Topological and Category-Theoretic Aspects of Abstract Elementary Classes

by

Michael Joseph Lieberman

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in The University of Michigan
2009

Doctoral Committee:

Professor Andreas R. Blass, Chair
Emeritus Professor Peter G. Hinman
Associate Professor James P. Tappenden
Assistant Professor François G. Dorais
Assistant Professor Carl B. Mummert

To my family, and to the memory of my grandmother, Irene Taylor.

ACKNOWLEDGEMENTS

I would like to express my gratitude to my advisor, Professor Andreas Blass, who first suggested that the interplay between abstract elementary classes and accessible categories might be a fruitful area of inquiry, and was happy to discuss what seemed initially to be an unrelated pet project of mine—topology and rank functions for Galois types—which now comprises the first sections of this thesis.

I also owe a debt to Professor Alexei Kolesnikov. Had it not been for a tossed-off question in one of his seminar talks, I might never have turned my attention to the problem of topologizing Galois types, nor, were it not for his continued interest and guidance, would the initial seed have borne the fruit on display here. I would also like to acknowledge the support and guidance I have received from Professor John Baldwin, who has been a tireless reader of the individual components that came to constitute this thesis.

Thanks are also due to the members of my committee—Professors Peter Hinman, James Tappenden, François Dorais, and Carl Mummert—and to the faculty and staff of the Department of Mathematics at the University of Michigan.

Last, but not least, I thank my parents, Jay and Sally, my sister, Hannah, and my grandparents (indeed, the rest of my extended family) for their ceaseless support and encouragement, and acknowledge the largely non-academic (but nonetheless essential) contributions of my friends and colleagues: my longtime interlocutor and musical co-conspirator Godfre Leung, whose influence on my intellectual and musical life cannot be overstated; Mark Iwen and Craig Spencer, friends with whom I happily whiled away many of my non-mathematical hours; Ryan Kinser and Kelly McCoy, for their unstinting hospitality and companionship; and Dave Allen, Dave Anderson, Arvind Baskaran, Jonathan Bober, Hualong Feng, Leo Goldmakher, Jose Gómez, Jasun Gong, Hester Graves, Wendy Grus, Marc Krawitz, Felipe Ramirez, Sourya Shrestha, Daniel Spoth, Liz Vivas, and many others, for making my experience at Michigan such a pleasant and interesting one.

TABLE OF CONTENTS

CHAPTER 1

Introduction

Classical model theory is the study of elementary classes: classes of structures which, like groups, rings, and partially ordered sets, can be characterized by theories in first-order logic; that is, by sets of first-order sentences. While the field might be said, in retrospect, to have had its beginnings in the downward Löwenheim-Skolem theorem (1915), the completeness and compactness theorems (1930), or, more conventionally, in the work of Tarski in the 50s, it did not emerge in its modern form until 1965, with the advent of Morley's Theorem ([21]). This result settled the question of categoricity for countable languages: if an elementary class (theory) in a countable language is categorical in one uncountable cardinality—if there is, up to isomorphism, only a single model of that size—then the class is categorical in all uncountable cardinalities. To put it another way, we associate with each theory T a spectrum $I(T,\kappa)$, with κ ranging over the infinite cardinals, which tells us for each κ the number of isomorphism classes of models of size κ. Morley's Theorem gives all possible spectra for countable categorical theories: they are categorical only in $\aleph_0$, or in precisely the uncountable cardinals, or in all cardinals. Using the machinery of stability theory, Saharon Shelah proved a generalization to uncountable theories, where categoricity in any cardinal above the cardinality of the theory implies categoricity in all such cardinals. The next natural question, "If T isn't categorical in κ, how bad can $I(T, \kappa)$ be?", spawned Shelah's project of classification theory (more or less completed for countable theories in [24]) which has been an overwhelming preoccupation of first-order model theory ever since.

There are reasons to move in a slightly different direction, though. While first-order model theory has wide-ranging applications in mainstream mathematics (for an early example, see Ax and Kochen's work on Diophantine equations over local fields, distilled in [19]), there are a number of perfectly natural mathematical structures that it cannot properly analyze, by virtue of the fact that they cannot be satisfactorily characterized in first-order logic. Naturally, any abstract theory of mathematical structure that fails to encompass Banach spaces and the complex numbers with exponentiation (to take two key examples, the the former being the subject of work by Itai Ben Yaakov, C. Ward Henson, and Jose Iovino, and the latter being the particular province of Boris Zilber, beginning in [27]) is poorer than one might like. We are led, then, to consider the model theory of nonelementary classes—those that must be described in, say, an infinitary logic (formulas may contain infinite conjunctions and disjunctions, or infinitely many quantifiers), or a logic incorporating the quantifier Q ("there exist uncountably many"). The driving questions here are similar to the classical case: can we obtain results about the categoricity spectra analogous to Morley's Theorem? What can we say about the proliferation of types over sets of each cardinality; that is, the stability spectra? To what extent do the tools and results of first-order stability theory generalize to these contexts? To address these logic-specific questions all at once, we consider classes of structures that simultaneously generalize classes of models of theories in all of the above logics. Of the notions currently in circulation, abstract elementary classes—AECs—seem to exhibit the best balance of generality and richness of structure.

Although a thorough treatment of the definitions and basic facts that we need in connection with AECs will be provided in Section 2.1, we give the roughest of introductions here. As we will doggedly maintain in the sequel, one would not go far wrong in thinking of AECs as the category-theoretic hulls of elementary classes: we abandon syntax entirely, but retain the essential properties of the elementary submodel relation. In particular, an AEC is a class of structures $\mathcal{K}$ in a given language, closed under isomorphism, that comes equipped with a strong substructure relation $\prec_{\mathcal{K}}$ (a partial order on $\mathcal{K}$), and satisfies (among other axioms—see [2], or Chapter 2 below):

- Closure under unions of $\prec_{\mathcal{K}}$-increasing chains.
- Coherence: If $M_0 \prec_{\mathcal{K}} M_2$ and $M_0 \subseteq M_1 \prec_{\mathcal{K}} M_2$, then $M_0 \prec_{\mathcal{K}} M_1$.
- Löwenheim-Skolem: There exists a cardinal $\mathrm{LS}(\mathcal{K})$ such that for any $M \in \mathcal{K}$ and subset $A \subseteq M$, there is an $M_0 \in \mathcal{K}$ with $A \subseteq M_0 \prec_{\mathcal{K}} M$ and $|M_0| \leq |A| + \mathrm{LS}(\mathcal{K})$.

An elementary class equipped with the elementary submodel relation is an AEC, as are classes of models of sentences of $L_{\omega_1,\omega}$, $L(Q)$, $L_{\omega_1,\omega}(Q)$, and so on, given a suitable choice of $\prec_{\mathcal{K}}$. Hence this is an ideal context in which to seek, among other things, categoricity and stability spectra for the nonelementary classes we have chosen to investigate. Accessible categories—the category theorist's answer to the need for such a context—are still more general but, as we will see, nonetheless provide us with useful information in the relevant special cases.

To discuss stability (and stability spectra) in the context of AECs, we require a notion of type. Given the purely algebraic nature of AECs, there is no chance of resorting to the syntactic types to which logicians are most accustomed. One defines an alternative notion—Galois type—as follows: in any AEC $\mathcal{K}$ with the amalgamation property, there is a monster model $\mathfrak{C}$. For all intents and purposes, we may regard any structure in $\mathcal{K}$ as being a substructure of $\mathfrak{C}$. A Galois (1-)type over a structure M is the orbit of an element of $\mathfrak{C}$ under the automorphisms of $\mathfrak{C}$ that fix M (the analogy with Galois theory is instructive). An AEC is said to be λ-Galois stable if there are at most λ Galois types over each structure of size λ. This notion is certainly consistent with the syntax-free spirit of AECs and, as it happens, one can still employ a near-classical toolkit of topological methods and rank functions to analyze Galois types and Galois stability.

In Chapter 3, I define, for each structure M and cardinal λ, a topology on the set of Galois types over M in which types over small substructures of M (i.e. size at most λ) take on the role played by formulas in topologizing sets of syntactic types. In particular, a basic clopen set consists of all types over M that extend a given type over a small substructure. The topological properties of the resulting space, X_M^λ, are closely connected to the model-theoretic properties of M and $\mathcal{K}$. Critically, the property of χ-tameness of Galois types in a class $\mathcal{K}$—for any distinct types p and q over a structure M, there is a substructure $M' \prec_{\mathcal{K}} M$ of size at most χ on which their restrictions differ—emerges as a separation principle, namely the Hausdorff axiom, for the appropriate spaces. Beyond giving a new way of understanding purely model-theoretic properties of AECs, these spaces, peculiar as they are (the intersection of λ open sets is open in X_M^λ, for example, meaning that it is astonishingly difficult to be an accumulation point, and that none of the X_M^λ are compact), immediately give rise to Cantor-Bendixson ranks, and motivate the definition of a family of Morley-like ranks.

We give the definitions of these rank functions in Chapter 4, and provide a few of the most interesting consequences of their application to questions involving Galois stability. Roughly speaking, and recalling the analogy between formulas and Galois types over small structures, we first define an ordinal-valued Morley-like rank RM^λ (for each cardinal λ) on "formulas," and then define it on "types" as a minimum of the ranks of the constituent "formulas." The RM^λ are monotonic, invariant under automorphisms of $\mathfrak{C}$, satisfy a (slightly weak) unique extension property, and relate nicely to the topological structure: for example, if RM^λ is ordinal valued on the types over M, isolated points are dense in X^λ_M. If RM^λ is ordinal-valued on all types in an AEC $\mathcal{K}$, we say that $\mathcal{K}$ is λ-totally transcendental, a notion that is the key to a series of results concerning Galois stability in tame AECs. We begin with a proof that stability in a cardinal λ satisfying $\lambda^{\aleph_0} > \lambda$ implies λ-total transcendence and, by means of a series of results that use total transcendence to estimate the numbers of Galois types over models, translate this into several upward stability transfer results. The most interesting of these results, Theorem 4.26, establishes, among other things, that a class $\mathcal{K}$ that is stable in a cardinal λ satisfying the condition above is stable in any μ such that $\mathrm{cf}(\mu) > \lambda$ provided that $\mathcal{K}$ is stable on an interval just below μ. This is a significant generalization of a state-of-the-art result from [5]. In fact, the theorem itself has a number of generalizations, many of which are of interest in their own right. The most general form requires only weak tameness (the tameness condition is required to hold only for types over saturated models), and a weaker form of stability below κ which, although less than natural in a model theoretic sense, does crop up naturally in a category theoretic context.

This is the sole thread remaining to be woven into the picture. As we have noted above and will note again, abstract elementary classes represent, in a sense, the most recent stage of a long, slow drift toward category-theoretic viewpoints and methods within abstract model theory. There is an answering movement from the other camp: the last two decades have seen a new subfield—categorial model theory—emerge from the work of a small group of individuals whose recent work is of a primarily category-theoretic bent: Michael Makkai and Robert Paré (see, in particular, [20]), and Jiří Adamek and Jiří Rosický (see [1]). Their investigations have led them to the notion of accessible category, and have brought them within a stone's throw of those working in our corner of abstract model theory. Indeed, if one considers the accessible categories with directed colimits described in [22] (as we will, briefly here and in detail in Chapter 5), it seems that we have at last found a point at which model theorists and category theorists may shake hands. That the deep affinity between AECs and accessible categories has as yet gone unremarked upon outside of the current author's work is surprising, and we will devote Chapter 5 to remedying this situation.

The barest of outlines: recall that a Scott domain is a poset that contains a set of compact elements (x is compact if and only if for every directed system $\{x_i \,|\, i \in D\}$, $x \leq \bigvee_{i\in D} x_i$ only if $x \leq x_i$ for some $i \in D$), is closed under directed joins, and has the property that every element can be obtained as a directed join of compact elements. A λ-accessible category (λ a regular cardinal) is a generalization: it contains a set of λ-presentable objects (X is λ-presentable if and only if each map from X into a λ-directed colimit factors through one of the cocone maps, and the factorization through any such map is essentially unique, in the sense that any two factorizations through a particular cocone map are eventually identified in the λ-directed system), it is closed under λ-directed colimits, and every object can be obtained as a λ-directed colimit of λ-presentables. For any AEC $\mathcal{K}$, we regard it as a category in the only natural way: the objects are the structures in $\mathcal{K}$ and the morphisms are the $\prec_{\mathcal{K}}$-embeddings. The Löwenheim Skolem property for AECs guarantees that every structure is the $\mathrm{LS}(\mathcal{K})^+$-directed union of its substructures of size $\mathrm{LS}(\mathcal{K})$, and closure under unions of chains implies, by an easy exercise in category theory, that $\mathcal{K}$ is, in fact, closed under all directed colimits. In particular, then, any AEC $\mathcal{K}$ is $\mathrm{LS}(\mathcal{K})^+$-accessible and, as we will see, a structure $M \in \mathcal{K}$ is μ-presentable for regular $\mu > \mathrm{LS}(\mathcal{K})$ if and only if $|M| < \mu$.

In [22], Jiří Rosický devotes considerable attention to accessible categories with directed colimits (without mentioning AECs) and, in so doing, arrives at several notions and results that are of interest in our context. The full details of the correspondence between his category-theoretic notions and the model theoretic properties of AECs can be found in Chapter 5—here we focus only on what he calls weak λ-stability. For our purposes, if an AEC with amalgamation $\mathcal{K}$ is weakly λ-stable, then any $M \in \mathcal{K}$ of cardinality λ has a saturated extension of size λ. As it happens, this is precisely the weaker stability condition required in the stability transfer result for weakly tame AECs mentioned above. The remarkable thing is that any accessible category is weakly stable (in the sense of Rosický) in infinitely many, and arbitrarily large, cardinals. We postpone a detailed discussion of the result until Chapter 5. Suffice it to say, for now, that our result, in conjunction with that of Rosický, gives the beginnings of a stability spectrum for weakly tame AECs.

We close with an astonishing result which, although merely a curiosity at present, may have the potential to help in the analysis of categorical AECs—perhaps even in generating new categoricity transfer results. Taking a key result from [22] and translating into the AEC context, we will see that in any categorical AEC (categorical in cardinality λ, say), the sub-AEC consisting of structures of cardinality at least λ is equivalent (in the category-theoretic sense) to a highly structured subcategory of the category of sets with actions of the monoid of endomorphisms of the categorical model in $\mathcal{K}$; that is, the monoid $M = \mathrm{Hom}_{\mathcal{K}}(K, K)$, where K is the unique model in $\mathcal{K}$ of size λ. The straightforwardness of the correspondence between models and M-sets suggests that this peculiar representation theorem may be put to some use in the future, if only as a translation of familiar—but hard—problems in categorical AECs into radically new ones involving what are, theoretically at least, far simpler algebraic entities.

CHAPTER 2

Preliminaries

We begin with a brief tour of the background material—basic definitions and results—which will be used throughout this thesis. Section 2.1 lays out the model-theoretic context: abstract elementary classes, the associated notion of Galois type, and the host of properties thereof that will be invoked freely beginning in Chapter 3. Readers already familiar with these details may wish to skip to Section 2.2, which introduces accessible categories—material that will be essential in understanding the proceedings in Chapter 5.

2.1 AECs and Galois Types: Definitions and Basic Results

Naturally this exposition will not be exhaustive; readers interested in further details may wish to consult [2], and, for a presentation emphasizing the context in which to situate these details, [10]. As suggested in the introduction, AECs can be seen as a fundamentally category-theoretic generalization of elementary classes, where we excise syntactic considerations, and retain the purely diagrammatic, category-theoretic properties of the elementary submodel relation. In particular:

Definition 2.1. Let L be a fixed finitary signature (one-sorted, for simplicity). A class of L-structures equipped with a strong submodel relation, $(\mathcal{K}, \prec_{\mathcal{K}})$, is an *abstract elementary class (AEC)* if it satisfies the following axioms:

A0 The relation $\prec_{\mathcal{K}}$ is a partial order.

A1 For all M, N in $\mathcal{K}$, if $M \prec_{\mathcal{K}} N$, then $M \subseteq_L N$.

A2 (Closure Under Isomorphism)

1. If $M \in \mathcal{K}$, N is an L-structure, and $M \cong_L N$, then $N \in \mathcal{K}$.
2. If $M_i, N_i \in \mathcal{K}$ for $i = 1, 2$, and there are L-structure isomorphisms $f_i : M_i \to N_i$ with $f_1 \subset f_2$, then if $N_1 \prec_{\mathcal{K}} N_2$, $M_1 \prec_{\mathcal{K}} M_2$.

A3 (Unions of Chains) Let $(M_\alpha | \alpha < \delta)$ be a continuous $\prec_{\mathcal{K}}$-increasing sequence.

1. $\bigcup_{\alpha<\delta} M_\alpha \in \mathcal{K}$.
2. For all $\alpha < \delta$, $M_\alpha \prec_{\mathcal{K}} \bigcup_{\alpha<\delta} M_\alpha$.
3. If $M_\alpha \prec_{\mathcal{K}} M$ for all $\alpha < \delta$, then $\bigcup_{\alpha<\delta} M_\alpha \prec_{\mathcal{K}} M$.

A4 (Coherence) If $M_0, M_1 \prec_{\mathcal{K}} M$ in $\mathcal{K}$, and $M_0 \subseteq_L M_1$, then $M_0 \prec_{\mathcal{K}} M_1$.

A5 (Downward Löwenheim-Skolem) There exists an infinite cardinal $\mathrm{LS}(\mathcal{K})$ with the property that for any $M \in \mathcal{K}$ and subset A of M, there exists $M_0 \in \mathcal{K}$ with $A \subseteq M_0 \prec_{\mathcal{K}} M$ and $|M_0| \leq |A| + \mathrm{LS}(\mathcal{K})$.

The prototypical example, of course, is the case in which $\mathcal{K}$ is an elementary class—the class of models of a particular first order theory T—and $\prec_{\mathcal{K}}$ is the elementary submodel relation. Given that the axioms above were designed to capture the fundamental properties of elementary classes and elementary submodel, it comes as no surprise that $(\mathcal{K}, \prec_{\mathcal{K}})$ is an AEC, with $\mathrm{LS}(\mathcal{K}) = |T|$. Naturally, the notion of AEC is far more general. Given a sentence $\phi \in L_{\infty,\omega}$, $\mathcal{K} = \mathrm{Mod}(\phi)$, and $\prec_{\mathcal{K}}$ the relation of elementarity with respect to a fragment $\mathcal{A}$ of $L_{\infty,\omega}$ containing ϕ, $(\mathcal{K}, \prec_{\mathcal{K}})$ is an AEC with $\mathrm{LS}(\mathcal{K}) = |\mathcal{A}|$. Similarly for models of sentences in $L(Q)$, provided we are careful in our choice of $\prec_{\mathcal{K}}$, in which case we have $\mathrm{LS}(\mathcal{K}) = \aleph_1$. Notice that, unlike in the elementary case, classes arising from such non-first-order logics need not contain arbitrarily large models, even if they contain models in some infinite cardinalities. For more concrete but nonetheless nonelementary examples: Artinian commutative rings with unit form an AEC (under elementary submodel), as do the classes of modules described in [3]. The class of Banach spaces is not quite an AEC, not being closed under unions of countable chains, but the class of Banach spaces and their metric subspaces (that is, the normed linear spaces) is an AEC.

We return to the task of laying out the basic terminology and notation. For any infinite cardinal λ, we denote by $\mathcal{K}_\lambda$ the subclass of $\mathcal{K}$ consisting of all models of cardinality λ (with the obvious interpretations for such notations as $\mathcal{K}_{\leq\lambda}$ and $\mathcal{K}_{>\lambda}$). We say that $\mathcal{K}$ is λ-categorical if $\mathcal{K}_\lambda$ contains only a single model up to isomorphism. For $M, N \in \mathcal{K}$, we say that a map $f : M \to N$ is a $\mathcal{K}$-embedding (or, more often, a strong embedding) if f is an injective homomorphism of $L(\mathcal{K})$-structures, and $f[M] \prec_{\mathcal{K}} N$; that is, f induces an isomorphism of M onto a strong submodel of N. In that case, we write $f : M \hookrightarrow_{\mathcal{K}} N$. Notice that we may now think of $\mathcal{K}$ as a concrete category—a subcategory of the category of L-structures—an observation that first appears in [12], is dealt with explicitly in [18], and will be the focus of Section 5.1 of this thesis. In the meantime, we notice that the union of chains axiom, A3, gives us closure of $\mathcal{K}$ not merely under unions of increasing sequences, but also under colimits (direct limits) of chains of $\mathcal{K}$-embeddings. By an easy exercise in universal algebra (see, for example, Corollary 1.7 in [1]), this implies that $\mathcal{K}$ is closed under arbitrary directed colimits.

Definition 2.2. Let $\mathcal{K}$ be an AEC.

1. We say that an $\mathcal{K}$ has the *joint embedding property (JEP)* if for any $M_1, M_2 \in$

$\mathcal{K}$, there is an $M \in \mathcal{K}$ that admits strong embeddings of both M_1 and M_2, $f_i : M_i \hookrightarrow_{\mathcal{K}} M$ for $i = 1, 2$.

2. We say that an $\mathcal{K}$ has the *amalgamation property (AP)* if for any $M_0 \in \mathcal{K}$ and strong embeddings $f_1 : M_0 \hookrightarrow_{\mathcal{K}} M_1$ and $f_2 : M_0 \hookrightarrow_{\mathcal{K}} M_2$, there are strong embeddings $g_1 : M_1 \hookrightarrow_{\mathcal{K}} N$ and $g_2 : M_2 \hookrightarrow_{\mathcal{K}} N$ such that the following diagram commutes:

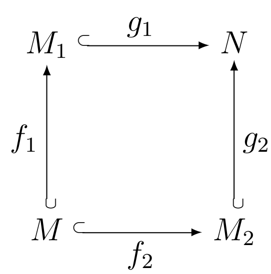

In category theoretic terms, the joint embedding property guarantees that the AEC contains cocones on all diagrams of the form ($\bullet \quad \bullet$), while the amalgamation property guarantees that it contains cocones on all diagrams of the form ($\bullet \leftarrow \bullet \rightarrow \bullet$). Notice that both properties hold in elementary classes, as a consequence of the compactness of first order logic. In this more general context, devised to subsume classes of models in logics without any compactness to fall back on, both appear as additional (and nontrivial) assumptions on the class. Indeed, it is often preferable to assume less, working with classes that, for example, admit amalgamation only in particular cardinalities. In our present investigations, however, we will assume the unrestricted amalgamation property described above, as well as the joint embedding property (unless otherwise indicated).

It is not immediately clear what we might embrace as a suitable notion of type in AECs, as we have dispensed with syntax, and removed ourselves to a world of abstract embeddings and diagrams thereof. The best candidate—the Galois type—has its origins in the work of Shelah (first appearing in [25] and [26]). The allusion to

Galois theory built into the terminology is instructive: just as the type of a complex number x over a given number field F (or, to be precise, its set of realizations) is the orbit of x under $\mathrm{Aut}_F(\mathbb{C})$, the group of automorphisms of $\mathbb{C}$ that fix F, we will come to identify Galois types (in AECs with amalgamation) with orbits in a large model under automorphisms that fix particular submodels. First, though, we present the general definition, which makes sense in any AEC.

Definition 2.3. For triples (M, a_1, N_1), and (M, a_2, N_2), where $M \prec_{\mathcal{K}} N_i$ and $a_i \in N_i$ for $i = 1, 2$, we say that $(M, a_1, N_1) \sim (M, a_2, N_2)$ if there are strong embeddings $f_i : N_i \hookrightarrow_{\mathcal{K}} N$ such that $f_1(a_1) = f_2(a_2)$ and f_1 and f_2 agree on M; that is, the following diagram commutes:

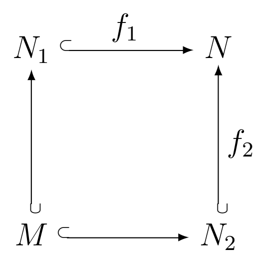

Our aim is to define Galois types as equivalence classes under $\sim$. Here we are faced with a potential difficulty: the relation $\sim$ need not be transitive for a general AEC. One could simply work with the transitive closure of $\sim$, but we instead call to the attention of the reader the easily-verified fact that the relation $\sim$ is transitive when restricted to the class of triples (M, a, N) where N is a model over which we are able to amalgamate. Since we are assuming AP, this is no restriction at all: $\sim$ itself is transitive and, moreover, an equivalence relation.

Definition 2.4. Let $M, N \in \mathcal{K}$, $M \prec_{\mathcal{K}} N$ and $a \in N$. The Galois type of a over M in N, denoted ga-tp$(a/M, N)$, is the set of all triples $\sim$-equivalent to (M, a, N).

A simpler and more concrete characterization of Galois types is possible in our context, but we first require a bit more model-theoretic background. In particular, we need a suitable notion of monster model.

Definition 2.5. 1. A model $M \in \mathcal{K}$ is *λ-model homogeneous* if for any $N \prec_{\mathcal{K}} M$ and $N' \in \mathcal{K}_{<\lambda}$ with $N \prec_{\mathcal{K}} N'$, there is an embedding of N' into M that fixes N. We say that M is *model homogeneous* if it is $|M|$-model homogenous.

2. A model $M \in \mathcal{K}$ is *strongly λ-model homogeneous* if it is λ-model homogeneous and for any $N, N' \prec_{\mathcal{K}} M$ with $|N|, |N'| < \lambda$, any isomorphism from N to N' extends to an automorphism of M. We say that M is *strongly model homogeneous* if it is strongly $|M|$-model homogeneous.

Assuming amalgamation and joint embedding, we are assured of a supply of large model homogeneous structures in $\mathcal{K}$. Under additional assumptions about cardinal arithmetic, the following result (Theorem 8.5 from [2]) suffices:

Theorem 2.6. *For any μ satisfying $\mu^{<\mu} = \mu$ and $\mu > 2^{LS(\mathcal{K})}$, there is a strongly model homogeneous (hence also model homogeneous) model of cardinality μ.*

In fact, this argument may be reworked so as to avoid any extra set-theoretic assumptions. A sketch of such a modified argument may be found in [2], immediately following the statement of the theorem above. As in first order logic, we fix a single such large strongly model homogeneous model—the monster model of $\mathcal{K}$—which we denote by $\mathfrak{C}$. We regard all models $M \in \mathcal{K}$ as submodels of $\mathfrak{C}$ with $|M| < |\mathfrak{C}|$. Now, we may effect the promised reinterpretation of the notion of Galois type. First:

Lemma 2.7. *Triples (M, a_1, N_1) and (M, a_2, N_2) are $\sim$-equivalent if and only if there is an automorphism of $\mathfrak{C}$ that fixes M pointwise and maps a_1 to a_2.*

Proof: ($\Rightarrow$) There exist embeddings $g_i : N_i \hookrightarrow_{\mathcal{K}} N$ for $i = 1, 2$ that agree on M, and satisfy $g_1(a_1) = g_2(a_2)$. Consider the map $g_2^{-1} \circ g_1$: it fixes M pointwise, and maps a_1 to a_2. By strong model homogeneity of $\mathfrak{C}$, there is an automorphism of $\mathfrak{C}$ that extends it.

($\Leftarrow$) Let f be such an automorphism. Let $N \prec_{\mathcal{K}} \mathfrak{C}$ be any model containing the set $N_2 \cup f[N_1]$. The embedding $f \restriction N_1 : N_1 \hookrightarrow_{\mathcal{K}} N$ and the inclusion $N_2 \hookrightarrow_{\mathcal{K}} N$ witness that $(M, a_1, N_1) \sim (M, a_2, N_2)$. $\square$

We may now give the simpler (but still equivalent) definition of Galois type:

Definition 2.8. Let $M \in \mathcal{K}$, and $a \in \mathfrak{C}$. The *Galois type of a over M*, denoted ga-tp(a/M), is the orbit of a in $\mathfrak{C}$ under $\mathrm{Aut}_M(\mathfrak{C})$, the group of automorphisms of $\mathfrak{C}$ that fix M. We denote by ga-S(M) the set of all Galois types over M.

In this interpretation, the above-indicated connection with Galois theory is somewhat clearer. In case $\mathcal{K}$ is an elementary class with $\prec_{\mathcal{K}}$ as elementary submodel, the Galois types over M correspond to the complete syntactic types over M: map each complete type to its set of realizations in $\mathfrak{C}$, and each orbit under $\mathrm{Aut}_M(\mathfrak{C})$ to the complete type over M of one of its members. In general, however, Galois types and syntactic types do not match up, even in cases when the logic underlying the AEC is clear (say, $\mathcal{K} = \mathrm{Mod}(\psi)$, with $\psi \in L_{\omega_1,\omega}$). A crucial definition:

Definition 2.9. We say that an AEC $\mathcal{K}$ is λ-*Galois stable* if for every $M \in \mathcal{K}_\lambda$, $|\text{ga-S}(M)| = \lambda$.

To put Galois stability in its proper context, we note that $|\text{ga-S}(M)| \geq |M|$ for all M, since each element $a \in M$ gives rise to a distinct orbit under automorphisms of $\mathfrak{C}$ fixing M (namely, the set $\{a\}$), hence also to a distinct Galois type over M. Moreover,

Proposition 2.10. *Let $\mathcal{K}$ be an AEC with amalgamation. For any $M \in \mathcal{K}$ with $|M| \geq LS(\mathcal{K}) + |L(\mathcal{K})|$, $|ga\text{-}S(M)| \leq 2^{|M|}$.*

Proof: Let $M \in \mathcal{K}$ be of cardinality $\lambda \geq \text{LS}(\mathcal{K}) + |L(\mathcal{K})|$, and let q be a Galois type over M. The type q is the orbit of some element $a \in \mathfrak{C}$ under $\text{Aut}_M(\mathfrak{C})$. By the Downward Löwenheim-Skolem Property, a is contained in a strong extension M' of M of cardinality λ. As this holds for all types over M, we have

$$|\text{ga-S}(M)| \leq |\{M' \in \mathcal{K}_\lambda \mid M' \succ_{\mathcal{K}} M\}| \leq |\mathcal{K}_\lambda|$$

Because the signature $L(\mathcal{K})$ is finitary and $\lambda \geq |L(\mathcal{K})|$, there can be at most 2^λ $L(\mathcal{K})$-structures of size λ, let alone elements of $\mathcal{K}_\lambda$. Hence $|\text{ga-S}(M)| \leq 2^\lambda$. $\square$

To summarize, for $M \in \mathcal{K}$ of sufficiently large size, $|M| \leq |\text{ga-S}(M)| \leq 2^{|M|}$. We run through a few more basic definitions and notations:

Definition 2.11. Let $\mathcal{K}$ be an AEC with amalgamation, and $\mathfrak{C}$ its monster model.

1. For any M, $a \in \mathfrak{C}$, and $N \prec_{\mathcal{K}} M$, the *restriction of ga-tp(a/M) to N*, denoted ga-tp$(a/M) \upharpoonright N$, is the orbit of a under $\text{Aut}_N(\mathfrak{C})$. This notion is well-defined: the restriction depends only on ga-tp(a/M), not on a itself.

2. Let $N \prec_{\mathcal{K}} M$ and $p \in \text{ga-S}(N)$. We say that M *realizes* p if there is an element $a \in M$ such that ga-tp$(a/M) \upharpoonright N = p$. Equivalently, M realizes p if the orbit in $\mathfrak{C}$ corresponding to p meets M.

3. We say that a model M is λ-*Galois-saturated* if for every $N \prec_{\mathcal{K}} M$ with $|N| < \lambda$ and every $p \in \text{ga-S}(N)$, p is realized in M.

Henceforth, the word "type" should be understood to mean "Galois type," unless otherwise indicated. Moreover, when there is no risk of confusion (namely in Chapters 3 and 4 below) we will omit the word "Galois" altogether, speaking simply of types, λ-stability, and λ-saturation. We retain it for the present, however. Now, we will have occasion to use the following fact (Theorem 8.14 in [2]), which holds under the assumption of AP:

Proposition 2.12. *For $\lambda > LS(\mathcal{K})$, a model $M \in \mathcal{K}$ is λ-Galois-saturated if and only if it is λ-model homogeneous.*

Proof: We sketch the argument, to give a sense of the extent to which certain classical arguments (back and forth, unions of chains) carry over to AECs. The "if" direction is easy: Suppose M is λ-model homogeneous, $N \prec_{\mathcal{K}} M$ with $|N| < \lambda$, and consider a type $p \in \text{ga-S}(N)$, say $\text{ga-tp}(a/N)$. We must check that p is realized in M. By the Löwenheim-Skolem Property (and the fact that $\lambda > \text{LS}(\mathcal{K})$) N has an extension $N' \prec_{\mathcal{K}} \mathfrak{C}$ with $N' \supseteq N \cup \{a\}$ and $|N'| < \lambda$. By λ-model homogeneity of M, there is an embedding $f : N' \hookrightarrow_{\mathcal{K}} M$ that fixes N. One can easily see that $\text{ga-tp}(f(a)/N) = \text{ga-tp}(a/N)$, meaning that $f(a)$ is the desired realization of p in M.

Now, suppose that M is λ-Galois-saturated, that $N \prec_{\mathcal{K}} M$, and $N' \succ_{\mathcal{K}} N$ with $|N'| = \mu < \lambda$. To establish the result, we must construct an embedding of N' into M over N. Enumerate $N' \setminus N$ as $\{a_i \,|\, i < \mu\}$. We construct an increasing sequence of maps $f_i : N_i \simeq M_i$, $i < \mu$, where $N \prec_{\mathcal{K}} N_i$, $N \prec_{\mathcal{K}} M_i \prec_{\mathcal{K}} M$, and $a_i \in N_{i+1}$ for each $i < \mu$. The union of the f_i, restricted to N', will be the desired embedding. Set $N_0 = N$, and let f_0 be the identity on N_0. Take unions at limit stages. Suppose we have defined $f_i : N_i \hookrightarrow_{\mathcal{K}} M_i$. If $a_i \in N_i$, define $N_{i+1} = N_i$, $M_{i+1} = M_i$, and $f_{i+1} = f_i$. Otherwise, proceed as follows: By strong model homogeneity of $\mathfrak{C}$, f_i extends to an automorphism $\bar{f}_i$ of $\mathfrak{C}$. Notice that the type of $\bar{f}_i(a_i)$ over $\bar{f}_i[N_i] = M_i \prec_{\mathcal{K}} M$ must be realized by some $b \in M$, by λ-saturation, in which case there is an automorphism $g \in \text{Aut}_{M_i}(\mathfrak{C})$ that takes $\bar{f}_i(a)$ to b. Let M_{i+1} be a strong submodel of M of cardinality λ that contains M_i and b, let $N_{i+1} = \bar{f}_i^{-1} \circ g^{-1}[M_{i+1}]$, and let $f_{i+1} = g \circ \bar{f}_i \restriction N_{i+1}$.
$\square$

As in the result above, we will typically work at or above $\text{LS}(\mathcal{K})$. Classical results linking stability to the existence of saturated models that arise from union of chains arguments also transfer to this context. For example:

Proposition 2.13. *If λ is a regular cardinal, $\lambda > LS(\mathcal{K})$, and for every $N \in \mathcal{K}_{<\lambda}$, $|ga\text{-}S(N)| < \lambda$, then every $M \in \mathcal{K}$ of cardinality λ has a saturated extension $M' \in \mathcal{K}$ which is also of cardinality λ.*

Proof: Let $M \in \mathcal{K}_\lambda$, and take a filtration of M of length λ; that is, a continuous $\prec_\mathcal{K}$-increasing sequence $(M_i \,|\, i < \lambda)$ with $M = \bigcup_{i<\lambda} M_i$ and $|M_i| < \lambda$. We define a $\prec_\mathcal{K}$-increasing sequence $(M'_i \,|\, i < \lambda)$ such that, for all $i < \lambda$, $M_i \prec_\mathcal{K} M'_i$ and M'_{i+1} realizes all types over M'_i. The model $M' = \bigcup_{i<\lambda} M'_i$ will be the desired extension. The construction proceeds as follows: set $M'_0 = M_0$. At stage $i+1$, let M''_i be a model of cardinality less than λ containing $M'_i \cup M_{i+1}$. By our assumption about the number of types over models in $\mathcal{K}_\lambda$, there is a model of cardinality less than λ that extends M''_i and realizes all types in ga-S(M_i). Make this M'_{i+1}. At limit stages, take the union: $M'_i = \bigcup_{j<i} M'_j$. By regularity of λ, the union M'_i will be of cardinality less than λ. That this works is clear: given any $N \prec_\mathcal{K} M' := \bigcup_{i<\lambda} M'_i$ of cardinality less than λ, $N \subseteq M'_i$ for some i by regularity of λ. By coherence (Axiom A4 of AECs), $N \prec_\mathcal{K} M'_i$, meaning that for any $p \in$ ga-S(N), all of its extensions to types over M'_i are realized in M'_{i+1}. As there is at least one such, p is itself realized in M'_i, and therefore in M'. $\square$

As in first order classification theory, one of the central preoccupations of those working with AECs is the identification of "dividing lines," in the sense of Shelah: properties of classes which, when present, guarantee nice structural properties that disappear—at least in particular examples—when one finds oneself on the wrong side of a dividing line, only to be replaced by new misbehaviors; that is, "nonstructure." Certain of these properties for AECs (excellence, as described in [11], and the existence of good or semi-good frames, as considered in [23] and [16], respectively) echo classical dividing lines (simplicity, stability, and superstability), but we here concern ourselves primarily with a property that has no particularly natural classical analogue, first arising in the middle of a proof in [26] as a mysterious but necessary hypothesis: tameness. Roughly speaking, tameness of an AEC means that types are completely determined by their restrictions to small submodels, a condition slightly reminiscent of the locality properties of syntactic types. A possible intuition (an important one, in the sense that it motivates many of the definitions in Chapters 3 and 4) is the following: we may regard types over small models as playing a role analogous to formulae in first order model theory, in which case tameness guarantees that types are determined by their constituent formulae.

As usual, there are multiple formulations. We begin with the most general version:

Definition 2.14. 1. We say that $\mathcal{K}$ is (χ,λ)-*tame* if for every $M \in \mathcal{K}_\lambda$, if p and p' are distinct types over M, then there is an $N \prec_{\mathcal{K}} M$, $|N| \leq \chi$, such that $p \upharpoonright N \neq p' \upharpoonright N$.

2. We say that $\mathcal{K}$ is *weakly* (χ,λ)-*tame* if the condition above holds for every saturated $M \in \mathcal{K}_\lambda$.

As one would expect, we say that a class is $(\chi, \leq \lambda)$-tame if it is (χ,κ)-tame for all $\kappa \leq \lambda$, and so on for other such variations. We say that a class is (χ,∞)-tame if it is (χ,λ)-tame for all λ. The case of (χ,∞)-tameness is of sufficient importance to warrant a notation of its own:

Definition 2.15. 1. We say that $\mathcal{K}$ is χ-*tame* if for every $M \in \mathcal{K}$, if p and p' are distinct types over M, there is an $N \prec_{\mathcal{K}} M$, $|N| \leq \chi$, such that $p \upharpoonright N \neq p' \upharpoonright N$.

2. We say that $\mathcal{K}$ is *weakly* χ-*tame* if the condition above holds for every saturated $M \in \mathcal{K}$.

A few words about the significance of the assumption of tameness: although it holds automatically in elementary classes (as well as in homogeneous classes, among other settings), it fails in some natural mathematical contexts—[4] and [6], for example, construct examples of non-tame AECs from Abelian groups. Moreover, it is independent from the other properties one associates with well-behaved AECs. It is established in [6], for example, that given an AEC with amalgamation, there is an AEC without amalgamation with precisely the same tameness spectrum. Its utility more than justifies the loss of generality involved in its assumption, though: it plays an indispensable role in establishing all of the existing upward categoricity transfer results (such as those of [13], [15], and [26]), and the downward transfer result included in Chapter 14 of [2] hinges on weak tameness. The dependence is, if anything, more pronounced in the analysis of the stability spectra of AECs: tameness appears as a hypothesis for nearly all of the partial stability spectrum results of [5], [14], and this thesis, although weak tameness occasionally suffices here and in [5]. At present, attempts to produce results of this nature for non-tame classes have not met with any success.

We close with a very brief introduction to two additional locality properties of Galois types (readers interested in a more detailed discussion than that provided here may wish to consult [2]).

Definition 2.16. Galois types in $\mathcal{K}$ are (κ, λ)-*local* if for every continuous $\prec_{\mathcal{K}}$-increasing sequence $(M_i \,|\, i < \kappa)$ with $M = \bigcup_{i<\kappa} M_i$, $|M| = \lambda$, if $q, q' \in \text{ga-S}(M)$ are distinct, there exists an $i < \kappa$ such that $q \upharpoonright M_i \neq q' \upharpoonright M_i$.

The final property we introduce is the following:

Definition 2.17. Galois types are (κ, λ)-*compact* in $\mathcal{K}$ if for every continuous $\prec_{\mathcal{K}}$-

increasing sequence $(M_i \,|\, i < \kappa)$ with $M = \bigcup_{i<\kappa} M_i$, $|M| = \lambda$, and increasing sequence of types $p_i \in M_i$ for $i < \kappa$, there exists a type $p \in \text{ga-S}(M)$ with $p \upharpoonright M_i = p_i$ for all $i < \kappa$.

In a sense, this notion provides a measure of the extent of compactness in AECs. Whereas types in elementary classes are (∞, ∞)-compact, by the compactness of first order logic, far less is true in a general AEC. It is true, though, that in any AEC with amalgamation, types are $(\aleph_0, \infty)$-compact (Theorem 11.1 in [2]). We will make use of this result in Section 4.2.

2.2 Accessible Categories

Of the basic properties that we retain in passing to abstract elementary classes from classes of structures born of syntactic considerations (classes of models of first order theories, sentences in $L_{\kappa,\omega}$, $L_{\omega_1,\omega}(Q)$, and so on), two stand out as being of particular importance. First, the union axioms ensure that the class is closed under unions of chains, giving us the structure needed to run certain nearly-classical model-theoretic arguments, a few examples of which we saw in the previous section. Moreover, the Downward Löwenheim-Skolem Property for AECs guarantees that any structure $M \in \mathcal{K}$ can be approximated by submodels of small size (can be obtained, in particular, as the directed union of its submodels of cardinality at most $\text{LS}(\mathcal{K})$) meaning that an AEC $\mathcal{K}$ is, in fact, generated from the set of all such small models, $\mathcal{K}_{\text{LS}(\mathcal{K})}$. Although accessible categories—the category theorists' preferred generalization of classes of structures, both elementary and nonelementary (see [1] and [20])—involve a slightly greater degree of abstraction and hence greater generality, they are also characterized by precisely these two traits: each accessible category is closed under certain highly directed colimits (if not arbitrary directed colimits), and is generated from a set of "small" objects.

To flesh out what we mean by "small," we require a notion of size that makes sense in an arbitrary category. Since, in particular, we do not wish to restrict ourselves to categories of structured sets, our notion will need to be more subtle than mere cardinality. The solution to this quandary—presentability—first appeared in [8], and has subsequently been treated in a more accessible fashion in [1] and [20], the latter being a particularly good source of concrete examples. We begin with the simplest and most mathematically natural case:

Definition 2.18. An object N in a category $\mathbf{C}$ is said to be *finitely presentable* if the corresponding hom-functor $\mathrm{Hom}_{\mathbf{C}}(N, -)$ preserves directed colimits.

Less cryptically, N is finitely presentable if for any directed poset I and diagram $D : (I, \leq) \to \mathbf{C}$ (that is, for any functor D whose domain is I regarded as a category, where the objects are precisely the elements of I and there is a morphism $i \to j$ precisely when $i \leq j$) with colimit cocone $(\phi_i : D(i) \to M)_{i \in I})$, any map $f : N \to M$ factors through one of the maps in the colimit cocone: $f = \phi_i \circ g$ for some $i \in I$ and $g : N \to D(i)$, as in the diagram below.

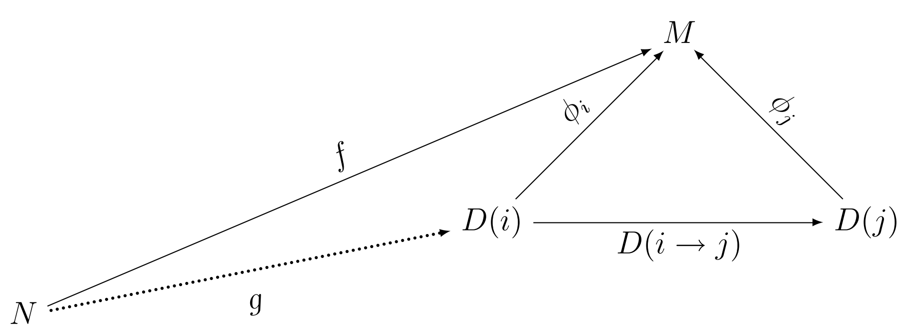

Moreover, this factorization must be essentially unique, in the sense that for any two such, say g and g' from N to $D(i)$ with $f = \phi_i \circ g = \phi_i \circ g'$, there is a $j \geq i$ in I such that $D(i \to j) \circ g = D(i \to j) \circ g'$.

Examples:

1. Consider the Scott posets mentioned in the introduction of this thesis. We may regard any such poset as a category (in the manner of I above) in which, as one can easily see, colimits correspond to joins, and vice versa. In this case, the defining characteristic of a compact element—x is compact if and only if whenever $x \leq \bigvee_{i \in I} x_i$ with the latter a directed join, $x \leq x_i$ for some $i \in I$—guarantees that we have precisely the kind of factorization described above, meaning that any compact element is finitely presentable. Clearly, the converse also holds.

2. In **Set**, the category of sets, an object X is finitely presentable if and only if it is a finite set.

3. Let Σ be a finitary relational signature, and $\mathbf{Rel(\Sigma)}$ the category of Σ-structures and maps that preserve the relations $R \in \Sigma$. An object M in $\mathbf{Rel(\Sigma)}$ is finitely presentable if and only if $|M|$ is finite and there are only finitely many Σ-edges in M: $\sum_{R \in \Sigma} |R^M| < \aleph_0$.

4. In **Grp**, the category of groups and group homomorphisms, an object G is finitely presentable if and only if it is finitely presented in the usual sense: G has finitely many generators subject to finitely many relations. To see that any finitely presented group G (with finite generating set $X \subseteq G$, say) is finitely presentable, let $G' = \mathrm{Colim}_{i \in I} G_i$ be a directed colimit in **Grp** with colimit cocone $(\phi_i : G_i \to G')$, and let $f : G \to G'$ be a group homomorphism. The set $f[X] \subseteq G'$ is finite and, since G' is a directed union of the images $\phi_i[G_i]$, we must have $f[X] \subseteq \phi_i[G_i]$ for some $i \in I$. For each $x \in X$, choose an element $a_x \in \phi_i^{-1}(f(x)) \subseteq G_i$, and define a map from X to G_i that takes each $x \in X$ to the associated a_x. As the generators of G are subject to only finitely many

relations, directedness of the colimit ensures that there exists a $j \geq i$ such that the images of the a_x under the diagram map $G_i \to G_j$ satisfy all the relations appearing in the presentation of G. The corresponding map of X into G_j extends to a homomorphism $g : G \to G_j$ with the property that $\phi_j \circ g = f$, as required. This factorization can be shown to be essentially unique. The converse is left as an exercise.

As shown in [1], the same holds in any variety of finitary algebras.

Many more examples can be found in [1]. One more word about the category **Grp**: every object of **Grp**—every group—can be obtained as the directed union (colimit) of its finitely generated subgroups, hence as a directed colimit of finitely presentable objects. Moreover, **Grp** is closed under arbitrary directed colimits. This means, in short, that **Grp** is a finitely accessible category. The precise definition:

Definition 2.19. A category **C** is *finitely accessible* if

- **C** contains only a set of finitely presentable objects up to isomorphism, and every object in **C** is a directed colimit of finitely presentable objects.
- **C** is closed under directed colimits.

Finitely accessible categories abound in mainstream mathematics: the category **Grp**; **Rel**($\boldsymbol{\Sigma}$), under the conditions described above; **Set**, the category of sets; and **Pos**, the category of posets and monotone functions. It should be noted that these examples are not merely closed under directed colimits, but are also cocomplete; that is, they are closed under all colimits. Categories satisfying this stronger closure condition are called locally finitely presentable but, given that most of the categories we consider in the sequel are not cocomplete, we will have little to say about this special case.

The notions of finite presentability and finite accessibility generalize in a natural fashion. Let λ be an infinite regular cardinal. We first recall:

Definition 2.20. 1. A poset I is said to be λ-*directed* if for every subset $X \subseteq I$ of cardinality less than λ, there is an element $i \in I$ such that for every $x \in X$, $x \leq i$.

2. A colimit in a category $\mathbf{C}$ is λ-directed if it is the colimit of a λ-directed diagram; that is, a diagram of the form $D : (I, \leq) \to \mathbf{C}$, where I is a λ-directed poset.

Generalizing finitely presentable objects, we define:

Definition 2.21. An object N is said to be λ-*presentable* if the corresponding functor $\mathrm{Hom}(N, -)$ preserves λ-directed colimits.

We may unravel this definition just as we did when considering finitely presentable objects: N is λ-presentable if for any λ-directed poset I and diagram $D : (I, \leq) \to \mathbf{C}$ with colimit cocone $(\phi_i : D(i) \to M)_{i \in I}$, any map $f : N \to M$ factors through one of the maps in the colimit cocone: $f = \phi_i \circ g$ for some $i \in I$ and some $g : N \to D(i)$ (as in the diagram following Definition 2.18 above). Moreover, this factorization must be essentially unique, in the same sense as before.

For any category $\mathbf{C}$ and infinite regular cardinal λ, we denote by $\mathbf{Pres}_\lambda(\mathbf{C})$ a full subcategory of $\mathbf{C}$ consisting of one representative of each isomorphism class of λ-presentable objects; that is, $\mathbf{Pres}_\lambda(\mathbf{C})$ is a skeleton of the full subcategory consisting of all λ-presentable objects.

One should note that it is customary—and sometimes advantageous—to phrase things in terms of λ-filtered (rather than λ-directed) diagrams and colimits, but the two characterizations are fundamentally equivalent. See, in particular, Remark 1.21 in [1]. Now, the crucial definition:

Definition 2.22. 1. Let λ be an infinite regular cardinal. A category $\mathbf{C}$ is λ-*accessible* if

- $\mathbf{C}$ is closed under λ-directed colimits
- $\mathbf{C}$ contains only a set of λ-presentable objects up to isomorphism, and every object in $\mathbf{C}$ is a λ-directed colimit of λ-presentables.

2. We say that a category $\mathbf{C}$ is *accessible* if it is λ-accessible for some λ.

A natural question: If a category is λ-accessible, will it be accessible in regular cardinals $\mu \geq \lambda$ and, if so, in which of these cardinals? As it happens, there is a sufficient condition for upward transfer of accessibility, although it is rather subtle. The following result appears as Theorem 2.11 in [1]:

Theorem 2.23. *For regular cardinals $\lambda < \mu$, the following are equivalent:*

1. *Each λ-accessible category is μ-accessible.*
2. *The category of λ-directed posets with order embeddings (which is λ-accessible) is μ-accessible.*
3. *For each set X of less than μ elements the poset of subsets of size less than λ, $P_{<\lambda}(X)$, has a cofinal set of cardinality less than μ.*
4. *In each λ-directed poset, every subset of less than μ elements is contained in a λ-directed subset of less than μ elements.*

Definition 2.24. For regular cardinals λ and μ, we say that λ is sharply less than μ, denoted $\lambda \triangleleft \mu$, if they satisfy the equivalent conditions of the theorem above.

A few examples to give a sense of the relation $\triangleleft$:

1. $\omega \triangleleft \mu$ for every uncountable regular cardinal μ.

2. For every regular λ, $\lambda \lhd \lambda^+$.

3. For any regular cardinals λ and μ with $\lambda \leq \mu$, $\lambda \lhd (2^\mu)^+$.

4. Whenever μ and λ are regular cardinals with $\beta^\alpha < \mu$ for all $\beta < \mu$ and $\alpha < \lambda$, then $\lambda \lhd \mu$.

See 2.13 in [1] for more examples. The critical point, perhaps, is that for each set of regular cardinals L, there are arbitrarily large regular cardinals μ with the property that $\lambda \lhd \mu$ for all $\lambda \in L$.

Of course, we are chiefly interested in accessible categories in their capacity as very general environments in which to study abstract model theory. It is high time, then, that we zero in on the connection between accessible categories and more conventional classes of models, and consider the additional axioms we may wish to impose on accessible categories to ensure that they suit our fundamentally model-theoretic purposes. To begin, we notice:

Proposition 2.25. *(Rosický, [22]) Given a first order theory T in language $L(T)$ and* **Elem**(T) *the category with objects the models of T and morphisms the elementary embeddings, then for any regular $\mu > |L(T)| + \aleph_0$,* **Elem**(T) *is μ-accessible, and $M \in \mathcal{K}$ is μ-presentable if and only if $|M| < \mu$.*

Accessible categories are far more general than this, of course, and encompass categories of models of basic sentences in any infinitary logic $L_{\kappa,\lambda}$, where we say that a sentence is basic if it is a conjunction of sentences of the form $\forall x(\phi \to \psi)$, where ϕ and ψ are positive existential (that is, built from atomic formulas using only conjunction, disjunction, and existential quantification). Moreover, they also encompass many of the natural classes of mathematical structures that provided the initial impetus for the study of nonelementary classes: for example, the category of

Banach spaces with contractions as morphisms is an accessible category.

A certain amount of structural information may be lost in thinking of our categories of models as being simply accessible. In particular, if we regard a category of the form $\mathbf{Elem}(T)$—with T a first order theory in language $L(T)$—as being accessible in a regular cardinal $\lambda > |L(T)| + \aleph_0$, we are accounting explicitly only for the closure of $\mathbf{Elem}(T)$ under λ-directed colimits. Since $\lambda \geq \aleph_1$, then, we are not accounting for closure under directed colimits nor, consequently, for closure under unions of chains. As the latter property is essential for our purposes, we will henceforth focus our attention on those accessible categories that do possess arbitrary directed colimits. Categories of this nature were first considered in [22], and most of the notions presented in the remainder of this section have their origins in that paper.

As a start, [22] introduces category-theoretic analogues of two familiar properties: joint embedding and amalgamation. Joint embedding for a category $\mathbf{C}$ amounts to little more than a restatement of the corresponding property for AECs given in Section 2.1: for any objects A and B in $\mathbf{C}$, there is an object C with morphisms $A \to C$ and $B \to C$. The amalgamation property is precisely the same as well: for any object A of $\mathbf{C}$ and morphisms $f_1 : A \to B_1$ and $f_2 : A \to B_2$, there is an object C with morphisms $g_1 : B_1 \to C$ and $g_2 : B_2 \to C$ such that the following diagram commutes:

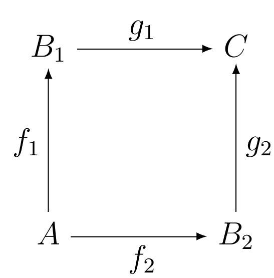

There follow a sequence of definitions from [22], for which we endeavor to provide suitable motivation of a more conventional model-theoretic nature. When we return to these notions in Chapter 5, we will work out explicitly what they mean in the case that interests us most: abstract elementary classes. First,

Definition 2.26. A category **C** is λ-*categorical* if it contains a unique (up to isomorphism) object N which is λ^+-presentable, but not μ-presentable for any $\mu < \lambda^+$. **C** is said to be *strongly* λ-*categorical* if it contains a unique (up to isomorphism) λ^+-presentable object.

Consider the elementary case, where $\mathbf{C} = \mathbf{Elem}(T)$. Recall that for $\mu > |L(T)| + \aleph_0$, an object of $\mathbf{Elem}(T)$ is μ-presentable if and only if it is of cardinality less than μ, by Proposition 2.25, meaning that $\mathbf{Elem}(T)$ is λ-categorical if and only if it contains (up to isomorphism) exactly one model of cardinality less than or equal to λ, but not less than μ for any regular $\mu < \lambda$. This implies, readily enough, that $\mathbf{Elem}(T)$ is λ-categorical if and only if it contains exactly one model of cardinality λ; that is, it is λ-categorical in the usual sense. The category $\mathbf{Elem}(T)$ is strongly λ-categorical if and only if it contains exactly one model of cardinality at most λ (up to isomorphism). More abstractly (and crucially, for the work in Section 5.4 below), a category **C** is strongly λ-categorical when $\mathbf{Pres}_\lambda(\mathbf{C})$ is a one object category; that is, a monoid.

Another notion that is obviously of model-theoretic provenance:

Definition 2.27. Let λ be a regular cardinal. An object M in a category **C** is said to be λ-*saturated* if for any λ-presentable objects N, N' and morphisms $f : N \to M$ and $g : N \to N'$, there is a morphism $h : N' \to M$ such that the following diagram commutes:

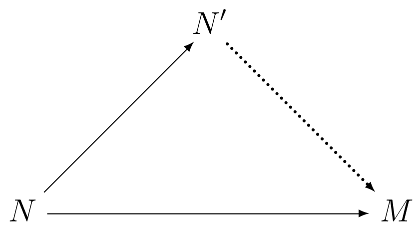

This looks more like a homogeneity condition (and, to give a preview of what is to come in Chapter 5, a great deal like the λ-model homogeneity defined above). As one would expect, though, in an elementary class $\mathbf{Elem}(T)$, for any $\lambda > |L(T)| + \aleph_0$, an object in $\mathbf{Elem}(T)$ is λ-saturated in this sense precisely when it is λ-saturated in the classical sense—it realizes all complete types over subsets of size less than λ.

Definition 2.28. Let λ be a regular cardinal. A morphism $f : M \to N$ in a category $\mathbf{C}$ is said to be λ-*pure* if for any commutative square

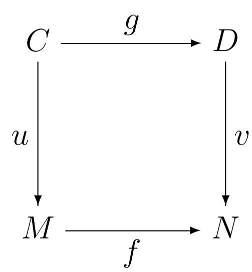

in which C and D are λ-presentable, there is a morphism $h : D \to M$ such that $h \circ g = u$.

Abstractly, any split monomorphism is λ-pure. The converse, while true in **Set**, does not hold in general. A necessary and sufficient condition for λ-purity can be given provided we are working in a category of structures: let Σ be a λ-ary signature, and $\Sigma-$**Struct** the category of Σ-structures and homomorphisms (preserving relations only in the forward direction). Given N in $\Sigma-$**Struct** and M a Σ-substructure, M is a λ-pure substructure if for every positive primitive formula ϕ in $L_{\lambda,\lambda}$ (that is, for every ϕ of the form $\exists\bar{y}\psi(\bar{x},\bar{y})$, where ψ is a conjunction of atomic formulas) and

$\bar{a} \in M$, $M \models \phi[\bar{a}]$ if and only if $N \models \phi[\bar{a}]$. The reader may wish to consult [22] for the justification of this claim, and for more concrete examples. In the finitary case, this notion bears a certain resemblance to the concept of purity familiar from, say, group theory, where one says that a subgroup H of a group G is a pure subgroup if every element of H that has an nth root in G also has an nth root in H (of course, if it has an nth root in H, it has one in G as well). More suggestively, H is a pure subgroup of G if for each positive primitive formula of the form $\phi = \exists y(y^n = x)$ and $a \in H$, $H \models \phi[a]$ if and only if $G \models \phi[a]$. A priori, ω-purity is a more stringent condition, but the essential flavor is the same.

More generally, we consider what, in a category of the form **Elem**(T), it might mean for an elementary inclusion $M \prec N$ to be λ-pure. First, notice that any complete type p over a submodel C of M is realized somewhere (if not in M itself), and we may take a model D containing a realization a of p which is also of size less than λ. If there is an embedding of D in N that makes the diagram in Definition 2.28 above commute (with the other maps in the square being the inclusions), then the image of a in N is a realization of the original type p. By λ-purity of the inclusion $M \prec N$, there is an embedding of D in M that makes the upper triangle commute—this again means that the image of a in M under this embedding is a realization of p. That is, λ-purity implies that any type over a submodel of M of size less than λ that is realized in N is already realized in M. This is, of course, only a necessary condition, but may be useful as intuition. We will return to this issue in Chapter 5, at which time we will carry out the argument (for AECs) in full detail.

We introduce another concept, which is of interest in certain contexts:

Definition 2.29. Let λ be a regular cardinal. An object M in a category **C** is said to be λ-*closed* if every map with domain M is λ-pure.

We note, however, that in any accessible category satisfying the amalgamation property, saturation and closure coincide. That is, we have:

Theorem 2.30. *(Rosicky, [22]) If* $\mathbf{C}$ *is a category with the amalgamation property, an object* M *is* λ*-closed if and only if it is* λ*-saturated.*

As we are now able to speak meaningfully of λ-saturation, and have access to the notion of λ-purity which we may think of, in rough terms, as a condition involving the realization of types, it is reasonable to think that we may be able to come up with an analogue of stability for accessible categories (or, in fact, for arbitrary categories). One such analogue, introduced in [22], will be the subject of a great deal of attention in Chapter 5:

Definition 2.31. Let λ be a regular cardinal. A category $\mathbf{C}$ is said to be *weakly* λ*-stable* if for any λ^+-presentable N and morphism $f : N \to M$, f factors as $N \xrightarrow{g} M' \xrightarrow{h} M$, where M' is λ^+-presentable, and h is λ-pure.

To get a sense of what this involves, we again consider the category of models of a first order theory, $\mathbf{Elem}(T)$. Weak λ-stability of $\mathbf{Elem}(T)$ would mean that for any N with $|N| \leq \lambda$ (λ^+-presentable, that is) and elementary extension M, there is an intermediate extension N', $N \prec N' \prec M$, with $|N'| \leq \lambda$ and such that the inclusion $N' \prec M$ is λ-pure; that is, every type over a subset of N' of cardinality less than λ that is realized in M is also realized in N'. In other words, N has an extension of size at most λ that is λ-saturated relative to M. It is relatively easy to see, then, that if the theory T is λ-stable in the standard type-counting sense of the term, the category $\mathbf{Elem}(T)$ is weakly λ-stable.

We will have a good deal more to say about weak λ-stability in Chapter 5. For now, we call to the reader's attention the following remarkable fact (Proposition 2 in [22]):

Proposition 2.32. *Let* $\mathbf{C}$ *be a* λ*-accessible category, and* μ *a regular cardinal such that* $\lambda \trianglelefteq \mu$ *and* $|\mathbf{Pres}_\lambda(\mathbf{C})^{mor}| < \mu$ *(where* $\mathbf{Pres}_\lambda(\mathbf{C})^{mor}$ *denotes the set of morphisms in* $\mathbf{Pres}_\lambda(\mathbf{C})$*, the skeleton of the subcategory of* λ*-presentable objects in* $\mathbf{C}$*). Then* $\mathbf{C}$ *is weakly* $\mu^{<\mu}$*-stable.*

From the discussion of the relation of sharp inequality above, this means that any accessible category is weakly stable in infinitely many, arbitrarily large cardinals. As we will see in Section 5.3, this surprising result, in conjuction with one of the theorems from Section 4.3, will give us the beginnings of a stability spectrum for weakly tame AECs.

CHAPTER 3

Galois Types and Topology

Let $\mathcal{K}$ be an abstract elementary class. We assume throughout that $\mathcal{K}$ contains a monster model, say $\mathfrak{C}$. All work is done in, and all types are orbits of elements of, the model $\mathfrak{C}$. For any $\lambda \geq \mathrm{LS}(\mathcal{K})$ and $M \in \mathcal{K}$, consider the set

$$\mathcal{S}_M^\lambda = \{N \prec_{\mathcal{K}} M \,:\, |N| \leq \lambda\}$$

We define a topology on ga-S(M) as follows: for each $N \in \mathcal{S}_M^\lambda$ and $p \in$ ga-S(N), let $U_{p,N} \subseteq$ ga-S(M) be given by

$$U_{p,N} = \{q \in \text{ga-S}(M) \,:\, q \restriction N = p\}$$

Claim 3.1. *The $U_{p,N}$ form a basis for a topology on ga-S(M).*

Proof: It is easy to see that any $q \in$ ga-S(M) is contained in one of the $U_{p,N}$. In particular, for any $N \in \mathcal{S}_M^\lambda$, $q \in U_{q \restriction N, N}$.

Suppose now that $q \in U_{p_1,N_1} \cap U_{p_2,N_2}$; that is, that $q \restriction N_1 = p_1$ and $q \restriction N_2 = p_2$. As $\lambda \geq \mathrm{LS}(\mathcal{K})$, there is a model $N \prec_{\mathcal{K}} M$ containing $N_1 \cup N_2$ which is of cardinality at most λ, and which is therefore contained in $\mathcal{S}_M^\lambda$. Obviously $q \in U_{q \restriction N, N}$. Moreover, for any $q' \in U_{q \restriction N, N}$ (hence satisfying $q' \restriction N = q \restriction N$),

$$q' \restriction N_1 = (q' \restriction N) \restriction N_1 = (q \restriction N) \restriction N_1 = q \restriction N_1 = p_1$$

and, similarly, $q' \restriction N_2 = p_2$. Hence $q \in U_{q \restriction N, N} \subseteq U_{p_1,N_1} \cap U_{p_2,N_2}$, and we are done. $\square$

Notation 3.2. We denote by X_M^λ the set ga-S(M) endowed with the topology generated by the sets $U_{p,N}$.

Consider the special case in which $\mathcal{K}$ is an elementary class equipped with the elementary submodel relation. Recall that the Galois types over a model $M \in \mathcal{K}$ correspond to the complete first order types over M. Hence we have a more familiar topology to which to compare those just defined: the syntactic topology. Given the profoundly abstract context in which we defined our topologies, one would not necessarily expect any relation whatsoever. As it happens, there is a relationship. Moreover, it is a remarkably simple one:

Proposition 3.3. *The topology of X_M^λ refines the syntactic topology on S(M)*

Proof: Take any basic open set in the syntactic topology, say U_ϕ, the set of types over M containing the formula ϕ. We wish to show that this set is also open in X_M^λ. To that end, let q be a type in U_ϕ. Let N be an elementary submodel of M of size at most λ that contains all the parameters in ϕ, and consider the basic open set $U_{q\restriction N,N}$. Certainly $q \in U_{q\restriction N,N}$. For any $q' \in U_{q\restriction N,N}$, q' has the same restriction to N as q, and must therefore contain the same formulas with parameters in N as q does. Hence $\phi \in q'$, and $q' \in U_\phi$. Since this holds for arbitrary $q' \in U_{q\restriction N,N}$, we have $U_{q\restriction N,N} \subseteq U_\phi$, thereby establishing the proposition. □

There is actually another topology on syntactic types to be considered in this context, which has been little considered but is implicit in the notion of $\aleph_0$-isolation defined in Section 3.1 of [7]. In this case the topology is induced not by formulas, but by restrictions to finite subsets of the domain. In particular, for any set A, we define a basis of open sets for S(A) as follows: for each finite $B \subseteq A$ and type $p \in$ S(B), take

$$U_p = \{q \in \mathrm{S}(A) \mid q \restriction B = p\}$$

where we here mean restriction in the syntactic sense. The induced topology is, on its face, more closely related to those of the X_M^λ. In fact, by a pair of arguments much like the one given above,

Proposition 3.4. *The topology induced by restriction to finite subsets refines the syntactic topology, and is refined by the topologies on the spaces X_M^λ.*

It is very nearly a special case of our topologies, in fact, but not quite: the purported bases for the X_M^λ are sure to be bases only if we work with $\lambda \geq \mathrm{LS}(\mathcal{K})$ (to ensure that the intersection axiom holds), hence with λ infinite. As a result, we cannot—in our framework—generate a topology from the types with finite domains.

Turning away from the elementary case, and refocusing on the spaces that will be our primary concern here, notice that if $|M| \leq \lambda$, the space X_M^λ is discrete. We will be concerned exclusively with the more interesting case, when $|M| > \lambda$. Ideally, there would be some connection between the spaces X_M^λ obtained as we vary the parameters M and λ. In reality, there is: this scheme for topologizing sets of types over models $M \in \mathcal{K}$—the assignment $(M, \lambda) \mapsto X_M^\lambda$—is functorial in both arguments. We begin with functoriality in the first argument:

Lemma 3.5. *Given any $M, M' \in \mathcal{K}$, $M \prec_{\mathcal{K}} M'$, the natural restriction map $\mathrm{r}_{M',M}$: ga-S(M') $\to$ ga-S(M) is a continuous surjection from $X_{M'}^\lambda$ to X_M^λ.*

Proof: To see that the map $\mathrm{r}_{M',M}$: ga-S$(M') \to$ ga-S(M) is surjective, notice that for any $p \in$ ga-S(M), $p =$ ga-tp(a/M) for some $a \in \mathfrak{C}$. Hence $p =$ ga-tp$(a/M) = \mathrm{r}_{M',M}($ga-tp$(a/M'))$. To verify continuity, take any basis element $U_{p,N} \subseteq X_M^\lambda$. We then have

$$\begin{aligned} \mathrm{r}^{-1}_{M',M}(U_{p,N}) &= \{q \in \text{ga-S}(M') : q \upharpoonright M \upharpoonright N = p\} \\ &= \{q \in \text{ga-S}(M') : q \upharpoonright N = p\} \\ &= U_{p,N} \end{aligned}$$

where the latter is now the basic open set in $X^\lambda_{M'}$ (since $N \in \mathcal{S}^\lambda_M \subseteq \mathcal{S}^\lambda_{M'}$, this makes sense). At any rate, $\mathrm{r}^{-1}_{M',M}(U_{p,N})$ is certainly open in $X^\lambda_{M'}$, so we are done. $\square$

One can easily see that $\mathrm{r}_{M',M} : \text{ga-S}(M') \to \text{ga-S}(M)$ is not merely a continuous surjection, but in fact a quotient map.

We now turn to the assertion of full functoriality in the first argument. Let $M, M' \in \mathcal{K}$, and let $f : M \to M'$ be a $\mathcal{K}$-embedding. We obtain an induced map $f^* : \text{ga-S}(M') \to \text{ga-S}(M)$ as follows: the map f extends to an automorphism $\bar{f} : \mathfrak{C} \to \mathfrak{C}$. The map f^* is given by the following composition:

$$\begin{array}{ccccc} \text{ga-S}(M') & \xrightarrow{\mathrm{r}_{M',f[M]}} & \text{ga-S}(f[M]) & \xrightarrow{\phi} & \text{ga-S}(M) \\ q = (a, M') & \mapsto & (a, f[M]) & \mapsto & (\bar{f}^{-1}(a), M) \end{array}$$

It is simple enough to verify that ϕ is independent of the choice of $\bar{f}$, and that it is defined on types.

Proposition 3.6. *The map f^* is a continuous surjection from $X^\lambda_{M'}$ to X^λ_M.*

Proof: Naturally, $f[M] \prec_{\mathcal{K}} M'$, so Lemma 3.5 above implies that the first map in this composition is a continuous function from $X^\lambda_{M'}$ onto $X^\lambda_{f[M]}$, so it suffices to consider only the second, which we are calling ϕ.

The map ϕ has an obvious inverse on representatives and, by an argument essentially indistinguishable from that given for ϕ, this ϕ^{-1} is a well-defined map from ga-S(M) to ga-S$(f[M])$. It follows that ϕ is bijective.

To verify continuity, consider an arbitrary basic open set $U_{p,N} \subseteq X^\lambda_M$. Take $q \in U_{p,N}$ with representative, say, (a, M). Notice that $f[N] \prec_{\mathcal{K}} f[M]$, $|f[N]| = |N| \le \lambda$, and thus $f[N] \in \mathcal{S}^{\lambda}_{f[M]}$. Hence $q' = \phi^{-1}(q)$, the type represented by $(\bar{f}(a), f[M])$, is contained in the basic open set $U_{q'|f[N],f[N]} \subseteq X^{\lambda}_{f[M]}$. That is, $\phi^{-1}(U_{p,N}) \subseteq U_{q'|f[N],f[N]}$. Given any $q'' \in U_{q'|f[N],f[N]}$, on the other hand, we have $q'' \restriction f[N] = q' \restriction f[N]$, which is to say that given any representative $(b, f[M])$ of q'', there is a $g \in \mathrm{Aut}_{f[N]}(\mathfrak{C})$ with $g(b) = \bar{f}(a)$. Notice that the images $\phi(q') = q$ and $\phi(q'')$ have representatives (a, M) and $(\bar{f}^{-1}(a), M)$, respectively. By the now familiar yoga, $\bar{f}^{-1}g\bar{f} \in \mathrm{Aut}_N(\mathfrak{C})$ and

$$\bar{f}^{-1}g\bar{f}(\bar{f}^{-1}(b)) = \bar{f}^{-1}g(b) = \bar{f}^{-1}(\bar{f}(a)) = a$$

meaning that $(\bar{f}^{-1}(b), M)$ and (a, M) are equivalent over N, from which it follows that $\phi(q') \restriction N = q \restriction N = p$. So $\phi(q') \in U_{p,N}$, $q' \in \phi^{-1}(U_{p,N})$, and ultimately $U_{q'|f[N],f[N]} \subseteq \phi^{-1}(U_{p,N})$. Coupled with the reverse inclusion established above, we have $\phi^{-1}(U_{p,N}) = U_{q'|f[N],f[N]}$. Finally, we conclude that ϕ is a continuous map from $X^{\lambda}_{f[M]}$ to X^{λ}_M.

Now, the map $f^* : X^{\lambda}_N \to X^{\lambda}_M$, being a composition of continuous surjections (actually, a quotient map and a homeomorphism), is itself a continuous surjection (quotient map).

Notice that when $f : M \to N$ is simply an inclusion, f^* is precisely the restriction map $r_{N,M}$. In particular, for any identity map $Id_M : M \to M$, Id^*_M is the identity map on X^{λ}_M. Moreover, it takes minimal effort to see that for any $\mathcal{K}$-embeddings $f : M \to N$ and $g : L \to M$, $(fg)^* = g^* f^*$. Hence we do indeed have functoriality of the sort claimed above. □

As it happens, the far easier fact of functoriality in the cardinal parameter proves to be considerably more interesting. First:

Proposition 3.7. *For any $\mu > \lambda$, the set-theoretic identity map $Id_{\mu,\lambda} : X^{\mu}_M \to X^{\lambda}_M$ is continuous.*

Proof: Consider a basic open set $U_{p,N} \subseteq X^\lambda_M$, where $N \in \mathcal{S}^\lambda_M$ and $p \in$ ga-S(N). Since $\lambda < \mu$, $\mathcal{S}^\lambda_M \subseteq \mathcal{S}^\mu_M$, and we therefore have $N \in \mathcal{S}^\mu_M$. Hence $U_{p,N}$ is in the basis of X^μ_M, and certainly open. □

What this means is that for each $M \in \mathcal{K}$, we obtain a well-behaved spectrum of topological spaces ranging from $X^{\mathrm{LS}(\mathcal{K})}_M$ to $X^{|M|}_M$ (at which point the spaces become discrete, as we noted earlier), all of which are connected by continuous bijections; that is, we have

$$X^{\mathrm{LS}(\mathcal{K})}_M \longleftarrow X^{\mathrm{LS}(\mathcal{K})^+}_M \longleftarrow \cdots \longleftarrow X^\lambda_M \longleftarrow \cdots \longleftarrow X^{|M|}_M \longleftarrow$$

As we will see, there is a close correspondence between topological properties of the spaces in the spectrum and model-theoretic properties of the model M and of $\mathcal{K}$, the ambient AEC.

3.1 Isolated Types and λ-Saturation

We begin with a very simple example of such a correspondence between topology and model theory which, while far less significant than those found in the ensuing sections, captures some of the charm of this setup. In particular, we discover a sufficient condition for isolated points to be dense in X^μ_M: saturation of M in a cardinal $\lambda > \mu$. We also extract a partial converse. To begin,

Claim 3.8. *A type $q \in X^\lambda_M$ is an isolated point if and only if there is an $N \in \mathcal{S}^\lambda_M$ and $p \in$ ga-S(N) such that q is the unique extension of p to a type over M.*

Proof: If q is isolated, $\{q\}$ is open. Thus $q \in U_{p,N} \subseteq \{q\}$ for some basis element $U_{p,N}$. Naturally, we then have $\{q\} = U_{p,N}$; that is, q is the unique extension of p, as claimed. The converse is trivial. □

We might say that in the above situation, p isolates q. By the final section of this chapter, we will be able to give a far more interesting (and far less trivial) characterization of isolated points in the spaces X^{λ}_{M}, but this will do for the moment. We turn, then, to the connection between types isolated in X^{λ}_{M} and types realized in M, and draw out a few simple consequences:

Claim 3.9. *A type of the form ga-tp(a/M) with $a \in M$ is isolated in every X^{μ}_{M}.*

Proof: Say $q = \text{ga-tp}(a/M)$ with $a \in M$. Take $N \in \mathcal{S}^{\mu}_{M}$ such that $a \in N$, and consider the type $p = q \upharpoonright N$. We must verify that q is the unique extension of p in ga-S(M). Clearly q is such an extension. Given any other, say $q = (a', M)$, we have $q' \upharpoonright N = q \upharpoonright N$, and thus there is an automorphism $h \in \text{Aut}_N(\mathfrak{C})$ such that $h(a') = a$. Since $a \in N$, though, this means that $a = a'$. Hence $q' = q$, and we are done. □

Proposition 3.10. *If M is λ-saturated, isolated points are dense in X^{μ}_{M} for all $\mu < \lambda$.*

Proof: Suppose M is λ-saturated. Consider a basic open neighborhood $U_{p,N}$ in X^{μ}_{M} with $\mu < \lambda$, where $N \in \mathcal{S}^{\mu}_{M}$ and $p \in \text{ga-S}(N)$. By λ-saturation of M, p is realized in M, meaning that there is an element $a \in M$ such that $\text{ga-tp}(a/M) \upharpoonright N = p$. Hence $\text{ga-tp}(a/M) \in U_{p,N}$ and, from the claim above, $\text{ga-tp}(a/M)$ is an isolated point of X^{λ}_{M}. We have shown, then, that any basic open set in X^{μ}_{M} contains an isolated point, from which it follows that isolated types are dense, as claimed. □

In words: if a model M is λ-saturated, then isolated points are dense in all of the spaces in the spectrum for M below X^{λ}_{M}. There is a partial converse to this result: if the types of elements of M are dense in the spaces X^{μ}_{M} on an interval just below λ, M is λ-saturated. To be precise:

Proposition 3.11. *If there is a cardinal $\kappa < \lambda$ such that $\{\text{ga-tp}(a/M) \,|\, a \in M\}$ is a dense subset of each X^{μ}_{M} with $\kappa \leq \mu < \lambda$, M is λ-saturated.*

Proof: Let $p \in \text{ga-S}(N)$, $N \prec_{\mathcal{K}} M$ with $|N| = \chi < \lambda$. Take $\mu \geq \chi$ with $\kappa \leq \mu < \kappa$. By density of $\{\text{ga-tp}(a/M) \,|\, a \in M\}$ in X^{μ}_{M}, there is a type $q \in U_{p,N}$ of the form $q = \text{ga-tp}(a/M)$, where $a \in M$. The fact that $q \in U_{p,N}$ implies that $q \restriction N = p$, which is to say that $\text{ga-tp}(a/N) = p$. This means, of course, that p is realized in M. $\Box$

As a special case:

Corollary 3.12. *If $\{ga\text{-}tp(a/M) \,|\, a \in M\}$ is a dense subset of X^{λ}_{M}, M is λ^{+}-saturated.*

Whereas Proposition 3.10 allowed us to translate a model-theoretic condition into information about the spectrum associated with the model M, these results go in the opposite direction: given particular information about the spectrum (here, that the types of elements of M are dense in the spaces just below X^{λ}_{M}) we are able to deduce something about the structure of the model itself (here, that M is λ-saturated). In the ensuing sections, we will further analyze the ways in which the spaces X^{λ}_{M} encode information about $\mathcal{K}$ and the models contained therein, obtaining correspondences a good deal deeper than the toy results above.

3.2 Separation and Tameness

In fact, we turn immediately to the correspondence that forms the centerpiece of this chapter, and which motivated the definition of the topologies on the spaces X^{λ}_{M}. Recalling that types over small models play the role of formulas in generating the aforementioned topologies, and also recalling our intuition that tameness means that distinct types over potentially large models may be distinguished by such "formulas," it should come as little surprise that tameness translates into a separation condition. In particular:

Claim 3.13. *The AEC $\mathcal{K}$ is (χ,λ)-tame if and only if for all $M \in \mathcal{K}_\lambda$, X^χ_M is Hausdorff.*

Proof: ($\Rightarrow$) Suppose that $\mathcal{K}$ is (χ,λ)-tame, and that $M \in \mathcal{K}_\lambda$. Let $q, q' \in X^\chi_M$, $q \neq q'$. By tameness, there is an $N \in \mathcal{S}^\chi_M$ with the property that $q \restriction N \neq q' \restriction N$. Certainly $q \in U_{q\restriction N,N}$ and $q' \in U_{q'\restriction N,N}$. Furthermore, $U_{q\restriction N,N} \cap U_{q'\restriction N,N} = \emptyset$: if we had $q'' \in U_{q\restriction N,N} \cap U_{q'\restriction N,N}$, then we would have $q'' \restriction N = q \restriction N$ and $q'' \restriction N = q' \restriction N$, whence $q \restriction N = q' \restriction N$, contradicting our choice of N.

($\Leftarrow$) Suppose X^χ_M is Hausdorff for all $M \in \mathcal{K}_\lambda$. If $q, q' \in$ ga-S(M), $q \neq q'$, then there are open sets $U \ni q$ and $U' \ni q'$ with $U \cap U' = \emptyset$. Given the way the topology was defined, we have, moreover, that for some $N_1, N_2 \in \mathcal{S}^\chi_M$, $p_1 \in$ ga-S(N_1), and $p_2 \in$ ga-S(N_2), $q \in U_{p_1,N_1}$ and $q' \in U_{p_2,N_2}$, with $U_{p_1,N_1} \cap U_{p_2,N_2} = \emptyset$. In particular, $q' \notin U_{p_1,N_1}$, meaning that, whereas $q \restriction N_1 = p_1$, $q' \restriction N_1 \neq p_1$. So $q \restriction N_1 \neq q' \restriction N_1$. Since the same argument works for any $M \in \mathcal{K}$ of size λ, the class $\mathcal{K}$ must be (χ,λ)-tame. $\square$

In reality, though, the separation axioms T0 and T1 are equivalent to the Hausdorff axiom in all X^λ_M, regardless of the semantic properties of the class:

Proposition 3.14. *Let $\mathcal{K}$ be an AEC, $M \in \mathcal{K}$, and $\lambda \geq LS(\mathcal{K})$. The space X^λ_M is Hausdorff if and only if it is $T0$.*

Proof: ($\Rightarrow$) Trivial.

($\Leftarrow$) Suppose X^λ_M is $T0$. Given distinct $q, q' \in X^\lambda_M$, then, there is an open set $U \subseteq X^\lambda_M$ containing q but not q'. Indeed, we may take U to be a basis element, $U_{p,N}$. That is, there is an $N \in \mathcal{S}^\lambda_M$ and $p \in$ ga-S(M) such that $q \restriction N = p$ but $q' \restriction N \neq p$. Hence $U_{p,N}$ and $U_{q'\restriction N,N}$ are disjoint open neighborhoods of q and q', respectively. $\square$

It follows, of course, that in such spaces the T1 axiom amounts to the same thing as the two in the previous proposition. So,

Corollary 3.15. *The AEC $\mathcal{K}$ is (χ,λ)-tame if and only if for all $M \in \mathcal{K}_\lambda$, X^χ_M is T0 (T1, Hausdorff).*

As a consequence, we have, among other things, the following topological spectrum result:

Claim 3.16. *If $\mathcal{K}$ is (χ,λ)-tame, X^μ_M is Hausdorff for all $M \in \mathcal{K}_\lambda$ and $\mu \geq \chi$.*

Proof: The antecedent holds if and only if for any $M \in \mathcal{K}_\lambda$, X^χ_M is Hausdorff. For any $\mu > \chi$, the map $\mathrm{Id}_{\mu,\chi} : X^\mu_M \to X^\lambda_M$ is a continuous injection, from whose existence it follows that X^μ_M is also a Hausdorff topological space. □

Notice that the less parametrized notion of χ-tameness corresponds to the property that for every $M \in \mathcal{K}$, X^χ_M is Hausdorff (T0, T1). Functoriality and basic topology (as in the proof of Claim 3.16 above) guarantee that χ-tameness implies that the spaces X^μ_M are Hausdorff for all $\mu \geq \chi$. This also means that $\mathcal{K}$ is μ-tame for all $\mu \geq \chi$, as one would hope.

Actually, we can say a bit more along these lines. The basic open sets $U_{p,N}$ are, in fact, clopen:

$$\begin{aligned}(U_{p,N})^C &= \text{ga-S}(M) \setminus U_{p,N} = \{q \in \text{ga-S}(M) : q \restriction N \neq p\} \\ &= \bigcup_{p' \in (ga-S(N)\setminus\{p\})} U_{p',N}\end{aligned}$$

where the final union is, quite obviously, open in X^λ_M. This means, of course, that every X^λ_M has a basis of clopen sets. Hence, by Claim 3.13,

Proposition 3.17. *An AEC $\mathcal{K}$ is (χ,λ)-tame if and only if X^χ_M is totally disconnected for every $M \in \mathcal{K}_\lambda$.*

By the same token,

Proposition 3.18. *If $\mathcal{K}$ is (χ,λ)-tame, X^{μ}_{M} is totally disconnected for every $M \in \mathcal{K}_{\lambda}$ and $\mu \geq \chi$.*

Naturally, there is an analogous result for χ-tameness, where we simply remove the restriction on the cardinality of the model M. Here again we have found an elegant correspondence between model-theoretic properties and the properties of spaces in the topological spectra for individual models: tameness ensures that the spaces are totally disconnected. In short, they are very nearly Stone spaces, with compactness being the only (potentially) missing ingredient. The natural question—whether they are in fact compact—will be addressed in Section 3.4 below. First, though, we need a bit more information about the spaces X^{λ}_{M}.

3.3 Uniform Structure

Up to now, we have studiously ignored the fact that the sets ga-S(M) support a natural notion of closeness: types with the same restriction to $N \in \mathcal{S}^{\lambda}_{M}$ might be said to be N-close, and might be reckoned to be closer still if they are N'-close for some $N' \in \mathcal{S}^{\lambda}_{M}$ with $N \prec_{\mathcal{K}} N'$. In fact, this way of thinking gives a uniform structure on ga-S(M), with respect to which we may speak sensibly of Cauchy nets, convergence, and completeness. The great virtue of this perspective is that it provides a formal grounding for the intuition that (roughly speaking) (κ,λ)-compactness and (κ,λ)-locality relate to the convergence behavior of (Cauchy) sequences of types. Indeed, since the topology induced by this uniform structure is precisely the same as the one employed in the previous section, nothing is lost by these considerations. In fact, a considerable amount of additional information about the X^{λ}_{M} is gained.

Given $M \in \mathcal{K}_{\kappa}$, and $\lambda \geq \mathrm{LS}(\mathcal{K})$, we define a uniform structure on ga-S(M) as follows: for each $N \in \mathcal{S}_M^\lambda$ define

$$U_N = \{(q, q') \in \text{ga-S}(M) \times \text{ga-S}(M) \,|\, q \upharpoonright N = q' \upharpoonright N\}$$

Notice that $N \prec_{\mathcal{K}} N'$ implies $U_{N'} \subseteq U_N$.

Let B be the set of all such U_N. We now define:

$$U_M^\lambda = \{U \subseteq \text{ga-S}(M) \times \text{ga-S}(M) \,|\, U \supseteq U_N \text{ for some } N \in \mathcal{S}_M^\lambda\}$$

That is, U_M^λ is the uniformity generated by B. Of course, we must verify

Claim 3.19. *B is the basis for a uniformity on ga-S(M)*

Proof: Each $U_N \in B$, obviously contains the diagonal $\Delta_M \subseteq \text{ga-S}(M) \times \text{ga-S}(M)$. Moreover, the converse of any U_N is simply U_N itself, and is therefore included in B. We must also check that for any U_N there is a $V \in B$ with the property that $V \circ V \subseteq U_N$—notice that $U_N \circ U_N = U_N$, meaning that we can simply take $V = U_N$.

Finally, let U_N and $U_{N'}$ be elements of B. We must show that there is an $N'' \in \mathcal{S}_M^\lambda$ such that $U_{N''} \subseteq U_N \cap U_{N'}$. Let N'' be any submodel of cardinality at most λ that contains the union of N and N'. Whenever $(q, q') \in U_{N''}$, $q \upharpoonright N'' = q' \upharpoonright N''$, from which it follows that $q \upharpoonright N = q' \upharpoonright N$ and $q \upharpoonright N' = q' \upharpoonright N'$, in which case $(q, q') \in U_N \cap U_{N'}$. □

The uniform topology on ga-S(M) is defined as follows: V is open only if for every $q \in V$ there is a $U \in U_M^\lambda$ such that the set $U[q] = \{q' \in \text{ga-S}(M) \,|\, (q, q') \in U\}$ is contained in V. For any such U there is an $N \in \mathcal{S}_M^\lambda$ such that $U_N \subseteq U$, and we then have $U_N[q] \subseteq V$. Noticing that $U_N[q]$ is precisely $U_{q \upharpoonright N, N}$, we have found a basic open set of X_M^λ containing q and completely contained in V. As we can do the same for any $q \in V$, it must be the case that V is open in X_M^λ. Reversing the argument, one discovers that the uniform topology is exactly the same as the one defined above. By a familiar result (see [17], for example), we have

Corollary 3.20. *For any $M \in \mathcal{K}$ and $\lambda \geq LS(\mathcal{K})$, X^{λ}_{M} is completely regular.*

Now that we have a notion of closeness in X^{λ}_{M}, we may discuss Cauchyness and convergence: a net $\{q_n \,|\, n \in D\}$ (D a directed set) is Cauchy if and only if for every $U \in U^{\lambda}_{M}$ there is an $n \in D$ such that for all $i, j \geq n$, $(q_i, q_j) \in U$. Equivalently, such a net is Cauchy if and only if for every N in $\mathcal{S}^{\lambda}_{M}$, there is an $n \in D$ such that for all $i, j \geq n$, $(q_i, q_j) \in U_N$; that is, $q_i \restriction N = q_j \restriction N$. Similarly, a net $\{q_n \,|\, n \in D\}$ converges to $q \in X^{\lambda}_{M}$ if and only if for every $N \in \mathcal{S}^{\lambda}_{M}$, there is an $n \in D$ such that for all $i \geq n$, $(q_i, q) \in U_N$; that is, $q_i \restriction N = q \restriction N$. We now turn to the promised examination of (κ, λ)-compactness in this framework, after a quick lemma:

Lemma 3.21. *Suppose $\mathcal{K}$ is $(\chi, \leq \lambda)$-tame. If for any $M \in \mathcal{K}_{\lambda}$ the space X^{μ}_{M} is complete for some μ in the range $\chi \leq \mu < cf(\kappa)$, then for any increasing chain $(M_i)_{i<\kappa}$ with union M and increasing chain of $p_i \in$ ga-S(M_i), there is a type $q \in$ ga-S(M) with $q \restriction M_i = p_i$ for all $i < \kappa$.*

Proof: For each $i < \kappa$, choose $q_i \in$ ga-S(M) extending p_i. This gives a net $\{q_i \,|\, i < \kappa\}$. Indeed, it is Cauchy: for any $N \in \mathcal{S}^{\mu}_{M}$, we must have $N \subseteq M_i$ for some i (since $|N| \leq \mu < cf(\kappa)$). For any $m, n \geq i$,

$$
\begin{aligned}
q_m \restriction M_i &= (q_m \restriction M_m) \restriction M_i \\
&= p_m \restriction M_i \\
&= p_i \\
&= p_n \restriction M_i \\
&= (q_n \restriction M_n) \restriction M_i \\
&= q_n \restriction M_i
\end{aligned}
$$

Hence for any $m, n \geq i$, $q_m \restriction N = q_n \restriction N$, so $(q_m, q_n) \in U_N$, and the net is Cauchy. By completeness of X^{μ}_{M}, it must therefore converge in X^{μ}_{M}, meaning that there is some $q \in X^\mu_M$ with the property that for all $N \in \mathcal{S}^\mu_M$, there is $i < \kappa$ such that for all $n \geq i$, $q \restriction N = q_n|N$. We wish to show that $q \restriction M_i = p_i$ for all $i < \kappa$.

Fix $i < \kappa$. Given any $N \prec_{\mathcal{K}} M_i$ of size at most μ, $N \in \mathcal{S}^\mu_M$ and we therefore know that for large enough $n < \kappa$ and, in particular, for some $n > i$,

$$q \restriction N = q_n \restriction N = ((q_n \restriction M_n) \restriction M_i) \restriction N = (p_n \restriction M_i) \restriction N = p_i \restriction N$$

Since this holds for absolutely any $N \prec_{\mathcal{K}} M_i$ of size at most μ, tameness implies that $q \restriction M_i = p_i$. □

From the lemma, we have

Theorem 3.22. *[Completeness, (κ, λ)-Compactness] If $\mathcal{K}$ is $(\chi, \leq \lambda)$-tame and for all $M \in \mathcal{K}_\lambda$ there exists μ with $\chi \leq \mu < cf(\kappa)$ such that X^μ_M is complete, then $\mathcal{K}$ is (κ, λ)-compact.*

The fact of the matter is that there does not seem to be any semblance of a converse (a measure of completeness following from a measure of type-theoretic compactness). This is attributable to the fact that completeness of X^λ_M is related to collatibility of arbitrary directed systems of types, not chains. To be precise,

Proposition 3.23. *The space X^μ_M is complete if and only if for any compatible system of types $p_N \in$ ga-S(N) over substructures N of size no greater than μ, there is a type $q \in$ ga-S(M) with the property that $q \restriction N = p_N$ for all N.*

Proof: ($\Leftarrow$) Suppose that $\{q_n \,|\, n \in D\}$ is a Cauchy net. Given any $N \in \mathcal{S}^\mu_M$, the fact that the net is Cauchy implies that the q_n eventually have the same restriction to N. Let p_N be this common restriction. I claim the types obtained in this fashion constitute a compatible system. This is simple enough: let $N, N' \in \mathcal{S}^\mu_M$, $N \prec_{\mathcal{K}} N'$. The q_n eventually restrict to $p_{N'}$ on N', hence also eventually restrict to $p_{N'} \restriction N$

on N. By our initial choice, it must be the case that $p_{N'} \restriction N = p_N$. There must, therefore, be a type $q \in$ ga-S(M) with $q \restriction N = p_N$ for all $N \in \mathcal{S}^{\mu}_{M}$. For any such N, it is clear that q eventually agrees with the net there, so $(q, q_n) \in U_N$ for sufficiently large n. Hence q is a limit of the net.

($\Rightarrow$) Suppose we have a compatible family of types p_N. For each p_N, choose an extension $q_N \in$ ga-S(M). The q_N form a net (the index set here is not merely directed, but μ^+-directed) which, I claim, is Cauchy. The proof is identical to the one given in the chain case above, and is therefore omitted. By completeness, the net has a limit $q \in$ ga-S(M). The proof that this q satisfies $q \restriction N = p_N$ for all $N \in \mathcal{S}^{\mu}_{M}$ is also similar to the one above, although we needn't concern ourselves with matters of tameness: since q is a limit of the net there is, for every $N \in \mathcal{S}^{\mu}_{M}$, an $N' \in \mathcal{S}^{\mu}_{M}$ with the property that for all $N'' \succ_{\mathcal{K}} N'$, $q \restriction N'' = q_N \restriction N''$. In particular, we can take such an N'' that contains both N and N'. We then have

$$\begin{aligned} q \restriction N &= q_{N''} \restriction N \\ &= (q_{N''} \restriction N'') \restriction N \\ &= p_{N''} \restriction N \\ &= p_N \end{aligned}$$

So we are done. $\square$

The connection between topology and locality is a good deal simpler. To begin,

Lemma 3.24. *Let $M \in \mathcal{K}_\lambda$. If there is a $\mu < cf(\kappa)$ such that X^{μ}_{M} is T0, then for any increasing chain $(M_i)_{i<\kappa}$ with union M and increasing chain of $p_i \in$ ga-S(M_i), if two types $q, q' \in$ ga-S(M) satisfy $q \restriction M_i = q' \restriction M_i$ for all $i < \kappa$, $q = q'$.*

Proof: Take any such increasing chain of structures, and corresponding chain of types. Given q and q' distinct, there is a basic open set $U_{p,N}$ in X^{μ}_{M} that contains q

but not q'. That is, $q \restriction N \neq q' \restriction N$. Naturally, $N \subseteq M_i$ for some $i < \kappa$, meaning that $q \restriction M_i \neq q' \restriction M_i$ as well. □

With no effort at all,

Theorem 3.25. *[Separation, Locality] If for every $M \in \mathcal{K}_\lambda$ there exists a $\mu < cf(\kappa)$ such that X^μ_M is T0, $\mathcal{K}$ is (κ, λ)-local.*

Once again, there isn't really a nice converse to the theorem above, for the same reason as before. As before, it seems that satisfaction of the separation axioms in individual X^μ_M is a matter of the unique or nonunique collatibility of directed systems of types over structures $N \in \mathcal{S}^\mu_M$. To be precise,

Proposition 3.26. *The space X^μ_M is T0 if and only if for any compatible system of types $p_N \in$ ga-S(N) over structures $N \in \mathcal{S}^\mu_M$ there is at most one type $q \in$ ga-S(M) with the property that $q \restriction N = p_N$ for all N.*

Proof: This proposition is a trivial rewriting of the earlier result regarding the equivalence between separation and tameness. □

To further drive home a point made in the first section—that the assignment

$$(M, \lambda) \mapsto X^\lambda_M$$

is nicely functorial—and to provide additional evidence that (κ, λ)-compactness and (κ, λ)-locality have implications on the topological side, we have

Theorem 3.27. *Let $\mathcal{K}$ be (κ, λ)-compact and (κ, λ)-local. If $M = \bigcup_{i<\kappa} M_i$ is a continuous increasing chain with $|M| = \lambda$, then for any $\mu < cf(\kappa)$,*

$$X^\mu_M = \mathrm{Lim}_{i<\kappa} X^\mu_{M_i}$$

in the category of topological spaces (uniform spaces).

As we will have no further use for this result, we omit the argument (which is both long and entirely straightforward). It is noteworthy, perhaps, that if one removes the locality assumption, one still obtains X^{μ}_M as a weak limit of the diagram.

3.4 Noncompactness

As one would expect, given the paucity of compactness available in AECs, the spaces X^{λ}_M need not be compact. As it happens, the picture is a good deal worse: more as an artifact of this particular way of topologizing types than as a consequence of any lack of logical or Galois type-theoretic compactness, we have a wholesale failure of compactness in the X^{λ}_M. Indeed, as we will see, they cannot even be countably compact if the ambient AEC is sufficiently tame. To begin, we notice:

Proposition 3.28. *Let $\mathcal{K}$ be an arbitrary AEC, $M \in \mathcal{K}$, and $\lambda \geq LS(\mathcal{K})$. The space X^{λ}_M is not compact.*

Proof: Suppose that X^{λ}_M is compact, and consider the open cover $\{U_{p,N} \,|\, p \in \text{ga-S}(N)\}$ for some fixed infinite $N \in \mathcal{S}^{\lambda}_M$. We must have a finite subcover, say

$$\{U_{p_1,N}, \ldots, U_{p_r,N}\}$$

Notice that for any $p \in \text{ga-S}(N)$, there is at least one type $q \in X^{\lambda}_M$ with the property that $q|N = p$. Notice also that this q lies in $U_{p,N}$, but not in $U_{p',N}$ for any other $p' \in \text{ga-S}(N)$. It follows that, if the set above is to be a cover, every $p \in \text{ga-S}(N)$ must appear in the list $\{p_1, \ldots, p_r\}$; that is, ga-S(N) must be finite. But of course $|\text{ga-S}(N)| \geq |N| \geq \aleph_0$. Contradiction. □

In light of the foregoing, we must abandon all hope that (κ, λ)-compactness will translate in a straightforward way to compactness of any particular X^{μ}_M. As it happens, though, things are less dire than one might think. In particular,

Theorem 3.29. *[Topological Compactness, (κ, λ)-Compactness] Suppose $\mathcal{K}$ is (κ, λ)-compact, and let $M \in \mathcal{K}_{\geq \lambda}$. Given any family of basic open sets $U_{p_i,N_i} \subseteq X^\mu_M$, $i < \kappa$, where the N_i form a continuous $\prec_\mathcal{K}$-increasing sequence with union of cardinality λ, if the family has the finite intersection property, $\bigcap_{i<\kappa} U_{p_i,N_i}$ is nonempty.*

Proof: The argument hinges on the easily established fact that if the U_{p_i,N_i} satisfy the finite intersection property, it must be the case that $p_j \restriction N_i = p_i$ for all $i \leq j < \kappa$. By (κ, λ)-compactness, there is a type p over $N = \bigcup_{i<\kappa} N_i$ such that $p \restriction N_i = p_i$ for all $i < \kappa$. Naturally, $N \prec_\mathcal{K} M$, by one of the union of chains axioms. Let $q \in$ ga-S(M) be an extension of p. Clearly, $q \restriction N_i = p_i$ for all $i < \kappa$, so $q \in \bigcap_{i<\kappa} U_{p_i,N_i}$. □

The condition in the conclusion of the result above is not particularly natural, topologically speaking. Consequently, if we seek a result in the other direction (say, a condition resembling compactness on the various spaces X^μ_M which is sufficient to guarantee (κ, λ)-compactness of $\mathcal{K}$) what we get is also less than natural. Recall, though, that we have, at the very least, the implication captured in Theorem 3.22, from completeness of the X^μ_M to compactness of Galois types in $\mathcal{K}$.

We now turn to the matter of countable compactness, or the failure thereof. For this purpose, we make extensive use of the following extraordinary property of the spaces at hand:

Proposition 3.30. *For any $M \in \mathcal{K}$ and $\lambda \geq LS(\mathcal{K})$, the intersection of any λ open subsets of X^λ_M is open.*

Proof: Let $\{U_i \,|\, i < \lambda\}$ be a family of open subsets of X^λ_M, and let $q \in \bigcap_{i<\lambda} U_i$. We wish to show that there is a basic open set containing q that is itself contained in $\bigcap_{i<\lambda} U_i$. For any $i < \lambda$, there is a basic open set U_{p_i,N_i} with $q \in U_{p_i,N_i} \subseteq U_i$. Since $\bigcap_{i<\lambda} U_{p_i,N_i} \subseteq \bigcap_{i<\lambda} U_i$, it suffices to find a basic open set containing q that is contained in $\bigcap_{i<\lambda} U_{p_i,N_i}$. The union of the N_i is of cardinality λ so, by the Downward Löwenheim Skolem Property, there is a submodel $N \hookrightarrow_{\mathcal{K}} M$ that contains it. By coherence, $N_i \hookrightarrow_{\mathcal{K}} N$ for all i. Consider the basic open neighborhood $U_{q\restriction N,N}$. It certainly contains q. For any $q' \in U_{q\restriction N,N}$, moreover,

$$q' \restriction N_i = q' \restriction N \restriction N_i = q \restriction N \restriction N_i = q \restriction N_i = p_i$$

Hence q' is contained in $\bigcap_{i<\lambda} U_{p_i,N_i}$ and, since this holds for any $q' \in U_{q\restriction N,N}$, $U_{q\restriction N,N} \subseteq \bigcap_{i<\lambda} U_{p_i,N_i}$, and we are done. $\square$

Recalling from Corollary 3.20 that the spaces X_M^λ are completely regular, we are led to consider the following generalization of the notion of p-space, familiar from general topology:

Definition 3.31. For any infinite cardinal $\lambda \geq \aleph_0$, we say that a topological space X is a *p_λ-space* if it is completely regular and the intersection of any collection of up to λ many open sets is open.

The normal notion of a p-space (covered in [9], for example) occurs in the special case that $\lambda = \aleph_0$. Now, as we pile separation conditions on a p_λ-space, interesting things happen. The general topology first, then implications in the current context:

Proposition 3.32. *Let X be a $T1$ p_λ-space. Any set $S \subseteq X$ with $|S| \leq \lambda$ is closed and discrete.*

Proof: We begin by showing that such an S is closed. This is easy: from the assumption that X is $T1$, it follows that for every $x \in S$, the singleton $\{x\}$ is closed in X. As S is a union of at most λ such closed sets, it is itself closed in X. To see that S is discrete, fix an element $x \in S$. Since X is $T1$, for every $y \in (X \setminus \{x\})$ there is an open neighborhood U_y of x in X that does not contain y. Take the intersection of all such U_y; that is, $U := \bigcap_{y \in (S \setminus \{x\})} U_y$. Because X is a p_λ-space, U is open in X. Naturally, it contains x, but misses all the other elements of S. As $\{x\} = S \cap U$, $\{x\}$ is open in the subspace topology on S. The result follows. □

More or less immediately, we have:

Corollary 3.33. *Let X be a $T1$ p_λ-space. No subset of X of size less than or equal to λ has an accumulation point in X.*

Proof: Let $S \subseteq X$, with $|S| \leq \lambda$. Suppose that S has an accumulation point $x \in X$. By Proposition 3.32, S is closed, meaning that the accumulation point x must lie in S. By the same result, though, S is discrete. Contradiction. □

Corollary 3.34. *Let X be a $T1$ p_λ-space. Then X is countably compact only if it is finite.*

Proof: Recall that a topological space is countably compact only if it is limit point compact (see, for example, [17]). If X is infinite, it contains a subset S of cardinality $\aleph_0$. Since $|S| \leq \lambda$, the proposition immediately above ensures that S has no accumulation point in X. Hence X is not limit point compact and, consequently, not a countably compact space. □

Now, fix an AEC $\mathcal{K}$. We know that X_M^λ is a p_λ-space for all $M \in \mathcal{K}$ and $\lambda \geq$ LS($\mathcal{K}$). Given what we have discerned of the relationship between tameness and the separation axioms—including T1—in the spaces of types, we may now infer:

Proposition 3.35. *Let $\mathcal{K}$ be (χ,λ)-tame. Then for any $M \in \mathcal{K}_\lambda$ and any $\mu \geq \chi$, X_M^μ is not countably compact.*

Proof: Suppose that X_M^μ is countably compact for some $M \in \mathcal{K}_\lambda$ and $\mu \geq \chi$. We learned in Section 3.2 that (χ,λ)-tameness of $\mathcal{K}$ implies that the space X_M^μ is T1. By

Corollary 3.34, then, X^μ_M is finite. But this is absurd: $|X^\mu_M| = |\text{ga-S}(M)| \geq |M| = \lambda$, and is therefore infinite. $\square$

We know, then, that in the topological spectrum for any model of size λ, we will not find any countably compact spaces above the tameness cardinal χ. It is strange that our desire for tameness of Galois types should preclude the possibility of even a modicum of topological compactness, but this is undoubtedly a sacrifice worth making.

We close with another consequence of Proposition 3.32:

Proposition 3.36. *Let $\mathcal{K}$ be (χ,λ)-tame. Then for any $M \in \mathcal{K}_\lambda$ and any $\mu \geq \chi$, a type $q \in X^\mu_M$ is an accumulation point of a set $S \subseteq X^\mu_M$ only if every neighborhood of q contains strictly more than λ elements of S.*

Proof: Let U be a neighborhood of q with $|U \cap S| \leq \lambda$. The intersection is discrete, meaning that there is a neighborhood V of q that does not meet $U \cap S$ in any other point. But then $U \cap V$ is a neighborhood of q with $(U \cap V) \cap S = \{q\}$, meaning that q cannot be an accumulation point of S. $\square$

In short, it is very difficult to be an accumulation point of types. This fact motivates the definition of the family of ranks for Galois types with which we will occupy ourselves in the next chapter.

CHAPTER 4

Morley-like Ranks for Tame AECs

We have seen that if an AEC $\mathcal{K}$ is χ-tame, then for any $\lambda \geq \chi$ (and, of course, $\lambda \geq \mathrm{LS}(\mathcal{K})$), the criterion for a type $q \in \text{ga-S}(M)$ to be an accumulation point of a family of types in X_M^λ is quite stringent: every neighborhood must contain more than λ types in the given family. This condition inspires the definition of the following Morley-like ranks:

Definition 4.1. [RM^λ] Assume $\mathcal{K}$ is χ-tame for some $\chi \geq \mathrm{LS}(\mathcal{K})$. For $\lambda \geq \chi$, we define RM^λ by the following induction: for any $q \in \text{ga-S}(M)$ with $|M| \leq \lambda$,

- $\mathrm{RM}^\lambda[q] \geq 0$.
- $\mathrm{RM}^\lambda[q] \geq \alpha$ for limit α if $\mathrm{RM}^\lambda[q] \geq \beta$ for all $\beta < \alpha$.
- $\mathrm{RM}^\lambda[q] \geq \alpha + 1$ if there exists a structure $M' \succ_{\mathcal{K}} M$ such that q has strictly more than λ extensions to types q' over M' with the property that

$$\mathrm{RM}^\lambda[q' \restriction N] \geq \alpha \text{ for all } N \in \mathcal{S}_{M'}^\lambda$$

For types q over M of arbitrary size, we define

$$\mathrm{RM}^\lambda[q] = \min\{\mathrm{RM}^\lambda[q \restriction N] : N \prec_{\mathcal{K}} M, |N| \leq \lambda\}.$$

Before we proceed, a bit of motivation for this longwinded definition: in topologizing ga-S(M) as X^λ_M, we essentially allowed the types over substructures of size at most λ to play a role analogous to that typically played by formulas in the definition of the topology on syntactic types. In defining the ranks RM^λ, we first define the rank of these "formulas" and subsequently define the rank on types extending them as the minimum of the ranks of their constituent "formulas," just as in the definition of Morley rank in classical model theory.

There is something to check here, though. The types over models of cardinality at most λ whose ranks were defined by the inductive clause are also covered by the second clause. We must ensure that there is no possibility of conflict between the two. Let RM^λ_* denote the ranks assigned to such types using only the inductive clause.

Lemma 4.2. *The ranks RM^λ_* are monotone: for any $N, N' \in \mathcal{K}_{\leq\lambda}$, $N' \prec_{\mathcal{K}} N$, and any $q \in$ ga-S(N), $RM^\lambda_*[q] \leq RM^\lambda_*[q \upharpoonright N']$.*

Proof: We show that for any ordinal α, $\mathrm{RM}^\lambda_*[q] \geq \alpha$ implies $\mathrm{RM}^\lambda_*[q \upharpoonright N'] \geq \alpha$. We proceed by an easy induction on α. The zero case is trivial, by the first bullet point. The limit cases follow from the induction hypothesis. For the successor case, notice that if $\mathrm{RM}^\lambda_*[q] \geq \alpha + 1$, the extension $M \succ_{\mathcal{K}} N$ that witnesses this fact is also an extension of N' and witnesses that $\mathrm{RM}^\lambda_*[q \upharpoonright N'] \geq \alpha + 1$. $\square$

So we needn't worry that the rank assigned to a type $q \in$ ga-S(M) with $M \in \mathcal{K}_{\leq\lambda}$ under the second clause ($\mathrm{RM}^\lambda[q] = \min\{\mathrm{RM}^\lambda[q \upharpoonright N] : N \prec_{\mathcal{K}} M, |N| \leq \lambda\}$) will differ from the rank of p under the first clause. Hence the ranks RM^λ are well-defined. In fact, we can now restate the third bullet point above in a more compact form: for $q \in$ ga-S(M) with $M \in \mathcal{K}_{\leq\lambda}$,

- $\mathrm{RM}^\lambda[q] \geq \alpha+1$ if there exists a structure $M' \succ_{\mathcal{K}} M$ such that q has strictly more than λ many extensions to types q' over M' with $\mathrm{RM}^\lambda[q'] \geq \alpha$.

We will invariably use this formulation in the sequel.

4.1 Properties of $\mathbf{RM}^\lambda$

We now consider the basic properties of the ranks RM^λ, beginning with an extension of our partial monotonicity result, Lemma 4.2.

Proposition 4.3. *[Monotonicity] If $M \prec_{\mathcal{K}} M'$ and $q \in$ ga-S(M'), then $RM^\lambda[q] \leq RM^\lambda[q \restriction M]$.*

Proof: Trivial, from the second clause of the definition. $\square$

Bear in mind that "p is a restriction of q" corresponds to the first order "$q \vdash p$," so this is indeed an appropriate analogue of the monotonicity result for the classical Morley rank.

Proposition 4.4. *[Invariance] Let f be an automorphism of the monster model $\mathfrak{C}$, let $q \in$ ga-S(M), and let $M' = f[M]$. Then the type $f[q] \in$ ga-S(M') satisfies $RM^\lambda[f[q]] = RM^\lambda[q]$.*

Proof: We again proceed by induction, and once again the zero and limit cases are trivial. It remains to show that whenever $\mathrm{RM}^\lambda[q] \geq \alpha + 1$, $\mathrm{RM}^\lambda[f[q]] \geq \alpha + 1$, with the converse following by replaying the argument with f^{-1} in place of f. As before, if $\mathrm{RM}^\lambda[q] \geq \alpha + 1$, then each $N \in \mathcal{S}^\lambda_M$ has an extension M_N over which $q \restriction N$ has more than λ many extensions of rank at least α. Notice that for any $N' \in \mathcal{S}^\lambda_{M'}$, $N' = f[N]$ for some $N \in \mathcal{S}^\lambda_M$ and, moreover, that

$$f[q] \restriction N' = f[q] \restriction f[N] = f[q \restriction N]$$

The extensions of $q \upharpoonright N$ to M_N of rank at least α map bijectively under the automorphism f to extensions of $f[q] \upharpoonright N'$ to types over $f[M_N]$ and, by the induction hypothesis, these image types are also of rank at least α. As a result, for any $N' \in \mathcal{S}^\lambda_{M'}$, $\mathrm{RM}^\lambda[f[q] \upharpoonright N'] \geq \alpha + 1$. Thus $\mathrm{RM}^\lambda[q] \geq \alpha + 1$, as desired. $\square$

Trivially,

Proposition 4.5. *[$\leq \lambda$-Local Character] For any $q \in$ ga-S(M), there is an $N \in \mathcal{S}^\lambda_M$ with $RM^\lambda[q \upharpoonright N] = RM^\lambda[q]$.*

Slightly less trivially,

Proposition 4.6. *[Relations Between RM^λ] Whenever $\lambda \leq \mu$, $RM^\mu[q] \leq RM^\lambda[q]$.*

Proof: By induction, with the zero and limit cases still trivial. If $\mathrm{RM}^\mu[q] \geq \alpha + 1$, then for every $N \in \mathcal{S}^\mu_M$, $q \upharpoonright N$ has more than μ extensions to types of RM^μ-rank at least α over some $M_N \succ_{\mathcal{K}} N$. In particular, there are more than λ such extensions and, by the induction hypothesis, they are all of RM^λ-rank at least α. Noticing that $\mathcal{S}^\lambda_M \subseteq \mathcal{S}^\mu_M$, one sees that the same holds for all $N \in \mathcal{S}^\lambda_M$, meaning that, ultimately, $\mathrm{RM}^\lambda[q] \geq \alpha + 1$. $\square$

The ranks $\mathrm{RM}^\lambda[-]$ also have the following attractive property: letting $\mathrm{CB}^\lambda[q]$ denote the topological Cantor-Bendixson rank of a type $q \in X^\lambda_M$, where M is the domain of q,

Proposition 4.7. *For any q (say $q \in$ ga-S(M)), $CB^\lambda[q] \leq RM^\lambda[q]$.*

Proof: We show by induction that $\mathrm{CB}^\lambda[q] \geq \alpha$ implies $\mathrm{RM}^\lambda[q] \geq \alpha$. The zero and limit cases are not interesting. If $\mathrm{CB}^\lambda[q] \geq \alpha + 1$, q is an accumulation point of types of CB^λ-rank at least α. This means, by Proposition 3.36, that every basic open neighborhood of q contains more than λ such types, which is to say that each $q \upharpoonright N$ with $N \in \mathcal{S}^\lambda_M$ has more than λ many extensions to types over M of CB^λ-rank at least α. By the induction hypothesis, these types are of RM^λ-rank at least α, and for each $N \in \mathcal{S}_M^\lambda$ witness the fact that $\mathrm{RM}^\lambda[q \restriction N] \geq \alpha + 1$. Naturally, it follows that $\mathrm{RM}^\lambda[q] \geq \alpha + 1$. □

A quick fact about CB^λ, incidentally:

Proposition 4.8. *If every $q \in X_M^\lambda$ has ordinal CB^λ-rank, isolated points are dense in X_M^λ.*

Proof: Consider a basic open set $U_{p,N}$ in X_M^λ. Let $q \in U_{p,N}$ be of minimal rank, say $\mathrm{CB}^\lambda[q] = \alpha$. There exists a neighborhood $U_{p',N'}$ of q that contains no more than λ types of CB^λ-rank at least α (otherwise, the rank of q would be strictly larger than α). Let N'' be a structure of size λ containing $N \cup N'$, and let $p'' = q \restriction N''$. We then have $U_{p'',N''} \subseteq U_{p,N}$ containing no more than λ types of rank $\geq \alpha$ (because $U_{p'',N''} \subseteq U_{p',N'}$) and, since these were the only kinds of types in $U_{p,N}$ to begin with, $|U_{p'',N''}| \leq \lambda$. It is therefore discrete (by Proposition 3.32), and q must be an isolated point. □

In particular, then, if all of the types over a model $M \in \mathcal{K}$ have ordinal RM^λ-rank, isolated points are dense in X_M^λ. Now, one of the great virtues of the classical Morley rank is that types have unique extensions of same Morley rank: if $p \in \mathrm{S}(A)$ has $\mathrm{RM}[p] = \alpha$ and $B \supseteq A$, there is at most one type $q \in \mathrm{S}(B)$ that extends p and has $\mathrm{RM}[q] = \alpha$. While we do not have a unique extension property for the ranks RM^λ, we do have a very close approximation, Proposition 4.10 below. First, we notice:

Lemma 4.9. *Let $M \prec_{\mathcal{K}} M'$, $q \in$ ga-$S(M)$ with an ordinal RM^λ-rank. There are at most λ extensions $q' \in$ ga-$S(M')$ with $RM^\lambda[q'] = RM^\lambda[q]$.*

Proof: For the sake of notational convenience, say $\mathrm{RM}^\lambda[q] = \alpha$. Let $S = \{q' \in$

ga-S$(M') \,|\, q' \restriction M = q,$ and $\mathrm{RM}^\lambda[q'] = \alpha\}$, and suppose that $|S| > \lambda$. For any $N \in \mathcal{S}^\lambda_M$, each q' is an extension of $q \restriction N$, meaning that $\mathrm{RM}^\lambda[q \restriction N] \geq \alpha + 1$ for all such N. But then $\mathrm{RM}^\lambda[q] \geq \alpha + 1$. Contradiction. □

It is worth noting here that we have made no use of tameness in the discussion above—not in the definition of the ranks RM^λ, nor in the proofs of the properties thereof—and have merely made use of the fact that $\lambda \geq \mathrm{LS}(\mathcal{K})$. From now on, however, it will often be critically important that we have the ability to separate distinct types, leading us to restrict attention to those RM^λ with $\lambda \geq \chi$, the tameness cardinal of the AEC at hand. We will be very explicit in making this assumption whenever it is necessary, beginning with the following essential property of the RM^λ:

Proposition 4.10. *[Quasi-unique Extension] Let $\mathcal{K}$ be χ-tame with $\chi \geq$ LS($\mathcal{K}$), and let $\lambda \geq \chi$. Let $M \prec_{\mathcal{K}} M'$, $q \in$ ga-S(M), and say that RM$^\lambda[q] = \alpha$. Given any rank α extension q' of q to a type over M', there is an intermediate structure M'', $M \prec_{\mathcal{K}} M'' \prec_{\mathcal{K}} M'$, $|M''| \leq |M| + \lambda$, and a rank α extension $p \in$ ga-S(M'') of q with $q' \in$ ga-S(M') as its unique rank α extension.*

Proof: Again, let $S = \{q'' \in \text{ga-S}(M') \,|\, q'' \restriction M = q,$ and $\mathrm{RM}^\lambda[q''] = \alpha\}$. From the lemma, $|S| \leq \lambda$, meaning that S is discrete (by tameness and Corollary 3.35). Thus there is some $N \in \mathcal{S}^\lambda_{M'}$ with the property that $q' \restriction N$ satisfies $U_{q' \restriction N, N} \cap S = \{q'\}$. Take $M'' \in \mathcal{K}$ containing both M and N and with $|M''| \leq |M| + |N| + \mathrm{LS}(\mathcal{K}) \leq |M| + \lambda$, and set $p = q' \restriction M''$. Notice that, by monotonicity,

$$\alpha = \mathrm{RM}^\lambda[q'] \leq \mathrm{RM}^\lambda[p] \leq \mathrm{RM}^\lambda[q] = \alpha$$

so $\mathrm{RM}^\lambda[p] = \alpha$. Naturally, q' is a rank α extension of p. Any such extension q'' of p is also a rank α extension of q, hence in S, and also an extension of $q' \restriction N$. By choice of N, then, we must have $q'' = q'$. □

A possible gloss for this complicated-looking result: while a type $p \in$ ga-S(M) of RM$^\lambda$-rank α may have many extensions of rank α over a model $M' \succ_{\mathcal{K}} M$, we need only expand its domain ever so slightly (adding at most λ elements of M') to guarantee that the rank α extension is unique. As innocuous as this proposition may seem, it is the linchpin of the analysis of stability in this framework, and will see extensive use in Section 4.3. Of course, quasi-unique extension applies only to types with ordinal RM$^\lambda$-rank. To take the fullest possible advantage of this property, then, we would be wise to confine our attention to AECs where RM$^\lambda$ is ordinal-valued on all types associated with the class.

4.2 Stability and λ-total transcendence

In classical model theory, we call a first order theory totally transcendental if all types over subsets of the monster model are assigned ordinal values by the Morley rank. We now define analogous notions for Galois types in AECs, one for each Morley-like rank RM$^\lambda$.

Definition 4.11. We say that $\mathcal{K}$ is λ-*totally transcendental* if for every $M \in \mathcal{K}$ and $q \in$ ga-S(M), RM$^\lambda[q]$ is an ordinal.

Proposition 4.12. *If $\mathcal{K}$ is λ-totally transcendental, it is μ-totally transcendental for all $\mu \geq \lambda$.*

Proof: See Proposition 4.6, the result concerning the relationship between the various Morley ranks. □

Putting this proposition together with Propositions 4.7 and 4.8, we have:

Proposition 4.13. *If $\mathcal{K}$ is λ-totally transcendental, then for any M and $\mu \geq \lambda$, isolated points are dense in X^μ_M.*

As a nice special case, if $\mathrm{LS}(\mathcal{K}) = \aleph_0$, and $\mathcal{K}$ is $\aleph_0$-tame and $\aleph_0$-totally transcendental, then isolated points are dense in each space $X_M^{\aleph_0}$. This bears a certain resemblance to the classical result linking total transcendence to the density of isolated points in the spaces of syntactic types.

Under the assumption of total transcendence, we also get an interesting (and peculiar) new characterization of λ-saturation:

Proposition 4.14. *If $\mathcal{K}$ is μ-totally transcendental and $\lambda > \mu$, then M is λ-saturated iff the following condition holds for all κ with $\mu \leq \kappa < \lambda$:*

$$q \in X_M^\kappa \text{ is isolated if and only if } q = \text{ga-tp}(a/M) \text{ for some } a \in M$$

Proof: ($\Rightarrow$) Suppose that M is λ-saturated. Let q be isolated in X_M^κ with $\mu \leq \kappa < \lambda$, say by the basic open set $U_{p,N}$. From the proof of Proposition 3.10, $\{\text{ga-tp}(a/M) \,|\, a \in M\}$ is a dense subset of X_M^κ, meaning that there is one such type in $U_{p,N} = \{q\}$. Hence $q = \text{ga-tp}(a/M)$ for some $a \in M$. The converse holds, by Claim 3.9.
($\Leftarrow$) In light of Proposition 4.13, isolated points are dense in X_M^κ for $\kappa \geq \mu$. Since every isolated point is of the form $\text{ga-tp}(a/M)$ for some $a \in M$, the set $\{\text{ga-tp}(a/M) \,|\, a \in m\}$ is dense in X_M^κ. As this holds for all κ with $\mu \leq \kappa < \lambda$, we conclude (by Proposition 3.11) that M is λ-saturated. $\square$

In short, a model M is λ-saturated if in any space X_M^κ between X_M^χ and X_M^λ in the spectrum for M, the only types isolated in X_M^κ are those that come by their isolation honestly—they are actually the types of elements of M.

As interesting as these results may be, the true power of the notion of λ-total transcendence lies in the leverage it provides in rather a different area: bounding the number of types over models. We will take advantage of this in Section 4.3 below to prove a variety of results related to Galois stability. Before we give any further thought to λ-totally transcendental AECs, though, it seems natural to inquire whether the notion is, in fact, a meaningful one; that is, whether one can actually find λ-totally transcendental AECs. One can, as the following proposition suggests:

Theorem 4.15. *If $\mathcal{K}$ is χ-tame for some $\chi \geq LS(\mathcal{K})$, and is λ-stable with $\lambda \geq \chi$ and $\lambda^{\aleph_0} > \lambda$, then $\mathcal{K}$ is λ-totally transcendental.*

Proof: Suppose that RM^λ is unbounded; that is, suppose that there is a type p over some $N \prec_{\mathcal{K}} \mathfrak{C}$ with $\mathrm{RM}^\lambda[p] = \infty$. Indeed, we may assume that $|N| \leq \lambda$ (if there is a type of rank ∞ over a larger structure, consider one of its restrictions). We proceed to construct a $\mathcal{K}$-structure $\bar{N}$ of size λ over which there are $\lambda^{\aleph_0}$ many types, thereby establishing that $\mathcal{K}$ is not λ-stable. We first need the following:

Claim 4.16. *For any λ, there exists an ordinal $\alpha_{\mathfrak{C}}$ such that for any q over $M \prec_{\mathcal{K}} \mathfrak{C}$, $RM^\lambda[q] = \infty$ just in case $RM^\lambda[q] \geq \alpha$.*

Proof: The number of all such types can be bounded in terms of $|\mathfrak{C}|$, meaning that the set of their ranks cannot be cofinal in the class of all ordinals. □ (Claim)

We construct $\bar{N}$ as follows:

Step 1: Set $p_\emptyset = p$. Since $p_\emptyset$ satisfies $\mathrm{RM}^\lambda[p_\emptyset] = \infty$, $\mathrm{RM}^\lambda[p_\emptyset] > \alpha$. Hence there is an extension $M \succ_{\mathcal{K}} N$ over which $p_\emptyset$ has more than λ many extensions of rank at least α and thus of rank ∞. Take any λ many of them, say $\{q_i \in \text{ga-S}(M) \,|\, i < \lambda\}$. This set is discrete in X^λ_M, meaning that for each $i < \lambda$, there is an $N_i^{'} \in \mathcal{S}^\lambda_M$ with the property that $q_j \restriction N_i^{'} \neq q_i \restriction N_i^{'}$ whenever $j \neq i$. Let $N_\emptyset$ be a $\mathcal{K}$-substructure of M containing $N \cup (\bigcup_{i<\lambda} N_i)$, $|N_\emptyset| = \lambda$. For each $i < \lambda$, define $p_i = q_i \restriction N_\emptyset$. Notice that:

- $\mathrm{RM}^\lambda[p_i] \geq \mathrm{RM}^\lambda[q_i] = \infty$
- We have $p_i \restriction N = (q_i \restriction N_\emptyset) \restriction N = q_i \restriction N = p_\emptyset$.

• If $p_i = p_j$, then

$$q_i \restriction N_i = (q_i \restriction N_\emptyset) \restriction N_i = p_i \restriction N_i = p_j \restriction N_i = (q_j \restriction N_\emptyset) \restriction N_i = q_j \restriction N_i,$$

in which case we must have $i = j$. That is, the p_i are distinct.

So we have a family $\{p_i \in \text{ga-S}(N_\emptyset) \,|\, i < \lambda\}$ of distinct extensions of $p_\emptyset$, all of rank ∞.

Step $n+1$: For each $\sigma \in \lambda^n$, we have a structure N_σ of size at most λ and a family of distinct types $\{p_{\sigma i} \in \text{ga-S}(N_\sigma) \,|\, i < \lambda\}$, all of rank ∞. For each $i < \lambda$, apply the process of step 1 to each $p_{\sigma i}$. By this method, we obtain (for each $i < \lambda$) an extension $N_{\sigma i} \succ_{\mathcal{K}} N_\sigma$ with $|N_{\sigma i}| = \lambda$ and a family $\{p_{\sigma i j} \in \text{ga-S}(N_{\sigma i}) \,|\, j < \lambda\}$ of distinct extensions of $p_{\sigma i}$, all of rank ∞.

Step ω: Notice that for each $\tau \in \lambda^\omega$, the sequence $N_\emptyset, N_{\tau\restriction 1}, \ldots, N_{\tau\restriction n}, \ldots$ is increasing and, moreover, $\prec_{\mathcal{K}}$-increasing (by coherence and the fact that everything embeds strongly in $\mathfrak{C}$). It follows that $N_\tau := \bigcup_{n<\omega} N_{\tau\restriction n}$ is in $\mathcal{K}$. Notice also that the $p_{\tau\restriction(n+1)} \in \text{ga-S}(N_{\tau\restriction n})$ for $n < \omega$ form an increasing sequence of types over the structures in this union. By (ω, ∞)-compactness of AECs (see the remark at the end of Section 2.1), there is a type $p_\tau \in \text{ga-S}(N_\tau)$ with the property that for all $n < \omega$, $p_\tau \restriction N_{\tau\restriction n} = p_{\tau\restriction(n+1)}$. Let $\bar{N}$ be a structure of size λ containing the union

$$\bigcup_{\sigma \in \lambda^{<\omega}} N_\sigma$$

(which is possible, since the union is of size at most $\lambda \cdot \lambda^{<\omega} = \lambda \cdot \lambda = \lambda$). For each $\tau \in \lambda^\omega$, let q_τ be an extension of p_τ to a type over $\bar{N}$. There are $\lambda^{\aleph_0}$ many such types, all over $\bar{N}$—it remains only to show that they are distinct. To that end, suppose that $\tau, \tau' \in \lambda^\omega$, $\tau \neq \tau'$. It must be the case that for some $\sigma \in \lambda^{<\omega}$, $\tau = \sigma i \ldots$ and $\tau' = \sigma j \ldots$ with $i \neq j$. Then

$$q_\tau \restriction N_\sigma = (q_\tau \restriction N_\tau) \restriction N_\sigma = p_\tau \restriction N_\sigma = p_{\sigma i}$$

Similarly, we have

$$q_{\tau'} \upharpoonright N_\sigma = (q_{\tau'} \upharpoonright N_{\tau'}) \upharpoonright N_\sigma = p_{\tau'} \upharpoonright N_\sigma = p_{\sigma j}$$

Since $p_{\sigma i}$ and $p_{\sigma j}$ are distinct by construction, $q_\tau \neq q_{\tau'}$, and the types are distinct as claimed. By our assumption that $\lambda^{\aleph_0} > \lambda$, then, $\mathcal{K}$ is not stable in λ. $\Box$

So the notion of λ-total transcendence is far from vacuous: totally transcendental AECs exist. Indeed, they exist in great abundance.

4.3 Transfer Results

As promised, we will now employ λ-total transcendence as a tool to analyze the proliferation of types over models in AECs. In fact, we introduce two methods of bounding the number of types over large models, the first of which is captured by the following result:

Theorem 4.17. *Let $\mathcal{K}$ be χ-tame for some $\chi \geq LS(\mathcal{K})$, and λ-totally transcendental with $\lambda \geq \chi + |L(\mathcal{K})|$. Then for any structure M with $|M| > \lambda$, $|ga\text{-}S(M)| \leq |M|^\lambda$.*

Proof: Take $q \in \text{ga-S}(M)$, $\text{RM}^\lambda[q] = \alpha$. It must be the case that $\text{RM}^\lambda[q \upharpoonright N] = \alpha$ for some $N \in \mathcal{S}^\lambda_M$ whence, by Proposition 4.10, there is an $N' \succ_{\mathcal{K}} N$ with $|N'| \leq \lambda$ such that $q \upharpoonright N'$ has rank α and has a unique rank α extension q over M. Hence

$$|\{q \in \text{ga-S}(M) \,|\, \text{RM}^\lambda[q] = \alpha\}| \leq |\{p \text{ over } N \in \mathcal{S}^\lambda_M \,|\, \text{RM}^\lambda[p] = \alpha\}|$$

and thus

$$\begin{aligned} |\text{ga-S}(M)| &= |\textstyle\bigcup_\alpha \{q \in \text{ga-S}(M) \,|\, \text{RM}^\lambda[q] = \alpha\}| \\ &\leq |\textstyle\bigcup_\alpha \{p \text{ over } N \in \mathcal{S}^\lambda_M \,|\, \text{RM}^\lambda[p] = \alpha\}| \\ &= |\{p \text{ over } N \in \mathcal{S}^\lambda_M\}| \end{aligned}$$

By Proposition 2.10, for any $N \in \mathcal{S}^{\lambda}_{M}$, $|\text{ga-S}(N)| \leq 2^{\lambda}$. Hence

$$\begin{aligned} |\text{ga-S}(M)| &\leq 2^{\lambda} \cdot |\mathcal{S}^{\lambda}_{M}| \\ &\leq 2^{\lambda} \cdot |M|^{\lambda} \\ &= |M|^{\lambda} \end{aligned}$$

So we are done. □

A major consequence of this theorem is:

Corollary 4.18. *If $\mathcal{K}$ is χ-tame for some $\chi \geq LS(\mathcal{K})$, and λ-totally transcendental with $\lambda \geq \chi + |L(\mathcal{K})|$, then for any μ satisfying $\mu^{\lambda} = \mu$, $\mathcal{K}$ is μ-stable.*

In light of Theorem 4.15, we may eliminate any mention of λ-total transcendence and translate the corollary above into a pure upward stability transfer result:

Proposition 4.19. *If $\mathcal{K}$ is χ-tame for some $\chi \geq LS(\mathcal{K})$, and λ-stable with $\lambda \geq \chi + |L(\mathcal{K})|$ and $\lambda^{\aleph_0} > \lambda$, then $\mathcal{K}$ is stable in every μ with $\mu^{\lambda} = \mu$.*

Proof: Under the assumptions on λ, stability in λ implies λ-total transcendence of $\mathcal{K}$ (by the theorem invoked above). We are then in the situation of Corollary 4.18, and the result follows. □

It should be noted that this result is weaker than one found in [14], obtained by other methods, which dispenses with the assumption that $\lambda^{\aleph_0} > \lambda$, inferring stability in μ with $\mu^{\lambda} = \mu$ from stability in any λ larger than the tameness cardinal. The potential advantages of our method become slightly clearer if we reinvent Proposition 4.19 in such a way as to make clear the larger project for which it should be the jumping off point:

Theorem 4.20. *Let $\mathcal{K}$ be χ-tame for some $\chi \geq LS(\mathcal{K})$, and λ-stable with $\lambda \geq \chi$ and $\lambda^{\aleph_0} > \lambda$. Then there exists $\kappa \leq \lambda$ such that $\mathcal{K}$ is μ-stable in every μ satisfying $\mu^{\kappa} = \mu$.*

Recall that if $\mathcal{K}$ is stable in such a λ, it is λ-totally transcendental, and we may take

$$\kappa = \min\{\nu \,|\, \mathcal{K} \text{ is } \nu\text{-totally transcendental}\}.$$

For this to be an actual improvement on existing results, of course, one would need to be able to give a meaningful bound on κ, the cardinal at which total transcendence first occurs, that places it well below λ. Given that there is something of an analogy between this cardinal and the cardinal parameter of nonforking in elementary classes, there is some hope that this can be achieved, transforming the theorem above into a stability spectrum result reminiscent of those familiar from first order model theory.

There is a different tack to be taken, though. We introduce a second method for bounding the number of types over models of size larger than the cardinal of total transcendence:

Theorem 4.21. *Let $\mathcal{K}$ be χ-tame for some $\chi \geq LS(\mathcal{K})$, and λ-totally transcendental with $\lambda \geq \chi$. For any $M \in \mathcal{K}$ with $cf(|M|) > \lambda$,*

$$|ga\text{-}S(M)| \leq |M| \cdot \sup\{|ga\text{-}S(N)| \,|\, N \prec_{\mathcal{K}} M, |N| < |M|\}.$$

Proof: Take a filtration of M:

$$M_0 \prec_{\mathcal{K}} M_1 \prec_{\mathcal{K}} \ldots \prec_{\mathcal{K}} M_\alpha \prec_{\mathcal{K}} \ldots \prec_{\mathcal{K}} M$$

for $\alpha < |M|$, with $|M_\alpha| < |M|$ for all α and $M = \bigcup_{\alpha<|M|} M_\alpha$. Let $q \in$ ga-S(M). By λ-total transcendence, $\text{RM}^\lambda[q] = \beta$ for some ordinal β and, moreover, this is witnessed by a restriction to a small submodel of M. That is, there is a submodel $N \in \mathcal{S}_M^\lambda$ with $\text{RM}^\lambda[q \upharpoonright N] = \beta$. By Proposition 4.10, there is an intermediate extension $N \prec_{\mathcal{K}} N' \prec_{\mathcal{K}} M$ with $|N'| = \lambda$ such that $q \upharpoonright N'$ has unique rank β extension q over M. Since $\text{cf}(|M|) > \lambda$, $N' \subseteq M_\alpha$ for some α (and, by coherence, $N' \prec_{\mathcal{K}} M_\alpha$).

Clearly, the type $q \upharpoonright M_\alpha$ has q as its only rank β extension over M. By a computation similar to that found in the proof of Proposition 4.17, then, we have

$$|\text{ga-S}(M)| \leq |M| \cdot \sup\{|\text{ga-S}(M_\alpha)| \,|\, \alpha < |M|\}$$

The inequality in the statement of the theorem is a trivial consequence. $\square$

In essence, the theorem asserts that for models M of cardinality μ with $\text{cf}(\mu) > \lambda$, the equality $|\text{ga-S}(M)| = \mu$ fails only if there is a submodel $N \prec_{\mathcal{K}} M$ with $|N| < \mu$ and $|\text{ga-S}(N)| > \mu$. To ensure that this does not occur—that we have, in short, stability in μ—we need considerably less than full stability in the cardinals below μ. To be precise:

Theorem 4.22. *Let $\mathcal{K}$ be χ-tame for some $\chi \geq LS(\mathcal{K})$ and λ-totally transcendental with $\lambda \geq \chi$, and let μ satisfy cf(μ) $> \lambda$. If for every $M \in \mathcal{K}_{<\mu}$, $|$ga-S(M)$| \leq \mu$, $\mathcal{K}$ is μ-stable.*

Proof: For any $M \in \mathcal{K}_\mu$,

$$|\text{ga-S}(M)| \leq |M| \cdot \sup\{|\text{ga-S}(N)| \,|\, N \prec_{\mathcal{K}} M, |N| < |M|\} \leq |M| \cdot |M| = |M|. \quad \square$$

We may tighten the statement up a bit, once we notice the following:

Lemma 4.23. *If there is a set S of cardinals cofinal in an interval $[\kappa, \mu)$ which has the property that for every $M \in \mathcal{K}$ with $|M| \in S$, $|$ga-S(M)$| \leq \mu$, then $|$ga-S(M)$| \leq \mu$ for every $M \in \mathcal{K}_{<\mu}$.*

Proof: Let $M \in \mathcal{K}_{<\mu}$. If $|M| \in S$, we are done. If not, take $\nu \in S$ with $\nu > |M|$, and let M' be a strong extension of cardinality ν, $M \hookrightarrow_{\mathcal{K}} M' \hookrightarrow_{\mathcal{K}} \mathfrak{C}$. By assumption, $|\text{ga-S}(M')| \leq \mu$. Every type over M has an extension to a type over M', and, of course, distinct types must have distinct extensions. Hence

$$|\text{ga-S}(M)| \leq |\text{ga-S}(M')| \leq \mu. \quad \square$$

So Theorem 4.22 becomes

Theorem 4.24. *Let $\mathcal{K}$ be χ-tame for some $\chi \geq LS(\mathcal{K})$ and λ-totally transcendental with $\lambda \geq \chi$, and let μ satisfy cf(μ) $> \lambda$. If there is a set S of cardinals cofinal in an interval $[\kappa, \mu)$ which has the property that for every $M \in \mathcal{K}$ with $|M| \in S$, $|ga\text{-}S(M)| \leq \mu$, $\mathcal{K}$ is μ-stable.*

The assumption on the number of types over models of cardinality less than μ in the theorem above is very weak—it is certainly more than sufficient to assume that $\mathcal{K}$ is ν-stable for each ν in a cofinal set below μ. In particular, we have:

Corollary 4.25. *If $\mathcal{K}$ is χ-tame for some $\chi \geq LS(\mathcal{K})$ and λ-totally transcendental with $\lambda \geq \chi$, μ satisfies cf(μ) $> \lambda$, and $\mathcal{K}$ is stable on a set of cardinals cofinal in an interval $[\kappa, \mu)$, then $\mathcal{K}$ is μ-stable.*

Using Theorem 4.15 again, we may recast the corollary as a stability transfer result that makes no mention of total transcendence. Though the product of this rewriting is a trivial consequence of the foregoing discussion (and, indeed, an extremely special case of Theorem 4.24), it nonetheless generalizes a state-of-the-art result: Theorem 2.1 in [5]. First, the statement of the theorem:

Theorem 4.26. *If $\mathcal{K}$ is χ-tame for some $\chi \geq LS(\mathcal{K})$ and λ-stable with $\lambda \geq \chi$ and $\lambda^{\aleph_0} > \lambda$, then for any μ with cf(μ) $> \lambda$, if $\mathcal{K}$ is stable on a set of cardinals cofinal in an interval $[\kappa, \mu)$, $\mathcal{K}$ is μ-stable.*

If we restrict our attention to the special case in which the AEC is $\aleph_0$-stable, and assume more stability below μ than we need, strictly speaking, we have:

Corollary 4.27. *If $\mathcal{K}$ has $LS(\mathcal{K}) = \aleph_0$ and is $\aleph_0$-tame and $\aleph_0$-stable, then for any μ with cf(μ) $> \aleph_0$, if $\mathcal{K}$ is κ-stable for all $\kappa < \mu$, $\mathcal{K}$ is μ-stable.*

In other words: in an AEC satisfying the hypotheses of the corollary, a run of stability will never come to an end at a cardinal μ of uncountable cofinality; it will, in fact, include μ, then μ^+, then μ^{++}, and so on. This remarkable fact is the aforementioned result of [5] (which also appears as Theorem 11.11 in [2]). What we have discerned by our method—and distilled in Theorem 4.26—is that it is not stability in $\aleph_0$ which is critical, but rather stability in any cardinal λ satisfying $\lambda^{\aleph_0} > \lambda$. Moreover, it is not necessary to have stability in every cardinality below the cardinal of interest, μ, but rather to have stability (or, indeed, somewhat less than stability) in a cofinal sequence of cardinals below μ.

We now shift our frame of reference. Thus far we have limited ourselves to AECs that are tame in some cardinal χ. The time has come to consider how this picture may differ in an AEC that is weakly χ-tame, rather than χ-tame. The most essential difference is that the argument for Theorem 4.15 no longer works, meaning that we can no longer infer total transcendence from stability. In a sense, then, we must treat total transcendence as a property in its own right, independent from the more conventional properties of AECs. On the other hand, the argument for Theorem 4.21 still goes through, provided the model in question is saturated. That is,

Proposition 4.28. *Let $\mathcal{K}$ be weakly χ-tame for some $\chi \geq LS(\mathcal{K})$, and μ-totally transcendental with $\mu \geq \chi$. If $M \in \mathcal{K}$ is a saturated model with $cf(|M|) > \mu$, then*

$$|ga\text{-}S(M)| \leq |M| \cdot \sup\{|ga\text{-}S(N)| \mid N \prec_{\mathcal{K}} M, |N| < |M|\}$$

We can use this to get bounds on the number of types over more general (that is, not necessarily saturated) models, provided that there are enough small saturated extensions in $\mathcal{K}$. In particular, if we can replace a model $M \in \mathcal{K}_\lambda$ with a saturated extension $M' \in \mathcal{K}_\lambda$, we have

$$|\text{ga-S}(M)| \leq |\text{ga-S}(M')|$$

where the latter is governed by the bound in the proposition above. Given an adequate supply of such extensions, then, we can prove results similar to those found above, but now for weakly tame AECs. The existence of saturated extensions also has the following consequence:

Lemma 4.29. *If every model in $\mathcal{K}_\lambda$ has a saturated extension in $\mathcal{K}_\lambda$, then for any $M \in \mathcal{K}_{<\lambda}$, $|\text{ga-S}(M)| \leq \lambda$.*

Proof: Take $M \in \mathcal{K}_{<\lambda}$. Let M' be a strong extension of M of cardinality λ. By assumption, M' has a saturated extension $\bar{M}$ of cardinality λ. This model, $\bar{M}$, realizes all types over submodels of size less than λ, hence realizes all types over the original model, M. It follows that $|\text{ga-S}(M)| \leq |\bar{M}| = \lambda$. □

Notice that this is precisely the type-counting condition from which we are able to infer stability using the bound in Proposition 4.28. Hence we have

Theorem 4.30. *Let $\mathcal{K}$ be weakly χ-tame for some $\chi \geq LS(\mathcal{K})$, and μ-totally transcendental with $\mu \geq \chi$. Suppose that λ is a cardinal with $cf(\lambda) > \mu$, and that every $M \in \mathcal{K}_\lambda$ has a saturated extension $M' \in \mathcal{K}_\lambda$. Then $\mathcal{K}$ is λ-stable.*

Proof: Let $M \in \mathcal{K}_\lambda$. By assumption, we may replace M with a saturated extension M' of cardinality λ, over which there can be at most λ types, by Proposition 4.28 and Lemma 4.29. As noted above, $|\text{ga-S}(M)| \leq |\text{ga-S}(M')|$, and we are done. □

As a simple special case:

Corollary 4.31. *Suppose $\mathcal{K}$ is weakly χ-tame for some $\chi \geq LS(\mathcal{K})$ and μ-totally transcendental with $\mu \geq \chi$, and suppose that λ is a regular cardinal with $\lambda > \mu$. If $\mathcal{K}$ is stable on an interval $[\kappa, \lambda)$, $\mathcal{K}$ is λ-stable.*

Proof: Let $M \in \mathcal{K}$ be of cardinality λ. By Proposition 2.13, M has a saturated extension $M' \in \mathcal{K}$ which is also of cardinality λ. Using the bound in Proposition 4.28, we get $|\text{ga-S}(M')| = \lambda$, whence $\lambda = |M| \leq |\text{ga-S}(M)| \leq \lambda$, and thus $|\text{ga-S}(M)| = \lambda$. The result follows. □

In particular,

Corollary 4.32. *Suppose $\mathcal{K}$ is weakly χ-tame for some $\chi \geq LS(\mathcal{K})$, μ-totally transcendental with $\mu \geq \chi$, and λ-stable with $\lambda \geq \mu$. Then $\mathcal{K}$ is λ^+-stable.*

This is weaker than a result of [5], which infers λ^+-stability from λ-stability in any AEC that is weakly tame in $\chi \leq \lambda$, with no additional assumptions. The methods of [5]—splitting, limit models—are better adapted to the task of proving local transfer results of this form in the context of weakly tame AECs. The machinery of ranks and total transcendence provides us with leverage of a different sort, well suited to the task of producing partial spectrum results of a more global nature.

At any rate, we saw in Theorem 4.30 that in a μ-totally transcendental AEC, even if merely weakly tame, we may infer stability in λ given an adequate supply of saturated extensions of cardinality λ. As it happens, the category-theoretic notion of weak λ-stability covered in Section 2.2 ensures that this requirement is met. Hearkening back to Proposition 2.32 (and anticipating the discussion in Chapter 5), we note that any AEC is weakly λ-stable for many cardinals λ. If it is totally transcendental, it is in fact λ-stable—λ-Galois stable—for many such λ. As we will see, these considerations yield a partial stability spectrum for weakly tame AECs.

CHAPTER 5

Accessible Categories and AECs

Given an AEC $\mathcal{K}$ (with or without the amalgamation and joint embedding properties), we regard it as a category in the only natural way: the objects are the models $M \in \mathcal{K}$, and the morphisms are precisely the strong embeddings. Since there is no serious risk of confusion, we will also refer to the category thus obtained as $\mathcal{K}$. In the next few pages, we spare a thought for the ways in which certain category-theoretic notions—especially ones associated with the study of accessible categories—manifest themselves in such categories, and the ways in which they interact with model-theoretic properties of the associated AEC.

5.1 AECs as Accessible Categories

The first step in our analysis involves showing the following:

Theorem 5.1. *Let $\mathcal{K}$ be an AEC. Then $\mathcal{K}$ is μ-accessible for every regular cardinal $\mu > LS(\mathcal{K})$. In particular, $\mathcal{K}$ is $LS(\mathcal{K})^+$-accessible.*

Our task, then, is to show that for each regular cardinal $\mu > \mathrm{LS}(\mathcal{K})$, $\mathcal{K}$ contains a set (up to isomorphism) of μ-presentable objects, every model in $\mathcal{K}$ can be obtained as a μ-directed colimit of μ-presentable objects, and $\mathcal{K}$ is closed under μ-directed colimits. We accomplish this through a series of easy lemmas. First:

Lemma 5.2. *Let $M \in \mathcal{K}$. For any regular $\mu > LS(\mathcal{K})$, M is a μ-directed union of its strong submodels of size less than μ.*

Proof: Consider the diagram consisting of all submodels of M of size less than μ and with arrows the strong inclusions. To check that this diagram is μ-directed, we must show that any collection of fewer than μ many such submodels have a common extension also belonging to the diagram. Let $\{M_\alpha \,|\, \alpha < \nu\}$, $\nu < \mu$, be such a collection. Since μ is regular, $\sup\{|M_\alpha| \,|\, \alpha < \nu\} < \mu$, whence

$$\Big|\bigcup_{\alpha<\nu} M_\alpha\Big| \leq \nu \cdot \sup\{|M_\alpha| \,|\, \alpha < \nu\} < \nu \cdot \mu = \mu$$

This set will be contained in a submodel $M' \prec_{\mathcal{K}} M$ of cardinality less than μ, by the Downward Löwenheim Skolem Property. For each $\alpha < \nu$, $M_\alpha \prec_{\mathcal{K}} M$ and $M_\alpha \subseteq M'$. Since $M' \prec_{\mathcal{K}} M$, coherence implies that $M_\alpha \prec_{\mathcal{K}} M'$. So we are done. $\square$

Moreover,

Lemma 5.3. *For any regular cardinal $\mu > LS(\mathcal{K})$, a model $M \in \mathcal{K}$ is μ-presentable if and only if $|M| < \mu$. In particular, M is $LS(\mathcal{K})^+$-presentable if and only if $|M| \leq LS(\mathcal{K})$.*

Proof: ($\Rightarrow$) Suppose that M is μ-presentable, and consider the identity map $M \hookrightarrow_{\mathcal{K}} M$. As we saw in the previous lemma, M is a μ-directed union of its submodels of size strictly less than μ. By μ-presentability of M, the identity map must factor through one of the inclusions $M' \hookrightarrow_{\mathcal{K}} M$ in the colimit cocone. Since all maps in the category are injective, M can have cardinality no greater than that of the model M'. Hence $|M| < \mu$.

($\Leftarrow$) Suppose $|M| = \nu < \mu$. Let M' be a μ-directed colimit, say

$$M' = \mathrm{Colim}_{i \in I} M_i$$

with I a μ-directed poset, connecting maps $\phi_{ij} : M_i \hookrightarrow_{\mathcal{K}} M_j$ for $i \leq j$, and colimit cocone maps $\phi_i : M_i \hookrightarrow_{\mathcal{K}} M'$. That is, for each $i \leq j$ in I, we have the commutative triangle

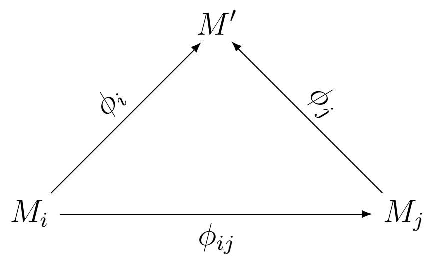

Consider an embedding $f : M \hookrightarrow_{\mathcal{K}} M'$. The image $f[M]$ is a strong submodel of M, and is of cardinality $\nu < \mu$. Since $\mathcal{K}$ is a concrete category, the submodels $\phi_i[M_i]$ of M' cover M', meaning that for each $m \in f[M]$ we may choose a $\phi_{i_m}[M_{i_m}]$ containing it. By μ-directedness of I, there is a $j \in I$ with $j \geq i_m$ for all $m \in f[M]$. By the commutativity condition above, one can see that $\phi_{i_m}[M_{i_m}] \subseteq \phi_j[M_j]$ for all m, meaning that $f[M] \subseteq \phi_j[M_j]$ and, by coherence, $f[M] \prec_{\mathcal{K}} \phi_j[M_j]$. Hence the embedding $f : M \hookrightarrow_{\mathcal{K}} M'$ factors through $\phi_j : M_j \hookrightarrow_{\mathcal{K}} M'$ as

$$M \xrightarrow{\phi_j^{-1} \circ f} M_j \xrightarrow{\phi_j} M'.$$

This factorization is unique: for any other factorization map $g : M \to M_j$, we have $\phi_j \circ (\phi_j^{-1} \circ f) = \phi_j \circ g$ and, since ϕ_j is a monomorphism, it follows that $\phi_j^{-1} \circ f = g$. This means, of course, that M is μ-presentable. $\square$

The punchline of all this is:

Lemma 5.4. *For any regular $\mu >$ LS($\mathcal{K}$), $\mathcal{K}$ contains a set of μ-presentables, namely $\mathcal{K}_{<\mu}$, and every model in $\mathcal{K}$ is a μ-directed colimit of objects in $\mathcal{K}_{<\mu}$.*

Recall the following fact, which we remarked upon in Section 2.1 above:

Lemma 5.5. *$\mathcal{K}$ is closed under directed colimits.*

Since every μ-directed diagram is, in particular, directed, we can complete the proof the theorem:

Lemma 5.6. *For any regular cardinal $\mu > LS(\mathcal{K})$, $\mathcal{K}$ is closed under μ-directed colimits.*

One often encounters assertions (here and elsewhere) to the effect that AECs are the result of extracting the purely category-theoretic content of elementary classes, preserving the essence of the elementary submodel relation while dispensing with syntax and certain properties—such as compactness—that are typically derived from the ambient logic. We obtain very definite confirmation of this claim if we compare Theorem 5.1 above with the result of Rosický cited in Section 2.2, which is, for reference:

Proposition 5.7. *(Rosický, [22]) Given a first order theory T in language $L(T)$ and $\mathbf{Elem}(T)$ the category with objects the models of T and morphisms the elementary embeddings, then for any regular $\mu > |L(T)|$, $\mathbf{Elem}(T)$ is μ-accessible, and $M \in \mathcal{K}$ is μ-presentable if and only if $|M| < \mu$.*

At the most fundamental level, then, AECs and elementary classes do have the same category-theoretic structure.

There is still more to the story, as we must also consider the way in which an AEC $\mathcal{K}$ sits inside the ambient category of $L(\mathcal{K})$-structures, whose objects are $L(\mathcal{K})$-structures and whose morphisms are precisely the injective $L(\mathcal{K})$-homomorphisms (which both preserve and reflect the relations in $L(\mathcal{K})$). The goal is to produce a category-theoretic axiomatization that, in any category $L-\mathbf{Struct}$, picks out all the subcategories corresponding to AECs in the signature L. A very elegant axiomatization of this form appears in [18], although the scope of that piece is slightly broader—one considers realizations of the axioms in base categories that generalize (but still closely resemble) categories of the form $L-\mathbf{Struct}$, with the aim of generalizing not only AECs, but also abstract metric classes. We note that, while we will not pursue this line of inquiry, our axiomatization is equally well suited for this purpose and is, in fact, perfectly equivalent. It has the added benefit, though, of condensing a number of axioms from [18] under the heading of accessibility, and thereby making clear the connection between AECs and the existing body of work on accessible categories. This perspective also clarifies, for example, that the abstract notion of size laid out in the aforementioned piece corresponds to the pre-existing and well-understood notion of presentability.

For our axiomatization, we introduce two definitions, the second of which is drawn from [18]:

Definition 5.8. Fix a category $\mathbf{B}$ and subcategory $\mathbf{C}$.

- We say that $\mathbf{C}$ is a *replete* subcategory of $\mathbf{B}$ if for every M in $\mathbf{C}$ and every isomorphism $f : M \to N$ in the larger category $\mathbf{B}$, both f and N are in $\mathbf{C}$.
- We say that $\mathbf{C}$ is a *coherent* subcategory of $\mathbf{B}$ if for every commutative diagram

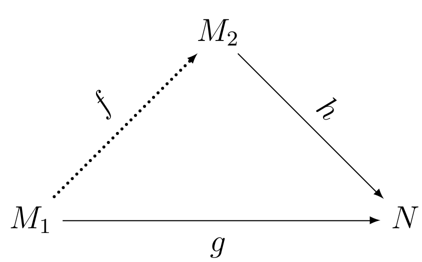

  with h and g (hence also their domains and codomains) in $\mathbf{C}$ and with f in $\mathbf{B}$, then in fact f is in $\mathbf{C}$.

Purely from Theorem 5.1 and the axioms for AECs,

Proposition 5.9. *An AEC $\mathcal{K}$ is a replete, coherent subcategory of $L(\mathcal{K})-$**Struct** which is μ-accessible for all $\mu > LS(\mathcal{K})$ and has all directed colimits. Moreover, the directed colimits are computed as in $L(\mathcal{K})-$**Struct**.*

Now, consider a category $L-$**Struct**, L a finitary signature, consisting of L-structures and injective L-homomorphisms, as before. The natural question: given a replete, coherent subcategory of $L-$**Struct** with all directed colimits (computed as in $L-$**Struct**) that is μ-accessible for all μ strictly larger than some cardinal λ, can it be regarded as an AEC? The answer is yes: for any such subcategory $\mathbf{C}$, consider the class consisting of its objects (call it $\mathbf{C}$ as well), with relation $\prec_{\mathbf{C}}$ defined by the condition that $M \prec_{\mathbf{C}} N$ if and only if $M \subseteq_L N$ and the inclusion map is a $\mathbf{C}$-morphism.

Claim 5.10. *The class $\mathbf{C}$ is an AEC.*

Proof: As in the proof of Lemma 4 in [18], the verification of most of the AEC axioms is trivial. The relation $\prec_{\mathbf{C}}$ is transitive, and certainly refines the substructure relation. Coherence and closure under isomorphism hold by assumption, and the union of chains axioms are easily verified as well. As for the Löwenheim-Skolem property, let $M \in \mathbf{C}$, and let $A \subseteq M$. Consider $\mu = |A| + \lambda$. The cardinal μ^+ is regular and $\mu^+ > \lambda$, meaning that $\mathbf{C}$ is μ^+-accessible. This means, in turn, that every object M is a μ^+-filtered colimit of μ^+-presentable objects, and thus the μ^+-directed union of the images of these μ^+-presentable objects under the cocone maps. All of these images are, of course, strong submodels of M. Since $|A| \leq |A| + \lambda < \mu^+$, the μ^+-directedness of the union implies that A is contained in one of the structures in the union, say N. As N is μ^+-presentable, it is, by the proof of the "only if" direction of Lemma 5.3 above, of cardinality at most $\mu = |A| + \lambda$. One can see, then, that λ will do as $\mathrm{LS}(\mathbf{C})$. $\square$

If we replace $L-\mathbf{Struct}$ with a particular category of metric L-structures (as in [18]), our axioms completely describe the abstract metric classes in the signature L.

5.2 Model Theory and Category Theory: Correspondences

We turn our attention now to the task of providing a dictionary between the language of accessible categories and that of AECs. We will primarily be interested in examining the translations of the category theoretic notions introduced in Section 2.2, but originally defined in [22]. The latter piece is, of course, concerned with accessible categories with directed colimits—almost AECs, as we now know. That piece contains no explicit mention of AECs, however, so our work is very much still to be done.

First, we note that the amalgamation and joint embedding properties for AECs (see Definition 2.1) and the amalgamation and joint embedding properties for arbitrary categories introduced in Section 2.2 are identical: $\mathcal{K}$ has the amalgamation property as an AEC if and only if it satisfies the category-theoretic analogue, and similarly for joint embedding. We noticed as much when presenting these analogues in Section 2.2.

Before we proceed to more interesting correspondences, we lay out two basic facts that will come in handy in simplifying the diagrams that crop up in our investigations and will, in particular, allow us to replace (without loss of generality) certain strong embeddings by strong inclusions.

Remark 5.11. 1. Any strong embedding $f : M_0 \hookrightarrow_{\mathcal{K}} M$ factors as an isomorphism $M_0 \hookrightarrow_{\mathcal{K}} f[M_0]$ followed by the strong inclusion $f[M_0] \hookrightarrow_{\mathcal{K}} M$.

2. Given a strong embedding $f : M_0 \hookrightarrow_{\mathcal{K}} M$, there is an extension M_1 of M_0

isomorphic to M. Moreover, we may take the isomorphism $g : M \to M_1$ to be inverse to f on $f[M_0]$; that is, $g \circ f : M_0 \hookrightarrow_{\mathcal{K}} M_1$ fixes M_0.

Now we may begin. Unless otherwise specified, λ is understood to be a regular cardinal. We first consider λ-saturation of the sort introduced in Section 2.2.

Proposition 5.12. *For any AEC $\mathcal{K}$ and $M \in \mathcal{K}$, M is λ-saturated if and only if it is λ-model homogeneous.*

Proof: ($\Rightarrow$) Let $N \prec_{\mathcal{K}} M$ and $N \prec_{\mathcal{K}} N'$, with $|N|$ and $|N'|$ strictly less than λ. Notice that, by Lemma 5.3, N and N' are λ-presentable. Then, by λ-saturation of M, there is a strong embedding $h : N' \hookrightarrow_{\mathcal{K}} M$ such that the following diagram commutes:

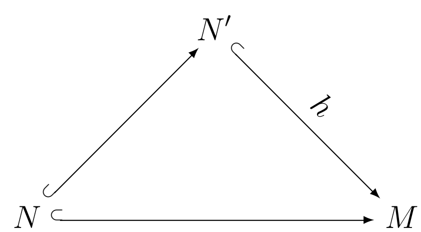

with $N \hookrightarrow_{\mathcal{K}} N'$ and $N \hookrightarrow_{\mathcal{K}} M$ the inclusions. This says precisely that h is an embedding of N' into M fixing N. So M is λ-model homogeneous.

($\Leftarrow$) Using one application of each of the facts in Remark 5.11, one can see that it suffices to consider diagrams of strong inclusions

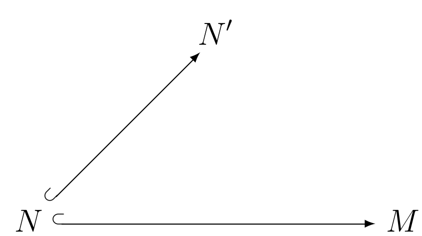

with N, and N' λ-presentable. Then $N, N' \in \mathcal{K}_{<\lambda}$ and λ-model homogeneity of M guarantees the existence of a strong embedding $h : N' \hookrightarrow_{\mathcal{K}} M$ fixing N, and therefore making the relevant diagram commute. Hence M is λ-saturated. $\square$

Recalling that λ-model homogeneity implies λ-Galois-saturation in any AEC, and that the converse holds in AECs with amalgamation (see Proposition 2.12), we get

Corollary 5.13. *For any AEC $\mathcal{K}$, $\lambda > LS(\mathcal{K})$, and $M \in \mathcal{K}$, if M is λ-saturated, then M is λ-Galois-saturated. Moreover, if $\mathcal{K}$ has the amalgamation property, M is λ-saturated if and only if M is λ-Galois-saturated.*

In Section 2.2, we considered the manifestations of λ-purity in several contexts, but most notably in elementary classes, where we observed that an elementary inclusion of a model M in a model M' is λ-pure only if for every $A \subseteq M$ with $|A| < \lambda$ and every $p \in \mathrm{S}(A)$, if p is realized in M', then it also realized in M. That is, M is λ-saturated relative to M'. We will obtain a similar result for AECs, but first note that relative λ-model homogeneity is a more obvious analogue in our context. In particular:

Proposition 5.14. *For $\lambda \geq LS(\mathcal{K})$, a strong inclusion $M \hookrightarrow_{\mathcal{K}} M'$ is λ-pure if and only if M is λ-model homogeneous relative to M': for any $N \prec_{\mathcal{K}} M$ and $N \prec_{\mathcal{K}} N' \prec_{\mathcal{K}} M'$ with $N, N' \in \mathcal{K}_{<\lambda}$, there is an embedding of N' into M fixing N.*

Proof: ($\Rightarrow$) Suppose that $M \hookrightarrow_{\mathcal{K}} M'$ is a λ-pure inclusion. Let $N \prec_{\mathcal{K}} M$ with $|N| < \lambda$ and N' with $N \prec_{\mathcal{K}} N' \prec_{\mathcal{K}} M'$ and $|N'| < \lambda$. The various inclusions yield the commutative square

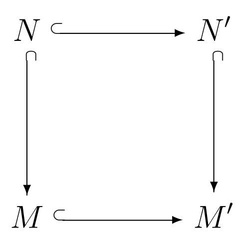

By λ-purity of the bottom inclusion, there is a strong embedding $h : N' \hookrightarrow_{\mathcal{K}} M$ that makes the upper triangle of

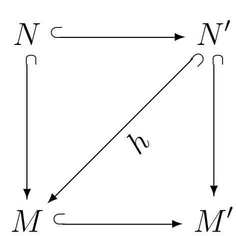

commute. This commutativity condition simply means that $h : N' \hookrightarrow_{\mathcal{K}} M$ fixes N, so we are done.

($\Leftarrow$) First, a little diagram wrangling: in light of Remark 5.11, given a diagram of the form in the definition of λ-purity, we may reinvent it as

$$\begin{array}{ccc} u[N] & \xhookrightarrow{v \circ g \circ u^{-1}} & v[N'] \\ \big\downarrow & & \big\downarrow \\ M & \hookrightarrow & M' \end{array}$$

where the vertical maps are strong inclusions (as is the bottom map, by assumption). If we can produce a map $h' : v[N'] \hookrightarrow_{\mathcal{K}} M$ that makes the upper left hand triangle commute—that is, such that $h' \circ (v \circ g \circ u^{-1})$ is just the inclusion of $u[N]$ in M—the map $h' \circ v$ will serve the same purpose in the original diagram. So, really, it suffices to consider diagrams of the form

$$\begin{array}{ccc} N & \longrightarrow & N' \\ \big\downarrow & & \big\downarrow \\ M & \hookrightarrow & M' \end{array}$$

with the bottom and vertical morphisms strong inclusions. In that case, of course, the upper map must be an inclusion as well (and strong, by coherence), so we may as well take all maps to be strong inclusions.

With this reduction, the proof becomes trivial: for any $N \prec_{\mathcal{K}} M$ with $|N| < \lambda$ and N' with $N \prec_{\mathcal{K}} N' \prec_{\mathcal{K}} M'$ and $|N'| < \lambda$, if M is λ-model homogeneous relative to M', there is an embedding $h : N' \to M$ that fixes N. But, as noted above, this is equivalent to making the upper left triangle of

$$\begin{array}{ccc} N & \hookrightarrow & N' \\ \big\downarrow & \swarrow_{h} & \big\downarrow \\ M & \hookrightarrow & M' \end{array}$$

commute. Hence the inclusion $M \hookrightarrow_{\mathcal{K}} M'$ is λ-pure. $\square$

Proposition 5.15. *For $\lambda > LS(\mathcal{K})$, a strong inclusion $M \hookrightarrow_{\mathcal{K}} M'$ is λ-pure only if M is λ-Galois-saturated relative to M': every type over $N \prec_{\mathcal{K}} M$ with $|N| < \lambda$ that is realized in M' is realized in M.*

Proof: By Proposition 5.14, λ-purity of $M \hookrightarrow_{\mathcal{K}} M'$ implies that M is λ-model homogeneous relative to M'. Let $N \prec_{\mathcal{K}} M$, $N \in \mathcal{K}_{<\lambda}$, and let p be any type over N that is realized in M', say by a. Take $N' \prec_{\mathcal{K}} M'$ containing $N \cup \{a\}$, $N' \in \mathcal{K}_{<\lambda}$. By relative λ-model homogeneity, there is an embedding $h : N' \hookrightarrow_{\mathcal{K}} M$ that fixes N. Thus any extension of h to an automorphism of $\mathfrak{C}$ lies in $\mathrm{Aut}_N(\mathfrak{C})$, and witnesses that a and $h(a)$ have the same Galois type over N. Since $h(a) \in M$, we are done. $\square$

Thus far we have only characterized λ-purity of inclusions. As a first step in generalizing to arbitrary strong embeddings, we show:

Proposition 5.16. *A strong embedding $f : M \hookrightarrow_{\mathcal{K}} M'$ is λ-pure if and only if the inclusion $f[M] \hookrightarrow_{\mathcal{K}} M'$ is λ-pure.*

Proof: The abbreviated version of the proof is: compositions (pre- or post-) of isomorphisms with λ-pure maps are λ-pure. If we consider suitable factorizations of the morphisms involved, the proposition follows. The longer and more concrete proof amounts to more diagram wrangling.

($\Rightarrow$) Suppose $f : M \hookrightarrow_{\mathcal{K}} M'$ is λ-pure. Consider a commutative square

$$\begin{array}{ccc} N & \xrightarrow{\;g\;} & N' \\ {\scriptstyle u}\downarrow & & \downarrow{\scriptstyle v} \\ f[M] & \hookrightarrow & M' \end{array}$$

with $N, N' \in \mathcal{K}_{<\lambda}$. Then the following diagram also commutes:

$$\begin{array}{ccc} N & \xrightarrow{\;g\;} & N' \\ {\scriptstyle f^{-1}\circ u}\downarrow & & \downarrow{\scriptstyle v} \\ M & \underset{f}{\hookrightarrow} & M' \end{array}$$

By λ-purity of f, there is a map $h' : N' \hookrightarrow_{\mathcal{K}} M$ such that $h' \circ g = f^{-1} \circ u$. Define $h = f \circ h' : N' \hookrightarrow_{\mathcal{K}} f[M]$. Then $h \circ g = u$, and h is the desired diagonal map for the original commutative square.

($\Leftarrow$) Suppose $f[M] \hookrightarrow_{\mathcal{K}} M'$ is λ-pure, and that $N, N' \in \mathcal{K}_{<\lambda}$. If we have a commutative diagram as below on the left, we can transform it into the one on the right

$$\begin{array}{ccc} N & \xrightarrow{\;g\;} & N' \\ {\scriptstyle u}\downarrow & & \downarrow{\scriptstyle v} \\ M & \underset{f}{\longrightarrow} & M' \end{array} \qquad \qquad \begin{array}{ccc} N & \xrightarrow{\;g\;} & N' \\ {\scriptstyle f\circ u}\downarrow & & \downarrow{\scriptstyle v} \\ f[M] & \hookrightarrow & M' \end{array}$$

By λ-purity of $f[M] \hookrightarrow_{\mathcal{K}} M'$, there is a map $h' : N' \hookrightarrow_{\mathcal{K}} f[M]$ such that $h' \circ g = f \circ u$. Set $h = f^{-1} \circ h' : N' \hookrightarrow_{\mathcal{K}} M$. One can see that this map witnesses λ-purity of $f : M \hookrightarrow_{\mathcal{K}} M'$. $\square$

We may now characterize λ-purity for arbitrary strong embeddings.

Corollary 5.17. *For $\lambda > LS(\mathcal{K})$, a strong embedding $f : M \hookrightarrow_{\mathcal{K}} M'$ is λ-pure if and only if $f[M]$ is λ-model homogeneous relative to M'. A strong embedding $f : M \hookrightarrow_{\mathcal{K}} M'$ is λ-pure only if $f[M]$ is λ-Galois-saturated relative to M'.*

Unfortunately, there seems to be no reason to think that the converse of the second statement holds for general M', but as an easy special case (and assuming amalgamation):

Proposition 5.18. *Let $\mathcal{K}$ be an AEC with amalgamation. For $\lambda > LS(\mathcal{K})$, if M' is λ-model homogeneous (or, equivalently, λ-Galois-saturated), a strong embedding $M \hookrightarrow_{\mathcal{K}} M'$ is λ-pure if and only if M is λ-model homogeneous (or, equivalently, λ-Galois-saturated).*

In particular,

Corollary 5.19. *Let $\mathcal{K}$ be an AEC with amalgamation. For $\lambda > LS(\mathcal{K})$, a strong embedding $M \hookrightarrow_{\mathcal{K}} \mathfrak{C}$ is λ-pure if and only if M is λ-model homogeneous (or, equivalently, λ-Galois-saturated).*

On a related note, recall that an object M is said to be λ-closed if every morphism with domain M is λ-pure.

Corollary 5.20. *If $\mathcal{K}$ is an AEC with amalgamation and $\lambda > LS(\mathcal{K})$, $M \in \mathcal{K}$ is λ-closed if and only if M is λ-model homogeneous (λ-Galois-saturated).*

Proof: The "if" direction is clear, by perusal of the diagram defining λ-purity. For the "only if" clause, notice that embeddings of M into $\mathfrak{C}$ are λ-pure, so λ-model homogeneity follows by Corollary 5.18. □

Hence λ-closed and λ-saturated objects coincide in AECs with amalgamation. This will not come as a surprise: as we noted in Section 2.2, this occurs in any category satisfying the amalgamation property.

We turn now to the category-theoretic property most indispensable for our purposes: weak λ-stability. Recall from Definition 2.31 that a category is said to be weakly λ-stable if every morphism $f : M \to M'$ with M a λ^+-presentable object factors as

$$M \longrightarrow \bar{M} \xrightarrow{\;g\;} M'$$

where $\bar{M}$ is λ^+-presentable and g is λ-pure. We saw that in the elementary case, weak λ-stability of $\mathbf{Elem}(T)$ follows from λ-stability of the first order theory T. We have yet to see how weak λ-stability relates to the proliferation of Galois types over models in AECs. First, we notice the following:

Proposition 5.21. *Let $\mathcal{K}$ be an AEC with amalgamation. For any $\lambda \geq LS(\mathcal{K})$, if $\mathcal{K}$ is weakly λ-stable, then any $M \in \mathcal{K}$ of cardinality λ has a saturated extension $\bar{M}$ which is also of cardinality λ.*

Proof: If $M \in \mathcal{K}_\lambda$, it is λ^+-presentable. Hence the embedding $M \hookrightarrow_{\mathcal{K}} \mathfrak{C}$ factors as $M \hookrightarrow_{\mathcal{K}} \bar{M} \hookrightarrow_{\mathcal{K}} \mathfrak{C}$, where $\bar{M}$ is λ^+-presentable (hence $\bar{M} \in \mathcal{K}_{\leq\lambda}$, and since $|\bar{M}| \geq |M| = \lambda$, $\bar{M} \in \mathcal{K}_\lambda$) and the second map in the factorization is λ-pure. By Corollary 5.19, the claim follows. $\square$

Proposition 5.22. *Let $\mathcal{K}$ be an AEC with amalgamation. For any $\lambda \geq LS(\mathcal{K})$, if $\mathcal{K}$ is weakly λ-stable, then for any $M \in \mathcal{K}_{<\lambda}$, $|ga\text{-}S(M)| \leq \lambda$.*

Proof: This follows from the previous proposition and from Lemma 4.29. $\square$

One would hope that λ-Galois stability of an AEC would imply weak λ-stability of the associated category, but there appears to be no reason to think that this is the case. The chief problem is that λ-purity of a strong inclusion $M \hookrightarrow_{\mathcal{K}} M'$ implies λ-Galois saturation of M relative to M', but the converse need not be true. It is possible to give a reasonably model-theoretic condition sufficient to guarantee weak λ-stability (and we will, in Theorem 5.24 below), but it is vaguely unsatisfying and bears no obvious relation to Galois stability. On the other hand, notice that the condition in the consequent of Proposition 5.21—the existence of saturated extensions of size λ—is precisely the condition from which we are able, by Theorem 4.30, to conclude λ-stability in a weakly χ-tame and μ-totally transcendental AEC, provided that $\lambda \geq \chi$

and $\mathrm{cf}(\lambda) > \mu$. That is, weak λ-stability actually implies full λ-Galois stability in this context.

We will return to this point in a moment, but pause to give a sufficient condition for weak λ-stability, as promised.

Lemma 5.23. *Let λ be a regular cardinal with $\lambda > LS(\mathcal{K})$, and assume that for all $N \in \mathcal{K}_{<\lambda}$, N has fewer than λ strong extensions of size less than λ (up to isomorphism over N). Then for any $M \in \mathcal{K}_{\leq\lambda}$ and $M' \in \mathcal{K}$ with $M \prec_{\mathcal{K}} M'$, there is an intermediate extension $\bar{M} \in \mathcal{K}_{\leq\lambda}$ that is λ-model homogeneous relative to M'. That is, $M \hookrightarrow_{\mathcal{K}} M'$ factors as*

$$M \hookrightarrow_{\mathcal{K}} \bar{M} \hookrightarrow_{\mathcal{K}} M'$$

where the latter inclusion is λ-pure.

Proof: Enumerate M as $\{a_i \mid i < |M|\}$. We construct $\bar{M}$ as the union of a continuous $\prec_{\mathcal{K}}$-increasing chain of length λ consisting of models of size less than λ, each of which collects small extensions of the preceding ones. In detail, we proceed as follows:

- $i = 0$: Take $N_0 \prec_{\mathcal{K}} M$, $|N_0| < \lambda$.
- $i = j + 1$: We have $N_j \prec_{\mathcal{K}} M'$, $|N_j| < \lambda$. By our assumption, N_j has fewer than λ many strong extensions in M' of size strictly less than λ, up to isomorphism over N_j. Select a representative from each isomorphism class. The union of N_j and all such representatives is of size strictly less than λ (by regularity). Hence we may take a model $N_i \prec_{\mathcal{K}} M'$ with $|N_i| < \lambda$ that contains the aforementioned union and the element a_j. Notice that $N_j \prec_{\mathcal{K}} N_i$, by coherence.
- i limit: Define $N_i = \bigcup_{j<i} N_j$. Notice that $|N_\alpha| < \lambda$, by regularity of λ.

Let $\bar{M} = \bigcup_{i<\lambda} N_i$. Notice that $\bar{M} \subseteq M'$, whence $\bar{M} \prec_{\mathcal{K}} M'$ by coherence.

In order to see that $\bar{M}$ has the desired property, let $N \prec_{\mathcal{K}} M$, $|N| < \lambda$, and let N' be a model of size less than λ with $N \prec_{\mathcal{K}} N' \prec_{\mathcal{K}} M'$. By regularity of λ, $N \subseteq N_i$ for some $i < \lambda$. Take a model $N'' \prec_{\mathcal{K}} M'$ of size less than λ that contains $N_i \cup N'$. This is a small extension of N_i in M', and so, by construction, there is an N''' collected in the next model in the chain, N_{i+1}, that is isomorphic to N'' over N_i, say via $f : N'' \hookrightarrow_{\mathcal{K}} N'''$. In particular, the isomorphism f fixes $N \prec_{\mathcal{K}} N_i$. Now, $N''' \subseteq N_{i+1} \prec_{\mathcal{K}} \bar{M}$ and, by coherence, $N''' \prec_{\mathcal{K}} \bar{M}$. The composition

$$N' \hookrightarrow_{\mathcal{K}} N'' \overset{f}{\hookrightarrow}_{\mathcal{K}} N''' \hookrightarrow_{\mathcal{K}} \bar{M}$$

is the desired embedding of N' in $\bar{M}$ fixing N. $\square$

Theorem 5.24. *If λ is a regular cardinal and for all $N \in \mathcal{K}_{<\lambda}$, N has fewer than λ strong extensions of size less than λ (up to isomorphism over N), $\mathcal{K}$ is weakly λ-stable.*

Proof: Let M be λ^+-presentable (hence an element of $\mathcal{K}_{\leq\lambda}$), and let $f : M \hookrightarrow_{\mathcal{K}} M'$ be an arbitrary embedding. Recall that f factors as

$$M \overset{f}{\simeq} f[M] \hookrightarrow_{\mathcal{K}} M'$$

Naturally, $|f[M]| \leq \lambda$. By the lemma, there is an intermediate extension $f[M] \prec_{\mathcal{K}} \bar{M} \prec_{\mathcal{K}} M'$ with $|\bar{M}| \leq \lambda$ (hence $\bar{M}$ λ^+-presentable) such that the inclusion $\bar{M} \hookrightarrow_{\mathcal{K}} M'$ is λ-pure. But then

$$M \overset{f}{\hookrightarrow}_{\mathcal{K}} \bar{M} \hookrightarrow_{\mathcal{K}} M'$$

is precisely the kind of factorization of f needed to witness λ-stability. $\square$

5.3 Implications for Galois Stability

We now return to the subject broached after Propositions 5.21 and 5.22: in weakly tame and totally transcendental AECs with amalgamation and joint embedding, for certain cardinals λ, weak λ-stability suffices to ensure λ-Galois stability. What makes this result so interesting is that, thanks to the result of Rosický presented here as Proposition 2.32, we have, for each AEC $\mathcal{K}$, an infinite list of cardinals λ in which it is weakly λ-stable. We work this out in detail below. For reference, though, the result in question was:

Proposition 5.25. *Let* $\mathbf{C}$ *be a* λ*-accessible category, and* μ *a regular cardinal such that* $\mu \trianglerighteq \lambda$ *and* $\mu > |\mathbf{Pres}_\lambda(\mathbf{C})^{mor}|$*. Then* $\mathbf{C}$ *is weakly* $\mu^{<\mu}$*-stable.*

We now analyze the import of this proposition in the context of AECs. To simplify the notation, for any AEC $\mathcal{K}$ and cardinal λ we replace the bulky $\mathbf{Pres}_\lambda(\mathcal{K})$ with $\mathcal{A}_{<\lambda}$; that is, we denote by $\mathcal{A}_{<\lambda}$ a full subcategory of $\mathcal{K}$ consisting of one representative of each isomorphism class of models in $\mathcal{K}_{<\lambda}$.

Corollary 5.26. *Let* $\mathcal{K}$ *be an AEC,* $\lambda > LS(\mathcal{K})$ *a regular cardinal, and* μ *a regular cardinal with* $\mu \trianglerighteq \lambda$ *and* $\mu > |(\mathcal{A}_{<\lambda})^{mor}|$*. Then* $\mathcal{K}$ *is weakly* $\mu^{<\mu}$*-stable.*

Proof: By Theorem 5.1, $\mathcal{K}$ is λ-accessible. The result then follows directly from the proposition above. □

We are now finally in a position to apply Theorem 4.30:

Theorem 5.27. *Let* $\mathcal{K}$ *be weakly* χ*-tame for some* $\chi \geq LS(\mathcal{K})$*, and* κ*-totally transcendental with* $\kappa \geq \chi$*. If* $\lambda > LS(\mathcal{K})$ *is a regular cardinal, and* μ *is a regular cardinal with* $\mu > \chi + \kappa$*,* $\mu \trianglerighteq \lambda$*, and* $\mu > |(\mathcal{A}_{<\lambda})^{mor}|$*, then* $\mathcal{K}$ *is* $\mu^{<\mu}$*-stable.*

Proof: By the assumptions on μ, $\mathcal{K}$ is weakly $\mu^{<\mu}$-stable by Corollary 5.26. We show that the conditions of Theorem 4.30 are satisfied, thereby concluding that $\mathcal{K}$ is not merely weakly $\mu^{<\mu}$-stable, but in fact $\mu^{<\mu}$-Galois stable.

Since μ is regular and $\mu > \kappa$, $\mathrm{cf}(\mu^{<\mu}) \geq \mu > \kappa$. Moreover, $\mu^{<\mu} > \chi$. From Proposition 5.21, we know that every $M \in \mathcal{K}_{(\mu^{<\mu})}$ has a saturated extension $M' \in$

$\mathcal{K}_{(\mu^{<\mu})}$. By the aforementioned theorem, then, we can indeed infer Galois stability in $\mu^{<\mu}$. $\Box$

The trick is that we cannot turn this into an upward stability result in the vein of Theorem 4.26, as the proof of the implication from stability to total transcendence in Section 4.2 requires that we have full tameness in the background—the argument breaks down if we attempt to generalize to weakly tame AECs. And, to be clear, the theorem is interesting only in the context of weakly tame AECs: in the special case in which $\mathcal{K}$ is χ-tame and we are able to make use of the aforementioned result, the proposition becomes:

Corollary 5.28. *Let $\mathcal{K}$ be χ-tame for some $\chi \geq LS(\mathcal{K})$ and κ-Galois stable with $\kappa > \chi$ and $\kappa^{\aleph_0} > \kappa$, and let λ be a regular cardinal, $\lambda > LS(\mathcal{K})$. If μ is a regular cardinal such that $\mu > \chi$, $\lambda \trianglelefteq \mu$ and $|(\mathcal{A}_{<\lambda})^{mor}| < \mu$, then $\mathcal{K}$ is $\mu^{<\mu}$-Galois stable.*

This is not particularly interesting. Consider the cardinals $\mu^{<\mu}$ in which we are guaranteed stability by the corollary above. Since $\mu > \kappa$, it must be the case that $(\mu^{<\mu})^\kappa = \mu^{<\mu}$. But we already have stability in all cardinals satisfying this exponential condition, by Proposition 4.18 (or the more general result in [14]).

Although it is slightly awkward that total transcendence cannot be eliminated as an assumption, we nonetheless have a partial stability spectrum result for weakly tame AECs and, moreover, the only partial spectrum result that is not limited to local transfer of the kind covered in [5] and in Corollary 4.32 above. What is most remarkable, perhaps, is the fact that it was derived by almost entirely category-theoretic means, and the way in which it reveals that the proliferation of types over large structures is controlled by the structure of $(\mathcal{A}_{<\lambda})^{mor}$ (or, if $\mathcal{A}_{<\lambda}$ contains only a single isomorphism class, say with representative C, by the structure of its monoid of endomorphisms, $\mathrm{Hom}_{\mathcal{K}}(C, C)$). This seems, in fact, to be one of the strengths of the accessible category viewpoint: it provides new ways of analyzing classes in terms of their smallest structures and the mappings between them. Indeed, this analysis leads to a structure theorem for AECs (categorical ones, at least), to which we now turn.

5.4 A Structure Theorem For Categorical AECs

We have saved one correspondence between category-theoretic and model-theoretic properties until the end. Recall the notions of λ-categoricity and strong λ-categoricity for arbitrary categories that were introduced in Definition 2.26. Their analogues are precisely what one would expect:

Corollary 5.29. *For any AEC $\mathcal{K}$, λ-categoricity of the corresponding category is equivalent to λ-categoricity in the usual sense. $\mathcal{K}$ is strongly λ-categorical if and only if it contains only a single model of size less than λ^+ (up to isomorphism).*

In [22], Jiri Rosický proves a structure theorem for strongly λ-categorical λ^+-accessible categories, which has, as an interesting special case, a structure theorem for large models in categorical AECs. For the sake of concreteness, we work things out only in this special case, and point interested readers to [22] for the general result (although, in reality, it differs only notationally from our work here).

Let $\mathcal{K}$ be a λ-categorical AEC. Denote by $\mathcal{K}'$ the class $\mathcal{K}_{\geq\lambda}$, with $\prec_{\mathcal{K}'}$ simply the restriction of $\prec_{\mathcal{K}}$. Notice that $\mathcal{K}'$ is still an AEC, albeit with $\mathrm{LS}(\mathcal{K}') = \lambda$. It is also worth noting that $(\mathcal{K}', \prec_{\mathcal{K}'})$ gives rise to precisely the same category as $(\mathcal{K}_{\geq\lambda}, \prec_{\mathcal{K}})$: for any $M, M' \in \mathcal{K}'$ (that is, $\mathcal{K}_{\geq\lambda}$), $\mathrm{Hom}_{\mathcal{K}'}(M, M') = \mathrm{Hom}_{\mathcal{K}}(M, M')$. It is λ^+-accessible (by Theorem 5.1), and strongly λ-categorical in the sense of Definition 2.26. Now, let C be a representative of the unique isomorphism class of models of cardinality λ, and note that $(\mathcal{K}')_\lambda$ is equivalent to the one object category consisting of C and the set of its endomorphisms. We use M to refer both to this one object category and to the corresponding monoid, where the operation on $\mathrm{Hom}_{\mathcal{K}'}(C, C)$ is simply composition: $f \cdot g = f \circ g$. We will show that $\mathcal{K}'$, the class of large models in $\mathcal{K}$, is equivalent to a highly structured subcategory of the category of sets with M-actions.

First, we fix our terminology:

Definition 5.30. Let M be a monoid. An *M-set* is a pair (X, ρ), where X is a set and $\rho : M \times X \to X$ is an action (which we typically write using product notation) satisfying the following conditions for all $a, b \in M$ and $x \in X$:

$$1 \cdot x = x \qquad\qquad (ab) \cdot x = a \cdot (b \cdot x)$$

A map $h : (X_1, \rho_1) \to (X_2, \rho_2)$ is an *M-set homomorphism* if for all $a \in M$ and $x \in X$,

$$h(a \cdot x) = a \cdot h(x)$$

where the actions on the left and right hand sides of the equation are ρ_1 and ρ_2, respectively.

Definition 5.31. Let M be a monoid. We denote by M**-Set** the category of M-sets and M-set homomorphisms. For any regular cardinal λ, we denote by (M, λ)**-Set** the full subcategory of M**-Set** consisting of all λ-directed colimits of copies of M, where the latter is considered as an M-set in the obvious way.

Recall also the notion of equivalence with which we will be working:

Definition 5.32. An *equivalence* between categories $\mathbf{C}$ and $\mathbf{D}$ is given by a pair of functors $F : \mathbf{C} \to \mathbf{D}$ and $G : \mathbf{D} \to \mathbf{C}$ with natural isomorphisms $F \circ G \simeq 1_{\mathbf{D}}$ and $G \circ F \simeq 1_{\mathbf{C}}$. Under these conditions, the functors F and G are referred to as *equivalences of categories*.

Any equivalence of categories $F : \mathbf{C} \to \mathbf{D}$ is full and faithful (that is, bijective on Hom-sets), and is essentially surjective, in the sense that for any object D in $\mathbf{D}$, there is an object C in $\mathbf{C}$ with $F(C) \simeq D$. In short, equivalent categories are structurally identical, as long as we are interested in objects only up to isomorphism.

We now produce the desired equivalence of categories. Recall that for any category $\mathbf{C}$, the category of presheaves on $\mathbf{C}$, denoted $\mathbf{Set}^{\mathbf{C}^{op}}$, consists of all contravariant $\mathbf{Set}$-valued functors on $\mathbf{C}$ and all natural transformations between them. First, we show:

Lemma 5.33. *The AEC $\mathcal{K}'$ is equivalent to the full subcategory of $\mathbf{Set}^{M^{op}}$ consisting of λ^+-directed colimits of $Hom_{\mathcal{K}'}(-, C)$.*

Proof (sketch): We define a functor $F : \mathcal{K}' \to \mathbf{Set}^{M^{op}}$ as follows: for any N in $\mathcal{K}'$,

$$F(N) = \mathrm{Hom}_{\mathcal{K}'}(-, N),$$

the functor that takes C to $F(N)(C) = \mathrm{Hom}_{\mathcal{K}'}(C, N)$ and takes any endomorphism $g : C \hookrightarrow_{\mathcal{K}} C$ to the set map $F(N)(g) : \mathrm{Hom}_{\mathcal{K}'}(C, N) \to \mathrm{Hom}_{\mathcal{K}'}(C, N)$ that sends each $h \in \mathrm{Hom}_{\mathcal{K}'}(C, N)$ to $h \circ g$. The equivalence F takes any any strong embedding $f : N \to N'$ to the map $F(f) : \mathrm{Hom}_{\mathcal{K}'}(-, N) \to \mathrm{Hom}_{\mathcal{K}'}(-, N')$, where $F(f)(g) = f \circ g$ for any $g \in \mathrm{Hom}_{\mathcal{K}'}(C, N)$. Every object in the image of F is (isomorphic to) a λ^+-directed colimit of copies of $\mathrm{Hom}_{\mathcal{K}'}(-, C)$ for the following reason: Any $N \in \mathcal{K}'$ is a λ^+-directed colimit of copies of C, say $N = \mathrm{Colim}_{i \in I}\, C$. By λ^+-presentability of C,

$$\mathrm{Hom}_{\mathcal{K}'}(C, N) = \mathrm{Hom}_{\mathcal{K}'}(C, \mathrm{Colim}_{i \in I}\, C) \simeq \mathrm{Colim}_{i \in I}\, \mathrm{Hom}_{\mathcal{K}'}(C, C)$$

meaning that

$$\mathrm{Hom}_{\mathcal{K}'}(-, N) \simeq \mathrm{Colim}_{i \in I} \mathrm{Hom}_{\mathcal{K}'}(-, C)$$

as functors on the category M, which has C as its only object.

Similar considerations yield the functor G in the other direction, which forms the second part of the equivalence. Any H in the subcategory of $\mathbf{Set}^{M^{op}}$ in which we are interested is a λ^+-directed colimit of copies of $\mathrm{Hom}_{\mathcal{K}'}(-, C)$, say

$$H = \mathrm{Colim}_{i \in I} \mathrm{Hom}_{\mathcal{K}'}(-, C)$$

where the maps in the I-indexed diagram (which lives in $\mathbf{Set}^{M^{op}}$) are natural transformations $\phi_{ij} : \mathrm{Hom}_{\mathcal{K}'}(-, C) \to \mathrm{Hom}_{\mathcal{K}'}(-, C)$ for $i \leq j$ in I. We show that this diagram arises (morphisms and all) from a diagram in $\mathcal{K}'$. By the Yoneda Lemma, the functor F is full, meaning that each ϕ_{ij} is $F(f_{ij})$ for some $f_{ij} : C \to C$. To see that the maps f_{ij} form a valid I-indexed diagram, we must verify that for $i \leq j \leq k$ in I, $f_{jk} \circ f_{ij} = f_{ik}$. To that end, notice that

$$F(f_{jk} \circ f_{ij}) = F(f_{jk}) \circ F(f_{ij}) = \phi_{jk} \circ \phi_{ij} = \phi_{ik} = F(f_{ik}).$$

The Yoneda Lemma also guarantees that F is faithful, meaning that $f_{jk} \circ f_{ij} = f_{ik}$, as desired.

By λ^+-presentability of C, again,

$$\mathrm{Colim}_{i \in I} \mathrm{Hom}_{\mathcal{K}'}(C, C) \simeq \mathrm{Hom}_{\mathcal{K}'}(C, \mathrm{Colim}_{i \in I} C)$$

where the latter colimit is that of the diagram in $\mathcal{K}'$ mentioned above. Since $\mathcal{K}'$ is closed under λ^+-directed colimits, $N = \mathrm{Colim}_{i \in I} C$ is in $\mathcal{K}'$, and we define $G(H) = N$. The proof that the compositions of F and G are naturally isomorphic to the identity functors on $\mathbf{C}$ and $\mathbf{D}$—which amounts to no more than the checking of details—is left as an exercise. $\square$

As an aside, for absolutely any AEC $\mathcal{K}$ and regular cardinal $\lambda > \mathrm{LS}(\mathcal{K})$, $\mathcal{K}$ is equivalent to the category of presheaves on $\mathcal{A}_{<\lambda}$ that are λ-directed colimits of representable functors (that is, λ-directed colimits of functors of the form $\mathrm{Hom}_{\mathcal{K}}(-, N)$, where N is an object of $\mathcal{A}_{<\lambda}$). The categoricity assumption under which we are currently operating merely guarantees that $\mathcal{A}_{<\lambda^+}$ is a monoid, allowing us to conclude the following:

Theorem 5.34. *Under the hypothesis above, the AEC $\mathcal{K}'$, regarded as a category in the usual way, is equivalent to the category (M^{op}, λ^+)-**Set**.*

Proof: We first note an equivalence of categories between $\mathbf{Set}^{M^{op}}$ and M^{op}-**Set** (which holds for any monoid M), in which any functor $H : M^{op} \to \mathbf{Set}$ is sent to the M^{op}-set $(H(C), \rho_H)$, where for any $a \in M$ and $x \in H(C)$, the action is given by $a \cdot x = H(a)(x)$. Explicitly, we have already shown that $\mathcal{K}'$ is equivalent to the full subcategory of $\mathbf{Set}^{M^{op}}$ consisting of all λ^+-directed colimits of $\mathrm{Hom}_{\mathcal{K}'}(-, C)$. Under the equivalence at hand, $\mathrm{Hom}_{\mathcal{K}'}(-, C)$ maps to the set $\mathrm{Hom}_{\mathcal{K}'}(C, C)$ (roughly speaking: M^{op}), with M^{op} acting by precomposition, whereas for arbitrary $N \in \mathcal{K}'$, $\mathrm{Hom}_{\mathcal{K}'}(-, N)$ maps to the set $\mathrm{Hom}_{\mathcal{K}'}(C, N)$, again with action of $\mathrm{Hom}_{\mathcal{K}'}(C, C)$ by precomposition. One can easily see that the image is precisely the full subcategory consisting of λ^+-directed colimits of M^{op} considered as an M^{op}-**Set** in the way indicated above. □

To emphasize, the equivalence between $\mathcal{K}'$ and (M^{op}, λ^+)-**Set** is given by:

$$N \in \mathcal{K}' \quad \mapsto \quad (\mathrm{Hom}_{\mathcal{K}'}(C, N), \rho_N)$$

where the action ρ_N is given, for any $a \in \mathrm{Hom}_{\mathcal{K}'}(C, C)$ and $x \in \mathrm{Hom}_{\mathcal{K}'}(C, N)$, by

$$a \cdot x = x \circ a$$

A strong embedding $f : N \hookrightarrow_{\mathcal{K}} N'$ is mapped to the M^{op}-set homomorphism $f^* : \mathrm{Hom}_{\mathcal{K}'}(C, N) \to \mathrm{Hom}_{\mathcal{K}'}(C, N')$ that takes any $g \in \mathrm{Hom}_{\mathcal{K}'}(C, N)$ to $f \circ g$. That the map f^* thus defined is in fact a homomorphism of M^{op}-sets is easily checked: for any $h \in M^{op} = \mathrm{Hom}_{\mathcal{K}'}(C, C)$,

$$f^*(h \cdot g) = f^*(g \circ h) = f \circ (g \circ h) = (f \circ g) \circ h = h \cdot (f \circ g) = h \cdot f^*(g).$$

The upshot is this: for any λ-categorical AEC, we may identify $\mathcal{K}' = \mathcal{K}_{\geq\lambda}$ with a category of relatively simple algebraic objects, representing each model by a set equipped with an action of $M^{op} = \mathrm{Hom}_{\mathcal{K}'}(C, C)$, the monoid of endomorphisms of the unique structure in cardinality λ, and replacing the abstract embeddings of $\mathcal{K}$ with concrete homomorphisms between such sets. This gives us a radically different—and perhaps simpler—context in which to consider questions originally posed in relation to AECs. The conjectures concerned with the upward transfer of categoricity, in particular, involve an analysis of the sub-AEC consisting of the structures whose cardinalities are greater than or equal to the cardinal at which categoricity first occurs; that is, a suitable $\mathcal{K}'$ of the form described above. The hope is that the translation described above provides a simplification not merely in appearance, but in the sense of providing genuine traction in addressing such problems. Given that we have reduced something as complex and general as an AEC to a category whose properties are determined entirely by the structure of the monoid $\mathrm{Hom}_{\mathcal{K}'}(C, C)$ (which is just $\mathrm{Hom}_{\mathcal{K}}(C, C)$, remember), there is some reason to believe that this will be the case.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] Jiří Adámek and Jiří Rosický. *Locally presentable and accessible categories.* Number 189 in London Mathematical Society Lecture Notes. Cambridge, New York, 1994.

[2] John Baldwin. Categoricity in abstract elementary classes. Available at http://www.math.uic.edu/~jbaldwin/pub/AEClec.pdf, August 2008.

[3] John Baldwin, Paul Eklof, and Jan Trlifaj. $N^{\perp}$ as an AEC. Available at http://www.math.uic.edu/~jbaldwin/model.html, June 2007.

[4] John Baldwin and Alexei Kolesnikov. Categoricity, amalgamation and tameness. To appear in: Israel Journal of Math, November 2007.

[5] John Baldwin, David Kueker, and Monica VanDieren. Upward stability transfer for tame abstract elementary classes. *Notre Dame Journal of Formal Logic*, 47(2):291–298, 2006.

[6] John Baldwin and Saharon Shelah. Examples of non-locality. To appear in Journal of Symbolic Logic. Available at http://www.math.uic.edu/~jbaldwin/model.html, November 2007.

[7] Steven Buechler. *Essential stability theory.* Perspectives in Mathematical Logic. Springer-Verlag, New York, 1996.

[8] Peter Gabriel and Friedrich Ulmer. *Lokal präsentierbare kategorien.* Number 221 in Lecture Notes in Mathematics. Springer-Verlag, Berlin, 1971.

[9] Leonard Gillman and Meyer Jerison. *Rings of continuous functions.* University Series in Higher Mathematics. Van Nostrand, Princeton, 1960.

[10] Rami Grossberg. Classification theory for abstract elementary classes. In Yi Zhang, editor, *Logic and Algebra*, volume 302 of *Contemporary Mathematics*, pages 165–204. AMS, 2002.

[11] Rami Grossberg and Alexei Kolesnikov. Superior abstract elementary classes are tame. Submitted. Available at http://www.math.cmu.edu/~rami/, September 2005.

[12] Rami Grossberg and Saharon Shelah. On universal locally finite groups. *Israel Journal of Mathematics*, 44(4):289–302, December 1983.

[13] Rami Grossberg and Monica VanDieren. Categoricity from one successor cardinal in tame abstract elementary classes. *Journal of Mathematical Logic*, 6(2):181–201, 2006.

[14] Rami Grossberg and Monica VanDieren. Galois-stability in tame abstract elementary classes. *Journal of Mathematical Logic*, 6(1):25–49, 2006.

[15] Rami Grossberg and Monica VanDieren. Shelah's categoricity conjecture from a successor for tame abstract elementary classes. *Journal of Symbolic Logic*, 71(2):553–568, 2006.

[16] Adi Jarden and Saharon Shelah. Good frames minus stability. Available at http://shelah.logic.at/files/875.pdf, September 2008.

[17] John L. Kelley. *General topology.* Number 27 in Graduate Texts in Mathematics. Springer-Verlag, New York, 2nd edition, 1975.

[18] Jonathan Kirby. Abstract elementary categories. Available at `http://people.maths.ox.ac.uk/~kirby/pdf/aecats.pdf`, August 2008.

[19] Angus Macintyre. Twenty years of p-adic model theory. In *Logic Colloquium '84*, number 120 in Studies in Logic and the Foundations of Mathematics, pages 121–153, Amsterdam, 1986. North-Holland.

[20] Michael Makkai and Robert Paré. *Accessible categories: the foundations of categorical model theory*, volume 104 of *Contemporary Mathematics*. AMS, Providence, RI, 1989.

[21] Michael Morley. Categoricity in power. *Transactions of the AMS*, 114:514–538, 1965.

[22] Jiří Rosický. Accessible categories, saturation and categoricity. *Journal of Symbolic Logic*, 62:891–901, 1997.

[23] Saharon Shelah. Categoricity in abstract elementary classes: going up inductive step. Preprint 600. Available at http://shelah.logic.at/short600.html.

[24] Saharon Shelah. *Classification theory and the number of nonisomorphic models*, volume 92 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, New York, 1st edition, 1978.

[25] Saharon Shelah. Classification theory for nonelementary classes, II. In *Classification Theory*, volume 1292 of *Lecture Notes in Mathematics*, pages 419–497. Springer-Verlag, Berlin, 1987.

[26] Saharon Shelah. Categoricity for abstract classes with amalgamation. *Annals of Pure and Applied Logic*, 98:261–294, 1999.

[27] Boris Zilber. Pseudo-exponentiation on algebraically closed fields of characteristic zero. *Annals of Pure and Applied Logic*, 132(1):67–95, 2004.