct is a bit murky. It appears that the first use of a metric ultraproduct (in disguise) is Wright's work [184] on

AW* algebras from 1954. Their use became more common in the 1960s and 1970s in the functional analysis communities. (We will be more explicit with references in the notes for Chapter 14.) The nonstandard hull construction is due to Luxemburg [115]. Continuous logic as presented here is due to Ben Yaacov and Usvyastov [10] and elaborated on considerably by Ben Yaacov, Berenstein, Henson, and Usvyatsov in the monograph [9]. A much earlier predecessor of this continuous logic was the one presented by Chang and Keisler in their book [27]. Henson's positive bounded logic [88] was also very popular in the community of those applying model-theoretic techniques in Banach space theory.

Asymptotic cones and Gromov's theorem

In this chapter, we present one of the more popular applications of the metric ultraproduct construction, namely van den Dries and Wilkie's metric ultraproduct take [181] on Gromov's asymptotic cone construction, which he used in proving his spectacular theorem on groups of polynomial growth. In Section 12.1 and Section 12.2, we describe some group-theoretic preliminaries and define the notion of the growth rate of a group, a crucial concept in explaining Gromov's result. In Section 12.3 we state Gromov's theorem on polynomial growth and give a detailed sketch of the proof based on the existence of a metric space satisfying certain properties; this metric space will end up being a particular asymptotic cone of the group. In Section 12.4 we give the definition of an asymptotic cone of a metric space and that of a group, both as metric ultraproducts and as nonstandard hulls. The latter construction was really the one used by van den Dries and Wilkie and is a bit more convenient for what follows. Section 12.5 and Section 12.6 complete the proof of Gromov's theorem by detailing the properties of the asymptotic cone needed in the proof; the properties in the former section are general while the properties in the latter need assumptions on the growth rate of the group. Section 12.7 presents a recent result of Hrushovski and Sapir, answering a question left open by van den Dries and Wilkie in [181], while Section 12.8 discusses the effect on the homeomorphism type of the asymptotic cone when changing some of the parameters involved in its construction.

12.1. Some group-theoretic preliminaries

In this section, we recall some definitions from group theory that we will need in this chapter. Given a group Γ and $X, Y \subseteq \Gamma$, we set $[X, Y]$ to be the subgroup of Γ generated by commutators $aba^{-1}b^{-1}$, with $a \in X$ and $b \in Y$. The **derived subgroup** of Γ is the group $\Gamma' := [\Gamma, \Gamma]$. Γ' is a normal subgroup of Γ and the **abelianization** of Γ , $\Gamma^{\text{ab}} = \Gamma/[\Gamma, \Gamma]$, is abelian. In fact, if N is a normal subgroup of Γ , then Γ/N is abelian if and only if Γ' is a subgroup of N , in which case the surjection $\Gamma \rightarrow \Gamma/N$ factors through Γ^{ab} . Note that Γ is abelian if and only if $\Gamma' = \{e\}$ if and only if $\Gamma^{\text{ab}} = \Gamma$.

To generalize being abelian, we consider groups such that the process of taking derived groups eventually trivializes: setting $\Gamma^{(1)} := \Gamma'$ and $\Gamma^{(n)} := (\Gamma^{(n-1)})'$, we say that Γ is **solvable** if $\Gamma^{(n)} = \{e\}$ for some $n \in \mathbb{N}$. Clearly abelian groups are solvable.

Example 12.1.1. $\Gamma = S_3 \times \mathbb{Z}$ is an infinite, nonabelian solvable group. Indeed, $\Gamma' = A_3 \times \{0\}$, which is abelian, whence $\Gamma^{(2)} = \{e\}$.

One can also consider the subgroups Γ_n of Γ defined by setting $\Gamma_1 := \Gamma'$ and $\Gamma_n := [\Gamma_{(n-1)}, \Gamma]$. Note that $\Gamma^{(n)}$ is a subgroup of Γ_n . We say that Γ is **nilpotent** if $\Gamma_n = \{e\}$ for some $n \in \mathbb{N}$. Note that abelian groups are nilpotent and nilpotent groups are solvable. Note also that the group from Example 12.1.1 is not nilpotent. (Indeed, it can be verified that $\Gamma_2 = \Gamma_1$ in that case.)

Example 12.1.2. Let $H := \left\{ \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} : a, b, c \in \mathbb{R} \right\}$. H is easily seen

to be a group under multiplication, the so-called **Heisenberg group**. It is clear that H is infinite and nonabelian. To see that H is nilpotent, we ask the

reader to verify that the map $\begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} \mapsto (a, b) : H \rightarrow \mathbb{R}^2$ is a surjective

group homomorphism with kernel $Z(H)$. It follows that $H/Z(H) \cong \mathbb{R}^2$ is abelian, whence H is nilpotent.

Notation 12.1.3. For a group Γ with subgroup Δ , we write $\Delta \leq_f \Gamma$ if Δ is a finite index subgroup of Γ .

If P is a property of groups, we say that Γ is **virtually** P if some finite index subgroup of Γ has property P . Note that virtually finite is the same as finite.

Definition 12.1.4. A group Γ is called **indicible** if there is a surjective homomorphism $\Gamma \rightarrow \mathbb{Z}$. If Γ is indicible, a **kernel** of Γ is the kernel of some homomorphism from Γ onto \mathbb{Z} .

The following fact will play a role later on:

Lemma 12.1.5. *Suppose that Γ is an infinite, solvable, finitely generated group. Then Γ is virtually indicable.*

Proof. Suppose first that Γ' has infinite index in Γ . Then since Γ/Γ' is an infinite finitely generated abelian group, by the Fundamental Theorem of Finitely Generated Abelian Groups, Γ/Γ' has a direct summand that is isomorphic to \mathbb{Z} , whence there is a surjective homomorphism from Γ/Γ' onto \mathbb{Z} , and thus from Γ itself.

If Γ' has finite index in Γ but $\Gamma^{(2)}$ has infinite index in Γ' , then the same argument shows that Γ' is indicable. Eventually, some $\Gamma^{(n)}$ must have infinite index in $\Gamma^{(n-1)}$ (whence $\Gamma^{(n-1)}$ is indicable), for otherwise $\{e\}$ will have finite index in Γ , whence Γ is finite, a contradiction. \square

12.2. Growth rates of groups

In this section, Γ is a group generated by a finite set X which we assume, without loss of generality, is closed under inverses. We then can define the **word-length** function $|\cdot|_X : \Gamma \rightarrow \mathbb{N}$ by $|g|_X :=$ the length of the shortest word in X representing g . When there is no possible source of confusion, we simply write $|g|$ instead of $|g|_X$. The following properties of the word-length function are easy to check:

- $|g| = 0$ if and only if $g = e$;
- $|g| = |g^{-1}|$;
- $|gh| \leq |g| + |h|$.

We can use the word-length function to define a metric d_X on Γ by defining $d_X(g, h) := |g^{-1}h|$. Once again, we write d instead of d_X if no confusion arises. Note also that this metric on Γ is *left-invariant*: for all $a, g, h \in \Gamma$, we have $d(ag, ah) = d(g, h)$.

Exercise 12.2.1. Suppose that X and X' are finite generating sets for a group Γ . Show that the identity map is a bi-Lipschitz homeomorphism, that is, show that there is $K \in \mathbb{N}$ such that, for all $g, h \in \Gamma$, we have

$$\frac{1}{K}d_{X'}(g, h) \leq d_X(g, h) \leq Kd_{X'}(g, h).$$

For each $n \in \mathbb{N}$ and $g \in \Gamma$, let $B_{\Gamma, X}(g, n)$ denote the closed ball of radius n with center g with respect to the metric d_X . We will write $B_{\Gamma}(g, n)$ if X is clear from context. It is fairly easy to see that all closed balls of radius n have the same *finite* size, denoted $G_X(n)$, or simply $G(n)$. The associated function $G_X : \mathbb{N} \rightarrow \mathbb{N}$ is called the **growth function** of the group Γ with respect to the generating set X . In this chapter, we will be concerned with

the rate of growth of the growth function G_X . Some examples might make things clearer.

Example 12.2.2. Let $\Gamma = \mathbb{Z}$ with generating set $\{\pm 1\}$. It is relatively clear that $G(n) = 2n + 1$ in this case.

Exercise 12.2.3. Let $\Gamma = \mathbb{Z}^2$ with generating set $\{(\pm 1, 0), (0, \pm 1)\}$. Show that $G(n) = 2n^2 + 2n + 1$.

The previous two examples are examples of groups with polynomial growth:

Definition 12.2.4. Γ has **growth degree** $\leq d$ if there is $c > 0$ and $d \in \mathbb{N}$ such that $G(n) \leq cn^d$ for all $n \geq 1$. If Γ has growth degree $\leq d$ for some d , then we say that Γ has **polynomial growth**.

The above definition is a little awkward in that the adjectives “growth degree $\leq d$ ” and “polynomial growth” apply to the group itself even though the definition involves the growth function with respect to a particular generating set. This is not an issue:

Exercise 12.2.5. Suppose that X and X' are finite generating sets for a group Γ . Show that there is $K \in \mathbb{N}$ such that $G_X(n) \leq G_{X'}(Kn)$ for all $n \in \mathbb{N}$. Conclude that the notion growth degree $\leq d$ (and hence the notion polynomial growth) is independent of the choice of generating set.

Exercise 12.2.6. Show that any finitely generated abelian group has polynomial growth.

The previous exercise can be significantly generalized:

Theorem 12.2.7 (Wolf). *Any finitely generated nilpotent group has polynomial growth.*

Proof of a special case. We only treat the case that Γ is not abelian but $\Gamma_2 = [\Gamma, \Gamma_1] = \{e\}$. Fix a finite generating set $\{g_1, \dots, g_m\}$ for Γ . Suppose that we have an element $g \in B_\Gamma(e, n)$. If we have adjacent occurrences of g_i and g_j in g where $j < i$, then we replace that with $g_j g_i [g_i^{-1}, g_j^{-1}]$. If we continue to switch like that, we end up rewriting $g = g_1^{a_1} \cdots g_m^{a_m} h$, where $h \in \Gamma_1$. A conservative upper bound for the number of “prefixes” is n^m . How many possible h 's are there? A conservative upper bound for the number of times switched is n^2 , so h is a product of at most n^2 commutators of generators and their inverses. Let S be a finite generating set for Γ_1 and let p be the maximal length of a commutator of generators of Γ in terms of S . Then $h \in B_{\Gamma_1}(e, pn^2)$. Since Γ_1 is abelian (by hypothesis), Γ_1 has polynomial growth, so there are c, d such that $|B_{\Gamma_1}(e, pn^2)| \leq c(pn^2)^d = cp^d n^{2d}$. Thus, there are at most $n^m \cdot cp^d n^{2d} = (cp^d)n^{m+2d}$ possible such g 's. It follows that Γ has growth degree at most $m + 2d$. \square

We now consider an example at the opposite extreme:

Example 12.2.8. Let $\Gamma = \mathbb{F}_2$, the free group on the two generators a and b . Note then that $\{a, a^{-1}, b, b^{-1}\}$ is a symmetric generating set. If $n \geq 1$, $|g| = n + 1$ and g starts with an a , then $a^{-1}g$ is a word with length n that starts with a , b , or b^{-1} . It follows that $G(n + 1) = 4 \cdot (\frac{3}{4}G(n)) = 3G(n)$, whence an inductive argument shows that $G(n) = 2 \cdot 3^n - 1$ for all n .

The previous example motivates:

Definition 12.2.9. Γ has **exponential growth** if there is $c > 1$ such that $G(n) \geq c^n$ for all n .

Exercise 12.2.10. Show that having exponential growth is independent of the generating set.

Exercise 12.2.11. Suppose that Γ is a finitely generated group and that Δ is a subgroup of Γ .

- (1) If Δ is finitely generated and X is such that $X \cap \Delta$ generates Δ , prove that $G_{X \cap \Delta} \leq G_X$. Conclude that Δ is of polynomial growth if Γ is of polynomial growth and Γ is of exponential growth if Δ is of exponential growth.
- (2) Suppose that Δ has finite index in Γ . Let X be a finite generating set for Γ obtained from taking a finite generating set Y for Δ and adjoining a (finite) set of coset representatives for Γ/Δ . Prove that there is $m \in \mathbb{N}$ such that $G_X(n) \leq |X| \cdot G_Y(mn)$ for all n . Conclude that Γ has polynomial (resp., exponential) growth if and only if Δ does.

By the second part of the previous exercise, the notions “virtually of polynomial growth” and “polynomial growth” are the same, as are “virtually of exponential growth” and “exponential growth.”

Remark 12.2.12. Although it will not play a role in this chapter, there are groups of so-called **intermediate growth**, that is, their growth functions grow faster than all polynomials but slower than exponentials. This was a major achievement due to Grigorchuk [70].

It will prove useful to consider G as a function $G : \mathbb{R}^{>0} \rightarrow \mathbb{N}$ (even though $G(r) = G(n)$ for $n \leq r < n + 1$). The following fact will be used in Section 12.6.

Lemma 12.2.13. $\lim_{r \rightarrow \infty} G(r)^{\frac{1}{r}}$ exists.

Proof. It is enough to show that, for any increasing sequence $(r_n)_{n \in \mathbb{N}}$ with $\lim_{n \rightarrow \infty} r_n = \infty$, we have $\lim_{n \rightarrow \infty} G(r_n)^{\frac{1}{r_n}}$ exists.

Fix $m, n \in \mathbb{N}$ and set $k := \lfloor \frac{r_m}{r_n} \rfloor + 1$. Then

$$G(r_m) \leq G(kr_n) \leq G(r_n)^k \leq G(r_n)^{\frac{r_m}{r_n} + 1}.$$

It follows that $\limsup_{m \rightarrow \infty} G(r_m)^{\frac{1}{r_m}} \leq G(r_n)^{\frac{1}{r_n}}$. Since n is arbitrary, we have that

$$\limsup_{m \rightarrow \infty} G(r_m)^{\frac{1}{r_m}} \leq \liminf_{n \rightarrow \infty} G(r_n)^{\frac{1}{r_n}},$$

yielding the desired result. \square

Corollary 12.2.14. Γ has exponential growth if and only if $\lim_{r \rightarrow \infty} G(r)^{\frac{1}{r}} > 1$.

Proof. If Γ has exponential growth, then there is $c > 1$ such that $G(n) \geq c^n$ for all n , whence $\lim_{n \rightarrow \infty} G(n)^{\frac{1}{n}} \geq c > 1$. Conversely, suppose that $c := \lim_{r \rightarrow \infty} G(r)^{\frac{1}{r}} > 1$. By the proof of the previous lemma, for any fixed n , we have $G(n)^{\frac{1}{n}} \geq c$, so $G(n) \geq c^n$, as desired. \square

12.3. Gromov's theorem on polynomial growth

In the last section, we saw that finitely generated nilpotent groups have polynomial growth. It follows from Exercise 12.2.11 that finitely generated virtually nilpotent groups have polynomial growth. Under the assumption that the group is solvable, the converse in fact holds, even under the more general assumption that the group does not have exponential growth:

Fact 12.3.1 (Milnor-Wolf [136]). If Γ is a finitely generated solvable group, then either Γ has exponential growth or else Γ is virtually nilpotent.

The above fact is really a combination of two results: Milnor proved that all finitely generated solvable groups are either **polycyclic** or have exponential growth while Wolf proved that all polycyclic groups are either virtually nilpotent or have exponential growth. (Polycyclic groups can be defined to be solvable groups for which every subgroup is finitely generated; while an efficient definition, it has the unfortunate downside that it does not explain the terminology.)

To summarize: amongst the finitely generated solvable groups, the groups with polynomial growth are precisely the virtually nilpotent groups. However, assuming that the group is solvable is a somewhat severe assumption. A remarkable theorem of Gromov is that the previous equivalence holds in general, without any solvability assumption:

Theorem 12.3.2 (Gromov's theorem on polynomial growth [71]). *If Γ is a finitely generated group of polynomial growth, then Γ is virtually nilpotent.*

What is striking about this theorem is that the “geometric” property of the group, namely the growth rate of the balls around the identity, capture precise algebraic information, namely virtual nilpotence.

In the rest of this section, we outline the proof of Gromov's theorem. Our proof sketch will be woefully incomplete, but we will make sure to emphasize the portions of the proof that will ultimately rely heavily on a certain metric ultraproduct construction, the so-called *asymptotic cone* construction; the rest of the chapter will be devoted to the study of asymptotic cones and the verification of the properties needed of it in the forthcoming proof sketch of Gromov's theorem.

From now on, we assume that Γ is a finitely generated group. A hint as how to proceed with the proof of Gromov's theorem is given by the following algebraic lemma, whose proof we omit and can be found in [181].

Lemma 12.3.3. *Suppose that Γ is an indicable group with kernel K . Suppose further that Γ does not have exponential growth. Then:*

- (1) K is finitely generated.
- (2) If Γ has growth degree $\leq d + 1$, then K has growth degree $\leq d$.
- (3) If K is virtually solvable, then so is Γ .

Thus, in order to prove Gromov's Theorem by induction on the growth degree of the group, we need a way to produce surjective homomorphisms onto \mathbb{Z} . Another source of inspiration is the following deep theorem. Recall that $\mathrm{GL}_n(\mathbb{C})$ is the group of invertible $n \times n$ matrices over the complex numbers.

Fact 12.3.4 (Tits alternative). Suppose that Γ is a finitely generated subgroup of $\mathrm{GL}_n(\mathbb{C})$. Then either Γ contains an isomorphic copy of the free group \mathbb{F}_2 or else Γ is virtually solvable.

Thus, for finitely generated subgroups of $\mathrm{GL}_n(\mathbb{C})$, polynomial growth implies virtual solvability (and hence virtual nilpotency by the Milnor-Wolf theorem). Thus, it would be great if one could get our group Γ (or some finite index subgroup of Γ) to be embedded in some $\mathrm{GL}_n(\mathbb{C})$. Now, $\mathrm{GL}_n(\mathbb{C})$ is the group of symmetries of a continuous object (namely \mathbb{C}^n), whilst Γ is the group of symmetries of a discrete object, namely Γ itself, equipped with its word metric. (Γ acts on itself by left multiplication and this action is by isometries.) Gromov's idea was that the metric space Γ , when viewed from “far away,” actually looks continuous and that the action of Γ on this “zoomed out” version of Γ is still by isometries. Provided that the zoomed out version of Γ has some nice properties, its group of isometries will be a so-called *Lie group*.

To be more precise, suppose that Y is a metric space and let $\text{Isom}(Y)$ denote the group of isometries of Y . $\text{Isom}(Y)$ has a natural topology on it, where a typical subbasic neighborhood of the identity function id_Y is a set of the form

$$U_{k,\epsilon} := \{f \in \text{Isom}(Y) : d(f(a), a) < \epsilon \text{ for all } a \in Y \text{ with } d(a, e) \leq k\}.$$

In the above, $e \in Y$ is simply some distinguished point fixed in advance. With this topology, $\text{Isom}(Y)$ is in fact a topological group.

The following important fact is a consequence of some deep work on Hilbert's fifth problem (see, for example, the book [137]).

Fact 12.3.5. Suppose that Y satisfies the following properties:

- (1) Y is *homogeneous*: given any $x, y \in Y$, there is $f \in \text{Isom}(Y)$ such that $f(x) = y$.
- (2) Y is connected and locally connected.
- (3) Y is proper.
- (4) Y is "finite-dimensional."

Then $\text{Isom}(Y)$ is a Lie group.

Even though we use quotation marks in item (4), the term finite-dimensional has a precise meaning. Since this definition would take us too far afield, we will be intentionally vague here.

What is a **Lie group**? Roughly speaking, it is a group that you can do calculus on. More precisely, a Lie group is a group that is also a *smooth manifold* such that the group operations are smooth. It is not entirely crucial that the reader understand the precise definition of a Lie group, for the following two facts about Lie groups will be all that we will need to know. We include very brief justifications for these facts mainly so that the reader can consult the literature for the missing details.

Lie group Fact #1. Suppose that G is a Lie group. Then there is $L \leq_f G$ and $n > 0$ such that $L/Z(L)$ embeds into $\text{GL}_n(\mathbb{C})$.

Brief justification. One associates to any Lie group G its **Lie algebra** \mathfrak{g} , which, amongst other things, is a finite-dimensional vector space. One has the **adjoint representation** $\text{Ad} : G \rightarrow \text{Aut}(\mathfrak{g})$, which is an injective group homomorphism. If G is connected, then the kernel of Ad is $Z(G)$. In general, the connected component G_e of the identity in G is a finite-index subgroup that is a connected Lie group, whence $G_e/Z(G_e)$ embeds into $\text{Aut}(\mathfrak{g}) \cong \text{GL}_n(\mathbb{C})$, where n is the dimension of \mathfrak{g} .

Lie group Fact #2. Suppose that G is a Lie group. Then for any $n > 0$, there is a neighborhood U of the identity in G such that all elements of $U \setminus \{e\}$ have order greater than n .

Brief justification. Lie groups have the **no small subgroups property**, meaning that there is a neighborhood V of e containing no subgroup of G other than $\{e\}$. (In fact, for a locally compact group, having the no small subgroups property is equivalent to being a Lie group; this is part of the deep work associated to the solution of Hilbert's fifth problem referred to above.) Now given $n > 0$, there is a neighborhood U of e such that the product of any n elements of U lies in V . (This is a basic property of topological groups.) It follows that any element of $U \setminus \{e\}$ has order greater than n .

Returning to Fact 12.3.5, in the next section, we will show how to construct a metric space $\text{Cone}(\Gamma)$ from Γ such that $\text{Cone}(\Gamma)$ is always homogeneous, connected, and locally connected. ($\text{Cone}(\Gamma)$ will be the zoomed out version of Γ referred to above.) In order to build $\text{Cone}(\Gamma)$, one will have to fix a positive infinite element R of \mathbb{R}^* (the level of "zoom") and so we actually write $\text{Cone}(\Gamma; R)$. In Section 12.6, we will also show the following:

Theorem 12.3.6. *If Γ has polynomial growth, then there is a positive infinite $R \in \mathbb{R}^*$ such that $\text{Cone}(\Gamma; R)$ is proper and finite-dimensional.*

Thus, the following theorem will be applicable.

Theorem 12.3.7. *Suppose that $\text{Cone}(\Gamma; R)$ is proper and finite-dimensional. Then Γ is virtually indicable.*

Before we outline the proof of Theorem 12.3.7, let us see how to readily deduce Gromov's theorem from the results already announced in this section.

Proof of Theorem 12.3.2. We prove, by induction on d , that if Γ has growth degree $\leq d$, then Γ is virtually nilpotent. When $d = 0$, we have that Γ is finite, so virtually nilpotent (as $\{e\}$ is of finite index in Γ and nilpotent).

Inductively assume that the theorem is true for d and that Γ has growth degree $\leq d + 1$. By Theorems 12.3.6 and 12.3.7, there is a finite index indicable subgroup Δ of Γ . Let K be a kernel of Δ . By Exercise 12.2.11 and Lemma 12.3.3(2), K has growth degree $\leq d$, whence, by induction, K is virtually nilpotent, and hence virtually solvable. By Lemma 12.3.3(3), we have that Δ is virtually solvable. By Milnor-Wolf, Δ is virtually nilpotent, whence Γ is also virtually nilpotent. \square

We now return to the proof of Theorem 12.3.7. Note that the theorem is easy to prove if Γ is virtually abelian, so we suppose from now on that

this is not the case. Suppose that $Y := \text{Cone}(\Gamma; R)$ is proper and finite-dimensional. By Fact 12.3.5 and Lie group Fact #1, there is $n > 0$ and $L \leq_f \text{Isom}(Y)$ such that $L/Z(L)$ embeds into $\text{GL}_n(\mathbb{C})$. As alluded to earlier, the elements of Γ can be construed as isometries of Y ; more precisely, there is a group homomorphism $l : \Gamma \rightarrow \text{Isom}(Y)$. The motivation behind bringing in $\text{Isom}(Y)$ was to somehow get a finite index subgroup of Γ which has as a homomorphic image an infinite subgroup of some $\text{GL}_n(\mathbb{C})$. A first step is the following:

Claim. *There is $\Delta \leq_f \Gamma$ such that, for any n , there is a homomorphism $\Delta \rightarrow L$ with image of size at least n .*

Proof of Claim. The claim is easy to prove when $l(\Gamma)$ is infinite, for then $\Delta := l^{-1}(L)$ has finite index in Γ and $l(\Delta)$ is an infinite subset of L . We thus assume that $l(\Gamma)$ is finite. Set $\tilde{\Gamma} := \ker(l)$, a finite-index subgroup of Γ . By Lie Group Fact #2, for each $n > 0$, there is a neighborhood U of id_Y such that U contains no elements of order $< n$. As a consequence of Theorem 12.5.2 below, there is a homomorphism $\tilde{\Gamma} \rightarrow \text{Isom}(Y)$ whose image intersects $U \setminus \{\text{id}_Y\}$. Consequently, for any $n > 0$, there is a homomorphism $\tilde{\Gamma} \rightarrow \text{Isom}(Y)$ whose image has size $\geq n$. A counting argument gives $\Delta \leq_f \Gamma'$ such that, for any $n > 0$, there is a homomorphism $\Delta \rightarrow L$ whose image has size $\geq n$, proving the claim.

Unfortunately, it is $L/Z(L)$, and not L itself, which embeds into $\text{GL}_n(\mathbb{C})$. Thus, if we consider what happens when we compose the homomorphisms $\Delta \rightarrow L$ given by the claim with the quotient map $L \rightarrow L/Z(L)$, three cases can occur:

- (i) there is $q \in \mathbb{N}$ such that all homomorphisms $\Delta \rightarrow L/Z(L)$ have image of size at most q ;
- (ii) all homomorphisms $\Delta \rightarrow L/Z(L)$ have finite image, but of unbounded size;
- (iii) there is a homomorphism $\Delta \rightarrow L/Z(L)$ with infinite image $\overline{\Delta}$.

Case (iii) is the one that is easiest to deal with given the earlier discussion of the Tits Alternative. Indeed, $\overline{\Delta}$, being isomorphic to an infinite, finitely generated subgroup of $\text{GL}_n(\mathbb{C})$, either contains a copy of \mathbb{F}_2 as a subgroup, or else is virtually solvable. The former alternative would imply that Γ has exponential growth by Exercise 12.2.11. In Section 12.6, we will prove the following weak converse to Theorem 12.3.6:

Theorem 12.3.8. *If $\text{Cone}(\Gamma; R)$ is proper for some $R \in \mathbb{R}^*$, then Γ does not have exponential growth.*

Recalling the hypotheses of Theorem 12.3.7, it must be the case that $\overline{\Delta}$ is virtually solvable. The proof of the theorem now follows from Lemma 12.1.5.

We now treat Case (i). The assumption of Case (i) implies that $[\Delta : \ker(h)] \leq q$ for any homomorphism $h : \Delta \rightarrow L/Z(L)$. If we set $\Delta_1 := \bigcap \{\ker(h) : h : \Delta \rightarrow L/Z(L)\}$, it follows that $\Delta_1 \leq_f \Delta$ and for any $n > 0$ there is a homomorphism $\Delta_1 \rightarrow Z(L)$ of cardinality $\geq n$. (The fact that $\Delta_1 \leq_f \Delta$ uses the elementary group theory fact that a finitely generated group has only finitely many subgroups of a given finite index, whence Δ_1 is actually a finite intersection.). Since $Z(L)$ is abelian and any such homomorphism must factor through Δ_1^{ab} , it follows that Δ_1^{ab} is infinite, and thus indicable, whence Δ_1 itself is indicable.

It remains to prove Case (ii). The proof of Case (ii) is nearly identical to the proof of Case (i) once one knows the following theorem of Jordan: there is $q \in \mathbb{N}$ such that every finite subgroup of $\text{GL}_n(\mathbb{C})$ has an abelian subgroup of index at most q .

This concludes the proof sketch of Theorem 12.3.7 and consequently the sketch of the proof of Gromov's Theorem 12.3.2.

12.4. Definition of asymptotic cones

We start with the general definition of an asymptotic cone of a metric space, although we specialize quickly to the case of groups.

Definition 12.4.1. Suppose that (X, d) is a metric space, \mathcal{U} an ultrafilter on \mathbb{N} , $o \in X^{\mathbb{N}}$ a sequence from X , and $r \in (\mathbb{R}^{>0})^{\mathbb{N}}$ is such that $\lim_{\mathcal{U}} r(n) = \infty$. Then the **asymptotic cone** with respect to all of this data, denoted $\text{Cone}(X; \mathcal{U}, o, r)$, is the metric ultraproduct

$$\text{Cone}(X; \mathcal{U}, o, r) := \prod_{\mathcal{U}} \left(X, \frac{d}{r(n)}, o(n) \right).$$

In other words, an asymptotic cone of a metric space is simply a metric ultraproduct of a family of metric spaces obtained by “scaling down” the metric of the original space.

The notation above is admittedly cumbersome, but we will shortly restrict our attention to the case of groups, where we will omit some of the data from the notation. We start with two results which will explain why this aforementioned omission will be harmless.

Lemma 12.4.2. *Suppose that (X, d) is a homogeneous metric space and fix $o, o' \in X^{\mathbb{N}}$. Then for any $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$ and any sequence $r \in (\mathbb{R}^{>0})^{\mathbb{N}}$ with $\lim_{\mathcal{U}} r(n) = \infty$, the asymptotic cones $\text{Cone}(X; \mathcal{U}, o, r)$ and $\text{Cone}(X; \mathcal{U}, o', r)$ are isometric.*

Proof. For each $n \in \mathbb{N}$, let $\sigma_n \in \text{Isom}(X)$ be such that $\sigma(o_n) = o'(n)$ and define $\sigma : X^{\mathbb{N}} \rightarrow X^{\mathbb{N}}$ by setting $\sigma(x)(n) := \sigma_n(x(n))$. Define

$$\sigma' : \text{Cone}(X; \mathcal{U}, o, r) \rightarrow \text{Cone}(X; \mathcal{U}, o', r)$$

by setting $\sigma'([x]_{\mathcal{U}}) := [\sigma(x)]_{\mathcal{U}}$. We leave it to the reader to verify that this map is a well-defined isometry. \square

Lemma 12.4.3. *Suppose that $f : (X, d) \rightarrow (Y, d)$ is a K -bi-Lipschitz homeomorphism. Then for any $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$, any $o \in X^{\mathbb{N}}$, and any $r \in (\mathbb{R}^{>0})^{\mathbb{N}}$ with $\lim_{\mathcal{U}} r(n) = \infty$, the ultraproduct map $f^{\mathcal{U}} : \text{Cone}(X; \mathcal{U}, o, r) \rightarrow \text{Cone}(Y; \mathcal{U}, f(o), r)$ is a K -bi-Lipschitz homeomorphism.*

Exercise 12.4.4. Prove the previous lemma.

Throughout the rest of this section, we fix a finitely generated group Γ .

Definition 12.4.5. For any ultrafilter \mathcal{U} on \mathbb{N} and any $r \in (\mathbb{R}^{>0})^{\mathbb{N}}$ such that $\lim_{\mathcal{U}} r(n) = \infty$, we set

$$\text{Cone}(\Gamma; \mathcal{U}, r) := \text{Cone}(\Gamma; \mathcal{U}, e, r),$$

where Γ is treated as a metric space equipped with the word metric associated to any finite generating set.

By Lemma 12.4.3, the choice of word metric on d does not change the bi-Lipschitz homeomorphism type of the cone, and for our purposes, that is sufficient (whence we feel comfortable removing mention of the particular metric involved). It turns out that the choice of \mathcal{U} and the choice of r do affect the homeomorphism type of the asymptotic cone; see Section 12.8 for more on this. However, for the purposes of proving Gromov's theorem, the choice of ultrafilter is not important.

In what follows, it will be useful to consider the nonstandard hull presentation of asymptotic cones as discussed in Section 11.2. We let (M, d, e) be the classical ultraproduct of the family $(\Gamma, \frac{d}{r(n)}, e)$, which we view as an internal metric space in the nonstandard universe, with $d([\gamma]_{\mathcal{U}}, [\eta]_{\mathcal{U}}) := [n \mapsto \frac{d(\gamma(n), \eta(n))}{r(n)}]_{\mathcal{U}} \in \mathbb{R}^*$. In this case, setting $R := [r]_{\mathcal{U}} \in \mathbb{R}^*$, we have that R is a positive infinite element of \mathbb{R}^* and that $M_{\text{fin}} := \{\gamma \in \Gamma^* : \frac{|\gamma|}{R} \in \mathbb{R}_{\text{fin}}\}$.

Consequently, in what follows, we will often simply consider a nonstandard extension Γ^* of our group Γ and a positive infinite element R of \mathbb{R}^* . We have the set

$$\Gamma_R := \left\{ \gamma \in \Gamma^* : \frac{d(\gamma, e)}{R} \in \mathbb{R}_{\text{fin}} \right\}$$

and a pseudo-metric d on Γ_R given by $d(\gamma, \eta) := \text{st}(\frac{d(\gamma, \eta)}{R})$. If we set

$$\mu = \mu_R := \{\gamma \in \Gamma_R : d(\gamma, e) = 0\},$$

we then have that μ_R is a subgroup of Γ_R and the set of left cosets Γ_R/μ_R is the metric space obtained from Γ_R by identifying elements that are of distance 0. We refer to this metric space as $\text{Cone}(\Gamma; R)$ and also call it an asymptotic cone. It agrees with the definition given above in the case that our nonstandard extension arose as an ultrapower and $R = [r]_{\mathcal{U}}$. Technically, $\text{Cone}(\Gamma; R)$ also depends on the choice of nonstandard extension (that is, the choice of the ultrafilter), but let us not further complicate the notation.

The advantage to the nonstandard hull presentation of asymptotic cones is that often we can reason about the cone using the algebra of the group Γ^* and the methods of nonstandard analysis such as transfer and overspill. We will see several examples of this in the next few sections.

We now consider some examples of asymptotic cones:

Example 12.4.6. Suppose that $\Gamma = \mathbb{Z}$ and let d be the word metric associated with the standard generating set $\{1\}$ for Γ . Fix any positive infinite $R \in \mathbb{R}^*$. We then have that $\Gamma_R := \{k \in \mathbb{Z}^* : \frac{k}{R} \in \mathbb{R}_{\text{fin}}\}$ and $\mu_R := \{k \in \mathbb{Z}^* : \frac{k}{R} \approx 0\}$. Note then that the map $k + \mu_R \mapsto \text{st}(\frac{k}{R}) : \text{Cone}(\mathbb{Z}; R) \rightarrow \mathbb{R}$ is an isometry. In this way, when zoomed out, the discrete space \mathbb{Z} turned into the continuous space \mathbb{R} . Note in this case that the asymptotic cone did not depend on the choice of R .

Exercise 12.4.7. For $\Gamma = \mathbb{Z}^2$ and the word metric d associated with the standard generating set $\{(1, 0), (0, 1)\}$, show that, for any positive infinite R , $\text{Cone}(\Gamma; R)$ is isometric to \mathbb{R}^2 equipped with the Paris metric, that is, equipped with the metric

$$d((x_1, x_2), (y_1, y_2)) := |x_1 - y_1| + |x_2 - y_2|.$$

The following exercise will be useful in Section 12.6:

Exercise 12.4.8. Fix a positive infinite $R \in \mathbb{R}^*$ and set $Y := \text{Cone}(\Gamma; R)$. For each $n \in \mathbb{N}$, prove that the map

$$(\gamma_1, \dots, \gamma_n) \mapsto \gamma_1 \cdots \gamma_n : B_{\Gamma^*}(e, R) \times \cdots \times B_{\Gamma^*}(e, R) \rightarrow B_{\Gamma^*}(e, nR)$$

induces a continuous surjection $B_Y(e\mu, 1) \times \cdots \times B_Y(e\mu, 1) \rightarrow B_Y(e\mu, n)$.

12.5. General properties of asymptotic cones

In this section, we fix a finitely generated group Γ with word metric d associated to some finite generating set. We also fix a positive infinite $R \in \mathbb{R}^*$ and we set $Y := \text{Cone}(\Gamma; R)$.

Theorem 12.5.1. *Y has the following properties:*

- (1) *Y is homogeneous.*

- (2) Y is a geodesic metric space: given $x, y \in Y$ with $r := d(x, y)$, there is an isometric embedding $g : [0, r] \rightarrow Y$ with $g(0) = x$ and $g(r) = y$.
- (3) Y is complete.

Proof. For the proof of (1), notice that, for all $\gamma, \eta_1, \eta_2 \in \Gamma_R$, we have

$$d(\gamma\eta_1\mu, \gamma\eta_2\mu) = d(\eta_1\mu, \eta_2\mu),$$

whence the map $\eta\mu \mapsto \gamma\eta\mu$ is an isometry of Y , denoted by $l(\gamma)$. It remains to note that, given $\eta_1\mu, \eta_2\mu \in Y$, we have $l(\eta_2\eta_1^{-1})(\eta_1\mu) = \eta_2\mu$.

We now prove (2). Using (1), we may assume that $x = e\mu$ and $y = \gamma\mu$. Set $m := |\gamma|$ so $r = d(x, y) = \text{st}(\frac{m}{R})$. Since $\frac{[rR]}{R} \approx r$, by replacing γ with an element of word length $[rR]$, we may assume that $m = [rR]$. Write $\gamma = \gamma_1 \cdots \gamma_{[rR]}$. We claim that the desired isometric embedding is given by $g : [0, r] \rightarrow Y$, $g(t) := \gamma_1 \cdots \gamma_{[tR]}\mu$. Indeed, given $0 \leq s < t \leq r$, we have that

$$d(g(s), g(t)) = \text{st} \left(\frac{[tR] - [sR]}{R} \right) = \text{st} \left(\frac{(t-s)R}{R} \right) = |s - t|,$$

as desired.

Finally, completeness follows from Theorem 11.3.1. \square

Note that geodesic metric spaces are connected and locally connected. We have thus verified that $\text{Cone}(\Gamma; R)$ has properties (1) and (2) in Fact 12.3.5 for any choice of positive infinite $R \in \mathbb{R}^*$.

As introduced in the proof of Theorem 12.5.1(1), there is a function $l : \Gamma_R \rightarrow \text{Isom}(Y)$ defined by $l(\gamma)(\eta\mu) := \gamma\eta\mu$. It is clear that l is a group homomorphism. As in our proof of Theorem 12.3.7, we set $\tilde{\Gamma} := \ker(l \upharpoonright \Gamma)$.

In the rest of this subsection, we work toward establishing the following theorem, which was used in the proof of Theorem 12.3.7 above:

Theorem 12.5.2. *Suppose that $\tilde{\Gamma}$ has finite index in Γ but $Z(\tilde{\Gamma})$ does not have finite index in Γ (e.g., when Γ has no finite index abelian subgroup). Then for any neighborhood U of id_Y in $\text{Isom}(Y)$, there is a group homomorphism $\tilde{\Gamma} \rightarrow \text{Isom}(Y)$ whose image intersects $U \setminus \{\text{id}_Y\}$.*

We first need an elementary group theory lemma:

Lemma 12.5.3. *Suppose that Δ is a group with finite generating set S . Suppose that $Z(\Delta)$ is not of finite index in Δ . Then $\{\gamma^{-1}s\gamma : \gamma \in \Delta, s \in S\}$ is infinite.*

Proof. We argue by contrapositive: suppose, for each $s \in S$, that $\{\gamma^{-1}s\gamma : \gamma \in \Delta\}$ is finite; since the latter set is in bijection with the set of cosets

$\Delta/C(s)$, where $C(s)$ is the centralizer of s in Δ , we see that $C(s) \leq_f \Delta$. Since this is true for each $s \in S$, we have $Z(\Delta) = \bigcap_{s \in S} C(s)$ has finite index in Δ . \square

In the proof below, we will need the following notion. Given $\gamma \in \Gamma^*$ and $r \in (\mathbb{R}^*)^{>0}$, we define $\delta(\gamma, r) := \max\{d(\gamma a, a) : |a| \leq r\}$.

Exercise 12.5.4. For $\gamma, \eta \in \Gamma^*$ and $r \in (\mathbb{R}^*)^{>0}$, prove that $\delta(\eta^{-1}\gamma\eta, r) \leq \delta(\gamma, r) + 2|\eta|$.

Proof of Theorem 12.5.2. Let S be a finite generating set for $\tilde{\Gamma}$ closed under inverses. We will actually prove the following statement: for any neighborhood U of id_Y in $\text{Isom}(Y)$, there is $\beta \in \tilde{\Gamma}^*$ and $s \in S$ such that $\beta^{-1}\tilde{\Gamma}\beta \subseteq \Gamma_R$ and $l(\beta^{-1}s\beta) \in U \setminus \{\text{id}_Y\}$. The desired homomorphism $\tilde{\Gamma} \rightarrow \text{Isom}(Y)$ is simply obtained by mapping $\gamma \in \tilde{\Gamma}$ to $l(\beta^{-1}\gamma\beta)$.

Without loss of generality, we may assume that $U = U_{k,\epsilon}$ for some k and ϵ . By Lemma 12.5.3 and transfer, there is $\gamma \in \tilde{\Gamma}^*$ and $s \in S$ such that $|\gamma^{-1}s\gamma| > \epsilon R$. Write $\gamma = s_1 \cdots s_t$ with $t \in \mathbb{N}^*$. We consider the initial products $\gamma_i := s_1 \cdots s_i$ for $0 \leq i \leq t$, with the convention that $\gamma_0 = e$. Let $M_i := \max\{\delta(\gamma_i^{-1}s\gamma_i, kR) : s \in S\}$. Finally, set $C := \max\{|s| : s \in S\} \in \mathbb{N}$.

Claim. *The following inequalities hold:*

- (1) $M_0 < \epsilon R$;
- (2) $M_t > \epsilon R$;
- (3) $|M_{i+1} - M_i| \leq 2C$ for $0 \leq i \leq t - 1$.

Proof of Claim. To see (1), observe that if $s \in S$ and $|a| \leq kR$, then since $s \in \ker(l)$, we have that $\frac{d(sa, a)}{R} \approx 0$, whence $d(sa, a) < \epsilon R$. It follows that $\delta(s, kR) < \epsilon R$, whence (1) follows. (2) follows from the definition of $\gamma_t = \gamma$. Finally, (3) follows from Exercise 12.5.4.

Thus, by the Claim, there is $i \in \{0, 1, \dots, t\}$ such that $|M_i - \epsilon R| \leq 2C$. Our desired β is γ_i . To see that $\beta^{-1}\tilde{\Gamma}\beta \subseteq \Gamma_R$, observe that any element of $\tilde{\Gamma}$ is of the form $s_1 \cdots s_n$ with each $s_i \in S$, whence it suffices to observe that $\beta^{-1}s\beta \in \Gamma_R$ for each $s \in S$. However, we have that

$$|\beta^{-1}s\beta| = d(\beta^{-1}s\beta \cdot e, e) \leq \delta(\beta^{-1}s\beta, kR) \leq M_i \leq \epsilon R + 2C,$$

whence it follows that $\beta^{-1}s\beta \in \Gamma_R$.

Take $s \in S$ such that $\delta(\beta^{-1}s\beta, kR) = M_i$; we claim that $l(\beta^{-1}s\beta) \in U \setminus \{\text{id}_Y\}$. Note that this choice of s gives us that $|\delta(\beta^{-1}s\beta, kR) - \epsilon R| \leq 2C$. In particular, there is $a \in \Gamma^*$ with $|a| \leq kR$ such that $|d(\beta^{-1}s\beta a, a) - \epsilon R| \leq 2C$, whence it follows that $d(\beta^{-1}s\beta a \mu, a \mu) = \epsilon$, and thus $l(\beta^{-1}s\beta)$ is not

the identity on Y . Moreover, for arbitrary $b\mu \in Y$ with $d(b\mu, e\mu) \leq k$, we have, by replacing b with an element of word length at most kR , that $d(\beta^{-1}s\beta b\mu, b\mu) = \text{st}(\frac{d(\beta^{-1}s\beta b, b)}{R}) \leq \epsilon$, whence $l(\beta^{-1}s\beta) \in U$. \square

12.6. Growth functions and properness of the asymptotic cones

Once again, we fix a finitely generated group Γ with word metric d associated to some finite generating set. In this section, we establish the two facts concerning growth rates and properness of asymptotic cones made in Section 12.3, namely Theorems 12.3.6 and 12.3.8. The first follows from a lemma that necessitates some notation:

For $R \in \mathbb{R}^*$ and $i, d \in \mathbb{N}$, let $P_{i,d}(R)$ be the statement: if $t \in \mathbb{N}^*$ and $g_1, \dots, g_t \in B_{\Gamma^*}(e, \frac{R}{4})$ are such that $B_{\Gamma^*}(g_1, \frac{R}{i}), \dots, B_{\Gamma^*}(g_t, \frac{R}{i})$ are pairwise disjoint, then $t \leq i^{d+1}$.

Lemma 12.6.1. *Suppose that $c \in \mathbb{R}^{>0}$ and R_0 is a positive infinite element of \mathbb{R}^* are such that $G(R_0) \leq cR_0^d$. Then there is positive infinite $R \leq R_0$ such that, for every $i \in \mathbb{N}$, $P_{i,d}(R)$ holds.*

Before we prove Lemma 12.6.1, let us see how it accomplishes what we want.

Theorem 12.6.2. *If Γ has polynomial growth, then there is a positive infinite R such that, setting $Y := \text{Cone}(\Gamma; R)$, we have: for every (standard) $k \in \mathbb{N}$, $B_Y(e\mu, 1)$ is covered by at most $(4k)^{d+1}$ many closed balls of radius $\frac{2}{k}$.*

Proof. Since Γ is of polynomial growth, there are c and d such that $G(n) \leq cn^d$ for all $n \in \mathbb{N}$. Then, by overspill, there is a positive infinite R_0 such that $G(R_0) \leq cR_0^d$. Let R' be as guaranteed to exist in Lemma 12.6.1 and set $R := \frac{R'}{4}$. We claim that this R works. Indeed, fix $k \in \mathbb{N}$ and let $t \in \mathbb{N}^*$ be maximal such that there are $g_1, \dots, g_t \in B_{\Gamma^*}(e, R)$ with $B_{\Gamma^*}(g_i, \frac{R}{k})$ pairwise disjoint. Since $P_{4k,d}(R')$ holds, it must be that $t \leq (4k)^{d+1}$. It follows that $B_{\Gamma^*}(g_i, \frac{2R}{k})$, $i = 1, \dots, t$, cover $B_{\Gamma^*}(e, R)$: if $g \in B_{\Gamma^*}(e, R)$ does not belong to $B_{\Gamma^*}(g_i, \frac{2R}{k})$ for any $i = 1, \dots, t$, then $B_{\Gamma^*}(g, \frac{R}{k})$ is disjoint from $B_{\Gamma^*}(g_i, \frac{R}{k})$, contradicting the maximality of t . Thus, in Y , the balls $B_Y(g_i\mu, \frac{2}{k})$ cover $B_Y(e\mu, 1)$, as desired. \square

Remark 12.6.3. The previous proof works under the seemingly weaker hypothesis that Γ has **near polynomial growth** in the sense that there are c and d such that $G(n) \leq cn^d$ for infinitely many n (rather than for all n). Thus, groups with near polynomial growth are virtually nilpotent. A substantial improvement of this latter result was obtained by Hrushovski in

[86], where it was shown that Γ is virtually nilpotent provided that $G(n) \leq cn^d$ for a *single* sufficiently large n .

Proof of Theorem 12.3.6. Suppose that Γ is of polynomial growth. Let R be as in Theorem 12.6.2 and set $Y := \text{Cone}(\Gamma; R)$. We show that Y is proper and finite-dimensional.

To see that Y is proper, first note that, by Theorem 12.6.2, $B_Y(e\boldsymbol{\mu}, 1)$ is totally bounded. Since $B_Y(e\boldsymbol{\mu}, 1)$ is also complete, it follows that it is compact. By Exercise 12.4.8, we have that $B_Y(e\boldsymbol{\mu}, n)$ is compact for all $n \in \mathbb{N}$. It follows that all closed balls are compact, whence Y is proper.

To see that Y is finite-dimensional, we show that the **Hausdorff dimension** of Y , denoted $d_{\text{Haus}}(Y)$, is finite. (We have been intentionally vague about our notion of dimension, but suffice it to say that having finite Hausdorff dimension implies having finite dimension.) By definition, given $X \subseteq Y$ and $s \in \mathbb{R}^{>0}$, we define $d_{\text{Haus}}(X) \leq r$ if: for every $\epsilon > 0$, there are balls B_i of radius $\text{rad}(B_i)$ covering X such that $\sum_i \text{rad}(B_i)^r < \epsilon$. It is clear that this definition is local in the sense that if Y is covered by countably many sets with Hausdorff dimension $\leq r$, then Y itself has Hausdorff dimension $\leq r$. Thus, it suffices to show that $d_{\text{Haus}}(B)$ is finite, where $B := B_Y(e\boldsymbol{\mu}, 1)$. However, given any $k \in \mathbb{N}$, by the choice of R , we have that B is covered by at most $(4k)^{d+1}$ many balls of radius $\frac{2}{k}$; since $(4k)^{d+1}(\frac{2}{k})^{d+2} = \frac{4^{d+1}2^{d+2}}{k}$ goes to 0 as $k \rightarrow \infty$, we see that $d_{\text{Haus}}(B) \leq d + 2$, as desired. \square

Proof of Lemma 12.6.1. Since all balls in this proof are balls in Γ^* , we simplify notation by omitting the subscript Γ^* . Assume, toward a contradiction, that the Lemma is false. In particular, for each $R \in [\log(R_0), R_0]$, there is $i \in \mathbb{N}$ such that $P_{i,d}(R)$ fails. The function $R \mapsto$ the least $i \in \mathbb{N}^*$ for which $P_{i,d}(R)$ fails is an internal function that takes values in \mathbb{N} , whence it must actually take values in $[1, K]$ for some $K \in \mathbb{N}$.

Let i_1 be the minimal i for which $P_{i,d}(R_0)$ fails, so $4 \leq i_1 \leq K$, and there exist $g(1, j)$, $1 \leq j \leq t_1$ (where $t_1 := i_1^{d+1} + 1$) such that each $g(1, j) \in B(e, \frac{R_0}{4})$ and the $B(g(1, j), \frac{R_0}{i_1})$ are pairwise disjoint. If $\log R_0 \leq \frac{R_0}{i_1}$, then, setting i_2 to be the least i for which $P_{i,d}(\frac{R_0}{i_1})$ fails, there exist $g(2, j)$, $1 \leq j \leq t_2$ (where, once again, $t_2 := i_2^{d+1} + 1$) such that each $g(2, j) \in B(e, \frac{R_0}{4i_1})$ and the $B(g(2, j), \frac{R_0}{i_1 i_2})$ are pairwise disjoint. Keep going in this fashion (using *internal induction*) until you reach the point where $\frac{R_0}{(i_1 i_2 \cdots i_u)} < \log R_0 \leq \frac{R_0}{(i_1 i_2 \cdots i_{u-1})}$, having thus constructed, for $1 \leq l \leq u$ and $1 \leq j \leq t_l$ (with $t_l := i_l^{d+1} + 1$), elements $g(l, j) \in B(e, \frac{R_0}{4i_1 \cdots i_{l-1}})$ such that $B(g(l, j), \frac{R_0}{i_1 \cdots i_l})$ are pairwise disjoint.

We now show how to construct more than cR_0^d many elements of $B(e, R_0)$, contradicting the fact that $G(R_0) \leq cR_0^d$. Indeed, let T denote the set of tuples (s_1, \dots, s_u) with $1 \leq s_l \leq t_l$ for $l = 1, \dots, u$. For each $s = (s_1, \dots, s_u) \in T$, define $g_s := g(1, s_1) \cdots g(u, s_u) \in \Gamma^*$. Note that

$$|T| \geq \prod_{l=1}^u i_l^{d+1} > \left(\frac{R_0}{\log R_0} \right)^{d+1} > cR_0^d,$$

where the second inequality follows from the definition of u and the third inequality follows from the simple limit statement $\lim_{x \rightarrow \infty} \frac{x}{(\log x)^{d+1}} = \infty$, whence $\frac{R_0}{(\log R_0)^{d+1}}$ is infinite, and thus larger than c . Consequently, if we are able to show that the g_s , for $s \in T$, are distinct elements of $B(e, R_0)$, it would follow that $G(R_0) > cR_0^d$, yielding the desired contradiction.

We first show that each $g_s \in B(e, R_0)$. Indeed, we have

$$|g_s| \leq \sum_{l=1}^u |g(l, s_l)| \leq \sum_{l=1}^u \frac{R_0}{4i_1 \cdots i_{l-1}} \leq \sum_{l=1}^u \frac{R_0}{4^l} < R_0.$$

To complete the proof, suppose that $p, s \in T$ are distinct; we show that $g_p \neq g_s$. Toward a contradiction, suppose that $g_p = g_s$. There must then exist $v < u$ such that $p_v \neq s_v$ and yet $g(v, p_v) \cdots g(u, p_u) = g(v, s_v) \cdots g(u, s_u)$, whence we have

$$g(v, s_v)^{-1} g(v, p_v) = g(v+1, s_{v+1}) \cdots g(u, s_u) g(u, p_u)^{-1} \cdots g(v+1, p_{v+1})^{-1}.$$

Since $p_v \neq s_v$, we have that $g(v, s_v) \notin B(g(v, p_v), \frac{R_0}{4i_1 \cdots i_v})$, whence

$$|g(v, p_v)^{-1} g(v, s_v)| > \frac{R_0}{4i_1 \cdots i_v}.$$

However, by the previous display,

$$\begin{aligned} & |g(v, p_v)^{-1} g(v, s_v)| \\ & \leq 2 \sum_{l=v+1}^u \frac{R_0}{4i_1 \cdots i_{l-1}} \leq \frac{R_0}{2i_1 \cdots i_v} \left(1 + \sum_{l=v+2}^u \frac{1}{4^{l-v-1}} \right) < \frac{R_0}{i_1 \cdots i_v}. \end{aligned}$$

The absurd conclusion $\frac{R_0}{4i_1 \cdots i_v} < \frac{R_0}{i_1 \cdots i_v}$ leads to the desired contradiction. \square

We now turn to the proof of the other fact concerning growth rates and properness:

Proof of Theorem 12.3.8. Suppose that Γ has exponential growth. Fix positive infinite $R \in \mathbb{R}^*$; we must show that $\text{Cone}(\Gamma; R)$ is not proper. Toward that end, set $r := \lim_{s \rightarrow \infty} G(s)^{\frac{1}{s}}$, which exists by Lemma 12.2.13. It

follows that $G(R)^{\frac{1}{R}} \approx G(2R)^{\frac{1}{2R}} \approx r$. Thus,

$$\left(\frac{G(2R)}{G(R)}\right)^{\frac{1}{R}} = \left(\frac{(G(2R)^2)^{1/2R}}{G(R)^{\frac{1}{R}}}\right) \approx \frac{r^2}{r} = r,$$

whence it follows that $\frac{G(2R)}{G(R)}$ is infinite. In particular, this means that $B_{\Gamma^*}(e, 2R)$ cannot be covered by finitely many balls of radius R , whence, in $\text{Cone}(\Gamma; R)$, finitely many balls of radius 1 cannot cover the closed ball around $e\mu$ of radius 2 and thus $\text{Cone}(\Gamma; R)$ is not proper. \square

Remark 12.6.4. In many parts of the literature, one often finds the discussion of locally compact asymptotic cones rather than proper asymptotic cones. While proper metric spaces are always locally compact, the converse is in general not true. However, for geodesic spaces, the converse is true and is an easy exercise.

We have now established all of the facts needed during the proof sketch of Gromov's theorem.

12.7. Properness of asymptotic cones revisited

We showed that having polynomial growth implies having a proper asymptotic cone. In [181], van den Dries and Wilkie asked if, conversely, having a proper asymptotic cone implied polynomial growth. This was only recently answered in the positive through a combination of results by Hrushovski and Sapir, which we now explain.

Definition 12.7.1. Suppose that Γ is a group, F is a subset of Γ , and $k \in \mathbb{N}$. We say that F is a **k -approximate subgroup** of Γ if there is a finite set $X \subseteq \Gamma$ such that $|X| \leq k$ and for which $F \cdot F \subseteq X \cdot F$.

The following definition is not standard but is convenient:

Definition 12.7.2. Γ has the **approximate subgroup property** if there is $k \in \mathbb{N}$ and arbitrarily large balls $B(e, r)$ such that $B(e, r)$ is a k -approximate subgroup of Γ .

Exercise 12.7.3. Γ has the approximate subgroup property if and only if there is $k \in \mathbb{N}$ and positive infinite $R \in (\mathbb{R}^*)^{>0}$ such that $B_{\Gamma^*}(e, R)$ is a k -approximate subgroup of Γ^* .

The following result is due to Sapir [152].

Proposition 12.7.4. *If there is R such that $\text{Cone}(\Gamma; R)$ is proper, then Γ has the approximate subgroup property.*

Proof. Set $Y := \text{Cone}(\Gamma; R)$. Take $g_1, \dots, g_k \in \Gamma^*$, $k \in \mathbb{N}$, such that $d(g_i \mu, e \mu) \leq \frac{7}{8}$ for $i = 1, \dots, k$ and such that

$$B_Y(e \mu, 1) \subseteq \bigcup_{i=1}^k B_Y(g_i \mu, \frac{1}{4}). \quad (*)$$

Note that $d(g_i, e) < R$ for $i = 1, \dots, k$. Let $P \in \mathbb{N}^*$ be smallest such that $2P \geq R$. Note that $P \notin \mathbb{N}$ and $2P < R + 2$.

Claim. $B_{\Gamma^*}(e, 2P) \subseteq \bigcup_{i=1}^k B_{\Gamma^*}(g_i, P)$.

Proof of Claim. Suppose, toward a contradiction, that there is $y \in \Gamma^*$ such that $d(y, e) \leq 2P$ and yet $d(y, g_i) > P$ for all $i = 1, \dots, k$. On the one hand, for all $i = 1, \dots, k$, we have

$$d(y \mu, g_i \mu) = \text{st} \left(\frac{d(y, g_i)}{R} \right) \geq \text{st} \left(\frac{P}{R} \right) \geq \frac{1}{2},$$

whence $d(y \mu, e \mu) > 1$ by (*). On the other hand,

$$d(y \mu, e \mu) = \text{st} \left(\frac{d(y, e)}{R} \right) \leq \text{st} \left(\frac{2P}{R} \right) \leq \text{st} \left(1 + \frac{2}{R} \right) = 1.$$

We have thus arrived at a contradiction and the claim is proven.

The proof of the proposition is completed (using Exercise 12.7.3) once we realize that the claim implies that

$$B_{\Gamma^*}(e, P) \cdot B_{\Gamma^*}(e, P) \subseteq \bigcup_{i=1}^k g_i B_{\Gamma^*}(e, P). \quad \square$$

Exercise 12.7.5. Verify that the previous proof goes through under the weaker assumption that there is R such that, in $Y = \text{Cone}(\Gamma; R)$, $B_Y(e \mu, M)$ is compact for some $M > 0$.

By Theorem 12.3.6 and Proposition 12.7.4, groups of polynomial growth have the approximate subgroup property. Thus, the following fact of Hrusovski [86, Theorem 7.1] generalizes Gromov's theorem:

Fact 12.7.6. If Γ has the approximate subgroup property, then Γ is virtually nilpotent.

Corollary 12.7.7. *If there is R such that $\text{Cone}(\Gamma; R)$ is proper, then Γ is virtually nilpotent.*

Since virtually nilpotent groups have polynomial growth, the previous corollary indeed yields a positive answer to the question of van den Dries and Wilkie mentioned at the beginning of this section.

12.8. Nonhomeomorphic asymptotic cones

In this section, we revert to the metric ultraproduct viewpoint of asymptotic cones. When $r(n) = n$ for all n , we refer to the asymptotic cone as **standard** (a non-standard bit of terminology) and omit the scaling sequence from the notation, thus simply writing $\text{Cone}(\Gamma; \mathcal{U})$.

It is reasonable to ask if the homeomorphism type of the asymptotic cone depends on the choice of \mathcal{U} and $r \in (\mathbb{R}^{>0})^{\mathbb{N}}$. The first result shows that when we are considering the groups under discussion in Gromov's theorem, there is a unique asymptotic cone:

Theorem 12.8.1 (Pansu [142]). *If Γ is virtually nilpotent, then all of its asymptotic cones are isometric.*

It becomes reasonable to ask if the previous fact holds for all groups. This is not the case:

Theorem 12.8.2 (Thomas-Velickovic [176]). *There is a group Γ and $\mathcal{U}, \mathcal{V} \in \beta\mathbb{N} \setminus \mathbb{N}$ such that $\text{Cone}(\Gamma; \mathcal{U})$ is simply connected while $\text{Cone}(\Gamma; \mathcal{V})$ is not simply connected (whence these standard asymptotic cones are not homeomorphic).*

The group Γ in the previous theorem is not finitely presented. Ol'shanskii and Sapir [139] were later able to prove the same result using a finitely presented group.

Since there are $2^{\mathfrak{c}}$ many nonisomorphic ultrafilters on \mathbb{N} , it becomes natural to ask whether or not there is a group with $2^{\mathfrak{c}}$ many nonhomeomorphic standard asymptotic cones. Surprisingly (or perhaps not so surprisingly -see Chapter 15), it turns out that the answer to this depends on set theory:

Theorem 12.8.3 (Kramer, Shelah, Tent, and Thomas [109]).

- (1) *If CH is true, then for all Γ , there are at most \mathfrak{c} many pairwise nonhomeomorphic standard asymptotic cones of Γ .*
- (2) *If CH is false, then there is a group Γ with $2^{\mathfrak{c}}$ many nonhomeomorphic standard asymptotic cones.*

Incidentally, the group in the second item of the previous fact has a unique asymptotic cone (up to bi-Lipschitz homeomorphism) when CH holds.

The previous fact still does not give an example of a group where one can prove, in ZFC, that it has the maximal number of nonhomeomorphic asymptotic cones. This was soon established:

Theorem 12.8.4 (Drutu-Sapir [47]). *There is a group Γ with \mathfrak{c} many non-homeomorphic standard asymptotic cones.*

We leave the realm of standard asymptotic cones and mention one result which also shows that the homeomorphism type of the asymptotic cone can also be affected by the choice of scaling sequence:

Theorem 12.8.5 (Osin-Ould Houcine [140]). *There is a finitely presented group Γ such that, for all primes p , there is $r_p \in (\mathbb{R}^{>0})^{\mathbb{N}}$ with $\lim_{n \rightarrow \infty} r_p(n) = \infty$ such that, for all distinct primes p and q and all $\mathcal{U}, \mathcal{V} \in \beta\mathbb{N}$, $\text{Cone}(\Gamma; \mathcal{U}, r_p)$ is not homeomorphic to $\text{Cone}(\Gamma; \mathcal{V}, r_q)$.*

12.9. Notes and references

Our treatment of Gromov's theorem and the asymptotic cone follows [181] very closely. A recent spectacular application of ultraproduct methods is to the recent classification of finite approximate groups by Breuillard, Green, and Tao [20] following Hrushovski's breakthrough in [86]. The ideas are very similar in spirit as those used in the proof of Gromov's theorem and an excellent explanation of both and their connection to Hilbert's fifth problem can be found in Tao's book [172].

Sofic groups

In this chapter, we introduce the class of sofic groups, which is the class of all groups that can be embedded in a metric ultraproduct of symmetric groups endowed with their Hamming metric. In Section 13.1, we discuss the general topic of a metric ultraproduct of bi-invariant metric groups, of which the aforementioned ultraproduct of symmetric groups is a special case. In Section 13.2 we define the class of sofic groups in terms of almost homomorphisms and then give the above ultraproduct reformulation. In Section 13.3, we present numerous examples of sofic groups and list some closure properties of this class. Finally, in Section 13.4, we prove that the so-called Kervaire-Laudenbach conjecture on equations in groups holds for the class of sofic groups.

13.1. Ultraproducts of bi-invariant metric groups

By a **metric group**, we mean a triple (G, \cdot, d) such that (G, \cdot) is a group, (G, d) is a metric space, and \cdot is a continuous function with respect to the metric. Since these are examples of metric spaces with extra structure, we would like to consider taking ultraproducts of metric groups. As we learned in Section 11.4, we need to consider families of metric groups for which the group operation is uniformly continuous and such that a single modulus works for all of the elements in the family. Following Enflo [52], a metric group (G, \cdot, d) is called **uniform** if the multiplication is uniformly continuous with respect to d . There is a natural source of uniform metric groups:

Definition 13.1.1. A metric d on a group (G, \cdot) is said to be **bi-invariant** if, for every $a, b, c \in G$, we have $d(ab, ac) = d(ba, ca) = d(b, c)$. A metric

group (G, \cdot, d) where d is a bi-invariant metric for the group (G, \cdot) will be referred to as a **bi-invariant metric group**.

Lemma 13.1.2. *If (G, \cdot, d) is a bi-invariant metric group, then $\Delta(\epsilon) = \frac{\epsilon}{2}$ is a modulus of uniform continuity for multiplication.*

Proof. We simply calculate: if $d(a, c), d(b, d) < \frac{\epsilon}{2}$, then

$$d(ab, cd) \leq d(ab, cb) + d(cb, cd) = d(a, c) + d(b, d) < \epsilon.$$

□

Exercise 13.1.3. If (G, \cdot, d) is a bi-invariant metric group, then $\Delta(\epsilon) = \epsilon$ is a modulus of uniform continuity for the function $a \mapsto a^{-1}$.

Given any metric group (G, \cdot, d) , we always view it as a pointed metric space by taking the identity e_G as the basepoint. We now see that we can take the metric ultraproduct $\prod_{\mathcal{U}}(G_i, d_i)$ of a family of bi-invariant metric groups and the resulting family has a well-defined binary operation on it given by pointwise multiplication: $[a]_{\mathcal{U}} \cdot [b]_{\mathcal{U}} := [i \mapsto a(i) \cdot b(i)]_{\mathcal{U}}$. It is easy to verify that this binary operation is indeed a group operation on $\prod_{\mathcal{U}} G_i$ and the usual metric on $\prod_{\mathcal{U}} G_i$ is bi-invariant.

Exercise 13.1.4. Show that the completion of a bi-invariant metric group is once again a bi-invariant metric group. (*Hint.* Consider the ultrapower description of the completion given in Section 11.3.)

As we mentioned at the beginning of this section, we know that we can take the metric ultraproduct of a family of uniform metric groups (provided that they are uniform in a uniform way). We will not do that in the rest of this chapter and instead deal only with the special class of bi-invariant metric groups. Let us briefly give a reason why morally this is not too big a loss of generality.

A topological group G is a **SIN group** if there is a neighborhood base of the identity consisting of sets closed under conjugation. Note that every bi-invariant metric group is a SIN group. By [149, 2.17], being SIN is equivalent to: for every open neighborhood U of the identity e in G , there is an open neighborhood V of e such that $gVg^{-1} \subseteq U$ for all $g \in G$. In particular, every uniform metric group is SIN. However, for metric groups, SIN is equivalent to admitting a compatible bi-invariant metric by a result of Klee [107]. Thus, in summary: a metric group is uniform if and only if it admits a compatible bi-invariant metric.

We end this section with a list of examples of bi-invariant metric groups.

Example 13.1.5. Every group equipped with the discrete metric is a bi-invariant metric group.

Example 13.1.6. Every abelian metric group admits a compatible bi-invariant metric. Indeed, every metric group admits a compatible left-invariant metric (Birkhoff-Kakutani), which is necessarily bi-invariant since the group is abelian.

Example 13.1.7. Every compact metric group admits a compatible bi-invariant metric. Indeed, by compactness, the group must be uniform.

The next two examples are critical to the rest of the chapter.

Example 13.1.8. The **normalized Hamming metric on S_n** is the function

$$d_{\text{Ham},n}(\sigma, \tau) := d_{\text{Ham}}(\sigma, \tau) := \frac{1}{n} |\{k : \sigma(k) \neq \tau(k)\}|.$$

Exercise 13.1.9. Prove that $d_{\text{Ham},n}$ is a bi-invariant metric on S_n .

For the next example, recall that U_n is the collection of unitary $n \times n$ matrices (say over the complex numbers), where a matrix u is unitary if $u^* = u^{-1}$ (u^* being the conjugate transpose of u). U_n is then a subgroup of $\text{GL}_n(\mathbb{C})$ under multiplication.

Example 13.1.10. U_n is a bi-invariant metric group under the **Hilbert-Schmidt metric**:

$$d_{\text{HS}}(u, v) := \sqrt{\frac{1}{n} \sum_{i,j=1}^n |u_{ij} - v_{ij}|^2}.$$

Exercise 13.1.11. Prove that d_{HS} is a bi-invariant metric and that

$$d_{\text{HS}}(u, v)^2 = \text{tr}((u - v)(u - v)^*).$$

Here, tr is the normalized trace on $n \times n$ matrices, that is, $\text{tr}(a) = \frac{1}{n} \sum_{i=1}^n a_{ii}$.

13.2. Definition of sofic groups

A classical result in finite group theory is *Cayley's theorem*: if G is a finite group, then there is $n > 0$ for which there is an injective group homomorphism $G \rightarrow S_n$. The idea behind sofic groups is to ask, for an infinite group G , whether or not every finite subset of G admits an “almost injective group morphism” into some S_n . How do we quantify “almost”? For example, what would it mean for $\phi(a)\phi(b)$ to be “almost” the same as $\phi(ab)$? It is here that we will use the notion of Hamming distance on S_n introduced in the previous section.

Here is the precise definition of sofic group:

Definition 13.2.1. A group G is **sofic** if: for every finite $F \subseteq G$ and every $\epsilon > 0$, there is $n > 0$ and a function $\phi : F \rightarrow S_n$ such that:

- (1) for all $g, h \in F$, if $gh \in F$, then $d_{\text{Hamm}}(\phi(g)\phi(h), \phi(gh)) < \epsilon$;
- (2) if $e \in F$, then $d_{\text{Hamm}}(\phi(e), \text{id}) < \epsilon$; and
- (3) for all distinct $g, h \in F$, $d_{\text{Hamm}}(\phi(g), \phi(h)) \geq \frac{1}{2}$.

Remarks 13.2.2.

- (1) The idea behind the third requirement in the previous definition is that we want our functions ϕ to be injective in a way that is uniformly bounded away from 0 (or else we could satisfy the first two requirements in an artificial way). The precise choice of $\frac{1}{2}$ is, however, irrelevant, for we arrive at the same class of groups if we choose any number in $(0, 1)$. This will follow from our ultraproduct characterization of soficity appearing in Theorem 13.2.7; see Corollary 13.2.9.
- (2) Sometimes we will define ϕ as above to be a total function, that is, defined on all of G , even though its behavior on F is what is relevant. In these situations, we might refer to ϕ as a (F, ϵ) -**morphism**.

Example 13.2.3. By the *proof of* Cayley's theorem, every finite group is sofic. Indeed, if G is a finite group, then one gets a homomorphism $\phi : G \rightarrow S_G$, where S_G is the symmetric group on the set G , given by $\phi(g)(h) := gh$. In what follows, we simply identify S_G with S_n , where $n = |G|$. Since ϕ is an actual homomorphism, the first two requirements in the definition of soficity are satisfied. To see the third requirement, notice that $\phi(g)(h) = \phi(g')(h)$ if and only if $gh = g'h$, that is, if and only if $g = g'$. It follows that $d_{\text{Hamm}}(\phi(g), \phi(g')) = 1$ for distinct $g, g' \in G$.

Exercise 13.2.4.

- (1) For any $m, n \geq 1$, prove that the embedding $\eta : S_m \hookrightarrow S_{mn}$ given by $\eta(\sigma)(im + j) = \sigma(j)$ for any $i = 0, \dots, n-1$ and $j = 1, \dots, m-1$ is a group homomorphism that is also an isometric embedding with respect to the normalized Hamming metrics.
- (2) For any $m \geq 1$, prove that the embedding $\iota : S_n \times S_n \hookrightarrow S_{2n}$ given by $\iota(\sigma, \tau)(i) = \sigma(i)$ if $i = 1, \dots, n$ while $\iota(\sigma, \tau)(i) = \tau(i - n)$ if $i = n + 1, \dots, 2n$, is a group homomorphism such that

$$d_{\text{Hamm}}(\iota(\sigma_1, \tau_1), \iota(\sigma_2, \tau_2)) = \frac{1}{2} [d_{\text{Hamm}}(\sigma_1, \tau_1) + d_{\text{Hamm}}(\sigma_2, \tau_2)].$$

- (3) Use parts (1) and (2) to show that if G_1 and G_2 are both sofic groups, then so is $G_1 \times G_2$.

We will see many more examples of sofic groups later on. The reason we bring up the topic of soficity in this book is that we can characterize sofic

groups in terms of metric ultraproducts. First, we need a quick detour concerning the so-called *amplification trick*. Given $\sigma \in S_n$ and $k \geq 1$, let $\sigma^{\otimes k} \in S_{n^k}$ be the permutation defined by $\sigma^{\otimes k}(i_1, \dots, i_k) := (\sigma(i_1), \dots, \sigma(i_k))$; here, we view S_{n^k} as the set of permutations of tuples (i_1, \dots, i_k) , where each $i_j \in \{1, \dots, n\}$. We need the following lemma:

Lemma 13.2.5. $1 - d_{\text{Hammm}}(\sigma^{\otimes k}, \tau^{\otimes k}) = (1 - d_{\text{Hammm}}(\sigma, \tau))^k$.

Exercise 13.2.6. Prove the previous lemma.

We can now prove the ultraproduct reformulation of soficity:

Theorem 13.2.7. *The group G is sofic if and only if there is a set I , an ultrafilter \mathcal{U} on I , a family $(n_i)_{i \in I}$ of natural numbers, and an injective group homomorphism $\phi : G \rightarrow \prod_{\mathcal{U}} S_{n_i}$, where the latter ultraproduct is the metric ultraproduct with each S_{n_i} equipped with their respective Hamming metric.*

Proof. First suppose that G is sofic. Let I be the set of all pairs $(F, \frac{1}{n})$, where $F \subseteq G$ is finite and $n > 0$. For such a pair, let $X_{(F, \frac{1}{n})} := \{(F', \frac{1}{m}) \in I : F' \supseteq F, m \geq n\}$. Note that the family $(X_{(F, \frac{1}{n})})_{(F, \frac{1}{n}) \in I}$ has the FIP, whence we may take an ultrafilter \mathcal{U} on I containing each $X_{(F, \frac{1}{n})}$. For each $(F, \frac{1}{n}) \in I$, let $\phi_{(F, \frac{1}{n})} : G \rightarrow S_{(F, \frac{1}{n})}$ be a $(F, \frac{1}{n})$ -morphism. For each $g \in G$, define $a_g \in \prod_{(F, \frac{1}{n}) \in I} S_{(F, \frac{1}{n})}$ by $a_g(F, \frac{1}{n}) := \phi_{(F, \frac{1}{n})}(g)$. We can then define $\phi : G \rightarrow \prod_{\mathcal{U}} S_{(F, \frac{1}{n})}$ by $\phi(g) = [a_g]_{\mathcal{U}}$. We leave it to the reader to check that ϕ is an injective group morphism.

We now prove the converse. We will proceed with the naïve idea and see that a modification will be needed. Suppose that $\phi : G \rightarrow \prod_{\mathcal{U}} S_{n_i}$ is an injective group morphism. Fix $F \subseteq G$ finite and $\epsilon > 0$. We can then find $n > 0$ and an injective function $\theta : F \rightarrow S_n$ such that items (1) and (2) of the definition are met. Indeed, the map $\theta : F \rightarrow S_{n_i}$ given by $\theta(g) = \phi(g)(i)$ works for \mathcal{U} -almost all i . However, since we only know that $\phi(x) \neq \phi(y)$ for all distinct $x, y \in F$, this merely tells us that $d_{\text{Hammm}}(\theta(x), \theta(y)) \geq \delta$ for some (potentially) small $\delta > 0$, not the $\frac{1}{2}$ needed in the definition of soficity. However, we can consider $\theta^{\otimes k} : F \rightarrow S_{n^k}$ given by $\theta^{\otimes k}(g) := \theta(g)^{\otimes k}$. Then $1 - d_{\text{Hammm}}(\theta^{\otimes k}(x), \theta^{\otimes k}(y)) \leq (1 - \delta)^k < \frac{1}{2}$ if k is large enough.

Now unfortunately, after this amplification process, we have ruined the ϵ -almost homomorphism property (that is, requirements (1) and (2) in the definition of soficity). Nevertheless, not all is lost. Indeed, note that, given $F \subseteq G$, we can calculate the δ above (independently of ϵ) and from δ we calculate k . Thus, if we choose θ to be an ϵ' -almost homomorphism for sufficiently small $\epsilon' > 0$, then after amplification by k , the new almost-homomorphism will be an ϵ -almost homomorphism. \square

Exercise 13.2.8. Verify in detail the conclusion of the previous proof.

Corollary 13.2.9. *If, in the definition of soficity, we replace $\frac{1}{2}$ with any number in $(0, 1)$, we get the same class of sofic groups.*

Proof. The proof of the previous theorem goes through with any number in $(0, 1)$ instead of $\frac{1}{2}$. \square

Exercise 13.2.10. Prove that a group G is sofic if and only if there is a set I , an ultrafilter \mathcal{U} on I , a family $(n_i)_{i \in I}$ of natural numbers, and a *distance-preserving* group homomorphism $\phi : G \rightarrow \prod_{\mathcal{U}} S_{n_i}$, where G is equipped with the discrete metric. (*Hint.* Mimic the proof of Theorem 13.2.7 using Corollary 13.2.9.)

Note that the proof of Theorem 13.2.7 shows that a countable sofic group can be embedded into a metric ultraproduct of sofic groups indexed over a countable set. One can in fact say more. Since the proof would take us too far afield, we omit it.

Theorem 13.2.11. *Let G be a countable group. The following are equivalent:*

- (1) G is sofic.
- (2) G embeds into $\prod_{\mathcal{U}} S_n$ for all $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$.
- (3) G embeds into $\prod_{\mathcal{U}} S_n$ for some $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$.

In the next section, we will gather a large collection of examples of sofic groups. Surprisingly, the following question is open and is arguably one of the major open problems in modern group theory:

Question 13.2.12. Are all groups sofic?

It is natural to wonder: why metric ultraproducts? That is, is it possible that every group embeds into a *discrete* ultraproduct of symmetric groups? It turns out the answer to this question is no, but is not as simple as one might expect. Note that, by Cayley's theorem, a group embeds into a discrete ultraproduct of symmetric groups if and only if it embeds into a discrete ultraproduct of finite groups.

Theorem 13.2.13. *There is a group that is not embeddable into any discrete ultraproduct of finite groups. In fact, one can even find a finitely generated such group.*

We offer two proofs of this theorem.

Proof 1. Let σ be the sentence $(\exists x)(\exists y)(\exists z)[x^2y = yx^2 \wedge xy \neq yx \wedge z^{-1}xz = x^2]$. Suppose that $G \models \sigma$ and take $a, b, c \in G$ such that $a^2b = ba^2$, $ab \neq ba$, and $c^{-1}ac = a^2$. We claim that c must have infinite order. To see this,

for any $g \in G$, let $C(g) := \{h \in G : gh = hg\}$, the centralizer of g in G . We note that, for any $g, h \in G$, that $C(g)$ is a subgroup of $C(g^2)$ and $h^{-1}C(g)h = C(h^{-1}gh)$. In our context, $C(a)$ is a proper subgroup of $C(a^2)$, denoted $C(a) < C(a^2)$, and $c^{-1}C(a)c = C(c^{-1}ac) = C(a^2)$. Consequently, we have $C(a) < c^{-1}C(a)c$. Iterating this, we see that $C(a) < c^{-1}C(a)c < c^{-2}C(a)c^2 < \dots$, whence it follows that c has infinite order.

Consequently, any finite group models $\neg\sigma$, and thus any ultraproduct of finite groups models $\neg\sigma$. Since $\neg\sigma$ is a universal sentence, it follows that any subgroup of an ultraproduct of finite groups also models $\neg\sigma$.

Now let $X = \mathbb{N} \times \mathbb{Z}$ and let $G = S_X$, the group of permutations of X . We show that $G \models \sigma$, whence G is not a subgroup of an ultraproduct of finite groups. Let $f \in S_X$ be defined by $f(i, j) = (i, j + 1)$. It is routine to check that f and f^2 are conjugate in G . Now let $g \in G$ be such that $g(0, 2j) = (0, 2j + 2)$ for all $j \in \mathbb{Z}$ whilst g fixes all other elements of X . Note that $f^2g = gf^2$ but $fg \neq gf$. Thus $G \models \sigma$.

To get a finitely generated counterexample, just take the subgroup of G generated by f, g , and the element that conjugates f to f^2 . \square

Proof 2. Let G be an infinite, simple, finitely presented group. (For example, one can take Thompson's group V [21].) Recall that for G to be finitely presented means that there are $a_1, \dots, a_n \in G$ and words $w_1(x), \dots, w_m(x)$, where $x = (x_1, \dots, x_n)$, such that:

- G is generated by a_1, \dots, a_n ,
- $w_i(a) = e$ for $i = 1, \dots, m$ (where $a = (a_1, \dots, a_n)$), and
- given any other group H and $b_1, \dots, b_n \in H$ for which $w_i(b) = e$ for all $i = 1, \dots, m$, the map $\phi : \{a_1, \dots, a_n\} \rightarrow H$ defined by $\phi(a_i) = b_i$ for all $i = 1, \dots, n$ extends to a group homomorphism $\phi : G \rightarrow H$.

Now suppose, toward a contradiction, that $\phi : G \rightarrow \prod_{\mathcal{U}} H_j$ is an injective group homomorphism of G into a discrete ultraproduct of finite groups. Take $a_1, \dots, a_n \in G$ and words w_1, \dots, w_m as above. Set $[b_i]_{\mathcal{U}} := \phi(a_i)$ for $i = 1, \dots, m$. Then $w_i([b_1]_{\mathcal{U}}, \dots, [b_n]_{\mathcal{U}}) = e$ for all $i = 1, \dots, m$, whence, by Loś's theorem, for \mathcal{U} -almost all j , we have $w_i(b_1(j), \dots, b_n(j)) = e$ for all $i = 1, \dots, m$. Note also that, for \mathcal{U} -almost all j , we have $b_1(j) \neq e$. (This assumes that $a_1 \neq e$, which is a harmless assumption.) Thus, for any j satisfying the previous two properties, we have group homomorphisms $\phi_j : G \rightarrow H_j$ defined by $\phi_j(a_i) = b_i(j)$ for $i = 1, \dots, n$ and for which ϕ_j is not constantly equal to the identity (as $b_1(j) \neq e$). Since G is simple, the kernel of ϕ_j is either all of G or $\{e\}$; since we just showed that the former

possibility is not the case, this implies that ϕ_j is injective, contradicting the fact that G is infinite and H_j is finite. \square

13.3. Examples of sofic groups

In this section, we focus on examples of sofic groups. We can give a large source of examples of sofic groups by establishing that every *amenable* group is sofic.

Definition 13.3.1. A group G is **amenable** if, for every finite $K \subseteq G$ and $\epsilon > 0$, there is a nonempty finite $F \subseteq G$ such that, for all $g \in K$, we have $|gF \Delta F| < \epsilon|F|$. We call such a set F a (K, ϵ) -**Følner set**.

Some exercises might help us become acquainted with this notion:

Exercise 13.3.2. Prove that every finite group is amenable.

Exercise 13.3.3. Prove that \mathbb{Z} is amenable.

Exercise 13.3.4. Prove that amenability is a local property in the sense that a group G is amenable if and only if every finitely generated subgroup of G is amenable.

Exercise 13.3.5. Suppose that G is a *countable* group. Prove that G is amenable if and only if G has a **Følner sequence**, namely a sequence of nonempty finite sets $F_n \subseteq G$ such that, for every $g \in G$, we have

$$\lim_{n \rightarrow \infty} \frac{|gF_n \Delta F_n|}{|F_n|} = 0.$$

Amenability is an incredibly robust notion in that it has countless, seemingly very different, reformulations. We offer one such reformulation here:

Theorem 13.3.6. G is amenable if and only if there is a left-invariant finitely-additive probability measure on G .

Proof Sketch. We only prove the forward direction under the simplifying assumption that G is countable. In this case, let $(F_n)_{n \in \mathbb{N}}$ be a Følner sequence for G and, fixing $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$, set

$$\mu(A) := \lim_{\mathcal{U}} \frac{|A \cap F_n|}{|F_n|}$$

for every $A \subseteq G$. It is straightforward to check that μ is a left-invariant finitely additive probability measure on G .

The other direction is more difficult and requires some nontrivial functional analysis. \square

Exercise 13.3.7. Suppose that G is amenable.

- (1) Prove that every subgroup of G is amenable.
- (2) Suppose that $f : G \rightarrow H$ is a surjective group homomorphism. Prove that H is also amenable.

The converse of Exercise 13.3.7 is also true in the following sense:

Proposition 13.3.8. *Suppose that $f : G \rightarrow H$ is a surjective group homomorphism such that $\ker(f)$ and H are amenable. Then G is also amenable.*

Proof sketch. Let $K \subseteq G$ be finite. Take a finite $F \subseteq G$ such that $f \upharpoonright F$ is injective and such that $|f(g)f(F)\Delta f(F)| < \epsilon|f(F)|$ for all $g \in K$; this is possible since H is amenable. For $g \in K$ and $x \in F$ such that $f(gx) \in f(F)$, take $y(g, x) \in \ker(f)$ such that $f(gx) \in f(F)y(g, x)$. Let T be the set of such $y(g, x)$'s, a finite subset of $\ker(f)$. Since $\ker(f)$ is amenable, there is finite $L \subseteq \ker(f)$ such that $|yL\Delta L| < \epsilon|L|$ for all $y \in T$. Let $M := FL$, a finite subset of G . Note that for $g \in K$ and “most” $x \in F$ and $z \in L$, we have $g(xz) \in FL$, as $gx \in Fy(g, x)$ and so $gxz \in Fy(g, x)L \subseteq F$. More precisely, we leave it to the reader to verify that M is a $(F, 4\epsilon)$ -Følner set. \square

Exercise 13.3.9. Verify the claim made at the end of the proof of the previous proposition.

The next proposition shows that many groups are amenable.

Proposition 13.3.10. *If G is virtually solvable, then G is amenable.*

Proof. Since amenability is local (Exercise 13.3.4), it suffices to assume that G is countable. Since virtually solvable groups contain *normal* solvable groups of finite index, by Exercises 13.3.2 and 13.3.8, it suffices to assume that G itself is solvable. By Exercise 13.3.8 again, we may induct on the derived length of G , allowing us to assume that G is actually abelian. Since G is the union of its finitely generated subgroups, using the locality of amenability again, we can assume that G is finitely generated abelian. By the fundamental theorem of finitely generated abelian groups, G is isomorphic to the direct sum of finitely many cyclic groups. Exercise 13.3.8 shows that the direct sum of finitely many amenable groups is amenable, whence we reduce to the case that G is cyclic. The fact that all cyclic groups are amenable follows from Exercises 13.3.2 and 13.3.3. \square

Nevertheless, some important groups are not amenable:

Example 13.3.11. \mathbb{F}_2 is not amenable. While one can prove that \mathbb{F}_2 is not amenable directly from the definition of amenability given above, we choose to use the characterization given in Theorem 13.3.6. Suppose, toward a

contradiction, that μ is a left-invariant, finitely additive probability measure on G . Note immediately that $\mu(\{e\}) = 0$, for if $\mu(\{e\}) = r > 0$, then $\mu(\{g\}) = r$ for every $g \in \mathbb{F}_2$ by left-invariance of μ , whence $\mu(F) = |F|r$ for any finite subset $F \subseteq \mathbb{F}_2$ by finite additivity; taking F such that $|F|r > 1$ contradicts the fact that $\mu(\mathbb{F}_2) = 1$.

Let a and b denote the two generators of \mathbb{F}_2 . Let X_a denote the set of all elements of \mathbb{F}_2 consisting of reduced words beginning with the letter a . Define the sets X_b , $X_{a^{-1}}$, and $X_{b^{-1}}$ similarly. Note that

$$a^{-1}X_a = X_a \cup X_b \cup X_{b^{-1}} \cup \{e\},$$

which, by left-invariance of μ , leads to the equation

$$\mu(X_a) = \mu(a^{-1}X_a) = \mu(X_a) + \mu(X_b) + \mu(X_{b^{-1}}) + \mu(\{e\}),$$

whence we conclude that $\mu(X_b) = \mu(X_{b^{-1}}) = 0$. In a symmetric fashion, one concludes that $\mu(X_a) = \mu(X_{a^{-1}}) = 0$. Consequently, one arrives at the ridiculous conclusion that $\mu(\mathbb{F}_2) = 0$.

Returning to soficity, we have:

Theorem 13.3.12. *Amenable groups are sofic.*

Proof. Let G be an amenable group. Since both amenability and soficity are local, we may as well assume that G is countable. Let $(F_n)_{n \in \mathbb{N}}$ be a Følner sequence for G . Thus, each element of G “almost” acts like a permutation of F_n . More precisely, for each $g \in G$ and $n \in \mathbb{N}$, let $\phi_n(g)$ be the partial function from F_n to F_n defined by $\phi_n(g)(x) := gx$ if $gx \in F_n$. Clearly $\phi_n(g)$ is an injection, so we can extend $\phi_n(g)$ to a permutation of F_n in an arbitrary way, thus getting a function $\phi_n : G \rightarrow S_{F_n}$. For each $g \in G$, let $a_g \in \prod_{n \in \mathbb{N}} S_{F_n}$ be given by $a_g(n) := \phi_n(g)$. We leave it to the reader to verify that the map $\phi : G \rightarrow \prod_{\mathcal{U}} S_{F_n}$ given by $\phi(g) := [a_g]_{\mathcal{U}}$ is an injective group homomorphism. \square

Unfortunately, the previous theorem does not show that free groups are sofic. Nevertheless, free groups are indeed sofic, as we now explain. We first notice the following fact:

Proposition 13.3.13. *If a group is residually sofic, then it is actually sofic.*

Proof. Suppose that G is a residually sofic group and take $F \subseteq G$ finite and $\epsilon > 0$. Without loss of generality, we may assume that $e \in F$. By Exercises 10.3.5 and 13.2.4, we have that G is fully residually sofic. Let H be a sofic group and let $\phi : G \rightarrow H$ be a group homomorphism such that $\phi(x) \neq e$ for any $x \in FF^{-1} \setminus \{e\}$. It follows that $\phi \upharpoonright F$ is injective. Let $\psi : H \rightarrow S_n$ be a $(\phi(F), \epsilon)$ -morphism. It is routine to verify that the map $\psi \circ \phi : G \rightarrow S_n$ is a (F, ϵ) -morphism. \square

Note that, in particular, we have that every residually amenable (and thus every residually finite) group is sofic.

It remains now to note:

Proposition 13.3.14. *Free groups are residually finite.*

Proof. Let a_1, \dots, a_m be the generators of the free group \mathbb{F}_m . Consider a nontrivial reduced word $g := a_{i_n}^{e_n} \cdots a_{i_2}^{e_2} a_{i_1}^{e_1}$, with each $e_i = \pm 1$. Set $[n+1] := \{1, \dots, n+1\}$. Define partial functions f_1, \dots, f_m from $[n+1]$ to $[n+1]$ by $f_{i_k}(k) = k+1$ if $e_k = 1$ while $f_{i_k}(k+1) = k$ if $e_k = -1$. (Note that these definitions do not conflict with one another since the word is reduced.). Now each f_{i_k} is only a partially defined injection, so we can extend it to an actual element of S_{n+1} in any way we like. Note that $f_{i_k}^{e_k}(k) = k+1$ for all $k = 1, \dots, n$. Thus, the group homomorphism $\phi : \mathbb{F}_m \rightarrow S_{n+1}$ obtained by mapping each a_i to f_i is such that $\phi(g)(1) = n+1$, whence $\phi(g)$ is not the identity. The desired result follows. \square

Corollary 13.3.15. *Residually free groups are sofic. In particular, universally free groups are sofic.*

The previous corollary lends itself to a curious result:

Corollary 13.3.16. *All groups are sofic if and only if the class of sofic groups is closed under taking quotients.*

Proof. The forward direction is obvious. To prove the backward direction, assume that the class of sofic groups is closed under taking quotients and suppose that G is an arbitrary group. By locality of soficity, it suffices to assume that G is finitely generated. If G is generated by n elements, then there is a surjective group homomorphism $\pi : \mathbb{F}_n \rightarrow G$. Since \mathbb{F}_n is sofic, the standing assumption implies that G is sofic, as desired. \square

Proposition 13.3.8 above may be recast as follows: if G is a group with a normal subgroup N such that N and G/N are both amenable, then G is also amenable. Surprisingly, the sofic version of this result is unknown and appears to be difficult. Nevertheless, there is a very interesting partial result in this direction:

Theorem 13.3.17 (Elek-Szabo [48]). *Suppose that G is a group with a normal subgroup N such that N is sofic and G/N is amenable. Then G is sofic.*

Before proving this result, let us set up some notation that will make the proof run more smoothly. For $g \in G$, let $\bar{g} \in G/N$ denote the coset gN . For any $X \subseteq G$, we set $\bar{X} := \{\bar{g} : g \in X\} \subseteq G/N$. We fix any function $r : G/N \rightarrow G$ such that $r(\bar{g}) = \bar{g}$. In other words, r is a section of the

quotient map: it picks out one element from each coset. We set $f : G \rightarrow G$ to be $f(g) = r(\bar{g})$ (so f just assigns to a group element the distinguished element of its coset) and $s : G \rightarrow N$ to be $s(g) := f(g)^{-1}g$. Here are some easy facts about these functions:

Exercise 13.3.18. Suppose that $g, h \in G$. With the notation from the previous paragraph, prove the following:

- (1) $\overline{f(g)} = \bar{g}$.
- (2) $f(r(\bar{g})) = r(\bar{g})$.
- (3) $s(r(\bar{g})) = e$.
- (4) $f(gh) = f(gf(h))$.
- (5) $s(gh) = s(gf(h))s(h)$.

To avoid subscripts that are hard to parse in the following proof, we write $\text{Sym}(X)$ as a synonym for S_X , the group of permutations of a set X .

Proof of Theorem 13.3.17. Fix $\epsilon > 0$ and finite $F \subseteq G$. Since G/N is amenable, we may take an (\bar{F}, ϵ) -Følner set \bar{A} for G/H . In other words, given any $g \in F$, we have, for at least $(1 - \epsilon)|\bar{A}|$ many $\bar{a} \in \bar{A}$, that $\overline{g\bar{a}} \in \bar{A}$. Without loss of generality, we may assume that A is contained in the range of r or, in other words, that $r(\bar{g}) = g$ for all $g \in A$. In this case, for $g \in G$, if $\overline{g\bar{a}} \in \bar{A}$, say $\overline{g\bar{a}} = \bar{a}'$ for some $a' \in A$, then $f(ga) = a'$. Consequently, for each $g \in G$, the partial map from A to A given by mapping a to $f(ga)$ whenever $\overline{g\bar{a}} \in \bar{A}$ is injective, whence we may extend it to a bijection $l(g) : A \rightarrow A$. Since N is sofic, we may also fix a (\hat{F}, ϵ) -morphism $\psi : N \rightarrow S_n$, where \hat{F} is a finite subset of N to be determined soon.

Claim 1. For each $g \in G$, the function $\phi(g) : [n] \times A \rightarrow [n] \times A$ given by $\phi(g)(i, a) := (\psi(s(ga))(i), f(ga))$ if $\overline{g\bar{a}} \in \bar{A}$ while $\phi(g)(i, a) = (i, l(g)(a))$ otherwise, is a permutation of $[n] \times A$.

Proof of Claim 1. Given $i' \in [n]$ and $a' \in A$, we see that either there is $g \in G$ such that $f(ga) = a'$, in which case we take $i \in [n]$ such that $\psi(s(ga))(i) = i'$ and then $\phi(g)(i, a) = (i', a')$; otherwise, since $l(g)$ is surjective, we can find $a \in A$ such that $l(g)(a) = a'$ and then $\phi(i', a) = (i', a')$.

We thus have a map $\phi : G \rightarrow \text{Sym}([n] \times A)$. This map ϕ will be a $(F, 3\epsilon)$ -morphism; since ϵ is arbitrary, this will verify that G is sofic. We break this verification up into smaller bits.

Claim 2. If $e \in F$, then $d_{\text{Hamm}}(\phi(e), \text{id}) < \epsilon$.

Proof of Claim 2. Note that $\phi(e)(i, a) = (\psi(e)(i), a)$ for all $(i, a) \in [n] \times A$; since $\psi(e)(i) = i$ for at least $(1 - \epsilon)n$ many $i \in [n]$, we get that $\phi(e)(i, a) = (i, a)$ for at least $(1 - \epsilon)n|A|$ many $(i, a) \in n \times A$.

Claim 3. *If $g, h \in F$ are distinct, then $d_{\text{Hamm}}(\phi(g), \phi(h)) \geq 1 - 2\epsilon$.*

Proof of Claim 3. We distinguish between two cases. First suppose that $\bar{g} = \bar{h}$. If $\overline{ga} \in \bar{A}$ (which happens for at least $(1 - \epsilon)|A|$ many $a \in A$), then $f(ga) = f(ha)$, whence $s(ga) = f(ga)^{-1}ga$ and $s(ha) = f(ha)^{-1}ha$ are distinct elements of $A^{-1}FA$. Thus, setting $\hat{F} := A^{-1}FA$, we see that for at least $(1 - \epsilon)n$ many elements of $[n]$, $\psi(s(ga))(i) \neq \psi(s(ha))(i)$. Summarizing, in this case, we have that for at least $(1 - \epsilon)^2n|A| \geq (1 - 2\epsilon)n|A|$ many $(i, a) \in [n] \times A$, we have that $\phi(g)(i, a) \neq \phi(h)(i, a)$.

Now suppose that $\bar{g} \neq \bar{h}$. Set $X := \{a \in A : \overline{ga} \notin \bar{A}\}$ and $Y := \{a \in A : \overline{ha} \notin \bar{A}\}$. If $a \notin X \cup Y$, then $l(g)(a) = f(ga)$ and $l(h)(a) = f(ha)$, which implies $l(g)(a) \neq l(h)(a)$. Since $|X \cup Y| \leq 2\epsilon|A|$, we see that for at least $(1 - 2\epsilon)n|A|$ many $(i, a) \in [n] \times A$, we have $\phi(g)(i, a) \neq \phi(h)(i, a)$, as desired.

Claim 4. *If $g, h \in F$ are such that $gh \in F$, then $d_{\text{Hamm}}(\phi(g)\phi(h), \phi(gh)) < \epsilon$.*

Proof of Claim 4. Suppose that $a \in A$ is such that $\overline{ha}, \overline{gha} \in A$; note that this happens for at least $(1 - 2\epsilon)|A|$ many $a \in A$. We then have

$$\phi(gh)(i, a) = (\psi(s(gha))(i), f(gha)) = (\psi(s(gf(ha))s(ha))(i), f(gf(ha))).$$

On the other hand,

$$\begin{aligned} \phi(g)(\phi(h)(i, a)) &= \phi(g)(\psi(s(ha))(i), f(ha)) \\ &= (\psi(s(gf(ha)))(\psi(s(ha))(i), f(gf(ha))). \end{aligned}$$

Note that $s(gf(ha)), s(ha) \in \hat{F}$. Thus, for at least $(1 - \epsilon)n$ many $i \in [n]$, the above two displays are equal. Summarizing, for at least $(1 - 3\epsilon)n|A|$ many (i, a) , we have that $\phi(gh)(i, a) = \phi(g)(\phi(h)(i, a))$, which is what we wanted to prove.

This concludes the proof of the theorem. □

13.4. An application of sofic groups

An important reason for considering the class of sofic groups is that there are many instances of famous conjectures about groups that can be solved under the further assumption that the group in question is sofic. In this section, we treat one such example.

For a given group G , consider the equation

$$(\ddagger\ddagger) \quad x^{k_1}g_1x^{k_2}g_2 \cdots x^{k_s}g_s = e,$$

where x is a variable, $g_1, \dots, g_s \in G$, and $k_1, \dots, k_s \in \mathbb{Z}$. Unlike the case of polynomial equations over a field, where one can always find a solution in

an extension field, we need not always be able to solve equation $(\dagger\dagger)$ in a group extending G :

Example 13.4.1. If the equation $x^{-1}axb^{-1} = e$ has a solution in a group extending G , then a and b are conjugate in that extension. Thus, if, for example, a and b have different orders, then the equation has no solution in an extension group.

Equation $(\dagger\dagger)$ is called **regular** if $k_1 + \cdots + k_s \neq 0$. Notice the equation in the previous example is *not* regular.

Conjecture 13.4.2 (Kervaire-Laudenbach). *For any regular equation $(\dagger\dagger)$ with coefficients from G , $(\dagger\dagger)$ has a solution in a group extending G .*

Let us say that G is a **KL-group** if G satisfies the Kervaire-Laudenbach conjecture. If, in addition, the solution can actually be found in G , we say that G is a **strong KL-group**. The main theorem of this section is that sofic groups are KL-groups. The first piece of the puzzle is the following:

Theorem 13.4.3 (Gerstenhaber-Rothaus [65]). *For each n , U_n is a strong KL-group.*

The proof of the previous fact uses some notions from algebraic topology, so we only sketch the main idea. Consider equation $(\dagger\dagger)$ over U_n . Let $f : U_n \rightarrow U_n$ be given by $f(x) := x^{k_1}g_1x^{k_2}g_2 \cdots x^{k_s}g_s$. It is enough to show that f is onto. Let $k := k_1 + \cdots + k_s$. If each $g_i = 1$, then f is simply the function $f_k(x) := x^k$, which is onto. (Exercise.) Since U_n is path connected, we can continuously move each g_i toward the identity, showing that the functions f and f_k are *homotopic*. This implies that f and f_k have the same *degree*, namely n^k , which roughly means that f “wraps” U_n around itself n^k times. In particular, this implies that f is also onto.

Corollary 13.4.4. *For any nonprincipal ultrafilter \mathcal{U} on \mathbb{N} , the metric ultraproduct $\prod_{\mathcal{U}} U_n$ is a strong KL-group.*

Exercise 13.4.5. Prove the previous corollary.

Now to every permutation $\sigma \in S_n$ we have the corresponding permutation matrix $A_\sigma \in U_n$ given by $(A_\sigma)_{ij} = 1$ if $\sigma(i) = j$, and 0 otherwise.

Exercise 13.4.6. The map $\sigma \mapsto A_\sigma : S_n \rightarrow U_n$ is a group homomorphism.

Now unfortunately this map is not distance preserving, when S_n is equipped with its Hamming distance and when U_n is equipped with its Hilbert-Schmidt distance. However, we do have the following:

Lemma 13.4.7. *For all $\sigma, \tau \in S_n$, we have $d_{\text{Hamm}}(\sigma, \tau) = \frac{1}{2}(d_{\text{HS}}(A_\sigma, A_\tau))^2$.*

Proof. Let $X = \{i : \sigma(i) \neq \tau(i)\} = \{i : (\sigma^{-1}\tau)(i) \neq i\} = \{i : (\tau^{-1}\sigma)(i) \neq i\}$. Note then that

$$\operatorname{tr}(I - A_{\sigma\tau^{-1}}) = \operatorname{tr}(I - A_{\tau\sigma^{-1}}) = \frac{|X|}{n},$$

so

$$d_{\text{Hammm}}(\sigma, \tau) = \frac{|X|}{n} = \operatorname{tr}((A_\sigma - A_\tau)(A_\sigma - A_\tau)^*).$$

The result now follows from Exercise 13.1.11. \square

As a result, we have:

Corollary 13.4.8. *For any ultrafilter \mathcal{U} on \mathbb{N} , there is an injective group homomorphism $\prod_{\mathcal{U}} S_n \rightarrow \prod_{\mathcal{U}} U_n$, where the ultraproducts are the metric ultraproducts and the factors are equipped with the Hamming metrics and Hilbert-Schmidt metrics, respectively.*

Proof. Given $\sigma \in \prod_{n \in \mathbb{N}} S_n$, define $A_\sigma \in \prod_{n \in \mathbb{N}} U_n$ by $A_\sigma(n) := A_{\sigma(n)}$. The desired homomorphism is defined by sending $[\sigma]_{\mathcal{U}}$ to $[A_\sigma]_{\mathcal{U}}$. For this to be well defined and injective, we need to know that $\lim_{\mathcal{U}} d_{\text{Hammm}}(\sigma(n), \text{id}) = 0$ if and only if $\lim_{\mathcal{U}} d_{\text{HS}}(A_{\sigma(n)}, I) = 0$, which follows immediately from the previous lemma. \square

We can now conclude:

Theorem 13.4.9 (Pestov [143]). *Sofic groups are KL-groups.*

Proof. By Corollary 13.4.4, $\prod_{\mathcal{U}} U_n$ is a strong KL-group. By Corollary 13.4.8, any sofic group is isomorphic to a subgroup of $\prod_{\mathcal{U}} U_n$; since a subgroup of a strong KL-group is clearly a KL-group, the theorem now follows. \square

Definition 13.4.10. A group G is **hyperlinear** if there is an injective group homomorphism from G into a metric ultraproduct $\prod_{\mathcal{U}} U_{n_i}$ of unitary groups.

Corollary 13.4.4 showed that hyperlinear groups are KL-groups. Corollary 13.4.8 immediately implies:

Theorem 13.4.11. *Every sofic group is hyperlinear.*

The following questions are still open:

Question 13.4.12.

- (1) Is every group hyperlinear?
- (2) Is every hyperlinear group sofic?

Question (1) is related to a famous open problem in operator algebras, namely the *Connes embedding problem*, which will be discussed in Section 14.5.

13.5. Notes and references

The notion of sofic group was introduced by Gromov in [72], although the term “sofic” was introduced by Weiss in [183]. Our treatment of sofic groups follows the manuscript of Lupini and Capraro [22], although this manuscript contains much more information and further applications of the notion of soficity. Another nice source is Pestov’s survey [143], which also contains a connection with the von Neumann algebra ultraproducts that will be discussed in the next chapter. Our discussion on amenable groups comes from Tao’s blog post

<https://terrytao.wordpress.com/2009/04/14/some-notes-on-amenability/>

while a more comprehensive introduction can be found in [151]. Proof 1 of Theorem 13.2.13 comes from Macpherson’s article [117, Theorem 6.0.23].

Functional analysis

In this chapter, we present a selected assortment of applications of the metric ultraproduct to functional analysis. In Section 14.1, we introduce the Banach space ultraproduct, and we apply this construction in Section 14.2 to study the local geometry of Banach spaces. In Section 14.3 we study C^* -algebras, proving that the category of commutative unital C^* -algebras is dually equivalent to the category of compact Hausdorff spaces (thus extending the Stone duality presented in Section 3.4) and consequently that the ultraproduct of a family of commutative unital C^* -algebras corresponds to the ultracoproduct of the corresponding compact Hausdorff spaces (as introduced in Section 6.10). In Section 14.4, we switch gears and study von Neumann algebras and an ultraproduct construction for the subclass of tracial von Neumann algebras. Finally, in Section 14.5, we discuss the famous Connes embedding problem, which asks if every tracial von Neumann algebra embeds into a tracial ultraproduct of matrix algebras, and present a proof of a theorem of Radulescu connecting this problem with the problem of determining if every group is hyperlinear, a problem raised at the end of Section 13.4.

14.1. Banach space ultraproducts

In this chapter, all vector spaces will be over the complex numbers. We first recall some basic definitions concerning Banach spaces.

Definition 14.1.1. If X is a vector space, a **norm** on X is a function $\|\cdot\| : X \rightarrow \mathbb{R}$ satisfying, for all $x, y \in X$ and all $\lambda \in \mathbb{C}$:

- (1) $\|x\| \geq 0$, with $\|x\| = 0$ if and only if $x = 0$;
- (2) $\|x + y\| \leq \|x\| + \|y\|$;
- (3) $\|\lambda x\| = |\lambda| \|x\|$.

A **normed space** is a pair $(X, \|\cdot\|)$, where X is a vector space and $\|\cdot\|$ is a norm on X .

If $\|\cdot\|$ is a norm on X , then we get an induced metric on X given by $d(x, y) := \|x - y\|$. The normed space $(X, \|\cdot\|)$ is said to be a **Banach space** if the induced metric on X is complete.

In what follows, we view normed spaces as pointed metric spaces (see Definition 11.1.1) with 0 as the distinguished basepoint.

Exercise 14.1.2. Suppose that $(X_i)_{i \in I}$ is a family of normed spaces and \mathcal{U} is an ultrafilter on I . Prove that the (metric) ultraproduct $\prod_{\mathcal{U}} X_i$ carries a natural normed space structure, where the vector space operations are those induced by the pointwise operations and where the norm is given by $\|[x]_{\mathcal{U}}\| := \lim_{\mathcal{U}} \|x(i)\|$.

In the next exercise, we give an alternate, algebraic characterization of an ultraproduct of normed spaces. First, recall that, given a family $(X_i)_{i \in I}$ of normed spaces, there is a natural norm on $\ell^\infty(X_i)$ given by $\|x\| := \sup_{i \in I} \|x(i)\|$. Second, given a normed space X and a closed subspace Y , we can turn the quotient vector space X/Y into a normed space by setting $\|x + Y\| := \inf_{y \in Y} \|x + y\|$.

Exercise 14.1.3. Suppose that $(X_i)_{i \in I}$ is a family of normed spaces and \mathcal{U} is an ultrafilter on I . Let $Y := \{x \in \ell^\infty(X_i) : \lim_{\mathcal{U}} x(i) = 0\}$. Prove that Y is a closed subspace of $\ell^\infty(X_i)$ and that $\ell^\infty(X_i)/Y$ is isomorphic (as a normed space) to $\prod_{\mathcal{U}} X_i$.

By Theorems 11.3.1 and 11.3.2, we have:

Theorem 14.1.4. *Suppose that $(X_i)_{i \in I}$ is a family of normed spaces and \mathcal{U} is an ultrafilter on I . If \mathcal{U} is countably incomplete or each X_i is a Banach space, then $\prod_{\mathcal{U}} X_i$ is a Banach space.*

Ultraproducts of normed spaces often possess similar properties to the constituent spaces. Here is one example of such a phenomenon. First, we say that a normed space X is **uniformly convex** if, for every $0 < \epsilon \leq 2$, there is $\delta > 0$ such that, for all $x, y \in X$ with $\|x\| = \|y\| = 1$, if $\|x - y\| \geq \epsilon$, then $\|\frac{x+y}{2}\| \leq 1 - \delta$.

Exercise 14.1.5. Prove that if X is a uniformly convex normed space, then so is any ultrapower $X^{\mathcal{U}}$ of X .

The following lemma will prove useful in the next section:

Lemma 14.1.6. *Suppose that $(X_i)_{i \in I}$ is a family of normed spaces and \mathcal{U} is an ultrafilter on I . Suppose further that $[x_1]_{\mathcal{U}}, \dots, [x_n]_{\mathcal{U}} \in \prod_{\mathcal{U}} X_i$ are*

linearly independent. Then for \mathcal{U} -almost all i , we have that $x_1(i), \dots, x_n(i) \in X_i$ are also linearly independent.

Proof. Suppose, toward a contradiction, that, for \mathcal{U} -almost all i , there are $c_1(i), \dots, c_n(i) \in \mathbb{C}$, not all equal to 0, such that $c_1(i)x_1(i) + \dots + c_n(i)x_n(i) = 0$. Take $j \in \{1, \dots, n\}$ such that, for \mathcal{U} -almost all i , we have $|c_k(i)| \leq |c_j(i)|$ for all $k \in \{1, \dots, n\}$. For ease of exposition, let us assume that $j = 1$. Then, setting $d_k(i) := \frac{c_k(i)}{c_1(i)}$, we have, for \mathcal{U} -almost all i , that $x_1(i) + d_2(i)x_2(i) + \dots + d_n(i)x_n(i) = 0$. Since $|d_k(i)| \leq 1$ for \mathcal{U} -almost all i and all $k = 2, \dots, n$, we have that $d_k := \lim_{\mathcal{U}} d_k(i)$ exists. We leave it to the reader to verify that $[x_1]_{\mathcal{U}} + d_2[x_2]_{\mathcal{U}} + \dots + d_n[x_n]_{\mathcal{U}} = 0$, yielding a contradiction. \square

Exercise 14.1.7.

- (1) If X is finite dimensional, prove that, for any ultrafilter \mathcal{U} , we have $X^{\mathcal{U}} \cong X$.
- (2) Suppose that $(X_i)_{i \in I}$ is a family of finite-dimensional normed spaces and that \mathcal{U} is an ultrafilter on I . Prove that $\prod_{\mathcal{U}} X_i$ is finite dimensional if and only if $\lim_{\mathcal{U}} \dim(X_i) < \infty$ and, in this case, $\dim(\prod_{\mathcal{U}} X_i) = \lim_{\mathcal{U}} \dim(X_i)$.

We now prove that ultrapowers commute with quotients. In what follows, given a closed subspace Y of a normed space X , a set I , and an element $x \in \ell^{\infty}(X)$, we slightly abuse notation by writing $x + Y$ for the element of $\ell^{\infty}(X/Y)$ given by $(x + Y)(i) := x(i) + Y$. Note that $x + Y$ does indeed belong to $\ell^{\infty}(X/Y)$ as $\|x(i) + Y\| \leq \|x(i)\|$ for all $i \in I$.

Proposition 14.1.8. *Let Y be a closed subspace of the Banach space X . Then the map $\phi : X^{\mathcal{U}}/Y^{\mathcal{U}} \rightarrow (X/Y)^{\mathcal{U}}$ given by $\phi([x]_{\mathcal{U}} + Y^{\mathcal{U}}) := [x + Y]_{\mathcal{U}}$ is an isomorphism of Banach spaces.*

Proof. Fix $[x]_{\mathcal{U}} \in X^{\mathcal{U}}$ and set $r := \|[x]_{\mathcal{U}} + Y^{\mathcal{U}}\|$ and $s := \|[x + Y]_{\mathcal{U}}\| = \lim_{\mathcal{U}} \|x(i) + Y\|$. Observe that

$$r = \inf_{[y]_{\mathcal{U}} \in Y^{\mathcal{U}}} \|[x]_{\mathcal{U}} + [y]_{\mathcal{U}}\| = \inf_{[y]_{\mathcal{U}} \in Y^{\mathcal{U}}} \lim_{\mathcal{U}} \|x(i) + y(i)\|.$$

Fix $\epsilon > 0$ and take $[y]_{\mathcal{U}} \in Y^{\mathcal{U}}$ such that $\lim_{\mathcal{U}} \|x(i) + y(i)\| \leq r + \epsilon$. Then we have that

$$s = \lim_{\mathcal{U}} \|x(i) + Y\| \leq \lim_{\mathcal{U}} \|x(i) + y(i)\| \leq r + \epsilon.$$

Letting $\epsilon \rightarrow 0$, we see that $\|[x + Y]_{\mathcal{U}}\| \leq r$. Conversely, fix $\epsilon > 0$ and take $A \in \mathcal{U}$ such that $\|x(i) + Y\| < s + \epsilon$ for all $i \in A$. For each $i \in A$, take $y(i) \in Y$ such that $\|x(i) + y(i)\| \leq \|x(i) + Y\| + \epsilon$. Set $y(i) := 0$ for $i \notin A$. Note that $y \in \ell^{\infty}(Y)$. Moreover, we have that $\lim_{\mathcal{U}} \|x(i) + y(i)\| \leq \lim_{\mathcal{U}} \|x(i) + Y\| + \epsilon \leq s + 2\epsilon$. It follows that $r \leq s + 2\epsilon$; letting $\epsilon \rightarrow 0$, we have $r \leq s$ and thus $r = s$.

This calculation shows that ϕ is a well-defined linear map that is isometric. It remains to show that ϕ is surjective. To see this, take $w \in \ell^\infty(X/Y)$. For each $i \in I$, take $x(i) \in X$ and $y(i) \in Y$ such that $w(i) = x(i) + Y$ and $\|w(i)\| = \|x(i) + y(i)\|$. Note then that $x + y \in \ell^\infty(X)$, whence $[x + y]_{\mathcal{U}} \in X^{\mathcal{U}}$ and that $\phi([x + y]_{\mathcal{U}}) = [w]_{\mathcal{U}}$. \square

An important class of Banach spaces arises from studying Hilbert spaces. We first recall the relevant definitions.

Definition 14.1.9. An **inner product** on a vector space X is a map $\langle \cdot, \cdot \rangle : X \times X \rightarrow \mathbb{C}$ satisfying, for all $x, y, z \in X$ and $\lambda \in \mathbb{C}$:

- (1) $\langle x + y, z \rangle = \langle x, z \rangle + \langle y, z \rangle$,
- (2) $\langle \lambda x, y \rangle = \lambda \langle x, y \rangle$, and
- (3) $\langle y, x \rangle = \overline{\langle x, y \rangle}$.

An **inner product space** is a pair $(X, \langle \cdot, \cdot \rangle)$, where X is a vector space and $\langle \cdot, \cdot \rangle$ is an inner product on X . Any inner product $\langle \cdot, \cdot \rangle$ on X induces a norm $\|\cdot\|$ on X defined by $\|x\| := \sqrt{\langle x, x \rangle}$. An inner product space $(X, \langle \cdot, \cdot \rangle)$ is called a **Hilbert space** if the induced normed space is a Banach space (that is, if the metric induced by the norm induced by the inner product is complete).

Exercise 14.1.10. Suppose that $(X_i)_{i \in I}$ is a family of inner product spaces and \mathcal{U} is an ultrafilter on I . Prove that there is a natural inner product on $\prod_{\mathcal{U}} X_i$ given by $\langle [x]_{\mathcal{U}}, [y]_{\mathcal{U}} \rangle = \lim_{\mathcal{U}} \langle x(i), y(i) \rangle$. Consequently, if \mathcal{U} is countably incomplete or each X_i is a Hilbert space, then $\prod_{\mathcal{U}} X_i$ is also a Hilbert space.

Finally, we discuss the notion of bounded operators. The following fact is easy and quite standard:

Fact 14.1.11. Suppose that X and Y are normed spaces and $T : X \rightarrow Y$ is a linear transformation. Then the following are equivalent:

- (1) T is continuous at some point in X .
- (2) T is continuous.
- (3) T is uniformly continuous.
- (4) T is **bounded**, meaning that it maps bounded sets to bounded sets.

Definition 14.1.12. Suppose that $T : X \rightarrow Y$ is a bounded linear transformation. The **operator norm of T** is the quantity $\|T\| := \sup\{\|T(x)\| : x \in X, \|x\| \leq 1\}$.

It is easy to check that $\|T\| = \sup\{\|T(x)\| : x \in X, \|x\| = 1\}$ and that, whether using the unit ball or the unit sphere to calculate $\|T\|$, both supremums are actually achieved. Another characterization of $\|T\|$ is that it is the smallest nonnegative real number M such that $\|T(x)\| \leq M\|x\|$ for all $x \in X$.

One lets $\mathcal{B}(X, Y)$ denote the set of bounded linear transformations from X to Y . When $X = Y$, we write $\mathcal{B}(X)$ instead of $\mathcal{B}(X, X)$. It is fairly easy to verify that $\mathcal{B}(X, Y)$ is a subspace of the vector space of all linear transformations from X to Y . Moreover, $\mathcal{B}(X, Y)$ is a normed space under the operator norm. If Y is a Banach space, then it can be shown that $\mathcal{B}(X, Y)$ is also a Banach space.

Suppose that X, Y , and Z are all normed spaces and $T \in \mathcal{B}(X, Y)$ and $S \in \mathcal{B}(Y, Z)$. Then it is fairly easy to check that $S \circ T \in \mathcal{B}(X, Z)$ and that $\|S \circ T\| \leq \|S\| \cdot \|T\|$. In particular, $\mathcal{B}(X)$ is closed under composition, giving it the structure of a unital normed algebra in the sense of the following definition:

Definition 14.1.13. A **normed algebra** is a normed space A equipped with a binary operation \cdot for which the following hold for all $x, y, z \in A$ and $\lambda \in \mathbb{C}$:

- (1) $(x \cdot y) \cdot z = x \cdot (y \cdot z)$;
- (2) $(x + y) \cdot z = (x \cdot z) + (y \cdot z)$ and $x \cdot (y + z) = (x \cdot y) + (x \cdot z)$;
- (3) $\lambda(x \cdot y) = (\lambda x) \cdot y = x \cdot (\lambda y)$;
- (4) $\|x \cdot y\| \leq \|x\| \cdot \|y\|$.

Furthermore, the normed algebra is called **unital** if there is $1 \in A$ such that $1 \cdot x = x \cdot 1 = x$ for all $x \in A$. A (unital) normed algebra for which the underlying normed space is a Banach space is called a **(unital) Banach algebra**.

Thus, to reiterate, we have:

Proposition 14.1.14. *For any normed space X , $\mathcal{B}(X)$ is a unital normed algebra. If X is a Banach space, then $\mathcal{B}(X)$ is a unital Banach algebra.*

Exercise 14.1.15. Suppose that $(A_i)_{i \in I}$ is a family of (unital) Banach algebras and \mathcal{U} is an ultrafilter on I . Prove that $\prod_{\mathcal{U}} A_i$ is also a (unital) Banach algebra when equipped with the pointwise operations.

Exercise 14.1.16. Suppose that $(X_i)_{i \in I}$ and $(Y_i)_{i \in I}$ are families of normed spaces and \mathcal{U} is an ultrafilter on I . Prove that there is a linear isometric embedding $\iota : \prod_{\mathcal{U}} \mathcal{B}(X_i, Y_i) \rightarrow \mathcal{B}(\prod_{\mathcal{U}} X_i, \prod_{\mathcal{U}} Y_i)$ given by

$$\iota([T]_{\mathcal{U}})([x]_{\mathcal{U}}) := [T(i)x(i)]_{\mathcal{U}}.$$

14.2. Applications to local geometry of Banach spaces

In this section, we give a few applications of Banach space ultraproducts to the study of the local geometry of Banach spaces.

Definition 14.2.1. An invertible linear map $T : X \rightarrow Y$ between normed spaces is an ϵ -**isomorphism** if for all $x \in X$, we have

$$(1 - \epsilon)\|x\| \leq \|T(x)\| \leq (1 + \epsilon)\|x\|.$$

We say that X is ϵ -**isomorphic** to Y , written $X \cong_\epsilon Y$, if there is an ϵ -isomorphism $X \rightarrow Y$.

Note that \cong_ϵ is not a symmetric relation. However, if $X \cong_\epsilon Y$, then $Y \cong_\delta X$ for $\delta = \frac{\epsilon}{1-\epsilon}$.

Exercise 14.2.2. Prove that an invertible linear map $T : X \rightarrow Y$ between Banach spaces is an ϵ -isomorphism if and only if: for all $x \in X$ with $\|x\| = 1$, we have $1 - \epsilon \leq \|T(x)\| \leq 1 + \epsilon$.

Proposition 14.2.3. *If F is a finite-dimensional subspace of $\prod_{\mathcal{U}} X_i$, then for each $\epsilon > 0$, there are \mathcal{U} -many i for which there are finite-dimensional subspaces F_i of X_i satisfying $F \cong_\epsilon F_i$.*

Proof. Take a basis $[x_1]_{\mathcal{U}}, \dots, [x_n]_{\mathcal{U}}$ of F and consider the linear mappings $\phi_i : F \rightarrow X_i$ that send $[x_j]_{\mathcal{U}}$ to $x_j(i)$ for each $j = 1, \dots, n$. Set $F_i := \phi_i(F)$. By Lemma 14.1.6, F_i has dimension n for \mathcal{U} -almost all i . We leave it to the reader to check that $\phi_i : F \rightarrow F_i$ is an ϵ -isomorphism for \mathcal{U} -almost all i . \square

Exercise 14.2.4. Verify the end of the proof of the previous proposition.

The following is a central notion in the local geometry of Banach spaces:

Definition 14.2.5. We say that a Banach space Y is **finitely representable** in a Banach space X if, for every finite-dimensional subspace F of Y and every $\epsilon > 0$, there is a (finite-dimensional) subspace F' of X such that $F \cong_\epsilon F'$.

Proposition 14.2.3 immediately yields:

Corollary 14.2.6. *For any ultrafilter \mathcal{U} , $X^{\mathcal{U}}$ is finitely representable in X .*

It is clear from the definitions that if Y is finitely representable in X and Z is a closed subspace of Y , then Z is also finitely representable in X . It follows that every closed subspace of $X^{\mathcal{U}}$ is finitely representable in X . We now prove that the converse holds. In what follows, by an ϵ -**isomorphic embedding** $T : X \rightarrow Y$, we mean an injective linear map that is an ϵ -isomorphism between X and $T(X)$.

Theorem 14.2.7. *If X and Y are Banach spaces, then Y is finitely representable in X if and only if there is a linear isometric embedding of Y into $X^{\mathcal{U}}$ for some ultrafilter \mathcal{U} .*

Proof. Suppose that Y is finitely representable in X . For each finite-dimensional subspace F of Y and each $\epsilon > 0$, let $T_{(F,\epsilon)} : F \rightarrow X$ be an ϵ -isomorphic embedding. Let I denote the set of all such pairs (F, ϵ) and let \mathcal{U} be an ultrafilter on I containing all sets $A_{(F_0, \epsilon_0)} := \{(F, \epsilon) \in I : F_0 \subseteq F, \epsilon \leq \epsilon_0\}$. For $x \in Y$, let $\psi_x \in X^I$ be defined by $\psi_x(F, \epsilon) = T_{(F,\epsilon)}(x)$ if $x \in F$ and $\psi_x(F, \epsilon) = 0$ otherwise.

We claim that the map $T : Y \rightarrow X^{\mathcal{U}}$ given by $T(x) := [\psi_x]_{\mathcal{U}}$ is the desired linear isometric embedding. We first show that it is linear. Given $x, y \in Y$, set $F_0 := \text{span}(x, y)$. Then for $(F, \epsilon) \in A_{(F_0, 1)}$, we have $\psi_x(F, \epsilon) = T_{(F,\epsilon)}(x)$ and $\psi_y(F, \epsilon) = T_{(F,\epsilon)}(y)$; since $A_{(F_0, 1)} \in \mathcal{U}$ and each $T_{(F,\epsilon)}$ is linear, we see that T is linear.

We finish by showing that T is isometric. Take $x \in Y$ with $\|x\| = 1$ and $\epsilon > 0$. Take $A \in \mathcal{U}$ such that $|\|T(x)\| - \|\psi_x(F, \delta)\|| < \epsilon$ for all $(F, \delta) \in A$. Suppose now that $(F, \delta) \in A_{(\text{span}(x), 1)}$; then $\psi_x(F, \delta) = T_{(F,\delta)}(x)$, whence $(1 - \epsilon) \leq \|\psi_x(F, \delta)\| \leq (1 + \epsilon)$. By considering $(F, \delta) \in A \cap A_{(\text{span}(x), 1)}$, it follows that $(1 - 2\epsilon) \leq \|T(x)\| \leq (1 + 2\epsilon)$. Since ϵ was arbitrary, we get the desired result. \square

Exercise 14.2.8. Suppose that Y is a separable Banach space that is finitely representable in the Banach space X . Prove that, for any $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$, Y linearly isometrically embeds into $X^{\mathcal{U}}$.

Exercise 14.2.9. Prove that any Banach space finitely representable in a Hilbert space is also a Hilbert space, that is, is the Banach space naturally obtained from a Hilbert space. (Try doing this directly from the definition as well to gain an appreciation for Theorem 14.2.7.)

We apply Theorem 14.2.7 to some issues around reflexivity for Banach spaces. We first recall the requisite definitions.

Given any Banach space X , set $X^* := \mathcal{B}(X, \mathbb{C})$, which is the set of all continuous linear functions $\varphi : X \rightarrow \mathbb{C}$. By our discussion in the previous section, X^* is itself a Banach space when equipped with the operator norm, called the **dual space** of X .

One can consider the dual X^{**} of the dual space X^* , which is once again a Banach space. We will use letters like Φ to denote elements of X^{**} . Given $x \in X$, we have an element $\Phi_x \in X^{**}$ given by $\Phi_x(\varphi) := \varphi(x)$. This gives a linear map $\Delta : X \rightarrow X^{**}$ defined by $\Delta(x) := \Phi_x$. Moreover, $\|\Phi_x\| = \sup_{\|\varphi\| \leq 1} \|\varphi(x)\| = \|x\|$, where the second equality follows from the

Hahn-Banach theorem (see [35, III.6]). It follows that Δ is an isometric embedding.

Definition 14.2.10. Let X be a Banach space and let $\Delta : X \rightarrow X^{**}$ be as above.

- (1) We say that X is **reflexive** if Δ is surjective.
- (2) We say that X is **super-reflexive** if every space that is finitely representable in X is reflexive.

Not every Banach space is reflexive (for example, see [35, III.11]). However, many familiar spaces are (super-) reflexive:

Exercise 14.2.11.

- (1) Prove that any finite-dimensional Banach space is super-reflexive.
- (2) Prove that any Hilbert space is super-reflexive.

Fact 14.2.12. Let X be a Banach space and Y a closed subspace. Then X is reflexive if and only if both Y and X/Y are reflexive.

Proof. See, for example, [35, V.4]. □

Corollary 14.2.13. *If X is a Banach space, then X is super-reflexive if and only if every ultrapower $X^{\mathcal{U}}$ of X is reflexive.*

Proof. The forward direction follows from the fact that $X^{\mathcal{U}}$ is finitely representable in X . For the converse, take Y finitely representable in X and embed Y as a closed subspace of some ultrapower of X . By hypothesis, this ultrapower is reflexive, whence so is Y by Fact 14.2.12. □

Example 14.2.14. The *Milman-Pettis theorem* states that uniformly convex Banach spaces are reflexive (see [29, Theorem 1.17]). By Exercise 14.1.5 and Corollary 14.2.13, they are in fact super-reflexive.

We can now prove the “super” version of Fact 14.2.12:

Corollary 14.2.15. *Let X be a Banach space and Y a closed subspace. Then X is super-reflexive if and only if both Y and X/Y are super-reflexive.*

Proof. To prove the backward direction, assume that both Y and X/Y are super-reflexive and fix an ultrafilter \mathcal{U} . By Corollary 14.2.13, $Y^{\mathcal{U}}$ is reflexive. By Proposition 14.1.8 and Corollary 14.2.13, $(X/Y)^{\mathcal{U}} \cong X^{\mathcal{U}}/Y^{\mathcal{U}}$ is reflexive. By Fact 14.2.12, $X^{\mathcal{U}}$ is reflexive. Since \mathcal{U} was an arbitrary ultrafilter, Corollary 14.2.13 implies that X is super-reflexive.

We leave the proof of the forward direction as an exercise. □

Exercise 14.2.16. Prove the forward direction of the previous corollary.

For not necessarily reflexive spaces, we now show that X^{**} is finitely representable in X . In fact, we show something stronger. First, a definition.

Definition 14.2.17. If Y is a closed subspace of a Banach space X , we say that Y is a **complemented subspace** of X if there is a closed subspace Z of X such that $X = Y \oplus Z$.

In a general Banach space, a closed subspace need not be complemented. For a general discussion on this topic, see [35, III.13]. In many familiar Banach spaces, all closed subspaces are complemented:

Exercise 14.2.18.

- (1) Prove that every subspace of a finite-dimensional Banach space is complemented.
- (2) Prove that every closed subspace of a Hilbert space is complemented.

The following is an alternate formulation of a subspace of a Banach space being complemented:

Exercise 14.2.19. If X is a Banach space and Y is a closed subspace of X , then Y is a complemented subspace of X if and only if there is a continuous linear map $P : X \rightarrow X$ such that $P(X) = Y$ and $P \circ P = P$.

Here is the promised connection between biduals and ultrapowers for general Banach spaces:

Theorem 14.2.20. For any Banach space X , there is an ultrafilter \mathcal{U} and a linear isometric embedding $T : X^{**} \rightarrow X^{\mathcal{U}}$ such that:

- (1) for all $x \in X$, $T(\Phi_x)$ is the diagonal image of x in $X^{\mathcal{U}}$.
- (2) $T(X^{**})$ is a complemented subspace of $X^{\mathcal{U}}$.

The proof of Theorem 14.2.20 uses the so-called *Principle of Local Reflexivity* [131]:

Fact 14.2.21. For every pair of finite-dimensional subspaces $E \subseteq X^{**}$ and $F \subseteq X^*$ and every $\epsilon > 0$, there is an ϵ -isomorphic embedding $T_{E,F,\epsilon} : E \rightarrow X$ such that:

- $T_{E,F,\epsilon}(\Phi_x) = x$ for all $x \in X$ such that $\Phi_x \in E$.
- For all $\Phi \in E$ and all $\varphi \in F$, we have $\varphi(T_{E,F,\epsilon}(\Phi)) = \Phi(\varphi)$.

Before discussing the proof, we need a quick digression on the weak*-topology:

Definition 14.2.22. Let X be a normed space. The **weak*-topology** on X^* is the topology induced on X^* by viewing it as a subset of the space \mathbb{C}^X , which is equipped with the product topology.

Alternatively, the weak*-topology on X^* is the weakest topology on X^* making each evaluation maps $\varphi \mapsto \varphi(x) : X^* \rightarrow \mathbb{C}$ (for $x \in X$) continuous.

Exercise 14.2.23. Suppose that X is a normed space, $(\varphi_i)_{i \in I}$ is a family from X^* , φ is another element in X^* , and \mathcal{U} is an ultrafilter on I . Prove that $\lim_{\mathcal{U}} \varphi_i = \varphi$ in the weak*-topology if and only if $\lim_{\mathcal{U}} (\varphi_i(x)) = \varphi(x)$ for all $x \in X$, where the latter ultralimit is calculated in \mathbb{C} .

Exercise 14.2.24. Suppose that X is a normed space. Prove that every (operator-norm) bounded subset of X^* that is weak*-closed is actually weak*-compact.

In what follows, we let I denote the set of triples (E, F, ϵ) as in the statement of Fact 14.2.21 and we fix an ultrafilter \mathcal{U} on I containing all elements

$$A_{(E_0, F_0, \epsilon_0)} := \{(E, F, \epsilon) \in I : E_0 \subseteq E, F_0 \subseteq F, \epsilon \leq \epsilon_0\}.$$

For each $\Phi \in X^{**}$, we define an element $x_\Phi \in X^I$ by $x_\Phi(E, F, \epsilon) = T_{E, F, \epsilon}(\Phi)$ if $\Phi \in E$ and $x_\Phi(E, F, \epsilon) = 0$ otherwise.

Lemma 14.2.25. *For any $\Phi \in X^{**}$, we have $\lim_{(E, F, \epsilon) \rightarrow \mathcal{U}} \Phi_{x_\Phi(E, F, \epsilon)} = \Phi$, where the ultralimit is calculated in the weak* topology on X^{**} .*

Proof. To prove the lemma, fix $\varphi \in X^*$. By Exercise 14.2.23, it suffices to show that

$$\left(\lim_{\mathcal{U}} \Phi_{x_\Phi(E, F, \epsilon)}\right)(\varphi) = \Phi(\varphi).$$

Set $E_0 := \text{span}(\Phi)$ and $F_0 := \text{span}(\varphi)$. Suppose that $(E, F, \epsilon) \in A_{(E_0, F_0, 1)}$. Then $x_\Phi(E, F, \epsilon) = T_{E, F, \epsilon}(\Phi)$ and $\Phi_{x_\Phi(E, F, \epsilon)}(\varphi) = \varphi(T_{E, F, \epsilon}(\Phi)) = \Phi(\varphi)$. The proof of the lemma is now complete. \square

Proof of Theorem 14.2.20. The desired linear isometric embedding is given by $T(\Phi) := [x_\Phi]_{\mathcal{U}}$. As in the proof of Theorem 14.2.7, this map is a linear isometric embedding. Moreover, if $x \in X$, setting $E_0 := \text{span}(\Phi_x)$, then whenever $(E, F, \epsilon) \in A_{(E_0, \{0\}, 1)}$, we have $x_\Phi(E, F, \epsilon) = T_{E, F, \epsilon}(\Phi_x) = x$, whence $T(\Phi_x)$ coincides with the diagonal image of x in $X^{\mathcal{U}}$.

In order to prove that $T(X^{**})$ is a complemented subspace of $X^{\mathcal{U}}$, by Exercise 14.2.19, it suffices to find a linear map $P : X^{\mathcal{U}} \rightarrow X^{\mathcal{U}}$ that is the identity on $T(X^{**})$ and for which $\|P\| \leq 1$. We define this map as $P([x]_{\mathcal{U}}) := T(\lim_{\mathcal{U}} \Phi_x)$, where the ultralimit is calculated in the weak* topology on X^{**} , which exists by Exercise 14.2.24.

We first show that P is the identity on $T(X^{**})$. Toward that end, fix $\Phi \in X^{**}$. We then have

$$P(T(\Phi)) = P([x_\Phi]_{\mathcal{U}}) = T(\lim_{\mathcal{U}} \Phi_{x_\Phi(E, F, \epsilon)}) = T(\Phi),$$

where the last equality follows from Lemma 14.2.25.

To finish the proof, we show that $\|P\| \leq 1$. Toward that end, fix $[x]_{\mathcal{U}} \in X^{\mathcal{U}}$. Using that T is an isometric embedding, we then have

$$\begin{aligned} \|P([x]_{\mathcal{U}})\| &= \|\lim_{\mathcal{U}} \Phi_{x(i)}\| = \sup_{\|\varphi\| \leq 1} \|(\lim_{\mathcal{U}} \Phi_{x(i)})(\varphi)\| \\ &= \sup_{\|\varphi\| \leq 1} \|\lim_{\mathcal{U}} \varphi(x(i))\| = \sup_{\|\varphi\| \leq 1} \lim_{\mathcal{U}} \|\varphi(x(i))\| \\ &\leq \sup_{\|\varphi\| \leq 1} \lim_{\mathcal{U}} \|x(i)\| = \|[x]_{\mathcal{U}}\|. \quad \square \end{aligned}$$

Corollary 14.2.26. *If a class of Banach spaces is closed under ultraproducts and complemented subspaces, then it is closed under biduals.*

14.3. Commutative C^* -algebras and ultracoproducts of compact spaces

In the remainder of this chapter, we briefly sketch some appearances of Banach space ultraproducts in the area of functional analysis known as **operator algebras**. Ultraproduct techniques play a pivotal role in the modern study of operator algebras, but since the mathematics behind the scenes is beyond the scope of this text, we merely discuss a few matters along these lines. In this section, we discuss a class of operator algebras known as C^* -algebras.

Definition 14.3.1. A **Banach $*$ -algebra** is a Banach algebra A equipped with a map, usually denoted $*$: $A \rightarrow A$, satisfying, for all $x, y \in A$ and $\lambda \in \mathbb{C}$:

- (1) $(x^*)^* = x$
- (2) $(x + y)^* = x^* + y^*$
- (3) $(\lambda x)^* = \overline{\lambda}x^*$
- (4) $(xy)^* = y^*x^*$
- (5) $\|x^*\| = \|x\|$.

It is easy to verify that if A is a unital Banach $*$ -algebra, then $1^* = 1$.

Exercise 14.3.2. Suppose that $(A_i)_{i \in I}$ is a family of (unital) Banach $*$ -algebras and \mathcal{U} is an ultrafilter on I . Prove that $\prod_{\mathcal{U}} A_i$ is also a (unital) Banach $*$ -algebra when equipped with the pointwise operations.

Example 14.3.3. If H is a Hilbert space and $T \in \mathcal{B}(H)$, then there is a unique $T^* \in \mathcal{B}(H)$ satisfying $\langle Tx, y \rangle = \langle x, T^*y \rangle$ for all $x, y \in H$. T^* is often referred to as the **adjoint** of T . It is then routine to check that $\mathcal{B}(H)$ is a unital Banach $*$ -algebra. (The verification of axiom (5) will appear in the proof of the next lemma.)

The Banach $*$ -algebra $\mathcal{B}(H)$ in the previous example also satisfies a further property, which turns out to be the crucial axiom that turns a Banach $*$ -algebra into a C^* -algebra:

Lemma 14.3.4. *Suppose that H is a Hilbert space. Then for all $T \in \mathcal{B}(H)$, we have $\|T^*T\| = \|T\|^2$.*

Proof. Fix $x \in H$ with $\|x\| \leq 1$. We then have

$$\|Tx\|^2 = \langle Tx, Tx \rangle = \langle T^*Tx, x \rangle \leq \|T^*Tx\| \cdot \|x\| \leq \|T^*T\| \leq \|T^*\| \cdot \|T\|.$$

It follows that $\|T\|^2 \leq \|T^*T\| \leq \|T^*\| \cdot \|T\|$. In particular, $\|T\| \leq \|T^*\|$. Since $T^{**} = T$, applying the previous inequality to T^* yields that $\|T^*\| \leq \|T\|$. Thus, $\|T\| = \|T^*\|$, whence $\|T\|^2 = \|T^*T\|$, as desired. \square

Definition 14.3.5. A Banach $*$ -algebra A is called a **C^* -algebra** if it further satisfies the **C^* -identity** $\|x^*x\| = \|x\|^2$ for all $x \in A$.

Thus, $\mathcal{B}(H)$ is a unital C^* -algebra for any Hilbert space H . Here is another important example:

Example 14.3.6. Let X be a compact space and let $C(X)$ denote the set of continuous functions $f : X \rightarrow \mathbb{C}$. We can equip $C(X)$ with the structure of a $*$ -algebra by considering pointwise addition, multiplication, and scalar multiplication, and by setting, for $f \in C(X)$, $f^*(x) := \overline{f(x)}$. Moreover, we equip $C(X)$ with the norm $\|f\| := \sup_{x \in X} |f(x)|$, which exists since X is compact and f is continuous. It is fairly clear that $C(X)$ is then a unital Banach $*$ -algebra, the identity being the function that is constantly 1. Moreover, $C(X)$ is actually a C^* -algebra, for

$$\begin{aligned} \|f^*f\| &= \sup_{x \in X} |(f^*f)(x)| = \sup_{x \in X} |\overline{f(x)}f(x)| \\ &= \sup_{x \in X} |f(x)|^2 = \left(\sup_{x \in X} |f(x)| \right)^2 = \|f\|^2. \end{aligned}$$

We note that $C(X)$ as in the previous example is a **commutative** C^* -algebra, meaning that the multiplication operation is commutative. It is actually the case that *all* commutative unital C^* -algebras are isomorphic to a unique algebra of the form $C(X)$ as in the previous example. In fact, even more is true, namely that the above functor C is actually one half of a duality between the categories of compact Hausdorff spaces and unital

commutative C^* -algebras that “extends” Stone duality in a way that we will now explain. The discussion that follows is far from complete and we freely omit many details as we are merely trying to outline the general idea.

We first remark that the category of C^* -algebras has as its objects C^* -algebras and with its morphisms $*$ -homomorphisms, which are linear maps which also preserve the multiplication and $*$ -operations. It is worth noting that such maps are automatically contractive, that is, they have operator norm at most 1. (This is peculiar to C^* -algebras and does not necessarily hold for $*$ -homomorphisms between arbitrary Banach $*$ -algebras.) The category of commutative unital C^* -algebras forms a subcategory of the category of C^* -algebras where the morphisms are required to be unital, that is, they must map the identity to the identity. (Thus, the category of commutative unital C^* -algebras is not a full subcategory of the category of commutative C^* -algebras.)

Next note that C is indeed a contravariant functor between the category of compact Hausdorff¹ spaces and the category of commutative unital C^* -algebras, for if $f : X \rightarrow X'$ is continuous, we get a $*$ -algebra homomorphism $C(f) : C(X') \rightarrow C(X)$ given by $C(f)(g) = g \circ f$

We now ponder how we might get a functor going in the other direction. We take our cue from Stone duality, namely we consider the space of morphisms from our given commutative unital C^* -algebra to the “simplest” such algebra \mathbb{C} ; this space of morphisms can be viewed as $C(*)$, where $*$ is a one-point space. More precisely, we consider the functor Spec from commutative unital C^* -algebras to compact Hausdorff spaces given by taking $\text{Spec}(A)$ to be the set of unital $*$ -homomorphisms $A \rightarrow \mathbb{C}$. (The “ultrafilter” perspective in this context is to view $\text{Spec}(A)$ as the set of maximal ideals of A .) It is fairly straightforward to see that $\text{Spec}(A)$ is a weak*-closed subset of A^* . Since elements of $\text{Spec}(A)$ have operator norm at most 1 (in fact exactly 1), $\text{Spec}(A)$ is weak*-compact by Exercise 14.2.24. Thus $\text{Spec}(A)$ is indeed an object in the category of compact Hausdorff spaces. Here is how Spec acts on morphisms: If $h : A \rightarrow A'$ is a unital $*$ -homomorphism, then we set $\text{Spec}(h) : \text{Spec}(A') \rightarrow \text{Spec}(A)$ to be given by $\text{Spec}(h)(i) := i \circ h$. It is straightforward to verify that $\text{Spec}(h)$ is indeed weak*-continuous.

Theorem 14.3.7 (Gelfand duality). *The functors C and Spec witness that the category of compact Hausdorff spaces and the category of commutative unital C^* -algebras are dually equivalent.*

We will not discuss the proof in too much detail, but we at least want to explain how to witness that the compositions of the two functors are naturally isomorphic to the respective identity functors. On the one hand, given

¹We require Hausdorff here so as to be able to recover X from $C(X)$.

a compact Hausdorff space X , one can show that $\text{Spec}(C(X))$ is homeomorphic to X by showing that, given any homomorphism $h : C(X) \rightarrow \mathbb{C}$, there is a unique $x_h \in X$ such that $h(f) = f(x_h)$ for all $f \in C(X)$. Thus, the mapping $h \mapsto x_h : \text{Spec}(C(X)) \rightarrow X$ is the desired homeomorphism. On the other hand, one can show that $C(\text{Spec}(A))$ is isomorphic to A by showing that every continuous function $f : \text{Spec}(A) \rightarrow \mathbb{C}$ is of the form $f(h) = h(a_f)$ for a unique $a_f \in A$. Consequently, the map $f \mapsto a_f : C(\text{Spec}(A)) \rightarrow A$ is the desired isomorphism. Of course, one must show that everything here is “natural”, but this is all fairly routine.

We now briefly discuss why Gelfand duality is an “extension” of Stone duality (see Section 3.4). Let X be a Stone space. Then Stone duality associates to X a Boolean algebra $\text{Cl}(X)$, the collection of clopen subsets of X , while Gelfand duality associates to it a commutative unital C^* -algebra $C(X)$. How do we see that these algebraic associations “agree”? It turns out that the key is to identify $C(X)$ with its Boolean algebra of projections, which we now define.

Definition 14.3.8. A **projection** in a C^* -algebra A is an element p such that $p^2 = p^* = p$. We let $P(A)$ denote the set of projections in A .

The terminology comes from the case that the C^* -algebra is $\mathcal{B}(H)$ for some Hilbert space H . Recall that in this case, given any closed subspace K of H and any $x \in H$, one can write $x = y + z$, with $y \in K$ and $z \in K^\perp$, where K^\perp denotes the closed subspace of H consisting of elements orthogonal to all elements of K ; moreover, such a decomposition is unique.

Definition 14.3.9. If H is a Hilbert space and K is a closed subspace of H , then the linear map $P = P_K : H \rightarrow H$ for which $P(x)$ is the unique element of K such that $x - P(x) \in K^\perp$ for all $x \in H$, is called the **orthogonal projection of H onto K** .

It is clear that an orthogonal projection map as in the previous definition is a projection in the C^* -algebra $\mathcal{B}(H)$. Conversely, it can be shown that any projection in $\mathcal{B}(H)$ is an orthogonal projection onto some closed subspace.

Note that in any unital C^* -algebra, 0 and 1 are projections.

Exercise 14.3.10. Given any compact space X , prove that the projections in $C(X)$ are precisely those continuous functions $f : X \rightarrow \mathbb{C}$ such that $f(x) \in \{0, 1\}$ for all $x \in X$.

Exercise 14.3.11. Given any unital C^* -algebra A , prove that $P(A)$ is a Boolean algebra under the operations $p \vee q := p + q - pq$ and $p \wedge q := pq$. Moreover, prove that P is actually a functor from the category of unital C^* -algebras to the category of Boolean algebras.

Until further notice, C denotes the restriction of the functor above to the subcategory of Stone spaces. We now have two functors from Stone spaces to Boolean algebras, namely Cl and $\text{P} \circ C$.

Theorem 14.3.12. *Cl and $\text{P} \circ C$ are naturally isomorphic functors.*

Proof sketch. Given a Stone space X , define $\eta_X : \text{Cl}(X) \rightarrow \text{P}(C(X))$ by $\eta_X(U)$ is the function which is 1 on U and 0 on U^c . The collection of η_X 's can be verified to witness the conclusion of the theorem. \square

We next want to give a C^* -algebraic characterization of the commutative unital C^* -algebras that arise as $C(X)$ for X a Stone space. Here is the relevant definition:

Definition 14.3.13. A C^* -algebra A is said to have **real rank 0** if the closed linear span of $\text{P}(A)$ in A is equal to all of A .

Fact 14.3.14. If A is a commutative unital C^* -algebra, then A has real rank 0 if and only if $\text{Spec}(A)$ is a Stone space.

Until further notice Spec denotes the restriction of the functor above to the subcategory of real rank 0 commutative unital C^* -algebras. We now have two functors from the aforementioned category to the category of Stone spaces, namely Spec and $S \circ \text{P}$. (Recall that if A is a Boolean algebra, then $S(A)$ is the Stone space of all ultrafilters on A .)

Theorem 14.3.15. *Spec and $S \circ \text{P}$ are naturally isomorphic functors.*

Proof. Given a real rank 0 commutative unital C^* -algebra A , define $\epsilon_A : \text{Spec}(A) \rightarrow S(\text{P}(A))$ by $\epsilon_A(h) := h \upharpoonright \text{P}(A)$. Note that $h \upharpoonright \text{P}(A)$ does indeed take values in $\{0, 1\}$ and the restriction is a Boolean algebra homomorphism. The collection of ϵ_A 's can be verified to witness the conclusion of the theorem. \square

From the above, one can deduce the following:

Theorem 14.3.16. *$\text{Cl} \circ \text{Spec}$ and $C \circ S$ witness that the category of real rank 0 commutative unital C^* -algebras and the category of Boolean algebras are equivalent categories.*

It is with all of the above natural identifications that Gelfand duality restricted to the categories of Stone spaces and real rank 0 commutative unital C^* -algebras “is” just Stone duality.

We now leave the abstract nonsense and return to ultraproduct related matters.

Exercise 14.3.17. If $(A_i)_{i \in I}$ is a family of C^* -algebras and \mathcal{U} is an ultrafilter on I , prove that the Banach space ultraproduct $\prod_{\mathcal{U}} A_i$ is also a C^* -algebra with the pointwise operations. If \mathcal{U} -almost all of the factors are commutative (resp., unital), prove that $\prod_{\mathcal{U}} A_i$ is also commutative (resp., unital).

We recall from Example 6.10.8 that the category of compact Hausdorff spaces admits ultraproducts. Combining this with Exercise 14.3.17 and Exercise 6.10.6, we arrive at:

Theorem 14.3.18. *Suppose that $(A_i)_{i \in I}$ is a family of commutative unital C^* -algebras and $(X_i)_{i \in I}$ is a family of compact Hausdorff spaces. Further suppose that \mathcal{U} is an ultrafilter on I . Then*

$$\text{Spec} \left(\prod_{\mathcal{U}} A_i \right) \cong \prod_{\mathcal{U}} \text{Spec}(A_i) \quad \text{and} \quad C \left(\prod_{\mathcal{U}} X_i \right) \cong \prod_{\mathcal{U}} C(X_i).$$

Here is an example of how this perspective can be useful when trying to identify topological properties of ultraproducts.

Exercise 14.3.19.

- (1) If X is a compact Hausdorff space, prove that X is connected if and only if $P(C(X)) = \{0, 1\}$.
- (2) Suppose that $(X_i)_{i \in I}$ is a family of compact Hausdorff spaces and \mathcal{U} is an ultrafilter on I . Suppose further that $[f]_{\mathcal{U}} \in \prod_{\mathcal{U}} C(X_i)$ is a projection. Prove that there are projections $g(i) \in C(X_i)$ such that $[f]_{\mathcal{U}} = [g]_{\mathcal{U}}$.
- (3) Suppose that $(X_i)_{i \in I}$ is a family of compact Hausdorff spaces and \mathcal{U} is an ultrafilter on I . Suppose further that each X_i is connected. Prove that $\prod_{\mathcal{U}} X_i$ is also connected.

14.4. The tracial ultraproduct construction

In this section, we discuss a different class of operator algebras, the class of von Neumann algebras, and an ultraproduct construction appropriate for studying a large class of them. In order to introduce von Neumann algebras, it behooves us to consider a different perspective on C^* -algebras.

Definition 14.4.1. A **concrete C^* -algebra** is a $*$ -subalgebra of $\mathcal{B}(H)$ (for H some Hilbert space) that is closed in the operator norm topology.

It is clear that a concrete C^* -algebra is a C^* -algebra in the sense from before. On the other hand:

Fact 14.4.2 (Gelfand-Neimark). Every C^* -algebra is isomorphic to a concrete C^* -algebra.

Taking this “concrete” perspective on operator algebras, we arrive at a new class of operator algebras by considering a different topology on $\mathcal{B}(H)$:

Definition 14.4.3. The **strong operator topology** (SOT) on $\mathcal{B}(H)$ is the topology of pointwise convergence. More precisely, given $T \in \mathcal{B}(H)$, a basic SOT-open neighborhood of T is of the form $\{S \in \mathcal{B}(H) : \|S\xi - T\xi\| < \epsilon \text{ for all } \xi \in F\}$, where F ranges over the finite subsets of H and ϵ ranges over positive real numbers.

Note that we will start using Greek letters for elements of H , reserving Roman letters for elements of operator algebras.

Definition 14.4.4. A **von Neumann algebra** is a unital $*$ -subalgebra of $\mathcal{B}(H)$ closed in the strong operator topology.

Remark 14.4.5. Note that if $T_i \rightarrow T$ in operator norm, then $T_i \rightarrow T$ in the strong operator topology: if $\xi \in H$, then $\|T_i\xi - T\xi\| \leq \|T_i - T\|\|\xi\| \rightarrow 0$ since $\|T_i - T\| \rightarrow 0$. Consequently, every von Neumann algebra is a unital C^* -algebra.

Remark 14.4.6. In the spirit of “concrete” versus “abstract,” what we have defined above is a concrete von Neumann algebra. There is also an abstract characterization: a C^* -algebra M is a von Neumann algebra if there is a Banach space X such that $X^* \cong M$ (as Banach spaces).

Here are some examples of von Neumann algebras. We start with a rather trivial (but still important) example:

Example 14.4.7. $\mathcal{B}(H)$ is a von Neumann algebra. In particular, $M_n(\mathbb{C})$, the algebra of $n \times n$ matrices over \mathbb{C} , is a von Neumann algebra, as it is $\mathcal{B}(H)$ for H an n -dimensional Hilbert space.

Example 14.4.8. Suppose that (X, μ) is a probability space. Recall that $L^2(X, \mu)$ is a Hilbert space under $\langle f, g \rangle := \int_X f(x)\overline{g(x)}d\mu(x)$. Given $f \in L^\infty(X, \mu)$, define $M_f : L^2(X, \mu) \rightarrow L^2(X, \mu)$ to be $M_f(g) := fg$. Then M_f is a bounded operator on $L^2(X, \mu)$ with $\|M_f\| = \|f\|_\infty$. Let $\pi : L^\infty(X, \mu) \rightarrow \mathcal{B}(L^2(X, \mu))$ be given by $\pi(f) := M_f$. We claim that $\pi(L^\infty(X, \mu))$ is a commutative von Neumann algebra. (One often abuses notation and simply says that $L^\infty(X, \mu)$ is a von Neumann algebra.) To see this, suppose that T is in the strong operator closure of $\pi(L^\infty(X, \mu))$, say $T = \lim_i M_{f_i}$ (strong operator convergence). Set $f := T(1)$. Then for any $g \in L^\infty(X, \mu)$, we have

$$T(g) = \lim_i M_{f_i}g = \lim_i M_{f_i}M_g(1) = M_g(\lim_i M_{f_i}(1)) = gf.$$

It follows that $f \in L^\infty(X, \mu)$. Moreover, since $L^\infty(X, \mu)$ is dense in $L^2(X, \mu)$, it follows that $T(g) = fg$ for all $g \in L^2(X, \mu)$, that is, $T = M_f$.

Remark 14.4.9. The converse of the previous example is also true, that is, every commutative von Neumann algebra is isomorphic to one of the form $L^\infty(X, \mu)$ for some probability space (X, μ) . Thus, von Neumann algebra theory is often referred to as “noncommutative measure theory” (just as C*-algebra theory is referred to as “noncommutative topology”).

The next example is important in representation theory and will also be important in the next section.

Example 14.4.10. Suppose that G is a group. Let $\ell^2(G)$ be the Hilbert space formally generated by an orthonormal basis ζ_h for all $h \in G$. For any $g \in G$, define $u_g := f(g)$ to be the linear operator on $\ell^2(G)$ determined by $u_g(\zeta_h) = \zeta_{gh}$ for all $h \in G$. Notice that $f(g)$ is unitary for all $g \in G$ (since $u_g^* = u_g^{-1} = u_{g^{-1}}$) and so f is a unitary representation of G ; it is called the **left regular representation**. The **group von Neumann algebra of G** , denoted $L(G)$, is the von Neumann subalgebra of $\mathcal{B}(\ell^2(G))$ generated by $f(G)$.

In order to understand the difference between operator norm convergence and strong operator convergence, consider the following:

Exercise 14.4.11. Suppose that $(P_n)_{n \in \mathbb{N}}$ is a sequence of projections from $\mathcal{B}(H)$ such that $P_n < P_{n+1}$ for all $n \in \mathbb{N}$, that is, such that $P_n(H) \subsetneq P_{n+1}(H)$. Set $H_n := P_n(H)$, $H_\infty := \overline{\bigcup_n H_n}$, and P to be the projection onto H_∞ . Prove that P_n converges in the strong operator topology to P but does not converge in the norm topology. Moreover, prove that P is the least upper bound of the sequence $(P_n)_{n \in \mathbb{N}}$ (in the Boolean algebra sense of the ordering).

We now turn to the issue of taking ultraproducts of von Neumann algebras. Since von Neumann algebras are also C*-algebras, one might wonder if we can simply take the usual C*-algebra ultraproduct of a family of von Neumann algebras and the result will be a von Neumann algebra again. Sadly, this is not the case, as the next proposition shows. Here, $\ell^\infty(n)$ is the von Neumann algebra associated to the probability space on n points where each point has measure $\frac{1}{n}$.

Proposition 14.4.12. *For any $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$, $\prod_{\mathcal{U}} \ell^\infty(n)$ is not a von Neumann algebra.*

Proof. Set $A := \prod_{\mathcal{U}} \ell^\infty(n)$. For each $m \in \mathbb{N}$, define $f_m \in \prod_{n \in \mathbb{N}} \ell^\infty(n)$ by setting $f_m(n)(i) = 1$ if $i \leq m$ and 0 otherwise. Note that each $f_m(n)$ is a projection in $\ell^\infty(n)$, whence $[f_m]_{\mathcal{U}}$ is a projection in A . Moreover, $[f_m]_{\mathcal{U}} < [f_{m+1}]_{\mathcal{U}}$ for all $m \in \mathbb{N}$. (Here, we are using the fact that in von Neumann algebras $L^\infty(X)$, the ordering on projections coincides with the

usual ordering on functions.) We claim that this sequence has no least upper bound in A , whence A is not a von Neumann algebra by Exercise 14.4.11. Indeed, suppose, toward a contradiction, that $[f]_{\mathcal{U}}$ is the least upper bound of the sequence $([f_m]_{\mathcal{U}})_{m \in \mathbb{N}}$. For each $m \in \mathbb{N}$, set

$$A_m := \{n \geq m : f(n)(i) \geq 1 \text{ for all } i \leq m\}.$$

By assumption, each $A_m \in \mathcal{U}$. For each n , set $m(n) := \max\{m : n \in A_m\}$. Note that $\lim_{\mathcal{U}} m(n) = \infty$. Let $g \in \prod_{n \in \mathbb{N}} \ell^\infty(n)$ be given by $g(n)(i) = f(n)(i)$ for $i < m(n)$ while $g(n)(i) = 0$ for $i \geq m(n)$. Then $g(n) < f(n)$ for all n , so $[g]_{\mathcal{U}} < [f]_{\mathcal{U}}$. However, given any $m \in \mathbb{N}$, if $n \in A_m$, then since $m - 1 < m(n)$, for all $i = 1, \dots, m - 1$, we have $g(n)(i) \geq 1 = f_{m-1}(n)(i)$. Consequently, if $n \in A_m$, then $g(n) \geq f_{m-1}(n)$, whence $[g]_{\mathcal{U}} \geq [f_{m-1}]_{\mathcal{U}}$. It follows that $[g]_{\mathcal{U}}$ is also a upper bound for the sequence $([f_m]_{\mathcal{U}})_{m \in \mathbb{N}}$, a contradiction to the choice of $[f]_{\mathcal{U}}$. \square

There is, however, an ultraproduct construction that works for a large family of von Neumann algebras and which yields ultraproducts that are von Neumann algebras again. The trick is to consider a metric ultraproduct with respect to a different metric (that is, not the metric that arises from the operator norm). Such metrics will arise from traces, which we now define. In what follows, all C^* -algebras will be assumed to be unital.

Definition 14.4.13. If A is a C^* -algebra, then a **tracial state** on A is a linear functional $\tau : A \rightarrow \mathbb{C}$ such that

- (1) $\tau(1) = 1$;
- (2) τ is positive, that is, for all $a \in A$, $\tau(a^*a) \geq 0$;
- (3) $\tau(ab) = \tau(ba)$ for all $a, b \in A$.

The tracial state is called **faithful** if, in addition, we have:

- (4) $\tau(a^*a) = 0$ if and only if $a = 0$.

Before we show how the notion of tracial states helps us in defining our new ultraproduct construction, we give a few examples:

Example 14.4.14. The normalized trace τ_n on $M_n(\mathbb{C})$ is given by defining $\tau_n(a) = \frac{1}{n} \sum_i a_{ii}$. It is straightforward to check that τ_n is a faithful tracial state on $M_n(\mathbb{C})$.

However, if H is infinite-dimensional, then there is no faithful tracial state on $\mathcal{B}(H)$. Indeed, let H_1 be a proper closed subspace of H with the same dimension as H . Let $v : H \rightarrow H_1$ be a surjective linear isometry, which is possible since H and H_1 have the same dimension (whence are isomorphic) and view v as an element of $\mathcal{B}(H)$. Let $p \in \mathcal{B}(H)$ denote the projection onto H_1 . It is straightforward to verify that $v^*(x) = v^{-1}(p(x))$,

whence $v^*v = 1$ and $vv^* = p$. Consequently, if τ were a faithful tracial state on $\mathcal{B}(H)$, then we would have that

$$1 = \tau(1) = \tau(v^*v) = \tau(vv^*) = \tau(p)$$

and thus $\tau(1-p) = 0$. Since $(1-p)^*(1-p) = 1-p$, faithfulness of τ implies that $1-p = 0$, that is, $p = 1$, which is a contradiction to the fact that $H_1 \neq H$.

Example 14.4.15. Given any probability space (X, μ) , there is a faithful tracial state τ on $L^\infty(X, \mu)$ given by $\tau(f) := \int_X f d\mu$. For this reason, traces are often considered the noncommutative version of integrals.

Example 14.4.16. For any group G , there is a faithful tracial state τ on $L(G)$ defined by setting $\tau(x) := \langle x\zeta_e, \zeta_e \rangle$.

The previous example might seem a bit strange at first. Here is an exercise to help understand where this comes from:

Exercise 14.4.17. Suppose that G is a finite group of size n , whence $L(G)$ is a von Neumann subalgebra of $\mathcal{B}(\ell^2(G)) \cong M_n(\mathbb{C})$. Show that the above trace on $L(G)$ is the restriction of the normalized trace τ_n on $M_n(\mathbb{C})$ to $L(G)$.

We now return to how tracial states help us define new metrics. Indeed, if τ is a faithful tracial state on a C^* -algebra A , then τ induces an inner product on A given by $\langle x, y \rangle_\tau := \tau(y^*x)$. This inner product on A in turn induces a norm $\|\cdot\|_\tau$ on A defined by $\|x\|_\tau := \sqrt{\langle x, x \rangle_\tau} = \sqrt{\text{tr}(x^*x)}$. As usual, the norm $\|\cdot\|_\tau$ gives rise to a metric d_τ on A given by $d_\tau(x, y) := \|x - y\|_\tau$.

Lemma 14.4.18. *Suppose that τ is a faithful tracial state on A . Then for all $x \in A$, we have $|\tau(x)| \leq \|x\|_\tau$.*

Proof. Note that $|\tau(x)| = \langle x, 1 \rangle_\tau$. By the Cauchy-Schwarz inequality, we have $|\tau(x)| \leq \|x\|_\tau \cdot \|1\|_\tau = \|x\|_\tau$, as desired. \square

Example 14.4.19. Consider the normalized trace τ_n on $M_n(\mathbb{C})$. It is straightforward to verify that, for any $a \in M_n(\mathbb{C})$, we have $\tau_n(a^*a) = \frac{1}{n} \sum_{i,j=1}^n |a_{ij}|^2$, whence $\|a\|_\tau = \sqrt{\frac{1}{n} \sum_{i,j=1}^n |a_{ij}|^2}$. Note that the induced metric d_τ on $M_n(\mathbb{C})$, when restricted to U_n , is exactly the Hilbert-Schmidt metric introduced Exercise 13.1.11. We will have more to say about this connection in the next section.

Example 14.4.20. If τ is the tracial state on $L^\infty(X, \mu)$ given by $\tau(f) = \int f d\mu$, then the corresponding norm on $L^\infty(X, \mu)$ is the norm more commonly denoted $\|\cdot\|_2$, that is,

$$\|f\|_\tau = \|f\|_2 = \sqrt{\int_X |f|^2 d\mu}.$$

Remark 14.4.21. Motivated by the previous example, if the trace τ is clear from context, then the norm $\|\cdot\|_\tau$ is often written as $\|\cdot\|_2$ and referred to as the 2-norm on A .

Example 14.4.22. Consider the projections $[f_m]_{\mathcal{U}} \in \prod_{\mathcal{U}} \ell^\infty(n)$ in the proof of Proposition 14.4.12. Recall that $f_m(n)(i) = 1$ if $i \leq m$ and $f_m(n)(i) = 0$ otherwise. In particular, letting μ_n denote the normalized counting measure on $\ell^\infty(n)$, we have $\int_{\ell^\infty(n)} |f_m|^2 d\mu_n = \frac{m+1}{n}$. Consequently, letting τ_n denote the corresponding trace on $\ell^\infty(n)$, we have $\lim_{\mathcal{U}} \|f_m\|_{\tau_n} = 0$ for all $m \in \mathbb{N}$.

Recall that the projections $[f_m]_{\mathcal{U}}$ from the previous example were an issue when trying to conclude that $\prod_{\mathcal{U}} \ell^\infty(n)$ is a von Neumann algebra. However, these problematic projections “disappear in the limit” if we switch focus to the metric d_τ instead of the metric induced from the operator norm. That being said, there is an issue with this approach as, in general, multiplication need not be uniformly continuous with respect to d_τ , even on d_τ -bounded subsets. There is, however, a compromise that turns out to work and which relies on the following inequality:

Lemma 14.4.23. *Suppose that τ is a tracial state on the C^* -algebra A . Then for $x, y \in A$, we have $\|xy\|_\tau \leq \|x\| \|y\|_\tau$ and $\|xy\|_\tau \leq \|x\|_\tau \|y\|$.*

Proof. We only prove the former statement, the proof of the latter being nearly identical. We begin by computing

$$\|xy\|_\tau^2 = \tau((xy)^*(xy)) = \tau(y^*x^*xy) = \tau(x^*xyy^*).$$

Next, a little bit of functional analysis tells us that $(\|x^*x\| \cdot 1) - x^*x$ is a positive element of A . (It helps to think of the function spaces $C(X)$ here: for any function $f \in C(X)$, if $M = \sup_{x \in X} |f(x)|$, then the constant function M dominates f . This special case and Gelfand duality is what is behind the proof of the general case.) Since the product of two positive elements is positive, we have that $\|xy\|_\tau^2 = \tau(x^*xyy^*) \leq \tau(\|x\|^2 yy^*) = \|x\|^2 \|y\|_\tau^2$. \square

Lemma 14.4.23 immediately implies:

Corollary 14.4.24. *If τ is a tracial state on the C^* -algebra A , then for all $x \in A$, we have $\|x\|_\tau \leq \|x\|$.*

Lemma 14.4.23 shows that, when restricted to subsets of A that are bounded in operator norm, multiplication becomes uniformly continuous (even Lipschitz) with respect to d_τ . This leads us to the definition of the tracial ultraproduct. Suppose that $(A_i, \tau_i)_{i \in I}$ is a family of C^* -algebras equipped with faithful tracial states and \mathcal{U} is an ultrafilter on I . We set

$$\ell^\infty(A_i) := \left\{ f \in \prod_{i \in I} A_i : \sup_{i \in I} \|f(i)\| < \infty \right\}$$

and

$$\mathcal{I}_{\mathcal{U}}^{\text{tr}} := \left\{ f \in \ell^\infty(A_i) : \lim_{\mathcal{U}} \|f(i)\|_{\tau_i} = 0 \right\}.$$

Proposition 14.4.25. $\mathcal{I}_{\mathcal{U}}^{\text{tr}}$ is a closed two-sided ideal of the C^* -algebra $\ell^\infty(A_i)$.

Proof. $\mathcal{I}_{\mathcal{U}}^{\text{tr}}$ is a vector subspace of $\ell^\infty(A_i)$ since each $\|\cdot\|_{\tau_i}$ is a norm on A_i . To see that $\mathcal{I}_{\mathcal{U}}^{\text{tr}}$ is a two-sided ideal, fix $f \in \ell^\infty(A_i)$ and $g \in \mathcal{I}_{\mathcal{U}}^{\text{tr}}$. Let $M := \sup_{i \in I} \|f(i)\|$. Then by Lemma 14.4.23, we have $\|(fg)(i)\|_{\tau_i} \leq M\|g(i)\|_{\tau_i}$, whence $\lim_{\mathcal{U}} \|(fg)(i)\|_{\tau_i} = 0$, and $fg \in \mathcal{I}_{\mathcal{U}}^{\text{tr}}$. The proof that $gf \in \mathcal{I}_{\mathcal{U}}^{\text{tr}}$ is identical.

It remains to see that $\mathcal{I}_{\mathcal{U}}^{\text{tr}}$ is closed. Toward that end, fix a Cauchy sequence $(f_n)_{n \in \mathbb{N}}$ from $\mathcal{I}_{\mathcal{U}}^{\text{tr}}$ with limit $f \in \ell^\infty(A_i)$; we must show that $f \in \mathcal{I}_{\mathcal{U}}^{\text{tr}}$. Fix $\epsilon > 0$ and take $n \in \mathbb{N}$ such that $\|f - f_n\| < \epsilon$. Given $i \in I$, we have

$$\begin{aligned} \|f(i)\|_{\tau_i} &\leq \|f(i) - f_n(i)\|_{\tau_i} + \|f_n(i)\|_{\tau_i} \\ &\leq \|f(i) - f_n(i)\| + \|f_n(i)\|_{\tau_i} < \epsilon + \|f_n(i)\|_{\tau_i}, \end{aligned}$$

where the second inequality sign uses Corollary 14.4.24. Taking ultralimits, we have $\lim_{\mathcal{U}} \|f(i)\| \leq \epsilon$; since ϵ is arbitrary, we have that $f \in \mathcal{I}_{\mathcal{U}}^{\text{tr}}$. \square

It is a standard fact that the quotient of a C^* -algebra by a closed two-sided ideal is a C^* -algebra again (with the obvious quotient operations), allowing us to make the following:

Definition 14.4.26. Using the notation above, the **tracial ultraproduct** of the family $(A_i, \tau_i)_{i \in I}$ is the C^* -algebra $\ell^\infty(A_i)/\mathcal{I}_{\mathcal{U}}^{\text{tr}}$. To distinguish this ultraproduct from the usual C^* -ultraproduct, we write $\prod_{\mathcal{U}}^{\text{tr}}(A_i, \tau_i)$, or simply $\prod_{\mathcal{U}}^{\text{tr}} A_i$.

In the rest of this section, given $f \in \ell^\infty(A_i)$, we will continue to denote its coset in $\prod_{\mathcal{U}}^{\text{tr}} A_i$ by $[f]_{\mathcal{U}}$. If we were behaving, we would decorate it differently, such as $[f]_{\mathcal{U}}^{\text{tr}}$. However, this notation would become very cumbersome and we hope this abuse causes no confusion.

Proposition 14.4.27. *Using the notation above, there is a well-defined faithful tracial state τ on $\prod_{\mathcal{U}}^{\text{tr}} A_i$ given by $\tau([f]_{\mathcal{U}}) := \lim_{\mathcal{U}} \tau_i(f(i))$.*

Proof. First note that, by Lemma 14.4.18 and Corollary 14.4.24, given $f \in \ell^\infty(A_i)$, we have

$$\sup_{i \in I} |\tau_i(f(i))| \leq \sup_{i \in I} \|f(i)\|_{\tau_i} \leq \sup_{i \in I} \|f(i)\| < \infty,$$

whence $\lim_{\mathcal{U}} |\tau_i(f(i))|$ exists. To see that the function is well-defined, suppose $[f]_{\mathcal{U}} = [g]_{\mathcal{U}}$, that is, $\lim_{\mathcal{U}} \|f(i) - g(i)\|_{\tau_i} = 0$. By Lemma 14.4.18, we have $\lim_{\mathcal{U}} |\tau_i(f(i) - g(i))| \leq \lim_{\mathcal{U}} \|f(i) - g(i)\|_{\tau_i} = 0$, whence $\lim_{\mathcal{U}} \tau_i(f(i)) = \lim_{\mathcal{U}} \tau_i(g(i))$, as desired. It is fairly clear that τ is a tracial state. To see that τ is faithful, note that

$$\tau([f]^*[f]) = \lim_{\mathcal{U}} \tau_i(f(i)^*f(i)) = \lim_{\mathcal{U}} \|f(i)\|_{\tau_i}.$$

Consequently, if $\tau([f]^*[f]) = 0$, then $f \in \mathcal{I}_{\mathcal{U}}^{\text{tr}}$, that is, $[f]_{\mathcal{U}} = 0$. □

We started this discussion by pondering an ultraproduct construction for von Neumann algebras that yielded a von Neumann algebra again. It turns out that if we take the tracial ultraproduct of a family of tracial von Neumann algebras satisfying an extra continuity condition, then the resulting tracial ultraproduct will indeed be a von Neumann algebra (and the ultraproduct trace will also satisfy that extra continuity condition). In the remainder of this section, for a C^* -algebra A , we denote its operator norm unit ball by $A_{\leq 1}$.

Definition 14.4.28. If M is a von Neumann algebra, then a tracial state τ on M is **normal** if $\tau \upharpoonright M_{\leq 1}$ is continuous with respect to the strong operator topology. A **trace** on M is a faithful normal tracial state. A **tracial von Neumann algebra** is a pair (M, τ) , where M is a von Neumann algebra and τ is a trace on M .

Remark 14.4.29. The tracial states defined in Examples 14.4.15 and 14.4.16 are normal.

We devote the rest of this section to proving the following:

Theorem 14.4.30. *Suppose that $(M_i, \tau_i)_{i \in I}$ is a family of tracial von Neumann algebras and \mathcal{U} is an ultrafilter on I . Then $\prod_{\mathcal{U}}^{\text{tr}}(M_i, \tau_i)$ is a tracial von Neumann algebra when equipped with the ultralimit trace introduced in Proposition 14.4.27.*

Recall that, at the moment, even if we are considering a tracial ultraproduct of a family of von Neumann algebras equipped with tracial states, the tracial ultraproduct is merely a C^* -algebra equipped with a tracial state. We thus need a criteria for knowing when such a C^* -algebra is in fact a von Neumann algebra. The key is the so-called **GNS construction** (named after Gelfand, Naimark, and Segal), which we now describe.

Suppose that A is a C^* -algebra equipped with a faithful tracial state τ . We define $L^2(A, \tau)$ to be the completion of A with respect to the metric d_τ . It is a standard fact that the inner product $\langle \cdot, \cdot \rangle_\tau$ extends naturally to an inner product on $L^2(A, \tau)$, which we continue to denote by $\langle \cdot, \cdot \rangle_\tau$, whence $L^2(A, \tau)$ is a Hilbert space. For $x \in A$, we write \hat{x} if we wish to emphasize its role as a vector in the Hilbert space $L^2(A, \tau)$. Lemma 14.4.23 states that, for $x, y \in A$, we have

$$\|x\hat{y}\|_\tau \leq \|x\| \|\hat{y}\|_\tau.$$

Since the vectors \hat{y} , for $y \in A$, are dense in $L^2(A, \tau)$, we see that the map $\hat{y} \mapsto x\hat{y} : A \rightarrow A$ extends to a bounded operator $\pi_\tau(x) \in \mathcal{B}(L^2(A, \tau))$ with $\|\pi_\tau(x)\| = \|x\|$. Consequently, we have a faithful representation $\pi_\tau : A \rightarrow \mathcal{B}(L^2(A, \tau))$.

Example 14.4.31. If $A = L^\infty(X, \mu)$ and $\tau(f) = \int_X f d\mu$, then $L^2(A, \tau) = L^2(X, \mu)$ and π_τ is an isomorphism between A and $\pi_\tau(A)$.

In general, if M is a von Neumann algebra with tracial state τ , then $\pi_\tau(M)$ need not be a von Neumann subalgebra of $\mathcal{B}(L^2(M, \tau))$, that is, $\pi_\tau(M)$ may not be an SOT-closed subalgebra of $\mathcal{B}(L^2(M, \tau))$. However:

Fact 14.4.32. If τ is a tracial state on the von Neumann algebra M , then τ is normal if and only if $\pi_\tau(M)$ is a von Neumann subalgebra of $\mathcal{B}(L^2(M, \tau))$.

Proof. See [1, Prop 2.5.12, Theorem 2.6.1, and Proposition 2.6.4]. \square

We return to the general case, that is, τ is a faithful tracial state on the C^* -algebra A . Set M to be the closure of $\pi_\tau(A)$ in the strong operator topology, so M is a von Neumann subalgebra of $\mathcal{B}(L^2(A, \tau))$. We note that τ extends to a trace on M given by $\tau(x) = \langle x\hat{1}, \hat{1} \rangle_\tau$. (Note that, for $a \in A$, we have $\tau(\pi_\tau(a)) = \tau(a)$, whence τ really is an extension of its restriction to A .)

The next proposition will be the key to the proof of Theorem 14.4.30:

Proposition 14.4.33. *Continuing with the notation from above, we have that $\pi_\tau(A_{\leq 1}) = M_{\leq 1}$ if and only if $A_{\leq 1}$ is d_τ -complete. In this case, A is a von Neumann algebra and τ is a trace on A (that is, τ is normal).*

Proof sketch. First suppose that $\pi_\tau(A_{\leq 1}) = M_{\leq 1}$. To show that $A_{\leq 1}$ is d_τ -complete, fix a d_τ -Cauchy sequence $(x_n)_{n \in \mathbb{N}}$ from $A_{\leq 1}$. By Lemma 14.4.23, given $y \in A$, we have $\|x_n\hat{y} - x_m\hat{y}\|_\tau \leq \|y\| \|x_n - x_m\|_\tau$, whence $(x_n\hat{y})_{n \in \mathbb{N}}$ is a Cauchy sequence in $L^2(A, \tau)$. It is readily verified that the map $y \mapsto \lim_{n \rightarrow \infty} x_n\hat{y} : A \rightarrow L^2(A, \tau)$ extends to an element $x \in \mathcal{B}(L^2(A, \tau))$ with $\|x\| = \lim_{n \rightarrow \infty} \|x_n\|_\tau$. By definition, $\pi_\tau(x_n) \rightarrow x$ in the strong operator topology. Since $x \in M_{\leq 1}$, our assumption tells us that $x \in \pi_\tau(A_{\leq 1})$, say

$x = \pi_\tau(x')$ with $x' \in A_{\leq 1}$. It remains to show that $\lim_{n \rightarrow \infty} d_\tau(x_n, x') = 0$, which follows from the fact that

$$\|x_n - x'\|_\tau = \|\pi_\tau(x_n)\hat{1} - \pi_\tau(x')\hat{1}\| = \|\pi_\tau(x_n)\hat{1} - x\hat{1}\| \rightarrow 0.$$

Now suppose that $A_{\leq 1}$ is d_τ -complete and suppose that $x \in M_{\leq 1}$; we wish to show that $x \in \pi_\tau(A_{\leq 1})$. By the *Kaplansky Density Theorem*, there is a sequence $(x_n)_{n \in \mathbb{N}}$ from $A_{\leq 1}$ such that $\pi_\tau(x_n) \rightarrow x$ in the strong operator topology. (This may seem trivial from the fact x is in the strong operator closure of $\pi_\tau(A)$, but getting the vectors to have operator norm at most 1 is the content of the theorem.) In particular, $\pi_\tau(x_n)(\hat{1}) \rightarrow x\hat{1}$, whence $\|x_n\|_\tau = \|\pi_\tau(x_n)\hat{1}\|_\tau \rightarrow \|x\hat{1}\|_\tau$. It follows that the sequence $(x_n)_{n \in \mathbb{N}}$ is d_τ -Cauchy, whence there is $y \in A_{\leq 1}$ such that $\lim_{n \rightarrow \infty} d_\tau(x_n, y) = 0$. By Lemma 14.4.23, it follows that, for $z \in A$, we have

$$\|(x_n - y)\hat{z}\|_\tau \leq \|z\| \|(x_n - y)\hat{1}\|_\tau = \|z\| \|x_n - y\|_\tau.$$

Since A is dense in $L^2(A, \tau)$, it follows that $\pi_\tau(x_n) \rightarrow \pi_\tau(y)$ in the strong operator topology. We conclude that $x = \pi_\tau(y)$, as desired.

The moreover part follows immediately from the first part and Fact 14.4.32. □

We are now ready to prove our main result:

Proof of Theorem 14.4.30. Recall that we have a family $(M_i, \tau_i)_{i \in I}$ of tracial von Neumann algebras and an ultrafilter \mathcal{U} on I . We wish to show that the tracial ultraproduct $M := \prod_{\mathcal{U}}^{\text{tr}} M_i$, which, a priori, is merely a C^* -algebra equipped with its ultraproduct trace τ , is actually a von Neumann algebra. To do this, we use the criteria developed in Proposition 14.4.33, that is, we show that $M_{\leq 1}$ is d_τ -complete. The proof is nearly identical to the proof of Theorem 11.3.1 with a small wrinkle.

Suppose that $([x_n]_{\mathcal{U}})_{n \in \mathbb{N}}$ is a d_τ -Cauchy sequence from $M_{\leq 1}$; we wish to show that it has a limit in $M_{\leq 1}$. By rescaling if necessary, we may assume that $\|[x_n]_{\mathcal{U}}\| < 1$ for each $n \in \mathbb{N}$. The benefit of making this slight change is that we may then assume that $x_n(i) \in (M_i)_{\leq 1}$ for all $n \in \mathbb{N}$ and all $i \in I$. We are thus considering a Cauchy sequence in the metric ultraproduct $\prod_{\mathcal{U}}((M_i)_{\leq 1}, d_{\tau_i})$, where each metric space $((M_i)_{i \leq 1}, d_{\tau_i})$ is complete by Fact 14.4.32 and Proposition 14.4.33. Consequently, Theorem 11.3.1 implies that the sequence $([x_n]_{\mathcal{U}})_{n \in \mathbb{N}}$ has a limit in $\prod_{\mathcal{U}}((M_i)_{\leq 1}, d_{\tau_i})$, which is of course contained in M . □

14.5. The Connes embedding problem

In this section, all ultraproducts will be tracial ultraproducts, whence we simplify our notation and simply use the usual ultraproduct notation. In

addition, when the trace on a von Neumann algebra is the “canonical” one, we omit mention of the trace.

In his seminal paper [34], Connes made the following comment. (Don’t worry, we will translate into our terminology right afterward.)

“We now construct an approximate imbedding of N in \mathcal{R} . Apparently such an imbedding ought to exist for all II_1 factors because it does for the regular representation of free groups. However, the construction below relies on condition 6.”

Connes was considering a particular tracial von Neumann algebra N satisfying some particular extra conditions. The first sentence, in our terminology, indicates that he is about to construct an injective trace-preserving $*$ -homomorphism from N to a tracial ultrapower of a particularly important tracial von Neumann algebra known as the **hyperfinite II_1 factor \mathcal{R}** . Rather than define \mathcal{R} here², it suffices to replace this statement with the desire to embed N into $\prod_{\mathcal{U}} M_n(\mathbb{C})$ for some nonprincipal ultrafilter \mathcal{U} on \mathbb{N} . His next sentence then asserts that such an embedding ought to exist for any tracial von Neumann algebra N (he writes “ II_1 factor” which is a particular kind of tracial von Neumann algebra, but this does not affect the statement) since it does for $L(\mathbb{F}_2)$, the von Neumann algebra associated to the free group \mathbb{F}_2 . He does not attempt to prove this fact and he simply mentions that his embedding uses the fact that N satisfies some particular extra condition.

This seemingly innocuous remark has since turned into one of the most famous problems in operator algebras. Some people call it the Connes Embedding Conjecture, but “ought to” does not seem to indicate a strong enough belief that the result is actually true, whence many downgrade “Conjecture” to “Problem”:

Conjecture 14.5.1 (Connes embedding problem (CEP)). *For any separable tracial von Neumann algebra M , there is an embedding of tracial von Neumann algebras $M \hookrightarrow \prod_{\mathcal{U}} M_n(\mathbb{C})$.*

Here, a tracial von Neumann algebra M is called separable if $M_{\leq 1}$ is d_{τ} -separable.

The Connes Embedding Problem has proven to be one of the most interesting open problems in the theory of operator algebras. It has highly

²By a serious theorem of Connes (see [34, Corollary 7.2]), $\mathcal{R} \cong L(G)$ for any amenable group G in which all conjugacy classes (except the trivial one) are infinite. An example of such a group is $\bigcup_{n \geq 2} S_n$. This is a perfect example of how the group von Neumann algebra can “forget” a lot of the algebraic structure of the original group.

nontrivial reformulations in terms of C^* -algebras, quantum information theory, logic, etc. . . . In this section, we only mention a connection with the class of hyperlinear groups introduced in the previous chapter.

Definition 14.5.2. Suppose that A is a unital C^* -algebra. An element $u \in A$ is called a **unitary** if $uu^* = u^*u = 1$.

It is fairly easy to see that the collection of unitaries in A form a group under multiplication, which we denote by $U(A)$.

Exercise 14.5.3. Suppose that τ is a faithful tracial state on A . Prove that the metric d_τ restricted to $U(A)$ is a bi-invariant metric (see Section 13.1) on $U(A)$.

In particular, when $A = M_n(\mathbb{C})$, the metric d_τ on $U(M_n(\mathbb{C})) = U_n$ is the Hilbert-Schmidt metric introduced in the previous chapter.

The next fact indicates that “almost” unitaries are “near” actual unitaries:

Fact 14.5.4. Given $\epsilon > 0$, there is $\delta > 0$ such that, for any tracial von Neumann algebra M and any $x \in M_{\leq 1}$, if $\|xx^* - 1\|_\tau, \|x^*x - 1\|_\tau < \delta$, then there is a unitary $u \in U(M)$ such that $\|x - u\|_\tau < \epsilon$.

Exercise 14.5.5. Suppose that $(M_i)_{i \in I}$ is a family of tracial von Neumann algebras and \mathcal{U} is an ultrafilter on I . Prove that $U(\prod_{\mathcal{U}} M_i) = \prod_{\mathcal{U}} U(M_i)$ as bi-invariant metric groups. (*Hint.* Use the previous fact.)

By the previous exercise, we see that if G is a countable group such that $L(G)$ embeds into $\prod_{\mathcal{U}} M_n(\mathbb{C})$, then G is hyperlinear (as defined in Section 13.4), for G embeds as a subgroup of $U(L(G))$, which in turn embeds in $U(\prod_{\mathcal{U}} M_n(\mathbb{C})) = \prod_{\mathcal{U}} U(M_n(\mathbb{C})) = \prod_{\mathcal{U}} U_n$. The main result in this section is the converse of this observation:

Theorem 14.5.6 (Radulescu[145]). *For any countable group G , G is hyperlinear if and only if there is an embedding of tracial von Neumann algebras $L(G) \hookrightarrow \prod_{\mathcal{U}} M_n(\mathbb{C})$. Consequently, a positive solution to CEP implies that all groups are hyperlinear.*

The following result explains how one obtains embeddings from a group von Neumann algebra $L(G)$ into an arbitrary tracial von Neumann algebra:

Proposition 14.5.7. *Suppose that G is a group and M is a tracial von Neumann algebra. Then a group homomorphism $i : G \rightarrow U(M)$ extends to an embedding of tracial von Neumann algebras $\tilde{i} : L(G) \rightarrow M$ if and only if $\tau(i(g)) = 0$ for all $g \in G \setminus \{e\}$.*

Proof sketch. The forward direction is clear. For the backward direction, note that such an embedding i extends to an embedding of the span of $\{\zeta_g : g \in G\}$ in $L(G)$ into M that preserves the $*$ -algebra operations and the trace. It remains to see that this embedding extends to all of $L(G)$. The key to this is that the strong operator topologies on $L(G)_{\leq 1}$ and $M_{\leq 1}$ are the same as those induced by the metrics arising from the trace (as shown in the proof of Proposition 14.4.33), which are preserved by assumption. \square

At first glance, given a hyperlinear group G , we only know that there is a group homomorphism $i : G \rightarrow \prod_{\mathcal{U}} U_n = U(\prod_{\mathcal{U}} M_n(\mathbb{C}))$ such that $i(g) \neq 1$ for all $g \in G \setminus \{e\}$. In order to obtain an embedding as in Proposition 14.5.7, one needs to perform some cosmetic surgery on this initial embedding i , as we now explain.

Lemma 14.5.8. *Suppose that A is a C^* -algebra and $u \in U(A)$ is such that $|\tau(u)| = 1$. Then $u = \tau(u) \cdot 1$.*

Proof. The result follows from the calculation

$$\|u - \tau(u) \cdot 1\|_{\tau}^2 = \tau((u^* - \overline{\tau(u)}) \cdot 1)(u - \tau(u) \cdot 1) = \tau(u^*u) - |\tau(u)|^2 = 0.$$

\square

Lemma 14.5.9. *Given $u \in U_n$, set $u' := \begin{pmatrix} u & 0 \\ 0 & 1 \end{pmatrix} \in U_{2n}$. (The 1 in the bottom right corner indicates the $n \times n$ identity matrix.). If $u \neq 1$, then $|\text{tr}(u')| < 1$.*

Proof. By Lemma 14.4.18, $|\text{tr}(u')| \leq \|u'\|_{\tau} = 1$. If $|\text{tr}(u')| = 1$, then by Lemma 14.5.8, $u' = \text{tr}(u') \cdot 1$. In this case, $\text{tr}(u') = 1$, which implies that $u' = 1$ and hence $u = 1$. \square

We now sketch the ‘‘amplification trick’’ pertinent to hyperlinear groups. (See Section 13.2 above for the sofic version of this trick.) Given matrices A and B of size $m \times n$ and $p \times q$ respectively, we define their tensor product $A \otimes B$ to be the $mp \times nq$ matrix given by

$$A \otimes B := \begin{pmatrix} A_{11}B & A_{12}B & \dots \\ \vdots & \ddots & \\ A_{m1}B & & A_{mn}B \end{pmatrix}$$

In particular, if $A \in M_n(\mathbb{C})$, then $A \otimes A \in M_{n^2}(\mathbb{C})$.

Exercise 14.5.10. Suppose that $A \in M_n(\mathbb{C})$ and τ_n is the normalized trace on $M_n(\mathbb{C})$. Prove that $\tau_{n^2}(A \otimes A) = \tau_n(A)^2$.

Exercise 14.5.11. Suppose that G is a hyperlinear group. For any $F \subseteq G$ and any $\epsilon > 0$, prove that there is $n \in \mathbb{N}$ and a function $\phi : G \rightarrow U_n$ satisfying:

- (1) For all $g, h \in F$, if $gh \in F$, then $d_\tau(\phi(g)\phi(h), \phi(gh)) < \epsilon$;
- (2) If $e \in F$, then $d_\tau(\phi(e), 1) < \epsilon$;
- (3) For all distinct $g \in F \setminus \{e\}$, $|\tau(g)| < \epsilon$.

(*Hint.* Use Lemma 14.5.9 and Exercise 14.5.10.)

Exercise 14.5.12. Suppose that G is a countable hyperlinear group. Prove that there is a group homomorphism $i : G \rightarrow \prod_{\mathcal{U}} U_n$ such that $\tau(i(g)) = 0$ for all $g \in G \setminus \{e\}$. (*Hint.* Use Exercise 14.5.11.)

Note that Proposition 14.5.7 and Exercise 14.5.12 immediately yield the other direction of Theorem 14.5.6.

14.6. Notes and references

Banach space ultraproducts seem to appear for the first time in Krivine's thesis [110]. One would be remiss if one did not mention two of the more spectacular uses of ultraproduct techniques in Banach space, namely *Krivine's theorem* on block finite representability of ℓ_p and c_0 and the *Krivine-Maurey theorem* stating that stable Banach spaces contain some ℓ_p almost isometrically. See Iovino's book [87] for an in-depth discussion of these results. Our treatment in Section 14.2 follows Coleman [32]. Corollary 14.2.26 was used by Dineen in [44] to show that the bidual of a so-called JB* triple is once again a JB* triple. The work of Wright in [184] seems to be the earliest use of the tracial ultraproduct construction. It is no exaggeration to say that this construction has been of incredible use in the theory of von Neumann algebras since then. As mentioned in Section 14.5, Connes' work in [34] makes serious use of this ultraproduct as does McDuff's work on central sequences in [134]. Our treatment of the tracial ultraproduct construction relies on the book [1] in many ways. Proposition 14.4.12 is inspired by a similar example in [143].

During the writing of this book, a group of computer scientists have claimed to give a negative solution to the Connes Embedding Problem [91]. Since the paper involved is over 160 pages long, it may take quite a while before we have a consensus on the correctness of the solution. The connection with computer science may seem quite strange, but it relies on earlier work by Kirchberg [106], Fritz [60], Junge et. al [93], and Ozawa [141], showing that the CEP is equivalent to a problem in quantum information theory known as Tsirelson's problem. It is worth noting that even if the proof in

[91] is correct, it does not seem to (at this moment) give any insight into the problem of whether or not there is a group that is not hyperlinear.

Part 4

Advanced topics

Does an ultrapower depend on the ultrafilter?

In this chapter, we show that the question of whether or not an ultrapower of a countable structure in a countable language with respect to an ultrafilter on a countable set depends on the ultrafilter may depend on whether or not CH is true, and whether or not it depends on CH is characterized by the model-theoretic notion of stability. In Section 15.1, we give the background for the discussion, define the notion of stability, and state the previous claim precisely. In Section 15.2, we give a more-or-less complete proof of the unstable case while in Section 15.3 we sketch a proof of the stable case, taking a lot of model-theoretic facts about stability for granted.

15.1. Statement of results

In this chapter, we address the question appearing in the title, that is, given a structure \mathcal{M} and nonprincipal ultrafilters \mathcal{U} and \mathcal{V} , must $\mathcal{M}^{\mathcal{U}}$ be isomorphic to $\mathcal{M}^{\mathcal{V}}$?

Stated in this generality, we see that the answer is obviously: no! For example, if \mathcal{U} and \mathcal{V} are regular ultrafilters on sets I and J , respectively, then $|\mathcal{M}^{\mathcal{U}}| = |M|^{|I|}$ while $|\mathcal{M}^{\mathcal{V}}| = |M|^{|J|}$; as long as $|M|^{|I|} \neq |M|^{|J|}$, $\mathcal{M}^{\mathcal{U}}$ and $\mathcal{M}^{\mathcal{V}}$ cannot be isomorphic simply for cardinality reasons.

To avoid such cardinality-theoretic trivialities, let us refine the question by assuming that \mathcal{U} and \mathcal{V} are on index sets of the same size, say they are both ultrafilters on a cardinal κ . However, even in this case, we have seen

instances when $\mathcal{M}^{\mathcal{U}}$ and $\mathcal{M}^{\mathcal{V}}$ need not be isomorphic. For example, as long as $\text{Th}(\mathcal{M})$ is not minimal in Keisler's order (see Section 8.6), we see that there is a cardinal κ and a regular nongood ultrafilter \mathcal{U} on κ such that $\mathcal{M}^{\mathcal{U}}$ is not κ^+ -saturated. If \mathcal{V} is a good ultrafilter on κ , then $\mathcal{M}^{\mathcal{V}}$ is κ^+ -saturated, whence $\mathcal{M}^{\mathcal{U}}$ and $\mathcal{M}^{\mathcal{V}}$ are not isomorphic.

We refine our question one last time: suppose that \mathcal{M} is a countable structure in a countable language and \mathcal{U} and \mathcal{V} are nonprincipal ultrafilters on \mathbb{N} . Must $\mathcal{M}^{\mathcal{U}}$ be isomorphic to $\mathcal{M}^{\mathcal{V}}$? The reason for insisting on ultrafilters on \mathbb{N} is that such ultrafilters are automatically good, whence the phenomenon from the previous paragraph cannot appear.

The question in the previous paragraph is surprisingly complicated to answer. Note that if \mathcal{M} is countable and \mathcal{U} is a nonprincipal ultrafilter on \mathbb{N} , then $\mathcal{M}^{\mathcal{U}}$ is an \aleph_1 -saturated structure of cardinality \mathfrak{c} . If CH holds, then $\mathcal{M}^{\mathcal{U}}$ is in fact saturated. Since $\mathcal{M}^{\mathcal{V}}$ will also be saturated for any other nonprincipal ultrafilter \mathcal{V} on \mathbb{N} , Theorem 8.1.9 immediately implies:

Theorem 15.1.1. *Assume that CH holds. Suppose that \mathcal{M} is a countable structure in a countable language. Then for any nonprincipal ultrafilters $\mathcal{U}, \mathcal{V} \in \beta\mathbb{N} \setminus \mathbb{N}$, we have $\mathcal{M}^{\mathcal{U}} \cong \mathcal{M}^{\mathcal{V}}$.*

So what happens if we assume that CH fails? Well, sometimes the conclusion of the previous theorem still holds. For example, as mentioned in Section 8.6, if T is an uncountably categorical theory (such as ACF_p) and $\mathcal{M} \models T$ is countable, then $\mathcal{M}^{\mathcal{U}}$ is saturated for any nonprincipal ultrafilter \mathcal{U} on \mathbb{N} , whence all such ultrapowers are once again isomorphic.

However, there are seemingly simple structures such as $(\mathbb{N}, <)$ which, assuming that CH fails, can have nonisomorphic ultrapowers corresponding to ultrafilters on \mathbb{N} (we will prove this fact in the next section). It turns out that the key to the proof of Theorem 15.1.1 is the ordering on \mathbb{N} : whether or not the structure \mathcal{M} possesses any definable relation that even resembles an ordering is the deciding factor as to whether or not it has nonisomorphic ultrapowers (again, assuming that CH fails). Here is the precise definition of what it means for a structure to possess a definable relation “resembling an ordering”:

Definition 15.1.2. Fix a language \mathcal{L} and an \mathcal{L} -structure \mathcal{M} .

- (1) If $\varphi(x; y)$ is an \mathcal{L} -formula and $a, b \in M$, write $a <_{\varphi} b$ if $\mathcal{M} \models \varphi(a; b) \wedge \neg\varphi(b; a)$.
- (2) We say that $\varphi(x; y)$ is **unstable in \mathcal{M}** if, for every $n \in \mathbb{N}$, there is a sequence $a_1, \dots, a_n \in M$ with $a_i <_{\varphi} a_j$ for $1 \leq i < j \leq n$. We call such a sequence a **φ -chain of length n in \mathcal{M}** .
- (3) If φ is not unstable in \mathcal{M} , we say that φ is **stable in \mathcal{M}** .

- (4) We say that \mathcal{M} is **stable** if all formulae in \mathcal{M} are stable. If \mathcal{M} is not stable, we say that \mathcal{M} is **unstable**.

Note that φ being unstable in \mathcal{M} means that there are arbitrarily long finite sequences from M that are linearly ordered under $<_{\varphi}$. Consequently, \mathcal{M} is stable if no formula in \mathcal{M} linearly orders arbitrarily long finite sequences.

Exercise 15.1.3. Suppose that \mathcal{M} is stable and \mathcal{N} is elementarily equivalent to \mathcal{M} . Prove that \mathcal{N} is also stable.

By the previous exercise, it makes sense to speak of a complete theory as being stable or unstable.

Example 15.1.4.

- (1) Uncountably categorical theories are stable; this is part of the proof of Morley's categoricity theorem (see [126]).
- (2) The theory of any module over any ring is stable; this is a result due to Fisher [58].
- (3) The theory of nonabelian free groups is stable; this is a difficult result of Sela [156].
- (4) $(\mathbb{N}, <)$ is unstable.
- (5) $(\mathbb{R}, +, \cdot)$ is unstable.

The following exercise asks you to fulfill a promise made in Proposition 8.6.14, namely that every theory with the nfcp (see Definition 8.6.6) is stable:

Exercise 15.1.5. Prove that every theory with the nfcp is stable. (*Hint.* If $\varphi(x, y)$ is unstable, prove that the formula $\psi(x, y_1, y_2, y_3, y_4)$ given by $(\varphi(x, y_1) \leftrightarrow \neg\varphi(x, y_2)) \wedge (\varphi(x, y_3) \leftrightarrow \varphi(x, y_4))$ has the fcp.)

We can now precisely state the connection between stability and isomorphic ultrapowers.

Theorem 15.1.6. *Assume that CH fails.*

- (1) If \mathcal{M} is unstable, then there are $\mathcal{U}, \mathcal{V} \in \beta\mathbb{N} \setminus \mathbb{N}$ such that $\mathcal{M}^{\mathcal{U}} \not\cong \mathcal{M}^{\mathcal{V}}$.
- (2) If \mathcal{M} is stable, then $\mathcal{M}^{\mathcal{U}} \cong \mathcal{M}^{\mathcal{V}}$ for all $\mathcal{U}, \mathcal{V} \in \beta\mathbb{N} \setminus \mathbb{N}$.

In the next section, we prove Theorem 15.1.6(1) above in full detail. The proof of Theorem 15.1.6(2) is a bit beyond the scope of this book, but in Section 15.3 we briefly outline why it holds.

15.2. The case when \mathcal{M} is unstable

In this section, we prove Theorem 15.1.6(1). Toward that end, we suppose that \mathcal{M} is unstable as witnessed by the unstable formula $\varphi(x; y)$.

Definition 15.2.1. For $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$ and an infinite cardinal λ , a φ - λ -gap in $\mathcal{M}^{\mathcal{U}}$ consists of two sequences $\vec{a} := ([a_m]_{\mathcal{U}})_{m < \omega}$ and $\vec{b} := ([b_\gamma]_{\mathcal{U}})_{\gamma < \lambda}$ from $\mathcal{M}^{\mathcal{U}}$ such that:

- (1) $[a_m]_{\mathcal{U}} <_{\varphi} [a_n]_{\mathcal{U}}$ for all $m < n < \omega$ and $[b_\gamma]_{\mathcal{U}} <_{\varphi} [b_\delta]_{\mathcal{U}}$ for all $\delta < \gamma < \lambda$;
- (2) $[a_m]_{\mathcal{U}} <_{\varphi} [b_\gamma]_{\mathcal{U}}$ for all $m < \omega$ and $\gamma < \lambda$;
- (3) there does *not* exist $[c]_{\mathcal{U}} \in \mathcal{M}^{\mathcal{U}}$ such that $[a_m]_{\mathcal{U}} <_{\varphi} [c]_{\mathcal{U}} <_{\varphi} [b_\gamma]_{\mathcal{U}}$ for all $m < \omega$ and $\gamma < \lambda$.

We may write (\vec{a}, \vec{b}) as an abbreviation for the above gap and call λ the length of the gap.

Exercise 15.2.2. Show that any gap in $\mathcal{M}^{\mathcal{U}}$ must have uncountable length.

We define $\kappa(\varphi, \mathcal{M}^{\mathcal{U}})$ to be the minimal cardinality of a gap in $\mathcal{M}^{\mathcal{U}}$. The idea behind the proof of (1) is to find $\mathcal{U}, \mathcal{V} \in \beta\mathbb{N} \setminus \mathbb{N}$ for which $\kappa(\varphi, \mathcal{M}^{\mathcal{U}}) \neq \kappa(\varphi, \mathcal{M}^{\mathcal{V}})$, for then clearly $\mathcal{M}^{\mathcal{U}} \not\cong \mathcal{M}^{\mathcal{V}}$. Precisely, we will prove the following:

Theorem 15.2.3. *For each regular cardinal κ with $\aleph_1 \leq \kappa \leq \mathfrak{c}$, there is $\mathcal{U} \in \beta\mathbb{N}$ such that $\kappa(\varphi, \mathcal{M}^{\mathcal{U}}) = \kappa$.*

Note that Theorem 15.2.3 indeed proves Theorem 15.1.6(1), for when CH fails, $\aleph_1 < \aleph_2 \leq \mathfrak{c}$, so there are \mathcal{U} and \mathcal{V} for which $\kappa(\varphi, \mathcal{M}^{\mathcal{U}}) = \aleph_1$ and $\kappa(\varphi, \mathcal{M}^{\mathcal{V}}) = \aleph_2$ and thus $\mathcal{M}^{\mathcal{U}} \not\cong \mathcal{M}^{\mathcal{V}}$.

Remark 15.2.4. By a result of Solovay [166], it is consistent with ZFC that there are continuum many regular cardinals below \mathfrak{c} . Consequently, Theorem 15.2.3 shows us that it is consistent with ZFC that there are continuum many nonisomorphic ultrapowers of \mathcal{M} with respect to ultrafilters on \mathbb{N} . However, since there are $2^{\mathfrak{c}}$ many nonisomorphic nonprincipal ultrafilters on \mathbb{N} , it is a priori possible that there are $2^{\mathfrak{c}}$ many nonisomorphic ultrapowers of \mathcal{M} with respect to nonprincipal ultrafilters on \mathbb{N} . Farah and Shelah showed that this is indeed the case [55].

We split the proof of Theorem 15.2.3 into two parts. We first show that $\kappa(\varphi, \mathcal{M}^{\mathcal{U}})$ coincides with a cardinality that depends only on \mathcal{U} . First, some terminology and notation.

Fix $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. We let $<_{\mathcal{U}}$ denote the ordering on $\mathbb{N}^{\mathbb{N}}$ given by $f <_{\mathcal{U}} g$ if and only if $f(n) < g(n)$ for \mathcal{U} -almost all $n \in \mathbb{N}$. Note that $<_{\mathcal{U}}$ induces an ordering on $\mathbb{N}^{\mathcal{U}}$, also denoted $<_{\mathcal{U}}$, which is simply the interpretation of the

ordering symbol in the structure $(\mathbb{N}, <)^{\mathcal{U}}$. By Łoś's theorem, $<_{\mathcal{U}}$ is a linear ordering on $\mathbb{N}^{\mathcal{U}}$. We call $f \in \mathbb{N}^{\mathbb{N}}$ \mathcal{U} -infinite if $[f]_{\mathcal{U}}$ is an infinite element of $\mathbb{N}^{\mathcal{U}}$, that is, if for each $m \in \mathbb{N}$, we have that $f(n) > m$ for \mathcal{U} -almost all $n \in \mathbb{N}$. We also write $\lim_{\mathcal{U}} f = \infty$ in this situation. We let $\mathbb{N}_{\mathcal{U},\infty}^{\mathbb{N}}$ denote the set of \mathcal{U} -infinite elements of $\mathbb{N}^{\mathbb{N}}$. Note that $<_{\mathcal{U}}$ restricted to $\mathbb{N}_{\mathcal{U},\infty}^{\mathbb{N}}$ is not bounded below.

Definition 15.2.5. For $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$, we define the **lower cofinality** of \mathcal{U} , denoted $\text{lcf}(\mathcal{U})$, to be the coinitality of the ordering $<_{\mathcal{U}}$ restricted to $\mathbb{N}_{\mathcal{U},\infty}^{\mathbb{N}}$, that is, the minimal cardinality of a subset X of $\mathbb{N}_{\mathcal{U},\infty}^{\mathbb{N}}$ that is coinital in $\mathbb{N}_{\mathcal{U},\infty}^{\mathbb{N}}$, meaning that, for each $f \in \mathbb{N}_{\mathcal{U},\infty}^{\mathbb{N}}$, there is $g \in X$ such that $g <_{\mathcal{U}} f$.

In other words, viewing $\mathbb{N}^{\mathcal{U}}$ as a nonstandard model of \mathbb{N} , $\text{lcf}(\mathcal{U})$ is the coinitality of the infinite part of $\mathbb{N}^{\mathcal{U}}$. The first step toward proving Theorem 15.2.3 is the following:

Theorem 15.2.6. For each ultrafilter \mathcal{U} on \mathbb{N} , we have $\kappa(\varphi, \mathcal{M}^{\mathcal{U}}) = \text{lcf}(\mathcal{U})$.

Fix $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. In order to prove Theorem 15.2.6, we consider the following:

Situation (*)

- We have sets $Y_m \in \mathcal{U}$ such that $Y_0 = \mathbb{N}$, $Y_m \supseteq Y_{m+1}$, and $\bigcap_{m \in \mathbb{N}} Y_m = \emptyset$.
- $\Phi : \mathbb{N} \rightarrow \mathbb{N}$ is given by $\Phi(i) = m$ if $i \in Y_m \setminus Y_{m+1}$. Note that $\Phi \in \mathbb{N}_{\mathcal{U},\infty}^{\mathbb{N}}$.
- For each $i \in \mathbb{N}$, we have a φ -chain $a_0^i, \dots, a_{\Phi(i)}^i$ in \mathcal{M} of length $\Phi(i) + 1$.

We will really only be interested in the following two instances of Situation (*):

Example 15.2.7.

- (1) $Y_m = \{i \in \mathbb{N} : i \geq m\}$. In this case, $\Phi(i) = i$ and we let a_0^i, \dots, a_i^i be any φ -chain in \mathcal{M} of length $i+1$, which exists since φ is unstable.
- (2) Suppose we have fixed a φ -chain $([a_m]_{\mathcal{U}} : m \in \mathbb{N})$ in $\mathcal{M}^{\mathcal{U}1}$. In this case, we let

$$Y_m = \{i \in \mathbb{N} : i \geq m \text{ and } a_0(i), \dots, a_m(i) \text{ is a } \varphi\text{-chain in } \mathcal{M}\}.$$

It is clear that $Y_0 = \mathbb{N}$. By Łoś's theorem, each $Y_m \in \mathcal{U}$ and it is clear that $Y_m \supseteq Y_{m+1}$ and $\bigcap_{m \in \mathbb{N}} Y_m = \emptyset$. We then set $a_j^i := a_j(i)$ for $j \leq \Phi(i)$.

¹Technically, we only defined finite φ -chains, but the definition extends to infinite φ -chains in the obvious way

Now suppose that we are in Situation $(*)$ and $h \in \mathbb{N}^{\mathbb{N}}$. We define an element $a_h \in M^{\mathbb{N}}$ by declaring $a_h(i) := a_{h(i)}^i$ if $h(i) \leq \Phi(i)$, and otherwise $a_h(i) = a_{\Phi(i)}^i$. When h is the function that is constantly m , we denote the corresponding element of $M^{\mathbb{N}}$ by a_m .

Remark 15.2.8. There could potentially be some confusion in the notation of the previous paragraph in the case that we are in Example 15.2.7(2) above. As we start with a sequence $([a_m]_{\mathcal{U}})$, use this to define the elements a_k^i for $k \leq \Phi(i)$, and then use this latter sequence to define a sequence in $\mathcal{M}^{\mathcal{U}}$ also denoted $([a_m]_{\mathcal{U}})$. However, it is clear that both definitions of $a_m(i)$ agree when $m \leq \Phi(i)$; since $\lim_{\mathcal{U}} \Phi = \infty$, we see that both definitions of a_m are equivalent modulo \mathcal{U} , whence no confusion can arise.

In the next two lemmas, suppose that we are in Situation $(*)$ and the elements a_h are defined as above.

Lemma 15.2.9. *Suppose that $g, h \in \mathbb{N}^{\mathbb{N}}$ are such that $g, h \leq_{\mathcal{U}} \Phi$ and $g \not\leq_{\mathcal{U}} h$. Then $g <_{\mathcal{U}} h$ if and only if $[a_g]_{\mathcal{U}} <_{\varphi} [a_h]_{\mathcal{U}}$.*

Proof. First suppose that $g <_{\mathcal{U}} h$. Consider $i \in \mathbb{N}$ such that $g(i) < h(i) \leq \Phi(i)$. It follows that $a_g(i) = a_{g(i)}^i <_{\varphi} a_{h(i)}^i = a_h(i)$; since there is a large set of such i 's, we get $[a_g]_{\mathcal{U}} <_{\varphi} [a_h]_{\mathcal{U}}$.

Now suppose that $g \not<_{\mathcal{U}} h$. Then since $<_{\mathcal{U}}$ is a linear order, we have that $h <_{\mathcal{U}} g$. By the first paragraph, we have that $[a_h]_{\mathcal{U}} <_{\varphi} [a_g]_{\mathcal{U}}$, whence we have that $[a_g]_{\mathcal{U}} \not<_{\varphi} [a_h]_{\mathcal{U}}$. □

Lemma 15.2.10. *Suppose that $[b]_{\mathcal{U}} \in \mathcal{M}^{\mathcal{U}}$ is such that $[a_m]_{\mathcal{U}} <_{\varphi} [b]_{\mathcal{U}}$ for all $m \in \mathbb{N}$. Then there is $h \in \mathbb{N}_{\mathcal{U}, \infty}^{\mathbb{N}}$ such that $[a_m]_{\mathcal{U}} <_{\varphi} [a_h]_{\mathcal{U}} <_{\varphi} [b]_{\mathcal{U}}$ for all $m \in \mathbb{N}$. Moreover, if $g \in \mathbb{N}^{\mathbb{N}}$ is such that $g \leq_{\mathcal{U}} \Phi$ and $[b]_{\mathcal{U}} <_{\varphi} [a_g]_{\mathcal{U}}$, then $[a_h]_{\mathcal{U}} <_{\varphi} [a_g]_{\mathcal{U}}$ as well.*

Proof. For each $m \in \mathbb{N}$, set

$$X_m := \{i \in Y_m : (\forall k \leq m)(a_k(i) <_{\varphi} b(i))\},$$

which belongs to \mathcal{U} by assumption. Note also that $X_{m+1} \subseteq X_m$ and $\bigcap_m X_m = \emptyset$. For $i \in X_m \setminus X_{m+1}$, set $h(i) = m$. If $i \notin X_0$, set $h(i) = 0$. Note that, for $i \in X_m$, we have $h(i) \geq m$, whence it follows that $\lim_{\mathcal{U}} h = \infty$. Note also that $h \leq \Phi$.

We claim that this h is as desired. Since $\lim_{\mathcal{U}} h = \infty$, we have that $m <_{\mathcal{U}} h$ and thus $[a_m]_{\mathcal{U}} <_{\varphi} [a_h]_{\mathcal{U}}$ by Lemma 15.2.9. Moreover, for $i \in X_0$, $a_h(i) = a_{h(i)}^i <_{\varphi} b(i)$, whence $[a_h]_{\mathcal{U}} < [b]_{\mathcal{U}}$.

Finally, suppose that $g \in \mathbb{N}^{\mathbb{N}}$ is such that $g \leq_{\mathcal{U}} \Phi$ and $[b]_{\mathcal{U}} <_{\varphi} [a_g]_{\mathcal{U}}$ and yet, toward a contradiction, we have that $[a_g]_{\mathcal{U}} <_{\varphi} [a_h]_{\mathcal{U}}$. By Lemma 15.2.9,

we have that $g <_{\mathcal{U}} h$ and so $a_{g(i)}(i) <_{\varphi} b(i)$ for \mathcal{U} -almost all i , contradicting $b(i) <_{\varphi} a_{g(i)}(i)$ for \mathcal{U} -almost all i . \square

Proof of Theorem 15.2.6. We first show that $\kappa(\varphi, \mathcal{M}^{\mathcal{U}}) \leq \text{lcf}(\mathcal{U})$. Let $(h_{\gamma})_{\gamma < \text{lcf}(\mathcal{U})}$ be a decreasing coinital sequence for $<_{\mathcal{U}}$. Since \mathcal{M} is unstable, we know that we are in Situation (*) by Example 15.2.7(1). By Lemma 15.2.9, $([a_m]_{\mathcal{U}})_{m < \omega}$ and $([a_{h_{\gamma}}]_{\mathcal{U}})_{\gamma < \text{lcf}(\mathcal{U})}$ satisfy the first two properties of being a φ - $\text{lcf}(\mathcal{U})$ -gap; such sequences are often referred to as **pregaps**. To verify that this pregap is indeed a gap, suppose that $[a_m]_{\mathcal{U}} <_{\varphi} [c]_{\mathcal{U}}$. By Lemma 15.2.10, there is $h \in \mathbb{N}_{\mathcal{U}, \infty}^{\mathbb{N}}$ such that $[a_m]_{\mathcal{U}} <_{\varphi} [a_h]_{\mathcal{U}} <_{\varphi} [c]_{\mathcal{U}}$. Since our sequence is coinital, there is γ such that $h_{\gamma} <_{\mathcal{U}} h$. By Lemma 15.2.9, $[a_m]_{\mathcal{U}} <_{\varphi} [a_{h_{\gamma}}]_{\mathcal{U}} <_{\varphi} [a_h]_{\mathcal{U}}$. By the proof of Lemma 15.2.10, we have that $[a_{h_{\gamma}}]_{\mathcal{U}} <_{\varphi} [c]_{\mathcal{U}}$. It follows that $([a_m]_{\mathcal{U}})_{m < \omega}$ and $([a_{h_{\gamma}}]_{\mathcal{U}})_{\gamma < \text{lcf}(\mathcal{U})}$ form a φ - $\text{lcf}(\mathcal{U})$ -gap, whence $\kappa(\varphi, \mathcal{M}^{\mathcal{U}}) \leq \text{lcf}(\mathcal{U})$.

For the other direction, suppose that $([a_m]_{\mathcal{U}})_{m < \omega}$ and $([b_{\gamma}]_{\mathcal{U}})_{\gamma < \kappa(\varphi, \mathcal{M}^{\mathcal{U}})}$ form a gap. We use these $[a_m]_{\mathcal{U}}$'s to put us in Situation (*) as in Example 15.2.7(2). We define, by recursion on $\gamma < \kappa(\varphi, \mathcal{M}^{\mathcal{U}})$, a $<_{\mathcal{U}}$ -decreasing sequence h_{γ} from $\mathbb{N}_{\mathcal{U}, \infty}^{\mathbb{N}}$ as follows: Let $h_0 \in \mathbb{N}_{\mathcal{U}, \infty}^{\mathbb{N}}$ be any function with $h \leq_{\mathcal{U}} \Phi$. Now suppose that h_{δ} has been defined for all $\delta < \gamma$. By minimality of $\kappa(\varphi, \mathcal{M}^{\mathcal{U}})$, there is $[b]_{\mathcal{U}} \in \mathcal{M}^{\mathcal{U}}$ such that $[a_m]_{\mathcal{U}} <_{\varphi} [b]_{\mathcal{U}} < [a_{h_{\delta}}]_{\mathcal{U}}$ for all $m \in \mathbb{N}$ and $\delta < \gamma$. By Lemma 15.2.10, there is $g_1 \in \mathbb{N}_{\mathcal{U}, \infty}^{\mathbb{N}}$ such that $[a_{g_1}]_{\mathcal{U}} < [b]_{\mathcal{U}}$. By Lemma 15.2.10 again, there is $g_2 \in \mathbb{N}_{\mathcal{U}, \infty}^{\mathbb{N}}$ such that $[a_{g_2}]_{\mathcal{U}} <_{\varphi} [b_{\gamma}]_{\mathcal{U}}$. Let $h_{\gamma} := \min(g_1, g_2)$. Then $h_{\gamma} \in \mathbb{N}_{\mathcal{U}, \infty}^{\mathbb{N}}$ is such that $h_{\gamma} <_{\mathcal{U}} h_{\delta}$ for all $\delta < \gamma$ by Lemma 15.2.10, and $[h_{\gamma}]_{\mathcal{U}} <_{\varphi} [b_{\gamma}]_{\mathcal{U}}$.

We claim that the sequence $(h_{\gamma})_{\gamma < \kappa(\varphi, \mathcal{M}^{\mathcal{U}})}$ is coinital. Indeed, suppose that $h \in \mathbb{N}_{\mathcal{U}, \infty}^{\mathbb{N}}$ is such that $h <_{\mathcal{U}} h_{\gamma}$ for all $\gamma < \kappa(\varphi, \mathcal{M}^{\mathcal{U}})$. By construction, that implies that $[a_h]_{\mathcal{U}} <_{\varphi} [b_{\gamma}]_{\mathcal{U}}$ for all γ , contradicting that $([a_m]_{\mathcal{U}})_{m < \omega}$ and $([b_{\gamma}]_{\mathcal{U}})_{\gamma < \kappa(\varphi, \mathcal{M}^{\mathcal{U}})}$ formed a gap. It follows that $\text{lcf}(\mathcal{U}) \leq \kappa(\varphi, \mathcal{M}^{\mathcal{U}})$, as desired. \square

This completes the first part of the proof of Theorem 15.2.3. Before moving onto the second part of the proof, we make a quick digression, returning to the subject of Keisler's order from Section 8.6 and making good on a promise made after Corollary 8.6.29.

First, we need:

Exercise 15.2.11. Suppose that T is unstable and \mathcal{U} is an ultrafilter (not necessarily on \mathbb{N}). Prove that, for any $\mathcal{M} \models T$, we have that $\mathcal{M}^{\mathcal{U}}$ is not $\text{lcf}(\mathcal{U})^+$ -saturated. (*Hint.* The proof of Theorem 15.2.6 above works for any ultrafilter \mathcal{U} , not just those on \mathbb{N} .)

We also need the following generalization of Theorem 8.6.25:

Theorem 15.2.12 (Shelah). *For any infinite cardinals κ , λ , and ν satisfying $\lambda^{\aleph_0} = \lambda \leq 2^\kappa$ and $\aleph_0 < \nu \leq \lambda$, there is a regular ultrafilter \mathcal{U} on κ such that $\text{pcf}(\mathcal{U}) = \lambda$ and $\text{lcf}(\mathcal{U}) = \nu$.*

Here is the promised improvement of Corollary 8.6.29:

Corollary 15.2.13. *Suppose that T_1 and T_2 are countable theories such that T_1 is stable and T_2 is unstable. Then $T_1 \triangleleft T_2$.*

Proof. Without loss of generality, we may assume that T_1 has the fcp. Fix an uncountable cardinal κ and let \mathcal{U} be a regular ultrafilter on κ such that $\text{pcf}(\mathcal{U}) = 2^\kappa$ and such that $\text{lcf}(\mathcal{U}) \leq \kappa$; this is possible by Theorem 15.2.12. Then by Theorem 8.6.20, we have that \mathcal{U} saturates T_1 . However, by Exercise 15.2.11, \mathcal{U} does not saturate T_2 . \square

In order to complete the proof of Theorem 15.2.3, it suffices to prove:

Theorem 15.2.14. *For each regular cardinal κ with $\aleph_1 \leq \kappa \leq \mathfrak{c}$, there is $\mathcal{U} \in \beta\mathbb{N}$ such that $\text{lcf}(\mathcal{U}) = \kappa$.*

Remark 15.2.15. In Theorem 15.2.14, κ must be assumed to be regular as coinitality is always a regular cardinal.

The reader is warned that the proof of Theorem 15.2.14 is fairly technical. The key to proving this theorem is to build the ultrafilter slowly as in the construction of good ultrafilters.

Let $F \subseteq \omega^\omega$ be of large oscillation modulo the cofinite filter with $|F| = \kappa$ (see Definition 8.5.2); this is possible by Lemma 8.5.6. We enumerate $F = (f_\alpha)_{\alpha < \kappa}$. Given $f \in F$, we define $g_f : \beta\omega \rightarrow \omega + 1$ by declaring $g_f(\mathcal{U}) = n$ if $f(m) = n$ for \mathcal{U} -almost all m ; if no such n exists (that is, if f is \mathcal{U} -infinite), then we declare $g_f(\mathcal{U}) = \omega$. We put all of these functions into a single function $G : \beta\omega \rightarrow (\omega + 1)^F$, that is, $G(\mathcal{U})(f) := g_f(\mathcal{U})$.

Exercise 15.2.16.

- (1) Prove that G is continuous.
- (2) Prove that $G(\beta\omega \setminus \omega) = (\omega + 1)^F$. (*Hint.* This uses the assumption that F is of large oscillation modulo the cofinite filter.)

By Proposition 3.1.15 and Exercise 15.2.16, there is a closed subset K of $\beta\omega \setminus \omega$ such that $G(K) = (\omega + 1)^F$ but $G(K') \neq (\omega + 1)^F$ for any proper closed subset K' of K . In the rest of this section, we fix such a closed subset K of $\beta\omega \setminus \omega$.

For $C \subseteq \kappa$ and $x, y \in (\omega + 1)^F$, we write $x \equiv_C y$ if $x(f_\alpha) = y(f_\alpha)$ for all $\alpha \in C$.

Definition 15.2.17. For $A \subseteq \omega$, a **support** for A is a subset $C \subseteq \kappa$ such that, for any $x, y \in (\omega + 1)^F$, if $x \equiv_C y$, then $x \in G(U_A \cap K)$ if and only if $y \in G(U_A \cap K)$.

Given the graph s of a finite partial function from F to $\omega + 1$, we let $[s]$ denote all those elements of $(\omega + 1)^F$ which extend s . Given $A \subseteq \omega$, let $\text{supp}(A)$ denote the set of $\alpha < \kappa$ for which there are s and n such that $[s] \not\subseteq G(U_A \cap K)$ but $[s \cup (f_\alpha, n)] \subseteq G(U_A \cap K)$.

Lemma 15.2.18. $\text{supp}(A)$ is the smallest support for A .

Proof. We first show that $\text{supp}(A)$ is contained in every support for A . Suppose that $C \subseteq \kappa$ is a support for A and yet, toward a contradiction, that $\text{supp}(A) \not\subseteq C$. Take $\alpha \in \text{supp}(A) \setminus C$. Let s and n witness that $\alpha \in \text{supp}(A)$. Take $y \in [s] \setminus G(U_A \cap K)$. Let $x \in (\omega + 1)^F$ be such that $x(f_\alpha) = n$ and $x(f_\beta) = y(f_\beta)$ for all $\beta \neq \alpha$. Since $s \cup \{(f_\alpha, n)\} \subseteq x$, we have that $x \in G(U_A \cap K)$. Since $\alpha \notin C$, we have $x \equiv_C y$, whence $y \in G(U_A \cap K)$ since C is a support for A , yielding a contradiction.

We now show that $\text{supp}(A)$ is a support for A .

Claim. Suppose that s witnesses that $\alpha \in \text{supp}(A)$. If $f_\beta \in F$ is such that $\beta \notin \text{supp}(A)$, then $s \setminus \{(f_\beta, s(f_\beta))\}$ also witnesses that $\alpha \in \text{supp}(A)$.

Proof of Claim. Note that $[s \setminus \{(f_\beta, s(f_\beta))\}] \not\subseteq G(U_A \cap K)$. If

$$[s \setminus \{(f_\beta, s(f_\beta))\} \cup (f_\alpha, n)] \not\subseteq G(U_A \cap K),$$

then since $[s \cup (f_\alpha, n)] \subseteq G(U_A \cap K)$, we would get that $\beta \in \text{supp}(A)$. Consequently, $[s \setminus \{(f_\beta, s(f_\beta))\} \cup (f_\alpha, n)] \subseteq G(U_A \cap K)$, as desired.

By the claim, given any witness s to the fact that $\alpha \in \text{supp}(A)$, one can always find a subset $s' \subseteq s$ such that s' also witnesses that $\alpha \in \text{supp}(A)$ and with $\text{dom}(s') \subseteq \text{supp}(A)$.

Now suppose that $x \equiv_{\text{supp}(A)} y$ and $x \in G(U_A \cap K)$. Since $\bigcap_{s \subseteq x} [s] \setminus G(U_A \cap K) = \emptyset$, by compactness there is some $s \subseteq x$ such that $[s] \setminus G(U_A \cap K) \neq \emptyset$. Since s is finite and $[\emptyset] \setminus G(U_A \cap K) = \emptyset$, we may as well assume that s is such that there is some $\alpha < \kappa$ for which $[s \cup (f_\alpha, x(f_\alpha))] \subseteq G(U_A \cap K)$. By the above fact, we may also assume that $\text{dom}(s) \subseteq \text{supp}(A)$. Consequently, $y \in G(U_A \cap K)$, as desired. \square

Lemma 15.2.19. Given $A \subseteq \omega$, there is a countable support for A . Consequently, $\text{supp}(A)$ is countable.

Proof. Let $S := \bigcup_{H \in \mathcal{P}_f(F)} (\omega + 1)^H$. Let $T \subseteq S$ be maximal such that $T' := \{[t] : t \in T\}$ is a maximal collection of pairwise disjoint subsets of $G(U_A \cap K)$. Note that T is countable.

Claim 1. $\bigcup T'$ is dense in $(\omega + 1)^F \setminus G(K \setminus U_A)$.

Proof of Claim 1. Fix $f \in (\omega + 1)^F \setminus G(K \setminus U_A)$ and take $s \in S$ with $s \subseteq f$ such that $[s] \subseteq (\omega + 1)^F \setminus G(K \setminus U_A)$. Note that $[s] \subseteq G(U_A \cap K)$, otherwise, there is $g \in [s]$ such that $g \notin G(U_A \cap K)$, whence $g \in G(K \setminus U_A)$, contradicting that $[s] \subseteq (\omega + 1)^F \setminus G(K \setminus U_A)$. It follows that $[s] \cap \bigcup T' \neq \emptyset$, otherwise we contradict the maximality of T' .

Claim 2. $G(G^{-1}(\overline{\bigcup T'}) \cap K) \cup G(K \setminus U_A) = (\omega + 1)^F$.

Proof of Claim 2. Suppose $f \in (\omega + 1)^F \setminus G(K \setminus U_A)$. By Claim 1, there is a net (g_i) from $\bigcup T'$ such that $f = \lim g_i$. Let $\mathcal{U}_i \in K$ be such that $G(\mathcal{U}_i) = g_i$. Since K is compact, after passing to a subnet, we might as well assume that $\mathcal{U}_i \rightarrow \mathcal{U} \in K$. It follows that $g_i = G(\mathcal{U}_i) \rightarrow G(\mathcal{U}) \in \overline{\bigcup T'}$, whence $\mathcal{U} \in G^{-1}(\overline{\bigcup T'}) \cap K$. Since $G(\mathcal{U}) = f$, the claim is proven.

Claim 3. $G(U_A \cap K) = \overline{\bigcup T'}$.

Proof of Claim 3. By Claim 2, $G((G^{-1}(\overline{\bigcup T'}) \cap K) \cup (K \setminus U_A)) = (\omega + 1)^F$. Since $G^{-1}(\overline{\bigcup T'}) \cap K$ and $K \setminus U_A$ are closed, by the choice of K , we have $G^{-1}(\overline{\bigcup T'}) \cap K \cup (K \setminus U_A) = K$. Now since $\bigcup T' \subseteq (\omega + 1)^F \setminus G(K \setminus U_A)$, we have that $\overline{\bigcup T'} \subseteq G(U_A \cap K)$. On the other hand, if $\mathcal{U} \in U_A \cap K$, then $\mathcal{U} \in G^{-1}(\overline{\bigcup T'}) \cap K$, whence $G(\mathcal{U}) \in \overline{\bigcup T'}$, as desired.

For each $t \in T$, let $H_t := \text{dom}(t)$. Set $\text{supp}_T(A) := \{\alpha < \kappa : f_\alpha \in \bigcup_{t \in T} H_t\}$. Since T is countable and each H_t is finite, we have that $\text{supp}_T(A)$ is countable. It remains to prove:

Claim 4. $\text{supp}_T(A)$ is a support for A .

Proof of Claim 4. Suppose that $x \equiv_{\text{supp}_T(A)} y$ and $x \in G(U_A \cap K)$. By Claim 3, it suffices to show that $y \in \overline{\bigcup T'}$. Fix $\alpha < \kappa$ and set $s := y|_{\{f_\alpha\}}$. We want to find an element of $\bigcup T'$ which agrees with y on f_α . First suppose that $\alpha \in \text{supp}_T(A)$, say $\alpha \in H_t$. Since $x(f_\alpha) = y(f_\alpha)$ and $x \in G(U_A \cap K) = \overline{\bigcup T'}$, we are done in this case. Otherwise, suppose $\alpha \notin \text{supp}_T(A)$. Let $z \in \bigcup T'$ be such that $z(f_\alpha) = x(f_\alpha)$. Let $z^* \in (\omega + 1)^F$ be such that $z^*(f_\alpha) = y(f_\alpha)$ while $z^*(f_\beta) = z(f_\beta)$ for all $\beta \neq \alpha$. Since $\alpha \notin \text{supp}_T(A)$, if $t \subseteq z$ with $t \in T$, then we still have that $t \subseteq z^*$, whence $z^* \in \bigcup T'$, and again we are done. \square

Before explaining the construction of \mathcal{U} , we need a bit more notation and terminology. First, we set $F_\alpha := (f_\delta)_{\alpha \leq \delta < \kappa}$. For a filter \mathcal{F} on ω and $h \in \omega^\omega$, we say that h is \mathcal{F} -infinite if, for every $m \in \omega$, we have that $\{n \in \omega : h(n) > m\} \in \mathcal{F}$; this extends our earlier terminology of \mathcal{U} -infinite to the case of arbitrary filters. We also write $f <_{\mathcal{F}} g$ if $\{n \in \omega : f(n) < g(n)\} \in \mathcal{F}$.

We proceed by constructing an increasing sequence of filters \mathcal{F}_α on ω for $\alpha < \kappa$ satisfying the following properties.

- (1) If $A \in \mathcal{F}_\alpha$, then $\text{supp}(A) \subseteq \alpha$.
- (2) If $A \subseteq \omega$ is such that $\text{supp}(A) \subseteq \alpha$, then either $A \in \mathcal{F}_{\alpha+1}$ or $\omega \setminus A \in \mathcal{F}_{\alpha+1}$.
- (3) F_α is of large oscillation modulo \mathcal{F}_α .
- (4) f_α is $\mathcal{F}_{\alpha+1}$ -infinite.
- (5) For every $h \in H_\alpha := \{h \in \omega^\omega : h \text{ is } \mathcal{F}_\alpha\text{-infinite}\}$, we have $f_\alpha <_{\mathcal{F}_{\alpha+1}} h$.

We start the construction by defining $\mathcal{F}_0 := \{A \subseteq \omega : K \subseteq U_A\}$. We leave it to the reader to verify that properties (1) and (3) are indeed satisfied.

Now suppose that \mathcal{F}_γ has been defined for all $\gamma < \alpha$. It is easy to see that if α is a limit ordinal, then setting $\mathcal{F}_\alpha := \bigcup_{\gamma < \alpha} \mathcal{F}_\gamma$ is as desired.

Now suppose that \mathcal{F}_α has been constructed; we show how to construct $\mathcal{F}_{\alpha+1}$. We first let $\mathcal{F}'_{\alpha+1}$ denote the filter generated by \mathcal{F}_α together with the sets of the form $\bigcup_{n > m} (f_\alpha^{-1}(n) \cap h^{-1}(n, \omega))$, for $m \in \omega$ and $h \in H_\alpha$.

Exercise 15.2.20. Verify the following facts about $\mathcal{F}'_{\alpha+1}$:

- (1) $\mathcal{F}'_{\alpha+1}$ is a proper filter.
- (2) f_α is $\mathcal{F}'_{\alpha+1}$ -infinite.
- (3) $f_\alpha <_{\mathcal{F}'_{\alpha+1}} h$ for all $h \in H_\alpha$.
- (4) For every $A \in \mathcal{F}'_{\alpha+1}$, $\alpha + 1$ is a support for A , whence $\text{supp}(A) \subseteq \alpha + 1$.

We now let $\mathcal{F}_{\alpha+1}$ be an extension of $\mathcal{F}'_{\alpha+1}$ maximal with respect to the property that every element has its minimal support contained in $\alpha + 1$. By the previous exercise, it follows that all of the properties needed of $\mathcal{F}_{\alpha+1}$ are true except for property (3), which we verify now.

Claim. $F_{\alpha+1}$ is of large oscillation modulo $\mathcal{F}_{\alpha+1}$.

Proof of Claim. Take $A \in \mathcal{F}_{\alpha+1}$, distinct $\delta_1, \dots, \delta_n \geq \alpha + 1$, and $m_1, \dots, m_n \in \omega$. Since $\mathcal{F}_0 \subseteq \mathcal{F}'_{\alpha+1}$, there is $x \in G(U_A \cap K)$. Let $y \in (\omega + 1)^F$ be defined by setting $y(f_\delta) = x(f_\delta)$ for $\delta \leq \alpha$ and $y(f_{\delta_i}) = m_i$ for $i = 1, \dots, n$ (and defined on the other members of F arbitrarily). Since $\text{supp}(A) \subseteq \alpha + 1$, it follows that $y \in G(U_A \cap K)$. Also, $y \in G(\bigcap_{i=1}^n U_{f_{\delta_i}^{-1}(m_i)})$, and it is clear that this then implies that $G^{-1}(y) \subseteq \bigcap_{i=1}^n U_{f_{\delta_i}^{-1}(m_i)}$. It follows that $U_A \cap \bigcap_{i=1}^n U_{f_{\delta_i}^{-1}(m_i)} \neq \emptyset$, as desired.

Set $\mathcal{U} := \bigcup_{\alpha < \kappa} \mathcal{F}_\alpha$. By Exercise 15.2.20(2), \mathcal{U} is actually an ultrafilter on ω . (Note that this uses the regularity of κ : for any $A \subseteq \omega$, $\text{supp}(A) \subseteq \alpha$ for some $\alpha < \kappa$ and then $\mathcal{F}_{\alpha+1}$ decides A .) We claim that $\text{lcf}(\mathcal{U}) = \kappa$.

First suppose that $H \subseteq \omega^\omega$ is such that $|H| < \kappa$ and each $h \in H$ is \mathcal{U} -infinite. Then by the regularity of κ again, there is some $\alpha < \kappa$ such that, for each $h \in H$ and each $n \in \omega$, $\text{supp}(h^{-1}(n, \omega)) \subseteq \alpha$. (This uses that each such minimal support is countable.) Since $\mathcal{F}_{\alpha+1}$ must then decide each $\text{supp}(h^{-1}(n, \omega))$ and h is \mathcal{U} -infinite, we must have that h is $\mathcal{F}_{\alpha+1}$ -infinite. Since $f_{\alpha+1}$ is $\mathcal{F}_{\alpha+2}$ -infinite and $f_{\alpha+1} <_{\mathcal{F}_{\alpha+2}} h$ for every $h \in H_{\alpha+1}$, it follows that H is not coinital, whence $\text{lcf}(\mathcal{U}) \geq \kappa$.

On the other hand, it is clear that $f_{\alpha+1} <_{\mathcal{U}} f_\alpha$ and the proof of the previous paragraph (in the case that $|H| = 1$) shows that F forms a coinital sequence, whence $\text{lcf}(\mathcal{U}) \leq \kappa$, and thus $\text{lcf}(\mathcal{U}) = \kappa$. This concludes the proof of Theorem 15.2.14 and thus the proof of Theorem 15.2.3.

15.3. The case when \mathcal{M} is stable

In this section, we sketch a proof of Theorem 15.1.6(2), freely making use of important facts from the area of model theory known as *stability theory*. Clearly, Theorem 15.1.6(2) follows from Corollary 6.8.4, Theorem 8.1.9, and the following theorem:

Theorem 15.3.1. *Suppose that \mathcal{M} is a stable structure in a countable language and $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. Then $\mathcal{M}^{\mathcal{U}}$ is \mathfrak{c} -saturated.*

By Theorem 8.2.1, $\mathcal{M}^{\mathcal{U}}$ is always \aleph_1 -saturated. The importance of the previous theorem is that, in case \mathcal{M} is stable, the level of saturation can be improved to \mathfrak{c} -saturation (which is only an actual improvement if CH fails).

To prove Theorem 15.3.1, we will need to use one of the most important facts of stability theory, namely that a stable theory possesses a well-behaved notion of independence generalizing linear independence in vector spaces and algebraic independence in algebraically closed fields.

In the rest of this section, we assume that \mathcal{M} is a stable structure and we fix a \mathfrak{c} -saturated elementary extension \mathbb{M} of $\mathcal{M}^{\mathcal{U}}$. By a *small* subset of \mathbb{M} , we mean a subset of \mathbb{M} of size $< \mathfrak{c}$.

For small subsets A and B of \mathbb{M} , we consider the **type of A over B** , which is the set

$$\text{tp}(A/B) := \{\varphi(x) : \varphi \text{ is an } L_B\text{-formula and } \mathbb{M} \models \varphi(A)\}.$$

This notation is a bit sloppy: we really should fix an enumeration $A = (a_i)_{i < \lambda}$, where $\lambda < \mathfrak{c}$, and then $\varphi(x_{i_1}, \dots, x_{i_k})$ is a formula with parameters from B and $i_1, \dots, i_k < \lambda$, and $\mathbb{M} \models \varphi(a_{i_1}, \dots, a_{i_k})$.

When $B = \emptyset$, we simply write $\text{tp}(A)$ instead of $\text{tp}(A/\emptyset)$.

Fact 15.3.2. There is a relation \perp defined on triples of small subsets of \mathbb{M} satisfying the following properties:

- (1) *Invariance.* If $\text{tp}(ABC) = \text{tp}(A'B'C')$, then $A \perp_C B$ if and only if $A' \perp_{C'} B'$.
- (2) *Symmetry.* $A \perp_C B$ if and only if $B \perp_C A$.
- (3) *Transitivity.* $A \perp_C BD$ if and only if $A \perp_C B$ and $A \perp_{BC} D$.
- (4) *Finite character.* $A \perp_C B$ if and only if $a \perp_C B$ for all finite tuples a from A .
- (5) *Extension.* For all A, B, C , there is A' such that $\text{tp}(A/C) = \text{tp}(A'/C)$ and $A' \perp_C B$.
- (6) *Local character.* For any finite tuple a and any B , there is a countable $B_0 \subseteq B$ such that $a \perp_{B_0} B$.
- (7) *Stationarity of types.* For all A, A', B , and all small elementary submodels N , if $\text{tp}(A/N) = \text{tp}(A'/N)$, $A \perp_N B$, and $A' \perp_N B$, then $\text{tp}(A/BN) = \text{tp}(A'/BN)$.

Here, we use the usual model-theoretic convention of writing BC instead of $B \cup C$.

Remark 15.3.3. One should think of $A \perp_C B$ as saying that one learns no new information about A using BC than one already had using C alone. With this in mind, most of the above properties should be more or less intuitive, except perhaps for stationarity of types. That being said, it is stationarity of types that, in some sense, differentiates the stable theories from other theories which possess relations satisfying the other six axioms.

Besides the previous fact, we will also need to know that certain nice kinds of sequences exist in \mathbb{M} .

Definition 15.3.4. Suppose that I is a linearly ordered set and $(a_i)_{i \in I}$ is a sequence from \mathbb{M} . Suppose also that A is a small subset of \mathbb{M} . We say that $(a_i)_{i \in I}$ is:

- (1) an **indiscernible sequence over A** (or **A -indiscernible**) if, for all $n \in \mathbb{N}$ and all $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$ from I , we have $\text{tp}(a_{i_1} \dots a_{i_n}/A) = \text{tp}(a_{j_1} \dots a_{j_n}/A)$;
- (2) an **independent sequence over A** (or **A -independent**) if, for all $i \in I$, we have $a_i \perp_A a_{<i}$, where $a_{<i} := \{a_j : j < i\}$;
- (3) a **Morley sequence over A** if it is both an indiscernible sequence over A and an independent sequence over A . Moreover, if $p(x)$ is the common type of the elements of $(a_i)_{i \in I}$ over A , we say that $(a_i)_{i \in I}$ is a **Morley sequence in p** .

Here is the crucial fact that we will need:

Fact 15.3.5. Suppose that I is a small linearly ordered set, A is a small subset of \mathbb{M} , and $p(x)$ is a 1-type over A . Then there is a sequence $(a_i)_{i \in I}$ from \mathbb{M} that is a Morley sequence in p .

Exercise 15.3.6. Suppose that $(a_i)_{i \in I}$ is a Morley sequence in p and

$$\text{tp}((a_i)_{i \in I}/A) = \text{tp}((b_i)_{i \in I}/A).$$

Then $(b_i)_{i \in I}$ is also a Morley sequence in p .

Exercise 15.3.7. Suppose that $(a_i)_{i \in I}$ is an independent sequence over A .

- (1) Suppose that $J, K \subseteq I$ are such that $J < K$, that is, $j < k$ for all $j \in J$ and $k \in K$. Prove that $a_K \perp_A a_J$, where $a_J := \{a_j : j \in J\}$ and $a_K := \{a_k : k \in K\}$.
- (2) Prove that $(a_i)_{i \in I}$ is also an independent set over A , that is, for any $i \in I$, we have $a_i \perp_A a_{\neq i}$, where $a_{\neq i} := \{a_j : j \neq i\}$.

Exercise 15.3.8. Suppose that $A \subseteq M^{\mathcal{U}}$ is countable and $(a_n)_{n \in \mathbb{N}}$ is a Morley sequence over A whose common type is p . Prove that there is a set $D \subseteq M^{\mathcal{U}}$ that is an A -independent set of realizations of p of size \mathfrak{c} . (*Hint.* By Exercise 15.3.7(2), $(a_n)_{n \in \mathbb{N}}$ is an A -independent set. Let $A_n \subseteq M$ be finite of increasing size realizing larger fragments of $\text{tp}((a_n)_{n \in \mathbb{N}}/A)$. Let $D = \prod_{\mathcal{U}} A_n$ and argue that D is an A -independent set of realizations of p . Use Theorem 6.8.4 to conclude that $|D| = \mathfrak{c}$.)

Proof of Theorem 15.3.1. Suppose that $B \subseteq M^{\mathcal{U}}$ is such that $|B| < \mathfrak{c}$ and suppose that $q(x)$ is a set of \mathcal{L}_B -formulae that is finitely satisfiable in $M^{\mathcal{U}}$. We need to show that q can be realized in $M^{\mathcal{U}}$. Using upward Löwenheim-Skolem (and enlarging B if necessary), we may assume that B is an elementary submodel of $M^{\mathcal{U}}$.

Since \mathbb{M} is \mathfrak{c} -saturated, we may find $c \in \mathbb{M}$ which realizes q . By local character, there is a countable $B_0 \subseteq B$ such that $c \perp_{B_0} B$. Let B_1 be a countable elementary submodel of B containing B_0 . By transitivity, $c \perp_{B_1} B$. Set p to be the restriction of q to B_1 . By Fact 15.3.5, there is a countable sequence from \mathbb{M} that is a Morley sequence in p . By \aleph_1 -saturation of $M^{\mathcal{U}}$ and Exercise 15.3.6, there is a sequence $(a_n)_{n \in \mathbb{N}}$ from $M^{\mathcal{U}}$ that is a Morley sequence in p . By Exercise 15.3.8, there is $D \subseteq M^{\mathcal{U}}$ that is a B_1 -independent set and $|D| = \mathfrak{c}$.

The rest of the proof proceeds by “forking calculus.” By finite character, local character, transitivity, and the fact that $|B| < \mathfrak{c}$, there is $D_0 \subseteq D$ with $|D_0| < \mathfrak{c}$ such that $B \perp_{D_0} D$. Fix $a \in D \setminus D_0$; we show that a realizes q , finishing the proof of the theorem. By transitivity and symmetry, we have

that $a \downarrow_{D_0} B$. Using that D is a B_1 -independent set and transitivity, we also have $a \downarrow_{B_1} D_0$. By transitivity again, we have that $a \downarrow_{B_1} B$. Recall also that $c \downarrow_{B_1} B$ and $\text{tp}(a/B_1) = \text{tp}(c/B_1) = p$. By stationarity of types, we have that $\text{tp}(a/B) = \text{tp}(c/B)$. Recalling that c realized q , we see that a realizes q , as desired. \square

15.4. Notes and references

Our proof of Theorem 15.2.6 follows that developed in [53] (which was written for operator algebraists). Theorem 15.2.14 is due to Dow [46], and we follow that article quite closely. Our proof of Theorem 15.3.1 follows the proof given in [54, Theorem 5.6(1)]. The requisite stability theory needed in the proof of Theorem 15.3.1 can be found in the textbook [174].

The Keisler-Shelah theorem

In this chapter, we prove the Keisler-Shelah theorem, which states that two structures (in the same language) are elementarily equivalent if and only if they have isomorphic ultrapowers. This was proven by Keisler using the GCH and we gave this argument in Section 8.4. The proof presented in Section 16.1, which makes no extra set-theoretic assumptions, is due to Shelah [157]. In the remaining sections we give a few sample applications of the Keisler-Shelah theorem: Section 16.2 gives a soft criteria for when a collection of structures in a given language is axiomatizable, Section 16.3 gives a quick proof of Robinson's joint consistency theorem, and Section 16.4 presents a theorem describing exactly when two matrix rings over a pair of fields are elementarily equivalent.

16.1. The Keisler-Shelah theorem

The purpose of this section is to prove the Keisler-Shelah theorem, which gives an ultrapower-theoretic reformulation of elementary equivalence:

Theorem 16.1.1 (Keisler-Shelah theorem). *Let \mathcal{M} and \mathcal{N} be \mathcal{L} -structures. Then \mathcal{M} and \mathcal{N} are elementarily equivalent if and only if there is an ultrafilter \mathcal{U} such that $\mathcal{M}^{\mathcal{U}} \cong \mathcal{N}^{\mathcal{U}}$.*

Of course, the hard direction is the forward direction. The idea of the proof is really a combination of two techniques we have already seen thus far: we construct the ultrafilter as an increasing chain of filters by transfinite recursion, at each stage taking care of requirements such as ensuring that

a particular set or its complement is put in the filter. (This technique was done in the section on constructing good ultrafilters.) As we are doing so, we will also enumerate the direct power of the structures in such a way that the resulting map induces a bijection of the corresponding ultrapowers. Naturally, as we are building the filter, we will need to make sure that sets are put in that capture the truth of formulas of the previously defined elements, and so balancing both of these things is a delicate act.

We first prove Theorem 16.1.1 under a technical simplifying assumption. Afterward, we will explain a simple “reduct trick” which allows us to deduce the general case.

For a cardinal λ , we define the cardinal $\mu(\lambda)$ to be the least cardinal η for which $\lambda^\eta > \lambda$. The following facts are clear:

Lemma 16.1.2.

- (1) $\mu(\aleph_0) = \aleph_0$.
- (2) For any cardinal λ , $\mu(\lambda) \leq \lambda$.
- (3) For any cardinal λ , $\mu(2^\lambda) \geq \lambda^+$.

Here is the version of Theorem 16.1.1 we will prove first:

Theorem 16.1.3. *Fix a cardinal λ . Suppose that \mathcal{M} and \mathcal{N} are elementarily equivalent \mathcal{L} -structures with $|\mathcal{L}| \leq \lambda$ and $\max(|M|, |N|) < \mu(\lambda)$. Then there is an ultrafilter \mathcal{U} on λ such that $\mathcal{M}^{\mathcal{U}} \cong \mathcal{N}^{\mathcal{U}}$.*

We now work toward proving Theorem 16.1.3. First, for a cardinal λ , by a λ -**pair**, we mean a pair (F, \mathcal{F}) , where

- $F \subseteq \mu(\lambda)^\lambda$ and
- \mathcal{F} is a filter on λ .

As in the chapter on good ultrafilters, we need a notion of a filter having many possible extensions. For our purposes, a slightly more general notion of consistency is needed. We describe this notion now.

Definition 16.1.4. We say that a filter \mathcal{F} is κ -**generated** if it has a base of cardinality at most κ .

Here is the notion of consistency that will be useful for us. It is a natural strengthening of the notion of a set of functions having large oscillation modulo a filter presented in Definition 8.5.2.

Definition 16.1.5. If κ and λ are cardinals, we say that a λ -pair (F, \mathcal{F}) is κ -**consistent** if:

- (1) \mathcal{F} is κ -generated; and

- (2) Given a set $X \in \mathcal{F}$, a cardinal $\beta < \mu(\lambda)$, a sequence $(f_\rho)_{\rho < \beta}$ of distinct elements of F , a sequence $(\sigma_\rho)_{\rho < \beta}$ of ordinals less than $\mu(\lambda)$, we have that

$$X \cap \{\zeta < \lambda : f_\rho(\zeta) = \sigma_\rho \text{ for all } \rho < \beta\} \neq \emptyset.$$

We will need some coherence properties of the notion of consistent pairs that follow easily from the definition and we thus leave them as an exercise.

Exercise 16.1.6.

- (1) If (F, \mathcal{F}) is a κ -consistent λ -pair and $\kappa < \gamma$, then (F, \mathcal{F}) is also γ -consistent.
- (2) If (F, \mathcal{F}) is a κ -consistent λ -pair and $F' \subseteq F$, then (F', \mathcal{F}) is also a κ -consistent λ -pair.
- (3) Suppose that δ is a cardinal and for every $\zeta < \delta$, we have a κ_ζ -consistent λ -pair $(F_\zeta, \mathcal{F}_\zeta)$. Suppose further that $F_\zeta \supseteq F_\rho$ and $\mathcal{F}_\zeta \subseteq \mathcal{F}_\rho$ for $\zeta < \rho < \delta$. Suppose that each $\kappa_\zeta \leq \kappa$ and the cofinality of δ is $\leq \kappa$. Then $(\bigcap_{\zeta < \delta} F_\zeta, \bigcup_{\zeta < \delta} \mathcal{F}_\zeta)$ is a κ -consistent λ -pair.

The first lemma helps us get started building “large” consistent pairs. The proof is almost identical to the proof of Lemma 8.5.6 (and is in fact easier due to the fact that we are only working with the filter $\{\lambda\}$). We thus leave it as an exercise.

Lemma 16.1.7. *Given a cardinal λ , there is $F \subseteq \mu(\lambda)^\lambda$ with $|F| = 2^\lambda$ such that $(F, \{\lambda\})$ is a λ -consistent λ -pair.*

Exercise 16.1.8. Prove Lemma 16.1.7.

Note, of course, that $(F, \{\lambda\})$ as in the previous lemma is a κ -consistent λ -pair for any $\kappa \geq 1$.

The next lemma is the analogue of Lemma 8.5.7 to the current situation. As before, we ask that we can extend our filter to decide a given set at the cost of losing a “small” number of extensions.

Some notation will be useful: given a filter \mathcal{F} on λ and $A \subseteq \lambda$, let $\mathcal{F}[A]$ be the filter generated by \mathcal{F} and $\{A\}$. Note that if \mathcal{F} is κ -generated, so is $\mathcal{F}[A]$.

Lemma 16.1.9. *Suppose that (F, \mathcal{F}) is a κ -consistent λ -pair and $A \subseteq \lambda$. Then there is $F' \subseteq F$ with $|F \setminus F'| < \mu(\lambda)$ such that either $(F', \mathcal{F}[A])$ is a κ -consistent λ -pair or $(F', \mathcal{F}[\lambda \setminus A])$ is a κ -consistent λ -pair.*

Proof. Suppose that $(F, \mathcal{F}[A])$ is not κ -consistent. By definition, there is $X \in \mathcal{F}$, $\beta < \mu(\lambda)$, distinct f_ρ from F for $\rho < \beta$ and ordinals $\sigma_\rho < \mu(\lambda)$ for $\rho < \beta$ such that $X \cap A \cap B = \emptyset$, where $B := \{\zeta < \lambda : f_\rho(\zeta) = \sigma_\rho\}$. Let

$F' := F \setminus \{f_\rho : \rho < \beta\}$. We claim that $(F', \mathcal{F}[\lambda \setminus A])$ is κ -consistent. To see this, fix $Y \in \mathcal{F}$, $\beta' < \mu(\lambda)$, distinct f'_ρ from F' with $\rho < \beta'$, and ordinals $\sigma'_\rho < \mu(\lambda)$ for $\rho < \beta'$. Let $B' := \{\zeta < \lambda : f'_\rho(\zeta) = \sigma'_\rho \text{ for all } \rho < \beta'\}$. We need to show that $Y \cap (\lambda \setminus A) \cap B' \neq \emptyset$. Since (F, \mathcal{F}) is κ -consistent, we have that $B \cap B' \cap X \cap Y \neq \emptyset$. Since $X \cap A \cap B = \emptyset$, we have that $Y \cap (\lambda \setminus A) \cap B' \neq \emptyset$, as desired. \square

A simple induction argument applied to the preceding lemma gives the following:

Lemma 16.1.10. *Suppose that (F, \mathcal{F}) is a κ -consistent λ -pair with $\mu(\lambda) \leq \kappa$. Suppose that, for each $\zeta < \kappa$, $A_\zeta \subseteq \lambda$ is given. Then there is $F' \subseteq F$ with $|F \setminus F'| \leq \kappa$, and a filter \mathcal{F}' on λ extending \mathcal{F} such that (F', \mathcal{F}') is a κ -consistent λ -pair for which, given any $\zeta < \kappa$, either $A_\zeta \in \mathcal{F}'$ or $(\lambda \setminus A_\zeta) \in \mathcal{F}'$.*

When proving Theorem 16.1.3, we will find ourselves in the situation where we have already built a filter that thinks some existential statements with parameters are true most of the time and that filter is consistent. We would like to find an actual witness to the existential statements while maintaining consistency. The next lemma aids us in this endeavor.

Lemma 16.1.11. *Suppose that \mathcal{M} is an \mathcal{L} -structure with $|\mathcal{M}| < \mu(\lambda)$ and that (F, \mathcal{F}) is a κ -consistent λ -pair. Let $(\varphi_\zeta(x, y_1, \dots, y_{n(\zeta)}))_{\zeta < \kappa}$ be a set of \mathcal{L} -formulas closed under conjunction. For each $\zeta < \kappa$, fix a tuple $a^\zeta := (a^\zeta_1, \dots, a^\zeta_{n(\zeta)})$ of functions from M^λ . Suppose that, for each $\zeta < \kappa$, we have*

$$\{\nu < \lambda : \mathcal{M} \models \exists x \varphi_\zeta(x, a^\zeta(\nu))\} \in \mathcal{F}.$$

Then there are $b \in M^\lambda$, $F' \subseteq F$ with $|F \setminus F'| \leq \kappa$, and a filter $\mathcal{F}' \supseteq \mathcal{F}$ such that (F', \mathcal{F}') is a κ -consistent λ -pair and such that, for every $\zeta < \kappa$, we have

$$\{\nu < \lambda : \mathcal{M} \models \varphi_\zeta(b(\nu), a^\zeta(\nu))\} \in \mathcal{F}'.$$

In order to prove the previous lemma, we will actually need to introduce a more general version of consistent pairs.

Definition 16.1.12. A λ -triple is a triple (F, G, \mathcal{F}) , where (F, \mathcal{F}) is a λ -pair and G is a collection of functions $g : \lambda \rightarrow \beta(g)$ for some cardinal $\beta(g) < \mu(\lambda)$. A λ -triple is κ -consistent if item (2) of the definition of κ -consistent λ -pair is replaced by

- (2') given $X \in \mathcal{F}$, a cardinal $\beta < \mu(\lambda)$, a sequence $(f_\rho)_{\rho < \beta}$ of distinct elements of F , a sequence $(\sigma_\rho)_{\rho < \beta}$ of ordinals less than $\mu(\lambda)$, $f \in F$ distinct from the f_ρ 's, and $g \in G$, we have that

$$X \cap \{\zeta < \lambda : f_\rho(\zeta) = \sigma_\rho \text{ for all } \rho < \beta\} \cap \{\zeta < \lambda : f(\zeta) = g(\zeta)\} \neq \emptyset.$$

We need to show that κ -consistent λ -pairs can always be expanded to κ -consistent λ -triples at the expense of removing a small number of elements from F .

Lemma 16.1.13. *Suppose that (F, \mathcal{F}) is a κ -consistent λ -pair and that $g : \lambda \rightarrow \gamma$ is a function, where γ is a cardinal below $\mu(\lambda)$. Suppose also that $\mu(\lambda) \leq \kappa$. Then there is $F' \subseteq F$ with $|F \setminus F'| \leq \kappa$ such that $(F', \{g\}, \mathcal{F})$ is a κ -consistent λ -triple.*

Proof. Suppose that the lemma is false, in other words:

(*) for every $\tilde{F} \subseteq F$, if $|\tilde{F}| \leq \kappa$, the λ -triple $(F \setminus \tilde{F}, \{g\}, \mathcal{F})$ is not κ -consistent.

This allows us to define sets F_ζ, \tilde{F}_ζ for $\zeta < \kappa^+$ for which:

- $F_0 = F$;
- $\tilde{F}_\zeta \subseteq F_\zeta$, $|\tilde{F}_\zeta| \leq \kappa$, and the λ -triple $(F_\zeta \setminus \tilde{F}_\zeta, \{g\}, \mathcal{F})$ is not κ -consistent;
- $F_{\zeta+1} := F_\zeta \setminus \tilde{F}_\zeta$;
- $F_\eta := \bigcap_{\zeta < \eta} F_\zeta$ if η is a limit ordinal.

Moreover, the \tilde{F}_ζ 's may be chosen so that, for each $\zeta < \kappa^+$, we may find cardinals $\beta_\zeta < \mu(\lambda)$, distinct functions f_ρ^ζ from \tilde{F}_ζ for $\rho < \beta_\zeta$ and ordinals $\sigma_\rho^\zeta < \mu(\lambda)$ for $\rho < \beta_\zeta$, and $f^\zeta \in \tilde{F}_\zeta$ distinct from the f_ρ^ζ 's such that, setting

$$A_\zeta := \{\nu < \lambda : f_\rho^\zeta(\nu) = \sigma_\rho^\zeta \text{ for all } \rho < \beta_\zeta\} \cap \{\nu < \lambda : f^\zeta(\nu) = g(\nu)\},$$

we have that $\mathcal{F}[A_\zeta]$ is not a proper filter on λ , whence there are sets $X_\zeta \in E$ (where E is a generating set for \mathcal{F}) such that $A_\zeta \cap X_\zeta = \emptyset$. Since $|E| \leq \kappa$, there is some $X \in E$ and κ^+ -many ζ for which $X = X_\zeta$. Since $\mu(\lambda) \leq \kappa$, there is $\beta < \mu(\lambda)$ and κ^+ -many of the aforementioned ζ for which $\beta_\zeta = \beta$. By re-indexing, we may as well assume that $\beta_\zeta = \beta$ and $X_\zeta = X$ for all $\zeta < \kappa^+$.

Suppose that $\gamma < \mu(\lambda)$ is such that $g \in \gamma^\lambda$. Consider the set

$$A := \{\nu < \lambda : f_\rho^\zeta(\nu) = \sigma_\rho^\zeta \text{ for all } \zeta < \gamma \text{ and } \rho < \beta\} \\ \cap \{\nu < \lambda : f^\zeta(\nu) = \zeta \text{ for all } \zeta < \gamma\}.$$

Since we are considering $|\gamma| + |\beta| < \mu(\lambda)$ many distinct functions and (F, \mathcal{F}) is κ -consistent, we see that $A \cap X \neq \emptyset$. Fix $\nu \in A \cap X$. Set $\zeta := g(\nu)$; then $\nu \in A_\zeta \cap X$, a contradiction to the fact that $A_\zeta \cap X = \emptyset$. \square

Lemma 16.1.14. *Suppose that (F, \mathcal{F}) is a κ -consistent λ -pair and that G as above is such that $\mu(\lambda) + |G| \leq \kappa$. Then there is $F' \subseteq F$ with $|F \setminus F'| \leq \kappa$ such that (F', G, \mathcal{F}) is a κ -consistent λ -triple.*

Proof. For each $g \in G$, let $F_g \subseteq F$ satisfy the conclusion of the previous lemma. Then $F' := \bigcap_{g \in G} F_g$ is as desired. \square

Proof of Lemma 16.1.11. For the convenience of the reader, we recall the hypotheses of the lemma:

- \mathcal{M} is an \mathcal{L} -structure with $|M| < \mu(\lambda)$.
- (F, \mathcal{F}) is a κ -consistent λ -pair.
- $(\varphi_\zeta(x, y_1, \dots, y_{n(\zeta)}))_{\zeta < \kappa}$ is a set of \mathcal{L} -formulas closed under conjunction.
- For each $\zeta < \kappa$, $a^\zeta := (a_1^\zeta, \dots, a_{n(\zeta)}^\zeta)$ is a tuple of functions from M^λ .
- For each $\zeta < \kappa$, we have

$$\{\nu < \lambda : \mathcal{M} \models \exists x \varphi_\zeta(x, a^\zeta(\nu))\} \in \mathcal{F}.$$

We seek to find $b \in M^\lambda$, $F' \subseteq F$ with $|F \setminus F'| \leq \kappa$ and a filter $\mathcal{F}' \supseteq \mathcal{F}$ such that (F', \mathcal{F}') is a κ -consistent λ -pair and such that, for every $\zeta < \kappa$, we have

$$\{\nu < \lambda : \mathcal{M} \models \varphi_\zeta(b(\nu), a^\zeta(\nu))\} \in \mathcal{F}'.$$

Set $\delta := |M|$ and let $(a_\zeta)_{\zeta < \delta}$ enumerate M . We want to expand our κ -consistent pair to a κ -consistent triple using the set of functions $G := \{g_\zeta : \zeta < \kappa\}$, where $g_\zeta(\nu) =$ the first ordinal η such that $\mathcal{M} \models \varphi_\zeta(a_\eta, a^\zeta(\nu))$ if such an η exists; otherwise, set $g_\zeta(\nu) = 0$. Since $\mu(\lambda) + |G| \leq \kappa$, Lemma 16.1.14 allows us to find $\tilde{F} \subseteq F$ with $|F \setminus \tilde{F}| \leq \kappa$ for which $(\tilde{F}, G, \mathcal{F})$ is a κ -consistent λ -triple.

Fix $f \in \tilde{F}$ and define $b \in M^\lambda$ by $b(\nu) = a_{f(\nu)}$ unless $f(\nu) \geq \delta$, in which case $b(\nu) := a_0$. For each $\zeta < \kappa$, set $B_\zeta := \{\nu < \lambda : \mathcal{M} \models \varphi_\zeta(b(\nu), a^\zeta(\nu))\}$. Let \mathcal{F}' be the filter generated by \mathcal{F} and the B_ζ 's. Note that \mathcal{F}' is still κ -generated. Finally, set $F' := \tilde{F} \setminus \{f\}$. We show that these choices are as in the conclusion of the lemma.

The only item that needs checking is that (F', \mathcal{F}') is κ -consistent. Toward this end, fix $X \in \mathcal{F}$, $\beta < \mu(\lambda)$, a sequence of distinct elements f_ρ from F' , and ordinals σ_ρ less than $\mu(\lambda)$. Let $B := \{\nu < \lambda : f_\rho(\nu) = \sigma_\rho \text{ for all } \rho < \beta\}$. Since the formulae are closed under taking conjunction, it suffices to show that, for each $\zeta < \kappa$, we have that $B \cap X \cap B_\zeta \neq \emptyset$. Set $U := \{\nu < \lambda : f(\nu) = g_\zeta(\nu)\}$ and $V := \{\nu < \lambda : \mathcal{M} \models \exists x \varphi_\zeta(x, a^\zeta(\nu))\}$. Note that $U \cap V \subseteq B_\zeta$ and $V \in \mathcal{F}$. Set $B' := B \cap U$. Since $(\tilde{F}, G, \mathcal{F})$ is a κ -consistent λ -triple, we know that $B' \cap V \cap X \neq \emptyset$. Since $B' \cap V \cap X \subseteq B \cap X \cap B_\zeta$, we have the desired conclusion. \square

We are now ready to prove the technical simplification of the Keisler-Shelah theorem:

Proof of Theorem 16.1.3. Suppose that λ is a cardinal and \mathcal{M} and \mathcal{N} are elementarily equivalent \mathcal{L} -structures with $|\mathcal{L}| \leq \lambda$ and $\max(|M|, |N|) < \mu := \mu(\lambda)$. Let $(A_\rho)_{\rho < 2^\lambda}$ enumerate the subsets of λ . We now construct, for $\rho < 2^\lambda$, sequences (F_ρ) , (\mathcal{F}_ρ) , (a_ρ) , and (b_ρ) , such that:

- (1) each F_ρ is a collection of functions $\lambda \rightarrow \mu$ with $F_\rho \supseteq F_{\rho+1}$;
- (2) each \mathcal{F}_ρ is a filter on λ with $\mathcal{F}_\rho \subseteq \mathcal{F}_{\rho+1}$;
- (3) $|F_0 \setminus F_\rho| \leq \lambda + |\rho|$;
- (4) $(F_\rho, \mathcal{F}_\rho)$ is a $(\lambda + |\rho|)$ -consistent λ -pair;
- (5) either $A_\rho \in \mathcal{F}_{\rho+1}$ or $\lambda \setminus A_\rho \in \mathcal{F}_{\rho+1}$;
- (6) $M^\lambda = \{a_\rho : \rho < 2^\lambda\}$ and $N^\lambda = \{b_\rho : \rho < 2^\lambda\}$;
- (7) for any $\zeta < 2^\lambda$, any formula $\varphi(x_1, \dots, x_n)$, and any $\rho_1, \dots, \rho_n < \zeta$, either

$$X(\varphi, \rho_1, \dots, \rho_n, \mathcal{M}) \in \mathcal{F}_\zeta \text{ or } X(\neg\varphi, \rho_1, \dots, \rho_n, \mathcal{M}) \in \mathcal{F}_\zeta;$$

- (8) for any $\zeta < 2^\lambda$, any formula $\varphi(x_1, \dots, x_n)$, and any $\rho_1, \dots, \rho_n < \zeta$, we have

$$X(\varphi, \rho_1, \dots, \rho_n, \mathcal{M}) \in \mathcal{F}_\zeta \Leftrightarrow X(\varphi, \rho_1, \dots, \rho_n, \mathcal{N}) \in \mathcal{F}_\zeta.$$

Here, $X(\varphi, \rho_1, \dots, \rho_n, \mathcal{M}) := \{\nu < \lambda : \mathcal{M} \models \varphi(a_{\rho_1}(\nu), \dots, a_{\rho_n}(\nu))\}$, and analogously for $X(\varphi, \rho_1, \dots, \rho_n, \mathcal{N})$. If we can successfully carry out the construction of these sequences, then item (5) says that $\mathcal{U} := \bigcup_{\rho < 2^\lambda} \mathcal{F}_\rho$ is an ultrafilter on λ while item (8) says that the map $[a_\rho]_{\mathcal{U}} \mapsto [b_\rho]_{\mathcal{U}}$ is an isomorphism $\mathcal{M}^{\mathcal{U}} \rightarrow \mathcal{N}^{\mathcal{U}}$, yielding the desired result.

We can start by Lemma 16.1.7: there is $F_0 \subseteq \mu^\lambda$ with $|F_0| = 2^\lambda$ such that $(F_0, \{\lambda\})$ is a λ -consistent λ -pair, so we may set $\mathcal{F}_0 := \{\lambda\}$.

It is also clear how to define F_η and \mathcal{F}_η for limit ordinals η : set $F_\eta := \bigcap_{\rho < \eta} F_\rho$ and $\mathcal{F}_\eta := \bigcup \mathcal{F}_\rho$. By Exercise 16.1.6, these are as desired.

Assume now that we have defined F_ρ , \mathcal{F}_ρ for $\rho \leq \sigma$ and a_ρ and b_ρ for $\rho < \sigma$. We show how to define $F_{\sigma+1}$, $\mathcal{F}_{\sigma+1}$, a_σ , and b_σ .

To do a proper back-and-forth, we need to explain how to put each element of M^λ in the domain and each element of N^λ in the range, but since the arguments are the same, let us only show how to put a particular element of M^λ in the domain. So let a_σ be an element of M^λ not already listed as some a_ρ . Since $(F_\sigma, \mathcal{F}_\sigma)$ is a $(\lambda + |\sigma|)$ -consistent λ -pair and there are at most $\lambda + |\sigma|$ -many sets $X(\varphi, \sigma, \rho_1, \dots, \rho_n, \mathcal{M})$ of formulae $\varphi(x, y_1, \dots, y_n)$ and $\rho_1, \dots, \rho_n < \sigma$, Lemma 16.1.10 implies that we can find $F' \subseteq F_\sigma$ and $\mathcal{F}' \supseteq \mathcal{F}_\sigma$ such that $|F_\sigma \setminus F'| \leq \lambda + |\sigma|$, (F', \mathcal{F}') is $\lambda + |\sigma|$ -consistent, and, either $X(\varphi, \sigma, \rho_1, \dots, \rho_n, \mathcal{M}) \in \mathcal{F}'$ or $\lambda \setminus X(\varphi, \sigma, \rho_1, \dots, \rho_n, \mathcal{M}) \in \mathcal{F}'$. Let

$$\Gamma := \{\varphi(x, a_{\rho_1}, \dots, a_{\rho_n}) : X(\varphi, \sigma, \rho_1, \dots, \rho_n, \mathcal{M}) \in \mathcal{F}'\}.$$

It is clear that if $\varphi(x, a_{\rho_1}, \dots, a_{\rho_n}) \notin \Gamma$, then $\neg\varphi(x, a_{\rho_1}, \dots, a_{\rho_n}) \in \Gamma$. Note also that if $\varphi(x, a_{\rho_1}, \dots, a_{\rho_n}) \in \Gamma$, then $X(\exists x\varphi, \rho_1, \dots, \rho_n, \mathcal{N}) \in \mathcal{F}'$. Indeed, if this were not the case, then we would have $X(\exists x\varphi, \rho_1, \dots, \rho_n, \mathcal{N}) \notin \mathcal{F}_\sigma$, whence by (8) we have that $X(\exists x\varphi, \rho_1, \dots, \rho_n, \mathcal{M}) \notin \mathcal{F}_\sigma$ and thus, by (7) we have $X(\neg\exists x\varphi, \rho_1, \dots, \rho_n, \mathcal{M}) \in \mathcal{F}_\sigma \subseteq \mathcal{F}'$, yielding a contradiction. Thus, by Lemma 16.1.11, there is $b_\sigma : \lambda \rightarrow N$ and $F'' \subseteq F'$ and $\mathcal{F}'' \supseteq \mathcal{F}'$ such that $|F' \setminus F''| \leq \lambda + |\sigma|$ and (F'', \mathcal{F}'') is a $(\lambda + |\sigma|)$ -consistent λ -pair and, for each $\varphi(x, a_{\rho_1}, \dots, a_{\rho_n}) \in \Gamma$, we have $\{\nu < \lambda : \mathcal{N} \models \varphi(b_\sigma(\nu), b_{\rho_1}(\nu), \dots, b_{\rho_n}(\nu))\} \in \mathcal{F}''$. Finally, apply Lemma 16.1.9 to get $F_{\sigma+1} \subseteq F''$ with $|F'' \setminus F_{\sigma+1}| \leq \lambda + |\sigma|$ and such that either $(F_{\sigma+1}, \mathcal{F}''[A_\sigma])$ is $\lambda + |\sigma|$ -consistent or $(F_{\sigma+1}, \mathcal{F}''[\lambda \setminus A_\sigma])$ is $\lambda + |\sigma|$ -consistent, and let $\mathcal{F}_{\sigma+1}$ be whichever one it is. This completes the construction and the proof of the theorem. \square

We now explain how to deduce the general Keisler-Shelah theorem, Theorem 16.1.1, from the technical simplification Theorem 16.1.3. The main idea is the following simple exercise:

Exercise 16.1.15. Suppose that T is a complete \mathcal{L} -theory which has a model of size κ . Prove that there is a sublanguage $\mathcal{L}' \subseteq \mathcal{L}$ with $|\mathcal{L}'| \leq 2^\kappa$ with the property that any model of $T \upharpoonright \mathcal{L}'$ has a unique expansion to a model of T .

We are now ready to explain the general Keisler-Shelah theorem:

Proof of Theorem 16.1.1. Suppose \mathcal{M} and \mathcal{N} are elementarily equivalent \mathcal{L} -structures and let $T := \text{Th}(\mathcal{M}) = \text{Th}(\mathcal{N})$. Let $\kappa := \max(|M|, |N|)$. Let $\mathcal{L}' \subseteq \mathcal{L}$ be as in Exercise 16.1.15. Set $\lambda := 2^\kappa$. By Lemma 16.1.2, we have that $\mu(\lambda) \geq \kappa^+ > \max(|M|, |N|)$. Thus, by Theorem 16.1.3, there is an ultrafilter \mathcal{U} on λ such that $(\mathcal{M} \upharpoonright \mathcal{L}')^\mathcal{U} \cong (\mathcal{N} \upharpoonright \mathcal{L}')^\mathcal{U}$. Since taking ultrapowers commutes with taking reducts, the defining property of \mathcal{L}' implies that $\mathcal{M}^\mathcal{U} \cong \mathcal{N}^\mathcal{U}$, as desired. \square

We end this section with one observation comparing the proof of the Keisler-Shelah theorem from GCH given in Section 8.4 and the above proof. Suppose that \mathcal{M} and \mathcal{N} are elementarily equivalent \mathcal{L} -structures and set $\kappa := \max(|M|, |N|, |\mathcal{L}|)$. The proof of Theorem 16.1.1 given above constructs an ultrafilter \mathcal{U} on 2^κ such that $\mathcal{M}^\mathcal{U} \cong \mathcal{N}^\mathcal{U}$. On the other hand, the proof from GCH given in Section 8.4 yielded an ultrafilter \mathcal{V} on the smaller cardinal κ itself for which $\mathcal{M}^\mathcal{V} \cong \mathcal{N}^\mathcal{V}$.

It is natural to wonder if there is a construction that leads to a ZFC proof of the Keisler-Shelah theorem where one can take the ultrafilter \mathcal{U} to be on κ itself. Shelah showed that is not the case. In fact, he showed the following stronger result:

Theorem 16.1.16 (Shelah [160]). *It is consistent with ZFC that there are countable elementarily equivalent graphs Γ_1 and Γ_2 such that, for any ultrafilters \mathcal{U} and \mathcal{V} on \mathbb{N} , $\Gamma_1^{\mathcal{U}} \not\cong \Gamma_2^{\mathcal{V}}$.*

16.2. Application: Elementary classes

Definition 16.2.1. Let \mathcal{L} be a language, and let \mathcal{K} be a class of \mathcal{L} -structures. We say that \mathcal{K} is an **elementary class** (or an **axiomatizable class**) if there is a set T of \mathcal{L} -sentences such that $\mathcal{K} = \text{Mod}(T)$, the class of models of T . In this case, we call T a set of **axioms** for \mathcal{K} .

In “nature” one often encounters a class of \mathcal{L} -structures that one believes forms an elementary class, even though one may not be able to come up with a set of axioms for the class. In this section, we show how to use the Keisler-Shelah theorem to give an equivalent characterization of axiomatizable class that does not make mention of any specific set of axioms, but rather checks that the class \mathcal{K} satisfies three very natural closure properties.

Definition 16.2.2. Given a class \mathcal{K} of \mathcal{L} -structures, we let the **theory of \mathcal{K}** be

$$\text{Th}(\mathcal{K}) := \{\sigma : \mathcal{M} \models \sigma \text{ for all } \mathcal{M} \in \mathcal{K}\}.$$

If $\mathcal{N} \models \text{Th}(\mathcal{K})$, we say that \mathcal{N} is **pseudo- \mathcal{K}** . We let $\text{Pseudo}\mathcal{K}$ denote the class of pseudo- \mathcal{K} structures.

Exercise 16.2.3. Prove the following:

- (1) $\text{Pseudo}\mathcal{K}$ is the smallest elementary class containing \mathcal{K} .
- (2) $\mathcal{N} \in \text{Pseudo}\mathcal{K}$ if and only if, whenever $\mathcal{N} \models \sigma$, then there is $\mathcal{M} \in \mathcal{K}$ such that $\mathcal{M} \models \sigma$.

Lemma 16.2.4. $\mathcal{N} \in \text{Pseudo}\mathcal{K}$ if and only if there is a set I , an ultrafilter \mathcal{U} on I , and a family $(\mathcal{M}_i)_{i \in I}$ of structures from \mathcal{K} such that $\mathcal{N} \equiv \prod_{\mathcal{U}} \mathcal{M}_i$.

Proof. The backward direction is obvious. For the forward direction, suppose that $\mathcal{N} \in \text{Pseudo}\mathcal{K}$. Let I be the set of finite subsets of $\text{Th}(\mathcal{N})$. For each $\sigma \in \text{Th}(\mathcal{N})$, let $X_\sigma := \{i \in I : \sigma \in i\}$. Note that the family of $(X_\sigma)_{\sigma \in \text{Th}(\mathcal{N})}$ has the FIP. Let \mathcal{U} be an ultrafilter on I containing each X_σ . By Exercise 16.2.3(2), for each $i \in I$, there is $\mathcal{M}_i \in \mathcal{K}$ such that $\mathcal{M}_i \models \bigwedge_{\sigma \in i} \sigma$. Then for each $\sigma \in \text{Th}(\mathcal{N})$, since $X_\sigma \in \mathcal{U}$, we have that $\prod_{\mathcal{U}} \mathcal{M}_i \models \sigma$. It follows that $\prod_{\mathcal{U}} \mathcal{M}_i \models \text{Th}(\mathcal{N})$, whence it follows that $\mathcal{N} \equiv \prod_{\mathcal{U}} \mathcal{M}_i$. \square

Corollary 16.2.5. *For a class \mathcal{K} of \mathcal{L} -structures, the following are equivalent:*

- (1) \mathcal{K} is an elementary class.
- (2) \mathcal{K} is closed under ultraproducts and elementary equivalence.

(3) $\mathcal{K} = \text{Pseudo } \mathcal{K}$.

Using the Keisler-Shelah theorem, we can do even better. First, we say that a class \mathcal{K} of \mathcal{L} -structures is closed under ultraroots if, for any \mathcal{L} -structure \mathcal{M} and any ultrafilter \mathcal{U} , $\mathcal{M}^{\mathcal{U}}$ belongs to \mathcal{K} , then so does \mathcal{M} .

Corollary 16.2.6. *For a class \mathcal{K} of \mathcal{L} -structures, we have that \mathcal{K} is an elementary class if and only if it is closed under isomorphism, ultraproducts, and ultraroots.*

Proof. The forward direction is clear. For the backward direction, suppose that \mathcal{K} is closed under isomorphism, ultraproduct, and ultraroots. We show that \mathcal{K} is an elementary class by verifying item (2) of Corollary 16.2.5. Since we are already assuming that \mathcal{K} is closed under ultraproducts, we just need to show that it is closed under elementary equivalence. Toward this end, suppose that $\mathcal{M} \equiv \mathcal{N}$ and $\mathcal{M} \in \mathcal{K}$; we must show that $\mathcal{N} \in \mathcal{K}$. By the Keisler-Shelah theorem, there is an ultrafilter \mathcal{U} such that $\mathcal{M}^{\mathcal{U}} \cong \mathcal{N}^{\mathcal{U}}$. Since \mathcal{K} is closed under ultraproducts, we have that $\mathcal{M}^{\mathcal{U}}$ belongs to \mathcal{K} , whence so does $\mathcal{N}^{\mathcal{U}}$ since \mathcal{K} is closed under isomorphisms. Finally, since \mathcal{K} is closed under ultraroots, it follows that $\mathcal{N} \in \mathcal{K}$, as desired. \square

16.3. Application: Robinson's joint consistency theorem

We can use the Keisler-Shelah theorem to give a short proof of the following classical theorem of Robinson:

Theorem 16.3.1 (Robinson's joint consistency theorem). *Suppose that \mathcal{L}_1 and \mathcal{L}_2 are first-order languages and T_1 and T_2 are consistent \mathcal{L}_1 and \mathcal{L}_2 theories, respectively. Let $\mathcal{L} := \mathcal{L}_1 \cap \mathcal{L}_2$ and suppose that there is a complete \mathcal{L} -theory $T \subseteq T_1 \cap T_2$. Then $T_1 \cup T_2$ is consistent.*

Proof. Let $\mathcal{M}_1 \models T_1$ and $\mathcal{M}_2 \models T_2$. For $i = 1, 2$, let \mathcal{N}_i denote the reduct of \mathcal{M}_i to \mathcal{L} . Then $\mathcal{N}_1, \mathcal{N}_2 \models T$, whence, since T is complete, we have that $\mathcal{N}_1 \equiv \mathcal{N}_2$. By the Keisler-Shelah theorem, there is an ultrafilter \mathcal{U} such that $\mathcal{N}_1^{\mathcal{U}} \cong \mathcal{N}_2^{\mathcal{U}}$. Let $f : \mathcal{N}_1^{\mathcal{U}} \rightarrow \mathcal{N}_2^{\mathcal{U}}$ be an isomorphism. Since $\mathcal{M}_i^{\mathcal{U}} = \mathcal{N}_i^{\mathcal{U}}$ for $i = 1, 2$, we can expand $\mathcal{N}_2^{\mathcal{U}}$ to an \mathcal{L}_1 -structure \mathcal{P} by interpreting the symbols in $\mathcal{L}_1 \setminus \mathcal{L}$ so that $f : \mathcal{M}_1^{\mathcal{U}} \rightarrow \mathcal{P}$ is an isomorphism. Note then that $\mathcal{P} \models T_1$. One can then extend \mathcal{P} further to an $\mathcal{L}_1 \cup \mathcal{L}_2$ structure \mathcal{Q} by interpreting the symbols in $\mathcal{L}_2 \setminus \mathcal{L}$ in the same manner as $\mathcal{M}_2^{\mathcal{U}}$. It follows that $\mathcal{Q} \models T_1 \cup T_2$, as desired. \square

Exercise 16.3.2. Show that the conclusion of Robinson's joint consistency theorem may fail if one does not assume that there is a complete \mathcal{L} -theory contained in $T_1 \cap T_2$.

There are certainly more elementary proofs of Robinson's joint consistency theorem (see [28, Section 2.2]), but when you have a hammer, why not use it?

16.4. Application: Elementary equivalence of matrix rings

In this section, we use the Keisler-Shelah theorem to give a quick proof of a fact concerning elementary equivalence of matrix rings. We first need the following algebraic result:

Proposition 16.4.1. *If K and L are fields and $m, n \in \mathbb{N}$, then $M_m(K) \cong M_n(L)$ (as rings) if and only if $m = n$ and $K \cong L$.*

Proof Sketch. First, note that K^m is a simple $M_m(K)$ -module and is in fact the only one: if M is also a simple $M_m(K)$ -module, then, fixing $v \in M \setminus \{0\}$, we have that $Av \mapsto Ae_1$ is a module morphism that is injective as M is simple and surjective as $M_m(K)$ is simple.

Next note that the map $a \mapsto (\vec{x} \mapsto a\vec{x}) : K \rightarrow \text{End}_{M_m(K)} K^m$ is an isomorphism; this follows from the fact that the only matrices that commute with all matrices are the diagonal ones.

Thus, if $M_m(K) \cong M_n(L)$, then their unique simple modules are isomorphic, whence so are the corresponding endomorphism rings, that is, $K \cong L$. But then once $K \cong L$, we have that $m = n$ by dimension considerations. \square

The Keisler-Shelah theorem allows us to obtain the analogous fact for elementary equivalence immediately. First, we need:

Exercise 16.4.2. For any ring R and any $n \in \mathbb{N}$, prove that $M_n(R)^{\mathcal{U}} \cong M_n(R^{\mathcal{U}})$ as rings.

Theorem 16.4.3. *If K and L are fields and $m, n \in \mathbb{N}$, then $M_m(K) \equiv M_n(L)$ as rings if and only if $m = n$ and $K \equiv L$ (as rings).*

Proof. This follows immediately from the Keisler-Shelah theorem, Proposition 16.4.1, and Exercise 16.4.2. \square

16.5. Notes and references

Our proof of the Keisler-Shelah theorem given in Section 16.1 follows [28] quite closely. For a proof of the Keisler-Shelah theorem in Henson's positive bounded logic (a precursor to the continuous logic presented in Section 11.4), see [88]. For a proof of the Keisler-Shelah theorem for modern continuous logic from the Keisler-Shelah theorem for classical logic, see the recent preprint [69]. The soft test for axiomatizability of a class of structures presented in Section 16.2 is due to Keisler [97]. Robinson's joint consistency

theorem is due to Robinson (of course) [36]. The result in Section 16.4 is a special case of a much more general class of results due to Mal'cev [118].

Large cardinals

In this chapter, we investigate what happens when some definitions and results from throughout this book involving \aleph_0 are asked to hold for uncountable cardinals instead. It turns out that the resulting cardinals are so large that they cannot even be proven to exist in ZFC. The discussion of these large cardinals begins with the notion of worldly cardinals in Section 17.1 and inaccessible cardinals in Section 17.2. Section 17.3 discusses the notion of measurable cardinals in depth, which are cardinals that possess a maximally complete ultrafilter. Section 17.4 begins with a discussion of infinitary logic, ultimately leading to the notions of weakly and strongly compact cardinals. The last kind of large cardinal introduced is a Ramsey cardinal in Section 17.5. The modern viewpoint of large cardinals as critical points of elementary embeddings is presented in Section 17.6. Finally, in Section 17.7, we present a theorem of Martin showing that the existence of a Ramsey cardinal allows one to prove analytic determinacy (as introduced in Section 5.3), which in turn implies that all Σ_2^1 sets of reals are both Lebesgue and Baire measurable.

We advise the reader that this chapter assumes a fair amount more set theory than the previous chapters. The uninitiated reader may wish to consult Appendix B before proceeding.

17.1. Worldly cardinals

Any discussion of large cardinals should really start from the notion of worldly cardinals. To motivate these cardinals, we first mention the following:

Fact 17.1.1. For any uncountable cardinal κ , V_κ is a model of all the axioms for ZFC except for possibly the replacement axiom.

Definition 17.1.2. An uncountable cardinal κ is called **worldly** if $V_\kappa \models$ ZFC.

In other words, κ is worldly if and only if V_κ satisfies the replacement axiom: whenever $x \in V_\kappa$ and $f : x \rightarrow V_\kappa$ is a function definable over V_κ , then the image of f also belongs to V_κ .

Theorem 17.1.3. ZFC does not prove that there exist worldly cardinals.

Proof. If ZFC proved that there exist worldly cardinals, then ZFC proves $\text{Con}(\text{ZFC})$, contradicting Gödel's second incompleteness theorem. \square

In fact, the following stronger version of the previous theorem is true:

Theorem 17.1.4. $\text{ZFC} \vdash \text{Con}(\text{ZFC}) \rightarrow \text{Con}(\text{ZFC} + \neg \exists \kappa \text{ worldly})$.

Proof. Suppose that λ is the least worldly cardinal. We claim that V_λ is a model of ZFC without any worldly cardinals. Indeed, if $\kappa < \lambda$ is a cardinal such that V_λ believes that κ is worldly, then by absoluteness, κ is a worldly cardinal in V itself, contradicting that λ is the least worldly cardinal. \square

On the other hand:

Theorem 17.1.5. $\text{ZFC} \not\vdash \text{Con}(\text{ZFC}) \rightarrow \text{Con}(\text{ZFC} + \exists \kappa \text{ worldly})$.

Proof. If $\text{ZFC} \vdash \text{Con}(\text{ZFC}) \rightarrow \text{Con}(\text{ZFC} + \exists \kappa \text{ worldly})$, then since $\text{ZFC} + \exists \kappa \text{ worldly} \vdash \text{Con}(\text{ZFC})$, we get

$$\text{ZFC} + \exists \kappa \text{ worldly} \vdash \text{Con}(\text{ZFC} + \exists \kappa \text{ worldly}),$$

contradicting Gödel's second incompleteness theorem again. \square

We phrase the conclusion of the previous theorem as: the theory $\text{ZFC} + \exists \kappa \text{ worldly}$ has **larger consistency strength** than ZFC. Note the difference between this and, say, the independence of CH. In this latter situation, $\text{Con}(\text{ZFC})$ and $\text{Con}(\text{ZFC} + \text{CH})$ are equivalent in ZFC, whence they have the same consistency strength (and likewise for $\text{Con}(\text{ZFC})$ and $\text{Con}(\text{ZFC} + \neg \text{CH})$).

Remark 17.1.6. One can view the axiom “there exists a worldly cardinal” as a beefed-up version of the axiom of infinity. Indeed, let $\text{ZFC} - \text{Inf}$ denote the axioms of ZFC without the axiom of infinity. Then since V_ω is a model of $\text{ZFC} - \text{Inf}$, we see that $\text{ZFC} \vdash \text{Con}(\text{ZFC} - \text{Inf})$, and thus ZFC has larger consistency strength than $\text{ZFC} - \text{Inf}$.

In what follows, we always consider ZFC as our “base theory” and measure consistency strength relative to ZFC. We will thus omit ZFC from our theories and so, for example, refer to the existence of a worldly cardinal when we really mean ZFC together with the existence of a worldly cardinal.

17.2. Inaccessible cardinals

As mentioned in the previous section, we will consider large cardinal notions motivated by asking that uncountable cardinals satisfy some properties satisfied by \aleph_0 . The first of these is motivated by considering the fact that \aleph_0 is a regular limit cardinal. In fact, \aleph_0 is a regular *strong limit* cardinal in the sense of the following definition.

Definition 17.2.1. Let κ be a cardinal. We say that κ is a **strong limit cardinal** if, for every cardinal $\lambda < \kappa$, we have that $2^\lambda < \kappa$.

Of course, GCH implies that the notions of limit cardinal and strong limit cardinal are the same. Uncountable (strong) limit cardinals are not hard to find:

Example 17.2.2.

- (1) \aleph_ω is a limit cardinal.
- (2) \beth_ω is a strong limit cardinal, where $\beth_0 := \aleph_0$, $\beth_{\alpha+1} := 2^{\beth_\alpha}$, and $\beth_\gamma := \sup_{\beta < \gamma} \beth_\beta$ when γ is a limit ordinal.

Note that the cardinals in the previous example are singular. It becomes more difficult to find an example of a (strong) limit cardinal that is also regular. We give such cardinals a name:

Definition 17.2.3. Let κ be an uncountable cardinal. We say that κ is a **strongly inaccessible cardinal** (resp., **weakly inaccessible cardinal**) if it is both regular and a strong limit cardinal (resp., limit cardinal).

Exercise 17.2.4. Suppose that κ is strongly inaccessible.

- (1) For each $\alpha < \kappa$, $|V_\alpha| < \kappa$.
- (2) For each $x \subseteq V_\kappa$, we have $x \in V_\kappa$ if and only if $|x| < \kappa$.

Armed with Exercise 17.2.4, we can now establish:

Theorem 17.2.5. *Strongly inaccessible cardinals are worldly.*

Proof. Suppose that $x \in V_\kappa$ and $f : x \rightarrow V_\kappa$ is a function. Let $y \subseteq V_\kappa$ be the image of f . By Exercise 17.2.4, $|x| < \kappa$, whence $|y| \leq |x| < \kappa$, and thus $y \in V_\kappa$ by Exercise 17.2.4 again. It follows that V_κ satisfies the replacement axiom. \square

As a consequence of the previous theorem, ZFC cannot prove that strongly inaccessible cardinals exist, and, in fact, the existence of a strongly inaccessible cardinal has higher consistency strength than ZFC alone.

Remark 17.2.6. The existence of a strongly inaccessible cardinal has higher consistency strength than the existence of a worldly cardinal. Indeed, one can show that the smallest worldly cardinal, should it exist, is singular (in fact has cofinality \aleph_0), and thus, in particular, is not strongly inaccessible. Thus, if κ is a strongly inaccessible cardinal, then the smallest worldly cardinal is below κ and thus belongs to V_κ . It follows that $\text{ZFC} + \exists\kappa \text{ strongly inaccessible} \vdash \text{Con}(\text{ZFC} + \exists\kappa \text{ worldly})$.

As mentioned in the preceding remark, worldly cardinals can be singular. However, for regular cardinals, they are the same as strongly inaccessible:

Proposition 17.2.7. *A worldly cardinal is a strong limit cardinal. In particular, a regular cardinal is worldly if and only if it is strongly inaccessible.*

Proof. Suppose that κ is a worldly cardinal. Fix $\lambda < \kappa$; we show that $2^\lambda < \kappa$. For this, it suffices to show that $2^\lambda \in V_\kappa$. Note that $\mathcal{P}(\lambda) \in V_{\lambda+2} \subseteq V_\kappa$ and $\mathcal{P}(\lambda)$ is the same whether it is calculated in V or V_κ . Since $V_\kappa \models \text{ZFC}$, there is a minimal $\gamma \in V_\kappa$ for which there is a bijection $f : \gamma \rightarrow \mathcal{P}(\lambda)$ that belongs to V_κ . Consequently, $2^\lambda = \gamma$ belongs to V_κ , as desired. \square

Remark 17.2.8. The existence of a weakly inaccessible cardinal is equiconsistent with the existence of a strongly inaccessible cardinal, for a weakly inaccessible cardinal in V becomes strongly inaccessible in L .

We can posit a stronger large cardinal axiom by asserting that there are two strongly inaccessible cardinals, for if $\kappa_1 < \kappa_2$ are strongly inaccessible cardinals, then V_{κ_2} is a model of ZFC with a strongly inaccessible cardinal. In general, for any ordinals $\beta < \alpha$, asserting that there is a sequence of strongly inaccessible cardinals of length α is stronger than asserting the existence of such a sequence of length β .

Of course, one can then posit the existence of a proper class of strongly inaccessible cardinals, and such a theory has stronger consistency strength than the existence of a sequence of strongly inaccessible cardinals of any given ordinal length.

A strongly inaccessible cardinal κ is said to be **2-inaccessible** if κ is a limit of strongly inaccessible cardinals. Equivalently, κ is 2-inaccessible if κ is the κ th strongly inaccessible cardinal. It is a fact that the existence of a 2-inaccessible cardinal has higher consistency strength than the existence of a proper class of strongly inaccessible cardinals. Then there are 3-inaccessible cardinals, which are strongly inaccessible limits of 2-inaccessible cardinals.

One can continue in this manner, leading to the definition of α -**inaccessible cardinals** for every ordinal α . The interested reader can consult [94] for a proper definition.

We now turn to a large cardinal notion that is even stronger than 2-inaccessible. To do so, we need to introduce some terminology.

Definition 17.2.9. Suppose that κ is a regular uncountable cardinal and $C \subseteq \kappa$. We say that C is:

- (1) **closed in κ** if, for every $X \subseteq C$ with $\sup X < \kappa$, we have that $\sup X \in C$;
- (2) **unbounded in κ** if, for every $\alpha < \kappa$, there is $\beta \in C$ such that $\alpha < \beta$;
- (3) **club in κ** if it is both closed in κ and unbounded in κ .

Exercise 17.2.10. Suppose that κ is a strongly inaccessible cardinal. Prove that the following subsets of κ are club subsets of κ :

- (1) $\{\lambda < \kappa : \lambda \text{ is a cardinal}\}$,
- (2) $\{\lambda < \kappa : \lambda \text{ is a limit cardinal}\}$,
- (3) $\{\lambda < \kappa : \lambda \text{ is a strong limit cardinal}\}$.

Lemma 17.2.11. *If C and D are club subsets of κ , then so is $C \cap D$.*

Proof. It is fairly clear that $C \cap D$ is closed in κ . To see that $C \cap D$ is unbounded in κ , fix $\alpha < \kappa$ and take $\alpha_1 \in C$ with $\alpha < \alpha_1$; this is possible since C is unbounded in κ . Then take $\alpha_2 \in D$ with $\alpha_1 < \alpha_2$; this is possible since D is unbounded in κ . Continuing in this fashion, we construct an increasing sequence $(\alpha_n)_{n < \omega}$ of ordinals less than κ with $\alpha_n \in C$ for n odd and $\alpha_n \in D$ for n even. Let $\alpha' := \sup_n \alpha_n$. Since $\text{cof}(\kappa) > \omega$, we have that $\alpha' < \kappa$, whence $\alpha' \in C \cap D$ since C and D are both closed in κ . Since $\alpha < \alpha'$, we see that $C \cap D$ is unbounded in κ , as desired. \square

Exercise 17.2.12. If $\lambda < \kappa$ and $(C_\alpha)_{\alpha < \lambda}$ are all club subsets of κ , then $\bigcap_{\alpha < \lambda} C_\alpha$ is also a club subset of κ .

Definition 17.2.13. The **club filter on κ** , denoted $\text{Club}(\kappa)$, is the filter on κ generated by the club subsets of κ .

It follows that $D \in \text{Club}(\kappa)$ if and only if there is a club subset C of κ such that $C \subseteq D$. The previous exercise shows that $\text{Club}(\kappa)$ is a κ -complete filter on κ . $\text{Club}(\kappa)$ is never an ultrafilter on κ ; see [90, Theorem 12.5]. (Interestingly enough, the proof for $\kappa = \aleph_1$ necessarily uses AC, for it is consistent with ZF that $\text{Club}(\aleph_1)$ is an ultrafilter; see [90, Theorem 12.12].)

Definition 17.2.14. $S \subseteq \kappa$ is **stationary** if the filter generated by $\text{Club}(\kappa)$ and S is a proper filter on κ .

In other words, S is stationary if $S \cap C \neq \emptyset$ for all club subsets C of κ . Note that all club subsets of κ are stationary.

Definition 17.2.15. A cardinal κ is **Mahlo** if it is strongly inaccessible and

$$\{\lambda < \kappa : \lambda \text{ is a strongly inaccessible cardinal}\}$$

is a stationary subset of κ .

Since we think of stationary subsets of κ as large subsets of κ , a Mahlo cardinal is a strongly inaccessible cardinal such that there are a lot of strongly inaccessible cardinals below κ .

Exercise 17.2.16. Mahlo cardinals are 2-inaccessible.

One can show that the converse of the previous exercise is not true: if κ is the least 2-inaccessible cardinal, then κ is not Mahlo (see [89, Exercise 8.6]). In fact, the existence of a Mahlo cardinal has higher consistency strength than the existence of a 2-inaccessible cardinal.

Perhaps positing the existence of cardinals such as Mahlo cardinals seems a bit ridiculous. However, it turns out that such large cardinals are “small” in a quasi-technical sense. Moreover, starting with the next section, we will turn our attention to large cardinal notions that appear rather organically in terms of the themes presented throughout this book and yet these cardinals are indeed “larger” than the larger cardinals presented thus far.

17.3. Measurable cardinals

We recall the following definition from Remark 6.6.10:

Definition 17.3.1. An uncountable cardinal κ is **measurable** if there is a nonprincipal κ -complete ultrafilter on κ .

Since all ultrafilters are \aleph_0 -complete, the previous definition is indeed the uncountable generalization of the fact that nonprincipal ultrafilters on ω exist.

Exercise 17.3.2. Measurable cardinals are regular.

Measurable cardinals are large:

Theorem 17.3.3. *Measurable cardinals are strongly inaccessible.*

Proof. By the previous exercise, we know that measurable cardinals are regular. It remains to show that they are strong limit cardinals. We prove the contrapositive, whence we assume there is some $\lambda < \kappa$ such that $\kappa \leq 2^\lambda$. We may thus find $S \subseteq \mathcal{P}(\lambda)$ such that $|S| = \kappa$. Suppose that \mathcal{U} is a κ -complete ultrafilter on S ; it suffices to show that \mathcal{U} is principal. For each

$\alpha \in \lambda$, let $\mathcal{X}_\alpha := \{X \in S : \alpha \in X\}$, and let $\epsilon_\alpha \in \{-1, 1\}$ be such that $\mathcal{X}_\alpha^{\epsilon_\alpha} \in \mathcal{U}$, where $\mathcal{X}_\alpha^1 := \mathcal{X}_\alpha$ and $\mathcal{X}_\alpha^{-1} := S \setminus \mathcal{X}_\alpha$. By κ -completeness, $\bigcap_{\alpha < \lambda} \mathcal{X}_\alpha^{\epsilon_\alpha} \in \mathcal{U}$. It remains to note that $\bigcap_{\alpha < \lambda} \mathcal{X}_\alpha^{\epsilon_\alpha}$ contains at most one element, namely $\{\alpha < \lambda : \epsilon_\alpha = 1\}$. It follows that \mathcal{U} is principal. \square

The next result shows that the existence of a countably complete nonprincipal ultrafilter is equiconsistent with a measurable cardinal, a result which was alluded to in Remark 6.6.10.

Proposition 17.3.4. *Suppose that κ is the least cardinal that possesses a countably complete nonprincipal ultrafilter. Then κ is measurable.*

Proof. Let \mathcal{U} be a countably complete nonprincipal ultrafilter on κ . Suppose, toward a contradiction, that κ is not measurable, whence there are $\lambda < \kappa$ and sets $(X_i)_{i < \lambda}$, with $X_i \notin \mathcal{U}$ for each $i < \lambda$ and yet $\bigcup_{i < \lambda} X_i \in \mathcal{U}$. We may also assume that the X_i 's are pairwise disjoint by replacing each X_i by $X_i \setminus \bigcup_{j < i} X_j$. We now define $\mathcal{V} \subseteq \mathcal{P}(\lambda)$ by declaring, for $Y \subseteq \lambda$, that $Y \in \mathcal{V}$ if and only if $\bigcup_{i \in Y} X_i \in \mathcal{U}$. We complete the proof by verifying that \mathcal{V} is a countably complete nonprincipal ultrafilter on λ ; since $\lambda < \kappa$, this contradicts the defining property of κ .

It is clear that $\emptyset \notin \mathcal{V}$ and $\lambda \in \mathcal{V}$ by the choice of the sets X_i . It is also clear that \mathcal{V} is closed under superset. We now verify that \mathcal{V} is closed under countable intersections. Suppose that $Y_n \in \mathcal{V}$ for each $n < \omega$. Since the sets $(X_i)_{i < \lambda}$ are pairwise disjoint, we have that $\bigcup_{i \in \bigcap_{n < \omega} Y_n} X_i = \bigcap_{n < \omega} \bigcup_{i \in Y_n} X_i$; since \mathcal{U} is countably complete, we see that $\bigcap_{n < \omega} \bigcup_{i \in Y_n} X_i \in \mathcal{U}$, whence $\bigcap_{n < \omega} Y_n \in \mathcal{V}$, as desired. Finally, since each $X_i \notin \mathcal{U}$, we have that \mathcal{V} is nonprincipal. \square

We now show that every measurable cardinal κ possesses a κ -complete nonprincipal ultrafilter that has a nice extra technical property that we will use in what follows.

Lemma 17.3.5. *Suppose that κ is an uncountable cardinal and \mathcal{U} is a κ -complete nonprincipal ultrafilter on κ . Furthermore, set $\mathcal{L} := \{R\}$, with R a binary relation and consider the \mathcal{L} -structure $\mathcal{M} := (\kappa, \epsilon)$. Then:*

- (1) $R^{\mathcal{M}^{\mathcal{U}}}$ is a well-ordering.
- (2) $d(\kappa)$ is an initial segment of $R^{\mathcal{M}^{\mathcal{U}}}$, that is, for each $\alpha < \kappa$, $d(\alpha)$ is the α th element of $\mathcal{M}^{\mathcal{U}}$.
- (3) The order type of $R^{\mathcal{M}^{\mathcal{U}}}$ is strictly larger than κ .

Proof. (1). It is clear from Łoś's theorem that $R^{\mathcal{M}^{\mathcal{U}}}$ is a linear ordering. To see that it is a well-ordering, suppose, toward a contradiction, that $[f_0]_{\mathcal{U}} > [f_1]_{\mathcal{U}} > \dots$ is a strictly decreasing sequence. For each $n \in \omega$, let

$X_n := \{i \in \kappa : f_n(i) > f_{n+1}(i)\}$. By assumption, each $X_n \in \mathcal{U}$, whence, by countable completeness, there is $i \in \bigcap_{n < \omega} X_n$. We then have that $f_0(i) > f_1(i) > f_2(i) > \dots$, contradicting that κ is an ordinal.

One proves (2) by induction on $\alpha < \kappa$. Suppose this is true for all $\beta < \alpha$. We must show that the predecessors of $d(\alpha)$ are precisely the $d(\beta)$ for $\beta < \alpha$. Since d is an embedding, we have $R^{\mathcal{M}^{\mathcal{U}}}(d(\beta), d(\alpha))$, so each such $d(\beta)$ is indeed a predecessor of $d(\alpha)$. Now suppose that $[f]_{\mathcal{U}} \in \mathcal{M}^{\mathcal{U}}$ is such that $R^{\mathcal{M}^{\mathcal{U}}}([f]_{\mathcal{U}}, d(\alpha))$. For each $\beta < \alpha$, let $X_\beta := \{i \in \kappa : f(i) = \beta\}$. Then by assumption $\bigcup_{\beta < \alpha} X_\beta \in \mathcal{U}$. By κ -completeness, there is some $\beta < \alpha$ such that $X_\beta \in \mathcal{U}$, whence $[f]_{\mathcal{U}} = d(\beta)$.

To prove (3), we first note that, by (2), the order type of $R^{\mathcal{M}^{\mathcal{U}}}$ is at least κ . However, by Exercise 6.6.14, since \mathcal{U} is not κ^+ -complete, d is not onto, which completes the proof. \square

Let us refer to the order type of $R^{\mathcal{M}^{\mathcal{U}}}$ above as $\kappa^{\mathcal{U}}$. The question now arises: what is the κ th element of $\kappa^{\mathcal{U}}$?

Definition 17.3.6. A κ -complete nonprincipal ultrafilter \mathcal{U} on the uncountable cardinal κ is called **normal** if $[\text{id}]_{\mathcal{U}}$ is the κ th element of $\kappa^{\mathcal{U}}$, where $\text{id} : \kappa \rightarrow \kappa$ is the identity map.

Given an ultrafilter \mathcal{U} on κ and a function $g : \kappa \rightarrow \kappa$, we say that g is **\mathcal{U} -regressive** if $g(\alpha) < \alpha$ for \mathcal{U} -many $\alpha \in \kappa$. Here is a useful reformulation of normality:

Exercise 17.3.7. A κ -complete nonprincipal ultrafilter \mathcal{U} on the uncountable cardinal κ is normal if and only if whenever $g : \kappa \rightarrow \kappa$ is \mathcal{U} -regressive, then g is constant on a set in \mathcal{U} .

Here is an example of the utility of normal ultrafilters:

Proposition 17.3.8. *Suppose that κ is a cardinal of uncountable cofinality and \mathcal{U} is a normal ultrafilter on κ . Then $\text{Club}(\kappa) \subseteq \mathcal{U}$. In particular, every element of \mathcal{U} is stationary.*

Proof. Suppose toward a contradiction that $C \subseteq \kappa$ is a club subset with $\kappa \setminus C \in \mathcal{U}$. For $\alpha < \kappa$, define $f(\alpha) := \sup\{\beta \in C : \beta < \alpha\}$. Note that, since C is closed, we have that $f(\alpha) < \alpha$ for $\alpha \in \kappa \setminus C$, whence f is \mathcal{U} -regressive. It follows that there is $X \in \mathcal{U}$ and $\xi < \kappa$ such that $f(\alpha) = \xi$ for all $\alpha \in X$. Since C is unbounded in κ , there must exist $\beta \in C$ such that $\xi < \beta$. Also, since \mathcal{U} is nonprincipal and κ -complete, X is unbounded in κ , whence there is $\alpha \in X$ such that $\beta < \alpha$. Thus, $f(\alpha) \geq \beta > \xi$, which is a contradiction. \square

We now show that a measurable cardinal always possesses a normal ultrafilter:

Proposition 17.3.9. *If κ is a measurable cardinal, then there is a normal ultrafilter on κ .*

Proof. Fix an arbitrary nonprincipal κ -complete ultrafilter \mathcal{U} on κ . Fix $f : \kappa \rightarrow \kappa$ such that $[f]_{\mathcal{U}}$ is the κ th element of $\kappa^{\mathcal{U}}$. We claim that $f(\mathcal{U})$ is a normal ultrafilter on κ . By Exercise 6.6.8, $f(\mathcal{U})$ is κ -complete. Note that $f(\mathcal{U})$ is also nonprincipal: for $\alpha < \kappa$, $\{\alpha\} \in f(\mathcal{U})$ if and only if $f^{-1}(\alpha) \in \mathcal{U}$, which implies that $[f]_{\mathcal{U}} = d(\alpha)$, contradicting the fact that $[f]_{\mathcal{U}}$ is the κ th element of $\kappa^{\mathcal{U}}$ while $d(\alpha)$ is the α th element. Let \mathcal{M} be as in Lemma 17.3.5. As discussed in Exercise 6.7.3, $\mathcal{M}^{\mathcal{U}}[f] \cong \mathcal{M}^{f(\mathcal{U})}$ and the isomorphism sends $[f]_{\mathcal{U}}$ to $[\text{id}]_{f(\mathcal{U})}$; since $[f]_{\mathcal{U}}$ is clearly the κ th element of $\mathcal{M}^{\mathcal{U}}[f]$, we are done. \square

Corollary 17.3.10. *Measurable cardinals are Mahlo.*

Proof. Suppose that κ is measurable. We already know that κ is strongly inaccessible. To complete the proof, it suffices to show that the set F of regular cardinals below κ is stationary. Indeed, by Exercise 17.2.10, the set F' of strong limit cardinals below κ is club, whence $F \cap F'$ is stationary, as desired.

Let \mathcal{U} be a normal ultrafilter on κ . By Proposition 17.3.8, it suffices to show that $F \in \mathcal{U}$. Suppose, toward a contradiction, that $\kappa \setminus F \in \mathcal{U}$. It follows that the function $\alpha \mapsto \text{cof}(\alpha)$ is \mathcal{U} -regressive, whence, by Exercise 17.3.7, there is $\lambda < \kappa$ such that $E_{\lambda} \in \mathcal{U}$, where $E_{\lambda} := \{\alpha \in \kappa : \text{cof}(\alpha) = \lambda\}$. For each $\alpha \in E_{\lambda}$, there is an increasing sequence $(x_{\alpha,\xi})_{\xi < \lambda}$ such that $\sup_{\xi < \lambda} x_{\alpha,\xi} = \alpha$. For each $\xi < \lambda$, there is some $y_{\xi} < \kappa$ and some $A_{\xi} \in \mathcal{U}$ such that $x_{\alpha,\xi} = y_{\xi}$ for all $\alpha \in A_{\xi}$; this is because the map $\alpha \mapsto x_{\alpha,\xi}$ is \mathcal{U} -regressive. By κ -completeness, $\bigcap_{\xi < \lambda} A_{\xi} \in \mathcal{U}$. However, $\bigcap_{\xi < \lambda} A_{\xi} := \{\sup_{\xi < \lambda} y_{\xi}\}$, contradicting that \mathcal{U} is nonprincipal. \square

It is a fact that the existence of a measurable cardinal has higher consistency strength than the existence of a Mahlo cardinal.

We end this section with a digression on how large cardinals need not be so large if one does not assume the axiom of choice. For example:

Theorem 17.3.11. *If there is a measurable cardinal, then it is consistent with ZF that \aleph_1 is measurable.*

The previous theorem is beyond the scope of this book (see [90, Theorem 12.2] for a proof), but we can prove something of a similar nature.

In the rest of this section, we work in ZF. Below, AD refers to the axiom of determinacy as introduced in Section 5.3.

Theorem 17.3.12 (Solovay [165]). *AD implies that \aleph_1 is measurable.*

We give Martin's proof [127] of Theorem 17.3.12, which relies on the theory of *computability*. (The reader unfamiliar with computability theory can refer to the friendly introduction [51].)

Definition 17.3.13. Given $x, y \in \mathbb{N}^{\mathbb{N}}$, we say that x is **computable from** y (or that x is **Turing reducible to** y), denoted $x \leq_T y$, if there is a computer program with access to y such that, given $n \in \mathbb{N}$, after a finite period of time it will halt and return the value $x(n)$.

Of course the previous definition is not mathematically precise, relying on the intuitive notion of computer program. There are, however, several formalizations of this notion that one can pursue. For our purposes, we will rely on this intuitive definition and some of its basic properties, listed below:

Facts 17.3.14. The relation \leq_T on $\mathbb{N}^{\mathbb{N}}$ satisfies the following properties:

- (1) \leq_T is reflexive and transitive.
- (2) For every $x \in \mathbb{N}^{\mathbb{N}}$, the set $\{y \in \mathbb{N}^{\mathbb{N}} : y \leq_T x\}$ is countable.
- (3) If $(x_n)_{n \in \mathbb{N}}$ is a countable sequence from $\mathbb{N}^{\mathbb{N}}$, then there is $x \in \mathbb{N}^{\mathbb{N}}$ such that $x_n \leq_T x$ for all $n \in \mathbb{N}$.
- (4) For every $x \in \mathbb{N}^{\mathbb{N}}$, there is $y \in \mathbb{N}^{\mathbb{N}}$ such that $x \leq_T y$ but $y \not\leq_T x$.

Proof idea. (1) follows immediately from the definition of \leq_T , and (2) follows from the fact that there are only countably many programs that have access to a particular $x \in \mathbb{N}^{\mathbb{N}}$.

For (3), take any $x \in \mathbb{N}^{\mathbb{N}}$ such that, for all $m, n \in \mathbb{N}$, we have $x(p_n^{m+1}) := x_n(m)$, where p_n is the $(n+1)$ -st prime.

Finally, (4) can be proven by recursion using a diagonalization argument over all possible functions computable by x to ensure that the function constructed is not computable from x . \square

By Fact 17.3.14(1), we get an equivalence relation \equiv_T on $\mathbb{N}^{\mathbb{N}}$ by declaring $x \equiv_T y$ if and only if $x \leq_T y$ and $y \leq_T x$. The equivalence class of x is called the **(Turing) degree** of x and is denoted \mathbf{x} . The relation \leq_T descends to a partial order, also denoted \leq_T , on the set \mathcal{D} of degrees by setting $\mathbf{x} \leq_T \mathbf{y}$ if and only if $x \leq_T y$.

Definition 17.3.15. Given a degree $\mathbf{x} \in \mathcal{D}$, we set the **cone of** \mathbf{x} to be the set $c(\mathbf{x}) := \{\mathbf{y} : \mathbf{x} \leq_T \mathbf{y}\}$. A **cone in** \mathcal{D} is a set of the form $c(\mathbf{x})$ for some $\mathbf{x} \in \mathcal{D}$.

The key to Martin's proof of Theorem 17.3.12 is the following:

Theorem 17.3.16 (Martin). *Suppose that AD holds and $\mathcal{A} \subseteq \mathcal{D}$. Then either \mathcal{A} contains a cone or $\mathcal{D} \setminus \mathcal{A}$ contains a cone.*

Proof. Let $A := \{x \in \mathbb{N}^{\mathbb{N}} : \mathbf{x} \in \mathcal{A}\}$. We claim that if player I has a winning strategy in $\mathcal{G}(A)$, then \mathcal{A} contains a cone while if player II has a winning strategy in $\mathcal{G}(A)$, then $\mathcal{D} \setminus \mathcal{A}$ contains a cone. We will only prove the former case, leaving the latter to the reader.

Suppose that $\sigma : \mathbb{N}^{<\mathbb{N}} \rightarrow \mathbb{N}^{<\mathbb{N}}$ is a winning strategy for player I. By using familiar Gödel numbering of finite sequences, we may view σ as an element of $\mathbb{N}^{\mathbb{N}}$ and thus can consider its degree σ .

We claim $c(\sigma) \subseteq \mathcal{A}$, finishing the proof of the theorem. Toward this end, suppose that $\mathbf{x} \in c(\sigma)$. Consider a play of the game where player II plays x and player I plays according to σ . This produces a sequence $y \in \mathbb{N}^{\mathbb{N}}$ such that $\mathbf{y} \in \mathcal{A}$ (as player I plays according to its winning strategy). It remains to note that $\mathbf{y} = \mathbf{x}$. Indeed, $x \leq_T y$ (as it is just the “even part” of y) while y is computable from x and σ ; but since σ is computable from x , we have that y is computable from x alone, that is, $y \leq_T x$. \square

We will need one other ingredient:

Lemma 17.3.17. *AD implies that \aleph_1 is a regular cardinal.*

Proof. It is standard that the regularity of \aleph_1 follows from the *countable axiom of choice for \mathbb{R}* , which states that every countable family of nonempty subsets of $\mathbb{N}^{\mathbb{N}}$ has a choice function. We show that this latter statement follows from AD.

Let X_0, X_1, \dots be countably many nonempty subsets of $\mathbb{N}^{\mathbb{N}}$. We set

$$\mathcal{X} := \{a \in \mathbb{N}^{\mathbb{N}} : (a(1), a(3), a(5), \dots) \notin X_{a(0)}\}.$$

By AD, $G_{\mathcal{X}}$ is determined. However, player I cannot have a winning strategy in $G_{\mathcal{X}}$, for if the strategy told player I to start with n , then player II can win by just playing some sequence $(a(1), a(3), a(5), \dots) \in X_n$, contradicting that the strategy was winning for player I. Consequently, it must be player II who has the winning strategy. We can thus define a choice function for the above family by setting $F(n)$ to be the sequence of moves player II makes according to the winning strategy if player I plays $(n, 0, 0, \dots)$. \square

Remark 17.3.18. The previous lemma might seem strange at first, but in models of ZF without choice, \aleph_1 need not be regular.

Proof of Theorem 17.3.12. Let $\mathcal{V} := \{\mathcal{A} \subseteq \mathcal{D} : \mathcal{A} \text{ contains a cone}\}$.

Claim 1. *\mathcal{V} is a nonprincipal countably complete ultrafilter on \mathcal{D} .*

Proof of Claim 1. It is clear that $\emptyset \notin \mathcal{V}$. To see that $\mathcal{D} \in \mathcal{V}$, note that the cone determined by any computable sequence (such as a constant sequence) is all of \mathcal{D} . It also clear that \mathcal{V} is closed under supersets. Theorem 17.3.16 shows that, given any $\mathcal{A} \subseteq \mathcal{D}$, either $\mathcal{A} \in \mathcal{V}$ or $\mathcal{D} \setminus \mathcal{A} \in \mathcal{V}$. \mathcal{V} is nonprincipal

since there is more than one cone by Fact 17.3.14(4). It remains to prove closure under countable intersections. Suppose that $\mathcal{A}_n \in \mathcal{V}$ for all $n \in \mathbb{N}$. By the countable axiom of choice for \mathbb{R} (see the proof of Lemma 17.3.17), for each $n \in \mathbb{N}$, we can choose $x_n \in \mathbb{N}^{\mathbb{N}}$ such that $c(\mathbf{x}_n) \subseteq \mathcal{A}_n$. By Fact 17.3.14(3), there is $x \in \mathbb{N}^{\mathbb{N}}$ such that $x_n \leq_T x$ for all $n \in \mathbb{N}$. It follows that $c(\mathbf{x}) \subseteq \bigcap_{n \in \mathbb{N}} \mathcal{A}_n$, whence $\bigcap_n \mathcal{A}_n \in \mathcal{V}$.

Let $G : \mathbb{N}^{\mathbb{N}} \rightarrow \aleph_1$ be a surjective map (which exists without needing to use AC) and define $F : \mathcal{D} \rightarrow \aleph_1$ by setting $F(\mathbf{x}) := \sup\{G(y) : y \leq_T x\}$. By Fact 17.3.14(2), we are indeed taking the supremum of a countable set of ordinals, which, by Lemma 17.3.17, is still a countable ordinal, whence F really does take values in \aleph_1 .

We now set $\mathcal{U} := F(\mathcal{V})$, an ultrafilter on \aleph_1 . Since $\mathcal{U} \leq_{RK} \mathcal{V}$, it follows that \mathcal{U} is countably complete. To finish the proof of the theorem, it suffices to establish:

Claim 2. \mathcal{U} is nonprincipal.

Proof of Claim 2. It suffices to show that, for any $\alpha < \aleph_1$, we have that $\{\mathbf{x} : F(\mathbf{x}) \geq \alpha\} \in \mathcal{V}$. Take $y \in \mathbb{N}^{\mathbb{N}}$ such that $G(y) = \alpha$. By definition, $c(\mathbf{y}) \in \mathcal{V}$. However, $\mathbf{y} \leq_T \mathbf{x}$ implies $F(\mathbf{x}) \geq G(y) = \alpha$, that is, $\{\mathbf{x} : F(\mathbf{x}) \geq \alpha\} \supseteq c(\mathbf{y})$, finishing the proof. \square

17.4. Strongly and weakly compact cardinals

In this section, we consider the following “infinitary” extensions of first-order logic:

Definition 17.4.1. Given a language \mathcal{L} and a cardinal κ , the set of \mathcal{L}_κ -formulae are obtained by allowing two additional ways of constructing new formulae:

- If $\alpha < \kappa$ and $(\varphi_\beta)_{\beta < \alpha}$ is a family of \mathcal{L}_κ -formulae, then so is $\bigwedge_{\beta < \alpha} \varphi_\beta$.
- If $\alpha < \kappa$, $(x_\beta)_{\beta < \alpha}$ is a sequence of variables and φ is an \mathcal{L}_κ -formula, then so is $(\exists x_\beta)_{\beta < \alpha} \varphi$.

Note that the \mathcal{L}_{\aleph_0} -formulae are the same as the ordinary first-order \mathcal{L} -formulae. One can extend the semantics of ordinary first-order logic in the obvious way to define the truth of an \mathcal{L}_κ -formula in an \mathcal{L} -structure.

Here is an example of the greater expressive power of infinitary logics:

Exercise 17.4.2. Let $\mathcal{L} = \{R\}$, where R is a binary relation. Show that there is an \mathcal{L}_{\aleph_1} -sentence σ such that, for any \mathcal{L} -structure \mathcal{M} , $\mathcal{M} \models \sigma$ if and only if $R^{\mathcal{M}}$ is a well-ordering of M .

Given the obvious connection between intersections and conjunctions, the following extension of Łoś's theorem should not be too surprising:

Theorem 17.4.3 (Łoś's theorem for \mathcal{L}_κ). *Suppose that $(\mathcal{M}_i)_{i \in I}$ is a family of \mathcal{L} -structures and \mathcal{U} is a κ -complete ultrafilter on I . Further suppose that $\alpha < \kappa$, $\varphi((x_\beta)_{\beta < \alpha})$ is an \mathcal{L}_κ -formula, and $[a_\beta]_{\mathcal{U}} \in \prod_{\mathcal{U}} M_i$ for $\beta < \alpha$. Then*

$$\prod_{\mathcal{U}} \mathcal{M}_i \models \varphi((a_\beta]_{\mathcal{U}})_{\beta < \alpha} \Leftrightarrow \{i \in I : \mathcal{M}_i \models \varphi((a_\beta(i))_{\beta < \alpha})\} \in \mathcal{U}.$$

Exercise 17.4.4. Prove Theorem 17.4.3.

One of the first issues that arises with the infinitary logics \mathcal{L}_κ is that the compactness theorem no longer holds!

Exercise 17.4.5. Let $\mathcal{L} = \{c, c_0, c_1, \dots\}$, and set

$$\Sigma := \{c \neq c_n : n < \omega\} \cup \{\forall x \bigvee_{n < \omega} x = c_n\}.$$

Show that Σ is finitely satisfiable but not satisfiable.

The compactness theorem for first-order logic can be viewed as the statement: whenever Σ is a set of sentences of \mathcal{L}_{\aleph_0} such that every subset of size $< \aleph_0$ has a model, then Σ has a model. The literal extension of this statement to uncountable cardinals leads to the following notion:

Definition 17.4.6. Suppose that κ is an uncountable cardinal. We say that κ is **strongly compact** if: whenever Σ is a set of sentences of \mathcal{L}_κ such that every subset of size $< \kappa$ has a model, then Σ has a model.

Exercise 17.4.7. Prove that strongly compact cardinals are regular.

The usual equivalence between the compactness theorem and the ultrafilter theorem can be extended to give an ultrafilter characterization of strongly compact cardinals:

Theorem 17.4.8. *Suppose that κ is an uncountable cardinal. Then κ is strongly compact if and only if, for every set I and every κ -complete filter \mathcal{F} on I , there is a κ -complete ultrafilter \mathcal{U} on I extending \mathcal{F} .*

Before proving this theorem, we establish some notation. Fix a set I with $|I| \geq \kappa$. We set $\mathcal{P}_\kappa(I) := \{x \subseteq I : |x| < \kappa\}$. For each $x \in \mathcal{P}_\kappa(I)$, set $\hat{x} := \{y \in \mathcal{P}_\kappa(I) : x \subseteq y\}$.

Exercise 17.4.9. Verify that the set

$$\{X \subseteq \mathcal{P}_\kappa(I) : \hat{x} \subseteq X \text{ for some } x \in \mathcal{P}_\kappa(I)\}$$

is a κ -complete filter on $\mathcal{P}_\kappa(I)$.

Proof of Theorem 17.4.8. We leave the proof of the forward direction as an exercise. We now prove the backward direction. Suppose that Σ is a set of \mathcal{L}_κ sentences such that every subset of Σ of size $< \kappa$ has a model. Let $\mathcal{F} := \{X \subseteq \mathcal{P}_\kappa(\Sigma) : \hat{x} \subseteq X \text{ for some } x \in \mathcal{P}_\kappa(\Sigma)\}$, a κ -complete filter on $\mathcal{P}_\kappa(\Sigma)$ by Exercise 17.4.9. By assumption, there is a κ -complete ultrafilter \mathcal{U} on $\mathcal{P}_\kappa(\Sigma)$ extending \mathcal{F} . For each $x \in \mathcal{P}_\kappa(\Sigma)$, let \mathcal{M}_x be an L -structure such that $\mathcal{M}_x \models x$. Set $\mathcal{M} := \prod_{\mathcal{U}} \mathcal{M}_x$. We claim that $\mathcal{M} \models \Sigma$. Indeed, given $\sigma \in \Sigma$, set $x := \{\sigma\}$. Since $\hat{x} \in \mathcal{F} \subseteq \mathcal{U}$ and $\mathcal{M}_y \models \sigma$ for all $y \in \mathcal{P}_\kappa(\Sigma)$ with $\sigma \in y$, we obtain that $\mathcal{M} \models \sigma$ from Łoś's theorem for \mathcal{L}_κ . \square

Exercise 17.4.10. Prove the forward direction of the previous theorem. (Note that this direction does not require κ to be uncountable. When $\kappa = \aleph_0$, the theorem states that the by-now familiar fact that the compactness theorem implies that every filter can be extended to an ultrafilter.)

Definition 17.4.11. For any cardinal κ and set I with $|I| \geq \kappa$, an ultrafilter \mathcal{U} on $\mathcal{P}_\kappa(I)$ is called **fine** if it is κ -complete and extends the filter $\{X \subseteq \mathcal{P}_\kappa(I) : \hat{x} \subseteq X \text{ for some } x \in \mathcal{P}_\kappa(I)\}$.

Thus, the proof of Theorem 17.4.8 actually showed:

Corollary 17.4.12. κ is strongly compact if and only if for any set I with $|I| \geq \kappa$, there is a fine ultrafilter on $\mathcal{P}_\kappa(I)$.

Corollary 17.4.13. Strongly compact cardinals are measurable.

Proof. Suppose that κ is strongly compact. Let $\mathcal{F} := \{X \subseteq \kappa : |\kappa \setminus X| < \kappa\}$. Note that \mathcal{F} is a κ -complete filter on κ . By strong compactness, there is a κ -complete ultrafilter \mathcal{U} on κ extending \mathcal{F} . It is immediate that \mathcal{U} is nonprincipal, whence it follows that κ is measurable. \square

The existence of a strongly compact cardinal has higher consistency strength than the existence of a measurable cardinal.

As we saw in the previous section, having normal κ -complete ultrafilters yielded desirable consequences. It thus makes sense to consider a notion of normality in this context as well. Given a cardinal κ , a set I with $|I| \geq \kappa$, a function $f : \mathcal{P}_\kappa(I) \rightarrow I$, and an ultrafilter \mathcal{U} on $\mathcal{P}_\kappa(I)$, we say that f is \mathcal{U} -regressive if $f(x) \in x$ for \mathcal{U} -almost all $x \in \mathcal{P}_\kappa(I)$.

Definition 17.4.14. Given a cardinal κ and a set I with $|I| \geq \kappa$, a fine ultrafilter \mathcal{U} on $\mathcal{P}_\kappa(I)$ is **normal** if whenever $f : \mathcal{P}_\kappa(I) \rightarrow I$ is \mathcal{U} -regressive, then f is constant on a set in \mathcal{U} .

Unlike the previous section, it is not the case that, whenever κ is strongly compact, there is a normal ultrafilter on $\mathcal{P}_\kappa(I)$. We need to give this stronger property a name. But what is stronger than “strong”? Why, “super”, of course!

Definition 17.4.15. If κ is an uncountable cardinal, then κ is called **supercompact** if, whenever $|I| \geq \kappa$, there is a normal ultrafilter on $\mathcal{P}_\kappa(I)$.

Remark 17.4.16. Why is it we say that being supercompact is stronger than strongly compact? Well, it is consistent that there is a model with a strongly compact cardinal that is not supercompact. It is also consistent that all strongly compact cardinals are supercompact. It is an interesting open question whether or not the existence of a supercompact cardinal is equiconsistent with the existence of a strongly compact cardinal.

We previously saw that strongly compact cardinals are measurable. It turns out measurable cardinals have a reformulation in terms of a compactness-like principle as well. In order to explain this, we first need an exercise, which should remind the reader of the hyperfinite generators of ultrafilters from Section 9.7.

Exercise 17.4.17. Fix a cardinal κ and let $\mathcal{L} := \{P_S : S \subseteq \kappa\}$, where each P_S is a unary relation symbol. Let \mathcal{M} be the \mathcal{L} -structure with universe κ such that $P_S^{\mathcal{M}} = S$ for all $S \subseteq \kappa$. Let \mathcal{N} be an elementary extension of \mathcal{M} in the sense of \mathcal{L}_κ . For any $b \in N$, let $\mathcal{U}_b := \{S \subseteq \kappa : b \in S^{\mathcal{N}}\}$. Prove that \mathcal{U}_b is a κ -complete ultrafilter on κ . Moreover, prove that \mathcal{U}_b is principal if and only if $b \in \kappa$.

Theorem 17.4.18. *Suppose that κ is an uncountable cardinal. Then the following are equivalent:*

- (1) κ is measurable.
- (2) For every increasing sequence $(\Sigma_\alpha)_{\alpha < \kappa}$ of \mathcal{L}_κ -sentences, if each Σ_α has a model, then $\bigcup_{\alpha < \kappa} \Sigma_\alpha$ has a model.
- (3) If \mathcal{M} is an \mathcal{L} -structure of cardinality κ , then \mathcal{M} has a proper \mathcal{L}_κ -elementary extension.

Proof. (1) \Rightarrow (2): Let \mathcal{U} be a nonprincipal κ -complete ultrafilter on κ . For $\alpha < \kappa$, let $\mathcal{M}_\alpha \models \Sigma_\alpha$. We claim $\prod_{\mathcal{U}} \mathcal{M}_\alpha \models \bigcup_{\alpha < \kappa} \Sigma_\alpha$. To see this, fix $\sigma \in \Sigma_\alpha$. Then $\mathcal{M}_\beta \models \sigma$ for all $\beta \geq \alpha$. By nonprincipality and κ -completeness, we have that $\{\beta < \kappa : \alpha \leq \beta\} \in \mathcal{U}$. It follows that $\mathcal{M}_\beta \models \sigma$ for \mathcal{U} -almost all β , whence, by Loś's theorem for \mathcal{L}_κ , we have that $\prod_{\mathcal{U}} \mathcal{M}_\alpha \models \sigma$, as desired.

(2) \Rightarrow (3): Let $(a_\beta)_{\beta < \kappa}$ be an enumeration of M . Let $T := \text{Th}_{\mathcal{L}_\kappa}(\mathcal{M}_M)$ denote the \mathcal{L}_κ -elementary diagram of \mathcal{M} . Let c be a new constant symbol. For any $\alpha < \kappa$, set $\Sigma_\alpha := T \cup \{c \neq a_\beta : \beta < \alpha\}$. For any $\alpha < \kappa$, note that $(\mathcal{M}_M, a_\alpha) \models \Sigma_\alpha$. Thus, by (2), $\bigcup_{\alpha < \kappa} \Sigma_\alpha$ has a model, whose reduct to \mathcal{L} is a proper \mathcal{L}_κ -elementary extension of \mathcal{M} .

(3) \Rightarrow (1) follows immediately from Exercise 17.4.17. □

We now consider one further compactness theorem for \mathcal{L}_κ , the weakest of them all (whence the name):

Definition 17.4.19. An uncountable cardinal κ is **weakly compact** if it is strongly inaccessible, and whenever Σ is a set of sentences of \mathcal{L}_κ with $|\Sigma| = \kappa$ such that every subset of size $< \kappa$ has a model, then Σ has a model.

Note that the difference between the compactness statements in the definitions of strongly compact and weakly compact cardinals is that weakly compact cardinals put a size limitation on the set of sentences we are trying to show has a model.

Remark 17.4.20. It is slightly unattractive to have to insert strong inaccessibility into the definition of weakly compact cardinals. However, if \mathcal{L}_κ satisfies the weak compactness theorem in the definition of weakly compact cardinals, then one can only show in ZFC that κ is weakly inaccessible (see [89, Exercises 17.17 and 17.18]). An unpublished result of Kunen (see Boos' article [17]) shows that if μ is a measurable cardinal, then there is a forcing extension of the universe in which μ becomes \mathfrak{c} and almost all (in the sense of a normal ultrafilter on \mathcal{U}) $\kappa < \mu$ are such that \mathcal{L}_κ satisfy the weak compactness theorem. Since no uncountable cardinal below \mathfrak{c} is a strong limit cardinal, we see that these κ show that it is consistent that satisfying the weak compactness theorem need not imply strong inaccessibility. To summarize: if you insist that all large cardinals be at least worldly, then one needs to insert that into the definition of the weakly compact cardinal. An even more compelling reason is that there are a number of other interesting reformulations of a weakly compact cardinal that only hold if one assumes that the cardinal is strongly inaccessible to begin with. See [89, Section 17] for more on this.

Theorem 17.4.18 above gives:

Corollary 17.4.21. *Measurable cardinals are weakly compact.*

Proof. Suppose that κ is a measurable cardinal. We already know that κ is strongly inaccessible. Now suppose that Σ is a set of sentences of \mathcal{L}_κ with $|\Sigma| = \kappa$ and such that every subset of size $< \kappa$ has a model. We need to show that Σ has a model. Enumerate Σ as $(\sigma_\beta)_{\beta < \kappa}$ and let $\Sigma_\alpha := \{\sigma_\beta : \beta < \alpha\}$. Then $\Sigma_\beta \subseteq \Sigma_\alpha$ for $\beta < \alpha$ and each Σ_α has a model as $|\Sigma_\alpha| = |\alpha| < \kappa$. Thus, by Theorem 17.4.18, Σ has a model. \square

Later, we will see that if κ is measurable, then for any κ -complete non-principal ultrafilter \mathcal{U} on κ , there are \mathcal{U} -many weakly compact cardinals below κ . In particular, the existence of a measurable cardinal has a higher consistency strength than the existence of a weakly compact cardinal.

One can also show that weakly compact cardinals are Mahlo. However, any such proof of this result often involves a detour through one of the many other reformulations of weak compactness.

17.5. Ramsey cardinals

We now seek to generalize the infinite version of Ramsey's theorem to uncountable cardinals. In order to do this, we first introduce the convenient (although slightly confusing) arrow notation. Suppose that X is a set and η is a cardinal with $|X| \geq \eta$. We extend our earlier notation by letting $X^{[\eta]}$ denote the subsets of X of cardinality η . If κ , λ , and μ are also cardinals, we write $\kappa \rightarrow (\lambda)_\mu^\eta$ to mean any coloring $f : \kappa^{[\eta]} \rightarrow \mu$ of $\kappa^{[\eta]}$ with μ colors has a homogeneous subset H of size λ , that is, f restricted to $H^{[\eta]}$ is constant. (Of course, for this to even be possible, one needs $\lambda, \eta \leq \kappa$.)

With this notation in hand, one can now state the infinite Ramsey theorem as for any $m, n < \aleph_0$, we have $\aleph_0 \rightarrow (\aleph_0)_m^n$.

When attempting to generalize the above statement to uncountable cardinals, we will not simply replace all occurrences of \aleph_0 by some uncountable cardinal κ and see what happens (this leads to something different). We instead first consider the property: for all $n < \aleph_0$ and all $\mu < \kappa$, $\kappa \rightarrow (\kappa)_\mu^n$. Unfortunately, this leads us to a large cardinal notion that we have already considered (see [28, Exercises 7.3.21 and 7.3.22]):

Theorem 17.5.1. *Suppose that κ is an uncountable cardinal. Then the following are equivalent:*

- (1) κ is weakly compact.
- (2) For every $n < \aleph_0$ and $\mu < \kappa$, $\kappa \rightarrow (\kappa)_\mu^n$.
- (3) $\kappa \rightarrow (\kappa)_2^2$.

However, there is a (seemingly) slight tweak that one can make that leads to a new large cardinal notion.

Definition 17.5.2. A cardinal κ is a **Ramsey cardinal** if, for any coloring c of $\mathcal{P}_f(\kappa)$ with $\mu < \kappa$ colors, there is $A \subseteq \kappa$ with $|A| = \kappa$ such that, for any $n \in \omega$, c restricted to $A^{[n]}$ is constant.

We extend the arrow notation to cover the conclusion of the definition: $\kappa \rightarrow (\kappa)_\mu^{<\omega}$.

Remark 17.5.3. It turns out that κ is Ramsey if and only if $\kappa \rightarrow (\kappa)_2^{<\omega}$. (See [89].)

By Theorem 17.5.1, Ramsey cardinals are weakly compact. One can show that Ramsey cardinals have a higher consistency strength than weakly compact cardinals.

It is worth pointing out that, unlike the other large cardinal notions, the defining property does not hold for \aleph_0 :

Proposition 17.5.4. $\aleph_0 \not\rightarrow (\aleph_0)_2^{<\omega}$.

Proof. Define a coloring c on the finite subsets of ω by $c(s) = 0$ if $|s| \in s$ and $c(s) = 1$ otherwise. We show that there is no infinite homogeneous subset. Let $A \subseteq \omega$ be infinite with $n := \min(A)$. Enumerate $A = \{a_0, a_1, \dots\}$ in increasing order. Then $c(\{a_0, a_1, \dots, a_{n-1}\}) = 0$ and $c(\{a_1, \dots, a_n\}) = 1$. \square

Theorem 17.5.5. *Measurable cardinals are Ramsey cardinals.*

Proof. Suppose that κ is a measurable cardinal and fix a normal ultrafilter \mathcal{U} on κ .

Claim. *For any $n \in \omega$ and any coloring $c : \kappa^{[n]} \rightarrow \mu$ with $\mu < \kappa$, there is a set $H \in \mathcal{U}$ that is homogeneous for c .*

Proof of Claim. We prove the claim by induction on n . The case $n = 1$ follows immediately from κ -completeness. We now suppose it is true for n and fix a coloring $c : \kappa^{[n+1]} \rightarrow \mu$ with $\mu < \kappa$. Each $s \in \kappa^{[n]}$ induces a coloring c_s on $\kappa \setminus s$ by defining $c_s(\alpha) := c(\alpha \cup \{s\})$. By the base case of the induction, we may fix $H_s \in \mathcal{U}$ monochromatic for c_s with color q_s . The function $s \mapsto q_s$ is a coloring of $\kappa^{[n]}$, so by induction there is a homogeneous set $\bar{H} \in \mathcal{U}$ for that coloring with color q .

We set $H := \{\alpha \in \bar{H} : \alpha \in H_s \text{ for all } s \in (\alpha \cap \bar{H})^{[n]}\}$. Note first that c is homogeneous on H with color q . Indeed, suppose that $t \in H^{[n+1]}$ and set $\alpha := \max(t)$ and $s := t \setminus \{\alpha\}$. Then $s \in (\alpha \cap \bar{H})^{[n]}$, whence $\alpha \in H_s$, and hence $c_s(\alpha) = q$, that is, $c(t) = q$. It remains to show that $H \in \mathcal{U}$. To see this, suppose that $\alpha \in \bar{H} \setminus H$. Then there is $s \in (\alpha \cap \bar{H})^{[n]}$ such that $\alpha \notin H_s$. Pick such an s and enumerate $s = \{f_1(\alpha), \dots, f_n(\alpha)\}$. For $\alpha \notin \bar{H} \setminus H$, define $f_i(\alpha)$ arbitrarily. Suppose, toward a contradiction, that $H \notin \mathcal{U}$. Then $\bar{H} \setminus H \in \mathcal{U}$ and $f_i(\alpha) < \alpha$ for all $\alpha \in \bar{H} \setminus H$ and $i = 1, \dots, n$. In other words, each f_i is \mathcal{U} -regressive. By normality, there are $X_1, \dots, X_n \in \mathcal{U}$ such that $f_i \upharpoonright X_i$ is constant. Setting $X := X_1 \cap \dots \cap X_n$, we see that there is a single s such that, for all $\alpha \in X$, we have $s = \{f_1(\alpha), \dots, f_n(\alpha)\}$. It remains to note that $X \cap H_s = \emptyset$, a contradiction. This proves the claim.

Now fix a coloring $c : \mathcal{P}_f(\kappa) \rightarrow \mu$ with $\mu < \kappa$ and let c_n denote its restriction to $\kappa^{[n]}$. By the claim, for each $n \in \omega$, there is a set $H_n \in \mathcal{U}$ which

is homogeneous for c_n . Set $H := \bigcap_{n \in \omega} H_n$. By countable completeness, $H \in \mathcal{U}$. Since \mathcal{U} is nonprincipal and κ -complete, $|H| = \kappa$, and is thus as desired. \square

We will see an application of Ramsey cardinals in Section 17.7.

We have now introduced all of the large cardinals that we plan on discussing in this book. (There are many, many more interesting large cardinal notions, but we have focused on those connected to the themes already discussed in this book. See [89] and [94] for more information on large cardinals.) To summarize, we have

$$\begin{aligned} \text{supercompact} &\Rightarrow \text{strongly compact} \Rightarrow \text{measurable} \Rightarrow \text{Ramsey} \\ &\Rightarrow \text{weakly compact} \\ &\Rightarrow \text{Mahlo} \Rightarrow \text{2-inaccessible} \\ &\Rightarrow \text{strongly inaccessible} \Rightarrow \text{worldly}. \end{aligned}$$

Moreover, all of these implications represent a strict increase in consistency strength except for perhaps the first implication (which is still an open question).

17.6. Measurable cardinals as critical points of elementary embeddings

In this section, we discuss a more modern viewpoint of large cardinals as so-called critical points of elementary embeddings of the set-theoretic universe.

First, suppose that \mathcal{U} is an ultrafilter on an index set I . We would like to consider the ultrapower $V^{\mathcal{U}}$ of the set-theoretic universe V . Unfortunately, the naïve idea of considering equivalence classes of functions $f : I \rightarrow V$ under equality modulo \mathcal{U} does not work as then equivalence classes would be proper classes. However, an idea of Scott is to consider those elements of the equivalence class of minimal rank. More precisely, given $f : I \rightarrow V$, we define the **restricted equivalence class of f with respect to $\sim_{\mathcal{U}}$** to be

$$[f]_{\mathcal{U}}^r := \{g \in V^I : f \sim_{\mathcal{U}} g \text{ and for any } h \in V^I \text{ with } f \sim_{\mathcal{U}} h \\ \text{we have } \text{rank}(g) \leq \text{rank}(h)\}.$$

Note now that $[f]_{\mathcal{U}}^r$ is a set. We define the **ultrapower of V with respect to \mathcal{U}** to be $V^{\mathcal{U}} := \{[f]_{\mathcal{U}}^r : f : I \rightarrow V\}$. We consider the relation E on $V^{\mathcal{U}}$ given by membership: $[f]_{\mathcal{U}}^r E [g]_{\mathcal{U}}^r$ if and only if $f(i) \in g(i)$ for \mathcal{U} -almost all $i \in I$.

Exercise 17.6.1. Verify that Łoś's theorem holds in the following context: if $\varphi(x_1, \dots, x_n)$ is a formula in the language of set theory and $[f_1]_{\mathcal{U}}^r, \dots, [f_n]_{\mathcal{U}}^r$

are elements of $V^{\mathcal{U}}$, show that

$$(V^{\mathcal{U}}, E) \models \varphi([f_1]_{\mathcal{U}}^r, \dots, [f_n]_{\mathcal{U}}^r) \Leftrightarrow \{i \in I : (V, \epsilon) \models \varphi(f_1(i), \dots, f_n(i))\} \in \mathcal{U}.$$

Conclude that the diagonal embedding $d : (V, \epsilon) \rightarrow (V^{\mathcal{U}}, E)$ is an elementary embedding.

The *Mostowski collapse* procedure allows one to find a unique inner model $M \subseteq V$ with $(V^{\mathcal{U}}, E) \cong (M, \epsilon)$ if:

- (1) for each $[f]_{\mathcal{U}}^r \in V^{\mathcal{U}}$, the collection of all $[g]_{\mathcal{U}}^r$ with $[g]_{\mathcal{U}}^r E [f]_{\mathcal{U}}^r$ is a set; and
- (2) E is well founded.

Item (1) holds for any \mathcal{U} by the way we defined $V^{\mathcal{U}}$. In connection with item (2), we have:

Exercise 17.6.2. Prove that E is well founded if and only if \mathcal{U} is countably complete.

Based on the previous exercise, we assume from now on that \mathcal{U} is at least countably complete, whence $(V^{\mathcal{U}}, E) \cong (M, \epsilon)$ for a unique inner model $M \subseteq V$. For simplicity, we often identify $V^{\mathcal{U}}$ with M .

Before we go any further, we gather some basic facts about elementary embeddings:

Exercise 17.6.3. Suppose that M is an inner model, $j : V \rightarrow M$ is an elementary embedding, and α is an ordinal. Prove that:

- (1) $j(\alpha)$ is an ordinal.
- (2) $\alpha \leq j(\alpha)$.
- (3) If $j(\beta) = \beta$ for all $\beta \in \alpha$, then $j(x) = x$ for all $x \in V_\alpha$.

It is worth emphasizing that in the previous exercise, one only needs to assume that M is a transitive class and that there is an elementary embedding $j : V \rightarrow M$, for then it follows that M is an inner model of ZFC.

In the rest of this section, M denotes an inner model of ZFC.

Lemma 17.6.4. *Suppose that \mathcal{U} is a nonprincipal κ -complete ultrafilter on κ . Then $d(\gamma) = \gamma$ for all $\gamma < \kappa$.*

Proof. Suppose that $[f]_{\mathcal{U}}^r < d(\gamma)$. Then $f(\alpha) < \gamma$ for \mathcal{U} -almost all $\alpha \in \kappa$. By κ -completeness, there is $\eta < \gamma$ such that $f(\alpha) = \eta$ for \mathcal{U} -almost all $\alpha \in \kappa$, whence $[f]_{\mathcal{U}}^r = d(\eta)$. \square

Proposition 17.6.5. *Suppose that $j : V \rightarrow M$ is a nontrivial elementary embedding. Then the least ordinal α such that $\alpha < j(\alpha)$ is a cardinal.*

Proof. Let α be the least ordinal such that $\alpha < j(\alpha)$. Fix $\beta < \alpha$ and $f : \beta \rightarrow \alpha$; it suffices to show that f is not a surjection. First note that, for $\xi < \beta$, we have that $j(f)(\xi) = j(f)(f(\xi)) = j(f(\xi)) = f(\xi)$. It remains to note that if f were surjective, then by elementarity, $j(f)$ would also be surjective; since the codomain of $j(f)$ is $f(\alpha) > \alpha$, this contradicts the fact that $j(f)(\xi) = f(\xi)$ for all $\xi < \beta$. \square

Definition 17.6.6. If $j : V \rightarrow M$ is a nontrivial elementary embedding, then the **critical point of j** , denoted $\text{crit}(j)$, is the least ordinal κ such that $\kappa < j(\kappa)$.

Theorem 17.6.7. *Suppose that κ is a measurable cardinal and \mathcal{U} is a nonprincipal κ -complete ultrafilter on κ . If $d : V \rightarrow V^{\mathcal{U}}$ is the corresponding diagonal embedding, then d is a nontrivial elementary embedding with $\text{crit}(d) = \kappa$.*

Proof. By Lemma 17.6.4, we already know that $\text{crit}(d) \geq \kappa$. Let $\text{id} : \kappa \rightarrow \kappa$ be the identity map. Suppose $\gamma < \kappa$. Then $\{\alpha < \kappa : \gamma < \alpha\} \in \mathcal{U}$ by κ -completeness, whence $\gamma = d(\gamma) < [\text{id}]_{\mathcal{U}}^r$ and thus $\kappa \leq [\text{id}]_{\mathcal{U}}^r$. On the other hand, $[\text{id}]_{\mathcal{U}}^r < d(\kappa)$. It follows that $\kappa < d(\kappa)$ and thus $\text{crit}(d) = \kappa$. \square

Exercise 17.6.8. Suppose that \mathcal{U} is a nonprincipal κ -complete ultrafilter on the uncountable cardinal κ . Then the following are equivalent:

- (1) \mathcal{U} is normal.
- (2) In $V^{\mathcal{U}}$, $\kappa = [\text{id}]_{\mathcal{U}}^r$.
- (3) For every $X \subseteq \kappa$, $X \in \mathcal{U}$ if and only if $\kappa \in d(X)$.

We now prove the converse of Theorem 17.6.7. The previous exercise gives us a hint of how to proceed:

Theorem 17.6.9. *Suppose that $j : V \rightarrow M$ is a nontrivial elementary embedding with $\text{crit}(j) = \kappa$. Then κ is measurable. In fact, $\mathcal{U}_j := \{X \subseteq \kappa : \kappa \in j(X)\}$ is a normal ultrafilter on κ .*

Proof. Since $j(\omega) = \omega$, we have that κ is uncountable. We now verify the latter claim:

- Since $\kappa < j(\kappa)$, we have that $\kappa \in j(\kappa)$, so $\kappa \in \mathcal{U}_j$.
- Since $j(\emptyset) = \emptyset$, $\emptyset \notin \mathcal{U}_j$.
- Suppose that $X \in \mathcal{U}_j$ and $X \subseteq Y$. Since $j(X) \subseteq j(Y)$, we have that $\kappa \in j(Y)$, so $Y \in \mathcal{U}_j$.
- Suppose that $X, Y \in \mathcal{U}_j$. Since $\kappa \in j(X) \cap j(Y) = j(X \cap Y)$, we have that $X \cap Y \in \mathcal{U}_j$.

- Suppose that $X \subseteq \kappa$. If $X \notin \mathcal{U}_j$, then $\kappa \in j(\kappa) \setminus j(X) = j(\kappa \setminus X)$, so $\kappa \setminus X \in \mathcal{U}_j$.
- Suppose $\alpha < \kappa$. Then $j(\{\alpha\}) = \{j(\alpha)\} = \{\alpha\}$ by Lemma 17.6.3, whence $\kappa \notin j(\{\alpha\})$ and so $\mathcal{U}_j \neq \mathcal{U}_\alpha$ and thus \mathcal{U}_j is not principal.
- κ -complete: Suppose that $\gamma < \kappa$ and $\mathcal{X} = (X_\alpha)_{\alpha < \gamma}$ is a collection from \mathcal{U}_j . By elementarity, in M , $j(\mathcal{X})$ is a sequence of length $j(\gamma)$ of subsets of $j(\kappa)$; these facts remain true in V . By Lemma 17.6.3, $j(\mathcal{X}) = (j(X_\alpha) : \alpha < \gamma)$. Then $\kappa \in \bigcap_{\alpha < \gamma} j(X_\alpha) = j(\bigcap_{\alpha < \gamma} X_\alpha)$, and hence $\bigcap_{\alpha < \gamma} X_\alpha \in \mathcal{U}_j$.
- Suppose $f : \kappa \rightarrow \kappa$ is \mathcal{U}_j -regressive. Let $X := \{\gamma < \kappa : f(\gamma) < \gamma\}$. Then $X \in \mathcal{U}_j$, so $\kappa \in j(X) = \{\gamma < j(\kappa) : j(f)(\gamma) < \gamma\}$, so $j(f)(\kappa) < \kappa$, say $j(f)(\kappa) = \alpha$. Let $Y = \{\gamma < \kappa : f(\gamma) = \alpha\}$. Then $\kappa \in j(Y)$, so $Y \in \mathcal{U}_j$ and so f is \mathcal{U}_j -constant. \square

Using the notation of the previous theorem, setting $d : V \rightarrow V^{\mathcal{U}}$ to be the diagonal embedding, Exercise 17.6.8 states that $\mathcal{U}_d = \mathcal{U}$ if and only if \mathcal{U} is normal. In general, we have the following result:

Theorem 17.6.10. *Suppose that $j : V \rightarrow M$ is a nontrivial elementary embedding with $\text{crit}(j) = \kappa$. Define \mathcal{U}_j as in the previous theorem. Then there is an elementary embedding $j' : V^{\mathcal{U}_j} \rightarrow M$ such that $j' \circ d_{\mathcal{U}_j} = j$.*

Proof. Given $[f]_{\mathcal{U}_j}^r \in V^{\mathcal{U}_j}$, set $j'([f]_{\mathcal{U}_j}^r) := j(f)(\kappa)$. (Note that this makes sense as $f : \kappa \rightarrow V$ so $j(f) : j(\kappa) \rightarrow M$ and $\kappa < j(\kappa)$.) Of course, one must check that this definition is independent of representatives. Toward this end, suppose that $[f]_{\mathcal{U}_j}^r = [g]_{\mathcal{U}_j}^r$. It follows that $X := \{\alpha < \kappa : f(\alpha) = g(\alpha)\} \in \mathcal{U}_j$, whence $\kappa \in j(X) = \{\alpha < j(\kappa) : j(f)(\alpha) = j(g)(\alpha)\}$, as desired.

Why is j' elementary? Suppose $V^{\mathcal{U}_j} \models \varphi([f]_{\mathcal{U}_j}^r)$. Then $X = \{\alpha < \kappa : \varphi(f(\alpha))\} \in \mathcal{U}_j$, so $\kappa \in j(X) = \{\alpha < j(\kappa) : M \models \varphi(j(f)(\alpha))\}$. Thus $M \models \varphi(j(f)(\kappa))$, i.e. $M \models \varphi(j'([f]_{\mathcal{U}_j}^r))$, as desired.

Finally, we show that $j' \circ d_{\mathcal{U}_j} = j$. Fix $a \in V$. Then $d_{\mathcal{U}_j}(a) = [c_a]_{\mathcal{U}_j}^r$, where $c_a : \kappa \rightarrow V$ is the function constantly equal to a . Then $j(c_a)$ is the function on $j(\kappa)$ constantly equal to $j(a)$. In particular, $j(c_a)(\kappa) = j(a)$. By the definition of j' , we have $j'(d_{\mathcal{U}_j}(a)) = j(a)$, as desired. \square

Before moving on, we give a nice application of this perspective:

Theorem 17.6.11 (Scott [154]). *If $V = L$, then there are no measurable cardinals. Consequently, measurable cardinals do not exist in L .*

Proof. Suppose that $V = L$ and yet, toward a contradiction, that there is a measurable cardinal. Let κ be the least measurable cardinal and let $d : V \rightarrow V^{\mathcal{U}}$ be the diagonal embedding corresponding to some nonprincipal

κ -complete ultrafilter on κ . Since $V = L$, we have that $V^{\mathcal{U}} = V$. Since d is elementary, we have that $d(\kappa)$ is the least measurable cardinal in $V^{\mathcal{U}} = V$. Since $d(\kappa) > \kappa$, we have reached a contradiction. \square

Large cardinal notions stronger than measurable cardinals can often be rephrased in terms of the existence of elementary embeddings $j : V \rightarrow M$ with M 's that “resemble” V more closely. One way in which we can measure this is to ask what subsets of M actually belong to M .

Definition 17.6.12. For a cardinal λ , we say that an elementary embedding $j : V \rightarrow M$ is λ -**supercompact** if $M^\lambda \subseteq M$.

We now give a criteria for λ -supercompactness. Recall that for a function $f : X \rightarrow Y$ and a subset $Z \subseteq X$, we set $f''Z := \{f(x) : x \in Z\}$ to be the image of Z under f .

Proposition 17.6.13. *Suppose that \mathcal{U} is a countably complete ultrafilter on a set S , set $M := V^{\mathcal{U}}$, and let $d : V \rightarrow M$ be the diagonal embedding. Then for any cardinal λ , d is λ -supercompact if and only if $d''\lambda \in M$.*

Proof. The forward direction is clear. Now suppose that $d''\lambda \in M$ and suppose that $Y \subseteq M$ is such that $|Y| \leq \lambda$. Write $Y = \{[f_\alpha]_{\mathcal{U}}^r : \alpha < \lambda\}$. Let $h : S \rightarrow \mathcal{P}(\lambda)$ be such that $[h]_{\mathcal{U}}^r = d''\lambda$. Define $g : S \rightarrow V$ by declaring $g(i)$ is a function with domain $h(i)$ and such that $g(i)(\alpha) = f_\alpha(i)$. It follows that $[g]_{\mathcal{U}}^r(d(\alpha)) = [f_\alpha]_{\mathcal{U}}^r$ for every $\alpha < \lambda$ and the domain of $[g]_{\mathcal{U}}^r = d''\lambda$, whence the range of $[g]_{\mathcal{U}}^r$ is Y and thus $Y \in M$. \square

The next result follows immediately from Proposition 17.6.13:

Corollary 17.6.14. *If κ is a measurable cardinal and \mathcal{U} is a nonprincipal κ -complete ultrafilter on κ , then the diagonal embedding d is κ -supercompact.*

Before moving on, here is an application:

Theorem 17.6.15. *Suppose that \mathcal{U} is a normal ultrafilter on κ . Then for \mathcal{U} -almost all $\lambda < \kappa$, λ is weakly compact.*

Proof. By Corollary 17.6.14 and Theorem 17.5.1, we see that κ remains weakly compact in $V^{\mathcal{U}}$. The result now follows Exercise 17.6.8. \square

There is a limit to how supercompact the diagonal embedding can be. In the context of Proposition 17.6.13, we have:

Proposition 17.6.16. *If d is λ -supercompact, then $\lambda \leq |S|$.*

Proof. We show that $d''(|S|^+) \notin M$. Fix $[f]_{\mathcal{U}}^r \in M$; we show that $d''(|S|^+) \neq [f]_{\mathcal{U}}^r$. Set $A := \{i \in S : |f(i)| \leq |S|\}$.

Case 1. $A \in \mathcal{U}$. Since $|\bigcup_{i \in A} f(i)| \leq |S|$, we may take $\alpha \in |S|^+ \setminus \bigcup_{i \in A} f(i)$. We claim that $d(\alpha) \notin [f]_{\mathcal{U}}^r$; since $d(\alpha) \in d''(|S|^+)$, this suffices to finish the proof in this case. Suppose, toward a contradiction, that $d(\alpha) \in [f]_{\mathcal{U}}^r$. Then $\alpha \in f(i)$ for \mathcal{U} -almost all i . In particular, there is $i \in A$ such that $\alpha \in f(i)$, a contradiction to the choice of α .

Case 2. $S \setminus A \in \mathcal{U}$. In this case, since $|S| < |f(i)|$ for all $i \in S \setminus A$, we may find an injective function $h : S \rightarrow V$ such that $h(i) \in f(i)$ for all $i \in S \setminus A$. It follows that $[h]_{\mathcal{U}}^r \in [f]_{\mathcal{U}}^r$. If $[h]_{\mathcal{U}}^r \in d''(|S|^+)$, then there would be $\alpha < |S|^+$ such that $h(i) = \alpha$ for \mathcal{U} -almost all i . We would then have that h is both injective and constant on a \mathcal{U} -large set, a contradiction. \square

In order to get stronger large cardinal notions, we are thus searching for cardinals κ that admit λ -supercompact elementary embeddings with $\lambda \geq \kappa$. We have to make a choice between the relationship between λ and $j(\kappa)$. If we choose that $j(\kappa) \leq \lambda$, then the resulting cardinals will be **huge**, which is a different story. We thus focus on $\lambda < j(\kappa)$.

Definition 17.6.17. We say that κ is λ -**supercompact** if there is a λ -supercompact embedding $j : V \rightarrow M$ such that $\text{crit}(j) = \kappa$ and $j(\kappa) > \lambda$.

Since $d''\kappa = \kappa$ for a nonprincipal κ -complete ultrafilter on κ , we are motivated to define an ultrafilter \mathcal{U}'_d by setting $X \in \mathcal{U}'_d$ if and only if $d''\lambda \in d(X)$. But what is the ultrafilter defined on? Notice that $M \models |d''\lambda| < d(\kappa)$, so $d''\lambda \in \mathcal{P}_{d(\kappa)}^M(d(\lambda))$, whence $\mathcal{P}_\kappa(\lambda) \in \mathcal{U}'_d$. This hints that maybe we should have that be the index set for \mathcal{U}'_d .

To summarize: for $\lambda \geq \kappa$, we set \mathcal{U}'_j to be the ultrafilter on $\mathcal{P}_\kappa(\lambda)$ given by $X \in \mathcal{U}'_j$ if and only if $j''\lambda \in j(X)$.

Theorem 17.6.18. *If $\lambda \geq \kappa$ and κ is λ -supercompact as witnessed by the embedding $j : V \rightarrow M$, then \mathcal{U}'_j is a normal ultrafilter on $\mathcal{P}_\kappa(\lambda)$.*

Exercise 17.6.19. Prove the previous theorem.

The previous theorem is actually one half of the following:

Theorem 17.6.20. *There is a normal ultrafilter on $\mathcal{P}_\kappa(\lambda)$ if and only if κ is λ -supercompact.*

Proof. Suppose that \mathcal{U} is a normal ultrafilter on $\mathcal{P}_\kappa(\lambda)$. Note that it follows that κ is uncountable (Exercise). Let $d : V \rightarrow V^{\mathcal{U}}$ be the corresponding diagonal embedding. We show that d is a λ -supercompact embedding with $\text{crit}(d) = \kappa$ and $d(\kappa) > \lambda$.

Let $\text{id} : \mathcal{P}_\kappa(\lambda) \rightarrow \mathcal{P}_\kappa(\lambda)$ be the identity map. By Proposition 17.6.13, in order to show that d is λ -supercompact, it suffices to show that $d''\lambda = [\text{id}]_{\mathcal{U}}^r$. First suppose that $\alpha < \lambda$. Then $\alpha \in x$ for \mathcal{U} -almost all x by fineness, whence

$d(\alpha) \in [\text{id}]_{\mathcal{U}}^r$. It follows that $d''\lambda \subseteq [\text{id}]_{\mathcal{U}}^r$. For the other direction, suppose that $[f]_{\mathcal{U}}^r \in [\text{id}]_{\mathcal{U}}^r$. Then $f(x) \in x$ for \mathcal{U} -almost all x . By normality, f is constant on a set in \mathcal{U} , whence $[f]_{\mathcal{U}}^r = d(\alpha)$ for some $\alpha < \kappa$.

Arguing as in Lemma 17.6.4, it can be seen that $\text{crit}(d) \geq \kappa$. It remains to show that $d(\kappa) > \lambda$ (which also establishes that $\text{crit}(d) = \kappa$). Indeed, this follows from the fact that the order-type of $[\text{id}]_{\mathcal{U}}^r$ is less than $d(\kappa)$ in $V^{\mathcal{U}}$ (as the order-type of x is less than κ for \mathcal{U} -almost all x) together with the fact that the order-type of $[\text{id}]_{\mathcal{U}}^r$ is the order-type of $d''\lambda$ (by the previous paragraph), which in turn is λ . \square

Corollary 17.6.21. *κ is supercompact if and only if κ is λ -supercompact for all $\lambda \geq \kappa$.*

Exercise 17.6.22. Suppose that κ is supercompact. Prove that there is a normal ultrafilter \mathcal{U} on κ such that α is measurable for \mathcal{U} -almost all $\alpha < \kappa$.

While not quite as easy to state as Corollary 17.6.21, there is also an elementary embedding characterization of strongly compact cardinals. The reader may consult [89] for a proof.

Theorem 17.6.23. *κ is strongly compact if and only if, for all $\lambda \geq \kappa$, there is an elementary embedding $j : V \rightarrow M$ with $\text{crit}(j) = \kappa$ satisfying for any $X \subseteq M$ with $|X| \leq \lambda$, there is $Y \in M$ such that $X \subseteq Y$ and $M \models |Y| < j(\kappa)$.*

17.7. An application of large cardinals

In this chapter, we have spent a lot of time talking about various kinds of large cardinals. But what if you do not care about large cardinals? What if you only care about, say, the real numbers? Well, it turns out that large cardinals can help you prove things about the real numbers as well.

Recall in Section 5.2 we mentioned that all analytic sets of real numbers are both Lebesgue and Baire measurable, but that one cannot prove, in ZFC, the same fact for the next level of the projective hierarchy. However, you can prove this result using large cardinals:

Theorem 17.7.1 (Solovay [167]). *If there is a measurable cardinal, then all Σ_2^1 subsets of \mathbb{R} are Lebesgue and Baire measurable.*

Solovay's proof uses a lot of set theory that we are not assuming in this book. We instead give Martin's proof of Solovay's theorem that proceeds via determinacy. The proof given here also reduces the large cardinal assumption:

Theorem 17.7.2 (Martin [128]). *Suppose that there exists a Ramsey cardinal. Then analytic determinacy holds.*

Recall from Section 5.3 that analytic determinacy implies that all Σ_2^1 subsets of \mathbb{R} are Lebesgue and Baire measurable, whence we really do obtain Solovay's result as a corollary.

We devote the rest of this section to proving Theorem 17.7.2. We need three bits of preparation: a lemma on Ramsey cardinals, a theorem on determinacy of open games, and a discussion on tree-representations of analytic sets.

We start with the lemma on Ramsey cardinals.

Lemma 17.7.3. *Suppose that κ is a Ramsey cardinal and $f_i : \kappa^{[m_i]} \rightarrow \mathbb{N}$ is a coloring for each $i \in \mathbb{N}$. Then there is $X \subseteq \kappa$ of cardinality κ that is homogeneous for each f_i .*

Proof. Let $f : \mathcal{P}_f(\kappa) \rightarrow \mathbb{N}^{\mathbb{N}}$ be given by $f(X)(i) := f_i(X')$, where X' is obtained from X by either adding elements to X or deleting elements from X so that $|X'| = m_i$. Since κ is strongly inaccessible, f is a coloring with “few” colors, so since κ is a Ramsey cardinal, there is $A \subseteq \kappa$ with $|A| = \kappa$ that is homogeneous for f . It follows that A is homogeneous for each f_i . \square

Remark 17.7.4. In our proof of Martin's theorem, we will really only need there to be a homogeneous subset of size \aleph_1 , whence the large cardinal hypothesis really needed in the proof is that the **Erdős cardinal** η_{ω_1} exists.

We now prove that all open games are determined:

Theorem 17.7.5 (Gale-Stewart). *For any set X , if $A \subseteq X^{\mathbb{N}}$ is open, then $\mathcal{G}(A)$ is determined.*

Proof. Suppose that player I does not have a winning strategy. We show that player II does have a winning strategy. The strategy can be summarized as “play not to lose”. Suppose that player I opens with $a_0 \in X$. Since player I does not have a winning strategy, this means that there must be some $a_1 \in X$ that does not lead to an automatic loss for player II. Suppose that player I responds with $a_2 \in X$. Again, since player I does not have a winning strategy, there must be some $a_3 \in X$ such that, if player II responds with a_3 , then they have not automatically lost the game. Player II continues to play in this fashion.

Why is this strategy winning for player II? Suppose, toward a contradiction, that it is not a winning strategy, whence there is some play $(a_0, a_1, \dots) \in X^{\mathbb{N}}$ played according to this strategy such that player I wins the game, that is, $(a_0, a_1, \dots) \in A$. Since A is open, there is some $k \in \mathbb{N}$ with k even such that all infinite extensions of this finite sequence belong to A . This means that the position (a_0, \dots, a_k) is winning for player I, contradicting the fact that they do not have a winning strategy. \square

Our final bit of preparation is a short discussion on trees on \mathbb{N} and the connection with analytic sets. Given a set X , a **tree on X** is a set $T \subseteq X^{<\mathbb{N}}$ such that if $u \in T$ and v is an initial segment of u , then $v \in T$. In particular, $\emptyset \in T$ (if $T \neq \emptyset$) and is called the **root** of T . An **infinite branch in T** is a function $f \in X^{\mathbb{N}}$ such that $f \upharpoonright n = (f(0), \dots, f(n-1)) \in T$ for all $n \in \mathbb{N}$. We let $[T]$ denote the set of branches through T . T is said to be **well founded** if $[T] = \emptyset$ and **ill founded** otherwise.

In connection with analytic sets, we will be concerned with trees on sets of the form $X \times Y$. If T is a tree on $X \times Y$, then technically $[T]$ is a subset of $(X \times Y)^{\mathbb{N}}$. But $(X \times Y)^{\mathbb{N}} \cong X^{\mathbb{N}} \times Y^{\mathbb{N}}$, whence we often write infinite branches of T as (f, g) , with $f \in X^{\mathbb{N}}$ and $g \in Y^{\mathbb{N}}$. Under this identification, we have that $(f(0), g(0), \dots, f(n-1), g(n-1)) \in T$ for all $n \in \mathbb{N}$. For $x \in X^{\mathbb{N}}$, we write

$$T(x) = \{(u_0, \dots, u_{n-1}) \in Y^{<\mathbb{N}} : (x(0), u_0, \dots, x(n-1), u_{n-1}) \in T\}.$$

It is easy to check that $T(x)$ is a tree on Y .

Fact 17.7.6 (Tree representation of analytic sets). Given $A \subseteq \mathbb{N}^{\mathbb{N}}$, we have that A is analytic if and only if there is a tree T on $\mathbb{N} \times \mathbb{N}$ such that $x \in A$ if and only if $T(x)$ is ill founded.

Given $u, v \in \mathbb{N}^{<\mathbb{N}}$, we set $u \leq_{KB} v$ if and only if:

- $u \supseteq v$, or
- u and v are incompatible and $u(n) < v(n)$ where n is the first place that they differ.

It is clear that \leq_{KB} is a linear ordering on $\mathbb{N}^{<\mathbb{N}}$, called the **Kleene-Brouwer ordering** on $\mathbb{N}^{<\mathbb{N}}$.

Exercise 17.7.7. Given a tree T on \mathbb{N} , show that the following are equivalent:

- (1) T is well founded;
- (2) (T, \leq_{KB}) is a well-ordering;
- (3) $(T, <_{KB})$ embeds (as a linear ordering) into $(\omega_1, <)$.

Finally, we fix an enumeration $(u_n)_{n \in \mathbb{N}}$ of $\mathbb{N}^{<\mathbb{N}}$ such that $|u_n| \leq n$ for each $n \in \mathbb{N}$.

We are now ready to prove Theorem 17.7.2.

Proof of Theorem 17.7.2. Let $A \subseteq \mathbb{N}^{\mathbb{N}}$ be co-analytic and let $T \subseteq (\mathbb{N} \times \mathbb{N})^{<\mathbb{N}}$ be a tree such that $x \in A$ if and only if $T(x)$ is well founded.

In order to show that $\mathcal{G}(A)$ is determined, we introduce a new game, $\mathcal{G}'(A)$, which we know is determined, and for which a strategy for each

player in $\mathcal{G}'(A)$ can be used to yield a strategy for the corresponding player in $\mathcal{G}(A)$.

In the game $\mathcal{G}'(A)$, player II continues to play elements of \mathbb{N} while player I plays elements of $\mathbb{N} \times \kappa$. So a play of the game looks like

$$(x_0, \eta_0, x_1, x_2, \eta_1, x_3, \dots),$$

where each $x_i \in \mathbb{N}$ and each $\eta_i \in \kappa$. Let x denote the sequence $(x_n)_{n \in \mathbb{N}}$. We say that player I wins $\mathcal{G}'(A)$ if and only if the following two conditions hold:

- If $u_i \notin T(x)$, then $\eta_i = 0$.
- If $u_i, u_j \in T(x)$, then $u_i <_{KB} u_j$ if and only if $\eta_i < \eta_j$.

Note that if player I wins $\mathcal{G}'(A)$, then $<_{KB}$ is well founded on $T(x)$ (otherwise we would obtain an infinite descending chain of ordinals below κ), and thus $x \in A$. Thus, if player I has a winning strategy for winning $\mathcal{G}'(A)$, then player I also has a winning strategy for winning $\mathcal{G}(A)$ (just follow the strategy for $\mathcal{G}'(A)$ but do not actually play the elements of κ).

Claim. $\mathcal{G}'(A)$ is determined.

Proof of Claim. Define $\mathcal{G}''(A)$ to be just like $\mathcal{G}'(A)$ except that player II also has to play elements of κ . Note then that $\mathcal{G}''(A)$ is an open game on $\mathbb{N} \times \kappa$, whence, by the Gale-Stewart theorem, $\mathcal{G}''(A)$ is determined. It is straightforward to see that it follows that $\mathcal{G}'(A)$ is also determined. \square

By the claim and the paragraph preceding it, we are left to show if player II has a winning strategy for $\mathcal{G}'(A)$, then player II also has a winning strategy for $\mathcal{G}(A)$. Suppose that $s \in \mathbb{N}^{2n+1}$, so $s = (s_0, s_1, \dots, s_{2n})$, which we think of as the first $2n+1$ moves of the game $\mathcal{G}(A)$. How should player II respond? In order to take advantage of the fact that player II has a winning strategy in $\mathcal{G}'(A)$, we should somehow simulate a play of the game $\mathcal{G}'(A)$ where the s_{2k} 's are the first coordinates of player I's moves and the s_{2k+1} 's are player II's moves. Let us assume that player I is playing this simulated version of $\mathcal{G}'(A)$ to win. More precisely, let D_s consist of those $i < n$ for which $(s \upharpoonright k, u_i) \in T$ for some k (namely $k = |u_i|$). Set $m_s := |D_s|$; note that $m_s \leq n$. We then want to consider a partial play of the game $\mathcal{G}'(A)$, namely $(s_0, \eta_0, s_1, s_2, \eta_1, s_3, \dots, s_{2n}, \eta_n)$, where:

- If $i \notin D_s$, then $\eta_i = 0$.
- If $i, j \in D_s$, then $u_i <_{KB} u_j$ if and only if $\eta_i < \eta_j$.

Note that the η_i 's in the second clause can be any m_s elements of κ , as long as they are ordered appropriately.

Thus, given any $Q \in \kappa^{[m_s]}$, we consider the play s_Q of the game $\mathcal{G}'(A)$ which is as above and with $Q = \{\eta_i : i \in D_s\}$. We can then define $f_s(Q)$ to be what the winning strategy for player II tells us the next move for player II should be if s_Q has been played thus far.

Since κ is a Ramsey cardinal, by Lemma 17.7.3, there is $H \subseteq \kappa$ with $|H| = \kappa$ such that H is homogeneous for each f_s with $s \in \mathbb{N}^{2n+1}$, $n \in \mathbb{N}$. We can now define a strategy for player II in $\mathcal{G}(A)$ to be given by $\sigma(s) := f_s(Q)$ for any $Q \subseteq H^{[m_s]}$.

We claim that σ is a winning strategy for player II in $\mathcal{G}(A)$. Suppose, toward a contradiction, that this is not the case. Then there is a play of the game $x \in \mathbb{N}^{\mathbb{N}}$ played according to σ such that player II lost, that is, $x \in A$, that is, $T(x)$ is well founded. It follows that there is an embedding $f : (T(x), <_{KB}) \rightarrow (H, <)$ (this is where all we needed is that $|H| \geq \aleph_1$). We thus consider the play of the game $\mathcal{G}'(A)$ where $\eta_i := f(u_i)$ when $u_i \in T(x)$ and $\eta_i = 0$ otherwise. It is clear that player I wins this play of the game $\mathcal{G}'(A)$. However, it is also clear that player II played according to their winning strategy in this play of the game, a contradiction. This concludes the proof of Theorem 17.7.2.

Martin's theorem was later generalized by Martin and Steel [130]:

Theorem 17.7.8. *Fix $n \in \omega$ and assume that there are n Woodin cardinals with a measurable cardinal above them. Then Π_{n+1}^1 -determinacy holds.*

We will not mention what Woodin cardinals are except to say that they can be defined using elementary embeddings as above and that they are important in the connection between large cardinals and descriptive set theory (as evidenced by the previous theorem).

We conclude by mentioning a different result of Solovay [168] along these lines that also involves large cardinals:

Theorem 17.7.9 (Solovay). *Suppose that there is a strongly inaccessible cardinal. Then it is consistent that all sets of reals are Lebesgue and Baire measurable.*

Note the differences between the two Solovay theorems. The latter one, while using a weaker large cardinal hypothesis, only gives a consistency result, although it is for all sets of reals.

17.8. Notes and references

The study of large cardinals forms a huge part of modern set theory, and it can be quite overwhelming to the uninitiated. We highly recommend the introductory articles

<https://plato.stanford.edu/entries/independence-large-cardinals/>

and

<https://plato.stanford.edu/entries/large-cardinals-determinacy/#LarCarA>

We follow Jech [89] and Kanamori [94] in most places, but occasionally also borrow from Neeman’s survey article [138]. The term “wordly cardinal” is due to Hamkins. Our proof of Theorem 17.3.12 follows Jech’s treatment in [90]. We follow [28] in our discussion of weakly and strongly compact cardinals. Our proof that measurable cardinals are Ramsey follows Neeman [138] as does much of Section 17.6, although we also use [89] to help fill in some details. Our proof of Martin’s theorem, Theorem 17.7.2, was heavily influenced by some lecture notes of Rosendal, which can be found at

<http://homepages.math.uic.edu/~rosendal/WebpagesMathCourses/MATH511-notes/DST%20notes%20-%20AnalyticDeterminacy03.pdf>

Part 5

Appendices

Logic

A.1. Languages and structures

Definition A.1.1. A language \mathcal{L} consists of three types of symbols: **relation** symbols, **function** symbols, and **constant** symbols. Moreover, each relation and function symbol comes equipped with a natural number, called its **arity**.

Note that some authors lump constant symbols in with the function symbols as 0-ary function symbols.

Example A.1.2. One may consider the language $\mathcal{L} = \{R, F, G, c\}$, where R is a relation symbol of arity 2 (otherwise known as a binary relation symbol), F is a binary function symbol, G is a function symbol of arity 1 (otherwise known as a unary function symbol), and c is a constant symbol.

Definition A.1.3. Let \mathcal{L} be a language. An \mathcal{L} -**structure** \mathcal{M} consists of the following data:

- (1) A nonempty set M , called the **universe** of the structure.
- (2) For each n -ary relation symbol R , a subset $R^{\mathcal{M}} \subseteq M^n$.
- (3) For each n -ary function symbol F , a function $F^{\mathcal{M}} : M^n \rightarrow M$.
- (4) For each constant symbol c , an element $c^{\mathcal{M}} \in M$.

The relation $R^{\mathcal{M}}$ is called the **interpretation of R in \mathcal{M}** . One uses the same nomenclature for $F^{\mathcal{M}}$ and $c^{\mathcal{M}}$.

Example A.1.4. We return to the language introduced in Example A.1.2. One may consider the structure \mathcal{M} whose universe consists of \mathbb{R} , whose interpretation of R is the usual ordering on \mathbb{R} , whose interpretation of F is

addition of real numbers, whose interpretation of G is the additive inverse of a real number, and whose interpretation of c is 0. One might denote this structure as $\mathcal{M} = (\mathbb{R}; <, +, -, 0)$.

Remark A.1.5. If one's intention was to study the previous structure using first-order logic, then it is common to label the symbols using the same symbol as their intended interpretation. In other words, one might replace F by $+$, G by $-$, c by 0, and R by $<$. (One might even then write $2 < 3$ instead of $(2, 3) \in <^{\mathcal{M}}$.)

However, if one does indeed adopt this convention, one must be careful to not conflate the symbol with its interpretation. Indeed, one may consider the \mathcal{L} -structure \mathcal{N} whose universe is the set $\{6, 7, 8\}$, whose interpretation of $<$ is equality, whose interpretation of $+$ is the constant function 6, whose interpretation of $-$ is the constant function 8, and whose interpretation of 0 is also 8. This is still a legitimate \mathcal{L} -structure even though the interpretations of the symbols no longer resemble the “intended” use.

Definition A.1.6. If $\mathcal{L}_1 \subseteq \mathcal{L}_2$ are both languages and \mathcal{M} is an \mathcal{L}_2 -structure, the **reduct of \mathcal{M} to \mathcal{L}_1** is the \mathcal{L}_1 -structure obtained from \mathcal{M} by “forgetting” to interpret the symbols in $\mathcal{L}_2 \setminus \mathcal{L}_1$. If \mathcal{N} is the reduct of \mathcal{M} to \mathcal{L}_1 , one also says that \mathcal{M} is an **expansion** of \mathcal{N} to \mathcal{L}_2 .

Example A.1.7. $(\mathbb{R}; +, 0)$ is the reduct of $(\mathbb{R}; <, +, -, 0)$ to the language consisting just of F and c .

Example A.1.8. Suppose that \mathcal{M} is an \mathcal{L} -structure and $A \subseteq M$ is a subset of the universe. One defines the language \mathcal{L}_A to be the language obtained by adding to \mathcal{L} new constant symbols c_a , one for every $a \in A$. There is then a natural expansion of \mathcal{M} to an \mathcal{L}_A -structure, denoted \mathcal{M}_A , where $c_a^{\mathcal{M}_A} = a$ for every $a \in A$.

A.2. Syntax and semantics

Fix a language \mathcal{L} . We also fix a collection (v_i) of variables.

Definition A.2.1. The set of **\mathcal{L} -terms** is defined by recursion:

- (1) Each variable v_i is an \mathcal{L} -term.
- (2) Each constant symbol c is an \mathcal{L} -term.
- (3) If t_1, \dots, t_n are previously defined \mathcal{L} -terms and F is an n -ary function symbol, then $Ft_1 \cdots t_n$ is also an \mathcal{L} -term.

One thinks of terms as the “nouns” in our set-up as they are intended to name elements of our universe. However, if a term has variables, then it does not name anything until the variables are replaced by elements of

the universe. To define this precisely, given an \mathcal{L} -structure \mathcal{M} , by an **assignment of variables in \mathcal{M}** , we mean a map $v_i \mapsto b_i$ which maps every variable to an element $b_i \in M$. We write \vec{b} as shorthand for this assignment of variables.

Definition A.2.2. Given an \mathcal{L} -structure \mathcal{M} and an assignment of variables \vec{b} , we define the **interpretation of t in \mathcal{M} with respect to the assignment \vec{b}** , denoted $t^{\mathcal{M}}[\vec{b}]$, by recursion on the complexity of t :

- (1) $v_i^{\mathcal{M}}[\vec{b}] = b_i$.
- (2) $c^{\mathcal{M}}[\vec{b}] = c^{\mathcal{M}}$.
- (3) $(Ft_1 \cdots t_n)^{\mathcal{M}}[\vec{b}] = F^{\mathcal{M}}(t_1^{\mathcal{M}}[\vec{b}], \dots, t_n^{\mathcal{M}}[\vec{b}])$.

Example A.2.3. Returning to Example A.1.2, an example of a term is $t = +v_2 0$. For the sake of sanity, we might write this term as $v_2 + 0$. If \vec{b} is an assignment of variables in $\mathcal{M} = (\mathbb{R}; <, +, -, 0)$, then $t^{\mathcal{M}}[\vec{b}] = b_2 + 0$, where this latter 0 is the number 0, as opposed to the symbol 0.

Remark A.2.4. If the v_i 's that occur in t are amongst v_1, \dots, v_n , then we write $t(v_1, \dots, v_n)$ and denote the interpretation $t^{\mathcal{M}}[\vec{b}]$ as $t^{\mathcal{M}}(b_1, \dots, b_n)$ (for the other values of the variables are irrelevant).

We now describe the “assertions” in our set-up:

Definition A.2.5. We define the **\mathcal{L} -formulae** by recursion:

- (1) If t_1 and t_2 are \mathcal{L} -terms, then $t_1 = t_2$ is an \mathcal{L} -formula.
- (2) If t_1, \dots, t_n are \mathcal{L} -terms and R is an n -ary relation symbol, then $Rt_1 \cdots t_n$ is an \mathcal{L} -formula.
- (3) If φ is an \mathcal{L} -formula, then so is $\neg\varphi$.
- (4) If φ_1 and φ_2 are \mathcal{L} -formulae, then so is $\varphi_1 \wedge \varphi_2$.
- (5) If φ is an \mathcal{L} -formula, then so is $\exists v_i \varphi$ for any variable v_i .

The first two kinds of formulae are called **atomic \mathcal{L} -formulae**. A formula is called **quantifier-free** if it is built up using the first four clauses in the previous definition, that is, it does not have any appearances of the quantifier \exists .

We now explain how to define the truth of a formula in a structure:

Definition A.2.6. Given an \mathcal{L} -structure \mathcal{M} and an assignment of variables \vec{b} , one defines the relation $\mathcal{M} \models \varphi[\vec{b}]$ by recursion on the complexity of φ :

- (1) $\mathcal{M} \models t_1 = t_2[\vec{b}]$ if and only if $t_1^{\mathcal{M}}[\vec{b}] = t_2^{\mathcal{M}}[\vec{b}]$.
- (2) $\mathcal{M} \models Rt_1 \cdots t_n[\vec{b}]$ if and only if $(t_1^{\mathcal{M}}[\vec{b}], \dots, t_n^{\mathcal{M}}[\vec{b}]) \in R^{\mathcal{M}}$.
- (3) $\mathcal{M} \models \neg\varphi[\vec{b}]$ if and only if $\mathcal{M} \not\models \varphi[\vec{b}]$.

- (4) $\mathcal{M} \models (\varphi_1 \wedge \varphi_2)[\vec{b}]$ if and only if $[\mathcal{M} \models \varphi_1[\vec{b}] \text{ and } \mathcal{M} \models \varphi_2[\vec{b}]]$.
- (5) $\mathcal{M} \models \exists v_i \varphi[\vec{b}]$ if and only if there is $a \in M$ such that $\mathcal{M} \models \varphi[\vec{b}^{v_i/a}]$, where $\vec{b}^{v_i/a}$ is the assignment of variables defined exactly as \vec{b} except v_i is mapped to a .

An occurrence of a variable in a formula is **free** if it is not bound by a quantifier. (This is a bit vague but can be formalized by a definition by recursion.) If the free variables of φ are amongst v_1, \dots, v_n , then we write $\varphi(v_1, \dots, v_n)$ and then write $\mathcal{M} \models \varphi(b_1, \dots, b_n)$ instead of $\mathcal{M} \models \varphi[\vec{b}]$. An **\mathcal{L} -sentence** is an \mathcal{L} -formula without free variables. If σ is a sentence, then the choice of variable assignment is irrelevant, and we just write $\mathcal{M} \models \sigma$. An **\mathcal{L} -theory** is a set of \mathcal{L} -sentences. If T is an \mathcal{L} -theory, we write $\mathcal{M} \models T$ if $\mathcal{M} \models \sigma$ for all $\sigma \in T$. If there is an \mathcal{L} -structure \mathcal{M} such that $\mathcal{M} \models T$, then we say that T is **satisfiable**. We write $T \models \sigma$ if for every $\mathcal{M} \models T$, we have $\mathcal{M} \models \sigma$. Note that if T is not satisfiable, then $T \models \sigma$ for all \mathcal{L} -sentences σ . \mathcal{L} -theories T_1 and T_2 are called **equivalent**, denoted $T_1 \equiv T_2$, if they have the same class of models.

Definition A.2.7. \mathcal{L} -structures \mathcal{M} and \mathcal{N} are called **elementarily equivalent**, denoted $\mathcal{M} \equiv \mathcal{N}$ if, for all \mathcal{L} -sentences σ , we have $\mathcal{M} \models \sigma$ if and only if $\mathcal{N} \models \sigma$.

Definition A.2.8. A satisfiable \mathcal{L} -theory T is called **complete** if whenever $\mathcal{M}, \mathcal{N} \models T$, we have $\mathcal{M} \equiv \mathcal{N}$.

Definition A.2.9. Given an \mathcal{L} -structure \mathcal{M} , the **theory of \mathcal{M}** , is the theory

$$\text{Th}(\mathcal{M}) := \{\sigma : \mathcal{M} \models \sigma\}.$$

Note that $\text{Th}(\mathcal{M})$ is a complete theory. Moreover, this is the only example of a complete theory: a satisfiable theory T is complete if and only if $T \equiv \text{Th}(\mathcal{M})$ for some (equivalently any) $\mathcal{M} \models T$.

A.3. Embeddings

Fix a language \mathcal{L} .

Definition A.3.1. Suppose that \mathcal{M} and \mathcal{N} are \mathcal{L} -structures and $f : M \rightarrow N$ is a function. We say that:

- (1) f is a **homomorphism from \mathcal{M} to \mathcal{N}** , denoted $f : \mathcal{M} \rightarrow \mathcal{N}$, if:
- (a) for all n -ary relation symbols R and $\vec{a} \in M^n$, we have $\vec{a} \in R^{\mathcal{M}} \Rightarrow f(\vec{a}) \in R^{\mathcal{N}}$;
 - (b) for all n -ary function symbols F and $\vec{a} \in M^n$, we have $f(F^{\mathcal{M}}(\vec{a})) = F^{\mathcal{N}}(f(\vec{a}))$;

- (c) for all constant symbols c , we have $f(c^{\mathcal{M}}) = c^{\mathcal{N}}$.
- (2) If $f : \mathcal{M} \rightarrow \mathcal{N}$ is a homomorphism, we further say that f is an **embedding from \mathcal{M} into \mathcal{N}** if f is injective and (a) above is strengthened to:
- (a') for all n -ary relation symbols R and $\vec{a} \in M^n$, we have $a \in R^{\mathcal{M}} \Leftrightarrow f(\vec{a}) \in R^{\mathcal{N}}$.
- (3) f is an **isomorphism** if it is a surjective embedding.
- (4) An isomorphism $f : \mathcal{M} \rightarrow \mathcal{M}$ is called an **automorphism of \mathcal{M}** .

Definition A.3.2. \mathcal{L} -structures \mathcal{M} and \mathcal{N} are called **isomorphic**, denoted $\mathcal{M} \cong \mathcal{N}$, if there is an isomorphism $f : \mathcal{M} \rightarrow \mathcal{N}$.

Note that isomorphism is an equivalence relation on the class of \mathcal{L} -structures.

Exercise A.3.3. Prove that isomorphic structures are elementarily equivalent.

Definition A.3.4. Given \mathcal{L} -structures \mathcal{M} and \mathcal{N} , we say that \mathcal{M} is a **substructure** of \mathcal{N} , denoted $\mathcal{M} \subseteq \mathcal{N}$, if $M \subseteq N$ and the inclusion map $i : M \rightarrow N$ is an embedding $i : \mathcal{M} \rightarrow \mathcal{N}$.

Definition A.3.5. Suppose that \mathcal{M} and \mathcal{N} are \mathcal{L} -structures and $f : M \rightarrow N$ is a function. We say that f is an **elementary embedding** if, for all \mathcal{L} -formulae $\varphi(v_1, \dots, v_n)$ and all $a_1, \dots, a_n \in M$, we have

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

If $M \subseteq N$ and the inclusion map $i : \mathcal{M} \rightarrow \mathcal{N}$ is an elementary embedding, then we say that \mathcal{M} is an **elementary substructure** of \mathcal{N} , denoted $\mathcal{M} \preceq \mathcal{N}$.

Example A.3.6. Suppose that $\mathcal{L} = \{<\}$, $\mathcal{M} = (\mathbb{N}, <)$, and $\mathcal{N} = (\mathbb{Z}, <)$. Then $\mathcal{M} \subseteq \mathcal{N}$ but $\mathcal{M} \not\preceq \mathcal{N}$, for if $\varphi(x)$ is the formula $\exists y(y < x)$, then $\mathcal{M} \not\models \varphi(0)$ but $\mathcal{N} \models \varphi(0)$.

Definition A.3.7. Suppose that \mathcal{M} is an \mathcal{L} -structure. We define the **diagram** of \mathcal{M} to be the set of quantifier-free \mathcal{L}_M -sentences σ such that $\mathcal{M}_M \models \sigma$. The **elementary diagram** of \mathcal{M} is defined in the same manner except one does not require that σ be quantifier-free.

Theorem A.3.8. *The models of the diagram (resp., elementary diagram) of \mathcal{M} are those \mathcal{L}_M -structures of the form $(\mathcal{N}, (h(a))_{a \in M})$, where $h : \mathcal{M} \rightarrow \mathcal{N}$ is an embedding (resp., elementary embedding).*

Theorem A.3.9 (Löwenheim-Skolem theorems). *Suppose that \mathcal{M} is an \mathcal{L} -structure.*

- (1) Upward Löwenheim-Skolem. *For every cardinal $\kappa \geq \max(|\mathcal{L}|, |M|)$, there is an elementary extension \mathcal{N} of \mathcal{M} with $|N| = \kappa$.*
- (2) Downward Löwenheim-Skolem. *Given any subset $X \subseteq M$, there is an elementary substructure \mathcal{N} of \mathcal{M} such that $X \subseteq N$ and with $|N| = \max(|X|, |\mathcal{L}|)$.*

Definition A.3.10. Suppose that $(I, <)$ is a linearly ordered set and for each $i \in I$, \mathcal{M}_i is an \mathcal{L} -structure. We say that the family $(\mathcal{M}_i)_{i \in I}$ forms a **chain** (resp., **elementary chain**) of \mathcal{L} -structures if, for all $i, j \in I$ with $i < j$, we have that $\mathcal{M}_i \subseteq \mathcal{M}_j$ (resp., $\mathcal{M}_i \preceq \mathcal{M}_j$).

Theorem A.3.11. *If $(\mathcal{M}_i)_{i \in I}$ is a chain of \mathcal{L} -structures, then the union $\bigcup_{i \in I} M_i$ is naturally the universe of a structure $\bigcup_{i \in I} \mathcal{M}_i$ such that $\mathcal{M}_j \subseteq \bigcup_{i \in I} \mathcal{M}_i$ for all $j \in I$. Moreover, if the chain is an elementary chain, then $\mathcal{M}_j \preceq \bigcup_{i \in I} \mathcal{M}_i$ for all $j \in I$.*

A.4. References

A standard undergraduate textbook in logic is Enderton's book [50]. Two good graduate-level textbooks in model theory are Marker's [126] and Tent and Ziegler's [174]. A classic text, which includes many things covered in the rest of this book, is Chang and Keisler's [28]. A recent text in model theory aimed at undergraduates is Kirby's book [105].

Set theory

B.1. The axioms of ZFC

In this book (and in most mathematics textbooks), our underlying axiomatic set theory will be ZFC, which stands for **Zermelo-Fraenkel set theory with choice**. While we will not go into extreme depth into these axioms (see [38] for a fantastic introduction), we would like to at least give a short explanation of the axioms. This will become useful in parts of the book where we are highlighting certain foundational issues concerning ultrafilters.

First, there is an axiom which states the “obvious” property that a set is determined by its members. More precisely, the **axiom of extensionality** states that, given any two sets x and y , if, for all sets z , we have $z \in x$ if and only if $z \in y$, then $x = y$.

Next, there is a group of axioms that allow one to construct new sets from pre-existing sets. For example, the **pairset axiom** allows one to conclude, from sets x and y that are known to exist, that there is a set z such that $x \in z$ and $y \in z$. Similarly, given a set x , the **unionset axiom** allows one to conclude the existence of a set y such that, for all $w \in x$ and all $z \in w$, we have that $z \in y$. The **powerset axiom** allows one to conclude, given a set x , the existence of a set y that contains all the subsets of x as members.

Notice that in the above axioms we did not say, for example, that the set z in the pairset axiom consisted *precisely* of x and y , that is, $z = \{x, y\}$, but rather that merely that $x \in z$ and $y \in z$. However, one can infer the existence of such a set using the **comprehension axiom**, which states that if x is a set and $\varphi(v)$ is a formula (with parameters) in the first-order language of set theory (which consists solely of a binary relation symbol, interpreted as membership), then there is a set y that consists of precisely

those sets in x that satisfy φ . Thus, if one lets $\varphi(v)$ be $v = x \vee v = y$, then one can form $w = \{a \in z : \varphi(a)\}$, and it is clear then that $w = \{x, y\}$.

There is one final “set-creation” axiom (besides the axiom of choice, which we will describe in detail later), called the **axiom of replacement**, which is a little harder to motivate (but is ultimately very useful) and which states: if x is a set and $\varphi(v_1, v_2)$ defines a function on x , that is, for all $w \in x$, there is a unique z such that $\varphi(w, z)$ holds, then there is a set y consisting of precisely those z for which there is $w \in x$ such that $\varphi(w, z)$. In other words, one is allowed to “replace” all the elements of x by their “images” under the function whose graph is φ and collect them into a new set.

At this point, we do not actually have any guarantee that there are any sets at all! We need an axiom that guarantees that there is a set, and the most minimalistic way of achieving this is the **nullset axiom**, which simply states that there is a set x with no members, that is, for all sets y , we have $y \notin x$. By the axiom of extensionality, this set is unique and is denoted by \emptyset .

Now that we have some set at our disposal, with the aid of our set-creating axioms, we can make some new sets, such as $\{\emptyset, \emptyset\} = \{\emptyset\}$. With more effort, we can make all sorts of *finite* sets. Mathematics would be fairly mundane if we only had finite sets, so we need an axiom asserting that infinite sets exist. The **axiom of infinity** states that there is a set x such that $\emptyset \in x$ and, for all $y \in x$, one also has $\{y\} \in x$. It is a fairly easy exercise to see that such a set is infinite.

We have now listed all but two of the axioms of ZFC. The **axiom of foundation** asserts that there is no infinite, decreasing sequence of sets $x_1 \ni x_2 \ni x_3 \ni \dots$. The real utility of this axiom is that the hierarchical perspective on the set-theoretic universe given in Section B.4 is accurate.

The axioms mentioned thus far constitute the axioms of ZF. There is one axiom which turns ZF into ZFC, and that is the **axiom of choice** (AC), which states the following: Suppose that x is a set whose elements are nonempty and pairwise disjoint. Then there is a set y that contains exactly one element of each element of x . In other words, there is a *choice function* for x which assigns, to each $w \in x$, an element $f(w) \in w$. The set y above is just the image of this function.

While AC is an axiom in the spirit of the unionset, pairing, replacement, and comprehension axioms in that it creates new sets from old, it is not quite as “intuitive” or “constructive” as the others. Nevertheless, nearly almost all mathematicians treat AC as an acceptable axiom and most do not make a fuss when using it. In fact, AC is actually equivalent, in ZF, to familiar statements of mathematics, such as Tychonoff’s theorem, which

states that a product of compact spaces is compact again. Tychonoff's theorem is ubiquitous throughout mathematics and thus if one is willing to accept it as a valid result, then one must also accept AC as well.

On the other hand, AC does yield some pathological consequences, such as the existence of non-Lebesgue measurable sets. For this reason, some mathematicians make a point of explicitly noting when AC is used in a proof. For a very interesting and thorough discussion of AC, we refer the reader to [90].

As we will discuss in Section B.5, AC does not follow from the axioms of ZF, and thus we really are obtaining a stronger theory in ZFC by adding it as an axiom.

While the version of AC that we presented is the easiest to define (and it explains the name), in practice, one often uses AC in one of a variety of other avatars, as we now explain.

Definition B.1.1. A **partially ordered set** (or **poset**) is a set P equipped with a binary relation \leq satisfying, for all $x, y, z \in P$:

- (1) $x \leq x$;
- (2) $x \leq y$ and $y \leq x$ implies $x = y$;
- (3) $x \leq y$ and $y \leq z$ implies $x \leq z$.

The poset (P, \leq) is called a **linearly ordered set** if it further satisfies, for all $x, y \in X$:

- (4) $x \leq y$ or $y \leq x$.

If \leq is a partial ordering, then the associated **strict ordering** is the binary relation $<$ defined by $x < y$ if $x \leq y$ but $x \neq y$.

Definition B.1.2. Suppose that (P, \leq) is a poset.

- (1) A **chain** in (P, \leq) is a linearly ordered subset of P .
- (2) For $A \subseteq P$, an **upper bound for A** is an element $x \in P$ such that $a \leq x$ for all $a \in A$.
- (3) An element $x \in P$ is called **maximal** if there does not exist $y \in P$ with $x \leq y$ and $x \neq y$. The notion of a **minimal** element is defined analogously.

Definition B.1.3. **Zorn's lemma** is the statement: for any poset (P, \leq) , if every chain in (P, \leq) has an upper bound, then (P, \leq) has a maximal element.

Theorem B.1.4. In ZF, AC is equivalent to Zorn's lemma.

Definition B.1.5. A linear ordering \leq on a set I is called a **well-ordering** of I if every nonempty subset of I has a minimal element. If \leq is a well-ordering on I , we call the pair (I, \leq) a **well-ordered set**.

Theorem B.1.6. In ZF, AC is equivalent to the **well-ordering principle**, which states that every set can be well ordered, that is, for any set x , there is a well-ordering \leq on x .

B.2. Ordinals

Definition B.2.1. A set x is **transitive** if for all $y \in x$ and $z \in y$, we have $z \in x$.

Definition B.2.2. An **ordinal** is a set x such that x is transitive and (x, \in) is a well-ordered set. The class of ordinals is denoted On .

Lemma B.2.3. The class of ordinals is transitive and well ordered by \in .

One often uses Greek letters such as α and β for ordinals. By the previous lemma, we sometimes write $\alpha < \beta$ instead of $\alpha \in \beta$.

Lemma B.2.4. If α and β are ordinals, then $(\alpha, \in) \cong (\beta, \in)$ if and only if $\alpha = \beta$.

Lemma B.2.5. If (x, \leq) is a well-ordering, then there is a unique ordinal α such that $(x, \leq) \cong (\alpha, \in)$.

Lemma B.2.6.

- (1) \emptyset is an ordinal. If α is any ordinal, then $\emptyset \leq \alpha$.
- (2) If α is an ordinal, then $\alpha \cup \{\alpha\}$ is also an ordinal, denoted $\alpha + 1$. Moreover, if β is an ordinal, then either $\beta \leq \alpha$ or else $\alpha + 1 \leq \beta$. For this reason, $\alpha + 1$ is called the **successor of α** .
- (3) If C is a set of ordinals, then $\sup C$ is also an ordinal and is equal to $\bigcup_{\alpha \in C} \alpha$.

Example B.2.7. One often denotes \emptyset by 0. Note then that $\{0\}$ is an ordinal, which we denote by 1. It follows then that $\{0, 1\}$ is an ordinal, which we denote by 2. In this way, every natural number n is an ordinal when identified with $\{0, \dots, n - 1\}$. One then sets $\omega = \bigcup_{n \in \mathbb{N}} n = \{0, 1, 2, \dots\}$, which is the first infinite ordinal. But then one can keep going, considering $\omega + 1 = \omega \cup \{\omega\}$, which is also an ordinal, and so on.

A **successor ordinal** is one of the form $\alpha + 1$ for some ordinal α . An ordinal that is not a successor ordinal is called a **limit ordinal**.

One can perform proofs by “transfinite induction” on ordinals:

Theorem B.2.8. *Suppose that P is a statement about ordinals. Suppose that:*

- P is true for 0.
- For any ordinal α , if P is true for α , then P is true for $\alpha + 1$.
- For any limit ordinal α , if P is true for β for all $\beta < \alpha$, then P is true for α .

Then P is true for all ordinals.

B.3. Cardinals

Definition B.3.1. If A is a set, the **cardinality of A** , denoted $|A|$, is the least ordinal α such that there is a well-ordering $<$ on A such that $(A, \leq) \cong (\alpha, \in)$.

Thus, the cardinality of a set is an ordinal. Note that by the axiom of choice, every set admits a well-ordering, so every set has a cardinality. Ordinals that are cardinalities of sets are called **cardinals**. One often uses κ, λ to denote cardinals.

Exercise B.3.2. For any two sets A and B , we have that $|A| = |B|$ if and only if there is a bijection $A \rightarrow B$.

Exercise B.3.3. Every element of ω is a cardinal. ω is also a cardinal.

It is customary to write \aleph_0 when thinking of ω as a cardinal.

Exercise B.3.4. For any set X , $|X| < |\mathcal{P}(X)|$.

Theorem B.3.5 (Schroder-Berstein). $|A| = |B|$ if and only if $|A| \leq |B|$ and $|B| \leq |A|$.

Definition B.3.6. A set is **countable** if $|A| \leq \aleph_0$.

Set $\omega_1 := \{\alpha : \alpha \text{ is a countable ordinal}\}$.

Exercise B.3.7. Prove the following:

- (1) ω_1 is an ordinal.
- (2) ω_1 is uncountable.
- (3) ω_1 is a cardinal.

We usually write \aleph_1 when we think of ω_1 as a cardinal. More generally:

Definition B.3.8.

- (1) $\aleph_{\alpha+1} = \{\delta : |\delta| = \aleph_\alpha\}$.
- (2) If α is a limit ordinal, then $\aleph_\alpha = \sup_{\beta < \alpha} \aleph_\beta$.

Lemma B.3.9.

- (1) Each \aleph_α is a cardinal.
- (2) $\alpha < \beta$ if and only if $\aleph_\alpha < \aleph_\beta$.
- (3) If κ is an infinite cardinal, then $\kappa = \aleph_\alpha$ for some α .

Definition B.3.10. κ^+ is the least cardinal bigger than κ .

Note that κ^+ exists by Exercise B.3.4.

Exercise B.3.11. $\aleph_{\alpha+1} = \aleph_\alpha^+$.

A **successor cardinal** is one of the form κ^+ for some κ . Nonzero cardinals that are not successor cardinals are called **limit cardinals**.

Definition B.3.12. Given any ordinal α , the **cofinality** of α , denoted $\text{cof}(\alpha)$, is the least cardinal λ such that there is a function $f : \lambda \rightarrow \alpha$ for which $f(\lambda)$ is unbounded in α .

Exercise B.3.13. Prove that $\text{cof}(\aleph_0) = \aleph_0$. More generally, prove that $\text{cof}(\aleph_\alpha) = \text{cof}(\alpha)$.

Definition B.3.14. An infinite cardinal κ is called **regular** if $\text{cof}(\kappa) = \kappa$; otherwise, it is called **singular**.

Lemma B.3.15. Any infinite successor cardinal is regular.

Definition B.3.16. If κ and λ are cardinals, we define:

- (1) $\kappa + \lambda$ to be the cardinality of the disjoint union of κ and λ .
- (2) $\kappa \cdot \lambda$ to be the cardinality of $\kappa \times \lambda$.

Lemma B.3.17. Suppose that κ and λ are cardinals:

- (1) If $\kappa, \lambda \in \omega$, then $\kappa + \lambda$ and $\kappa \cdot \lambda$ agree with usual natural number addition and multiplication.
- (2) Suppose that $\kappa \neq 0$, $\lambda \neq 0$, and one of κ or λ are infinite. Then $\kappa + \lambda = \kappa \cdot \lambda = \max(\kappa, \lambda)$.

Definition B.3.18. Suppose that κ and λ are cardinals. We define κ^λ to be the cardinality of the set of functions $\lambda \rightarrow \kappa$.

Another name for 2^{\aleph_0} is \mathfrak{c} .

Exercise B.3.19. Suppose that κ and λ are finite cardinals. Prove that the above definition of κ^λ agrees with the usual exponentiation of natural numbers.

Lemma B.3.20. Suppose that κ , λ , and μ are cardinals. Then:

- (1) $(\kappa^\lambda)^\mu = \kappa^{\lambda \cdot \mu}$.

- (2) If λ is infinite and $2 \leq \kappa < \lambda$, then $2^\lambda = \kappa^\lambda = \lambda^\lambda$.
 (3) If κ is regular and $\lambda < \kappa$, then $\kappa^\lambda = \max(\kappa, \sup_{\mu < \kappa} \mu^\lambda)$.

Proposition B.3.21 (König's theorem). *If κ is infinite, then $\kappa^{\text{cof}(\kappa)} > \kappa$.*

Definition B.3.22. The **continuum hypothesis** (CH) is the statement that $2^{\aleph_0} = \aleph_1$. The **generalized continuum hypothesis** (GCH) is the statement that, for all ordinals α , we have $2^{\aleph_\alpha} = \aleph_{\alpha+1}$.

B.4. V and L

The ZFC axioms allow us to paint a nice picture of the way the set-theoretic universe is constructed from the bottom up. One starts with the 0th level $V_0 = \emptyset$. Supposing the α th level V_α has been constructed (where α is some ordinal), one constructs the next level $V_{\alpha+1} := V_\alpha \cup \mathcal{P}(V_\alpha)$, that is, one adds to V_α all subsets of V_α . If α is a limit ordinal, then one sets $V_\alpha := \bigcup_{\beta < \alpha} V_\beta$. (There is a healthy use of both the unionset and replacement axioms here.) In this way, one can continue this process through all of the ordinals and one sets $V := \bigcup_\alpha V_\alpha$, which is then the *universe* of all sets (which is a proper class, that is, is not a set). This presentation of the set-theoretic universe is called the **Zermelo hierarchy**.

A useful construction, due to Gödel, is to reconsider what to do at successor stages in the previous process, namely, instead of adding all subsets of V_α to $V_{\alpha+1}$, one instead should only add “definable” subsets. The hierarchy thus obtained is then denoted $L = \bigcup_\alpha L_\alpha$ and is called the **constructible universe**. To be clear, at stage $\alpha + 1$, one only adds those sets of the form $\{x \in L_\alpha : \varphi(x) \text{ holds}\}$, where φ is a formula with parameters from L_α , and where “holding” means that if we interpret all quantifiers to range only over L_α , then the resulting formula is true. (See Appendix B.6 for a precise definition.)

One can show that L is also a model of ZFC. The **axiom of constructibility** is the statement $V = L$ and is an interesting extension of ZFC. For example, while ZFC cannot prove nor disprove CH (see the next section), ZFC, together with the axiom of constructibility, can prove CH.

L is a prototypical example of an **inner model** of ZFC, where an inner model of ZFC is a transitive class M such that (M, \in) is a model of ZFC containing all ordinals.

B.5. Relative consistency statements

By Gödel's completeness theorem for first-order logic, ZF is consistent (that is, does not prove both some sentence σ and its negation $\neg\sigma$) if and only if it has a model, which would be a structure $\mathcal{M} = (M, E)$ in the language of

set theory such that all axioms of ZF are true in \mathcal{M} . (We expand on this somewhat in the next section.) Unfortunately, Gödel's second incompleteness theorem implies that *if* ZF is consistent, then it cannot prove its own consistency. More precisely, there is a sentence $\text{Con}(\text{ZF})$ in the language of set theory which asserts that ZF is consistent and that $\text{ZF} \not\vdash \text{Con}(\text{ZF})$ *unless* ZF is itself inconsistent (in which case it proves all assertions). Of course, if ZF is consistent, then $\text{ZF} \not\vdash \neg \text{Con}(\text{ZF})$ either, whence $\text{Con}(\text{ZF})$ is an example of a sentence in the language of set theory which is *independent* of ZF.

In this book, we adopt the philosophy adopted by (nearly) all mathematicians, namely that ZF is in fact consistent.

In the last section, we mentioned that the constructible universe L was a model of ZF. This use of the word model is somewhat imprecise in that L is a class, not a set, and thus cannot be a model in the sense we are using here. To be more precise, working in ZF, if one starts with a model M of ZF, then by mimicking the construction of L , but now inside of M , then one obtains a submodel N of M that is now a model of ZFC. Consequently, the statement $\text{Con}(\text{ZF}) \rightarrow \text{Con}(\text{ZFC})$ is an axiom of ZF.

One refers to this kind of a result as a **relative consistency result** in that ZFC is consistent if one assumes that ZF itself is consistent, and this implication is actually provable from the axioms of ZF. Since $\text{Con}(\text{ZFC})$ clearly implies $\text{Con}(\text{ZF})$, we say that ZF and ZFC are *equiconsistent* in ZF. Thus, our standing assumption that ZF is consistent also implies that ZFC is consistent.

It turns out that the negation of the axiom of choice is also relatively consistent with ZF, namely that if one starts with a model M of ZF, then one can construct an *outer model* $M[G]$ obtained using the method of *forcing*, such that $M[G]$ is a model of $\text{ZF} + \neg \text{AC}$. Consequently, ZF and $\text{ZF} + \neg \text{AC}$ are also **equiconsistent** in ZF.

Note also that, if $\text{ZF} \vdash \text{AC}$, then $\text{ZF} + \neg \text{AC}$ would be an inconsistent theory, whence so would ZF since it is equiconsistent with $\text{ZF} + \neg \text{AC}$. Thus, $\text{ZF} \not\vdash \text{AC}$. For the same reason, $\text{ZF} \not\vdash \neg \text{AC}$ and AC is another statement independent of ZF.

Another example of this kind of discussion concerns the continuum hypothesis CH. In the model L , CH is true; consequently $\text{ZFC} + \text{CH}$ is equiconsistent with ZFC in ZFC. As in the case of AC, one can use forcing to produce an outer model of $\text{ZFC} + \neg \text{CH}$, whence $\text{ZFC} + \neg \text{CH}$ is also equiconsistent with ZFC. As a result, CH is independent of ZFC.

B.6. Relativization and absoluteness

Throughout this discussion, we fix a class M and a binary relation E on M . Given a formula $\varphi(x_1, \dots, x_n)$, one defines its **relativization to** (M, E) , denoted $\varphi^{(M, E)}$, by recursion on complexity of formulae:

- $(x = y)^{(M, E)}$ is simply the formula $x = y$;
- $(x \in y)^{(M, E)}$ is simply the formula $x \in y$;
- $(\neg\varphi)^{(M, E)}$ is the formula $\neg\varphi^{(M, E)}$;
- $(\varphi \wedge \psi)^{(M, E)}$ is the formula $\varphi^{(M, E)} \wedge \psi^{(M, E)}$;
- $(\exists x\varphi)^{(M, E)}$ is the formula $\exists x(x \in M \wedge \varphi^{(M, E)})$.

When E is just membership in M , we write φ^M instead of $\varphi^{(M, E)}$. Intuitively speaking, φ^M is the statement φ relativized to M in the sense that all quantifiers are restricted to varying over M instead of the entire universe V .

Given a sentence σ in the language of set theory, we say that (M, E) is a model of σ if $\sigma^{(M, E)}$ is a true statement. Thus, when we spoke of a model (M, E) of ZF in the previous sections, we really meant that for every sentence σ in the axioms for ZF, $\sigma^{(M, E)}$ is a true statement.

From now on, we speak only of models (M, \in) . A formula $\varphi(x_1, \dots, x_n)$ is **absolute for** M if, for any $a_1, \dots, a_n \in M$, we have $\varphi^M(a_1, \dots, a_n)$ holds if and only if $\varphi(a_1, \dots, a_n)$ holds. In other words, truth in M and truth in V coincide for instances of φ with parameters from M .

It is important to note that a large class of formulae are absolute for any transitive model. First, we let Δ_0 denote the smallest class of formulae containing the quantifier-free formulae, closed under connectives, and closed under **bounded quantification**, that is, if φ is a Δ_0 formula, then so are $(\exists x)(x \in y \wedge \varphi)$ and $(\forall x)(x \in y \rightarrow \varphi)$. It is routine to show that all Δ_0 formulae are absolute for any transitive model. Moreover, many familiar notions can be expressed using Δ_0 -formulae, such as “ x is a subset of y ”, “ x is transitive”, and “ x is an ordinal”. Building an inventory of the formulae that are Δ_0 and, more generally, the formulae that are absolute, is important in considering various models of set theory. It is also important to note that statements involving cardinals are often not absolute (and thankfully so, for that is what allows one to prove independence of statements such as CH).

Given any set x , let $\text{rank}(x)$ be the least ordinal α such that $x \in V_{\alpha+1}$. It is a fact that the rank function is absolute for transitive models of ZF.

B.7. References

As mentioned above, Devlin's book [38] is a nice, casual introduction to set theory. Another book aimed at undergraduates is Enderton's book [49]. An encyclopedic treatment of set theory (and a source for many set-theoretic facts occurring throughout this book) is Jech's monograph [89]. For an in depth discussion of the axiom of choice, see Jech's book [90]. A treatment of forcing aimed at mathematicians (as opposed to logicians) is Weaver's book [182].

Category theory

C.1. Categories

In this section, we develop the small bit of category theory needed throughout the book. The basic idea behind category theory is that almost always, when one encounters some new mathematics, they are presented first with the type of objects that will be studied and the appropriate functions between these objects. A category axiomatizes this common thread:

Definition C.1.1. A **category** \mathcal{C} consists of the following data:

- (1) a class of **objects** $\text{Obj}(\mathcal{C})$,
- (2) for each pair of objects X and Y , a collection $\text{Hom}_{\mathcal{C}}(X, Y)$ of **morphisms** from X to Y , and
- (3) a **composition operation** that assigns, to each triple of objects X, Y , and Z , a map $\circ : \text{Hom}_{\mathcal{C}}(Y, Z) \times \text{Hom}_{\mathcal{C}}(X, Y) \rightarrow \text{Hom}_{\mathcal{C}}(X, Z)$

satisfying the following two axioms:

- (i) for each quadruple of objects X, Y, Z , and W and morphisms $f \in \text{Hom}_{\mathcal{C}}(X, Y)$, $g \in \text{Hom}_{\mathcal{C}}(Y, Z)$, and $h \in \text{Hom}_{\mathcal{C}}(Z, W)$, we have $h \circ (g \circ f) = (h \circ g) \circ f$, and
- (ii) for each object X , there is a morphism $1_X \in \text{Hom}_{\mathcal{C}}(X, X)$ with the property that, for any object Y , we have
 - for any morphism $f \in \text{Hom}_{\mathcal{C}}(X, Y)$, we have $f \circ 1_X = f$, and
 - for any morphism $g \in \text{Hom}_{\mathcal{C}}(Y, X)$, we have $1_X \circ g = g$.

Item (i) in the previous definition simply asks that composition be associative, while item (ii) asks for the existence of an identity morphism. Note

that there is no requirement that $\text{Hom}_{\mathcal{C}}(X, Y)$ be nonempty for each pair of objects X and Y .

Examples C.1.2.

- (1) The category of sets, denoted Set , is defined to be the category whose objects are sets X , whose morphisms $f \in \text{Hom}_{\text{Set}}(X, Y)$ are simply functions $f : X \rightarrow Y$, and whose composition is given by usual function composition.
- (2) The category of vector spaces over some field F , denoted $F\text{-vs}$, is defined to be the category whose objects are vector spaces over F , whose morphisms $T \in \text{Hom}_{F\text{-vs}}(V, W)$ are linear transformations $T : V \rightarrow W$, and whose composition is given by usual function composition.
- (3) The category of groups, denoted Group , is defined to be the category whose objects are groups, whose morphisms $f \in \text{Hom}_{\text{Group}}(G, H)$ are group homomorphisms $f : G \rightarrow H$, and whose composition is given by usual function composition.
- (4) The category of topological spaces, denoted Top , is defined to be the category whose objects are topological spaces X , whose morphisms $f \in \text{Hom}_{\text{Top}}(X, Y)$ are continuous functions $f : X \rightarrow Y$, and whose composition is given by usual function composition.
- (5) Given a first-order language \mathcal{L} , the category of \mathcal{L} -structures has as its objects \mathcal{L} -structures and whose morphisms are homomorphisms between \mathcal{L} -structures (see Section 6.10).

We often use the suggestive notation $f : X \rightarrow Y$ to denote an element $f \in \text{Hom}_{\mathcal{C}}(X, Y)$ as morphisms are usually genuine functions. So as to not give the reader the idea that morphisms in categories are always functions and composition is always composition of functions, we present the following example:

Example C.1.3. Let $\mathbb{P} := (P, \leq)$ be a partially ordered set. We can consider this partially ordered set as a category (which we also denote by \mathbb{P}) whose objects are the elements x of P and where, for each $x, y \in P$, we have that $\text{Hom}_{\mathbb{P}}(x, y) \neq \emptyset$ if and only if $x \leq y$, in which case there is a unique morphism in this set. (There is a unique choice of composition map in this case.)

Definition C.1.4. Suppose that \mathcal{C} is a category. We say that \mathcal{C} is **locally small** if, for all objects X and Y , $\text{Hom}_{\mathcal{C}}(X, Y)$ is a set (as opposed to a proper class). If, in addition, $\text{Obj}(\mathcal{C})$ is a set, we say that \mathcal{C} is **small**.

Definition C.1.5. A morphism $f : X \rightarrow Y$ is an **isomorphism** if there is a morphism $g : Y \rightarrow X$ such that $g \circ f = 1_X$ and $f \circ g = 1_Y$.

Definition C.1.6. A morphism $f : X \rightarrow Y$ is a(n):

- (1) **epimorphism** if for all morphisms $g_1, g_2 : Y \rightarrow Z$, $g_1 \circ f = g_2 \circ f$ implies $g_1 = g_2$;
- (2) **monomorphism** if for all morphisms $g_1, g_2 : Z \rightarrow X$, $f \circ g_1 = f \circ g_2$ implies $g_1 = g_2$.

Definition C.1.7. If \mathcal{C} and \mathcal{D} are categories, we say that \mathcal{D} is a **subcategory** of \mathcal{C} if:

- (1) $\text{Obj}(\mathcal{D}) \subseteq \text{Obj}(\mathcal{C})$;
- (2) for all $X, Y \in \text{Obj}(\mathcal{D})$, we have $\text{Hom}_{\mathcal{D}}(X, Y) \subseteq \text{Hom}_{\mathcal{C}}(X, Y)$;
- (3) for all $X \in \text{Obj}(\mathcal{D})$, $1_X \in \text{Hom}_{\mathcal{D}}(X, X)$; and
- (4) if $f \in \text{Hom}_{\mathcal{D}}(X, Y)$ and $g \in \text{Hom}_{\mathcal{D}}(Y, Z)$, then the composition of f and g in the categories \mathcal{C} and \mathcal{D} coincide.

\mathcal{D} is further said to be a **full subcategory** of \mathcal{C} if, for any two objects X and Y in \mathcal{D} , $\text{Hom}_{\mathcal{C}}(X, Y) = \text{Hom}_{\mathcal{D}}(X, Y)$.

Example C.1.8. The category of compact Hausdorff spaces, equipped with continuous functions as morphisms, is a full subcategory of the category of all topological spaces.

Definition C.1.9. Given a category \mathcal{C} , the **opposite category**, denoted \mathcal{C}^{op} , has the same objects as \mathcal{C} but for which one defines $\text{Hom}_{\mathcal{C}^{op}}(Y, X) := \text{Hom}_{\mathcal{C}}(X, Y)$, that is, the arrows of the morphisms are “reversed”.

C.2. Functors, natural transformations, and equivalences of categories

Now that we have defined categories, we now describe the appropriate maps between categories. There are actually two kinds of maps that one considers, namely those that “preserve” the directions of arrows and those that “reverse” the directions of arrows.

Definition C.2.1. Suppose that \mathcal{C} and \mathcal{D} are categories. A **covariant functor** F from \mathcal{C} to \mathcal{D} , denoted $F : \mathcal{C} \rightarrow \mathcal{D}$, is described by the data

- (1) each object X from \mathcal{C} is mapped to an object $F(X)$ from \mathcal{D} , and
- (2) each morphism $f : X \rightarrow Y$ from \mathcal{C} is mapped to a morphism $F(f) : F(X) \rightarrow F(Y)$,

satisfying the axioms

- (i) $F(1_X) = 1_{F(X)}$ for all objects X from \mathcal{C} , and
- (ii) for all morphisms $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ from \mathcal{C} , one has $F(g \circ f) = F(g) \circ F(f)$.

One defines a **contravariant functor** $F : \mathcal{C} \rightarrow \mathcal{D}$ in the same way except that (2) above is replaced with

$$(2') \text{ each morphism } f : X \rightarrow Y \text{ from } \mathcal{C} \text{ is mapped to a morphism } F(f) : F(Y) \rightarrow F(X),$$

and item (ii) is replaced with

$$(ii') \text{ for all morphisms } f : X \rightarrow Y \text{ and } g : Y \rightarrow Z \text{ from } \mathcal{C}, \text{ one has } F(g \circ f) = F(f) \circ F(g).$$

Various examples of functors will be encountered throughout the text.

Exercise C.2.2. Verify that the category whose objects are categories and whose morphisms are covariant functors is actually a category. In particular, verify that the composition of two functors is a functor.

We finally need to explain what the appropriate notion of isomorphism is for categories. The definition does not simply state that there is a functor that has an inverse but rather that there is a pair of functors whose composition in either order are “naturally the same” as the identity functor. This leads us first to the definition of natural transformation.

Definition C.2.3. Suppose that $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{C} \rightarrow \mathcal{D}$ are two covariant functors. A **natural transformation** η from F to G is an assignment to each object X from \mathcal{C} , a morphism $\eta_X : F(X) \rightarrow G(X)$ in \mathcal{D} so that, for each morphism $f : X \rightarrow Y$ from \mathcal{C} , we have $\eta_Y \circ F(f) = G(f) \circ \eta_X$. We say that the natural transformation is a **natural isomorphism** if each η_X is an isomorphism.

Definition C.2.4. We say that \mathcal{C} and \mathcal{D} are **equivalent** if there are covariant functors $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$ such that $G \circ F$ is naturally isomorphic to $1_{\mathcal{C}}$ and $F \circ G$ is naturally isomorphic to $1_{\mathcal{D}}$. We say that \mathcal{C} and \mathcal{D} are **dually equivalent** if the same is true but with contravariant functors.

C.3. Limits

Definition C.3.1. Let \mathcal{I} and \mathcal{C} be categories. Then a **diagram in \mathcal{C} of shape \mathcal{I}** is a functor $F : \mathcal{I} \rightarrow \mathcal{C}$. The diagram is **small** if \mathcal{I} is small.

Definition C.3.2. Let $F : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram.

- (1) A **cone** for F is an object $C \in \mathcal{C}$ and a collection of morphisms $f_A : C \rightarrow F(A)$, one for each object A of \mathcal{I} , such that, for all morphisms $f : A \rightarrow B$ in \mathcal{I} , we have $F(f) \circ f_A = f_B$.

- (2) A **limit** of F is a universal cone (L, f) of F , meaning that if (C, g) is any other cone for F , then there is a unique morphism $h : C \rightarrow L$ such that $f_A \circ h = g_A$ for all objects A of \mathcal{I} .

Using the universal property, if a limit of F exists, then it is unique, whence we may refer to it as *the* limit of F . In regards to our discussion in Section 6.10, we will only need the following special cases of limits.

First suppose that \mathcal{I} consists of a set of objects I and whose only morphisms are the identity morphisms. (Such a category is called a **discrete category**.) In this case, a diagram $F : \mathcal{I} \rightarrow \mathcal{C}$ with shape \mathcal{I} is simply a family $(A_i)_{i \in I}$ of objects in \mathcal{C} . In this case, the limit of F , should it exist, is called the **product** of the family $(A_i)_{i \in I}$, denoted $\prod_{i \in I} A_i$. The cone property, in this case, simply states that, for every $i \in I$, there is a morphism $f_i : \prod_{i \in I} A_i \rightarrow A_i$, while the universality property implies that, whenever B is an object from \mathcal{C} equipped with morphisms $g_i : B \rightarrow A_i$ for each $i \in I$, then there is a unique morphism $h : B \rightarrow \prod_{i \in I} A_i$ such that $f_i \circ h = g_i$.

In the category of sets, products always exist and are simply cartesian products with the usual projection maps. Likewise, products exist in the category of groups and the category of topological spaces (with the usual product topology) as well as the full subcategory of compact (Hausdorff) spaces (by Tychonoff's theorem).

Another special case of limits is when the objects of \mathcal{I} are the elements in a *directed* poset (I, \leq) , where the morphisms are as in Example C.1.3. (Recall that a poset (I, \leq) is directed if, for all $a, b \in I$, there is $c \in I$ such that $a \leq c$ and $b \leq c$.) In this case, we consider diagrams $F : \mathcal{I}^{op} \rightarrow \mathcal{C}$, and the limit, if it exists, is called the **inverse limit** or **projective limit** of the directed family $(A_i)_{i \in I}$, denoted $\varprojlim A_i$. The cone property now implies that the maps $f_i : \varprojlim A_i \rightarrow A_i$ satisfy $p_{ij} f_i = f_j$ when $i \leq j$ (here, the unique morphism from $A_i \rightarrow A_j$ is denoted p_{ij}).

Inverse limits exist in the category of sets: $\varprojlim A_i = \{(a_i)_{i \in I} \in \prod_{i \in I} A_i : a_j = p_{ij}(a_i) \text{ for all } i \leq j\}$. The same underlying set works in the case of the groups and compact topological spaces.

The dual notion of **colimit** is defined by turning the arrows around. More precisely, a **cocone** of F is an object C with morphisms $f_A : F(A) \rightarrow C$ such that, for all morphisms $f : A \rightarrow B$, we have $f_B \circ F(f) = f_A$. A **colimit** of F is a universal cocone (L, f) of F , meaning that for every cone (C, g) of F , there is a unique morphism $h : L \rightarrow C$ such that, for all objects $A \in \mathcal{I}$, we have $h \circ f_A = g_A$. Once again, if the colimit of F exists, it is unique.

In the case of diagrams stemming from discrete categories, colimits are called **coproducts** and are denoted $\coprod_{i \in I} A_i$. In the case of the category of sets, the coproduct is just the disjoint union with inclusion maps. In

the category of topological spaces, the coproduct is the direct sum. (The category of compact Hausdorff spaces is more complicated and is discussed in Section 6.10.)

In the case of digrams stemming from directed sets, the colimit, if it exists, is called the **direct limit** of the family and is denoted $\varinjlim A_i$. The direct limit in the case of the category of sets is the directed union.

Suppose that $G : \mathcal{C} \rightarrow \mathcal{D}$ is a covariant functor and $F : \mathcal{I} \rightarrow \mathcal{C}$ is a diagram in \mathcal{C} . If (C, f) is a cone for F , then (GC, Gf) is easily seen to be a cone for the diagram $G \circ F$. If (C, f) is the limit of F , it need not be the case that (GC, Gf) is the limit of $G \circ F$; if it is, we say that G preserves the limit (C, f) . We say that G is **continuous** if it preserves all limits. The same discussion can be made for colimits, leading to the notion of a **co-continuous** functor. If G is contravariant, then it takes cones (resp., co-cones) of F to co-cones (resp., cones) of $(G \circ F)$ and is called continuous (resp., co-continuous) if it takes limits (resp., colimits) to co-limits (resp., limits).

A category \mathcal{C} is called **complete** if all limits of small diagrams in \mathcal{C} exist and similarly for **co-complete**. It is called **bi-complete** if it is both. Examples of bicomplete categories include sets, groups, and (compact Hausdorff) topological spaces.

C.4. References

A standard resource for category theory is MacLane's book [116].

Hints and solutions to selected exercises

Chapter 1. Ultrafilter basics

Exercise 1.1.11. To see that $\emptyset \notin \mathcal{U}$ and $S \in \mathcal{U}$, apply the hypothesis with the partition $S = S \cup \emptyset$. To see that \mathcal{U} is closed under supersets, suppose that $A \in \mathcal{U}$ and $A \subseteq B$. By considering the partition $S = (S \setminus B) \cup (B \setminus A) \cup A$, one sees that $S \setminus B \notin \mathcal{U}$; by considering the partition $S = (S \setminus B) \cup B$, we infer that $B \in \mathcal{U}$. Finally, if $A, B \in \mathcal{U}$, then by considering the partition $S = (A \cap B) \cup (A \Delta B) \cup (S \setminus (A \cup B))$ and using closure under supersets, one infers that $A \cap B \in \mathcal{U}$.

Exercise 1.1.16. For the “if” direction, suppose that \mathcal{F} is an ultrafilter and \mathcal{F}' is a filter containing \mathcal{F} . Suppose, toward a contradiction, that there is $A \in \mathcal{F}' \setminus \mathcal{F}$. Since \mathcal{F} is an ultrafilter, we have that $S \setminus A \in \mathcal{F} \subseteq \mathcal{F}'$, whence $\emptyset \in \mathcal{F}'$, which is a contradiction. For the “only if” direction, if \mathcal{F} is a maximal filter and $A \subseteq S$ is such that $A \notin \mathcal{F}$, then by Exercise 1.1.7 and the fact that $\mathcal{F} \cup \{A\}$ does not have the FIP (else it would generate a filter extending \mathcal{F}), we see that $S \setminus A \in \mathcal{F}$.

Exercise 1.1.20. Let $\mathcal{F} := \{A \subseteq \mathbb{N} : \sum_{n \in \mathbb{N} \setminus A} \frac{1}{n} \text{ converges}\}$. Verify that \mathcal{F} is a filter on \mathbb{N} and that any ultrafilter \mathcal{U} on \mathbb{N} extending \mathcal{F} is as desired.

Exercise 1.2.4. For (1), take $C \in \mathcal{U}$ and set $A := (\mathbb{N} \setminus C) \times C$. For (2), it is straightforward to show that both statements are equivalent to the statement that $\{s \in \mathbb{N} : (s, n) \in B\} \in \mathcal{U}$.

Exercise 1.3.8. Both parts follow from Exercise 1.3.3.

Exercise 1.4.6. If \mathcal{U} is an ultrafilter on κ , show that its isomorphism class has size at most 2^κ .

Exercise 1.5.2. The upper bound is obvious. For the lower bound, one must show that no countable collection \mathcal{B} of subsets of \mathbb{N} can be a base for an ultrafilter. To see this, enumerate $\mathcal{B} = (B_n)_{n \in \mathbb{N}}$ and recursively define a set $C \subseteq \mathbb{N}$ such that $B_n \not\subseteq C$ and $B_n \not\subseteq \mathbb{N} \setminus C$ for any $n \in \mathbb{N}$.

Exercise 1.6.2. By assumption, we have that $[g_1 \circ f]_{\mathcal{U}} = [g_2 \circ f]_{\mathcal{U}}$, that is, $g_1(f(s)) = g_2(f(s))$ for \mathcal{U} -almost all $s \in S$. Since the range of f belongs to \mathcal{V} , we see that $g_1(t) = g_2(t)$ for \mathcal{V} -almost all $t \in T$, that is, $[g_1]_{\mathcal{V}} = [g_2]_{\mathcal{V}}$.

Chapter 2. Arrow's theorem on fair voting

Exercise 2.1.4. Suppose that $\pi(v)(i) > \pi(v)(j)$ for all $v \in V$. Fix $v \in V$ and define a new state of the election π' such that $\pi'(w) = \pi(v)$ for all $w \in V$. Apply (IA) to π and π' and (U) to π' to conclude that $f(\pi)(i) > f(\pi)(j)$.

Exercise 2.2.8. Fix a state of the election π . By Theorem 2.2.7, there is a unique $\sigma \in S_n$ such that $\{v \in V : \pi(v) = \sigma\} \in \mathcal{U}_f$. By definition, this set is decisive for f , whence $f(\pi) = \sigma$. By definition, it follows that $f = f_{\mathcal{U}_f}$.

Chapter 3. Ultrafilters in topology

Exercise 3.1.4. Fix $\epsilon > 0$. Then the set $\{n \in \mathbb{N} : d(x_n, x) < \epsilon\}$ is cofinite, and thus belongs to \mathcal{U} since \mathcal{U} is nonprincipal.

Exercise 3.1.5(3). First note that the assumption implies that $y_n \neq 0$ for a \mathcal{U} -large set of n , whence $\frac{x_n}{y_n}$ is defined for a \mathcal{U} -large set of n , which is good enough for considering the ultralimit of this latter sequence. Fix $0 < \delta < |y|$ and $M \in \mathbb{N}$ such that $|x|, |y| \leq M$. Then for \mathcal{U} -almost all $n \in \mathbb{N}$, we have

$$\left| \frac{x_n}{y_n} - \frac{x}{y} \right| \leq \frac{|x_n - x||y| + |x||y_n - y|}{|y||y_n|} \leq \frac{M}{(|y| - \delta)|y|} (|x_n - x| + |y_n - y|).$$

Since the \mathcal{U} -ultralimit of the right-hand side is 0, the result follows.

Exercise 3.1.12. Use Theorem 3.1.6 and Theorem 3.1.9.

Exercise 3.2.4. Unabusing the notation might help, that is, show that $\lim_{x, \mathcal{U}} \iota(x) = \mathcal{U}$.

Exercise 3.3.13. For the forward direction, show that $\mathcal{U} \cup \{C\}$ is once again a z -filter on X , and thus $C \in \mathcal{U}$ by maximality of \mathcal{U} . For the backward direction, suppose that \mathcal{F} is a z -filter on X properly extending \mathcal{U} . Take $C \in \mathcal{F} \setminus \mathcal{U}$ and obtain a contradiction by showing that $C \cap Z \neq \emptyset$ for all $Z \in \mathcal{U}$.

Exercise 3.4.13. First show that for any given $b \in \mathbb{B}$, the evaluation map $x \mapsto x(b) : 2^{\mathbb{B}} \rightarrow \mathbf{2}$ is continuous. Then use that preimages under continuous maps of points are closed.

Exercise 3.4.34. It is clear that the ultrafilter theorem for Boolean algebras implies the usual ultrafilter theorem. Conversely, suppose that the ultrafilter theorem holds. Note that this assumption yields Tychonoff's theorem for compact Hausdorff spaces without having to assume the axiom of choice. Consequently, we also have that $S(\mathbb{B})$ is compact. Now given a filter \mathcal{F} on \mathbb{B} , the set $\bigcap_{a \in \mathbb{B}} U_a$ is nonempty as it is the intersection of a family of nonempty closed subsets with the FIP. Any \mathcal{U} in this intersection is an ultrafilter on \mathbb{B} extending \mathcal{F} .

Chapter 4. Ramsey theory and combinatorial number theory

Exercise 4.1.2. Suppose, inductively, that a sequence x_1, \dots, x_{d-1} has been constructed in increasing order such that, for all $i = 1, \dots, d-1$, we have

$$x_i \in \bigcap_{1 \leq j < k < d} A_{(x_j, x_k)} \cap \bigcap_{1 \leq j < d} B_{x_j} \cap C.$$

We wish to extend this sequence by a new element x_d so that the extended sequence satisfies the analogous properties. Since \mathcal{U} is nonprincipal, we may choose $x_d > x_{d-1}$ for which

$$x_d \in \bigcap_{1 \leq j < k \leq d-1} A_{(x_j, x_k)} \cap \bigcap_{1 \leq j \leq d-1} B_{x_j} \cap C,$$

continuing the construction as desired. We note that all of the sets in the above display are in \mathcal{U} by the inductive construction.

Exercise 4.1.3. Fix $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. For $j = 0, \dots, k-1$ and $x_1, \dots, x_{k-j-1} \in \mathbb{N}$, we inductively define the sets $A_{(x_1, \dots, x_{k-j-1})}^j$. For each $x_1, \dots, x_{k-1} \in \mathbb{N}$, let

$$A_{(x_1, \dots, x_{k-1})}^0 := \{z \in \mathbb{N} : (x_1, \dots, x_{k-1}, z) \in C_1\}.$$

Without loss of generality, we may assume that

$$(\mathcal{U}x_1) \cdots (\mathcal{U}x_{k-1}) A_{(x_1, \dots, x_{k-1})}^0 \in \mathcal{U}.$$

We inductively define $A_{(x_1, \dots, x_{k-j-1})}^j$ by setting

$$A_{(x_1, \dots, x_{k-j-1})}^j := \{x_{k-j} \in \mathbb{N} : A_{(x_1, \dots, x_{k-j})}^{j-1} \in \mathcal{U}\}.$$

Note that $A^{k-1} \in \mathcal{U}$, so we may take $x_1 \in A^{k-1}$. Now take $x_2 \in A^{k-1} \cap A_{x_1}^{k-2}$ with $x_2 > x_1$. Then take $x_3 \in A^{k-1} \cap A_{x_1}^{k-2} \cap A_{x_2}^{k-2} \cap A_{(x_1, x_2)}^{k-3}$ with $x_3 > x_2$. Continue in this manner, making the inductive step precise as in the solution of Exercise 4.1.2.

Exercise 4.3.4. It is clear that if A is thick, then $\text{BD}(A) = 1$. Conversely, suppose that A is not thick, say there is $m \in \mathbb{N}$ such that A does not contain any intervals of length m . Let $N \in \mathbb{N}$ be a sufficiently large natural number and let I be an interval of length N such that $\text{BD}(A)$ is approximately equal to $\frac{|A \cap I|}{N}$. Divide I into N/m intervals $I_1, \dots, I_{N/m}$ of length m with a small left-over interval. Now each interval I_k contains an element of $\mathbb{N} \setminus A$ by assumption. Consequently, $\text{BD}(A)$ is approximately less than or equal to $\frac{N - \frac{N}{m}}{N} = 1 - \frac{1}{m}$.

Exercise 4.3.6. Let $(I_n)_{n \in \mathbb{N}}$ be a sequence of intervals with $|I_n| = n$ such that $\delta(A \cup B, n) = \frac{|(A \cup B) \cap I_n|}{n}$. Since $|(A \cup B) \cap I_n| \leq |A \cap I_n| + |B \cap I_n|$, we see that $\delta(A \cup B, n) \leq \frac{|A \cap I_n|}{n} + \frac{|B \cap I_n|}{n} \leq \delta(A, n) + \delta(B, n)$.

Exercise 4.3.13. Fix $f \in B(\mathbb{Z})$ and $k \in \mathbb{Z}$. Let $M \in \mathbb{N}$ be such that $|f(x)| \leq M$ for all $x \in \mathbb{Z}$. Next note that, writing $I_n = [a_n, b_n]$ (and assuming, without loss of generality, that $a_n \geq k$), we have

$$\left| \sum_{x \in I_n} f(x) - \sum_{x \in I_n} (k \cdot f)(x) \right| \leq \sum_{i=1}^k |f(a-i)| + \sum_{i=0}^{k-1} |f(b-i)| \leq 2Mk.$$

By dividing both sides by $|I_n|$ and taking ultralimits (and using the fact that $\lim_{\mathcal{U}} |I_n| = \infty$), we obtain the desired result.

Exercise 4.4.8. First show that if I is a sufficiently long interval, then I contains a Δ_r -set. Thus, if A is not syndetic, then there are sufficiently long intervals contained in $\mathbb{N} \setminus A$, and consequently there will be a Δ_r -set contained in the complement of A , whence A is not a Δ_r^* -set.

Chapter 5. Foundational concerns

Exercise 5.1.12. Fix a filter \mathcal{F} on X . Note that

$$F := \{f \in 2^{2^X} : f \text{ is a filter on } X \text{ extending } \mathcal{F}\}$$

is closed in 2^{2^X} . Also, for any $A \subseteq X$, we have that

$$D_A := \{f \in 2^{2^X} : f(A) = 1 \text{ or } f(X \setminus A) = 1\}$$

is closed in 2^{2^X} . Notice that $\{F\} \cup \{D_A : A \subseteq X\}$ has the FIP. Hence, by the compactness of 2^{2^X} , $F \cap \bigcap_{A \subseteq X} D_A$ is nonempty; any element of this intersection is an ultrafilter on X extending \mathcal{F} .

Exercise 5.2.14. Take $X \in \mathcal{U}$ and $Y \in 2^{\mathbb{N}}$ which are eventually equal. Take $k \in \mathbb{N}$ so that, for every $n \geq k$, $n \in X \iff n \in Y$. Then Y contains $X \cap [k, \omega)$, which belongs to \mathcal{U} .

Exercise 5.3.1. Let $(\sigma_\alpha)_{\alpha < 2^\omega}$ be an enumeration of the strategies for player I and let $(\tau_\alpha)_{\alpha < 2^\omega}$ be an enumeration of the strategies for player II. Inductively build sequences $(x_\alpha)_{\alpha < 2^\omega}$ and $(y_\alpha)_{\alpha < 2^\omega}$ of sequences from $\mathbb{N}^{\mathbb{N}}$ such that:

- $\{x_\alpha : \alpha < 2^\omega\} \cap \{y_\alpha : \alpha < 2^\omega\} = \emptyset$;
- for every $\alpha < 2^\omega$, y_α is a play of the game where player I plays according to σ_α ;
- for every $\alpha < 2^\omega$, x_α is a play of the game where player II plays according to τ_α .

If this can be done, it follows that $\{x_\alpha : \alpha < 2^\omega\}$ is not determined. To see that this construction can be carried out, assume $(x_\alpha)_{\alpha < \gamma}$ and $(y_\alpha)_{\alpha < \gamma}$ are constructed as above for some $\gamma < 2^\omega$. Choose y_γ such that y_γ is a play of a game where player I plays according to σ_γ (of which there are 2^{\aleph_0}), and $y_\gamma \notin \{x_\alpha : \alpha < \gamma\}$. Similarly, choose x_γ such that x_γ is a play of a game where player II plays according to τ_γ , and $x_\gamma \notin \{y_\alpha : \alpha \leq \gamma\}$.

Exercise 5.3.7. For any $a \in \mathbb{N}^{\mathbb{N}}$ which is not strictly increasing, let $I(a)$ be the restriction of a to the largest initial segment for which a is strictly increasing. Extend the notion of A_a to finite sequences in the obvious way. Define $f : \mathbb{N}^{\mathbb{N}} \rightarrow 2^{\mathbb{N}}$ by $f(a) = A_a$ if a is strictly increasing, $f(a) = (\mathbb{N} ([0, \max A_{I(a)}] A_{I(a)}))$ if a is not strictly increasing and odd breaks, and $f(a) = A_{I(a)}$ if a is not strictly increasing and even breaks. Then $f^{-1}(\mathcal{U}) = D_{\mathcal{U}}$ and f is continuous.

Exercise 5.4.2. Assume first that \mathcal{U} is selective and let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a function. If $\text{range}(f)$ is finite, then f is constant on a set in \mathcal{U} . Otherwise, we write $\mathbb{N} = \bigsqcup_{a \in \text{range}(f)} f^{-1}(\{a\})$. If f is not constant on a set in \mathcal{U} , then $f^{-1}(\{a\}) \notin \mathcal{U}$ for any $a \in \text{range}(f)$. Since \mathcal{U} is selective, there is $B \in \mathcal{U}$ such that $|B \cap f^{-1}(\{a\})| = 1$ for all $a \in \text{range}(f)$, whence f is injective on B .

For the converse, let $(A_i)_{i \in \mathbb{N}}$ be a partition of \mathbb{N} with $A_i \notin \mathcal{U}$ for all $i \in \mathbb{N}$. Define $f : \mathbb{N} \rightarrow \mathbb{N}$ by $f(n) = i$ iff $n \in A_i$. Note f is not constant on a set in \mathcal{U} as $f^{-1}(\{i\}) = A_i \notin \mathcal{U}$. Thus, there exists $B \in \mathcal{U}$ such that f is injective on B . Hence $|B \cap A_i| \leq 1$ for all $i < \omega$, as desired.

Exercise 5.4.7. If $\text{cf}(\alpha) = 1$, then the last member of the sequence is an infinite pseudo-intersection. Otherwise, $\text{cf}(\alpha) = \omega$ and there exists a sequence $(\alpha_n)_{n < \omega}$ of ordinals such that $\sup_{n < \omega} \alpha_n = \alpha$. Note that, for any finite $I \subseteq \alpha$, $\bigcap_{i \in I} B_i$ is infinite since $B_{\max I} \setminus \bigcap_{i \in I} B_i = \bigcup_{i \in I} B_{\max I} \setminus B_i$ is finite as $B_{\max I} \subseteq^* B_i$ for all $i \in I$.

We define $(x_n)_{n < \omega}$ inductively as follows. Choose $x_0 \in B_{\alpha_0}$. If $(x_n)_{n \leq k}$ has been constructed, choose $x_{k+1} \in (\bigcap_{n \leq k+1} B_{\alpha_n}) \setminus \{x_n : n \leq k\}$. Then $X := \{x_n : n < \omega\}$ is an infinite subset of \mathbb{N} . For any $\gamma < \alpha$, let $n < \omega$ be

such that $\gamma < \alpha_n$, and observe $X \subseteq^* B_{\alpha_n} \subseteq^* B_\gamma$, so $X \subseteq^* B_\gamma$. Consequently, X is an infinite pseudo-intersection of $(B_\beta)_{\beta < \alpha}$.

Exercise 5.4.13. Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be such that $f(\mathcal{U}) = \mathcal{V}$ and let $\mathbb{N} = \bigsqcup_{n < \omega} A_n$ be a partition with $A_n \notin \mathcal{V}$ for each $n < \omega$. Set $B_n := f^{-1}(A_n)$ and observe that $B_n \notin \mathcal{U}$ for any $n < \omega$. Since \mathcal{U} is weakly selective, there is $B \in \mathcal{U}$ such that $B \cap B_n$ is finite for all $n < \omega$. Set $A := f(B)$, so $A \in \mathcal{V}$ as $f^{-1}(A) \supseteq B \in \mathcal{U}$. Now, if $z \in A \cap A_n$, then there exists $b \in B$ such that $f(b) = z \in A_n$, so $b \in B \cap f^{-1}(A_n) = B \cap B_n$. Hence $A \cap A_n \subseteq f(B \cap B_n)$ is finite.

Exercise 5.4.16. For (1), for any $i < \omega$ and $n < \omega$, we see $X_n \setminus [i, \omega) \subseteq [0, i)$ is finite, so $X_n \subseteq^* [i, \omega)$, whence $[i, \omega) \in \mathcal{F}$. In particular, $\mathbb{N} \in \mathcal{F}$. Also, $\emptyset \notin \mathcal{F}$ since the X_n 's are infinite. If $A, B \in \mathcal{F}$, then for all but finitely many $n < \omega$, we see $X_n \setminus (A \cap B) = (X_n \setminus A) \cup (X_n \setminus B)$ is finite, so $X_n \subseteq^* A \cap B$. Hence $A \cap B \in \mathcal{F}$. Finally, if $A \in \mathcal{F}$ and $A \subseteq B \subseteq \mathbb{N}$, then for all but finitely many $n < \omega$, we see $X_n \setminus B \subseteq X_n \setminus A$ is finite, so $X_n \subseteq^* B$. Hence $B \in \mathcal{F}$. Thus, \mathcal{F} is a filter containing the Frèchet filter.

For (2), let \mathcal{U} be an ultrafilter extending \mathcal{F} . In particular, \mathcal{U} is non-principal. Consider the partition $\mathbb{N} = \bigsqcup_{n < \omega} X_n$, and observe that, for every $n < \omega$, $X_n \notin \mathcal{U}$ since $\mathbb{N} \setminus X_n = \bigcup_{k \neq n} X_k \in \mathcal{F}$. For any $Z \subseteq \mathbb{N}$, if $Z \cap X_n$ is finite for all $n < \omega$, then for any $n < \omega$, $X_n \setminus (\mathbb{N} \setminus Z) = X_n \cap Z$ is finite, so $X_n \subseteq^* \mathbb{N} \setminus Z$. Hence $\mathbb{N} \setminus Z \in \mathcal{F} \subseteq \mathcal{U}$. Thus, \mathcal{U} is not a P-point.

Chapter 6. Classical ultraproducts

Exercise 6.3.6. It is clear that $N_{\mathcal{F}}$ contains the identity. Suppose that $a, b \in N_{\mathcal{F}}$. Since filters are closed under intersections, we have that $a(i) = b(i) = e$ for \mathcal{F} -many $i \in I$, whence $a(i)b(i) = e$ for \mathcal{F} -many $i \in I$, and thus $ab \in N_{\mathcal{F}}$. That $N_{\mathcal{F}}$ is closed under inversion and conjugation is straightforward, so one has that $N_{\mathcal{F}}$ is a normal subgroup of $\prod_{i \in I} G_i$.

Now let $\phi : \prod_{i \in I} G_i / N_{\mathcal{F}} \rightarrow \prod_{\mathcal{F}} G_i$ be given by $\phi(gN_{\mathcal{F}}) := [g]_{\mathcal{F}}$. Note indeed that ϕ is well defined: if $gN_{\mathcal{F}} = hN_{\mathcal{F}}$, then $h^{-1}g \in N_{\mathcal{F}}$, that is, $h(i)^{-1}g(i) = e_{G_i}$ for \mathcal{F} -many $i \in I$, that is, $g(i) = h(i)$ for \mathcal{F} -many $i \in I$, and thus $[g]_{\mathcal{F}} = [h]_{\mathcal{F}}$. Reversing the previous reasoning shows that ϕ is injective. It is clear that ϕ is surjective, whence ϕ is a bijection. It follows immediately from the definitions that ϕ is an isomorphism of structures.

Exercise 6.4.9. For each $\sigma \in T$, let $X_\sigma = \{\Delta \in I : \sigma \in \Delta\}$. Given finitely many \mathcal{L} -sentences $\sigma_1, \dots, \sigma_n \in T$, we have that $\{\sigma_1, \dots, \sigma_n\} \in \bigcap_{i=1}^n X_{\sigma_i}$, whence the collection $(X_\sigma)_{\sigma \in T}$ has the FIP. Let \mathcal{U} be an ultrafilter on I containing each X_σ . We claim that $\prod_{\mathcal{U}} \mathcal{M}_\Delta \models T$. Indeed, given $\sigma \in T$, $\mathcal{M}_\Delta \models \sigma$ whenever $\sigma \in \Delta$, that is, whenever $\Delta \in X_\sigma$; since $X_\sigma \in \mathcal{U}$, we have that $\prod_{\mathcal{U}} \mathcal{M}_\Delta \models \sigma$ by Łós's theorem.

Exercise 6.5.5. The exercise hinges on the fact that given any proper filter \mathcal{F}' extending \mathcal{F} and any $A \subseteq I$, either $\mathcal{F}' \cup \{A\}$ generates a proper filter or $\mathcal{F}' \cup \{I \setminus A\}$ generates a proper filter.

Exercise 6.6.2. It is clear that d is a bijection when \mathcal{U} is principal. Conversely, fix $a \in M^{\mathbb{N}}$ such that $a(m) \neq a(n)$ for all distinct $m, n \in \mathbb{N}$. Take $b \in M$ such that $d(b) = [a]_{\mathcal{U}}$, whence $b = a_n$ for \mathcal{U} -almost all n . It follows that there is a unique $n \in \mathbb{N}$ such that $\{n\} \in \mathcal{U}$, whence \mathcal{U} is principal.

Exercise 6.6.7. The backward direction is clear. For the forward direction, suppose that \mathcal{U} is countably incomplete and let $(F_n)_{n \in \mathbb{N}}$ be a countable collection of elements of \mathcal{U} such that $\bigcap_{n \in \mathbb{N}} F_n = \emptyset$. Set $E_0 := I$ and for each $n \in \mathbb{N}$, set $E_{n+1} := F_0 \cap \dots \cap F_n$. It is clear that the sequence $(E_n)_{n \in \mathbb{N}}$ is as desired.

Exercise 6.6.11. First suppose that \mathcal{U} is countably incomplete. Following the hint, take $(E_n)_{n \in \mathbb{N}}$ from \mathcal{U} such that $I = E_0 \supseteq E_1 \supseteq E_2 \supseteq \dots$ and such that $\bigcap_{n \in \mathbb{N}} E_n = \emptyset$. Define $f : I \rightarrow \mathbb{N}$ by setting $f(x) =$ the maximal $n \in \mathbb{N}$ such that $x \in E_n$. We must show that $f(\mathcal{U})$ is nonprincipal. If, toward a contradiction, $f(\mathcal{U}) = \mathcal{U}_n$, then $f(x) = n$ for \mathcal{U} -many $x \in I$. However, $f(x) \geq n + 1$ for $x \in E_{n+1}$, whence $f(x) \geq n + 1$ for \mathcal{U} -many $x \in I$, yielding a contradiction. Conversely, suppose that there is a nonprincipal ultrafilter \mathcal{V} on \mathbb{N} such that $\mathcal{V} \leq_{RK} \mathcal{U}$. By Lemma 6.6.9, \mathcal{V} is countably incomplete. By Exercise 6.6.8, it follows that $f(\mathcal{U})$ is also countably incomplete.

Exercise 6.7.3. We only prove the first part of (3). While one can do this by hand (by induction on complexity of formulae), the shortest proof uses the Tarski-Vaught test for elementary substructures. To use this test, we suppose that $\varphi(x_1, \dots, x_n, y)$ is a \mathcal{L} -formula and $g_1, \dots, g_n : M \rightarrow M$ are functions such that $\mathcal{M}^{\mathcal{U}} \models \exists y \varphi([g_1 \circ f]_{\mathcal{U}}, \dots, [g_n \circ f]_{\mathcal{U}}, y)$. We must show that such a witness can be found in $\mathcal{N}[f]$. Take $X \in \mathcal{U}$ such that, for all $i \in X$, we have that $\mathcal{M} \models \exists y (g_1(f(i)), \dots, g_n(f(i)), y)$ as witnessed by $a_i \in M$. We may of course assume that if $f(i) = f(j)$ for some $i, j \in X$, then $a_i = a_j$. In this case, this allows us to define, for $i \in X$, define $h : M \rightarrow M$ so that $h(f(i)) = a_i$. (One defines h on $M \setminus f(X)$ in an arbitrary fashion.) It follows that $\mathcal{M}^{\mathcal{U}} \models \varphi([g_1 \circ f]_{\mathcal{U}}, \dots, [g_n \circ f]_{\mathcal{U}}, [h \circ f]_{\mathcal{U}})$, whence the Tarski-Vaught test is successful and $\mathcal{N}[f] \preceq \mathcal{M}^{\mathcal{U}}$.

Exercise 6.10.3. One treats function symbols just as in the case of groups. For example, if F is a unary function symbol, then for all $i \in I$ and $a \in M_i$, if $g_i : M_i \hookrightarrow M$ is the canonical inclusion, we define $F^{\mathcal{M}}(g_i(a)) := g_i(F^{M_i}(a))$. Note that this is well defined (that is, independent of representative) by the fact that the system of structures is directed. The interpretation of predicates is a little more interesting: if there is *some* $i \in I$ and $a \in M_i$ such that $a \in P^{M_i}$, then we declare that $g_i(a) \in P^{\mathcal{M}}$.

Exercise 6.10.4. The proof is exactly like the proof of Theorem 6.10.1, noting that the maps in the abstract case are in fact homomorphisms of \mathcal{L} -structures.

Exercise 6.11.9. Let σ be a sentence such that $\prod_{\mathcal{F}} \mathcal{M}_i \models \sigma$. Suppose that σ is determined by $(\gamma; \psi_1, \dots, \psi_m)$ and let $X_j := \{i \in I : \mathcal{M}_i \models \psi_j\}$, whence $\mathcal{P}(I)/\mathcal{F} \models \gamma([X_1]_{\mathcal{F}}, \dots, [X_m]_{\mathcal{F}})$. Since $\mathcal{M}_i \equiv \mathcal{N}_i$ for each $i \in I$, we have that $X_j = X'_j := \{i \in I : \mathcal{N}_i \models \psi_j\}$, whence $\prod_{\mathcal{F}} \mathcal{N}_i \models \sigma$ as well.

Chapter 7. Applications to geometry, commutative algebra, and number theory

Exercise 7.1.4. This follows from Loś's theorem and the fact that being an algebraically closed field is a first-order property. To see this latter fact, note that a field K is algebraically closed if and only if $K \models \sigma_n$ for all $n \geq 1$, where σ_n is the sentence

$$\forall a_0 \cdots \forall a_n (a_n \neq 0 \rightarrow \exists x (a_0 + a_1 x + \cdots + a_n x^n = 0)).$$

Exercise 7.2.16. Suppose the statement of the exercise fails for some n and d . For each $t \in \mathbb{N}$, let K_t be a field, and $f_1(t), \dots, f_m(t) \in K_t[X_1, \dots, X_n]$ all have degree at most d satisfying:

- $gh \in (f_1(t), \dots, f_m(t))$ implies $g \in (f_1(t), \dots, f_m(t))$ or $h \in (f_1(t), \dots, f_m(t))$ for all $g, h \in K_t[X_1, \dots, X_n]$ of degree at most t , and
- $(f_1(t), \dots, f_m(t))$ is a proper ideal of $K_t[X_1, \dots, X_n]$ that is not prime.

Fix a nonprincipal ultrafilter \mathcal{U} on \mathbb{N} . For each $t \in \mathbb{N}$, take $g(t), h(t) \in K_t[X_1, \dots, X_n]$ such that $g(t)h(t) \in (f_1(t), \dots, f_m(t))$ but $g(t), h(t) \notin (f_1(t), \dots, f_m(t))$. It follows that $g, h \in \prod_{\mathcal{U}} K_t[X_1, \dots, X_n]$ are such that $gh \in (f_1, \dots, f_m)$ but $g, h \notin (f_1, \dots, f_m)$. Thus, f_1, \dots, f_m do not generate a prime ideal in $\prod_{\mathcal{U}} K_t[X_1, \dots, X_n]$. It is also clear that (f_1, \dots, f_m) generate a proper ideal of $\prod_{\mathcal{U}} K_t[X_1, \dots, X_n]$ (for otherwise $(f_1(t), \dots, f_m(t))$ generate $K_t[X_1, \dots, X_n]$ for \mathcal{U} -almost all $t \in \mathbb{N}$).

However, for each $i = 1, \dots, m$, we have that $f_i \in \prod_{\mathcal{U}} K_t[X_1, \dots, X_n]$ has degree $\leq d$, and thus belongs to $(\prod_{\mathcal{U}} K_t)[X_1, \dots, X_n]$. By Theorem 7.2.15, f_1, \dots, f_m do not generate a prime ideal in $(\prod_{\mathcal{U}} K_t)[X_1, \dots, X_n]$ either. Consequently, there are $a, b \in (\prod_{\mathcal{U}} K_t)[X_1, \dots, X_n]$ such that $ab \in (f_1, \dots, f_m)$ but $a, b \notin (f_1, \dots, f_m)$. Since both $a(t)$ and $b(t)$ have degree at most t for \mathcal{U} -almost all $t \in \mathbb{N}$, this is a contradiction to our choices.

Exercise 7.3.5. For (1), first suppose that R is a local ring with unique maximal ideal \mathfrak{m} . Suppose that $x, y \in R$ are noninvertible elements. Then the principal ideals generated by x and y must be proper and thus contained

in \mathfrak{m} . Consequently, $x+y \in \mathfrak{m}$, whence $x+y$ is also noninvertible. Conversely, if \mathfrak{m}_1 and \mathfrak{m}_2 are distinct maximal ideals of R , then taking $x \in \mathfrak{m}_2 \setminus \mathfrak{m}_1$, we see that the ideal generated by \mathfrak{m}_1 and x is all of R , whence there is $y \in \mathfrak{m}_1$ and $a \in R$ such that $ax + y = 1$. Thus, ax and y are noninvertible elements of R whose sum is invertible. (2) now follows from (1). For (3), note that, by (1), the unique maximal ideal in a local ring consists precisely of the noninvertible elements. For (4), the fact that R is a local ring follows directly from (2) and Łoś's theorem and the fact that $\mathfrak{m} \cong \prod_{\mathcal{U}} \mathfrak{m}_i$ follows directly from (3). It is readily verified that the obvious map $(\prod_{\mathcal{U}} R_i) / (\prod_{\mathcal{U}} \mathfrak{m}_i) \rightarrow \prod_{\mathcal{U}} (R_i / \mathfrak{m}_i)$ is an isomorphism.

Chapter 8. Ultraproducts and saturation

Exercise 8.1.14. The forward direction follows from Exercise 8.1.4 and Proposition 8.1.11. To prove the converse, suppose that \mathcal{M}_A is κ^+ -universal for every $A \subseteq M$ with $|A| < \kappa$. Let $\Sigma(x)$ be a finitely satisfiable set of \mathcal{L}_A -formulae; we seek a realization of Σ in M . Let \mathcal{N} be a model of $\text{Th}(\mathcal{M}_A)$ of cardinality $\leq \kappa$ containing a realization b of $\Sigma(x)$. By assumption, there is an elementary embedding $i : \mathcal{N}_A \hookrightarrow \mathcal{M}_A$. It follows that $i(b)$ is a realization of $\Sigma(x)$ in M .

Exercise 8.2.4. The backward direction follows immediately from Łoś theorem. To prove the forward direction, let $\{n_0, n_1, \dots\}$ be an enumeration of $N \setminus M$ and let $\Sigma(x_0, x_1, \dots)$ be the collection of quantifier-free \mathcal{L}_M -formulae $\varphi(x_0, \dots, x_k)$ (as k varies over ω) such that $\mathcal{N} \models \varphi(n_0, \dots, n_k)$. It suffices to show that there is a realization (a_0, a_1, \dots) of $\Sigma(x)$ in $\mathcal{M}^{\mathcal{U}}$, for then the map $n_i \mapsto a_i$ is the desired embedding. Since $\mathcal{M}^{\mathcal{U}}$ is \aleph_1 -saturated, it suffices to show that Σ is finitely satisfiable in \mathcal{M}_M , which follows from the fact that \mathcal{M} is e.c. in \mathcal{N} .

Exercise 8.3.6. First suppose that \mathcal{U} is κ -regular. Let E be a κ -regularizing set for \mathcal{U} and enumerate E as $(A_\alpha)_{\alpha < \kappa}$. For each $i \in I$, set $f(i) := \{\alpha_1^i, \dots, \alpha_m^i\}$, where $A_{\alpha_1^i}, \dots, A_{\alpha_m^i}$ denote the finitely many elements of \mathcal{U} containing i . (If i does not belong to any set in E , set $f(i) = \emptyset$.) This f is as desired.

For the converse, fix an f as in the statement of the exercise. For $\alpha < \kappa$, set $A_\alpha := \{i \in I : \alpha \in f(i)\}$. One readily checks that $\{A_\alpha : \alpha < \kappa\}$ has cardinality κ (although it may contain some repetitions) and is a κ -regularizing set for \mathcal{U} .

Exercise 8.3.7. Suppose that \mathcal{U} is κ -regular and $\mathcal{U} \leq_{RK} \mathcal{V}$. Suppose that the index sets of \mathcal{U} and \mathcal{V} are I and J , respectively. Take $g : J \rightarrow I$ such $g(\mathcal{V}) = \mathcal{U}$ and take $f : I \rightarrow \mathcal{P}_f(\kappa)$ as in Exercise 8.3.6 witnessing that \mathcal{U} is κ -regular. Exercise 8.3.6, now applied to $f \circ g$, shows that \mathcal{V} is also κ -regular.

Exercise 8.3.14. A standard compactness argument shows that if $\mathcal{M} \models \text{Th}_{\forall}(\mathcal{N})$, then there is $\mathcal{N}' \equiv \mathcal{N}$ such that $\mathcal{M} \subseteq \mathcal{N}'$. By downward Löwenheim-Skolem, we may assume that $|N'| \leq \kappa = \max(|M|, |\mathcal{L}|)$. If \mathcal{U} is κ -regular, then $\mathcal{N}^{\mathcal{U}}$ is κ^+ -universal, whence \mathcal{N}' embeds into $\mathcal{N}^{\mathcal{U}}$ and (2) follows. (2) implies (3) is trivial and (3) implies (1) follows from Łoś's theorem.

Exercise 8.4.1. Fix a multiplicative concern function $C : \mathcal{P}_f(\Sigma) \rightarrow \mathcal{U}$. First suppose that C is locally finite and fix $i \in I$. Let $u_1, \dots, u_k \in \mathcal{P}_f(\Sigma)$ enumerate those u such that $i \in C(u)$. Then $i \in C(u_1) \cap \dots \cap C(u_k) = C(u_1 \cup \dots \cup u_k)$. Set $C_i := u_1 \cup \dots \cup u_k$. To see that C_i is as desired, fix $u \in \mathcal{P}_f(\Sigma)$. If $i \in C(u)$, then $u = u_j$ for some j and hence $u \subseteq C_i$. Conversely, if $u \subseteq C_i$, then $C(C_i) \subseteq C(u)$ and hence $i \in C(u)$.

Next suppose that there is $C_i \in \mathcal{P}_f(\Sigma)$ with the stated properties. Fix $i \in I$ and suppose that $u \in \mathcal{P}_f(\Sigma)$ is such that $i \in C(u)$. Then $u \subseteq C_i$. Since C_i is finite, this leaves finitely many possibilities for u .

Exercise 8.4.19. Suppose that \mathcal{U} is a nonprincipal ultrafilter on ω and let $f : \mathcal{P}_f(\omega) \rightarrow \mathcal{U}$ be antimonotonic. Given $u \in \mathcal{P}_f(\omega)$, set $n_u \in \omega$ to be the minimal $n \in \omega$ such that $u \subseteq n$. (Recall that a natural number is defined as the set of natural numbers below it.) Define $g : \mathcal{P}_f(\omega) \rightarrow \mathcal{U}$ to be $g(u) := f(n_u)$. Since $u \subseteq n_u$ and f is antimonotonic, we have that $g(u) = f(n_u) \subseteq f(u)$, so g refines f . To see that g is multiplicative, note that $g(u \cup v) = f(n_{u \cup v}) = f(\max(n_u, n_v))$. Since f is antimonotonic, $f(\max(n_u, n_v)) = f(n_u) \cap f(n_v)$, whence we have $g(u \cup v) = g(u) \cap g(v)$.

Exercise 8.4.20. Suppose that \mathcal{V} is κ -regular. First suppose that $\mathcal{U} \times \mathcal{V}$ is κ^+ -good. To show that \mathcal{V} is κ^+ -good, fix a language \mathcal{L} with $|\mathcal{L}| \leq \kappa$ and an \mathcal{L} -structure \mathcal{M} ; it suffices to show that $\mathcal{M}^{\mathcal{V}}$ is κ^+ -saturated. Since \mathcal{V} is κ -regular, $\mathcal{M}^{\mathcal{V}}$ is κ^+ -saturated if and only if $(\mathcal{M}^{\mathcal{U}})^{\mathcal{V}} = \mathcal{M}^{\mathcal{U} \times \mathcal{V}}$ is κ^+ -saturated; this latter fact holds since $\mathcal{U} \times \mathcal{V}$ is assumed to be κ^+ -good. The converse can be argued in the same fashion.

Exercise 8.4.22. Let \mathcal{U} be a regular ultrafilter on κ that is not good (which exists since κ is uncountable) and let \mathcal{V} be a good ultrafilter on κ . By Exercise 8.4.20, $\mathcal{U} \times \mathcal{V}$ is good while $\mathcal{V} \times \mathcal{U}$ is not good. Let \mathcal{M} be as in Theorem 8.4.16. Then $\mathcal{M}^{\mathcal{U} \times \mathcal{V}}$ is κ^+ -saturated while $\mathcal{M}^{\mathcal{V} \times \mathcal{U}}$ is not. Thus $\mathcal{M}^{\mathcal{U} \times \mathcal{V}} \not\cong \mathcal{M}^{\mathcal{V} \times \mathcal{U}}$ and so $\mathcal{U} \times \mathcal{V} \not\equiv_{RK} \mathcal{V} \times \mathcal{U}$ by Theorem 6.7.1.

Chapter 9. Nonstandard analysis

Exercise 9.2.12. To see that there are no largest or smallest infinite galaxies, take even $N \in \mathbb{N}^* \setminus \mathbb{N}$ and note that $\gamma(\frac{N}{2}) < \gamma(N) < \gamma(2N)$. For the other statement, if $\gamma(M) < \gamma(N)$ with both M and N even, note that $\gamma(M) < \gamma(\frac{M+N}{2}) < \gamma(N)$.

Exercise 9.3.10. Define the function $S_f : \mathbb{N} \rightarrow \mathbb{R}$ by $S_f(n) = \sum_{i=0}^{n-1} f(\frac{1}{n}) \cdot \frac{1}{n}$, so $\int_0^1 f(x)dx = \lim_{n \rightarrow \infty} S_f(n)$. Now the right-hand side of the exercise is merely suggestive notation for $S_f(N)$ for infinite $N \in \mathbb{N}^*$.

Exercise 9.5.12. This follows essentially from Exercise 9.5.10 and the definition of iterated ultrapower.

Exercise 9.5.14. The forward direction follows immediately from Exercise 9.5.12 and Corollary 6.9.8. The reverse direction follows from the fact that limit ultrapowers are elementary extensions of the original structures.

Exercise 9.5.24. Show that the map $\tilde{\sigma} : \mathcal{M}^{\mathcal{U}|\mathcal{F}_Y} \rightarrow \mathcal{M}^{\mathcal{U}|\mathcal{F}_X}$ given by $\tilde{\sigma}([f]_Y) = [f \circ \sigma]_X$ is an induced embedding. Conclude using Example 9.5.22.

Exercise 9.6.11. Show that nonempty internal subsets of \mathbb{R}^* that are bounded above have least upper bounds.

Exercise 9.6.20. For (1), note that $A = \{x \in X^* : (\exists y \in X^*) (x, y) \in \Gamma(f)\}$; by the internal definition principle, A is thus internal. The proof for why $\text{range}(f)$ is internal is similar. (2) follows from an application of the overflow principle (Exercise 9.6.12).

Exercise 9.7.7. First suppose that T is an S -subsemigroup. Given $\alpha, \beta \in T$, take $\gamma \in T$ such that $\gamma \sim \alpha + \beta^*$. It follows that

$$\pi(\alpha) \oplus \pi(\beta) = \mathcal{U}_\alpha \oplus \mathcal{U}_\beta = \mathcal{U}_{\alpha+\beta^*} = \mathcal{U}_\gamma = \pi(\gamma),$$

i.e., $\pi(T)$ is closed under \oplus , so it is a subsemigroup of $\beta\mathbb{N}$. Reversing the above line of reasoning establishes the converse.

Exercise 9.7.10. First suppose that α is idempotent and take $A \subseteq \mathbb{N}$ such that $\alpha \in A^*$. Since $\alpha \sim \alpha + \alpha^*$, we have that $\alpha + \alpha^* \in A^{**}$. By transfer, it follows that $\alpha \in (A_\alpha)^*$, whence $\alpha \in A^* \cap (A_\alpha)^* = (\alpha A)^*$. For the converse, suppose that α is not idempotent, whence there is some $A \subseteq \mathbb{N}$ such that $\alpha \in A^*$ but $\alpha + \alpha^* \notin A^{**}$. By transfer again, it follows that $\alpha \notin (\alpha A)^*$.

Now suppose that α is idempotent and $s \in A_\alpha$. Then $s + \alpha \in A^*$, whence $\alpha \in (A - s)^*$. By the first part of the exercise, we have $\alpha \in (\alpha(A - s)^*)$, that is, $\alpha + \alpha^* \in (A - s)^{**}$, whence $s + \alpha + \alpha^* \in A^{**}$. It follows that $s + \alpha \in A_\alpha^*$. The case that $s \in {}_\alpha A$ is handled similarly.

Exercise 9.8.2. Suppose that \mathcal{U} is a Hausdorff ultrafilter on I and $\mathcal{V} \leq_{RK} \mathcal{U}$ as witnessed by $\mathcal{V} = h(\mathcal{U})$, where $h : I \rightarrow J$. Suppose that $f, g : J \rightarrow J$ are such that $f(\mathcal{V}) = g(\mathcal{V})$. Then $(f \circ h)(\mathcal{U}) = (g \circ h)(\mathcal{U})$. Let $k : J \rightarrow I$ be any function that is injective on both $(f \circ h)(I)$ and $(g \circ h)(I)$. Since $(k \circ f \circ h)(\mathcal{U}) = (k \circ g \circ h)(\mathcal{U})$ and \mathcal{U} is Hausdorff, we have that $k \circ f \circ h \equiv_{\mathcal{U}} k \circ g \circ h$; since k is injective on $(f \circ h)(I) \cup (g \circ h)(I)$, it follows that $f \equiv_{\mathcal{V}} g$, as desired.

Chapter 10. Limit groups

Exercise 10.2.5. (1) is trivial if the group is \mathbb{Z} ; otherwise, use the well-known fact that centralizers of elements in nonabelian free groups are all cyclic. (2) follows from (1) and the fact that commutative transitivity is a universally axiomatizable property.

Exercise 10.2.6. We prove (1) by contrapositive: if $a, b \in G \setminus \{e\}$ do not commute and $c \in Z(G) \setminus \{e\}$, then the facts that a and c commute and that b and c commute show that G is not commutative transitive. (2) follows from (1) since $G \times H$ is nonabelian (as it contains the nonabelian subgroup $G \times \{e\}$) and has nontrivial center (as the center contains $\{e\} \times H$). $\mathbb{F}_2 \times \mathbb{Z}$ shows that none of the three named classes are closed under direct products.

Exercise 10.2.11. Suppose that H is a maximal abelian subgroup of \mathbb{F}_2 . Then $H = \langle a \rangle$ for some $a \in H$. Suppose $b \in \mathbb{F}_2 \setminus H$. Since H is a maximal abelian subgroup, it follows that $\langle a, b \rangle$ is a free subgroup of \mathbb{F}_2 . If $ba^i b^{-1} = a^j$ for some $i, j \neq 0$, then $ba^i b^{-1} a^{-j} = e$ is a nontrivial relation in $\langle a, b \rangle$, yielding a contradiction.

Exercise 10.3.16. Suppose that G_1 and G_2 are universally free groups. Suppose that H is a finitely generated subgroup of $G_1 * G_2$; it suffices to show that H is fully residually free. Take finitely generated subgroups H_1 and H_2 of G_1 and G_2 , respectively, such that H is a subgroup of $H_1 * H_2$; it suffices to show that $H_1 * H_2$ is fully residually free. This is fairly straightforward using the fact that free products of free groups are free together with the normal form for elements of free products.

Exercise 10.4.11. Fix $k \in \mathbb{N}$; it suffices to show that $\nu((G, S), (H_i, S_i)) \geq k$ for \mathcal{U} -almost all $i \in I$. However, this follows from an immediate application of Łoś's theorem.

Exercise 10.4.14. Suppose that G is a nonabelian limit group generated by a and b . By Theorem 10.4.13, there are free groups G_i with generators a_i and b_i such that (G_i, S_i) converges to (G, S) in \mathcal{G}_2 . Since a and b do not commute, it follows that a_i and b_i do not commute for sufficiently large i , whence a_i and b_i freely generate G_i . It follows that a and b freely generate G as well.

Chapter 11. Metric ultraproducts

Exercise 11.3.3. Fix a sequence $(I_n)_{n \in \mathbb{N}}$ such that $I_n \supseteq I_{n+1}$ for all $n \in \mathbb{N}$ and such that $\bigcap_{n \in \mathbb{N}} I_n = \emptyset$. For each $i \in I$, set $n(i)$ to be the maximal $n \in \mathbb{N}$ such that $i \in I_n$. Fix $[f]_{\mathcal{U}} \in \prod_{\mathcal{U}} M_i$ and define $g \in \prod_{i \in I} X_i$ by letting $g(i)$ be any element of X_i such that $d_i(f(i), g(i)) < \frac{1}{n(i)}$. We claim that $[f]_{\mathcal{U}} = [g]_{\mathcal{U}}$, finishing the proof. Indeed, given any $n \geq 1$ and any

$i \in I_n$, we have that $n(i) \geq n$ and thus $d_i(f(i), g(i)) < \frac{1}{n(i)} \leq \frac{1}{n}$. Since $I_n \in \mathcal{U}$, we have that $d([f]_{\mathcal{U}}, [g]_{\mathcal{U}}) \leq \frac{1}{n}$. Letting n tend to ∞ , we get that $[f]_{\mathcal{U}} = [g]_{\mathcal{U}}$.

Exercise 11.4.7. One proves this result by induction on the complexity of formulae. If $\varphi(\vec{x}) = P(\vec{x})$ for some predicate symbol P , set $\Delta_{\varphi} := \Delta_P$. If $\varphi(\vec{x}) = u(\psi_1(\vec{x}), \dots, \psi_n(\vec{x}))$ for some continuous function $u : [0, 1]^n \rightarrow [0, 1]$, then set $\Delta_{\varphi}(\epsilon) := \min_{i=1, \dots, n} \Delta_{\psi_i}(\frac{\Delta_u(\epsilon)}{2})$. Finally, if $\varphi(\vec{x}) = \sup_y \psi(\vec{x}, y)$, then set $\Delta_{\varphi} := \Delta_{\psi}$.

Exercise 11.4.9. One proves the result by induction on the complexity of φ . As in the classical case, one has, for any term $t(x_1, \dots, x_n)$ and any $[\vec{a}]_{\mathcal{U}} = ([a_1]_{\mathcal{U}}, \dots, [a_n]_{\mathcal{U}}) \in M$, that $t^{\mathcal{M}}(\vec{a}) = [i \mapsto t^{\mathcal{M}_i}(\vec{a}(i))]_{\mathcal{U}}$. Consequently, if $\varphi(\vec{x}) = P(\vec{x})$ for some predicate symbol P , we have that

$$\begin{aligned} \varphi^{\mathcal{M}}([\vec{a}]_{\mathcal{U}}) &= P^{\mathcal{M}}(t_1^{\mathcal{M}}([\vec{a}]_{\mathcal{U}}), \dots, t_n^{\mathcal{M}}([\vec{a}]_{\mathcal{U}})) \\ &= \lim_{\mathcal{U}} P^{\mathcal{M}_i}(t_1^{\mathcal{M}_i}(\vec{a}(i)), \dots, t_n^{\mathcal{M}_i}(\vec{a}(i))) \\ &= \lim_{\mathcal{U}} \varphi^{\mathcal{M}_i}(\vec{a}(i)). \end{aligned}$$

The second equality follows from the aforementioned fact about interpretations of terms.

Next suppose that $\varphi(\vec{x}) = u(\psi_1(\vec{x}), \dots, \psi_n(\vec{x}))$, where each ψ_j is a formula and u is a continuous function. We then have

$$\begin{aligned} \varphi^{\mathcal{M}}([\vec{a}]_{\mathcal{U}}) &= u(\psi_1^{\mathcal{M}}([\vec{a}]_{\mathcal{U}}), \dots, \psi_n^{\mathcal{M}}([\vec{a}]_{\mathcal{U}})) \\ &= u(\lim_{\mathcal{U}} \psi_1^{\mathcal{M}_i}(\vec{a}(i)), \dots, \lim_{\mathcal{U}} \psi_n^{\mathcal{M}_i}(\vec{a}(i))) \\ &= \lim_{\mathcal{U}} u(\psi_1^{\mathcal{M}_i}(\vec{a}(i)), \dots, \psi_n^{\mathcal{M}_i}(\vec{a}(i))) \\ &= \lim_{\mathcal{U}} \varphi^{\mathcal{M}_i}(\vec{a}(i)). \end{aligned}$$

The second equality follows from the induction hypothesis and the third equality follows from Theorem 3.1.14.

We finally suppose that $\varphi(\vec{x}) = \sup_y \psi(\vec{x}, y)$. Set $r := \varphi^{\mathcal{M}}([\vec{a}]_{\mathcal{U}})$, $r_i := \varphi^{\mathcal{M}_i}(\vec{a}(i))$, and $s := \lim_{\mathcal{U}} r_i$. We want to prove that $s = r$. We first show that $r \leq s$. To see this, fix $[b]_{\mathcal{U}} \in M$ and note that $\psi^{\mathcal{M}_i}(\vec{a}(i), b(i)) \leq r_i$, whence $\lim_{\mathcal{U}} \psi^{\mathcal{M}_i}(\vec{a}(i), b(i)) \leq \lim_{\mathcal{U}} r_i = s$. By the induction hypothesis, we know that the left-hand side of the above inequality is $\psi^{\mathcal{M}}([\vec{a}]_{\mathcal{U}}, [b]_{\mathcal{U}})$, whence taking the supremum over b gives $r \leq s$. We now prove that $s \leq r$. To see this, fix $\epsilon > 0$ and take $b(i) \in M_i$ such that $\psi^{\mathcal{M}_i}(\vec{a}(i), b(i)) \geq r_i - \epsilon$. By the induction hypothesis we have that $\psi^{\mathcal{M}}([\vec{a}]_{\mathcal{U}}, [b]_{\mathcal{U}}) = \lim_{\mathcal{U}} \psi^{\mathcal{M}_i}(\vec{a}(i), b(i)) \geq \lim_{\mathcal{U}} (r_i - \epsilon) = s - \epsilon$. It follows that $r = \varphi^{\mathcal{M}}([\vec{a}]_{\mathcal{U}}) \geq s - \epsilon$. Letting ϵ tend to 0 gives $r \geq s$, as desired.

Exercise 11.5.2. Fix $a, b, c \in \prod_{i \in I} M_i$. It is clear that $d(a, a) = 0$ and $d(a, b) = d(b, a)$. We now check the triangle inequality. Fix $\epsilon > 0$. Take $J_1 \in \mathcal{F}$ such that $\sup_{i \in J_1} d_i(a(i), b(i)) < d(a, b) + \epsilon$. Similarly, take $J_2 \in \mathcal{F}$ such that $\sup_{i \in J_2} d_i(b(i), c(i)) < d(b, c) + \epsilon$. Then $\sup_{i \in J_1 \cap J_2} d_i(a(i), c(i)) \leq d(a, b) + d(b, c) + 2\epsilon$. Since $J_1 \cap J_2 \in \mathcal{F}$, we have that

$$d(a, c) = \inf_{J \in \mathcal{F}} \sup_{i \in J} d_i(a(i), c(i)) < d(a, b) + d(b, c) + 2\epsilon.$$

Letting ϵ tend to 0 establishes the triangle inequality.

Chapter 12. Asymptotic cones and Gromov's theorem

Exercise 12.2.6. This follows from the fundamental theorem of finitely generated abelian groups together with the fact that having polynomial growth is closed under taking finite direct sums.

Exercise 12.2.11. The first part of (1) is obvious and the second part of (1) follows from the fact that any finite generating set for Δ can be extended to a finite generating set for Γ . We now prove (2). Suppose that Y is a finite generating set for Δ and let T be a finite set of coset representatives for Δ in Γ . Let $m \in \mathbb{N}$ be big enough so that whenever we multiply $a \cdot b$ with $a \in T$ and $b \in Y \cup T$, we can write the resulting product as $y_1 \cdots y_k \cdot z$, where $y_1, \dots, y_k \in Y$, $z \in T$, and $k \leq m$. It follows that any element of $B_X(n)$ can be written as $y_1 \cdots y_l \cdot z$, where $y_1, \dots, y_l \in Y$, $z \in T$, and $l \leq mn$. Consequently, $G_X(n) \leq |X| \cdot G_Y(mn)$, as desired.

Exercise 12.4.4. By assumption, we know, for all $x, x' \in X$, that

$$\frac{1}{K} d_X(x, x') \leq d_Y(f(x), f(x')) \leq K d_X(x, x').$$

Consequently, given any $[x]_{\mathcal{U}}, [x']_{\mathcal{U}} \in \text{Cone}(X; \mathcal{U}, o, r)$, we have

$$\frac{1}{K} d_X(x(n), x'(n)) \leq d_Y(f(x(n)), f(x'(n))) \leq K d_X(x(n), x'(n)).$$

Dividing all sides by $r(n)$ and taking the \mathcal{U} -ultralimit yields that $f^{\mathcal{U}}$ is also a K -bi-Lipschitz homeomorphism.

Exercise 12.4.8. Set $f_n : B_{\Gamma^*}(e, R) \times \cdots \times B_{\Gamma^*}(e, R) \rightarrow B_{\Gamma^*}(e, nR)$ to be the map

$$f_n(\gamma_1, \dots, \gamma_n) := \gamma_1 \cdots \gamma_n.$$

Note that if $d(\gamma_i, \eta_i)/R \approx 0$ for all $i = 1, \dots, n$, then an inductive argument, using bi-invariance of the metric, shows that $d(\gamma_1 \cdots \gamma_n, \eta_1 \cdots \eta_n)/R \leq \sum_{i=1}^n d(\gamma_i, \eta_i)/R \approx 0$ as well. Consequently, we get an induced map $\hat{f}_n : B_Y(e\mu, 1) \times \cdots \times B_Y(e\mu, 1) \rightarrow B_Y(e\mu, n)$. To see that \hat{f}_n is continuous, fix

$\epsilon > 0$ and note that if $d_Y(\gamma_i \boldsymbol{\mu}, \eta_i \boldsymbol{\mu}) < \frac{\epsilon}{n}$ for all $i = 1, \dots, n$, then $\frac{d(\gamma_i, \eta_i)}{R} < \frac{\epsilon}{n}$ for all $i = 1, \dots, n$, and thus

$$d(\gamma_1 \cdots \gamma_n, \eta_1 \cdots \eta_n)/R \leq \sum_{i=1}^n d(\gamma_i, \eta_i)/R < \epsilon,$$

whence $d_Y(\gamma_1 \cdots \gamma_n \boldsymbol{\mu}, \eta_1 \cdots \eta_n \boldsymbol{\mu}) \leq \epsilon$, as desired. Finally, to see that \hat{f}_n is surjective, fix $\gamma \in \Gamma^*$ such that $d_Y(\gamma \boldsymbol{\mu}, e \boldsymbol{\mu}) < n$, that is, $\text{st}(d(\gamma, e)/R) < n$. This allows us to write $\gamma = \gamma_1 \cdots \gamma_n \cdot \eta$, where each $|\gamma_i|/R < n$ and $|\eta|/R \approx 0$. It follows that $\hat{f}_n(\gamma_1 \boldsymbol{\mu}, \dots, \gamma_n \boldsymbol{\mu}) = \gamma \boldsymbol{\mu}$.

Exercise 12.5.4. Take $a \in \Gamma^*$ with $|a| \leq r$ such that $\delta(\eta^{-1} \gamma \eta, r) = d(\eta^{-1} \gamma \eta a, a)$. We then note that

$$d(\eta^{-1} \gamma \eta a, a) = d(\gamma \eta a, \eta a) \leq d(\gamma \eta a, \gamma a) + d(\gamma a, a) + d(a, \eta a).$$

Using that the metric is bi-invariant, we have that the last expression equals

$$d(\eta, e) + d(\gamma a, a) + d(e, \eta) \leq \delta(\gamma, r) + 2|\eta|,$$

as desired.

Chapter 13. Sofic groups

Exercise 13.1.11. Note that

$$\text{tr}((u - v)(u - v)^*) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n (u - v)_{ij} (u - v)_{ji}^* = \frac{1}{n} \sum_{i,j=1}^n |u_{ij} - v_{ij}|^2.$$

The bi-invariance of d_{HS} actually follows from this formula. Indeed,

$$\begin{aligned} \text{tr}((wu - vw)(wu - vw)^*) &= \text{tr}(w(u - v)(u - v)^* w^*) \\ &= \text{tr}((u - v)(u - v)^* w^* w) = \text{tr}((u - v)(u - v)^*). \end{aligned}$$

Here, we have used the fact that $\text{tr}(ab) = \text{tr}(ba)$ for all $a, b \in M_n(\mathbb{C})$.

Exercise 13.2.4. Parts (1) and (2) are straightforward calculations. For (3), suppose that G_1 and G_2 are sofic groups, $F \subseteq G_1 \times G_2$ is finite, and $\epsilon > 0$ is given. Without loss of generality, we may assume that $F = F_1 \times F_2$, where F_1 and F_2 are finite subsets of G_1 and G_2 , respectively. For each $i = 1, 2$, let $\phi_i : F_i \rightarrow S_{n_i}$ be an (F_i, ϵ) -morphism. By part (1), we may assume that $n_1 = n_2$; call this common value n . Let $\phi : F \rightarrow S_{2n}$ be given by $\phi(a, b) = \iota(\phi_1(a), \phi_2(b))$, where $\iota : S_n \times S_n \hookrightarrow S_{2n}$ is the embedding defined in part (2). We claim that ϕ is a (F, ϵ) -morphism. Indeed, if $g = (g_1, g_2)$ and $h = (h_1, h_2)$ are elements of F such that $gh \in F$, then by part (2) of

the exercise, we have

$$\begin{aligned} d(\phi(g)\phi(h), \phi(gh)) &= d(\iota(\phi_1(g_1)\phi_1(h_1), \phi_2(g_2)\phi_2(h_2)), \iota(\phi_1(g_1h_1), \phi_2(g_2h_2))) \\ &= \frac{1}{2} \sum_{i=1}^2 d(\phi_i(g_i)\phi_i(h_i), \phi_i(g_ih_i)) < \epsilon. \end{aligned}$$

The proof of the second part of the definition of (F, ϵ) -morphism is similar. For the third part, suppose that $g, h \in F$ as above are distinct. Suppose, without loss of generality, that $g_1 \neq h_1$. Then $d(\phi(g), \phi(h)) \geq \frac{1}{2}d(\phi_1(g_1), \phi_1(h_1)) \geq \frac{1}{4}$. While this does not quite align with the definition of (F, ϵ) -morphism, this inequality is still good enough in light of Corollary 13.2.9, which is hinted at in Remarks 13.2.2(1).

Exercise 13.2.6. Fix $\sigma, \tau \in S_n$ and let $X = \{i : \sigma(i) = \tau(i)\}$. Then $1 - d_{\text{Hammm}}(\sigma, \tau) = \frac{|X|}{n}$. On the other hand, $\{(i_1, \dots, i_k) : \sigma^{\otimes k}(i_1, \dots, i_k) = \tau^{\otimes k}(i_1, \dots, i_k)\} = X^k$, so $1 - d_{\text{Hammm}}(\sigma^{\otimes k}, \tau^{\otimes k}) = \frac{|X|^k}{n^k}$. The desired result follows.

Exercise 13.2.10. In the proof of Theorem 13.2.7, choose $\phi = \phi_{(F, \frac{1}{n})}$ to be an $(F, \frac{1}{n})$ -morphism except that for distinct $g, h \in F$, we have $d_{\text{Hammm}}(\phi(g), \phi(h)) \geq 1 - \frac{1}{n}$, which is possible by Corollary 13.2.9. The resulting embedding $\phi : G \rightarrow \prod_{\mathcal{U}} S_{(F, \frac{1}{n})}$ is now distance preserving.

Exercise 13.3.5. The backward direction is clear. To prove the forward direction, enumerate $G = (g_n)_{n \in \mathbb{N}}$ and, for each $n \in \mathbb{N}$, set $G_n := (g_k)_{k \leq n}$. Letting F_n be a finite $(G_n, \frac{1}{n})$ -Følner set, we see that $(F_n)_{n \in \mathbb{N}}$ is a Følner sequence for G .

Exercise 13.3.7. For (1), suppose that H is a subgroup of G . We will show that H is amenable using Theorem 13.3.6. By Theorem 13.3.6, there is a finitely additive left-invariant probability measure μ on G . Fix a set T of right coset representatives for H . Define $\nu : \mathcal{P}(H) \rightarrow [0, 1]$ by $\nu(A) := \mu(AT)$. One can check that ν is a finitely additive left-invariant probability measure on H . Perhaps the trickiest part is finite additivity. To see this, suppose that A and B are disjoint subsets of H ; to show that $\nu(A \cup B) = \nu(A) + \nu(B)$, by finite additivity of μ , it suffices to show that AT and BT are disjoint subsets of G . If $ax = by$ for some $a \in A$, $b \in B$, and $x, y \in T$, then by the definition of T , we must have $x = y$; but then it follows that $a = b$, which contradicts that A and B are disjoint.

For (2), fix finite $K \subseteq H$ and $\epsilon > 0$. For each $y \in K$, take $x_y \in G$ such that $f(x_y) = y$ and set $K' := \{x_y : y \in K\}$. Fix a finite (K', ϵ) -Følner set $F \subseteq G$. Then $f(F)$ is a (K, ϵ) -Følner set.

Chapter 14. Functional analysis

Exercise 14.1.3. Consider the function $\Phi: \ell^\infty(X_i) \rightarrow \prod_{\mathcal{U}} X_i$ given by $\Phi(x) := [x]_{\mathcal{U}}$. It is clear that Φ is linear and surjective. Moreover, Φ is contractive, that is, $\|\Phi(x)\| \leq \|x\|$. Indeed, $\|\Phi(x)\| = \|[x]_{\mathcal{U}}\| = \lim_{\mathcal{U}} \|x(i)\| \leq \sup_{i \in I} \|x(i)\| = \|x\|$. Thus, by the FIP for Banach spaces, we have that $\hat{\Phi}: \ell^\infty(X_i)/\ker(\Phi) \rightarrow \prod_{\mathcal{U}} X_i$ given by $\hat{\Phi}(x + \ker(\Phi)) = \Phi(x)$ is an isomorphism. It remains to notice that $Y = \ker(\Phi)$.

Exercise 14.1.7. For (1), since the diagonal embedding of X into $X^{\mathcal{U}}$ is an isometric embedding, if we can show that $\dim(X^{\mathcal{U}}) = \dim(X)$, then the diagonal embedding is necessarily onto and thus an isomorphism between X and $X^{\mathcal{U}}$. To see this dimension equality, suppose that $[x_1]_{\mathcal{U}}, \dots, [x_n]_{\mathcal{U}} \in X^{\mathcal{U}}$ are linearly independent. Lemma 14.1.6 then implies that $x_1(i), \dots, x_n(i)$ are linearly independent elements of X for \mathcal{U} -almost all $i \in I$; as a result, we have that $n \leq \dim(X)$, whence $\dim(X^{\mathcal{U}}) \leq \dim(X)$.

We now prove (2). It suffices to check that if there is $n \in \mathbb{N}$ such that $\dim(X_i) = n$ for \mathcal{U} -almost all $i \in I$, then $\dim(\prod_{\mathcal{U}} X_i) = n$. Indeed, suppose that this statement is true and $\dim(\prod_{\mathcal{U}} X_i) = m < \infty$. Take a basis $[x_1]_{\mathcal{U}}, \dots, [x_m]_{\mathcal{U}}$ for $\prod_{\mathcal{U}} X_i$. By Lemma 14.1.6, $x_1(i), \dots, x_m(i)$ are linearly independent for \mathcal{U} -almost all $i \in I$, whence $\dim(X_i) \geq m$ for \mathcal{U} -almost all $i \in I$. Suppose, toward a contradiction, that $\dim(X_i) > m$ for \mathcal{U} -almost all $i \in I$ and let Z_i be an $(m+1)$ -dimensional subspace of X_i for these $i \in I$. Then by the above statement, $\dim(\prod_{\mathcal{U}} Z_i) = m+1$, contradicting the fact that $\prod_{\mathcal{U}} Z_i$ is a subspace of $\prod_{\mathcal{U}} X_i$.

We now prove the above statement. Suppose that there is $n \in \mathbb{N}$ such that $\dim(X_i) = n$ for \mathcal{U} -almost all $i \in I$. For these $i \in I$, let $x_1(i), \dots, x_n(i) \in X_i$ be a basis for X_i consisting of vectors of norm at most 1. By considering quotient spaces, one may even assume that, for these $i \in I$, we have $d(x_k(i), \text{span}(x_1(i), \dots, x_{k-1}(i))) \geq \frac{1}{2}$ for all $1 \leq k \leq n$. We show that $[x_1]_{\mathcal{U}}, \dots, [x_n]_{\mathcal{U}}$ form a basis for $\prod_{\mathcal{U}} X_i$. It is clear that $[x_1]_{\mathcal{U}}, \dots, [x_n]_{\mathcal{U}}$ span $\prod_{\mathcal{U}} X_i$. Moreover, if $[x_1]_{\mathcal{U}}, \dots, [x_n]_{\mathcal{U}}$ were linearly dependent, then we would have, without loss of generality, $[x_n]_{\mathcal{U}} \in \text{span}([x_1]_{\mathcal{U}}, \dots, [x_{n-1}]_{\mathcal{U}})$, whence $d(x_n(i), \text{span}(x_1(i), \dots, x_{n-1}(i))) < \frac{1}{2}$ for \mathcal{U} -almost all $i \in I$, contradicting our choice of bases.

Exercise 14.2.16. Suppose that X is a super-reflexive Banach space and Y is a closed subspace of X ; we wish to show that both Y and X/Y are super-reflexive. To see that Y is super-reflexive, by Corollary 14.2.13, it suffices to show that $Y^{\mathcal{U}}$ is reflexive for every ultrafilter \mathcal{U} . However, $Y^{\mathcal{U}}$ is a subspace of the Banach space $X^{\mathcal{U}}$, which is reflexive by Corollary 14.2.13, and hence is itself reflexive by Fact 14.2.12. Similarly, by Proposition 14.1.8, $(X/Y)^{\mathcal{U}}$

is isomorphic to $X^{\mathcal{U}}/Y^{\mathcal{U}}$, which is reflexive by Corollary 14.2.13 and Fact 14.2.12.

Exercise 14.2.23. First suppose that $\lim_{\mathcal{U}} \varphi_i = \varphi$ in the weak*-topology and fix $x \in X$. Since the evaluation at x map is continuous, it follows from the ultralimit characterization of continuity that $\lim_{\mathcal{U}} \varphi_i(x) = \varphi(x)$. Conversely, suppose that $\lim_{\mathcal{U}} \varphi_i(x) = \varphi(x)$ for all $x \in X$. We wish to show that $\lim_{\mathcal{U}} \varphi_i = \varphi$ in the weak*-topology. To see this, fix a basic open neighborhood U of φ in X^* . Consequently, $U = \{\psi \in X^* : \max_{j=1, \dots, n} |\varphi(x_j) - \psi(x_j)| < \epsilon\}$ for some $x_1, \dots, x_n \in X$ and $\epsilon > 0$. Given $j = 1, \dots, n$, since $\lim_{\mathcal{U}} \varphi_i(x_j) = \varphi(x_j)$, we have that $|\varphi_i(x_j) - \varphi(x_j)| < \epsilon$ for \mathcal{U} -almost all $i \in I$. Consequently, we have that $\psi_i \in U$ for \mathcal{U} -almost all $i \in I$, as desired.

Exercise 14.3.19. For part (1), first assume that X is connected and $f \in P(C(X))$. By Exercise 14.3.10, we have that $f(x) \in \{0, 1\}$ for all $x \in X$; since X is connected, f must be constant and thus $f = 0$ or $f = 1$. Conversely, if X is not connected, then we can find a nonempty, proper clopen subset Y of X and the function $f : X \rightarrow \mathbb{C}$ given by $f(x) = 0$ if $x \in Y$ and $f(x) = 1$ if $x \in X \setminus Y$ belongs to $P(C(X)) \setminus \{0, 1\}$.

For (2), we may assume, without loss of generality, that \mathcal{U} is countably incomplete. Fix a family $(I_n)_{n \in \mathbb{N}}$ of elements of \mathcal{U} such that $I_n \supseteq I_{n+1}$ for all $n \in \mathbb{N}$ and such that $\bigcap_{n \in \mathbb{N}} I_n = \emptyset$. Now fix $[f]_{\mathcal{U}} \in P(\prod_{\mathcal{U}} C(X_i))$. For each $n \in \mathbb{N}$, set

$$J_n := \left\{ i \in I_n : \max(\|f(i)^* - f(i)\|, \|f(i)^2 - f(i)\|) < \frac{1}{n} \right\}.$$

For each $i \in I$, let $n(i)$ be the maximal $n \in \mathbb{N}$ such that $i \in J_n$. Note that $\lim_{\mathcal{U}} n(i) = \infty$. For each $i \in I$, let $g(i) \in P(C(X_i))$ be such that $\|f(i) - g(i)\| < \frac{1}{n(i)}$. Since $\lim_{\mathcal{U}} n(i) = \infty$, it follows that $[f]_{\mathcal{U}} = [g]_{\mathcal{U}}$, as desired.

For (3), suppose that each X_i is connected; by part (1), in order to show that $\prod_{\mathcal{U}} X_i$ is also connected, it suffices to show that the only projections in $C(\prod_{\mathcal{U}} X_i)$ are the constantly 0 and 1 functions. Since $C(\prod_{\mathcal{U}} X_i) \cong \prod_{\mathcal{U}} C(X_i)$, it suffices to show that the only projections in $\prod_{\mathcal{U}} C(X_i)$ are the elements $[0]_{\mathcal{U}}$ and $[1]_{\mathcal{U}}$, which is indeed the case by parts (1) and (2) and the assumption that each X_i is connected.

Exercise 14.4.17. Note first that every element of $L(G)$ can be written in the form $x = \sum_{g \in G} a_g u_g$. Next observe that, for each $g \in G$, the matrix representation of u_g with respect to the basis $(\zeta_h)_{h \in G}$ is a permutation matrix and that all diagonal entries of this matrix representation of u_g are 0 whenever $g \neq e$. Conclude that for any $x = \sum_{g \in G} a_g u_g$, the normalized trace of the matrix representation of x is a_e , the coefficient of u_e .

Chapter 15. Does an ultrapower depend on the ultrafilter?

Exercise 15.2.16. For (1), fix distinct $f_1, \dots, f_m, f'_1, \dots, f'_n \in F$, subsets $U_1, \dots, U_m \subseteq \omega$, and ordinals $k_1, \dots, k_n \in \omega$. It suffices to show that

$$\bigcap_{i=1}^m \{U \in \beta\omega : g_{f_i}(U) \in U_i\} \cap \bigcap_{i=1}^n \{U \in \beta\omega : g_{f'_i}(U) \geq k_i\}$$

is an open subset of $\beta\omega$. However, $\{U \in \beta\omega : g_{f_i}(U) \in U_i\} = \{U \in \beta\omega : f_i^{-1}(U_i) \in U\}$ is a basic open set in $\beta\omega$ and $\{U \in \beta\omega : g_{f_i}(U) \geq k_i\}$ is the complement of the closed set $\bigcup_{j=1}^{k_i-1} \{U \in \beta\omega : f_i^{-1}(j) \in U\}$, whence also open.

For (2), fix a function $\Psi : F \rightarrow \omega + 1$; we see $U \in \beta\omega$ such that $G(U) = \Psi$. Fix $f \in F$. If $\Psi(f) \in \omega$, then set $X_f := \{n \in \omega : f(n) = \Psi(f)\}$. If $\Psi(f) = \omega$, then for each $m \in \omega$, set $X_{f,m} := \{n \in \omega : f(n) \geq m\}$. It suffices to show that the family $(X_f)_{f \in \Psi^{-1}(\omega)} \cup (X_{f,m})_{f \in \Psi^{-1}(\{\omega\}), m \in \omega}$, together with the Fréchet filter, has the FIP, for then any ultrafilter \mathcal{U} containing all of these sets is a nonprincipal ultrafilter such that $G(\mathcal{U}) = \Psi$. However, the fact that this family has the FIP follows immediately from the fact that F is of large oscillation modulo the Fréchet filter.

Exercise 15.3.7. For part (1), by finite character, it suffices to assume that K is finite. When $|K| = 1$, we are trying to show that $a_k \perp_A a_J$. Since $a_J \subseteq a_{<k}$, we conclude that $a_k \perp_A a_J$ by monotonicity. Now suppose that $|K| > 1$ and write $K = L \cup \{k\}$, where $k > L$. By the induction hypothesis, we know that $a_L \perp_A a_J$, whence $a_J \perp_A a_L$ by symmetry. By the definition of an A -independent sequence (and transitivity again), we have that $a_k \perp_A a_J a_L$. By transitivity, we have $a_k \perp_{Aa_L} a_J$, and thus by symmetry, we have $a_J \perp_{Aa_L} a_k$. By transitivity, we conclude that $a_J \perp_A a_L a_k$, that is, $a_J \perp_A a_K$, and by symmetry one last time, we have $a_K \perp_A a_J$, as desired.

For (2), fix $i \in I$ and write $a_{\neq i} = a_{<i} a_{>i}$. By the definition of an A -independent sequence, we have $a_i \perp_A a_{<i}$. By part (1) of the exercise, we have that $a_{>i} \perp_A a_{<i}$, whence $a_{>i} \perp_{Aa_{<i}} a_i$ by transitivity, and thus $a_i \perp_{Aa_{<i}} a_{>i}$ by symmetry. By transitivity, we have $a_i \perp_A a_{<i} a_{>i}$, as desired.

Chapter 16. The Keisler-Shelah theorem

Exercise 16.1.15. Take $\mathcal{M} \models T$ such that $|M| = \kappa$. Define an equivalence relation on \mathcal{L} , where symbols of the same type are related if their interpretations in \mathcal{M} are the same. Choose a representative from each equivalence class and let \mathcal{L}' be the collection of all such representatives. Observe $|\mathcal{L}'| \leq 2^{|\mathcal{L}|} = 2^\kappa$. Conclude using the completeness of T .

Exercise 16.2.3. For (1), it is clear that Pseudo- \mathcal{K} is an elementary class containing \mathcal{K} . If T is an \mathcal{L} -theory such that every element of \mathcal{K} is a model of T , then $T \subseteq \text{Th}(\mathcal{K})$, whence every element of Pseudo- \mathcal{K} is a model of T . (2) follows from the definition by considering negations.

Exercise 16.3.2. Let $\mathcal{L}_1 = \{P, Q\}$ and $\mathcal{L}_2 = \{Q, R\}$, where P, Q , and R are unary predicate symbols. Let $T_1 = \{\forall x(P(x) \wedge Q(x))\}$ and let $T_2 = \{\forall x(\neg Q(x) \wedge R(x))\}$. Then $T_1 \cup T_2$ is not satisfiable.

Exercise 16.4.2. It is readily verified that the map which sends the element $[A(k)]_{\mathcal{U}}$ of $M_n(R)^{\mathcal{U}}$ to the matrix over $R^{\mathcal{U}}$ whose (i, j) -entry is $[A(k)_{ij}]_{\mathcal{U}}$ is an isomorphism of rings, where $A(k)_{ij}$ is the (i, j) -entry of the matrix $A(k)$.

Chapter 17. Large cardinals

Exercise 17.2.4. One proves (1) by induction on $\alpha < \kappa$. Clearly, $|V_0| = 0 < \kappa$. If $|V_\alpha| < \kappa$, then $|V_{\alpha+1}| = 2^{|V_\alpha|} < \kappa$ since κ is a strong limit cardinal. If $|V_\xi| < \kappa$ for all $\xi < \alpha$, then $|V_\alpha| = \bigcup_{\xi < \alpha} |V_\xi| < \kappa$ since κ is regular.

For (2), take $x \subseteq V_\kappa$. If $x \in V_\kappa = \bigcup_{\alpha < \kappa} V_\alpha$, then $x \in V_\alpha$ for some $\alpha < \kappa$. Since V_α is transitive, $x \subseteq V_\alpha$, and by (1), $|V_\alpha| < \kappa$, so $|x| < \kappa$. On the other hand, if $|x| < \kappa$, by regularity of κ , there exists $\alpha < \kappa$ such that $x \subseteq V_\alpha$, whence $x \in V_{\alpha+1} \subseteq V_\kappa$.

Exercise 17.2.12. Let η be such that $\text{cof}(\eta) > \omega$. We proceed by induction on $\lambda < \text{cof}(\eta)$ to show that for any sequence $(C_\alpha)_{\alpha < \lambda}$ of club subsets of η , $\bigcap_{\alpha < \lambda} C_\alpha$ is a club subset of η . By Lemma 17.2.11, we see that if $\bigcap_{\alpha < \lambda} C_\alpha$ is a club subset of η , then $D \cap (\bigcap_{\alpha < \lambda} C_\alpha)$ is a club subset of η for any club subset D of η , whence the successor case holds. Now assume $D_\gamma := \bigcap_{\alpha < \gamma} C_\alpha$ is a club subset of η for any $\gamma < \lambda$. Then $(D_\gamma)_{\gamma < \lambda}$ is a decreasing sequence of club subsets of η , and $\bigcap_{\gamma < \lambda} D_\gamma = \bigcap_{\alpha < \lambda} C_\alpha$, whence it is enough to show $\bigcap_{\gamma < \lambda} D_\gamma$ is a club subset of η . Clearly, $\bigcap_{\gamma < \lambda} D_\gamma$ is closed. Now take $\xi < \eta$. We construct a sequence $(\xi_\gamma)_{\gamma < \lambda}$ as follows: Choose $\xi_0 \in D_0$ such that $\xi < \xi_0$. For each $0 < \gamma < \lambda$, choose $\xi_\gamma \in D_\gamma$ such that $\xi_\gamma > \sup\{\xi_{\gamma'} : \gamma' < \gamma\}$. Set $\xi' := \sup\{\xi_\gamma : \gamma < \lambda\} < \kappa$. Then $\xi' \in \bigcap_{\gamma < \lambda} D_\gamma$ and $\xi < \xi'$. Hence $\bigcap_{\gamma < \lambda} D_\gamma$ is unbounded.

Exercise 17.3.2. Let \mathcal{U} be a nonprincipal κ -complete ultrafilter on κ . Assume, toward a contradiction, that $\text{cof}(\kappa) < \kappa$. Let $\kappa_i < \kappa$ be such that $\kappa = \sup_{i < \text{cof}(\kappa)} \kappa_i$. For every $i < \text{cof}(\kappa)$, observe $[\kappa_i, \kappa) := \bigcap_{\xi < \kappa_i} (\kappa \setminus \{\xi\}) \in \mathcal{U}$. Consequently, $\emptyset = \bigcap_{i < \text{cof}(\kappa)} [\kappa_i, \kappa) \in \mathcal{U}$, a contradiction.

Exercise 17.3.7. First assume \mathcal{U} is normal and let $g : \kappa \rightarrow \kappa$ be \mathcal{U} -regressive. Let $X \in \mathcal{U}$ be such that g is regressive on X . Then for any $\alpha \in X$, we have that $g(\alpha) < \alpha = \text{id}(\alpha)$, so $[g]_{\mathcal{U}} < [\text{id}]_{\mathcal{U}}$. Since \mathcal{U} is normal, $[\text{id}]_{\mathcal{U}}$ is

the κ th element of $\kappa^{\mathcal{U}}$, so by Lemma 17.3.5, there exists $\alpha < \kappa$ such that $[g]_{\mathcal{U}} = d(\alpha)$. Thus, g is constant on a set in \mathcal{U} .

For the converse, assume that whenever $g : \kappa \rightarrow \kappa$ is \mathcal{U} -regressive, then g is constant on a set in \mathcal{U} . First observe that $d(\alpha) < [\text{id}]_{\mathcal{U}}$ for any $\alpha < \kappa$. Now let $g : \kappa \rightarrow \kappa$ be such that $[g]_{\mathcal{U}} < [\text{id}]_{\mathcal{U}}$. Hence g is \mathcal{U} -regressive, so by assumption there exists $\alpha < \kappa$ such that $[g]_{\mathcal{U}} = d(\alpha)$. Thus, the set of predecessors of $[\text{id}]_{\mathcal{U}}$ is simply $\{d(\alpha) : \alpha < \kappa\}$, and by Lemma 17.3.5, we conclude $[\text{id}]_{\mathcal{U}}$ is the κ th element of $\kappa^{\mathcal{U}}$ and thus \mathcal{U} is normal.

Exercise 17.4.10. Assume κ is strongly compact. Let \mathcal{F} be a κ -complete filter on I . Let \mathcal{L} be the language which has a unary relation symbol X for every $X \in \mathcal{P}(I)$. Let \mathcal{I} be the \mathcal{L} -structure on I where $X^{\mathcal{I}} := X$ for every $X \in \mathcal{P}(I)$. Set $\Sigma := \text{Th}_{\mathcal{L}, \kappa}(\mathcal{I}) \cup \{X(c) : X \in \mathcal{F}\}$ where c is a new constant symbol. Then every subset of Σ of size less than κ has a model since \mathcal{F} is κ -complete. Since κ is strongly compact, there is a model \mathcal{N} of Σ . Define $\mathcal{U} \subseteq \mathcal{P}(I)$ by declaring that, for every $X \in \mathcal{P}(I)$, $X \in \mathcal{U}$ iff $\mathcal{N} \models X(c)$. Show that \mathcal{U} is a κ -complete ultrafilter on I which extends \mathcal{F} .

Exercise 17.6.2. First assume that E is not well founded, whence there exists a sequence $([f_n]_{\mathcal{U}}^r)_{n < \omega}$ such that $[f_{n+1}]_{\mathcal{U}}^r E [f_n]_{\mathcal{U}}^r$ for any $n < \omega$. For $n < \omega$, set $X_n := \{i \in I : f_{n+1}(i) \in f_n(i)\}$ and observe that each $X_n \in \mathcal{U}$. If $i \in \bigcap_{n < \omega} X_n$, then $(f_n(i))_{n < \omega}$ is an infinite descending \in -sequence in V . Since V is well founded, we conclude that $\bigcap_{n < \omega} X_n = \emptyset$, whence \mathcal{U} is countably incomplete. The other direction follows from the fact that countably incomplete ultrafilters yield \aleph_1 -saturated ultrapowers.

Exercise 17.6.8. (1) \implies (2): If $g : \kappa \rightarrow V$ is such that $[g]_{\mathcal{U}}^r \in [\text{id}]_{\mathcal{U}}^r$, then there exists $X \in \mathcal{U}$ such that $g(\alpha) \in \alpha$ for any $\alpha \in X$. Define $f : \kappa \rightarrow \kappa$ by $f(\alpha) = g(\alpha)$ if $\alpha \in X$ and $f(\alpha) = 0$ otherwise. Then $[f]_{\mathcal{U}} < [\text{id}]_{\mathcal{U}}$. By Lemma 17.3.5, $[\text{id}]_{\mathcal{U}}$ is the κ th element of $\kappa^{\mathcal{U}}$, whence there exists $\gamma < \kappa$ such that $[f]_{\mathcal{U}} = \gamma = [\hat{\gamma}]_{\mathcal{U}}$ where $\hat{\gamma} : \kappa \rightarrow \kappa$ is such that $\text{range}(\hat{\gamma}) = \{\gamma\}$. Hence $g =_{\mathcal{U}} \hat{\gamma}$, whence $[g]_{\mathcal{U}}^r = d(\gamma) = \gamma < \kappa$ by Lemma 17.6.4. Thus $[\text{id}]_{\mathcal{U}}^r = \kappa$.

For (2) \implies (3), note that, for every $X \subseteq \kappa$, we have $X \in \mathcal{U}$ iff $[\text{id}]_{\mathcal{U}}^r \in d(X)$. Indeed, for any $X \subseteq \kappa$, we have $[\text{id}]_{\mathcal{U}}^r \in d(X)$ if and only if $\{\alpha < \kappa : \text{id}(\alpha) \in X\} \in \mathcal{U}$ if and only if $X \in \mathcal{U}$.

For (3) \implies (1), observe that for any functions $f, g : \kappa \rightarrow \kappa$, the following items are equivalent:

- $f =_{\mathcal{U}} g$,
- $\{\alpha < \kappa : f(\alpha) = g(\alpha)\} \in \mathcal{U}$,
- $\kappa \in d(\{\alpha < \kappa : f(\alpha) = g(\alpha)\})$,
- $\kappa \in \{\alpha < d(\kappa) : d(f)(\alpha) = d(g)(\alpha)\}$,
- $d(f)(\kappa) = d(g)(\kappa)$.

Now let $f : \kappa \rightarrow \kappa$ be \mathcal{U} -regressive. By definition, there exists some $X \in \mathcal{U}$ such that f is regressive on X . Since $\kappa \in d(X)$, and $d(f)$ is regressive on $d(X)$, there must exist $\gamma < \kappa$ such that $d(f)(\kappa) = \gamma$. At the same time, $d(\hat{\gamma})(\kappa) = \gamma$, where $\hat{\gamma} : \kappa \rightarrow \kappa$ is such that $\text{range}(\hat{\gamma}) = \{\gamma\}$. Hence, by the above equivalence, we have that $f =_{\mathcal{U}} \hat{\gamma}$, whence f is \mathcal{U} -constant. By Exercise 17.3.7, \mathcal{U} is normal.

Exercise 17.6.22. Let κ be 2^κ -supercompact, so there exists a transitive model M and an embedding $j : V \rightarrow M$ such that $\text{crit}(j) = \kappa$, $j(\kappa) > 2^\kappa$, and $M^{2^\kappa} \subseteq M$. By Theorem 17.6.9, $\mathcal{U}_j := \{X \subseteq \kappa : \kappa \in j(X)\}$ is a normal ultrafilter on κ . Observe that $M \models \text{ZF}$ and $\text{On} \subseteq M$. Since j is κ -supercompact, we have $\mathcal{P}^M(\kappa) = \mathcal{P}(\kappa)$. Since j is 2^κ -supercompact, we have $\mathcal{U}_j \in M$. It follows that $M \models [\mathcal{U}_j \text{ is a normal ultrafilter on } \kappa]$ and thus $M \models [\kappa \text{ is measurable}]$. By Theorem 17.6.10, there is an elementary embedding $j' : V^{\mathcal{U}_j} \rightarrow M$ defined by $j'([f]_{\mathcal{U}_j}^r) = j(f)(\kappa)$. Since $j'(\kappa) = j'([\text{id}]_{\mathcal{U}_j}^r) = j(\text{id})(\kappa) = \kappa$, we have that $V^{\mathcal{U}_j} \models [\kappa \text{ is measurable}]$ by elementarity of j' . By Exercise 17.6.8, $\{\alpha < \kappa : \alpha \text{ is measurable}\} \in \mathcal{U}_j$.

Bibliography

- [1] C. Anantharaman and S. Popa, *An introduction to II_1 factors*, book in preparation.
- [2] K. J. Arrow, *A difficulty in the concept of social welfare*, Journal of political economy, **58** (1950), pp. 328–346.
- [3] James Ax, *The elementary theory of finite fields*, Ann. of Math. (2) **88** (1968), 239–271, DOI 10.2307/1970573. MR229613
- [4] James E. Baumgartner, *A short proof of Hindman's theorem*, J. Combinatorial Theory Ser. A **17** (1974), 384–386, DOI 10.1016/0097-3165(74)90103-4. MR354394
- [5] James E. Baumgartner and Richard Laver, *Iterated perfect-set forcing*, Ann. Math. Logic **17** (1979), no. 3, 271–288, DOI 10.1016/0003-4843(79)90010-X. MR556894
- [6] Benjamin Baumslag, *Residually free groups*, Proc. London Math. Soc. (3) **17** (1967), 402–418, DOI 10.1112/plms/s3-17.3.402. MR215903
- [7] Gilbert Baumslag, Alexei Myasnikov, and Vladimir Remeslennikov, *Discriminating completions of hyperbolic groups*, Geom. Dedicata **92** (2002), 115–143, DOI 10.1023/A:1019687202544. Dedicated to John Stallings on the occasion of his 65th birthday. MR1934015
- [8] Mathias Beiglböck, *An ultrafilter approach to Jin's theorem*, Israel J. Math. **185** (2011), 369–374, DOI 10.1007/s11856-011-0114-5. MR2837141
- [9] Itai Ben Yaacov, Alexander Berenstein, C. Ward Henson, and Alexander Usvyatsov, *Model theory for metric structures*, Model theory with applications to algebra and analysis. Vol. 2, London Math. Soc. Lecture Note Ser., vol. 350, Cambridge Univ. Press, Cambridge, 2008, pp. 315–427, DOI 10.1017/CBO9780511735219.011. MR2436146
- [10] Itai Ben Yaacov and Alexander Usvyatsov, *Continuous first order logic and local stability*, Trans. Amer. Math. Soc. **362** (2010), no. 10, 5213–5259, DOI 10.1090/S0002-9947-10-04837-3. MR2657678
- [11] Vitaly Bergelson, *Combinatorial and Diophantine applications of ergodic theory*, Handbook of dynamical systems. Vol. 1B, Elsevier B. V., Amsterdam, 2006, pp. 745–869, DOI 10.1016/S1874-575X(06)80037-8. Appendix A by A. Leibman and Appendix B by Anthony Quas and Máté Wierdl. MR2186252
- [12] Vitaly Bergelson, *Ultrafilters, IP sets, dynamics, and combinatorial number theory*, Ultrafilters across mathematics, Contemp. Math., vol. 530, Amer. Math. Soc., Providence, RI, 2010, pp. 23–47, DOI 10.1090/conm/530/10439. MR2757532

- [13] Andreas Blass, *A model without ultrafilters* (English, with Russian summary), Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys. **25** (1977), no. 4, 329–331. MR476510
- [14] Andreas Blass, *Combinatorial cardinal characteristics of the continuum*, in Handbook of Set Theory, Springer, 2010, pp. 395–489.
- [15] Andreas Blass and Saharon Shelah, *Ultrafilters with small generating sets*, Israel J. Math. **65** (1989), no. 3, 259–271, DOI 10.1007/BF02764864. MR1005010
- [16] Andreas Raphael Blass, *Orderings of ultrafilters*, ProQuest LLC, Ann Arbor, MI, 1970. Thesis (Ph.D.)—Harvard University. MR2939947
- [17] William Boos, *Infinitary compactness without strong inaccessibility*, J. Symbolic Logic **41** (1976), no. 1, 33–38, DOI 10.2307/2272942. MR409184
- [18] David Booth, *Ultrafilters on a countable set*, Ann. Math. Logic **2** (1970/71), no. 1, 1–24, DOI 10.1016/0003-4843(70)90005-7. MR277371
- [19] N. Bourbaki, *Éléments de mathématique. Topologie algébrique. Chapitres 1 à 4* (French), Springer, Heidelberg, 2016. MR3617167
- [20] Emmanuel Breuillard, Ben Green, and Terence Tao, *The structure of approximate groups*, Publ. Math. Inst. Hautes Études Sci. **116** (2012), 115–221, DOI 10.1007/s10240-012-0043-9. MR3090256
- [21] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Introductory notes on Richard Thompson’s groups*, Enseign. Math. (2) **42** (1996), no. 3–4, 215–256. MR1426438
- [22] Valerio Capraro and Martino Lupini, *Introduction to sofic and hyperlinear groups and Connes’ embedding conjecture*, Lecture Notes in Mathematics, vol. 2136, Springer, Cham, 2015. With an appendix by Vladimir Pestov, DOI 10.1007/978-3-319-19333-5. MR3408561
- [23] H. Cartan, *Filtres et ultrafiltres*, Comptes Rendus de l’Académie des Sciences, **205** (1937), pp. 777–779.
- [24] H. Cartan, *Théorie des filtres*, Comptes Rendus de l’Académie des Sciences, **205** (1937), pp. 595–598.
- [25] Eduard Čech, *On bicomact spaces*, Ann. of Math. (2) **38** (1937), no. 4, 823–844, DOI 10.2307/1968839. MR1503374
- [26] Christophe Champetier and Vincent Guirardel, *Limit groups as limits of free groups*, Israel J. Math. **146** (2005), 1–75, DOI 10.1007/BF02773526. MR2151593
- [27] Chen-chung Chang and H. Jerome Keisler, *Continuous model theory*, Annals of Mathematics Studies, No. 58, Princeton Univ. Press, Princeton, N.J., 1966. MR0231708
- [28] Chen-chung Chang and H. Jerome Keisler, *Model theory*, Dover Publications, 1990.
- [29] Charles Chidume, *Geometric properties of Banach spaces and nonlinear iterations*, Lecture Notes in Mathematics, vol. 1965, Springer-Verlag London, Ltd., London, 2009. MR2504478
- [30] Ian Chiswell, *Introduction to Λ -trees*, World Scientific Publishing Co., Inc., River Edge, NJ, 2001, DOI 10.1142/4495. MR1851337
- [31] Matt Clay and Dan Margalit, *Groups*, Office hours with a geometric group theorist, Princeton Univ. Press, Princeton, NJ, 2017, pp. 3–20. MR3587214
- [32] E. Coleman, *Banach space ultraproducts*, Irish Math. Soc. Bull. **18** (1987), 30–39. MR885794
- [33] W. W. Comfort, *Ultrafilters: some old and some new results*, Bull. Amer. Math. Soc. **83** (1977), no. 4, 417–455, DOI 10.1090/S0002-9904-1977-14316-4. MR454893
- [34] A. Connes, *Classification of injective factors. Cases II_1 , II_∞ , III_λ , $\lambda \neq 1$* , Ann. of Math. (2) **104** (1976), no. 1, 73–115, DOI 10.2307/1971057. MR454659
- [35] John B. Conway, *A course in functional analysis*, 2nd ed., Graduate Texts in Mathematics, vol. 96, Springer-Verlag, New York, 1990. MR1070713
- [36] Abraham Robinson, *A result on consistency and its application to the theory of definition*, Nederl. Akad. Wetensch. Proc. Ser. A. **59** = Indag. Math. **18** (1956), 47–58. MR0078307

- [37] V. B. Dem'yanov, *On cubic forms in discretely normed fields* (Russian), Doklady Akad. Nauk SSSR (N.S.) **74** (1950), 889–891. MR0037836
- [38] Keith J. Devlin, *Fundamentals of contemporary set theory*, Universitext, Springer-Verlag, New York-Heidelberg, 1979. MR541746
- [39] Mauro Di Nasso, *Hypernatural numbers as ultrafilters*, Nonstandard analysis for the working mathematician, Springer, Dordrecht, 2015, pp. 443–474. MR3409522
- [40] Mauro Di Nasso, *Iterated hyper-extensions and an idempotent ultrafilter proof of Rado's theorem*, Proc. Amer. Math. Soc. **143** (2015), no. 4, 1749–1761, DOI 10.1090/S0002-9939-2014-12342-2. MR3314087
- [41] Mauro Di Nasso and Marco Forti, *Hausdorff ultrafilters*, Proc. Amer. Math. Soc. **134** (2006), no. 6, 1809–1818, DOI 10.1090/S0002-9939-06-08433-4. MR2207497
- [42] Mauro Di Nasso, Isaac Goldbring, and Martino Lupini, *Nonstandard methods in Ramsey theory and combinatorial number theory*, Lecture Notes in Mathematics, vol. 2239, Springer, Cham, 2019, DOI 10.1007/978-3-030-17956-4. MR3931702
- [43] Mauro Di Nasso and Eleftherios Tachtsis, *Idempotent ultrafilters without Zorn's lemma*, Proc. Amer. Math. Soc. **146** (2018), no. 1, 397–411, DOI 10.1090/proc/13719. MR3723149
- [44] Seán Dineen, *The second dual of a JB^* triple system*, Complex analysis, functional analysis and approximation theory (Campinas, 1984), North-Holland Math. Stud., vol. 125, North-Holland, Amsterdam, 1986, pp. 67–69. MR893410
- [45] Hans-Dieter Donder, *Regularity of ultrafilters and the core model*, Israel J. Math. **63** (1988), no. 3, 289–322, DOI 10.1007/BF02778036. MR969944
- [46] Alan Dow, *Saturated Boolean algebras and their Stone spaces*, Topology Appl. **21** (1985), no. 2, 193–207, DOI 10.1016/0166-8641(85)90104-X. MR813288
- [47] Cornelia Druțu and Mark Sapir, *Tree-graded spaces and asymptotic cones of groups*, Topology **44** (2005), no. 5, 959–1058, DOI 10.1016/j.top.2005.03.003. With an appendix by Denis Osin and Mark Sapir. MR2153979
- [48] Gábor Elek and Endre Szabó, *On sofic groups*, J. Group Theory **9** (2006), no. 2, 161–171, DOI 10.1515/JGT.2006.011. MR2220572
- [49] Herbert B. Enderton, *Elements of set theory*, Academic Press [Harcourt Brace Jovanovich, Publishers], New York-London, 1977. MR0439636
- [50] Herbert B. Enderton, *A mathematical introduction to logic*, Elsevier, 2001.
- [51] Herbert B. Enderton, *Computability theory: An introduction to recursion theory*, Academic Press, 2010.
- [52] P. Enflo, *Uniform structures and square roots in topological groups. I, II*, Israel J. Math. **8** (1970), 230–252; *ibid.* **8** (1970), 253–272, DOI 10.1007/bf02771561. MR263969
- [53] Ilijas Farah, Bradd Hart, and David Sherman, *Model theory of operator algebras I: stability*, Bull. Lond. Math. Soc. **45** (2013), no. 4, 825–838, DOI 10.1112/blms/bdt014. MR3081550
- [54] Ilijas Farah, Bradd Hart, and David Sherman, *Model theory of operator algebras II: model theory*, Israel J. Math. **201** (2014), no. 1, 477–505, DOI 10.1007/s11856-014-1046-7. MR3265292
- [55] Ilijas Farah and Saharon Shelah, *A dichotomy for the number of ultrapowers*, J. Math. Log. **10** (2010), no. 1-2, 45–81, DOI 10.1142/S0219061310000936. MR2802082
- [56] S. Feferman and R. L. Vaught, *The first order properties of products of algebraic systems*, Fund. Math. **47** (1959), 57–103, DOI 10.4064/fm-47-1-57-103. MR108455
- [57] Ben Fine, Anthony M. Gaglione, Gerhard Rosenberger, and Dennis Spellman, *On CT and CSA groups and related ideas*, J. Group Theory **19** (2016), no. 5, 923–940, DOI 10.1515/jgth-2016-0005. MR3545911
- [58] Edward R. Fisher, *Abelian structures. I*, Abelian group theory (Proc. Second New Mexico State Univ. Conf., Las Cruces, N.M., 1976), Springer, Berlin, 1977, pp. 270–322. Lecture Notes in Math., Vol. 616. MR0540014

- [59] T. Frayne, A. C. Morel, and D. S. Scott, *Reduced direct products*, Fund. Math. **51** (1962/63), 195–228, DOI 10.4064/fm-51-3-195-228. MR142459
- [60] Tobias Fritz, *Tsirelson's problem and Kirchberg's conjecture*, Rev. Math. Phys. **24** (2012), no. 5, 1250012, 67, DOI 10.1142/S0129055X12500122. MR2928100
- [61] Zdeněk Frolík, *Non-homogeneity of $\beta P - P$* , Comment. Math. Univ. Carolinae **8** (1967), 705–709. MR266160
- [62] Harry Furstenberg, *Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions*, J. Analyse Math. **31** (1977), 204–256, DOI 10.1007/BF02813304. MR498471
- [63] Anthony M. Gaglione and Dennis Spellman, *Even more model theory of free groups*, Infinite groups and group rings (Tuscaloosa, AL, 1992), Ser. Algebra, vol. 1, World Sci. Publ., River Edge, NJ, 1993, pp. 37–40. MR1377955
- [64] D. Galvin, *Ultrafilters, with applications to analysis, social choice and combinatorics*, unpublished notes, 2009.
- [65] Murray Gerstenhaber and Oscar S. Rothaus, *The solution of sets of equations in groups*, Proc. Nat. Acad. Sci. U.S.A. **48** (1962), 1531–1533, DOI 10.1073/pnas.48.9.1531. MR166296
- [66] Saeed Ghasemi, *Reduced products of metric structures: a metric Feferman-Vaught theorem*, J. Symb. Log. **81** (2016), no. 3, 856–875, DOI 10.1017/jsl.2016.20. MR3569108
- [67] Leonard Gillman and Meyer Jerison, *Rings of continuous functions*, Graduate Texts in Mathematics, No. 43, Springer-Verlag, New York-Heidelberg, 1976. Reprint of the 1960 edition. MR0407579
- [68] Robert Goldblatt, *Lectures on the hyperreals: An introduction to nonstandard analysis*, Graduate Texts in Mathematics, vol. 188, Springer-Verlag, New York, 1998, DOI 10.1007/978-1-4612-0615-6. MR1643950
- [69] I. Goldbring and H. J. Keisler, *Continuous sentences preserved under reduced products*, to appear in Journal of Symbolic Logic.
- [70] R. I. Grigorchuk, *Degrees of growth of finitely generated groups and the theory of invariant means* (Russian), Izv. Akad. Nauk SSSR Ser. Mat. **48** (1984), no. 5, 939–985. MR764305
- [71] Mikhael Gromov, *Groups of polynomial growth and expanding maps*, Inst. Hautes Études Sci. Publ. Math. **53** (1981), 53–73. MR623534
- [72] M. Gromov, *Endomorphisms of symbolic algebraic varieties*, J. Eur. Math. Soc. (JEMS) **1** (1999), no. 2, 109–197, DOI 10.1007/PL00011162. MR1694588
- [73] A. Grothendieck and J. A. Dieudonné, *Éléments de géométrie algébrique. I* (French), Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 166, Springer-Verlag, Berlin, 1971. MR3075000
- [74] A. Grothendieck, *Éléments de géométrie algébrique. II. Étude globale élémentaire de quelques classes de morphismes* (French), Inst. Hautes Études Sci. Publ. Math. **8** (1961), 222. MR217084
- [75] A. Grothendieck, *Éléments de géométrie algébrique. III. Étude cohomologique des faisceaux cohérents. I*, Inst. Hautes Études Sci. Publ. Math. **11** (1961), 167. MR217085
- [76] R. Gurevič, *On ultraproducts of compact Hausdorff spaces*, J. Symbolic Logic **53** (1988), no. 1, 294–300, DOI 10.2307/2274446. MR929393
- [77] Eric J. Hall, Kyriakos Keremedis, and Eleftherios Tachtsis, *The existence of free ultrafilters on ω does not imply the extension of filters on ω to ultrafilters*, MLQ Math. Log. Q. **59** (2013), no. 4-5, 258–267, DOI 10.1002/malq.201100092. MR3100753
- [78] C. Ward Henson, *Foundations of nonstandard analysis: a gentle introduction to nonstandard extensions*, Nonstandard analysis (Edinburgh, 1996), NATO Adv. Sci. Inst. Ser. C: Math. Phys. Sci., vol. 493, Kluwer Acad. Publ., Dordrecht, 1997, pp. 1–49. MR1603228

- [79] Horst Herrlich, Kyriakos Keremedis, and Eleftherios Tachtsis, *Remarks on the Stone spaces of the integers and the reals without AC*, Bull. Pol. Acad. Sci. Math. **59** (2011), no. 2, 101–114, DOI 10.4064/ba59-2-1. MR2852866
- [80] Edwin Hewitt, *Rings of real-valued continuous functions. I*, Trans. Amer. Math. Soc. **64** (1948), 45–99, DOI 10.2307/1990558. MR26239
- [81] Neil Hindman, *The existence of certain ultra-filters on N and a conjecture of Graham and Rothschild*, Proc. Amer. Math. Soc. **36** (1972), 341–346, DOI 10.2307/2039156. MR307926
- [82] Neil Hindman, *Finite sums from sequences within cells of a partition of N* , J. Combinatorial Theory Ser. A **17** (1974), 1–11, DOI 10.1016/0097-3165(74)90023-5. MR349574
- [83] Neil Hindman and Dona Strauss, *Algebra in the Stone-Čech compactification: Theory and applications*, De Gruyter Textbook, Walter de Gruyter & Co., Berlin, 2012. Second revised and extended edition [of MR1642231]. MR2893605
- [84] Paul Howard and Jean E. Rubin, *Consequences of the axiom of choice*, Mathematical Surveys and Monographs, vol. 59, American Mathematical Society, Providence, RI, 1998. With 1 IBM-PC floppy disk (3.5 inch; WD), DOI 10.1090/surv/059. MR1637107
- [85] Paul E. Howard, *Los’ theorem and the Boolean prime ideal theorem imply the axiom of choice*, Proc. Amer. Math. Soc. **49** (1975), 426–428, DOI 10.2307/2040659. MR384548
- [86] Ehud Hrushovski, *Stable group theory and approximate subgroups*, J. Amer. Math. Soc. **25** (2012), no. 1, 189–243, DOI 10.1090/S0894-0347-2011-00708-X. MR2833482
- [87] José Iovino, *Applications of model theory to functional analysis*, Dover Publications, Inc., Mineola, NY, 2014. Revised reprint of the 2002 original; With a new preface, notes, and an updated bibliography. MR3362124
- [88] C. Ward Henson, José Iovino, Alexander S. Kechris, and Edward Odell, *Analysis and logic*, London Mathematical Society Lecture Note Series, vol. 262, Cambridge University Press, Cambridge, 2002. Lectures from the mini-courses offered at the International Conference held at the University of Mons-Hainaut, Mons, August 25–29, 1997; Edited by Catherine Finet and Christian Michaux. MR1967837
- [89] T. Jech, *Set theory*, Springer Science & Business Media, 2002.
- [90] Thomas J. Jech, *The axiom of choice*, Studies in Logic and the Foundations of Mathematics, Vol. 75, North-Holland Publishing Co., Amsterdam-London; American Elsevier Publishing Co., Inc., New York, 1973. MR0396271
- [91] Z. Ji, A. Natarajan, T. Vidick, J. Wright, and H. Yuen, *MIP*=RE*, arXiv preprint arXiv:2001.04383, 2020.
- [92] Renling Jin, *The sumset phenomenon*, Proc. Amer. Math. Soc. **130** (2002), no. 3, 855–861, DOI 10.1090/S0002-9939-01-06088-9. MR1866042
- [93] M. Junge, M. Navasques, C. Palazuelos, D. Perez-Garcia, V. B. Scholz, and R. F. Werner, *Connes embedding problem and Tsirelson’s problem*, J. Math. Phys. **52** (2011), no. 1, 012102, 12, DOI 10.1063/1.3514538. MR2790067
- [94] Akihiro Kanamori, *The higher infinite: Large cardinals in set theory from their beginnings*, 2nd ed., Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2009. Paperback reprint of the 2003 edition. MR2731169
- [95] A. Karrass and D. Solitar, *On a theorem of Cohen and Lyndon about free bases for normal subgroups*, Canadian J. Math. **24** (1972), 1086–1091, DOI 10.4153/CJM-1972-112-0. MR320150
- [96] Alexander S. Kechris, *Classical descriptive set theory*, Graduate Texts in Mathematics, vol. 156, Springer-Verlag, New York, 1995, DOI 10.1007/978-1-4612-4190-4. MR1321597
- [97] H. Jerome Keisler, *Ultraproducts and elementary classes*, Nederl. Akad. Wetensch. Proc. Ser. A **64** = Indag. Math. **23** (1961), 477–495. MR0140396
- [98] H. Jerome Keisler, *Limit ultrapowers*, Trans. Amer. Math. Soc. **107** (1963), 382–408, DOI 10.2307/1993808. MR148547

- [99] H. Jerome Keisler, *Good ideals in fields of sets*, Ann. of Math. (2) **79** (1964), 338–359, DOI 10.2307/1970549. MR166105
- [100] H. Jerome Keisler, *Ultraproducts and saturated models*, Nederl. Akad. Wetensch. Proc. Ser. A **67** = Indag. Math. **26** (1964), 178–186. MR0168483
- [101] H. Jerome Keisler, *Ultraproducts which are not saturated*, J. Symbolic Logic **32** (1967), 23–46, DOI 10.2307/2271240. MR218224
- [102] H. Jerome Keisler, *The ultraproduct construction*, Ultrafilters across mathematics, Contemp. Math., vol. 530, Amer. Math. Soc., Providence, RI, 2010, pp. 163–179, DOI 10.1090/conm/530/10444. MR2757537
- [103] Juliette Kennedy, Saharon Shelah, and Jouko Väänänen, *Regular ultrapowers at regular cardinals*, Notre Dame J. Form. Log. **56** (2015), no. 3, 417–428, DOI 10.1215/00294527-3132788. MR3373611
- [104] Kyriakos Keremedis, *Tychonoff products of two-element sets and some weakenings of the Boolean prime ideal theorem*, Bull. Pol. Acad. Sci. Math. **53** (2005), no. 4, 349–359, DOI 10.4064/ba53-4-1. MR2214925
- [105] Jonathan Kirby, *An invitation to model theory*, Cambridge University Press, Cambridge, 2019, DOI 10.1017/9781316683002. MR3967730
- [106] Eberhard Kirchberg, *On nonsemissplit extensions, tensor products and exactness of group C^* -algebras*, Invent. Math. **112** (1993), no. 3, 449–489, DOI 10.1007/BF01232444. MR1218321
- [107] V. L. Klee Jr., *Invariant metrics in groups (solution of a problem of Banach)*, Proc. Amer. Math. Soc. **3** (1952), 484–487, DOI 10.2307/2031907. MR47250
- [108] Péter Komjáth and Vilmos Totik, *Ultrafilters*, Amer. Math. Monthly **115** (2008), no. 1, 33–44, DOI 10.1080/00029890.2008.11920493. MR2375774
- [109] Linus Kramer, Saharon Shelah, Katrin Tent, and Simon Thomas, *Asymptotic cones of finitely presented groups*, Adv. Math. **193** (2005), no. 1, 142–173, DOI 10.1016/j.aim.2004.04.012. MR2132762
- [110] J. L. Krivine, *Sous-espaces et cônes convexes dans les espaces L_p* , PhD thesis, Centre national de la recherche scientifique, 1967.
- [111] Kenneth Kunen, *Ultrafilters and independent sets*, Trans. Amer. Math. Soc. **172** (1972), 299–306, DOI 10.2307/1996350. MR314619
- [112] Kenneth Kunen, *Weak P -points in \mathbb{N}^** , Colloquia Mathematica Societatis Janos Bolyai, **23** (1978), pp. 741–749.
- [113] Jerzy Łoś, *Quelques remarques, théorèmes et problèmes sur les classes définissables d'algèbres* (French), Mathematical interpretation of formal systems, North-Holland Publishing Co., Amsterdam, 1955, pp. 98–113. MR0075156
- [114] D. J. Lewis, *Cubic homogeneous polynomials over p -adic number fields*, Ann. of Math. (2) **56** (1952), 473–478, DOI 10.2307/1969655. MR49947
- [115] W. A. J. Luxemburg, *A general theory of monads*, Applications of Model Theory to Algebra, Analysis, and Probability (Internat. Sympos., Pasadena, Calif., 1967), Holt, Rinehart and Winston, New York, 1969, pp. 18–86. MR0244931
- [116] Saunders Mac Lane, *Categories for the working mathematician*, 2nd ed., Graduate Texts in Mathematics, vol. 5, Springer-Verlag, New York, 1998. MR1712872
- [117] Dugald Macpherson, *Model theory of finite and pseudofinite groups*, Arch. Math. Logic **57** (2018), no. 1-2, 159–184, DOI 10.1007/s00153-017-0584-1. MR3749405
- [118] Anatolii Ivanovič Mal'cev, *The metamathematics of algebraic systems. Collected papers: 1936–1967*, Studies in Logic and the Foundations of Mathematics, Vol. 66, North-Holland Publishing Co., Amsterdam-London, 1971. Translated, edited, and provided with supplementary notes by Benjamin Franklin Wells, III. MR0349383

- [119] Maryanthe Elizabeth Malliaris, *Persistence and regularity in unstable model theory*, ProQuest LLC, Ann Arbor, MI, 2009. Thesis (Ph.D.)—University of California, Berkeley. MR2713926
- [120] M. E. Malliaris, *Hypergraph sequences as a tool for saturation of ultrapowers*, J. Symbolic Logic **77** (2012), no. 1, 195–223, DOI 10.2178/jsl/1327068699. MR2951637
- [121] M. Malliaris and S. Shelah, *Cofinality spectrum theorems in model theory, set theory, and general topology*, J. Amer. Math. Soc. **29** (2016), no. 1, 237–297, DOI 10.1090/jams830. MR3402699
- [122] M. Malliaris and S. Shelah, *An example of a new simple theory*, arXiv preprint arXiv:1804.03254, 2018.
- [123] Maryanthe Malliaris and Saharon Shelah, *Keisler’s order has infinitely many classes*, Israel J. Math. **224** (2018), no. 1, 189–230, DOI 10.1007/s11856-018-1647-7. MR3799754
- [124] M. Malliaris and S. Shelah, *Keisler’s order is not simple (and simple theories may not be either)*, arXiv preprint arXiv:1906.10241, 2019.
- [125] M. Malliaris and S. Shelah, *A new look at interpretability and saturation*, Ann. Pure Appl. Logic **170** (2019), no. 5, 642–671, DOI 10.1016/j.apal.2019.01.001. MR3926500
- [126] David Marker, *Model theory: An introduction*, Graduate Texts in Mathematics, vol. 217, Springer-Verlag, New York, 2002. MR1924282
- [127] Donald A. Martin, *The axiom of determinateness and reduction principles in the analytical hierarchy*, Bull. Amer. Math. Soc. **74** (1968), 687–689, DOI 10.1090/S0002-9904-1968-11995-0. MR227022
- [128] Donald A. Martin, *Measurable cardinals and analytic games*, in Mathematical Logic in the 20th Century, World Scientific, 2003, pp. 264–268.
- [129] D. A. Martin and R. M. Solovay, *Internal Cohen extensions*, Ann. Math. Logic **2** (1970), no. 2, 143–178, DOI 10.1016/0003-4843(70)90009-4. MR270904
- [130] Donald A. Martin and John R. Steel, *Projective determinacy*, Proc. Nat. Acad. Sci. U.S.A. **85** (1988), no. 18, 6582–6586, DOI 10.1073/pnas.85.18.6582. MR959109
- [131] Antonio Martínez-Abejón, *An elementary proof of the principle of local reflexivity*, Proc. Amer. Math. Soc. **127** (1999), no. 5, 1397–1398, DOI 10.1090/S0002-9939-99-04687-0. MR1476378
- [132] A. R. D. Mathias, *Solution of problems of Choquet and Puritz*, Conference in Mathematical Logic—London ’70 (Bedford Coll., London, 1970), Springer, Berlin, 1972, pp. 204–210. Lecture Notes in Math., Vol. 255. MR0363911
- [133] Hideyuki Matsumura, *Commutative ring theory*, 2nd ed., Cambridge Studies in Advanced Mathematics, vol. 8, Cambridge University Press, Cambridge, 1989. Translated from the Japanese by M. Reid. MR1011461
- [134] Dusa McDuff, *Central sequences and the hyperfinite factor*, Proc. London Math. Soc. (3) **21** (1970), 443–461, DOI 10.1112/plms/s3-21.3.443. MR281018
- [135] Arnold W. Miller, *There are no Q -points in Laver’s model for the Borel conjecture*, Proc. Amer. Math. Soc. **78** (1980), no. 1, 103–106, DOI 10.2307/2043048. MR548093
- [136] John Milnor, *Growth of finitely generated solvable groups*, J. Differential Geometry **2** (1968), 447–449. MR244899
- [137] Deane Montgomery and Leo Zippin, *Topological transformation groups*, Interscience Publishers, New York-London, 1955. MR0073104
- [138] Itay Neeman, *Ultrafilters and large cardinals*, Ultrafilters across mathematics, Contemp. Math., vol. 530, Amer. Math. Soc., Providence, RI, 2010, pp. 181–200, DOI 10.1090/conm/530/10445. MR2757538
- [139] Alexander Yu. Ol’shanskii and Mark V. Sapir, *A finitely presented group with two non-homeomorphic asymptotic cones*, Internat. J. Algebra Comput. **17** (2007), no. 2, 421–426, DOI 10.1142/S021819670700369X. MR2310154

- [140] D. Osin and A. O. Houcine, *Finitely presented groups with infinitely many non-homeomorphic asymptotic cones*, to appear in Algebraic and Geometric Topology.
- [141] Narutaka Ozawa, *Tsirelson's problem and asymptotically commuting unitary matrices*, J. Math. Phys. **54** (2013), no. 3, 032202, 8, DOI 10.1063/1.4795391. MR3059438
- [142] Pierre Pansu, *Croissance des boules et des géodésiques fermées dans les nilvariétés* (French, with English summary), Ergodic Theory Dynam. Systems **3** (1983), no. 3, 415–445, DOI 10.1017/S0143385700002054. MR741395
- [143] Vladimir G. Pestov, *Hyperlinear and sofic groups: a brief guide*, Bull. Symbolic Logic **14** (2008), no. 4, 449–480, DOI 10.2178/bsl/1231081461. MR2460675
- [144] Bedřich Pospíšil, *Remark on bicomact spaces*, Ann. of Math. (2) **38** (1937), no. 4, 845–846, DOI 10.2307/1968840. MR1503375
- [145] Florin Rădulescu, *The von Neumann algebra of the non-residually finite Baumslag group $\langle a, b | ab^3a^{-1} = b^2 \rangle$ embeds into R^ω* , Hot topics in operator theory, Theta Ser. Adv. Math., vol. 9, Theta, Bucharest, 2008, pp. 173–185. MR2436761
- [146] F. P. Ramsey, *On a problem of formal logic*, Proc. London Math. Soc. (2) **30** (1929), no. 4, 264–286, DOI 10.1112/plms/s2-30.1.264. MR1576401
- [147] V. N. Remeslennikov, *\exists -free groups* (Russian), Sibirsk. Mat. Zh. **30** (1989), no. 6, 193–197, DOI 10.1007/BF00970922; English transl., Siberian Math. J. **30** (1989), no. 6, 998–1001 (1990). MR1043446
- [148] Abraham Robinson, *Non-standard analysis*, Princeton Landmarks in Mathematics, Princeton University Press, Princeton, NJ, 1996. Reprint of the second (1974) edition; With a foreword by Wilhelmus A. J. Luxemburg, DOI 10.1515/9781400884223. MR1373196
- [149] Walter Roelcke and Susanne Dierolf, *Uniform structures on topological groups and their quotients*, Advanced Book Program, McGraw-Hill International Book Co., New York, 1981. MR644485
- [150] Mary Ellen Rudin, *Partial orders on the types in βN* , Trans. Amer. Math. Soc. **155** (1971), 353–362, DOI 10.2307/1995690. MR273581
- [151] Volker Runde, *Lectures on amenability*, Lecture Notes in Mathematics, vol. 1774, Springer-Verlag, Berlin, 2002, DOI 10.1007/b82937. MR1874893
- [152] Mark Sapir, *On groups with locally compact asymptotic cones*, Internat. J. Algebra Comput. **25** (2015), no. 1-2, 37–40, DOI 10.1142/S0218196715400020. MR3325875
- [153] Hans Schoutens, *The use of ultraproducts in commutative algebra*, Lecture Notes in Mathematics, vol. 1999, Springer-Verlag, Berlin, 2010, DOI 10.1007/978-3-642-13368-8. MR2676525
- [154] D. Scott, *Measurable cardinals and constructible sets*, in Mathematical Logic in the 20th Century, World Scientific, 2003, pp. 407–410.
- [155] Zlil Sela, *Diophantine geometry over groups. I. Makanin-Razborov diagrams*, Publ. Math. Inst. Hautes Études Sci. **93** (2001), 31–105, DOI 10.1007/s10240-001-8188-y. MR1863735
- [156] Zlil Sela, *Diophantine geometry over groups VIII: Stability*, Annals of Mathematics, **177** (2013), pp. 787–868.
- [157] Saharon Shelah, *Every two elementarily equivalent models have isomorphic ultrapowers*, Israel J. Math. **10** (1971), 224–233, DOI 10.1007/BF02771574. MR297554
- [158] Saharon Shelah, *Saturation of ultrapowers and Keisler's order*, Ann. Math. Logic **4** (1972), 75–114, DOI 10.1016/0003-4843(72)90012-5. MR294113
- [159] Saharon Shelah, *Proper forcing*, in Proper Forcing, Springer, 1982, pp. 73–113.
- [160] Saharon Shelah, *Vive la différence I: Nonisomorphism of ultrapowers of countable models*, in Set Theory of the Continuum, Springer, 1992, pp. 357–405.
- [161] Saharon Shelah, *Proper and improper forcing*, vol. 5 of Perspectives in Logic, Cambridge University Press, 2017.

- [162] S. Shelah and M. E. Rudin, *Unordered types of ultrafilters*, *Topology Proc.* **3** (1978), no. 1, 199–204 (1979). MR540490
- [163] W. Sierpiński, *Fonctions additives non complètement additives et fonctions non mesurables*, *Fundamenta Mathematicae*, **1** (1938), pp. 96–99.
- [164] T. Skolem, *Über die Nicht-charakterisierbarkeit der Zahlenreihe mittels endlich oder abzählbar unendlich vieler Aussagen mit ausschliesslich Zahlenvariablen*, *Fundamenta mathematicae*, **23** (1934), pp. 150–161.
- [165] R. Solovay, *Measurable cardinals and the axiom of determinateness*, *Lecture Notes of UCLA Summer Institute on Axiomatic Set Theory*, Los Angeles, (1967).
- [166] R. M. Solovay, 2^{\aleph_0} can be anything it ought to be, in *The Theory of Models*, *Proceedings of the 1963 International Symposium at Berkeley*, *Studies in Logic and the Foundations of Mathematics*, North-Holland, Amsterdam, 1965, p. 435.
- [167] R. M. Solovay, *On the cardinality of σ_2^1 sets of reals*, in *Foundations of Mathematics*, Springer, 1969, pp. 58–73.
- [168] Robert M. Solovay, *A model of set-theory in which every set of reals is Lebesgue measurable*, *Ann. of Math. (2)* **92** (1970), 1–56, DOI 10.2307/1970696. MR265151
- [169] M. H. Stone, *The theory of representations for Boolean algebras*, *Trans. Amer. Math. Soc.* **40** (1936), no. 1, 37–111, DOI 10.2307/1989664. MR1501865
- [170] M. H. Stone, *Applications of the theory of Boolean rings to general topology*, *Trans. Amer. Math. Soc.* **41** (1937), no. 3, 375–481, DOI 10.2307/1989788. MR1501905
- [171] E. Szemerédi, *On sets of integers containing no k elements in arithmetic progression*, *Acta Arith.* **27** (1975), 199–245, DOI 10.4064/aa-27-1-199-245. MR369312
- [172] Terence Tao, *Hilbert’s fifth problem and related topics*, *Graduate Studies in Mathematics*, vol. 153, American Mathematical Society, Providence, RI, 2014, DOI 10.1088/0253-6102/41/3/335. MR3237440
- [173] A. Tarski, *Une contribution à la théorie de la mesure*, *Fundamenta Mathematicae*, **15** (1930), pp. 42–50.
- [174] Katrin Tent and Martin Ziegler, *A course in model theory*, *Lecture Notes in Logic*, vol. 40, Association for Symbolic Logic, La Jolla, CA; Cambridge University Press, Cambridge, 2012, DOI 10.1017/CBO9781139015417. MR2908005
- [175] Guy Terjanian, *Un contre-exemple à une conjecture d’Artin* (French), *C. R. Acad. Sci. Paris Sér. A-B* **262** (1966), A612. MR197450
- [176] Simon Thomas and Boban Velickovic, *Asymptotic cones of finitely generated groups*, *Bull. London Math. Soc.* **32** (2000), no. 2, 203–208, DOI 10.1112/S0024609399006621. MR1734187
- [177] Stevo Todorćević, *Introduction to Ramsey spaces*, *Annals of Mathematics Studies*, vol. 174, Princeton University Press, Princeton, NJ, 2010, DOI 10.1515/9781400835409. MR2603812
- [178] Douglas Ulrich, *Keisler’s order is not linear, assuming a supercompact*, *J. Symb. Log.* **83** (2018), no. 2, 634–641, DOI 10.1017/jsl.2018.1. MR3835081
- [179] Lou van den Dries, *Lectures on the model theory of valued fields*, *Model theory in algebra, analysis and arithmetic*, *Lecture Notes in Math.*, vol. 2111, Springer, Heidelberg, 2014, pp. 55–157, DOI 10.1007/978-3-642-54936-6_4. MR3330198
- [180] L. van den Dries and K. Schmidt, *Bounds in the theory of polynomial rings over fields. A nonstandard approach*, *Invent. Math.* **76** (1984), no. 1, 77–91, DOI 10.1007/BF01388493. MR739626
- [181] L. van den Dries and A. J. Wilkie, *Gromov’s theorem on groups of polynomial growth and elementary logic*, *J. Algebra* **89** (1984), no. 2, 349–374, DOI 10.1016/0021-8693(84)90223-0. MR751150
- [182] Nik Weaver, *Forcing for mathematicians*, *World Scientific Publishing Co. Pte. Ltd.*, Hackensack, NJ, 2014, DOI 10.1142/8962. MR3184751

- [183] Benjamin Weiss, *Sofic groups and dynamical systems*, Sankhyā Ser. A **62** (2000), no. 3, 350–359. Ergodic theory and harmonic analysis (Mumbai, 1999). MR1803462
- [184] Fred B. Wright, *A reduction for algebras of finite type*, Ann. of Math. (2) **60** (1954), 560—570, DOI 10.2307/1969851. MR65037

Index

- Δ_r -set, 55
- Δ_r^* -set, 55
- α -inaccessible cardinal, 313
- ϵ -isomorphism, 254
- κ -regularizing set, 129
- κ -saturated structure, 124
- κ -universal structure, 126
- C^* -algebra, 104, 260
 - faithful tracial state on a , 267
 - projection in a , 262
 - real rank 0, 263
 - tracial state on a , 267
 - unitary in a , 275
- p -adic integers, 119
- z -filter, 36
- z -ultrafilter, 33, 36
- Łoś's theorem
 - for \mathcal{L}_κ , 321
 - for continuous logic, 206
 - for first-order logic, 87

- abelianization of a group, 212
- amenable group, 240
- approximate subgroup, 229
 - property, 229
- Arrow's theorem, 19, 20
- Artin's conjecture, 121
- asymptotic cone, 221
- Ax's theorem, 110
- Ax-Kochen-Ershov theorem, 121

- axiom of determinacy, 68
 - for a pointclass, 68

- Baire category theorem, 64
- Baire measurable set, 64
- Baire space, 63
- Banach $*$ -algebra, 259
- Banach algebra, 253
- Banach density, 50, 57
- Banach space, 250
 - complemented subspace of a , 257
 - dual space of a , 255
 - reflexive, 256
 - super-reflexive, 256
 - uniformly convex, 250
 - weak*-topology on a , 258
- block decisive set of voters, 23
- Boolean algebra, 39
 - concrete, 39
- Borel determinacy, 69
- Borel sets, 63
- bounded linear transformation, 252
 - adjoint of a , 260

- Cantor space, 63
- category, 357
- club filter, 313
- club subset of a cardinal, 313
- cofinality, 352
- comeager, 64
- commutative transitive group, 184

- compactification, 32
- compactness theorem, 89
- complete embedding, 164
- complete expansion, 164
- complete extension, 164
- complete structure, 164
- completely regular space, 34
- completely separated, 37
- computability theory, 318
- Connes embedding problem, 274
- consistency strength, 310
- consistent pair, 142, 298
- consistent triple, 300
- constructible set, 111, 116
- continuous language, 203
- continuous logic, 201
- continuum hypothesis, 13, 72
- critical point, 329
- CSA group, 185

- decisive set of voters, 22
- definable
 - function, 111
 - set, 111
- density character of a metric space, 200
- derived subgroup, 212
- descriptive set theory, 62
- determined game, 67
- diagonal embedding, 84
 - for metric ultrapowers, 197
- diagram, 345
 - elementary, 345
- dictator, 20
- distribution, 136
 - accurate, 137
- election, 19
 - state of a , 20
- election procedure, 20
 - fair, 20
- elementary class, 305
- elementary equivalence, 344
- embedding, 345
 - elementary, 345
- equivalent categories, 360
- Erdős cardinal, 334
- existentially closed structure, 128
- expansion, 342
- external set, 170

- Følner sequence, 240
- Følner set, 240
- faithfully flat ring extension, 115
- Feferman-Vaught theorem, 104
- filter, 3
 - base for a , 4
 - Fréchet, 4
 - generated by a base, 4
 - on a Boolean algebra, 39
- finite cover property (fcp), 147
- finite intersection property (FIP), 4
- finitely representable, 254
- flat ring extension, 112
- formulae, 343
 - atomic, 343
- Frayne's theorem, 131
- FS-set, 47
- fully residually free group, 186
- functor, 359
- Furstenberg's correspondence
 - principle, 53
- Furstenberg's recurrence theorem, 54

- galaxy, 161
- Gale-Stewart theorem, 334
- gap, 284
- Gelfand duality, 104, 261
- geodesic metric space, 224
- GNS construction, 271
- Greenleaf-Ax-Kochen theorem, 120
- Gromov's theorem on polynomial
 - growth, 216
- growth function of a group, 213
 - exponential growth, 215
 - near polynomial growth, 226
 - polynomial growth, 214

- Hausdorff dimension, 227
- Heisenberg group, 212
- Hilbert space, 252
- Hilbert-Schmidt metric, 235
- Hindman's theorem, 47
 - strong version, 49
- homogeneous metric space, 223
- hyperfinite II_1 factor, 274

- hyperfinite set, 173
 - internal cardinality of a , 173
- hyperlinear group, 247
- hyperreal numbers, 158
 - finite elements, 159
 - infinite elements, 159
 - infinitesimal elements, 158, 159
- ideal on a set, 64
 - σ -, 64
- idempotent hyperinteger, 175
- independence property, 153
- independence relation, 293
- independent sequence, 293
- indicable group, 212
 - kernel of, 212
- indiscernible sequence, 293
- induced embedding between
 - ultrapowers, 165
- inner product space, 252
- internal definition principle, 171
- internal set, 170, 198
- iterated ultrapower, 98, 99, 166
- Jin's sunset theorem, 55, 56
- Keisler's order, 146, 288
- Keisler-Shelah theorem, 100, 141, 297
- Kervaire-Lajudembach conjecture, 246
- KL-group, 246
 - strong, 246
- Kleene-Brouwer ordering, 335
- language, 341
- large oscillation, 142
- Lebesgue measurable set, 63
- Lie group, 218
- limit group, 182
- limit ultrapower, 166
- Lindenbaum algebra, 82
- local ring, 117
 - henselian, 119
- lower cofinality, 285
- Mahlo cardinal, 314
- malnormal subgroup, 185
- marked group(s), 189
 - isomorphic, 189
 - space of, 189
 - ultraproducts of, 190
- meager set, 64
- mean, 51
 - invariant, 52
- measurable cardinal, 69, 93, 314
- measure-preserving transformation, 53
- metric group, 233
 - bi-invariant, 234
 - uniform, 233
- metric structure, 201, 203
- Milnor-Wolf theorem, 216
- modulus of uniform continuity, 202
- Morley sequence, 293
- natural transformation, 360
- nfcf, 147, 283
- nilpotent group, 212
- nonstandard hull, 198
- normalized Hamming metric, 235
- normed algebra, 253
 - unital, 253
- normed space(s), 250
 - ultraproduct of, 250
- nowhere dense set, 64
- null set, 63
- operator norm, 252
- ordinal, 350
- orthogonal projection, 262
- overflow principle, 171
- P-point, 73
- partition regular property, 49, 56
- perfectly normal space, 35
- piecewise syndetic set, 55, 57
- point-homogeneous space, 75
- pointed metric space, 196
- Polish space, 62, 63
- polycyclic group, 216
- powerset algebra, 39
- principle of local reflexivity, 257
- projective set, 65
- proper metric space, 198
- pseudo-intersection, 72

- Ramsey cardinal, 325
- Ramsey's theorem
 - finite version, 89
 - infinite version, 46
- reduced power
 - of sets, 84
 - of structures, 85
- reduced product
 - of metric spaces, 208
 - of metric structures, 209
 - of sets, 84
 - of structures, 85
- reduct, 342
- regular cardinal, 352
- regular equation, 246
- residually finite group, 243
- residually free group, 186
- residue field, 117
- Robinson's joint consistency
 - theorem, 306
- Rudin-Keisler order, 14, 94, 98

- S-subsemigroup, 175
- S-topology, 174
- saturated structure, 124
- Scott's theorem, 330
- sentence, 344
- singular cardinal, 352
- sofic group, 235
- solvable group, 212
- stable
 - formula, 282
 - structure, 283
 - theory, 149, 283
- standard part, 160
- stationary subset of a cardinal, 313
- Stone duality theorem, 43
- Stone representation theorem, 38, 44
- Stone space, 38
 - of a Boolean algebra, 40
- Stone-Čech compactification, 31, 32
- strict order property, 153
- strong limit cardinal, 311
- strong operator topology, 265
- strongly compact cardinal, 321
- strongly inaccessible cardinal, 311
- structure, 341
- substructure, 345
 - elementary, 345
- supercompact cardinal, 323
- supercompact elementary
 - embedding, 331
- syndetic set, 54
- Szemerédi's theorem, 54

- tail set, 64
- terms, 342
- theory, 344
 - complete, 344
 - of a class of structures, 305
 - satisfiable, 344
- thick set, 50
- tracial ultraproduct, 270
- tree, 335
 - ill founded, 335
 - infinite branch in a , 335
 - root of a , 335
 - well founded, 335
- Turing degree(s), 318
 - cone of, 318
- Turing reducible, 318
- Tychonoff space, 34
- Tychonoff's theorem, 30, 61
- type, 292

- ultrafilter theorem, 6
 - for Boolean algebras, 40
 - intermediate version, 60
 - weak version, 60
- ultrafilter(s), 4
 - κ -regular, 129
 - κ^+ -good, 136
 - countably complete, 92
 - countably incomplete, 92
 - fine, 322
 - game, 66
 - good, 136
 - Hausdorff, 177
 - idempotent, 48, 61, 75
 - induced, 5, 11
 - isomorphic, 9
 - minimal, 70, 95
 - morphism between, 8
 - nonprincipal, 5
 - normal, 316, 322

- number, 13, 15
- on a Boolean algebra, 39
- principal, 5
- product of, 16, 98
- pushforward, 8, 41
- quantifier, 6
- quasi-normal, 71
- Ramsey, 70
- regular, 101, 129
- saturating a theory, 146
- selective, 70
- strongly summable, 75
- sum of, 47
- uniform, 6
- union, 76
- weakly selective, 73
- weakly summable, 75
- ultralimit, 27
- ultrapower
 - chain, 100
 - extension, 100
 - of sets, 84
 - of structures, 85
 - of the set-theoretic universe, 327
 - system, 166
- ultraproduct
 - embedding, 88
 - of metric structures, 204
 - of sets, 84
 - of structures, 85
- uncountably categorical theory, 146
- underflow principle, 171
- universal structure, 126
- universal theory, 132
- universally free group, 182

- virtual property of a group, 212
- von Neumann algebra, 265
 - normal tracial state on a, 271
 - of a group, 266
 - trace on a, 271
 - tracial, 271

- weakly compact cardinal, 324
- weakly inaccessible cardinal, 311
- Woodin cardinals, 337
- word-length function, 213
- worldly cardinal, 310

- Zariski closed set, 111
- zero-dimensional, 31
- zeroset, 35


Selected Published Titles in This Series

- 220 **Isaac Goldbring**, *Ultrafilters Throughout Mathematics*, 2022
219 **Michael Joswig**, *Essentials of Tropical Combinatorics*, 2021
218 **Riccardo Benedetti**, *Lectures on Differential Topology*, 2021
217 **Marius Crainic, Rui Loja Fernandes, and Ioan Mărcuț**, *Lectures on Poisson Geometry*, 2021
216 **Brian Osserman**, *A Concise Introduction to Algebraic Varieties*, 2021
215 **Tai-Ping Liu**, *Shock Waves*, 2021
214 **Ioannis Karatzas and Constantinos Kardaras**, *Portfolio Theory and Arbitrage*, 2021
213 **Hung Vinh Tran**, *Hamilton–Jacobi Equations*, 2021
212 **Marcelo Viana and José M. Espinar**, *Differential Equations*, 2021
211 **Mateusz Michałek and Bernd Sturmfels**, *Invitation to Nonlinear Algebra*, 2021
210 **Bruce E. Sagan**, *Combinatorics: The Art of Counting*, 2020
209 **Jessica S. Purcell**, *Hyperbolic Knot Theory*, 2020
208 **Vicente Muñoz, Ángel González-Prieto, and Juan Ángel Rojo**, *Geometry and Topology of Manifolds*, 2020
207 **Dmitry N. Kozlov**, *Organized Collapse: An Introduction to Discrete Morse Theory*, 2020
206 **Ben Andrews, Bennett Chow, Christine Guenther, and Mat Langford**, *Extrinsic Geometric Flows*, 2020
205 **Mikhail Shubin**, *Invitation to Partial Differential Equations*, 2020
204 **Sarah J. Witherspoon**, *Hochschild Cohomology for Algebras*, 2019
203 **Dimitris Koukoulopoulos**, *The Distribution of Prime Numbers*, 2019
202 **Michael E. Taylor**, *Introduction to Complex Analysis*, 2019
201 **Dan A. Lee**, *Geometric Relativity*, 2019
200 **Semyon Dyatlov and Maciej Zworski**, *Mathematical Theory of Scattering Resonances*, 2019
199 **Weinan E, Tiejun Li, and Eric Vanden-Eijnden**, *Applied Stochastic Analysis*, 2019
198 **Robert L. Benedetto**, *Dynamics in One Non-Archimedean Variable*, 2019
197 **Walter Craig**, *A Course on Partial Differential Equations*, 2018
196 **Martin Stynes and David Stynes**, *Convection-Diffusion Problems*, 2018
195 **Matthias Beck and Raman Sanyal**, *Combinatorial Reciprocity Theorems*, 2018
194 **Seth Sullivant**, *Algebraic Statistics*, 2018
193 **Martin Lorenz**, *A Tour of Representation Theory*, 2018
192 **Tai-Peng Tsai**, *Lectures on Navier-Stokes Equations*, 2018
191 **Theo Bühler and Dietmar A. Salamon**, *Functional Analysis*, 2018
190 **Xiang-dong Hou**, *Lectures on Finite Fields*, 2018
189 **I. Martin Isaacs**, *Characters of Solvable Groups*, 2018
188 **Steven Dale Cutkosky**, *Introduction to Algebraic Geometry*, 2018
187 **John Douglas Moore**, *Introduction to Global Analysis*, 2017
186 **Bjorn Poonen**, *Rational Points on Varieties*, 2017
185 **Douglas J. LaFountain and William W. Menasco**, *Braid Foliations in Low-Dimensional Topology*, 2017
184 **Harm Derksen and Jerzy Weyman**, *An Introduction to Quiver Representations*, 2017
183 **Timothy J. Ford**, *Separable Algebras*, 2017
182 **Guido Schneider and Hannes Uecker**, *Nonlinear PDEs*, 2017
181 **Giovanni Leoni**, *A First Course in Sobolev Spaces, Second Edition*, 2017

For a complete list of titles in this series, visit the
AMS Bookstore at www.ams.org/bookstore/gsmseries/.


Ultrafilters and ultraproducts provide a useful generalization of the ordinary limit processes which have applications to many areas of mathematics. Typically, this topic is presented to students in specialized courses such as logic, functional analysis, or geometric group theory. In this book, the basic facts about ultrafilters and ultraproducts are presented to readers with no prior knowledge of the subject and then these techniques are applied to a wide variety of topics. The first part of the book deals solely with ultrafilters and presents applications to voting theory, combinatorics, and topology, while also dealing also with foundational issues. The second part presents the classical ultraproduct construction and provides applications to algebra, number theory, and nonstandard analysis. The third part discusses a metric generalization of the ultraproduct construction and gives example applications to geometric group theory and functional analysis. The final section returns to more advanced topics of a more foundational nature.

The book should be of interest to undergraduates, graduate students, and researchers from all areas of mathematics interested in learning how ultrafilters and ultraproducts can be applied to their specialty.



© 2022 Laurel Hungerford Photography

ISBN 978-1-4704-6900-9



9 781470 469009

GSM/220.H



For additional information
and updates on this book, visit

www.ams.org/bookpages/gsm-220

