Categoricity

John Baldwin

August 19, 2008

Contents

1
Introduction

Modern model theory began with Morley's [110] categoricity theorem: A first order theory is categorical in one uncountable cardinal κ (has a unique model of that cardinality) if and only if it is categorical in all uncountable cardinals. This result triggered the change in emphasis from the study of logics to the study of theories. Shelah's taxonomy of first order theories by the stability classification established the background for most model theoretic researches in the last 35 years. This book lays out of some of the developments in extending this analysis to classes that are defined in non-first order ways. Inspired by [117, 80], we proceed via short chapters that can be covered in a lecture or two.

There were three streams of model-theoretic research in the 1970's. For simplicity in the discussion below I focus on vocabularies (languages) which contain only countably many relation and function symbols. In one direction workers in algebraic model theory melded sophisticated algebraic studies with techniques around quantifier elimination and developed connections between model theory and algebra. A second school developed fundamental model theoretic properties of a wide range of logics. Many of these logics were obtained by expanding first order logic by allowing longer conjunctions or longer strings of first order quantifiers; others added quantifiers for 'there exist infinitely many', 'there exist uncountably many', 'equicardinality', and many other concepts. This work was summarized in the Barwise-Feferman volume [24]. The use of powerful combinatorial tools such as the Erdös-Rado theorem on the one hand and the discovery that Chang's conjecture on two cardinal models for arbitrary first theories is independent of ZFC and that various two cardinal theorems are connected to the existence of large cardinals [34] caused a sense that pure model theory was deeply entwined both with heavy set-theoretic combinatorics and with (major) extensions of ZFC. In the third direction, Shelah made the fear of independence illusory for the most central questions by developing the stability hierarchy. He split all first order theories into 5 classes. Many interesting algebraic structures fall into the three classes (ω-stable, superstable, strictly stable) whose models admit a deep structural analysis. This classification is (set theoretically) absolute as are various fundamental properties of such theories. Thus, for stable theories, Chang's conjecture is proved in ZFC [93, 122]. Shelah focused his efforts on the test question: compute the function $I(T, \kappa)$ which gives the number of models of cardinality κ. He achieved the striking main gap theorem. Every theory T falls into one of two classes. T may be intractable, that is $I(T, \kappa) = 2^\kappa$, the maximum, for every sufficiently large κ. Or, every model of T is decomposed as a tree of countable models and the number of models in κ is bounded well below 2^κ. The description of this tree and the proof of the theorem required the development of a far reaching generalization of the Van der Waerden axiomatization of independence in vector spaces and fields. This is not the place for even a cursory survey of the development of stability theory over the last 35 years. However, the powerful tools of the Shelah's calculus of independence and orthogonality are fundamental to the applications of model theory in the 1990's to Diophantine geometry and number theory [30].

Since the 1970's Shelah has been developing the intersection of the second and third streams described above: the model theory of the class of models of a sentence in one of a number of 'non-elementary' logics. He builds on Keisler's work [80] for the study of $L_{\omega_1,\omega}$ but to extend to other logics he needs a more general framework and the Abstract Elementary Classes (AEC) we discuss below provide one. In the last ten years, the need for such a study has become more widely appreciated as a result of work on both such concrete problems as complex exponentiation and Banach spaces and programmatic needs to understand 'type-definable' groups and to understand an analogue to 'stationary types' in simple theories.

Our goal here is to provide a systematic and intelligible account of some central aspects of Shelah's work and related developments. We study some very specific logics (e.g. $L_{\omega_1,\omega}$) and the very general case of abstract elementary classes. The survey articles by Grossberg [48] and myself [9] provide further background and motivation for the study of AEC that is less technical than the development here.

An abstract elementary class (AEC) $\boldsymbol{K}$ is a collection of models and a notion of 'strong submodel' $\prec$ which satisfies general conditions similar to those satisfied by the class of models of a first order theory with $\prec$ as elementary submodel. In particular, the class is closed under unions of $\prec$-chains. A Löwenheim-Skolem number is assigned to each AEC: a cardinal κ such that each $M \in \boldsymbol{K}$ has a strong submodel of cardinality κ. Examples include the class of models of a $\forall\exists$ first order theory with $\prec$ as substructure, a complete sentence of $L_{\omega_1,\omega}$ with $\prec$ as elementary submodel in an appropriate fragment of $L_{\omega_1,\omega}$ and the class of submodels of a homogeneous model with $\prec$ as elementary submodel. The models of a sentence of $L_{\omega_1,\omega}(Q)$ (Q is the quantifier 'there exists uncountably many') fit into this context modulo two important restrictions. An artificial notion of 'strong submodel' must be introduced to guarantee the satisfaction of the axioms concerning unions of chains. More important from a methodological viewpoint, without still further and unsatisfactory contortions, the Löwenheim number of the class will be $\aleph_1$.

In general the analysis is not nearly as advanced as in the first order case. We have only approximations to Morley's theorem and only a rudimentary development of stability theory. (There have been significant advances under more specialized assumptions such as homogeneity or excellence [49, 71] and other works of e.g. Grossberg, Hyttinen, Lessmann, and Shelah.) The most dispositive result is Shelah's proof that assuming $2^{\aleph_n} < 2^{\aleph_{n+1}}$ for $n < \omega$, if a sentence of $L_{\omega_1,\omega}$ is categorical up to $\aleph_\omega$ then is categorical in all cardinals. Categoricity up to $\aleph_\omega$ is essential [56, 5].

The situation for AEC is even less clear. One would like at least to show that an AEC could not alternate indefinitely between categoricity and non-categoricity. The strongest result we show here is from [128]: if an AEC $\boldsymbol{K}$ is categorical on a proper class of successor cardinals and has the amalgamation property then it is eventually categorical.

This state of affairs in a major reason that this monograph is titled categoricity. Although a general stability theory for abstract elementary classes is the ultimate goal, the results here depend heavily on assuming categoricity in at least one cardinal.

There are several crucial aspects of first order model theory. By Lindström's theorem [102] we know they can be summarized as: first order logic is the only logic (of Lindström type) with Löwenheim number $\aleph_0$ that satisfies the compactness theorem. One corollary of compactness in the first order case plays a distinctive role here, the amalgamation property: two elementary extensions of a fixed model M have a common extension over M. In particular, the first order amalgamation property allows the identification (in a suitable monster model) of a syntactic type (the description of a point by the formulas it satisfies) with an orbit under the automorphism group (we say Galois type).

Some of the results here and many associated results were originally developed using considerable extensions to ZFC. However, later developments and the focus on AEC rather than $L_{\kappa,\omega}$ (for specific large cardinals κ) have reduced such reliance. All results in this book are proved in ZFC or in ZFC + $2^n < 2^{n+1}$ for finite n; we call this proposition the very weak generalized continuum hypothesis VWGCH and this assertion with n replaced by any cardinal μ is the WGCH. Without this assumption, some crucial results have not been proved in ZFC; the remarkable fact is that such a benign assumption as VWGCH is all that is required. Some of the uses of stronger set theory to analyze categoricity of $L_{\omega_1,\omega}$-sentences can be avoided by the assumption that the class of models considered contains arbitrarily large models.

We now survey the material with an attempt to convey the spirit and not the letter of various important concepts; precise versions are in the text. With a few exceptions that are mentioned at the time all the work expounded here was first discovered by Shelah in a long series of papers.

Part I (Chapters 2-4) contains a discussion of Zilber's quasiminimal excellent classes [150]. This is a natural generalization of the study of first order strongly minimal theories to the logic $L_{\omega_1,\omega}$ (and some fragments of $L_{\omega_1,\omega}(Q)$. It clearly exposes the connections between categoricity and homogeneous combinatorial geometries; there are natural algebraic applications to the study of various expansions of the complex numbers. We expound a very concrete notion of 'excellence' for a combinatorial geometry. Excellence describes the closure of an independent n-cube of models. This is a fundamental structural property of countable structures in a class $\boldsymbol{K}$ which implies that $\boldsymbol{K}$ has arbitrarily large models (and more). Zilber's contribution is to understand the connections of these ideas to concrete mathematics, to recognize the relevance of infinitary logic to these concrete problems, and to prove that his particular examples are excellent. These applications require both great insight in finding the appropriate formal context and substantial mathematical work in verifying the conditions laid down. Moreover, his work has led to fascinating speculations in complex analysis and number theory. As pure model theory of $L_{\omega_1,\omega}$, these results and concepts were all established in greater generality by Shelah [124] more than twenty years earlier. But Zilber's work extends Shelah's analysis in one direction by applying to some extensions of $L_{\omega_1,\omega}$. We explore the connections between these two approaches at the end of Chapter 26. Before turning to that work, we discuss an extremely general framework.

The basic properties of abstract elementary classes are developed in Part II (chapters 5-8). In particular, we give Shelah's presentation theorem which represents every AEC as a pseudoelementary class (class of reducts to a vocabulary L of a first order theory in an expanded language L') that omit a set of types. Many of the key results (especially in Part IV) depend on having Löwenheim number $\aleph_0$. Various successes and perils of translating $L_{\omega_1,\omega}(Q)$ to an AEC (with countable Löwenheim number) are detailed in Chapters 6-8 along with the translation of classes defined by sentences of $L_{\omega_1,\omega}$ to the class of atomic models of a first order theory in an expanded vocabulary. Chapter 8 contains Shelah's beautiful ZFC proof that a sentence of $L_{\omega_1,\omega}(Q)$ that is $\aleph_1$-categorical has a model of power $\aleph_2$.

In Part III (Chapters 10-18) we first study the conjecture that for 'reasonably well-behaved classes', categoricity should be either eventually true or eventually false. We formalize 'reasonably well-behaved' via two crucial hypotheses: amalgamation and the existence of arbitrarily large models. Under these assumptions, the notion of *Galois type over a model* is well-behaved and we recover such fundamental notions as the identification of 'saturated models' with those which are 'model homogeneous'. Equally important, we are able to use the omitting types technology originally developed by Morley to find Ehrenfeucht-Mostowski models for AEC. This leads to the proof that categoricity implies stability in smaller cardinalities and eventually, via a more subtle use of Ehrenfeucht-Mostowski models, to a notion of superstability. The first goal of these chapters is to expound Shelah's proof of a downward categoricity theorem for an AEC (satisfying the above hypothesis) and categorical in a successor cardinal. A key aspect of that argument is the proof that if $\boldsymbol{K}$ is categorical above the Hanf number for

AEC's, then two distinct Galois types differ on a 'small' submodel. Grossberg and VanDieren [46] christened this notion: tame.

We refine the notion of tame in Chapter 12 and discuss three properties of Galois types: tameness, locality, and compactness. Careful discussion of these notions requires the introduction of cardinal parameters to calibrate the notion of 'small'. We analyze this situation and sketch examples related to the Whitehead conjecture showing how non-tame classes can arise. Grossberg and VanDieren develop the theory for AEC satisfying very strong tameness hypotheses. Under these conditions they showed categoricity could be transferred upward from categoricity in two successive cardinals. Key to obtaining categoricity transfer from one cardinal λ^+ is the proof that the union of a 'short' chain of saturated models of cardinality λ is saturated. This is a kind of superstability consideration; it requires a further and still more subtle use of the Ehrenfeucht-Mostowski technology and a more detailed analysis of splitting; this is carried out in Chapter 16.

In Chapters 17 and 18 we conclude Part III and explore AEC without assuming the amalgamation property. We show, under mild set-theoretic hypotheses (weak diamond), that an AEC which is categorical in κ and fails the amalgamation property for models of cardinality κ has many models of cardinality κ^+.

In Part IV (Chapters 19-27) we return to the more concrete situation of atomic classes, which, of course, encompasses $L_{\omega_1,\omega}$. Using $2^{\aleph_0} < 2^{\aleph_1}$, one deduces from a theorem of Keisler [80] that an $\aleph_1$-categorical sentence ψ in $L_{\omega_1,\omega}$ is ω-stable. Note however that ω-stability is proved straightforwardly (Chapter 7) if one assumes ψ has arbitrarily large models. In Chapters 19-24, we introduce an independence notion and develop excellence for atomic classes. Assuming cardinal exponentiation is increasing below $\aleph_\omega$, we prove a sentence of $L_{\omega_1,\omega}$ that is categorical up to $\aleph_\omega$ is excellent. In Chapters 25-26 we report Lessmann's [99] account of proving Baldwin-Lachlan style characterizations of categoricity for $L_{\omega_1,\omega}$ and Shelah's analog of Morley's theorem for excellent atomic classes. We conclude Chapter 26, by showing how to deduce the categoricity transfer theorem for arbitrary $L_{\omega_1,\omega}$-sentences from a (stronger) result for complete sentences. Finally, in the last chapter we explicate the Hart-Shelah example of an $L_{\omega_1,\omega}$-sentence that is categorical up to $\aleph_n$ but not beyond and use it to illustrate the notion of tameness.

The work here has used essentially in many cases that we deal with classes with Löwenheim number $\aleph_0$. Thus, in particular, we have proved few substantive general results concerning $L_{\omega_1,\omega}(Q)$ (the existence of a model in $\aleph_2$ is a notable exception). Shelah has substantial not yet published work attacking the categoricity transfer problem in the context of 'frames'; this work does apply to $L_{\omega_1,\omega}(Q)$ and does not depend on Löwenheim number $\aleph_0$. We do not address this work [134, 131, 130] nor related work which makes essential use of large cardinals ([105, 88].

A solid graduate course in model theory is an essential prerequisite for this book. Nevertheless, the only quoted material is very elementary model theory (say a small part of Marker's book [108]), and two or three theorems from the Keisler book [80] including the Lopez-Escobar theorem characterizing well-orderings.

We include in Appendix A a full account of the Hanf number for omitting types. In Appendix B we give the Keisler technology for omitting types in uncountable models. The actual combinatorial principle that extends ZFC and is required for the results here is the Devlin-Shelah weak diamond. A proof of the weak diamond from weak GCH below $\aleph_\omega$ appears in Appendix C. In Appendix D we discuss a number of open problems. Other natural background reference books are [108, 60, 122, 34].

The foundation of all this work is Morley's theorem [110]; the basis for transferring this result to infinitary logic is [80]. In addition to the fundamental papers of Shelah, this exposition depends heavily on various works by Grossberg, Lessmann, Makowski, VanDieren, and Zilber and on conversations with Adler, Coppola, Dolich, Drueck, Goodrick, Hart, Hyttinen, Kesala, Kirby, Kolesnikov, Kueker, Laskowski, Marker, Shelah, and Shkop as well as these authors. The book would never have happened if not for the enthusiasm and support of Rami Grossberg, Monica VanDieren and Andres Villaveces. They brought the subject alive for me and four conferences in Bogota and the 2006 AIM meeting on Abstract Elementary Classes were essential to my understanding of the area. Grossberg, in particular, was a unending aid in finding my way. I thank the logic seminar at the University of Barcelona and especially Enriques Casanovas for the opportunity to present Part IV in the Fall of 2006 and for their comments.

Part I

Quasiminimal Excellence and Complex Exponentiation

In Part I of the book, we provide a concrete example of the applicability of the abstract methods to a mathematical problem. That is, we describe sufficient conditions for categoricity in all powers and elaborate a specific algebraic example. In Parts II and III we study more general contexts where only partial categoricity transfer methods are known. We close the book in Part IV by describing Shelah's solution of the categoricity transfer problem for $L_{\omega_1,\omega}$ and placing Part I in that context.

One reason to study AEC is that first order logic is inadequate to describe certain basic mathematical structures. One example is complex exponentiation; the ring of integers is interpretable in $(\mathcal{C}, \cdot, +, \exp)$ with domain the zeros of the sine function and so the resulting theory is unstable. Zilber suggested investigating this structure in $L_{\omega_1,\omega}$. A somewhat simpler object, the covers of the multiplicative group of $\mathcal{C}$, is a concrete example of a quasiminimal excellent classes. This is a specific example of the notion of excellence discussed in Part IV. A notion of excellence presupposes a notion of dependence. In the case of quasiminimality, this dependence notion determines a geometry. We prove that a quasiminimal excellent class is categorical in all uncountable powers. We expound several papers of Zilber and explain the fundamental notion of excellence in the 'rank one' case. We study the 'covers' situation in detail. It provides a 'real' example of categorical structures which are not homogenous and thus shows why the notion of excellence must be introduced. We briefly discuss (with references) at the end of Chapter 4 the extensions of these methods to other mathematical situations (complex exponentiation and semi-abelian varieties). We return to the more general consideration of categoricity of $L_{\omega_1,\omega}$ in Part IV and discuss the full notion of excellence. After Proposition 26.20, we examine which parts of Zilber's work are implied by Shelah [124, 125] and the sense in which Zilber extends Shelah.

2
Combinatorial Geometries and Infinitary Logics

In this chapter we introduce two of the key concepts that are used throughout the text. In the first section we define the notion of a combinatorial geometry and describe its connection with categoricity in the first order case. The second section establishes the basic notations for infinitary logics.

2.1 Combinatorial Geometries

Shelah's notion of forking provides a notion of independence which greatly generalizes the notion of Van der Waerden [143] by allowing one to study structures with a family of dimensions. (See Chapter II of [8].) But Van der Waerden's notion is the natural context for studying a single dimension; in another guise such a relation is called a combinatorial geometry. Most of this monograph is devoted to structures which are determined by a single dimension; in the simplest case, which we study first, the entire universe is the domain of a single combinatorial geometry.

One strategy for the analysis of $\aleph_1$-categorical first order theories in a countable vocabulary proceeds in two steps: a) find a definable set which admits a nice dimension theory b) show this set determines the model up to isomorphism. The sufficient condition for a set D to have a nice dimension theory is that D be strongly minimal or equivalently that algebraic closure on D forms a pregeometry in the sense described below. It is natural to attempt to generalize this approach to study categoricity in non-elementary contexts. In this chapter we review the first order case to set the stage. For more detail on this chapter consult e.g. [32].

In the next few chapters, we study the infinitary analog to the first order concept of strong minimality – quasiminimality. In Part IV we show that the strategy of reducing categoricity to a quasiminimal excellent subset can be carried out for sentences of $L_{\omega_1,\omega}$.

Definition 2.1.1. *A* pregeometry *is a set G together with a dependence relation*

$$cl : \mathcal{P}(G) \to \mathcal{P}(G)$$

satisfying the following axioms.
A1. $cl(X) = \bigcup\{cl(X') \colon X' \subseteq_{fin} X\}$
A2. $X \subseteq cl(X)$
A3. $cl(cl(X)) = cl(X)$
A4. *(Exchange) If $a \in cl(Xb)$ and $a \notin cl(X)$, then $b \in cl(Xa)$.*
If points are closed the structure is called a geometry.

Definition 2.1.2. *A geometry is* homogeneous *if for any closed $X \subseteq G$ and $a, b \in G - X$ there is a permutation of G which preserves the closure relation (i.e. an automorphism of the geometry), fixes X pointwise, and takes a to b.*

Exercise 2.1.3. *If G is a homogeneous geometry, X, Y are maximally independent subsets of G, there is an automorphism of G taking X to Y.*

The most natural examples of homogeneous geometries are vector spaces and algebraically closed fields with their usual notions of closure. The crucial properties of these examples are summarised in the following definition.

Definition 2.1.4.
1. *The structure M is* strongly minimal *if every first order definable subset of any elementary extension M' of M is finite or cofinite.*
2. *The theory T is strongly minimal if it is the theory of a strongly minimal structure.*
3. *$a \in \mathrm{acl}(X)$ if there is a first order formula with parameters from X and with finitely many solutions that is satisfied by a.*

Definition 2.1.5. *Let X, Y be subsets of a structure M. An elementary isomorphism from X to Y is* 1-1 *map from X onto Y such that for every first order formula $\phi(\mathbf{v})$, $M \models \phi(\mathbf{x})$ if and only if $M \models \phi(f\mathbf{x})$.*

Note that if M is the structure (ω, S) of the natural numbers and the successor function, then (M, S) is isomorphic to $(\omega - \{0\}, S)$. But this isomorphism is not elementary.

The next exercise illustrates a crucial point. The argument depends heavily on the exact notion of algebraic closure; the property is not shared by all combinatorial geometries. In the quasiminimal case, discussed in the next chapter, excellence can be seen as the missing ingredient to prove this extension property. The added generality of Shelah's notion of excellence is to expand the context beyond a combinatorial geometry to a more general dependence relation.

Exercise 2.1.6. *Let X, Y be subsets of a structure M. If f takes X to Y is an elementary isomorphism, f extends to an elementary isomorphism from* $\mathrm{acl}(X)$ *to* $\mathrm{acl}(Y)$. *(Hint: each element of* $\mathrm{acl}(X)$ *has a minimal description.)*

The content of Exercise 2.1.6 is given without proof on its first appearance [15]; a full proof is given in [108]. The next exercise recalls the use of combinatorial geometries to study basic examples of categoricity in the first order context.

Exercise 2.1.7. *Show a complete theory T is strongly minimal if and only if it has infinite models and*

1. *algebraic closure induces a pregeometry on models of T;*
2. *any bijection between acl-bases for models of T extends to an isomorphism of the models.*

Exercise 2.1.8. *A strongly minimal theory is categorical in any uncountable cardinality.*

2.2 Infinitary Logic

Infinitary logics $L_{\kappa,\lambda}$ arise by allowing infinitary Boolean operations (bounded by κ) and by allowing quantification over sequences of variables length $< \lambda$. Various results concerning completeness, compactness and other properties of these logics were established during the late 1960's and early 1970's. See for example [21, 80, 90, 24]. We draw on a few of these results as needed. Here we just fix the notation.

Notation 2.1. *For cardinals κ and λ and a vocabulary τ, $L_{\kappa,\lambda}(\tau)$ is the smallest collection Φ of formulas such that:*

1. *Φ contains all atomic τ-formulas in the variables v_i for $i < \lambda$.*
2. *Φ is closed under $\neg$.*
3. *Φ is closed under $\bigvee \Psi$ and $\bigwedge \Psi$ where Ψ is a set of fewer than κ formulas that contain strictly less than λ free variables.*
4. *Φ is closed under sequences of universal and existential quantifiers over less than λ variables v_i.*

Thus, the logic $L_{\omega_1,\omega}$ is obtained by extending the formation rules of first order logic to allow countable conjunctions and disjunctions; each formula of $L_{\omega_1,\omega}$ only finitely many free variable. $L_{\infty,\lambda}$ allows conjunctions of arbitrary length.

Definition 2.2. *A* fragment *Δ of $L_{\omega_1,\omega}$ is a countable subset of $L_{\omega_1,\omega}$ closed under subformula, substitutions of terms, finitary logical operations and such that: whenever $\Theta \subset \Delta$ is countable and $\phi, \bigvee \Theta \in \Delta$ then $\bigvee\{\exists x\theta : \theta \in \Theta\}$, $\bigvee\{\phi \wedge \theta : \theta \in \Theta\}$, and $\bigvee(\{\phi\} \cup \Theta)$ are all in Δ. Further, when dealing with theories with linearly ordered models, we require that if $\phi, \bigvee \Theta \in \Delta$ then* $\bigvee(\{\text{for arb large } x)\theta : \theta \in \Theta\}$.

The following semantic characterization of $L_{\omega_1,\omega}$ equivalence is an important tool.

Definition 2.3. *Two structures A and B are* back and forth equivalent *if there is a nonempty set I of isomorphisms of substructures A onto substructures of B such that:*

(forth) For every $f \in I$ and $a \in A$ there is a $g \in I$ such that $f \subseteq g$ and $a \in \operatorname{dom} g$.

(back) For every $f \in I$ and $b \in B$ there is a $g \in I$ such that $f \subseteq g$ and $b \in \operatorname{rg} g$.

We write $A \approx_p B$.

Proofs of the following theorem and related results can be found in e.g. [22] and [90].

Fact 2.4 (Karp)**.** *The following are equivalent.*

1. $A \approx_p B$
2. *A and B are $L_{\infty,\omega}$-elementarily equivalent.*

Either of these conditions implies that if A and B are both countable then $A \approx B$, i.e. they are isomorphic.

$L(Q)$ or $L_{\omega_1,\omega}(Q)$ is obtained by adding the further quantifier, 'there exists uncountably many' to the underlying logic. Truth for $L_{\omega_1,\omega}(Q)$ is defined inductively as usual; the key point is that $M \models (Qx)\phi(x)$ if and only if $|\{m : M \models \phi(m)\}| \geq \aleph_1$. There are other semantics for this quantifier (see [24]). There are other interpretations of the Q-quantifier, requiring for example that ϕ has κ solutions for some other infinite κ.

Fact 2.5. *[Löwenheim-Skolem theorems]* Unlike first order logic, the existence of models in various powers of an $L_{\omega_1,\omega}$-theory is somewhat complicated. The downward Löwenheim-Skolem to $\aleph_0$ holds for sentences (not theories) in $L_{\omega_1,\omega}$. For every $\alpha < \omega_1$, there is a sentence ϕ_α that has no model of cardinality greater than $\beth_\alpha$. (See Chapter 13 of Keisler book or [111].) Appendix A implies that if a sentence of $L_{\omega_1,\omega}$ has a model with cardinality at least $\beth_{\omega_1}$ then it has arbitrarily large models. Much of the difficulty of Part IV, of this text stems from dealing with this issue.

Remark 2.6 (Set theoretical notation)**.** We use with little reference such basic notations as $\operatorname{cf}(\lambda)$ for cofinality and (cub) for closed unbounded set. For background on such concepts see a set theory text such as [92].

3
Abstract Quasiminimality

In this chapter we introduce Zilber's notion [150] of an abstract quasiminimal-excellent class and prove Theorem 3.23: $L_{\omega_1,\omega}$-definable quasiminimal-excellent classes satisfying the countable closure condition are categorical in all powers. In the next chapter we expound Zilber's simplest concrete algebraic example. In Chapter 26, we will place this example in the context of Shelah's more general notion.

An abstract quasiminimal class is a class of structures in a countable vocabulary that satisfy the following two conditions, which we expound leisurely. The class is *quasiminimal excellent* if it also satisfies the key notion of *excellence* which is described in Assumption 3.15. A *partial monomorphism* is a 1-1 map which preserves quantifier-free formulas.

Assumption 3.1 (Condition I)**.** *Let $\boldsymbol{K}$ be a class of L-structures which admit a closure relation cl_M mapping $X \subseteq M$ to $\mathrm{cl}_M(X) \subseteq M$ that satisfies the following properties.*

1. *Each cl_M defines a pregeometry (Definition 2.1.1) on M.*
2. *For each $X \subseteq M$, $\mathrm{cl}_M(X) \in \boldsymbol{K}$.*
3. *If f is a partial monomorphism from $H \in \boldsymbol{K}$ to $H' \in \boldsymbol{K}$ taking $X \cup \{y\}$ to $X' \cup \{y'\}$ then $y \in \mathrm{cl}_H(X)$ iff $y' \in \mathrm{cl}_{H'}(X')$.*

Our axioms say nothing explicit about the relation between $\mathrm{cl}_N(X)$ and $\mathrm{cl}_M(X)$ where $X \subset M \subset N$. Note however that if $M = \mathrm{cl}_N(X)$, then monotonicity of closure implies $\mathrm{cl}_N(Y) \subseteq M$ for any $Y \subseteq M$. And, if closure is definable in some logic L^* (e.g. a fragment of $L_{\omega_1,\omega}$) and $X \subseteq M$ with $M \prec_{L^*} N$ then $\mathrm{cl}_N(X) = \mathrm{cl}_M(X)$. We will use this observation at the end of the chapter.

Condition 3.1.3) has an *a priori* unlikely strength: quantifier-free formulas determine the closure; in practice, the language is specifically expanded to guarantee this condition.

Remark 3.2. The following requirement is too strong: for any $M, N \in \boldsymbol{K}$ with $M \subseteq N$ and $X \subseteq M$,

$$\mathrm{cl}_N(X) = \mathrm{cl}_M(X).$$

Consider the first order theory of an equivalence relation with infinitely many infinite classes. Restrict to those models where every equivalence class is countable. With the closure of a set X defined to be the set of elements that are equivalent to some member of X, this class gives a quasiminimal excellent class with the countable closure property (Definition 3.8). But it does not satisfy the stronger condition proposed in this example. Note that in this example $\mathrm{cl}_M(\emptyset) = \emptyset$. This class is axiomatized in $L(Q)$.

Our development here is abstract; in general there is no syntactic conditions on the definability of $\boldsymbol{K}$. However, as presented here, the key arguments for the existence of models in arbitrarily large cardinalities (Lemma 3.23) use the assumption that $\boldsymbol{K}$ is definable in $L_{\omega_1,\omega}$.

Definition 3.3. *Let A be a subset of $H, H' \in \boldsymbol{K}$. A map from $X \subset H$ to $X' \subset H'$, which fixes $X \cap A$ is called a* partial A-monomorphism *if its union with the identity map on A preserves quantifier-free formulas (i.e. is an isomorphism).*

Frequently, but not necessarily, A will be the universe of member G of $\boldsymbol{K}$.

Definition 3.4. *Let $Ab \subset M$ and $M \in \boldsymbol{K}$. The (quantifier-free) type of b over A in M, written $\mathrm{tp}_{qf}(b/A; M)$, is the set of (quantifier-free)* first-order *formulas with parameters from A true of b in M.*

Exercise 3.5. *Why is M a parameter in Definition 3.4?*

Exercise 3.6. *Let $Ab \subset M$, $Ab' \subset M'$ with $M, M' \in \boldsymbol{K}$. Show there is a partial A-monomorphism taking a' to b' if and only if*

$$\mathrm{tp}_{qf}(b/A; M) = \mathrm{tp}_{qf}(b'/A; M').$$

The next condition connects the geometry with the structure on members of $\boldsymbol{K}$.

Assumption 3.7 (Condition II : $\aleph_0$-homogeneity over models)**.** *Let $G \subseteq H, H' \in \boldsymbol{K}$ with G empty or a countable member of $\boldsymbol{K}$ that is closed in H, H'.*

1. *If f is a partial G-monomorphism from H to H' with finite domain X then for any $y \in \mathrm{cl}_H(X)$ there is $y' \in H'$ such that $f \cup \{\langle y, y'\rangle\}$ extends f to a partial G-monomorphism.*
2. *If f is a bijection between $X \subset H \in \boldsymbol{K}$ and $X' \subset H' \in \boldsymbol{K}$ which are* separately cl-*independent (over G) subsets of H and H' then f is a G-partial monomorphism.*

Definition 3.8. *We say a closure operation satisfies the* countable closure property *if the closure of a countable set is countable.*

Of course, this implies that for any infinite $X \subseteq M \in \boldsymbol{K}$, $|\mathrm{cl}_M(X)| = |X|$. These semantic conditions yield syntactic consequences. A structure M is often called 'quasiminimal' if every $L_{\omega_1,\omega}$-definable subset of M is countable or cocountable; our formal definition includes homogeneity conditions.

Lemma 3.9. *Suppose* cl_M *on an uncountable structure* $M \in \boldsymbol{K}$ *satisfies Conditions I, II and the countable closure property.*

1. *For any finite set* $X \subset M$*, if* $a, b \in M - \mathrm{cl}_M(X)$*,* a, b *realize the same* $L_{\omega_1,\omega}$*-type over* X*.*
2. *Every* $L_{\omega_1,\omega}$*-definable subset of* M *is countable or cocountable. This implies that* $a \in \mathrm{cl}_M(X)$ *iff it satisfies some* ϕ *over* X*, which has only countably many solutions.*

Proof. Condition 1) follows from Condition II (Assumption 3.7) by constructing a back and forth: let $G = \mathrm{cl}_M(X)$. The collection $\mathcal{S}$ of finite G-monomorphisms that contain $\langle a, b\rangle$ is a back and forth (see Definition 2.3); $\{\langle a, b\rangle\} \in \mathcal{S}$ by Assumption 3.7.1. Go back and forth using Assumption 3.7.2 if the new element is independent and Assumption 3.7.1 if the extension is to a dependent element. To see condition 2), suppose both ϕ and $\neg\phi$ had uncountably many solutions with ϕ defined over X. Then there are a and b satisfying ϕ and $\neg\phi$ respectively and neither is in $\mathrm{cl}(X)$; this contradicts 1). $\square_{3.9}$

We immediately conclude a form of ω-stability. By an $L_{\omega_1,\omega}$-type of a sequence of $\boldsymbol{a}$ of length n in $M \in \boldsymbol{K}$ over A, we mean the set of $L_{\omega_1,\omega}$ with n free variables and finitely many parameters from A satisfied by $\boldsymbol{a}$. By Lemma 3.12, each such type is determined by the countable set of formulas giving the quantifier-free type of $\boldsymbol{a}$.

Corollary 3.10. *Suppose* cl *on an uncountable structure* $M \in \boldsymbol{K}$ *satisfies Conditions I, II, and the countable closure property. Then* $\boldsymbol{K}$ *is* syntactically *stable in every cardinality. That is, there are only* $|X|$*,* $L_{\omega_1,\omega}$*-types over any infinite set* $X \subseteq M$*.*

Exercise 3.11. *Consider the class of models of an equivalence relation such that each class has one or two elements. Let* $\mathrm{cl}_N(\boldsymbol{a})$ *be the algebraic closure in the normal first order sense. Show that this example fails both* ω*-homogeneity (Assumption 3.7.1) and Lemma 3.9. (This example was constructed by John Goodrick and Alice Medvedev.)*

The ω-homogeneity condition yields by an easy induction:

Lemma 3.12. *Suppose Conditions I and II hold. Let* $G \in \boldsymbol{K}$ *be* countable *and suppose* $G \subset H, H' \in \boldsymbol{K}$ *with* G *closed in* H, H' *if* $G \neq \emptyset$*.*

1. *If $X \subset H$, $X' \subset H'$ are finite and f is a G-partial monomorphism from X onto X' then f extends to a G-partial monomorphism from $\mathrm{cl}_H(GX)$ onto $\mathrm{cl}_{H'}(GX')$. Thus, for any M, N, $\mathrm{cl}_M(\emptyset) \approx \mathrm{cl}_M(\emptyset)$.*

2. *If X is an independent set of cardinality at most $\aleph_1$, and f is a G-partial monomorphism from X to X' then f extends to a G-partial monomorphism from $\mathrm{cl}_H(GX)$ onto $\mathrm{cl}_{H'}(GX')$.*

Proof. The first statement is proved by constructing a back and forth extending f from $\mathrm{cl}_H(X)$ to $\mathrm{cl}_H(X)$. Arrange the elements of $\mathrm{cl}_H(X)$ (of $\mathrm{cl}_H(X)$) in order type ω with X (X') first. Simultaneously construct f and f^{-1}, extending f at even stages and f^{-1} at odd stages using Assumption 3.7.1. The second follows by induction from the first (by replacing G by $\mathrm{cl}_H(GX_0)$ for X_0 a countable initial segment of X). $\square_{3.12}$

For algebraic closure the cardinality restriction on X is unnecessary. Assumption 3.15 in the definition of a quasiminimal excellent class, allows us to extend Lemma 3.12.2 to sets of arbitrary cardinality for classes that satisfy the countable closure property. In particular, Lemma 3.12 implies that $\boldsymbol{K}$ is $\aleph_1$-categorical. But we (apparently) need the extra condition of excellence to get categoricity in higher cardinalities.

Now we define the concept of *quasiminimal excellence*. It makes certain amalgamation requirements (Lemma 3.18) in the context of a geometry; Shelah's notion of *excellence* works in the context of a more general independence relation. We need two technical definitions to simplify the statement of the condition.

Definition 3.13. *Let $C \subseteq H \in \boldsymbol{K}$ and let X be a finite subset of H. We say $\mathrm{tp}_{\mathrm{qf}}(X/C)$ is* determined over *the finite C_0 contained in C if: for every partial monomorphism f mapping X into H', for every partial monomorphism f_1 mapping C into H', if $f \cup (f_1 \restriction C_0)$ is a partial monomorphism, $f \cup f_1$ is also a partial monomorphism. We write*

$$\mathrm{tp}_{\mathrm{qf}}(X/C_0) \models \mathrm{tp}_{\mathrm{qf}}(X/C)$$

to express this concept.

Definition 3.14. *Let $H \in \boldsymbol{K}$ and $C \subset H$. C is called special if there is a countable independent subset A and $A_0, \ldots A_{n-1} \subseteq A$ such that:*

$$C = \bigcup_{i<n} \mathrm{cl}_H(A_i).$$

Assumption 3.15 (Condition III: Quasiminimal Excellence)**.** *Let $H \in \boldsymbol{K}$ and suppose $C \subset H$ is special. Then for any finite $X \subset \mathrm{cl}_H(C)$, $\mathrm{tp}(X/C; H)$ is determined by a finite subset C_0 of C.*

One sort of special set is particularly important. In the following definition it is essential that $\subset$ be understood as *proper* subset.

Definition 3.16. *Let $Y \subset M \in \boldsymbol{K}$.*

1. *For any Y, $\mathrm{cl}^-_M(Y) = \bigcup_{X \subset Y} \mathrm{cl}_M(X)$.*
2. *We call C (the union of)* an n-dimensional cl_M-independent system *if $C = \mathrm{cl}^-_M(Z)$ and Z is an independent set of cardinality n.*

To visualize a 3-dimensional independent system think of a cube with the empty set at one corner A and each of the independent elements z_0, z_1, z_2 at the corners connected to A. Then each of $\mathrm{cl}(z_i, z_j)$ for $i < j < 3$ determines a side of the cube: $\mathrm{cl}^-(Z)$ is the union of these three sides; $\mathrm{cl}(Z)$ is the entire cube.

In particular Condition III, which is the central point of excellence, asserts (for dimension 3) that the type of any element in the cube over the union of the three given sides is determined by the type over a finite subset of the sides. The 'thumbtack lemma' of Chapter 4 verifies this condition in a specific algebraic context. Lemma 3.18 provides a less syntactic version of the excellence condition: there is a particularly strong amalgam over any special set. To state it we need a definition.

Definition 3.17. *We say $M \in \boldsymbol{K}$ is* prime *over the set $X \subset M$ if every partial monomorphism of X into $N \in \boldsymbol{K}$ extends to a monomorphism of M into N.*

We say a map is *closed* if it takes closed subsets of H to closed subsets of H'.

Lemma 3.18. *Suppose $\boldsymbol{K}$ has the countable closure property. Let $H \in \boldsymbol{K}$ and suppose $C \subset H$ is special. Then $\mathrm{cl}_H(C)$ is prime over C. Moreover, if f maps C onto $C' \subseteq H'$ is a closed map then f extends to $\hat{f}$ mapping $\mathrm{cl}_H(C)$ onto $\mathrm{cl}_{H'}(C')$.*

Proof. Let $\overline{C} = \mathrm{cl}_H(C)$ and let $\overline{C}' = \mathrm{cl}_{H'}(f(C))$. They are both countable, so choose an ordering of each of length ω. Inductively we construct partial embeddings f_n from H to H' for $n \in \omega$ such that for each n:

- $\mathrm{preim}(f_n)$ is finite,
- $f_n \subseteq f_{n+1}$, and
- $f_n \cup f$ is a partial embedding.

Take $f_0 = \emptyset$. We construct the f_n for $n > 0$ via the back and forth method, going forth for odd n and back for even n. For odd n, let a be the least element of $\overline{C} - \mathrm{preim} f_{n-1}$. The set $\mathrm{preim} f_{n-1} \cup \{a\}$ is a finite subset of $\overline{C}$, so by quasiminimal excellence and finite character there is a finite subset C_0 of C such that the quantifier-free type of $\mathrm{preim} f_{n-1} \cup \{a\}$ over C is determined over C_0 and $a \in \mathrm{cl}_H(C_0)$. Let $g = f_{n-1} \cup f \restriction C_0$. By induction, $f_{n-1} \cup f$ is a partial monomorphism, so g is a partial monomorphism. By Condition 3.7.1 ($\aleph_0$-homogeneity), there is $b \in H'$ such that $f_n \coloneqq g \cup \{(a, b)\}$ is a partial monomorphism. Since the type of $\mathrm{preim} f_n$ over C is determined over C_0, $f_n \cup f$ is a partial monomorphism, as required.

For even n, note that $f(C)$ is special in H' because f is a closed partial monomorphism. Also note that the inverse of a partial monomorphism is a partial monomorphism. Hence we can perform the same process as for odd steps, reversing the roles of H and H', to find f_n whose image contains the least element of $\overline{C}'$ not in the image of f_{n-1}. Let $\hat{f} = \bigcup_{n\in\mathbb{N}} f_n$. Then $\hat{f}$ is an monomorphism extending f, defined on all of $\overline{C}$, whose image is all of $\overline{C}'$. Hence $\hat{f}$ is a closed monomorphism. $\square_{3.18}$

Theorem 3.19 generalizes the 'extension of maps to algebraic closures' (Exercise 2.1.6) to 'extension of maps to quasiminimal closures'. It is instructive to contemplate why the argument for Exercise 2.1.6 does not generalize to the current context.

Theorem 3.19. *Let $\boldsymbol{K}$ be a quasiminimal excellent class. Suppose $H, H' \in \boldsymbol{K}$ satisfy the countable closure property. Let $\mathcal{A}, \mathcal{A}'$ be cl-independent subsets of H, H' with $\mathrm{cl}_H(\mathcal{A}) = H$, $\mathrm{cl}_{H'}(\mathcal{A}') = H'$, respectively, and ψ a bijection between $\mathcal{A}$ and $\mathcal{A}'$. Then ψ extends to an isomorphism of H and H'.*

Proof. First, we outline the argument. We have the obvious directed union $\{\mathrm{cl}(X) : X \subseteq \mathcal{A}; |X| < \aleph_0\}$ with respect to the partial order of finite subsets of X by inclusion. And $H = \bigcup_{X\subset\mathcal{A};|X|<\aleph_0} \mathrm{cl}(X)$. So the theorem follows immediately if for each finite $X \subset \mathcal{A}$ we can choose $\psi_X : \mathrm{cl}_H(X) \to H'$ so that $X \subset Y$ implies $\psi_X \subset \psi_Y$. We prove this by induction on $|X|$. If $|X| = 1$, the condition is immediate from $\aleph_0$-homogeneity and the countable closure property. Suppose $|Y| = n+1$ and we have appropriate ψ_X for $|X| < n+1$. We will prove two statements.

1. $\psi_Y^- : \mathrm{cl}^-(Y) \to H'$ defined by $\psi_Y^- = \bigcup_{X\subset Y} \psi_X$ is a monomorphism.

2. ψ_Y^- extends to ψ_Y defined on $\mathrm{cl}(Y)$.

Before completing the argument for Theorem 3.19, we describe the structure H (and H') in a little more detail.

Notation 3.20. *Fix a countably infinite subset $\mathcal{A}_0$ of $\mathcal{A}$ and write $\mathcal{A}$ as the disjoint union of $\mathcal{A}_0$ and a set $\mathcal{A}_1$; without loss of generality, we can assume ψ is the identity on $\mathrm{cl}_H(\mathcal{A}_0)$ and work over $G = \mathrm{cl}_H(\mathcal{A}_0)$. We may write $\mathrm{cl}^*_H(X)$ to abbreviate $\mathrm{cl}_H(\mathcal{A}_0 X)$.*

The first statement is proved in the next few paragraphs. The second step follows immediately by Lemma 3.18. We work with a fixed $X \subset \mathcal{A}_1$ with $|X| = n+1$ and write Z for $\psi(X)$ (and similarly for subsets of X).

Our first task is to prove Statement (1) of Theorem 3.19. By the induction hypothesis it suffices to consider n-element subsets of X. Let X_i for $i < m = n+1$ be the n-elements subsets of X and let x_k be the element of X not in X_k. Note that for each $J \subset n+1$, each union $\bigcup_{i\in J} \mathrm{cl}^*_H(X_i)$ is special. By induction we have for each $i < m$ a map ψ_{X_i} from $\mathrm{cl}^*_H(X_i)$ onto $\mathrm{cl}^*_{H'}(Z_i)$. To prove (1) it suffices to show that $\bigcup_{i<m} \psi_{X_i}$ is a G-monomorphism. The exchange axiom guarantees that $\mathrm{cl}^*_H(X_i) \cap \mathrm{cl}^*_H(X_j) = \mathrm{cl}^*_H(X_i \cap X_j)$. By the global induction, the maps ψ_{X_i} agree where more than one is defined so $\bigcup_{i<m} \psi_{X_i}$ is a 1-1 map. We prove by induction on $k \le m$ that $\psi_k = \bigcup_{i<k} \psi_{X_i}$ is a G-monomorphism. This is evident for $k = 1$. When $k = m$, we finish. At the induction step, we first prove there is a map $\hat{\psi}$ mapping $\mathrm{cl}^*_H(X_k)$ monomorphically onto $\mathrm{cl}^*_{H'}(Z_k)$ and extending $\psi_k \cup f$. For $k = 2$, this follows from Lemma 3.12. For $k \ge 2$, since $C_k = \bigcup_{i<k} \mathrm{cl}^*_H(X_i)$ is special, $\mathrm{cl}^*_H(C_k)$ is prime over $\bigcup_{i<k} \mathrm{cl}^*_H(X_i)$. So $\hat{\psi}$ exists, as $\mathrm{cl}^*_H(C_k) = \mathrm{cl}^*_H(X_k)$.

Note any $\mathbf{c} \in C_k$ can be written as $\boldsymbol{a}\mathbf{b}$ with $\boldsymbol{a} \in C_{k-1}$ and $\mathbf{b} \in \mathrm{cl}_H(GX_k)$. Take an arbitrary such tuple $\boldsymbol{a}\mathbf{b}$. We will construct a submodel $H_0 \subset H$ and a monomorphism τ from H_0 into H' such that ψ_k and τ agree on $\boldsymbol{a}$ and $\mathbf{b}$. This suffices to show ψ_k is a monomorphism.

Let B be a finite subset of $\mathcal{A}_0$ so that $\boldsymbol{a}, \mathbf{b} \in \mathrm{cl}_H(BX)$. Let $z \in \mathcal{A}_0 - B$, let $H_0 = \mathrm{cl}_H(BXz)$, and let $H'_0 = \mathrm{cl}_{H'}(\psi(BXz))$. By Lemma 3.12.2, there is an automorphism σ of H_0, fixing $\mathrm{cl}_H(BX_k)$ and swapping x_k and z.

The idea is to compare ψ_k on H_0 with the composite embedding $\tau = \sigma'^{-1}\psi_{X_k}\sigma$, where σ' is an automorphism of H'_0 which 'corresponds' to σ. We find σ' by conjugating σ by $\hat{\psi}: \sigma' = \hat{\psi}\sigma\hat{\psi}^{-1}$. Let

$$\tau = \sigma'^{-1}\psi_{X_k}\sigma = \hat{\psi}\sigma^{-1}\hat{\psi}^{-1}\psi_{X_k}\sigma.$$

Write X_{ik} for $X_i \cap X_k$. The tuple $\boldsymbol{a} \in C_{k-1} = \bigcup_{i<k-1} \mathrm{cl}_H(BX_i)$, so

$$\sigma(\boldsymbol{a}) \in \bigcup_{i<k-1} \mathrm{cl}_H(BX_{ik}z) \subset \bigcup_{i<k-1} \mathrm{cl}^*_H(X_{ik}) \subset \mathrm{dom}\,\psi_{k-1}.$$

By hypothesis, for $i < k$, ψ_{X_k} agrees with ψ_{X_i} on $\mathrm{cl}^*_H(X_{ik})$, hence ψ_{X_k} agrees with ψ_{k-1} on $\bigcup_{i<k} \mathrm{cl}^*_H(X_{ik})$. Also $\hat{\psi}$ and ψ_k both extend ψ_{k-1}, so

$$\tau(\boldsymbol{a}) = \hat{\psi}\sigma^{-1}\hat{\psi}^{-1}\psi_{X_k}\sigma(\boldsymbol{a}) = \psi_{k-1}\sigma^{-1}\psi_{k-1}^{-1}\psi_{k-1}\sigma(\boldsymbol{a}) = \psi_{k-1}(\boldsymbol{a}) = \psi_k(\boldsymbol{a}).$$

The tuple b is in $\mathrm{cl}_H(BX_k)$, so it is fixed by σ. The monomorphisms ψ_{X_k} and $\hat{\psi}$ preserve the closure, so $\hat{\psi}^{-1}\psi_{X_k}(\mathbf{b}) \in \mathrm{cl}_H(BX_k)$ is fixed by σ^{-1}. So

$$\tau(\mathbf{b}) = \hat{\psi}\sigma^{-1}\hat{\psi}^{-1}\psi_{X_k}\sigma(\mathbf{b}) = \hat{\psi}\hat{\psi}^{-1}\psi_{X_k}(\mathbf{b}) = \psi_{X_k}(\mathbf{b}) = \psi_k(\mathbf{b}).$$

Thus for any quantifier-free formula R,

$$H \models R(\boldsymbol{a}, \mathbf{b}) \leftrightarrow H' \models R(\tau(\boldsymbol{a}), \tau(\mathbf{b})) \leftrightarrow H' \models R(\psi_k(\boldsymbol{a}), \psi_k(\mathbf{b})).$$

This holds for any tuples $\boldsymbol{a}, \mathbf{b}$ (for a suitable choice of τ) and so ψ_k is a partial embedding. In particular, $\psi_X = \psi_{n+1}$ is a partial embedding. It is a union of finitely many closed partial embeddings, hence is a closed partial embedding. $\square_{3.19}$

This completes the proof of Theorem 3.19; We have shown that the isomorphism type of a structure in $\boldsymbol{K}$ with countable closures is determined by the cardinality of a basis for the geometry. If M is an uncountable model in $\boldsymbol{K}$ that satisfies the countable closure property, the size of M is the same as its dimension so there is at most one model in each uncountable cardinality which has countable closures. It remains to show that there is at least one.

Definition 3.21. *Define $M \prec_{\boldsymbol{K}} N$ on the quasiminimal excellent class $\boldsymbol{K}$ if $M, N \in \boldsymbol{K}$, M is a substructure of N, and $\mathrm{cl}_M(X) = \mathrm{cl}_N(X)$ for $X \subseteq M$.*

Lemma 3.22. *Let $\boldsymbol{K}$ be a quasiminimal excellent class such that the class $\boldsymbol{K}$ and the closure relation is definable in $L_{\omega_1,\omega}$. The class $\boldsymbol{K}'$ of models of $\boldsymbol{K}$ that have infinite dimension and satisfy the countable closure property is closed under unions of increasing $\prec_{\boldsymbol{K}}$-chains.*

Proof. Let H be a model with countable dimension. Let L^* be a countable fragment of $L_{\omega_1,\omega}$ containing the Scott sentence σ of H (see 7.1.6) and the formulas defining independence and the definition of the class $\boldsymbol{K}$. Note that if $H_0 \subseteq H_1$, with both in $\boldsymbol{K}'$, a back and forth shows H_0 is an L^*-elementary submodel of H_1. Now let $\langle H_i : i < \mu \rangle$ be a $\prec_{\boldsymbol{K}}$-increasing chain of members of $\boldsymbol{K}'$. Then the union is also a model $H_\mu \in \boldsymbol{K}$. For any countable $X' \subseteq H_\mu$ and $a \in \mathrm{cl}_{H_\mu}(X')$, there is a finite $X \subseteq X'$ with $a \in \mathrm{cl}_{H_\mu}(X)$. There is an i with $X \subseteq H_i$ and $\mathrm{cl}_{H_i}(X) = \mathrm{cl}_{H_\mu}(X)$; so since H_i satisfies the countable closure property, so does H_μ. $\square_{3.22}$

Now we get the main result.

Theorem 3.23. *Let $\boldsymbol{K}$ be a quasiminimal excellent class such that the class $\boldsymbol{K}$ and the closure relation is definable in $L_{\omega_1,\omega}$. If there is an $H \in \boldsymbol{K}$ which contains an infinite* cl*-independent set, then there are members of $\boldsymbol{K}$ of arbitrary cardinality which satisfy the countable closure property.*

Thus, the $M \in \boldsymbol{K}$ that have the countable closure property form a class that is categorical in all uncountable powers.

Proof. Fix H, L^*, σ and $\boldsymbol{K}'$ as in the proof of Lemma 3.22. Let X be a countable independent subset of H, H_1 the closure of Xa where a is independent from X and let $H_0 = \mathrm{cl}_{H_1}(X)$. Note that $H_0 \prec_{L^*} H_1$ by a back and forth. Since cl is L^*-definable, the argument after Assumption 3.1 yields $\mathrm{cl}_{H_1}(Y) = \mathrm{cl}_{H_0}(Y)$ for any $Y \subset H_0$. So we have a model which is isomorphic to a proper L^*-elementary and $\prec_{\boldsymbol{K}}$-extension. By Vaught's old argument (Theorem 5.3 of [145]) since all members of $\boldsymbol{K}$ with countably infinite dimension are isomorphic, one can construct a continuous L^*-elementary increasing chain of members of $\boldsymbol{K}$ for any $\alpha < \aleph_1$. Thus we get a model of power $\aleph_1$ which has countable closures by Lemma 3.22. Now by cardinal induction, we construct for every κ a model of cardinal κ that has the countable closure property. We use categoricity below κ (obtained by induction) to continue the chain at limit ordinals. The categoricity now follows by argument immediately before the statement of Definition 3.21. $\square_{3.23}$

Remark 3.24. Zilber omitted exchange in his original definition but it holds in the natural contexts he considers so we made it part of our definition of quasiminimal excellence. This decision was validated when Zilber (unpublished) later found an example showing the necessity of assuming exchange in his general context. The exact formulation of Conditions I-III refine Zilber's statement; they resulted from examples suggested by Goodrick-Medvedev and Kirby and further vital discussions with Kirby who contributed greatly to the final versions of the proofs of Lemma 3.18 and Theorem 3.19. He provides further variations and extensions in [83]. Our argument by direct limits for Theorem 3.19 is more conceptual than Zilber's and avoids his notion of a 'perfect' subset.

In the remainder of this chapter we foreshadow the connections of quasiminimal excellent classes with the more general concepts which are investigated in the remainder of the book.

We can glean from the proofs of Theorem 3.23 and Lemma 3.22 that the models of a quasiminimal class (definable in $L_{\omega_1,\omega}$) that have the countable closure property form an abstract elementary class (Definition 5.1) with $\prec_{\boldsymbol{K}}$ from Definition 3.21 as the notion of submodel and with arbitrarily large models. Note that the requirement that $\mathrm{cl}(X) \in \boldsymbol{K}$ combined with the countable closure property implies that $\boldsymbol{K}$ has Löwenheim-Skolem number $\aleph_0$.

Lemma 3.25. *Suppose class $\boldsymbol{K}$ is a quasiminimal excellent class and the closure relation is definable in $L_{\omega_1,\omega}$. Let $\boldsymbol{K}'$ be the class of those models of a quasiminimal excellent class $\boldsymbol{K}$ such that satisfy the countable closure property. Define $M \prec_{\boldsymbol{K}} N$ as in Definition 3.21. Then $(\boldsymbol{K}', \prec_{\boldsymbol{K}})$ is an abstract elementary class with the amalgamation property (Definitions 5.1, 5.10).*

Proof. Lemma 3.22 shows that Condition **A.3** is satisfied; $\mathrm{LS}(\boldsymbol{K}') = \aleph_0$ since the class is axiomatized in $L_{\omega_1,\omega}$; the other conditions of Definition 5.1 are routine for $\prec_{\boldsymbol{K}}$. To show amalgamation, choose $Z_0 \subset H_0$, $Z_1 \subset H_1 - H_0$, and $Z_2 \subset H_2 - H_0$ so that $H_1 = \mathrm{cl}_{H_1}(Z_0Z_1)$ and $H_2 = \mathrm{cl}_{H_2}(Z_0Z_2)$. Note that Z_0Z_1 and Z_0Z_2 are each independent. Now choose a model G that is the closure of the independent set X_0, X_1, X_2 where for each i, $|X_i| = |Z_i|$. Note that $H_1 \approx \mathrm{cl}_G(X_0X_1)$ and $H_2 \approx \mathrm{cl}_G(X_0X_2)$; this completes the proof. $\square_{3.25}$

Remark 3.26. In Lemma 3.25, we assumed the closure relation was definable in $L_{\omega_1,\omega}$. It follows using the Q-quantifier to express countable closure that $\boldsymbol{K}'$ is $L_{\omega_1,\omega}(Q)$-definable. Note however, that the notion of strong submodel is *not* $L_{\omega_1,\omega}(Q)$-elementary submodel; there are countable models. Kirby [83] proves a converse: if $\boldsymbol{K}$ is a quasi-excellent class as defined here and the class of models in $\boldsymbol{K}$ with the countable closure condition is closed under unions of $\prec_{\boldsymbol{K}}$-chains then $\boldsymbol{K}'$ is definable in $L_{\omega_1,\omega}(Q)$.

In the language introduced in Definition 9.6, the following lemma says that the Galois types in a quasiminimal excellent classes are the same as the syntactic types in $L_{\omega_1,\omega}$. And in the terminology of Chapter 12, this implies the class is 'tame'. One might hope for a straightforward observation that in the quasiminimal context, syntactic and Galois types are the same. But, the argument for the following corollary relies on the explicit consequence of excellence which is half of the categoricity argument. However, we don't need to know that $\boldsymbol{K}$ has arbitrarily large models nor that $\boldsymbol{K}$ is $L_{\omega_1,\omega}$-definable.

Corollary 3.27. *Let $\boldsymbol{K}$ be a quasiminimal excellent class, with $G \prec_{\boldsymbol{K}} H, H'$ all in $\boldsymbol{K}$. If $a \in H, a' \in H'$ realize the same quantifier free type over G (i.e. there is a G-monomorphism taking a to a') then there is a $\boldsymbol{K}$-isomorphism from $\mathrm{cl}(Ga)$ onto $\mathrm{cl}(Ga')$. Thus,*

1. *(G, a, H) and (G, a', H') realize the same Galois type.*
2. *Galois types are the same as $L_{\omega_1,\omega}$-types.*

Proof. For countable G the result follows from Lemma 3.12. But for uncountable G, we use Theorem 3.19 with Z as a basis for G and $\mathcal{A} = Za$, $\mathcal{A}' = Za'$. The first conclusion is now immediate; we have an embedding of $\mathrm{cl}_H(Ga)$ into H'. For the second, slightly stronger variant, just note that Galois types always refine syntactic types and we have just established the converse. $\square_{3.27}$

The following notion is useful for studying covers of semi-abelian varieties (Chapter 4, [151, 26, 41]). As in higher rank first order structures, we cannot define a combinatorial geometry on the whole structure. Here are the main features of the definition.

Definition 3.28. *Let $\boldsymbol{K}$ be a class of L-structures which admit a function cl_M mapping $X \subseteq M$ to $\mathrm{cl}_M(X) \subseteq M$ that satisfies the following properties.*

1. *cl_M satisfies conditions A1-A3 of the definition of a combinatorial geometry (Definition 2.1.1) (not exchange).*
2. *cl_M induces a quasiminimal excellent geometry on a distinguished sort U. Conditions 3.1 and 3.7.*
3. *$M = \mathrm{cl}_M(U)$.*

Exercise 3.29. *Formalize Definition 3.28 and prove almost quasiminimal classes with countable chain condition are categorical in all powers ([150]).*

We will extend the following basic example at several points including Chapter 27. Here the notion of 'almost' is particularly strong; the model is the algebraic closure of the quasiminimal set. The more general definition is motivated in Chapter 4.

Example 3.30. Consider a three-sorted structure with an infinite set I and G the collection of all functions with finite support from I into the third sort that contains only $\{0, 1\}$ and with the evaluation predicate: $E(i, g, a)$ holds if and only if $g(i) = a$. Note that this structure is almost quasiminimal; the I sort is quasiminimal; the G-sort is not. The natural addition on G and 2 is definable so in further considerations we may often add those functions.

4

Covers of the Multiplicative Group of $\mathcal{C}$

In this chapter, we expound a relatively simple algebraic example [151, 26] of a categorical quasiminimal excellent class and apply the results of Chapter 3 to conclude it is categorical in all uncountable powers. Even this example requires some significant algebraic information which is beyond the scope of this monograph. But as we discuss at the end of the chapter, it provides a concrete algebraic example of an $L_{\omega_1,\omega}$ sentence which is categorical in all powers without every model being $\aleph_1$-homogeneous. This example will further serve to illustrate a number of the complexities which we will investigate in the final part of the book, beginning with Chapter 19. At the end of the chapter we briefly survey two further directions that Zilber has laid out: replacing the multiplicative group of $\mathcal{C}$ treated here by a semi-abelian variety, and studying the full structure of the complex numbers under exponentiation.

The first approximation to finding an infinitary axiomatization of complex exponentiation considers short exact sequences of the following form.

$$0 \to \mathcal{Z} \to H \to F^* \to 0. \tag{4.1}$$

H is a torsion-free divisible abelian group (written additively), F is an algebraically closed field, and $\exp$ is the homomorphism from $(H, +)$ to $(F^*, \cdot)$, the multiplicative group of F. We can code this sequence as a structure for a language L that includes $(H, +, \pi, E, S)$ where π denotes the generator of $\ker \exp$, $E(h_1, h_2)$ iff $\exp(h_1) = \exp(h_2)$ and we pull back sum by defining $H \models S(h_1, h_2, h_3)$ iff $F \models \exp(h_1) + \exp(h_2) = \exp(h_3)$. Thus H now represents both the multiplicative and additive structure of F. We want to show that the class of all such H (with standard Z) is categorical in each uncountable cardinal.

Let $\exp : H \mapsto F^*$. To guarantee Assumption 3.1.3 we include the following symbols in L. For each affine variety over $\mathbb{Q}$, $\hat{V}(x_1, \ldots x_n)$, we add a relation symbol V interpreted by

$$H \models V(h_1, \ldots, h_n) \text{ iff } F \models \hat{V}(\exp(h_1), \ldots, \exp(h_n)).$$

This includes the definition of S mentioned above; we skip over some fuss to handle the pullback of relations which have 0 in their range.

Lemma 4.1. *There is an $L_{\omega_1,\omega}$-sentence Σ such that there is a 1-1 correspondence between models of Σ and sequences (4.1).*

Proof. The sentence asserts first that the quotient of H by E with the image of $+$ corresponding to $\times$ and the image of S to $+$ is an algebraically closed field. That is a first order condition; using $L_{\omega_1,\omega}$ we guarantee every element of the kernel is and integer multiple of the fixed element π. This same proviso insures that the relevant closure operation (Definition 4.2) has countable closures. $\Box_{4.1}$

Definition 4.2. *For $X \subset H \models \Sigma$,*

$$\mathrm{cl}(X) = \exp^{-1}(\mathrm{acl}(\exp(X)))$$

where acl *is the field algebraic closure in F.*

Using this definition of closure the key result of [151] asserts:

Theorem 4.3. *Σ is quasiminimal excellent with the countable closure condition and categorical in all uncountable powers.*

Our goal is this section is to prove this result modulo one major algebraic lemma. We will frequently work directly with the sequence (1) rather than the coded model of Σ. Note that (1) includes the field structure on F. That is, two sequences are isomorphic if there are maps H to H' etc. where the first two are group isomorphisms but the third is a field isomorphism, that commute with the homomorphism in the short exact sequence.

It is easy to check Condition I, cl gives a combinatorial geometry, and to see that the closure of finite sets is countable. We need more notation about the divisible closure (in the multiplicative group of the field) to understand the remaining conditions.

Definition 4.4. *By a* divisibly closed multiplicative subgroup *associated with $a \in \mathcal{C}^*$, $a^{\mathbb{Q}}$, we mean a* **choice** *of a multiplicative subgroup containing a and isomorphic to the additive group $\mathbb{Q}$.*

Definition 4.5. *We say $b_1^{\frac{1}{m}} \in b_1^{\mathbb{Q}}, \ldots b_\ell^{\frac{1}{m}} \in b_\ell^{\mathbb{Q}} \subset \mathcal{C}^*$,* determine the isomorphism type *of $b_1^{\mathbb{Q}}, \ldots b_\ell^{\mathbb{Q}} \subset \mathcal{C}^*$ over the subfield k of $\mathcal{C}$ if given subgroups of the form $c_1^{\mathbb{Q}}, \ldots c_\ell^{\mathbb{Q}} \subset \mathcal{C}^*$ and ϕ_m such that*

$$\phi_m : k(b_1^{\frac{1}{m}} \ldots b_\ell^{\frac{1}{m}}) \rightarrow k(c_1^{\frac{1}{m}} \ldots c_\ell^{\frac{1}{m}})$$

is a field isomorphism it extends to a field isomorphism

$$\phi_\infty : k(b_1^{\mathbb{Q}}, \dots b_\ell^{\mathbb{Q}}) \to k(c_1^{\mathbb{Q}}, \dots c_\ell^{\mathbb{Q}}).$$

To see the difficulty consider the following example.

Example 4.6. Let b_1, b_2 (respectively c_1, c_2) be pairs of linearly independent complex numbers such that each satisfies $(x_1 - 1)^2 = x_2$. Suppose ϕ, which maps $\mathbb{Q}(b_1, b_2)$ to $\mathbb{Q}(c_1, c_2)$, is a field isomorphism; ϕ need not extend to a field isomorphism of their associated multiplicative subgroups $\mathbf{b}^{\mathbb{Q}}$, $\mathbf{c}^{\mathbb{Q}}$. To see this fix a square root function on $\mathcal{C}$. Let b_2 be transcendental and $b_1 = 1 + \sqrt{b_2}$. Let c_2 be another transcendental and set $c_1 = 1 - \sqrt{c_2}$. Now $\mathbb{Q}(b_1, b_2)$ is isomorphic to $\mathbb{Q}(c_1, c_2)$, taking b_i to c_i. But suppose we also specify that $b_1^{1/2} = \sqrt{(1 + \sqrt{b_2})}$, $b_2^{1/2} = \sqrt{b_2}$, $c_1^{1/2} = \sqrt{(1 - \sqrt{c_2})}$, $c_2^{1/2} = \sqrt{c_2}$. There is no field isomorphism taking $\mathbb{Q}(b_1, b_2, b_1^{1/2}, b_2^{1/2})$ to $\mathbb{Q}(c_1, c_2, c_1^{1/2}, c_2^{1/2})$ with $b_i^{1/m}$ going to $c_i^{1/m}$ (for i and m being 1 or 2).

From another perspective we can see that both the b_1, b_2 and c_1, c_2 satisfy the irreducible variety W given by:

$$(X_1 - 1)^2 = X_2.$$

But there are two different choices for an irreducible variety $W^{1/2}$, which is mapped onto W by squaring in each coordinate. The first is given by

$$(Y_1^2 - 1)^2 = Y_1^2$$

$$Y_1^2 - Y_2 = 1$$

and the second by

$$(Y_1^2 - 1)^2 = Y_1^2$$

$$Y_1^2 + Y_2 = 1.$$

As in Chapter 3, for G a subgroup of H, H' and $H, H' \models \Sigma$, a partial function ϕ on H is called a G-monomorphism if it preserves L-quantifier-free formulas with parameters from G.

Lemma 4.7. *Suppose $b_1, \dots b_\ell \in H$ and $c_1, \dots c_\ell \in H'$ are each linearly independent sequences (over and from G) over $\mathbb{Q}$. Let $\hat{G}$ ($\hat{H}$) be the subfield generated by $\exp(G)$ ($\exp(H)$). If*

$$\hat{G}(\exp(b_1)^{\mathbb{Q}}, \dots \exp(b_\ell)^{\mathbb{Q}}) \approx \hat{G}(\exp(c_1)^{\mathbb{Q}}, \dots \exp(c_\ell)^{\mathbb{Q}})$$

as fields, then mapping b_i to c_i is a G-monomomorphism preserving each variety V.

Proof. Let $G \subset H$, $h_i \in H - G, g_i \in G$, with $\mathbf{b} = \mathbf{hg}$ and suppose q_i, r_i are rational numbers. Note

$$H \models V(q_1h_1, \dots q_\ell h_\ell, r_1g_1, \dots r_mg_m)$$

iff

$$\hat{H} \models \hat{V}(\exp(q_1h_1), \dots \exp(h_\ell b_\ell), \exp(r_1g_1), \dots \exp(r_mg_m))$$

iff

$$\hat{G}(\exp(h_1)^{\mathbb{Q}}, \dots \exp(h_\ell)^{\mathbb{Q}}, \exp(g_1)^{\mathbb{Q}}, \dots \exp(g_m)^{\mathbb{Q}})$$
$$\models \hat{V}(\exp(q_1h_1), \dots \exp(q_\ell h_\ell), \exp(r_1g_1), \dots \exp(r_mg_m)).$$

Now apply the main hypothesis and then retrace the equivalences for $\mathbf{c}$. $\square_{4.7}$

From this fact, it is straightforward to see Condition II.2 (Assumption 3.7.2); we need that there is only one type of a closure-independent sequence. But Fact 4.7 implies that for $\mathbf{b} \in H$ to be closure independent, the associated $\exp(\mathbf{b})$ must be algebraically independent and of course there is a unique type of an algebraically independent sequence.

For Condition II.1 and the excellence condition III we need an algebraic result. We call this result, Theorem 4.10, the thumbtack lemma based on the following visualization of Kitty Holland. The various nth roots of $b_1, \dots b_m$ hang on threads from the b_i. These threads can get tangled; but the theorem asserts that by sticking in a finite number of thumbtacks one can ensure that the rest of strings fall freely.

Remark 4.8. Let k be an algebraically closed subfield in $\mathcal{C}$ and let $\boldsymbol{a} \in \mathcal{C} - k$. A field theoretic description of the relation of $\boldsymbol{a}$ to k arises by taking the irreducible variety over k realized by $\boldsymbol{a}$: $\boldsymbol{a}$ is a generic realization of the variety given by the finite conjunction $\phi(\mathbf{x}, \mathbf{b})$ of the polynomials generating the ideal in $k[\mathbf{x}]$ of those polynomials which annihilate $\boldsymbol{a}$. From a model theoretic standpoint we can say, choose $\mathbf{b}$ so that the type of $\boldsymbol{a}/k$ is the unique nonforking extension of $\mathrm{tp}(\boldsymbol{a}/\mathbf{b})$. We use the model theoretic formulation below. See [8], page 39.

The ten page proof by Zilber and Bays of the next algebraic lemma is beyond the scope of this book. (See [151, 26]). The argument in [151] left out the crucial independence hypothesis we next describe. [26] corrects this difficulty, notes the necessity of naming π, and extends the result to arbitrary characteristic. The proof involves the theory of fractional ideals of number fields, Weil divisors, and the normalization theorem.

Definition 4.9. *A set of algebraically closed fields $L_1, \dots L_n$, $L_i \subset F$, each with finite transcendence degree over $\mathbb{Q}$ is said to be* from an independent system *if there is a subfield C of F and finite set B of elements that are algebraically independent over C and so that each $L_i = \mathrm{acl}(CB_i)$ for some $B_i \subseteq B$.*

The following general version of thumbtack lemma is applied for various sets of parameters to prove different conditions of quasiminimal excellence. In Theorem 4.10 we write $\sqrt{1}$ for the group of roots of unity. If any of the L_i are defined, the reference to $\sqrt{1}$ is redundant. We write $\mathrm{gp}(\boldsymbol{a})$ for the multiplicative subgroup generated by $\boldsymbol{a} \in \mathcal{C}$.

Theorem 4.10 (Thumbtack Lemma)**.** *[151, 26] Let $P \subset \mathcal{C}$ be a finitely generated extension of $\mathbb{Q}$ and $L_1, \ldots L_n$ be from an independent system of subfields of $\hat{P} = \mathrm{acl}(P)$. Fix divisibly closed subgroups $a_1^{\mathbb{Q}}, \ldots a_r^{\mathbb{Q}}$ with $a_1, \ldots a_r \in \hat{P}$ and $b_1^{\mathbb{Q}}, \ldots b_\ell^{\mathbb{Q}} \subset \mathcal{C}^*$. If $b_1 \ldots b_\ell$ are multiplicatively independent over $\mathrm{gp}(a_1, \ldots a_r) \cdot \sqrt{1} \cdot L_1^* \cdot \ldots L_n^*$ then for some m $b_1^{\frac{1}{m}} \in b_1^{\mathbb{Q}}, \ldots b_\ell^{\frac{1}{m}} \in b_\ell^{\mathbb{Q}} \subset \mathcal{C}^*$, determine the isomorphism type of $b_1^{\mathbb{Q}}, \ldots b_\ell^{\mathbb{Q}}$ over $P(L_1, \ldots L_n, \sqrt{1}, a_1^{\mathbb{Q}}, \ldots a_r^{\mathbb{Q}})$.*

The Thumbtack Lemma immediately yields:

Lemma 4.11. *Condition II.1 ($\aleph_0$-homogeneity over models, Assumption 3.7.1) of quasiminimal excellence holds.*

Proof. We first show the case $G = \emptyset$. This proceeds in two stages. Note that for any $H \models \Sigma$ the elements of $\mathrm{dcl}_H(\emptyset)$ have the form $q\pi_H$ and $\exp(q\pi_H)$. That is, $\mathrm{dcl}_H(\emptyset)$ is $(Q, Q(\sqrt{1}))$. Thus, any partial isomorphism of an H and H' must map $q\pi_H$ to $q\pi_{H'}$ and $\exp(q\pi_H)$ to $\exp(q\pi_{H'})$ for every q. We show this map is actually an L-isomorphism. That is, we show that for any finite sequence of rationals and $\mathbf{e} = \langle q_0\pi_H, \ldots, q_s\pi_H\rangle \in H$ and $\mathbf{e}' = \langle q_0\pi_{H'}, \ldots, q_s\pi_{H'}\rangle \in H'$, $\mathrm{qftp}_L(\exp(\mathbf{e})) = \mathrm{qftp}_L(\exp(\mathbf{e}'))$. For sufficiently large N, there are primitive Nth roots of unity, η, η', such that $\exp(\mathbf{e}) \subseteq \mathbb{Q}(\eta)$ and similarly $\exp(\mathbf{e}') \subseteq \mathbb{Q}(\eta')$. Thus each $q_i\pi_H = \eta^{m_i}$ for appropriate m_i and $q_i\pi_{H'} = \eta'^{m_i}$. Now mapping η to η' induces an isomorphism between two isomorphic copies of the field $\mathbb{Q}[x]/(x^N - 1)$. So there is a unique isomorphism between $\mathrm{dcl}_H(\emptyset)$ and $\mathrm{dcl}_{H'}(\emptyset)$, with domain Q. Now to verify Condition II.1 on the rest of H, apply Lemma 4.10 with $P = \mathbb{Q}$, no L's and the x_i as the a_i and y as b_1.

It remains to show: If $G \models \Sigma$ and $f = (\langle x_1, x_1'\rangle \ldots \langle x_r, x_r'\rangle)$ is a partial G-monomorphism from H to H' then for any $y \in H$ there is y' in some H'' with $H' \prec_{\boldsymbol{K}} H''$ such that $f \cup \{\langle y, y'\rangle\}$ extends f to a partial G-monomorphism. Since $G \models \Sigma$, $\exp(G) = \hat{G}$ is an algebraically closed field that is in the domain of any G-monomorphism. We work in the ω-stable theory of this field (See Remark 4.8.) For each i, let a_i denote $\exp(x_i)$ and similarly for x_i', a_i'. Choose a finite sequence $\mathbf{d} \in \exp(G)$ such that the sequences $(a_1, \ldots a_r, y)$ and $(a_1', \ldots a_r')$ are each independent (in the forking sense) from $\exp(G)$ over $\mathbf{d}$ and $\mathrm{tp}((a_1, \ldots a_r)/\mathbf{d})$ is stationary. Now we apply the thumbtack lemma. Let P_0 be $\mathbb{Q}(\mathbf{d})$. Let $n = 1$ and L_1 be the algebraic closure of P_0. We set $P_0(\mathbf{d}, a_1, \ldots a_r)$ as P. Take b_1 as $\exp(y)$ and set $\ell = 1$.

Now apply Lemma 4.10 to find m so that $b_1^{\frac{1}{m}}$ determines the isomorphism type of $b_1^{\mathbb{Q}}$ over $L_1(a_1^{\mathbb{Q}}, \ldots a_r^{\mathbb{Q}}) = P_0(L_1, a_1^{\mathbb{Q}}, \ldots a_r^{\mathbb{Q}})$. Let $\hat{f}$ denote the (partial) map f induces from $\exp(H)$ to $\exp(H')$ over $\exp(G)$. Choose $b_1'^{\frac{1}{m}}$ to satisfy the quantifier-free field type of $\hat{f}(\mathrm{tp}(b_1^{\frac{1}{m}}/L_1(a_1^{\mathbb{Q}}, \ldots a_r^{\mathbb{Q}}))$ and with $(a_1', \ldots a_\ell', b_1'^{\frac{1}{m}})$ independent from $\hat{G}$ over $\mathbf{d}$. Now by Lemma 4.10, $\hat{f}$ extends to field isomorphism between $L_1(a_1^{\mathbb{Q}}, \ldots a_r^{\mathbb{Q}}, b_1^{\mathbb{Q}})$ and $L_1((a_1')^{\mathbb{Q}}, \ldots (a_r')^{\mathbb{Q}}, (b_1')^{\mathbb{Q}})$. Since the sequence

$(a_1, \ldots a_r, b_1{}^{\frac{1}{m}})$ and $(a'_1, \ldots a'_r, b'_1{}^{\frac{1}{m}})$ are each independent (in the forking sense) from $\exp(G)$ over L_1, we can extend this map to take $\exp(G)(a_1^{\mathbb{Q}}, \ldots a_r^{\mathbb{Q}}, b_1^{\mathbb{Q}})$ to $\exp(G)((a'_1)^{\mathbb{Q}}, \ldots (a'_r)^{\mathbb{Q}}, (b'_1)^{\mathbb{Q}})$ and pull back to find y'; this suffices by Fact 4.7. $\square_{4.11}$

Note there is no claim that $y' \in H'$ and there can't be.

One of the key ideas discovered by Shelah in the investigation of non-elementary classes is that in order for types to be well-behaved one may have to make restrictions on the domain. (E.g., we may have few types over models but not over arbitrary substructures.) This principle is illustrated by the following result of Zilber, proving Condition III, excellence for covers of $(\mathcal{C}^*, \cdot)$. Note that in Lemma 4.12, $G_1, \ldots G_n$ form an independent system in the sense of the combinatorial geometry just if the $L_i = \exp(G_i)$ are independent in the sense of Definition 4.9.

Lemma 4.12. *Suppose $Z = \{G_1, \ldots G_n\}$ where each $G_i \subset H$ is an n-dimensional independent system. If $h_1, \ldots h_\ell \in G^- = \mathrm{cl}(G_1 \cup \ldots G_n)$ then there is finite set $A \subset G^-$ such that any ϕ taking $h_1, \ldots, h_\ell$ into H which is an A-monomorphism is also a G^--monomorphism.*

Proof. Let $L_i = \exp(G_i)$ for $i = 1, \ldots, n$; $b_j^q = \exp(qh_j)$ for $j = 1, \ldots, \ell$ and $q \in \mathbb{Q}$. We may assume the h_i are linearly independent over the vector space generated by the G_i; this implies the b_i are muliplicatively independent over $L_1^* \cdot L_2^* \cdot \ldots L_n^*$. Now apply the thumbtack lemma with $r = 0$. This gives an m such that the field theoretic type of $b_1^{\frac{1}{m}}, \ldots, b_\ell^{\frac{1}{m}}$ determines the quantifier free type of $(h_1, \ldots, h_\ell)$ over G^-. So we need only finitely many parameters from G^- and we finish. $\square_{4.12}$

Proof of Theorem 4.3. We just show that Σ defines a quasiminimal excellent class with the countable closure condition. The homogeneity conditions of Condition II.1 (Assumption 3.7) were established in Lemma 4.11. The fact that closure forms a combinatorial geometry (Assumption 3.1.1,2) with the countable closures is immediate from the definition of closure (4.2). Assumption 3.1.3 (quantifier elimination) holds by the first part of the argument for Lemma 4.11 and since we added to the language of Σ predicates for the pull-back of all quantifier-free relations on the field F. Corollary 4.12 asserts Excellence (Assumption 3.15). So we finish by Theorem 3.19; the quasimimimal excellence implies categoricity. $\square_{4.3}$

This completes the proof that covers of the multiplicative group form an example of a quasiminimal excellent class. We now investigate the inhomogeneity aspects of this example. In Part I, homogeneity generally means sequence homogeneity in the following sense. Parts II and III usually consider the more restricted notion of model homogeneity.

Definition 4.13. *A structure M is κ-sequence homogeneous if for any $\boldsymbol{a}, \mathbf{b} \in M$ of length less than κ, if $(M, \boldsymbol{a}) \equiv (M, \mathbf{b})$ then for every c, there exists d such* *that* $(M, \boldsymbol{a}c) \equiv (M, \mathbf{b}d)$. *Usually, the 'sequence' is omitted and one just says* κ*-homogeneous.*

Keisler[80] generalized Morley's categoricity theorem to sentences in $L_{\omega_1,\omega}$, assuming that the categoricity model was $\aleph_1$-homogeneous. This theorem is the origin of the study of homogeneous model theory which is well expounded in e.g. [31]. We now give simple model theoretic examples showing the homogeneity does not follow from categoricity. Marcus and later Julia Knight [107, 85] (details in Example 19.9) showed:

Fact 4.14. *There is a first order theory* T *with a prime model* M *such that*

1. M *has no proper elementary submodel.*
2. M *contains an infinite set of indiscernibles.*

Exercise 4.15. *Show that the* $L_{\omega_1,\omega}$*-sentence satisfied only by atomic models of the theory* T *in Fact 4.14 has a unique model.*

Example 4.16. Now construct an $L_{\omega_1,\omega}$-sentence ψ whose models are partitioned into two sets; on one side is an atomic model of T, on the other is an infinite set. Then ψ is categorical in all infinite cardinalities but no model is $\aleph_1$-homogeneous because there is a countably infinite maximal indiscernible set.

Example 4.17. Now we see that the example of this chapter has the same inhomogeneity property. Consider the basic diagram:

$$0 \to Z \to H \to F^* \to 0. \tag{4.2}$$

Let a be a transcendental number in F^*. Fix h with $\exp(h) = a$ and define $a_n = \exp(\frac{h}{n}) + 1$ for each n. Now choose h_n so that $\exp(h_n) = a_n$. Let $X_r = \{h_i : i \leq r\}$ and let $X = \bigcup_r X_r$. Note that $a_m = a^{\frac{1}{m}} + 1$ where we have chosen a specific mth root.

Claim 4.18. *For each* r, $\mathrm{tp}(h/X_r)$ *is a principal type, but,* $\mathrm{tp}_{L_{\omega_1,\omega}}(h/X)$ *is not implied by its restriction to any finite set.*

Proof. We make another application of the thumbtack Lemma 4.10 with $\mathbb{Q}(\exp(\mathrm{span}(X_r))$ as P , $a_1, \ldots a_r$ as themselves, all L_i are empty, and a as b_1. By the lemma there is an m such that $a^{\frac{1}{m}}$ determines the isomorphism type of $a^{\mathbb{Q}}$ over $P(a_1^{\mathbb{Q}}, \ldots a_r^{\mathbb{Q}})$. That is, if ϕ_m is the minimal polynomial of $a^{\frac{1}{m}}$ over P, $(\exists y)\phi_m(y) \wedge y^m = x$ generates $\mathrm{tp}(a/\exp(\mathrm{span}(X_r)))$. Pulling back by Lemma 4.7, we see $\mathrm{tp}(h/X_r)$ is principal and even complete for $L_{\omega_1,\omega}$. In particular, for any $m' \geq m$, any two m'th roots of a have the same type over $\exp(X_r)$. But for sufficiently large s, one of these m'th roots is actually in X_{rs} so for each r, $p_r = \mathrm{tp}(a/X_r)$ does not imply $p = \mathrm{tp}(a/X)$. Any realization c of p_r, behaves like a; i.e. $\mathrm{tp}(c/X)$ is not isolated. That is, $\mathrm{tp}(a/X)$ is not implied by its restriction to any finite set. And by Lemma 4.7, this implies $\mathrm{tp}_{L_{\omega_1,\omega}}(h/X)$ is not implied by its restriction to any finite set. $\square_{4.18}$

Now specifically to answer the question of Keisler [80], page 123, for this example we need to show there is a sentence ψ in a countable fragment L^* of $L_{\omega_1,\omega}$ such that ψ is $\aleph_1$-categorical but has a model which is not $(\aleph_1, L^*)$-homogeneous. Let L^* be a countable fragment containing the categoricity sentence for 'covers'. Fix H, X and p as in Claim 4.18. We have shown no formula of $L_{\omega_1,\omega}$ (let alone L^*) with finitely many parameters from X implies p. By the omitting types theorem for L^*, there is a countable model H_0 of ψ which contains an L^*-equivalent copy X' of X and omits the associated p'. By categoricity, H_0 imbeds into H. But H also omits p'. As, if $h' \in H$, realizes p', then $\exp(h') \in \mathrm{acl}(\exp(X')) \subseteq H_0$, as $\exp(h') = (a'_n - 1)^n$. So since the kernel of $\exp$ is standard, in any extension of H_0, any choice of a logarithm of $\exp(h')$ is in H_0. In particular, $h' \in H_0$, which is a contradiction. Thus the type p' cannot be realized and H is not homogeneous since the map from X to X' cannot be extended.

Exercise 4.19. *Show that X in Example 4.17 is a set of indiscernibles.*

Remark 4.20 (Extensions)**.** This work can be extended in at least two directions.

1. The first is to study covers of other algebraic groups. The first step in this work is to replace $(C^*, \cdot)$ by the multiplicative group of the algebraic closure of $\mathcal{Z}_p$; this was accomplished by [26]. A more ambitious goal is to replace the multiplicative group of the field by a general semi-abelian variety. Continuing work on elliptic curves is due to Gavrilovich [41] and Bays's forthcoming thesis. For higher dimensional varieties, one needs the notion of almost quasiminimality introduced in Chapter 3. But, defining $\mathrm{cl}_H(X)$ as $\exp^{-1}(\mathrm{acl}(\exp(X))$ fails exchange. The solution in [148] uses the inverse image of a rank 1 subvariety as the quasiminimal set whose closure is the universe. This work is particularly exciting as in [148], Zilber draws specific 'arithmetic' (in the sense of algebraic geometry) consequences for semi-abelian varieties from categoricity properties of the associated short exact sequence. These conclusions depend on Shelah's deduction of excellence from categoricity up to $\aleph_\omega$, which we expound in Part IV. The categoricity (or not) of various higher dimensional varieties is very open and will in many cases require extensions of the language.

2. The most ambitious aim of Zilber's program is to realize $(\mathcal{C}, +, \cdot, \exp)$ as a model of an $L_{\omega_1,\omega}$-sentence discovered by the Hrushovski construction. Here are two steps towards this objective.

Objective A. Model theory: Using a Hrushovski-like dimension function (Example 5.9) expand $(\mathcal{C}, +, \cdot)$ by a unary function f which behaves like exponentiation. Prove that the theory Σ of $(\mathcal{C}, +, \cdot, f)$ in an appropriate logic is well behaved.

This objective is realized in [149]. A summary connecting Zilber's program with this monograph appears in [9], updated in [3].

Objective B. Algebra and analysis: Prove $(\mathcal{C}, +, \cdot, \exp)$ is a model of the sentence Σ found in Objective A.

Attempting to resolve this objective raises enormous and interesting problems in both algebraic geometry and complex analysis; see [149, 147, 109].

Part II

Abstract Elementary Classes

This is page 34
Printer: Opaque this

'Non-elementary classes' is a general term for any logic other than first order. Some of the most natural extensions of first order logic arise by allowing conjunctions of various infinite lengths or cardinality quantifiers. In this chapter we introduce a precise notion of 'abstract elementary class' (AEC) which generalizes the notion of a definable class in some of these logics. In this monograph we pursue a dual track of proving certain results for general AEC and some for very specific logics, especially $L_{\omega_1,\omega}$ and $L_{\omega_1,\omega}(Q)$. A natural source of further examples is [24]. We show in this Part the relation between these various logics and the more general notion of an AEC. The main purpose of introducing the notion of AEC is to provide a common semantic framework to prove results for many (infinitary) logics without the distraction of specific syntax.

We introduce the notion of an AEC in Chapter 5 and provide a number of examples. We prove Shelah's presentation theorem which allows us to use the technology of Ehrenfeucht-Mostowski models. We give a detailed account of the role of Hanf numbers in AEC. These results and related terminology are fundamental for Part III. We survey the connections of AEC with other extensions of first order logic such as continuous logic and homogeneous model theory.

Chapter 6 contains the proof of the Lopez-Escobar theorem, Theorem 6.1.6. Although this result, which asserts that $\aleph_1$ is a bound on the length of well-orderings that can be defined in $L_{\omega_1,\omega}(Q)$, has independent interest, the principal application in this monograph is as a tool for further analysis as we mention in the next few paragraphs. We analyze the difficulties in representing the models of an $L_{\omega_1,\omega}(Q)$-sentence as an AEC with Löwenheim number $\aleph_0$. In the second section we present Keisler's fundamental result that a sentence of $L_{\omega_1,\omega}(Q)$ with less that $2^{\aleph_1}$ models of cardinality $\aleph_1$ has only countably many syntactic types over the empty set.

In Chapter 7, we explore in detail the notion of a 'complete theory' in an infinitary logic. We see that the easy remark in first order logic that 'categoricity implies completeness' becomes problematic. Indeed the whole notion that a structure uniquely determines 'its theory' becomes problematic. We apply the Lopez-Escobar Theorem 6.1.6 to prove Theorem 7.3.2, which asserts that a sentence in $L_{\omega_1,\omega}(Q)$ that has few models in power $\aleph_1$ is implied by a 'complete' sentence of $L_{\omega_1,\omega}(Q)$ that has an uncountable model. This is an essential step in preparing for the analysis of categoricity for sentences of $L_{\omega_1,\omega}$ that takes place in Part IV. And in Chapter 8, we apply Theorem 6.1.6 again to present Shelah's beautiful proof (in ZFC) that if a sentence in $L_{\omega_1,\omega}(Q)$ is $\aleph_1$-categorical then it has a model in $\aleph_2$. In Chapter 20, we will strengthen this result by replacing $\aleph_1$-categorical by ω-stable but at the cost of developing a certain amount of stability theory for $L_{\omega_1,\omega}$.

5
Abstract Elementary Classes

In this chapter we introduce the semantic notion of an Abstract Elementary Class (AEC) and give a number of examples to introduce the concept. One important feature distinguishing abstract elementary classes from elementary classes is that they need not have arbitrarily large models. We prove a surprising syntactic representation theorem for such classes. This representation provides a sufficient condition for an AEC to have arbitrarily large models.

We always work in the context of the class of all structures in a fixed vocabulary τ . Structure is taken in the usual sense: a set with specified interpretations for the relation, constant and function symbols of τ.

When Jónsson generalized the Fraïssé construction to uncountable cardinalities [76, 77], he did so by describing a collection of axioms, which might be satisfied by a class of models, that guaranteed the existence of a homogeneous-universal model; the substructure relation was an integral part of this description. Morley and Vaught [112] replaced substructure by elementary submodel and developed the notion of saturated model. Shelah [139, 140] generalized this approach in two ways. He moved the amalgamation property from a basic axiom to a constraint to be considered. (But this was a common practice in universal algebra as well.) He made the *substructure* notion a 'free variable' and introduced the notion of an *Abstract Elementary Class*: a class of structures and a 'strong' substructure relation which satisfied variants on Jonsson's axioms.

The most natural exemplar of the axioms below is the class of models of a complete first order theory with $\prec_{\boldsymbol{K}}$ as first order elementary submodel. In Part IV we will study AEC given as the models of a sentence ϕ in $L_{\omega_1,\omega}$ with $\prec_{\boldsymbol{K}}$ as elementary submodel in fragment generated by ϕ. The quasiminimal excellent classes of Part I are another example and we will discuss many other examples in the course of the book.

In the following a *chain* $\langle A_i : i < \delta\rangle$ is a sequence of members of $\boldsymbol{K}$ such that if $i < j, A_i \prec_{\boldsymbol{K}} A_j$; the chain is continuous if for each limit ordinal α, $A_\alpha = \bigcup_{\beta<\alpha} A_\beta$.

Definition 5.1. *A class of τ-structures equipped with a notion of 'strong submodel', $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$, is said to be an* abstract elementary class (*AEC*) *if the class $\boldsymbol{K}$ and class of pairs satisfying the binary relation $\prec_{\boldsymbol{K}}$ are each closed under isomorphism and satisfy the following conditions.*

- **A1**. *If $M \prec_{\boldsymbol{K}} N$ then $M \subseteq N$.*
- **A2**. *$\prec_{\boldsymbol{K}}$ is a partial order (i.e. a reflexive and transitive binary relation) on $\boldsymbol{K}$.*
- **A3**. *If $\langle A_i : i < \delta\rangle$ is a continuous $\prec_{\boldsymbol{K}}$-increasing chain:*
  1. *$\bigcup_{i<\delta} A_i \in \boldsymbol{K}$;*
  2. *for each $j < \delta$, $A_j \prec_{\boldsymbol{K}} \bigcup_{i<\delta} A_i$*
  3. *if each $A_i \prec_{\boldsymbol{K}} M \in \boldsymbol{K}$ then $\bigcup_{i<\delta} A_i \prec_{\boldsymbol{K}} M$.*
- **A4**. *If $A, B, C \in \boldsymbol{K}$, $A \prec_{\boldsymbol{K}} C$, $B \prec_{\boldsymbol{K}} C$ and $A \subseteq B$ then $A \prec_{\boldsymbol{K}} B$.*
- **A5**. *There is a* Löwenheim-Skolem number $\mathrm{LS}(\boldsymbol{K})$ *such that if $A \subseteq B \in K$ there is an $A' \in \boldsymbol{K}$ with $A \subseteq A' \prec_{\boldsymbol{K}} B$ and $|A'| \leq |A| + \mathrm{LS}(\boldsymbol{K})$.*

Exercise 5.2. *Show that the same $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ are AEC if the word continuous is deleted from the hypothesis of* **A3**.

We refer to **A4** as the coherence property; it is sometimes called 'the funny axiom' or, since it is easily seen to follow in the first order case as an application of the Tarski-Vaught test for elementary submodel, the Tarski-Vaught property. However, Shelah sometimes uses 'Tarski-Vaught' for the union axioms. We frequently write *Löwenheim-number* for Löwenheim-Skolem number as we are representing only the downward aspect of the Löwenheim-Skolem theorem. For simplicity, we usually assume the Löwenheim-number is at least $|\tau|$.

Notation 5.3.
1. *By a direct union we will mean a direct limit of a directed system $\langle M_i, f_{ij} : i, j \in I\rangle$ where each $M_i \in \boldsymbol{K}$ and each f_{ij} is a $\prec_{\boldsymbol{K}}$-inclusion (so $i < j$ implies M_i is a subset of M_j).*
2. *If $f : M \mapsto N$ is* 1-1 *and $fM \prec_{\boldsymbol{K}} N$, we call f a $\boldsymbol{K}$-*embedding *or* strong embedding *or sometimes carelessly just embedding.*
3. *A $\prec_{\boldsymbol{K}}$-direct limit of members of $\boldsymbol{K}$ is a direct limit of a directed system $\langle M_i, f_{ij} : i, j \in I\rangle$ with each $M_i \in \boldsymbol{K}$ and each f_{ij} is a $\boldsymbol{K}$-embedding.*

Note that while reflexivity and transitivity from **A2** are true for strong embeddings as well as strong submodels, antisymmetry need not be. The following exercise often simplifies notation.

Exercise 5.4. *Let $\boldsymbol{K}$ be an AEC. Show that if there is a strong embedding of $M \in \boldsymbol{K}$ into $N \in \boldsymbol{K}$, then there is an $M' \in \boldsymbol{K}$ which extends M and is isomorphic to N.*

Grossberg and Shelah pointed out in [52], an AEC can be viewed as a concrete category. A precise formulation of this idea involves describing the relations between two classes of morphism: the τ-embeddings and the morphisms that represent strong embeddings. The inductive argument that shows that closure under well-ordered direct limits implies closure under arbitrary direct limits (See e.g. Theorem 21.5 of [42].) translates in our situation to the following straightforward Lemma.

Lemma 5.5. *Any AEC is closed under directed unions and closed under $\prec_{\boldsymbol{K}}$-direct limits.*

Note however that Grätzer proves that the concrete construction of a direct limit satisfies the universal mapping property for direct limits of 1-1 homomorphisms. In our context **A.3.3** is the assertion that the well-ordered union satisfies the universal mapping property of direct limits of morphisms. So the categorical version of Lemma 5.5 is: a category that is closed under well-ordered 1-1 direct limits is closed under arbitrary 1-1 direct limits. In general one can not pass from closure under 1-1 direct limits to arbitrary direct limits; however, an interesting collection of AEC where this is possible is discussed in [4].

Exercise 5.6. *Show the class of well-orderings with $\prec_{\boldsymbol{K}}$ taken as end extension satisfies the first four properties of an AEC. Does it have a Löwenheim-number?*

Exercise 5.7. *Show that any model M in an AEC $\boldsymbol{K}$ admits a* filtration*; that is, it can be written as a continuous increasing chain of submodels M_i with $|M_i| \leq |i| + \mathrm{LS}(\boldsymbol{K})$, such that for $i < j$, $M_i \prec_{\boldsymbol{K}} M_j$ and all $M_i \prec_{\boldsymbol{K}} M$. Note that* **A.3.3** *is convenient but not essential for this decomposition.*

Exercise 5.8. *The models of a sentence of first order logic or any countable fragment of $L_{\omega_1,\omega}$ with the associated notion of elementary submodel as $\prec_{\boldsymbol{K}}$ gives an AEC with Löwenheim-number $\aleph_0$.*

Classes given by sentences of the logics $L(Q)$ and $L_{\omega_1,\omega}(Q)$ (defined in Chapter 2) are not immediately seen as AEC. We discuss the several approaches in Chapters 6, 7, and 8. One of the most fertile sources of examples is the Hrushovski construction.

Example 5.9 (Hrushovski Construction)**.** If $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ is derived from a dimension function in the standard way (e.g. [2, 9, 20]), the closure of the class $\boldsymbol{K}$ under unions of $\prec_{\boldsymbol{K}}$-chains is an AEC. More precisely, let $\langle K(N), \wedge, \vee\rangle$ be a lattice of substructures of a model N. For present purposes a rank is a function δ from

$K(N)$ to a discrete subgroup of the reals ($\mathcal{R}$), which is defined on each N in a class $\boldsymbol{K}$. We write $\delta(A/B) = \delta(A \vee B) - \delta(B)$ to indicate the relativization of the rank. We demand *only* that δ is lower semimodular:

$$\delta(A \vee B) - \delta(B) \leq \delta(A) - \delta(A \wedge B).$$

For $A, B \in \boldsymbol{K}$, we say A is a *strong substructure* of B and write $A \prec_{\boldsymbol{K}} B$ if for every $B' \in \boldsymbol{K}_{-1}$ with $B' \subseteq B$, $\delta(B'/B' \cap A) \geq 0$. Under reasonable conditions, this class is an AEC and the first order theory of the $\prec_{\boldsymbol{K}}$-homogeneous-universal for various classes $\boldsymbol{K}$ and choices of δ provide a series of interesting examples. Many are discussed in the surveys mentioned above. Zilber refined this construction by requiring certain sets to be fixed under strong extensions. In particular, he [149] exploits this construction to explore complex exponentiation. Further in [113], Zilber tries to explain these examples as *analytic Zariski structures*; he links this notion more closely to AEC in [146].

We will see below the importance of several properties that an AEC may possess.

Definition 5.10. *Let $\boldsymbol{K}$ be an AEC.*

1. *$\boldsymbol{K}$ has the* amalgamation property *if $M \prec_{\boldsymbol{K}} N_1$ and $M \prec_{\boldsymbol{K}} N_2 \in \boldsymbol{K}$ with all three in $\boldsymbol{K}$ implies there is an N_3 into which both N_1 and N_2 can be strongly embedded over M.*

2. *$\boldsymbol{K}$ has the* disjoint amalgamation property *if the images of the imbedding in part i) intersect in the image of N.*

3. Joint embedding *means any two members of $\boldsymbol{K}$ can be strongly embedded some member of $\boldsymbol{K}$.*

The following is an exercise in 'renaming'.

Exercise 5.11. *More formally, $\boldsymbol{K}$ has the* amalgamation property *if for every pair of strong embeddings from $M \in \boldsymbol{K}$ into N_1, N_2 there exists a model $N \in \boldsymbol{K}$ and g_1, g_2 with g_i a strong embedding of N_i into N such that $f_1 g_1 = f_2 g_2$. Show that an equivalent condition arises if f_1, f_2 are assumed to be identity maps and that one of the g_i can then be chosen as the identity.*

Remark 5.12. The generality of the notion of AEC can be seen by some rather strange ways that examples can be generated. On the one hand, taking a universal-existential first order theory T, and setting $\prec_{\boldsymbol{K}}$ as just substructure provides some mischievous counterexamples (Example 8.6). One can also restrict to the existentially closed models of T, also with substructure; in this case, if T has amalgamation, this is known as a Robinson theory [121, 62].

On the other hand if $\boldsymbol{K}$ is any class of models which is closed under elementary equivalence, taking $\prec_{\boldsymbol{K}}$ as elementary submodel yields an AEC. Thus, any class defined by an infinite disjunction of first order sentences or even an infinite disjunction of first order theories becomes an AEC. In particular, the class of

Artinian (descending chain condition) commutative rings with unit becomes an AEC under elementary submodel. (See [7].) This class is not sufficiently cohesive to really accord with the intuition of an AEC as generalizing the notion of a complete theory. This is not an accident. If $\prec_{\boldsymbol{K}}$ is elementary embedding, $\boldsymbol{K}$ has the amalgamation property by the standard first order argument. We can then use elementary equivalence to split the given AEC into a family of disjoint first order classes and restrict to the study of these. (Compare Lemma 17.14.)

Although in studying categoricity it is natural to think in terms of complete theories, AEC also are a natural generalization of varieties or universal Horn classes. Further interesting examples arise from the notion of cotorsion theories in the study of modules [4, 39, 53, 142]

Example 5.13. The set $\boldsymbol{K} = \{\alpha : \alpha \leq \aleph_1\}$ with $\prec_{\boldsymbol{K}}$ as initial segment is an AEC with $\aleph_1$ countable models. It is $\aleph_1$-categorical and satisfies both amalgamation and joint embedding.

We call the next result: the presentation theorem. Shelah's result [139] allows us to replace the entirely semantic description of an abstract elementary class by a syntactic one. Remarkably, the notion of an AEC, which is designed to give a version of the Fraïssé construction and thus saturated models, also turns out to allow via Theorem 5.14 the use of the second great model theoretic technique of the 50's: Ehrenfeucht-Mostowski models [38].

The proof of the presentation theorem can be thought of as having two stages. By adding $\mathrm{LS}(\boldsymbol{K})$ function symbols to form a language τ' we can regard each model of cardinality at most $\mathrm{LS}(\boldsymbol{K})$ as being finitely generated. If we look at finitely generated τ'-structures, the question of whether the structure is in $\boldsymbol{K}$ is a property of the quantifier free type of the generators. Similarly the question of whether one τ' finitely generated structure is strong in another is a property of the τ' type of the generators of the larger model. Thus, we can determine membership in $\boldsymbol{K}$ and strong submodel for finitely generated (and so all models of cardinality $\mathrm{LS}(\boldsymbol{K})$) by omitting types. But every model is a direct limit of finitely generated models so using the AEC axioms on unions of chains (and coherence) we can extend this representation to models of all cardinalities.

Theorem 5.14 (Shelah's Presentation Theorem)**.** *If K is an AEC with Löwenheim-number $\mathrm{LS}(\boldsymbol{K})$ (in a vocabulary τ with $|\tau| \leq \mathrm{LS}(\boldsymbol{K})$), there is a vocabulary $\tau' \supseteq \tau$ with cardinality $|\mathrm{LS}(\boldsymbol{K})|$, a first order τ'-theory T' and a set Γ of at most $2^{\mathrm{LS}(\boldsymbol{K})}$ partial types such that:*

$$\boldsymbol{K} = \{M' \restriction \tau : M' \models T' \text{ and } M' \text{ omits } \Gamma\}.$$

Moreover, the $\prec_{\boldsymbol{K}}$ relation satisfies the following conditions:

1. *if M' is a τ'-substructure of N' where M', N' satisfy T' and omit Γ then $M' \restriction \tau \prec_{\boldsymbol{K}} N' \restriction \tau$;*
2. *if $M \prec_{\boldsymbol{K}} N$ there is an expansion of N to a τ'-structure such that M is the universe of a τ'-substructure of N;*

3. *Finally, the class of pairs* (M, N) *with* $M \prec_{\boldsymbol{K}} N$ *forms a* $PC\Gamma(\mathrm{LS}(\boldsymbol{K}), 2^{\mathrm{LS}(\boldsymbol{K})})$*-class in the sense of Definition 5.26.*

   Without loss of generality we can guarantee that T' *has Skolem functions.*

Proof. Let τ' contain n-ary function symbols F_i^n for $n < \omega$ and $i < \mathrm{LS}(\boldsymbol{K})$. We take as T' the theory which asserts that for each $i < \lg(\boldsymbol{a})$, $F_i^n(\boldsymbol{a}) = a_i$. (This just says that the first n elements enumerated in the substructure generated by the n-tuple $\boldsymbol{a}$ are the components of $\boldsymbol{a}$.) For any τ'-structure M' and any $\boldsymbol{a} \in M'$, let $M'_{\boldsymbol{a}}$ denote the subset of M' enumerated as $\{F_i^n(\boldsymbol{a}) : i < \mathrm{LS}(\boldsymbol{K})\}$ where $n = \lg(\boldsymbol{a})$. We also write $M'_{\boldsymbol{a}}$ for the partial τ'-structure induced on this set from M'. Note that $M'_{\boldsymbol{a}}$ may not be either a τ' or even a τ-structure since it may not be closed under all the operations. The isomorphism type of $M'_{\boldsymbol{a}}$ (and thus whether M'_a is in fact a τ'-structure) is determined by the quantifier free τ'-type of $\boldsymbol{a}$. We use the notations: $M_{\boldsymbol{a}}$ and $M'_{\boldsymbol{a}}$ throughout the proof. Let Γ be the set of quantifier free τ'-types of finite tuples $\boldsymbol{a}$ such that $M'_{\boldsymbol{a}} \restriction \tau \notin \boldsymbol{K}$ or for some $\mathbf{b} \subset \boldsymbol{a}$, $M'_{\mathbf{b}} \restriction \tau \not\prec_{\boldsymbol{K}} M'_{\boldsymbol{a}} \restriction \tau$.

We claim T' and Γ suffice. That is, if

$$\boldsymbol{K}' = \{M' \restriction \tau : M' \models T' \text{ and } M' \text{ omits } \Gamma\}$$

then $\boldsymbol{K} = \boldsymbol{K}'$.

Let the τ'-structure M' omit Γ; in particular, each $M'_{\boldsymbol{a}} \restriction \tau$ is a τ-structure. Write M' as a direct limit of the finitely generated partial τ'-structures $M'_{\boldsymbol{a}}$. (These may not be closed under the operations of τ'.) By the choice of Γ, each $M'_{\boldsymbol{a}} \restriction \tau \in \boldsymbol{K}$ and if $\boldsymbol{a} \subseteq \boldsymbol{a}'$, $M'_{\boldsymbol{a}} \restriction \tau \prec_{\boldsymbol{K}} M'_{\boldsymbol{a}'} \restriction \tau$, and so by the unions of chains axiom (Definition 5.1 A3.1) $M' \restriction \tau \in \boldsymbol{K}$.

Conversely, if $M \in \boldsymbol{K}$ we define by induction on $|\boldsymbol{a}|$, τ-structures $M_{\boldsymbol{a}}$ for each finite subset $\boldsymbol{a}$ of M and expansions of $M_{\boldsymbol{a}}$ to partial τ'-structures $M'_{\boldsymbol{a}}$. Let $M_\emptyset$ be any $\prec_{\boldsymbol{K}}$-substructure of M with cardinality $\mathrm{LS}(\boldsymbol{K})$ and let the $\{F_i^0 : i < \mathrm{LS}(\boldsymbol{K})\}$ be constants enumerating the universe of $M_\emptyset$. Given a sequence $\mathbf{b}$ of length $n+1$, choose $M_{\mathbf{b}} \prec_{\boldsymbol{K}} M$ with cardinality $\mathrm{LS}(\boldsymbol{K})$ containing all the $M_{\boldsymbol{a}}$ for $\boldsymbol{a}$ a proper subsequence of $\mathbf{b}$. Let $\{F_i^{n+1}(\mathbf{b}) : i < \mathrm{LS}(\boldsymbol{K})\}$ enumerate the universe of $M_{\mathbf{b}}$. Now each $M_{\boldsymbol{a}} \restriction \tau \in \boldsymbol{K}$ and if $\mathbf{b} \subset \mathbf{c}$, $M_{\mathbf{b}} \prec_{\boldsymbol{K}} M_{\mathbf{c}}$ by the coherence property so M' omits Γ as required.

Now we consider the moreover clause. For 1) we have M' is a direct limit of finitely generated partial τ'-structures $M'_{\boldsymbol{a}}$ and N' is a $\prec_{\boldsymbol{K}}$-direct limit of $N'_{\boldsymbol{a}}$ where $M'_{\boldsymbol{a}} = N'_{\boldsymbol{a}}$ for $\boldsymbol{a} \in M$ because $M' \restriction \tau$ is a τ-substructure of $N' \restriction \tau$. Each $M'_{\boldsymbol{a}} \restriction \tau \prec_{\boldsymbol{K}} N' \restriction \tau$ so, using Definition 5.1 A3.3 and Lemma 5.5, the direct limit $M' \restriction \tau$ is a strong submodel of $N' \restriction \tau$. For 2), just be careful in carrying out the expansion of N to a τ' structure in the paragraph beginning 'conversely', that if $\boldsymbol{a} \in M'$, $M_{\boldsymbol{a}} \subseteq M$. The third clause is now evident. And since T' is a universal theory we can Skolemize without changing the class of universes of substructures of models. $\square_{5.14}$

Remark 5.15. *1. There is no use of amalgamation in this theorem.*

2. *The only penalty for increasing the size of the language or the Löwenheim number is that the size of L' and the number of types omitted may increase.*
3. *Clauses 1 and 2 of Theorem 5.14 are more useful than Clause 3; see Theorem 9.17 and its applications.*

Much of this book is concerned with the *spectrum* of an AEC: the function which tells us the number of members of $\boldsymbol{K}$ in each cardinality. Formally,

Notation 5.16. *For any class of models $\boldsymbol{K}$, $I(\boldsymbol{K}, \lambda)$ denotes the number of isomorphism types of members of $\boldsymbol{K}$ with cardinality λ.*

Using the presentation theorem we have a first basic fact about the spectrum. Before proving it, we introduce the general notion of a 'Hanf number'. Hanf [54] introduced the following extremely general and soft argument. $P(\boldsymbol{K}, \lambda)$ ranges over such properties as $\boldsymbol{K}$ has a model in cardinality λ, $\boldsymbol{K}$ is categorical in λ, or the type q is omitted in some model of $\boldsymbol{K}$ of cardinality λ.

Theorem 5.17 (Hanf)**.** *If there is a set of classes $\boldsymbol{K}$ of given kind (e.g. defined by sentences of $L_{\mu,\nu}$ for some fixed μ, ν) of a given similarity type then for any property $P(\boldsymbol{K}, \lambda)$ there is a cardinal κ such that if $P(\boldsymbol{K}, \lambda)$ holds for some $\lambda > \kappa$ then $P(\boldsymbol{K}, \lambda)$ holds for arbitrarily large λ.*

Proof. Let

$$\mu_{\boldsymbol{K}} = \sup\{\lambda : P(\boldsymbol{K}, \lambda) \text{ holds if there is such a max}\}$$

then

$$\kappa = \sup \mu_{\boldsymbol{K}}$$

as $\boldsymbol{K}$ ranges over the set of all classes of the given type. $\square_{5.17}$

Since there are a proper class of sentences in $L_{\infty,\omega}$, there are a proper class of aec with a given similarity type. So we have to modify this notion slightly to deal with AEC.

Notation 5.18. *For any aec $\boldsymbol{K}$, let $\kappa_{\boldsymbol{K}} = \sup(|\tau_{\boldsymbol{K}}|, \mathrm{LS}(\boldsymbol{K}))$.*

Now, for any cardinal κ, there are only a set AEC $\boldsymbol{K}$ with $\kappa_{\boldsymbol{K}} = \kappa$. As, the AEC is determined by its restriction to models below the Lówenheim number (just close the small models under union). (Compare Exercise 17.13.) And there are only a set of such restrictions. So we can apply Hanf's argument if we replace $|\tau|$ by $\kappa_{\boldsymbol{K}}$. And we can strengthen the result if we prove a crucial property.

Definition 5.19. *P is downward closed if there is a κ_0 such that if $P(\boldsymbol{K}, \lambda)$ holds with $\lambda > \kappa_0$, then $P(\boldsymbol{K}, \mu)$ holds if $\kappa_0 < \mu \leq \lambda$.*

The following is obvious.

Theorem 5.20. *If a property P is downward closed then for any κ there is a cardinal μ such for any AEC $\boldsymbol{K}$ with $\kappa_{\boldsymbol{K}} = \kappa$, if some model in $\boldsymbol{K}$ with cardinality $> \mu$ has property P, the property P holds in all cardinals greater than μ.*

That is, if AEC are downward closed for a property P there is a Hanf Number for P in the following stronger sense.

Definition 5.21. *[Hanf Numbers] The* Hanf number *for P, among AEC $\boldsymbol{K}$ with $\kappa_{\boldsymbol{K}} = \kappa$, is μ if:*
there is a model in $\boldsymbol{K}$ with cardinality $> \mu$ that has property P, implies the property P holds in all cardinals greater than μ.

In this sense it is obvious that there is a Hanf Number for existence using the Löwenheim Skolem downward property. We now give an explicit calculation of that number. The presentation theorem is a descendent of results [33, 34, 111] linking the Hanf number for various infinitary logics with the Hanf number for omitting types.

We will speak of several notions of type. The syntactic types of first order logic are fundamental and are used here. There is a natural syntactic notion of type with respect to any logic and the semantic notion of Galois type plays a fundamental role in Part III.

Definition 5.22. *For any logic $\mathcal{L}$, a* syntactic partial $\mathcal{L}$-type *p is a set of $\mathcal{L}$ consistent set of first order formulas in a fixed finite number of variables realized in a model. p is a complete type over a set A if every $\mathcal{L}$-formula with parameters from A (or its negation) is in p. If necessary we write $\mathcal{L}(\tau)$ to specify the vocabulary under consideration.*

The syntactic types of first order logic are fundamental and are used in this chapter.

Notation 5.23. *[Hanf Function]*

1. *A (first order syntactic) type p is a consistent set of first order formulas in a fixed finite number of variables. p is a* complete *type over a set A if every formula with parameters from A (or its negation) is in p. We sometimes say partial for emphasis.*

2. *Let $\eta(\lambda, \kappa)$ be the least cardinal μ such that if a first order theory T with $|T| = \lambda$ has models of every cardinal less than μ which omit each of a set Γ of types, with $|\Gamma| = \kappa$, then there are arbitrarily large models of T which omit Γ. We call $\eta(\lambda, \kappa)$ the* Hanf function *for omitting κ types for theories of size λ.*

3. *Write $\eta(\kappa)$ for $\eta(\kappa, \kappa)$.*

4. *We write $H(\kappa)$ for $\beth_{(2^\kappa)^+}$.*

5. *For a similarity type τ, $H(\tau)$ means $H(|\tau|)$. With a fixed $\boldsymbol{K}$, we write H_1 for $H(\kappa_{\boldsymbol{K}}) = H(\sup(\tau_{\boldsymbol{K}}, \mathrm{LS}(\boldsymbol{K})))$.*

Morley's omitting types theorem, Theorem A.3 of the appendix [111][33], is a crucial tool for this book. It computes the Hanf function for omitting types. Dealing with infinitary logics smears the clear distinction in the first order case between results for countable and uncountable similarity types. We use a strong version of the computation of Hanf numbers from Shelah's book, [122], VII.5.4, VII.5.5. It is an immediate corollary of Theorem A.3.

Corollary 5.24. $\eta(\kappa, 2^\kappa) \leq H(\kappa) = \beth_{(2^\kappa)^+}$.

Corollary 5.25. *If $\boldsymbol{K}$ is an AECand $\boldsymbol{K}$ has a model of cardinality at least $H_1 = H(\kappa_{\boldsymbol{K}})$ then $\boldsymbol{K}$ has arbitrarily large models.*

When $\kappa_{\boldsymbol{K}} = \max(\mathrm{LS}(\boldsymbol{K}), |\tau(\boldsymbol{K})|)$, $H(\kappa_{\boldsymbol{K}}) = H_1$ is sometimes called the Hanf number of $\boldsymbol{K}$. This is somewhat misleading because a single class can not have a Hanf number – a Hanf number is a maximum for all sentences and for all similarity types of a given cardinality (and in the case of AEC a fixed Löwenheim-Skolem number). As elaborated above, it is not the Hanf number of $\boldsymbol{K}$ but the Hanf number for all AEC with the same $\kappa_{\boldsymbol{K}}$. Crucially, the classes (which are not AEC) of models we consider next have the same property: for any model M with $|M| \geq \eta(\tau)$, (again $|\tau| = \mathrm{LS}(\boldsymbol{K})$) there are models of all cardinalities in the class that omit all types omitted in M.

We have represented $\boldsymbol{K}$ as a $PC\Gamma$ class in the following sense. Recall that a PC (*pseudoelementary*) class is the collection of reducts to a vocabulary τ of models of a theory T' in an expanded vocabulary τ'.

Definition 5.26. *Let Γ be a collection of first order types in finitely many variables over the empty set in a vocabulary τ'. A $PC(T, \Gamma)$* class *is the class of reducts to $\tau \subset \tau'$ of models of a first order theory τ'-theory which omit all members of the specified collection Γ of partial types.*

We write $PC\Gamma$ to denote such a class without specifying either T or Γ. And we write $\boldsymbol{K}$ is $PC(\lambda, \mu)$ if $\boldsymbol{K}$ can be presented as $PC(T, \Gamma)$ with $|T| \leq \lambda$ and $|\Gamma| \leq \mu$. (We sometimes write $PC\Gamma(\lambda, \mu)$ to emphasize the type omission. In the simplest case, we say $\boldsymbol{K}$ is λ-presented if $\boldsymbol{K}$ is $PC(\lambda, \lambda)$.)

In this language the Presentation Theorem 5.14 asserts any AEC $\boldsymbol{K}$ is $2^{\mathrm{LS(K)}}$-representable.

Keisler [80] proves a number of strong results for some special cases of $PC\Gamma$-classes. He calls a class of models a PC_δ-class if it is the class of reducts of a countable first order theory. Keisler has a different name for what Shelah calls $PC(\aleph_0, \aleph_0)$ and I call $PC\Gamma(\aleph_0, \aleph_0)$.

Definition 5.27. *$\boldsymbol{K}$ is a PC_δ class over $L_{\omega_1,\omega}$ if $\boldsymbol{K}$ is the class of reducts to $\tau(\boldsymbol{K})$ of the class of models of a sentence of $L_{\omega_1,\omega}$ in some expansion τ' of τ.*

Further he writes PC_δ over $L_{\omega_1,\omega}$ if it is the class of reducts of a countable set of $L_{\omega_1,\omega}$-sentences in a countable vocabulary; thus this notion corresponds to our $PC(\aleph_0, \aleph_0)$. In particular, he proves a categoricity transfer theorem between cardinals κ and λ of certain specific forms. (See Theorem 24 of [80] and Remark 15.10). The following example of Silver highlights the weakness of $PC\Gamma$-classes and the need to study AEC.

Example 5.28. *Let $\boldsymbol{K}$ be class of all structures (A, U) such that $|A| \leq 2^{|U|}$. Then $\boldsymbol{K}$ is actually a PC-class. But $\boldsymbol{K}$ is κ-categorical if and only if $\kappa = \beth_\alpha$ for a limit ordinal α. (i.e. $\mu < \kappa$ implies $2^\mu < \kappa$.) Thus there are PC-classes for which both the categoricity spectrum and its complement are cofinal in the class of all cardinals.*

One of the main problems in the study of AEC is Shelah's conjecture that there is a Hanf number for categoricity, a κ such that categoricity above κ implies categoricity in all larger cardinals. Silver's example shows such a conjecture fails for PC-classes and so certainly for PCΓ classes. Thus, AEC are a more reasonable candidate for a framework in which to study categoricity.

Exercise 5.29. *Show that the class $\boldsymbol{K}$ of Example 5.28 with $\prec_{\boldsymbol{K}}$ as elementary submodel satisfies all axioms for an AEC except unions of chains.*

We will see many problems concerning the spectrum of classes defined in $L_{\omega_1,\omega}$ can be reduced to classes of structures of the following sort. We specialize the notion of $PC\Gamma$ in two ways; we work in the original language instead of allowing reducts; the omitted types are complete types in finitely many variables over the empty set.

Definition 5.30. *1. A* finite diagram *or* $EC(T, \Gamma)$-class *is the class of models of a first order theory T which omit all types from a specified collection Γ of complete types in finitely many variables over the empty set.*

2. $EC(T, Atomic)$ denotes the class of atomic *models of T.*

Definition 5.30.2 abuses the $EC(T, \Gamma)$ notation, since for consistency, we really should write non-atomic. But atomic is shorter and emphasizes that we are restricting to the atomic models of T.

Exercise 5.31. *The models of an $EC(T, \Gamma)$ with the ordinary first order notion of elementary submodel as $\prec_{\boldsymbol{K}}$ gives an AEC with Löwenheim-number $|T|$.*

Some authors attach the requirement that $\boldsymbol{K}$ satisfy amalgamation over sets to the definition of finite diagram. We stick with the original definition from [119] and reserve the more common term, *homogeneous model theory* for the classes with set amalgamation.

Definition 5.32. *We say an AEC admits amalgamation over arbitrary sets if for any pair of embeddings f, g from a set A into $M, N \in \boldsymbol{K}$, there exist an $M' \in \boldsymbol{K}$ and* strong *embeddings f_1, g_1 from M, N into M' so that $f_1 f = g_1 g$.*

Definition 5.33. *If $\boldsymbol{K}$ is an abstract elementary class which admits amalgamation over arbitrary subsets of models, the study of all submodels of a monster model is called:* homogeneous model theory.

If T is a first order theory which admits elimination of quantifiers, then for any Γ, $EC(T,\Gamma)$ will have sequentially homogeneous universal domains just if it admits amalgamation over arbitrary sets.

There is extensive study of homogeneous model theory by Buechler, Grossberg, Hyttinnen, Lessmann, Shelah and others developing many analogs of the first order theory including the stability hierarchy, DOP, simplicity etc. ([67, 31, 51, 97, 68]. We briefly discuss two important examples.

Example 5.34. Banach spaces do *not* form an AEC because the union of complete spaces need not be complete. Indeed, the union of a countable chain of Banach spaces is not a Banach space, but its completion is a Banach space. From the perspective of Banach space model theory [57, 74, 72, 73, 29, 136, 58], there is no distinction between a Banach space and a dense subset of it. The theory of a Banach space X is determined uniquely by the theory of any dense subset of X. (This is a consequence of the fact that in Banach space model theory we deal with approximations.)

Thus, Banach space model theory can be thought of as the study not of Banach spaces, but of structures whose completion is a Banach Space. Alternatively, one can say the class of subspaces of Banach spaces forms an AEC.

Example 5.35 (Robinson Theories)**.** One important example of homogeneous model theory is the study of Robinson theories. This notion was developed first in [121] but rediscovered and baptized in the unpublished [62]. We follow the terminology of [115]. Let Δ be a family of (first order) formulas which is closed under conjunction, disjunction, negation, and subformula. A Robinson theory T is a theory with Δ-amalgamation that is axiomatized by the universal closures of Δ-formulas.

Abusing normal language, a Robinson theory is said to be *stable* or *simple* if appropriate translations of stable and simple are true for Δ-formulas on the *existentially closed* models of T. (Usually, this is phrased in terms of behavior on a universal domain.) In the interesting cases, the class of existentially closed models is not first order so this is an inherently infinitary setting.

Example 5.36 (CATS)**.** Robinson theories were further generalized by Pillay [114] and then by Ben Yaacov [27, 28, 135]. T is said to be positive Robinson if the requirement that Δ be closed under negation is dropped and for $\boldsymbol{a} \in M$, $\mathbf{b} \in N$, (M, N Δ-existentially closed model of T): if $\mathrm{tp}^M_\Delta(\boldsymbol{a}) \subseteq \mathrm{tp}^N_\Delta(\mathbf{b})$, $\boldsymbol{a}$ and $\mathbf{b}$ have the same Δ-existential type. Again simple and stable are defined in terms of the behavior of Δ-formulas on a universal domain. The existentially closed models of a positive Robinson theory form an AEC.

The universal domain of a cat is just a "monster of kind III" in [121] (i.e. it is a particular case of a big (D,λ) homogeneous model for good D as in [119] with the extra assumption of compactness for positive existential formulas. In this sense the universal domain of a cat is certainly a particular case of a "monster" for an abstract elementary class. Thus, there is no chance of interpreting an arbitrary AEC in a cat but in a weak sense one can regard a CAT as an AEC.

However, if one wants to recover a class of models satisfying Shelah's axioms of an AEC from the universal domain, the situation is different. Just as with Banach spaces, the union axiom must be adjusted by taking the completion of a union. To recover the class of the models and be ”faithful” to what people studying cats actually care about, would in certain cases (for example, Hausdorff cat, or any metric cat as in [137]) yield, e.g. an abstract metric class, not an AEC.

Thanks to Olivier Lessmann, Jose Iovino, and Alex Usvyatsov for outlining the issues above.

Example 5.37 (Metric AEC)**.** Hirvonen and Hyttinen [63] and Shelah and Usvyatsov [135] discuss the notion of a *metric abstract elementary class* . This is a study of many sorted structures with a complete metric on each sort. This provides a common generalization of AEC and CATS. As in AEC, there is no compactness hypothesis. [63] provide both a clear summary of the connections between the various context and develop the theory of metric AEC to prove categoricity transfer for *homogeneous* metric AEC. Kirby [84] provides a categorical treatment unifying the two cases. This axiomatization nicely encapsulates the notion of an AEC as the closure under direct limits of a set of a structures. (Compare Definition 17.10 and following exercises.)

Example 5.38. Still another approach to dealing with the problem of completing a metric space is provided by [18, 19]. Here the approach is to axiomatize the notion of prime model over a union.

6
Two basic results about $L_{\omega_1,\omega}(Q)$

This chapter has two parts. In the first section we extend the Lopez-Escobar/Morley theorem [103, 111] on the nondefinability of well-order from $L_{\omega_1,\omega}$ to $L_{\omega_1,\omega}(Q)$. This extension seems to be well-known to cognoscenti but I was unable to find it in the literature[1]. We begin the exploration of ways to fit $L_{\omega_1,\omega}(Q)$ into the AEC framework.

The second section contains the proof of Keisler's theorem if a sentence of $L_{\omega_1,\omega}$ with less than $2^{\aleph_1}$ models in $\aleph_1$ then for any countable fragment $L_{\mathcal{A}}$, then *every* member of $\boldsymbol{K}$ realizes only countably many $L_{\mathcal{A}}$-types over $\emptyset$. This is the key to showing (assuming $2^{\aleph_0} < 2^{\aleph_1}$ that an $\aleph_1$ categorical sentence is ω-stable.

6.1 Non-definability of Well-order in $L_{\omega_1,\omega}(Q)$

We begin with attempts to regard the models of an $L_{\omega_1,\omega}(Q)$ as an AEC. It is easy to see:

Exercise 6.1.1. *The models of a sentence of $L(Q)$ with the associated notion of elementary submodel as $\prec_{\boldsymbol{K}}$ does* not *give an AEC.*

So we want to consider some other notions of strong submodel.

Definition 6.1.2. *Let ψ be a sentence in $L_{\omega_1,\omega}(Q)$ in a countable vocabulary and let L^* be the smallest countable fragment of $L_{\omega_1,\omega}(Q)$ containing ψ. De-*

[1] In fact I later found Barwise[23] had given a quite different proof deducing the result on the well-ordering number of $L_{\omega_1,\omega}(Q)$ from the omitting types theorem for $L(Q)$.

fine a class $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ by letting $\boldsymbol{K}$ be the class of models of ψ in the standard interpretation. We consider several notions of strong submodel.

1. $M \prec^* N$ *if*
   (a) $M \prec_{L^*} N$ *and*
   (b) $M \models \neg(Qx)\theta(x, \boldsymbol{a})$ *then*
   $$\{b \in N : N \models \theta(b, \boldsymbol{a})\} = \{b \in M : N \models \theta(b, \boldsymbol{a})\}.$$
2. $M \prec^{**} N$ *if*
   (a) $M \prec_{L^*} N$,
   (b) $M \models \neg(Qx)\theta(x, \boldsymbol{a})$ *then*
   $$\{b \in N : N \models \theta(b, \boldsymbol{a})\} = \{b \in M : N \models \theta(b, \boldsymbol{a})\}$$
   and
   (c) $M \models (Qx)\theta(x, \boldsymbol{a})$ *implies* $\{b \in N : N \models \theta(b, \boldsymbol{a})\}$ *properly contains* $\{b \in M : N \models \theta(b, \boldsymbol{a})\}$.

The following exercises are easy but informative.

Exercise 6.1.3. $(\boldsymbol{K}, \prec^*)$ *is an AEC with Löwenheim Number* $\aleph_1$.

Exercise 6.1.4. $(\boldsymbol{K}, \prec^{**})$ *is* not *an AEC. (Hint: Consider the second union axiom* **A3.3** *in Definition 5.1 and a model with a definable uncountable set.)*

Remark 6.1.5. The Löwenhheim number of the AEC $(\boldsymbol{K}, \prec^*)$ defined in Definition 6.1.2 is $\aleph_1$. We would like to translate an $L_{\omega_1,\omega}(Q)$-sentence to an AEC with Löwenhheim number $\aleph_0$ and which has at least approximately the same number of models in each uncountable cardinality. This isn't quite possible but certain steps can be taken in that direction. This translation will require several steps. We begin here with a fundamental result about $L_{\omega_1,\omega}(Q)$; in Chapter 7, we will complete the translation.

Here are the background results in $L_{\omega_1,\omega}$. They are proved as Theorem 12 and Theorem 28 from [80]. In applications, we may add the linear order to the vocabulary to discuss structures which do not admit a definable order in the basic vocabulary.

Theorem 6.1.6 (Lopez-Escobar, Morley)**.** *Let ψ be an $L_{\omega_1,\omega}(\tau)$-sentence and suppose $P, <$ are a unary and a binary relation in τ. Suppose that for each $\alpha < \omega_1$, there is a model M_α of ψ such that $<$ linear orders $P(M_\alpha)$ and α imbeds into $(P(M_\alpha), <)$. Then there is a (countable) model M of ψ such that $(P(M), <)$ contains a copy of the rationals.*

If N is linearly ordered, N is an *end extension* of M if every element of M comes before every element of $N - M$. We give some further explanation of the following result in Appendix B

Theorem 6.1.7 (Keisler). *Let L^* be a countable fragment of $L_{\omega_1,\omega}$. If a countable linearly-ordered model M has a proper L^*-elementary end extension, then it has one with cardinality $\aleph_1$.*

These two results can be combined to show that if a sentence in $L_{\omega_1,\omega}$ has a model that linearly orders a set in order type ω_1 then it has a model of cardinality $\aleph_1$ where the order is not well-founded. We imbed that argument in proving the same result for $L_{\omega_1,\omega}(Q)$.

Theorem 6.1.8. *Let τ be a similarity type which includes a binary relation symbol $<$. Suppose ψ is a sentence of $L_{\omega_1,\omega}(Q)$, $M \models \psi$, and the order type of $(M,<)$ is ω_1. There is a model N of ψ with cardinality $\aleph_1$ such that the order type of $(N,<)$ imbeds $\mathbb{Q}$.*

Proof. Extend the vocabulary τ to τ' by adding a function symbol $f_\phi(\mathbf{x}, y)$ for each formula $(Qy)\phi(y,\mathbf{x})$ in $L_{\omega_1,\omega}(Q)$ and a constant symbol for ω. Expand M to a τ'-structure M' by interpreting f_ϕ as follows:

1. If $M \models (Qy)\phi(y,\boldsymbol{a})$, $(\lambda y)f_\phi(y,\boldsymbol{a})$ is a partial function with domain the solution set of $\phi(y,\boldsymbol{a})$ *onto* M.

2. If $M \models \neg(Qy)\phi(y,\boldsymbol{a})$, $(\lambda y)f_\phi(y,\boldsymbol{a})$ is a partial function with domain the solution set of $\phi(y,\boldsymbol{a})$ into the imbedded copy of ω.

Now let L^* be a countable fragment in τ' of $L_{\omega_1,\omega}$ which contains every subformula of ψ which is in $L_{\omega_1,\omega}$ and a sentence θ expressing the properties of the Skolem functions for $L_{\omega_1,\omega}(Q)$ that we have just defined. Let ψ^* be the conjunction of θ with an L^*-sentence which asserts that 'ω is standard' and a translation of ψ obtained by replacing each subformula of ψ of the form $(Qy)\phi(y,\mathbf{z})$ by the formula $f_\phi(y,\mathbf{z})$ is onto and each formula of the form $(\neg(Qy)\phi(y,\mathbf{z}))$ by $(f_\phi(y,\mathbf{z})$ maps into ω. Then for any τ'-structure N of cardinality $\aleph_1$ which satisfies ψ^*, $N \restriction \tau$ is a model of ψ. Now expand τ' to τ'' by adding a new unary predicate P. Let the sentence χ assert M is an elementary *end* extension of $P(M)$. For every $\alpha < \omega_1$ there is a model M_α of $\psi^* \wedge \chi$ with order type of $(P(M),<)$ greater than α. (Start with P as α and alternately take an L^*-elementary submodel and close down under $<$. After ω steps we have the P for M_α.) Now by Theorem 6.1.6 there is a countable structure $(N_0, P(N_0))$ such that $P(N_0)$ contains a copy of $(Q,<)$ and N_0 is an elementary end extension of $P(N_0)$. By Theorem 6.1.7, N_0 has an L^*-elementary extension N of cardinality $\aleph_1$. Clearly, $P(N)$ contains a copy of $(Q,<)$ and, as observed, $N \models \psi$. $\square_{6.1.8}$

It is easy to modify the proof to obtain the conclusion by weakening the hypothesis from one model with order type $\aleph_1$ to a family of models that define arbitrarily long countable well-orderings.

6.2 The number of models in ω_1

The next theorem is the key tool for the main result announced in the introduction. It is proved as Theorem B.6.

Theorem 6.2.1. *Fix a countable fragment $L_{\mathcal{A}}$ of $L_{\omega_1,\omega}$, a theory T in $L_{\mathcal{A}}$ such that $<$ is a linear order of each model of T. For each $p(\mathbf{x})$ an $L_{\mathcal{A}}$-type (possibly incomplete) over the empty set, there is a sentence $\theta_p \in L_{\omega_1,\omega}$ satisfying the following conditions.*

1. *If p is omitted in an uncountable model $(B,<)$ of T then for any countable $(A,<)$ such that $(B,<)$ is an end $L_{\mathcal{A}}$-elementary extension of $(A,<)$, $(A,<) \models \theta_p$.*
2. *θ_p satisfies:*
   (a) *If $B \models \theta_p$ then B omits p.*
   (b) *θ_p is preserved under unions of chains of $L_{\mathcal{A}}$-elementary end extensions;*
   (c) *for any family X of $L_{\mathcal{A}}$-types $\langle p_m : m < \omega\rangle$ over $\emptyset$ and any countable A, if $A \models \theta_{p_m}$ for each M then A has a proper $L_{\mathcal{A}}$-elementary extension that satisfies each θ_{p_m}.*
3. *Let X be a collection of complete $L_{\mathcal{A}'}(\tau')$-types (for some $\mathcal{A}' \subseteq \mathcal{A}$ and $\tau' \subseteq \tau$) over the empty set that are realized in every uncountable model of T. Then, X is countable.*

Note the following immediate corollary.

Corollary 6.2.2. *Let $(A,<)$ be countable and suppose for each $m < \omega$ $A \models \theta_{p_m}$. Then there is an uncountable end extension B of A omitting all the p_m.*

Proof. By condition Theorem 6.2.1 2c) there is a proper elementary end extension of A_1 of A satisfying all the θ_{p_m}. Iterate this construction through ω_1 using Theorem 6.2.1 2b) at limit stages. By Theorem 6.2.1 2a) the limit model omits all the p_m. $\square_{6.2.2}$

The following example (Baldwin/Marker) shows the significance of *end* extension in the statement above.

Remark 6.2.3. The fragment is first order logic. In the base model M we have points of two colors–say red and blue and the red points and blue points each form a copy of (Z, s). Let p be the type which says there are no new red points and q the type which says no new blue points. Of course it is true that M has elementary extensions of cardinality $\aleph_1$ omitting p and extensions of size $\aleph_1$ omitting q but none omitting both.

But–this observation does not take into account the *end* extension. We also have some linear order $>$ of A. One of the following holds: a) there exists x all $y > x$ are the same color b) for all x there are $y, z > x$ of different colors.

If a) holds then only one of p or q can be omitted in an elementary end extension.

If b) holds then neither p nor q can be omitted in an elementary end extension.

The remaining results do not assume there is a linear ordering in the language; we will add one in order to apply Theorem B.6. That is why the generality of PC_δ class over $L_{\omega_1,\omega}$-classes (our $PC\Gamma(\aleph_0, \aleph_0)$ from Definition 5.27) is necessary. The notion of an $L_{\mathcal{A}}$-type is defined in Definition 5.22.

Theorem 6.2.4. *If a PC_δ over $L_{\omega_1,\omega}$ class $\boldsymbol{K}$ has an uncountable model then for any countable fragment $L_{\mathcal{A}}$, there are only countably many $L_{\mathcal{A}}$-types over $\emptyset$ realized in* every *uncountable member of $\boldsymbol{K}$.*

Proof. Let ϕ be a τ'-sentence of $L_{\omega_1,\omega}$ such that $\boldsymbol{K}$ is the class of τ reducts of models of ϕ. Let $L_{\mathcal{A}}(\tau')$ be the smallest fragment that contains ϕ. Let X be the collection of $L_{\mathcal{A}}(\tau)$-types over $\emptyset$ realized in *every* uncountable model of ϕ.

Well-order an uncountable model of ϕ in order type ω_1 to get $(B, <)$. Let T' be the $L_{\mathcal{A}}(\tau')$-theory of $(B, <)$. We can construct $(A, <)$ such that $(B, <)$ is an uncountable $L_{\mathcal{A}}$-end extension of $(A, <)$. If p is realized in every uncountable model in $\boldsymbol{K}$ then p is realized in every uncountable model of T'. So applying Theorem B.6.3 to T' and X we have the result. $\square_{6.2.4}$

Theorem 6.2.5 (Keisler)**.** *If a PC_δ over $L_{\omega_1,\omega}$ class $\boldsymbol{K}$ has an uncountable model but less than 2^{ω_1} models of power $\aleph_1$ then for any countable fragment $L_{\mathcal{A}}$, then* every *member of $\boldsymbol{K}$ realizes only countably many $L_{\mathcal{A}}$-types over $\emptyset$.*

Proof. Let $T' = \{\phi\}$ be a τ'-sentence of $L_{\omega_1,\omega}$ such that $\boldsymbol{K}$ is the class of τ reducts of models of ϕ. Let $L_{\mathcal{A}}(\tau')$ be the smallest fragment that contains ϕ.

If the conclusion fails for some natural number p there are uncountably many $L_{\mathcal{A}}(\tau)$-types over the empty set realized in some model $B \in \boldsymbol{K}$; wolog $|B| = \aleph_1$. First note that we can expand the language with a unary predicate and functions so that there is a set U of p-tuples that realize distinct p-types and U has the same cardinality as the universe. This can be expressed by a sentence of $L_{\omega_1,\omega}$, so we have a PC_δ over $L_{\omega_1,\omega}$-class $\boldsymbol{K}''$ such that *every* uncountable model realizes uncountably many types. We will show $\boldsymbol{K}''$ has $2^{\aleph_1}$ models of cardinality $\aleph_1$. Indeed their reducts to τ are pairwise not mutually embeddible.

Suppose that $\boldsymbol{K}''$ is axiomatized in the fragment $L_{\mathcal{A}''}$ and let B'' be an uncountable model. Now fix $A'' = A_\emptyset$ as a countable submodel so that B'' is an $L_{\mathcal{A}''}$-end extension of A'' and $p_\emptyset$ as any $L_{\mathcal{A}}$-p-type over $\emptyset$ realized in A''. We construct a family of countable τ''-models A_s for $s \in 2^{<\omega_1}$ and $L_{\mathcal{A}}$-types p_s over the empty set to satisfy the following conditions:

1. if $s < t$ then A_t is an $L_{\mathcal{A}''}$-end extension of A_s;
2. if $s \leq t$ then A_t realizes p_s;
3. if $s < t$ and $s\frown i \not\leq t$ (for $i \in \{0, 1\}$) then $A_t \models \theta_{p_{s\frown i}}$ and so omits $p_{s\frown i}$.

Then if $\sigma \in 2^{\omega_1}$ and $s \in 2^{<\omega_1}$, $M_\sigma = \bigcup_{s\subset\sigma} M_s$ realizes p_s iff $s < \sigma$. This clearly suffices as $\sigma \neq \tau \in 2^{\omega_1}$ implies the τ-reduct M_σ cannot be embedded in the τ-reduct M_τ.

Now for the construction. For the limit stage we need to know that if we have an increasing chain M_i such that for $i_0 < j < \alpha$, $M_j \models \theta_{p_{i_0}}$ then so does M_α. This is immediate from Theorem B.6.2b.

Now for the successor stage. We have an A_s satisfying the conditions. That is, A_s realizes p_t if $t \leq s$ and $A_s \models \theta_{p_{t\frown i}}$ if $t < s$ and $t\frown i \not\leq s$. Let $\boldsymbol{K}^3$ be the class of all τ-reducts of $L_{\mathcal{A}''}$-end extensions of A_s that omit $p_{t\frown i}$ if $t < s$ and $t\frown i \not\leq s$. Corollary 6.2.2 gives us an uncountable $L_{\mathcal{A}''}$-end extension B_s of A_s in $\boldsymbol{K}^3$. By Theorem 6.2.4 only countably many $L_{\mathcal{A}}$ types over $\emptyset$ are realized in all models in $\boldsymbol{K}'''$. So we can choose $p_{s\frown 0}$ that is realized in B_s but omitted in some uncountable $L_{\mathcal{A}''}$-end extension of A_s, $B_1 \in \boldsymbol{K}^3$. Choose $A_{s\frown 0}$ as a countable $L_{\mathcal{A}''}$-end extension of A_s that realizes $p_{s\frown 0}$ and with

$$A_s \prec_{L_{\mathcal{A}''}} A_{s\frown 0} \prec_{L_{\mathcal{A}''}} B_s.$$

By Theorem B.6.1, $A_{s\frown 0} \models \theta_{p_t}$ for $t < s$ and $t\frown i \not\leq s$.

To choose $p_{s\frown 1}$ and $A_{s\frown 1}$, we now apply Theorem 6.2.4 to $\boldsymbol{K}^4$ obtained by requiring in addition to $\boldsymbol{K}^3$ that $p_{s\frown 0}$ is omitted. We know B_1 is one $L_{\mathcal{A}''}$-end extension of A_s that is in $\boldsymbol{K}^4$. Since B_1 realizes $\aleph_1$-types there must be a type $p_{s\frown 1}$ realized in B_1 and omitted in some uncountable $L_{\mathcal{A}''}$-end extension of $A_{s\frown 0}$; thus $A_{s\frown 0} \models \theta_{p_{s\frown 1}}$. Let $A_{s\frown 1}$ be a countable $L_{\mathcal{A}''}$-end extension of A_s with B_1 an $L_{\mathcal{A}''}$-end extension of $A_{s\frown 1}$ so that $A_{s\frown 1}$ realizes $p_{s\frown 1}$; by Theorem B.6.1 $A_{s\frown 1}$ satisfies condition 3). This completes the construction.

Remark 6.2.6. *We have chosen to proven Theorem 6.2.5 only for $L_{\omega_1,\omega}$ but Keisler (Corollary 5.10 of [79]) proves the result for $L_{\omega_1,\omega}(Q)$ and Kaufmann [78] asserts that result extends to $L_{\omega_1,\omega}(aa)$.*

7
Categoricity implies Completeness

We defined the logics $L_{\omega_1,\omega}$ and $L_{\omega_1,\omega}(Q)$ in Chapter 2.2. We now give full arguments for understanding the relationship between categoricity and completeness in $L_{\omega_1,\omega}$ and the resulting alternative descriptions of categorical classes in $L_{\omega_1,\omega}$. We introduce the role of complete sentences in $L_{\omega_1,\omega}$ and $L_{\omega_1,\omega}(Q)$ in Section 1. In Sections 2 and 3, we explain how to obtain complete sentences in $L_{\omega_1,\omega}$ from the hypothesis of arbitrarily large models and few models in $\aleph_1$, respectively. In those three sections we consider certain arguments that work for both $L_{\omega_1,\omega}$ and $L_{\omega_1,\omega}(Q)$ but restrict to $L_{\omega_1,\omega}$ when necessary. We discuss finishing the arguments for $L_{\omega_1,\omega}(Q)$ in Section 4.

This chapter is a prelude for Part IV; the following example of David Kueker shows dealing with $L_{\omega_1,\omega}$ is a real restriction. There are categorical AEC that are not closed under $L_{\omega_1,\omega}$-equivalence (and so certainly are not $L_{\omega_1,\omega}$-axiomatizable).

Example 7.1. There is an AEC $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ satisfying AP and JEP, with $\mathrm{LS}(\boldsymbol{K}) = \omega$, that is κ-categorical for all infinite κ, but which is not closed under $L_{\infty,\omega}$-elementary equivalence. The vocabulary L consists of a single unary predicate P. $\boldsymbol{K}$ is the class of all L-structures M such that $|P^M| = \omega$ and $|(\neg P)^M|$ is infinite. Define $M \prec_{\boldsymbol{K}} N$ if and only if $P^M = P^N$ and $M \subseteq N$. Clearly $\boldsymbol{K}$ is categorical in all infinite powers. Since $\boldsymbol{K}$-extensions can't expand P, it is possible to keep the cardinality of P fixed while still satisfying the unions of chains axioms.

7.1 Completeness

We have a natural definition of completeness for $L_{\omega_1,\omega}$ (and analogously for $L_{\omega_1,\omega}(Q)$).

Definition 7.1.1. *A sentence ψ in $L_{\omega_1,\omega}$ is called* complete *if for every sentence ϕ in $L_{\omega_1,\omega}$, either $\psi \models \phi$ or $\psi \models \neg\phi$.*

In first order logic, the theory of a structure is a well-behaved object; here such a theory is not so nice. An infinite conjunction of first order sentences behaves very much like a single sentence; in particular it satisfies both the upward and downward Löwenheim Skolem theorems. In contrast, the conjunction of all $L_{\omega_1,\omega}$-sentences true in an uncountable model may not have a countable model. In its strongest form Morley's theorem asserts: Let T be a first order theory having only infinite models. If T is categorical in some uncountable cardinal then T is complete and categorical in every uncountable cardinal. This strong form does not generalize to $L_{\omega_1,\omega}$; take the disjunction of a sentence which is categorical in all cardinalities with one that has models only up to, say, $\beth_2$. Since $L_{\omega_1,\omega}$ fails the upwards Löwenheim-Skolem theorem, the categoricity implies completeness argument, which holds for first order sentences, fails in this context. However, if the $L_{\omega_1,\omega}$-sentence ψ is categorical in κ, then, applying the downwards Löwenheim-Skolem theorem, for every sentence ϕ that is consistent with ψ either all models of $\psi \wedge \phi$ or all models of $\psi \wedge \neg\phi$ have cardinality less than κ. So if ϕ and ψ are κ-categorical sentences with a common model of power κ they are equivalent on models of cardinality $\geq \kappa$.

This is a real obstacle to downward categoricity transfer. Let $\alpha < \omega_1$. If ϕ is a sentence categorical in all powers and ϕ_α is a sentence of $L_{\omega_1,\omega}$ that has no model of cardinality greater than $\beth_\alpha$ (see Fact 2.5), $\phi \vee \psi_\alpha$ is categorical in all cardinals above $\beth_\alpha$ but not (in general) below. Thus, any conjecture concerning downward categoricity transfer will have to add some feature (e.g. joint embedding in the appropriate category).

Exercise 7.1.2. *Suppose $\phi \in L_{\omega_1,\omega}$ is categorical in $\kappa \geq \beth_{\omega_1}$. Use the fact that the Hanf number of $L_{\omega_1,\omega}$ is $\beth_{\omega_1}$ (Fact 2.5) to show all models of ϕ of cardinality greater than $\beth_{\omega_1}$ satisfy the same sentences of $L_{\omega_1,\omega}$.*

The fundamental tool for discussing complete sentences gives them their name: Scott sentences. We reproduce a proof (which Keisler attributes to Chang) of the Scott isomorphism theorem from the first chapter of Keisler's book [80]. The first step of the proof is extended to $L_{\omega_1,\omega}(Q)$ in Lemma 7.1.5. We present a weaker version (due to Kueker) of the main Scott theorem for $L_{\omega_1,\omega}(Q)$ that is adequate for our purpose in Theorem 7.4.2. Of course, a direct translation of Scott's theorem to $L_{\omega_1,\omega}(Q)$ makes no sense as there are no countable models.

We will show that under certain conditions any categorical sentence is implied by a complete sentence. The following example, based on one given by David

Marker, shows this *replacement* is significant. The straightforward translation of the Baldwin-Lachlan theorem to $L_{\omega_1,\omega}$ is false.

Example 7.1.3. There is a sentence ϕ of $L_{\omega_1,\omega}$ that is $\aleph_1$ categorical and satisfies the amalgamation property but has $2^{\aleph_0}$ countable models and does not satisfy the joint embedding property. Let the vocabulary contain a binary relation σ, a constant symbol 0, and a binary symbol S. We require that if $M \models \phi$, $(M, S, 0)$ is a model of the first order theory of $(\mathcal{Z}, S)$. Further, let ϕ assert that if P is non-empty every element is a finite predecessor or successor of 0.

We need some definitions. By L^* type here, as in Definition 5.22, we mean 'maximal (over given domain) satisfiable set of L^*-formulas'.

Definition 7.1.4. *1. A τ-structure M is L^*-small for L^* a countable fragment of $L_{\omega_1,\omega}(Q)(\tau)$ if M realizes only countably many L^*-types.*

2. A τ-structure M is $L_{\omega_1,\omega}(Q)$-small if M realizes only countably many $L_{\omega_1,\omega}(Q)(\tau)$-types.

3. We have the analogous notions for $L_{\omega_1,\omega}$.

The following Lemma holds, *mutatis mutandis* for $L_{\omega_1,\omega}$.

Lemma 7.1.5. *If the structure M is $L_{\omega_1,\omega}(Q)$-small, then there is a countable fragment L^* of $L_{\omega_1,\omega}(Q)$ such that for every tuple $\boldsymbol{a}$ in M there is a formula $\phi_{\boldsymbol{a}}(\mathbf{x}) \in L^*$ such that $M \models \phi_{\boldsymbol{a}}(\boldsymbol{a})$ and $M \models \phi_{\boldsymbol{a}}(\mathbf{x}) \rightarrow \psi(\mathbf{x})$ for each $L_{\omega_1,\omega}(Q)$-formula true of $\boldsymbol{a}$.*

Proof. Choose a countable subset A of M that is 'relatively ω-saturated in the following sense. Every $L_{\omega_1,\omega}(Q)$-type over a finite sequence from A that is realized in M is realized in A. For each finite tuple $\boldsymbol{a} \in A$ and each countable ordinal β, we define a formula $\phi^{\alpha}_{\boldsymbol{a}}(\mathbf{x})$.

$\phi^0_{\boldsymbol{a}}(\mathbf{x})$ is the conjunction of basic formulas satisfied by $\boldsymbol{a}$

If β is a limit ordinal,

$$\phi^{\beta}_{\boldsymbol{a}}(\mathbf{x}) \text{ is } \bigwedge_{\gamma<\beta} \phi^{\gamma}_{\boldsymbol{a}}(\mathbf{x})$$

For successor ordinals we need some preliminary notation. For each pair $\boldsymbol{a}, a$ let $\psi^{\alpha}_{\boldsymbol{a},a}$ be that one of the two formulas $(Qx)\phi^{\alpha}_{\boldsymbol{a}a}(\boldsymbol{a}x)$, $\neg(Qx)\phi^{\alpha}_{\boldsymbol{a}a}(\boldsymbol{a}x)$ that is true in M. For each pair $\boldsymbol{a}, \phi^{\beta}_{\boldsymbol{a}}$, let

$$A_{\boldsymbol{a},\phi^{\beta}_{\boldsymbol{a}}} = \{a \in A : M \models (\exists x)\phi^{\beta}_{\boldsymbol{a},a}(\boldsymbol{a}x)\}.$$

(This notation is finer than needed here; A would suffice. But we use $A_{\boldsymbol{a},\phi^{\beta}_{\boldsymbol{a}}}$ in proving Theorem 7.4.2.) Now, $\phi^{\beta+1}_{\boldsymbol{a}}(\mathbf{x})$ is:

$$\phi^{\beta}_{\boldsymbol{a}}(\mathbf{x}) \wedge \bigwedge_{a\in A} \psi^{\beta}_{\boldsymbol{a},a} \wedge \bigwedge_{a\in A_{\boldsymbol{a},\phi^{\beta}_{\boldsymbol{a}}}} (\exists x)\phi^{\beta}_{\boldsymbol{a},a}(\mathbf{x}, x) \wedge (\forall x) \bigvee_{a\in A_{\boldsymbol{a},\phi^{\beta}_{\boldsymbol{a}}}} \phi^{\beta}_{\boldsymbol{a},a}(\mathbf{x}, x).$$

Note that for all $\boldsymbol{a} \in A$ and all β, $M \models \phi^{\beta}_{\boldsymbol{a}}(\boldsymbol{a})$. If $\alpha < \beta$,

$$M \models (\forall \mathbf{x})[\phi^{\beta}_{\boldsymbol{a}}(\mathbf{x}) \rightarrow \phi^{\alpha}_{\boldsymbol{a}}(\mathbf{x})].$$

Thus, since A is countable and since all types realized in M are also realized in A, for every $\boldsymbol{a}$ there is a countable α such that for all $\beta \geq \alpha$,

$$M \models (\forall \mathbf{x})[\phi^{\beta}_{\boldsymbol{a}}(\mathbf{x}) \leftrightarrow \phi^{\alpha}_{\boldsymbol{a}}(\mathbf{x})].$$

Choosing $\phi_{\boldsymbol{a}}$ as this $\phi^{\alpha}_{\boldsymbol{a}}$ we have that $\phi_{\boldsymbol{a}}$ decides all formulas in M. That is, for every $\lg(\boldsymbol{a})$-ary formula χ either $M \models [\phi_{\boldsymbol{a}} \rightarrow \chi]$ or $M \models [\phi_{\boldsymbol{a}} \rightarrow \neg\chi]$; the answer depends on whether $\boldsymbol{a}$ satisfies χ.

Since A is countable, there is a countable α such for all n and all n-tuples $\boldsymbol{a}$, for all $\beta \geq \alpha$,

$$M \models (\forall \mathbf{x})\phi^{\alpha}_{\boldsymbol{a}} \leftrightarrow \phi^{\beta}_{\boldsymbol{a}}.$$

So we choose L^* to contain all the $\phi^{\alpha}_{\boldsymbol{a}}$ for this α. Note that although we have formulas $\phi^{\alpha}_{\boldsymbol{a}}$ only for $\boldsymbol{a} \in A$, the 'ω-saturation' of A guarantees that for each $\mathbf{b} \in M$, there is an $\boldsymbol{a}$ with $\phi_{\boldsymbol{a}}(\mathbf{b})$. $\square_{7.1.5}$

A Scott sentence for a countable model M is a complete sentence satisfied by M; it characterizes M up to isomorphism among countable models. The Scott sentence for an uncountable small model is the Scott sentence for a countable L^*-submodel of M, where L^* is the smallest fragment containing a formula for each type realized in M.

Lemma 7.1.6 (Scott's Isomorphism Theorem)**.** *Let M be a τ-structure for some countable τ that is small for $L_{\omega_1,\omega}$. There is an $L_{\omega_1,\omega}$-sentence ϕ_M such that $M \models \phi_M$ and all countable models of ϕ_M are isomorphic. This implies that ϕ_M is complete for $L_{\omega_1,\omega}$.*

Proof. Apply the notation from and let α be the bound from the proof of Lemma 7.1.5 (but omitting the reference to the Q-quantifier in the construction). Let ϕ_M be the sentence

$$\phi^{\alpha}_{\emptyset} \wedge \bigwedge_{n<\omega;\boldsymbol{a}\in A} (\forall \mathbf{x})[\phi^{\alpha}_{\boldsymbol{a}} \rightarrow \phi^{\alpha+1}_{\boldsymbol{a}}].$$

Let L^* be the fragment generated by ϕ_M and the $\phi^{\gamma}_{\boldsymbol{a}}$ for $\boldsymbol{a} \in A$ and $\gamma \leq \alpha+1$. Let M' be a countable L^*-elementary submodel of M.

Now we show by a back and forth argument that if N is a countable model of ϕ_M, then $N \approx M'$. It suffices to show that for finite tuples of the same length $\mathbf{m} \in M', \mathbf{n} \in N$, if $N \models \phi^{\alpha}_{\mathbf{m}}(\mathbf{n})$ then for every $m \in M'$ there is an $n \in N$ such that (forth)

$$N \models \phi^{\alpha}_{\mathbf{m}m}(\mathbf{n}n)$$

and (back) for each $n \in N$ there is an $m \in M'$ such that:

$$N \models \phi^{\alpha}_{\mathbf{m}n}(\mathbf{m}m).$$

For the 'forth', note that since $N \models \phi$, $N \models \phi^{\alpha+1}_{\mathbf{m}}(\mathbf{n})]$. This implies $N \models (\exists x)\phi^{\alpha}_{\mathbf{m}m}(\mathbf{n}x)]$ and this exactly what is needed. The 'back' is similar.

Now the downward Löwenheim-Skolem theorem for $L_{\omega_1,\omega}$ yields that ϕ_M is complete and $M \models \phi_M$ by construction. $\square_{7.1.6}$

A complete sentence in $L_{\omega_1,\omega}$ is necessarily $\aleph_0$-categorical (using downward Löwenheim-Skolem). Moreover, every countable structure is characterized by a complete sentence – its Scott sentence. So if a model satisfies a complete sentence, it is $L_{\infty,\omega}$-equivalent to a countable model.

Let M be the only model of power κ of an $L_{\omega_1,\omega}$-sentence ψ. We want to find sufficient conditions so that there is a complete sentence ψ' which implies ψ and is true in M. We will consider two such conditions: ψ has arbitrarily large models; ψ has few models of cardinality $\aleph_1$. Scott's theorem and the downward Löwenheim-Skolem theorem easily yield the next result.

Exercise 7.1.7. *If ψ is a complete sentence in $L_{\omega_1,\omega}$ in a countable vocabulary τ then every model M of ψ realizes only countably many $L_{\omega_1,\omega}$-types over the empty set. I.e., every model of ψ is small.*

One key tool for our study is a different representation for the class of models of complete $L_{\omega_1,\omega}$-sentences, as reducts of the atomic models of a complete first order theory (Theorem 7.1.12). We begin with a somewhat weaker characterization: as models omitting a set of partial types. This result stems from Chang, Scott, and Lopez-Escobar (see e.g. [33]).

Theorem 7.1.8. *Let ψ be a sentence in $L_{\omega_1,\omega}$ in a countable vocabulary τ. Then there is a countable vocabulary τ' extending τ, a first order τ'-theory T, and a countable collection of τ'-types Γ such that reduct is a* 1-1 *map from the models of T which omit Γ onto the models of ψ.*

Proof. Expand τ to τ' by adding a predicate $P_\phi(\mathbf{x})$ for each $L_{\omega_1,\omega}$ formula ϕ which is a subformula of ψ. This includes a nullary predicate symbol P_χ, taken as a propositional constant, for any subformula χ of ψ that is a sentence. Form the L'-theory T by adding axioms saying that for atomic formula ϕ, $P_\phi \leftrightarrow \phi$ and the P_ϕ represent each finite Boolean connective and quantification faithfully: E.g.

$$P_{\neg\phi(\mathbf{x})} \leftrightarrow \neg P_\phi(\mathbf{x}),$$

and

$$P_{(\forall \mathbf{x})\phi(\mathbf{x})} \leftrightarrow (\forall \mathbf{x}) P_\phi(\mathbf{x}),$$

and that, as far as first order logic can, the P_ϕ preserve the infinitary operations: for each i,

$$P_{\phi_i(\mathbf{x})} \rightarrow P_{\bigvee_i \phi_i(\mathbf{x})}$$

and

$$P_{\bigwedge_i \phi_i(\mathbf{x})} \rightarrow P_{\phi_i(\mathbf{x})}.$$

Consider the set Γ of partial types

$$p_{\bigwedge_i \phi_i(\mathbf{x})} = \{\neg P_{\bigwedge_i \phi_i(\mathbf{x})}\} \cup \{P_{\phi_i(\mathbf{x})} : i < \omega\}$$

where $\bigwedge_i \phi_i(\mathbf{x})$ is a subformula of ψ and

$$p_{\bigvee_i \phi_i(\mathbf{x})} = \{P_{\bigwedge_i \phi_i(\mathbf{x})}\} \cup \{\neg P_{\phi_i(\mathbf{x})} : i < \omega\}$$

where $\bigvee_i \phi_i(\mathbf{x})$ is a subformula of ψ.

Finally add to T the axiom: P_ψ. Since Γ contains a type for each infinite conjunction that is a subformula if M is a model of T which omits all the types in Γ, $M \restriction \tau \models \psi$. Moreover, each model of ψ has a *unique* expansion to a model of T which omits the types in Γ (since this is an expansion by definitions in $L_{\omega_1,\omega}$). $\Box_{7.1.8}$

Since all the new predicates in the reduction described above are $L_{\omega_1,\omega}$-definable this is a natural extension of Morley's procedure of replacing each first order formula ϕ by a predicate symbol P_ϕ. In the first order case, this quantifier elimination guarantees amalgamation over *arbitrary* sets for first order categorical T; the amalgamation does *not* follow in this case. (Compare Chapter 13 of [117].)

We have represented the models of ψ as a $PC\Gamma$ class in the sense of Definition 5.26. Reformulating Theorem 7.1.8 in this language, we have shown the following result for $L_{\omega_1,\omega}$.

Corollary 7.1.9. *Every $L_{\omega_1,\omega}$-sentence in a countable language is ω-presented. That is, the class of models of ψ is a $PC(\aleph_0, \aleph_0)$-class.*

Exercise 7.1.10. *Show that ψ is a sentence in $L_{\lambda^+,\omega}$ in a language of cardinality κ, ψ is μ-presented where μ is the larger of κ and λ.*

Exercise 7.1.11. *In general a $PC\Gamma$ class will not be an AEC class of τ structures. Why?*

More strongly, since there is a 1-1 correspondence between models of ψ and models of T that omit Γ, we can reduce spectrum considerations for an $L_{\omega_1,\omega}$-sentence ψ to the study of an $EC(T,\Gamma)$-class (Definition 5.30). In general there may be uncountably many complete L^*-types over the empty set consistent with ψ where L^* is the minimum fragment containing ψ. If ψ is complete, then there are only countably many $L_{\omega_1,\omega}$-types realized in a model of ψ. We can modify the proof of Theorem 7.1.8 to guarantee that the models of T omitting Γ are *atomic*.

Theorem 7.1.12. *Let ψ be a* complete *sentence in $L_{\omega_1,\omega}$ in a countable vocabulary τ. Then there is a countable vocabulary τ' extending τ and a complete first order τ'-theory T such that reduct is a* 1-1 *map from the* atomic *models of T onto the models of ψ.*

Proof. Fix a model M of ψ. Let L^* be a countable fragment of $L_{\omega_1,\omega}$ that contains ψ and the formulas $\phi^\alpha_{\boldsymbol{a}}$ from Lemma 7.1.5. Thus, *every* $L_{\omega_1,\omega}$-type consistent with ψ is represented by an L^*-formula. Now modify the proof of Theorem 7.1.8 by adding predicates $P_\chi(\mathbf{x})$ for every $\chi(\mathbf{x}) \in L^*$. In particular each $\boldsymbol{a}$

realizes the isolated type generated by $P_{\phi^{\alpha}_{\boldsymbol{a}}}$. The resulting theory T and collection of types Γ is as required. $\square_{7.1.12}$

So in particular, *any complete sentence of $L_{\omega_1,\omega}$ can be replaced (for spectrum purposes) by considering the atomic models of a first order theory*.

The crux is that to get an *atomic* and so $\aleph_0$-categorical class rather than just a $PC\Gamma(\aleph_0, \aleph_0)$-class, we have to first have a complete sentence. This will be our task in the next two sections (Theorem 7.2.12 and Theorem7.3.2). But first we give an example (derived from more specific examples of Marker and Kueker) showing that this a real task.

Example 7.1.13. Let ϕ be a sentence of $L_{\omega_1,\omega}$ with $2^{\aleph_0}$ countable models and no uncountable models. Let ψ be a sentence that is categorical (with models) in every infinite cardinal. Then $\phi \vee \psi$ is categorical in all powers and has $2^{\aleph_0}$ countable models. No such example with joint embedding is known.

7.2 Arbitrarily Large Models

To show a categorical sentence with arbitrarily large models is implied by a complete sentence we use Ehrenfeucht-Mostowski models. Morley's omitting types theorem (Section 7.2 of [34], Appendix A) is a fundamental technique in first order model theory. It is essential for the foundations of simplicity theory and for the construction of indiscernibles in infinitary logic. We use the first order version here; in Lemma 15.2 we prove the analog for abstract elementary classes.

Notation 7.2.1.
1. *For any linearly ordered set $X \subseteq M$ where M is a τ'-structure and $\tau' \supseteq \tau$, we write $\mathbf{D}_\tau(X)$ (diagram) for the set of τ-types of finite sequences (in the given order) from X. We will omit τ if it is clear from context.*
2. *Such a diagram of an order indiscernible set, $\mathbf{D}_\tau(X) = \Phi$, is called 'proper for linear orders' .*
3. *If X is a sequence of τ-indiscernibles with diagram $\Phi = \mathbf{D}_\tau(X)$ and any τ model of Φ has built in Skolem functions, then for any linear ordering I, $EM(I, \Phi)$ denotes the τ-hull of a sequence of order indiscernibles realizing Φ.*
4. *If $\tau_0 \subset \tau$, the reduct of $EM(I, \Phi)$ to τ_0 is denoted $EM_{\tau_0}(I, \Phi)$.*

What we here call the 'diagram' of a set indiscernibles is variously referred to as a template or an EM-set.

Exercise 7.2.2. *Suppose τ 'contains Skolem functions' and $X \subset M$ is sequence of order indiscernibles with diagram Φ. Show that for any linearly ordered set Z, $EM(Z, \Phi)$ is a model that is τ-elementarily equivalent to M.*

Lemma 7.2.3. *If $(X, <)$ is a sufficiently long linearly ordered subset of a τ-structure M, for any τ' extending τ (the length needed for X depends on $|\tau'|$) there is a countable set Y of τ'-indiscernibles (and hence one of arbitrary order type) such that $\mathbf{D}_\tau(Y) \subseteq \mathbf{D}_\tau(X)$. This implies that the only (first order) τ-types realized in $EM(X, \mathbf{D}_{\tau'}(Y))$ were realized in M.*

The phrase 'sufficiently long' is evaluated in Theorem 9.17 using Theorem A.3.1 (which implies Lemma 7.2.3). At this stage, we stress the overall outline of the argument in the easiest case. We need a little background on orderings; see [116] for a fuller treatment.

Definition 7.2.4. *A linear ordering $(X, <)$ is k-transitive if every map between increasing k-tuples extends to an order automorphism of $(X, <)$.*

Exercise 7.2.5. *Show any 2-transitive linear order is k-transitive for all finite k.*

Exercise 7.2.6. *Show there exist 2-transitive linear orders in every cardinal; hint: take the order type of an ordered field.*

Exercise 7.2.7. *If $\Phi(Y)$ is the diagram of a sequence of τ-order indiscernibles, show any order automorphism of Y extends to an automorphism of the τ-structure $EM(Y, \Phi)$.*

Definition 7.2.8. *For any model M and a, B contained in M, the* Galois-type *of a over B in M is the orbit of a under the automorphisms of M which fix B.*

Remark 7.2.9 (Warning)**.** The notion of Galois type requires an ambient model M. To link the notion of Galois-type, or rather saturation with respect to Galois types, with a notion of homogeneity, one must make an amalgamation hypotheses. Then one can find a 'monster model' in which to work. We develop Galois types over models under appropriate assumptions in Chapter 9 and the sequel. We will speak indiscriminately of the number of Galois types over M as an upper bound on the number of Galois n-types for any finite n.

Exercise 7.2.10. *If Y is a 2-transitive linear ordering, then for any τ and Φ is proper for linear orders, $EM(Y, \Phi)$ has $|\tau|$ Galois types over the empty set.*

The following easy exercise uses the notion of $\mathcal{L}$-type from Definition 5.22.

Exercise 7.2.11. *For any reasonable logic $\mathcal{L}$ (i.e. a logic such that truth is preserved under isomorphism) and any model M the number of $\mathcal{L}$-types over the empty set in M is at most the number of Galois types over the empty set in M.*

Using Ehrenfeucht-Mostowski models we now show any sentence in $L_{\omega_1,\omega}$ with arbitrarily large models has arbitrarily large models that are $L_{\omega_1,\omega}$-equivalent to a countable model. We have used Galois types in the proof of Corollary 7.2.12 to prefigure our work in Chapter 9, but this was not essential; see the problem set in Chapter 13 of [80].

Corollary 7.2.12. *Suppose an $L_{\omega_1,\omega}(\tau)$-sentence ψ has arbitrarily large models.*

1. *In every infinite cardinality ψ has a model that realizes only countably many $L_{\omega_1,\omega}(\tau)$-types over the empty set.*
2. *Thus, if ψ is categorical in some cardinal κ, ψ is implied by a consistent complete sentence ψ', which has a model of cardinality κ.*

Proof. By Theorem 7.1.8, we can extend τ to τ' and choose a first order theory T and a countable set of types Γ such $\mathrm{mod}(\psi) = PC_\tau(T, \Gamma)$. Since ψ has arbitrarily large models we can apply Theorem 7.2.3 to find τ''-indiscernibles for a Skolemization of T in an extended language τ''. Now take an Ehrenfeucht-Mostowski τ''-model M for the Skolemization of T over a set of indiscernibles ordered by a 2-transitive dense linear order. Then for every n, M has only countably many orbits of n-tuples and so realizes only countably many complete types over the empty set in any logic where truth is preserved by automorphism – in particular in $L_{\omega_1,\omega}$. So the τ-reduct of M realizes only countably many $L_{\omega_1,\omega}(\tau)$-types. If ψ is κ-categorical, let ψ' be the Scott sentence of this Ehrenfeucht-Mostowski model with cardinality κ. $\square_{7.2.12}$

We will extend this result to $L_{\omega_1,\omega}(Q)$ in Section 7.4. We do not actually need the categoricity for Corollary 7.2.12; we included it to make the statement parallel to Theorem 7.3.2.

Exercise 7.2.13. *Show that the ψ' chosen in Theorem 7.2.12 is unique.*

7.3 Few models in small cardinals

Recall that the goal of this chapter is to show that a categorical sentence in $L_{\omega_1,\omega}$ ($L_{\omega_1,\omega}(Q)$) is implied by a complete sentence, which has an uncountable model. In section 7.2, we assumed the existence of arbitrarily large models, now we assume $I(\aleph_1, \psi) < 2^{\aleph_1}$. We rely on the undefinability of well-order in $L_{\omega_1,\omega}(Q)$, which we treated in Chapter 6. We treat the two cases together in the beginning of this section, then finish $L_{\omega_1,\omega}$; we complete the argument for $L_{\omega_1,\omega}(Q)$ in Section 7.4. For the first lemma analogous arguments work both $L_{\omega_1,\omega}(Q)$ and $L_{\omega_1,\omega}$.

Theorem 7.3.1. *If the $L_{\omega_1,\omega}(Q)$-τ-sentence ψ has a model of cardinality $\aleph_1$ which is L^*-small for every countable τ-fragment L^* of $L_{\omega_1,\omega}(Q)$, then ψ has a $L_{\omega_1,\omega}(Q)$-small model of cardinality $\aleph_1$.*

Proof. Add to τ a binary relation $<$, interpreted as a linear order of M with order type ω_1. Using that M realizes only countably many types in any τ-fragment, define an continuous increasing chain of countable fragments L_α for $\alpha < \aleph_1$ such that each type in L_α that is realized in M is a formula in $L_{\alpha+1}$. Extend the similarity type further to τ' by adding new $2n+1$-ary predicates $E_n(x, \mathbf{y}, \mathbf{z})$ and $n+1$-ary functions f_n. Let M satisfy $E_n(\alpha, \boldsymbol{a}, \mathbf{b})$ if and only if $\boldsymbol{a}$ and $\mathbf{b}$ realize the same L_α-type and let f_n map M^{n+1} into the initial ω elements of the order, so that $E_n(\alpha, \boldsymbol{a}, \mathbf{b})$ if and only if $f_n(\alpha, \boldsymbol{a}) = f_n(\alpha, \mathbf{b})$. Note:

1. $E_n(\beta, \mathbf{y}, \mathbf{z})$ refines $E_n(\alpha, \mathbf{y}, \mathbf{z})$ if $\beta > \alpha$;
2. $E_n(0, \boldsymbol{a}, \mathbf{b})$ implies $\boldsymbol{a}$ and $\mathbf{b}$ satisfy the same quantifier free τ-formulas;
3. If $\beta > \alpha$ and $E_n(\beta, \boldsymbol{a}, \mathbf{b})$, then for every c_1 there exists c_2 such that
   (a) $E_{n+1}(\alpha, c_1\boldsymbol{a}, c_2\mathbf{b})$ and
   (b) if there are uncountably many c such that $E_{n+1}(\alpha, c\boldsymbol{a}, c_1\boldsymbol{a})$ then there are uncountably many c such that $E_{n+1}(\alpha, c\mathbf{b}, c_2\mathbf{b})$.
4. f_n witnesses that for any $a \in M$ each equivalence relation $E_n(a, \mathbf{y}, \mathbf{z})$ has only countably many classes.

All these assertions can be expressed by an $L_{\omega_1,\omega}(Q)(\tau')$ sentence ϕ. Let L^* be the smallest τ'-fragment containing $\phi \wedge \psi$. Now by Theorem 6.1.8 there is a structure N of cardinality $\aleph_1$ satisfying $\phi \wedge \psi$ such that there is an infinite decreasing sequence $d_0 > d_1 > \ldots$ in N. For each n, define $E_n^+(\mathbf{x}, \mathbf{y})$ if for some i, $E_n(d_i, \mathbf{x}, \mathbf{y})$. Now using 1), 2) and 3) prove by induction on the quantifier rank of ϕ that $N \models E_n^+(\boldsymbol{a}, \mathbf{b})$ implies $N \models \phi(\boldsymbol{a})$ if and only if $N \models \phi(\mathbf{b})$ for every $L_{\omega_1,\omega}(Q)(\tau)$-formula ϕ. (Suppose the result holds for all n and all θ with quantifier rank at most γ. Suppose $\phi(\boldsymbol{a})$ is $(\exists x)\psi(\boldsymbol{a}, x)$ with $n = \lg(\boldsymbol{a})$, ψ has quantifier rank γ, and $E_n^+(\boldsymbol{a}, \mathbf{b})$. So for some i, $E_n(d_i, \boldsymbol{a}, \mathbf{b})$ and for some a, $N \models \psi(\boldsymbol{a}, a)$. By the conditions on the E_n, there is a b such that $E_{n+1}(d_{i+1}, \boldsymbol{a}, a, \mathbf{b}, b)$. By induction we have $N \models \psi(\mathbf{b}, b)$ and so $N \models \phi(\mathbf{b})$. Use 3b) for the Q-quantifier.) For each n, $E_n(d_0, \mathbf{x}, \mathbf{y})$ refines $E_n^+(\mathbf{x}, \mathbf{y})$ and by 4) $E_n(d_0, \mathbf{x}, \mathbf{y})$ has only countably many classes; so N is small. $\square_{7.3.1}$

Now we show that sentences of $L_{\omega_1,\omega}(Q)$ ($L_{\omega_1,\omega}$) that have few models can be extended to complete sentences in the same logic that have uncountable models. We rely on Theorem 6.2.5 for $L_{\omega_1,\omega}$ and Corollary 5.10 of [79] for $L_{\omega_1,\omega}(Q)$). We get the main result of this chapter for $L_{\omega_1,\omega}$ now; we extend to $L_{\omega_1,\omega}(Q)$ in Section 7.4

Theorem 7.3.2. *If an $L_{\omega_1,\omega}$-sentence ψ has fewer than $2^{\aleph_1}$ models of cardinality $\aleph_1$ then there is a complete $L_{\omega_1,\omega}$-sentence ψ_0 that implies ψ and has a model of cardinality $\aleph_1$.*

Proof. By Theorem 6.2.5 every model of ψ of cardinality $\aleph_1$ is L^*-small for every countable fragment L^*. By Theorem 7.3.1 ψ has a model of cardinality $\aleph_1$ which is τ-small. By Lemma 7.1.6, we finish. $\square_{7.3.2}$

A weaker version of this result (requiring fewer than $2^{\aleph_0}$ models of cardinality $\aleph_1$) was obtained by Makkai using admissible model theory in [104]. We used the proof of Theorem 7.3.2 to provide a more direct proof that any counterexample to Vaught's conjecture has a small model of size $\aleph_1$ in [12].

So to study categoricity of an $L_{\omega_1,\omega}$-sentence ψ, we have established the following reduction. If ψ has arbitrarily large models, without loss of generality, ψ is complete and therefore small. If ψ has few models of power $\aleph_1$, we can study a subclass of the models of ψ defined by a complete $L_{\omega_1,\omega}$-sentence ψ'. We will in fact prove sufficiently strong results about ψ' to deduce a nice theorem for ψ. Note that since ψ' is complete, the models of ψ' form an $EC(T, Atomic)$-class in an extended similarity type τ'. We study these atomic classes in Chapters 19-27.

Conclusion 7.3.3. *The uncountable models of $\aleph_1$-categorical sentences in $L_{\omega_1,\omega}$ can be studied by proving sufficiently strong results about atomic classes.*

By the Scott isomorphism theorem, the reduction to complete sentences destroys any study of the countable models of an incomplete sentence in $L_{\omega_1,\omega}$. In particular then this reduction is not helpful for studying Vaught's conjecture that no sentence of $L_{\omega_1,\omega}$ can have exactly $\aleph_1$-countable models. The notion of a finitary AEC [65, 64, 66] may provide some tools for this problem.

7.4 Categoricity and Completeness for $L_{\omega_1,\omega}(Q)$

In this section we explore formulating Theorem 7.2.12 and Conclusion 7.3.3 for $L_{\omega_1,\omega}(Q)$. We will see that obtaining a complete sentence proceeds in the same way. But then there is much more to the story.

Definition 7.4.1. *A sentence ψ in $L_{\omega_1,\omega}(Q)$ is called* complete *for $L_{\omega_1,\omega}(Q)$ if for every sentence ϕ in $L_{\omega_1,\omega}(Q)$, either $\psi \models \phi$ or $\psi \models \neg\phi$.*

We first note that a 'small' model has an $L_{\omega_1,\omega}(Q)$-Scott sentence. The details of this argument are due to David Kueker; the assertion is implicit in [120].

Theorem 7.4.2. *Suppose the τ-structure M realizes only countably many $L_{\omega_1,\omega}(Q)$ types. Then there is a* complete *sentence σ_M of $L_{\omega_1,\omega}(Q)$ such that $M \models \sigma_M$.*

Proof. Choose a subset A of M as in Lemma 7.1.5. We showed in Lemma 7.1.5 that for each $\boldsymbol{a} \in M$ there is a formula $\phi_{\boldsymbol{a}}$ that decides all formulas in M with $|\boldsymbol{a}|$ free variables. Now let σ_M be the conjunction of the following sentences:

1. $(\forall x) \bigvee_{a\in A} \phi_a(x)$
2. $\bigwedge_{a\in A} (\exists x)\phi_a(x)$
3. $(\forall \mathbf{x})[\phi_{\boldsymbol{a}}(\mathbf{x}) \to (\forall x) \bigvee_{a\in A_{\boldsymbol{a},\phi_{\boldsymbol{a}}}} \phi_{\boldsymbol{a},a}(\mathbf{x}, x)]$
4. $(\forall \mathbf{x})[\phi_{\boldsymbol{a}}(\mathbf{x}) \to \bigwedge_{a\in A_{\boldsymbol{a},\phi_{\boldsymbol{a}}}} (\exists x)\phi_{\boldsymbol{a},a}(\mathbf{x}, x)]$
5. $(\forall \mathbf{x})[\phi_{\boldsymbol{a}}(\mathbf{x}) \to \lambda(\mathbf{x})$ for any basic formula λ satisfied by $\boldsymbol{a}$.
6. $(Qx)\phi_a(x)$ if this sentence is true in M.
7. $\neg(Qx)\phi_a(x)$ if this sentence is true in M.

8. $(\forall \mathbf{x})[\phi_{\boldsymbol{a}}(\mathbf{x}) \rightarrow (Qx)\phi_{\boldsymbol{a}a}(\mathbf{x}x)]$ if this sentence is true in M.
9. $(\forall \mathbf{x}[\phi_{\boldsymbol{a}}(\mathbf{x}) \rightarrow \neg(Qx)\phi_{\boldsymbol{a}a}(\mathbf{x}x)]$ if this sentence is true in M.

Claim 7.4.3. *If $N \models \sigma_M$ then for all $\boldsymbol{a} \in A$ and all $\mathbf{b} \in N$, if $N \models \phi_{\boldsymbol{a}}(\mathbf{b})$, then for any formula ψ,*

$$M \models \psi(\boldsymbol{a}) \text{ if and only if } N \models \psi(\mathbf{b}).$$

Proof. We prove this result by induction on ψ. Suppose $N \models \phi_{\boldsymbol{a}}(\mathbf{b})$. If $\psi(\mathbf{x})$ is quantifier-free then $\phi_{\boldsymbol{a}}(\mathbf{x}) \rightarrow \psi(\mathbf{x})$ by definition. The interesting cases are when ψ has the form $(\exists x)\chi(\mathbf{x}, x)$, $(\forall x)\chi(\mathbf{x}, x)$, or $(Qx)\chi(\mathbf{x}, x)$. By the induction hypothesis, for every $a \in A$ and $b \in N$, if $N \models \phi_{\boldsymbol{a}a}(\mathbf{b}b)$ then for any $\chi(\mathbf{x}, y)$, $M \models \chi(\boldsymbol{a}a)$ if and only if $N \models \chi(\mathbf{b}b)$. For $(\exists x)$ apply condition 4); for $(\forall x)$, apply condition 3); Q is slightly more involved. By our choice of $\phi_{\boldsymbol{a}a}$ to decide all formulas on M, we know either $N \models \phi_{\boldsymbol{a}a}(\mathbf{x}x) \rightarrow \chi(\mathbf{x}x)$ or $N \models \phi_{\boldsymbol{a}a}(\mathbf{x}x) \rightarrow \neg\chi(\mathbf{x}x)$. Now since A is countable and using the last observation, the following four statements are equivalent.

- $M \models (Qx)\chi(\boldsymbol{a}, x)$.
- There is an $a \in A$ such that $\chi(\boldsymbol{a}, a)$ and $M \models (Qx)\phi_{\boldsymbol{a}a}(\boldsymbol{a}x)$.
- There is an $a \in A$ and $b \in N$ such that $N \models \phi_{\boldsymbol{a}a}(\mathbf{b}b)$, $N \models \chi(\mathbf{b}, b)$ and $N \models (Qx)\phi_{\mathbf{b}b}(\mathbf{b}x)$.
- $N \models (Qx)\chi(\mathbf{b}, x)$.

The equivalence of the second pair uses 8) and 9). We conclude M and N agree on all sentences of $L_{\omega_1,\omega}(Q)$ by the claim. $\square$7.4.2

Just as in proving Theorem 7.3.2 we can invoke Theorem 6.2.5 and Theorem 7.3.1 to conclude:

Theorem 7.4.4. *If an $L_{\omega_1,\omega}(Q)$-sentence ψ has fewer than $2^{\aleph_1}$ models of cardinality $\aleph_1$ then there is a complete* small *$L_{\omega_1,\omega}(Q)$-sentence ψ_0 that implies ψ and has a model of cardinality $\aleph_1$ and such that every model of ψ_0 is small.*

Theorem 7.4.4 differs from Theorem 7.3.2 because the sentence ψ_0 in Theorem 7.4.4 does *not* have a countable model (it implies $Qx(x = x)$). We further transform ψ_0 to a ψ' as in the following definition.

Definition 7.4.5. *Let ψ_0 be a $L_{\omega_1,\omega}(Q)$-complete sentence with vocabulary τ in the countable fragment L^* of $L_{\omega_1,\omega}(Q)$ such that every model of ψ_0 is small. Form τ' by adding predicates for formulas as in Theorem 7.1.8 but also add for each formula $(Qx)\phi(x, \mathbf{y})$ a predicate $R_{(Qx)\phi(x,\mathbf{y})}$ and add the axiom*

$$(\forall x)[(Qx)\phi(x, \mathbf{y}) \leftrightarrow R_{(Qx)\phi(x,\mathbf{y})}].$$

Let ψ' be the conjunction of ψ_0 with the $L_{\omega_1,\omega}(Q)$-τ'-axioms encoding this expansion. Let $\boldsymbol{K}_1$ be the class of atomic τ'-models of $T(\psi)$, the first order τ'-theory containing all first order consequences of ψ'.

Notation 7.4.6. *1. Let $\leq^*$ be the relation on $\boldsymbol{K}_1$: $M \leq^* N$ if $M \prec_{\tau'} N$ and for each formula $\phi(x,\mathbf{y})$ and $\mathbf{m} \in M$, if $M \models \neg R_{(Qx)\phi(x,\mathbf{m})}$ then $R_{\phi(x,\mathbf{m})}$ has the same solutions in M and N.*

2. Let $\leq^{**}$ be the relation on $\boldsymbol{K}_1$: $M \leq^{**} N$ if $M \prec_{\tau'} N$ and for each formula $\phi(x,\mathbf{y})$ and $\mathbf{m} \in M$, $M \models \neg R_{(Qx)\phi(x,\mathbf{m})}$ if and only if $R_{\phi(x,\mathbf{m})}$ has the same solutions in M and N.

It is easy to check that $(\boldsymbol{K}_1, \leq^*)$ is an AEC, but $(\boldsymbol{K}_1, \leq^{**})$ need not be an AEC. It can easily happen that each of a family of models $M_i \leq^{**} M$ but $\bigcup_i M_i \not\leq^{**} M$.

To bring the Löwenheim-Skolem number down to $\aleph_0$, we allowed countable models of $\boldsymbol{K}_1$ that are not models of ψ; unfortunately, we may also have gained uncountable models of $\boldsymbol{K}_1$ that are not models of ψ. Working with $(\boldsymbol{K}_1, \leq^*)$, one cannot show that many models for $\boldsymbol{K}_1$ implies many models of ψ. We will see in Chapter 8, one example of using $(\boldsymbol{K}_1, \leq^{**})$ to work around this difficulty.

If we are willing to drop our demand that the target class have Löwenheim number $\aleph_0$, then we can construct an AEC from each sentence ϕ of $L_{\omega_1,\omega}(Q)$; namely consider the uncountable members of $\boldsymbol{K}_1$ with $\prec_{\boldsymbol{K}}$ as $\leq^*$. (Note that the restriction to complete sentences was not really needed to define $\leq^*$.) Thus the work in this monograph on AEC in Part III applies to sentences of $L_{\omega_1,\omega}(Q)$; but the analysis of categoricity of $L_{\omega_1,\omega}(Q)$ in Part IV, which uses Löwenheim number $\aleph_0$, requires very different arguments. For a much more sophisticated analysis see such not yet published articles such as [134, 131, 130].

Conclusion 7.4.7. *There is no known straightforward translation of (categorical) sentences of $L_{\omega_1,\omega}(Q)$ into AEC with Löwenheim number $\aleph_0$. There are two procedures for continuing the analysis of categorical $L_{\omega_1,\omega}(Q)$-sentences:*

1. *Work with various notions of substructure on an associated class $\boldsymbol{K}_1$ of atomic models.*
2. *Regard classes defined by sentences of $L_{\omega_1,\omega}(Q)$ as an AEC with Löwenheim number $\aleph_1$.*

8
A model in $\aleph_2$

A first order sentence with an infinite model has models in all cardinalities; in particular no sentence is *absolutely categorical*, has exactly one model. Sentences of $L_{\omega_1,\omega}$ may be absolutely categorical but only if the unique model has cardinality $\aleph_0$. In the early 70's, I asked whether a sentence of $L(Q)$ could have exactly one model and that model have cardinality $\aleph_1$. Shelah [120] showed the answer was no for $(L_{\omega_1,\omega}(Q))$ with some additional set theoretic hypotheses that he removed in [139]. In this chapter we introduce methods of getting structural properties on the models in an AEC that have cardinality λ by restricting the number of models of cardinality λ^+. And from these conditions on models of cardinality λ and λ^+, we show the existence of a model of power λ^{++}. Most concretely, we present Shelah's proof [139] that if a sentence of $L_{\omega_1,\omega}(Q)$ is categorical in $\aleph_1$ then it has a model of cardinality $\aleph_2$.

An example of Julia Knight [85] shows that there are sentences of $L_{\omega_1,\omega}$ that have models only in cardinality $\aleph_0$ and $\aleph_1$ so the restriction on the number of models is necessary. Hjorth [59] has extended this to find a sentence ϕ_α of $L_{\omega_1,\omega}$ that has models of cardinality only up $\aleph_\alpha$ for any countable ordinal α.

The general setting here will be an AEC. We show first that if an AEC is categorical in λ and λ^+ and has no 'maximal triple' in power λ then it has a model in power λ^{++}. We then restrict our attention to $L_{\omega_1,\omega}$. We have shown in Chapter 7 that there is no loss of generality in assuming an $\aleph_1$-categorical sentence in $L_{\omega_1,\omega}$ is also $\aleph_0$-categorical. Then we show for such sentences in $L_{\omega_1,\omega}$ there are no maximal triples in $\aleph_0$. We finish by massaging the proof to handle $L_{\omega_1,\omega}(Q)$.

Definition 8.1. *We say (M, N) is a* proper pair *in λ, witnessed by a, if $M \prec_{\boldsymbol{K}} N$ and $a \in N - M$ and $|M| = |N| = \lambda$.*

Shelah [139] describes this concept in terms of the class of triples $\boldsymbol{K}^3_\lambda = \{(M, N, a) : (M, N)$ is proper pair witnessed by $a\}$. The fixed a is not used in the next Lemma but plays a central role in the proof of Lemma 8.4.

Lemma 8.2. *If an AEC $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ is categorical in λ and has a proper pair (M, N) in λ then there is a model in $\boldsymbol{K}$ with cardinality λ^+.*

Proof. Let $M_0 = M$. For any α, given M_α, choose $M_{\alpha+1}$ so that $(M, N) \approx (M_\alpha, M_{\alpha+1})$ and take unions at limits. The union of M_α for $\alpha < \lambda^+$ is as required. $\square_{S8.2}$

Definition 8.3. *A* maximal triple *is a triple (M, a, N) such that a witnesses that (M, N) is a proper pair and there is no proper pair (M', N') witnessed by a such that $M \prec_{\boldsymbol{K}} M'$, $M \neq M'$, $N \prec_{\boldsymbol{K}} N'$.*

In Shelah's language a maximal triple is a maximal element when $\boldsymbol{K}^3_\lambda$ ordered by the relation

$$(M, a, N) \leq (M', a', N')$$

if $a = a'$, $N \prec_{\boldsymbol{K}} N'$, $M \prec_{\boldsymbol{K}} M'$ and this last inclusion is proper.

Lemma 8.4. *If there are no maximal triples of cardinality λ and there is a proper pair of cardinality λ then there is a proper pair of cardinality λ^+.*

Proof. Let a witness that (M_0, N_0) is a proper pair in λ. Since there are no maximal triples, for $i < \lambda^+$ we can construct proper pairs (M_i, N_i) such that M_{i+1} is a proper $\prec_{\boldsymbol{K}}$ extension of M_i and N_{i+1} is a $\prec_{\boldsymbol{K}}$ extension of N_i but no M_i contains a; that is, the properness of each (M_i, N_i) is witnessed by the same a. So $(\bigcup_{i<\lambda^+} M_i, \bigcup_{i<\lambda^+} N_i)$ is the required proper pair. $\square_{8.4}$

We have shown that if there are no maximal triples in λ and $\boldsymbol{K}$ is λ^+-categorical then there is a model in λ^{++}. We will show there are no maximal triples in $\aleph_0$ if $\boldsymbol{K}$ is $\aleph_0$-categorical and has few models in $\aleph_1$. For this, we need another definition.

Definition 8.5. *$M \prec_{\boldsymbol{K}} N$ is a* cut-pair *in λ if $|M| = |N| = \lambda$ and there exist models N_i for $i < \omega$ such that $M \prec_{\boldsymbol{K}} N_{i+1} \prec_{\boldsymbol{K}} N_i \prec_{\boldsymbol{K}} N$ with N_{i+1} a proper submodel of N_i and $\bigcap_{i<\omega} N_i = M$.*

Let $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ be the collection of dense linear orders with elementary submodel and let $(\mathbb{Q}, <)$ be the rational order. Then $((-\infty, \sqrt{2}), (-\infty, \infty))$ is a cut-pair. Tapani Hyttinen has provided the following barebones example of a maximal triple.

Example 8.6 (Maximal Triple)**.** Let τ contain unary predicate symbols P and S and infinitely many constant symbols c_i. Let $\boldsymbol{K}$ be the class of τ structures M such that.

1. P and S partition M.
2. There is at most one element in P.
3. The c_i are distinct and lie in S.

4. If there is an element in S not named by one of the c_i then there is an element in P.

It is easy to check that $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ (where $\prec_{\boldsymbol{K}}$ is interpreted as substructure) is an AEC.

Now let P be empty in M and suppose the c_i exhaust $S(M)$. Let N extend M by adding just a point a in P. Then (M, a, N) is a maximal triple. Of course, $\boldsymbol{K}$ is not $\aleph_0$-categorical.

We rely on the following classical result of Ulam (see e.g. 6.12 of [92] or II.4.12 of [40]) and an easy consequence.

Fact 8.7. *For any regular $\kappa > \omega$,*

1. *For any regular $\kappa > \omega$ there is family of κ stationary subsets of κ that are pairwise disjoint.*
2. *Moreover, there exist a family of stationary set S_i for $i < 2^\kappa$ such if $i \neq j$, $S_i - S_j$ is stationary.*

We give the idea of the following proof; the details are clear in both [139] and [48].

Lemma 8.8. *Suppose $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ is λ-categorical. If $\boldsymbol{K}$ has a cut-pair in cardinality λ and it has a maximal triple in λ, then $I(\lambda^+, \boldsymbol{K}) = 2^{\lambda^+}$.*

Proof. Let (M, N) be a cut-pair and let (M', a, N') be a maximal triple. For S a stationary subset of λ^+, define M_i^S for $i < \lambda^+$ so that the universe of M_i^S is a subset of λ^+ and the union of the M_i^S has universe λ^+. We demand that (M_i^S, M_{i+1}^S) is isomorphic to (M, N) if i is 0 or a successor ordinal. But if i is a limit ordinal, let (M_i^S, M_{i+1}^S) be a cut-pair if $i \notin S$ and for some a_i, let $(M_i^S, a_i, M_{i+1}^S) \approx (M', a, N')$ if $i \in S$. Then, let $M^S = \bigcup_{i<\lambda^+} M_i^S$. Now, if $S_1 - S_2$ is stationary, $M^{S_1} \not\approx M^{S_2}$. If f is an isomorphism between them, we find a contradiction by intersecting $S_1 - S_2$ with the cub E consisting of those $\delta < \lambda^+$ such that $M_\delta^{S_1}$ and $M_\delta^{S_2}$ both have domain δ and $i < \delta$ if and only if $f(i) < \delta$. If δ is in the intersection, as $\delta \in S_1$, $a_\delta^{S_1} \in M_{\delta+1}^{S_1} - M_\delta^{S_1}$; $f(a_\delta^{S_1}) \in M^{S_2} - M_\delta^{S_2}$. But, $M_\delta^{S_2} = \bigcap_{n<\omega} M_\delta^{S_2,n}$ for appropriate $M_\delta^{S_2,n}$, since $(M_\delta^{S_2}, M_{\delta+1}^{S_2})$ is a cut pair. So $f(a_\delta^{S_1}) \notin M_\delta^{S_2,n}$ for some n. Let N denote $f^{-1}(M_\delta^{S_2,n})$. Then for some $\gamma \in E \cap (S_1 - S_2)$, $N \prec_{\boldsymbol{K}} M_\gamma^{S_1}$. But then $(N, a_\delta^{S_1}, M_\gamma^{S_1})$ properly extends $(M_\delta^{S_1}, a_\delta^{S_1}, M_{\delta+1}^{S_1})$ and this contradiction yields the theorem. $\square_{8.8}$

Now we need the following result, which depends heavily on our restricting λ to be $\aleph_0$ and also requiring the AEC to be a $PC\Gamma(\aleph_0, \aleph_0)$ class. Some extensions to other cardinalities are mentioned in [139] and a more detailed argument appears in [133] and as Theorem 7.11 in [48]. See also [106].

Recall from Corollary 7.1.9 an $\aleph_1$-categorical sentence in $L_{\omega_1,\omega}$ can be represented as an AEC which is a $PC\Gamma(\aleph_0, \aleph_0)$ class.

Lemma 8.9. *If $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ is an $\aleph_0$-categorical $PC\Gamma(\aleph_0, \aleph_0)$ class that is also an AEC and has a model of power $\aleph_1$, then there is a cut pair in $\aleph_0$.*

Proof. By definition, $\boldsymbol{K}$ is the class of τ-reducts of models of a first order τ_1-theory T, which omit a countable set Γ of types. Let $M \in \boldsymbol{K}$ be a model with universe $\aleph_1$; write M as $\bigcup_{i<\aleph_1} M_i$ with the M_i countable. For simplicity, assume the universe of M_0 is $\aleph_0$. Expand M to a τ^*-structure M^* by adding to τ_1, an order $<$ and a binary function g. Interpret $<$ as the natural order on $\aleph_1$ and g so that $g(i, x)$ is a τ-isomorphism from M_0 to M_i witnessing that each M_i is countable. Note that a unary predicate P naming M_0 and a binary relation $R(x, y)$ such that $R(a, i)$ if and only $a \in M_i$ are easily definable from g. Moreover, for each i, $\{x : R(x, i)\}$ is closed under the functions of τ_1.

Let ψ be a sentence in $L_{\omega_1,\omega}(\tau^*)$ describing this situation; the existence of ψ follows since $\boldsymbol{K}$ is a $PC\Gamma(\aleph_0, \aleph_0)$ class. By Theorem 6.1.6 (Theorem 12 of [80]), there is a model N^* of ψ with cardinality $\aleph_0$ in which $<$ is not well-founded. For any $b \in N^*$, let

$$N_b = \{x \in N^* : R(x, b)\}.$$

Let a_i for $i < \omega$ be a properly descending chain. Then N_{a_i} has universe

$$\{x \in N^* : R(x, a_i)\}$$

and

$$N_{a_i} \restriction \tau \prec_{\boldsymbol{K}} N^* \restriction \tau$$

by the moreover clause of the presentation theorem, Theorem 5.14. Because of g, each N_{a_i} is τ-isomorphic to $P(N^*)$. Let I be the set of $b \in N^*$ such that for every i, $b < a_i$. By the union axiom, Definition 5.1 A3.3,

$$N_I = \bigcup_{b \in I} N_b \restriction \tau \prec_{\boldsymbol{K}} N^* \restriction \tau.$$

Our required cut-pair is (N_I, N_0). $\square_{8.9}$

Theorem 8.10. *If $\boldsymbol{K}$ is a $\aleph_0$-categorical $PC\Gamma(\aleph_0, \aleph_0)$ class that is also an AEC and has a unique model of power $\aleph_1$, then there is a model of power $\aleph_2$.*

Proof. By Lemma 8.9, there is a cut-pair in $\aleph_0$. Since ψ is $\aleph_1$-categorical, Lemma 8.8 implies there is no maximal triple in $\aleph_0$. So by Lemma 8.4 there is a proper pair in $\aleph_1$ and then by Lemma 8.2, there is a model of power $\aleph_2$. $\square_{8.10}$

Corollary 8.11. *An $\aleph_1$-categorical sentence ψ in $L_{\omega_1,\omega}$ has a model of power $\aleph_2$.*

Proof. By Theorem 7.3.2, we may assume ψ has the form of Theorem 8.10. $\square_{8.11}$

We want to extend Corollary 8.11 from $L_{\omega_1,\omega}$ to $L_{\omega_1,\omega}(Q)$. The difficulty is to find an appropriate AEC. By Theorem 7.4.4, we can find a small $L_{\omega_1,\omega}(Q)$-complete sentence ψ_0 which is satisfied by the model of cardinality $\aleph_1$. Since ψ_0 is small, the associated $T(\psi)$ can be taken $\aleph_0$-categorical.

Recall the associated classes $(\boldsymbol{K}_1, \leq^*)$ and $(\boldsymbol{K}_1, \leq^{**})$ from Chapter 7, Notation 7.4.6. $\boldsymbol{K}_1$ is a class of models of a first order theory $T(\psi)$. (cf. Definition 7.4.5.) We call these 'nonstandard models' of ψ while standard models are models of the original $L_{\omega_1,\omega}(Q)$ sentence ψ. We would like to transfer our attention to $\boldsymbol{K}_1$, But, $\boldsymbol{K}_1$ may have $2^{\aleph_1}$ (nonstandard) models of cardinality $\aleph_1$. Although $(\boldsymbol{K}_1, \leq^*)$ is an AEC, we can't just work with it as in that class we don't generate $2^{\aleph_1}$ (standard) models of ψ. And we can't just apply our results to $(\boldsymbol{K}_1, \leq^{**})$ as it does not satisfy Definition 5.1 A3.3. However, by working with $(\boldsymbol{K}_1, \leq^{**})$, we are able to apply the idea of the proof of Lemma 8.8 and get the maximal number of standard models in $\aleph_1$. The failure to satisfy Definition 5.1 A3.3 is evaded by proving the following variant on Lemma 8.9.

Lemma 8.12. *$(\boldsymbol{K}_1, \leq^{**})$ has a cut pair in $\aleph_0$.*

Proof. The argument is identical to Lemma 8.9 except for two key points.

In Lemma 8.9, we described an expansion of a model of power $\aleph_1$ in the first paragraph. In the second paragraph we formalized this description in $L_{\omega_1,\omega}(\tau^*)$, relying on the assumption $\boldsymbol{K}$ is a $PC\Gamma(\aleph_0, \aleph_0)$ class. For the $L_{\omega_1,\omega}(Q)$-case, just write the description in $L_{\omega_1,\omega}(Q)$ and apply Theorem 6.1.8 rather than Theorem 6.1.6. We can use a sentence of $L_{\omega_1,\omega}(Q)$ since from the $\aleph_1$-categoricity we know our sentence is small.

The penultimate sentence of the proof of Lemma 8.9 read: by the union axiom: Definition 5.1 A3.3,

$$N_I = \bigcup_{b \in I} N_b \restriction \tau \prec_{\boldsymbol{K}} N^* \restriction \tau.$$

This is precisely the union axiom that fails for $\leq^{**}$. But in this situation, for any $i < \omega$ we have $N_I \leq^* N_{i+1} \leq^{**} N_i$ so $N_I \leq^{**} N_i$, which is exactly what we need. $\square_{8.12}$

Now we can get the main result. We are working with the class $(\boldsymbol{K}_1, \leq^{**})$, which is not an AEC. So we must check that each time we use an argument about AEC's, we do not rely on the condition that fails for $(\boldsymbol{K}_1, \leq^{**})$. We can naturally define cut pairs and maximal triples in this context.

Corollary 8.13. *An $\aleph_1$-categorical sentence ψ in $L_{\omega_1,\omega}(Q)$ has a model of power $\aleph_2$.*

Proof. Since Lemma 8.2 uses only A3.1 of Definition 5.1, which holds of $(\boldsymbol{K}_1, \leq^{**})$, it suffices to show there is a $(\boldsymbol{K}_1, \leq^{**})$ proper pair of cardinality $\aleph_1$ that are standard models of ψ. By the proof of Lemma 8.12, there is a $(\boldsymbol{K}_1, \leq^{**})$ cut-pair in $\aleph_0$. Lemma 8.8 does not depend on A3.3. So by the same proof, we get $2^{\aleph_1}$ non-isomorphic models of $\boldsymbol{K}_1$ and since we took $\leq^{**}$-extensions $\aleph_1$-times, each is actually a standard model of ψ. But this contradicts the categoricity of ψ so there must no maximal triple in $(\boldsymbol{K}_1, \leq^{**})$. By Lemma 8.4 (which again does not depend on A3.3) we have a standard $(\boldsymbol{K}_1, \leq^{**})$ proper pair of cardinality $\aleph_1$ and we finish. $\square_{8.13}$

[139], [48] and [106] were the major sources for this Chapter. Coppola [35] has provided a notion of Q-abstract elementary class which allows one to give an axiomatic proof including the case of $L_{\omega_1,\omega}(Q)$.

Part III

Abstract Elementary Classes with Arbitrarily Large Models

Parts Three and Four addresses two distinct themes in the study of non-elementary classes. In Part 3, we study abstract elementary classes under a few additional conditions and try to determine the eventual behavior of the class. The motivating problem is Shelah's conjecture that if a class is categorical in arbitrarily large cardinalities, then it is categorical for all sufficiently large cardinals. Thus, we assume that each $\boldsymbol{K}$ has arbitrarily large models. Part Four addresses the conjecture that if there are few models in each cardinal then the class must have arbitrarily large models; it proceeds in a much more specific vein by studying classes defined in $L_{\omega_1,\omega}$.

In chapters 9-16 we work under the following strong assumption. We will see that strong analogs of Morley's theorem can be proved in this context.

Assumption. $\boldsymbol{K}$ is an abstract elementary class.

1. $\boldsymbol{K}$ has arbitrarily large models.
2. $\boldsymbol{K}$ satisfies the amalgamation property and the joint embedding property.

Part III largely expounds the work of Shelah in [128] but includes extensions by the Grossberg, VanDieren, Hyttinen and the author.

Recall that we say $\boldsymbol{K}$ has the *amalgamation property* if $M \le N_1$ and $M \le N_2 \in K$ with all three in $\boldsymbol{K}$ implies there is a common strong extension N_3 completing the diagram. *Joint embedding* means any two members of $\boldsymbol{K}$ have a common strong extension. Crucially, we amalgamate only over members of $\boldsymbol{K}$; this distinguishes this context from the context of homogeneous structures.

As we note in Lemma 17.14, it is easy to partition an AEC with the amalgamation property into a collection of AEC which each have the joint embedding property as well. This is an essential, albeit trivial, reduction; it is easy to construct counterexamples (with sporadic small models) to the main results in this analysis if the joint embedding hypothesis is dropped.

In Chapter 9, we introduce the crucial notion of a Galois-type. Although this notion can be defined without assuming amalgamation, it is best-behaved under that hypothesis. With amalgamation, we have a monster (homogeneous-universal with respect to strong embeddings) model and the Galois types are the orbits within this model. Even for classes defined in very nice logics (e.g. $L_{\omega_1,\omega}$) and which satisfy amalgamation, the Galois type may differ from the natural syntactic notion (see Chapter 27). We define saturation and stability for Galois types in Chapter 9 and prove that categoricity in λ implies Galois-μ-stability for $\mu < \lambda$.

The notion of a brimful model introduced in Chapter 10 is used in Chapter 11 to discuss the relationship between special, limit, saturated models and models generated by sequences of indiscernibles. Limit models are an alternative which apply in cases where Galois-saturation is vacuous (in $\mathrm{LS}(\boldsymbol{K})$) or where there are not enough saturated models. They refine the notion of saturation in some (superstable) situations. We discuss three properties (tameness, locality, compactness) that are automatic for first order syntactic types but problematic for Galois types in Chapter 12. We defined in Chapter 5 the Hanf number H_1 for AEC with vocabulary of a fixed cardinality. The most important result in Chapter 12 is: if an

AEC satisfying our general assumption is categorical in $\lambda > H_1$ then for any $\mu < \mathrm{cf}(\lambda)$ it is (χ_μ, μ) (weakly) tame for some $\chi_\mu < H_1$. (The proof due to Baldwin-Hyttinen substantially simplifies the argument advanced in [128].)

In Chapter 13, we introduce the appropriate notion of independence for this context, based on non-splitting and develop its basic properties. In Chapter 14, we prove the theorem of Grossberg-VanDieren that categoricity in a cardinal and its successor transfer upwards for *tame* AEC. Chapter 15 generalizes the techniques of Keisler and Morley to AEC and invokes the notion of splitting to prove that if $\boldsymbol{K}$ is categorical in some successor $\lambda > H_2$ (a larger Hanf number explained in the introduction to Chapter 15) then it is categorical on the interval $[H_2, \lambda]$. (This slightly weakens the claim in [128].) Finally, in Chapter 16, we use the methods of Chapter 10, non-splitting and E-M models to analyze a kind of superstability for this context. We prove that that if $\boldsymbol{K}$ is categorical in λ then for μ less than λ, the union of $< \mu^+$ models that are μ-saturated is also μ-saturated. Assuming tameness, this yields directly upward categoricity transfer from a successor λ^+ with $\lambda > \mathrm{LS}(\boldsymbol{K})$. This yields Conclusion 16.13: There is a cardinal μ depending on κ such that $\boldsymbol{K}$ is an AEC with $\kappa_{\boldsymbol{K}} = \kappa$, and $\boldsymbol{K}$ is categorical in some successor cardinal $\lambda^+ > \mu$, then $\boldsymbol{K}$ is categorical in all cardinals greater than μ.

In the last two Chapters of this Part, we drop our global assumption of amalgamation. Chapter 17 explores some results that can be proved without assuming amalgamation and describes some weakenings of that hypothesis. We show in Chapter 18 that, assuming the weak diamond, if $\boldsymbol{K}$ has few ($< 2^{\lambda^+}$) models of cardinality λ^+ then it has amalgamation in λ. This argument plays a key role in Part IV. The result also provides a justification for the hypotheses of this Part. Chapter 18 shows eventual categoricity implies eventual amalgamation; the main theme of Part III is a partial converse.

9
Galois types, Saturation, and Stability

In this chapter we take advantage of the standing assumptions in Part III, joint embedding and amalgamation to find a monster model. That is, we give an abstract account of Morley-Vaught [112]. We then define the fundamental notion of Galois type in terms of orbits of stabilizers of submodels. This allows an identification of 'model-homogeneous' with 'saturated'. We further show that a λ-categorical AEC is μ-stable for μ with $\mathrm{LS}(\boldsymbol{K}) \leq \mu < \lambda$ and that the model in the categoricity cardinal is Galois-saturated.

Definition 9.1. *1. M is* μ-model homogeneous *if for every $N \prec_{\boldsymbol{K}} M$ and every $N' \in \boldsymbol{K}$ with $|N'| < \mu$ and $N \prec_{\boldsymbol{K}} N'$ there is a $\boldsymbol{K}$-embedding of N' into M over N.*

2. M is strongly μ-model homogeneous *if it is μ-model homogeneous and for any $N, N' \prec_{\boldsymbol{K}} M$ and $|N|, |N'| < \mu$, every isomorphism f from N to N' extends to an automorphism of M.*

3. M is strongly model homogeneous *if it is strongly $|M|$-model homogeneous.*

To emphasize, this differs from the homogeneous context (Definition 5.33) because the N must be *in* $\boldsymbol{K}$. Note that if M is μ-model homogenous, it embeds every model in $\boldsymbol{K}$ of cardinality $\leq \mu$. It is easy to show:

Lemma 9.2. *If M_1 and M_2 have cardinality μ and are μ-model homogeneous with $\mu > \mathrm{LS}(\boldsymbol{K})$ then $M_1 \approx M_2$.*

Proof. If M_1 and M_2 have a common submodel N of cardinality $< \mu$, this is an easy back and forth. Now suppose N_1, (N_2) is a small submodel of M_1, (M_2) respectively. By the joint embedding property there is a small common extension N of N_1, N_2 and by model homogeneity N is embedded in both M_1 and M_2. $\square_{9.2}$

Note that in the absence of joint embedding, to get uniqueness we would (as in [139]) have to add to the definition of 'M is model homogeneous' that all models of cardinality $< \mu$ are embedded in M.

Exercise 9.3. *Suppose M is μ-model homogeneous with cardinality μ, $N_0, N_1, N_2 \in \boldsymbol{K}$ with $N_0 \prec N_1, N_2 \prec M$, for each i, $|N_i| < \mu$, and f is an isomorphism between N_1 and N_2 over N_0. Then f extends to an automorphism of M.*

An easy induction shows:

Theorem 9.4. *If $\mu{*}^{<\mu*} = \mu*$ (which implies $\mu*$ is regular) and $\mu* \geq 2^{\mathrm{LS}(\boldsymbol{K})}$ then there is a model $\mathbb{M}$ of cardinality $\mu*$ which is strongly model homogeneous and in particular model homogeneous.*

We call the model $\mathbb{M}$ constructed in Theorem 9.4, the *monster* model. From now on all, structures considered are substructures of $\mathbb{M}$ with cardinality $< \mu*$. The standard arguments for the use of a monster model in first order model theory ([61, 32] apply here. Although the existence of a monster model as described above requires some extension of set theory (e.g. the existence of arbitrarily large strongly inaccessible cardinals or the GCH), all the arguments can be reworked to avoid this assumption. One method of reworking is to show any model M can be embedded in a strongly $|M|$-homogeneous model. This construction is fairly routine but the result is so basic to later discussions that we give a lengthy sketch of the argument in the following exercise.

Exercise 9.5. *Recall that we assume $\boldsymbol{K}$ has the amalgamation property. Show that for each λ, and each model M of size λ there is a strongly λ-model homogeneous model N containing M. As in Proposition 2.2.7 of [32] follow the next few steps.*

(0) Without loss of generality, we can apply Exercise 9.3 iteratively to replace M by a possibly larger model which is λ-model homogeneous to guarantee the first clause of strong λ-model homogeneity.

*(1) Given **M** and f mapping M_1 into M, with M_1 a substructure of M of size λ, there is an N extending M and g extending f such that $g : N \mapsto N$.*

(That N is the domain of g is not a misprint; construct g and N via an ω-chain using amalgamation as formulated in Exercise 5.11.)

(2) Given M and $f : M_1 \mapsto M$, with $M_1 \prec_{\boldsymbol{K}} M$ of size λ, there is N extending M and g an automorphism of N extending f.

This is a back-and-forth ω-chain using (1)

(3) Given M there is N such that any $f : M_1 \mapsto M$ of size λ extends to an automorphism of N. Enumerate all such functions f and apply (2) repeatedly.

(4) The full result is then obtained by an ω-chain from (3).

We now define the notion of a Galois type; the most general definition is an equivalence relation on triples (M, a, N) where M is a base model and $a \in N - M$. We will quickly show that in the presence of amalgamation, the classes of this equivalence relation are represented as orbits of subgroups of the automorphism group of the monster model. We use our own notation but the relation to that of Shelah in [133, 128] should be clear.

Definition 9.6. *1. For $M \prec_{\boldsymbol{K}} N_1 \in \boldsymbol{K}$, $M \prec_{\boldsymbol{K}} N_2 \in \boldsymbol{K}$ and $a \in N_1 - M$, $b \in N_2 - M$, write $(M, a, N_1) \sim_{AT} (M, b, N_2)$ if there exist strong embeddings f_1, f_2 of N_1, N_2 into some N^* which agree on M and with $f_1(a) = f_2(b)$.*

2. Let $\sim$ be the transitive closure of $\sim_{AT}$ (as a binary relation on triples).

3. We say the Galois type a over M in N_1 is the same as the Galois type a over M in N_2 if $(M, a, N_1) \sim (M, b, N_2)$.

Exercise 9.7. *If $\boldsymbol{K}$ has amalgamation, $\sim_{AT}$ is an equivalence relation and $\sim = \sim_{AT}$.*

And now the homogeneity of the monster model yields a more concrete representation.

Lemma 9.8. *Suppose $\boldsymbol{K}$ has amalgamation and joint embedding. $(M, a, N_1) \sim (M, b, N_2)$ if and only if there are embeddings g_1 and g_2 of N_1, N_2 into $\mathbb{M}$ that agree on M and such that for some $\alpha \in \mathrm{aut}_{\mathrm{M}}(\mathbb{M})$, $\alpha(g_1(a)) = g_2(b)$.*

Proof: For the non-trivial direction, apply strong $|M|$-model homogeneity to $g_1^{-1} \circ g_2$. $\square_{9.8}$

Lemma 9.8 justifies the following equivalent definition of Galois type (when $\boldsymbol{K}$ has amalgamation and joint embedding).

Definition 9.9. *Let $M \in \boldsymbol{K}$, $M \prec_{\boldsymbol{K}} \mathbb{M}$ and $a \in \mathbb{M}$. The* Galois type *of a over M ($\in \mathbb{M}$) is the orbit of a under the automorphisms of $\mathbb{M}$ which fix M. Denote it by $\mathrm{tp}(a/M)$.*

We freely use the phrase, 'Galois type of a over M', dropping the ($\in \mathbb{M}$) since $\mathbb{M}$ is fixed and any M and a are contained in $\mathbb{M}$. We usually use Galois type in the sense of Definition 9.9. But examples are sometimes easier to verify with Definition 9.6; we study the situation when amalgamation has not been assumed in Chapter 17.

Definition 9.10. *The set of Galois types over M is denoted $\mathbb{S}(M)$.*

Note that since $\mathbb{M}$ is homogeneous, $\mathbb{S}(M)$ depends only on the isomorphism type of M and not on the particular strong embedding of M into $\mathbb{M}$.

Definition 9.11. *Let $M \subseteq N \subset \mathbb{M}$ and $a \in \mathbb{M}$. The restriction of $\mathrm{tp}(a/N)$ to M, denoted $\mathrm{tp}(a/N) \restriction M$ is the orbit of a under $\mathrm{aut}_{\mathrm{M}}(\mathbb{M})$.*

We say a Galois type p over M is realized in N with $M \prec_{\boldsymbol{K}} N \prec_{\boldsymbol{K}} \mathbb{M}$ if $p \cap N \neq \emptyset$.

Definition 9.12. *The model M is μ-*Galois saturated *if for every $N \prec_{\boldsymbol{K}} M$ with $|N| < \mu$ and every Galois type p over N, p is realized in M.*

The following model-homogeneity=saturativity theorem was announced with an incomplete proof in [140]. Full proofs are given in Theorem 6.7 of [48] and .26 of [133]. Here, we give a simpler argument making full use of the amalgamation hypothesis. In Chapter 17, we discuss what can be done with weaker amalgamation hypotheses.

Theorem 9.13. *For $\lambda > \mathrm{LS}(\boldsymbol{K})$, The model M is λ-Galois saturated if and only if it is λ-model homogeneous.*

Proof. It is obvious that λ-model homogeneous implies λ-Galois saturated. Let $M \prec_{\boldsymbol{K}} \mathbb{M}$ be λ-saturated. We want to show M is λ-model homogeneous. So fix $M_0 \prec_{\boldsymbol{K}} M$ and N with $M_0 \prec_{\boldsymbol{K}} N \prec_{\boldsymbol{K}} \mathbb{M}$. Say, $|N| = \mu < \lambda$. We must construct an embedding of N into M over M_0. Enumerate $N - M_0$ as $\langle a_i : i < \mu\rangle$. We will define f_i for $i < \mu$ an increasing continuous sequence of maps with domain N_i and range M_i so that $M_0 \prec_{\boldsymbol{K}} N_i \prec_{\boldsymbol{K}} \mathbb{M}$, $M_0 \prec_{\boldsymbol{K}} M_i \prec_{\boldsymbol{K}} M$ and $a_i \in N_{i+1}$. The restriction of $\bigcup_{i<\mu} f_i$ to N is the required embedding. Let $N_0 = M_0$ and f_0 the identity. Suppose f_i has been defined. Choose the least j such that $a_j \in N - N_i$. By the model homogeneity of $\mathbb{M}$, f_i extends to an automorphism $\hat{f}_i$ of $\mathbb{M}$. Using the saturation, let $b_j \in M$ realize the Galois type of $\hat{f}_i(a_j)$ over M_i. So there is an $\alpha \in \operatorname{aut}\mathbb{M}$ which fixes M_i and takes b_j to $\hat{f}_i(a_j)$. Choose $M_{i+1} \prec_{\boldsymbol{K}} M$ with cardinality μ and containing $M_i b_j$. Now $\hat{f}_i^{-1} \circ \alpha$ maps M_i to N_i and b_j to a_j. Let $N_{i+1} = \hat{f}_i^{-1} \circ \alpha(M_{i+1})$ and define f_{i+1} as the restriction of $\alpha^{-1} \circ \hat{f}_i$ to N_{i+1}. Then f_{i+1} is as required. $\square_{9.13}$

Remark 9.14. Analysis of the preceding argument shows that the notion of Galois type has been defined precisely so that assuming amalgamation, Galois-saturated models of the same cardinality are isomorphic. If one has amalgamation over sets, then one can formulate this condition in terms of first order types– the classical notion of homogeneity. But once the domains are restricted to models, trying to prove with no further hypotheses that models saturated for such types are isomorphic leads one (or at least led Shelah) to the current formulation of Galois types.

The identification in Theorem 9.13 allows us to develop many properties of Galois-saturated models similar to those in first order model theory. In particular they are homogenous with respect to Galois types over models. David Kueker noted the necessity for later arguments of making these points explicit and provided the proof below of Corollary 9.16.

Exercise 9.15. *Suppose M is μ-Galois-saturated.*

1. *Every $N \in \boldsymbol{K}$ of cardinality at most μ can be strongly embedded in M.*

2. *If $M_0 \prec_{\boldsymbol{K}} M$, $M_0 \prec_{\boldsymbol{K}} N$, $|M_0| < \mu$, and $|N| \leq \mu$ there is a strong embedding of N into M over M_0.*

Corollary 9.16. *Let M be Galois- saturated,* $\mathrm{LS}(\boldsymbol{K}) < |\mathrm{M}|$. *Let $M' \prec_{\boldsymbol{K}} M$, with* $\mathrm{LS}(\boldsymbol{K}) < |\mathrm{M}'| < |\mathrm{M}|$. *Assume $a, b \in M$ have the same galois-type over M'. Then there is an automorphism of M fixing M' and taking a to b.*

Proof. Using Definition 9.9, let h be an automorphism of the monster fixing M' and taking a to b. Pick M'' such that $M' \prec_{\boldsymbol{K}} M'' \prec_{\boldsymbol{K}} M$, $a \in M''$, and $|M''| < |M|$. Let $N = h[M'']$. (Note that $b \in N$ and $M' \prec_{\boldsymbol{K}} N$.) It suffices to show that there is a $\boldsymbol{K}$-embedding g of N into M fixing M' and b. As, then we can let f be the restriction to M'' of the composition of h followed by g; f fixes M' and maps a to b, and by Exercise 9.3 f extends to the desired automorphism of M. Now, to get g, pick M^* such that $M' \prec_{\boldsymbol{K}} M^* \prec_{\boldsymbol{K}} M$, $(M \cap N) \subseteq M^*$, $|M^*| = |N|$ and pick N^* such that $N \prec_{\boldsymbol{K}} N^*$, $M^* \subseteq N^*$, and $|N^*| = |N|$. Then the proof of Theorem 9.13 yields g^* embedding N^* into M over M^* whose restriction to N is the required g. $\square_{9.16}$

We turn now to defining Galois stability and deriving stability from categoricity. Using the Presentation Theorem and Morley's omitting types theorem, we can find Skolem models over sets of indiscernibles in an AEC. Recall Notation 7.2.1.

Theorem 9.17. *If $\boldsymbol{K}$ is an abstract elementary class in the vocabulary τ, which is represented as a $PC\Gamma$ class witnessed by τ', T', Γ that has arbitrarily large models, there is a τ'-diagram Φ such that for every linear order $(I, <)$ there is a τ'-structure $M = EM(I, \Phi)$ such that:*

1. *$M \models T'$.*

2. *I is a set of τ'-indiscernibles in M.*

3. *$M \restriction \tau$ is in $\boldsymbol{K}$.*

4. *If $I' \subset I$ then $EM_\tau(I', \Phi) \prec_{\boldsymbol{K}} EM_\tau(I, \Phi)$.*

Proof. The first three clauses are a direct application of Lemma A.3.1 , Morley's theorem on omitting types. See also problem 7.2.5 of Chang-Keisler [34]. It is automatic that $EM(I', \Phi)$ is an τ'-substructure of $EM(I, \Phi)$. The moreover clause of Theorem 5.14 allows us to extend this to $EM_\tau(I', \Phi) \prec_{\boldsymbol{K}} EM_\tau(I, \Phi)$. $\square_{9.17}$

Definition 9.18. 1. *Let $N \subset \mathbb{M}$. N is λ-*Galois-stable *if for every $M \subset N$ with cardinality λ, only λ Galois-types over M are realized in N.*

2. *$\boldsymbol{K}$ is λ-Galois-stable if $\mathbb{M}$ is. That is $\mathrm{aut}_M(\mathbb{M})$ has only λ orbits for every $M \subset \mathbb{M}$ with cardinality λ.*

Since we are usually working in an AEC, we will frequently abuse notation and write stable rather than Galois-stable. This definition does not depend on a choice of monster model. For a given λ any pair of strongly λ'-model homogeneous models with $\lambda' > \lambda$ will agree on the number of λ-Galois types.

Claim 9.19. *If $\boldsymbol{K}$ is λ-categorical, the model M with $|M| = \lambda$ is σ-Galois stable for every $\sigma < \lambda$.*

Proof. Represent M as $EM(\lambda, \Phi)$. For any $\mu < \lambda$ and any submodel N of size μ, we show M realizes only μ-types over N. For, $N \subset EM(K, \Phi)$ for some K, and any $a \in M$ is given as a term $\sigma(\mathbf{k}, \mathbf{j})$. There are only μ choices for the isomorphism type of $K\mathbf{j}$ because λ is well ordered. Suppose $\mathbf{j}$ and $\mathbf{j}'$ have the same order type over K, $a = \sigma(\mathbf{k}, \mathbf{j})$ and $a' = \sigma(\mathbf{k}, \mathbf{j}')$. Then

$$EM(K, \Phi), a, EM(K\mathbf{j}, \Phi) \sim EM(K, \Phi), a', EM(K\mathbf{j}', \Phi).$$

By Definition 9.6, a and a' realize the same Galois type over N. $\square_{9.19}$

Theorem 9.20. *If $\boldsymbol{K}$ is categorical in λ, then $\boldsymbol{K}$ is σ-Galois-stable for every σ with $\mathrm{LS}(\boldsymbol{K}) \le \sigma < \lambda$.*

Proof. Suppose $\boldsymbol{K}$ is not σ-stable for some $\sigma < \lambda$. Then by Löwenheim-Skolem and amalgamation, there is a model N of cardinality σ^+ which is not σ-stable. Let M be the σ-stable model with cardinality λ constructed in Claim 9.19. Categoricity and joint embedding imply N can be embedded in M. The resulting contradiction proves the theorem. $\square_{9.20}$

Remark 9.21. *We don't need the full assumption that $\boldsymbol{K}$ has amalgamation; it would suffice to assume amalgamation on $\boldsymbol{K}_{<\lambda}$.*

While we state the following result for $\lambda \ge \mathrm{LS}(\boldsymbol{K})$, note that in the Löwenheim number, saturation may be a trivial notion (if there are no smaller models).

Corollary 9.22. *Suppose $\boldsymbol{K}$ is stable in $\lambda \ge \mathrm{LS}(\boldsymbol{K})$.*

1. *If λ is regular, there is a saturated model of power λ and so model homogeneous.*
2. *For general λ, there is $\mathrm{cf}(\lambda)$-saturated model of power λ.*

In particular, if $\boldsymbol{K}$ is categorical in a regular κ, the categoricity model is saturated.

Proof. Choose in $M_i \prec_{\boldsymbol{K}} \mathbb{M}$ using λ-stability and Löwenheim-Skolem, for $i < \lambda$ so that each M_i has cardinality $\le \lambda$ and M_{i+1} realizes all types over M_i. By regularity, it is easy to check that M_λ is saturated; a standard modification does part 2. $\square_{9.22}$

This result is from [128]; a different argument is provided in [10].

10
Brimful Models

We continue to assume the amalgamation property and the existence of arbitrarily large models and thus can work in the monster model. We introduce the notion of a brimful model and prove some elementary properties of such models. They will play an important role in Chapters 11, 13, and 16

Since we deal with reducts, we will consider several structures with the same universe; it is crucial to keep the vocabulary of the structure in mind. The AEC under consideration has vocabulary τ; it is presented as reducts of models of theory T' (which omit certain types) in a vocabulary τ'. In addition, we have the class of linear orderings (LO) in the background.

We really have three AEC's: $(LO, \subset)$, $(\boldsymbol{K}', \prec_{\boldsymbol{K}'})$, which is the models of T' with $\prec_{\boldsymbol{K}'}$ as τ'-closed subset, and $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$. We are describing the properties of the EM-functor between $(LO, \subset)$ and $\boldsymbol{K}'$ or $\boldsymbol{K}$. $\boldsymbol{K}'$ is only a tool that we are singling out to see the steps in the argument. The following definitions hold for any of the three classes and I write $\leq$ for the notion of substructure. In this chapter I am careful to use $\leq$ when discussing all three cases versus $\prec_{\boldsymbol{K}}$ for the AEC.

Definition 10.1. *M_2 is (relatively) σ-*universal over *M_1* in *N if $M_1 \leq M_2 \leq N$ and whenever $M_1 \leq M_2' \leq N$, with $|M_1| \leq |M_2'| \leq \sigma$, there is a $\leq$-embedding fixing M_1 and taking M_2' into M_2.*

Remark 10.2. *The restriction to $M_2' \leq N$ is crucial. We introduce in Definition 11.4 a stronger notion, σ-universal, where this restriction is omitted. The notions are equivalent if N is saturated. But we use the weaker notion as a tool to establish stability and to find saturated models.*

I introduce one term for shorthand. It is related to Shelah's notion of *brimmed* in [134] but here the brimful model is bigger than the models it is universal over while brimmed models may have the same cardinality. (Brimful comes from the English expression, 'brim full'; it refers to a bucket which is completely full (to the brim).)

Definition 10.3. *M is* brimful *if for every $\sigma < |M|$, and every $M_1 \leq M$ with $|M_1| = \sigma$, there is an $M_2 \leq M$ with cardinality σ that is σ-universal over M_1 in M.*

The next notion makes it easier to write the proof of the lemma that follows it.

Notation 10.4. *Let $J \subset I$ be linear orders. We say a and b in I realize the same cut over J and write $a \sim_J b$ if for every $j \in J$, $a < j$ if and only if $b < j$.*

Exercise 10.5. *Prove the order type (*otp*) of the lexicographic linear order on $I = \lambda^{<\omega}$ is neither dense nor well-founded.*

Claim 10.6 (Lemma 3.7 of [88])**.** *The lexicographic linear order on $I = \lambda^{<\omega}$ is brimful.*

Proof. Let $J \subset I$ have cardinality $\theta < \lambda$. Without loss of generality we can assume $J = A^{<\omega}$ for some $A \subset \lambda$. For any $\sigma \in I$ let σ^* be $\sigma \restriction n$ for the least n such that $\sigma \sim_J \sigma \restriction n$. Then for any τ, $\sigma \sim_J \tau$ if and only if $\sigma^* \sim_J \tau^*$. Then the cut of A is determined by $\sigma \restriction (n-1)$ and the place of $\sigma(n-1)$ in A. Thus there are only θ cuts over J realized in I. For each cut C_α, $\alpha < \theta$, we choose a representative $\sigma_\alpha \in I - J$ of length n such that $\sigma_\alpha \restriction (n-1) \in J$, so C_α contains $\{\sigma_\alpha\widehat{\ }\tau : \tau \in \lambda^{<\omega}\}$. We can assume any J^* extending J is $J^* = B^{<\omega}$ for some $B \subset \lambda$, say with $\mathrm{otp}(B) = \gamma < \theta^+$. Thus, the intersection of J^* with a cut in J is isomorphic to a subset of $\gamma^{<\omega}$. We finish by noting for any ordinal with $|\gamma| = \theta$, $\gamma^{<\omega}$ can be embedded in $\theta^{<\omega}$. Thus, the required θ-universal set over J is $J \cup \{\sigma_\alpha\widehat{\ }\tau : \tau \in \theta^{<\omega}, \alpha < \theta\}$.

Qing Zhang has provided the following elegant argument for the last claim. First show by induction on γ there is a map g embedding γ in $\theta^{<\omega}$. (E.g. if $\gamma = \lim_{i<\theta} \gamma_i$, and g_i maps γ_i into $\theta^{<\omega}$, let for $\beta < \gamma$, $g(\beta) = i\widehat{\ }g_i(\beta)$ where $\gamma_i \leq \beta < \gamma_{i+1}$.) Then let h map $\gamma^{<\omega}$ into $\theta^{<\omega}$ by, for $\sigma \in \gamma^{<\omega}$ of length n, setting $h(\sigma) = \langle g(\sigma(0)), \ldots, g(\sigma(n-1))\rangle$. $\square_{10.6}$

The argument for Claim 10.6 yields:

Corollary 10.7. *Suppose $\mu < \lambda$ are cardinals. Then for any $X \subset \mu^{<\omega}$ and any Y with $X \subseteq Y \subset \lambda^{<\omega}$ and $|X| = |Y| < \mu$, there is an order embedding of Y into $\mu^{<\omega}$ over X.*

Exercise 10.8. *For an ordinal γ, let $\gamma^{\omega *}$ denote the functions from ω to γ with only finitely many non-zero values ([116]). Show $\gamma^{\omega *}$ is a dense linear order and so is not isomorphic to $\gamma^{<\omega}$. Vary the proof above to show $\gamma^{\omega *}$ is brimful.*

Exercise 10.9. *Show that every decreasing chain in $\lambda^{<\omega}$ is countable.*

Since every τ'-substructure N of $EM(I, \Phi)$ is contained in a substructure $EM(I_0, \Phi)$ for some subset I_0 of I with $|I_0| = |N|$, we have immediately:

Claim 10.10. *If I is brimful as a linear order, $EM(I, \Phi)$ is brimful as an τ'-structure.*

Now there are some subtle uses of the 'coherence axiom': $M \subseteq N \prec_{\boldsymbol{K}} N_1$ and $M \prec_{\boldsymbol{K}} N_1$ implies $M \prec_{\boldsymbol{K}} N$.

Claim 10.11. *If I is brimful as linear order, $EM_\tau(I, \Phi)$ is brimful as a member of $\boldsymbol{K}$.*

Proof. Let $M = EM(I, \Phi)$; we must show $M \restriction \tau$ is brimful as a member of $\boldsymbol{K}$. Suppose $M_1 \prec_{\boldsymbol{K}} M \restriction \tau$ with $|M_1| = \sigma < |M|$. Then there is $N_1 = EM(I', \Phi)$ with $|I'| = \sigma$ and $M_1 \subseteq N_1 \leq M$. By Theorem 9.17.5, $N_1 \restriction \tau \prec_{\boldsymbol{K}} M \restriction \tau$. So $M_1 \prec_{\boldsymbol{K}} N_1 \restriction \tau$ by the coherence axiom. Let M_2 have cardinality σ and $M_1 \prec_{\boldsymbol{K}} M_2 \prec_{\boldsymbol{K}} M \restriction \tau$. Choose a τ'-substructure N_2 of M with cardinality σ containing N_1 and M_2. Now, N_2 can be embedded by a map f into the σ-universal τ'-structure N_3 containing N_1 which is guaranteed by Claim 10.10. But $f(N_2) \restriction \tau \prec_{\boldsymbol{K}} N_3 \restriction \tau$ by the coherence axiom so $N_3 \restriction \tau$ is the required σ-universal extension of M_1. $\square_{10.11}$

Note that this argument does not address extensions M_2 of M_1 unless they are already embedded in $M \restriction \tau$; we weaken that hypothesis in the next chapter.

11
Special, limit and saturated models

In this chapter we consider the relations among special, limit, saturated and Ehrenfeucht-Mostowski models. Making heavy use of the Ehrenfeucht-Mostowski models we establish the uniqueness of limit models in μ (though not over their base) if $\mathrm{LS}(\boldsymbol{K}) < \mu < \lambda$ and λ is a regular categoricity cardinal. These are the key tools for erecting a theory of superstability in Chapter 16.

Even for stable first order theories there are no saturated models in cardinalities with cofinality $\kappa(T)$. The notion of a special model as a union of saturated models was developed in [112] (see also [34, 61]) to address the need to use GCH to find saturated models. For regular cardinals this problem is easily solved for stable theories but special models provide a substitute for saturation in limit cardinals. For stable AEC with amalgamation limit models, [139, 141, 144] are an analog to saturation. If $\mathrm{LS}(\boldsymbol{K}) < \mu < \lambda$ and λ is a regular categoricity cardinal, we will see that limit models in μ are saturated and thus unique. Note however, that even in the case of a countable first order stable theory; when $\lambda^{\omega} = \lambda$, there are saturated models but there are limit models which are not saturated. A saturated model M is $|M|$-homogeneous but not $|M|^{+}$-homogeneous; in even the simplest situations, a saturated model can be embedded in itself in various ways. For example, if T is the first order theory of successor on the integers, a model consists of many copies of the integers. The theory is categorical in every uncountable power. Let M be the model of power $\aleph_1$; it can be embedded in itself with any 'codimension' up to $\aleph_1$. Limit models will allow us to capture some of the properties of such an embedding with maximal codimension without introducing a notion of dimension of types.

Finally, the notion of saturated model is useless for studying models of cardinality $\mathrm{LS}(\boldsymbol{K})$; every model may be trivially saturated (if there are no smaller models). Limit models provide a substitute for saturation in $\mathrm{LS}(\boldsymbol{K})$. We begin with the definition of a special model and develop some properties of special models.

Definition 11.1. *1. Let $|N| = |M| = \mu$. For any ordinal α, we say N is an α-*special extension *of M if $N = \bigcup M_i$ where $\langle M_i : i < \mu \times \alpha \rangle$ is a continuous increasing chain with $M_0 = M$, and M_{i+1} is a strong extension of M_i which realizes all 1-Galois types over M_i.*

2. We say M is λ-special where λ is a singular cardinal if M is union of a continuous increasing chain $\langle M_i : i < \lambda \rangle$ so that each M_i has cardinality $< \lambda$ and M_{i+1} realizes all types over M_i.

Shelah [128] writes $M <^1_{\mu,\alpha} N$ for our n is α-special over M. He also writes $M <^0_{\mu,\alpha} N'$ if $M \leq^0_{\mu,\alpha} N$ and $N \prec_{\boldsymbol{K}} N'$. We noted in Theorem 9.22 that if $\boldsymbol{K}$ is categorical in a regular cardinal, the categoricity model is saturated; for singular $\lambda > \mathrm{LS}(\boldsymbol{K})$ we have the obvious weakening.

Corollary 11.2. *Suppose $\boldsymbol{K}$ is categorical in λ and λ is singular. The model of power λ is λ-special.*

Proof. Choose in $M_i \prec_{\boldsymbol{K}} \mathbb{M}$ using $< \lambda$-stability and Löwenheim-Skolem, for $i < \lambda$ so that each M_i has cardinality $< \lambda$ and M_{i+1} realizes all types over M_i. $\square_{11.2}$

Using our standard amalgamation hypothesis, it is routine to show:

Lemma 11.3. *If $\boldsymbol{K}$ is stable in μ and $\alpha < \mu^+$, every M of cardinality μ has an α-special extension.*

To study limit models we replace the relative notion of σ-universality from Chapter 10 by the following unrestricted notion of σ-universality, which was also introduced by Shelah (cf. [128]).

Definition 11.4. *Let $M_1, M_2 \in \boldsymbol{K}$. M_2 is σ-*universal over *M_1 if $M_1 \prec_{\boldsymbol{K}} M_2$ and whenever $M_1 \prec_{\boldsymbol{K}} M_2'$, with $|M_1| \leq |M_2'| \leq \sigma$, there is a $\boldsymbol{K}$-embedding fixing M_1 and taking M_2' into M_2.*

The uniqueness of α-special models (Lemma 11.5) is a bit more intricate than their existence. The argument is a back and forth, similar to Theorem 9.13; the chain has length $\mu \times \alpha$ precisely to give enough space for this back and forth. Lemma 11.5.1 is asserted without proof in 1.15 of [134]; another exposition of the result is in [45].

Lemma 11.5. *1. If $M \in \boldsymbol{K}_\kappa$ and M^1 is 1-special over M, M^1 is κ-universal over M.*

2. If α is a limit ordinal and N and N' are both α-special over M then they are isomorphic over M.

Proof. Part 1). As M^1 is 1-special over M, Definition 11.1 tells us M^1 is a continuous union for $i < \kappa$ of M_i with $M_0 = M$, and each M_{i+1} realizes all Galois types over M_i. Now fix any strong extension N of M with $|N| = \kappa$. We will construct a $\boldsymbol{K}$-isomorphism f from M^1 into an extension N^1 of N with $N \subset \overline{N} \prec_{\boldsymbol{K}} N^1$, where $\overline{N}$ denotes the range of f. By the coherence axiom $f^{-1} \restriction N$ is the required map.

To construct f, enumerate $N - M$ as $\langle a_i : i < \kappa \rangle$. We construct a continuous increasing sequence of maps f_i. Let $f_0 = 1_M$. Suppose we have defined f_i, N_i and f_i taking M_i onto N_i. (Note N_i may not be a submodel of N.) Let j be least with $a_j \notin N_i$. By the definition of special $f_i^{-1}(\mathrm{tp}(a_j/N_i)) \in \mathbb{S}(M_i)$ is realized by some $b \in M_{i+1}$. So there is an automorphism g of $\mathbb{M}$ extending f_i and mapping b to a_j. Let f_{i+1} be $g \restriction M_{i+1}$ and $N_{i+1} = f(M_{i+1})$. Finally f is the union of the f_i and N^1 is the union of the N_i'.

For part 2), use the universality to construct a back and forth mapping M_γ into $N_{\gamma+1}$ for even successor γ and vice versa for odd γ. $\Box_{11.5}$

Now we introduce limit models as in [139, 141, 144]. The following notion is most useful when δ is a limit ordinal.

Definition 11.6. *Suppose $\overline{M} = \langle M_i : i < \delta \rangle$ is a sequence of models of cardinality μ, and $\delta \leq \mu^+$ with M_{i+1} μ-universal over M_i. We call such a sequence a (μ, δ)-chain and, when δ is a limit ordinal, M_δ a (μ, δ)-limit model.*

Note a simple back and forth argument shows that the isomorphism type of a limit model over M depends only on the cofinality of δ. In particular,

Lemma 11.7. *If N and N' are respectively a δ and a $|M| \times \delta$ limit model over M, where δ is a limit ordinal less than μ^+, then N and N' are isomorphic over M.*

We will show quickly that limit models are the same as special models and two towers of the same limit length are isomorphic. More strongly we will show that for μ less than the categoricity cardinal if $\alpha, \alpha' < \mu^+$ are limit ordinals then an α-limit and an α'-limit over a saturated M are isomorphic, although our argument does not show isomorphism over M. Our proof proceeds by showing they can both be represented as EM models and giving a simple test for the isomorphism of EM-models.

Note that Theorem 11.8 applies with $\kappa = \mathrm{LS}(\boldsymbol{K})$; later results involving saturation are proved only above the Löwenheim number.

Theorem 11.8. *Let $\boldsymbol{K}$ be κ-Galois stable with the amalgamation property. Suppose $|M| = \kappa$ and $\alpha \leq \kappa^+$.*

1. *An α-special model over M is an $(|M|, \alpha)$-limit model.*
2. *If N and N' are α-special and $(|M|, |M| \times \alpha)$-limit over M respectively, then they are isomorphic over M.*
3. *If α is a limit ordinal and N is α-special over both M and M' and f is an isomorphism from M onto M', then f extends to an automorphism of N.*

Proof. Part 1) is an easy induction from Lemma 11.5.1 and Part 2) follows from Lemma 11.5 and Lemma 11.7. Part 3) is just a translation of Lemma 11.5.2). $\square_{11.8}$

Now we show that we can also represent limit models as Ehrenfeucht-Mostowski models and this will enable us to show the uniqueness of limit models.

Notation 11.9. *For any diagram Ψ in the vocabulary τ with Ehrenfeucht-Mostowski τ-models $EM(I,\Psi)$ and any $\tau' \subseteq \tau$, $EM_{\tau'}(I,\Psi)$ denotes the reduct of $EM(I,\Phi)$ to τ'. Let τ_0 be the vocabulary of $\boldsymbol{K}$. Fix the EM-diagram Φ as in Theorem 9.17 and Definition 7.2.1.*

We need one key technical fact asserted by Shelah [128] and given a clear proof by Hyttinen; it provides further sufficient conditions for finding saturated models.

Lemma 11.10. *Suppose Φ satisfies for every linear order I, $EM_{\tau_0}(I,\Phi) \in \boldsymbol{K}$. If $\boldsymbol{K}$ is λ-categorical for a cardinal λ then for every J with $|\tau(\Phi)| \leq \mathrm{LS}(\boldsymbol{K}) < |\mathrm{J}| < \mathrm{cf}(\lambda)$ such that for each $\theta < |J|$, J contains an increasing sequence of cardinality θ^+, $M = EM_{\tau_0}(J,\Phi)$ is Galois-saturated.*

Proof. For each $\theta < |J|$, we prove M is θ^+-saturated. Choose an increasing sequence J_0 of J of length θ^+ and form J' by inserting a copy of λ immediately after the sequence. By the categoricity and since $\theta^+ \leq \mathrm{cf}(\lambda)$, applying Corollary 9.22.2, $N = EM_{\tau_0}(J',\Phi)$ is θ^+-saturated. Now let $M_0 \prec_{\boldsymbol{K}} M$ have cardinality θ and let $q \in \mathbb{S}(M_0)$. Then q is realized in N by some $\rho(\mathbf{j},\mathbf{j}')$ where $\mathbf{j} \in J$ and $\mathbf{j}'$ is in the new copy of λ. Choose $K \subset J$ with $|K| = \theta$ and $M_0\mathbf{j} \subseteq EM_{\tau_0}(K,\Phi)$. Since J_0 has length θ^+, we can find $\mathbf{j}'' \in J_0$ such that $K\mathbf{j}\mathbf{j}'$ and $K\mathbf{j}\mathbf{j}''$ have the same order type. Thus $\rho(\mathbf{j},\mathbf{j}'') \in M$ realizes q. $\square_{11.10}$

The following slightly more general result is stated in [129].

Exercise 11.11. *Show the same result holds if we assume that J has descending sequences of length θ^+ for each θ.*

Exercise 11.12. *Show the argument generalizes to show that if J is a linear order that contains an increasing sequence of cardinality θ^+ for each $\theta < \mu$ for some $\mu \leq |J|$, then $M = EM_{\tau_0}(J,\Phi)$ is μ-Galois-saturated.*

Here is an immediate application of Lemma 11.10. Many results in the next few chapters have the hypothesis that $\boldsymbol{K}$ is categorical in a *regular* cardinal λ. The role of the regularity is to guarantee that the unique model in the categoricity cardinal is saturated (Lemma 9.22).

Corollary 11.13. *Suppose $\boldsymbol{K}$ is an AEC with vocabulary τ that is categorical in λ and λ is regular.*

1. *If $\mathrm{LS}(\boldsymbol{K}) < \mu \leq \lambda$ then $M_\mu = EM_\tau(\mu^{<\omega},\Phi)$ is saturated and, in particular, it is μ-model homogeneous.*

2. *If $EM_\tau(I_0, \Phi)$ is μ-saturated for some ordering I_0 that satisfies the hypotheses of Lemma 11.10 and I'_0 is an extension of I, then $EM_\tau(I'_0, \Phi)$ is μ-saturated.*

3. *Suppose $\boldsymbol{K}$ is λ-categorical for regular λ. If κ is a cardinal with $\kappa \leq \lambda$, then $EM_\tau(\kappa, \Phi)$ is saturated.*

Proof. Just note that each of the orderings satisfies the conditions of Lemma 11.10 or Exercise 11.12 . $\square_{11.13}$

In [10], we proved the existence of saturated models in limit cardinals below a regular categoricity cardinal using the brimful models studied in Chapter 10. Now we turn to the connections between limit models and EM-models. A straightforward back and forth yields.

Lemma 11.14. *Let δ be a limit ordinal with $\delta < \mu^+$. Suppose $\overline{M}$ and $\overline{N}$ are $(\mu, \mu \times \delta)$-limits with $M_0 \approx N_0$. Then $\overline{M}$ and $\overline{N}$ are term by term isomorphic on the subsequence of models indexed by the $\mu \times i$ for $i < \delta$.*

Now we get a stronger representation of limit models below the categoricity cardinal. Note that here and below, we sometimes state a result as true all $\mu < \lambda$ where λ is regular while the proof yields the stronger statement, true for for arbitrary λ and all $\mu < \mathrm{cf}(\lambda)$; we may quote the stronger version.

Lemma 11.15. *Suppose $\boldsymbol{K}$ is λ-categorical (with λ regular and $\mu < \lambda$). Let δ be a limit ordinal with $\delta < \mu^+$ and $I = \mu^{<\omega}$.*

1. *$\langle EM(I \times \alpha, \Phi) : \alpha < \delta \rangle$ is a (μ, δ)-chain over $EM(I, \Phi)$.*

2. *For any $\alpha < \mu^+$, $EM(I \times \alpha, \Phi) \approx EM(I \times \mu \times \alpha, \Phi)$.*

3. *Every (μ, δ)-chain over $M_0 \approx EM(I, \Phi)$ is isomorphic to $EM(I \times \delta, \Phi)$.*

4. *Finally, if M is δ-special over $M_0 \approx EM(I, \Phi)$, $M \approx EM(I \times \delta, \Phi)$.*

Proof. First we show 1). For each α, $I \times \alpha$ is an initial segment of $I \times \lambda$ so if $X \subset I \times \alpha$, and $X \subset Y \subset I \times \lambda$, with $|Y| = \mu$, $Y - X$ is the disjoint union of a subset Y_0 of $(I \times \alpha) - X$ and a subset Y_1 with cardinality μ of $I \times \lambda$ which lies entirely above $I \times \alpha$. By the universality over the empty set of $\mu^{<\omega}$ (see the proof of Claim 10.6), the second set can be embedded in the $\alpha + 1$st copy of $\mu^{<\omega}$. But since $EM(I \times \lambda, \Phi)$ is saturated (by categoricity) this implies $EM(I \times (\alpha + 1), \Phi)$ is actually μ-universal over $EM(I \times \alpha, \Phi)$. This strengthens the conclusion of Claim 10.11.

Now we consider part 2). For each $i < \mu$, let $I_i = \{i\frown\sigma : \sigma \in \mu^{<\omega}\}$. Since $\mu^{<\omega}$ is isomorphic to I_i by the map $\sigma \mapsto i\frown\sigma$; so $I \approx I \times \mu$. This implies that for any $\alpha < \mu^+$, $EM(I \times \alpha, \Phi) \approx EM(I \times \mu \times \alpha, \Phi)$. Part 3) is now immediate from Lemma 11.14 and Part 1). And by part 3) and part 2) again we see that if M is α-special over M_0, then $M \approx EM(I \times \mu \times \alpha, \Phi) \approx EM(I \times \alpha, \Phi)$. $\square_{11.15}$

We conclude with two very useful applications of this analysis.

Theorem 11.16. *Suppose $\boldsymbol{K}$ is λ-categorical for a regular cardinal λ. If $\mathrm{LS}(\boldsymbol{K}) < \mu < \lambda$, $|M| = \mu$ and M' is a (μ, α) limit model over M where $\alpha < \mu^+$ is a limit ordinal. Then M' is saturated.*

Consequently, N, N' are (μ, α) and (μ, α') limits over M where $\alpha, \alpha' < \mu^+$ are limit ordinals, then N and N' are isomorphic.

Proof. Without loss of generality, invoking Lemma 11.7, we replace α by $\mu \times \alpha$. The first three terms of the sequence are M, M_1, M_2. By Corollary 11.13, there is a saturated model $M_0' \approx EM(I_\mu, \Phi)$ of cardinality μ, which again without loss of generality extends M. Since M_1 is universal over M, we can take $M_0' \prec_{\boldsymbol{K}} M_1$. Moreover, we can assume M_2 is universal over M_0'. (Let N be any strong extension of M_0'. By amalgamation let N' be a (literal) extension of M_1 with N strongly embedded in N'. The universality of M_2 over M_1 gives an embedding of N into M_2 over M_0'. Now we replace the original sequence by one beginning with M_0', M_2; clearly the limit model is the same. By Lemma 11.15.3, M' can be represented as $EM(I \times \alpha, \Phi)$. Since I satisfies the hypothesis of Lemma 11.10, so does $I \times \alpha$. Thus, M' is saturated. We conclude the isomorphism by the uniqueness of saturated models. $\square_{11.16}$

Shelah (6.5.3 of [128]) asserts that any two limit models over a model M_0 are isomorphic over M_0. But the proof in [128] only shows saturation as above.

Proving the uniqueness of limit models (independent of the length of chain) without invoking directly EM-models is a very tricky business; see [141, 144, 47]. These authors do not assume categoricity but assume certain "superstability' conditions. Their proof uses notions of strong types, reduced towers, splitting etc. and does not explicitly invoke EM-models. Moreover, their argument gets isomorphism over M which we do not have here. However, the crucial assumptions can only be deduced from categoricity by the use of even more sophisticated uses of EM-models than those in this chapter; compare with Chapter 16. And we get an easy proof for a special case of: unions of saturated models are saturated.

Theorem 11.17. *Suppose χ is a limit cardinal with $\mathrm{LS}(\boldsymbol{K}) < \chi < \lambda^+$ and $\boldsymbol{K}$ is categorical in λ^+. Then the increasing union of $\delta < \chi^+$ saturated models of cardinality χ is saturated.*

Proof. Note that there are saturated models of cardinality χ by Lemma 11.13. Let $N = \bigcup_{i<\delta} N_i$ where each N_i is χ-saturated. Let $M \prec_{\boldsymbol{K}} N$ have cardinality $\kappa < \chi$ and fix $p \in \mathbb{S}(M)$. We must show p is realized in N. Let $\mu = \kappa^+$. Let $X = \{i < \delta : (N_{i+1} \cap M) - N_i \neq \emptyset\}$. Then $|X| \leq \kappa$ and we can list X as $\langle x_i : i < \delta' \rangle$ where $\delta' < \max(\mu, \mathrm{cf}(\delta)^+) < \chi$. Construct a continuous increasing chain M_γ of μ-saturated models, each with cardinality μ, for $\gamma < \delta'$ so that $(M \cap N_{x_\gamma}) \subseteq M_\gamma \subset N_{x_\gamma}$ and $M_{\gamma+1}$ is μ-universal over M_γ. This is possible by the μ^+-saturation of N_i for each i. If $\kappa < \mathrm{cf}(\delta)$ then p is realized in $M_{\gamma+1}$ for some $\gamma < \delta'$. If $\mathrm{cf}(\delta) \leq \kappa$, $\langle M_\gamma : \gamma < \delta' \rangle$ is a (μ, δ')-chain. Now by Lemma 11.15.3 and Corollary 11.13.2 $M_{\delta'} \prec_{\boldsymbol{K}} N$ is μ-saturated. So p is realized in N and we finish. $\square_{11.17}$

12
Locality and Tameness

In this chapter we discuss some important ways in which Galois types behave differently from 'syntactic types' (types given by sets of formulas). In first order logic or even homogeneous model theory, the union of an increasing chain of types is a type. Moreover, in the same contexts, if two types are different, this difference is witnessed by their restriction to a small subset of the domain. However, the notion of the union of an increasing chain of Galois-types is not well-behaved in general. We can find an upper bound for an ω-chain of types. But it may not be unique. A chain of length κ, with $\kappa > \omega$ may not even have an upper bound. We continue to work in an AEC with amalgamation and $\mathbb{M}$ denotes the monster model.

Tameness is an incredibly important property; all categoricity transfer results proved for AEC in Part III depend on tameness in some way. In Theorem 12.15 we derive a weak-tameness below the categoricity cardinal lemma which is a crucial tool for the downward categoricity transfer in Chapter 15.

In this chapter we explore the variations of 'locality phenomena'. The following argument is subsumed below but for concreteness we separate it out.

Theorem 12.1. *If $M = \bigcup_{i<\omega} M_i$ is an increasing chain of members of $\boldsymbol{K}$ and $\{p_i : i < \omega\}$ satisfies $p_{i+1} \restriction M_i = p_i$, there is a $p_\omega \in \mathbb{S}(M)$ with $p_\omega \restriction M_i = p_i$ for each i.*

Proof. Let a_i realize p_i. By hypothesis, for each $i < \omega$, there exists f_i which fixes M_{i-1} and maps a_i to a_{i-1}. Let g_i be the composition $f_1 \circ \ldots f_i$. Then g_i maps a_i to a_0, fixes M_0 and $g_i \restriction M_{i-1} = g_{i-1} \restriction M_{i-1}$. Let M'_i denote $g_i(M_i)$ and M' their union. Then $\bigcup_{i<\omega} g_i$ is an isomorphism between M and M'. So by model-homogeneity there exists an automorphism h of $\mathbb{M}$ with $h \restriction M_i = g_i \restriction M_i$

for each i. Let $a_\omega = h^{-1}(a_0)$. Now $g_i^{-1} \circ h$ fixes M_i and maps a_ω to a_i for each i. This completes the proof. $\square_{12.1}$

The following condition on an increasing chain allows us to find upper bounds in certain cases.

Definition 12.2. *Let $\langle M_i : i < \gamma \rangle$ be an increasing $\prec_{\boldsymbol{K}}$-chain of submodels of $\mathbb{M}$. A* coherent chain of Galois types *of length γ is an increasing chain of types $p_i \in \mathbb{S}(M_i)$ equipped with realizations a_i of p_i and for $i < j < \gamma$ functions $f_{ij} \in \operatorname{aut}(\mathbb{M})$ such that f_{ij} fixes M_i, $f_{ij}(a_j) = a_i$ and (this is the coherence) for $i < j < k < \gamma$, $f_{ij} \circ f_{jk} = f_{ik}$.*

The following easy result has a number of nice consequences. It could easily be extended to direct limits.

Theorem 12.3. *1. If $p_i \in \mathbb{S}(M_i)$ for $i < \delta$ is a coherent chain of Galois types, there is a $p_\delta \in \mathbb{S}(M_\delta)$ that extends each p_i so that $\langle p_i : i \leq \delta \rangle$ is a coherent sequence.*

2. Conversely, if $p_\delta \in \mathbb{S}(M_\delta)$ extends $p_i \in \mathbb{S}(M_i)$ for $i < \delta$, there is a choice of $f_{i,j}$ for $i \leq j \leq \delta$ that witness $\langle p_i : i \leq \delta \rangle$ is a coherent sequence.

Proof. Note that the commutativity conditions on the f_{ij} imply that the sequence of maps $f_{0j} \restriction M_j$ for $j \leq k < \delta$ is an increasing sequence. I.e.,

$$f_{0k} \restriction M_j = (f_{0j} \restriction M_j) \circ (f_{jk} \restriction M_j) = f_{0j} \restriction M_j.$$

Let $f_\delta = \bigcup_j f_{0j} \restriction M_j$ and let h be an extension of f_δ to an automorphism of $\mathbb{M}$. Choose $h^{-1}(a_0)$ as a_δ and let $p_\delta = \operatorname{tp}(a_\delta / M_\delta)$. Let $f_{i,\delta}$ be $f_{i,0}^{-1} \circ h$. Check that this is the required coherent system.

For the converse, if a_δ realizes p_δ and a_i realizes p_i we have maps $f_{i,\delta}$ which fix M_i and map a_δ to a_i. The required maps $f_{i,j}$ for $i \leq j$ are $f_{i,\delta} \circ f_{j,\delta}^{-1}$. $\square_{12.3}$

Now we introduce two specific notions of locality.

Definition 12.4. *1. $\boldsymbol{K}$ has* (κ, λ)-local galois types *if for every continuous increasing chain $M = \bigcup_{i<\kappa} M_i$ of members of $\boldsymbol{K}$ with $|M| = \lambda$ and for any $p, q \in \mathbb{S}(M)$: if $p \restriction M_i = q \restriction M_i$ for every i then $p = q$.*

2. Galois types are (κ, λ)-compact *in $\boldsymbol{K}$ if for every model M with $|M| = \lambda$ and for every continuous increasing chain $M = \bigcup_{i<\kappa} M_i$ of members of $\boldsymbol{K}$ and every increasing chain $\{p_i : i < \kappa\}$ of members $\mathbb{S}(M_i)$ there is a $p \in \mathbb{S}(M)$ with $p \restriction M_i = p_i$ for every i.*

As usual with such parameterized notions we can write abbreviations such as $(< \kappa, \leq \lambda)$ to mean any cardinality $< \kappa$ in the first coordinate and any cardinality $\leq \lambda$ in the second coordinate. In this notation, Theorem 12.1 says every AEC is $(\aleph_0, \infty)$-compact. Lemma 3.10 of [17] provides an AEC that is not $(\aleph_0, \aleph_0)$-local (see Example 12.10). The following result was stated by Shelah in e.g. [128]; we give a straightforward proof.

Lemma 12.5. *For any λ, if $\boldsymbol{K}$ has $(<\kappa, \leq \lambda)$-local Galois types, then Galois types are $(\leq \kappa, \leq \lambda)$-*compact *in $\boldsymbol{K}$.*

Proof. Consider a continuous increasing chain $M = \bigcup_{i<\kappa} M_i$ of members of $\boldsymbol{K}$ with cardinality at most λ and an increasing chain $\{p_i : i < \kappa\}$ of members $\mathbb{S}(M_i)$. It suffices to find a coherent system representing the chain. Define the a_i and $f_{i,j}$ by induction. At successor stages define $f_{j,j+1}$ as any function fixing M_j with $f_{j,j+1}(a_{j+1}) = a_j$. Define the other $f_{i,j+1}$ by composition to meet the conditions. At limit stages apply Theorem 12.3 to find an a_δ and $f_{i,\delta}$. Note that since the a_δ satisfies each p_i for $i < \delta$, the locality guarantees that it satisfies the given p_δ. $\square_{12.5}$

Just reformulating the notation, Lemma 12.5 implies that if $\boldsymbol{K}$ is $(<\kappa, \infty)$-local then it is (κ, ∞)-compact.

Now we turn to the notion of tameness. The property was first isolated in [128] in the midst of a proof. Grossberg and VanDieren [45] focused attention on the notion as a general property of AEC's. We introduce a parameterized version in hopes of deriving tameness from categoricity by an induction. And weakly tame is the version that can actually be proved (Theorem 12.15).

Definition 12.6.
1. *We say $\boldsymbol{K}$ is (χ, μ)-*weakly tame *if for any* saturated $N \in \boldsymbol{K}$ *with $|N| = \mu$ if $p, q, \in \mathbb{S}(N)$ with $q \neq p$ then for some $N_0 \leq N$ with $|N_0| \leq \chi$, $p \restriction N_0 \neq q \restriction N_0$.*
2. *We say $\boldsymbol{K}$ is (χ, μ)-*tame *if the previous condition holds for all N with cardinality μ.*

Finally, we say $\boldsymbol{K}$ is κ (weakly)-tame if it is (κ, λ)-(weakly)-tame for every λ greater than κ. Note that for any μ, any $\boldsymbol{K}$ is (μ, μ)-tame. There are a few relations between tameness and locality.

Exercise 12.7. *Prove the following algebraic version of tameness. $\boldsymbol{K}$ is (χ, μ)-tame if for any model M of cardinality μ and any $a, b \in \mathbb{M}$:*

If for every $N \prec_{\boldsymbol{K}} M$ with $|N| \leq \chi$ there exists $\alpha \in \mathrm{aut_N}(\mathbb{M})$ with $\alpha(a) = b$, then there exists $\alpha \in \mathrm{aut_M}(\mathbb{M})$ with $\alpha(a) = b$.

Lemma 12.8. *If $\lambda \geq \kappa$ and $\mathrm{cf}(\kappa) > \chi$, then (χ, λ)-tame implies (κ, λ)-local.*

Proof. Suppose $\langle M_i, p_i : i < \kappa \rangle$ is an increasing chain with $\bigcup_i M_i = M$ and $|M| = \lambda$. If both $p, q \in \mathbb{S}(M)$ extend each p_i, by (χ, λ)-tameness, there is a model N of cardinality χ on which they differ. Since $\mathrm{cf}(\kappa) > \chi$, N is contained in some M_i. $\square_{12.8}$

Lemma 12.9. *If $\boldsymbol{K}$ is $(<\mu, <\mu)$-local then M is $(\mathrm{LS}(\boldsymbol{K}), \mu)$-tame.*

Proof. We prove the result by induction on μ and it is clear for $\mu = \mathrm{LS}(\boldsymbol{K})$. Suppose it holds for all $\kappa < \mu$. Let p, q be distinct types in $\mathbb{S}(M)$ where $|M| = \mu$ and write M as an increasing chain $\langle M_i : i < \mu \rangle$ with $|M_i| \leq |i| + \mathrm{LS}(\boldsymbol{K})$. Let p_i, respectively q_i denote the restriction to M_i. Since $p \neq q$, locality gives an M_j with $p_j \neq q_j$ and $|M_j| < \mu$. By induction there exists an $N \prec_{\boldsymbol{K}} M_j$

with $|N| = \mathrm{LS}(\boldsymbol{K})$ and $p_j \restriction N \neq q_j \restriction N$. But then, $p \restriction N \neq q \restriction N$ and we finish. $\square_{12.9}$

Example 12.10. In Chapter 27.5 we give an example of a type which is not $(\aleph_0, \aleph_1)$-tame. By Lemma 12.8, it is not $(\aleph_1, \aleph_1)$-local.

Nontameness can arise in natural mathematical settings. An Abelian group is $\aleph_1$-free if every countable subgroup is free. An Abelian group H is *Whitehead* if every extension of Z by H is free. Shelah constructed an Abelian group of cardinality $\aleph_1$ which is $\aleph_1$-free but not a Whitehead group. (See [39] Chapter VII.4.) Baldwin and Shelah [17] code this into an example of nontameness. Essentially certain points in models of the AEC $\boldsymbol{K}$ code an extension of Z by an abelian group H. Every short exact sequence

$$0 \to \mathcal{Z} \to V \to H' \to 0. \tag{12.1}$$

where H' is a countable submodel of H splits but the short exact sequence ending in H does not. Thus the AEC is not $(\aleph_1, \aleph_1)$-local nor $(\aleph_1, \aleph_0)$-tame.

The following argument by Baldwin, Kueker, and VanDieren [14] shows the strength of assuming tameness and the even greater strength of assuming locality as well.

Theorem 12.11. *Suppose* $\mathrm{LS}(\boldsymbol{K}) = \aleph_0$. *If* $\boldsymbol{K}$ *is* $(\aleph_0, \infty)$*-tame and* μ*-Galois-stable for all* $\mu < \kappa$ *and* $\mathrm{cf}(\kappa) > \aleph_0$ *then* $\boldsymbol{K}$ *is* κ*-Galois-stable.*

Proof. For purposes of contradiction suppose there are more than κ types over some model M^* in $\boldsymbol{K}$ of cardinality κ. We may write M^* as the union of a continuous chain $\langle M_i \mid i < \kappa \rangle$ under $\prec_{\boldsymbol{K}}$ of models in $\boldsymbol{K}$ which have cardinality $< \kappa$. We say that a type over M_i *has many extensions* to mean that it has $\geq \kappa^+$ distinct extensions to a type over M^*. For every i, there is some type over M_i with many extensions. To see this note that each type over M^* is the extension of some type over M_i and, by our assumption, there are less than κ many types over M_i, so at least one of them must have many extensions.

More strongly, for every i, if the type p over M_i has many extensions, then for every $j > i$, p has an extension to a type p' over M_j with many extensions. This follows as, every extension of p to a type over M^* is the extension of some extension of p to a type over M_j. By our assumption there are less than κ many such extensions to a type over M_j, so at least one of them must have many extensions.

We can further show, for every i, if the type p over M_i has many extensions, then for all sufficiently large $j > i$, p can be extended to two types over M_j each having many extensions. By the preceding paragraph, it suffices to establish the result for some $j > i$. So assume that there is no $j > i$ such that p has two extensions to types over M_j each having many extensions. Again, the preceding paragraph tells us that for every $j > i$, p has a unique extension to a type p_j over M_j with many extensions. Let S^* be the set of all extensions of p to a type over M^* – so $|S^*| \geq \kappa^+$. Then S^* is the union of S_0 and S_1, where S_0 is the set of all q in S^* such that $p_j \subseteq q$ for all $j > i$, and S_1 is the set of all q in S^* such that q does not extend p_j for some $j > i$. Now if q_1 and q_2 are different types in S^* then, since $\boldsymbol{K}$ is $(\aleph_0, \infty)$-tame and $\mathrm{cf}(\kappa) > \aleph_0$, their restrictions to some $M_i \prec_{\boldsymbol{K}} M^*$ with $i < \kappa$ must differ. Hence their restrictions to all sufficiently large M_j must differ. Therefore, S_0 contains at most one type. On the other hand, if q is in S_1 then, for some $j > i$, $q \upharpoonright M_j$ is an extension of p to a type over M_j which is different from p_j, hence has at most κ extensions to a type over M^*. Since there are $< \kappa$ types over each M_j (by our stability assumption) and just κ models M_j there can be at most κ types in S_1. Thus S^* contains at most κ types, a contradiction.

Now we can conclude that there is a countable $M \prec_{\boldsymbol{K}} M^*$ such that there are $2^{\aleph_0}$ types over M. Let p be a type over M_0 with many extensions. By the preceding paragraph, there is a $j_1 > 0$ such that p has two extensions p_0, p_1 to types over M_{j_1} with many extensions. Iterating this construction we find a sequence of models M_{j_n} and a tree p_s of types for $s \in 2^\omega$ with the 2^n types p_s (where s has length n) all over M_{j_n} and each p_s has many extensions. Invoking $\aleph_0$-tameness, we can replace each M_{j_n} by a countable M'_{j_n} and p_s by p'_s over M'_{j_n} while preserving the tree structure on the p'_s. Let $\hat{M}$ be the union of the M'_{j_n}. Now for each $\sigma \in 2^\omega$, $p_\sigma = \bigcup_{s \subset \sigma} p_s$ is a Galois-type, by Remark

We have now contradicted ω-Galois-stability: this establishes Theorem 12.11. $\square_{12.11}$

By adding locality we get a much stronger result. The article [14] gets somewhat stronger results assuming only weak-tameness.

Corollary 12.12. *Suppose* $\mathrm{LS}(\boldsymbol{K}) = \aleph_0$ *and* $\boldsymbol{K}$ *has the amalgamation property. If* $\boldsymbol{K}$ *is* $(\aleph_0, \infty)$*-tame,* (ω, ∞)*-local, and* ω*-Galois-stable then* $\boldsymbol{K}$ *is Galois-stable in all cardinalities.*

Proof. We prove the result by induction on cardinals. Theorem 12.11 extends to limit cardinals of uncountable cofinality with no difficulty. For limit cardinals of cofinality ω, at the stage where we called upon $\aleph_0$-tameness to show each type over M_i with many extensions has two extensions over M_j for $j > i$, we now use the hypothesis of (ω, ∞)-locality. $\square_{12.12}$

We now provide a surprising sufficient condition for tameness. Note that EM models built on sequences of finite sequences have the following property. Suppose that we can find a sequence of countably many order indiscernibles $\mathbf{b}^i$ where each $\mathbf{b}^i$ is a sequence $b^i_0, \ldots b^i_{k-1}$ and Φ is an Ehrenfeucht-Mostowski template for models built on indiscernible sequences of k-tuples. If J is k isomorphic linear orders placed one after the other then $EM_\tau(J, \Phi) \in \boldsymbol{K}$.

Recall from Definition 7.2.4 and the following exercises that a linear order L is *transitive* if for any $k < \omega$ and any increasing sequences of length k, $\boldsymbol{a}, \mathbf{b}$, there is an automorphism of L taking $\boldsymbol{a}$ to $\mathbf{b}$. In particular, any two intervals (including intervals of the form (a, ∞) or $(-\infty, b)$) are required to be isomorphic. Moreover there are k-transitive linear orderings of all infinite cardinalities.

Notation 12.13. *Recall from Definition 5.21 that* $H(\kappa)$ *denotes* $\beth_{(2^\kappa)^+}$ *and that we write* H_1*, for* $H(|\tau_0|)$*, the Hanf number for AECs whose vocabulary has car-*

dinality $|\tau_0|$. With a fixed $\boldsymbol{K}$, we write H_1 for $H(\kappa_{\boldsymbol{K}}) = H(\sup(\tau_{\boldsymbol{K}}, \mathrm{LS}(\boldsymbol{K})))$. We established in Chapter 7.2 the convention that the vocabulary of $EM(I,\Phi)$ is the vocabulary $\tau(\Phi)$ of Φ, while for any $\tau' \subseteq \tau(\Phi)$, $EM_{\tau'}(I,\Phi)$ denotes the reduct of $EM(I,\Phi)$ to τ'. We exploit the subtle relationship among such reducts in the next proof and this notation must be examined carefully to understand the argument.

Lemma 12.14. *Let $H_1 \leq \mu < \lambda = |I|$. Suppose $N = EM_{\tau_0}(J,\Phi)$ where J is a transitive linear order that contains an increasing sequence of length $|J| = \mu$. Suppose further that J is an initial segment of I and that $M = EM_{\tau_0}(I,\Phi)$ is saturated. Now suppose $a, b \in M - N$ and realize the same Galois type over N_0 for any $N_0 \prec_{\boldsymbol{K}} N$ with $|N_0| < H_1$. Then a, b realize the same Galois type over N.*

Proof. For some terms σ, ρ: $a = \sigma(\mathbf{s},\mathbf{t})$ and $b = \rho(\mathbf{s},\mathbf{t})$ where $\mathbf{s} \in J$ and $\mathbf{t} \in I - J$. Say the length of $\mathbf{s}$ is k so $\mathbf{s} = \langle s_0, \ldots s_{k-1}\rangle$. Since J is an initial segment the elements of $\mathbf{t}$ are all greater than all elements of J.

For each cardinal $\chi < H_1$ we find a sublinear order K_χ of J whose Skolem hull, $N^1_\kappa = EM(K_\chi,\Phi)$ is a prototype model of cardinality χ where a and b realize the same Galois type over N^1_κ. We then use Theorem A.3.1 to transfer this picture to an arbitrary linear order L. Finally, we get the desired result that a and b have the same type over N if we take L as $J - \mathbf{s}$. After this sketch of the argument, we begin in earnest.

By the transitivity we may write $J - \{\mathbf{s}\}$ as the union of k intervals J_i, separated by the s_i such that each interval contains an increasing sequence of length $|J|$. Let K_χ be a subset of J composed of k isomorphic linear orders $K^\chi_i \subset J_i$, each of which contains an increasing sequence of length χ. Since, $|K_\chi \mathbf{s}| < H_1$, a and b realize the same Galois type over $EM(K_\chi \mathbf{s},\Phi)$. So, there is an automorphism α of M fixing $EM(K_\chi,\Phi)$ and $\mathbf{s}$ and taking a to b. Let N^χ_1 be the structure whose universe is the closure of $K_\chi \mathbf{st}$ under the $\tau_1 = \tau(\Phi)$ functions and α. If τ_2 is obtained by adding constant symbols $\mathbf{s}^*, \mathbf{t}^*$ for $\mathbf{s}, \mathbf{t}$ and a function symbol F to τ_1 and interpreting F as α, N^χ_1 is naturally a τ_2-structure of cardinality χ satisfying:

1. $N^\chi_1 \restriction \tau_0$ is in $\boldsymbol{K}$.
2. $\mathrm{cl}_{\tau_1}(K_\chi \mathbf{st}, M) \restriction \tau_0 \prec_{\boldsymbol{K}} \mathrm{cl}_{\tau_1,\alpha}(K_\chi \mathbf{st}, M) \restriction \tau_0$.
3. α fixes $\mathrm{cl}_{\tau_1}(K_\chi \mathbf{s}, M)$ and takes $a = \sigma(\mathbf{s},\mathbf{t})$ to $b = \rho(\mathbf{s},\mathbf{t})$.

The first two conditions hold by the characterization of submodels in the Presentation Theorem 5.14, since each of N^χ_1, $\mathrm{cl}_{\tau_1}(K_\chi \mathbf{st}, M)$, $\mathrm{cl}_{\tau_1,\alpha}(K_\chi \mathbf{st}, M)$ are τ_1 substructures of $EM(I,\Phi)$. The third holds by the choice of α. The first two are determined by the omission of τ_2-types and so are preserved in other EM-models.

By Theorem A.3.1 we can obtain a countable sequence $\mathbf{b}^i$ of τ_2-indiscernibles where each $\mathbf{b}^i$ is a sequence $b^i_0, \ldots b^i_{k-1}$ and the type of the $\langle \mathbf{b}^i_j : i < n\rangle$ is realized in K^χ_j for each n. Let Φ_2 be the τ_2-diagram of these indiscernibles. Now for any linear order L consisting of k isomorphic segments we can form τ_2-Skolem hulls of L by interpreting $\mathbf{s}$ as points separating the segments and $\mathbf{t}$ as a sequence at the end. Then, we have:

a. $N_L = EM_{\tau_0}(L, \Phi_2)$ is in $\boldsymbol{K}$,

b. $EM_{\tau_0}(L\mathbf{st}, \Phi) \prec_{\boldsymbol{K}} EM_{\tau_0}(L, \Phi_2)$,

c. and F is an automorphism of N_L which takes $\sigma(\mathbf{s}, \mathbf{t})$ to $\rho(\mathbf{s}, \mathbf{t})$ and fixes $EM_{\tau_0}(L\mathbf{s}, \Phi)$ pointwise.

In particular, if we take $J - \mathbf{s}$ for L, c) yields that $N_{J-\mathbf{s}} = EM_{\tau_0}(J - \mathbf{s}, \Phi_2)$ has an automorphism which takes $\sigma(\mathbf{s}, \mathbf{t})$ to $\rho(\mathbf{s}, \mathbf{t})$ and fixes $EM_{\tau_0}(J, \Phi) = N$ pointwise.

Let τ_2^- be τ_2 without the constant symbols $\mathbf{s}^*, \mathbf{t}^*$ and Φ_2^- the corresponding restriction of Φ_2. Then $EM_{\tau_2^-}(J\mathbf{t}, \Phi_2^-) \approx EM_{\tau_2^-}(J - \mathbf{s}, \Phi_2)$ by mapping $\mathbf{s}$ to the interpretation of $\mathbf{s}^*$ and $\mathbf{t}$ to the interpretation of $\mathbf{t}^*$ in $EM(J - \mathbf{s}, \Phi_2)$. Restricting further to τ_0, $N_{J-\mathbf{s}} = EM_{\tau_0}(J - \mathbf{s}, \Phi_2) \approx EM_{\tau_0}(J\mathbf{t}, \Phi_2^-)$.

By conditions a) and b) $EM_{\tau_0}(J\mathbf{t}, \Phi) \prec M$ and $EM_{\tau_0}(J\mathbf{t}, \Phi) \prec EM_{\tau_0}(J\mathbf{t}, \Phi_2^-)$. So we can embed $N_{J-\mathbf{s}} \approx EM_{\tau_0}(J\mathbf{t}, \Phi_2^-)$ into M over $EM_{\tau_0}(J\mathbf{t}, \Phi)$ by some g. The image under g of the interpretation of F in $EM(J\mathbf{t}, \Phi_2^-)$ is an automorphism β of $gN_{J-\mathbf{s}}$ which fixes N and takes a to b. Invoking the homogeneity of M, β extends to an automorphism of M and proves that a and b realize the same Galois type over N as required. $\square_{12.14}$

Now we apply this result and the fact that each element of $N = EM(J, \Phi)$ is given by one of small number of terms to conclude the required weak-tameness.

Theorem 12.15. *If $\boldsymbol{K}$ is λ-categorical for $\lambda \geq H_1$, then for any μ with $H_1 \leq \mu < \mathrm{cf}(\lambda)$, for some $\chi < H_1$, $\boldsymbol{K}$ is (χ, μ)-weakly tame.*

Proof. Let M be the categoricity model and let N be a saturated submodel of M of smaller cardinality μ. By Lemma 11.10, we can represent N as $EM_{\tau_0}(J, \Phi)$ where τ_0 is the vocabulary of $\boldsymbol{K}$ and J is a transitive linear order that contains an increasing sequence of length $|J| = \mu$. Now write M as $EM(I, \Phi)$ where J is an initial segment of I. We want to choose χ so that if two types over N disagree they disagree on a model of cardinality less than χ. Suppose $a \in \mathbb{M}$ realizes $p \in \mathbb{S}(N)$. Since M is μ^+-homogenous universal there is an automorphism of $\mathbb{M}$ fixing N and mapping a into M. So we reduce to considering $a, b \in M - N$. For some terms σ, τ: $a = \sigma(\mathbf{s}, \mathbf{t})$ and $b = \rho(\mathbf{s}, \mathbf{t})$ where $\mathbf{s} \in J$ and $\mathbf{t} \in I - J$. Say the length of $\mathbf{s}$ is k. By the contrapositive of Lemma 12.14, there is $\chi < H_1$ and an $N_0 \prec_{\boldsymbol{K}} N$ with $|N_0| = \chi$ so that $\mathrm{tp}(a/N_0) \neq \mathrm{tp}(b/N_0)$. Since J is an initial segment of I and J is k-transitive, this also holds for any a', b' given by $\sigma(\mathbf{s}', \mathbf{t}')$ and $\rho(\mathbf{s}', \mathbf{t}')$. Since there are only $|\tau_1|$ pairs of terms and $|\tau_1|$ is not cofinal in H_1, we have the result. $\square_{12.15}$

This result is stated in [128] and [129]; the straightforward proof of Lemma 12.14 is due to Tapani Hyttinen after reading [129] and various comments thereon by both Baldwin and Shelah. The argument given in [129] contains many ideas that are not necessary for the exact result here, but may be useful for generalizations.

Question 12.16. *The argument for Theorem 12.15 does not immediately give that there is a cardinal $\chi(\Phi) < H_1$ such that $\boldsymbol{K}$ is $(\chi(\Phi), [H_1, < \lambda))$-tame. Can we obtain that result? Of course it is $(H_1, [H_1, < \lambda))$-tame. In the downward categoricity argument, we actually use the exact formulation of Theorem 12.15: for each $\mu < \lambda$ there is a $\chi_\mu < H_1$ such that $\boldsymbol{K}$ is (χ_μ, μ) tame.*

Locality and tameness provide key distinctions between the general AEC case and homogenous structures. In homogeneous structures, types are syntactic objects and locality is trivial. Thus, as pointed out by Shelah, Hyttinen, and Buechler-Lessmann, Lemma 12.5 applies in the homogeneous context.

13
Splitting and Minimality

In this chapter we introduce the notion of splitting for AEC. This is an approximation to an independence notion such as forking. In this chapter we work in a a $\boldsymbol{K}$ with arbitrarily large models and the amalgamation property which is categorical in a successor cardinal λ.

We use splitting to construct types with the following properties. For a model M with cardinality μ, we want to find $p \in \mathbb{S}(M)$ which is a) nonalgebraic (i.e. not realized) b) minimal (Definition 13.21), and c) extendible (has a non-algebraic extension to the model of cardinality λ). For $\mu < \lambda$, any M with cardinality μ, and any nonalgebraic $p \in \mathbb{S}(M)$ we will always be able to find $p \in \mathbb{S}(M')$ for an extension M' of M that is minimal. And if M is saturated with $|M| > \mathrm{LS}(\boldsymbol{K})$, we will be able to extend any nonalgebraic $p \in \mathbb{S}(M)$. Thus, if $\boldsymbol{K}$ is μ-categorical and $\mu > \mathrm{LS}(\boldsymbol{K})$ (so saturation makes sense), the problem is solved in this chapter. But in parallel we will be preparing to handle the situation where $\boldsymbol{K}$ is not μ-categorical (Chapter 16) or $\mu = \mathrm{LS}(\boldsymbol{K})$ ([101] for $\mathrm{LS}(\boldsymbol{K}) = \aleph_0$).

Now, we introduce the key notion of splitting.

Definition 13.1. *A type $p \in \mathbb{S}(N)$ μ-*splits *over $M \prec_{\boldsymbol{K}} N$ if and only if there exist $N_1, N_2 \in \boldsymbol{K}_{\leq\mu}$ and h, a $\boldsymbol{K}$-embedding from N_1 into N with $h(N_1) = N_2$ such that $M \prec_{\boldsymbol{K}} N_l \prec_{\boldsymbol{K}} N$ for $l = 1, 2$, h fixes M, and $p \restriction N_2 \neq h(p \restriction N_1)$.*

In the next lemma we isolate the key tool for finding submodels over which a type does not split. Note the implicit use of amalgamation to get the monster model required in the next proof. We are given a type p and a sequence of models in the monster model; just as in the first order case we make copies of this situation to form a tree of types.

Lemma 13.2. *Let κ be the least cardinal with $2^\kappa > \mu$. Suppose there is an increasing chain of models M_i with $|M_i| \le \mu$ for $i < \kappa$ and $p \in \mathbb{S}(M_\kappa) = \mathbb{S}(\bigcup_{i<\kappa} M_i)$ such that for each i, $p \restriction M_{i+1}$ μ-splits over M_i. Then $\boldsymbol{K}$ is not stable in μ.*

Proof. In this proof we discuss various elementary maps h between structures of size μ; for any such h, $\hat{h}$ denotes an extension of h to an automorphism of the monster model which remains fixed throughout the argument.

For each i, there exist $M_i^1, M_i^2 \prec_{\boldsymbol{K}} M_{i+1}$ with $M_i \prec_{\boldsymbol{K}} M_i^1, M_i^2$ and a $\boldsymbol{K}$-embedding f_i which fixes M_i, maps M_i^1 to M_i^2 but $f_i(p \restriction M_i^1) \neq p \restriction M_i^2$.

For each $\eta \in 2^{\le\kappa}$ we define a map h_η with domain $M_{\lg(\eta)}$, a model N_η and a type $p_\eta \in \mathbb{S}(N_\eta)$.

Let $N_\emptyset = M_0$, $h_\emptyset = \mathrm{id}_{M_0}$, $p_\emptyset = p \restriction M_0$. Take $h_0 = \mathrm{id}_{M_1}$, $N_0 = M_1$, $p_0 = p \restriction M_1$; and set $h_1 = \hat{f}_0$, $N_1 = \hat{f}_0(M_1)$, $p_1 = \hat{f}_0(p \restriction M_1)$. Take unions at limits. At successors, let $h_{\eta\frown 0} = \hat{h}_\eta \restriction M_{\lg(\eta)+1}$ and $h_{\eta\frown 1} = \hat{h}_\eta \circ \hat{f}_{\lg(\eta)} \restriction M_{\lg(\eta)+1}$.

Then, let $N_\eta = h_\eta(M_{\lg(\eta)})$ and $p_\eta = h_\eta(p \restriction M_{\lg(\eta)})$. Let M_κ^* be a model of cardinality μ, chosen by the Lowenheim-Skolem property to contain all the $h_\eta(M_\gamma) = N_\eta$ for $\gamma < \kappa$ and $\eta \in 2^\gamma$. We can easily extend each p_η to a type $\hat{p}_\eta$ over M_κ^*. But if $\eta \neq \eta'$ then $\hat{p}_\eta \neq \hat{p}_{\eta'}$. To see this, let $\eta \wedge \eta' = \nu$ (where $\wedge$ is 'meet'). Then $\hat{p}_\eta \restriction N_{\nu\frown 0} \neq \hat{p}_{\eta'} \restriction N_{\nu\frown 1}$ since the domain of each contains $h_\nu(M^2_{\lg(\nu)})$ and $h_\nu(p \restriction M^2_{\lg(\nu)}) \neq h_\nu(f_{\lg(\nu)}(p \restriction M^1_{\lg(\nu)}))$ as $p \restriction M^2_{\lg(\nu)} \neq f_{\lg(\nu)}(p \restriction M^1_{\lg(\nu)})$. $\square_{13.2}$

Notice three easy facts about μ-splitting.

Exercise 13.3. *1. If $p \in \mathbb{S}(M)$ χ-splits over M_0 then for any $\chi' \ge \chi$, $p \in \mathbb{S}(M)$, p χ'-splits over M_0.*

2. If $p \in \mathbb{S}(M)$ does not χ-split over M_0 and $M_0 \prec_{\boldsymbol{K}} M_1 \prec_{\boldsymbol{K}} M$ with $|M_1| = \chi$, then

   (a) $p \in \mathbb{S}(M)$ does not χ-split over M_1.

   (b) $p \restriction M_1$ does not χ-split over M_0.

Exercise 13.4. *Construct an example showing the converse to Exercise 13.3.2 does not hold. (Hint: Consider Example 20.13. A weaker 'converse' is Exercise 13.9.)*

We now establish existence, uniqueness and extension for nonsplitting in this context. More precisely, existence requires only stability; uniqueness (except in the same cardinal) requires tameness as well; and the ability to make non-splitting extensions to models of greater cardinality depends on categoricity in a still larger cardinal.

Lemma 13.5. *[Existence] Suppose $|M| \ge \mu \ge \mathrm{LS}(\boldsymbol{K})$, and $\boldsymbol{K}$ is μ-stable. Then for every $p \in \mathbb{S}(M)$, there is an $N_p \prec_{\boldsymbol{K}} M$ with $|N_p| = \mu$ such that p does not μ-split over N_p.*

Proof. If $|M| = \mu$, the result is trivially true with $N_p = M$. Fix κ minimal so that $\mu^\kappa > \mu$. If the conclusion fails, we can construct a sequence of models $\langle N_i : i < \kappa\rangle$ such that $p \restriction N_{i+1}$ μ-splits over N_i for each i (so $|N_i| \leq \mu$). That is, there exist $N_i^1, N_i^2 \prec_{\boldsymbol{K}} M$ with $N_i \prec_{\boldsymbol{K}} N_i^1, N_i^2$ and a $\boldsymbol{K}$-embedding f_i which fixes N_i, maps N_i^1 to N_i^2 but $f_i(p \restriction N_i^1) \neq p \restriction N_i^2$. Choose $N_{i+1} \prec_{\boldsymbol{K}} M$ to contain N_i^1, N_i^2. Take unions at limits. Either at some stage we find an N_i of cardinality μ such that p does not split over N_i or we find a chain of length κ as in Lemma 13.2. The union of this chain has cardinality μ. By Lemma 13.2, we contradict μ-stability. $\square_{13.5}$

The trivial case of Lemma 13.5, p does not μ-split over its domain if that domain has cardinality μ, is not sufficient to gain non-splitting extensions. We provide two ways to find extensions. We can derive the existence of a non-splitting extension to models of size μ by requiring that N is universal over N_p; later we do even more if N is saturated. The following argument is based on one of VanDieren [144].

Lemma 13.6. *If $N \prec_{\boldsymbol{K}} M \prec_{\boldsymbol{K}} M'$ all have cardinality μ, M is μ-universal over N and $p \in \mathbb{S}(M)$ does not μ-split over N, there is a $p' \in \mathbb{S}(M')$ that does not μ-split over N.*

Proof. Let α be a $\boldsymbol{K}$-embedding of M' into M over N. By monotonicity, Exercise 13.3.2.a, α extends to an automorphism of the monster model that fixes N. Fix a realization a of p. Let $N_1 = \alpha^{-1}(M)$. So we have $N \prec_{\boldsymbol{K}} \alpha(M') \prec_{\boldsymbol{K}} M \prec_{\boldsymbol{K}} M' \prec_{\boldsymbol{K}} N_1$. Since $\mathrm{tp}(a/M)$ does not split over N, by invariance $p' = \mathrm{tp}(\alpha^{-1}(a)/N_1)$ does not μ-split over N. As, $\alpha(N_1) = M$, the non-splitting guarantees that $\alpha(p') = p' \restriction \alpha(N_1)$. That is, $\mathrm{tp}(a/M) = \mathrm{tp}(\alpha^{-1}(a)/M)$. So, p' is a non-splitting extension of p to N_1 and its restriction to M' meets the requirements. $\square$13.6

The following uniqueness of non-splitting extensions result is stated without proof in [128]; we follow the idea of a variant in [144].

Theorem 13.7. *[Uniqueness] Assume $\boldsymbol{K}$ is $(\chi, |N|)$-tame for some $\mathrm{LS}(\boldsymbol{K}) \leq \chi \leq \mu$. Let $M_0 \prec_{\boldsymbol{K}} M \prec_{\boldsymbol{K}} N$ with $\chi \leq |M_0| \leq \mu$ and suppose M is μ-universal over M_0. Then if $p \in \mathbb{S}(M)$ does not μ-split over M_0, p has at most one extension to $\mathbb{S}(N)$ that does not μ-split over M_0.*

If N is saturated, we need only $(\chi, |N|)$-weak tameness.

Proof. Suppose for contradiction that r and q in $\mathbb{S}(N)$ are distinct non-splitting extensions of p. By tameness, there is an N_1 with $M_0 \prec_{\boldsymbol{K}} N_1 \prec_{\boldsymbol{K}} N$ with $|N_1| = \max(\chi, |M_0|)$ and $q \restriction N_1 \neq r \restriction N_1$. Since M is μ-universal over M_0, there is an embedding f of N_1 into M over M_0; let N_2 denote $f(N_1)$. Since r and q don't μ-split over M_0, $r \restriction N_2 = f(r \restriction N_1)$ and $q \restriction N_2 = f(q \restriction N_1)$. But since $r \restriction N_1 \neq q \restriction N_1$, $f(r \restriction N_1) \neq f(q \restriction N_1)$. By transitivity of equality, $r \restriction N_2 \neq q \restriction N_2$. Since $N_2 \prec_{\boldsymbol{K}} M$, this contradicts that r and q both extend $p \in \mathbb{S}(M)$. $\square_{13.7}$

The tameness hypothesis in Lemma 13.7 is needed only to move to larger cardinalities; specifically:

Exercise 13.8. *Show that the 'tameness' hypothesis in Lemma 13.7 is not needed if the conclusion is weakened to consider only $|M| = |N| = \mu$.*

Use the uniqueness to show.

Exercise 13.9. *If $M_0 \prec_{\boldsymbol{K}} M_1 \prec_{\boldsymbol{K}} M_2$, $|M_0| = |M_2|$ each of the extensions is $|M_0|$-universal, $p \in \mathbb{S}(M_2)$ does not split over M_1 and $p \upharpoonright M_1$ does not split over M_0, then p does not split over M_0. Show that the restriction on the cardinality of M_2 can be replaced by assuming $\boldsymbol{K}$ is $(< |M_0|, |M_2|)$-tame.*

Before considering the strongest versions of the extension property, we note that Grossberg and VanDieren partially calculated the stability spectrum of $\boldsymbol{K}$.

Theorem 13.10. *[45] Let $\mu \geq \mathrm{LS}(\boldsymbol{K})$. If $\boldsymbol{K}$ is (μ, ∞)-tame and μ-stable then $\boldsymbol{K}$ is stable in all κ with $\kappa^\mu = \kappa$.*

Proof. We show $\mathbb{S}(M)$ has cardinality κ if $|M| = \kappa$. Fix any $N \prec_{\boldsymbol{K}} M$ with cardinality μ. Applying μ-stability, Lemmas 11.3 and 11.5 extend N to a model N' which is μ-universal over N. Then let M' be an amalgam of M and N' over N with cardinality κ; M' is μ-universal over N. Repeating this process $\kappa^\mu = \kappa$ times we get a model N^1 such that every submodel N_0 of M with cardinality μ has a μ-universal extension N_0' in N^1. Now iterate that construction $\mu^+ \leq \kappa$ times to get a model $\widehat{M}$ of cardinality κ such that every submodel N_0 of $\widehat{M}$ with cardinality μ has a μ-universal extension $N_0' \prec_{\boldsymbol{K}} \widehat{M}$. Since every type over M has at least one extension (possibly algebraic) to $\widehat{M}$, it suffices to prove $\mathbb{S}(\widehat{M}) = \kappa$. Clearly $|M| > \mu$. So, if $p \in \mathbb{S}(\widehat{M})$, by Lemma 13.5, p does not μ-split over a submodel N_p of size μ. By construction there is an $N_p' \prec_{\boldsymbol{K}} \widehat{M}$, which is μ-universal over N_p. By tameness and Theorem 13.7, p is the unique extension of $p \upharpoonright N_p'$ to $\widehat{M}$ which does not split over N_p. So the number of nonsplitting extensions of $p \upharpoonright N_p$ to $\widehat{M}$ is bounded by $|\mathbb{S}(N_p')| = \mu$. Then $|\mathbb{S}(\widehat{M})| \leq \kappa^\mu \cdot \mu = \kappa$ and we finish. $\square_{13.10}$

The following result of Baldwin, Kueker and VanDieren [14] tells us bit more about the stability spectrum under tameness hypotheses.

Theorem 13.11. *Let $\boldsymbol{K}$ be an abstract elementary class with the amalgamation property that has Löwenheim-Skolem number $\leq \kappa$ and is (κ, κ^+)-weakly-tame. Then if $\boldsymbol{K}$ is Galois-stable in κ it is also Galois-stable in κ^+.*

Proof We proceed by contradiction and assume that M^* is a model of cardinality κ^+ with more than κ^+ types over it. By Lemmas 11.3 and 11.8, we can extend M^* to a (κ, κ^+)-limit model which is saturated. Since M^* has at least as many types as the original we just assume that M^* is a saturated, (κ, κ^+)-limit model witnessed by $\langle M_i : i < \kappa^+ \rangle$.

Let $\{p_\alpha : \alpha < \kappa^{++}\}$ be a set of distinct types over M^*. Stability in κ and Lemma 13.5 yield that for every p_α there exists $i_\alpha < \kappa^+$ such that p_α does not κ-split over M_{i_α}. By the pigeon-hole principle there exists $i^* < \kappa^+$ and $A \subseteq \kappa^{++}$ of cardinality κ^{++} such that for every $\alpha \in A$, $i_\alpha = i^*$.

Now apply the argument for Theorem 12.11 to the p_α for $\alpha \in A$ to conclude there exist $p, q \in S(M^*)$ and $i < i' \in A$, such that neither p nor q κ-splits over M_i or $M_{i'}$ but $p \restriction M_{i'} = q \restriction M_{i'}$. By weak tameness, there exists an ordinal $j > i'$ such that $p \restriction M_j \neq q \restriction M_j$. Notice that neither $p \restriction M_j$ nor $q \restriction M_j$ κ-split over M_i. This contradicts Lemma 13.7 by giving us two distinct extensions of a non-splitting type to the model M_j which by construction is universal over $M_{i'}$. $\square_{13.11}$

If we assume (κ, ∞)-tameness as well, an immediate induction extends Theorem 13.11 to deduce stability in κ^{+n} for every $n < \omega$ from stability in κ.

By $\mathrm{aut}_{\mathrm{M}}(\mathbb{M})$ we mean the set of automorphisms of $\mathbb{M}$ which fix M pointwise. Now we establish a sufficient condition M for finding a non-splitting extension of $p \in \mathbb{S}(M)$. Note that the requirement in Lemma 13.13 that p be realized in N is guaranteed by the saturation of N if $\boldsymbol{K}$ is λ-categorical.

Definition 13.12. *A type $p \in \mathbb{S}(N)$ is* nonalgebraic *if it is not realized in N.*

Lemma 13.13. *[Extension] Let $\mu < \lambda$, $M = EM(\mu^{<\omega}, \Phi) = EM(I, \Phi)$, and suppose $p \in \mathbb{S}(M)$ is nonalgebraic. Suppose p does not $|M_0|$-split over $M_0 = EM(I_0, \Phi)$ with $I_0 \subset I$ and $|I_0| < |I|$. If p is realized in $N = EM(J, \Phi)$ $(J = \lambda^{<\omega})$, then there is an extension $\hat{p}$ of p to N so that for some M_0' with $|M_0'| = |M_0|$, $\hat{p}$ does not $|M_0|$-split over M_0'.*

Proof. Let $\tau(\boldsymbol{a}) \in N$ realize p (for some term τ). We can increase I_0 to I_0' so that $(\boldsymbol{a} \cap M) \subset I_0'$ and let $M_0' = EM(I_0', \Phi)$. Now extend J to J' by adding a finite sequence $\boldsymbol{a}'$ with a_i' in the same I-cut as a_i but a_i' greater than any element of J in that cut. Now since $\boldsymbol{a}$ and $\boldsymbol{a}'$ realize the same type in the language of orders over I, a compactness argument shows there is an extension J'' of J' and an automorphism f of J'' which fixes I and maps $\boldsymbol{a}$ to $\boldsymbol{a}'$. Thus, $\alpha = \tau(\boldsymbol{a}')$ realizes p in $EM(J'', \Phi)$. We now show $p' = \mathrm{tp}(\tau(\boldsymbol{a}')/N)$ does not $|M_0|$-split over M_0'. Consider any N_1, N_2 in N with $|N_1| = |N_2| = |M_0|$ and a map h over M_0' which maps N_1 to N_2; by homogeneity of the monster, h extends to an $\eta \in \mathrm{aut}_{\mathrm{M_0}}(\mathbb{M})$. We will show there is a map $\hat{\eta} \in \mathrm{aut}_{\mathrm{N_2}}(\mathbb{M})$ witnessing that $h(p' \restriction N_1) = p' \restriction N_2$; that is $\hat{\eta}$ maps $\eta(\alpha)$ to α. Choose $K \subseteq J$ with $|K| < |M|$ such that $N_1, N_2 \prec_{\boldsymbol{K}} N_3 = EM(K, \Phi)$. By the 'relative saturation' (Corollary 10.7) of I in J there is an order isomorphism g fixing I_0' and mapping K to $K' \subseteq I$. Moreover, without loss of generality, we take the domain of g to be $I_0'K\boldsymbol{a}'$ and g fixes $\boldsymbol{a}'$. (To see this, note that every decreasing chain in $\mu^{<\omega}$ is countable. For each a_i', choose a countable $X_i \subset I$ with a_i' as the infimum of X_i. Then require $\bigcup_i X_i$ as well as I_0' to be fixed by g.) Thus g extends to a $\gamma \in \mathrm{aut}(\mathbb{M})$, such that γ fixes $EM(I_0'\boldsymbol{a}', \Phi)$. In particular, γ fixes $M_0'\alpha$. Let N_1' denote $\gamma(N_1)$ and N_2' denote $\gamma(N_2)$. Then $\eta' = \gamma\eta\gamma^{-1} \in \mathrm{aut}_{\mathrm{M_0'}}(\mathbb{M})$ takes N_1' to N_2'. Since p does not split over M_0 (and thus not over M_0') there is an $\hat{\eta}' \in \mathrm{aut}_{\mathrm{N}_2'}(\mathbb{M})$, taking $\eta'(\alpha)$ to α. We claim $\gamma^{-1}\hat{\eta}'\gamma$ is the required η; clearly this map fixes N_2. But since γ fixes α,

$$\gamma^{-1}\hat{\eta}'\gamma(\eta(\alpha)) = \gamma^{-1}\hat{\eta}'\gamma(\eta(\gamma^{-1}\alpha))$$

which is $\gamma(\hat{\eta}'(\eta'(\alpha))) = \alpha$, by the choice of $\hat{\eta}'$ and since γ fixes α. $\square_{13.13}$

Using very heavily the representation of saturated models as specific kinds of EM-models we can conclude:

Lemma 13.14. *Suppose the AEC $\boldsymbol{K}$ is categorical in the regular cardinal λ. If $M \prec_{\boldsymbol{K}} N$ is saturated for $\mathrm{LS}(\boldsymbol{K}) < |\mathrm{M}| \leq |\mathrm{N}| \leq \lambda$ and p is a nonalgebraic type in $\mathbb{S}(M)$, then for some $N_p \prec_{\boldsymbol{K}} M$, p has an extension $\hat{p}$ to $\mathbb{S}(N)$ which does not $|N_p|$-split over N_p. (The cardinality of N_p can be chosen with $\mathrm{LS}(\boldsymbol{K}) \leq |\mathrm{N_p}| < |\mathrm{N}|$.)*

Proof. Using Exercise 13.3.1 it suffices to prove the result for the model N of cardinality λ. Note that categoricity in λ implies that $N \approx EM(\lambda^{<\omega}, \Phi)$. By the saturation of M and N and using Exercise 9.3 we may assume M has the form $EM(\mu^{<\omega}, \Phi)$. Choose χ with $\mathrm{LS}(\boldsymbol{K}) \leq \chi < \mu$; then $\boldsymbol{K}$ is χ-stable. By Lemma 13.5 (with χ playing the role of μ), each $p \in \mathbb{S}(M)$ does not split over an N_p with $|N_p| \leq \chi < \mu$. Since M is saturated, M is $|N_p|$-universal over N_p. We can extend N_p to have the form $EM(I_0, \Phi)$ for some $I_0 \subseteq \chi^{<\omega}$ by Exercise 13.3.2. We finish by Lemma 13.13. $\square_{13.14}$

With some further work and assuming tameness, we can show that this non-splitting extension is in fact non-algebraic. We show the existence of non-algebraic extensions in two stages. In first order logic it is immediate that a non-algebraic type over a model M has many realizations. David Kueker provided the following example showing that if M is not saturated this 'obvious' fact may fail.

Example 13.15. *Let the language L contain equality and one unary predicate symbol P. Let $\boldsymbol{K}$ be the collection of L-structures M such that $|P(M)| \leq 1$. Then $(\boldsymbol{K}, \subseteq)$ is an AEC with the amalgamation property. Let $M \in \boldsymbol{K}$ with $|P(M)| = 0$ and let $M \subset N$ with $a \in P(N)$. Then $\mathrm{tp}(a/M)$ is non-algebraic but never has more than one realization. Moreover, $\mathrm{tp}(a/N)$ does not $|M|$-split over M.*

We will give two arguments showing that a non-algebraic type over a saturated model has a non-algebraic extension. The first is a lemma of Lessmann that uses tameness; the second uses the hypothesis that we are below a categoricity cardinal. We will obtain the saturated model M in the hypothesis of Lemma 13.16 from categoricity in some regular $\lambda \geq |M|$. Then, we will invoke tameness to convert the existence of non-splitting extensions obtained in Lemma 13.14 to the existence of non-algebraic extensions to models of larger cardinality.

Lemma 13.16. *Suppose $\boldsymbol{K}$ is (χ, μ)-weakly tame with $\mathrm{LS}(\boldsymbol{K}) \leq \chi < \mu$ and that $\boldsymbol{K}$ is χ-stable. Let M be a saturated model of cardinality μ. If $p \in \mathbb{S}(M)$ is nonalgebraic and $M \prec_{\boldsymbol{K}} N$ with $|N| = |M|$ then p has a nonalgebraic extension to N. In particular, p has at least μ^+ realizations in $\mathbb{M}$.*

Proof. Choose $M_0 \prec_{\boldsymbol{K}} M$ with $|M_0| = \chi$ so that p does not χ-split over M_0. (Apply Lemma 13.5 with χ here as μ there.) Since $\boldsymbol{K}$ is χ-stable and $\chi < \mu$ we can choose M_1 with $M_0 \prec_{\boldsymbol{K}} M_1 \prec_{\boldsymbol{K}} M$ and M_1 is 1-special over M_0 (so $|M_1| = \chi$). By Theorem 11.8.1, M_1 is χ-universal over M_0. Since there is a saturated model of cardinality μ, there is a proper $\boldsymbol{K}$-elementary extension N' of N which is saturated (See Exercise 9.15). Let f be an isomorphism between M and N' which fixes M_1. Then both p and $f(p) \upharpoonright M$ are non-χ-splitting extensions of $p \upharpoonright M_1$; by Lemma 13.7, $f(p)$ extends p. We need only weak tameness because M is saturated. Clearly, $f(p)$ is nonalgebraic. We obtain the 'in particular' by choosing N to realize p and iterating. $\square_{26.2}$

Now we find μ^+ realizations without invoking tameness but uses we are below a categoricity cardinal to quote Lemma 11.3.

Theorem 13.17. *Suppose $\boldsymbol{K}$ is λ-categorical for a regular λ with $\lambda > \mu >$ $\mathrm{LS}(\boldsymbol{K})$, M is saturated of cardinality μ, and $p \in \mathbb{S}(M)$. Then any non-algebraic $p \in \mathbb{S}(M)$ is realized μ^+ times.*

Proof. Let κ be the least cardinal with $2^\kappa > \mu$ and let α be the ordinal $\kappa + \omega$. By Lemma 11.3 we can extend M to an M_1 that is α-special over M. But by Lemma 11.16, M_1 is saturated so we can take M as M_1 and know that it is α-special over a model of cardinality μ, say by $\langle M_j : j < \mu \times \alpha \rangle$.

Now for any $p \in M$, apply the actual construction in Lemma 13.5, choosing the N_i contained in the $M_{\mu\times i}$ and N_i^1, N_i^2 in the $M_{\mu\times i+1}$ element and N_{i+2} in the $M_{\mu\times i+2}$ to find an $N_p \prec_{\boldsymbol{K}} M$ with M μ-universal over N_p and p does not μ-split over N_p. Now, applying Lemma 13.6 iteratively, we have the result. $\square_{13.17}$

Now we show that, assuming categoricity, there are nonalgebraic extensions up to the categoricity cardinal. The following exercise is key. Note that Example 13.15 showed the possibility that $\hat{p}$ in the next exercise is realized.

Exercise 13.18. *Show that if $p \in \mathbb{S}(M)$ extends to $\hat{p} \in \mathbb{S}(N)$ and $\hat{p}$ is realized in N but $\hat{p}$ does not split over M, then $\hat{p}$ has a unique realization in N.*

Theorem 13.19. *Suppose $\boldsymbol{K}$ is λ-categorical and λ is regular. Suppose further that $\boldsymbol{K}$ is (χ, λ)-weakly tame for some χ with $\mathrm{LS}(\boldsymbol{K}) \leq \chi < \lambda$. Every non-algebraic type p over a saturated model M with $\lambda \geq |M| > \mathrm{LS}(\boldsymbol{K})$ has a nonalgebraic extension to any extension N of M with $|N| \leq \lambda$.*

Proof. Without loss of generality N is a saturated model and $|N| = \lambda$. If $|M| = \lambda$, the result is given by Lemma 13.16. If $|M| < \lambda$, apply Lemma 13.14 to obtain a $\hat{p} \in \mathbb{S}(N)$ which does not $|N_p|$-split over N_p. Using the 'in particular' from Lemma 13.16, p is realized at least twice in N. Now, the $\hat{p} \in \mathbb{S}(N)$ from Lemma 13.14 is nonalgebraic. As, if $\hat{p}$ is realized by some $a \in N$, then, by Exercise 13.18, $\hat{p}$ splits over M. $\square_{13.19}$

Definition 13.20. *A type $p \in \mathbb{S}(M)$ is* big *if p is realized at least $|M|^+$ times in $\mathbb{M}$.*

Now we introduce the notion of a minimal type and show that such types exist. Note that the definition deals only with types of the same cardinality; we will use tameness to establish the uniqueness of nonalgebraic extensions of minimal types over saturated models to models of larger cardinality.

Definition 13.21. *A type $p \in \mathbb{S}(M)$ is minimal if p is big and there is at most one big extension of p to any model of the* same *cardinality as M.*

Note that if $p \in \mathbb{S}(N)$ and $p' \in \mathbb{S}(M)$ is a non-algebraic extension of p and $|N| = |M|$, then p' is minimal. But this is no longer obvious if $|M| > |N|$. The following observation is immediate. But note that it depends on our defining minimality for big types rather than just non-algebraic types. Example 13.15 show that there are non-algebraic types that cannot be properly extended to non-algebraic types. Minimality is sometimes [44] (and perhaps better) called quasiminimal.

Lemma 13.22. *If M is saturated, $p \in \mathbb{S}(M)$ is minimal, and $M \prec_{\boldsymbol{K}} N$ with $|N| = |M|$, then p has a unique big extension to $\mathbb{S}(N)$.*

Theorem 13.23. *If $\boldsymbol{K}$ is categorical in some regular $\lambda > \mu$ then for some (any) saturated M with cardinality μ, there is a minimal type over M.*

Proof. Of course, $\boldsymbol{K}$ is stable in μ. The proof proceeds by attempting to construct a tree of models of height κ for the least κ such that $2^\kappa \geq \mu$. Each model will have cardinality μ. Stability guarantees that this construction fails and gives us a minimal type over a saturated model. Construct this tree along with functions and realizations witnessing that each path is a coherent chain of nonalgebraic Galois types in the sense of Definition 12.2. At successor ordinals this is easy (as in Lemma 12.5). At limit ordinals we use a construction as in Lemma 12.3 to find the union of the chain of types constructed at that point and to define functions preserving the coherence of the system. More precisely, as in 9.7 part 5 of [128] we construct the tree as follows. For each $\alpha < \kappa$ and for each $\eta \in 2^\alpha$ we choose M_η, a_η and $h_{\eta,\eta\restriction\beta}$ for $\beta \leq \alpha$ such that the $M_{\eta\restriction\beta}$ for $\beta \leq \alpha$ are a continuous increasing $\prec_{\boldsymbol{K}}$-chain of models of cardinality μ satisfying:

1. $M_{\eta\restriction(\beta+1)}$ is ω-special over $M_{\eta\restriction\beta}$;
2. $h_{\eta,\eta\restriction\beta}$ maps $a_{\eta\restriction\beta}$ to a_η while fixing $M_{\eta\restriction\beta}$;
3. if $\gamma \leq \beta \leq \alpha$ and $\eta \in 2^\alpha$,
$$h_{\eta,\eta\restriction\gamma} = h_{\eta,\eta\restriction\beta} \circ h_{\eta\restriction\beta,\eta\restriction\gamma};$$
4. and for each η, $M_{\eta\frown 0} = M_{\eta\frown 1}$, but $\mathrm{tp}(a_{\eta\frown 0}/M_{\eta\frown 0}) \neq \mathrm{tp}(a_{\eta\frown 1}/M_{\eta\frown 0})$.

We are able to choose the models ω-special for clause 1) because of Lemma 13.22. Each model is the construction is a limit model; by Lemma 11.16, if M_0 is saturated, each model is saturated. Since $\boldsymbol{K}$ is stable there is a minimal type over some M_η. The type is big as every non-algebraic type over a saturated model is big by Lemma 13.19. We need the categoricity to invoke Lemma 11.16 and Lemma 13.19. $\square_{13.23}$

If $\boldsymbol{K}$ is categorical in λ, the use of limit models is unnecessary and one can obtain the following result just noting that under the categoricity hypothesis all models in the construction of Lemma 13.23 are saturated.

Corollary 13.24. *If $\boldsymbol{K}$ is categorical in λ and stable in λ then there is a minimal type $\hat{p} \in \mathbb{S}(M)$ where $|M| = \lambda$.*

Corollary 13.25. *Suppose $\boldsymbol{K}$ is λ-categorical and λ is regular. Suppose that $\boldsymbol{K}$ is $(\chi, \leq \lambda)$-weakly tame. Every minimal type p over a saturated model M with $\lambda \geq |M| > \mathrm{LS}(\boldsymbol{K})$ (and $|M| \geq \chi$) has a nonalgebraic extension, which is minimal, to any saturated extension N of M with $|N| \leq \lambda$.*

Proof. If $|M| = \lambda$, the result is given by Lemma 13.16. If $|M| < \lambda$, apply Theorem 13.19 to extend p to a non-algebraic type $\hat{p}$ over N. If $\hat{p}$ is not minimal, this is witnessed by extensions q_1, q_2 in some $\mathbb{S}(N_1)$ where $N \prec_{\boldsymbol{K}} N_1$. By weak tameness q_1 and q_2 differ on a submodel N_1' of size χ. By saturation we can choose $N_1' \subset N$ to contradict the minimality of p. $\square_{13.25}$

Using the observation that an increasing chain of extensions of a minimal type is coherent, Lessmann [101] provides the following extension.

Exercise 13.26. *Assume $\boldsymbol{K}$ is (χ, λ)-tame for some $\chi < \lambda$ and categorical for all κ with $\lambda^+ \leq \kappa \leq \mu$ and M is a saturated model of cardinality λ. If $p \in \mathbb{S}(M)$ is minimal then p has a unique nonalgebraic extension to any model of size μ.*

Our organization of the argument in Chapter 14 avoids the necessity of constructing rooted minimal types as in [46].

Remark 13.27. We are particularly grateful to David Kueker for his careful reading of this chapter. He not only corrected some arguments and filled some gaps but contributed significantly to organizing the argument more clearly.

The arguments in this chapter are applied below a categoricity cardinal and need the categoricity assumption. At several points we have had to invoke the embedding of models in an Ehrenfeucht-Mostowski model, which is guaranteed by categoricity.

14
Upward Categoricity Transfer

In this chapter we assume $\boldsymbol{K}$ has the amalgamation property and joint embedding and arbitrarily large models. Under the additional hypothesis that $\boldsymbol{K}$ is (χ, ∞)-tame for some $\chi < \lambda$, and $\lambda > \mathrm{LS}(\boldsymbol{K})$ we show that if $\boldsymbol{K}$ is λ and λ^+-categorical then it is categorical in all larger cardinalities. This result is due to Grossberg and VanDieren [46] although many elements are from Shelah [128]. We further prepare for the same result but weakening the hypothesis to just categoricity in λ^+ (Chapter 16) and to use these techniques for proving categoricity in a large enough λ^+ yields categoricity on an interval below λ^+ (Chapter 15).

Much of this chapter requires a tameness hypothesis. We are careful to use only weak tameness so the results apply in Chapter 15, where we work only from categoricity and apply Theorem 12.15.

Remark 14.1. *[Shelah's Categoricity Conjecture]* We would like to prove Shelah's conjecture that for every AEC $\boldsymbol{K}$, there is a cardinal κ such that categoricity in κ implies categoricity in all cardinalities greater than κ. In this monograph we restrict our ambition and further assume that $\boldsymbol{K}$ has the amalgamation property and arbitrarily large models. (As explained in the introduction to Part III and after Lemma 17.14 without significant loss of generality, we further assume $\boldsymbol{K}$ has the joint embedding property.) Under these hypotheses we prove some cases of the conjecture. The first result in this context is Shelah's proof [1] in [128] that categoricity in λ^+ with $\lambda \geq H_2$ (see introduction to Chapter 15) implies categoricity on the interval $[H_2, \lambda^+]$. We see this argument as taking place in two steps. 1) If $\boldsymbol{K}$ is λ^+ categorical then $\boldsymbol{K}$ is H_2-categorical. 2) If $\boldsymbol{K}$ is μ-categorical with

[1] Shelah asserted a marginally stronger result; see Remark 15.15.

$\mu > \mathrm{LS}(\boldsymbol{K})$ and (χ, μ)-weakly tame for some $\chi < \mu$ then $\boldsymbol{K}$ is μ^+-categorical. Following [128], we establish 1) in Chapter 15. But first we follow the lead of Grossberg and VanDieren [46], prove 2) and use 2) and induction on cardinality to obtain upward categoricity transfer. It is crucial for this argument to show that there is no (p, λ)-Vaughtian pair (Definition 14.2) with p a minimal type over a model of size λ. Using weak tameness and categoricity in λ as well as λ^+, this is not too difficult (Lemma 14.12). The nonexistence of (even true) (p, λ)-Vaughtian pairs is also fairly straightforward if we replace the categoricity in λ assumption by: the union of less than λ^+ saturated models of cardinality λ is saturated. And we proved this union theorem for singular λ in Lemma 11.17. Thus we are able to conclude this chapter by deducing upward categoricity transfer from categoricity in two successive cardinals or in the successor of a singular. The more difficult case of the successor of a regular $\lambda > \mathrm{LS}(\boldsymbol{K})$ is postponed to Chapter 16 where by two more applications of Ehrenfeucht-Mostowski models we are able to obtain the unions of chains of saturated models hypothesis.

A crucial tool here, and the reason that we must assume categoricity in a *successor cardinal* is the notion of a Vaughtian pair. I write $p(N)$ for the set of solutions of p in N.

Definition 14.2. *1. A* (p, λ)-Vaughtian pair *is a pair of models* $M \prec_{\boldsymbol{K}} N$ *and a* $p \in \mathbb{S}(M_0)$ *for a submodel* M_0 *of* M *such that p has a nonalgebraic extension to* $\mathbb{S}(M)$ *while* $p(M) = p(N)$, $M \neq N$, *and* $|M| = |N| = \lambda$.

2. A true (p, λ)-*Vaughtian pair is one where both* M *and* N *are saturated.*

When we have categoricity in successive cardinals the saturation of M and N is free; we have to make it explicit to handle categoricity from a single cardinal. The term 'true' Vaughtian pair was introduced in [44]; we follow the argument of [101] in moving between arbitrary and true Vaughtian pairs. We need to transfer the existence of a Vaughtian pair between cardinalities; downwards with the same base type p is very easy. But to get a true Vaughtian pair in the lower cardinality requires a further assumption.

Definition 14.3. $\boldsymbol{K}$ admits λ-saturated unions *if the union of less than* λ^+ *saturated models of cardinality* λ *is saturated.*

To admit λ-saturated unions is an approximation to superstability. We will show in Theorem 16.8 that in the context of Part III, if $\boldsymbol{K}$ is categorical in a regular λ then $\boldsymbol{K}$ admits μ-saturated unions for $\mathrm{LS}(\boldsymbol{K}) < \mu < \lambda$. We will flag our use of this hypothesis in Chapters 14 and 15.

Fact 14.4. *Let* $\boldsymbol{K}$ *be* $\leq \lambda$*-stable. Suppose there is a* (p, λ)*-Vaughtian pair.*

1. *For any* μ *with* $|\operatorname{dom} p| \leq \mu \leq \lambda$ *there is a* (p, μ)*-Vaughtian pair.*
2. *If the* (p, λ) *pair is true and* $\boldsymbol{K}$ *admits* μ*-saturated unions then the* (p, μ)*-Vaughtian pair can be chosen true.*

Proof. i) If (M, N) is a (p, λ)-Vaughtian pair, alternately choose $M_i \prec_{\boldsymbol{K}} M$ and $N_i \prec_{\boldsymbol{K}} N$ with $\operatorname{dom} p \subseteq M_0$, $N_0 - M \neq \emptyset$, and with $p(M_{i+1}) \supset p(N_i)$ for μ steps. Then (M_μ, N_μ) is a (p, μ)-Vaughtian pair.

ii) Repeat the same proof guaranteeing that each M_i is μ-saturated; p has a non-algebraic extension to M_μ by Lemma 13.14. $\square_{14.4}$

With somewhat more difficulty we transfer the existence of a Vaughtian pair of a non-splitting extension $\hat{p}$ of a minimal type p to the existence of a Vaughtian pair for p. Recall from Chapter 13 that a type $p \in \mathbb{S}(M)$ is minimal if there is at most one big extension of p to any model of the *same* cardinality. If $\boldsymbol{K}$ is (χ, λ^+)-weakly tame for some $\chi < \lambda^+$ and λ^+-categorical, Corollary 13.19 implies any minimal type over a saturated model M of cardinality $\leq \lambda$ has a minimal nonsplitting extension to any larger model *with cardinality at most* λ^+. The following result evolved from a lemma of VanDieren.

Lemma 14.5. *Let $\boldsymbol{K}$ be κ'-categorical for some regular $\kappa' \geq \kappa$. Suppose $|M|$ is a saturated model of cardinality $\mu > \mathrm{LS}(\boldsymbol{K})$, $M \prec_{\boldsymbol{K}} N$ and $|N| = \kappa$ with $\kappa \geq \mu$. Suppose $\boldsymbol{K}$ is $(\chi, [\mu, \kappa'])$-weakly tame for some $\chi < \mu$. Let $p \in \mathbb{S}(M)$ be minimal and $\hat{p} \in \mathbb{S}(N)$ be an extension of p which does not μ-split over M. If there is a $(\hat{p}, \kappa')$ Vaughtian pair then*

1. *there is a (p, κ') Vaughtian pair and*

2. *there is a (p, μ) Vaughtian pair.*

Proof. Suppose $N \prec_{\boldsymbol{K}} N_1 \prec_{\boldsymbol{K}} N_2$ are models of cardinality κ' and N is saturated. By Lemma 13.25, and the categoricity of $\boldsymbol{K}$ in κ', there is a nonsplitting and nonalgebraic extension q of $\hat{p}$ to N_1. If some $b \in N_2 - N_1$ realizes p but not $\hat{p}$, then, by weak tameness, $\hat{p}$ and $\mathrm{tp}(b/N)$ violate the minimality of p. So if (N_1, N_2) witness the existence of $(\hat{p}, \kappa)$-Vaughtian pair, p is omitted in $N_2 - N_1$. But then Fact 14.4 yields a (p, μ)-Vaughtian pair and we finish. $\square_{14.5}$

Lemma 14.6. *Suppose $\boldsymbol{K}$ is categorical in λ^+, $\boldsymbol{K}$ is (χ, λ)-weakly tame for some $\chi < \lambda$, and $\boldsymbol{K}$ admits λ-saturated unions. Then for any saturated $M \in \boldsymbol{K}$ with cardinality λ, and for any minimal $p \in \mathbb{S}(M)$, there is no true (p, λ)-Vaughtian pair.*

Proof. For contradiction, let $M \prec_{\boldsymbol{K}} N_0 \prec_{\boldsymbol{K}} N_1$, N_0 and N_1 saturated models of size λ, and $p_0 \in \mathbb{S}(N_0)$ be a nonalgebraic extension of p which is not realized in N_1. Then p_0 is also minimal.

Choose N, N' of size χ with $N \prec_{\boldsymbol{K}} N' \prec_{\boldsymbol{K}} N_0$ such that p_0 does not $|N|$-split over N and N' is universal over N. (N exists by Lemma 13.5, N' by stability, Lemma 11.3, and Lemma 11.5; we can assume $N' \prec_{\boldsymbol{K}} N_0$ by saturation of N_0, since $\lambda > |N'|$). We now construct an increasing and continuous sequence of saturated models $(N_i : i < \lambda^+)$ of size λ and types $p_i \in S(N_i)$ such that p_i is an N'-automorphic image of p_0 which is not realised in N_{i+1}.

This is possible: Having constructed N_i, choose $f_i \in Aut(\mathbb{M}/N')$ such that $f_i(N_0) = N_i$ (this is possible since $\lambda > LS(\boldsymbol{K})$). Let $p_i = f_i(p_0)$ and $N_{i+1} =$

$f_i(N_1)$. At limits, let N_i be the union of the N_j for $j < i$, which is saturated by the hypothesis on unions of chains of saturated models.

Notice that since each p_i is an N'-automorphic image of p_0, each p_i does not split over N and $p_i \restriction N' = p_0 \restriction N'$. Thus, by uniqueness of nonsplitting extensions (Lemma 13.7), $p_i \restriction N_0 = p_0$. Thus, each p_i is the unique nonalgebraic extension of p_0 in $\mathbb{S}(N_i)$.

This contradicts λ^+-categoricity: Let $M^* = \bigcup_{i<\lambda^+} N_i$. Then M^* does not realise p_0: Otherwise, there is $a \in M^*$ realising p_0. Since $a \notin N_0$, there is $i < \lambda^+$ such that $a \in N_{i+1} \setminus N_i$. Then $\mathrm{tp}(a/N_i)$ is a nonalgebraic extension of the minimal type p_0. Hence $\mathrm{tp}(a/N_i) = p_i$ which is a contradiction, since p_i is not realised in N_{i+1}. $\square_{14.6}$

Of course the 'union of chains' hypothesis of Lemma 14.6 is immediate if $\boldsymbol{K}$ is λ-categorical so there is no need to introduce 'true' to study transfer of categoricity from two successive cardinals. The weaker hypothesis stated for Lemma 14.6 plays a crucial role in extending the argument here to assuming categoricity in a single successor cardinal after we establish that $\boldsymbol{K}$ admits λ-saturated unions in Chapter 16. We need closure under unions of chains of saturated models rather than just uniqueness of limit models because N_1 is an arbitrary saturated extension of N_0. Indeed N_1 cannot be universal over N_0 since they are a Vaughtian pair; so the saturation of the limit models in the N_i sequence cannot be deduced from uniqueness of limit models (which are the limit of a chain of *universal* extensions).

Lemma 14.7. *Suppose $\boldsymbol{K}$ is (χ, λ)-weakly tame for some $\chi < \lambda$. Let $p \in \mathbb{S}(M)$, $|M| = \lambda$ where p is minimal and M is saturated and suppose there is no (p, λ)-Vaughtian pair. Then any N with $|N| = \lambda^+$ and $M \prec_{\boldsymbol{K}} N$ has λ^+ realizations of p. Moreover, if $\boldsymbol{K}$ admits λ-saturated unions, 'no true (p, λ)-Vaughtian pair' is a sufficient hypothesis.*

Proof. Since M is saturated there is a non-algebraic extension of p to any extension of cardinality λ (Theorem 13.19). If N has fewer than λ^+ realizations of p there is a (p, λ^+) Vaughtian pair and thus, by Fact 14.4.1, a (p, λ) Vaughtian pair. For the moreover, use Fact 14.4.2. $\square_{14.7}$

The previous result shows that if we have categoricity up to κ then all *minimal* types over a model of cardinality κ are realized in any extension of cardinality κ^+. The crucial step, due to Grossberg and VanDieren, (Theorem 4.1 of [46]), extends this result to all types. They introduce the following notion.

Definition 14.8. *N admits a (p, λ, α)-resolution over M if $|N| = |M| = \lambda$ and there is a continuous increasing sequence of models M_i with $M_0 = M$, $M_\alpha = N$ and a realization of p in $M_{i+1} - M_i$ for every i.*

Lemma 14.9. *Assume $p \in \mathbb{S}(M)$ is minimal and $\boldsymbol{K}$ does not admit a (p, λ)-Vaughtian pair. If N admits a (p, λ, α)-resolution over M, with $\alpha = \lambda \cdot \alpha$ then N realizes every $q \in \mathbb{S}(M)$.*

Proof. Fix a (p, λ, α)-resolution $\langle N_i : i \leq \alpha \rangle$ of N. We construct an increasing continuous chain of models M_i and M'_i for $i \leq \alpha$ such that M'_0 realizes q and we show N is isomorphic to M_α by constructing isomorphisms f_i from N_i into M'_i. We denote the image of f_i by M_i. In fact, by a Vaughtian pair argument we will get $M_\alpha \approx M'_\alpha$. Let $M_0 = N_0 = M$ and f_0 be the identity on M. Take unions at limits.

Let M'_0 be any extension of M_0 which contains realizations of both p and q; amalgamation guarantees the existence of M'_0.

The following coding trick allows us to catch up. We can write α as the union of disjoint sets $\langle S_i : i < \alpha \rangle$ such that $|S_i| = \lambda$ and $\min(S_i) > i$ (except at 0), i.e. $0 \in S_0$ and $\min S_i > i$ for $0 < i < \alpha$. At stage i, let $\widehat{S}_i = \langle b_\zeta : \zeta \in S_i \rangle$ enumerate $p(M'_i)$, (possibly with repetitions).

At successor stages, $i + 1$, we guarantee b_i is in the range of f_{i+1}. We have an isomorphism f_i taking N_i onto $M_i \prec_{\boldsymbol{K}} M'_i$. We have defined $\{b_j : j \in S_k\}$ for $k \leq i$. Since $\min S_t \geq t$ for all t, $b_i \in \bigcup_{t \leq i} \widehat{S}_t$. By the amalgamation property there exists a model M^* which contains M'_i and f extending f_i and mapping N_{i+1} into M^*.

Now if $b_i \in M_i$, our goal is met. Just let M'_{i+1} be M^* and take f_{i+1} to be f and M_{i+1} its image.

Suppose $b_i \notin M_i$. By assumption there is a $c \in N_{i+1} - N_i$ that realizes p. Note that $\mathrm{tp}(f(c)/M_i) = f(\mathrm{tp}(c/N_i))$ is a non-algebraic extension of p. We know $b_i \in M'_i - M_i$ realizes another nonalgebraic extension of p. So $\mathrm{tp}(f(c)/M_i) = \mathrm{tp}(b_i, M_i)$. By the definition of Galois type we can choose M'_{i+1} is an extension of M'_i and g mapping M^* into M'_{i+1} with $g(f(c)) = b_i$. Then $f_{i+1} = g \circ f$ as an embedding of N_{i+1} into M'_{i+1} with $f_{i+1}(c) = b_i$, as required.

Now, we claim f_α maps N onto M'_α. By the construction the image M_α of f_α satisfies $p(M_\alpha) = p(M'_\alpha)$ and $M_\alpha \prec_{\boldsymbol{K}} M'_\alpha$ by **A3.3** But then $M_\alpha = M'_\alpha$ lest (M_α, M'_α) be a (p, λ) Vaughtian pair. Since f_α fixes M, q is realized in N, as required. $\square_{14.9}$

From this result we easily deduce the less technical statement in Theorem 14.11. The assumption that there is a saturated model in μ is needed when μ is singular. In the proof of categoricity on $[H_2, \lambda^+]$, we will guarantee the saturation by induction and ground the induction with Theorem 15.12.

Exercise 14.10. *Suppose $M \prec_{\boldsymbol{K}} N$ and $p \in \mathbb{S}(M)$ is realized μ^+ times in N. Show there is N' with $M' \prec_{\boldsymbol{K}} N' \prec_{\boldsymbol{K}} N$ such that N' admits a (p, μ, α)-decomposition.*

Theorem 14.11. *Assume $\boldsymbol{K}$ is (χ, μ)-weakly tame for some $\chi < \mu$. Suppose $\boldsymbol{K}$ is μ-categorical and there is a saturated model M in cardinality μ such that there is $p \in \mathbb{S}(M)$ that is minimal, and there is no (p, μ)-Vaughtian pair. Then every model of cardinality μ^+ is saturated.*

Proof. Let $N \in \boldsymbol{K}$ have cardinality μ^+. Choose any $M' \prec_{\boldsymbol{K}} N$ with cardinality μ. By categoricity in μ, M' is saturated. We will show every type over M' is realized in N. Fix α with $\alpha \cdot \mu = \mu$. By Lemma 14.7, p is realized μ^+ times in N. By Exercise 14.10, there is an N' with $M' \prec_{\boldsymbol{K}} N' \prec_{\boldsymbol{K}} N$ such that N' admits a (p, μ, α)-decomposition. By Lemma 14.9, N' and *a fortiori* N realize every type over M' and we finish. $\square_{14.11}$

We next show the hypothesis of Theorem 14.11 easily transfers to larger categoricity cardinals; this fact will fuel the induction step of Theorem 14.13.

Lemma 14.12. *Suppose there is a saturated model M in cardinality μ such that there is $p \in \mathbb{S}(M)$ that is minimal, and there is no (p, μ)-Vaughtian pair. Now fix $\kappa \geq \mu^+$ such that all models of $\boldsymbol{K}$ with cardinality κ are saturated and $\boldsymbol{K}$ is (χ, κ)-weakly tame for some $\chi < \mu$. For any $N \in \boldsymbol{K}_\kappa$, there is a minimal non-algebraic type $\hat{p} \in \mathbb{S}(N)$ such that there is no $(\hat{p}, \kappa)$-Vaughtian pair.*

Proof. By Lemma 13.25, which applies by weak tameness, there is a nonsplitting nonalgebraic extension $\hat{p}$ of p to $\mathbb{S}(N)$ that is minimal. Applying the contrapositive of Lemma 14.5, there is no $(\hat{p}, \kappa)$-Vaughtian pair. *This is a crucial application of tameness.* $\square_{14.12}$

We now conclude by induction the result (Theorem 14.13.ii) of Grossberg and VanDieren. Our formulation of Theorem 14.13.i) prepares for the proof from categoricity in a single successor.

Theorem 14.13. *Suppose $\mathrm{LS}(\boldsymbol{K}) < \lambda$, and that $\boldsymbol{K}$ has arbitrarily large models, the amalgamation property, and joint embedding.*

1. *Suppose $\boldsymbol{K}$ is λ^+-categorical, (χ, κ')-tame for some $\chi < \lambda$, and $\boldsymbol{K}$ admits λ-saturated unions. Further, suppose there is a saturated model M in cardinality λ such that there is $p \in \mathbb{S}(M)$ that is minimal, and there is no (p, λ)-Vaughtian pair. Then every model of cardinality at least λ^+ is saturated and so $\boldsymbol{K}$ is categorical in all cardinals κ with $\kappa' \geq \kappa \geq \lambda^+$.*
2. *In particular, if $\boldsymbol{K}$ is (χ, ∞)-tame for some $\chi < \lambda$, and $\boldsymbol{K}$ is both λ and λ^+-categorical then $\boldsymbol{K}$ is categorical in all cardinals $\kappa \geq \lambda$.*

Proof. For i) we prove by induction on $\kappa \geq \lambda^+$ that every model in $\boldsymbol{K}$ of cardinality at least λ^+ is saturated. Suppose the result is true below κ.

If κ is a limit cardinal, we show an arbitrary $M \in \boldsymbol{K}$ with $|M| = \kappa$ is Galois saturated. Let $N \prec_{\boldsymbol{K}} M$ with $|N| < |M|$. By Löwenheim-Skolem and hypothesis choose N' with $N \prec_{\boldsymbol{K}} N' \prec_{\boldsymbol{K}} M$, $|N| < |N'| < |M|$, $|N'| \geq \lambda^+$, and with $|N'|$ is regular, By induction $\boldsymbol{K}$ is $|N'|$-categorical and so N' is saturated. So every type over N is realized in M. I.e., M is saturated as required.

If κ is a successor cardinal, say $\kappa = \mu^+$, by Lemma 14.12 and the induction hypothesis, the hypotheses of Theorem 14.11 hold at μ. So, we finish by applying Theorem 14.11 to conclude κ-categoricity.

For ii) we show that the categoricity in λ and λ^+ implies the hypotheses of i) hold in λ^+. Let M_0 be the saturated model of cardinality λ. By Theorem 13.23 and categoricity, there is a minimal type $p \in \mathbb{S}(M_0)$. By Lemma 14.6, there is no (p, λ)-Vaughtian pair. Now apply Lemma 14.12 to get this condition in λ^+. $\square_{14.13}$

Categoricity in λ is used twice in this argument: to find a minimal type over a saturated model (which is thus extendible) and to show there is no Vaughtian pair in λ. Both of these results follow easily once we know the closure of the class of saturated models under unions of chains. So using Lemma 11.17, we have:

Corollary 14.14. *Suppose* $\mathrm{LS}(\boldsymbol{K}) < \lambda$, $\boldsymbol{K}$ *is* (χ, ∞)*-weakly tame for some* $\chi < \lambda$*, and that* $\boldsymbol{K}$ *has arbitrarily large models, the amalgamation property, and joint embedding. If* $\boldsymbol{K}$ *is* λ^+*-categorical and* λ *is singular then* $\boldsymbol{K}$ *is categorical in all cardinals* $\mu > \lambda$.

Proof. Choose a saturated M_0 of cardinality λ by Lemma 11.16. By Lemma 13.23) there is a minimal type in $\mathbb{S}(M_0)$. By Lemma 13.25, this minimal type is extendible. Finally, we show there is no Vaughtian pair, by adapting the proof of Lemma 14.6, replacing the use of categoricity in λ by the closure of saturated models under unions of chains. $\square_{14.14}$

We will improve this result to categoricity in one sucessor cardinal in Chapter 16.

Remark 14.15. These arguments were refined from Shelah [128] and Grossberg-Vandieren [46] by the author and others. In particular, the idea that assuming tameness would allow the proof of upwards categoricity and the proof of Lemma 14.9 is due to Grossberg and VanDieren. The current argument for Lemma 14.6 is due to Lessmann. Lessmann [101] (see also [16] has also extended the result to the case $\lambda = \mathrm{LS}(\boldsymbol{K}) = \aleph_0$. This is a crucial extension since as noted in Chapters 3and 4, there are actual mathematical situations in countable languages where $\aleph_1$-categoricity is easier to establish; however, the tameness remains open in these interesting cases. The notation (p, λ)-Vaughtian pair was introduced in early versions of this book. This is the first explicit statement of Theorem 14.14 although it is implicit in [128].

15
Omitting types and Downward Categoricity

Our goal is to show that if an AEC is categorical on a sufficiently large successor cardinal, it is in fact categorical for all sufficiently large cardinals. (See Conclusion 16.13.) In view of Theorem 5.20, the key to this is to prove a downwards categoricity theorem. For this we need to make precise some notation. In Notation 5.23 we introduced the function $H(\kappa) = \beth_{(2^\kappa)^+}$. We noted in Corollary 5.24 that the Hanf number for omitting types in a $PC\Gamma$-class with vocabulary τ is bounded by $H(|\tau|)$. We showed in Corollary 5.25 the Hanf number for the existence of models in an AEC with fixed $\kappa_{\boldsymbol{K}}$ is bounded by $H_1 = H(\kappa_{\boldsymbol{K}})$. We write H_2 for $H(H(\kappa_{\boldsymbol{K}}))$. The significance of H_2 is a major technical point of this chapter that we expound in Remark 15.7.

In this chapter we prove that if an AEC $\boldsymbol{K}$ satisfies AP and JEP and is categorical in λ and λ^+ for some $\lambda \geq H_2$ then it is categorical in all $\mu \in [H_2, \lambda^+]$. Assuming $\boldsymbol{K}$ is (χ, ∞)-tame for some $\chi < H_1$ allows one to extend this conclusion to categoricity in all $\mu \geq H_2$. (We discuss the relation to the somewhat stronger result asserted by Shelah in Remark 15.15.) The argument for downward transfer proceeds in several steps. We showed in Theorem 12.15 that for each κ with $H_1 \leq \kappa < \lambda^+$, there is a χ_κ such that $\boldsymbol{K}$ is (χ_κ, κ)-weakly tame for some $\chi_\kappa < H_1$. The fundamental fact in Chapter 14 is that if for some μ greater than the tameness cardinal χ (e.g. if $\boldsymbol{K}$ is (χ, ∞)-tame) $\boldsymbol{K}$ is μ-categorical and has a model of M with $|M| = \mu$ and a minimal type $p \in \mathbb{S}(M)$ such that there is no (p, μ)-Vaughtian pair then $\boldsymbol{K}$ is categorical in all larger cardinals. The content of this chapter is to move from categoricity in λ and λ^+ with $\lambda > H_2$ to getting the conditions of the preceding sentence with $\mu = H_2$.

Theorem A.3 is a general statement of 'Morley's omitting types theorem. We make several applications of that result in this chapter. The first (Lemma 15.2) uses Theorem A.3.1 and allows us to conclude: each model in $\boldsymbol{K}$ of cardinality at least H_1 is $\mathrm{LS}(\boldsymbol{K})$-saturated and each model of cardinality at least H_2 is H_1-saturated; the second (Theorem 15.9, which depends on Theorem 15.8), using Theorem A.3.2b, with the tameness, extends this to: each model in $\boldsymbol{K}$ of cardinality H_2 is H_2-saturated. Thus we have transferred λ^+-categoricity down to H_2-categoricity; now we work our way back up. As in Chapter 14, we are able to find a minimal type p based on a model M^* of cardinality H_2. A third variant, Theorem A.3.2a, on the omitting types theorem shows (Theorem 15.12) that if there is a (p, H_2)-Vaughtian pair then there is a (p, λ)-Vaughtian pair. We showed there was no (p, λ)-Vaughtian pair in Lemma 14.6. Finally applying the induction from Lemma 14.13 beginning at H_2 we conclude that $\boldsymbol{K}$ is categorical up to λ^+ and, assuming full tameness, in all larger cardinals.

In the next chapter we will use a far more sophisticated investigation of non-splitting to show that if $\boldsymbol{K}$ is λ^+-categorical (with no categoricity hypothesis in λ), there is still no (p, λ)-Vaughtian pair. Thus we obtain the full strength of both [128] and [46].

Notation 15.1. *We fix this notation for the chapter. $\boldsymbol{K}$ is an AEC in a vocabulary τ and $|\tau| = \mathrm{LS}(\boldsymbol{K})$. τ_1 is the expansion of τ given by the presentation theorem. $\mu = (2^{|\tau|})^+$.*

First we extend Morley's omitting types theorem, Theorem A.3.1 to Galois types.

Lemma 15.2. *[II.1.5 of 394] Let $\boldsymbol{K}$ be an AEC in a vocabulary τ of cardinality κ. If $M_0 \prec_{\boldsymbol{K}} M$ and $|M| \geq H(|M_0| + \kappa)$, we can find an EM-set Φ such that the following hold.*

1. *The τ-reduct of the Skolem closure of the empty set is M_0.*
2. *For every I, $M_0 \prec_{\boldsymbol{K}} EM_\tau(I, \Phi)$.*
3. *If I is finite, $EM_\tau(I, \Phi)$ can be embedded in M over M_0 as a τ_1'-structure.*
4. *$EM_\tau(I, \Phi)$ omits every galois type over M_0 which is omitted in M.*

Proof. Let τ_1 be the Skolem language given by the presentation theorem and consider M as the reduct of a τ_1-structure M^1. Add constants for M_0 to form τ_1'. Extend T_1 from the presentation theorem to a Skolemized T_1' by fixing the values of the functions on M_0. Now apply Lemma 7.2.3 to find an EM-diagram Φ (in τ_1') with all τ-types of finite subsets of the indiscernible sequence realized in M. Now 1) and 2) are immediate. 3) is easy since we chose Φ so all finite subsets of the indiscernible set (and so their Skolem closures) are realized in M.

The omission of Galois types is more tricky. Consider both M and $N = EM_\tau(I, \Phi)$ embedded in $\mathbb{M}$. Let N^1 denote the τ_1'-structure $EM(I, \Phi)$. We need to show that if $a \in N$, $p = \mathrm{tp}(a/M_0)$ is realized in M. For some $\mathbf{e} \in I$, a is in the τ_1-Skolem hull $N^{\mathbf{e}}$ of $\mathbf{e}$. By 3) there is an embedding α of $N^{\mathbf{e}}$ into M^1 over M_0. α is also an isomorphism of $N^{\mathbf{e}} \restriction \tau$ into M. Now, by the model homogeneity, α extends to an automorphism of $\mathbb{M}$ fixing M_0 and $\alpha(a) \in M$ realizes p. $\square_{15.2}$

Now we can rephrase this result as

Corollary 15.3. *The Hanf number for omitting Galois types over a model in any AEC with a vocabulary and Löwenheim-Skolem number bounded by κ is at most $H(\kappa)$.*

This result has immediate applications in the direction of transferring categoricity.

Theorem 15.4. *If $\boldsymbol{K}$ is categorical in a regular cardinal $\lambda \geq H(|M_0|)$ then every $M \in \boldsymbol{K}$ with $|M| \geq H(|M_0|)$ is $|M_0|$-saturated. Thus, if $|M| \geq H_1$, M is $\mathrm{LS}(\boldsymbol{K})^+$-saturated and if $|M| \geq H_2$, M is H_1-saturated.*

Proof. By the last lemma, if M omits a type p over a model of size $|M_0|$, there is a model $N \in \boldsymbol{K}$ with cardinality λ which omits p. But, the unique model of power λ is saturated. $\square_{15.4}$

But we need to show that the model in H_2 is saturated. We must use the 'two cardinal' aspect laid out in part 2 of Theorem A.3. The argument depends essentially on the fact that H_2 is a μ-collection cardinal in the following sense.

Definition 15.5. *For any μ, κ, κ is a* μ-collection cardinal *if for every $\chi < \kappa$, $\beth_\mu(\chi) \leq \kappa$.*

Thus in our case where μ is fixed as $(2^{LS(\boldsymbol{K})})^+$; for any non-zero α, cardinals of the form $H_2 \times \alpha$ are μ-collection cardinals.

We need the following lemma because we are only able to deduce weak-tameness, as opposed to tameness, from categoricity. The proof is patterned on that of Lemma 11.17.

Lemma 15.6. *Suppose $\boldsymbol{K}$ is λ-categorical, $M \in \boldsymbol{K}$, $|M| < \mathrm{cf}(\lambda)$ and M is μ-saturated for some μ with $\mathrm{LS}(\boldsymbol{K}) < \mu < |\mathrm{M}|$. Then if $N \prec_{\boldsymbol{K}} M$ with $|N| = \mu$, there is an N' with $N \prec_{\boldsymbol{K}} N' \prec_{\boldsymbol{K}} M$, $|N'| = \mu$ and N' is saturated.*

Proof. Write $N = \bigcup_{i<\mathrm{cf}(\mu)} N_i$ where $|N_i| = \mu_i < \mu$.

If μ is regular the result is easy to see. Construct a continuous increasing chain N'_i with $|N'_i| = \mu_i$ for $i < \mu$ with $N_i \prec_{\boldsymbol{K}} N'_i$ and each N'_{i+1} realizes all types over N'_i. Using regularity it is easy to see that $N' = \bigcup_{i<\mu} N'_i$ is saturated.

For singular μ the argument is more complicated; we work by induction on μ. So let μ be the least cardinal where the conclusion fails. Now applying the induction hypothesis, construct a continuous increasing chain N'_i with $|N_i| = \mu_i$ for $i < \mu$ with $N_i \prec_{\boldsymbol{K}} N'_i$ and each N'_i is μ_i-saturated. To see that $N' = \bigcup_{i<\mu} N'_i$ is saturated, choose $N^* \prec_{\boldsymbol{K}} N'$; say, with $|N^*| = \kappa < \mu$. If $\kappa < \mathrm{cf}(\mu)$, the usual argument for the regular case shows every type over N^* is realized in N'. If $\mathrm{cf}(\mu) \leq \kappa$, we will construct $\hat{N}$ with $N^* \prec_{\boldsymbol{K}} \hat{N} \prec_{\boldsymbol{K}} N$ so that $\hat{N}$ is κ^+-saturated. Let $X = \{i < \mu : (N'_{i+1} \cap N^*) - N'_i \neq \emptyset\}$. Enumerate X as $\langle x_i : i < \delta\rangle$ where $\delta < \kappa^+$. Now choose a continuous increasing chain $\hat{N}_\gamma$ so that for sufficiently large γ, $(N^* \cap N'_{x_\gamma}) \subseteq \hat{N}_\gamma \subset N'_{x_\gamma}$ and $\hat{N}_{\gamma+1}$ is κ^+-universal over $\hat{N}_\gamma$. This is possible by the κ^{++}-saturation of N'_i for sufficiently large i. Now,

$\langle \hat{N}_\gamma : \gamma < \delta\rangle$ is a (κ^+, δ)-tower. So by Lemma 11.15.3 and Corollary 11.13.2 $\hat{N}_\delta \prec_{\boldsymbol{K}} N'$ is κ^+-saturated and is the required $\hat{N}$. $\Box_{15.6}$

Remark 15.7. The following proof revolves around three cardinals: κ, κ^1, χ; a fourth cardinal θ appears in the statement to help calibrate the others. We will start with a model M of cardinality κ that is not saturated. We say $N^1 \prec_{\boldsymbol{K}} M$ witnesses the non-saturation if some type q over N^1 is omitted in M. We choose $\kappa_1 < \kappa$ so there is a model of N^1 of size κ_1 witnessing non-saturation. We will also need that $\boldsymbol{K}$ is (χ, κ^1)-weakly tame and $\chi < \kappa^1$. To apply *weak tameness*, we need that N^1 is itself saturated; to guarantee this, we assume M is θ-saturated. We describe the value of θ in a moment. We will expand M to a τ^+-structure in an extended language τ^+ with $|\tau^+| = \chi$. The properties of the expansion are described in the first displayed list in the following proof. Then we apply Theorem A.3.2b to construct an EM-diagram Φ in a language extending τ^+ by an additional unary predicate P such that for any $M^* = EM(I, \Phi)$, there a sequence of submodels N_i for $i < \omega$ (which are isomorphic to submodels of the original N^1) so that q induces a type over $N = \bigcup N_i$ that is 'strongly omitted'. The 'strong omitting' guarantees that in any cardinality $> \chi$ there is a model that is not χ^+-saturated. This contradicts categoricity.

We apply Theorem 15.8 in Theorem 15.9 with $\theta = H_1$. To guarantee that M is H_1-saturated, we need $|M| \geq H_2$. And $\kappa_1 \geq \theta$ had to be at least H_1 so that we could assert that $\boldsymbol{K}$ is (χ, κ_1)-weakly tame for some $\chi < H_1$.

The following argument uses both first order types and Galois types. The notation tp refers to Galois types and we use English for first order types. Recall our conventions in Notation 15.1. Below we discuss the structures $EM(\omega, \Phi)$ and $EM(n, \Phi)$; note that this structure is given by Φ and ω or n; it does not depend on embedding into a larger index set.

The following Lemma showing the transfer of non-saturation will be used by contraposition to show the transfer of categoricity.

Theorem 15.8. *[II.1.6 of 394] Suppose $\boldsymbol{K}$ is λ-categorical and $|M| = \kappa$ where κ is a μ-collection cardinal. Suppose further that for some θ, $\boldsymbol{K}$ is $(<\chi, [\theta, \kappa])$-tame. If M is not saturated but is θ-saturated then there is a τ^+ with $|\tau^+| = \chi < \theta$ and a τ^+-diagram Φ such that for every ordered set I: there is a Galois-type q' over a model $N' \prec_{\boldsymbol{K}} M$ with $|N'| = \chi$ that is omitted in $EM_\tau(I, \Phi)$. In particular, for every λ', there is a model of cardinality λ' that is not χ-saturated.*

Proof. If M is not saturated, there is an $N^1 \prec_{\boldsymbol{K}} M$ with $|N^1| < \kappa$ and $q \in \mathbb{S}(N^1)$ which is omitted in M. We have M is θ-saturated, so by Lemma 15.6, by choosing the cardinality κ^1 of N^1 to be minimal with N_1 witnessing the non-saturation, we may assume that N^1 is itself saturated. Since $\kappa^1 \in [\theta, \kappa]$, by (χ, κ^1) weak-tameness, for every $c \in M$, there is an $M_c \prec_{\boldsymbol{K}} N^1$ with $|M_c| = \chi$ such that $\mathrm{tp}(c/M_c) \neq q \restriction M_c$.

We first expand τ_1 to a language τ^+ by adding a unary predicate P and χ additional functions of each finite arity. Then we expand M to a τ^+-structure with the following properties. Fix a model $N^0 \prec_{\boldsymbol{K}} N^1$ with $|N^0| = \chi$.

1. $P(M) = N^1$;
2. $\mathrm{cl}_{\tau^+}(\emptyset) \restriction \tau = N^0$;
3. For every $\boldsymbol{a} \in M$,
   (a) $\mathrm{cl}_{\tau^+}(\boldsymbol{a}) \restriction \tau \prec_{\boldsymbol{K}} M$;
   (b) $(\mathrm{cl}_{\tau^+}(\boldsymbol{a}) \cap N^1) \restriction \tau \prec_{\boldsymbol{K}} N^1$;
   (c) if $\boldsymbol{a} \subset \mathbf{b}$ then $\mathrm{cl}_{\tau^+}\boldsymbol{a} \restriction \tau \prec_{\boldsymbol{K}} \mathrm{cl}_{\tau^+}\mathbf{b} \restriction \tau$;
   (d) $q \restriction (\mathrm{cl}_{\tau^+}(\boldsymbol{a}) \cap N^1) \restriction \tau$ is omitted in $\mathrm{cl}_{\tau^+}(\boldsymbol{a}) \restriction \tau$.

To make this expansion we mimic the proof of the presentation theorem writing N^1 and M as direct unions of finitely generated models in the expanded vocabulary τ^+ and meeting certain additional requirements. We interpret the symbols of τ^+ in stages. On the first pass define by induction on the length of $\boldsymbol{a}$ for each finite $\boldsymbol{a} \in M$, partial τ^+-structures $M^0_{\boldsymbol{a}}$ such that $M^0_{\boldsymbol{a}} = \mathrm{cl}_{\tau_1}(\boldsymbol{a})$, $|M^0_{\boldsymbol{a}} \cap N^1| = \chi$ such that if $\boldsymbol{a} \in N^1$ then $(M^0_{\boldsymbol{a}} \cap N^1) \restriction \tau \prec_{\boldsymbol{K}} N^1 \restriction \tau$; of course we require that if $\boldsymbol{a} \subset \mathbf{b}$ then $M_{\boldsymbol{a}} \restriction \tau \prec_{\boldsymbol{K}} M_{\mathbf{b}} \restriction \tau$ and for $\boldsymbol{a} \in N^1$, $M^0_{\boldsymbol{a}} \prec_{\boldsymbol{K}} N^1$.

On a second pass, we interpret, again in stages, the remaining symbols from τ^+ to guarantee that for $M_{\boldsymbol{a}} = \mathrm{cl}_{\tau^+}(\boldsymbol{a})$, $q \restriction (M_{\boldsymbol{a}} \cap N^1)$ is not realized in $M_{\boldsymbol{a}}$. We can guarantee this in countably many steps. We view τ^+ as having $\chi \times \omega$ functions beyond τ_1 and P; at each step χ of them are interpreted and $M^i_{\boldsymbol{a}}$ denotes the closure of $\boldsymbol{a}$ under the first $\chi \times i$ functions. For each $c \in M^i_{\boldsymbol{a}}$ form $M^{i+1}_{\boldsymbol{a}}$ by adding points to $M_{\boldsymbol{a}} \cap N^1$ to guarantee c doesn't realize $q \restriction M^{i+1}_{\boldsymbol{a}}$. By tameness we need add at most χ points; we interpret the next χ functions in τ^+ so that when evaluated at $\boldsymbol{a}$ they enumerate these points. Again we require that $M^i_{\boldsymbol{a}} \restriction \tau \prec_{\boldsymbol{K}} M^i_{\mathbf{b}} \restriction \tau$, for $\boldsymbol{a} \subseteq \mathbf{b}$, $M^i_{\boldsymbol{a}} \restriction \tau \prec_{\boldsymbol{K}} M^i_{\mathbf{b}} \restriction \tau$, and that for $\boldsymbol{a} \in N^1$, $M^i_{\boldsymbol{a}} = \prec_{\boldsymbol{K}} N^1$. Further, demand in the construction that if $\boldsymbol{a} \in N^1$, $\mathrm{cl}_{\tau^+}(\boldsymbol{a}) \subseteq N^1$.

Now we construct a τ^+-diagram Φ so that for any $M^* = EM(I, \Phi)$, we can find $N_j \prec_{\boldsymbol{K}} M^*$ for $j \leq \omega$ so that if $N = \bigcup_j N_j$ and $q' = q \restriction N$:

1. $EM_\tau(n, \Phi)$ is isomorphic to a submodel of M, so $EM_\tau(I, \Phi) \in \boldsymbol{K}$.
2. $P(M^*) \restriction \tau \prec_{\boldsymbol{K}} M^* \restriction \tau$.
3. $N_0 \prec_{\boldsymbol{K}} P(M^*)$ has cardinality χ; $N_0 \approx N^0 \prec_{\boldsymbol{K}} N^1$.
4. N_j for $j \leq \omega$ is a continuous chain of τ-structures with cardinality χ, beginning with N_0 and all strong in $P(M^*)$.
5. The τ-reduct of the Skolem closure of the empty set is N_0.
6. $N_n \prec_{\boldsymbol{K}} EM_\tau(n, \Phi)$.

7. For every I, $EM_\tau(I,\Phi)$ omits q' in the strong sense that if $c \in EM_\tau(I,\Phi)$, then for some $j < \omega$,
$$q' \restriction N_j \neq (\mathrm{tp}(c/N_j), EM_\tau(I,\Phi)).$$

If we can establish these claims, $EM_\tau(I,\Phi)$ satisfies the conclusions of the Theorem. Namely, $EM_\tau(\omega,\Phi) \prec_{\boldsymbol{K}} EM_\tau(I,\Phi)$, $|EM_\tau(\omega,\Phi)| = \chi$ and the copy of q' over $EM_\tau(\omega,\Phi)$ is omitted in $EM_\tau(I,\Phi)$.

To establish the seven conditions, we apply Theorem A.3.2b and for this we need $\beth_{(2^{|\tau|})^+}(|N^1|) < \kappa$. This is exactly what it means for κ to be a μ-collection cardinal.

Apply Theorem A.3.2b, taking M for all the M_α and $P(M)$ as N^1 to satisfy the first two conditions. With the resulting Φ, in any $M^* = EM(I,\Phi)$ there is a countable set J of τ^+-indiscernibles over $P(M^*)$. (The cardinality of $Q(M^*)$ is not important here.) So for any pair of increasing n-tuples, $\boldsymbol{a}, \mathbf{b}$ from J, $\mathrm{cl}_{\tau^+}(\boldsymbol{a}) \cap P(M^*) = \mathrm{cl}_{\tau^+}(\mathbf{b}) \cap P(M^*)$; we call this common intersection N_n. Note that $N_0 = \mathrm{cl}_{\tau^+}(\emptyset) \prec_{\boldsymbol{K}} P(M^*)$ and $N_0 \approx N^0$. Each N_n has cardinality χ. By definition, each N_n is a τ^+-substructure of $P(M^*)$. Since $P(M^*)$ is a τ^+-substructure of M^* and the moreover clause of the presentation theorem, $N_n \prec_{\boldsymbol{K}} EM(I,\Phi)$. Note that $N_n \subset EM(n,\Phi)$. So $N_n \prec_{\boldsymbol{K}} EM_\tau(n,\Phi)$ by the coherence axiom. For condition 7), let $c \in EM(I,\Phi)$. Say $c \in \mathrm{cl}_{\tau^+}(\boldsymbol{a})$ for $\boldsymbol{a} \in I$ and $|\boldsymbol{a}| = n$. Then $EM(\boldsymbol{a},\Phi) \approx_{P(M^*)} EM(n,\Phi)$. So if c realizes q', there is a $c' \in EM(n,\Phi)$ realizing q'. But this directly contradicts clause 3d) describing the τ^+-closure. $\square_{15.8}$

Our application of the following result will be when λ is a successor cardinal. We state a slightly more general form.

Theorem 15.9. *If $\boldsymbol{K}$ is categorical in some λ, with $\mathrm{cf}(\lambda) > H_2$ then $\boldsymbol{K}$ is categorical in H_2 and indeed in any μ-collection cardinal κ between H_2 and λ.*

Proof. By Theorem 12.15, $\boldsymbol{K}$ is (χ,κ)-weakly tame for some $\chi < H_1$. Let M be a model of cardinality κ. Since $\kappa > H_2$, Lemma 15.4 implies M is H_1-saturated. If M is not saturated, applying Theorem 15.8 with $\theta = H_1$ shows that there is a model in cardinality λ that is not χ-saturated. But we know (Lemma 11.3) the model in λ is special and thus ρ-saturated for any $\rho < \mathrm{cf}(\lambda)$. $\square_{15.9}$

Remark 15.10. Keisler [80] deals with PC_δ-classes; classes defined as the reducts of a theory in a countable expansion of a first order language. He proves that if a PC_δ-class has a model in an ω_1-collection cardinal then it has a non-ω_1-saturated model (in the classical first order sense) in every cardinal λ. And for PC_δ classes he establishes that a categoricity model is saturated. So categoricity in any cardinal implies categoricity on all ω_1-collection cardinals. The difficulty in applying this result to categoricity transfer in our context is that for abstract elementary classes and more generally for $PC\Gamma$-classes, in interesting cases the categoricity model is not saturated in the classical sense. Keisler's answer to this difficulty [80], which was extended in [119] was to require homogeneity in the categoricity cardinal. We saw in Chapter 4 that this was insufficiently general. Shelah's solution for abstract elementary classes is to deal with Galois-types and prove that if $\boldsymbol{K}$ is categorical in λ then the model is $\mathrm{cf}(\lambda)$-Galois-saturated. The fly in the ointment is that now the transfer of non-saturation must be reproved to show that a Galois type can be omitted. This requires an additional tameness hypothesis and so we no longer have an easy way to get categoricity on a class of cardinals from categoricity in a single cardinal.

This completes the proof that categoricity transfers from λ down to H_2. To ground the induction carrying categoricity through the interval $[H_2, \lambda^+]$, we show how to transfer the existence of a Vaughtian pair of Galois types from λ to H_2. The result would be standard with better bounds for syntactic types; the issue is coding the Galois types. Recall again our notational conventions in the chapter.

We use in this argument the following notation, simply to find a small universe of sets that contains all the objects that we need in our construction.

Notation 15.11. *Let* $\mathrm{trcl}(x)$ *be the transitive closure of* x *[92]. For any cardinal* θ*, (the sets of hereditary cardinality* $< \theta$*)*

$$\mathcal{H}(\theta) = \{x : |\mathrm{trcl}(x)| < \theta\}.$$

Theorem 15.12. *If* $p \in \mathbb{S}(M^*)$ *with* $|M^*| = \chi < H_1 \leq \theta$*, and there is a true* (p, θ)*-Vaughtian pair then for any* $\lambda \geq \theta$*, there is a true* (p, λ) *Vaughtian pair.*

Proof. Let us state the claim more precisely. Suppose $p \in \mathbb{S}(M^*)$, $M^* \prec_{\boldsymbol{K}} M_0 \prec_{\boldsymbol{K}} N_0$ with $\theta = |M_0| = |N_0| > \beth_\mu(|M^*|)$, M_0 and N_0 are θ-Galois saturated, and p is omitted in $N_0 - M_0$. Say $|M^*| = \chi$. Then for any $\lambda \geq \chi$ there are a pair of models M_1, N_1 of cardinality λ with $M^* \prec_{\boldsymbol{K}} M_1 \prec_{\boldsymbol{K}} N^1$, M_1 and N_1 are λ-Galois saturated, and such that p is omitted in $N_1 - M_1$.

By the presentation theorem, there is a vocabulary τ' such that N_0 is the reduct to τ of a τ'-structure omitting a designated family of types. Without loss of generality the universe of the structure N_0 is a subset with cardinality θ of $\mathcal{H}(\theta)$ that contains the elements we now describe. Enrich the structure $\langle \mathcal{H}(\theta), \epsilon \rangle$ to a $\tau'' \supseteq \tau \cup \{\in\}$-structure by adding predicates P_0 for M_0, P_1 for N_0, naming the elements in M^* by the odd ordinals in χ, and including names for χ and each of its elements; call the resulting structure B. To express that M_0 and N_0 are θ-Galois saturated, recall from Lemma 11.10 a model is Galois-saturated if and only if it can be represented in the form $EM(I, \Phi)$ for a linear order with $|I| = \theta$ and with certain cofinality properties easily expressed in set theory. Thus the first order theory of B expresses that $P_0(B)$ and $P_1(B)$ are $|P_0(B)|$-saturated. In constructing the expansion to B, we require that $P_0(B) \restriction \tau \prec_{\boldsymbol{K}} P_1(B) \restriction \tau \prec_{\boldsymbol{K}} B \restriction \tau$. Apply Theorem A.3.2a to find for any $\lambda > \theta$ an elementary extension B' of B which realizes no τ''-types over the empty set not realized in B and with $|P_1(B')| = |P_0(B')| = |B'| = \lambda$. Since we required the representation of $P_i(B)$

as appropriate EM-models this transfer to λ and $|P_1(B')|, |P_0(B')|$ are λ-Galois saturated. Suppose for contradiction that some $c \in P_1(B') - P_0(B')$ realizes $p \in \mathbb{S}(M^*)$. By the Löwenheim-Skolem property there is a $d \in B'$ such that B' satisfies, '$d \prec_{\tau''} B'$' and each element of M^* is in d and $c \in d - M_0$ and d has cardinality χ (witnessed by a function f from χ to d which fixes M^*)'. As each element m of M^* has a name $m' \in \tau''$, $\{f(\alpha) \neq x : \alpha \in M^*\}$ is contained in $\mathrm{tp}(c/M^*)$. Further by the Löwenheim-Skolem property in $\boldsymbol{K}$, we can require '$d \prec_{\boldsymbol{K}} B'$'. Since no first order type, which is omitted in B, is realized in B' there are elements $c^-, d^-, f^- \in B$ satisfying the same type as cdf over the empty set. In particular, $\{f^-(\alpha) \neq c^- : \alpha \in M^*\}$ and $c^- \notin M_0$. Although neither f nor f^- respects τ'-structure, since $e^- f^-$ and ef realize the same type for $e \in \chi$, the map h from d^- to d defined by $h(x) = y$ if and only if for some $e \in \chi$, $B \models f^-(e) = x$ and $B' \models f(e) = y$ is a τ'' and in particular a τ-isomorphism, taking c to c^- and fixing M^* pointwise. Therefore, c and c^- realize the same Galois type over M^*. This contradicts that p is omitted in B. So, we have constructed a true (p, λ)-Vaughtian pair. $\square_{15.12}$

Remark 15.13. Many omitting types theorems can be proved by this 'expansion of set theory' technique. We have avoided the technique elsewhere as it seemed to only add complication; here there are important reasons for using it.

1. The element d is chosen in the extended model and a copy is chosen below. But we can't just add a predicate for d in advance to B because we have to cover every choice.
2. For the proof of categoricity transfer from one cardinal we can only show the non-existence of *true* Vaughtian pairs so we must guarantee in the construction for Theorem 15.12 that the larger pair of models is saturated; we are able to write this in set theory.

Finally, the proof outlined in the second paragraph of this chapter yields the following result which includes both a weaker version (categoricity assumed in two cardinals instead of one) of the main result of [128] and the extension to upward categoricity of Grossberg and VanDieren [46]. In the next chapter we recover Shelah's full assertion (modulo fine points about the lower bound) by deducing $\boldsymbol{K}$ omits λ-saturated unions from categoricity in λ^+.

Theorem 15.14. *Suppose $\boldsymbol{K}$ is categorical in cardinal λ^+ with $\lambda \geq H_2$ and either*

1. *$\boldsymbol{K}$ is λ-categorical, or*
2. *$\boldsymbol{K}$ omits λ-saturated unions,*

then

1. *$\boldsymbol{K}$ is categorical in every θ with $H_2 \leq \theta \leq \lambda^+$;*
2. *if, in addition $\boldsymbol{K}$ is (H_2, ∞)-weakly tame $\boldsymbol{K}$ is categorical in every θ with $H_2 \leq \theta$.*

Proof. By Theorem 12.15, there is a $\chi < H_1$ such that $\boldsymbol{K}$ is $(\chi, \leq \lambda)$-tame. Theorem 15.9 tells us $\boldsymbol{K}$ is categorical in H_2 and the model M_0 with cardinality H_2 is saturated. Since H_2 is a limit of limit cardinals, without loss of generality, χ is a limit cardinal. By Lemma 11.17, $< \chi^+$ unions of saturated models of cardinality χ are saturated. By Lemma 13.23, there is a minimal type p over a saturated model M^* of cardinality χ. Take a non-splitting extension of it to get a minimal type $\hat{p}$ over M_0. By Lemma 14.5.1, if there is a $(\hat{p}, H_2)$-Vaughtian pair, there is a (p, H_2)-Vaughtian pair and it is true since we have established categoricity in H_2. But then Theorem 15.12 gives us a true (p, λ)-Vaughtian pair over some M_1 with $|M_1| = \lambda$. But since M^* is saturated we can extend p to a non-splitting extension $p' \in \mathbb{S}(M_1)$ which is also minimal. The same N^1 constructed in Theorem 15.12 witnesses a true (p', λ)-Vaughtian pair. But by Theorem 14.6 there is no true (p', λ)-Vaughtian pair. Now use Lemmas 14.11 and 14.12 as in the proof of Theorem 14.13 to get categoricity on the interval $[H_2, \lambda]$. Of course, if we assume (χ, ∞)-tameness the induction continues through all cardinals. $\square_{15.14}$

Note that $\beth_{\omega_1}$ is the best possible lower bound for when an arbitrary infinitary sentence (hence an AEC) can become categorical. For any $\alpha < \aleph_1$ there is a sentence ψ_α of $L_{\omega_1,\omega}$ with no models in cardinals greater than $\beth_\alpha$ and with many models in each cardinal below $\beth_\alpha$ (e.g. page 69 [80], [111]). If ψ asserts the language has only equality, then $\psi_\alpha \vee \psi$ is categorical on $(\beth_\alpha, \infty)$. But this class clearly fails the amalgamation and joint embedding properties.

Remark 15.15. In the abstract of [128], Shelah announces the result of Theorem 15.14 with lower bound, 'a suitable Hanf number'. He appears to assert in the paper, using our language, that there is a $\chi < H_1$ such that $\boldsymbol{K}$ is $(\chi, [H_1, \lambda])$-tame; this is a more uniform result than we have obtained. In Claim II.1.6 and Theorem II.2.7 of [128], Shelah identifies the suitable Hanf number as $H(\chi) = \beth_{(2^\chi)^+}$. (He writes $\chi(\Phi)$.) The cardinal χ does depend on $\boldsymbol{K}$ (not the cardinality of the language of $\boldsymbol{K}$). We have weakened Shelah's stated result since $H_1 = \beth_{(2^{\mathrm{LS}(\boldsymbol{K})})^+} < \beth_{(2^\chi)^+} < H_2$ unless one is in a case when $\mathrm{LS}(\boldsymbol{K}) = \chi$; this result seems to be all that Shelah's methods actually yield as we explain in Remark 15.7. It remains open to determine the best lower bound on categoricity; e.g. whether H_2 can be reduced to Shelah's claim, or to H_1, or even lower.

We thank Chris Laskowski for several very useful conversations on this topic and for contributing the apt term 'collection cardinal'.

16
Unions of Saturated Models

In this chapter we show that if $\boldsymbol{K}$ is λ-categorical and $\mu < \lambda$ then any union of less than μ^+ μ-saturated models (of cardinality μ) is μ-saturated. We begin by providing a 'superstablity' condition and show it follows from λ-categoricity. Then we invoke our detailed connections between limit models and EM-models to show that if the main theorem fails $\boldsymbol{K}$ is not stable in μ. We conclude by deducing transfer of categoricity from an arbitrary successor cardinal.

To extend Theorem 11.17 (on unions of saturated models) to regular λ requires an interweaving of splitting arguments with the construction of Ehrenfeucht-Mostowski models. This chapter is based on some material in Sections 5 and 6 of [128]. Theorem 16.3 is a kind of superstability condition; it is analogous to saying $\kappa(T)$ is finite in the first order case. But while the first order assertion is proved by compactness and counting types, we need here to really exploit the properties of EM-models over well-chosen linear orders. Recall I denotes $\mu^{<\omega}$. We studied (μ, δ)-chains and limit models in Chapter 11. The following notion is patterned on $\kappa(T)$ (e.g. [8] in the first order case.

Definition 16.1. *$\kappa(\boldsymbol{K}, \mu)$ is the least β such that there is no (μ, β)-chain, $\langle M_i : i < \beta\rangle$, over a saturated model M_0 with a $p_\omega \in M_\omega$ that μ-splits over M_i for every $i < \beta$.*

Exercise 16.2. *Use the trick of Theorem 11.16 to show that it would not change the value of $\kappa(\boldsymbol{K}, \mu)$ if we strengthened the requirements on the chain by requiring M_{i+1} to be a limit model over M_i rather than merely universal.*

In the following argument we take products of I with various ordinals; for consistency with the usual conventions for ordinal multiplication, the order on the product is inverse lexicographic.

Theorem 16.3. *Assume $\boldsymbol{K}$ is λ-categorical (λ regular) and $\mathrm{LS}(\boldsymbol{K}) < \mu < \lambda$. Then $\kappa(\boldsymbol{K}, \mu) = \omega$*

Proof. Represent each M_i as $EM(I \times i, \Phi)$ by Lemma 11.15.3. Extend the chain to one of length μ^+; this is possible by stability in μ. For $\gamma \leq \mu^+$, let M_γ denote $EM(I \times \gamma, \Phi)$ and write L for $I \times \mu^+$. By Corollary 11.13.1 and Lemma 11.15.2, M_ω is saturated. By Lemma 13.14, p_ω has an extension to $q \in \mathbb{S}(M_{\mu^+})$ which does not split over some $N \prec_{\boldsymbol{K}} M_\omega$ with $|N| < \mu$ (and so does not split over M_ω).

Lemma 11.10 implies $EM(J, \Phi)$ is saturated for $J = I \times (\mu^+ + \mu^+)$. So we can embed M_{μ^+} into $EM(J, \Phi)$ over M_ω and realize q by a term $\sigma(\boldsymbol{a}, \mathbf{b}, \mathbf{c})$ where $\boldsymbol{a} \in I \times \omega$, $\mathbf{b} \in L - (I \times \omega)$ and $\mathbf{c} \in J - L$. Let $J_1 = I \times (\mu^+ + \mu^+ + \mu^+)$. Now replace each $b_i = \langle t_i, \gamma_i \rangle$ (where $t_i \in I$ and $\gamma_i < \mu^+$) by $b'_i = \langle t_i, \mu^+ + \gamma_i \rangle$ and each $c_i = \langle t_i, \gamma_i \rangle$ (where $t_i \in I$ and $\mu^+ \leq \gamma_i < \mu^+ + \mu^+$) by $c'_i = \langle t_i, \mu^+ + \mu^+ + \gamma_i \rangle$. Then $\mathbf{m} = \sigma(\boldsymbol{a}, \mathbf{b}', \mathbf{c}')$ also realizes p_ω since $\mathbf{bc}$ and $\mathbf{b}'\mathbf{c}'$ are order isomorphic over I_ω. Now we assume $q' = \mathrm{tp}(\mathbf{m}/M_{\mu^+})$ μ-splits over each M_i for $i < \omega$ and derive a contradiction. Fix n so that $q' = \mathrm{tp}(\mathbf{m}/M_{\mu^+})$ μ-splits over M_n and $\boldsymbol{a} \in M_n$. We will show a more precise version of: for any α with $\kappa \leq \alpha < \mu^+$:

$$(*) \qquad \mathrm{tp}(\mathbf{m}/M_{\mu^+}) \text{ splits over } M_\alpha;$$

this result contradicts Lemma 13.2.

Suppose that N_0, N_1 are submodels of M_{μ^+} which witness the splitting of q' over M_n. Choose $v \subset \mu^+$ with $|v| = \mu$ such that both N_0 and N_1 are contained in $EM(I \times v, \Phi)$. Fix γ with $\sup(v) < \gamma < \mu^+$ and let $\gamma' = \gamma + \omega$. Now let

$$M^- = EM(I \times (\omega \cup v \cup [\gamma', \mu^+)), \Phi).$$

Then N_0, N_1 also witness that $\mathrm{tp}(\mathbf{m}/M^-)$ μ-splits over M_n. Abbreviate $I \times n$ as I_n.

Claim 16.4. *There is an automorphism of $I \times (\mu^+ + \mu^+ + \mu^+)$ such that i) f fixes I_n, ii) f fixes $I \times [\gamma', \mu^+ + \mu^+ + \mu^+)$ and iii) f maps $I \times v$ into $I_{n+1} - I_n$. Thus, tp$(\mathbf{m}/EM(I_{n+1}, \Phi)$ splits over $EM(I_n, \Phi)$*

Proof. Since $I \times v$ has cardinality μ, it can be mapped (see Claim 10.5) into $I_{n+1} - I_n$ while fixing I_n. Then by going back and forth ω times this map can be extended to an automorphism of $I \times \gamma'$ and then just fix the rest. $\square_{16.4}$

Note this is not quite the argument given in 6.3 in [128]. Now we need another claim about the ordering I from Section 6 of [128] (although Shelah doesn't specify the order). Think of $I = \mu^{<\omega}$ as a tree. Observe that you could think of I as $I \times \mu$ by for $i < \mu$ taking the ith component of $I \times \mu$ to be all the extensions of

$\langle i\rangle$ in $\mu^{<\omega}$. Keeping that picture in mind, one can think of $I \approx I \times (\gamma+1)$ by keeping the first γ components the same and identifying the last one with right hand side of the picture. We formalize these arguments as part of the following proof.

Claim 16.5. *Fix $I = \mu^{<\omega}$. For any ordinal $\gamma < \mu^+$, $I \approx I \times (\gamma+1)$.*

Proof. We showed $I \approx I \times \mu$ in the proof of Lemma 11.15. We first show that for any $\gamma < \mu$, $I \times (\gamma+1) \approx I$. For this, let $\mathbf{i} = \langle i_0, \ldots i_n\rangle \in I \times (\gamma+1)$ with $i_0 \in \gamma+1$ and $\langle \ldots i_n\rangle \in I$. Map $\mathbf{i}$ to $\mathbf{i}' = \langle i'_0, \ldots i'_n\rangle \in I$ as follows. If $i_0 < \gamma$, $\mathbf{i}' = \mathbf{i}$. Map those sequences with $i_0 = \gamma$ isomorphically to the $\mathbf{i}' \in I$ with $i'_0 \geq \gamma$ using the second paragraph of the proof of Lemma 11.15. (In the notation of that argument we have $I_\gamma \approx I \approx I \times \mu \approx I \times [\gamma, \mu)$.)

Now for $\mu < \gamma < \mu^+$. Let α be the least counterexample. The first paragraph shows $\alpha \geq \mu$ so by division of ordinals there exist $\xi < \mu$ and δ so that $\alpha = \mu \times \delta + \xi$; moreover $\delta + \xi < \alpha$. Then

$$I \times (\alpha+1) = I \times ((\mu \times \delta + \xi) + 1).$$

Since the arithmetic of ordinals satisfies left distributivity and associativity of multiplication and addition, invoking $I \approx I \times \mu$ and then induction, we have:

$$I \times (\mu \times \delta + \xi + 1) = I \times (\mu \times \delta) + I \times (\xi+1) = I \times \delta + I \times (\xi+1) = I \times ((\delta+\xi)+1) = I.$$

$\square_{16.5}$

On the basis of Claim 16.5, we finish the proof of Theorem 16.3. Choose an increasing sequence α_i, for $i < \mu^+$, of successor ordinals with $\omega < \alpha_i < \mu^+$. With the claim we can easily construct, for each $i < \mu^+$ an automorphism f_i of $I \times (\mu^+ + \mu^+ + \mu^+)$ such that

1. f_i fixes $I \times (n-1)$,
2. f_i fixes $I \times [\alpha_i + 2, \mu^+ + \mu^+ + \mu^+)$,
3. f_i maps $I \times n$ onto $I \times [n, \alpha_i)$, and
4. f_i maps $I \times n + 1$ into $I \times (\alpha_i + 1)$.

Then $\mathrm{tp}(\mathbf{m}/M_{\alpha_i+1})$ μ-splits over M_{α_i}.

By Lemma 13.2, using Lemma 13.13 if κ is least such that $2^\kappa > \mu^+$, q does not μ-split over M_κ. We have a contradiction when $\alpha_i > \kappa$. $\square_{16.3}$

Corollary 16.6. *Assume $\boldsymbol{K}$ is λ-categorical (λ regular) and $\mathrm{LS}(\boldsymbol{K}) < \mu \leq \lambda$. Suppose $\langle M_i : i < \delta\rangle$ is a (μ, δ)-chain. If $p_\delta \in M_\delta$ then p_δ does not μ-split over M_α for some $\alpha < \delta$.*

Proof. If there is a counterexample to the Corollary, there is a counterexample to the theorem. $\square_{16.3}$

We have established that categoricity implies a kind of superstability related to the finiteness of $\kappa(T)$. But unlike the first order case, our results apply only below the categoricity cardinal. The next result is easy.

Exercise 16.7. *Suppose $\langle M_i : i < \delta\rangle$ is a $\boldsymbol{K}$-increasing chain of μ-saturated models and $\delta > \mathrm{cf}(\mu)$. Show $M = \bigcup_{i<\delta} M_i$ is δ-saturated.*

Now we give a self-contained argument for Theorem 6.7 in [128]:

Theorem 16.8. *Suppose $\boldsymbol{K}$ has joint embedding, the amalgamation property and arbitrarily large models. If $\boldsymbol{K}$ is λ-categorical for the regular cardinal λ and $\mathrm{LS}(\boldsymbol{K}) < \mu < \lambda$ then $\boldsymbol{K}$ admits μ-saturated unions (Definition 14.3).*

Proof. Let $N = \bigcup_{i<\delta} N_i$ where each N_i is μ-saturated and $\delta < \mu^+$. Let $M \prec_{\boldsymbol{K}} N$ have cardinality $\chi < \mu$ and fix $p \in S(M)$. We must show p is realized in N. If not, fix $\hat{p}$ as a non-algebraic extension of p to $\mathbb{S}(N)$ and let $d \in \mathbb{M}$ realize $\hat{p}$. By Exercise 16.7, we can assume $\delta < \mathrm{cf}(\mu)$. Without loss of generality $\chi \geq \delta, \mathrm{LS}(\boldsymbol{K})$.

Possibly enlarging M slightly we can construct a (χ, δ)-chain, $\langle M_i : i < \delta\rangle$, such that $\hat{p}$ does not χ-split over some M_i. For this we construct M_i as follows. Choose an increasing continuous chain $\bigcup_{i<\delta} M_i$, $M_i \prec_{\boldsymbol{K}} N_i$, each M_{i+1} is χ-universal over M_i, and each M_i is χ-saturated of cardinality χ. Further choose for each i a model M_i^+ such that $M_i^+ \prec_{\boldsymbol{K}} N$ and if $\hat{p}$ χ-splits over M_i, $\hat{p} \restriction M_i^+$ χ-splits over M_i and if $j < i$, $N_i \cap M_j^+ \subseteq M_{i+1}$. By Corollary 16.6 for some $i_0 < \delta$, $\hat{p} \restriction M_\delta$ does not χ-split over M_{i_0}. But now the choice of $M_{i_0}^+$ guarantees $\hat{p}$ itself does not χ-split over M_{i_0}. Since the models are saturated, by Lemma 11.13 and the uniqueness of saturated models, we can represent N_i as $EM(\mu, \Phi)$ and M_i as $EM(\chi, \Phi)$. Let $\mathbf{c}$ enumerate the universe of $M_\delta - M_i$; $|\mathbf{c}| = \chi$.

For some $\gamma < \chi^+$, there is a $\mathbf{c}'b' \in EM(\mu+\gamma, \Phi)$ which realize $\mathrm{tp}(\mathbf{c}, N_i)$ and $\mathrm{tp}(d, N_i)$ respectively. Without loss of generality $\mathbf{c}' = \sigma'(\mathbf{z}_0\mathbf{z}_1)$ where $\mathbf{z}_0 \subseteq \chi$ and $\mathbf{z}_1 \subseteq [\mu, \mu+\gamma)$, $|\mathbf{z}_0\mathbf{z}_1| = \chi$, and σ' is a sequence of terms in the expanded language and $b' = \sigma''(\mathbf{w}_0\mathbf{w}_1)$ (where $\mathbf{w}_i$ are finite sequences with $\mathbf{w}_0 \subseteq \chi$ and $\mathbf{w}_1 \subseteq [\mu, \mu+\gamma)$). Note that for each $\beta < \chi^+$ there is a canonical isomorphism from $\beta \cup [\mu, \mu+\gamma)$ onto $\beta + \gamma$. Now define by induction on $\alpha < \delta$ maps g_α. g_0 maps $\chi \cup [\mu, \mu+\gamma)$ to $\chi + \gamma = \gamma_0$. If g_α is defined, $g_{\alpha+1}$ maps $\gamma_\alpha \cup [\mu, \mu+\gamma)$ to $\gamma_\alpha + \gamma = \gamma_{\alpha+1}$. Let $\mathbf{c}_\alpha$ denote $\sigma'(\mathbf{z}_0, g_\alpha(\mathbf{z}_1))$ and b_α denote $\sigma''(\mathbf{w}_0, g_\alpha(\mathbf{w}_1))$. Note that for each α, $M_i\mathbf{c}_\alpha \approx M_i\mathbf{c} \in K_\chi$. We have constructed for $\alpha < \delta$, $M_{i,\alpha}, g_\alpha, b_\alpha$ such that each $M_{i,\alpha} = EM(\gamma_\alpha, \Phi)$ has cardinality χ, $M_i \prec_{\boldsymbol{K}} M_{i,\alpha} \prec_{\boldsymbol{K}} N_i$ and g_α is an isomorphism from M_δ to a submodel of $M_{i,\alpha+1}$ over $M_{i,\alpha}$. By the choice of σ'', for each α, b_α realizes $\hat{p} \restriction M_{i,\alpha}$.

If some b_α realizes p, we finish (since $M_{i,\alpha} \prec_{\boldsymbol{K}} N_i \prec_{\boldsymbol{K}} N$). If not, recall $d \in \mathbb{M}$ realizes $\hat{p}$. Now

(*) If $\alpha < \beta < \delta$, $\mathbf{c}_\beta b_\alpha$ does not realize the same type over M_i as $\mathbf{c}d$.

For this, note that since b_α not realize p it certainly does not realize $\hat{p} \restriction M_i\mathbf{c}$. But d realizes $\hat{p} \restriction M_i\mathbf{c}$. But, using Exercise 13.8,

(**) if $\beta < \alpha < \delta$, $\mathbf{c}_\beta b_\alpha$ does realize the same type over M_i as $\mathbf{c}d$ since neither $\mathrm{tp}(b_\alpha/\mathbf{c}_\beta)$ nor $\mathrm{tp}(d/\mathbf{c})$ split over M_i.

Let $\mathbf{e}_\alpha$ denote the δ-sequence $g_\alpha(\mathbf{z}_1), g_\alpha(\mathbf{w}_1)$. Then for any $\alpha, \beta < \delta$, $\mathbf{e}_\alpha$ is order isomorphic to $\mathbf{e}_\beta$ over χ and if $\alpha < \beta$, all elements of $\mathbf{e}_\alpha$ precede all elements of $\mathbf{e}_\beta$. Each $\mathbf{e}_\alpha$ is isomorphic to an ordinal $\epsilon = \epsilon_0 + \epsilon_1 \approx \mathbf{z}_1 + \mathbf{w}_1$.

Fix a linear order J such that $|J| < \lambda$ and there are $2^{|J|}$ cuts in J. Let $\hat{J}$ be an extension of J of cardinality $2^{|J|}$ which realizes those cuts. Let $I_J, I_{\hat{J}}$ be extensions of $J, \hat{J}$ respectively, which are χ^+-homogeneous. Let

$$\mathbf{c}'_i = \sigma'(\chi + \epsilon \times i, \chi + \epsilon \times i + \epsilon_0)$$

and

$$b'_i = \sigma''(\chi + \epsilon \times i + \epsilon_0, \chi + \epsilon \times (i+1)),$$

for $i \in \hat{J}$.

Now in $I_{\hat{J}}$ for any $i < j$, $i' < j'$ there is an automorphism fixing χ and taking $[\chi + \epsilon \times i, \chi + \epsilon \times i + \epsilon_0)$ to $[\chi + \epsilon \times i', \chi + \epsilon \times i' + \epsilon_0)$ and $[\chi + \epsilon \times j + \epsilon_0, \chi + \epsilon \times (j+1)$ to $[\chi + \epsilon \times j' + \epsilon_0, \chi + \epsilon \times (j'+1))$. Thus, in $EM(I_{\hat{J}}, \Phi)$, all tuples $\mathbf{c}'_i, b'_j$ with $i < j$ realize the same Galois type over $EM(\chi, \Phi)$. Thus in particular, conditions (*) and (**) hold for the sequence $\mathbf{c}'_i, b'_i$ for $i \in \hat{J}$. In particular, if c, c' realize distinct cuts in J $\mathrm{tp}(b_c/EM(J, \Phi)) \neq \mathrm{tp}(b_{c'}/EM(J, \Phi))$. Thus $\boldsymbol{K}$ is not $|J|$-stable and we finish. $\square_{16.8}$

This completes the proof. Here is a little more context. The argument above combines the idea of the proof of Lemma 6.7 of [128] with a revised finish. Shelah refers to 4.8.1, which as we now outline gives the result modulo a black box. The last four paragraphs of our proof of Theorem 16.8 give a special case of 4.8.2 of [128]; Shelah's proof is, "straight".

We now deduce from (*) and (**), that $\boldsymbol{K}$ satisfies the $(\chi, 1, \chi)$-order property, where, following Definition 4.3.2 of [128],

Definition 16.9. *$\boldsymbol{K}$ has the* $(\chi, 1, \chi)$-order property *if for every ordinal α there are a set A of cardinality χ and tuples $\mathbf{c}_\alpha$ of length χ and singletons b_i such that: $i_0 < j_0 < \alpha, i_1 < j_1 < \alpha$ implies there is no automorphism of $\mathbb{M}$ fixing A with $f(\mathbf{c}_{i_0}) = \mathbf{c}_{j_1}$ and $f(b_{j_0}) = b_{i_1}$. If A is empty, we just say* χ-order property.

Lemma 16.10. *If the hypothesis of Theorem 16.8 holds but the conclusion fails, then $\boldsymbol{K}$ has the $(\chi, 1, \chi)$-order property.*

Proof. Now for any cardinal $\theta \geq \chi$, form first $I'_\theta = \chi + \epsilon \times \theta$ and let I_θ be a linear order (possibly long) which extends θ and is χ^+-homogeneous. Then for $i < \theta$, let $\mathbf{c}'_i = \sigma'(\chi + \epsilon \times i, \chi + \epsilon \times i + \epsilon_0)$ and $b'_i = \sigma'(\epsilon \times i + \epsilon_0, \chi + \epsilon \times (i+1))$

Now in I_μ for any $i < j$, $i' < j'$ there is an automorphism fixing χ and taking $[\chi + \epsilon \times i, \chi + \epsilon \times i + \epsilon_0)$ to $[\chi + \epsilon \times i', \chi + \epsilon \times i' + \epsilon_0)$ and $[\epsilon \times j + \epsilon_0, \chi + \epsilon \times (j+1))$ to $[\epsilon \times j' + \epsilon_0, \chi + \epsilon \times (j'+1))$. Thus, in $EM(I_\mu, \Phi)$, all tuples $\mathbf{c}'_i, b'_j$ with $i < j$

realize the same Galois type over $EM(\chi, \Phi)$. Thus in particular, conditions (*) and (**) hold for the sequence $\mathbf{c}'_i, b'_i$ for $i < \theta$. $\square_{16.10}$

Finally we can finish Shelah's version based on Claim 4.8.1 of [128]: If an AEC has the χ-order property then it has the maximal number of models in all sufficiently large cardinalities. For this one must first note that by adding constants one can pass from the $(\chi, 1, \chi)$-order property to the χ-order property. Let me just note two points about Claim 4.8.1. The order property is for χ-types not finite tuples; it is defined in terms of automorphisms, not syntactically. Both of these extensions to the methods of [122, 140] are plausible, but they are not in print.

We return to the main line. Since we can deduce the nonexistence of Vaughtian pairs from the union of chains of saturated models lemma, we can apply Theorem 15.14.i) to deduce:

Theorem 16.11. *Suppose* $\mathrm{LS}(\boldsymbol{K}) < \lambda$*, and that* $\boldsymbol{K}$ *has arbitrarily large models, the amalgamation property, and joint embedding. Suppose* $\boldsymbol{K}$ *is* λ^+*-categorical.*

1. $\boldsymbol{K}$ *is categorical in all cardinals* $H_2 \leq \mu \leq \lambda^+$.
2. *If* $\boldsymbol{K}$ *is* (χ, ∞)*-tame for some* $\chi < \lambda^+$*, then* $\boldsymbol{K}$ *is* μ*-categorical for all* $\mu \geq \lambda^+$.

Remark 16.12. Our proof of Theorem 16.11 is based on [128]; another argument for part 2) was given by [44]. Lessmann [101] proved an analogous result to Theorem 16.11.2 for the case where λ is $\mathrm{LS}(\boldsymbol{K})$, provided $\mathrm{LS}(\boldsymbol{K}) = \aleph_0$.

The following result is implicit in [128].

Conclusion 16.13. *There is a cardinal* μ *depending on* κ *such that if* $\boldsymbol{K}$ *is an AEC with* $\kappa_{\boldsymbol{K}} = \kappa$*, and* $\boldsymbol{K}$ *is categorical in some successor cardinal* $\lambda^+ > \mu$*, then* $\boldsymbol{K}$ *is categorical in all cardinals greater than* μ.

Proof. There are only a set of AEC $\boldsymbol{K}$ with $\kappa_{\boldsymbol{K}} = \kappa$. Let

$$\mu_{\boldsymbol{K}} = \sup\{\lambda^+ : \boldsymbol{K} \text{ is } \lambda^+\text{-categorical}\}$$

if such a supremum exists. Then let μ be the maximum of H_2 $(= H_2(\kappa))$ and $\sup\{\mu_{\boldsymbol{K}} : \kappa_{\boldsymbol{K}} = \kappa\}$. Now if $\kappa_{\boldsymbol{K}} = \kappa$ and $\boldsymbol{K}$ is categorical in some successor cardinal greater than μ, $\boldsymbol{K}$ is categorical in arbitrarily large successor cardinals and therefore by Theorem 16.11 in all cardinals greater than H_2.

Note that while we can calculate H_2, there is only an existence claim for μ. There remains the possibility that as with Silver's Example 5.28, there is an AEC that is categorical only on a class of limit cardinals.

17
Life without Amalgamation

One general approach to proving Shelah's conjecture for abstract elementary classes (expressed for example in Conjecture 2.3 of [48] and Remark 14.1) is to deduce amalgamation from categoricity and then apply the results for AEC with the amalgamation property. Another approach is to work by induction, interweaving results about categoricity and amalgamation; this is somewhat closer to [131]. Still another is to work under some weakening of amalgamation and joint embedding such as 'no maximal models'. Under any of these approaches it is necessary to establish some notion of Galois type without assuming amalgamation and develop some its properties.

In Chapter 9, we assumed the amalgamation property and developed the notion of Galois type. Here, we show these notions make sense without any amalgamation hypotheses. Of course the notions defined here yield the previous concepts when amalgamation holds. The main technical goal of this chapter is to establish 'model-homogeneity = saturation' without assuming amalgamation. In the process, we obtain one embedding property, Lemma 17.4, which we needed even when assuming amalgamation.

Then we discuss two approaches to working in classes without amalgamation. First we outline the situation when amalgamation holds only up to (or at) a certain cardinal. This situation is further developed in, for example, [131]. Then, we describe some of the results that have been obtained when the amalgamation property is replaced by the hypothesis: $\boldsymbol{K}$ has no maximal models ([144] and [141]).

Surprisingly the proof that model homogeneity is equivalent to saturation has a formulation whose truth does not depend on the amalgamation property. In this chapter we only assume:

Assumption 17.1. $\boldsymbol{K}$ *is an abstract elementary class.*

The goal is to derive properties on embedding models from the realization of Galois types. We want to show that if M^1 realizes 'enough' types over M then any small extension N of M can be embedded into M^1. The idea is first published as 'saturation = model-homogeneity' in 3.10 of [139] (Theorem 17.5 below), where the proof is incomplete. Successive expositions in [133, 48], and by Baldwin led to this version, where the key lemma was isolated by Kolesnikov. In contrast to various of the expositions and like Shelah, we make no amalgamation hypothesis.

Whether we really gain anything concerning the 'model-homogeneity = saturation' theorem by not assuming amalgamation is unclear. I know of no example where either λ-saturated or λ-model homogeneous structures are proved to exist without using amalgamation, at least in λ.

The key idea of the construction is that to embed N into M^2, we construct a $M^1 \prec_{\boldsymbol{K}} M^2$ and a $\boldsymbol{K}$-isomorphism f from an $N^1 \in \boldsymbol{K}$ onto M^1 with $N \subseteq N_2$. Then the coherence axiom tells us restricting f^{-1} to N, gives the required embedding. We isolate the induction step of the construction of f in Lemma 17.4. We will apply the lemma in two settings. In one case $\overline{M}$ has the same cardinality as M and is presented with a filtration M_i. Then $\hat{M}$ will be one of the M_i. In the second, $\overline{M}$ is a larger saturated model and $\hat{M}$ will be chosen as a small model witnessing the realization of a type.

In this chapter we do *not* assume amalgamation and we work with Definition 9.6 of a Galois type. As before, we denote the set of Galois types over M by $\mathbb{S}(M)$.

Definition 17.2.

1. *We say the Galois type of* a *over* M *in* N_1 *is* strongly realized *in* N *with* $M \prec_{\boldsymbol{K}} N$ *if for some* $b \in N$, $(M, a, N_1) \sim_{AT} (M, b, N)$.
2. *We say the Galois type of* a *over* M *in* N_1 *is* realized *in* N *with* $M \prec_{\boldsymbol{K}} N$ *if for some* $b \in N$, $(M, a, N_1) \sim (M, b, N)$.

Now we need a crucial form of the definition of saturated from [133].

Definition 17.3. *The model* M *is* μ-Galois saturated *if for every* $N \prec_{\boldsymbol{K}} M$ *with* $|N| < \mu$ *and every Galois type* p *over* N, p *is* strongly *realized in* M.

Under amalgamation we define saturation using realization rather than strong realization and we have an equivalent notion. Without amalgamation, the notion here is obviously more restricted. Shelah writes in Definition 22 of [133] that in all 'interesting situations' we can use the strong form of saturation.

We use in this construction without further comment two basic observations. If f is a $\boldsymbol{K}$-isomorphism from M onto N and $N \prec_{\boldsymbol{K}} N_1$ there is an M_1 with $M \prec_{\boldsymbol{K}} M_1$ and an isomorphism f_1 (extending f) from M_1 onto N_1. (The dual holds with extensions of M.) Secondly, whenever $f_1 \circ f_2 : N \mapsto M$ and $g_1 \circ g_2 : N \mapsto M$ are maps in a commutative diagram, there is no loss of generality in assuming $N \prec_{\boldsymbol{K}} M$ and $f_1 \circ f_2$ is the identity.

Of course, under amalgamation of models of size $|M|$, we can delete the strongly in the hypothesis of the following lemma.

Lemma 17.4 (Kolesnikov)**.** *Suppose $M \prec_{\boldsymbol{K}} \overline{M}$ and $\overline{M}$* strongly realizes *all Galois-types over M. Let f be a $\boldsymbol{K}$-isomorphism from M onto N and suppose $\tilde{N}$ is a $\boldsymbol{K}$-extension of N. For any $a \in \tilde{N} - N$ there is a $b \in \overline{M}$ such that for any $\widehat{M}$ with $Mb \subseteq \widehat{M} \prec_{\boldsymbol{K}} \overline{M}$ and $|M| = |\widehat{M}| = \lambda$, there is an N^* with $\tilde{N} \prec_{\boldsymbol{K}} N^*$ and an isomorphism $\hat{f}$ extending f and mapping $\widehat{M}$ onto $\widehat{N} \prec_{\boldsymbol{K}} N^*$ with $\hat{f}(b) = a$.*

Proof. Choose $\tilde{M}$ with $M \prec_{\boldsymbol{K}} \tilde{M}$ and extend f to an isomorphism $\tilde{f}$ of $\tilde{M}$ and $\tilde{N}$. Let $\tilde{a}$ denote $\tilde{f}^{-1}(a)$. Choose $b \in \overline{M}$ to strongly realize the Galois type of $\tilde{a}$ over M in $\tilde{M}$. Fix any $\hat{M}$ with $Mb \subseteq \hat{M} \prec_{\boldsymbol{K}} \overline{M}$ and $|M| = |\hat{M}| = \lambda$. By the definition of strongly realize, we can choose an extension M^* of $\tilde{M}$ and $h : \hat{M} \mapsto M^*$ with $h(b) = \tilde{a}$. Lift $\tilde{f}$ to an isomorphism f^* from M^* to an extension N^* of $\tilde{N}$. Then $\hat{f} = (f^* \circ h) \restriction \hat{M}$ and $\hat{N}$ is the image of $\hat{f}$. $\square_{17.4}$

A key point in both of the following arguments is that while the N_i eventually exhaust N, they are not required to be submodels (or even subsets) of N.

Theorem 17.5. *Assume $\lambda > \mathrm{LS}(\boldsymbol{K})$. A model M^2 is λ-Galois saturated if and only if it is λ-model homogeneous.*

Proof. It is obvious that λ-model homogeneous implies λ-Galois saturated. Let M^2 be λ-saturated. We want to show M^2 is λ-model homogeneous. So fix $M_0 \prec_{\boldsymbol{K}} M^2$ and N with $M_0 \prec_{\boldsymbol{K}} N$. Say, $|N| = \mu < \lambda$. We will construct M^1 with $M_0 \prec_{\boldsymbol{K}} M^1 \prec_{\boldsymbol{K}} M^2$, N^1 with $M_0 \prec_{\boldsymbol{K}} N^1$ and $N_0 \subset N^1$, and an isomorphism f between N^1 and M^1. The restriction of f to N is the required map. We construct M^1 as a union of strong submodels M_i of M^2. At the same time we construct N^1 as the union of N_i with $|N_i| < \lambda$, which are strong extensions of N, and f_i mapping M_i onto N_i. As an auxiliary we will also construct a increasing chain of N'_i with $N_i \prec_{\boldsymbol{K}} N'_i$. Enumerate $N - M_0$ as $\langle a_i : i < \mu \rangle$. Let $N_0 = M_0$, $N'_0 = N$ and f_0 be the identity. At stage i, f_i, N_i, M_i, N'_i, are defined; we will construct N'_{i+1}, f_{i+1}, N_{i+1}, M_{i+1}. Apply Lemma 17.4 with a_j as a for the least j with $a_j \notin N_i$; take M_i for M; M_{i+1} is any submodel of M^2 with cardinality μ that realizes $\mathrm{tp}(f_i(a_j))/M_i)$ in M^2 by some b_j and plays the role of $\hat{M}$ in the lemma; N'_i is $\tilde{N}$ and N_i is N. The role of $\overline{M}$ is taken by M^2 at all stages of the induction. We obtain f_{i+1} as $\hat{f}$, N_{i+1} as $\hat{N}$ and N'_{i+1} as N^*. Finally f is the union of the f_i and N^1 is the union of the N_i. $\square_{17.5}$

Just how general is Theorem 17.5? It asserts the equivalence of 'M is λ-model homogeneous' with 'M is λ-saturated' and we claim to have proved this without assuming amalgamation. But the existence of either kind of model is near to implying amalgamation on $\boldsymbol{K}_{<\lambda}$. But it is only close. Let ψ be a sentence of $L_{\omega_1,\omega}$ which has saturated models of all cardinalities and ϕ be a sentence of $L_{\omega_1,\omega}$ which does not have the amalgamation property over models. Now let $\boldsymbol{K}$ be the AEC defined by $\psi \vee \phi$ (where we insist that on each model either the $\tau(\psi)$-relations or the $\tau(\phi)$-relations are trivial but not both). Then $\boldsymbol{K}$ has λ-model homogeneous models of every cardinality (which are saturated) but does not have either the joint embedding or the amalgamation property (or any restriction thereof). However, with some mild restrictions we see the intuition is correct. First an easy back and forth gives us:

Lemma 17.6. *If $\boldsymbol{K}$ has the joint embedding property and $\lambda > \mathrm{LS}(\boldsymbol{K})$ then any two λ-model homogeneous models M_1, M_2 of power λ are isomorphic.*

Proof. It suffices to find a common strong elementary submodel of M_1 and M_2 with cardinality $< \lambda$ but this is guaranteed by joint embedding and $\lambda > \mathrm{LS}(\boldsymbol{K})$. $\square_{17.6}$

Definition 17.7. *For any AEC $\boldsymbol{K}$, and $M \in \boldsymbol{K}$ let $\boldsymbol{K}^M$ be the AEC consisting of all direct limits of strong substructures of M.*

Lemma 17.8. *Suppose M is a λ-model homogeneous member of $\boldsymbol{K}$.*

1. *$\boldsymbol{K}^M_{<\lambda}$ has the amalgamation property.*
2. *If $\boldsymbol{K}$ has the joint embedding property $\boldsymbol{K}_{<\lambda}$ has the amalgamation property.*

Proof. The first statement is immediate and the second follows since then by Lemma 17.6 we have $\boldsymbol{K}^M_{<\lambda} = \boldsymbol{K}_{<\lambda}$. $\square_{17.8}$

Now by Lemma 17.8 and Theorem 17.5 we have:

Corollary 17.9. *If $\boldsymbol{K}$ has a λ-saturated model and has the joint embedding property then $\boldsymbol{K}_{<\lambda}$ has the amalgamation property.*

The corollary, which is Remark 30 of [133], confirms formally the intuition that under mild hypotheses we need amalgamation on $\boldsymbol{K}_{<\lambda}$ to get saturated models of cardinality λ. But we rely on the basic equivalence, proved without amalgamation to establish this result.

The formulation of these results and arguments followed extensive discussions with Rami Grossberg, Alexei Kolesnikov and Monica VanDieren; Kolesnikov singled out Lemma 17.4.

Now we turn to two ways of studying AEC without assuming the amalgamation property.

Restricting an AEC to models of bounded cardinality or even to a single cardinal provides an important tool for studying the entire class. We introduce here two notions of this sort. In [134], the notion of λ-frame is a strengthening of what we call here a weak AEC by introducing an abstract notion of dependence on the class of models of a fixed cardinality λ. See also [50].

Definition 17.10.

1. *For any AEC, $\boldsymbol{K}$ we write, e.g. $\boldsymbol{K}_{\leq\mu}$ for the associated class of structures in $\boldsymbol{K}$ of cardinality at most μ.*
2. *Note that if $\mu \geq \mathrm{LS}(\boldsymbol{K})$, ($\boldsymbol{K}_{<\mu}$) and $\boldsymbol{K}_{\leq\mu}$ have all properties of an AEC except the union of chain axioms apply only to chains of length $\leq \mathrm{cf}(\mu)$ ($< \mathrm{cf}(\mu)$).*

3. *We call such a class of structures and embeddings a weak AEC.*

Exercise 17.11. *If* $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ *is an abstract elementary class then the restriction of* $\boldsymbol{K}$ *and* $\prec_{\boldsymbol{K}}$ *to models of cardinality* λ *gives a weak abstract elementary class.*

The next two exercises are worked out in detail in [134].

Exercise 17.12. *If* $\boldsymbol{K}_\lambda$ *is a weak abstract elementary class, show* $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ *is an AEC with Löwenheim number* λ *if* $\boldsymbol{K}$ *taken as the collection of all direct limits of* $\boldsymbol{K}_\lambda$ *and for directed partial orders* $I \subset J$, $M = \bigcup_{i\in I} M_i$ *and* $N = \bigcup_{j\in J} M_J$ *with the* $M_i, M_j \in \boldsymbol{K}_\lambda$, *define* $M \prec_{\boldsymbol{K}} N$ *if* $M_s \prec_{\boldsymbol{K}_\lambda} N$ *if* $s \in I, t \in J$ *and* $s \leq t$ *in the sense of* J.

Exercise 17.13. *Show that if the AEC's* $\boldsymbol{K}_1$ *and* $\boldsymbol{K}_2$ *have Löwenheim number* λ *and the same restriction to models of size* λ *they are identical above* λ.

The amalgamation property is a major assumption; if it has been made, the next Lemma shows that the joint embedding property is largely a convenience. We noted some specific applications of this technique in Remark 5.12.

Lemma 17.14. *If* $\boldsymbol{K}_{<\kappa}$ *has the amalgamation property, then* $\boldsymbol{K}_{<\kappa}$ *is partitioned into a family of weak-AEC's that each have the joint embedding property.*

Proof. Define $M \simeq N$ if they have a common strong extension. Since $\boldsymbol{K}_{<\kappa}$ has the amalgamation property, $\simeq$ is an equivalence relation. It is not hard to check that each class is closed under short unions and so is a weak-AEC. $\square_{17.14}$

If $\boldsymbol{K}$ is AEC with arbitrarily large models and amalgamation, then only one of the equivalence classes determined by Lemma 17.14, can have arbitrarily large members and we can restrict attention to that class.

In the context considered in this chapter it is desirable to introduce some further notation.

Definition 17.15. 1. M *is* κ-amalgamation base *if* $|M| = \kappa$ *and any two strong extensions of* M *of cardinality* κ *can be amalgamated over* M.

2. M *is an* amalgamation base *if any two strong extensions of* M *can be amalgamated over* M.

In contrast to the general model theoretic and universal algebra literature, in the AEC literature the term amalgamation base is often applied to the local notion and the κ is just suppressed. As we'll see below this is a natural notation when studying classes where only some models M are $|M|$-amalgamation bases. But if they all are, the two notions coalesce as can be seen by the following exercise, which is proved by induction on λ. (The trick is to amalgamate two extensions of cardinality λ^+ by fixing one and amalgamating it with every extension of the 'heart' by a model of size λ.)

Exercise 17.16. *If every* $M \in \boldsymbol{K}_{<\lambda}$ *is a* $|M|$*-amalgamation base, then* $\boldsymbol{K}_{<\lambda}$ *has the amalgamation property. In particular, if every model* M *in* $\boldsymbol{K}$ *is* $|M|$*-amalgamation base, then* $\boldsymbol{K}$ *has the amalgamation property.*

Definition 17.17. *A model $M \in \boldsymbol{K}$ is* maximal *if there is no $N \in \boldsymbol{K}$ which is a proper $\prec_{\boldsymbol{K}}$-extension of M.*

The investigation of AEC's with no maximal models began in [141]. This work was more fully developed in [144] which both expands and corrects some arguments and extends the results.

Assumption 17.18. *For the remainder of this section we assume GCH and the Devlin-Shelah weak diamond. (See Appendix C and [144].)*

In this context the notion of stability in μ is amended to count Galois types over μ-amalgamation bases. Shelah-Villaveces and VanDieren have achieved the following results in studying AEC without maximal models.

Theorem 17.19. *Let $\boldsymbol{K}$ be categorical in $\lambda > H_1$ and let $\mu < \lambda$.*

1. *For every $M \in \boldsymbol{K}_\mu$, there exists a μ-amalgamation base N with $M \prec_{\boldsymbol{K}} N$.*
2. *$\boldsymbol{K}$ is μ-stable.*

We discussed limit models in Chapter 11; recall:

Definition 17.20. *For $\sigma \leq \mu^+$, $N \in \boldsymbol{K}_\mu$ is a (μ, σ) limit over $M \in \boldsymbol{K}_\mu$ if N is the union of a continuous chain of amalgamation bases M_i such that M_{i+1} is μ-universal (Definition 11.4) over M_i.*

While it is straightforward that if σ_1 and σ_2 have the same cofinality any N_1 which is a (μ, σ_1)-limit over M and N_2 which is a (μ, σ_2)-limit over M are isomorphic; such a conclusion for differing cofinalities is extremely difficult. Under a serious additional model theoretic hypothesis, VanDieren [144] attains this result.

Theorem 17.21. *Suppose the class of μ-amalgamation bases is closed under increasing $\prec_{\boldsymbol{K}}$-chains of length less than μ^+. Then for any $\sigma_1, \sigma_2 < \mu^+$, if N_1 is a (μ, σ_1)-limit over M and N_2 is a (μ, σ_2)-limit over M, N_1 is isomorphic to N_2 over M.*

This is a complex argument involving such additional tools as splitting, various kinds of towers, and several notions of extensions of towers. One might hope to combine these ideas with the work on categoricity in classes with amalgamation to obtain a categoricity transfer for AEC with no maximal models. See [47] for further progress.

18
Amalgamation and Few Models

In this section we prove that if $\boldsymbol{K}$ is an AEC which is categorical in λ and does not have the amalgamation property in λ then there are 2^{λ^+} models in λ^+.

The argument fails in ZFC; it uses the hypothesis that $2^\lambda < 2^{\lambda^+}$. We begin by expounding (through Theorem 18.7) the counterexample; the basic idea appears in [139]. This example builds on an earlier argument of Baumgartner [25] that we discuss in Example 19.17. There is a clearer version in [138]; with Coppola we have made some further simplifications. In particular, we give a different definition of the notion of strong submodel.

Theorem 18.1 (Martin's Axiom)**.** *There is a sentence ψ in $L(Q)$ with the joint embedding property that is κ-categorical for every $\kappa < 2^{\aleph_0}$. In ZFC one can prove ψ is $\aleph_0$-categorical but has neither the amalgamation property in $\aleph_0$ nor is ω-stable.*

We will work in a model of set theory that satisfies MA and with $2^{\aleph_0} = \aleph_{\omega+1}$. See [92, 75] for the existence of such.

We consider the class $\boldsymbol{K}$ of models in a vocabulary with two unary relations P, Q and two binary relations E, R which satisfy the following conditions. (For any relation symbol and model M we write P^M for the interpretation of P in M. However, abusing notation, when only one model is being considered (as in the next definition) we omit the subscript.)

For any model $M \in \boldsymbol{K}$,

1. P and Q partition M.

2. E is an equivalence relation on Q.

3. P and each equivalence class of E is denumerably infinite.
4. R is a relation on $P \times Q$ that is extensional on P. That is, thinking of R as the 'element' relation, each member of Q denotes a subset of P.
5. For every set X of n elements X from P and every subset X_0 of X and each equivalence class in Q, there is an element of that equivalence class that is R-related to every element of X_0 and not to any element of $X - X_0$.
6. Similarly, for every set of n elements Y from Q and every subset Y_0 of Y, there is an element of P that is R-related to every element of Y_0 and not to any element of $Y - Y_0$.

Note that every member of $\boldsymbol{K}$ has cardinality at most the continuum. For each $y \in Q$, let A_y denote the set of x that are R-related to y. Then, the last two properties imply that the A_y form an independent family of sets (any finite Boolean combination of them is infinite) and, since each E-equivalence class is infinite, for any two finite disjoint subsets u, v of P, there are $|M|$ elements of Q that each 'contain' every element of u and no element of v.

We will use the following special case of Martin's axiom that is tailored for our applications.

Definition 18.2. Martin's Axiom

1. *MA_κ is the assertion: If $\mathcal{F}$ is a collection of partial isomorphisms, partially ordered by extension, between two structures M and N of the same cardinality less than $2^{\aleph_0}$ that satisfies the countable chain condition then for any set of $\kappa < 2^{\aleph_0}$ dense subsets of $\mathcal{F}$, C_a, there is a filter G on $\mathcal{F}$ which intersects all the C_a.*
2. *$\mathcal{F}$ satisfies the countable chain condition if there is no uncountable subset of pairwise incompatible members of $\mathcal{F}$.*
3. *Martin's axiom is: $(\forall \kappa < 2^{\aleph_0})(\mathrm{MA}_\kappa)$.*

Definition 18.3. *Fix the class $\boldsymbol{K}$ as above and for $M, N \in \boldsymbol{K}$, define $M \prec_{\boldsymbol{K}} N$ if $P^M = P^N$ and for each $m \in Q^M$, $\{n \in N : mEn\} = \{n \in M : mEn\}$ (equivalence classes don't expand).*

Lemma 18.4. *[138] Martin's axiom implies that the class $\boldsymbol{K}$ is $\aleph_1$-categorical.*

Proof. We first define the forcing conditions.

Definition 18.5. *Fix continuous filtrations $\langle M_i : i < \kappa\rangle$ of M and $\langle N_i : i < \kappa\rangle$ of N by $\prec_{\boldsymbol{K}}$-submodels such that $M_{i+1} - M_i$ ($N_{i+1} - N_i$) is countable.*

1. *$\mathcal{F}$ is the set of finite partial isomorphisms f such that for each $i < \kappa$, and each $x \in \mathrm{dom}\, f$, $x \in M_i$ if and only if $f(x) \in N_i$.*
2. *we order the conditions by $f \leq g$ if $f \subseteq g$.*

Let us prove that these forcing conditions have the ccc.

Lemma 18.6. *$\mathcal{F}$ satisfies the countable chain condition.*

Proof. Let $\langle f_\alpha : \alpha < \aleph_1\rangle$ be a sequence of elements of $\mathcal{F}$. Without loss of generality, fix m and k so that the domain of each f_α contains m elements of P and k of Q. Applying the Δ-system lemma to the domain and the range, we can find Y (Y') contained in M (N) so that for an uncountable subset S, if $\alpha, \beta \in S$, $\operatorname{dom} f_\alpha \cap \operatorname{dom} f_\beta \cap M = Y$ ($\operatorname{rg} f_\alpha \cap \operatorname{rg} f_\beta \cap N = Y'$). Note that, in fact we may restrict to an uncountable S_1 so that all the f_α for $\alpha \in S_1$ intersect in a single bijection f. For if there were some $b \in Y$ and some α with $f_\alpha(b) \notin Y'$, then (as Y' is the root for the range), the $\{f_\alpha(b) : \alpha \in S\}$ give uncountably many distinct images for b contrary to the choice of the filtration. (A similar argument for the domains and the fact that there are only finitely many maps from Y onto Y' yield the bijection f.)

The requirement that conditions preserve the filtration yields, when $|Y| = \ell$, that for some $\langle i_j : j < \ell\rangle$, $Y \subset \bigcup_{j<\ell}(M_{i_j+1} - M_{i_j})$ and $Y' \subset \bigcup_{j<\ell}(N_{i_j+1} - N_{i_j})$.

Applying the condition on filtrations, each element in the domain (range) has only countably many possible images (preimages); so we can demand that in restricting to S_1 we guarantee that no element of M_i (N_i) occurs in $\operatorname{dom} f_\alpha - Y$ ($\operatorname{rg} f_\alpha - Y'$) for more than one α. But then for any $\alpha, \beta \in S_1$, $f_\alpha \cup f_\beta$ is the required extension. $\square_{18.6}$

If for $a \in M$, we let D_a be the set of conditions with $a \in \operatorname{dom} f$ and define R_b analogously for the range, it is easy to see that all the D_a and R_b are dense. Thus by Martin's axiom there is an isomorphism between M and N. $\square_{18.4}$

It is now trivial to verify:

Theorem 18.7. *$(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ is an AEC which is categorical in $\aleph_0$ and does not satisfy amalgamation in $\aleph_0$. Martin's Axiom implies $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ is categorical in all $\kappa < 2^{\aleph_0}$ as well.*

Proof. The extensionality guarantees the failure of amalgamation. Extend a countable structure in two ways by adding a name for a subset of P and for its complement. $\square_{18.7}$

Now we turn to show that under appropriate set theoretic hypotheses categoricity in λ and few models in λ^+ does imply amalgamation.

Notation 18.8. *In this chapter κ always denotes λ^+. We assume throughout that $2^\lambda < 2^{\lambda^+}$.*

For the next result we need the WGCH, specifically proposition θ_{λ^+}, Definition C.4. The following simple argument due to [139] and nicely explained in [48] shows:

Lemma 18.9 (WGCH). *Suppose $\lambda \geq \mathrm{LS}(\boldsymbol{K})$, $2^\lambda < 2^{\lambda^+}$, and $\boldsymbol{K}$ is λ-categorical. If amalgamation fails in λ there are models in $\boldsymbol{K}$ of cardinality λ^+ but no universal model of cardinality λ^+.*

Proof. Let $N_0 \prec_{\boldsymbol{K}} N_1, N_2$ witness the failure of amalgamation. Then both N_1 and N_2 are proper extensions of N_0. For $\rho \in \lambda^{<\kappa}$, we write $\ell(\rho)$ for the domain of ρ. For each such ρ we define a model M_ρ with universe $\lambda(1+\ell(\rho))$ so if ρ is an initial segment of η, $M_\rho \prec_{\boldsymbol{K}} M_\eta$, and if $\ell(\rho)$ is a limit ordinal $M_\rho = \bigcup_{\delta<\ell(\rho)} M_{\rho\restriction\delta}$. Finally, the key point is that for each ρ, $M_{\rho\frown 0}$ and $M_{\rho\frown 1}$ cannot be amalgamated over M_ρ.

This construction is immediate by λ-categoricity; just copy over N_0, N_1, N_2. For $\eta \in 2^\kappa$, $M_\eta = \bigcup_{\delta<\kappa} M_{\eta\restriction\delta}$; clearly, $|M_\eta| = \kappa$. Suppose for contradiction that there is a model M of cardinality κ which is universal. Let f_η be the embedding of M_η into M. The set C of $\delta < \kappa$ of the form: $\delta = \lambda(1+\delta)$ contains a cub. Applying Θ_{λ^+}, which holds by Theorem C.6, we find $\delta \in C$ and distinct $\eta, \nu \in 2^\kappa$ which agree only up to δ. Say, $\eta(\delta) = 0$ and $\nu(\delta) = 1$. Denoting $\eta \wedge \nu$ by $\rho = \eta\restriction\delta$, we have that f_η and f_ν map $M_{\rho\frown 0}$ and $M_{\rho\frown 1}$ into M over M_ρ. By the Löwenheim-Skolem theorem we have amalgamated an isomorphic copy of N_0, N_1, N_2 in K_λ. This contradiction yields the theorem. $\square_{18.9}$

We do not rely Lemma 18.9, but prove directly that the failure of amalgamation in λ and categoricity in λ imply the existence of 2^κ models in κ. This requires the following variant on the Devlin-Shelah weak diamond. We write $\hat{\Theta}_\lambda$ rather than Φ_λ (as in Appendix C) to indicate the coding of triples.

Definition 18.10. *The principle $\hat{\Theta}_{\lambda^+}(S)$ holds if (letting κ denote λ^+):*
For every function
$$F\colon \kappa^{<\kappa} \times \kappa^{<\kappa} \times \kappa^{<\kappa} \to 2$$
there is an oracle $g\colon\kappa \to 2$ such for every $\eta, \nu, \sigma\colon \kappa \to \kappa$ the set
$$\{\delta \in S\colon F(\eta\restriction\delta, \nu\restriction\delta, \sigma\restriction\delta) = g(\delta)\}$$
is stationary in λ^+.

Devlin and Shelah (Theorem 3.1 of [36] and general properties of normal ideals) deduce:

Fact 18.11. $2^\lambda < 2^{\lambda^+}$ *implies there is a family of λ^+ disjoint stationary subsets S_i of λ^+ such that for each i, $\hat{\Theta}_{\lambda^+}(S_i)$ holds.*

We will use this set theoretic principle to prove the main result of this chapter. Note that Fact refsetfact2 does not assert $\Theta_{\lambda^+}(S)$ for all stationary S but only for many.

Theorem 18.12. *[WGCH] Suppose $\lambda \geq \mathrm{LS}(\boldsymbol{K})$ and $\boldsymbol{K}$ is λ-categorical. If amalgamation fails in λ there are 2^{λ^+} models in $\boldsymbol{K}$ of cardinality $\kappa = \lambda^+$.*

Proof. Let $N_0 \prec_{\boldsymbol{K}} N_1, N_2$ witness the failure of amalgamation. Then both N_1 and N_2 are proper extensions of N_0, so, as in Lemma 8.2, $\boldsymbol{K}$ has models of cardinality κ.

For $\rho \in \lambda^{<\kappa}$, we write $\ell(\rho)$ for the domain of ρ. For each such ρ we define a model M_ρ with universe $\lambda(1+\ell(\rho))$ so if ρ is an initial segment of η, $M_\rho \prec_{\boldsymbol{K}}$

M_η, and if $\ell(\rho)$ is a limit ordinal $M_\rho = \bigcup_{\delta<\ell(\rho)} M_{\rho\restriction\delta}$. Finally, the key point is that: for each ρ, $M_{\rho\frown 0}$ and $M_{\rho\frown 1}$ cannot be amalgamated over M_ρ. For $\eta \in 2^\kappa$, let $M_\eta = \bigcup_{\delta<\kappa} M_{\eta\restriction\delta}$; clearly, $|M_\eta| = \kappa$. We will prove that 2^κ of the M_η are pairwise non-isomorphic.

We divide the proof into two cases; in the first the failure of amalgamation is even stronger and the second is the negation of the first.

Case A. There exists $N, M \in \boldsymbol{K}_\lambda$ with $N \prec_{\boldsymbol{K}} M$ such that for every M' extending M of cardinality λ there are M^0 and M^1 extending M' which cannot be amalgamated *even* over N.

In this case we strengthen the construction by fixing M_0 as N and demanding that for every η,

$$M_{\eta\frown 0} \text{ and } M_{\eta\frown 1} \text{ cannot be amalgamated over } M_0.$$

Now if $\eta \neq \nu \in 2^\kappa$,

$$\langle M_\eta, a\rangle_{a\in N} \not\approx \langle M_\nu, a\rangle_{a\in N}.$$

If there is an isomorphism between $\langle M_\eta, a\rangle_{a\in N}$ and $\langle M_\nu, a\rangle_{a\in N}$, then, denoting $\eta \wedge \nu$ by ρ, M_ν contains an amalgam of cardinality λ of $M_{\rho\frown 0}$ and $M_{\rho\frown 1}$ over $M_0 = N$, contradiction. But since $2^\lambda < 2^{\lambda^+}$, if there are 2^{λ^+} models after naming λ constants there are 2^{λ^+} models, period. We finish case A.

Case B. For all N, M with $N \prec_{\boldsymbol{K}} M \in \boldsymbol{K}_\lambda$, there is an M' extending M such that any M^0 and M^1 extending M' can be amalgamated over N.

We construct models M_η for $\eta \in 2^\rho$ by induction on $\rho \leq \kappa$. At limit ordinals we take unions. In the successor stage, we know by categoricity that M_η is isomorphic to the given N_0. So we can choose M^* and M^{**} isomorphic to N_1 and N_2, which *cannot* be amalgamated over M_η. Using the fact that we are in Case B twice, choose $M_{\eta\frown 0}$ extending M^* ($M_{\eta\frown 1}$ extending M^{**}) so that

$$\text{any extensions of } M_{\eta\frown 0}\ (M_{\eta\frown 1}) \text{ can be amalgamated over } M_\eta. \tag{18.1}$$

Now for each $\eta \in 2^\kappa$, we have a model M_η by taking a union of the path. For each $X \subset \kappa$, we are going to choose a path η_X such that if $X \neq Y$, $M_{\eta_X} \not\approx M_{\eta_Y}$. For this choice we use our set theoretic principle.

Since $C = \{\delta < \lambda^+ : \delta = \lambda(1+\delta)\}$ is a cub, we can apply Fact 18.11 to C to get a family of disjoint stationary sets $\langle S_i : i < \kappa\rangle$, each $S_i \subset C$, satisfying $\hat{\Theta}_{\lambda^+}(S_i)$ for each i. We define the function F as follows. For each $\delta \in C$ and each function $h : \delta \to \delta$ and functions $\eta, \nu \in 2^\delta$,

1. $F(\eta, \nu, h) = 1$ if M_η and M_ν have universe δ, h is an isomorphism between them and the embeddings $id : M_\eta \to M_{\eta\frown 0}$ and $h : M_\eta \to M_{\nu\frown 0}$ can be amalgamated.

2. $F(\eta, \nu, h) = 0$ otherwise.

For each i, apply $\hat{\Theta}_{\lambda^+}(S_i)$, to choose $g_i : \kappa \to 2$ so that for every $\hat{h} : \kappa \to \kappa$, the set of $\delta \in S_i$ such that

$$F(\eta \restriction \delta, \eta \restriction \delta, \hat{h} \restriction \delta) = g_i(\delta)$$

is stationary.

Now for $X \subset \kappa$, define $\eta_X : \kappa \to 2$ by

1. $\eta_X(\delta) = g_i(\delta)$ if $\delta \in S_i$ and $i \in X$
2. and 0 otherwise.

Each δ is in at most one S_i so this is well-defined.

Claim 18.13. *If* $X \neq Y$, $M_{\eta_X} \not\approx M_{\eta_Y}$.

Proof of Claim: Suppose for contradiction that $X \neq Y$ and there is an isomorphism h_{XY} between M_{η_X} and M_{η_Y}. Now our choice of g_i by $\hat{\Theta}_{\lambda^+}(S_i)$ gives for each η_X, η_Y a stationary subset S'_i of S_i, which depends on (X, Y, h_{XY}) such that:

$$S'_i = \{\delta : F(\eta_X \restriction \delta, \eta_Y \restriction \delta, h_{XY} \restriction \delta) = g_i(\delta)\}.$$

Note that the set of δ mapped to itself by h_{XY} is a cub D and let $S''_i = S'_i \cap D$.

Without loss of generality there is an $i \in X - Y$. Now fix such an i and $\delta \in S''_i$.

For ease of notation below, we write η for $\eta_X \restriction \delta$, ν for $\eta_Y \restriction \delta$, and h for $h_{XY} \restriction \delta$. Since $\delta \in D$ then h is an isomorphism between M_η and M_ν. Since $i \notin Y$, the definition of η_Y implies $\eta_Y(\delta) = 0$; thus, $\nu \lhd \nu\hat{\ }0 \lhd \eta_Y$. (Here we write $\lhd$ for 'initial segment'). There are now two subcases depending on the value of $\eta_X(\delta)$.

Subcase 1: $\eta_X(\delta) = 1$**.** By the definition of η_X, $\eta_X(\delta) = g_i(\delta)$; since $\delta \in S'_i$, $F(\eta, \nu, h) = 1$. Thus, invoking the definition of F, we see the identity map from M_η into $M_{\eta\hat{\ }0}$ and $h : M_\eta \to M_{\nu\hat{\ }0}$ can be amalgamated over M_η by a model $M^1 \in \boldsymbol{K}_\lambda$.

On the other hand, the identity maps M_η into $M_{\eta\hat{\ }1}$ and $\eta\hat{\ }1 \lhd \eta_X$; h maps M_η into $M_{\nu\hat{\ }0}$ and $\nu\hat{\ }0 \lhd \eta_Y$; and we have the isomorphism h_{XY} between M_{η_X} and M_{η_Y}. Since $h_{XY} \restriction \delta = h$, $M_{\eta\hat{\ }1}$ and $M_{\nu\hat{\ }0}$ can be amalgamated over M_η by a model $M^2 \in \boldsymbol{K}_\lambda$ (since λ is above the Löwenheim number).

In the first paragraph of the subcase we constructed a map from $M_{\eta\hat{\ }0}$ into M^1; in the second we constructed a map from $M_{\eta\hat{\ }1}$ into M^2. Both of these maps were over M_η. Applying condition 18.1 (displayed in Case B) to $\nu\hat{\ }0$, we can amalgamate M^1 and M^2 over $M_{\nu\hat{\ }0}$ in some M^3 of cardinality λ. But then we have amalgamated $M_{\eta\hat{\ }0}$ and $M_{\eta\hat{\ }1}$ over M_η. This contradicts the construction of the M_η-sequence and finishes Subcase 1.

Subcase 2: $\eta_X(\delta) = 0$**.** In this situation, $\eta\hat{\ }0 \lhd \eta_X$. Since h is an isomorphism between M_{η_X} and M_{η_Y}, we see that M_{η_Y} is an amalgam of $M_{\eta\hat{\ }0}$ and $M_{\nu\hat{\ }0}$ over M_η. By the definition of F, $F(\eta, \nu, h) = 1$. Since $\delta \in S''_i$, $g_i(\delta) = 1$ and by the definition of η_X, we conclude $\eta_X(\delta) = 1$. This contradicts the choice of the case and concludes the proof of Claim 18.13. $\square_{18.13}$

We also have proved the theorem. $\square_{18.12}$

Recall from Exercise 6.1.3 that if τ is countable vocabulary, then if $\boldsymbol{K}$ is the class of models of a sentence in $L_{\omega_1,\omega}(Q)$, $(\boldsymbol{K}, \prec^*)$ is an AEC with Löwenheim number $\aleph_1$. Thus by the results of this section (under weak diamond) if$(\boldsymbol{K}, \prec^*)$ is $\aleph_1$-categorical and has few models in $\aleph_2$, it has amalgamation in $\aleph_1$. The situation concerning a sentence in $L_{\omega_1,\omega}(Q)$ that has few models in $\aleph_1$ is slightly more complicated, since there is no naturally occurring AEC with Löwenheim number $\aleph_0$. By applying the analysis from Chapter 8, one can show:

Lemma 18.14 (Devlin-Shelah diamond)**.** *Let $\boldsymbol{K}$ be the class of models of a sentence of $L_{\omega_1,\omega}(Q)$. If $\boldsymbol{K}$ has fewer than $2^{\aleph_1}$ models of cardinality $\aleph_1$ then there is an associated class $(\boldsymbol{K}^*, \prec^{**})$ which has the amalgamation property in $\aleph_0$.*

Remark 18.15. Theorem 18.12 is from [139] but our exposition depends very heavily on [48].

Theorem 18.12 yields (under WGCH) that if an AEC $\boldsymbol{K}$ is categorical in all cardinals above $\mathrm{LS}(\boldsymbol{K})$ then it satisfies the amalgamation property. One might hope to improve this to conclude that if $\boldsymbol{K}$ is eventually categorical then it has the amalgamation property. But amalgamation in different cardinalities can behave differently.

Definition 18.16. *$\boldsymbol{K}$ has the λ-amalgamation property if for every triple of models M_0, M_1, M_2 with $M_0 \prec_{\boldsymbol{K}} M_1, M_0 \prec_{\boldsymbol{K}} M_1$* and *with $|M_i| = \lambda$ for $i < 3$, there is an amalgam M_2.*

Example 18.17. There is an AEC that for all $\kappa \geq \aleph_1$ is categorical in κ, satisfies κ-amalgamation, the joint embedding property, and is $(\aleph_0, \infty)$-tame (Galois types equal syntactic types) but fails $\aleph_0$-amalgamation and so has no model-homogeneous models.

The example is in fact a finite diagram. Let the vocabulary contain unary predicates P_i for $i < \omega$ and one binary relation E. Let the axioms of T assert that E is an equivalence relation with two classes, that $P_{i+1} \subset P_i$ and that each equivalence class contains exactly one element in P_i and not in P_{i+1}. Define $\boldsymbol{K}$ to be the class of models of T that omit the 2-type of a pair of elements that are each in all the P_i but inequivalent. We write $\prec_{\boldsymbol{K}}$ for first order elementary submodel.

It is easy to check that all the conditions are satisfied. Let p_ω be the type of an element satisfying all the p_i. There is a model M_0 omitting p_ω. There are then two incompatible (over M_0) choices for realizing p_ω. Thus amalgamation over M_0 is impossible. But once a model realizes p_ω, the amalgamation class is determined. Moreover, each model is determined exactly by the cardinality of the set of realizations of p_ω.

In contrast we can modify this example to get an atomic class that is categorical in all powers and satisfies amalgamation. Namely, add a predicate P_ω which implies all the P_i and add axioms to T asserting that P_ω is realized by infinitely many elements and any two elements of P_ω are E-equivalent. Now the atomic models of T form a class $\boldsymbol{K}$ that is categorical in all powers and satisfies amalgamation over models.

Note that in either variant the uncountable model M is not $(D, \aleph_1)$-homogeneous in the sense of [119]. Choose an $a \in M_0$ such that a is not equivalent to a realization in M of p_ω. Then $p_\omega(x) \cup \{E(x,a)\}$ is in $S_D(\{a\})$ but is not realized in M. Nevertheless, M is $L_{\mathcal{A}}$-homogeneous in the sense of Keisler [80] where $L_{\mathcal{A}}$ is the minimal fragment of $L_{\omega_1,\omega}$ containing $\bigwedge p_\omega$. The difference is that Shelah's notion allows the type to be realized in another member of $\boldsymbol{K}$ while Keisler's is restricted to truth in M.

Part IV

Categoricity in $L_{\omega_1,\omega}$

This Part brings the book full-circle. We began with the study of categoricity transfer for quasiminimal excellent classes. This gave us a very concrete notion of excellence. In this Part we solve the proportion: strongly minimal sets is to first order categoricity as quasiminimal excellence is to what. The answer is: Shelah's theory of excellent atomic classes.

This Part differs from Part III in two crucial ways. It is more concrete: we study categoricity transfer for sentences in $L_{\omega_1,\omega}$ rather than arbitrary abstract elementary classes. The general model theoretic hypotheses are much weaker; we do not assume the sentence has arbitrarily large models, nor that the class of models has the amalgamation property. Rather, working from the assumption that there are a small number of models in a small cardinality we develop these hypotheses step by step. But, we extend ZFC by assuming weak diamond (see Appendix C).

The remainder of the book is devoted to the study of atomic models of a first order theory. The following assertion is Theorem 7.1.12.

Let ψ be a complete sentence in $L_{\omega_1,\omega}$ in a countable language L. Then there is a countable language L' extending L and a first order L'-theory T such that reduct is a 1-1 map from the atomic models of T onto the models of ψ. So in particular, *any complete sentence of $L_{\omega_1,\omega}$ can be replaced (for spectrum purposes) by considering the atomic models of a first order theory*.

The careful reader will recall from Chapter 7, that for $L_{\omega_1,\omega}$ it is by no means trivial to replace an arbitrary sentence by a *complete* sentence. At the end of Chapter 26, we will use the full strength of the analysis of the spectrum problem for *complete* sentences to extend the categoricity result to arbitrary sentences in $L_{\omega_1,\omega}$.

Our exposition here requires more familiarity with first order stability theory than the first three parts. The fact that in [120, 124, 125] Shelah was introducing a sophisticated stability theoretic analysis of infinitary logic, before the notions of first order stability theory had been absorbed partially explains the great difficulty the logic community has had in understanding these results.

This part centers on a notion of generalized amalgamation called 'excellence'. We try here to give a feel for the notion before beginning the formal development. One key contribution of excellence is to construct arbitrarily large models. We are able to reduce the existence of models in large cardinals to certain n-dimensional amalgamation properties of countable models. They generalize the notion in Chapter 3 by replacing the closure defined there by a notion of dependence based on splitting.

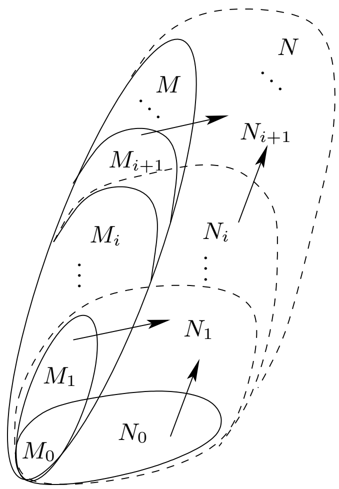

Why does this require n-dimensional amalgamation? If $\boldsymbol{K}$ is an AEC and every $M \in \boldsymbol{K}$ with cardinality $\aleph_0$ has a proper strong extension, it is easy to construct a model of cardinality $\aleph_1$. But how would we get a model of power $\aleph_2$? It would certainly suffice to know every model M of $\aleph_1$ has a proper strong extension. To show this, fix a filtration $\langle M_i : i < \aleph_1 \rangle$ of M by countable models. Now if every countable model has a proper extension there exists N_0 with $M_0 \prec_{\boldsymbol{K}} N_0$. We use a notion of independence to guarantee that N_0 is not contained in M. And then amalgamation allows us to construct N_1 as an amalgam of M_1 and N_0 over M_0. We can continue by induction. Higher dimensional amalgamation arises when we try to build a model of power $\aleph_3$.

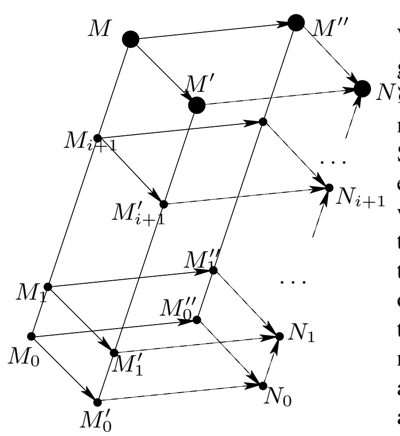

We just used amalgamation of pairs in $\aleph_0$ to get proper extensions in $\aleph_1$ and a model in $\aleph_2$. The same argument shows that amalgamation of pairs in $\aleph_1$ gives a model in $\aleph_3$. So given a model M of power $\aleph_1$ and two extensions M', M'' also of cardinality $\aleph_1$, we want to amalgamate them. Build filtrations $\langle M_i, M'_i, M''_i : i < \aleph_1 \rangle$ of each of the three models. Amalgamating M'_0 and M''_0 over M_0 gives us an N_0 to start the construction of the amalgam. But at the next stage we need that N_1 extend not just M'_1 and M''_1 but also N_0. The existence of such an N_1 is 3-amalgamation.

But excellence is not merely the existence of an n-amalgam but the existence of one that is closely connected (constructible over) its constituents. In Chapters 19 and 20, we lay out the key properties of atomic classes, define the notion of splitting, and develop non-splitting as the relevant notion of independence for this context. Chapter 21 is perhaps the most technical in the book; it concerns the basic properties of independent systems. The notion of excellence is finally formally defined in Chapter 22. Roughly, excellence asserts for each n, that there are primary models over independent systems of 2^{n-1} countable structures. We show that this n-amalgamation for all n on countable structures propagates to arbitrary cardinals. In Chapters 23 and 24 we show, assuming weak diamond, that 'few' models in each $\aleph_n$ (for $n < \omega$) implies excellence. The precise meaning of few is explored in Chapter 24 and Appendix C. We introduce the notion of $*$-excellence in Chapter 25; these are some simpler consequences of excellence which suffice to prove (in ZFC) the main result: If $\boldsymbol{K}$ is *-excellent and categorical in some uncountable power then it is categorical in all uncountable powers. This argument is carried out in Chapters 25 and 26. We complete our circle by describing the exact relationship of Part I with the categoricity result. In Chapter 27, we describe a variation on the example of [56] showing that categoricity up to $\aleph_\omega$ was needed for the main result.

We provide a complete proof for sentences of $L_{\omega_1,\omega}$; Shelah intimated in the introduction to [124] that similar methods would suffice for $L_{\omega_1,\omega}(Q)$. This seems to be overly optimistic unless one gives a very expansive meaning to 'similar'. However, he has developed an extensive theory to deal with this situation. See e.g. [134, 131, 130]. We will also see that Zilber's approach treats some cases in $L_{\omega_1,\omega}(Q)$.

The bulk of this Part originally appeared in [124, 125]. However, our exposition depends heavily on the work of Grossberg and Hart [49] for Chapter 21; the fact that what we have dubbed *-excellence are the actual working hypotheses for the categoricity transfer proof is due to Lessmann [100]. The counterexamples in Chapter 27 rely on analysis by Baldwin-Kolesnikov [5]. We thank Kolesnikov for the diagrams used in this introduction.

19
Atomic AEC

As expounded in Chapter 7 and summarized in the introduction to Part IV, our study of categoricity for $L_{\omega_1,\omega}$ reduces to the study of the atomic models of a complete first order theory. In this chapter we develop the basics of this approach. Unlike earlier sections of the book, this development is very closely analogous to first order stability theory and we have to develop a lot of technical background. This chapter is indirectly based on [120, 124, 125], where most of the results were originally proved. But our exposition owes much to [100, 99, 87, 49].

Definition 19.1. *Let T be a countable first order theory. A set A contained in a model M of T is* atomic *if every finite sequence in A realizes a principal (first order) type over the* empty set*.*

Thus if T is $\aleph_0$-categorical every model of T is atomic.

Notation 19.2. *We say an AEC, $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$, is* atomic *if $\boldsymbol{K}$ is the class of atomic models of a countable complete first order theory and $\prec_{\boldsymbol{K}}$ is first order elementary submodel.*

Part IV proceeds entirely in the following context. Of course the situation is only interesting when there are atomic models, but because we are starting with a complete sentence in $L_{\omega_1,\omega}$, this is automatic.

Assumption 19.3. $\boldsymbol{K}$ is the class of atomic models of a complete first order theory T in a countable vocabulary. Note that with $\prec_{\boldsymbol{K}}$ as $\prec$, first order elementary submodel, this is an abstract elementary class with Löwenheim number $\aleph_0$. Moreover, $\boldsymbol{K}$ is $\aleph_0$-categorical and every member of $\boldsymbol{K}$ is $\aleph_0$-homogeneous. We write $\mathbb{M}$ for the monster model of T; in interesting cases $\mathbb{M}$ is not in $\boldsymbol{K}$. *To stress this,* *we will write $\mathbb{M}'$ for the monster model of $\boldsymbol{K}$ – if we succeed in proving one exists. By the monster model for $\boldsymbol{K}$, we mean a model that is sufficiently homogeneous-universal for embeddings* over submodels*, not over arbitrary substructures.*

Note that the really interesting cases of this context have very complicated languages that arise by reduction from $L_{\omega_1,\omega}$-sentences in natural languages. Some of our examples are in very simple languages. They illustrate some particular points but are not representative of interesting atomic classes.

Unlike Part III, the default notion of type here is a first order type; but of a restricted sort. We will connect this development with Galois types in Lemma 25.13. There are several relevant notions of Stone space in this situation. Recall that $S(A)$ denotes the normal Stone space, all complete first order types over A consistent with $T(A)$, i.e., those realized in $\mathbb{M}$, the monster model of the ambient theory T. We write $S(A)$ with various sub(super)-scripts to mean certain subsets of the space of first order n-types. We will usually take the particular n to be clear from context.

Definition 19.4. 1. *Let A be an atomic set; $S_{\mathrm{at}}(A)$ is the collection of $p \in S(A)$ such that if $\boldsymbol{a} \in \mathbb{M}$ realizes p, $A\boldsymbol{a}$ is atomic.*

2. *Let A be an atomic set; $S^*(A)$ is the collection of $p \in S(A)$ such that p is realized in some $M \in \boldsymbol{K}$ with $A \subseteq M$.*

In speech, it is natural to refer to a $p \in S_{at}(A)$ as an 'atomic type'. Note that this does not imply that p is isolated but only that $p \restriction \boldsymbol{a}$ is isolated for each $\boldsymbol{a} \in A$. A key point is two types in $S^*(A)$ may not be simultaneously realizable in a member of $\boldsymbol{K}$.

Exercise 19.5. *If C is atomic and $\boldsymbol{a}\mathbf{c} \in C$ then $\mathrm{tp}(\boldsymbol{a}; \mathbf{c})$ is principal. Give an example of $p \in S_{at}(M)$ where p is not principal.*

Any atomic model of T is ω-homogeneous and so the unique countable atomic model is universal for countable atomic sets; thus, the following result is straightforward.

Lemma 19.6. *Any* countable *atomic subset of $\mathbb{M}$ can be embedded in a member of $\boldsymbol{K}$.*

Exercise 19.7. *Find an example of an atomic AEC and $M \in \boldsymbol{K}$ which is not strongly ω-homogenous. (Recall that a structure is* strongly ω-homogenous *if any two finite sequences that realize the same type can be mapped to one another by an automorphism of the structure.)*

The first sentence of the next lemma is immediate. The second follows from Lemma 19.6.

Lemma 19.8. *Let A be atomic; $S^*(A) \subseteq S_{at}(A)$. If, in addition, A is* countable*, $S^*(A) = S_{at}(A)$.*

But these identifications may fail for uncountable A; indeed unless $\boldsymbol{K}$ is homogeneous (Definition 5.33), there will be some atomic A and $p \in S_{\mathrm{at}}(A)$ which is not realized in a model in $\boldsymbol{K}$ ([98]). In such a case, even if $\boldsymbol{K}$ has amalgamation over models so there is a monster model $\mathbb{M}'$ for the atomic models there will be types in $S^*(A)$, even for countable A, that are not realized in some *choices* of $\mathbb{M}'$. See Example 19.28. When one has amalgamation over models but not over arbitrary subsets, monsters are unique only up to automorphisms over models.

Now we turn to a pair of serious examples that, because they are not homogenous, illustrate the real problems that confront us in Part IV. We listed the properties of these examples in Fact 4.14 and showed that they gave rise to an $\aleph_1$-categorical sentence of $L_{\omega_1,\omega}$ which had an inhomogeneous model. We now give more details of the construction [107]; we will call on them below to illustrate some of the subtleties of this chapter. A simpler version of Julia Knight is described first.

Example 19.9. There is a first order theory T with a prime model M such that M has no proper elementary submodel but M contains an infinite set of indiscernibles.

1. We begin with the simpler (and later) example of Julia Knight [86]. The vocabulary contains three unary predicates, W, F, I which partition the universe. Let W and I be countably infinite sets and fix an isomorphism f_0 between them. F is the collection of all bijections between W and I that differ from f_0 on only finitely many points. Add also a successor function on W so that (W, S) is isomorphic to ω under successor and the evaluation predicate $E(n, f, i)$ which holds if and only if $n \in W, f \in F, i \in I$ and $f(n) = i$.

   The resulting structure M is atomic and minimal. Since every permutation of I with finite support extends to an automorphism of M, I is a set of indiscernibles. But it is not ω-stable.

2. Now we expound the more complicated example of Marcus. We construct a family of countable sets B_i, languages L_i and groups H^i. The B_i are disjoint;
$$B^j = \cup_{i \leq j} B_i.$$
   Each group H^i will be a group of permutations (in fact of L_i-automorphisms of) B^i with 'almost finite' support: Each $f^i \in H^i$ is determined by its restriction f to B_0 and that restriction moves only finitely many elements of B_0. So in fact each group H^i is a representation of a fixed group H as permutations of B^i.

   (a) B_0 is an infinite set; L_0 is the single unary predicate P_0 (and equality); P_0 holds of all elements of B_0.

   (b) H is the collection of permutations of B_0 with finite support. We define H^i by induction so that each $f \in H$ has a unique extension to

f^i in H^i. Suppose we have defined H^i and B_i. B_{i+1} is the collection of elements $\{b^i_f : f \in H\}$. For each $g, f \in H$, we define g^i by $g^i(b^i_f) = b^i_{gf}$. Then, we set H^{i+1} to be $\{g^i : g \in H\}$.

(c) For $i > 0$, L_i contains symbols $\theta^i_c(u, v)$ for $c \in B_i$ such that for the distinguished element $b^i_1 = b^i_{id} \in B_{i+1}$, indexed by the identity permutation, and for each $c \in B_i$, $\theta^i_c(b^{i+1}_1, v)$ if and only if $v = c$. For other b_f, $\theta^i_c = \{\langle b^{i+1}_f, f(c)\rangle : f \in H^i\}$. Further, the symbol P_j is interpreted as B_j. Add a further predicate to L^i for each of the finitely many automorphism k-types realized in $\bigcup_{j\le i} B_j$.

(d) M is the $L = \bigcup_i L_i$-structure with universe $\bigcup_i B_i$.

Now a couple of observations: $\theta^i_{f^{-1}(c)}(b^{i+1}_f, x)$ iff $x = c$. Thus this formula θ^i defines $c \in B_i$ from $b_f \in B_{i+1}$. So M is minimal. But for each i and k, the automorphism group of M restricted to the group of L_i-automorphisms of $\bigcup_{j<i} B_i$ has only finitely many orbits on k-tuples. Thus M is atomic. And each permutation g of B_0 which moves only finitely many elements extends to an automorphism of M (taking b^i_1 to b^i_g) so B_0 is a set of indiscernibles.

Exercise 19.10. *Show that M has only countably many automorphisms; conclude not every permutation of B_0 extends to an automorphism of M.*

Example 19.11. There is an atomic class $\boldsymbol{K}$ and an uncountable set X which is atomic but is not contained in any member of $\boldsymbol{K}$. Let T be the first order theory of Example 19.9.2 and let $\mathbb{M}$ be its monster model. Embed the unique atomic model of T in $\mathbb{M}$ as M. Now by compactness add an uncountable set $Y \subseteq P_0(\mathbb{M})$ such that each n-tuple from B_0Y realizes a type realized in B_0. Then B_0Y is atomic but there is no atomic model of T which contains B_0Y. If Y is countable there is an atomic model containing B_0Y but it can *not* be embedded in $\mathbb{M}$ over B_0. Note $S^*(Y) \neq S_{at}(Y)$.

Definition 19.12. *The atomic class $\boldsymbol{K}$ is λ-stable if for every $M \in \boldsymbol{K}$ of cardinality λ, $|S_{\mathrm{at}}(M)| = \lambda$.*

To say $\boldsymbol{K}$ is ω-stable in this sense is strictly weaker than requiring $|S_{\mathrm{at}}(A)| = \aleph_0$ for arbitrary countable atomic A ([98]). See Example 19.27. The following example illustrates some further complexities.

Example 19.13. *Consider two structures $(\mathbb{Q}, <)$ and $(\mathbb{Q}, +, <)$. If $\boldsymbol{K}_1$ is the class of atomic models of the theory of dense linear order without endpoints, then $\boldsymbol{K}_1$ is not ω-stable; $\mathrm{tp}(\sqrt{2}; \mathbb{Q}) \in S_{\mathrm{at}}(\mathbb{Q})$. If $\boldsymbol{K}_2$ is the class of atomic models of the theory of the ordered Abelian group of rationals, then $\boldsymbol{K}_2$ is ω-stable; $\mathrm{tp}(\sqrt{2}; \mathbb{Q}) \notin S_{\mathrm{at}}(\mathbb{Q})$.*

Note that these notions are sensitive to naming constants. $\boldsymbol{K}_1$ is unstable; but if we name all the rationals (as the definable closure of a finite set), we get a class with exactly one model that is trivially ω-stable. By passing to an atomic class, we are coding $L_{\omega_1,\omega}$ and the infinite disjunctions mean that by being able to refer to countably many individuals we are essentially quantifying over a countable set.

Exercise 19.14. *Prove that if $\boldsymbol{K}$ is ω-stable then for every $M \in \boldsymbol{K}$ and every countable $A \subseteq M$, only countably many types over A are realized in M. Note that $S^*(A)$ may still be uncountable. (Again, we have not assumed amalgamation.)*

We apply on the following version of Theorem 6.2.5.

Fact 19.15 (Keisler)**.** *Let T be a consistent theory in $L_{\omega_1,\omega}(Q)$. If uncountably many types (in a countable fragment containing T) over the empty set are realized in some uncountable model of T, then T has 2^{ω_1} models of cardinality $\aleph_1$. (In particular, this applies to the class of atomic models of a first order theory).*

We use the weak continuum hypothesis, $2^{\aleph_0} < 2^{\aleph_1}$, *twice* in the following argument; the first time is by the reference to Chapter 18. We state the result for atomic classes but it extends to complete sentences of $L_{\omega_1,\omega}(Q)$.

Theorem 19.16. *Assume $2^{\aleph_0} < 2^{\aleph_1}$. If $\boldsymbol{K}$ is an atomic class with fewer than the maximal number of models in $\aleph_1$, then $\boldsymbol{K}$ is ω-stable.*

Proof. To apply Fact 19.15, fix a fragment L' of $L_{\omega_1,\omega}$ in which $\boldsymbol{K}$ is axiomatized. By Theorem 18.12, $\boldsymbol{K}$ has the amalgamation property in $\aleph_0$. So, if there is a countable model over which there are uncountably many L'-types, there is an uncountable model which realizes uncountably many types over a countable submodel M. But then if we expand the language by adding names for M we obtain by Fact 19.15, that the theory in the expanded language has $2^{\aleph_1}$ models. Since $2^{\aleph_0} < 2^{\aleph_1}$, the original theory does. This proves that for each countable M, $S^*(M)$ is countable and we finish by Lemma 19.8. $\square_{19.16}$

The assumption that $2^{\aleph_0} < 2^{\aleph_1}$ is essential in Theorem 19.16. Note, however, that If $\boldsymbol{K}$ has arbitrarily large models then the methods of Section 7.2 show (in ZFC) that categoricity in $\aleph_1$ and $\aleph_0$ implies ω-stability. Whether, $2^{\aleph_0} < 2^{\aleph_1}$ is essential in Theorem 19.16 depends on the choice of the logic. We sketch two examples and raise an open question.

Example 19.17. Let L contain a unary predicate P and a binary relation $<$. Let $\boldsymbol{K}$ be the models of an $L_{\omega_1,\omega}(Q)$-sentence asserting that M is an $\aleph_1$-dense linear order and P is a countable dense and co-dense subset. Define $M \prec_{\boldsymbol{K}} N$ if and only if $P(M) = P(N)$ and M is a first order elementary submodel of N. Baumgartner [25] proved that it is consistent with ZFC and $2^{\aleph_0} = 2^{\aleph_1} = \aleph_2$ that $\boldsymbol{K}$ is $\aleph_1$-categorical. But clearly it is not ω-stable. Abraham and Shelah [1] showed that Martin's axiom does not suffice for this argument.

Example 19.17 is neither an AEC (not closed under some unions) nor does it have countable Löwenheim number. Example 18.7, which we adduced in Chapter 18 to show the necessity of extending ZFC to prove failure of amalgamation in an AEC, also shows that '$\aleph_1$-categoricity implies ω-stability' requires such additional set theoretic axioms. That example provided a two sorted universe so that one sort Q codes subsets of the other P. Clearly, over any countable model (which necessarily contains all of P), there are uncountably many types of elements of Q which code distinct subsets of P. We would like to actually have a counterexample in $L_{\omega_1,\omega}$. However, although such an example was announced in [138], Laskowski proved that example had $2^{\aleph_0}$ models of cardinality $\aleph_1$. Thus the problem remains open.

We write $\mathrm{diag(A)}$ for the collection of first order formulas $\phi(\boldsymbol{a})$ which are true of sequences $\boldsymbol{a}$ from A. In other words, $\mathrm{diag(A)} = \mathrm{tp}(A/\emptyset)$.

Definition 19.18. *We say B is* atomic over *A if for every $\mathbf{b} \in B$, there is a formula $\phi(\mathbf{x}, \boldsymbol{a})$ such that $\mathrm{diag(A)} \models \phi(\mathbf{x}, \boldsymbol{a}) \rightarrow \psi(\mathbf{x}, \boldsymbol{a}')$ for every $\psi(\mathbf{x}, \boldsymbol{a}')$ with $\boldsymbol{a}' \in A$ such that $\psi(\mathbf{b}, \boldsymbol{a}')$. That is, for every $\mathbf{b} \in B$, $\mathrm{tp}(\mathbf{b}/A)$ is isolated.*

When A is atomic, we can further assume for any $\boldsymbol{a}' \in A$, if $\theta(\mathbf{w}, \mathbf{y})$ generates $\mathrm{tp}(\boldsymbol{a}'\boldsymbol{a}/\emptyset)$ then

$$[\theta(\mathbf{w}, \mathbf{y}) \wedge \phi(\mathbf{x}, \mathbf{y})] \rightarrow \psi(\mathbf{x}, \mathbf{y}).$$

The following result is proved by the standard combination of generating formulas.

Exercise 19.19. *If C is atomic over B and B is atomic over A then C is atomic over A.*

The notion of primary is sometimes called strictly prime or constructible.

Definition 19.20. *Given any sequence $\langle e_i : i < \lambda\rangle$, we write $E_{<j}$ for $\langle e_i : i < j\rangle$. If M can be written as $A \cup \langle e_i : i < \lambda\rangle$ such that $\mathrm{tp}(e_j/AE_{<j})$ is isolated for each j we say M is* primary *over A.*

Definition 19.21. *$M \in \boldsymbol{K}$ is* prime over *A if every elementary map from A into $N \in \boldsymbol{K}$ extends to an elementary map from M into N.*

Now it is an easy induction to show:

Exercise 19.22. *If M is primary over A, then M is atomic over A and prime over A.*

We want to identify those atomic sets (not even all countable ones) over which one can construct a primary model: in Lemma 19.26 we see the following condition suffices for countable A.

Definition 19.23. *The atomic set $A \subset \mathbb{M}$ is* good *if the isolated complete types over A are dense in $S_{at}(A)$.*

Lemma 19.24. *If A and $S_{\mathrm{at}}(A)$ are countable, then A is good.*

Proof. Suppose A is not good. Then there is a satisfiable $\phi(\mathbf{x}, \boldsymbol{a})$ with $\boldsymbol{a} \in A$, but no isolated type in $S_{at}(A)$ contains $\phi(\mathbf{x}, \boldsymbol{a})$. That is, for each $\psi(\mathbf{x}, \mathbf{b})$ with $\mathbf{b} \in A$, such that $\models (\forall \mathbf{x})\psi(\mathbf{x}, \mathbf{b}) \rightarrow \phi(\mathbf{x}, \boldsymbol{a})$, there is a $\mathbf{b}'$ such that $\psi(\mathbf{x}, \mathbf{b})$ has at least two extensions in $S_{at}(\boldsymbol{a}\mathbf{b}\mathbf{b}')$. We will use this to contradict the countability of $S_{at}(A)$.

Enumerate A as $\{a_i : i < \omega\}$. For $\eta \in 2^{<\omega}$, it is straightforward to define by induction $\psi_\eta(\mathbf{x}, \mathbf{b}_\eta)$ so that $\psi_\emptyset(\mathbf{x}, \mathbf{b}_\emptyset)$ is $\phi(\mathbf{x}, \boldsymbol{a})$, if $\eta \lhd \nu$ then $\psi_\eta(\mathbf{x}, \mathbf{b}_\eta) \rightarrow \psi_\nu(\mathbf{x}, \mathbf{b}_\nu)$, $\mathbf{b}_\eta \supseteq a_i$ if $\lg(\eta) > i$, and each $\psi_\eta(\mathbf{x}, \mathbf{b}_\eta)$ isolates a complete type over $\mathbf{b}_\eta$. Our non-density assumption allows us to guarantee that $\psi_{\eta\frown 0}(\mathbf{x}, \mathbf{b}_{\eta\frown 0})$ and $\psi_{\eta\frown 1}(\mathbf{x}, \mathbf{b}_{\eta\frown 1})$ are inconsistent. This gives uncountably many members of $S_{at}(A)$. $\square_{19.24}$

Moreover, if $A\boldsymbol{a}$ is atomic, $|S^n_{\mathrm{at}}(A\boldsymbol{a})| \leq |S^{n+m}_{\mathrm{at}}(A)|$ where $\lg(\boldsymbol{a}) = m$ and the superscript denotes the arity of the type. (Mapping $\mathrm{tp}(\mathbf{d}/A\boldsymbol{a})$ to $\mathrm{tp}(\boldsymbol{a}\mathbf{d}/A)$ is an injection.) So we have:

Lemma 19.25. *If A and $S_{\mathrm{at}}(A)$ are countable, then for any $p \in S_{\mathrm{at}}(A)$, and $\boldsymbol{a}$ realizing p, $A\boldsymbol{a}$ is good.*

Lemma 19.25 was almost immediate; after developing considerable machinery of ω-stable atomic classes in Chapter 20, we prove the converse in Lemma 25.5.

If A is good the isolated types are dense in the Stone space of types over the empty set in the theory obtained by naming the elements of A. Any prime model of a countable first order theory is primary over the empty set. Thus, we have the following result.

Lemma 19.26. *If A is countable and good (hence atomic), there is a primary model over A.*

Several authors [86, 91, 96] have constructed examples of good atomic sets A over which there is no atomic model; results of the same authors show such an A must have cardinality at least $\aleph_2$.

Example 19.27. Let B_0 be the set of indiscernibles in Example 19.9. Then B_0 is a countable atomic set that is not good, as the possible conjunctions of the formulas $\theta_c(x, d) \wedge P_1(x)$ (as d ranges through B_0) give rise to $2^{\aleph_0}$ non-isolated types (of elements in B_1) in $S_{at}(B_0)$. By Corollary 25.5, B_0 is not good.

Example 19.28. Let M be unique countable model of Example 19.9 and B_0 as in that example. Note that every $p \in S_{at}(M)$ is realized in M. (Otherwise M would have a proper extension.) Similarly every stationary type over a finite subset of M is algebraic. But there are $p \in S_{at}(B_0)$ that are not realized in M.

20
Independence in ω-stable Classes

We work in an atomic $\boldsymbol{K}$ which is ω-stable. We do not assume $\boldsymbol{K}$ has amalgamation; we will soon prove (Corollary 20.14) that it does have amalgamation for countable models. We define a notion of independence and in this Chapter we describe the basic properties of this relation. The most difficult of these is symmetry. As in Chapter II of [8], we can use these basic facts as properties of an abstract dependence relation to demonstrate other technical conditions.

We begin by defining an appropriate rank function. Condition 3b) is the crucial innovation to deal with atomic classes.

Definition 20.1. *Let $N \in \boldsymbol{K}$ and $\phi(\mathbf{x})$ be a formula with parameters from N. We define $R_N(\phi) \geq \alpha$ by induction on α.*

1. *$R_N(\phi) \geq 0$ if ϕ is realized in N.*
2. *For a limit ordinal δ, $R_N(\phi) \geq \delta$, if $R_N(\phi) \geq \alpha$ for each $\alpha < \delta$.*
3. *$R_N(\phi) \geq \alpha + 1$ if*
   - *(a) There is an $\boldsymbol{a} \in N$ and a formula $\psi(\mathbf{x}, \mathbf{y})$ such that both $\phi(x) \wedge \psi(x, \boldsymbol{a})$ and $\phi(x) \wedge \neg\psi(x, \boldsymbol{a})$ have rank at least α;*
   - *(b) for each $\mathbf{c} \in N$ there is a formula $\chi(x, \mathbf{c})$ isolating a complete type over $\mathbf{c}$ and $\phi(x) \wedge \chi(x, \mathbf{c})$ has rank at least α.*

We write $R_N(\phi)$ is -1 if ϕ is not realized in N. As usual the rank of a formula is the least α such that $R_N(\phi) \not\geq \alpha + 1$, and $R_N(\phi) = \infty$ if it is greater than or equal every ordinal. And we let the rank of a type be the minimum of the ranks of (finite conjunctions of) formulas in the type. For every type p, there is a formula ϕ in p with $R(p) = R(\phi)$.

Note that the rank of a formula $\phi(\mathbf{x}, \mathbf{b})$ in N depends only on the formula $\phi(\mathbf{x}, \mathbf{y})$ and the type of $\mathbf{b}$ over the empty set in the sense of N. We write 'in N' because there is no assumption that the amalgamation property holds; see Example 20.2. We will drop the subscript N when the context is clear. One can easily prove by induction that if $p \subset q$, $R(p) \geq R(q)$ and that there is a *countable* ordinal α such that if $R(p) > \alpha$ then $R(p) = \infty$.

Example 20.2. Let L contain a binary relation symbol E and three unary predicates for 'black', 'green' and 'red' . Let $\boldsymbol{K}$ be the class of structures such that E is an equivalence relation; each class has at most three elements; every point is colored by one of the three colors. Three elements in the same equivalence class must have distinct colors. (Note that there may be two element classes containing elements of the same color.) This is a universal first order theory. Let $\prec_{\boldsymbol{K}}$ be substructure, $\subseteq$. Then amalgamation fails. It is easy to check that $(\boldsymbol{K}, \subseteq)$ is an aec and in any reasonable sense of the word it is ω-stable. But it is not an atomic class; nor is it $\aleph_0$-categorical.

Theorem 20.3. *If $\boldsymbol{K}$ is ω-stable, then for any $M \in \boldsymbol{K}$ and any $p \in S_{at}(M)$, $R(p) < \infty$.*

Proof. Suppose not; i.e. there is a type $p \in S_{\mathrm{at}}(M)$ with M atomic and $R(p) = \infty$. Since the rank of a type is determined by a formula we can assume M is countable. Then for any finite subset C of M and any $c \in M$ there are finite C' containing cC and $p' \in S_{\mathrm{at}}(C')$ such that $p \restriction C = p' \restriction C$, but $p \restriction C'$ and $p' \restriction C'$ are contradictory and principal. (For this, note that if $\phi(x)$ generates $p \restriction C$, $R(\phi) = \infty \geq \omega_1 + 2$. Thus there is $\boldsymbol{a}$ and ψ witnessing 3a) of the definition of rank so both $\phi(x) \wedge \psi(x, \boldsymbol{a})$ and $\phi(x) \wedge \neg\psi(x, \boldsymbol{a})$ have rank at least $\omega_1 + 1$. Letting $C' = C\boldsymbol{a}c$ and applying 3b) we find a complete extension $p' \neq p$ over C' with rank at least ω_1 and so with infinite rank.)

Thus, we can choose by induction finite sets C_s and formulas ϕ_s for $s \in 2^{<\omega}$ such that:

1. If $s \subset t$, $C_s \subset C_t$ and $\phi_t \to \phi_s$.
2. For each $\sigma \in 2^\omega$, $\bigcup_{s \subset \sigma} C_s = M$.
3. $\phi_{s0}(x)$ and $\phi_{s1}(x)$ are over C_s and each generates a complete type over C_s.
4. ϕ_{s0} and ϕ_{s1} are contradictory.

In this construction the fact that we choose C' above to include an arbitrary $c \in M$ allows us to do 2) and the ϕ_{s0} and ϕ_{s1} generate appropriate choices of $p \restriction C_s, p' \restriction C_s$. Now, each p_σ generated by $\langle \phi_s : s \subset \sigma \rangle$ is in $S_{\mathrm{at}}(M)$ by conditions 2) and 3) so we contradict ω-stability. $\square_{20.3}$

Definition 20.4. *A complete type p over A* splits *over $B \subset A$ if there are $\mathbf{b}, \mathbf{c} \in A$ which realize the same type over B and a formula $\phi(\mathbf{x}, \mathbf{y})$ with $\phi(\mathbf{x}, \mathbf{b}) \in p$ and $\neg\phi(\mathbf{x}, \mathbf{c}) \in p$.*

We will want to work with extensions of sets that behave much like elementary extension.

Definition 20.5. *Let $A \subset B \subseteq \mathbb{M}$. We say A is* Tarski-Vaught *in B and write $A \leq_{\mathrm{TV}} B$ if for every formula $\phi(\mathbf{x}, \mathbf{y})$ and any $\boldsymbol{a} \in A$, $\mathbf{b} \in B$, if $\mathbb{M} \models \phi(\boldsymbol{a}, \mathbf{b})$ there is a $\mathbf{b}' \in A$ such that $\mathbb{M} \models \phi(\boldsymbol{a}, \mathbf{b}')$.*

Note that every atomic extension of a model is a Tarski-Vaught extension:

Exercise 20.6. *If $M \in \boldsymbol{K}$ and MB is atomic then $M \leq_{\mathrm{TV}} MB$.*

The next two lemmas allow us to find nonsplitting extensions of types over models and in certain cases over good sets.

Lemma 20.7 (Weak Extension)**.** *For any $p \in S_{\mathrm{at}}(A)$; if $A \leq_{\mathrm{TV}} B$, B is atomic and p does not split over some finite subset C of A, there is a unique extension of p to $\hat{p} \in S_{\mathrm{at}}(B)$ which does not split over C.*

Proof. Put $\phi(\mathbf{x}, \mathbf{b}) \in \hat{p}$ if and only if for some $\mathbf{b}'$ in A, which realizes the same type as $\mathbf{b}$ over C, $\phi(\mathbf{x}, \mathbf{b}') \in p$. It is easy to check that $\hat{p}$ is well-defined, consistent, and doesn't split over C, let alone A. By the Tarski-Vaught property, $\hat{p}$ is complete. Suppose for contradiction that $\hat{p} \notin S_{\mathrm{at}}(B)$. Then for some $\mathbf{e}$ realizing $\hat{p}$ and some $\mathbf{b} \in B$, $C\mathbf{be}$ is not an atomic set. By Tarski-Vaught again let $\mathbf{b}' \in A$ realize $\mathrm{tp}(\mathbf{b}/C)$; since $\mathbf{e}$ realizes $\hat{p} \restriction A = p \in S_{\mathrm{at}}(A)$, there is a formula $\theta(\mathbf{x}, \mathbf{y}, \mathbf{z})$ that implies $\mathrm{tp}(\mathbf{cb}'\mathbf{e}/\emptyset)$. By the definition of $\hat{p}$, $\theta(\mathbf{cb}, \mathbf{x}) \in \hat{p})$. Thus, $\theta(\mathbf{cbe})$ holds and $C\mathbf{be}$ is an atomic set after all.

For the uniqueness, suppose p_1, p_2 are distinct extensions of p to $S_{\mathrm{at}}(B)$ which do not split over C, but there is a sequence $\mathbf{b} \in B$ and a $\phi(\mathbf{x}, \mathbf{y})$ with $\phi(\mathbf{x}, \mathbf{b}) \in p_1$ but $\neg\phi(\mathbf{x}, \mathbf{b}) \in p_2$. Since $A \leq_{\mathrm{TV}} B$ and B is atomic, there is a $\mathbf{b}' \in A$ with $\mathbf{b}' \equiv_C \mathbf{b}$. Now p_1 a nonsplitting extension of p implies $\phi(x, \mathbf{b}') \in p$, while p_2 a nonsplitting extension of p implies $\neg\phi(x, \mathbf{b}') \in p$. This contradiction yields the theorem. $\square_{20.7}$

Lemma 20.8. *[Existence] Let $\boldsymbol{K}$ be ω-stable. Suppose $p \in S_{\mathrm{at}}(M)$ for some $M \in \boldsymbol{K}$. Then there is a finite $C \subset M$ such that p does not split over C.*

Proof. Choose finite $\mathbf{c} \in M$ and $\phi(\mathbf{x}, \mathbf{c})$ with $R_M(p) = R_M(\phi(\mathbf{x}, \mathbf{c})) = \alpha$. If p splits over $\mathbf{c}$, it is easy to construct two contradictory formulas over M which have rank α. Moreover, for any $\mathbf{e} \in M$, since $p \restriction \mathbf{ce}$ is principal, we can satisfy 2b) of Definition 20.1 and contradict $R_M(\phi(\mathbf{x}, \mathbf{c})) = \alpha$. Thus, p does not split over $\mathbf{c}$. $\square_{20.8}$

Theorem 20.9 (Unique Extension)**.** *If $p \in S_{\mathrm{at}}(M)$ (so p does not split over some finite subset C of M) and $M \subseteq B$, where B is atomic, then there is a unique extension of p to $\hat{p} \in S_{\mathrm{at}}(B)$ which does not split over C.*

Proof. By Lemma 20.8, there is a finite $C \subset M$ such that p does not split over C. Since M is a model $M \leq_{\mathrm{TV}} B$. By Lemma 20.7, p has a unique extension to $S_{\mathrm{at}}(B)$ which does not split over C. $\square_{20.9}$

Example 20.10. The assumption that the base in Theorem 20.9 is a model is essential. Consider the theory of expanding equivalence relations. That is, fix a language with infinitely many binary relations $E_n(x, y)$. Let T state that each E_n is an equivalence relation with infinitely many classes, all of which are infinite and with E_n finer than E_{n+1}. T is ω-stable with elimination of quantifiers. Let $\boldsymbol{K}$ be class of models of T which omit the type $p(x, y) = \{\neg E_n(x, y) : n < \omega\}$. Then $\boldsymbol{K}$ is excellent and even homogenous. But no type over the empty set has a non-splitting extension to a model.

Thus, even in ω-stable atomic classes, splitting may not be well behaved over arbitrary subsets. In particular, Example 20.10 is not simple in the sense of [70].

Definition 20.11. *Let ABC be atomic. We write $A \underset{C}{\downarrow} B$ and say A is free or independent from B over C if for any finite sequence $\boldsymbol{a}$ from A, $\mathrm{tp}(\boldsymbol{a}/B)$ does not split over some finite subset of C.*

Shelah [124] defines stable amalgamation only when the base C is good. Under his hypothesis, the types over the base are stationary.

The independence notion defined in Definition 20.11 satisfies many of the same properties as non-forking on an ω-stable first order theory but with certain restrictions on the domains of types. In some ways, the current setting is actually simpler that the first order setting. Every model is 'ω-saturated' in the sense that if A is a finite subset of $M \in \boldsymbol{K}$ and $p \in S_{\mathrm{at}}(A)$ then p is realized in M. In particular note that monotonicity and transitivity are immediate.

Fact 20.12.

- *[Monotonicity]*
  1. $A \underset{C}{\downarrow} B$ *implies* $A \underset{C'}{\downarrow} B$ *if* $C \subseteq C' \subseteq B$
  2. $A \underset{C}{\downarrow} B$ *implies* $A \underset{C}{\downarrow} B'$ *if* $C \subseteq B' \subseteq B$
- *[Transitivity] If $B \subseteq C \leq_{\mathrm{TV}} D$, $A \underset{C}{\downarrow} D$ and $A \underset{B}{\downarrow} C$ implies $A \underset{B}{\downarrow} D$.*

Defining a well-behaved independence relation from splitting is bit subtle.

Exercise 20.13. Consider the theory T of an equivalence relation with infinitely many classes, which are all infinite. Then T is atomic. Show that while the notion of independence defined in Definition 20.11 behaves well the property $\mathrm{tp}(a/B)$ does not split over C fails to satisfy transitivity – even if all the sets involved are models. (Let D be a model of T; let C add one new class and B add another class. Then let a be a new element of the last class.)

Note that there is no difficulty in applying Theorem 20.9 when p is not the type of a finite sequence but the type of an infinite set A, provided AM is atomic. We formalize this observation using the independence notation. Recall again that there is a global hypothesis of ω-stability in this chapter.

Corollary 20.14. *Suppose $M \subseteq B$, $M \subset A$ and that both A and B are atomic and $A \underset{M}{\downarrow} B$.*

1. *There exists $A' \supset M$, $A' \approx_M A$ and $A' \underset{M}{\downarrow} B$.*
2. *If $A' \approx_M A$ and $A' \underset{M}{\downarrow} B$ then $A' \approx_B A$.*
3. *In particular, amalgamation holds in $\aleph_0$.*

Proof. 1) and 2) Both existence and uniqueness are from Lemma 20.9; existence is evident and the uniqueness follows because the type of each finite subsequence of A does not split over some finite subset of M and so has a unique nonsplitting extension over B.

2) Suppose $M_1, M_2 \in \boldsymbol{K}$ are each countable extensions of M_0. By extension, Theorem 20.9 there exist $M_1', M_2' \in \boldsymbol{K}$ with $M_1' \underset{M_0}{\downarrow} M_2'$. But then $M_1'M_2'$ is atomic and so, by Lemma 19.6, is contained in a countable model as required. $\square_{20.14}$

However, if A' realizes the extension type as in Corollary 20.14, we know only that NBA' is atomic; we have not guaranteed that it is good or (except for countable sets) even embeddible in a member of $\boldsymbol{K}$. Moreover, $A' \approx_B A$ means the identity map on B extends to an isomorphism from AB onto $A'B$. But [6], this isomorphism may not extend to an automorphism of $\mathbb{M}$ (even when $\boldsymbol{K}$ is $\aleph_1$-categorical). We will obtain the stronger result as a corollary to $(\lambda, 2)$-uniqueness (Definition 23.5).

Apply the monotonicity property, the finite character of splitting, and Lemma 20.8 to show.

Exercise 20.15. *Suppose $M_1 \underset{M_0}{\downarrow} M_2$. For any $a, b \in M_1, M_2$, there exist countable M^a with $a \in M^a \prec_{\boldsymbol{K}} M_1$, countable M^b with $b \in M^b \prec_{\boldsymbol{K}} M_2$, and countable $M_0' \prec_{\boldsymbol{K}} M_0$ such that $M^a \underset{M_0'}{\downarrow} M^b$.*

Lemma 20.16 (Coextension)**.** *If M is countable and $M \prec N \in \boldsymbol{K}$ then for any finite $\boldsymbol{a}$ with $\mathrm{tp}(\boldsymbol{a}/M) \in S_{\mathrm{at}}(M)$, if $N \underset{M}{\downarrow} \boldsymbol{a}$, then $\mathrm{tp}(\boldsymbol{a}/N) \in S_{\mathrm{at}}(N)$.*

Proof. Since M is countable there is an M' in $\boldsymbol{K}$ with $M \prec_{\boldsymbol{K}} M'$ and $M\boldsymbol{a} \subset M'$. By the extension property for splitting choose a copy N' of N such that $N' \underset{M}{\downarrow} M'$. The extension property, Lemma 20.9, guarantees $N'M'$ and *a fortiori* $N'\boldsymbol{a}$ is atomic and $\mathrm{tp}(\boldsymbol{a}/N') \in S_{\mathrm{at}}(N')$. Since both $N' \underset{M}{\downarrow} a$ and $N \underset{M}{\downarrow} a$, Corollary 20.14 guarantees $N \equiv_{Ma} N'$ so $\mathrm{tp}(\boldsymbol{a}/N) \in S_{\mathrm{at}}(N)$ as required. $\square_{20.16}$

Definition 20.17. *1. $p \in S_{\rm at}(A)$ is* stationary *if there is a finite sequence $\boldsymbol{a} \in A$ and a formula $\phi(\mathbf{x}, \boldsymbol{a}) \in p$ and M containing A with $q \in S_{\rm at}(M)$ that extends p such that $\phi(\mathbf{x}, \boldsymbol{a}) \in q$ and $R(\phi(x, \boldsymbol{a})) = R(p) = R(q)$.*

2. If $p \in S_{\rm at}(B)$ does not split over $A \subseteq B$ and $p \restriction A$ is stationary then we say p is based *on A.*

The following equivalent property will be our working definition of stationary.

Lemma 20.18. *Suppose A is finite. If $p \in S_{\rm at}(A)$ is stationary then for any atomic B containing A, there is a unique non-splitting over A extension of p to $\hat{p} \in S_{\rm at}(B)$.*

Proof. Fix $M \in \boldsymbol{K}$ containing A let $q \in S_{at}(M)$ be as guaranteed in Definition 20.17.1. By the arguments in Lemmas 20.8 q does not split over A; by 20.7 (taking M as B) q is the unique non-splitting extension of p to $S_{\rm at}(M)$. Put $\phi(\mathbf{x}, \mathbf{b}) \in \hat{p}$ if and only if for some $\mathbf{b}'$ in M, which realizes the same type as $\mathbf{b}$ over A, $\phi(\mathbf{x}, \mathbf{b}') \in q$. Since M is ω-saturated, it is easy to check that $\hat{p}$ is well-defined, consistent, and doesn't split over A. Suppose for contradiction that $\hat{p} \notin S_{\rm at}(B)$. Then for some $\mathbf{e}$ realizing $\hat{p}$ and some $\mathbf{b} \in B$, $A\mathbf{be}$ is not an atomic set. Let $\mathbf{b}' \in M$ realize $\mathrm{tp}(\mathbf{b}/A)$. By definition, for any $\theta(\mathbf{x}, \mathbf{y}, \mathbf{z})$, $\theta(\mathbf{x}, \mathbf{b}, \boldsymbol{a}) \in \hat{p}$ if and only if $\theta(\mathbf{x}, \mathbf{b}', \boldsymbol{a}) \in q$. But $q \restriction \boldsymbol{a}\mathbf{b}'$ is principal so $\hat{p} \restriction \mathbf{b}\boldsymbol{a}$ is principal as required. To see $\hat{p}$ is unique, apply Lemma 20.7 again. $\square_{20.18}$

Note this does not imply for arbitrary atomic classes and large B that $\hat{p}$ is realized in a model. We have justified the following notation.

Notation 20.19. *If $p \in S_{\rm at}(A)$ is based on $\boldsymbol{a}$, for any $B \supseteq \boldsymbol{a}$ we denote by $p|B$ the unique nonsplitting extension of $p \restriction \boldsymbol{a}$ to $S_{\rm at}(B)$.*

Lemma 20.20. *1. If $p \in S_{\rm at}(M)$, then*

(a) p is stationary

(b) there is a finite $C \subset M$ such that p does not split over C and $p \restriction C$ is stationary.

2. If $\boldsymbol{K}$ is ω-stable, $\boldsymbol{K}$ is stable in every cardinality.

Proof. 1a) is immediate from Theorem 20.8; for 1b) choose a formula over a finite subset with the same rank. 2) is now easy since each type is based on a finite set. $\square_{20.20}$

Theorem 20.21. *Suppose $\boldsymbol{K}$ is ω-stable then $\mathbf{b} \underset{M}{\perp} \boldsymbol{a}$ implies $\boldsymbol{a} \underset{M}{\perp} \mathbf{b}$.*

Proof. Suppose for contradiction that $\mathbf{b} \underset{M}{\perp} \boldsymbol{a}$ and $\boldsymbol{a} \underset{M}{\not\perp} \mathbf{b}$. By stationarity of types over models, if $\boldsymbol{a}' \equiv_M \boldsymbol{a}$, $\mathbf{b}' \equiv_M \mathbf{b}$, and $\boldsymbol{a}' \in N$ with $\mathbf{b}' \underset{M}{\perp} N$, then $\mathbf{b}\boldsymbol{a} \equiv_M \mathbf{b}'\boldsymbol{a}'$; thus $\boldsymbol{a}' \underset{M}{\not\perp} \mathbf{b}'$. Construct by induction a continuous increasing chain of countable models M_i and elements $\boldsymbol{a}_i$ for $i < \aleph_1$ such that $M_0 = M$,

$\boldsymbol{a}_0 = \boldsymbol{a}$, $\boldsymbol{a}_i \equiv_M \boldsymbol{a}$, M_{i+1} is a countable model containing $M_i\boldsymbol{a}_i$ and $\boldsymbol{a}_i \underset{M}{\downarrow} M_i$. Let $N = \bigcup_{i<\aleph_1} M_i$. Without loss of generality, the universe of N is $\aleph_1$. Let J be an initial segment of $\aleph_1$. Note that if $i \notin J$, $\boldsymbol{a}_i \underset{M_0}{\downarrow} \bigcup_{i\in J} M_i$. Expand N by adding a name $<$ for the order on $\aleph_1$, names for the elements of M, and a relation symbol $R(x, y)$ such that $R(a, \alpha)$ holds if and only $a \in M_\alpha$. The properties of the construction described above (including $\boldsymbol{a}_i \underset{M}{\downarrow} M_i$) can now all be formalized in $L_{\omega_1,\omega}$, in the expanded language. By the Lopez-Escobar theorem (Theorem 6.1.6, we can choose a countable model $M_{\mathbb{Q}} = \bigcup_{i\in\mathbb{Q}} M_i$ with a new sequence of elements $\boldsymbol{a}_i$ for $i \in \mathbb{Q}$ such that for each i, $\boldsymbol{a}_i \equiv_M \boldsymbol{a}$, and if $i < j$, M_j is a countable model containing $M_i\boldsymbol{a}_i$ with $\boldsymbol{a}_i \underset{M_0}{\downarrow} M_i$. For each initial segment J of $\mathbb{Q}$, let M_J denote $\bigcup_{j\in J} M_j$. Since it holds in the well-ordered case, if $i \notin J$, $\boldsymbol{a}_i \underset{M_0}{\downarrow} M_J$.

Claim 20.22. *There is an element $\mathbf{b}_J$ such that if $i \notin J$, $\boldsymbol{a}_i \underset{M_0}{\downarrow} \mathbf{b}_J$, $\mathbf{b}_J \equiv_M \mathbf{b}$ and $\mathbf{b}_J \underset{M}{\downarrow} M_J$. This implies that $\boldsymbol{a}_i \underset{M_0}{\downarrow} \mathbf{b}_J$ if and only if $i \notin J$.*

Proof. First apply Lemma 20.9 to choose $\hat{\mathbf{b}}_J$ realizing $\mathrm{tp}(\mathbf{b}/M)$ with $\hat{\mathbf{b}}_J \underset{M}{\downarrow} M_J$ and $\mathrm{tp}(\hat{\mathbf{b}}_J/M_J) \in S_{\mathrm{at}}(M_J)$. By the second sentence of the proof of Theorem 20.21, we have $\boldsymbol{a}_i \underset{M}{\not\downarrow} \hat{\mathbf{b}}_J$ if $i \in J$. By Lemmas **??** and 19.26 we can choose $\hat{M}_J \in \boldsymbol{K}$ containing $M_J\hat{\mathbf{b}}_J$ and then by Lemma 20.9 choose $M^J_{\mathbb{Q}}$ to realize the same type as $M_{\mathbb{Q}}$ over M_J but so that $M^J_{\mathbb{Q}} \underset{M_J}{\downarrow} \hat{M}_J$. For $i \notin J$, denote by $\boldsymbol{a}^J_i$ the natural image of a_i in $M^J_{\mathbb{Q}}$. Since $M \leq_{TV} M_J \leq_{TV} \hat{M}_J$ and $\boldsymbol{a}^J_i \underset{M}{\downarrow} M_J$ if $i \notin J$, transitivity of nonsplitting yields $\boldsymbol{a}^J_i \underset{M}{\downarrow} \hat{\mathbf{b}}_J$. Now map $M^J_{\mathbb{Q}}$ onto $M_{\mathbb{Q}}$, fixing M_J and taking a^J_i to a_i for all $i \notin J$. The image of $\hat{\mathbf{b}}_J$ under this map is the required $\mathbf{b}_J$. Since $\boldsymbol{a}_i \in M_J$ for $i \in J$ and $\mathbf{b}_J \equiv_{M_J} \hat{\mathbf{b}}_J$, $\boldsymbol{a}_i \underset{M}{\not\downarrow} \mathbf{b}_J$ if $i \in J$. $\square_{20.22}$

We have found $2^{\aleph_0}$ types over $M_{\mathbb{Q}}$ (the $\mathrm{tp}(\mathbf{b}_J/M_{\mathbb{Q}})$ for J an initial segment of $\mathbb{Q}$). We must verify that they are in $S_{\mathrm{at}}(M_{\mathbb{Q}})$. We have constructed the $\hat{\mathbf{b}}_J \in \hat{M}_J$ so that for each initial segment J of $\mathbb{Q}$, $\mathrm{tp}(\hat{\mathbf{b}}_J/M_J)$ is in $S_{\mathrm{at}}(M_J)$ and $M^J_{\mathbb{Q}} \underset{M_J}{\downarrow} \hat{M}_J$. Then by the coextension lemma, Lemma 20.16, $\mathrm{tp}(\hat{\mathbf{b}}_J/M^J_{\mathbb{Q}}) \in S_{\mathrm{at}}(M^J_{\mathbb{Q}})$. So $\mathrm{tp}(\mathbf{b}_J/M_{\mathbb{Q}}) \in S_{\mathrm{at}}(M_{\mathbb{Q}})$; the result is proved. $\square_{20.21}$

We can extend symmetry to stationary types as follows. Note that the stationarity conditions are immediate if A is good.

Lemma 20.23. *Let A be a countable atomic set and suppose $p = \mathrm{tp}(a/A), q = \mathrm{tp}(b/A), r = \mathrm{tp}(ab/A) \in S_{\mathrm{at}}(A)$ are each based on finite subsets of A. Then $a \underset{A}{\downarrow} b$ if and only $b \underset{A}{\downarrow} a$.*

Proof. Suppose $a \underset{A}{\downarrow} b$. Let $M'' \subset A$ be a model. By stationarity we can choose b' realizing p with $b' \underset{A}{\downarrow} M''$. Mapping b' back to b over A yields an M' with $b \underset{A}{\downarrow} M'$. Since $a \underset{A}{\downarrow} b$, stationarity allows us to assume with $a \underset{A}{\downarrow} M'b$. By symmetry over models, we have $b \underset{A}{\downarrow} M'a$. Since M' is Tarski-Vaught in $M'a$, we can conclude by transitivity that $b \underset{A}{\downarrow} M'a$ and by monotonicity $b \underset{A}{\downarrow} a'$. $\square_{20.23}$

Question 20.24. *Does the 'pairs lemma' hold in this context (for stationary types)?*

Exercise 20.25. *If $\boldsymbol{a} \underset{M}{\downarrow} N$ then $\boldsymbol{a} \cap N \subset M$.*

In the next theorem, whose proof follows that of Conclusion 2.13 in [124], we weaken the hypothesis of $\aleph_1$-categoricity in Theorem 8.11 to ω-stability; we are still working in ZFC.

Theorem 20.26. *If an atomic class is ω-stable and has a model of power $\aleph_1$ then it has a model of power $\aleph_2$.*

Proof. As in Chapter 8, it suffices to show every model N in $\boldsymbol{K}$ of cardinality $\aleph_1$ has a proper elementary extension M in $\boldsymbol{K}$. Write N as a continuous increasing chain:
$\langle N_i : i < \aleph_1\rangle$. Now define an increasing sequence $\langle M_i : i < \aleph_1\rangle$ such that

1. M_0 is a proper $\prec_{\boldsymbol{K}}$-extension of N_0;
2. $M_i \underset{N_i}{\downarrow} N$

N_0 has a proper extension M_0' and by Lemma 20.9 we may choose a copy M_0 of M_0' that is independent from N over N_0. At successor stages take a proper extension $M'_{\alpha+1}$ of $M_\alpha N_{\alpha+1}$ (such an extension exists since $M_\alpha N_{\alpha+1}$ is countable and atomic) and by Lemma 20.9 there is a copy $M'_{\alpha+1}$ which is independent of N over $N_{\alpha+1}$. By Lemma 20.14.2, we can choose $M_{\alpha+1}$ to contain M_α and with $M_{\alpha+1} \underset{N_{\alpha+1}}{\downarrow} N$. Take unions at limits to complete the construction. M_{ω_1} is the required proper extension of N. $\square_{20.26}$

The hypothesis that there be a model with cardinality $\aleph_1$ is essential. The Marcus example, Example 4.14, is ω-stable as an atomic class and has exactly one model.

Now we can strengthen Theorem 8.11, replacing categoricity in $\aleph_1$ by few models in $\aleph_1$ at the cost of assuming $2^{\aleph_0} < 2^{\aleph_1}$. The following corollary is immediate since with this set-theoretic hypothesis, few models in $\aleph_1$ implies ω-stability (Lemma 19.16).

Corollary 20.27 ($2^{\aleph_0} < 2^{\aleph_1}$)**.** *If the atomic class $\boldsymbol{K}$ has at least one but fewer than $2^{\aleph_1}$ models of cardinality $\aleph_1$ then it has a model of power $\aleph_2$.*

Question 20.28. *Does Lemma 20.26 extend to $L_{\omega_1,\omega}(Q)$? I think Shelah claims only the analog of Lemma 20.27 in [120].*

21
Good Systems

In this chapter we introduce the notion of P-systems, systems of models indexed by partial orderings. This rather technical chapter will establish crucial tools for our later study of excellent classes. In particular, the generalized symmetry lemma, Theorem 21.21, will be quoted at crucial points in ensuing chapters. We will primarily be interested in systems indexed by the subsets of n, for arbitrary finite n. These systems generalize the independent n-dimensional systems of Chapter 3. They will enable us to analyze categorical structures where no geometry can be defined on the entire universe (i.e. has 'rank' greater than one). We work throughout in an ω-stable atomic class $\boldsymbol{K}$.

We describe the indexing of system of models by a partial order. Since in most of our specific applications, the partial order will be subset on $\mathcal{P}(n)$ for a finite set n, we will just use $\subseteq$ for the reflexive partial order.

Definition 21.1. *If $(P, \subseteq)$ is a partial order and $\mathcal{S} = \langle M_s : s \in P\rangle$ is a collection of models from $\boldsymbol{K}$ then we say $\mathcal{S}$ is a P-system if*

1. *$M_s \prec_{\boldsymbol{K}} M_t$ whenever $s \subseteq t$;*
2. *if $s \subset t$ then $M_s \neq M_t$;*
3. *If P admits a meet $\wedge$ then $M_{s\wedge t} = M_s \cap M_t$*

$\mathcal{S}$ is a (λ, P)-system if each $|M_s| = \lambda$.

Note that the distinction between subsets $\subseteq$ and proper subsets $\subset$ is crucial in the following definition. We will use the notations A_s and B_s repeatedly.

Notation 21.2. 1. *$\mathcal{S}$ is an* independent *P-system if for each $s \in P$,*

(a) $A_s = \bigcup_{t \subset s} M_t$ *is atomic;*

(b) $M_s \underset{A_s}{\downarrow} B_s$ *where* $B_s = \bigcup_{t \not\supseteq s} M_t$.

2. *An independent P-system is called good if in addition*

   (a) *each A_s is good and*

   (b) $A_s \leq_{\mathrm{TV}} B_s$ *for each* $s \in P$.

3. *We write (λ, n)-system (or $(\lambda, \mathcal{P}^-(n))$-system, for emphasis) for a P-system where*
$$P = \mathcal{P}^-(n) = \{s : s \subset n\}.$$

In order to deal with good P-systems, it is necessary to study an apparently weaker condition. Before describing it we need a generalization of the notion of Tarski-Vaught extension.

Definition 21.3. *Let $\overline{A} = \langle A_1, \ldots A_n \rangle$ and $\overline{B} = \langle B_1 \ldots B_n \rangle$ satisfy that each $A_i \subseteq B_i$. We say $\overline{A}$ is* Tarski-Vaught *in $\overline{B}$ and write $\overline{A} \leq_{\mathrm{TV}} \overline{B}$ if for every formula $\phi(\mathbf{x}, \mathbf{y})$ and any $\boldsymbol{a} \in \bigcup A_i$, $\mathbf{b} \in \bigcup B_i$, if $\mathbb{M} \models \phi(\boldsymbol{a}, \mathbf{b})$ there is a $\mathbf{b}' \in A$ such that $\mathbb{M} \models \phi(\boldsymbol{a}, \mathbf{b}')$ and if $b_s \in B_j$ then $b'_s \in A_j$.*

Note that if $\langle A_1, \ldots, A_{n-1} \rangle \leq_{\mathrm{TV}} \langle B_1, \ldots, B_{n-1} \rangle$, the same holds for any subsequence. Definition 21.3 is much stronger than the assertion: $\bigcup A_i \leq_{\mathrm{TV}} \bigcup B_i$; consider the following exercise.

Exercise 21.4. *Show: if $\langle A_1, A_2 \rangle \leq_{\mathrm{TV}} \langle B_1, A_2 \rangle$ then $B_1 \underset{A_1}{\downarrow} A_2$.*

Definition 21.5. *$\langle s_i : i < \alpha \rangle$ is an* enumeration *(or* linearization*) of P if it is a non-repeating list of the elements of P such that $s_i \subset s_j$ implies $i < j$.*

Our next goal is to show that a P-system is good if it is good under some enumeration (Definition 21.12). But we require some preliminary lemmas.

Lemma 21.6. *If $A \underset{C}{\downarrow} B$ and $C \leq_{\mathrm{TV}} B$ then*

1. $\langle C, A \rangle \leq_{\mathrm{TV}} \langle B, A \rangle$
2. *and* $AC \leq_{\mathrm{TV}} AB$.

Proof. We show 1) and 2) follows immediately. Since $A \underset{C}{\downarrow} B$, there is a $\mathbf{c}' \in C$ such that $\mathrm{tp}(\boldsymbol{a}/B)$ does not split over $\mathbf{c}'$. Suppose $\phi(\boldsymbol{a}, \mathbf{b}, \mathbf{c})$ with $\boldsymbol{a} \in A$, $\mathbf{b} \in B$, $\mathbf{c} \in C$ and without loss $\mathbf{c}' \subseteq \mathbf{c}$. Since $C \leq_{\mathrm{TV}} B$, there is a $\mathbf{b}' \in C$ that realizes the same type over $\mathbf{c}$ as $\mathbf{b}$. Then $\phi(\boldsymbol{a}, \mathbf{b}', \mathbf{c})$ or else $\mathrm{tp}(\boldsymbol{a}/B)$ splits over $\mathbf{c}$. $\square_{21.6}$

The same argument shows:

Exercise 21.7. *Suppose $\langle C_0, \ldots C_{n-1} \rangle \leq_{TV} \langle B_0, \ldots B_{n-1} \rangle$ and $A \underset{\bigcup_{i<n} C_i}{\downarrow} \bigcup_{i<n} B_i$, then $\langle C_0, \ldots C_{n-1}, A \rangle \leq_{TV} \langle B_0, \ldots B_{n-1}, A \rangle$.*

Exercise 21.8. *Show that if $A \leq_{\mathrm{TV}} B$ and $\mathrm{tp}(\mathbf{c}/A)$ is isolated then $A\mathbf{c} \leq_{\mathrm{TV}} B\mathbf{c}$*

The following partial Skolemization enables us to move goodness between cardinals.

Lemma 21.9. *There is a countable expansion L^* of L such that every atomic A can be expanded to an L^*-structure A^* so that if B^* is an L^*-substructure of A^* then letting $B = B^* \restriction L$:*

1. *$B \leq_{\mathrm{TV}} A$*
2. *A is good if and only if B is good.*

Proof. For every formula $\phi(\mathbf{x}, \mathbf{y})$ add a function G^ϕ from $\lg(\mathbf{y})$ tuples to $\lg(\mathbf{x})$ tuples. Expand A so that $\phi(G^\phi(\boldsymbol{a}), \boldsymbol{a})$ if $(\exists \mathbf{x})\phi(\mathbf{x}, \boldsymbol{a})$ and define G^ϕ arbitrarily otherwise. This guarantees the Tarski-Vaught condition.

Add also functions $F^{\phi,\psi}$ for each pair of formulas $\phi(\mathbf{x}, \mathbf{y})$, $\psi(\mathbf{x}, \mathbf{v})$. The domain of $F^{\phi,\psi}$ should be tuples of $\lg(\mathbf{y}) + \lg(\mathbf{v})$ tuples, and the range is finite sequences. (Technically, this is incorrect; by adding one more integer parameter to the name of the function we could fix the range properly; but the notation becomes too cumbersome.)

Expand A so that:

1. Suppose $(\exists \mathbf{x})\phi(\mathbf{x}, \boldsymbol{a})$:
   (a) if $\psi(\mathbf{x}, \mathbf{b})$ isolates a complete type over A implying $\phi(\mathbf{x}, \boldsymbol{a})$, then $F^{\phi,\psi}(\boldsymbol{a}) = \mathbf{b}$. (In this case, $F^{\phi,\psi}$ does not depend on the $\lg(\mathbf{v})$ argument).
   (b) if *no* $\psi(\mathbf{x}, \mathbf{b})$ isolates a complete type over A implying $\phi(\mathbf{x}, \boldsymbol{a})$, then for each ψ, there is a formula $\chi(\mathbf{x}, \mathbf{z})$ so that for any $\boldsymbol{a}$ with $\lg(\mathbf{y})$ and any $\mathbf{b}$ with $\lg(\mathbf{v})$, if $\psi(\mathbf{x}, \mathbf{b}) \rightarrow \phi(\mathbf{x}, \boldsymbol{a})$ then $(\exists \mathbf{x})\psi(\mathbf{x}, \mathbf{b}) \wedge \chi(\mathbf{x}, F^{\phi,\psi}(\boldsymbol{a}, \mathbf{b}))$ and $(\exists \mathbf{x})\psi(\mathbf{x}, \mathbf{b}) \wedge \neg\chi(\mathbf{x}, F^{\phi,\psi}(\boldsymbol{a}, \mathbf{b}))$.
2. If $\neg(\exists \mathbf{x})\phi(\mathbf{x}, \boldsymbol{a})$, the choice of $F^{\phi,\psi}(\boldsymbol{a})$ is arbitrary.

It is routine to check that these functions accomplish their goals. $\square_{21.9}$

We find bases for types over good sets A as Lemma 20.20 found bases for types over models. Ideally, we would find a subset of A on which the type is based. But this is impossible:

Example 21.10. Consider a vocabulary with a binary function $+$, and a projection π and an equivalence relation E. There are two sorts X and G. π maps X onto an ω-stable, $\aleph_0$-categorical group $(G, +)$, for definiteness Z_2^ω, such that each element of G is the image of infinitely many points. E splits $\pi^{-1}(g)$ into two infinite classes for each g. Suppose M_1 and M_2 are freely amalgamated over M_0 (i.e. $M_1 \underset{M_0}{\downarrow} M_2$). Then, by counting types $|S_{\mathrm{at}}(M_1 M_2)| = \aleph_0$ and so $M_1 M_2$ is a good set. But the type of an element in the fiber over $g_1 + g_2$ where $g_i \in M_i$, can only be made stationary by naming an element e of the fiber. Thus, when the following lemma is applied to this situation, g_1, g_2 is the $\mathbf{b}$ and $g_1 + g_2, e$ are the $\mathbf{c}$.

A key to the following argument is the observation that $\theta(x; \mathbf{c}) \vdash \mathrm{tp}(\boldsymbol{a}/A)$ means for any $\boldsymbol{a} \in A$, $\{\theta(x, \mathbf{b})\} \cup \mathrm{tp}(\boldsymbol{a}/\mathbf{b}) \vdash \mathrm{tp}(\boldsymbol{a}\mathbf{c}/\mathbf{b})$.

Lemma 21.11. *Suppose $p \in S_{\mathrm{at}}(A)$, and A is good.*

1. *There are $\mathbf{c}, \boldsymbol{a}$ such that:*

   (a) $p = \mathrm{tp}(\boldsymbol{a}/A)$*;*

   (b) $\mathrm{tp}(\boldsymbol{a}/A\mathbf{c}) \in S_{at}(A\mathbf{c})$ *is based on* $\mathbf{c}$*;*

   (c) and $\mathrm{tp}(\mathbf{c}/A)$ *is isolated.*

2. *There is a $\mathbf{b} \in A$ such that p does not split over $\mathbf{b}$.*

Proof. First assume that A is countable. Let $\boldsymbol{a}$ realize p with $A\boldsymbol{a}$ atomic. Since $A\boldsymbol{a}$ is countable there is a model $M \in \boldsymbol{K}$ with $A\boldsymbol{a} \subset M$. Lemma 19.25 implies $A\boldsymbol{a}$ is good. By Lemma 19.26, there is a primary model M' over A, which imbeds in M. Choose $\mathbf{c} \in M'$ with $\mathrm{tp}(\boldsymbol{a}/M')$, *a fortiori* $\mathrm{tp}(\boldsymbol{a}/A\mathbf{c})$, based on $\mathbf{c}$. Since M' is primary over A, $\mathrm{tp}(\mathbf{c}/A)$ is isolated. For 2), let $\mathbf{b} \in A$ be such that some $\theta(\mathbf{b}, \mathbf{x}) \vdash \mathrm{tp}(\mathbf{c}/A)$. Now if $\boldsymbol{a}, \boldsymbol{a}' \in A$ satisfy $\boldsymbol{a} \equiv_{\mathbf{b}} \boldsymbol{a}'$ then $\boldsymbol{a} \equiv_{\mathbf{c}} \boldsymbol{a}'$ and the result follows from part 1).

Now, we extend the result to A of arbitrary cardinality by applying Lemma 21.9. Let L^* be the expansion of L introduced in that Lemma and expand A to an L^*-structure A^*. If there is no such pair $\mathbf{b}, \mathbf{c}$, we can choose by induction $\mathbf{b}_\alpha, \mathbf{c}_\alpha, A_\alpha$ such that:

1. A_α is a countable L^*-substructure of A^*.
2. $\mathbf{b}_\alpha \in A_\alpha$; $\mathbf{c}_\alpha \in A_{\alpha+1}$; $\mathbf{b}_\alpha \subseteq \mathbf{c}_\alpha$.
3. $\mathrm{tp}(a/A_\alpha\mathbf{c}_\alpha)$ is based on $A_\alpha\mathbf{c}_\alpha$.
4. $\mathrm{tp}(a/A_\alpha)$ does not split over $\mathbf{b}_\alpha$.
5. $\mathrm{tp}(a/A_{\alpha+1})$ splits over $\mathbf{b}_\alpha$

Let A_0 be any countable L^*-substructure of A^*. Then choose $\mathbf{b}_0$ and $\mathbf{c}_0$ by the first paragraph of the proof to satisfy conditions 2)-4). By the hypothesis for contradiction there exist a pair of sequences $\mathbf{d}_0^0, \mathbf{d}_0^1$ witnessing that $\mathrm{tp}(\boldsymbol{a}/A)$ splits over $\mathbf{b}_0$. Let A_1 be the L^*-closure of $A_0\mathbf{d}_0^0\mathbf{d}_0^1\mathbf{c}_0$. Continue by induction taking unions at limits. Without loss of generality, we always insist that $\mathbf{b}_\alpha \subseteq \mathbf{c}_\alpha$. Now define f on the limit ordinals less than ω_1 by $f(\alpha)$ is the least γ such that $\mathbf{b}_\alpha \in A_\gamma$. Then f presses down on the limit ordinals so by Fodor's Lemma, there is a γ_0 and a stationary set S of limit ordinal such that if $\alpha \in S$, then $f(\alpha) = \gamma$. Now there are only countably many choices for $\mathbf{b}_\alpha \in A_\gamma$ so for cofinally many

α, $\mathbf{b}_\alpha = \mathbf{b}$ for a fixed $\mathbf{b} \in A_\gamma$. But then for cofinally many α we have $\mathrm{tp}(\boldsymbol{a}/A_\alpha)$ does not split over $\mathbf{b}$. But this contradicts clause 5) in construction. Thus the construction fails at some β and $\mathbf{b}_\beta, \mathbf{c}_\beta$ are as required. $\square_{21.11}$

Now we formalize the notion of being good for an enumeration and show that the choice of enumeration does not matter.

Definition 21.12. *Let $\mathcal{S} = \langle M_s : s \in P\rangle$ be a P-system for the finite partial order P. The system $\mathcal{S}$ is said to be* good with respect to *the enumeration $\langle s_i : i < \alpha\rangle$ if*

1. *A_s is good for each $s \in P$, and*
2. *for each $i < \alpha$, $\langle M_{s_i \wedge s_j} : j \leq i\rangle \leq_{\mathrm{TV}} \langle M_{s_j} : j \leq i\rangle$.*

By 'complication of notation' we can prove a more general version of Lemma 21.9.

Corollary 21.13. *There is a countable expansion L^* of L such that every atomic A can be expanded to an L^*-structure A^* so that if $\overline{C}$ is a sequence of atomic subsets of A (named by L^*-predicates P_i) and B^* is an L^*-substructure of A^* then letting $B = B^* \restriction L$ and $C_i' = P_i(B)$: $\overline{C}' \leq_{\mathrm{TV}} \overline{C}$. Moreover, if $\overline{C} = \mathcal{S}$ is an independent P-system then $\mathcal{S}'$ is good with respect to the given enumeration if and only if $\mathcal{S}$ is.*

Proof. As in the proof of Lemma 21.9 define functions on tuples from $\bigcup_i C_i$ into C_i to reflect goodness and to 'Skolemize'. To deal with the system, Skolemize the smallest models first and then extend. Then an L^*-substructure is Tarski-Vaught for the sequence. $\square_{21.13}$

Ostensibly, the notion of a good enumeration says nothing about independence; but it does imply some configurations are independent.

Lemma 21.14. *Let $\mathcal{S} = \langle M_s : s \in P\rangle$ be a P-system for the finite partial order P, which is enumerated $\langle s_i : i < \alpha\rangle$. If $\mathcal{S}$ is good with respect to this enumeration then for any $1 \leq i < \alpha$,*

1.

$$A_{s_i} \leq_{\mathrm{TV}} \bigcup_{j<i} M_{s_j}$$

and

2.

$$M_{s_i} \underset{A_{s_i}}{\downarrow} \bigcup_{j<i} M_{s_j}$$

Proof. The first conclusion is immediate from the definition of good with respect to the enumeration. For the second, fix $i < \alpha$ and choose $a \in M_{s_i}$. As A_{s_i} is good, by Lemma 21.11.2, there is a $b \in A_{s_i}$ so that $\mathrm{tp}(a/A_{s_i})$ does not split over b. Now suppose for contradiction that there are $c, c' \in \bigcup_{j<i} M_{s_j}$ such that $c \equiv_b c'$ but $c \not\equiv_a c'$. Since $ab \in M_{s_i} = M_{s_i \wedge s_i}$ and

$$\langle M_{s_i \wedge s_j} : j \leq i \rangle \leq_{\mathrm{TV}} \langle M_{s_j} : j \leq i \rangle,$$

there are $d, d' \in A_{s_i}$ so that $dd' \equiv_{ab} cc'$. (Note that for $j > i$, the definition of enumeration implies that for some k, $s_j \wedge s_i = s_k$ with $s_k \subset s_i$.) This contradicts that $\mathrm{tp}(a/A_{s_i})$ does not split over b. $\square_{21.14}$

Our goal, Theorem 21.21, is show to that goodness of a system indexed by a finite partial order is independent of the enumeration. We do this in two steps.

Theorem 21.15. *If P is a finite partial order and $\mathcal{S}$ is a P-system, the following are equivalent:*

1. *$\mathcal{S}$ is good.*
2. *$\mathcal{S}$ is good with respect to each enumeration.*

Proof. To show 2) implies 1) we need to show that for each s, A_s is good, $A_s \leq_{\mathrm{TV}} B_s$, and $M_s \underset{A_s}{\downarrow} B_s$. The first is clear regardless of the choice of enumeration. For the second, choose for each s an enumeration $\langle s_i : i < \alpha \rangle$ (depending on s) so that if $s = s_i$ then $\{s_j : j < i\} = \{t : t \not\supseteq s_i\}$. Lemma 21.14 shows that if $\mathcal{S}$ is good with respect to this enumeration then the remaining two conditions hold for s. Since we can apply this construction to any s, $\mathcal{S}$ is good.

For 1) implies 2) choose an arbitrary enumeration and suppose for induction that $S_j = \langle M_{s_j} : j < i \rangle$ is good. By Definition 21.17,

$$\langle M_{s_j \wedge s_i} : j < i \rangle \leq_{\mathrm{TV}} \langle M_{s_j} : j < i \rangle.$$

As $\mathcal{S}$ is good (and therefore independent)

$$M_{s_i} \underset{\bigcup_{t<s_i} M_t}{\downarrow} \bigcup_{j<i} M_{s_j}.$$

Thus by Lemma 21.6, we have

$$\langle M_{s_j \wedge s_i} : j \leq i \rangle \leq_{\mathrm{TV}} \langle M_{s_j} : j \leq i \rangle$$

as required. $\square_{21.15}$

We now need a variant of the argument for Lemma 20.7.

Lemma 21.16. *Suppose $p \in S_{\mathrm{at}}(A)$ is isolated by $\phi(\mathbf{x}, \boldsymbol{a})$.*

1. *p does not split over $\boldsymbol{a}$.*
2. *If $A \leq_{\mathrm{TV}} B$, where B is atomic and $q \in S_{\mathrm{at}}(B)$ extends p then q is isolated by $\phi(\mathbf{x}, \boldsymbol{a})$. Thus, q does not split over $\boldsymbol{a}$.*

Proof. If p splits there exist $\mathbf{b}, \mathbf{b}' \in A$ that realize the same type over $\boldsymbol{a}$ and a formula $\psi(\mathbf{x}, \mathbf{y})$ with $\psi(\mathbf{x}, \mathbf{b}) \in p$ and $\neg\psi(\mathbf{x}, \mathbf{b}') \in p$. But then $(\exists \mathbf{x})\phi(\mathbf{x}, \boldsymbol{a}) \wedge \neg\psi(\mathbf{x}, \mathbf{b})$ holds and $\phi(\mathbf{x}, \boldsymbol{a})$ does not generate p.

For part 2, if $\phi(\mathbf{x}, \boldsymbol{a})$ does not isolate q there are a $\mathbf{b} \in B$ and a formula, $\psi(\mathbf{x}, \mathbf{b})$ such that

$$(\exists \mathbf{x})\phi(\mathbf{x}, \boldsymbol{a}) \wedge \psi(\mathbf{x}, \mathbf{b}) \wedge (\exists \mathbf{x})\phi(\mathbf{x}, \boldsymbol{a}) \wedge \neg\psi(\mathbf{x}, \mathbf{b}).$$

By Tarski-Vaught, there is a $\mathbf{b}' \in A$ which realizes the same type over $\boldsymbol{a}$ as $\mathbf{b}$. But then $\phi(\mathbf{x}, \boldsymbol{a})$ does not generate p. Now to verify 'thus', apply part 1) to q. $\square_{21.16}$

Here is the most important lemma. The argument may appear unduly complicated. The difficulty is to find a base for the type of $\boldsymbol{a}_s$ over A_s. For this we need to use the description in Lemma 21.11.1. The existence of a point over which the type does not split would suffice for Case 1 in the argument below but we need a stationary type for Case 2.

Lemma 21.17. *Let $\mathcal{S} = \langle M_s : s \in P \rangle$ be a (λ, P)-system for the finite partial order P, which is enumerated $\langle s_i : i < \alpha \rangle$. If $\mathcal{S}$ is good for this enumeration, then for any $m < \alpha$,*

$$\langle M_{s \wedge s_m} : s \in P \rangle \leq_{\mathrm{TV}} \langle M_s : s \in P \rangle.$$

Proof. Fix m and fix $\boldsymbol{a}_s \in M_s$ for $s \in P$. We want to find $\boldsymbol{a}'_s \in M_{s \wedge s_m}$ so that $\boldsymbol{a}'_s = \boldsymbol{a}_s$ if $s = s_i$ with $i < m$ and $\langle \boldsymbol{a}_s : s \in P \rangle \equiv \langle \boldsymbol{a}'_s : s \in P \rangle$. To carry out the induction we expand the $\boldsymbol{a}'s$ by choosing $\mathbf{b}_s$ and $\mathbf{c}_s$ as follows:

Claim 21.18. *1. $\mathbf{b}_s \in A_s$ and $\mathbf{c}_s \in M_s$;*

2. $\mathrm{tp}(\boldsymbol{a}_s / A_s \mathbf{c}_s)$ is based on $\mathbf{c}_s$;

3. $\mathrm{tp}(\mathbf{c}_s / A_s)$ is isolated over $\mathbf{b}_s$;

4. $\mathbf{b}_s \subseteq \bigcup_{t \subset s} \boldsymbol{a}_t$.

Proof of Claim 21.18. Except for condition 4) this is straightforward from goodness using Lemma 21.11 for 3). In order to guarantee condition 4) we have to systematically construct the $\boldsymbol{a}_t$, $\mathbf{b}_t$ and $\mathbf{c}_t$. For simplicity of notation in this proof, we write $\boldsymbol{a}_m$ for $\boldsymbol{a}_{s_m}$ etc. For each $m < \alpha$ and $j \leq \alpha$, we construct $\mathbf{b}^j_m$ and $\mathbf{c}^j_m$ and an auxiliary $\mathbf{d}^j_m$. We do this by a downward induction on $j \leq \alpha$.

For each $m < \alpha$ and $j = \alpha$, let $\mathbf{c}^\alpha_m = \boldsymbol{a}^\alpha_m$ be the original $\boldsymbol{a}_m$. Let $\mathbf{b}^\alpha_m = \mathbf{d}^\alpha_m$ be the empty sequence. Suppose we have defined for $j + 1$; there are two cases for j.

Case 1: $s_i \not\subseteq s_j$: Then, nothing changes: for $m < \alpha$, $\mathbf{c}^j_m = \mathbf{c}^{j+1}_m$, $\mathbf{b}^j_m = \mathbf{b}^{j+1}_m$, $\mathbf{d}^j_m = \mathbf{d}^{j+1}_m$.

Case 2: $s_i \subseteq s_j$: Note $A_{s_j} \subset M_{s_j}$ and $\mathbf{c}^{j+1}_m {}^\frown \mathbf{d}^{j+1}_m \in M_{s_j}$. Apply Lemma 21.11 to choose $\mathbf{c}^*_j$, $\mathbf{b}^*_j$ and $\mathbf{d}^*_j$ so that $\mathrm{tp}(\mathbf{c}^*_j {}^\frown \mathbf{d}^*_j / A_{s_j}) = \mathrm{tp}(\mathbf{c}^{j+1}_m {}^\frown \mathbf{d}^{j+1}_m / A_{s_j})$, $\mathrm{tp}(\mathbf{d}^*_j / A_{s_j} \mathbf{c}^*_j)$ is based on $\mathbf{c}^*_j$ and $\mathrm{tp}(\mathbf{c}^*_j / A_{s_j})$ is isolated over $\mathbf{b}^*_j$. Then for $m < \alpha$, set $\mathbf{c}^j_m = \mathbf{c}^*_j$, $\mathbf{b}^j_m = \mathbf{b}^*_j$, $\mathbf{d}^j_m = \mathbf{d}^j_m \cup (M_{s_m} \cap \mathbf{b}^*_j)$.

This completes the construction and it is routine to see we obtain the claim by taking for $m < \alpha$, $\boldsymbol{a}_{s_m}$ as $\mathbf{d}^0_m$, $\mathbf{b}_{s_m}$ as $\mathbf{b}^0_m$, and $\mathbf{c}_{s_m}$ as $\mathbf{c}^0_m$. $\square_{21.18}$

Now, we define by induction on $i < \alpha$, $\boldsymbol{a}'_{s_i}, \mathbf{c}'_{s_i} \in M_{s_i \wedge s_m}$ so that for $k \leq \alpha$,

$$\langle \boldsymbol{a}_{s_i}\mathbf{c}_{s_i} : i < \alpha\rangle \equiv \langle \boldsymbol{a}'_{s_i}\mathbf{c}'_{s_i} : i < \alpha\rangle$$

and $\boldsymbol{a}_s = \boldsymbol{a}'_s$, $\mathbf{c}_s = \mathbf{c}'_s$ if $s \in I$. Moreover, we will require

$$\mathbf{b}_s \quad \subseteq \quad \bigcup_{t \subset s} \boldsymbol{a}'_t. \tag{21.1}$$

Suppose we have succeeded for $k < i$.

Let

$$\mathbf{d} = \langle \boldsymbol{a}_t : t < s_i\rangle, \quad \mathbf{e} = \langle \boldsymbol{a}_{s_j}, \mathbf{c}_{s_j} : j < i\rangle$$

and

$$\mathbf{d}' = \langle \boldsymbol{a}'_t : t < s_i\rangle, \quad \mathbf{e}' = \langle \boldsymbol{a}'_{s_j}, \mathbf{c}'_{s_j} : j < i\rangle.$$

Case 1. $s_i \in I$. We will show that $\mathbf{e} \equiv_{\mathbf{c}_{s_i} a_{s_i}} \mathbf{e}'$. Since

$$\mathbf{b}_{s_i} \subset \bigcup_{t \subset s_i} \boldsymbol{a}_t = \bigcup_{t \subset s_i} \boldsymbol{a}'_t \subset \bigcup_{j<i} \boldsymbol{a}_{s_j},$$

$\mathbf{e} \equiv_{\mathbf{b}_{s_i}} \mathbf{e}'$. The last containment holds since $t \subset s_i$ if and only if $t = s_j$ for some $j < i$.

By Lemma 21.14.1, $A_{s_i} \leq_{\mathrm{TV}} \bigcup_{j<i} M_{s_j}$; we chose $\mathbf{c}_{s_i}$ so that $\mathrm{tp}(\mathbf{c}_{s_i}/\mathbf{b}_{s_i}) \vdash \mathrm{tp}(\mathbf{c}_{s_i}/A_{s_i})$. Lemma 21.16.2 yields that the formula generating $\mathrm{tp}(\mathbf{c}_{s_i}/\mathbf{b}_{s_i})$ also generates $\mathrm{tp}(\mathbf{c}_{s_i}/\bigcup_{j<i} M_{s_j})$ and so $\mathbf{e} \equiv_{\mathbf{c}_{s_i}} \mathbf{e}'$.

Using again that $A_{s_i} \leq_{\mathrm{TV}} \bigcup_{j<i} M_{s_j}$, from Lemma 21.14.1, Exercise 21.8 gives

$$A_{s_i}\mathbf{c}_{s_i} \leq_{\mathrm{TV}} \bigcup_{j<i} M_{s_j}\mathbf{c}_{s_i}.$$

And $\mathrm{tp}(\boldsymbol{a}_{s_i}/A_{s_i}\mathbf{c}_{s_i})$ is based on $\mathbf{c}_{s_i}$ so, since $\mathbf{e} \equiv_{\mathbf{c}_{s_i}} \mathbf{e}'$, we have the final conclusion:

$$\mathbf{e} \equiv_{\mathbf{c}_{s_i}\boldsymbol{a}_{s_i}} \mathbf{e}'.$$

Case 2. $s_i \notin I$. Since $\mathbf{d}' \in M_{s_j \wedge s_m}$, we can choose $\mathbf{c}'_{s_i} \in M_{h_I(s_i)}$ with $\mathbf{c}'_{s_i}\mathbf{d}' \equiv \mathbf{c}_{s_i}\mathbf{d}'$. In particular, since $\mathbf{b}_{s_i} \subset \mathbf{d}'$ by eE quation 21.1, $\mathbf{c}'_{s_i}\mathbf{b}_{s_i} \equiv \mathbf{c}_{s_i}\mathbf{b}_{s_i}$. As $\mathrm{tp}(\mathbf{c}_{s_i}/\mathbf{b}_{s_i}) \vdash \mathrm{tp}(\mathbf{c}_{s_i}/A_{s_i})$ and as $A_{s_i} \leq_{\mathrm{TV}} \bigcup_{j<i} M_{s_j}$, Lemma 21.16.2 yields that the formula generating $\mathrm{tp}(\mathbf{c}_{s_i}/\mathbf{b}_{s_i})$ also generates $\mathrm{tp}(\mathbf{c}_{s_i}/\bigcup_{j<i} M_{s_j})$. In particular,

$$\mathrm{tp}(\mathbf{c}_{s_i}/\mathbf{b}_{s_i}) \vdash \mathrm{tp}(\mathbf{c}_{s_i}/\mathbf{e}).$$

By induction $\mathbf{ed} \equiv \mathbf{e}'\mathbf{d}'$ so $\mathbf{c}_{s_i}\mathbf{e} \equiv \mathbf{c}'_{s_i}\mathbf{e}'$.

Choose $\mathbf{f}$ with $\mathbf{c}'_{s_i} \subseteq \mathbf{f} \subseteq M_{s_j \wedge s_m}$ and $\mathrm{tp}(\mathbf{e}'/M_{s_j \wedge s_m})$ based on $\mathbf{f}$. Conjugate $p = \mathrm{tp}(\boldsymbol{a}_{s_i}/\mathbf{c}_{s_i})$ to p' over $\mathbf{c}'_{s_i}$; p and thus p' are both stationary. (It is to get this stationarity that we have to introduce the $\mathbf{c}_s$.) Now choose $\boldsymbol{a}'_{s_i} \in M_{s_j \wedge s_m}$ to realize $p'|\mathbf{f}$ (the nonsplitting extension of p to $\mathbf{f}$).

By the choice of $\mathbf{f}$, $\mathbf{e}' \models \mathrm{tp}(\mathbf{e}'/\mathbf{f})|\mathbf{f}\boldsymbol{a}'_{s_i}$, so by the symmetry of stationary types (Lemma 20.23), $\boldsymbol{a}'_{s_i} \models p'|\mathbf{f}\mathbf{e}'$ and so $\boldsymbol{a}'_{s_i} \models p'|\mathbf{c}'_{s_i}\mathbf{e}'$. We chose $\mathbf{c}'_{s_i}$ so that $\boldsymbol{a}_{s_i} \models p|\mathbf{c}_{s_i}\mathbf{e}$ and we verified $\mathbf{c}'_{s_i}\mathbf{e}' \equiv \mathbf{c}_{s_i}\mathbf{e}$ so

$$\boldsymbol{a}'_{s_i}\mathbf{c}'_{s_i}\mathbf{e}' \equiv \boldsymbol{a}_{s_i}\mathbf{c}_{s_i}\mathbf{e}$$

as required. $\square_{21.17}$

We need one further technical lemma.

Lemma 21.19. *If $\overline{S} = \langle s_i : i < n\rangle$ and $\overline{T} = \langle t_i : i < n\rangle$ are two enumerations of the partial order $(P, \subset)$, then it is possible to pass from $\overline{S}$ to $\overline{T}$ by a succession of transpositions so that each intervening step is also a enumeration*

Proof. We prove the result by induction on n. Consider a P with cardinality $n+1$. So $s_n = t_j$ for some $j < n$. Now successively move t_j up in the linear order by a sequence of transpositions. Each of them is permissible because $s_n = t_j$ is $(P, \subset)$-maximal. We then have two enumerations with last element s_n and t_j. By induction we can transform the first n elements and we finish. $\square_{21.19}$

Lemma 21.20. *If P is a finite partial order and $\mathcal{S}$ is a P-system, such that $\mathcal{S}$ is good with respect to some enumeration then $\mathcal{S}$ is good with respect to any enumeration.*

Proof. It suffices by means of transpositions to show that if $\mathcal{S}$ is good with respect to $\langle s_i : i < \alpha\rangle$, then it is good with respect to $\langle s'_i : i < \alpha\rangle$ where $s_i = s'_i$ unless $i = j$ or $j+1$ while $s'_{j+1} = s_j$ and $s'_j = s_{j+1}$. (We know by Lemma 21.19 that we use transpositions which preserve the property of being an enumeration.) Then

$$\langle M_{s_{j+1} \wedge s_k} : k \leq j+1\rangle \leq_{\mathrm{TV}} \langle M_{s_k} : k \leq j+1\rangle$$

by Lemma 21.17 (noting that the conclusion of Lemma 21.17 is preserved under subsequence). But this exactly what we need. $\square_{21.20}$

The following result is called Generalized symmetry for two reasons: it extends symmetry for independence of pairs to systems and the proof in a rough sense is iterating that independence.

Theorem 21.21 (Generalized Symmetry)**.** *If P is a finite partial order and $\mathcal{S}$ is a P-system, the following are equivalent:*

1. *$\mathcal{S}$ is good.*
2. *$\mathcal{S}$ is good with respect to some enumeration.*
3. *$\mathcal{S}$ is good with respect to any enumeration.*

Proof. The equivalence of 1) and 3) is Lemma 21.15; 3 implies 2) is trivial and 2) implies 3) is Lemma 21.20. $\square_{21.21}$

Remark 21.22. This material comes primarily from [124, 125] and [49] but with considerable rearrangement and adding of details. I thank Alf Dolich for a very careful critique of this chapter that found many errors. In particular, he pointed out the need for Lemma 21.19; Alice Medvedev pointed towards the simple argument for Lemma 21.19.

22
Excellence goes up

In this chapter we introduce the formal notion of excellence for an atomic class. This is a collection of requirements on $(\aleph_0, n)$-systems. We show that if these condition hold for all $(\aleph_0, n)$ (the class is excellent) then in fact the analogous properties hold for all (λ, n). We prove these result here for 'full'-systems; the arguments in this chapter just add a little complication to proving the result without 'full'. However, fullness plays an essential role in the next chapter showing that categoricity in enough small cardinals implies excellence.

There are two aspects of this analysis. The first is rather empty. All the conditions are defined conditionally: if there is a (λ, n)-independent system then something happens. We show how to move these conditional properties between cardinals. We then show that if there is an uncountable model the conditional properties have teeth since there are (λ, n)-independent systems and conclude that excellent classes have arbitrarily large models.

Definition 22.1. *1. The model $M \in \boldsymbol{K}$ is λ-full over the proper subset A of M if for every finite $B \subset M$, all stationary $p \in S_{at}(B)$, and any $C \subset M$ with $|C| < \lambda$, $p|ABC$ is realized in M.*

2. M is λ-full if M is λ-full over the empty set.

3. M is full if it is $|M|$-full.

4. The (λ, P)-system $\mathcal{S}$ is full if for each $s \in P$, M_s is λ-full over A_s.

Observe that if M is full over some A with $|A| \geq \lambda$, then M is λ-full.

Definition 22.2. $\boldsymbol{K}$ *satisfies the* (λ, n)-completeness property *if for any full* $(\lambda, \mathcal{P}^{-}(n))$*-diagram* $\mathcal{S}$*, there is a model* M_n *which completes* $\mathcal{S}$ *to a full* $(\lambda, \mathcal{P}(n))$*-diagram. That is,* M_n *is* λ*-full over* $A_n = \bigcup_{t \subset n} M_t$.

Remark 22.3.

1. It is tempting to call this notion (λ, n)-amalgamation. But completion is a stronger requirement; even $(\mu, 2)$-completion is stronger than amalgamation in μ, because the imbedding into the amalgam is the identity on both models.
2. $(\lambda, 0)$-completeness just abbreviates 'there is a full model of cardinality λ'.
3. Similarly $(\lambda, 1)$-completeness is an abbreviation for 'every full model of cardinality λ has a *proper* extension that is full over it. We have explicitly required the extension to be proper (Definition 21.1). So the Marcus-Shelah example, by Remark 19.28 satisfies $(\aleph_0, 0)$ but not $(\aleph_0, 1)$-completeness. Shelah [125] does not make this requirement but that leads to some technical inaccuracies. Note that this means some assumption (namely the existence of a model in $\aleph_1$) is necessary to obtain $(\aleph_0, 1)$-completeness. See Lemma 22.11.
4. Our (λ, n)-completeness plays the role of Shelah's (λ, n)-existence. We called it completeness rather than existence because the property itself does not imply existence. That is, for $n \geq 2$ it is possible that $\boldsymbol{K}$ is (λ, n)-complete but has no (λ, n)-independent system. We show that this cannot happen in the presence of $(\aleph_0, \leq n)$-goodness in Lemma 22.11. We removed Shelah's requirement that the system being completed is good.This emphasizes that goodness is guaranteed by induction, when we actually apply completeness.

Now we define the most important notion of this part of the book. In discussing independent systems, we introduced the notations $A_n = \bigcup_{s \subset n} M_s$; we now allow the independent system to vary and so naturally write $A_n^{\mathcal{S}}$ for a particular $\mathcal{S}$.

Definition 22.4.

1. $\boldsymbol{K}$ *is* (λ, n)-good *if for any full independent* $(\lambda, \mathcal{P}^{-}(n))$*-system* $\mathcal{S}$*,* $A_n^{\mathcal{S}}$ *is good.*
2. $\boldsymbol{K}$ *is* excellent *if it is* $(\aleph_0, n)$*-good for every* $n < \omega$.

Note that $(\aleph_0, 1)$-goodness implies ω-stability so we effectively have ω-stability in this chapter.

For countable models, this represents the idea described in the introduction to Part IV. If $A_n^{\mathcal{S}}$ is good, Lemma 19.26 yields a primary model over $A_n^{\mathcal{S}}$ and this is the amalgam which is tightly connected to the original diagram. We will see much later in Lemma 25.8 that the same conclusion follows for uncountable systems. The crucial tool for analyzing an uncountable model M of size λ is to construct a filtration $\langle M_i : i < \lambda \rangle$; then to construct a filtration of each M_i as $\langle M_{i,j} : j < |M_i| \rangle$ and to iterate. But one must guarantee independence of these filtrations. We do this by describing filtrations of independent systems.

Definition 22.5. *Suppose $\mathcal{S} = \langle M_s : s \in \mathcal{P}^-(n)\rangle$ is a (λ, n)-system. A* filtration *of $\mathcal{S}$ is a system $\mathcal{S}^\alpha = \langle M_s^\alpha : s \in \mathcal{P}^-(n), \alpha < \delta\rangle$ such that:*

1. *each $|M_s^\alpha| = \alpha^* = |\alpha| + \aleph_0$;*
2. *for each s in $\mathcal{P}^-(n)$, $\{M_s^\alpha : \alpha < \lambda\}$ is a filtration of M_s;*
3. *for each α, $\mathcal{S}^\alpha$ is an (α^*, n)-system.*

The filtration is full *or* independent *if each $\mathcal{S}^\alpha$ is a full good system.*

Shelah introduced an important tool for the analysis of models of size λ: work in a model of set theory of size λ, expanded by the vocabulary of our AEC (and the extra functions defined in Lemma 21.9). In applying the Lowenheim-Skolem theorem in this expanded language, we take advantage of the fact that our entire model theoretic analysis takes place in set theory.

Lemma 22.6. *Let λ be infinite and $n < \omega$. Suppose $\boldsymbol{K}$ is $(\aleph_0, n)$-good. Then, $\boldsymbol{K}$ is (λ, n)-good for every n.*

Proof. Suppose $\mathcal{S} = \langle M_s : s \in \mathcal{P}^-(n)\rangle$ is a (λ, n)-full independent system. Let $A_n = \bigcup_{s \subset n} M_s$. Fix a model V of a sufficiently large fragment of set theory with $\mathcal{S} \in V$ and $|V| = \lambda$. Let V' be an expansion of (V, ϵ) by naming the $\langle M_s : s \in \mathcal{P}^-(n)\rangle$ and add the predicates of L^* from Lemma 21.9 and from Lemma 21.13 to obtain a language L'. Then build a filtration $\langle V'_\alpha : \alpha < \lambda\rangle$ of V' by L'-elementary submodels with $|V'_\alpha| = |\alpha| + \aleph_0$. To get the filtration of $\mathcal{S}$, let $\mathcal{S}^\alpha$ be the interpretation of the M_s in V'_α. Since independence and $|M|$-fullness are expressible in set theory, each $\mathcal{S}^\alpha$ is $|V_\alpha|$-full and independent. Then A_n^0 is good since $\boldsymbol{K}$ is $(\aleph_0, n)$-good. And each A_n^α is then good by Lemma 21.9 and the choice of L^*. Moreover each $\mathcal{S}^\alpha \leq_{\mathrm{TV}} \mathcal{S}^\beta$ if $\alpha \leq \beta$ by Lemma 21.13. $\square_{22.6}$

We have attained the first aim of the chapter; we have moved from $(\aleph_0, n)$-good to (λ, n)-good. But our goal of constructing models of larger cardinality still needs work. Neither the hypothesis nor the conclusion of Lemma 22.6 actually asserts the existence of (λ, n)-independent systems; all is conditional. We first require one technical result. We show how to combine systems indexed by two copies of $\mathcal{P}(n)$ by using the lexicographical order on $\mathcal{P}(n) \times \{0, 1\}$.

Corollary 22.7. *Suppose $\boldsymbol{K}$ is $(\lambda, \leq n)$-good. If $\mathcal{S}_1 = \langle M_s^1 : s \subseteq n\}$ and $\mathcal{S}_2 = \langle M_s^2 : s \subseteq n\}$ are full-independent systems, $\mathcal{S}_1$ is good, and $\mathcal{S}^1 \leq_{\mathrm{TV}} \mathcal{S}^2$ then $\mathcal{S} = \mathcal{S}^1 \cup \mathcal{S}^2$ is a good full-independent system indexed by $P = \mathcal{P}(n) \times \{0, 1\}$.*

Proof. Enumerate $\mathcal{S}$ so that $\mathcal{S}_1$ is an initial segment of the enumeration. To show that $\mathcal{S}$ is good for this enumeration (Definition 21.12) and therefore good and independent by Theorem 21.21, we need only show that A_{s_i} is good for each i and that

$$\langle M_{s_i \wedge s_j} : j \leq i\rangle \leq_{\mathrm{TV}} \langle M_{s_j} : j \leq i\rangle. \tag{22.1}$$

For $i < 2^n$ the first requirement is immediate since $\mathcal{S}^1$ is good. To see A_{s_i} is good if s_i has the form $(x, 1)$ (i.e. $i > 2^n$), note that each such A_{s_i} is the union of a family indexed by a subset of $\mathcal{P}^-(n)$. So the goodness of A_{s_i} follows from the hypothesis that $\boldsymbol{K}$ is (λ, n)-good. Now 22.1 is immediate from $\mathcal{S}^1 \leq_{\mathrm{TV}} \mathcal{S}^2$, (interpolating a sequence from $\mathcal{S}^1$ between the two sides of 22.1). $\square_{22.7}$

Using this lemma, we are able to transfer completeness to larger cardinalities. But crucially the number of models that we can amalgamate drops.

Lemma 22.8. *Let λ be infinite and regular and $n < \omega$. Suppose $\boldsymbol{K}$ is $(< \lambda, \leq n+1)$-complete and $(\aleph_0, n)$-good. Then $\boldsymbol{K}$ is (λ, n)-complete.*

Proof. Note that $(< \lambda, \leq 1)$-completeness implies the existence of a model M in λ. Suppose $\mathcal{S}$ is a (λ, n)-full independent system. Choose a full-filtration $\mathcal{S}^\alpha$ (with respect to L^*) as in Lemma 22.6. We can further choose N^α for $\alpha < \lambda$ such that:

1. $|N^\alpha| = |\alpha| + \aleph_0$.
2. $N^\alpha \underset{A_n^\alpha}{\perp} A_n$
3. N^α is full over A_n^α.
4. Every stationary type over a finite subset C of N^α has a realization in $N^{\alpha+1}$ whose type over $N^\alpha A_n^{\alpha+1}$ does not split over C.

For the initial step we are given N^0 satisfying 1, 3, and 4) by $(\aleph_0, n)$-completeness and Condition 2) is obtained by the extension property for non-splitting (Corollary 20.14). For the induction step, note that using the $(\aleph_0, n)$-goodness, Lemma 22.6, and Lemma 22.7, the system $\langle M_s^\alpha, M_s^{\alpha+1}, N_\alpha : s \subseteq n\rangle$ is an $(|\alpha| + \aleph_0, \mathcal{P}^-(n+1))$ independent system. So we can find $N^{\alpha+1}$ satisfying 1,3, and 4) by $(|\alpha| + \aleph_0, n+1)$-completeness and Condition 2) is again obtained by the extension property for non-splitting. Take unions at limits. Then $N = \bigcup_{\alpha<\lambda} N_\alpha$ is the required completion of $\mathcal{S}$. Conditions 2 and 4 and the regularity of λ guarantee that N is full over A_n. $\square_{22.8}$

We use the regularity of λ in Lemma 22.9 only to show the limit model is full. Knowing the result for regular λ it easily extends to singular λ by applying Lemma 23.3, which could as easily have been proved at this point. (Note that if we drop full and do this chapter for arbitrary models, this difficulty with unions completely disappears.)

Corollary 22.9. *Let λ be infinite and $n < \omega$. Suppose $\boldsymbol{K}$ is $(< \lambda, \leq n+1)$-complete and $(\aleph_0, n)$-good. Then $\boldsymbol{K}$ is (λ, n)-complete.*

Now we want to prove that there actually are (λ, n)-independent systems. Note that the first sentence of the proof of Lemma 22.8 does yield the existence of a $(\lambda, 0)$-independent system. $(\aleph_0, 0)$ is not hard. However, if the A in Lemma 22.10 is a model, the constructed M may be A; so $(\aleph_0, 1)$-completion may fail. The subtlety is that we have required the models in P-systems to be *proper* extensions. The only requirements we must add for larger n are the existence of a model in $\aleph_1$ and enough goodness.

Lemma 22.10. *Let $\boldsymbol{K}$ be ω-stable. If A is a countable atomic set, there is an M with $A \subseteq M$ that is full over A.*

Proof. Note first that any countable atomic set can be embedded in a model (Lemma 19.6). Now construct M_i for $i < \omega$ so that every $p \in S_{\mathrm{at}}(M_i)$ is realized in M_{i+1} and let $M = \bigcup_i M_i$. Then if $p \in S(\mathbf{b})$ is stationary and $\mathbf{b}$ is finite, $A\mathbf{b} \subset M_i$ for some i and the nonsplitting extension of p to $A\mathbf{b}$ is realized in M_{i+1}. $\square_{22.10}$

We can't demand that M is proper over A, as Example 19.28 provides an $\aleph_0$-full model with *no* proper extension. With goodness and the generalized symmetry lemma in hand we get more.

Lemma 22.11. *If $\boldsymbol{K}$ has a $(\lambda, 1)$ independent system and $\boldsymbol{K}$ satisfies (λ, m)-goodness for $m < n$. then $\boldsymbol{K}$ satisfies (λ, n)-completeness and there is a (λ, n)-independent system.*

Proof. We prove the claim by induction on $n \geq 2$. Suppose there is an independent $(\lambda, \mathcal{P}(n-1))$ system $\mathcal{S} = \langle M_{s_i} : i < 2^{n-1}\rangle$. For $2^{n-1} < i < 2^n$, choose full M_{s_j} so that if $j = 2^{n-1} + i$, $s_j = s_i \cup \{n-1\}$ and $M_{s_j} \downarrow_{A_{s_j}} \bigcup_{t<j} M_{s_t}$ to create the sequence $\mathcal{S}'$. For $t < 2^n$, A_{s_t} is the union of a (λ, m) system for some $m < n$ so by the hypothesis of (λ, m) goodness, it is good. Now inductively applying Lemma 21.17 each initial segment of $\mathcal{S}'$ is a good sequence and so $\mathcal{S}'$ is a good sequence with respect to the given enumeration. So by Lemma 21.21, $\mathcal{S}'$ is a good $(\lambda, \mathcal{P}^-(n))$-independent system. $\square_{22.11}$

Note that for $(\aleph_0, 2)$-completeness, the $(\aleph_0, 1)$-goodness hypothesis is free.

We do *not* know that $A_n^{\mathcal{S}'} = \bigcup_{j<2^n} M_{s_j}$ is good. We will derive that from assuming few models in $\aleph_{n+2}$ in the next two chapters. Even worse, in general there may be (uncountable) good sets over which there is no prime model [96, 86, 91]. The requirement that there be an uncountable model is essential since the Marcus example (Example 4.14, 19.28) is ω-stable and has no $(\aleph_0, k)$-independent systems for $k \geq 1$.

To construct models of arbitrarily large cardinality, we first move excellence up to $\aleph_\omega$.

Theorem 22.12. *If $\boldsymbol{K}$ is $(\aleph_0, \leq n+1)$-good and $\boldsymbol{K}$ has an uncountable model, then there exists a full model in cardinality $\aleph_n$.*

Proof. We will prove by induction on m that if $m + k \leq n + 1$, there is an $(\aleph_m, k)$-independent system and $\boldsymbol{K}$ is $(\aleph_m, k)$-complete.

First consider $m = 0$. By $\aleph_0$-categoricity, the countable model M_0 is full. Since $\boldsymbol{K}$ has an uncountable model M_0 has a proper extension M_1. Apply Lemma 22.10 to get a full extension of M_1 (and so of M_0); so $\boldsymbol{K}$ is $(\aleph_0, 1)$-complete. By

Lemma 22.11, for every $k \leq n+1$, we have an $(\aleph_0, k)$-independent system. Since any countable atomic set is contained in a countable atomic model, for each $k \leq n$, $(\aleph_0, k)$-completeness follows by extending an $(\aleph_0, k)$-independent system to a model then applying Lemma 22.10 again.

Now to move from m to $m+1$. As long as $m \leq n-1$, by induction we have $(\aleph_m, 2)$-independent systems and $(\aleph_m, 2)$-completeness holds. By Lemma 22.8, $\boldsymbol{K}$ is $(\aleph_{m+1}, 1)$-complete. Lemma 22.11, the goodness assumption and Lemma 22.6 give us $(\aleph_{m+1}, k)$-independent systems for $k \leq n+1$. And Lemma 22.8 allows us to complete them if $m+1+k \leq n+1$. When $m = n$ and $k = 0$, we have the theorem. $\square_{22.12}$

Now, an easy cardinal induction yields that excellence (in $\aleph_0$) and an uncountable model implies (λ, n)-completeness for all λ and all n.

Corollary 22.13. *If $\boldsymbol{K}$ is excellent and has an uncountable model then $\boldsymbol{K}$ satisfies (λ, n)-completeness for every λ and so has full models in every cardinality.*

Proof. For the sake of contradiction, let λ be least such that $\boldsymbol{K}$ has $(<\lambda, <\omega)$-completeness but not $(\lambda, <\omega)$-completeness. By Lemma 22.11, $\lambda \geq \aleph_1$ and by Lemma 22.8, we see in fact that $\boldsymbol{K}$ has $(\lambda, <\omega)$-completeness. $\square_{22.13}$

23
Very few models implies Excellence

Our goal is:

Theorem 23.1 (Shelah)**.** *(For $n < \omega$, $2^{\aleph_n} < 2^{\aleph_{n+1}}$.) An atomic class $\boldsymbol{K}$ that has at least one uncountable model and that has very few models in $\aleph_n$ for each $n < \omega$ is excellent.*

Usually, one says that $\boldsymbol{K}$ has few models in λ if the number of models is less than 2^λ. However, exactly what can be proved in ZFC + $2^\lambda < 2^{\lambda^+}$, depends on exactly what we mean by 'few'. We define for $n < \omega$ what it means to have 'very few models in $\aleph_n$'; this will be actual hypothesis for the main results.

Definition 23.2. *We say*

1. *$\boldsymbol{K}$ has* few *models in power λ if $I(\boldsymbol{K}, \lambda) < 2^\lambda$.*
2. *$\boldsymbol{K}$ has* very few *models in power $\aleph_n$ if $I(\boldsymbol{K}, \aleph_n) \leq 2^{\aleph_{n-1}}$.*

So Theorem 23.1 is weaker than one might hope. The set theoretic hypothesis and various putative strengthenings (e.g. replacing very few by few) are discussed in Appendix C and in Chapter 24. In this chapter we reduce the theorem to one combinatorial/model theoretic result (Theorem 23.15), which asserts that under appropriate hypotheses, very few models in $\aleph_n$ implies $(\aleph_{n-2}, 2)$-systems are amalgamation bases. We defer the proof of Theorem 23.15 to Chapter 24.

The next lemma and its corollary provide a crucial amalgamation hypothesis.

Lemma 23.3. *Let δ be an infinite ordinal. If $\langle M_i : i < \delta\rangle$ is a continuous increasing chain of λ-full models with M_{i+1} λ-full over M_i then $M = \bigcup_i M_i$ is λ-full.*

Proof. Without loss of generality δ is a regular cardinal. The result is easy except when $\delta = \mathrm{cf}(\lambda) < \lambda$, so we deal with that case. Fix p stationary over a finite set B and let $C \subset M$ with $|C| < \lambda$. Again without loss of generality, $B \subset M_0$. Moreover, we may assume each finite sequence $\mathbf{c} \in C$ realizes a stationary type based in M_0. Let $C_j = C - M_j$. Define D_j so that $M_j \supset D_j \supseteq C \cap M_j$, $|D_j| < \lambda$ and $C_j \downarrow_{D_j} M_j$.

Now, using M_{j+1} is λ-full over M_j, construct a set X such that for each ordinal $j < \delta$, $X \cap (M_{j+1} - M_j)$ is a set of λ realizations of $p \restriction M_j D_{j+1}$. We claim any element $x \in X - C$ realizes $p \restriction C$. For some i, $x \in M_{i+1}$. Let $\mathbf{c}$ be any finite sequence from C. Then $\mathbf{c} = \mathbf{c}_1\mathbf{c}_2$ where $\mathbf{c}_1 \in M_{i+1}$ and $\mathbf{c}_2 \in M - M_{i+1}$. By the choice of X, we have $x \downarrow_B M_i D_{i+1}$. By the choice of D_{i+1}, $\mathbf{c}_2 \downarrow_{D_{j+1}} M_i D_{i+1} x$, whence $x \downarrow_{M_i D_{i+1}} c_2$. Now by transitivity we get $\mathbf{c}_1\mathbf{c}_2 \downarrow_B x$ as required. $\square_{23.3}$

There is an immediate intriguing corollary to Lemma 23.3. It establishes one of the hypotheses in the main induction Lemma 23.14. The existence and uniqueness of full models at the appropriate λ will follow from the global induction.

Corollary 23.4. *Let $(\boldsymbol{K}_\lambda, \prec_{\boldsymbol{K}})$ be the class of λ-full models of the atomic class $\boldsymbol{K}$ under first order elementary submodel.*

1. *For each λ, $\boldsymbol{K}_\lambda$ satisfies all axioms (Definition 5.1) for an AEC with Löwenheim number λ except* **A.3.3** *and coherence* **A4**.
2. *Suppose there is a unique λ-full model in $\boldsymbol{K}$ of cardinality λ and less than 2^{λ^+} of cardinality λ^+, then $\boldsymbol{K}_\lambda$ has the amalgamation property in λ.*

Proof. Item 1) is Lemma 23.3; item 2) is Theorem 18.12, noting that the two missing axioms are not used in the proof. $\square_{23.4}$

The crucial reason that we study full models is because they have uniqueness properties that do not hold for arbitrary models. Even in an $\aleph_0$-categorical situation (all models are isomorphic), we are looking for $\aleph_0$-categoricity of *systems* which does not hold without some condition like ‘full over’.

Definition 23.5. $\boldsymbol{K}$ *has* (λ, n)-uniqueness *if any two models which are full over a good full (λ, n)-independent system $\mathcal{S}$ are isomorphic over $A_n^{\mathcal{S}}$.*

The following interpretation of $(\lambda, 2)$-uniqueness will be used repeatedly. We have the existence of an isomorphism fixing M_2 mapping M_1 isomorphically onto M_1' from Corollary 20.14. But ([6]), this isomorphism may not extend to an automorphism (even when $\boldsymbol{K}$ is $\aleph_1$-categorical). Note, however, the following easy fact.

Lemma 23.6. *If $\boldsymbol{K}$ satisfies $(\aleph_0, 2)$-uniqueness for any independent good triple (M_0, M_1, M_2), of countable full models we have: if $M_1' \approx_{M_0} M_1$ and $M_1' \underset{M_0}{\downarrow} M_2$ then there exists $\alpha \in \mathrm{aut}_{M_2}\mathbb{M}$ mapping M_1 isomorphically onto M_1'.*

Proof. Apply Lemma 22.11 to get N a full extension of M_1M_2 and N' a full extension of $M_1'M_2$. By Corollary 20.14, there is an isomorphism f from M_1M_2

onto $M_1'M_2$ fixing M_2. Extending f^{-1} naturally gives an isomorphism from N' to a second full extension N'' of M_1M_2. By $(\aleph_0, 2)$-uniqueness there is an isomorphism g from N onto N'' that fixes M_1M_2. So fg is an isomorphism from N onto N'. Now by model homogeneity there is an extension of fg to an automorphism of $\mathbb{M}$. $\square_{23.6}$

We could extend this result to $(\lambda, 2)$-uniqueness provided we establish $(\lambda, 2)$-completeness. We want to show $(\aleph_0, n)$ uniqueness for all n. For $n = 1$, a somewhat stronger result holds.

Lemma 23.7. *Suppose that A is countable and good. If M and N are countable and full over A, then they are isomorphic over A.*

Proof. Write M as $A \cup \{a_i : i < \omega\}$ and N as $A \cup \{b_i : i < \omega\}$. We construct an increasing sequence of partial elementary maps f_i such that $\mathrm{dom}(f_i) = A_i$ and $\mathrm{rg}(f_i) = B_i$, where A_i, B_i add finitely many elements to A and such that $a_i \in A_{2i}$ and $b_i \in B_{2i+1}$.

We first construct f_0 and f_1. Since A is good, there is model M' that is primary over A. Without loss of generality $M' \prec_{\boldsymbol{K}} M$. Let f map M' into N over A. The stationary type $p = \mathrm{tp}(a_0/M')$ is the unique non-splitting extension of the stationary type $\mathrm{tp}(a_0/A\mathbf{c})$ for some $\mathbf{c} \in M'$ (by Lemma 21.11.1). By the fullness of N, $f(p)$ is realized by some $b \in N$ and $\mathrm{tp}(b/A) = \mathrm{tp}(a_0/A)$. Let $f_0 = 1_A \cup \langle a_0, b \rangle$. By Lemma 19.24, Ab is good so there is a primary model M'' over Ab. So there is g mapping M'' into M and extending the inverse of f_0. As before, the stationary type $\mathrm{tp}(b_0/M'')$ is the unique nonsplitting extension of some $q \in S_{at}(Ab\mathbf{d})$ with $\mathbf{d} \in M''$. So we set $f_1 = f_0 \cup \langle c, b_0 \rangle$ for some c realizing $g(q)$.

The induction step is the same argument noting that the domain or range of each f_i extends A by only finitely many points. $\square_{23.7}$

This yields a stronger version of uniqueness for countable full systems. We say two systems $\mathcal{S}, \mathcal{S}'$ are isomorphic if there is an isomorphism from $A_n^{\mathcal{S}}$ to $A_n^{\mathcal{S}'}$ that maps $M_s^{\mathcal{S}}$ to $M_s^{\mathcal{S}'}$ for each s.

Corollary 23.8. *Suppose $\boldsymbol{K}$ is $(\aleph_0, n-1)$-good.*

1. *$\boldsymbol{K}$ satisfies $(\aleph_0, n-1)$-uniqueness.*
2. *Moreover, any two countable full $(\aleph_0, \mathcal{P}^-(n))$-independent systems are isomorphic.*

Proof. The first claim is immediate from Lemma 23.7; the second follows by induction using that each M_s is full over A_s. Namely, let $\langle M_s : s \subset n \rangle$ and $\langle N_s : s \subset n \rangle$ be two full $(\lambda, \mathcal{P}^-(n))$-systems. By induction there is an f_n mapping the system $\langle M_s : s \subset n-1 \rangle$ with union A^M isomorphically onto $\langle N_s : s \subset n-1 \rangle$ with union A^N. Now $M_{\{n\}} \downarrow_{M_0} M_{\{0,1,\ldots n-1\}}$ and $N_{\{n\}} \downarrow_{N_0} N_{\{0,1,\ldots n-1\}}$ so f_n extends to $\hat{f}_n$ mapping $M_{\{n\}} A^M M_{\{0,1,\ldots n-1\}}$ onto $N_{\{n\}} A^N N_{\{0,1,\ldots n-1\}}$. Now by $(\aleph_0, n-1)$-uniqueness this map further extends by induction to each $M_{s \cup \{n\}}$ for $s \subset n-1$ and $|s| = n-2$. $\square_{23.8}$

We now want to show Theorem 23.1. We need three more technical lemmas. We inductively show (λ, n)-uniqueness.

Lemma 23.9. *Let λ be infinite and $n < \omega$. Suppose $\boldsymbol{K}$ is $(\aleph_0, n)$-good and satisfies $(< \lambda, \leq n+1)$-uniqueness; then $\boldsymbol{K}$ satisfies (λ, n)-uniqueness.*

Proof. Suppose $\mathcal{S} = \langle M_s : s \in \mathcal{P}^-(n)\rangle$ is a good full-(λ, n)-independent system with two completions $\mathcal{S}_1, \mathcal{S}_2$ obtained by adding $M = M_n$ and $N = N_n$ that are each full over $A_n = \bigcup_{s \subset n} M_s$. We construct a filtration $\mathcal{S}^\alpha$ of $\mathcal{S}$ and M^α, N^α completing it to $\mathcal{S}_1^\alpha$ and $\mathcal{S}_2^\alpha$ for $\alpha < \lambda$ such that;

1. $\mathcal{S}_1^\alpha = \langle M_s^\alpha : s \in \mathcal{P}^-(n)\rangle \cup \{M^\alpha\}$ ($\mathcal{S}_2^\alpha = \langle M_s^\alpha : s \in \mathcal{P}^-(n)\rangle \cup \{N^\alpha\}$) is a good $(|\alpha| + \aleph_0|, \mathcal{P}(n))$-full-independent system.
2. For $s \subseteq n$, M_s^α is a full-filtration of M_s and $|M_s^\alpha| = |\alpha| + \aleph_0$.
3. M^α and N^α are filtrations of M and N, respectively.
4. Each of M^α, N^α is full over A_n^α and
   (a) $M^\alpha \underset{A_n^\alpha}{\downarrow} A_n$
   (b) $N^\alpha \underset{A_n^\alpha}{\downarrow} A_n$
5. For $i = 1, 2$, $\mathcal{S}_i^\alpha \cup \mathcal{S}_i^{\alpha+1}$ is a good independent system indexed by $P = \mathcal{P}(n) \times \{1, 2\}$.

Conditions 1) and 2) are routine: just choose enough witnesses. Condition 3) is the extension axiom for splitting. Condition 4) is from Lemma 22.7.

Now by $(< \lambda, n)$ uniqueness there is an isomorphism f_0 from M^0 to N^0 over A_n^0. Suppose f_α is an isomorphism from M^α to N^α over A_n^α. Then, $\mathcal{S}^\alpha \cup \mathcal{S}^{\alpha+1} \cup \{M_\alpha\}$ and $\mathcal{S}^\alpha \cup \mathcal{S}^{\alpha+1} \cup \{N_\alpha\}$ are $\mathcal{P}^-(n+1)$-full-independent systems which extend to a $\mathcal{P}(n+1)$-full-independent system by $M^{\alpha+1}$ and $N^{\alpha+1}$. By $(< \lambda, n+1)$-uniqueness, f_α extends to an isomorphism $f_{\alpha+1}$ of $M^{\alpha+1}$ of $N^{\alpha+1}$. Take unions at limits. $\square_{23.9}$

We introduce some new terminology here; the concept of a P-system $\mathcal{S}$ being an amalgamation base. We will see that this is equivalent to goodness of $\mathcal{S}$ but our combinatorial methods are able to deal directly with finding amalgamation bases.

Definition 23.10. *Let $\mathcal{S} = \langle M_s : s \subset n\rangle$ be an independent (λ, n)-system. We say $A_n^{\mathcal{S}}$ or $\mathcal{S}$ is an* amalgamation base *if for every M_1, M_2 and embeddings f_1, f_2 of $A_n^{\mathcal{S}} = \bigcup_{s \subset n} M_s$ into M_1, M_2, there exist g_1, g_2 and N with $g_1 : M_1 \to N$, $g_2 : M_2 \to N$, and $g_1 f_1 \upharpoonright A_n^{\mathcal{S}} = g_2 f_2 \upharpoonright A_n^{\mathcal{S}}$.*

We establish a crucial link between goodness and amalgamation.

Corollary 23.11 ($2^{\aleph_0} < 2^{\aleph_1}$)**.** *Suppose $\boldsymbol{K}$ has less than $2^{\aleph_1}$ models of power $\aleph_1$. Let $\mathcal{S}$ be an $(\aleph_0, n)$-independent system. $\mathcal{S}$ is an amalgamation base if and only if $A_n^{\mathcal{S}}$ is good.*

Proof. If $A_n = A_n^{\mathcal{S}}$ is good, there is a primary model M' over A_n that can be embedded into M_1, M_2. Then amalgamating M_1 and M_2 over M' we see that $\mathcal{S}$ is an amalgamation base. For the converse, we adapt the argument for ω-stability of Lemma 19.16. $S^*(A_n) = S_{at}(A_n)$, because A_n is countable. As in the proof of Fact 19.15, since $\mathcal{S}$ is an amalgamation base, we can iteratively realize each element of $S^*(A_n)$, so if $S^*(A_n)$ is uncountable there is a single model realizing uncountably many elements of $S^*(A_n)$. Thus, $S^*(A_n)$ is countable by Fact 19.15 (applying $2^{\aleph_0} < 2^{\aleph_1}$ to name the elements of A_n). We finish the converse using Lemma 19.24. $\square_{23.11}$

We are proving by induction on n that if there are very few full models in each power up to $\aleph_n$ then $\boldsymbol{K}$ is $(\aleph_0, n)$-good. This induction must interweave goodness, uniqueness, and completeness. We now have the following information.

Remark 23.12. *Suppose $\boldsymbol{K}$ is $(\aleph_0, \leq n-1)$-good.*

1. *In $\aleph_0$:*
   (a) *It has $(\aleph_0, \leq n-1)$-uniqueness by Lemma 23.8.*
   (b) *It has $(\aleph_0, n)$-completeness by Lemma 22.11.*
2. *completeness above $\aleph_0$: By Lemma 22.8 it is $(\aleph_{n-k}, \leq k)$-complete, i.e.*
   (a) *$(\leq \aleph_{n-2}, \leq 2)$-complete,*
   (b) *$(\leq \aleph_{n-1}, \leq 1)$-complete, and*
   (c) *$(\leq \aleph_n, 0)$-complete.*
3. *uniqueness above $\aleph_0$: By Lemma 23.9, it is $(\aleph_{n-k}, \leq k-1)$-unique, i.e.*
   (a) *$(\leq \aleph_{n-2}, \leq 1)$-unique,*
   (b) *and $(\leq \aleph_{n-1}, 0)$-unique*

The crucial point is that we have uniqueness at one lower level than completeness. We now state two important lemmas for finding amalgamation bases. Using these lemmas we complete the proof of excellence. We then prove the first of the lemmas; proving the other comprises Chapter 24.

Lemma 23.13. *Let μ be a regular cardinal with $2^\mu < 2^{\mu^+}$. Suppose further that $(\mu, \leq n+1)$-completeness and $(\mu, \leq n)$-uniqueness hold. If*

1. *no full $(\mu, n+1)$-system $\mathcal{S}$ is an amalgamation base,*
2. *$\boldsymbol{K}$ satisfies $(\mu^+, 1)$-uniqueness,*

then no full (μ^+, n)-system $\mathcal{S}$ is an amalgamation base.

We now sketch the idea of the global induction; our initial outline is simpler, but defective. So we note the additional step needed. Suppose we have very few models in κ for $\kappa \leq \aleph_n$ and that from this we have deduced $\boldsymbol{K}$ is $(\aleph_0, \leq n-1)$-good. By Lemma 23.9 we have categoricity of full models in $\aleph_{n-1}$ ($(\aleph_{n-1}, 0)$-uniqueness) and by and Theorem 18.12 we have amalgamation in $\aleph_{n-1}$. In Remark 23.12 we have compiled the conclusions we can draw below $\aleph_n$ from $(\aleph_0, n-1)$-goodness. This is the information available to us at the induction step. For $k < n$, we have $(\aleph_{n-k}, \leq k)$-completeness and $(\aleph_{n-k}, \leq k-1)$-uniqueness. Suppose, by some miracle, we knew a bit more: that $\boldsymbol{K}$ was $(\leq \aleph_{n-1}, 1)$-unique. Then repeated application of Lemma 23.13 would yield every $(\aleph_0, n)$-system is an amalgamation base and we could finish by Lemma 23.11. But the results summarized in Remark 23.12 do not include $(\aleph_{n-1}, 1)$-uniqueness. By a much more technical argument, Theorem 23.15 tells us every $(\aleph_{n-2}, 2)$ system is an amalgamation base. This theorem uses the hypothesis of *very* few rather than just few models. Then we can apply the argument we just outlined, starting with $n-2$ instead of $n-1$; we do so now.

Lemma 23.14. *[Induction Step] Let $n \geq 2$. Suppose $\boldsymbol{K}$ is $(\aleph_0, n-1)$-good, has an uncountable model, and has very few models in $\aleph_n$. Then $\boldsymbol{K}$ is $(\aleph_0, n)$-good.*

Proof. First, $\boldsymbol{K}$ is $(\aleph_0, n-1)$-complete by Lemma 22.11. By Lemmas 23.9 and 22.8, Hypothesis 1 of Theorem 23.15 is true for $\lambda = \aleph_{n-2}$. Hypothesis 2) holds by Theorem 23.4. Since there are few models in $\aleph_n$, Theorem 23.15 implies there is an $(\aleph_{n-2}, 2)$-amalgamation base. As noted, Lemma 23.9 guarantees that we have $(\lambda, 1)$-uniqueness, so can apply Lemma 23.13 for $\lambda \leq \aleph_{n-2}$. Repeatedly applying the contrapositive of Lemma 23.13, we conclude that there is an $(\aleph_{n-(k+2)}, k+2)$ amalgamation base for each $k \leq n-2$. At $k = n-2$, we have an $(\aleph_0, n)$-full independent system that is an amalgamation base. Thus it is good by Lemma 23.11. But since all $(\aleph_0, n)$-full independent systems are isomorphic by Lemma 23.8 and $(\aleph_0, n-1)$-goodness, we have proved that they are all good. $\square_{23.14}$

Theorem 23.15. *(For $n < \omega$, $2^{\aleph_n} < 2^{\aleph_{n+1}}$.) Suppose $\boldsymbol{K}$ satisfies the following hypotheses.*

1. *completeness and uniqueness conditions:*
   - (a) *completeness:*
     - i. *$\boldsymbol{K}$ is $(\lambda, 0)$, $(\lambda, 1)$, $(\lambda, 2)$-complete;*
     - ii. *$\boldsymbol{K}$ is $(\lambda^+, 0)$, $(\lambda^+, 1)$-complete;*
   - (b) *uniqueness conditions:*
     - i. *$\boldsymbol{K}$ is $(\lambda, 0)$, $(\lambda, 1)$-unique;*
     - ii. *$\boldsymbol{K}$ is $(\lambda^+, 0)$-unique;*
2. *Amalgamation of full extensions in λ^+;*
3. *$\boldsymbol{K}$ has very few models in λ^{++};*

then some $(\lambda, 2)$-full system is an amalgamation base.

In Corollary 23.4.2, we showed condition 1b) and 3 implies Condition 2). We prove Theorem 23.15 in Chapter 24. To understand the following argument plug $n = 2$ into the proof of Lemma 23.14 and use Remark 23.12.

Proof of Main Theorem 23.1. $\boldsymbol{K}$ is $(\aleph_0, 1)$-good. The result follows by induction; Lemma 23.14 is the induction step. $\square_{23.1}$

It remains to prove Lemma 23.13. We want to show 'failure of amalgamation' in μ implies 'failure of amalgamation' in μ^+. In the first step we move from failure of certain amalgamations in μ to the non-existence of universal models in $\lambda = \mu^+$. The second step takes place at λ and moves from non-existence of universal models to failure of amalgamation.

Lemma 23.16. *Let μ be a regular cardinal with $2^\mu < 2^{\mu^+}$. Suppose further that $(\mu, \leq n+1)$-existence and $(\mu, \leq n)$-uniqueness hold. If no full $(\mu, n+1)$-system $\mathcal{S}$ is an amalgamation base, then if $\mathcal{S}$ is a full (μ^+, n)-system of cardinality μ^+, there is no universal model of cardinality μ^+ over $A = A_n^{\mathcal{S}}$.*

Proof. Let $\mathcal{S}$ be a $(\mu^+, \mathcal{P}^-(n))$-full independent system. Construct $\mathcal{S}^\alpha$ for $\alpha < \mu^+$ so that:

1. Each $\mathcal{S}^\alpha$ is a (μ, n)-full-independent system.
2. M_s^α is a filtration of M_s.

Without loss of generality we may assume the universe of $A^\alpha = \bigcup_s M_s^\alpha$ is the set of odd ordinals less than $\mu \times (1 + \alpha)$. Let $A = \bigcup_\alpha A^\alpha$. Now define by induction for each $\alpha < \mu^+$ and each $\eta \in 2^\alpha$, and each $s \in \mathcal{P}^-(n)$ models M^η such that:

1. For $\nu \in 2^{\mu^+}$, the sequence $\langle M_s^{\nu \restriction \alpha}; \alpha < \mu^+ \rangle$ is continuous and increasing.
2. For $\eta \in 2^\alpha$, and $\alpha = \beta + 1$.
   (a) M^η is full over A^α.
   (b) the universe of M^η is $\mu \times (1 + \alpha)$.
   (c) $M^\eta \underset{A^\alpha}{\downarrow} A$.
   (d) Every stationary type over a finite subset C of $M^{\eta \restriction \beta}$ is realized by an $\boldsymbol{a} \in M^\eta$ with $\boldsymbol{a} \underset{C}{\downarrow} M^{\eta \restriction \beta}$.
3. For $\beta < \mu^+$ there is no model M such that for $i = 0, 1$, there exist f_i which fix M_β and f_i maps $M^{\eta \frown i}$ into M.

The existence of M^η follows from $(\mu, n+1)$ completion; we satisfy condition 3) because no $(\mu, n+1)$-system is an amalgamation base.

We now conclude from the construction that there is no universal model over A. If N is universal over A then every M^η, for $\eta \in 2^{\mu^+}$ can be embedded by some f_η into N over M. By the Devlin-Shelah weak diamond (Appendix C), there are $\eta, \nu \in 2^{\mu^+}$ and $\alpha < \mu^+$ such that $\eta \restriction \alpha = \nu \restriction \alpha$, $\alpha = \mu \times \alpha$ and $f_\eta \restriction A^{\eta \restriction \alpha} = f_\nu \restriction A^{\nu \restriction \alpha}$. This contradicts condition 3. $\Box_{23.16}$

Remark 23.17. The next Lemma is a reformulation of 5.3 of [125]. A crucial point is that $(\lambda, 1)$-uniqueness is enough to show that if a (λ, n)-system is an amalgamation base, it has a universal extension *for every* n. At first glance it appears that Lemma 23.13 suffices to carry through the induction. We need Theorem 23.15 because at the key step in the induction we don't have $(\aleph_{n-1}, 1)$-uniqueness.

Lemma 23.18. *Suppose* $\boldsymbol{K}$ *has* $(\lambda, 1)$*-uniqueness (for every* M *with* $|M| = \lambda$*, there is unique* M' *full over* M*) and* (λ, n)*-completeness. Then if* A_n *is an amalgamation base there is a universal model over* A_n*.*

Proof. Let $\mathcal{S}$ be a $(\lambda, \mathcal{P}^-(n))$-full independent system and suppose that A_n is an amalgamation base. By (λ, n)-completeness there is at least one model N of cardinality λ that is full over A_n and so over the empty set. By $(\lambda, 1)$-completion there is an N' that is full over N with cardinality λ. We show N' is universal over A_n.

Suppose $A_n \subset M$ with $|M| = \lambda$. By $(\lambda, 1)$-completeness there is an M' that is full over M and *a fortiori* full over A_n. Since A_n is an amalgamation base, M' and N can be embedded in some M^*. Without loss of generality, by$(\lambda, 1)$-completion, M^* is full over N. But then N' and M^* are isomorphic over N by $(\lambda, 1)$-uniqueness and the isomorphism imbeds M into N' as required. $\Box_{23.18}$

Now putting together Lemmas 23.18 and 23.16, we have completed the proof of Lemma 23.13 and thus of the main theorem. $\Box_{23.13}$

24
Very few models implies amalgamation over pairs

We write 'amalgamation over pairs' as shorthand for 'some $(\lambda, 2)$-full system is an amalgamation base'. Our goal in this chapter is to show that if there are 'very few' models (i.e. $< 2^{\lambda^+}$) in λ^{++} for $\lambda < \aleph_\omega$ then one can amalgamate over at least one independent pair in λ:

Theorem 23.1 [Shelah] *(For $n < \omega$, $2^{\aleph_n} < 2^{\aleph_{n+1}}$.) An atomic class $\boldsymbol{K}$ that has at least one uncountable model and that has very few models in $\aleph_n$ ($I(\boldsymbol{K}, \lambda) \leq 2^{\aleph_{n-1}}$) for each $n < \omega$ is excellent.*

Theorem 23.15 *(For $n < \omega$, $2^{\aleph_n} < 2^{\aleph_{n+1}}$.) Suppose $\boldsymbol{K}$ satisfies the following hypotheses.*

1. *completeness and uniqueness conditions:*
   - (a) *completeness:*
     - i. *$\boldsymbol{K}$ is $(\lambda, 0)$, $(\lambda, 1)$, $(\lambda, 2)$-complete;*
     - ii. *$\boldsymbol{K}$ is $(\lambda^+, 0)$, $(\lambda^+, 1)$-complete;*
   - (b) *uniqueness:*
     - i. *$\boldsymbol{K}$ is $(\lambda, 0)$, $(\lambda, 1)$-unique;*
     - ii. *$\boldsymbol{K}$ is $(\lambda^+, 0)$-unique;*
2. *Amalgamation of full extensions in λ^+.*
3. *$\boldsymbol{K}$ has very few models in λ^{++}*

then some $(\lambda, 2)$-full system is an amalgamation base.

In Chapter 23, we reduced the proof of Theorem 23.1 to Theorem 23.15. We will discuss in Remark 24.13 the role of set theory in fixing the correct meaning of 'few' in this context. The form in Theorem 23.1 depends only on the WGCH holding below $\aleph_\omega$. Appendix C contains the required set theoretic notations and deduces the necessary combinatorial results from $2^{\aleph_n} < 2^{\aleph_{n+1}}$ for $n < \omega$.

Ambiguously, we may write $|M|$ for the universe of a model (rather than our usual M) when we want to emphasize that we are considering the domain without structure or as usual $|M|$ may denote the cardinality of the model.

We will build a tree of models indexed by $\eta, \nu \in 2^{\leq \lambda^{++}}$ with

$$|M^\eta| = \lambda^+ \times (1 + \lg(\eta))$$

and $\eta \lhd \nu$ implies $M^\eta \prec_{\boldsymbol{K}} M^\nu$. We will aim to ensure that if $\eta\widehat{\ }1 \lhd \nu$, $M^{\eta\widehat{\ }0}$ cannot be embedded in M^ν. In fact, we only obtain that result for η of size λ^+ and this leads to the weakening of the hypothesis to 'very few'.

Now we begin the actual argument. We will construct many (more than very few) non-isomorphic models of power λ^{++} from the hypothesis that there is no 2-amalgamation base in λ. Lemma 24.1 strengthens the conditions on models of cardinality λ. Notation 24.4 through Lemma 24.10 provide a tree of models in λ^+ so that if $M^{\eta\widehat{\ }0}$ imbeds in M^ν with $\eta \lhd \nu$ then $\eta\widehat{\ }0 \lhd \nu$. Then Construction 24.11 and an application of the Devlin-Shelah weak diamond yield the final result.

The following Lemma uses many of the inductive model theoretic hypotheses of Theorem 23.15. We show in Lemma 24.1 that if we cannot amalgamate over any independent pair then a certain stronger non-amalgamation result holds. This condition is applied in Construction 24.9.

Lemma 24.1. *Assume the hypotheses of Theorem 23.15 and $\lambda < \aleph_\omega$. Further assume that no $(\lambda, 2)$-independent pair is an amalgamation base. Suppose that M_0, M_1, M_1', M_2 with $M_0 \prec_{\boldsymbol{K}} M_1 \prec_{\boldsymbol{K}} M_1'$ and $M_0 \prec_{\boldsymbol{K}} M_2$ are models of size λ such that*

1. *M_0 is full, M_1, M_2 are full over M_0, M_1' is full over M_1*
2. *$M_2 \downarrow_{M_0} M_1'$*

Then for every $N \in \boldsymbol{K}$ of power λ that is an elementary extension of M_1 and of M_2 with $N \downarrow_{M_1} M_1'$, there are N^1, N^2 of power λ, which extend $M_1'N$ but cannot be amalgamated over $M_1'M_2$.

Proof. Suppose not and that N is a counterexample; without loss of generality, N is full over M_1M_2. By $(\lambda, 1)$-uniqueness and stationarity, N is a counterexample for all possible M_1'. It is straightforward to make the following construction.

Construction 24.2. *We define by induction on $\alpha < \lambda^+$, M_α^0, M_α^1 so that*

1. *$M_0^0 = M_0$, $M_0^1 = M_1$;*
2. *for $i = 0, 1$, $|M_\alpha^i| = \lambda$, M_α^i is full over M_α^0;*

3. *the sequences are continuous at limits;*

4. *the quadruples* $(M^0_\alpha, M^1_\alpha, M^0_{\alpha+1}, M^1_{\alpha+1})$ *and* (M_0, M_1, M_2, N) *are isomorphic.*

For $i = 0, 1$, let $M^i_* = \bigcup_{\alpha<\lambda^+} M^i_\alpha$ and without loss of generality fix the universe of M^0_α as $\{3i : i < \lambda(1+\alpha)\}$ and $|M^0_\alpha| \bigcup \{3i+1 : i < \lambda(1+\alpha)\}$ as the universe of M^1_α. Clearly M^0_* is full, M^1_* is full over M^0_*, and $M^0_\alpha \downarrow_{M^0_*} M^1_\alpha$.

Now using the assumption that no $(\lambda, 2)$-independent pair is an amalgamation base, we define by induction on $\alpha < \lambda^+$, for each $\eta \in 2^\alpha$, a model M_η with universe $\{3i, 3i+2 : i < \lambda(1+\alpha)\}$ such that:

Claim 24.3. *1. for* $\beta < \lg(\eta)$, $M_{\eta \restriction \beta} \prec_{\boldsymbol{K}} M_\eta$*;*

2. M_η *is full over* $M^0_{\lg(\eta)}$ *and* $M_\eta \downarrow_{M^0_{\lg(\eta)}} M^0_i$*;*

3. $M_{\eta\frown 0}$ *and* $M_{\eta\frown 1}$ *cannot be amalgamated over* $M_\eta M^0_{\lg(\eta)+1}$.

Having defined for $\nu \in 2^{<\lambda^+}$, we extend to $\eta \in 2^{\lambda^+}$; let $M_\eta = \bigcup_{\alpha<\lambda^+} M_{\eta\restriction\alpha}$. We have assumed amalgamation of full models in λ^+. So, for each η, M^1_* and M_η can be amalgamated over M^0_*. So there are models N_η with universe λ^+ and an elementary embedding f_η of M_η into N_η over M^0_*.

Now by the Devlin-Shelah weak diamond, Θ_{λ^+}, (Definition C.4) for some $\alpha = \lambda \times \alpha < \lambda^+, \eta, \nu$ and $N_\alpha =^{\mathrm{df}} N_{\eta\restriction\alpha} = N_{\nu\restriction\alpha}$, we have:

1. $\eta \restriction \alpha = \nu \restriction \alpha, \eta(\alpha) = 0, \nu(\alpha) = 1;$

2. $N_\alpha \prec N_\eta$ and $M^1_* \downarrow_{M^1_\alpha} N_\alpha$;

3. $f_\eta \restriction M_{\eta\restriction\alpha} = f_\nu \restriction M_{\eta\restriction\alpha}$ is into N_α.

We can choose $\beta = \lambda \times \beta$ so that $N_\alpha \cup M^1_{\alpha+1}$ is contained in both $N_\eta \restriction \beta \prec N_\eta$ and $N_\nu \restriction \beta \prec N_\nu$; $f_\eta \restriction M_{\eta\restriction(\alpha+1)}$ is into $N_\eta \restriction \beta$ and $f_\nu \restriction M_{\nu\restriction(\alpha+1)}$ is into $N_\nu \restriction \beta$. Now by 24.2.3 and the choice of M_0, M_1, M_2 and N, the models $N_\eta \restriction \beta$ and $N_\nu \restriction \beta$ can be jointly embedded over $N_\alpha \cup M^0_{\alpha+1}$ by say g_η, g_ν. But then $g_\eta f_\eta$ and $g_\nu f_\nu$ amalgamate $M_{\eta\restriction(\alpha+1)}$ and $M_{\nu\restriction(\alpha+1)}$ over $M_{\eta\restriction\alpha} \cup M^0_{\alpha+1}$ contradicting 24.3.3. Thus we have contradicted the existence of the counterexample N. $\square_{24.1}$

We now consider pairs $(\overline{M}, f)$ consisting of a sequence of models of cardinality λ and with universe a subset of λ^+ and a function f from λ^+ to itself. We will define a partial ordering on these pairs with respect to a function F which assigns a specific amalgamating model to an independent pair. We are able to find F with the following properties because of $(\lambda, 2)$-completeness and $(\lambda, 1)$-uniqueness.

Notation 24.4. *Fix a function* F *defined on tuples* (M_0, M_1, M_2, A) *where each* M_i *is a full model of cardinality* λ, $M_1 \cap M_2 = M_0$, *and* $A \supseteq M_1 M_2$ *with* $|A - M_1 M_2| = \lambda$ *so that:*

1. *$F(M_0, M_1, M_2, A)$ is a model M with universe A such that M_1 and M_2 are freely amalgamated over M_0 inside M.*
2. *If $(M_0^0, M_1^0, M_2^0, A^0)$ and $(M_0^1, M_1^1, M_2^1, A^1)$ are as in part 1) and f is* 1-1 *map from $M_1^0 M_2^0$ to $M_1^1 M_2^1$, which is an isomorphism when restricted to each of M_1^0 and M_2^0 then f can be extended to an isomorphism between $F(M_0^0, M_1^0, M_2^0, A^0)$ and $F(M_0^1, M_1^1, M_2^1, A^1)$.*

We now define a partial ordering of pairs $(\overline{M}, f)$, which depends on our choice of F.

Definition 24.5. 1. *$(\overline{M}, f)$ is a pair where $\overline{M} = \langle M_i : i < \lambda^+ \rangle$ is a continuous increasing sequence of models of size λ with $|M_i| \subset \lambda^+$ such that M_{i+1} is full over M_i and $f : \lambda^+ \to \lambda^+$.*

   (a) Let M_{λ^+} denote the union of $\langle M_i : i < \lambda^+ \rangle$.

   (b) Note that we may associate with any such f a closed unbounded set $C_f \subset \lambda^+$ such that $C_f \cap (\alpha, \alpha + f(\alpha)] = \emptyset$ for any $\alpha \in C_f$.

2. *Define a relation $\leq$ on pairs $(\overline{M}, f)$ by $(\overline{M}^1, f^1) \leq (\overline{M}^2, f^2)$ if the following hold.*

   (a) For $i \leq \lambda^+$, $M_i^1 \prec_{\boldsymbol{K}} M_i^2$.

   (b) $\{i < \lambda^+ : f^1(i) \leq f^2(i)\} \in \mathrm{cub}(\lambda^+)$.

   (c) For some $C \in \mathrm{cub}(\lambda^+)$ and all $\alpha \in C$, if $i \in [\alpha, \alpha + f^1(\alpha)]$ then

$$M_i^2 \cap M_{\lambda^+}^1 = M_i^1$$

   and

$$M_i^2 \underset{M_i^1}{\downarrow} M_{\lambda^+}^1.$$

3. *Now define $\leq_F$ by $(\overline{M}^1, g^1) \leq_F (\overline{M}^2, g^2)$ if $(\overline{M}^1, g^1) \leq (\overline{M}^2, g^2)$ and for some $\zeta < \lambda^{++}$, there is a sequence $\langle (\overline{M}^\xi, f^\xi) : \xi \leq \zeta \rangle$ which is continuous and $\leq$-increasing with $(\overline{M}, g^1) = (\overline{M}^0, f^0)$ and $(\overline{M}^2, g^2) = (\overline{M}^\zeta, f^\zeta)$ such that for every $\xi < \zeta$ there is a closed unbounded set C^ξ of α such that for each β with $\alpha \leq \beta < \alpha + f^\xi(\alpha)$:*

$$M_{\beta+1}^{\xi+1} = F(M_\beta^\xi, M_\beta^{\xi+1}, M_{\beta+1}^\xi, M_{\beta+1}^{\xi+1}).$$

The notion of $\leq_F$ extension will provide an equivalence relation on models which are the limits of pairs $(\overline{M}', f')$ that extend a fixed pair (M, f). We thank Rami Grossberg for pointing out the key to the next series of Claims: apply the idea of the proof (e.g. [48] or Lemma 18.9) that failure of amalgamation in μ implies there is no universal model in μ^+ with one of the equivalence classes playing the role of the universal model.

Claim 24.6. *We have:*

1. $\leq_F$ *is transitive and reflexive on pairs* $(\overline{M}, f)$*;*
2. *Any* $\leq_F$ *increasing sequence of strictly less than* λ^{++} *pairs* $(\overline{M}^i, f^i)$ *has a least upper bound.*

Proof. The first claim is immediate. For the second, suppose the sequence is $(\overline{M}^\xi, f^\xi)$ for ξ less than (without loss of generality) a regular cardinal $\mu \leq \lambda^+$. For $\xi < \zeta < \mu$ let $C_{\xi,\zeta}$ be a cub of λ^+ which witnesses 2b) and 2c) of Definition 24.5.

Suppose first that $\mu \leq \lambda$. Then certainly $\bigcap_{\xi<\zeta<\mu} C_{\xi,\zeta}$ contains a cub C such that if $\alpha \in C$, $\xi < \mu$ and $\beta \in [\alpha, \alpha + f^\xi(\alpha)]$, then $\beta \notin C$. When $\mu = \lambda^+$, let C be the diagonal intersection: $C = \{\gamma : \forall \alpha < \gamma(\gamma \in C_{\xi,\alpha})\}$.

Let $\langle \alpha_i : i < \lambda^+\rangle$ be an increasing enumeration of C. Now define by induction on i, a strictly increasing sequence of β_i by $\beta_0 = \alpha_0$, take unions at limits, and β_{i+1} is either $\beta_i + \sup_{\xi<\mu} f^\xi(\alpha_i)$ if it is bigger than β_i or $\beta_i + 1$ if not. Now we define a new $(\overline{M}, f)$ by letting $M_{\beta_i} = \bigcup_{\xi<\mu} M^\xi_{\alpha_i}$ and for j with $0 < j < \sup_{\xi<\mu} f^\xi(\alpha_i)$, $M_{\beta_i+j} = \bigcup_{\xi_0<\xi<\mu} M^\xi_{\alpha_i+j}$ where ξ_0 is chosen so that $\xi_0 < \xi < \mu$ implies $j < f^\xi(\alpha_i)$. Now let $\overline{M} = \langle M_\beta : \beta < \lambda^+\rangle$ and define f by $f(\beta_i) = \sup_{\xi<\mu} f^\xi(\alpha_i)$ while $f(\beta_i + j) = 0$ for $0 < j < f(\beta_i)$; $(\overline{M}, f)$ is as required, using a similar diagonal intersection to check condition 2d). $\square_{24.6}$

Notation 24.7. *Write* $f + 1$ *for the function defined by* $(f + 1)(i) = f(i) + 1$.

Claim 24.8. *Suppose* $(\overline{M}, f + 1) \leq_F (M^i, g^i)$ *for* $i = 1, 2$. *Then there is a cub* C *such that if* $\delta \in C$ *and* $\beta \in [\delta, \delta + f(\delta)]$, $M^1_\beta \approx_{M_\delta} M^2_\beta$.

Proof. Let C be the diagonal intersection of cubs witnessing each of the three components of the definition $\leq_F$ (including each of the C^ξ in Definition 24.5.3) for each pair. Then for each $\delta \in C$ and each $\beta \in [\delta, \delta + f(\delta)]$, we have

$$M^1_{\beta+1} = F(M_\beta, M^1_\beta, M_{\beta+1}, |M^1_{\beta+1}|)$$

and

$$M^2_{\beta+1} = F(M_\beta, M^2_\beta, M_{\beta+1}, |M^2_{\beta+1}|).$$

But $M^1_\beta \approx_{M_\beta} M^2_\beta$ since both are full extensions of M_β and this isomorphism extends over $M_{\beta+1}$ by stationarity and condition 2) of Notation 24.4 so $M^1_{\beta+1} \approx_{M_\delta} M^2_{\beta+1}$. By induction we have the result (using continuity at limits). $\square_{24.8}$

Now we construct for any pair $(\overline{M}, f)$ a family of extensions M_η of M_{λ^+} (for $\eta \in 2^{\leq\lambda^+}$) such that: there is an $\nu \in 2^{\lambda^+}$ such that for any $(\overline{M}', f')$ with $(\overline{M}, f + 1) \leq_F (\overline{M}', f')$, M_ν can not be embedded into M'_{λ^+} over M_{λ^+}.

Construction 24.9. *[Construction in λ^+] Fix $(\overline{M}, f)$ and $C_f = \{\alpha_i : i < \lambda^+\}$ a cub as in Definition 24.5.1. We define by induction on $\beta \leq \lambda^+$, for each $\eta \in 2^\beta$ a model M_η such that:*

1. *$M_\eta \cap M_{\lambda^+} = M_\beta$; $M_\eta \downarrow_{M_\beta} M_{\beta+1}$;*
2. *For $\gamma < \beta$, $\langle M_{\eta \restriction \gamma} : \gamma < \alpha \rangle$ is an increasing continuous sequence;*
3. *For each $\alpha \in C_f$,*
   (a) *For $\beta \in [\alpha, \alpha + f(\alpha))$, and $\eta \in 2^{\beta+1}$,*
   $$M_\eta = F(M_\beta, M_{\eta \restriction \beta}, M_{\beta+1}, |M_\eta|);$$
   (b) *For $\beta \in [\alpha, \alpha + f(\alpha)]$, $\eta, \nu \in 2^\beta$, if $\eta \restriction \alpha = \nu \restriction \alpha$ then $M_\eta = M_\nu$;*
   (c) *$|M_\eta| = \lambda^+ \times (1 + \beta)$;*
   (d) *For $\beta = \alpha + f(\alpha) + 1$, $\eta, \nu \in 2^\beta$, $\beta < \lambda^+$, with $\eta \restriction \alpha = \nu \restriction \alpha$, if $\eta(\alpha) = \nu(\alpha)$ then $M_\eta = M_\nu$ but if $\eta(\alpha) \neq \nu(\alpha)$ then M_η and M_ν cannot be amalgamated over $M_{\eta \restriction \alpha} M_\beta$.*
   (e) *For $\eta \in 2^{\lambda^+}$, $M_\eta = M^\eta_{\lambda^+}$.*

The construction is routine applying Lemma 24.1 to guarantee 3d).

Claim 24.10. *For any pair $(\overline{M}, f)$ there is an $\eta \in 2^{\lambda^+}$ such that for any $(\overline{M}', f')$ with $(\overline{M}, f+1) \leq_F (\overline{M}', f')$, the M_η from Construction 24.9 can not be embedded into M'_{λ^+} over M_{λ^+}.*

Proof. Suppose for contradiction that for each $\eta \in 2^{\lambda^+}$, there is a sequence $(\overline{N}^\eta, f^\eta)$ with $(\overline{M}, f) \leq_F (\overline{N}^\eta, f^\eta)$ and a function g_η mapping $M^\eta = M^\eta_{\lambda^+}$ into $N^\eta = N^\eta_{\lambda^+}$, fixing M_{λ^+}. By weak diamond, there is a $\delta \in C_f$ (C_f from Definition 24.5.1) and $\eta, \nu \in 2^{\lambda^+}$ such that $\eta \restriction \delta = \nu \restriction \delta$, $\eta(\delta) \neq \nu(\delta)$ but $g_\eta \restriction \delta = g_\nu \restriction \delta$.

Let C be the intersection of the cubs witnessing the $\leq_F$ with $\{\alpha : \alpha = |M_{\eta \restriction \alpha}|\}$ and with C_1. Choose any $\delta \in C$ and let $\gamma = \delta + f(\delta) + 1$. By Claim 24.8, g_η and g_ν witness that $M_{\eta \restriction \gamma}$ and $M_{\nu \restriction \gamma}$ can be amalgamated over $M_{\eta \restriction \delta} M_\gamma$ but this contradicts condition 3d) of Construction 24.9. $\square_{24.10}$

In the next step we build a tree $(\overline{M}^\nu, f^\nu)$ of length λ^{++} where Construction 24.9 provides the successor stage. Each $\overline{M}^\nu$ is a λ^+-sequence of models of cardinality λ. $M^\nu_{\lambda^+}$ is the union for $\beta < \lambda^+$ of the sequence $\overline{M}^\nu$.

Construction 24.11. *[Construction in λ^{++}] For every $\alpha < \lambda^{++}$ and $\nu \in 2^\alpha$ there is a pair $(\overline{M}^\nu, f^\nu)$ such that the following hold.*

1. *$\langle (\overline{M}^{\nu \restriction \gamma}, f^{\nu \restriction \gamma}) : \gamma < \alpha \rangle$ is continuous and increasing. Indeed $\nu \lhd \eta$ implies $(\overline{M}^\nu, f^\nu) \leq_F (\overline{M}^\eta, f^\eta)$.*

2. $|M^{\nu}_{\lambda^+}| = \lambda^+(1 + \lg(\nu))$.

3. *If* $(M^{\eta\frown 0}, f^{\eta\frown 0}) \leq_F (\overline{N}, f)$ *then* $M^{\eta\frown 1}_{\lambda^+}$ *cannot be elementarily embedded in* N_{λ^+} *over* M_{λ^+}. *In particular, if* $N_{\lambda^+} = M^{\nu}_{\lambda^+}$ *then* $\eta\frown 0 \lhd \nu$.

4. *For* $\eta \in 2^{\lambda^{++}}$, *let* M_η *denote* $\bigcup_{\alpha<\lambda^{++}} M^{\eta\restriction\alpha}_{\lambda^+}$.

Proof. Define the $(\overline{M}^{\nu}, f^{\nu})$ by induction on $\alpha = \lg(\nu) < \lambda^{++}$. For $\alpha = 0$, fix $M^{\langle\rangle}_{\lambda^+}$ as any full model of cardinality λ^+. Let $\overline{M}^{\langle\rangle} = M^{\langle\rangle}_i$ for $i < \lambda^+$ be a filtration of $M^{\langle\rangle}_{\lambda^+}$. And let $f^{\langle\rangle}$ be the constantly 0 function with domain λ^+. For α a limit apply Lemma 24.6.

For $\alpha = \beta + 1$, a successor, apply Claim 24.10 to the pair $(\overline{M}^{\nu}, f^{\nu})$ for each $\nu \in 2^{\beta}$. $\overline{M}^{\nu\frown 1}$ is the resulting $\overline{M}_\eta$; $f^{\nu\frown 1} = f^{\nu}$. We choose $\overline{M}^{\nu\frown 0}$ as an immediate F-successor of $(\overline{M}^{\nu}, f^{\nu} + 1)$ while $f^{\nu\frown 0} = f^{\nu} + 1$. $\square_{24.11}$.

$\Theta_{\chi,\lambda^{++}}$ is defined as Definition C.4 in Appendix C. We now prove the key lemma that forces us to make *very few* rather than just few the hypothesis of the main result.

Lemma 24.12. *If* $\Theta_{\chi,\lambda^{++}}$ *and no* $(\lambda, 2)$*-full system is an amalgamation base then* $I(\boldsymbol{K}, \lambda^{++}) > \chi$.

Proof. Fix representatives $\langle N_i : i < \chi\rangle$ on λ^{++} of the isomorphism types of models in $\boldsymbol{K}$ of cardinality λ^{++}. For each $\eta \in 2^{\lambda^{++}}$ fix f'_η that is an isomorphism between M^{η} from Construction 24.11.4 and the appropriate N_i. Now let f_η be the same as f'_η except $f_\eta(0) = i$ if and only if the image of f_η is N_i and for $1 < i < \omega$, $f_\eta(i) = f'_\eta(i - 1)$. Now apply $\Theta_{\chi,\lambda^{++}}$ using the cub C of α such that $\alpha = \lambda \times \alpha$ (so $|M^{\eta\restriction\alpha}_{\lambda^+}| = |M^{\nu\restriction\alpha}_{\lambda^+}| = \alpha$). We find $\delta \in C$ and $\eta, \nu \in 2^{\lambda^{++}}$ so that $\eta\restriction\delta = \nu\restriction\delta$, $\eta(\delta) \neq \nu(\delta)$, $f_\eta\restriction\delta = f_\nu\restriction\delta$. In particular $f_\eta(0) = f_\nu(0)$ so f_η and f_ν have isomorphic images. Without loss, $\eta(\delta) = 0$ and $\nu(\delta) = 1$. Then $f^{-1}_\eta(f_\nu \restriction M^{\nu\restriction(\delta+1)}_{\lambda^+})$ is an elementary embedding of $M^{\nu\restriction(\delta+1)}_{\lambda^+} = M^{(\nu\restriction\delta)\frown 1}_{\lambda^+}$ into M^{η} that is the identity on $M^{\nu\restriction\delta}_{\lambda^+} = M^{\eta\restriction\delta}_{\lambda^+}$. Since $M^{\nu\restriction(\delta+1)}_{\lambda^+}$ has cardinality λ^+ this map must be into M^{ρ} for some ρ with $\eta\restriction(\delta + 1) \lhd \rho \lhd \eta$ and $\lg(\rho) < \lambda^{++}$. This contradicts Claim 24.11.3. $\square_{24.12}$

Proof of Theorem 23.15 and Theorem 23.1. By Chapter 23, we need only prove Theorem 23.15. Theorem C.6 yields $\Theta_{\chi,\lambda^{++}}$, for $\chi = 2^{\lambda^+}$. We want to show that some $(\lambda, 2)$-full system is an amalgamation base if $I(\boldsymbol{K}, \lambda^{++}) \leq \lambda^{++}$ If not, Lemma 24.12 yields that $I(\boldsymbol{K}, \lambda^{++}) > \lambda^{++} \geq I(\boldsymbol{K}, \lambda^{++})$. With this contradiction we finish. $\square_{23.15}$

Remark 24.13. 1. It is erroneously asserted in [125] (Theorem 6.4 and footnote on page 265) that $\chi^{\aleph_0} < 2^{\lambda^{++}}$ is sufficient to obtain $I(\boldsymbol{K}, \lambda^{++}) > \chi$; this error was known at the time of publication but these instances of the earlier claim slipped through the editing process. The question of whether

$2^{\lambda^+} < 2^{\lambda^{++}}$ and $\chi^{\aleph_0} < 2^{\lambda^{++}}$ implies $\Theta_{\chi,\lambda^{++}}$ remains open. See Remark C.12 for more details.

2. Shelah has provided more model theoretic approaches to get around the set theoretic difficulties in [133] and later works which are not yet published.

3. Either GCH or $\neg O^{\#}$ imply that 'very few ' can be replaced by 'few'. See the last two pages of [125] or [133].

I particularly want to thank Rami Grossberg and Saharon Shelah for background on this section and Alexei Kolesnikov for his careful reading of this argument; the errors which remain are mine.

25
Excellence and *-Excellence

In this chapter we adapt the ideas of [100] to derive from Shelah's notion of excellence several consequences, which we call $*$-excellence. The proof of categoricity transfer will actually assume only $*$-excellence. We show that if the atomic class $\boldsymbol{K}$ is excellent (i.e. all countable independent *full* diagrams are good) then they same holds without the assumption that the models are full. We use without further remark the fact that Chapter 22 did not require that the models are full (although the full case was proved there).

Definition 25.1. *The atomic class* $\boldsymbol{K}$ *is* *-excellent *if*

1. $\boldsymbol{K}$ *has arbitrarily large models;*
2. $\boldsymbol{K}$ *is* ω*-stable;*
3. $\boldsymbol{K}$ *satisfies the amalgamation property;*
4. *Let* p *be a complete type over a model* $M \in \boldsymbol{K}$ *such that* $p \restriction C$ *is realized in* M *for each finite* $C \subset M$*, then there is a model* $N \in \boldsymbol{K}$ *with* N *primary over* $M\boldsymbol{a}$ *such that* p *is realized by* $\boldsymbol{a}$ *in* N*.*

Remark 25.2 (Excellence for first order theories)**.** Since ω-stable first order theories have primary models over arbitrary sets, and since the primary model over an atomic set is atomic if $(\boldsymbol{K}, \prec)$ is the collection of the atomic models of an ω-stable first order theory with elementary submodel as $\prec_{\boldsymbol{K}}$ then $\boldsymbol{K}$ is excellent. And in the specific way we have defined excellence (as an ω-stable class of atomic models), these are the only first order examples. If all models of T are atomic, then T is $\aleph_0$-categorical.

However, there is a more general sense in which excellence is used; fix a class $\boldsymbol{K}$ of models along with notions of independent, strong submodel, and primary. Drop the ω-stability requirement and call the class excellent just if this class admits primary models (which must stay in the class) over independent $\mathcal{P}^{-}(n)$-system of models. Since these three notions must each be specified there is no simple response to the question, 'what does excellence mean in the first order context?' In all the following examples of classes of models of countable first order theories, nonforking is the notion of independence; the others concepts vary. A unifying theme is that, unlike the infinitary case, if the class admits 'primary' models over independent pairs then n-amalgamation follows for all $n < \omega$. See [126, 55] for definitions, specifications of the notion of strong submodel, and for proofs of these claims. The class of a-models of a superstable first order theory is excellent if and only if T has ndop. To discuss arbitrary models of a superstable first order theory one must distinguish a particular notion of strong submodel and restrict to theories with ndop and notop. If we consider all models of an ω-stable theory, then since ω-stable theories admit prime models over sets, any ω-stable theory is excellent. But the models will not be sufficiently rich to support the development of the rest of the theory. Laskowski [95] and later unpublished work has observed that, if we take the notion of primary in the standard first order sense, the class of ω-saturated models of a countable ω-stable theory T is excellent if and only T has the eni-ndop. But, if we modify the language in our usual way by making each first order type over a finite set atomic, then by Theorem 4.18 of [122], then each class associated with ω-saturated models of an ω-stable theory is excellent.

Excellence is this generalized sense is a main contributor to defining the main gap. It is those classes which are excellent (and in addition) shallow that have a few models. And these models can all be realized as constructible (in an appropriate sense depending on the context) over independent trees of height ω. A more general study of excellent atomic classes appears in [49].

We return to our analysis of categoricity in atomic classes and take the following argument from [100]. Note the similarity to Lemma 23.7.

Lemma 25.3. *If M is a countable member of an ω-stable atomic class $\boldsymbol{K}$, there is a countable $N \in \boldsymbol{K}$ which is $\aleph_0$-universal over M (Definition 11.4).*

Proof. Let $\langle p_i : i < \omega\rangle$ enumerate $S_{at}(M)$. Define an increasing chain of models $\langle N_i : i < \omega\rangle$ with $N_0 = M$ and so that N_{i+1} is primary over $N_i c_i$ where $c_i \underset{N_0}{\downarrow} N_i$ and c_i realizes p_i. Let $N = \bigcup_i N_i$. Now let M' be an arbitrary countable member of $\boldsymbol{K}$ with $M \prec_{\boldsymbol{K}} M'$; enumerate $M' - M$ as $\langle a_i : i < \omega\rangle$. We will construct an increasing sequence of partial elementary maps f_i that are the identity on M and the domain of f_i is $M \bigcup\langle a_j : j < i\rangle$. Let $M_i \prec_{\boldsymbol{K}} M'$ be primary over $M \bigcup\langle a_j : j < i\rangle$. Given f_i, let k be least such that the range of f_i is contained in N_k. By the definition of primary extend f_i to an elementary embedding f'_i of M_i into N_k. Let $q = \operatorname{tp}(a_{i+1}/M_i)$. Then $f'_i(q) \in S_{at}(f'_i(M_k))$ has a nonsplitting extension $q' \in S_{at}(N_k)$, since the domain of f_i is good by Lemma 19.25. By construction, q' is realized by some $b \in N_{k+1}$. Let $f_{i+1} = f_i \cup \langle a_{i+1}, b\rangle$. This suffices. $\square_{25.3}$

If A is a countable good set, we can first take a primary model M over A and then a universal model over M to prove:

Corollary 25.4. *If A is a countable good set then there is a countable $N \in \boldsymbol{K}$ which is $\aleph_0$-universal over M*

Corollary 25.5. *The countable atomic set A is good if and only if $S_{at}(A)$ is countable.*

Proof. If A is good, there is a primary model M over A (Lemma 19.26). By Lemma 25.3, there is a countable N which is universal over M and therefore over A. So each $p \in S_{at}(A)$ is realized in N and $S_{at}(A)$ is countable. The converse is Lemma 19.24.1. $\square_{25.5}$

In Chapter 23 we deduced excellence for full countable models for an atomic class with few models in each power up to $\aleph_\omega$. We need to know arbitrary countable independent systems are good. So we now show: If every *full* independent system of countable models is good then every countable independent system of models is good.

We first show the strengthened form of stationarity in Lemma 23.6 holds for all countable models, not just full ones.

Lemma 25.6. *If $\boldsymbol{K}$ satisfies $(\aleph_0, 2)$-uniqueness for any independent triple (M_0, M_1, M_2), we have: if $M_1' \approx_{M_0} M_1$ and $M_1' \underset{M_0}{\downarrow} M_2$ then there exists $\alpha \in \mathrm{aut}_{\mathrm{M}_2}\mathbb{M}$ mapping M_1 isomorphically onto M_1'.*

Proof. By $\aleph_0$-categoricity M_0 is full. By $(\aleph_0, 1)$ extension, we can extend M_1, M_1' and M_2 to full models $\hat{M}_1, \hat{M}_1'$ and $\hat{M}_2$ with $\hat{M}_1 \approx_{M_0} \hat{M}_1'$ by some f. By Lemma 23.6, there is an automorphism of $\mathbb{M}$ extending f and fixing $\hat{M}_2$. This map suffices. $\square_{25.6}$

Before beginning the formal argument, we first consider the case of an independent pair. Let M_1 and M_2 be countable and independent over M_0. We would like to show that M_1M_2 is good (i.e. there are only countably many atomic types with domain M_1M_2).

We assume the conclusion holds if each M_i is full over M_0. The key observation (by Shelah of course) is that we prove the result in two steps.

Suppose first that M_1 is full over M_0 while M_2 may not be. M_0 is full by $\aleph_0$-categoricity. We want to show that there is a good countable set $A \supseteq M_1M_2$ so that for each $p \in S_{at}(M_1M_2)$, p extends to $\hat{p} \in S_{at}(A)$. Since $S_{at}(A)$ is countable this shows $S_{at}(M_1M_2)$ is countable and establishes the theorem.

Choose N_2 extending M_2 that is countable, full over M_0 and independent from M_1M_2 over M_2. By transitivity, M_1 is independent from N_2 over M_0. Then by hypothesis, $A = M_1N_2$ is good. Let N' be countable and universal over M_0.

Since there is a free amalgam of M_1 and N_2 over M_0, we may assume $M_1N_2 \subset N'$. Now choose N'' that is universal over M_2 and finally (by amalgamating N' and N'' over M_2) choose N with $N' \subset N$ and there is an embedding of N'' into N over M_2.

Now if $p \in S_{at}(M_1M_2)$ is realized by some a, we can, since M_1M_2a is countable, choose $M_a \in \boldsymbol{K}$ containing M_1M_2a. By universality over M_2, there is an α mapping M_a into N over M_2 and taking M_1 to some $M_1' \approx_{M_2} M_1$. By Lemma 25.6 and universality over M_2 there is a β, fixing N_2, taking M_1' to M_1 and mapping a into N, so $\beta\alpha(a)$ realizes p. Since N is atomic, $\hat{p} = \mathrm{tp}(\beta\alpha(a)/M_1N_2)$ is the required extension of p.

Now we know that M_1M_2 is good if M_2 is full over M_0 and M_1 is arbitrary; we want it for both M_1, M_2 arbitrary. Choose N_1 full over M_1 and independent from M_2 over M_1. We have just shown that N_1M_2 is good. Repeat the preceding argument, to conclude that M_1M_2 is good.

We pass to the general induction. Recall that in an $(\aleph_0, n)$-independent system $\mathcal{S}$, $A_s^{\mathcal{S}} = \bigcup_{t \subset s} M_t$ and in particular $A_n^{\mathcal{S}}$ is the union of all members of a $\mathcal{P}^-(n)$-system.

Lemma 25.7. *If every* full *independent system of countable models is good then every countable independent system of models is good.*

Proof. We prove by induction on n:

$(*_n)$ if each full $\mathcal{P}^-(n)$ countable independent system is good, an arbitrary countable $\mathcal{P}^-(n)$ independent system is good.

For each n there is a sub-induction on the invariant $k(\mathcal{S})$ that we now define. If $\mathcal{S}$ is an independent $\mathcal{P}^-(n)$-system, $k(\mathcal{S})$ is the number of $i < n$ such that for some $s \in \mathcal{P}^-(n)$, $M_{\{i\}} \prec_{\boldsymbol{K}} M_s$ (i.e. $i \in s$) and M_s is not full over A_s.

If $k(\mathcal{S}) = 0$, the result holds by the hypothesis of the theorem. We will show that if $(*_n)$ holds for each $\mathcal{S}$ with $k(\mathcal{S}) < m$ then it holds for each $\mathcal{S}$ with $k(\mathcal{S}) = m$. So fix $\mathcal{S} = \{M_s : s \subset n\}$ with $k(\mathcal{S}) = m$. To show $\mathcal{S}$ is good we must show every $p \in S_{at}(A_n^{\mathcal{S}})$ extends to a type over a fixed good set. Since $k(\mathcal{S}) > 0$ there is an i such that for some s, $M_{\{i\}} \prec_{\boldsymbol{K}} M_s$ (i.e. $i \in s$) and M_s is not full over A_s. Without loss of generality, say $i = n-1$. Let $X = \{s \subset n : n-1 \in s\}$; then, $|X| = 2^{n-1} - 1$. This is a $\mathcal{P}^-(n-1)$-independent system so by the main induction hypothesis $\bigcup\{M_s : s \in X\}$ is good so by Lemma 25.4 there is a countable universal model N' over $\bigcup\{M_s : s \in X\}$. Now for each s with $n-1 \notin s$, we let $M_s' = M_s$ and for each $s \in X$ we define a new M_s' so that

1. $\{M_s' : s \in X\}$ is a full independent system indexed by X;
2. $\bigcup_{s \in X} M_s' \downarrow_{\bigcup_{s \in X} M_s} A_n^{\mathcal{S}}$;
3. $\mathcal{S}' = \{M_s' : s \subset n\}$ is a $\mathcal{P}^-(n)$-independent system with $k(\mathcal{S}) < m$.

If we have completed the construction we see that $A_n^{\mathcal{S}} \cup \bigcup_{s \in X} M_s' = A_n^{\mathcal{S}'}$ is an atomic set by 2), Lemma 20.7, Lemma 21.14.1, and Lemma 21.11.2.

By induction on $k(\mathcal{S})$, $A_n^{\mathcal{S}'}$ is good. Moreover, we can imbed $A_n^{\mathcal{S}'}$ in N' over $\{M_s : s \in X\}$. By 2) this embedding is over $A_n^{\mathcal{S}}$ (applying now Lemma 25.6 and again Lemma 21.14.1) and N' realizes any $p \in S_{at}(A_n^{\mathcal{S}})$ and so p extends to a $\hat{p} \in S_{at}(A_n^{\mathcal{S}'})$ as required.

For the construction fix an enumeration r of $\mathcal{S}$ which has $\{M_s : n-1 \notin s\}$ as an initial segment. For $j < 2^{n-1}$, let $M'_{r(j)} = M_{r(j)}$. Then continue to define $M'_{r(j)}$ for $2^{n-1} < j < 2^n - 1$ to build an enumeration of $\mathcal{S}'$. At each stage demand that $M'_{r(j)}$ is full over $A^{\mathcal{S}'}_{r(j)} = \bigcup_{t \subset r(j)} M'_t$ and using Lemma 20.7, Lemma 21.14.1, and Lemma 21.11.2, that

$$M'_{r(j)} \underset{M_{r(j)} A^{\mathcal{S}'}_{r(j)}}{\downarrow} \bigcup_{i<j} M'_{s(i)}.$$

The independence calculus yields $\bigcup_{s \in X} M'_s \downarrow_{\{M_s s \in X\}} A_n^{\mathcal{S}}$ and thus Conditions 1), 2) and the first part of 3) are verified by the properties of independence; to see that $k(\mathcal{S}') < m$, note that if $m-1 \in s$, M'_s is full over $A_s^{\mathcal{S}'}$. $\square_{25.7}$

The following lemma is the remaining crucial step in showing $*$-excellence. We show that excellence implies that, like their countable counterparts, if a (λ, n)-system is good, there is a primary model over it.

Lemma 25.8. *Let λ be an infinite cardinal $n < \omega$. If $\boldsymbol{K}$ has primary models over A_n for $(\mu, n+1)$-independent systems for each $\mu < \lambda$, then there is a primary model over any (λ, n)-independent system.*

Proof. We know by Lemma 25.7 and Lemma 19.26 there are primary models over $(\aleph_0, m)$-independent systems for any m. Suppose $\mathcal{S}$ is a (λ, n)-independent system.

Choose a filtration $\mathcal{S}^\alpha$ (with respect to L^*) as in Lemma 22.6. We can further choose N^α for $\alpha < \lambda$ such that:

1. $|N^\alpha| = |\alpha| + \aleph_0$.
2. $A_n \underset{A_n^\alpha}{\perp} N^\alpha$
3. N^α is primary over A_n^α;
4. $N_i = \bigcup_{j<i} N_j$ for limit i.

For the initial step we are given N^0 satisfying 1) and 3); Condition 2) is obtained by the extension property for non-splitting. For the induction step, note that using the $(\aleph_0, n)$-goodness, Lemma 22.6, and Lemma 21.4, the system $\langle M_s^\alpha, M_s^{\alpha+1}, N_s^\alpha \subseteq n \rangle$ is an $(|\alpha| + \aleph_0, \mathcal{P}^-(n+1))$ independent system. So we can find $N^{\alpha+1}$ satisfying 1) and 3) by the hypothesis of the theorem and Condition 2) is again obtained by the extension property for non-splitting. Take unions at limits. Then $N = \bigcup_{\alpha<\lambda} N_\alpha$ is the required completion of $\mathcal{S}$. Conditions 2 and 4 guarantee that N is full over A_n. Since $|A_n| = \lambda$, N is full. $\square_{25.8}$

Lemma 25.9 (Dominance)**.** *Suppose $A \underset{M}{\downarrow} B$ where $M \in \boldsymbol{K}$ and ABM is an atomic set. If M' is primary over MA then $M' \underset{M}{\downarrow} B$*

Proof. By Lemma 21.6 (with M as C, MA as A and MAB as B, $MA \leq_{tv} MAB$. By Lemma 21.16.2, for any $\mathbf{c} \in M'$, $\mathbf{c} \downarrow_{M\boldsymbol{a}} MAB$ where some $\theta(\mathbf{x}, \boldsymbol{a})$ generates $\mathrm{tp}(\mathbf{c}/MA)$. But now transitivity of independence yields for each $\mathbf{b} \in B$, $\mathbf{b} \downarrow_M MA\mathbf{c}$. By the finite character of independence we finish. $\square_{25.9}$

Lemma 25.10. *If for all $\mu < \lambda$, there is a primary model over any independent pair of models of size μ, then for any model M of cardinality λ and any $\boldsymbol{a}$ such that $M\boldsymbol{a}$ is atomic, there is a primary model N over $M\boldsymbol{a}$.*

Proof. Write M as an increasing continuous chain of M_i with $|M_i| = |i| + \aleph_0$. Without loss of generality $\boldsymbol{a} \downarrow_{M_0} M$. Since M_0 is countable, there is a primary model N_0 over $M_0\boldsymbol{a}$. By the extension Theorem 20.9 and stationarity, we may assume $N_0 \underset{M_0}{\downarrow} M$. Suppose we have constructed an increasing continuous elementary chain N_i for $i < j$ with $N_i \underset{M_i}{\downarrow} M$. If j is a limit take $\bigcup_{i<j} N_i$ as N_j and note that by finite character $N_j \underset{M_j}{\downarrow} M$. If $j = i + 1$, note that by induction (M_i, N_i, M_j) is an $(|i| + \aleph_0, 2)$ system. Choose N_j primary over $N_i \cup M_j$ with $N_j \underset{M_j}{\downarrow} M$ by the $(|i|+\aleph_0, 2)$-existence property and Lemma 25.9. This completes the construction; it only remains to note that $N = \bigcup_{i<\lambda} N_i$ is primary over $M\boldsymbol{a}$. But this follows by induction using Exercise 25.11. $\square_{25.10}$

The following exercise completes the proof of Lemma 25.10.

Exercise 25.11. *Verify that N_δ is primary over $M_\delta\boldsymbol{a}$ for limit δ.*

Theorem 25.12. *If $\boldsymbol{K}$ is excellent then $\boldsymbol{K}$ is *-excellent.*

Proof. By definition $\boldsymbol{K}$ is ω-stable. We proved $\boldsymbol{K}$ has arbitrarily large models and amalgamation in Corollary 22.13. For the last condition, note first that if p is a complete type over a model $M \in \boldsymbol{K}$ such that $p \restriction C$ is realized in M for each finite $C \subset M$ then for any a realizing p, Ma is atomic. Then apply Lemma 25.10. $\square_{25.12}$

We now show that *-excellence implies that Galois-types are syntactic types in this context. This implies that $\boldsymbol{K}$ is $(\aleph_0, \infty)$-tame. This is straightforward from $*$-excellence but $*$-excellence was a non-immediate consequence of excellence.

In any atomic class types over sets make syntactic sense as in first order logic, but we have to be careful about whether they are realized. The definition of Galois type depended on the choice of a monster model. Since we have amalgamation over models, for any M and any $p \in S_{\mathrm{at}}(M)$, p is realized in our monster. But if M were an arbitrary subset A of $\mathbb{M}$, this might not be true; it depends on the embedding of A into $\mathbb{M}$.

Lemma 25.13. *If $\boldsymbol{K}$ is *-excellent then Galois types over a model M are the same as syntactic types in $S_{\mathrm{at}}(M)$.*

Proof. Equality of Galois types *over models* is always finer than equality of syntactic types. But if a, b realize the same $p \in S_{\mathrm{at}}(M)$, by 2) of Definition 25.1, we can map Ma into any model containing Mb and take a to b so the Galois types are the same. $\square_{25.13}$

Note however, we have more resources here than in a general AEC. The types in $S_{\mathrm{at}}(A)$ for A atomic played a crucial role in the analysis for $L_{\omega_1,\omega}$. I know of no use of Galois types over arbitrary subsets of models of $\boldsymbol{K}$ in Shelah's work. However, Hyttinen and Kerala use this concept in their insightful attempts to extend geometric stability theory to AEC [69, 65]. They introduce the notion of a finitary AEC. In Kueker's formulation,

Definition 25.14.
1. *An AEC $\boldsymbol{K}$ has* finite character *if for $M \subseteq N$ with $M, N \in \boldsymbol{K}$: if for every finite $\boldsymbol{a} \in M$ there is a $\boldsymbol{K}$-embedding of M into N fixing $\boldsymbol{a}$, then $M \prec_{\boldsymbol{K}} N$.*
2. *An AEC is* finitary *if $\boldsymbol{K}$ has arbitrarily large models, the amalgamation property, the joint embedding property, and has finite character.*

Hyttinen and Kesälä develop an extensive theory of finitary aec in [65, 64, 66]. They develop such tools as weak types and U-rank in this context. They prove strong categoricity transfer results (even for limit cardinals) for tame, simple, finitary AEC. For the class of omega-saturated models, they get the full Morley categoricity transfer theorem. The notion of *simple* here has a very different connotation from first order logic because the independence notion is defined to imply the existence of free extensions. Thus, it is possible to have stable and even (Example 27.1.2) totally categorial classes that are not simple in their sense.

Finitary AEC are much closer to $L_{\omega_1,\omega}$ than the general notion. But, in contrast to Part IV, they also include incomplete sentences of $L_{\omega_1,\omega}$ so study of countable models is possible. In addition to the work of Hyttinen and Kesala, see Kueker, [89], who clarifies the connection with specific infinitary logics and Trlifaj [142] who explores classes of modules that are finitary AECs.

26

Quasiminimal Sets and Categoricity Transfer

We work in an atomic class $\boldsymbol{K}$. That is, $\boldsymbol{K}$ is the class of atomic models of a first order theory T, which was obtained from a complete sentence in $L_{\omega_1,\omega}$ by adding predicates for all formulas in a countable fragment L^* of $L_{\omega_1,\omega}$.

In this chapter we first construct quasiminimal formulas in a *-excellent class, then we prove categoricity transfer in excellent classes. Finally we conclude categoricity transfer for arbitrary sentences of $L_{\omega_1,\omega}$. The notion of quasiminimality here will generalize Zilber's notion in the sense that his models are ones where the universe is quasiminimal This assertion is formalized in Proposition 26.20. This chapter is indirectly based on [120, 124, 125], where most of the results were originally proved. But our exposition owes a great deal to [100, 99, 87, 49].

We begin by introducing the notion of a big type. This is formally a different notion than a big Galois type introduced in Definition 13.20. But the intent is the same and eventually this can be seen to be specialization of the earlier notion to this context.

Definition 26.1. *The type $p \in S_{at}(A)$ that is contained in a model in $\boldsymbol{K}$, is* big *if for any $M' \supseteq A$ with $M' \in \boldsymbol{K}$ there exists an N' with $M' \prec_{\boldsymbol{K}} N'$ and with a realization of p in $N' - M'$.*

Lemma 26.2. *Let $A \subseteq M$ and $p \in S_{\mathrm{at}}(A)$. The following are equivalent.*

1. *There is an N with $M \prec N$ and $c \in N - M$ realizing p; i.e. p extends to a type in $S_{\mathrm{at}}(M)$.*

2. *For all M' with $M \prec M'$ there is an N', $M' \prec N'$ and some $d \in N' - M'$ realizing p.*

Proof. 2) implies 1) is immediate. For the converse, assume 1) holds. Without loss of generality, by amalgamation, M' contains N. Let $q = \mathrm{tp}(\mathbf{c}/M)$. By Theorem 20.9, there is a nonsplitting extension $\hat{q}$ of q to $S_{\mathrm{at}}(M')$; $\hat{q}$ is realized in $N' \in \boldsymbol{K}$ with $M' \prec_{\boldsymbol{K}} N'$. Moreover, it is not realized in M' because $\hat{q}$ does not split over M. $\square$26.2

For countable M', we will see below how to get N' via the omitting types theorem. But the existence of N' for uncountable cardinalities requires the use of n-dimensional cubes in $\aleph_0$. By iteratively applying Lemma 26.2, we can show:

Corollary 26.3. *Let $A \subseteq M$ and $p \in S_{\mathrm{at}}(A)$. If there is an N with $M \prec N$ and $c \in N - M$ realizing p then*

1. *p is big and*
2. *$\boldsymbol{K}$ has arbitrarily large models.*

Thus every nonalgebraic type over a model and every type with uncountably many realizations (check the hypothesis via Lowenheim-Skolem) is big. But if we consider a $\boldsymbol{K}$ with only one model: two copies of (Z, S), we see a type over a finite set can have infinitely many realizations without being big.

We introduce a notion of *quasiminimal set*; we will that the universe of an $L_{\omega_1,\omega}$-definable quasiminimal excellent class as in Chapter 3 is quasiminimal in the current sense. And the current notion specializes the notion of minimal type in Definition 13.21.

Definition 26.4. *The type $p \in S_{\mathrm{at}}(A)$ is* quasiminimal *if p is big and for any M containing A, p has a unique extension to a type over M which is not realized in M. We say $\phi(x\mathbf{c})$ is a* quasiminimal formula *if there is a unique $p \in S_{at}(\mathbf{c})$ with $\phi(x\mathbf{c}) \in p$ that is big. We then write that $\phi(x\mathbf{c})$ determines p.*

Note that whether $q(x, \boldsymbol{a})$ is big or quasiminimal is a property of $\mathrm{tp}(\boldsymbol{a}/\emptyset)$. Since every model is ω-saturated the minimal vrs strongly minimal difficulty of first order logic does not arise. Now almost as one constructs a minimal set in the first order context, we find a quasiminimal type; more details are in [99]

Lemma 26.5. *Let $\boldsymbol{K}$ be *-excellent. For any $M \in \boldsymbol{K}$, there is a $\mathbf{c} \in M$ and a formula $\phi(x, \mathbf{c})$ which is quasiminimal.*

Proof. It suffices to show the countable model has a quasiminimal formula $\phi(x, \mathbf{c})$ (since quasiminimality of depends on the type of $\mathbf{c}$ over the empty set). As in the first order case, construct a tree of formulas which are contradictory at each stage and are big. But as in the proof of Lemma 20.3 make sure the parameters in each infinite path exhaust M. Then, if we can construct the entire tree ω-stability is contradicted as in Lemma 20.3. So there is a quasiminimal formula. $\square_{26.5}$

Definition 26.6. *Let $\mathbf{c} \in M \in \boldsymbol{K}$ and suppose $\phi(x, \mathbf{c})$ determines a quasiminimal type over M. For any elementary extension N of M define* cl *on the set of realizations of $\phi(x, \mathbf{c})$ in N by $a \in \mathrm{cl}(A)$ if $\mathrm{tp}(a/A\mathbf{c})$ is not big.*

Equivalently, we could say $a \in \mathrm{cl}(A)$ if every realization of $\mathrm{tp}(a/A\mathbf{c})$ is contained in each $M' \in \boldsymbol{K}$ which contains $A\mathbf{c}$.

Lemma 26.7. *Let* $\mathbf{c} \in M \in \boldsymbol{K}$ *and suppose* $\phi(x, \mathbf{c})$ *determines a quasiminimal type over* M*. If the elementary extension* N *of* M *is full with* $|N| > |M|$*, then* cl *defines a pre-geometry on the realizations of* $\phi(x, \mathbf{c})$ *in* N*.*

Proof. Clearly for any a and A, $a \in A$ implies $a \in \mathrm{cl}(A)$. To see that cl has finite character note that if $\mathrm{tp}(a/A\mathbf{c})$ is not big, then it differs from the unique big type over $A\mathbf{c}$ and this is witnessed by a formula so a is in the closure of the parameters of that formula.

For idempotence, suppose $a \in \mathrm{cl}(B)$ and $B \subseteq \mathrm{cl}(A)$. Use the comment after Definition 26.6. Every $M \in \boldsymbol{K}$ which contains A contains B and every $M \in \boldsymbol{K}$ which contains B contains a; the result follows.

It is only to verify exchange that we need the fullness of N. Suppose $a, b \models \phi(x, \mathbf{c})$, each realizes a big type over $A \subseteq \phi(N)$ and $r = \mathrm{tp}(b/Aa\mathbf{c})$ is big. Since $r = \mathrm{tp}(a/A\mathbf{c})$ is big and N is full we can choose λ realizations a_i of r in N. Let $M' \prec N$ contain the a_i and let b' realize the unique big type over M' containing $\phi(x, \mathbf{c})$. Since $\mathrm{tp}(b/Aa\mathbf{c})$ is big, the uniqueness yields all pairs (a_i, b') realize the same type $p(x, y) \in S(A\mathbf{c})$ as (a, b). But then the a_i are uncountably many realization of $\mathrm{tp}(a/Ab\mathbf{c})$ so this type is big as well; this yields exchange by contraposition. $\square_{26.7}$

Exercise 26.8. *Show that closure relation defined in Lemma 26.7 satisfies the countable closure condition.*

So the dimension of the quasiminimal set is well-defined. To conclude categoricity, we must show that dimension determines the isomorphism type of the model; this is the topic of the next chapter.

The hypothesis that $\boldsymbol{K}$ is *-excellent is not needed for the definability of types that we discuss next.

Definition 26.9. *The type* $p(\mathbf{x}) \in S_{at}(M)$ *is definable over the finite set* $\mathbf{c}$ *if for each formula* $\phi(\mathbf{x}, \mathbf{y})$ *there is a formula* $(d_p x)\phi(\mathbf{x}, \mathbf{y})[\mathbf{y}, \mathbf{c}]$ *with free variable* $\mathbf{y}$ *such that* $(d_p \mathbf{x})\phi(\mathbf{x}, \mathbf{y})[\mathbf{m}, \mathbf{c}]$ *holds for exactly those* $\mathbf{m} \in M$ *such that* $\phi(\mathbf{x}, \mathbf{m}) \in p$*. This is a* defining schema *for* p*.*

The following result is asserted without proof (or even explicit mention) in the proof of Lemma 4.2 of [120]. In the proof we expand the language but in a way that does no harm.

Lemma 26.10. *There is an atomic class* $\boldsymbol{K}_1$ *in a vocabulary* τ_1*, whose models are in* 1-1 *correspondence with those of* $\boldsymbol{K}$ *such that: for each* τ_1*-formula* $\phi(\mathbf{x}, \mathbf{y})$ *and countable ordinal* α*, there is a* τ_1*-formula* $P_{\phi,\alpha}(\mathbf{y})$ *such that in any model* M *in* $\boldsymbol{K}_1$*,* $P_{\phi,\alpha}(\mathbf{m})$ *holds if and only if* $R_M(\phi(\mathbf{x}, \mathbf{m})) \geq \alpha$*.*

Proof. Define a sequence of classes and vocabularies $\tau^i, \boldsymbol{K}^i$ by adjoining predicates in τ^{i+1} which define rank for τ^i-formulas. Note that reduct is a 1-1 map from τ^ω structures to τ-structures. Then $\tau^\omega, \boldsymbol{K}^\omega$ are the required $\tau_1, \boldsymbol{K}_1$. $\square_{26.10}$

Henceforth, we assume $\boldsymbol{K}$ satisfies the conclusion of Lemma 26.10.

Lemma 26.11. *Let $\boldsymbol{K}$ be ω-stable. Every type over a model is definable.*

Proof. Let N be an atomic model of T and let $p \in S_{at}(N)$; choose $\phi(\mathbf{x}, \mathbf{c})$ so that $R(p) = R(\phi(\mathbf{x}, \mathbf{c})) = \alpha$. Now for any $\psi(\mathbf{x}, \mathbf{d})$, $\psi(\mathbf{x}, \mathbf{d}) \in p$ if and only if $R(\phi(\mathbf{x}, \mathbf{c}) \wedge \psi(\mathbf{x}, \mathbf{d})) = \alpha$. And, the collection of such $\mathbf{d}$ is defined by $P_{\phi(\mathbf{x},\mathbf{c}) \wedge \psi(\mathbf{x},\mathbf{y}),\alpha}(\mathbf{y}, \mathbf{c})$. $\square_{26.11}$

Note that if p doesn't split over C with $C \subset M \prec N$ and $\hat{p} \in S_{\mathrm{at}}(M)$ is a nonsplitting extension of p, $\hat{p}$ is defined by the same schema as p.

We adapt standard notation from e.g. [94] in our context. Note that we restrict our attention to **big** formulas. This will give us two cardinal transfer theorems that read exactly as those for first order but actually have different content because the first order versions refer to arbitrary infinite definable sets.

Definition 26.12. *1. A triple (M, N, ϕ) where $M \prec N \in \boldsymbol{K}$ with $M \neq N$, ϕ is defined over M, ϕ big, and $\phi(M) = \phi(N)$ is called a* Vaughtian triple.

2. We say $\boldsymbol{K}$ admits (κ, λ), witnessed by ϕ, if there is a model $N \in \boldsymbol{K}$ with $|N| = \kappa$ and $|\phi(N)| = \lambda$ and ϕ is big.

Of course, it is easy in this context to have definable sets which are countable in all models. But we'll show that this is really the only sense in which excellent classes differ from first order stable theories as far as two cardinal theorems are concerned.

The overall structure of the proof of the next result is based on Proposition 2.21 of [100]; but in the crucial type-omitting step we expand the argument of Theorem IX.5.13 in [8] rather than introducing nonorthogonality arguments at this stage.

Lemma 26.13. *Suppose $\boldsymbol{K}$ is *-excellent.*

1. If $\boldsymbol{K}$ admits (κ, λ) for some $\kappa > \lambda$ then $\boldsymbol{K}$ has a Vaughtian triple.

2. If $\boldsymbol{K}$ has a Vaughtian triple, for any (κ', λ') with $\kappa' > \lambda'$, $\boldsymbol{K}$ admits (κ', λ').

Proof. 1) Suppose $N \in \boldsymbol{K}$ with $|N| = \kappa$ and $|\phi(N)| = \lambda$. For notational simplicity we add the parameters of ϕ to the language. By Löwenheim-Skolem, we can embed $\phi(N)$ in a proper elementary submodel M and get a Vaughtian triple. 2) We may assume that M and N are countable. To see this, build within the given M, N countable increasing sequences of countable models M_i, N_i, fixing one element $b \in N - M$ to be in N_0 and choosing $M_i \prec M$, $N_i \prec N$, $M_i \prec N_i$ and $\phi(N_i) \subset \phi(M_{i+1})$. Then M_ω, N_ω are as required.

Now for any κ', we will construct a (κ', ω) model. Say $b \in N - M$ and let $q = \mathrm{tp}(b/M)$. Now repeatedly apply Lemma 25.10 to construct N_i for $i < \kappa'$ so that N_{i+1} is primary over the $N_i b_i$ where b_i realizes the non-splitting extension of q to $S_{\mathrm{at}}(N_i)$. Fix finite C contained in M so that q does not split over C. We prove by induction that each $\phi(N_i) = \phi(M)$. Suppose this holds for i, but there is an $e \in \phi(N_{i+1}) - \phi(M)$. Using Exercise 19.22, fix $\mathbf{m} \in N_i$ and $\theta(x, \mathbf{z}, y)$ such that $\theta(b_i, \mathbf{m}, y)$ isolates $\mathrm{tp}(e/N_i b_i)$. We will obtain a contradiction.

For every $\mathbf{n} \in N_i$, if $N_i \models (\exists y)(\theta(b_i, \mathbf{n}, y) \wedge \phi(y))$ then for some $d \in M$, $\theta(b_i, \mathbf{n}, d) \wedge \phi(d)$ holds, since $\phi(M) = \phi(N_i)$. Thus,

$$(\forall \mathbf{z})[(d_q x)((\exists y)\theta(x, \mathbf{z}, y) \wedge \phi(y))[\mathbf{z}, \mathbf{c}] \rightarrow (\exists y)\phi(y) \wedge (d_q x)\theta(x, \mathbf{z}, y)[\mathbf{z}, y, \mathbf{c}]].$$

We have $\theta(b_i, \mathbf{m}, e)$, so $N_i \models (d_q x)((\exists y)\theta(x, \mathbf{z}, y) \wedge \phi(y))[\mathbf{m}, \mathbf{c}]$. Thus by the displayed formula, $N_i \models (\exists y)\phi(y) \wedge (d_q x)\theta(x, \mathbf{z}, y)[\mathbf{m}, y, \mathbf{c}]$. That is, for some $d \in M$, $N_i \models (d_q x)(\theta(x, \mathbf{z}, y))[\mathbf{m}, d, \mathbf{c}]$. Since $\mathrm{tp}(b_i/N_i)$ is defined by d_q, we have $\theta(b_i, \mathbf{m}, d)$. But this contradicts the fact that $\theta(b_i, \mathbf{m}, y)$ isolates $\mathrm{tp}(e/N_i b_i)$m, as $\mathrm{tp}(e/N_i b_i) \models y \neq d$.

Thus, we have constructed a model $N_{\kappa'}$ of power κ' where ϕ is satisfied only countably many times. To construct a (κ', λ') model, iteratively realize the non-splitting extension of ϕ, λ' times. $\square_{26.13}$

We need one further corollary of Theorem 25.10. The details of the argument require several technical remarks. $\boldsymbol{a} \underset{M}{\downarrow} \mathbf{b}$

Lemma 26.14. *Suppose N is prime over $M\boldsymbol{a}$ and $\mathbf{d} \in N - M$. Then $\boldsymbol{a} \underset{M}{\downarrow} M\mathbf{d}$*

Proof. For some $\mathbf{c} \in M$, there is a formula $\phi(\mathbf{c}, \boldsymbol{a}, \mathbf{x})$, satisfied by $\mathbf{d}$, which implies $\mathrm{tp}(\mathbf{d}/M\boldsymbol{a})$. In particular, $\phi(\mathbf{c}, \boldsymbol{a}, \mathbf{x}) \rightarrow \mathbf{x} \neq \mathbf{m}$ for any $\mathbf{m} \in M$. Now let C be any finite subset of M, which contains $\mathbf{c}$. We show $\mathrm{tp}(\boldsymbol{a}/M\mathbf{d})$ splits over C. Namely, choose $\mathbf{d}' \in M$ with $\mathbf{d} \equiv_C \mathbf{d}'$. Then, we must have $\neg\phi(\mathbf{c}, \boldsymbol{a}, \mathbf{d}')$ as required. $\square_{26.14}$

Lemma 26.15. *If $p \in S(M_0)$ is quasiminimal and X is an independent set of realizations of p, there is a primary model over MX.*

Proof. Let $X = \{x_i : i < \lambda\}$. By Theorem 25.10 define M_{i+1} to be primary over $M_i x_i$, taking unions at limits. $\square_{26.15}$

Exercise 26.16. *Use the independence of X, Lemma 25.11, and Exercise 21.6 to verify that for limit δ, M_δ is in fact primary over X_δ.*

Now we conclude that categoricity transfers among uncountable powers for excellent classes. This applies also to sentences of $L_{\omega_1,\omega}$ that are categorical up to $\aleph_\omega$. They can be considered as atomic classes by Theorem 7.1.12 and are excellent by Theorem 23.1.

Theorem 26.17. *Suppose $\boldsymbol{K}$ is *-excellent. The following are equivalent.*

1. *$\boldsymbol{K}$ is categorical in some uncountable cardinality.*
2. *$\boldsymbol{K}$ has no Vaughtian triples.*

3. *$\boldsymbol{K}$ is categorical in every uncountable cardinal.*

Proof. We first show 1) implies 2). Suppose for contradiction that there is a two-cardinal model (M, N, ϕ) even though $\boldsymbol{K}$ is κ-categorical for some uncountable κ. By Theorem 26.13 $\boldsymbol{K}$ has a $(\kappa, \aleph_0)$-model. But by Theorem 22.13, if it is categorical there is a full model in the categoricity cardinal and every big definable subset of a full model has the same cardinality as the model.

3) implies 1) is obvious; it remains to show 2) implies 3). Let M_0 be the unique countable model. By Lemma 26.5, there is a quasiminimal formula $\phi(x, \mathbf{c})$ with parameters from M_0. For any λ, by Theorem 22.13, there is a full model N of $\boldsymbol{K}$ extending M_0 with cardinality λ. By Lemma 26.7, cl is a pregeometry on $\phi(N)$. Note that $\phi(M)$ is closed since by definition any element a of $\mathrm{cl}(\phi(M))$ both satisfies ϕ and is in every model containing $\phi(M)$, including M. Thus we can choose a basis X for $\phi(M)$. By Theorem 26.15, there is a prime model $M_{|X|}$ over MX. But $X \subset \phi(M_{|X|}) \subset \phi(M)$ so $\phi(M_{|X|}) = \phi(M)$; whence since we assume there are no two cardinal models, $M_{|X|} = M$ and M is prime and minimal over MX.

Now we show categoricity in any uncountable cardinality. If M, N are models of power λ, they are each prime and minimal over X, a basis for $\phi(M)$ and Y, a basis for $\phi(N)$, respectively. Now any bijection between X and Y is elementary by the moreover clause in Lemma 26.7. It extends to a map from M into N by primeness and it must be onto; otherwise there is a two cardinal model. $\square_{26.17}$

As in the first order case, we have the following easy corollary.

Corollary 26.18. *Suppose $\boldsymbol{K}$ is *-excellent. If $\boldsymbol{K}$ is not $\aleph_1$-categorical, then $\boldsymbol{K}$ has at least $n+1$ models of cardinality $\aleph_n$ for each $n < \omega$.*

Proof. There is a two-cardinal formula ϕ by Theorem 26.17. By Lemma 26.13 there is a model M_k of cardinality $\aleph_n$, such that $\phi(M_k)$ has cardinality $\aleph_k$ for each $k < n$. $\square_{26.18}$

We have completed the proof of 'the Morley theorem' for a *complete* sentence of $L_{\omega_1,\omega}$. But we want to extend the result to arbitrary sentences and prove:

Theorem 26.19. *(For $n < \omega$, $2^{\aleph_n} < 2^{\aleph_{n+1}}$.)*

1. *If $\phi \in L_{\omega_1,\omega}$ is categorical in $\aleph_n$ for all $n < \omega$, then ϕ is categorical in all uncountable cardinalities.*
2. *If ϕ has very few models in $\aleph_n$ for all $n < \omega$, then*
   - *(a) ϕ has arbitrarily large models and*
   - *(b) If ϕ is categorical in some uncountable κ, it is categorical in all uncountable κ*

Proof. In each case, ϕ has very few models in $\aleph_n$ for each $n < \omega$. By Lemma 7.3.2, choose a small model of ϕ with cardinality $\aleph_1$, say with Scott sentence ψ. Let $\boldsymbol{K}$ be the class of models of ψ. By Theorem 7.1.12, without loss of generality, we can view $\boldsymbol{K}$ as an atomic class. Theorem 23.1 implies $\boldsymbol{K}$ is excellent. By Theorem 25.12, $\boldsymbol{K}$ and thus ψ have arbitrarily large models. Since $\psi \rightarrow \phi$, this proves 2a). But it also gives 1) since the hypothesis of 1) implies ψ is categorical in all powers and if there is a model of ϕ and $\neg\psi$, by Löwenheim-Skolem there is one in $\aleph_1$ contradicting the categoricity of ϕ is $\aleph_1$.

For 2b) suppose ϕ is categorical in $\kappa > \aleph_0$. Then so is ψ whence, by Theorem 26.17, ψ is categorical in all uncountable powers. To show ϕ is categorical above κ note that by downward Löwenheim-Skolem all models of ϕ with cardinality at least κ satisfy ψ; the result follows by the categoricity of ψ. If ϕ is not categorical in some uncountable cardinal $\mu < \kappa$, there must be a sentence θ which is inconsistent with ψ but consistent with ϕ. Applying the entire analysis to $\phi \wedge \theta$, we find a complete sentence ψ' which has arbitrarily large models, is consistent with ϕ and contradicts ψ. But this is forbidden by categoricity in κ. $\square_{26.19}$

We conclude by connecting to Part I, showing that for sentences of $L_{\omega_1,\omega}$, Zilber's notion of quasiminimal excellence is the 'rank one' case of Shelah's theory excellence.

Proposition 26.20. *Suppose $(\boldsymbol{K}, \mathrm{cl}_M)$ is a class of L-structures with a closure relation cl_M satisfying the definition of a quasiminimal excellent class (Chapter 3) and the countable closure condition such that the class $\boldsymbol{K}$ and the closure relation is definable in $L_{\omega_1,\omega}$. Then $\boldsymbol{K}$ is an excellent atomic class and the formula $x = x$ is quasiminimal.*

Proof. Since $\boldsymbol{K}$ is $L_{\omega_1,\omega}$-definable and by Lemma 3.12.2 $\aleph_0$-categorical, the models of $\boldsymbol{K}$ are an atomic class (in some extended language). It is ω-stable by Corollary 3.10. So to show quasiminimal excellence specializes excellence it suffices to show that $a \in \mathrm{cl}(M)$ in the sense of Chapter 3 if and only if $a \not\downarrow M$ in sense of Chapter 20. If $a \in \mathrm{cl}(M)$ then $p = \mathrm{tp}(a/M)$ is realized by only countably many points. But if p does not split over some finite subset of M, it follows easily from Lemma 20.9 that p is realized by arbitrarily many points. But if $a \notin \mathrm{cl}(M)$, then $\mathrm{tp}(a/M)$ is realized uncountably many times; by quasiminimality there is a unique such type and it must be the non-splitting extension. $\square_{26.20}$

Thus, for sentences of $L_{\omega_1,\omega}$, the argument in Chapter 3 specializes the argument in this chapter. Assuming excellence, the first shows a certain geometrical condition (quasiminimality) implies categoricity in all cardinalities; the second shows categoricity in one cardinality implies the models are controlled by a set satisfying the geometric condition and thus categoricity in all powers. But there is another sense in which the argument of Part I is stronger than what is proved in Part IV. Part IV is formulated explicitly for atomic ($L_{\omega_1,\omega}$) classes. The argument in Part I applies to certain classes that are formulated in $L_{\omega_1,\omega}(Q)$ (most obviously, just say the closure is countable). Shelah's argument undoubtedly can be extended in a similar direction but the precise formulation of what can be obtained by these arguments has not been spelled out. Shelah's work on good frames (e.g.

[131]) aims to obtain categoricity transfer for $L_{\omega_1,\omega}(Q)$ but by more complicated methods than here.

27
Demystifying Non-excellence

We have shown (under weak GCH) that categoricity up to $\aleph_\omega$ of a sentence in $L_{\omega_1,\omega}$ implies categoricity in all uncountable cardinalities. Hart and Shelah [56] showed the necessity of the assumption by constructing sentences ϕ_k which were categorical up to some $\aleph_n$ but not eventually categorical. Since $(\aleph_0, \infty)$-tame classes with amalgamation are eventually categorical if they are categorical in one power, these examples are a natural location to look for classes that are categorical but not tame. Relying on [5], we present Kolesnikov's simplification of the example, correct some minor inaccuracies in [56], and provide examples of non-tameness. In particular, this paper answered a question posed by Shelah in [132] by specifying that categoricity fails exactly at $\aleph_{k-1}$; the proof is easy by combining the analysis in Parts III and IV of this monograph. We outline here the proof of the following theorem, referring to [5] for many of the technical arguments.

Theorem 27.1. *For each $2 \le k < \omega$ there is an $L_{\omega_1,\omega}$-sentence ϕ_k such that:*

1. *ϕ_k is categorical in μ if $\mu \le \aleph_{k-2}$;*
2. *ϕ_k is not $\aleph_{k-2}$-Galois stable;*
3. *ϕ_k is not categorical in any μ with $\mu > \aleph_{k-2}$;*
4. *ϕ_k has the disjoint amalgamation property;*
5. *For $k > 2$,*
   (a) *ϕ_k is $(\aleph_0, \aleph_{k-3})$-tame; indeed, syntactic first-order types determine Galois types over models of cardinality at most $\aleph_{k-3}$;*

(b) ϕ_k *is* $\aleph_m$*-Galois stable for* $m \leq k-3$*;*

(c) ϕ_k *is not* $(\aleph_{k-3}, \aleph_{k-2})$*-tame.*

In Section 27.1 we describe the example and define the sentences ϕ_k. In Section 27.2 we introduce the notion of a solution and prove lemmas about the amalgamation of solutions. From these we deduce in Section 27.4 positive results about tameness. In some sense, the key insight of [5] is that the amalgamation property holds in all cardinalities (Section 27.3) while the amalgamation of solutions is very cardinal dependent. We show in Section 27.5 that ϕ_k is not Galois stable in $\aleph_{k-2}$, hence not categorical above $\aleph_{k-2}$, and deduce the non-tameness.

27.1 The basic structure

This example is a descendent of the example in [15] of an $\aleph_1$-categorical theory which is not almost strongly minimal. That is, the universe is not in the algebraic closure of a strongly minimal set. Here is a simple way to describe such a model.

Example 27.1.1. Let G be a strongly minimal group and let π map X onto G. Add to the language a binary function $t : G \times X \to X$ for the fixed-point free action of G on $\pi^{-1}(g)$ for each $g \in G$. That is, represent $\pi^{-1}(g)$ as $\{ga : g \in G\}$ for some a with $\pi(a) = g$. This action of G is strictly 1-transitive. This guarantees that each fiber has the same cardinality as G and π guarantees the number of fibers is the same as $|G|$. Since there is no interaction among the fibers, categoricity in all uncountable powers is easy to check.

The next step in complexity combines the 'affine' idea with that of Example 3.30.

Example 27.1.2. Let G be the group of functions from I into $\mathcal{Z}_2$ (with evaluation) as in Example 3.30. Let G^* be an affine copy of G. That is, G^* is a copy of G but there is no addition on G^*. Rather, there is a function t_G mapping $G \times G^*$ to G^* by $t_G(g, h) = j$, just if j is the coordinate-wise sum of g and h. We will just write $h + g = j$. Now, as we work out in detail in a more complicated situation below, we have an example of totally categorical sentence in $L_{\omega_1,\omega}$. The current example is both excellent and homogeneous. (After naming a constant, each model is in the algebraic closure of a quasiminal excellent homogeneous set.) However, the notion of splitting does not behave well over arbitrary sets. we will show that for any $a \in G^*$, and any model M, $\mathrm{tp}(a/M)$ splits over $\emptyset$. Indeed, choose $c, d \in M \cap G^*$ that realize the same type over the $\emptyset$. For any g, we write S_g for the support of g. Since t_G codes a regular action, for any $a, b \in G^*$, there is a unique $g \in G$ with $a + g = b$.

Now choose

1. $g \in G$ such that $a = c + g$;

2. $h \in G \cap M$ such that $c = d + h$;

3. $h_0 \in G \cap M$ such that $S_g \cap S_h = S_{h_0}$.

Now let F fix $G \cup I \cup \mathcal{Z}_2$ and define F on G^* by $F(x) = x + h_0$. It is easy to check that F is an automorphism of the monster that fixes M setwise. If $\mathrm{tp}(a/M)$ does not split over $\emptyset$ then also $\mathrm{tp}(F(a)/M)$ does not split over $\emptyset$. But from the choices above, we see $F(a) = c + (g + h_0)$ and also $F(a) = d + (g + h_0 + h)$. A short computation shows $g + h_0$ and h have disjoint support. Thus, since h is not identically 0, the supports of $g + h_0$ and $g + h_0 + h$ have different cardinalities. As F fixes M, the formula $\theta(a, c)$ which specifies the cardinality of the support of the element x with $a = c + x$ does not hold of a and d.

This particular version of the example is due to Kolesnikov, simplifying ideas of Hyttinen and Kesala. In [64], Hyttinen and Kesala show that there is no good notion of independence over arbitrary sets for this example.

Now if we combine Examples 27.1.1 and 27.1.2, we will get an example that is not almost quasiminimal in the sense of Definition 3.28.

Notation 27.1.3. The formal language for the current example contains unary predicates I, K, G, G^*, H, H^*; a binary function e_G taking $G \times K$ to H; a function π_G mapping G^* to K, a function π_H mapping H^* to K, a 4-ary relation t_G on $K \times G \times G^* \times G^*$, a 4-ary relation t_H on $K \times H \times H^* \times H^*$. Certain other projection functions are in the language but not expressly described. These symbols form a vocabulary L'; we form the vocabulary L by adding a $k+1$-ary relation Q on $(G^*)^k \times H^*$, which will be explained in due course.

We start by describing the L'-structure $M(I)$ constructed from any set I with at least k elements. We will see that the L'-structure is completely determined by the cardinality of I. So we need to work harder to get failure of categoricity, and this will be the role of the predicate Q.

The structure $M(I)$ is a disjoint union of sets I, K, H, G, G^* and H^*. Let $K = [I]^k$ be the set of k-element subsets of I. H is a single copy of Z_2. Let G be the direct sum of K copies of Z_2. So G, K, and I have the same cardinality. We include K, G, and Z_2 as sorts of the structure with the evaluation function e_G: for $\gamma \in G$ and $k \in K$, $e_G(\gamma, k) = \gamma(k) \in Z_2$. So in $L'_{\omega_1,\omega}$ we can say that the predicate G denotes exactly the set of elements with finite support of ${}^K Z_2$.

Now, we introduce the sets G^* and H^*. We have a projection function π_G from G^* onto K. Thus, for $u \in K$, we can represent the elements of $\pi_G^{-1}(u)$ in the form $(u, x) \in G^*$; or alternatively, as $x \in G^*_u$. We refer to the set $\pi_G^{-1}(u)$ as the *G^*-stalk, or fiber over u.* Then we encode the affine action by the relation $t_G(u, \gamma, x, y) \subset K \times G \times G^* \times G^*$ which is the graph of a regular transitive action of G on G^*_u. (Of course, this can be expressed in $L'_{\omega,\omega}$). That is, for all $x = (u, x')$, $y = (u, y')$ there is a unique $\gamma \in G$ such that $t_G(u, \gamma, x, y)$ holds.

As a set, $H^* = K \times Z_2$. As before if $\pi_H(x) = v$ holds x has the form (v, x'), and we denote by H^*_v the preimage $\pi_H^{-1}(v)$. Finally, for each $v \in K$, $t_H(v, \delta, x, y) \subset K \times Z_2 \times H^* \times H^*$ is the graph of a regular transitive action of Z_2 on the stalk H^*_v.

$(*)$: Use additive notation for the action of G (of H) on the stalks of G^* (H^*).

1. For $\gamma \in G$, denote the action by $y = x + \gamma$ whenever it is clear that x and y come from the same G^*-stalk. It is also convenient to denote by $y - x$ the unique element $\gamma \in G$ such that $y = \gamma + x$.
2. For $\delta \in H$, denote the action by $y = x + \delta$, whenever it is clear that x and y come from the same H^*-stalk. Say that $\delta = y - x$.

Let ψ_k^1 be the Scott sentence for the L'-structure that we have described so far. This much of the structure is clearly categorical (and homogeneous). Indeed, suppose two such models have been built on I and I' of the same cardinality. Take any bijection between I and I'. To extend the map to G^* and H^*, fix one element in each partition class (stalk) in each model. The natural correspondence (linking those selected in corresponding classes) extends to an isomorphism. Thus we may work with a canonical L'-model; namely with the model that has copies of G (without the group structure) as the stalks G_u^* and copies of Z_2 (also without the group structure) as the stalks H_v^*. The functions t_G and t_H impose an affine structure on the stalks.

Notation 27.1.4. *The L-structure is imposed by a $(k+1)$-ary relation Q on $(G^*)^k \times H^*$, which has a local character. We will use only the following list of properties of Q, which are easily axiomatized in $L_{\omega_1,\omega}$:*

1. *Q is symmetric, with respect to all permutations, for the k elements from G^*;*
2. *$Q((u_1, x_1), \ldots, (u_k, x_k), (u_{k+1}, x_{k+1}))$ implies that $u_1, \ldots, u_{k+1}$ form all the k element subsets of a $k+1$ element subset of I. We call $u_1, \ldots, u_{k+1}$ a* compatible $(k+1)$-tuple*;*
3. *using the notation introduced at (*), Q is related to the actions t_G and t_H as follows:*
   (a) for all $\gamma \in G$, $\delta \in H$
   $$\begin{aligned} &Q((u_1, x_1), \ldots, (u_k, x_k), (u_{k+1}, x_{k+1})) \\ \Leftrightarrow \neg &Q((u_1, x_1 + \gamma), \ldots, (u_k, x_k), (u_{k+1}, x_{k+1})) \end{aligned}$$
   if and only if $\gamma(u_{k+1}) = 1$;
   (b)
   $$\begin{aligned} &Q((u_1, x_1), \ldots, (u_k, x_k), (u_{k+1}, x_{k+1})) \\ \Leftrightarrow \neg &Q((u_1, x_1), \ldots, (u_k, x_k), (u_{k+1}, x_{k+1} + \delta)) \end{aligned}$$
   if and only if $\delta = 1$.

Let ψ_k^2 be the conjunction of sentences expressing (1)–(3) above, and we let $\phi_k := \psi_k^1 \wedge \psi_k^2$.

It remains to show that such an expansion to $L = L' \cup \{Q\}$ exists. We do this by explicitly showing how to define Q on the canonical L'-structure. In fact, we describe $2^{|I|}$ such structures parameterized by functions ℓ.

Fact 27.1.5. *Let M be an L'-structure satisfying ψ^1_k. Let $I := I(M)$ and $K := K(M)$. Let $\ell\colon I \times K \to 2$ be an arbitrary function.*

For each compatible $k+1$ tuple $u_1, \ldots, u_{k+1}$, such that $u_1 \cup \ldots \cup u_{k+1} = \{a\} \cup u_{k+1}$ for some $a \in I$ and $u_{k+1} \in K$, define an expansion of M to L by

$$M \models Q((u_1, x_1), \ldots, (u_k, x_k), (u_{k+1}, x_{k+1}))$$

if and only if $x_1(u_{k+1}) + \ldots + x_k(u_{k+1}) + x_{k+1} = \ell(a, u_{k+1}) \mod 2$. Then M is a model of ϕ_k.

Indeed, it is straightforward to check that the expanded structure M satisfies ψ^2_k. We describe the interaction of G and Q a bit more fully. Using symmetry in the first k components, we obtain the following.

Fact 27.1.6. *For all $\gamma_1, \ldots, \gamma_k \in G$ and all $\delta \in H$ we have*

$$\begin{aligned} &Q((u_1, x_1), \ldots, (u_k, x_k), (u_{k+1}, x_{k+1})) \\ \Leftrightarrow\ &Q((u_1, x_1 + \gamma_1), \ldots, (u_k, x_k + \gamma_k), (u_{k+1}, x_{k+1} + \delta)) \end{aligned}$$

if and only if $\gamma_1(u_{k+1}) + \ldots + \gamma_k(u_{k+1}) + \delta = 0 \mod 2$.

In the next section, we show that ϕ_k is categorical in $\aleph_0, \ldots, \aleph_{k-2}$. So in particular ϕ_k is a complete sentence for all k.

Now we obtain abstract elementary classes $(\boldsymbol{K}_k, \prec_{\boldsymbol{K}})$ where $\boldsymbol{K}_k$ is the class of models of ϕ_k and for $M, N \models \phi_k$, $M \prec_{\boldsymbol{K}} N$ if M is a substructure of N. This is a stronger requirement than it seems; see Section 4 of [5]. Note that if $M \prec_{\boldsymbol{K}} N$ and $g \in G(M)$, since G is a group, the support of g in M equals the support of g in N. This allows us to verify that $(\boldsymbol{K}_k, \prec_{\boldsymbol{K}})$ is closed under unions of chains and satisfies the axioms for an AEC. In fact, a rather more detailed analysis in [5] shows that if $M \subset N$, where $M, N \models \phi_k$, then $M \prec_{L_{\omega_1,\omega}} N$.

27.2 Solutions and categoricity

As we saw at the end of the previous section, the predicate Q can be defined in somewhat arbitrary way. Showing categoricity of the L-structure amounts to showing that any model M, of an appropriate cardinality, is isomorphic to the model where all the values of ℓ are chosen to be zero; we call such a model a *standard model*. This motivates the following definition:

Definition 27.2.1. *Fix a model M. A* solution *for M is a selector f that chooses (in a compatible way) one element of the fiber in G^* above each element of K and one element of the fiber in H^* above each element of K. Formally, f is a pair of functions (g, h), where $g : K(M) \to G^*(M)$ and $h : K(M) \to H^*(M)$ such that $\pi_G g$ and $\pi_H h$ are the identity and for each compatible $(k+1)$ tuple $u_1, \ldots, u_{k+1}$:*

$$Q(g(u_1), \ldots, g(u_k), h(u_{k+1})).$$

Notation 27.2.2. *As usual* $k = \{0, 1, \dots k-1\}$ *and we write* $[A]^k$ *for the set of* k*-element subsets of* A.

We will show momentarily that if M and N have the same cardinality and have solutions f_M and f_N then $M \cong N$. Thus, in order to establish categoricity of ϕ_k in $\aleph_0, \dots, \aleph_{k-2}$, it suffices to find a solutions of each cardinality up to $\aleph_{k-2}$ in an arbitrary model of ϕ_k. Our approach is to build up the solutions in stages, for which we need to describe selectors over subsets of $I(M)$ (or of $K(M)$) rather than all of $I(M)$.

Definition 27.2.3. *We say that* (g, h) is a solution for the subset W of $K(M)$ *if for each* $u \in W$ *there are* $g(u) \in G^*_u$ *and* $h(u) \in H^*_u$ *such that if* $u_1, \dots, u_k, u_{k+1}$ *are a compatible* $k+1$ *tuple from* W*, then*

$$Q(g(u_1), \dots, g(u_k), h(u_{k+1})).$$

If (g, h) *is a solution for the set* W*, where* $W = [A]^k$ *for some* $A \subset I(M)$*, we say that* (g, h) *is a* solution over A.

Definition 27.2.4. *The models of* ϕ_k *have the* extension property for solutions *over sets of size* λ *if for every* $M \models \phi_k$*, any solution* (g, h) *over a set* A *with* $|A| = \lambda$*, and every* $a \in I(M) \setminus A$ *there is a solution* (g', h') *over* $A \cup \{a\}$*, extending* (g, h).

One can treat the element $g(u)$ as the image of the element $(u, 0)$ under the isomorphism between the standard model and M, where 0 represents the constantly zero function in the stalk G^*_u. Not surprisingly, we have the following:

Lemma 27.2.5. *If* M *and* N *have the same cardinality and have solutions* f_M *and* f_N *then* $M \cong N$.

Moreover, suppose $\boldsymbol{K}$ *has solutions and has extension of solutions for models of cardinality less than* $|M|$*; if* g *is an isomorphism between* L*-substructures* M', N' *of* M *and* N*, then the isomorphism* $\hat{g}$ *between* M *and* N *can be chosen to extend* g*. Finally, if* $f_{M'}$ *is a solution on* M' *which extends to a solution* f_M *on* M*, then* $\hat{g}$ *maps them to a similar extending pair on* N' *and* N.

Proof. We prove the 'moreover' clause; the first statement is a special case when g is empty and the 'finally' is included in the proof. Say, g maps M' to N'. Without loss of generality, $M \restriction L' = M(I), N \restriction L' = M(I')$. Let α be a bijection between I and I' which extends $g \restriction I$. Extend naturally to a map from $K(M)$ to $K(N)$ and from $G(M)$ to $G(N)$, which extends g on M'. By assumption there is a solution $f_{M'}$ on M'. It is clear that g maps $f_{M'}$ to a solution $f_{N'}$ on N'; by assumption $f_{N'}$ extends to a solution on N. (Note that if we do not have to worry about g, we let α be an arbitrary bijection from I to I' and let $\alpha(f_M(u))$ be $f_N(\alpha(u))$.) For $x \in G^*(M - M')$ such that $M \models \pi_G(x) = u$, there is a unique $a \in G(M)$ with $a = x - f_M(u)$ (the operation makes sense because a and $f_M(u)$ are in the same stalk).

Let $\alpha(x)$ be the unique $y \in N - N'$ such that

$$N \models t(\alpha(u), \alpha(a), f_N(\alpha(u)), y)$$

i.e., $y = \alpha(a) + f_N(\alpha(u))$ in the stalk $G^*_{\alpha(u)}(N)$.

Do a similar construction for H^* and observe that Q is preserved. $\square_{27.2.5}$

We prove the case $k = 2$.

Claim 27.2.6. *The models of ϕ_2 have the extension property for solutions over finite sets.*

Proof. Let $A := \{a_0, \ldots, a_{n-1}\}$ and (g, h) be a solution over A. For each $v = \{a, a_i\}$, let y_v be an arbitrary element of H^*_v. Now extend h to the function h' with domain $[A \cup \{a\}]^2$ by defining $h'(v) := y_v$.

It remains to define the function g' on each $\{a, a_i\}$, and we do it by induction on i. For $i = 0$, pick an arbitrary element $x \in G^*_{a,a_0}$. Let $\gamma_0 \in G$ be such that for $j = 1, \ldots, n-1$

$$\gamma_0(a, a_j) = 1 \text{ if and only if } M \models \neg Q((\{a, a_0\}, x), g(a_0, a_j), h'(a, a_j)).$$

It is clear that letting $g'(\{a, a_0\}) := (\{a, a_0\}, x + \gamma_0)$, we have a partial solution.

Suppose that $g'(\{a, a_j\})$, $j < i$, have been defined. Pick an arbitrary element $x \in G^*_{a,a_i}$. Let $\gamma_i \in G$ be such that for $j \in \{0, \ldots, n-1\} \setminus \{i\}$

$$\gamma_i(a, a_j) = 1 \text{ if and only if } M \models \neg Q((\{a, a_i\}, x), g(a_i, a_j), h'(a, a_j)).$$

Also let $\gamma'_i \in G$ be such that for $j < i$

$$\gamma'_i(a_i, a_j) = 1 \text{ if and only if } M \models \neg Q((\{a, a_j\}, x), g'(a, a_j), h(a_i, a_j)).$$

Now letting $g'(\{a, a_i\}) := (\{a, a_i\}, x + \gamma_i + \gamma'_i)$ yields a solution on $A \cup \{a\}$. $\square_{27.2.6}$

Corollary 27.2.7. *The sentence ϕ_2 is $\aleph_0$-categorical, and hence is a complete sentence.*

Proof. Let M be a countable model. Enumerate $I(M)$ as $\{a_i : i < \omega\}$. It is clear that a solution exists over the set $\{a_0, a_1\}$ (any elements in stalks $G^*_{a_0,a_1}$ and $H^*_{a_0,a_1}$ work). By the extension property for solutions over finite sets we get a solution defined over the entire $I(M)$. Hence ϕ_2 is countably categorical by Lemma 27.2.5. $\square_{27.2.7}$

We see that extension for solutions over finite sets translates into existence of solutions over countable sets. This is part of a general phenomenon that we describe below. For the general case $k \geq 2$; we state a couple of definitions and the main result but refer to [5] for the proofs.

Definition 27.2.8. *Let A be a subset of I of size λ, and consider an arbitrary n-element set $\{b_0, \ldots, b_{n-1}\} \subset I$. Suppose that, for each $(n-1)$-element subset w of $n = \{0, \ldots, n-1\}$, we have a solution (g_w, h_w) over $A \cup \{b_l : l \in w\}$ such that the solutions are compatible (i.e., $(\bigcup_w g_w, \bigcup_w h_w)$ is a function).*

We say that M has n-amalgamation for solutions over sets of size λ *if for every such set A, there is a solution (g, h) over $A \cup \{b_0, \ldots, b_{n-1}\}$ that simultaneously extends all the given solutions $\{(g_w, h_w) : w \in [n]^{n-1}\}$.*

For $n = 0$ the given system of solutions is empty, thus 0-amalgamation over sets of size λ is existence for solutions over sets of size λ. For $n = 1$, the initial system of solutions degenerates to just $(g_\emptyset, h_\emptyset)$, a solution on A; so the 1-amalgamation property corresponds to the extension property for solutions.

Remark 27.2.9. *Immediately from the definition we see that n-amalgamation for solutions over sets of a certain size implies m-amalgamation for solutions over sets of the same size for any $m < n$. Indeed, we can obtain m-amalgamation by putting $n - m$ elements of the set $\{b_0, \dots, b_{n-1}\}$ inside A.*

Again, the long argument for the next Lemma appears in Section 2 of [5].

Lemma 27.2.10. *The models of ϕ_k have the $(k-1)$-amalgamation property for solutions over finite sets.*

The next result is a rather simpler induction [5].

Lemma 27.2.11. *Let $M \models \phi_k$ for some $k \geq 2$ and let $n \leq k - 2$. If M has $(n+1)$-amalgamation for solutions over sets of size less than λ, then M has n-amalgamation for solutions over sets of size λ.*

Corollary 27.2.12. *Every model of ϕ_k of cardinality at most $\aleph_{k-2}$ admits a solution. Thus, the sentence ϕ_k is categorical in $\aleph_0, \dots, \aleph_{k-2}$.*

Proof. Let $M \models \phi_k$. By Lemma 27.2.10, M has $(k-1)$-amalgamation for solutions over finite sets. So M has $(k-2)$-amalgamation for solutions over countable sets, $(k-3)$-amalgamation for solutions over sets of size $\aleph_1$, and so on until we reach 0-amalgamation for solutions over sets of size $\aleph_{k-2}$ Since for $m < n$ and any λ, the n-amalgamation property for solutions over sets of cardinality λ implies m-amalgamation solutions over sets of cardinality λ, we have 0-amalgamation, that is, existence of solutions for sets of size up to and including $\aleph_{k-2}$. Now Lemma 27.2.5 gives categoricity in $\aleph_0, \dots, \aleph_{k-2}$. $\square_{27.2.12}$

Corollary 27.2.13. *For all $k \geq 2$, the sentence ϕ_k is complete.*

The following further corollary is used to establish tameness.

Corollary 27.2.14. *Let $M \models \phi_k$ for some $k \geq 2$ and $n \leq k-2$. Suppose M has 2-amalgamation for solutions over sets of cardinality λ . If $A_0 \subset A_1, A_2 \subset M$ have cardinality λ and $(g^1, h^1), (g^2, h^2)$ are solutions of A_1, A_2 respectively that agree on A, there is a solution (g, h) on $A_1 \cup A_2$ extending both of them.*

27.3 Disjoint amalgamation for models of ϕ_k

In contrast to the previous section, where we studied amalgamation properties of solutions, this section is about (the usual) amalgamation property for the class of models of ϕ_k. The amalgamation property is a significant assumption for the behavior and even the precise definition of Galois types, so it is important to establish that the class of models of our ϕ_k has it.

We claim that the class has the disjoint amalgamation property (Definition 5.10)in every cardinality.

Theorem 27.3.1 (Kolesnikov)**.** *Fix $k \geq 2$. The class of models of ϕ_k has the disjoint amalgamation property.*

Proof. Let $M_i = M_i(I_i)$, $i = 0, 1, 2$, where of course $I_0 \subset I_1, I_2$; K_0, K_1, K_2 are the associated sets of k-tuples. We may assume that $I_1 \cap I_2 = I_0$. Otherwise take a copy I_2' of $I_2 \setminus I_0$ disjoint from I_1, and build a structure M_2' isomorphic to M_2 on $I_0 \cup I_2'$.

We are building a model $M \models \phi_k$ on the set $I_1 \cup I_2$ making sure that it is a model of ϕ_k and that it embeds M_1 and M_2, where the embeddings agree over M_0. We start by building the L'-structure on $I_1 \cup I_2$. So let $I = I(M) := I_0 \cup I_2$; the set $K = [I]^k$ can be thought of as $K_1 \cup K_2 \cup \partial K$, where ∂K consists of the new k-tuples.

Let G be the direct sum of K copies of Z_2, notice that it embeds $G(M_1)$ and $G(M_2)$ in the natural way over $G(M_0)$. We will assume that the embeddings are identity embeddings.

Let G^* be the set of K many affine copies of G, with the action by G and projection to K defined in the natural way. Let H^* be the set of K many affine copies of Z_2, again with the action by Z_2 and the projection onto K naturally defined.

For $i = 1, 2$, we now describe the embeddings f_i of $G^*(M_i)$ and $H^*(M_i)$ into G^* and H^*. Later, we will define the predicate Q on M in such a way that f_i become embeddings of L-structures.

For each $u \in K_0$, choose arbitrarily an element $x_u \in G^*_u(M_0)$. Now for each $x' \in G^*_u(M_1)$, let γ be the unique element in $G(M_1)$ with $x' = x_u + \gamma$. Let $f_1(x') := (u, \gamma)$. Similarly, for each $x' \in G^*_u(M_2)$, let $\delta \in G(M_2)$ be the element with $x' = x_u + \delta$. Define $f_2(x') := (u, \delta)$. Note that the functions agree over $G^*_u(M_0)$: if $x' \in G^*_u(M_0)$, then the element $\gamma = x' - x_u$ is in $G(M_0)$. In particular, $f_1(x_u) = f_2(x_u) = 0$, the constantly zero function.

For each $u \in K_i \setminus K_0$, $i = 1, 2$, choose an arbitrary $x_u \in G^*_u(M_i)$, and for each $x' \in G^*_u(M_i)$ define $f_i(x') := (u, y - x_u)$. This defines the embeddings $f_i : G^*(M_i) \to G^*(M)$.

Embedding $H^*(M_i)$ into $H^*(M)$ is even easier: for each $v \in K_1$, pick an arbitrary $y_v \in H^*_v(M_1)$, and let $f_1(y_v) := (v, 0)$, $f_1(y_v + 1) := (v, 1)$. For each $v \in K_2$, if $v \in K_1$, define f_2 to agree with f_1. Otherwise choose an arbitrary $y_v \in H^*_v(M_2)$, and let $f_2(y_v) := (v, 0)$, $f_2(y_v + 1) := (v, 1)$.

This completes the construction of the disjoint amalgam for L'-structures. Now we define Q on the structure M so that f_i, $i = 1, 2$ become L-embeddings. The expansion is described in terms of the function ℓ that we discussed in Fact 27.1.5.

Let $u_1, \ldots, u_k, v$ be a compatible $k+1$ tuple of elements of K; $u_1 \cup \ldots \cup u_k \cup v = \{a\} \cup v$ for some $a \in I$.

Case 1. $u_1, \ldots, u_k, v \in K_1$ (or all in K_2). This is the most restrictive case. Each of the stalks $G^*_{u_i}(M_1)$ contains an element x_{u_i} defined at the previous stage; and the stalk H_v^* has the element $y_v \in M_1$. Define

$$\ell(a,v) := 0 \text{ if } M_1 \models Q((u_1, x_{u_1}), \ldots, (u_k, x_{u_k}), (v, y_v)),$$

and $\ell(a,v) := 1$ otherwise.

Case 2. At least one of the $u_1, \ldots, u_k, v$ is in ∂K. Then the predicate Q has not been defined on these $k+1$ stalks, and we have the freedom to define it in any way. So choose $\ell(a,v) := 0$ for all such compatible $k+1$ tuples.

Now define Q on M from the function ℓ as in Fact 27.1.5.

It is straightforward to check that f_1 and f_2 become L-embeddings into the L-structure M that we have built. $\square_{27.3.1}$

27.4 Tameness

Here we study the tameness properties for models of ϕ_k. We know that ϕ_k is categorical up to $\aleph_{k-2}$; so without loss of generality we may deal with the standard models of ϕ_k in powers $\aleph_0, \ldots, \aleph_{k-2}$.

In Section 27.5 we establish that ϕ_2 has continuum Galois types over a countable model; and that ϕ_3 is not $(\aleph_0, \aleph_1)$-tame. The first index where some tameness appears is $k = 4$.

Notation 27.4.1. *Let $M_0 \subset M \models \phi_k$. If $a \in M - M_0$ by a* submodel generated by *$M_0 \cup \boldsymbol{a}$, denoted M_0^a, we mean a structure constructed as follows. First take the definable closure of $M_0 \cup \boldsymbol{a}$ to obtain a set X. Then form X' by for any $u \in K(X)$ such that $G_u^*(M) \cap X$ is empty adding a single element from the fiber. Finally, take the definable closure of X'.*

In [5], this construction is refined to define 'full' submodels and minimal full submodels. That allows us to show the syntactic type described below is in fact existential and the example in 'model complete' in a precise generalization of the classical notion.

Lemma 27.4.2. *Suppose $M_0 \subset M \models \phi_k$ and $|M_0| \leq \aleph_{k-3}$. If $a, b \in M - M_0$ realize the same first order syntactic type over M_0 then there is an isomorphism f between M_0^a and M_0^b, fixing M_0 and mapping $\boldsymbol{a}$ to* $\mathbf{b}$.

Proof. Since a and b realize the same syntactic type the L-structures with universe $\mathrm{dcl}(M_0\boldsymbol{a})$ and $\mathrm{dcl}(M_0\mathbf{b})$ are isomorphic over M_0. Since $|M_0| \leq \aleph_{k-3}$, extension of solutions holds for models of cardinality $|M_0|$. We finish by applying the moreover clause of Lemma 27.2.5 to M_0, M_0^a and M_0^b. $\square_{27.4.2}$

It follows immediately that we can strengthen the hypothesis of the last lemma to $\boldsymbol{a}$ and $\mathbf{b}$ realize the same Galois type. We cannot go higher than $\aleph_{k-3}$ in Lemma 27.4.4 because we need the extension property for solutions, which we can only establish for models of size up to $\aleph_{k-3}$, to prove Lemma 27.4.2.

Lemma 27.4.3. *Let $k \geq 4$ and $\aleph_0 \leq \lambda \leq \aleph_{k-4}$. Then the class of models of ϕ_k is (λ, λ^+)-tame.*

Proof. Let M be a model of cardinality λ^+; and let $\boldsymbol{a}$, $\mathbf{b}$ have the same Galois types over all submodels of M of cardinality λ. By the disjoint amalgamation property, we may assume that M, $\boldsymbol{a}$, and $\mathbf{b}$ are inside some model N. Let $M_0 \prec M$ be of power λ; and let $\{M_i : i < \lambda^+\}$ be an increasing continuous chain of models beginning with M_0 and with union M. By Lemma 27.4.2 there is an isomorphism f_0 between M_0^a and M_0^b, fixing M_0 and mapping $\boldsymbol{a}$ to $\mathbf{b}$.

Let $\langle g_0, h_0 \rangle$ be a solution for M_0, and let $\langle g_0^a, h_0^a \rangle$ be a solution extending $\langle g_0, h_0 \rangle$ to the model M_0^a. As noted in the finally clause of Lemma 27.2.5, the induced solution $\langle g_0^b, h_0^b \rangle := \langle g_0^a, h_0^a \rangle^{f_0}$ for M_0^b extends the solution $\langle g_0, h_0 \rangle$ as well.

Now the extension property for solutions yields a chain $\{\langle g_i, h_i \rangle : i < \lambda^+\}$ of solutions for the models M_i, with $\langle g_i, h_i \rangle \subset \langle g_j, h_j \rangle$ for $i < j$. Using 2-amalgamation for solutions (which holds for $\lambda \leq \aleph_{k-4}$) and Corollary 27.2.14, we get increasing chains of solutions $\langle g_i^a, h_i^a \rangle$ and $\langle g_i^b, h_i^b \rangle$, $i < \lambda^+$, where $\langle g_{i+1}^a, h_{i+1}^a \rangle$ has domain M_{i+1}^a and is gotten by extension of solutions from the 2-amalgam of solutions $\langle g_i^a, h_i^a \rangle$ and $\langle g_{i+1}, h_{i+1} \rangle$ that has domain $M_i^a \cup M_{i+1}$. Further by repeated application of the strong form of Lemma 27.2.5 we get an increasing sequence isomorphisms f_i from M_i^a onto M_i^b which fix M_i and map a to b and preserve the solutions. The union of the f_i is the needed isomorphism between M^a and M^b that fixes M and sends $\boldsymbol{a}$ to $\mathbf{b}$. $\square_{27.4.3}$

We can establish an even better behavior for Galois types when $k \geq 3$;

Corollary 27.4.4. *Let $k \geq 3$. Then the class of models of ϕ_k is $(\aleph_0, \aleph_{k-3})$-tame. Moreover, the Galois types of finite tuples over a model of size up to $\aleph_{k-3}$ are determined by the syntactic types over that model.*

Proof. The first statement is an easy induction from the last lemma. We concentrate on the second, where it is enough to prove the claim for models of size $\aleph_0$. The proof will mimic the construction in the last lemma.

Fix $k \geq 3$ and suppose that $M \models \phi_k$ is a countable model and $\boldsymbol{a}$, $\mathbf{b}$ are finite tuples that have the same syntactic type over M. Find minimal (finite) X and Y such that $\boldsymbol{a}$ belongs to an L-structure generated by X, $\mathbf{b}$ to a structure generated by Y. Again, we let $X_0 := X \cap I(M)$. Let $M_0 \subset M$ be a finite L-substructure containing X_0; and let $\{M_i : i < \omega\}$ be an increasing chain of substructures converging to M.

Since $\boldsymbol{a}$ and $\mathbf{b}$ realize the same syntactic type over M_0 by Lemma 27.4.2, there is an isomorphism f between M_0^a and M_0^b over M_0.

The rest is a familiar argument: conjugate a solution on M_0^a by f to a solution on M_0^b; and use 2-amalgamation property on solutions (holds over finite sets for all $k \leq 3$ to extend the solutions to M^a and M^b. This gives the needed isomorphism fixing M and mapping $\boldsymbol{a}$ to $\mathbf{b}$. $\square_{27.4.4}$

27.5 Instability and Non-tameness

In this section we show that ϕ_k is not Galois stable in $\aleph_{k-2}$. We give details for the case: $k = 2$, showing there are continuum Galois types over a countable model of ϕ_2. The argument [5] for larger k involves a family of equivalence relations instead of just one and considerably more descriptive set theory.

Since for any u, the stalk G_u is affine (L')-isomorphic to the finite support functions from K to Z_2, without loss of generality we may assume each stalk has this form. We are working with models of cardinality $\leq \aleph_{k-2}$ so they admit solutions; thus, if we establish L'-isomorphisms they extend to L-isomorphisms. For any G^*-stalk G_u, the 0 in $(u, 0)$ denotes the identically 0-function in that stalk. But for a stalk in H^*, the 0 in $(u, 0)$ denotes the constant 0.

Claim 27.5.1. *Let M be the standard countable model of ϕ_2. There are $2^{\aleph_0}$ Galois types over M.*

Proof. Let E_0 be the equivalence relation of eventual equality on ${}^{\omega}2$; there are of course $2^{\aleph_0}$ equivalence classes.

Let $I(M) = \{a_0, \ldots, a_i, \ldots\}$. Pick a function $s \in {}^{\omega}2$, and define a model $M_s \succ M$ as follows. The L'-structure is determined by the set $I(M_s) = I(M) \cup \{b_s\}$. For the new compatible triples of the form $\{a_0, a_i\}$, $\{a_0, b_s\}$, $\{a_i, b_s\}$, define

$$M_s \models Q((\{a_0, a_i\}, 0), (\{a_0, b_s\}, 0), (\{a_i, b_s\}, 0))$$

if and only if $s(i) = 0$. The values of Q for any u_1, u_2, u_3 among the remaining new compatible triples is defined as:

$$M_s \models Q((u_0, 0), (u_1, 0), (u_2, 0)).$$

Note that 0 in the first two components of the predicate Q refer to the constantly zero functions in the appropriate G^*-stalks, and in the third component, 0 is a member of Z_2. A compact way of defining the predicate Q is:

$$(*) \qquad M_s \models Q((\{a_0, a_i\}, 0), (\{a_0, b_s\}, 0), (\{a_i, b_s\}, s(i))).$$

Note that by Notation 27.1.4, the definition of Q is determined on all of M.

Note that all the b_s realize the same *syntactic* type over M. Now we show that the E_0-class of s can be recovered from the structure of M_s over M; so the Galois types are distinct. Take two models M_s and M_t and suppose that the Galois types $\mathrm{tp}(b_s/M)$ and $\mathrm{tp}(b_t/M)$ are equal. Then there is an extension N of the model M_t and an embedding $f : M_s \to N$ that sends b_s to b_t. We work to show that in this case s and t are E_0-equivalent.

First, let us look at the stalks $G^*_{a_1,a_i}$, $G^*_{a_1,b_t}$, $H^*_{a_i,b_t}$ for $i > 1$. Since f fixes M, the constantly zero function $0 \in G^*_{a_0,a_i}$ is fixed by f. Let $x \in G^*_{a_1,b_t}$ be the image of $0 \in G^*_{a_1,b_s}$ under f. Then we have

$$M_t \models Q((\{a_1, a_i\}, 0), (\{a_1, b_t\}, x), (\{a_i, b_t\}, f(0))).$$

Since x is a finite support function, and we have defined

$$M_t \models Q((\{a_1, a_i\}, 0), (\{a_1, b_t\}, 0), (\{a_i, b_t\}, 0)),$$

for co-finitely many $i > 1$ we must have $f(0) = 0$ in the stalks $H^*_{a_i,b_t}$. In other words, f preserves all but finitely many zeros in $H^*_{a_i,b_t}$. In particular, by $(*)$ for any $s\!:\omega \to 2$ the functions s and $f(s)$ are E_0-equivalent.

We focus now on the stalks of the form $G^*_{a_0,a_i}$, $G^*_{a_0,b_t}$, $H^*_{a_i,b_t}$, $i \geq 1$. Again, since f fixes M, the constantly zero function $0 \in G^*_{a_0,a_i}$ is fixed by f. Letting $y \in G^*_{a_0,b_t}$ be the image of $0 \in G^*_{a_0,b_s}$ under f, we get

$$M_t \models Q((\{a_0, a_i\}, 0), (\{a_0, b_t\}, y), (\{a_i, b_t\}, f[s(i)])).$$

Since y is a finite support function, there is a natural number n such that $y(a_i, b_t) = 0$ for all $i > n$. Since we have defined

$$M_t \models Q((\{a_0, a_i\}, 0), (\{a_0, b_t\}, 0), (\{a_i, b_t\}, t(i))),$$

we get $t(i) = f(s(i))$ for all $i > n$, or $f(s)$ and t are E_0-equivalent. Combining this with the previous paragraph, we get that s is E_0-equivalent to t, as desired. $\square_{27.5.1}$

We refer the reader to [5] for the rather more complicated argument (it has a distinctive flavor of descriptive set theory) that many Galois types exist for a general k:

Theorem 27.5.2. *Let M be the standard model of ϕ_{k+2} of size $\aleph_k$. There are a family of elements b_s for $s \in 2^{\aleph_k}$ such the Galois types $(b_s/M; M_s)$ and $(b_t/M; M_t)$ are distinct. That is, b_s and b_t are in distinct orbits under automorpisms of the monster fixing M.*

We can now conclude, working with ϕ_k rather ϕ_{k+2}:

Proposition 27.5.3. *The class of models of ϕ_k is not $(\aleph_{k-3}, \aleph_{k-2})$-tame.*

Proof. Let s, t be sequences in ${}^{\omega_{k-2}}2$ with $\neg E_{k-2}(s, t)$. By Corollary 27.5.2, the Galois types of b_s, b_t over the standard model M of size $\aleph_{k-2}$ are different. But, by Corollary 27.4.4, the Galois type of b_s is the same as the Galois type of b_t over any submodel $N \prec M$, $\|N\| \leq \aleph_{k-3}$, as b_s and b_t have the same syntactic type over N. $\square_{27.5.3}$

This analysis shows the exact point that tameness fails. Grossberg pointed out that after establishing amalgamation in Section 27.3, non-tameness at some (μ, κ) could have been deduced from eventual failure of categoricity of the example and the known upward categoricity results [44, 101]. However, one could not actually compute the value of κ without the same technical work we used to show tameness directly. In addition, failure of categoricity is itself established using the Galois types constructed in Proposition 27.5.2.

In fact, the proof of Proposition 27.5.2 also yields:

Corollary 27.5.4. *Let $\chi_0, \ldots \chi_k$ be a strictly increasing sequence of infinite cardinals. Then there is a model of ϕ_{k+2} of cardinality χ_k over which there are 2^{χ_k} Galois types. In particular, ϕ_{k+2} is unstable in every cardinal greater than $\aleph_k$.*

We showed in Section 27.5 that ϕ_k is not Galois-stable in $\aleph_{k-2}$ and above; so ϕ_k is certainly not excellent. We have shown that the models of ϕ_k have disjoint amalgamation and it easy to see that ϕ_k has arbitrarily large models. Theorem 9.20 shows any Abstract Elementary class satisfying these conditions, that is categorical in λ, is Galois stable in μ for $\mathrm{LS}(\boldsymbol{K}) \leq \mu < \lambda$. Thus we can immediately deduce:

Theorem 27.5.5. *Let $k \geq 2$; ϕ_k is not $\aleph_{k-1}$-categorical.*

More refined arguments in [5], directly about the example, show that $\boldsymbol{K}$ has the maximal number of models in arbitrarily large cardinals.

Appendix A
Morley's Omitting Types Theorem

In this appendix, we give a general statement and proof of Morley's theorem for omitting types and finding two cardinal models. For the meaning of the 'arrow notation' below and the proof of the Erdos-Rado theorem see any model theory text. Marker [108] is particularly accessible and we have modelled this argument off his treatment of the countable case. However, it is crucial that we formulate the result for uncountable languages and omitting a large number of types. Shelah proves these general statements in Chapter VII.5 of [122] by distinctly different and in some ways more informative arguments. In particular, he explains the connection between the Hanf and well-ordering numbers.

Definition A.1 (Beth numbers)**.** $\beth_0(\kappa) = \kappa$. $\beth_{\alpha+1}(\kappa) = 2^{\beth_\alpha(\kappa)}$. $\beth_\delta(\kappa) = \Sigma_{\alpha<\delta}\beth_\alpha(\kappa)$. *If* $\kappa = \aleph_0$*, it is omitted.*

We need the Erdos-Rado theorem in a particular form.

Fact A.2 (Erdos-Rado)**.** *1.* $\beth_n(\kappa)^+ \rightarrow (\kappa^+)^n_\kappa$.

2. $\beth_{\alpha+n}(\kappa) \rightarrow (\beth_\alpha(\kappa)^+)^n_{\beth_\alpha(\kappa)}$

Proof. The first statement is exactly the Erdos-Rado theorem; the second follows by simple substitution and the identity $\beth_{\alpha+n}(\kappa) = \beth_n(\beth_\alpha(\kappa))$. $\square_{A.2}$

We work in a vocabulary τ and in a theory T which without loss of generality has Skolem functions. We write μ for $(2^{|\tau|})^+$ and H_1 for $\beth_\mu$. We introduced the notation used below in Chapter 5 and discussed it further in Chapter 7.2 and Chapter 9. We are sometimes cavalier in switching between constants and variables in the following two formulations. A τ-diagram is essentially a type in infinitely many variables; by increasing the number of variables while fixing the finite types realized, we are able to stretch a sequence of indiscernibles. To construct such a diagram we may, as in the following proof, pass to a theory in an inessential expansion τ^* of τ by replacing the variables by constants. Recall that we call a model M for vocabulary with a distinguished unary predicate P a (λ, κ)-model if $|M| = \lambda$ and $|P(M)| = \kappa$. The following theorem collects Morley's results on omitting types (part 1) with several variants on his two cardinal theorem for cardinals far apart (part 2). Note that both the hypothesis and conclusion those of 2a) are weaker than of 2b).

Theorem A.3. *Let T be a τ-theory, Γ a set of partial τ-types (in finitely many variables) over $\emptyset$ and $\mu = (2^{|\tau|})^+$, $H_1 = \beth_\mu$.*

1. *Suppose M_α for $\alpha < \mu$ are a sequence of τ-structures such that $|M_\alpha| > \beth_\alpha$ and M_α omits Γ.*

   Then, there is a countable sequence I of order indiscernibles such that the diagram of I is realized in each M_α and an extension Φ of T such that for every linear order J, $EM_\tau(J, \Phi) \models T$ and omits Γ.

   In particular, for every $\lambda \geq |\tau|$, there is a model N with $|N| = \lambda$ of T such that N omits Γ (i.e. omits each $p \in \Gamma$).

2. *Suppose P is a one-place τ-predicate. In addition to requiring that N omits Γ, we can demand.*

   (a) *Moreover, if for every $\alpha < \mu$, $|P(M_\alpha)| \geq \beth_\alpha$, then the indiscernibles in item 1) can be chosen inside P. Then, we can insure $|P(N)| = \lambda$.*

   (b) *Suppose for some κ and for every $\alpha < \mu$, $|M_\alpha| > \beth_\alpha(\kappa)$, $|P(M_\alpha)| = \kappa \geq |\tau|$, and M_α omits Γ.*
   Then, for every $\lambda \geq |\tau| + \kappa$, and any χ with $\kappa \leq \chi \leq \lambda$ there is a model N with $|N| = \lambda$ of T such that:

      i. *$|P(N)| = \chi$;*

      ii. *There is an infinite set $I \subseteq N$ of indiscernibles over $P(N)$.*

   (c) *If for each $\alpha < \mu$ both $P(|M_\alpha|) \geq H_1$ and $|M_\alpha| > \beth_\alpha(|P(M_\alpha)|)$ then for any χ with $|\tau| \leq \chi \leq \lambda$, we can choose N with $P(N) = \chi$.*

Proof. We will construct a sequence of theories Φ_i for $i < \mu$ such that their union Φ is an EM-template with certain specified properties. We first derive the results from the existence of Φ and then proceed to the much longer construction of Φ. Note that statement 1) is a special case of any of the formulations of the second. We prove only 2b) and then note the minor modifications needed for 2a) and 2c).

Clause ii) below guarantees the type is omitted; clause iii) asserts the new constants add no new elements to P; clause iv) asserts the new constants are indiscernible over P.

Let $C = \langle c_i : i < \omega\rangle$ be a sequence of new constant symbols and $\tau^* = \tau \cup C$. The τ^*-theory Φ will have the following properties:

i) $c_i \neq c_j$ if $i \neq j$;

ii) if $t(\mathbf{v})$ is a term then for each $p \in \Gamma$, there is a $\phi_p \in p$ such that if $i_1, \ldots i_n$ are strictly increasing:

$$\neg\phi_p(t(c_{i_1}, \ldots c_{i_n})) \in \Phi;$$

iii) if $t(v_1, \ldots, v_n, w_1, \ldots w_k)$ is a term and if $i_1, \ldots i_n$ and $j_1, \ldots j_n$ are strictly increasing,

$$(\forall \mathbf{w})((\bigwedge_{1\leq i\leq k} P(w_i) \wedge [P(t(c_{i_1}, \ldots c_{i_n}), \mathbf{w}) \\ \rightarrow (t(c_{i_1}, \ldots c_{i_n}, \mathbf{w}) = t(c_{j_1}, \ldots c_{j_n}, \mathbf{w}))] \tag{A.1}$$

is in Φ;

iv) for any $\psi(\mathbf{w}, \mathbf{v})$ if $i_1, \ldots i_n$ and $j_1, \ldots j_n$ are strictly increasing:

$$(\forall \mathbf{w}) \bigwedge P(w_i) \rightarrow [\psi(\mathbf{w}, c_{i_1}, \ldots c_{i_n}) \leftrightarrow \psi(\mathbf{w}, c_{j_1}, \ldots c_{j_n})] \in \Phi.$$

If we have such a Φ the result 2b) is clear. Form Φ' by adding χ distinct constants for members of P. By condition iv), the c_i remain indiscernible over the new constants. Let $N_0 = EM(J, \Phi')$ where $|J| = \chi$. The new constants guarantee that $|P(EM(J, \Phi))| = \kappa = |N_0|$. Extend J to an ordering I of cardinality λ and let $N = EM(I, \Phi')$. Then ii) guarantees that N omits Γ. and iv) shows the elements of I are actually indiscernible over $P(N)$. By iii), $P(EM(J, \Phi')) = P(EM(I, \Phi'))$ and so has cardinality χ.

2a) is even easier; we just choose set $\kappa = |\tau|$ and choose I inside P. Clause iii) and iv) are not needed. For 2c) under the additional hypothesis of 3) that $P(M_\alpha) \geq \beth_\alpha$ we can build a second sequence of indiscernibles inside P by the technique of Claim A.7. First take the Skolem hull M_μ of a sequence of indiscernibles in P of cardinality μ; M_μ contains a countable sequence J_0 outside of $P(M_\mu)$ that is indiscernible over P. Now add constants for $P(M_\mu)$ and then take the Skolem hull of a sequence of the same type as J_0 but with length λ; this model is as required.

We now prove a series of claims which will then support the induction step in the construction of Φ. We carry the following notation throughout.

Notation A.4. Let $(M_\alpha, P(M_\alpha))$ be a sequence of pairs satisfying the hypotheses of the theorem and let X_α be a subset of M_α with $|X_\alpha| > \beth_\alpha(\kappa)$. Fix a linear ordering of each M_α (we'll denote them all by $<$). In each of the following claims we construct from such a sequence a new sequence (M'_α, Y_α) meeting a specific condition, and which is a subsequence of the original sequence in the strong way that each $M'_\alpha = M_\beta$ for some $\beta \geq \alpha$. We guarantee that $Y_\alpha \subseteq X_\beta$ and $|Y_\alpha| > \beth_\alpha(\kappa)$.

Claim A.5 (Omit Γ). *Fix a term $t(\mathbf{x})$ The sequence (M'_α, Y_α) can be chosen to have the property:*

For each $p \in \Gamma$, there is a $\phi_p \in p$ such that if $i_1, \ldots i_n$ are strictly increasing and the y_{i_j} are in Y_α:

$$M'_\alpha \models \neg\phi_p(t(y_{i_1}, \ldots y_{i_n})).$$

Proof of Claim. Note $|\Gamma| \leq 2^{|\tau|}$. Let $N_\alpha = M_{\alpha+n}$. Define $F_\alpha : [X_{\alpha+n}]^n \to |\tau|^{|\Gamma|}$ by $F_\alpha(\mathbf{x})$ is the function f from Γ to the collection of 1-ary τ-formulas defined by $f(p)\{(t(x_1, \ldots x_n)\} = \phi_i$ is the least i (for some fixed enumeration of the 1-ary formulas) such that $\phi_i \in p$ and $\neg\phi_i(t(x_1, \ldots x_n))$.

Now $|X_{\alpha+n}| > \beth_{\alpha+n}(\kappa)$. By Fact A.2.2, noting that for $\alpha \geq 3$, $\beth_\alpha(\kappa) \geq |\tau|^{|\Gamma|}$:

$$\beth_{\alpha+n}(\kappa) \to (\beth_\alpha(\kappa)^+)^n_{|\tau|^{|\Gamma|}}$$

which is all we need to choose an appropriate $Y_\alpha \subseteq X_{\alpha+n}$. Let $\phi_{\alpha,p}$ denote the formula witnessing the omission of p on each tuple from $[Y_\alpha]^k$ and $Z_\alpha = \mathrm{rg}\, f_\alpha$. As α varies through $(2^{|\tau|})^+$, there are only $2^{|\tau|}$ choices for Z_α so we can choose a subsequence of the N_α as M'_α to obtain a fixed Z for the entire sequence. Let $\phi_p = \phi_i \in Z$ where i is least so that $\neg\phi_i(t(x_1, \ldots x_n))$ holds for all $\mathbf{x} \in Y_\alpha$ (for sufficiently large α). $\square_{A.5}$

Claim A.6 (Keep P small). *Fix a term $t(v_1, \ldots, v_n, w_1, \ldots, w_k)$. The sequence (M'_α, Y_α) has the property:*

If $\boldsymbol{a} = \langle a_1, \ldots a_k\rangle \in P(M'_\alpha)$, and $i_1, \ldots i_n$ and $j_1, \ldots j_n$ are strictly increasing and the $y_{i_j} \in Y_\alpha$:

$$(P(t(y_{i_1}, \ldots y_{i_n}, \boldsymbol{a})) \to (t(y_{i_1}, \ldots y_{i_n}, \boldsymbol{a}) = t(y_{j_1}, \ldots y_{j_n}, \boldsymbol{a}))).$$

Proof of Claim. Let $M'_\alpha = M_{\alpha+n}$. Fix $\eta \notin P(N_\alpha)$. For $x_1, \ldots x_n \in X_{\alpha+n}$, define $F_{\mathbf{x}} : [P(N_\alpha)]^k \to P(N_\alpha) \cup \{\eta\}$ by $F_{\mathbf{x}}(\boldsymbol{a}) = t(\mathbf{x}, \boldsymbol{a})$ if $t(\mathbf{x}, \boldsymbol{a}) \in P(N_\alpha)$ and η otherwise. Since $|P(N_\alpha)| = \kappa$, there are only $2^{\kappa+|\tau|}$ functions that can be $F_{\mathbf{x}}$. Thus, the map $\mathbf{x} \to F_{\mathbf{x}}$ partitions $[X_{\alpha+n}]^n$ into $2^{\kappa+|\tau|}$ pieces. Just as in Claim A.5, we obtain Y_α by Erdos-Rado; this time there is no need to refine the sequence. $\square_{A.6}$

Claim A.7 (Indiscernibility). *Fix a formula $\phi(v_1, \ldots, v_n, w_1, \ldots, w_k)$. The sequence (M'_α, Y_α) has the property for each α:*

If $\boldsymbol{a} = \langle a_1, \ldots a_k\rangle \in P(M'_\alpha)$, and $i_1, \ldots i_n$ and $j_1, \ldots j_n$ are strictly increasing and the $y_{i_j} \in Y_\alpha$:

$$\phi(y_{i_1}, \ldots y_{i_n}, \boldsymbol{a}) \leftrightarrow \phi(y_{j_1}, \ldots y_{j_n}, \boldsymbol{a}))).$$

Proof of Claim. Let $N_\alpha = M_{\alpha+n}$. For $x_1, \ldots x_n \in X_{\alpha+n}$, define $F_{\mathbf{x}} : [X_{\alpha+n}]^k \to \{0, 1\}$ by $F_{\mathbf{x}}(\boldsymbol{a}) = 0$ if $N_\alpha \models \phi(\mathbf{x}, \boldsymbol{a})$ and 1 otherwise. Again, $\mathbf{x} \to F_{\mathbf{x}}$ partitions $[X_{\alpha+n}]^n$ into 2^τ pieces. We obtain Y_α by Erdos-Rado; we refine the sequence to M'_α to make the same choice of ϕ for each α since 2 is not cofinal with $(2^{|\tau|})^+$. $\square_{A.7}$

Now we turn to the main construction and build Φ by induction. For $i < (2^{|\tau|})^+$, at each stage i of the construction we will have a collection of τ-sentences Φ_i and models $\langle (M_{i,\alpha}, X_{i,\alpha}) : \alpha < (2^{|\tau|})^+\rangle$, each linearly ordered by $<$ and $|X_{i,\alpha}| > \beth_\alpha(\kappa)$ so that if we interpret the c_i as an increasing sequence in $X_{i,\alpha}$, for each α:

$$M_{i,\alpha} \models \Phi_i.$$

Stage 0: Let Φ_0 be T along with the assertion that the c_i are distinct. For $\alpha < (2^{|\tau|})^+$ let $M_{0,\alpha}$ omit Γ and satisfies $|P(M_{0,\alpha})| = \kappa$. Let $X_{0,\alpha} = M_{0,\alpha}$. Clearly we can interpret the c_i to satisfy Φ_0.

Stage $i \equiv 1 \mod 3$**:** [omit Γ] Let $t_i = t(\mathbf{v})$. By Claim A.5, we can choose $\langle (M_{i+1,\alpha}, X_{i+1,\alpha}) : \alpha < (2^{|\tau|})^+\rangle$ with $|X_{i+1,\alpha}| > \beth_\alpha(\kappa)$, with $M_{i+1,\alpha} = M_{i,\beta}$ for some $\beta \geq \alpha$ and $X_{i+1,\alpha} \subseteq X_{i,\beta}$ and satisfying:

For all strictly increasing sequences $\mathbf{x}$ from $X_{i+1,\alpha}$ and all p,

$$M_{i+1,\alpha} \models \neg\phi_p(t(\mathbf{x})).$$

If $i_1, \ldots i_n$ are strictly increasing, for each p add: $\neg\phi_p(t(c_{i_1}, \ldots c_{i_n}))$ to Φ_i to form Φ_{i+1}.

Stage $i \equiv 2 \mod 3$**:** [keep P small] Let $t_i = t(v_1, \ldots, v_n)$. By Claim A.6, we can choose $\langle (M_{i+1,\alpha}, X_{i+1,\alpha}) : \alpha < (2^{|\tau|})^+\rangle$ with $|X_{i+1,\alpha}| > \beth_\alpha(\kappa)$, with $M_{i+1,\alpha} = M_{i,\beta}$ for some $\beta \geq \alpha$ and $X_{i+1,\alpha} \subseteq X_{i,\beta}$ and satisfying the conclusion of Claim A.6.

If $a_1, \ldots a_k \in P(M_{i+1,\alpha})$, and $i_1, \ldots i_n$ and $j_1, \ldots j_n$ are strictly increasing and the $c_{i_j} \in X_{i+1,\alpha}$, add the following to Φ_i to form Φ_{i+1}.

$$(\forall \mathbf{w})((\bigwedge_{1\leq i\leq k} P(w_i) \wedge [P(t(c_{i_1}, \ldots c_{i_n}), \mathbf{w}) \rightarrow (t(c_{i_1}, \ldots c_{i_n}, \mathbf{w}) = t(c_{j_1}, \ldots c_{j_n}, \mathbf{w}))].$$

Stage $i \equiv 3 \mod 3$**:** [Indiscernible over P] Let ϕ_i be the formula $\phi(v_1, \ldots, v_n, w_1, \ldots, w_k)$. By Claim A.7, we can choose $\langle (M_{i+1,\alpha}, X_{i+1,\alpha}) : \alpha < (2^{|\tau|})^+\rangle$ with $|X_{i+1,\alpha}| > \beth_\alpha(\kappa)$, with $M_{i+1,\alpha} = M_{i,\beta}$ for some $\beta \geq \alpha$ and $X_{i+1,\alpha} \subseteq X_{i,\beta}$ and satisfying the conclusion of Claim A.7. If (*) holds for $\boldsymbol{a}$ put $\phi(c_{i_1}, \ldots c_{i_k}, \boldsymbol{a}) \in \Phi_i$; otherwise $\neg\phi(c_{i_1}, \ldots c_{i_k}, \boldsymbol{a}) \in \Phi_i$.

If $a_1, \ldots a_k \in P(M_{i+1,\alpha})$, and $i_1, \ldots i_n$ and $j_1, \ldots j_n$ are strictly increasing and the $x_{i_j} \in X_{i+1,\alpha}$, put

$$(\forall \mathbf{w}) \bigwedge_{1\leq i\leq k} P(w_i) \rightarrow [\phi(c_{i_1}, \ldots c_{i_k}, \mathbf{w}) \leftrightarrow \phi(c_{j_1}, \ldots c_{j_k}, \mathbf{w})] \in \Phi_{i+1}.$$

Letting Φ be the union of the Φ_i, we have a consistent theory. To find a (λ, κ) model, first take a model M_0 of cardinality κ. There is a countable set I of indiscernibles over $P(M_0)$. Now we can stretch them to a set J size λ in a model M_1. Condition iii) guarantees that if $M_2 = EM(J, \Phi)$, $|P(M_2)| = \kappa$. $\square_{A.3}$

Note that as constructed we have not completely specified the diagram of the c_i. We can easily extend Φ to do so.

Remark A.8 (Morley's Method)**.** I believe that the phrase 'Morley's method' applies to the following observation (probably by Shelah). Suppose that in Theorem A.3.1 we fix a single set X (not necessarily a model) from which the approximations to the indiscernible set I are chosen. Then the finite types realized in I are those realized in X. A common application is to make X a long independent set; then to observe that I is also an independent set (e.g. [82].

Appendix B
Omitting types in Uncountable Models

In this appendix, we prove the key lemma for Keisler's argument that if a sentence ϕ of $L_{\omega_1,\omega}$ has few models in $\aleph_1$ then every model of ϕ realizes only countably many complete $L_{\mathcal{A}}$-types over $\emptyset$, where $L_{\mathcal{A}}$ is the smallest fragment including ϕ. This result, Theorem B.6, asserts that for any $L_{\mathcal{A}}$-type p over the $\emptyset$, there is a sentence θ_p such that a countable model $A \models \theta_p$ if and only A has an uncountable $L_{\mathcal{A}}$-elementary end extension omitting p. We rely on the omitting types theorem and a theorem on extendible models that are proved in [80]. We state Theorem B.6 rather differently than Keisler and rework the proof but the argument is really his. The sentence θ_p is not in $L_{\mathcal{A}}$, nor is it preserved under $L_{\mathcal{A}}$-elementary extension, but it is preserved by unions of $L_{\mathcal{A}}$-elementary chains.

The omitting types theorem for $L_{\omega_1,\omega}$ is proved in Chapter 11 of [80]. We will rely below on Keisler's version of *consistency property*, which appears in Chapters 3 and 4 of [80]. We defined the notion of a countable fragment of $L_{\omega_1,\omega}$ in Definition 2.2.

Theorem B.1 (Omitting types theorem)**.** *Let $L_{\mathcal{A}}$ be a countable fragment of $L_{\omega_1,\omega}$ and T a set of $L_{\mathcal{A}}$-sentences. Further, for each m, let p_m be a set of $L_{\mathcal{A}}$-formulas. Suppose*

1. *T has a model*

2. *and for each $m < \omega$ and any $\phi(x_1, \ldots x_{k_n}) \in L_{\mathcal{A}}$, if $T \cup (\exists \mathbf{x})\phi(\mathbf{x})$ has a model then so does $T \cup (\exists \mathbf{x})(\phi(\mathbf{x}) \wedge \neg\sigma_m)$ for some $\sigma_m \in p_m$.*

Then there is a model of T omitting all the p_m.

To formulate the next definition and theorem, we assume that there is a symbol $<$ in the vocabulary and that the theory T makes $<$ a linear order. In any linearly ordered structure we have the quantifier

$$(\exists \text{ arbitrarily large } u)\chi(u) : (\forall x)(\exists u)u > x \wedge \chi(u).$$

We extend the notion of fragment to require that if $\phi, \bigvee \Theta \in \Delta$ then $\bigvee(\{\text{for arb large } x)\theta : \theta \in \Theta\}$. Now we describe the notion of an *extendible* structure.

Definition B.2. *Let $(A, < \ldots)$ be a countable linearly ordered structure and $L_{\mathcal{A}}$ be a countable fragment of $L_{\omega_1,\omega}$. $(A, < \ldots)$ is $L_{\mathcal{A}}$-*extendible *if*

1. *$<$ has no last element;*
2. *($\exists$ arbitrarily large x) $\bigvee_n \phi_n \rightarrow \bigvee_n$ ($\exists$ arbitrarily large x) ϕ_n;*
3. *($\exists$ arbitrarily large x) $(\exists y)\phi(x, y) \rightarrow$*
   $(\exists y)$ ($\exists$ arbitrarily large x) $\phi(x, y) \vee$
   ($\exists$ arbitrarily large y)$(\exists x)\phi(x, y)$;

where $\bigvee_n \phi_n$ and $(\exists y)\phi$ are in $L_{\mathcal{A}}$.

Note that there is an $L_{\mathcal{A}}$-sentence over the empty set, θ^{ext}, such that for any countable $(B, <)$, $B \models \theta^{\text{ext}}$ if and only if $(B, <)$ is extendible.

Exercise B.3. *Verify that the conditions for extendibility are true in a model of power ω_1 that is ordered by ω_1.*

We use the next theorem in several places; it is proved as Theorem 28 of [80] and in [11].

Theorem B.4 (Keisler). *Let $(A, < \ldots)$ be a countable linearly ordered structure and $L_{\mathcal{A}}$ be a countable fragment of $L_{\omega_1,\omega}$. The following are equivalent:*

1. *$(A, <)$ has an $L_{\mathcal{A}}$-end-elementary extension;*
2. *$(A, <)$ has an $L_{\mathcal{A}}$-end-elementary extension with cardinality ω_1;*
3. *$(A, <)$ is $L_{\mathcal{A}}$-extendible;*
4. *$(A, <) \models \theta^{\text{ext}}$.*

Here are two notions that are used in the proof of B.4 that are essential for our proof of Theorem B.6.

Notation B.5. *We will write A' for the structure $(A, a)_{a \in A}$. Let T^A be the collection of all $L_{\mathcal{A}}(A)$-sentences true in A' along with all sentences $\theta(d)$ such that $\theta(x) \in L_{\mathcal{A}}$ and*

$$A \models (\exists y)(\forall x)[x > y \rightarrow \theta(x)].$$

We call an $L_{\mathcal{A}}(A)$ sentence ϕ an end extension sentence*, if it expresses that each type of a new element that is not above all members of A is omitted. Formally,*

$$\phi \coloneqq \bigwedge_{a \in A} (\forall y) \bigvee_{b \in A} (y = b \vee a < y).$$

The gist of the proof of Theorem B.4 uses the omitting types theorem to show $T^A \cup \{\phi\}$ is consistent for the end extension sentence ϕ. It is clear that

$$B \models T^A \cup \{\phi\}$$

implies B is a proper $L_{\mathcal{A}}$-elementary (by T^A), end (by ϕ) extension of A.

Now we come to the main result of this appendix, which restates Chapter 30 of [80].

Theorem B.6. *Fix a countable fragment $L_{\mathcal{A}}$ of $L_{\omega_1,\omega}$, a sentence ψ in $L_{\mathcal{A}}$ such that $<$ is a linear order of each model of ψ. For each $p(\mathbf{x})$ an $L_{\mathcal{A}}$-type (possibly incomplete) over the empty set, there is a sentence $\theta_p \in L_{\omega_1,\omega}$ satisfying the following conditions.*

1. *If p is omitted in an uncountable model $(B, <)$ of ψ then for any countable $(A, <)$ such that $(B, <)$ is an end $L_{\mathcal{A}}$-elementary extension of $(A, <)$, $(A, <) \models \theta_p$.*
2. *θ_p satisfies:*
   (a) *If $B \models \theta_p$ then B omits p.*
   (b) *θ_p is preserved under unions of chains of $L_{\mathcal{A}}$-elementary end extensions;*
   (c) *for any family X of $L_{\mathcal{A}}$-types $\langle p_m : m < \omega\rangle$ over $\emptyset$ and any countable A, if $A \models \theta_{p_m}$ for each m then A has a proper $L_{\mathcal{A}}$-elementary extension that satisfies each θ_{p_m}.*
3. *Let X be a collection of complete $L_{\mathcal{A}'}(\tau')$-types (for some $\mathcal{A}' \subseteq \mathcal{A}$ and $\tau' \subseteq \tau$) over the empty set that are realized in every uncountable model of ψ. Then, X is countable.*

Proof. We must first define the sentence θ_p. Let Γ be the set of all pairs $\langle \gamma, S\rangle$ where γ is a formula in $L_{\mathcal{A}}$ with all free variables displayed as $\gamma(\mathbf{x}, \mathbf{u}, \mathbf{y})$ and $S = (S\mathbf{u})$ is a finite sequence of $(\exists u_i)$ and 'for arbitrarily large u_j'. We fix the length of the $\mathbf{x}$-sequence but $\mathbf{u}$ and $\mathbf{y}$ can have any finite length.

Now for any type p in the variables $\mathbf{x}$, let θ_p be the conjunction of the extendability sentence θ^{ext} and the $L_{\omega_1,\omega}(\emptyset)$-sentence:

$$\bigwedge_{\langle \gamma, S\rangle \in \Gamma} (\forall \mathbf{y}) \left[(S\mathbf{u})(\exists \mathbf{x})\gamma(\mathbf{x}, \mathbf{u}, \mathbf{y}) \rightarrow \bigvee_{\sigma \in p} (S\mathbf{u})(\exists \mathbf{x})(\gamma(\mathbf{x}, \mathbf{u}, \mathbf{y}) \wedge \neg\sigma) \right].$$

Part 1). Given an uncountable model $(B,<)$ with order type ω_1 omitting p, we can choose a countable $(A,<)$ so that $(B,<)$ is an $L_{\mathcal{A}}$-elementary end extension of $(A,<)$. Then $(A,<)$ satisfies the extendibility sentence θ^{ext} by Theorem B.4. Applying Definition B.2.3, it is straightforward to prove the result by induction on the length of the string S.

Before proving part 2, we need some further notation.

Definition B.7 ($C_{p,A}$). *We say $C_{p,A}$ holds if for every $\langle\gamma(x,\mathbf{u},\mathbf{y}),S\rangle\in\Gamma$, every partition of $\mathbf{y}$ into $\mathbf{y}_1,\mathbf{y}_2$ and every substitution of an $\boldsymbol{a}\in A$ for $\mathbf{y}_2$ to yield $\gamma(\mathbf{x},\mathbf{u},\mathbf{y}_1,\boldsymbol{a})$ in $L_{\mathcal{A}}(A)$:*

$$A\models\left[(S\mathbf{u})(\exists\mathbf{x})\gamma(\mathbf{x},\mathbf{u},\mathbf{y}_1,\boldsymbol{a})\rightarrow\bigvee_{\sigma\in p}(S\mathbf{u})(\exists\mathbf{x})(\gamma(\mathbf{x},\mathbf{u},\mathbf{y}_1,\boldsymbol{a})\wedge\neg\sigma(\mathbf{x}))\right].$$

The following essential claim is easy to check using the observation that θ_p contains universal quantifiers while $C_{p,B}$ contains their instantiations over B.

Claim B.8. *For any B, $B\models\theta_p$ is equivalent to $C_{p,B}$ holds.*

Now, for part 2a) note that by Claim B.8, we may assume $C_{p,B}$ holds. But, if $C_{p,B}$ holds then B omits p since for any $\mathbf{b}\in B$ we can take $\gamma\in L_{\mathcal{A}}(B)$ from $C_{p,B}$ to be $\mathbf{x}=\mathbf{b}$, and then we have a direct statement that the type is omitted.

Part 2b) is immediate noting that if the hypothesis of one conjunct of θ_p is satisfied in some A_α, then a particular one of the disjuncts in the conclusion of the implication is true in A_α and is in $L_{\mathcal{A}}$ and so is true in every $L_{\mathcal{A}}$-elementary extension of A_α.

For part 2c), let A be a countable model of T and X a countable set of types over the empty set. We will concentrate on a single p, just noting at the key point that the omitting types theorem will allow us to lift all the θ_p for $p\in X$ to a single model.

Note that θ_p is actually in the form asserting a family of types is omitted (if we write $F\rightarrow\bigvee_i G_i$ as $\neg F\vee\bigvee_i G_i$). More precisely, define for each $\gamma(\mathbf{x},\mathbf{u},\mathbf{y})\in L_{\mathcal{A}}(A)$ and each string $(S\mathbf{u})$ the $L_{\mathcal{A}}(A)$-type:

$$\lambda_{p,\gamma,S}(\mathbf{y})=\{(S\mathbf{u})(\exists\mathbf{x})\gamma(\mathbf{x},\mathbf{u},\mathbf{y})\}\cup\{\neg(S\mathbf{u})(\exists\mathbf{x})\gamma(\mathbf{x},\mathbf{u},\mathbf{y})\wedge\neg\sigma:\sigma\in p\}.$$

Now, just checking the definition,

Lemma B.9. *For any B, B omits $\lambda_{p,\gamma,S}(\mathbf{y})$ for each γ,S with $\gamma\in L_{\mathcal{A}}(\emptyset)$ if and only if $B\models\theta_p$.*

Let T^A be the theory in $L_{\mathcal{A}}((A)\cup\{d\})$ introduced in Notation B.5 and let ϕ be the end extension sentence defined there. We want to show $T^A\cup\{\phi\}$ has a model (thus a proper $L_{\mathcal{A}}$-elementary end extension $(B,<)$ of $(A,<)$) omitting each $\lambda_{p,\gamma}$.

Lemma B.10. *Let A be countable and suppose $A\models\theta_p$. For any $L_{\mathcal{A}}(A)$-formula, $\gamma(\mathbf{x},\mathbf{u},\mathbf{y})$, any type $p(\mathbf{x})$ and any formula $\pi(d,\mathbf{y})\in L_{\mathcal{A}}(Ad)$, if $(\exists\mathbf{y})\pi(d,\mathbf{y})$ is consistent with T^A then*

- $(\exists \mathbf{y})(\pi(d,\mathbf{y}) \wedge \neg(S\mathbf{u})(\exists \mathbf{x})\gamma(\mathbf{x},\mathbf{u},\mathbf{y}))$ *is consistent with* T^A*, or*
- *for some* $\sigma \in p$*,*

$$(\exists \mathbf{y})(\pi(d,\mathbf{y}) \wedge (S\mathbf{u})(\exists \mathbf{x})\gamma(\mathbf{x},\mathbf{u},\mathbf{y}) \wedge \neg\sigma)$$

  is consistent with T^A.

Proof. Since $A \models \theta_p$, by Claim B.8, $C_{p,A}$ holds. Because of this observation we have suppressed additional parameters $\boldsymbol{a}$ which may occur in the formulas π and γ. Suppose $(\exists \mathbf{y})\pi(d,\mathbf{y})$ is consistent with T^A but

$$T^A \models \neg(\exists \mathbf{y})(\pi(d,\mathbf{y}) \wedge \neg(S\mathbf{u})(\exists \mathbf{x})\gamma(\mathbf{x},\mathbf{u},\mathbf{y})).$$

Then

$$T^A \models (\forall \mathbf{y})(\pi(d,\mathbf{y}) \rightarrow (S\mathbf{u})(\exists \mathbf{x})\gamma(\mathbf{x},\mathbf{u},\mathbf{y})).$$

Recall from the definition of T^A that the consistency of $\pi(d,\mathbf{y})$ means

$$A' \models (\exists \text{ arb large } x)(\exists \mathbf{y})\pi(x,\mathbf{y}).$$

Combining the last two,

$$A' \models (\exists \text{ arb large } x)(\exists \mathbf{y})(S\mathbf{u})(\exists \mathbf{x})(\gamma(\mathbf{x},\mathbf{u},\mathbf{y}) \wedge \pi(x,\mathbf{y})).$$

Let $S'x\mathbf{y}\mathbf{u}$ denote $(\exists \text{ arb large } x)(\exists \mathbf{y})(S\mathbf{u})$. With this notation, we have

$$A' \models (S'x\mathbf{y}\mathbf{u})(\exists \mathbf{x})(\gamma(\mathbf{x},\mathbf{u},\mathbf{y}) \wedge \pi(x,\mathbf{y})).$$

Now, since $C_{p,A}$ holds, for some $\sigma \in p$,

$$A' \models (S'x\mathbf{y}\mathbf{u})(\exists \mathbf{x})[\gamma(\mathbf{x},\mathbf{u},\mathbf{y}) \wedge \pi(x,\mathbf{y}) \wedge \neg\sigma(\mathbf{x})].$$

Again using the definition of T^A, we conclude

$$(\exists \mathbf{y})(S\mathbf{u})(\exists \mathbf{x})[\gamma(\mathbf{x},\mathbf{u},\mathbf{y}) \wedge \pi(d,\mathbf{y}) \wedge \neg\sigma(\mathbf{x})]$$

is consistent with T^A, whence

$$(\exists \mathbf{y})[\pi(d,\mathbf{y}) \wedge (S\mathbf{u})(\exists \mathbf{x})\gamma(\mathbf{x},\mathbf{u},\mathbf{y}) \wedge \neg\sigma(\mathbf{x})]$$

is consistent with T^A as required. $\square_{B.10}$

We conclude with the argument for Part 2c).

Lemma B.11. *For any countable model* $(A,<)$*, if* $A \models \theta_p$ *, for each* $p \in X$ *where* X *is countable, then* $(A,<)$ *has a proper* $L_{\mathcal{A}}$*-elementary end extension* $(B,<)$ *so that* $B \models \bigwedge_{p\in X}\theta_p$.

Proof. Since $(A,<)$ is extendible (this is implied by θ_p), each of the countably many types whose omission is encoded in ϕ is 'non-principal'. By Lemma B.10, for each $p \in X$, the same holds for each $\lambda_{p,\gamma,S}(\mathbf{y})$ with $\gamma \in L_{\mathcal{A}}(A)$. By the omitting types theorem, $T^A \cup \{\phi\}$ has a model (thus a proper $L_{\mathcal{A}}$-elementary end extension $(B,<)$ of $(A,<)$) omitting $\lambda_{p,\gamma,S}$ for each $p \in X$ and each $\gamma \in L_{\mathcal{A}}(A)$ and in particular for each $\gamma \in L_{\mathcal{A}}(\emptyset)$. Thus $B \models \bigwedge_{p\in X} \theta_p$. $\square B.11$

Finally, we prove Part 3) of Theorem B.6. Let X be a collection of complete $L_{\mathcal{A}'}(\tau')$-types (for some $\mathcal{A}' \subseteq \mathcal{A}$ and $\tau' \subseteq \tau$) over the empty set that are realized in every uncountable model of T. If $p \in X$, then for any such $(B,<)$ that is an $L_{\mathcal{A}}$-elementary end extension of a countable $(A,<)$,

$$(A,<) \models \neg\theta_p.$$

That is, there is a formula γ_p such that:

$$(A,<) \models (S\mathbf{u})(\exists \mathbf{x})\gamma_p,$$

but also for some $\sigma \in p$:

$$(A,<) \models \neg(S\mathbf{u})(\exists \mathbf{x})(\gamma_p \wedge \neg\sigma).$$

Note that while there are potentially continuum many formulas θ_p (infinite disjunction over p), there are only countably many possible formulas γ_p. So to conclude that there only countably many possible p, we need only show that if $p \neq q$ then $\gamma_p \neq \gamma_q$. Since X is a collection of *complete* $L_{\mathcal{A}'}(\tau')$-types, there is some $L_{\mathcal{A}}(\tau')$-formula σ with $\sigma \in p$ and $\neg\sigma \in q$.

$$(A,<) \models \neg(S\mathbf{u})(\exists \mathbf{x})\gamma_p \wedge \neg\sigma.$$

$$(A,<) \models \neg(S\mathbf{u})(\exists \mathbf{x})\gamma_q \wedge \neg\neg\sigma.$$

But if $\gamma_p = \gamma_q$, this contradicts that "arb large" and "there exists" distribute over disjunction. So we finish. $\square_{B.6}$

Appendix C
Weak Diamonds

In this appendix we state and prove the Devlin-Shelah weak diamond [36] and some variants. This is the only place the book goes beyond ZFC and we use only the axiom: $2^\lambda < 2^{\lambda^+}$. For the arguments here, one needs only rudimentary properties of closed unbounded sets (cub) and stationary sets, as found in many introductory set theory books, e.g. [92, 37]. In particular, $\mathrm{cub}(\lambda)$ denotes the filter of those sets which contain a closed unbounded subset of λ. More background on the role of weak diamonds and uniformization principles can be found in [39]. A more comprehensive exposition of the material will be in [43]. We introduce here a variant of the Devlin-Shelah diamond which we call $\Theta_{\chi,\lambda}$. $\Theta_{\lambda,\lambda}$ is precisely what is needed to show in Chapter 24 that excellence follows from the existence of very few models [125]. We show the slightly stronger fact that Θ_{χ,μ^+} holds if $2^\chi \leq 2^\mu$. Our proof is a slight variant of the original argument for the Devlin-Shelah diamond and quite different from the argument in [123]. $\Theta_{\chi,\lambda}$ is a simplified notation for what Shelah calls $\neg\mathrm{Unif}(\lambda,\chi,2,2)$. These results come from Chapter XIV of [123] which is reprinted with numerous corrections as an Appendix to [127]. It remains very hard to decipher.

We begin by stating the Devlin-Shelah weak diamond.

Definition C.1 (Φ_λ). *Φ_λ is the proposition; For any function $F : 2^{<\lambda} \to 2$ there exists $g \in 2^\lambda$ such that for every $f \in 2^\lambda$ the set $\{\delta < \lambda : F(f \restriction \delta) = g(\delta)\}$ is stationary.*

The function F partitions $2^{<\lambda}$ into two classes and Φ_λ asserts that g predicts for any f which member of the partition $f \restriction \delta$ lies in for stationarily many δ. Here is an ostensibly more general (but equivalent) formulation of this principle; we predict functions with range λ rather than 2.

Lemma C.2. Φ_λ *implies the proposition:*

For any function $F:\lambda^{<\lambda}\to 2$ *there exists* $g\in 2^\lambda$ *such that for every* $f\in\lambda^\lambda$ *the set* $\{\delta<\lambda : F(f\restriction\delta)=g(\delta)\}$ *is stationary.*

Proof. For $\alpha\le\lambda$, regard $f\in\alpha^\alpha$ as a subset of $\alpha\times\alpha$ and write χ_f for the characteristic function of the graph of f. Now given $F:\lambda^{<\lambda}\to 2$, define F^* on $2^{<\lambda}$ by $F^*(\eta)=F(f)$ if η is the characteristic function of the graph of the function f from α to α and 0 otherwise. By Φ_λ (applied to $\lambda\times\lambda$), there is an oracle g for F^*. For any $f\in\lambda^\lambda$, there is a cub C^f such that for $\delta\in C^f$, f maps δ into δ. Now $C^f\cap S$ where $S=\{\delta<\lambda : F^*(\chi_f\restriction\delta)=g(\delta)\}$ is the required stationary set. $\square_{C.2}$

For Remark 24.13.1 and to express clearly the argument in Lemma 24.12, we need a still more refined version of this principle. We prove this version with an additional parameter χ. The basic idea is that while it is crucial that g take on only two values, we can extend Lemma C.2 and allow the predicted function f to take values not only up to λ but up to certain $\chi>\lambda$ (on a small number of arguments). In the applications, failure of amalgamation will produce χ models in a larger cardinality.

We deal with functions with domain $\alpha\le\lambda$ and with certain restrictions on the range. We introduce some specific notation. $D(\alpha,\chi,\theta)$[1] is the collection of function with domain α whose first value is less than χ and all the rest are less than θ. Formally,

Definition C.3. *[$\Phi_{\chi,\lambda}$] Fix* $\chi<2^\lambda$.

1. *For* $\alpha\le\lambda$*, let* $D(\alpha,\chi,\theta)=\chi\times\theta^\alpha$*. Let* $D(<\lambda,\chi,\theta)=\bigcup_{\alpha<\lambda}D(\alpha,\chi,\theta)$. *Note* $D(\lambda,\lambda,2)=\lambda\times 2^\lambda\subset\lambda^\lambda$.
2. *Let* $\Phi_{\chi,\lambda}$ *be the following proposition.*

   For any function $F:D(<\lambda,\chi,2)\to 2$ *there exists* $g\in 2^\lambda$ *such that for every* $f\in D(\lambda,\chi,2)$ *the set* $\{\delta<\lambda : F(f\restriction\delta)=g(\delta)\}$ *is stationary.*
3. *In the more general notation of Definition C.11,* $\Phi_{\chi,\lambda}$ *is* $\neg\mathrm{Unif}(\lambda,\chi,2,2)$.

Then, $\Phi_{2,\lambda}$ is just another name for Φ_λ. By Lemma C.2, $\Phi_{2,\lambda}$ implies $\Phi_{\lambda,\lambda}$. Note that replacing $g\in 2^\lambda$ as written in Definition C.3.2 of $\Phi_{\chi,\lambda}$ by $g\in D(\lambda,\chi,2)$ yields an equivalent statement since we are only interested in the values of g on a stationary set.

For applications we need a restatement of the weak diamond.

Definition C.4. *[$\Theta_{\chi,\lambda}$] Let* $\Theta_{\chi,\lambda}$ *be the following proposition.*

For any collection of functions $\langle f_\eta:\eta\in 2^\lambda\rangle$*, with* $f_\eta\in D(\lambda,\chi,\lambda)$ *and any cub* $C\subset\lambda$ *there are* $\delta\in C$ *and* η,ν *such that:*

1. $\eta\restriction\delta=\nu\restriction\delta$,

[1]The relation D defined here is based on a more complicated notation in [127]; our notation is adequate for the cases we consider. But there is no direct translation.

2. $\eta(\delta) \neq \nu(\delta)$,

3. $f_\eta \restriction \delta = f_\nu \restriction \delta$,

Note in particular this implies $f_\eta(0) = f_\nu(0) < \chi$.

We observed that $\Phi_{2,\lambda} = \Phi_\lambda$ implies $\Phi_{\lambda,\lambda}$; thus a special case of Lemma C.5 is Φ_λ implies $\Theta_{\chi,\lambda}$ if $\chi \leq \lambda$. The additional step here (coming from [123] but with a different proof) is to allow $f(0)$ to attain values up to χ for still larger $\chi < 2^\lambda$. (Of course one could allow this freedom on any small number of arguments but we need only one so for ease of notation we give the freedom only at zero.) In Theorem C.6 we give a minor variant on the original proof of Φ_2 which suffices to prove $\Phi_{\chi,\lambda}$ if $\chi^\mu \leq 2^\mu$ (where $\lambda = \mu^+$).

Lemma C.5. *$\Phi_{\chi,\lambda}$ implies $\Theta_{\chi,\lambda}$.*

Proof. We must show that given $\{f_\eta \in D(\lambda,\chi,\lambda) : \eta \in 2^\lambda\}$ and a cub $C \subseteq \lambda$, there exist $\delta \in C$, and $\eta, \nu \in 2^\lambda$ such that: $\eta \restriction \delta = \nu \restriction \delta$, $\eta(\delta) \neq \nu(\delta)$, and $f_\eta \restriction \delta = f_\nu \restriction \delta$.

For each $\alpha < \lambda$, $\sigma \in 2^\alpha$ and $\tau \in D(\alpha,\chi,\lambda)$, let $F(\sigma,\tau) = 0$ if there is a $\rho \in 2^\lambda$ such that $\rho \restriction \alpha = \sigma$, f_ρ extends τ and $\rho(\alpha) = 0$; otherwise $F(\sigma,\tau) = 1$. As given, the arguments of F are pairs of functions from $D(< \lambda,\chi,\lambda)$. By coding pairs and applying the argument for Lemma C.2 we may assume that F is a map from $D(< \lambda,\chi,2)$ into 2. Thus, by $\Phi_{\chi,\lambda}$, there is an oracle $g \in D(\lambda,\chi,2)$ for F: for all $\eta \in 2^\lambda$ and $h \in D(\lambda,\chi,2)$,

$$\{\delta < \lambda : F(\eta \restriction \delta, h \restriction \delta) = g(\delta)\}$$

is stationary. The required $\eta \in 2^\lambda$ is defined by: for all $\alpha < \lambda$: $\eta(\alpha) = 1 - g(\alpha)$. Choose $\delta \in C$ so that:

$$F(\eta \restriction \delta, f_\eta \restriction \delta) = g(\delta).$$

First note that this implies $g(\delta) = 0$. For, if $g(\delta) = 1$, by the definition of η, $\eta(\delta) = 0$. Since $\eta \restriction \delta \lhd \eta$ and $f_\eta \restriction \delta \subset f_\eta$, we conclude from the definition of F that $F(\eta \restriction \delta, f_\eta \restriction \delta) = 0$. But then $g(\delta) = 0$.

Now, since $g(\delta) = 0$, the choice of δ guarantees that $F(\eta \restriction \delta, f_\eta \restriction \delta) = 0$. That is, there is an $\rho \in 2^\lambda$ such that $f_\eta \restriction \delta \subset f_\rho \in D(\lambda,\chi,\lambda)$, $\eta \restriction \delta = \rho \restriction \delta$, and $\rho(\delta) = 0$. As $\eta(\delta) = 1$, η, ρ and δ are as required. $\square_{C.5}$

Sometimes Φ_λ is written as Φ^2_λ to put it into a hierarchy where for Φ^k_λ the functions f and g map into k rather than 2. And it is known that even Φ^3_λ requires stronger set theoretic hypotheses. The difference between 3 and 2 is the ability to define the function f^g in the next proof. We have generalized the following argument from that in [36]. Shelah knew the stronger version long ago [123].

Theorem C.6. *If $\lambda = \mu^+$, $2^\mu < 2^\lambda$ and $\chi^\mu = 2^\mu$ then $\Phi_{\chi,\lambda}$ holds.*
In particular, if $2^{\aleph_{n-1}} < 2^{\aleph_n}$ then $\Phi_{2^{\aleph_{n-1}},\aleph_n}$

Proof. Suppose not. Then there is an $F : D(< \lambda,\chi,2) \to 2$ so that for every $g \in 2^\lambda$ there is an $f \in D(\lambda,\chi,2)$ for which the set $\{\delta < \lambda : F(f \restriction \delta) = g(\delta)\}$ is not stationary. Since g maps into 2, the set $\{\delta < \lambda : F(f \restriction \delta) = 1 - g(\delta)\}$ contains a cub. Replacing g by $1 - g$ we conclude: there is an $F : D(< \lambda, \chi, 2) \to 2$ so that for every $g \in 2^\lambda$ there is an $f = f^g \in D(\lambda, \chi, 2)$ for which the set $\{\delta < \lambda : F(f \restriction \delta) = g(\delta)\}$ contains a cub C^g.

Now we consider the set $\mathbf{T}$ of all structures $\langle \alpha, f_i, g_i \rangle_{i<\beta}$ where $\alpha, \beta < \lambda$, $g_i \in 2^\alpha$ and $f_i \in D(\alpha, \chi, 2)$. There are λ choices for α and only $\chi + 2^\mu$ choices for each f_i, g_i so there are only $\chi + 2^\mu$ such structures; by our hypothesis this is 2^μ. Let H be a $1-1$ function from $\mathbf{T}$ onto the set of functions from μ to 2.

Our goal is to define an equivalence relation on 2^λ and prove that the equivalence relation both has at most $2^\mu < 2^\lambda$ classes and is the identity. This contradiction will yield the theorem. For this purpose, we define by induction for each function $g \in 2^\lambda$, sequences of functions $\langle g_i, f_i : i < \mu \cdot \omega \rangle$ with $g_i \in 2^\lambda$ and $f_i \in D(\lambda, \chi, 2)$ and cubs on λ, C_n for $n < \omega$.

Fix $g \in 2^\lambda$. Let $C_0 = \lambda$ and for $i < \mu$, let $g_i = g$ and $f_i = f^g$. Then, let $C_1 = C^g$. Suppose we have constructed $\langle g_j, f_j : j < \mu \cdot n \rangle$ and C_n. We now define $\langle g_{\mu \cdot n+i}, f_{\mu \cdot n+i} : i < \mu \rangle$. Let $\beta_{\alpha,n} = \min\{\beta \in C_n : \beta > \alpha\}$ and to define the functions at α, let

$$T_{\alpha,n} = \langle \beta_{\alpha,n}, f_i \restriction \beta_{\alpha,n}, g_i \restriction \beta_{\alpha,n} \rangle_{i<\mu \cdot n}.$$

For each $i < \mu$ let

$$g_{\mu \cdot n+i}(\alpha) = H(T_{\alpha,n})(i).$$

Then let $f_{\mu \cdot n+i} = f^{g_{\mu \cdot n+i}}$ for $i < \mu$. Finally let C_{n+1} be the intersection of C_n with all the $C^{g_{\mu \cdot n+i}}$ for $i < \mu$; C_{n+1} is clearly a cub.

We require below the following easy observation.

Fact C.7. *Let $\langle C_i : i < \omega \rangle$ be a decreasing sequence of cubs and $C = \bigcap C_i$.*

1. *If $\delta_i = \min C_i$ and $\delta = \min C$ then $\delta = \sup \delta_i$.*
2. *If $\langle \gamma_i : i < \lambda \rangle$ is a continuous strictly increasing enumeration of C and*

$$\beta_{\gamma_j,n} = \min\{\alpha : \alpha > \gamma_j \wedge \alpha \in C_n\}$$

then $\gamma_{j+1} = \sup_{n<\omega} \beta_{\gamma_j,n}$.

Proof. The first is routine. For the second apply the first to the cubs

$$C_n^* = \{\alpha : \alpha > \gamma_j \wedge \alpha \in C_n\}.$$

$\square_{C.7}$

Now we define the equivalence relation.

Notation C.8. *For $g, g' \in 2^\lambda$ we have constructed C_n and C'_n for $n < \omega$ and also $\langle g_i, f_i : i < \mu \cdot \omega \rangle$, and $\langle g'_i, f'_i : i < \mu \cdot \omega \rangle$.*

1. *Let $C = \bigcap_{n<\omega} C_n$ and $C' = \bigcap_{n<\omega} C'_n$. Let $\langle \gamma_j : j < \lambda \rangle$, $\langle \gamma'_j : j < \lambda \rangle$ be strictly increasing continuous enumerations of C and C'.*

2. *We say gEg' if*

    (a) $\gamma_0 = \gamma_0'$;

    (b) *for every* $i < \mu \cdot \omega$, $g_i \restriction \gamma_0 = g_i' \restriction \gamma_0$;

    (c) *for every* $i < \mu \cdot \omega$, $f_i \restriction \gamma_0 = f_i' \restriction \gamma_0$;

Clearly E is an equivalence relation on 2^λ. Since there are only λ possible values for γ_0 and only 2^μ choices for 2b) and $(\chi + 2^\mu)^\mu$ possible representatives for 2c) there are there are less than 2^λ equivalence classes. We obtain our contradiction by proving:

Claim C.9. *For $g, g' \in 2^\lambda$, if gEg', then $g = g'$.*

Proof. We prove by induction on $j < \lambda$ that $\gamma_j = \gamma_j'$, and for each $i < \mu \cdot \omega$, $f_i \restriction \gamma_j = f_i' \restriction \gamma_j$ and $g_i \restriction \gamma_j = g_i' \restriction \gamma_j$. If we complete the construction, we have the lemma. Since C is a cub for any $\alpha < \lambda$ there is a $\gamma \in C$ with $\gamma > \alpha$. Then $g_0 \restriction \gamma = g_0' \restriction \gamma$. As $g_0 = g$ and $g_0' = g'$, g and g' agree on α.

For $j = 0$ this is immediate and at limits the result is routine. Consider a successor ordinal $j + 1$. For each $n < \omega$ and $i < \mu$, we defined the $g_{\mu \cdot n+i}$ so that $H(T_{\gamma_j,n}) = \langle g_{\mu\cdot n+i}(\gamma_j) : i < \mu\rangle$. But since $\gamma_j \in C_n$ and by induction,

$$\langle g_{\mu\cdot n+i}(\gamma_j) : i < \mu\rangle = \langle F(f_{\mu\cdot n+i} \restriction \gamma_j) : i < \mu\rangle = \langle F(f'_{\mu\cdot n+i} \restriction \gamma_j') : i < \mu\rangle.$$

But now applying the same argument to C_n', γ_j', we have

$$\langle F(f'_{\mu\cdot n+i} \restriction \gamma_j') : i < \mu\rangle = \langle g'_{\mu\cdot n+i}(\gamma_j') : i < \mu\rangle = H(T'_{\gamma_j',n}).$$

Since H is one-to-one we have for every $n < \omega$, that $T_{\gamma_j,n} = T'_{\gamma_j',n}$.

In particular, this implies that for all n, $\beta_{\gamma_j,n} = \beta'_{\gamma_j',n}$. Thus $\gamma_{j+1} = \gamma'_{j+1}$ by Fact C.7.2. But then induction and the fact that $T_{\gamma_j,n} = T'_{\gamma_j',n}$ implies that $g_i \restriction \beta_{\gamma_j,n} = g_i' \restriction \beta'_{\gamma_j,n}$ and $f_i \restriction \beta_{\gamma_j,n} = f_i' \restriction \beta'_{\gamma_j,n}$ for each $n < \omega$. Thus $g_i \restriction \gamma_{j+1} = g_i' \restriction \gamma'_{j+1}$ and $f_i \restriction \gamma_{j+1} = f_i' \restriction \gamma'_{j+1}$. This completes the proof of the claim. $\square_{C.9}$.

And so we have proved the main theorem. $\square_{C.6}$

Remark C.10. Our statement of Theorem C.6 was gleaned from Chapter XIV of [123] and the appendix of [127].

The argument here shows that one can allow the function from λ to have $\chi > \lambda$ values in a few places without disturbing the original argument. But the size of range of the functions enters the argument at a second point: defining the diagonalization H; in formulating Theorem C.6 we simply kept χ small enough for the argument for Φ_λ to continue to work.

We briefly describe the connections to some extensions by Shelah in [123] and [127]. He defines the term Unif with more parameters. I think the definition below is what his notion specifies in the relevant cases.

Definition C.11. *[*$\mathrm{Unif}(\lambda, \chi, \theta, \kappa)$*]* [2]

1. *Let* $\mathrm{Unif}(\lambda, \chi, \theta, \kappa)$ *be the following proposition.*

   There is a function $F : D(<\lambda, \chi, \theta) \to {}^{\lambda}\theta$ *such that for any* $g \in {}^{\lambda}\kappa$ *there is an* $f \in D(\lambda, \chi, \theta)$ *such that the set* $\{\delta < \lambda : F(f \restriction \delta) = g(\delta)\}$ *contains a cub.*

2. *Then* $\mathrm{Unif}(\lambda, \chi, 2^{<\lambda}, 2^{<\lambda})$ *is the following proposition. There is a function* $F : D(\lambda, \chi, 2^{<\lambda}) \to 2^{<\lambda}$ *such that for any* $g \in {}^{\lambda}({}^{<\lambda}2)$ *there is an* $f \in D(<\lambda, \chi, 2^{<\lambda})$ *such that the set* $\{\delta < \lambda : F(f \restriction \delta) = g(\delta)\}$ *contains a cub.*

Under this definition, using the trick in the first paragraph of the proof of Theorem C.6, Φ_λ is $\neg\mathrm{Unif}(\lambda, 2, 2, 2)$ and $\Phi_{\chi,\lambda}$ is $\neg\mathrm{Unif}(\lambda, \chi, 2, 2)$.

Remark C.12. Various statements in [125] are confusing and after consultation with Shelah I include some clarification.

1. In 2) on the bottom of page 270 of [125], the intent is to withdraw the claim: $I(\lambda^{++}, \boldsymbol{K})^{\aleph_0} = 2^{\lambda^{++}}$.

2. The difficulty in Theorem 6.4 and footnote on page 265 of [125] is the assertion that $I(\lambda, \boldsymbol{K}) > \chi$ (as in Lemma 24.12) can be obtained from $\neg\mathrm{Unif}(\lambda, \chi, 2^{\mu}, 2^{\mu})$. This is unproved; as we have shown, $\neg\mathrm{Unif}(\lambda, \chi, \lambda, 2)$ does suffice. Theorem 1.10 of the Appendix to [127] proves:

Theorem C.13. *If* $\lambda = \mu^+$, $2^{\mu} < 2^{\lambda}$ *and* $\chi^{\aleph_0} < 2^{\lambda}$, *then* $\mathrm{Unif}(\lambda, \chi, 2^{\mu}, 2^{\mu})$ *fail.*

The argument is based on mutually almost disjoint sets and is rather different from the one presented here or [36].

[2] These are simplifications of the notions defined by Shelah; we list all parameters throughout to avoid the confusion caused when one member of a string is suppressed.

Appendix D
Problems

These problems range from some that should be immediately accessible to major conjectures and methodological issues and areas in need of development. Many of them have been 'in the air' for some time. We have not attempted a detailed history of the problems but refer to relevant papers.

Problem D.1 (Eventual Categoricity)**.** *[Shelah's Categoricity Conjecture]* Calculate a cardinal $\mu(\kappa_{\boldsymbol{K}})$ such that if an AEC $\boldsymbol{K}$ is categorical in some cardinal $\lambda > \kappa_{\boldsymbol{K}}$ then $\boldsymbol{K}$ is categorical in all cardinals μ.

We noted in Chapter 16, that [128] proves the existence (with no method of calculation) of such a μ for successor cardinals λ when $\boldsymbol{K}$ satisfies amalgamation and joint embedding.

For the following problems suppose $\boldsymbol{K}$ has the amalgamation property, joint embedding, and arbitrarily large models. Suppose $\boldsymbol{K}$ is λ-categorical with $\lambda > H_1$.

1. If such an AEC is categorical in a limit cardinal greater than H_2 must it be categorical in all larger cardinals?
2. If $\boldsymbol{K}$ is categorical in a sufficiently large limit cardinal λ, must the model of cardinal λ be saturated.
3. Can one calculate $\mu(\kappa_{\boldsymbol{K}})$?
4. Is this problem any easier if you restrict to $L_{\omega_1,\omega}$?
5. Can the lower bound in Theorem 14.13 be brought below H_2? Is it as low as $\beth_\omega$ or $\beth_{\omega_1}$?

Can the results of Chapters 14 –16 be proved (possibly with further set theory) if the hypothesis of amalgamation is replaced by 'no maximal models'? (See Chapter 17, [141, 144].)

Problem D.2 (Tameness Spectrum)**.** Some related questions arise for tameness; the first three questions continue to assume amalgamation, joint embedding, and arbitrarily large models.

1. Is there any way to reduce the upper bound on χ in Theorem 12.15 (or find a lower bound above $\mathrm{LS}(\boldsymbol{K})$)?
2. Is there any way to replace weakly tame by tame or by local in Theorem 12.15?
3. Does categoricity in λ (sufficiently large) or even λ-stability imply $\boldsymbol{K}$ is (H_1, λ)-tame?
4. Find more, and more algebraic, examples of tameness appearing late [1].
5. Is there an eventually categorical AEC with Löwenheim number $\aleph_0$ that is not $(\aleph_0, \infty)$-tame?

Note that a positive solution of Problem D.2.3 would yield eventual categoricity. For nuances on this problem consult Chapters 14 – 16.

Problem D.3 (Other spectrum issues)**.** There are a number of question about the spectrum of amalgamation, tameness, stability and so on.

[Grossberg's Amalgamation Question] Calculate a cardinal $\mu(\kappa)$ such that if an AEC has amalgamation in some cardinal $\lambda > \kappa$ then $\boldsymbol{K}$ has amalgamation in all cardinals μ.

The promised [118] asserts that AEC may fail amalgamation for the first time at quite large cardinals, but below the first measureable.

1. Is there a short list of functions which give all possible stability spectrums for an AEC? Are there examples of mathematical interest in the various classes?
2. Are there similar functions for the spectrum of tameness, amalgamation, etc.?
3. In particular, is there a cardinal κ such that if an AEC $\boldsymbol{K}$ is categorical in $\lambda > \kappa$ then $\boldsymbol{K}$ is stable in all sufficiently large μ?
4. If an AEC fails amalgamation in λ (and λ is sufficiently large) must it have 2^{λ^+} models in λ^+? (No assumption is made on the number of models in λ.) ([13] has a counterexample with $\lambda = \mathrm{LS}(\boldsymbol{K}) = \aleph_0$.)

[1] Baldwin and Shelah [17] provide an example of an AEC with the amalgamation property in a countable language with Lowenheim-Skolem number $\aleph_0$ which is not $(\aleph_0, \aleph_1)$-tame but is $(2^{\aleph_0}, \infty)$-tame. But the language is artificial.

Problem D.4 (Atomic classes)**.** 1. Can one find an integer $m \geq 1$ and an atomic class $\boldsymbol{K}$ that is $\aleph_n$-categorical for $n \leq m$ but $\boldsymbol{K}$ does not have arbitrarily large models?

2. Is there an ω-stable atomic class with no model above $\aleph_2$?

3. Find an atomic class $\boldsymbol{K}$ that is ω-stable and tame but not excellent.

Problem D.5 (quasiminimal excellence)**.** 1. Determine if Zilber's conjecture that $(\mathcal{C}, +, \cdot, \exp)$ is a member of the quasiminimal class of pseudo-exponential fields is correct.

2. Can one show covers of semiabelian varieties (Chapter 4 and [151]) are tame by direct algebraic argument?

Problem D.6 ($L_{\omega_1,\omega}(Q)$)**.** 1. Show that the Shelah categoricity transfer theorem for $L_{\omega_1,\omega}$ also holds for $L_{\omega_1,\omega}(Q)$. More modestly, show this for a fragment of $L_{\omega_1,\omega}(Q)$ sufficient for the Zilber construction.

2. Is there an example of an AEC $\boldsymbol{K}$ with $\mathrm{LS}(\boldsymbol{K}) = \omega_1$ which is $(\aleph_0, \infty)$-tame and is not defined by a sentence in $L_{\omega_1,\omega}(Q)$?

3. Is it possible to represent an $\aleph_1$-categorical sentence in $L_{\omega_1,\omega}(Q)$ by an AEC with Löwenheim-Skolem number $\aleph_0$ and the same uncountable models?

Problem D.7 (Set theoretic issues)**.** 1. Is categoricity in $\aleph_1$ of an $L_{\omega_1,\omega}$-sentence absolute?[81]

2. More specifically, is there (in some model of set theory) an $\aleph_1$-categorical complete sentence of $L_{\omega_1,\omega}$ that is not ω-stable and/or does not have amalgamation in $\aleph_0$?

3. Can one prove the following assertion from VWGCH? If a sentence ϕ of $L_{\omega_1,\omega}$ has less than $2^{\aleph_n}$ models in $\aleph_n$ for each $n < \omega$ then ϕ defines an excellent class.

4. Is the Galois stability spectrum of an AEC absolute?

5. Find in ZFC examples of AEC where the Galois types are not compact in the sense of Definition 12.4[2].

Problem D.8 (Finitary Classes)**.** *See Definition 25.14.*

1. (Kueker) Is there a finitary class that is not axiomatizable in $L_{\infty,\omega}$.

2. (Kueker) If a finitary class is categorical in some $\kappa > \aleph_0$, must it be $L_{\omega_1,\omega}$ axiomatizable?

[2]The examples of non-compactness as in Problem D.7.5 in [17] use $\diamondsuit$ and $\square$.

3. Is there a sentence of $L_{\omega_1,\omega}$ which is $\aleph_1$-categorical and has both the amalgamation and joint embedding properties but has $2^{\aleph_0}$ countable models[3]?

4. Distinguish 'finite diagrams' (not necessarily homogeneous), finitary classes in the sense of [65], and 'atomic classes'.

5. Is Vaught's conjecture true for finite diagrams (or for finitary classes) that are ω-stable?

   Examples 7.1.3 and 5.13 illustrate the importance of joint embedding for Problem D.8.3 and ω-stability for Problem D.8.5.

Problem D.9 (Methodological issues)**.** 1. Are there examples of AEC with no syntax more natural that the $PC\Gamma$ characterization?

2. Find natural mathematical classes that are AEC but not axiomatizable in an infinitary logic. In particular, find such examples more connected with problems in geometry and algebra. Find such examples distinguished by natural AEC properties: tameness, amalgamation, stability etc.

3. Develop a theory of superstability that connects the stability spectrum function with properties in a fixed cardinal such as uniqueness of limit models and preservation of saturation under unions of chains. (This is in some sense a parallel to Shelah's work on good frames.) See also [50].

4. Prove categoricity transfer theorems for classes with no maximal models. See [144] and [141].

5. Develop geometric stability theory in the context of AEC [71]. This was the origin of the development of finitary classes ([65, 64, 66].

I thank David Kueker for significant clarifications and suggestions for this section.

Ollyantatambo, May 26, 2008.

[3]These question were raised by Marker and Kueker in the Fall of 2008 but Kierstead [81] had addressed them and provided conjectures in terms of admissible logics thirty years before.

References

[1] U. Abraham and S. Shelah. Martin's axiom does not imply that every two $\aleph_1$-dense sets of reals are isomorphic. *Israel J Math*, 38:161–176, 1981.

[2] John T. Baldwin. Rank and homogeneous structures. In Katrin Tent, editor, *Tits Buildings and the Theory of Groups Wurzburg Sept 14-17 2000*. Cambridge University Press, 2002.

[3] John T. Baldwin. The complex numbers and complex exponentiation: Why infinitary logic is necessary! In *Apuntes del XV Congreso Nacional*. Columbian Mathematical Society, 2006.

[4] John T. Baldwin, Paul Eklof, and Jan Trlifaj. $^\perp N$ as an abstract elementary class. *Annals of Pure and Applied Logic*, 149:25–39, 2007. preprint: `www.math.uic.edu/\~jbaldwin`.

[5] John T. Baldwin and Alexei Kolesnikov. Categoricity, amalgamation, and tameness. to appear: Israel Journal of Mathematics; also at www.math.uic.edu/˜ jbaldwin.

[6] John T. Baldwin and Alexei Kolesnikov. More counterexamples.

[7] J.T. Baldwin. Some EC_Σ classes of rings. *Zeitshrift Math. Logik. and Grundlagen der Math*, 24:489–492, 1978.

[8] J.T. Baldwin. *Fundamentals of Stability Theory*. Springer-Verlag, 1988.

[9] J.T. Baldwin. Notes on quasiminimality and excellence. *Bulletin of Symbolic Logic*, 10:334–367, 2004.

[10] J.T. Baldwin. Ehrenfeucht-Mostowski models in abstract elementary classes. In Yi Zhang, editor, *Logic and its Applications*, Contemporary Mathematics, pages 1–17. American Mathematical Society, 2005.

[11] J.T. Baldwin. Models in ω_1. lecture notes: see `http://www2.math.uic.edu/~jbaldwin/keisleruic6.pdf`, 2007.

[12] J.T. Baldwin. Vaught's conjecture, do uncountable models count? *Notre Dame Journal of Formal Logic*, pages 1–14, 2007.

[13] J.T. Baldwin, A. Kolesnikov, and S. Shelah. The amalgamation spectrum. in preparation.

[14] J.T. Baldwin, D.W. Kueker, and M. VanDieren. Upward stability transfer theorem for tame abstract elementary classes. *Notre Dame Journal of Formal Logic*, 47:291–298, 2006.

[15] J.T. Baldwin and A.H. Lachlan. On strongly minimal sets. *Journal of Symbolic Logic*, 36:79–96, 1971.

[16] J.T. Baldwin and O. Lessmann. Upward categoricity of very tame abstract elementary classes with amalgamation. *Annals of Pure and Applied Logic*, 143:29–42, 2006.

[17] J.T. Baldwin and S. Shelah. Examples of non-locality. submitted; on web page; Sh862.

[18] J.T. Baldwin and S. Shelah. The primal framework: I. *Annals of Pure and Applied Logic*, 46:235–264, 1990.

[19] J.T. Baldwin and S. Shelah. The primal framework II: Smoothness. *Annals of Pure and Applied Logic*, 55:1–34, 1991.

[20] J.T. Baldwin and S. Shelah. Randomness and semigenericity. *Transactions of the American Mathematical Society*, 349:1359–1376, 1997.

[21] J. Barwise, editor. *The syntax and semantics of infinitary languages*. LNM 72. Springer-Verlag, 1968.

[22] J. Barwise. Back and forth through infinitary logic. In M. Morley, editor, *Studies in Model Theory*, pages 1–34. Mathematical Association of America, 1973.

[23] J. Barwise. The role of omitting types in infinitary logic. *Archive fur Math. Log.*, 21:55–68, 1981.

[24] J. Barwise and S. Feferman, editors. *Model-Theoretic Logics*. Springer-Verlag, 1985.

[25] J. Baumgartner. On $\aleph_1$-dense sets of reals. *Fundamenta Mathematica*, LXXIX:101–106, 1973.

[26] M. Bays and B.I. Zilber. Covers of multiplicative groups of an algebraically closed field of arbitrary characteristic. preprint: arXive math.AC/0401301, 200?

[27] I. Ben-Yaacov. Positive model theory and compact abstract theories. *Journal of Mathematical Logic*, 3:85–118, 2003.

[28] I. Ben-Yaacov. Simplicity in compact abstract theories. *Journal of Mathematical Logic*, 3:163–191, 2003.

[29] A. Berenstein and S. Buechler. Simple stable homogeneous expansions of banach spaces. *Annals of Pure and Applied Logic*, 128:75–101, 2004.

[30] E. Bouscaren, editor. *Model Theory and Algebraic Geometry : An Introduction to E. Hrushovski's Proof of the Geometric Mordell-Lang Conjecture*. Springer-Verlag, 1999.

[31] S. Buechler and O. Lessmann. Simple homogeneous models. *Journal of the American Mathematical Society*, 16:91–121, 2003.

[32] Steven Buechler. *Essential Stability Theory*. Springer-Verlag, 1991.

[33] C. C. Chang. Some remarks on the model theory of infinitary languages. In J. Barwise, editor, *The syntax and semantics of infinitary languages*, pages 36–64. Springer-Verlag, 1968. LNM 72.

[34] C.C. Chang and H.J Keisler. *Model theory*. North-Holland, 1973. 3rd edition 1990.

[35] Andrew Coppola. *The theory of Q-abstract elementary classes*. PhD thesis, University of Illinois, Chicago, 2005.

[36] K. Devlin and S. Shelah. A weak version of $\diamond$ which follows from $2^{\aleph_0} < 2^{\aleph_1}$. *Israel Journal of Mathematics*, 28:239–247, 1978.

[37] Keith Devlin. *Constructibility*. Springer-Verlag, 1984.

[38] A. Ehrenfeucht and A. Mostowski. Models of axiomatic theories admitting automophisms. *Fund. Math.*, 43:50–68, 1956.

[39] P. Eklof and Alan Mekler. *Almost Free Modules: Set theoretic Methods*. North Holland, 1990.

[40] P. Eklof and Alan Mekler. *Almost Free Modules: Set theoretic Methods*. North Holland, 2002. 2nd edition.

[41] Misha Gavrilovich. *Model Theory of universal covering spaces of complex analytic varieties*. PhD thesis, Balliol college Oxford, 2006.

[42] George Grätzer. *Universal Algebra*. Van Nostrand, 1968.

[43] R. Grossberg. A course in model theory. preprint.

[44] R. Grossberg and M. VanDieren. Categoricity from one successor cardinal in tame abstract elementary classes. *The Journal of Mathematical Logic*, 6:181–201, 2006.

[45] R. Grossberg and M. VanDieren. Galois stability for tame abstract elementary classes. *Journal of Mathematical Logic*, 6:1–24, 2006.

[46] R. Grossberg and M. VanDieren. Shelah's categoricity conjecture from a successor for tame abstract elementary classes. *The Journal of Symbolic Logic*, 71:553–568, 2006.

[47] R. Grossberg, M. VanDieren, and Andrés Villaveces. Limit models in stable classes with amalgamation. preprint.

[48] Rami Grossberg. Classification theory for non-elementary classes. In Yi Zhang, editor, *Logic and Algebra*, pages 165–204. AMS, 2002. Contemporary Mathematics 302.

[49] Rami Grossberg and Bradd Hart. The classification theory of excellent classes. *The Journal of Symbolic Logic*, 54:1359–1381, 1989.

[50] Rami Grossberg and Alexei Kolesnikov. Excellent abstract elementary classes are tame. preprint.

[51] Rami Grossberg and Olivier Lessmann. Shelah's stability spectrum and homogeneity spectrum in finite diagrams. *Archive for Mathematical Logic*, 41:1–31, 2002.

[52] Rami Grossberg and Saharon Shelah. On universal locally finite groups. *Israel J Math*, 44:289–302, 1983.

[53] R. Gruber and J. Trlifaj. *Approximation and Endomorphism Algebras of Modules*. De Gruyter, 2006.

[54] William Hanf. Models of languages with infinitely long expressions. In *Abstracts of Contributed papers from the First Logic, Methodology and Philosopy of Science Congress, Vol.1*, page 24. Stanford University, 1960.

[55] Bradd Hart. An exposition of OTOP. In J. Baldwin, editor, *Classification Theory: Chicago, 1985*. Springer-Verlag, 1987.

[56] Bradd Hart and Saharon Shelah. Categoricity over P for first order T or categoricity for $\phi \in \mathrm{l}_{\omega_1\omega}$ can stop at $\aleph_k$ while holding for $\aleph_0, \cdots, \aleph_{k-1}$. *Israel Journal of Mathematics*, 70:219–235, 1990.

[57] W. Henson and Jose Iovino. Ultraproducts in analysis. In *Analysis and Logic*, volume 262 of *London Mathematical Society Lecture Notes*, pages 1–115. Cambridge University Press, 2002.

[58] Åsa Hirvonen. *Categoricity in homogeneous complete metric spaces*. PhD thesis, University of Helsinki, 2006.

[59] Greg Hjorth. Knight's model, its automorphism group, and characterizing the uncountable cardinals. *Notre Dame Journal of Formal Logic*, 2007.

[60] W. Hodges. What is a structure theory? *Bulletin of the London Mathematics Society*, 19:209–237, 1987.

[61] W. Hodges. *Model Theory*. Cambridge University Press, 1993.

[62] E. Hrushovski. Simplicity and the Lascar group. preprint.

[63] T. Hyttinen and Åsa Hirvonen. Categoricity in homogeneous complete metric spaces. preprint, 200?

[64] T. Hyttinen and M. Kesälä. Categoricity transfer in finitary abstract elementary classes. preprint, 200?

[65] T. Hyttinen and M. Kesälä. Independence in finitary abstract elementary classes. *Ann. Pure Appl. Logic*, 143(1-3):103–138, 2006.

[66] T. Hyttinen and M. Kesälä. Superstability in simple finitary AECs. *Fund. Math.*, 195(3):221–268, 2007.

[67] T. Hyttinen and S. Shelah. Strong splitting in stable homogeneous models. *Annals of Pure and Applied Logic*, 103:201–228, 2000.

[68] T. Hyttinen and S. Shelah. Main gap for locally saturated elementary submodels of a homogeneous structure. *J Symbolic Logic*, 66:1286–1302, 2001.

[69] Tapani Hyttinen. Types in abstract elementary classes. *Notre Dame J. Formal Logic*, 45(2):99–108 (electronic), 2004.

[70] Tapani Hyttinen and Olivier Lessmann. Simplicity and uncountable categoricity in excellent classes. *Ann. Pure Appl. Logic*, 139(1-3):110–137, 2006.

[71] Tapani Hyttinen, Olivier Lessmann, and Saharon Shelah. Interpreting groups and fields in some nonelementary classes. *J. Math. Log.*, 5(1):1–47, 2005.

[72] Jose Iovino. Stable Banach spaces and Banach space structures i. In *Models, algebras, and proofs*, volume 203 of *Lecture Notes in Pure and Applied Math.*, pages 77–95. Dekker, New York, 1999.

[73] Jose Iovino. Stable Banach spaces and Banach space structures ii. In *Models, algebras, and proofs*, volume 203 of *Lecture Notes in Pure and Applied Math.*, pages 97–117. Dekker, New York, 1999.

[74] Jose Iovino. *Applications of Model Theory to Functional Analysis*. Escuela Venezolana de Matemáticas, 2002.

[75] T. Jech. *Multiple Forcing*, volume 88 of *Cambridge Topics in Mathematics*. Cambridge University Press, 1987.

[76] B. Jónsson. Universal relational systems. *Mathematica Scandinavica*, 4:193–208, 1956.

[77] B. Jónsson. Homogeneous universal relational systems. *Mathematica Scandinavica*, 8:137–142, 1960.

[78] M. Kaufmann. The quantifier 'there exists uncountably many' and some of its relatives. In J. Barwise and S. Feferman, editors, *Model-Theoretic Logics*, pages 121–177. Springer-Verlag, 1985.

[79] H.J. Keisler. Logic with quantifier "there exists uncountably many". *Annals of Math. Logic*, 1:1–93, 1970.

[80] H.J Keisler. *Model theory for Infinitary Logic*. North-Holland, 1971.

[81] H. Kierstead. Countable models of ω_1-categorical theories in admissible languages. *Annals of Pure and Applied Logic*, 19:127–175, 1980.

[82] Byunghan Kim. Forking in simple theories. *Journal of the London Mathematical Society*, 57:257–267, 1998.

[83] Jonathan Kirby. Definability of quasiminimal excellent classes. `http://arxiv.org/PS_cache/arxiv/pdf/0707/0707.4496v3.pdf`, 2007.

[84] Jonathan Kirby. Abstract elementary categories. `http://people.maths.ox.ac.uk/~kirby/pdf/aecats.pdf`, 2008.

[85] J.F. Knight. A complete $l_{\omega_1,\omega}$-sentence characterizing $\aleph_1$. *Journal of Symbolic Logic*, 42:151–161, 1977.

[86] J.F. Knight. Prime and atomic models. *Journal of Symbolic Logic*, 43:385–393, 1978.

[87] Alexei Kolesnikov. Dependence relations in non-elementary classes. In *Logic and its Applications*, Contemporary Mathematics, pages 203–231. AMS, 2005.

[88] Oren Kolman and Saharon Shelah. Categoricity of theories in $L_{\kappa\omega}$, when κ is a measureable cardinal, part I. *Fundamenta Mathematica*, 151:209–240, 1996.

[89] D. W. Kueker. AEC and infinitary logic. preprint.

[90] D. W. Kueker. Back and forth arguments and infinitary model theory. In D. W. Kueker, editor, *Infinitary Logic: In Memoriam Carol Karp*, pages 17–72. Springer-Verlag, 1970. LNM 492.

[91] D. W. Kueker. Uniform theorem in infinitary logic. In A. Macintyre, L. Pacholski, and J. Paris, editors, *Logic Colloquium 77*. North Holland, 1978.

[92] K. Kunen. *Set Theory, An Introduction to Independence Proofs*. North Holland, 1980.

[93] A.H. Lachlan. A property of stable theories. *Fundamenta Mathematicae*, 77:9–20, 1972.

[94] D. Lascar. *Stability in Model Theory*. Longman, 1987. originally published in French as Stabilité en Théorie des Modèles (1986).

[95] Michael C. Laskowski. An old friend revisited: Countable models of ω-stable theories. *Notre Dame Journal of Formal Logic*, 48:133–141, 2007.

[96] Michael C. Laskowski and Saharon Shelah. On the existence of atomic models. *J. Symbolic Logic*, 58:1189–1194, 1993.

[97] Olivier Lessmann. Ranks and pregeometries in finite diagrams. *Journal of Pure and Applied Logic*, 106:49–83, 2000.

[98] Olivier Lessmann. Homogeneous model theory: Existence and categoricity. In Yi Zhang, editor, *Logic and Algebra*, pages 165–204. AMS, 2002. Contemporary Mathematics 302.

[99] Olivier Lessmann. Categoricity and U-rank in excellent classes. *Journal of Symbolic Logic*, pages 1317–1336, 2003.

[100] Olivier Lessmann. An introduction to excellent classes. In Yi Zhang, editor, *Logic and its Applications*, Contemporary Mathematics, pages 231–261. American Mathematical Society, 2005.

[101] Olivier Lessmann. Upward categoricity from a successor cardinal for an abstract elementary class with amalgamation. *Journal of Symbolic Logic*, 70:639–661, 2005. MR2140051.

[102] P. Lindstrom. On extensions of elementary logic. *Theoria*, 35:1–11, 1969.

[103] E. Lopez-Escobar. On definable well-orderings. *Fundamenta Mathematica*, 59:13–21, 299–300, 1966.

[104] M. Makkai. An ‘admissible generalization of a theorem on countable σ_1^1 sets of reals with applications. *Annals of Mathematical Logic*, 11:1–30, 1977.

[105] Michael Makkai and Saharon Shelah. Categoricity of theories in $l_{\kappa\omega}$, with κ a compact cardinal. *Annals of Pure and Applied Logic*, 47:41–97, 1990.

[106] J. A. Makowsky. Abstract embedding relations. In J. Barwise and S. Feferman, editors, *Model-Theoretic Logics*, pages 747–792. Springer-Verlag, 1985.

[107] L. Marcus. A minimal prime model with an infinite set of indiscernibles. *Israel Journal of Mathematics*, 11:180–183, 1972.

[108] D. Marker. *Model Theory: An introduction*. Springer-Verlag, 2002.

[109] D. Marker. A remark on Zilber's pseudo-exponentiation. *The Journal of Symbolic Logic*, 71:791–798, 2006.

[110] M. Morley. Categoricity in power. *Transactions of the American Mathematical Society*, 114:514–538, 1965.

[111] M. Morley. Omitting classes of elements. In Addison, Henkin, and Tarski, editors, *The Theory of Models*, pages 265–273. North-Holland, Amsterdam, 1965.

[112] M. Morley and R.L. Vaught. Homogeneous universal models. *Mathematica Scandinavica*, 11:37–57, 1962.

[113] N. Peatfield and B. Zilber. Analytic Zariski structures and the Hrushovski construction. preprint, 2003.

[114] A. Pillay. Forking in existentially closed structures. In A. Macintyre, editor, *Connections between Model Theory and Algebraic and Analytic Geometry*, quaderni di matematica. University of Naples, 2001.

[115] M. Pourmahdian and F. Wagner. A simple positive Robinson theory with $Lstp \neq Stp$. In Ali Enayat, Iraj Kalantari, and Mojtaba Moniri, editors, *Logic in Tehran Proceedings of the workshop and conference on logic, algebra, and arithmetic, held October 18-22, 2003*, Lecture Notes in Logic, 26. Association of Symbolic Logic, 2006.

[116] J. Rosenstein. *Linear Orderings*. Academic Press, New York, 1982.

[117] G. Sacks. *Saturated Model Theory*. Benjamin, Reading, Mass., 1972.

[118] S. Shelah. Amalgamation may be late in failing. announced.

[119] S. Shelah. Finite diagrams stable in power. *Annals of Mathematical Logic*, 2:69–118, 1970.

[120] S. Shelah. Categoricity in $\aleph_1$ of sentences in $L_{\omega_1,\omega}(Q)$. *Israel Journal of Mathematics*, 20:127–148, 1975. paper 48.

[121] S. Shelah. The lazy model-theoretician's guide to stability. *Logique et Analyse*, 18:241–308, 1975.

[122] S. Shelah. *Classification Theory and the Number of Nonisomorphic Models*. North-Holland, 1978.

[123] S. Shelah. *Proper Forcing*. Springer-Verlag, 1980. Springer Lecture Notes 940.

[124] S. Shelah. Classification theory for nonelementary classes. I. the number of uncountable models of $\psi \in L_{\omega_1\omega}$ part A. *Israel Journal of Mathematics*, 46:3:212–240, 1983. paper 87a.

[125] S. Shelah. Classification theory for nonelementary classes. I. the number of uncountable models of $\psi \in L_{\omega_1\omega}$ part B. *Israel Journal of Mathematics*, 46;3:241–271, 1983. paper 87b.

[126] S. Shelah. *Classification Theory and the Number of Nonisomorphic Models*. North-Holland, 1991. second edition.

[127] S. Shelah. *Proper Forcing*. Springer-Verlag, 1998. 2nd edition.

[128] S. Shelah. Categoricity for abstract classes with amalgamation. *Annals of Pure and Applied Logic*, 98:261–294, 1999. paper 394. Consult Shelah for post-publication revisions.

[129] S. Shelah. Completing 394. privately circulated F659, 200?

[130] S. Shelah. On categorical A.E.C. paper 734 in preparation, 200?

[131] S. Shelah. Toward classification theory of good λ frames and abstract elementary classes. preprint 705, 200?

[132] S. Shelah. On what I do not understand (and have something to say), model theory. *Mathematica Japonica*, 51:329377, 2000.

[133] S. Shelah. Categoricity of abstract elementary class in two successive cardinals. *Israel Journal of Mathematics*, 126:29–128, 2001. paper 576.

[134] S. Shelah. Categoricity of abstract elementary classes: going up inductive step. preprint 600, 200x.

[135] S. Shelah and A. Usvyatsov. Model theoretic stability and categoricity for complete metric spacess.

[136] S. Shelah and A. Usvyatsov. Banach spaces and groups - order properties and universal models. *Israel J Math*, 152:245–270, 2006.

[137] S. Shelah and A. Usvyatsov. Banach spaces and groups - order properties and universal models. *Israel J. Math.*, 152:245–270, 2006.

[138] Saharon Shelah. Abstract elementary classes near $\aleph_1$ sh88r. revision of Classification of nonelementary classes II, Abstract elementary classes; on the Shelah archive.

[139] Saharon Shelah. Classification of nonelementary classes II, abstract elementary classes. In J.T. Baldwin, editor, *Classification theory (Chicago, IL, 1985)*, pages 419–497. Springer, Berlin, 1987. paper 88: Proceedings of the USA–Israel Conference on Classification Theory, Chicago, December 1985; volume 1292 of *Lecture Notes in Mathematics*.

[140] Saharon Shelah. Universal classes, part I. In J.T. Baldwin, editor, *Classification theory (Chicago, IL, 1985)*, pages 264–419. Springer, Berlin, 1987. paper 300:Proceedings of the USA–Israel Conference on Classification Theory, Chicago, December 1985; volume 1292 of *Lecture Notes in Mathematics*.

[141] Saharon Shelah and Andrés Villaveces. Toward categoricity for classes with no maximal models. *Annals of Pure and Applied Logic*, 97:1–25, 1999. revised version : `www.math.rutgers.edu/~shelah`.

[142] J. Trlifaj. Cotilting classes as abstract elementary classes of finite character. preprint.

[143] B. L. Van der Waerden. *Modern Algebra*. Frederick Ungar Publishing, 1949.

[144] Monica VanDieren. Categoricity in abstract elementary classes with no maximal models. *Annals of Pure and Applied Logic*, 97:108–147, 2006.

[145] R.L. Vaught. Denumerable models of complete theories. In *Infinitistic Methods, Proc. Symp. Foundations of Math., Warsaw, 1959*, pages 303–321. Państwowe Wydawnictwo Naukowe, Warsaw, 1961.

[146] B.I. Zilber. Analytic Zariski structures, predimensions and non-elementary stability. preprint, 200?

[147] B.I. Zilber. Exponential sums equations and the schanuel conjecture. *J. London Math. Soc.*, 2002. includes preprint: Intersecting varieties with tori.

[148] B.I. Zilber. Model theory, geometry and arithmetic of universal covers of a semi-abelian variety. In *Model Theory and Applications*, Quaterna di matematica, pages 427–458, 2003.

[149] B.I. Zilber. Pseudo-exponentiation on algebraically closed fields of characteristic 0. *Annals of Pure and Applied Logic*, 132:67–95, 2004.

[150] B.I. Zilber. A categoricity theorem for quasiminimal excellent classes. In *Logic and its Applications*, Contemporary Mathematics, pages 297–306. AMS, 2005.

[151] B.I. Zilber. Covers of the multiplicative group of an algebraically closed field of characteristic 0. *J. London Math. Soc.*, pages 41–58, 2006.

Index