ACCESSIBLE CATEGORIES AND ABSTRACT ELEMENTARY CLASSES

MICHAEL J. LIEBERMAN*

April 28, 2010

Abstract

We highlight connections between two parallel areas of research: accessible categories and abstract elementary classes (AECs). In particular, we show that AECs are accessible categories, and provide a dictionary for translating properties and results between the two contexts. The category-theoretic viewpoint yields two surprising results: a partial stability spectrum for weakly tame AECs, and a structure theorem for categorical AECs under which the large models in such classes are represented as sets with monoid actions.

1 Introduction

There has been a recent surge of interest in the model theory of nonelementary classes, spurred by the appearance of a number of important applications in mainstream mathematics, including recent work on Banach spaces and the complex numbers with exponentiation. Two well-established frameworks for the analysis of such classes recommend themselves on the basis of the balance they strike between generality and richness of structure: abstract elementary classes and accessible categories. Although these notions were generated in the course of independent lines of investigation—in model theory and categorical logic, respectively—they exhibit striking similarities. Abstract elementary classes, on the one hand, were introduced by Shelah as a broad framework in which to carry out the project of classification theory for a wide array of nonelementary classes. In contrast with earlier work on, say, the model theory of $L_{\omega_1,\omega}$, where the methods were closely tied to the structure of the ambient logic and types associated with satisfiable sets of formulas, in AECs one dispenses with syntax, retaining only the essential, fundamentally category-theoretic structure carried by the strong embeddings. Accessible categories, on the other hand, may be regarded as an outgrowth of categorical logic, the program in which logical theories are associated with categories that capture their essential structure (elementary topoi, say) and classical models are identified with structure-preserving **Set**-valued functors on the associated categories. In [1], [4], and [13] one sees, in parallel with the story for AECs, a distinct shift in emphasis away from the category associated with a theory, and a focus on the abstract properties of the category of models in itself—in this way, one arrives at the notion of an accessible category.

The goal of the present inquiry (alongside independent work of Beke and Rosický) is to begin to fill in the details of the connection between AECs and accessible categories and to illustrate a few ways in which results from the world of accessible categories can be translated into novel results for AECs. As this paper represents an attempts at a rapprochement between model-theoretic

*The contents of this paper are drawn from the author's doctoral dissertation, completed at the University of Michigan under the supervision of Andreas Blass. The author also acknowledges useful conversations with John Baldwin.

and category-theoretic perspectives, we have endeavored to provide enough background detail to accomodate readers whose experience tends to place them squarely on one side or the other of the divide. In particular, Section 2 includes a few model-theoretic preliminaries, and outlines the state of play in the study of AECs—readers familiar with this material may wish to skim, focusing solely on a handful of recent results presented at the end (particularly Theorem 2.8). Section 3, in turn, comprises a brief introduction to accessible categories, which specialists should feel free to breeze through, if not skip entirely. Section 4 begins the process of reconciliation, realizing AECs as highly structured accessible categories, and includes a complete category-theoretic axiomatization of AECs in a finitary signature L as subcategories of the ambient category of L-structures. Section 5 translates a number of of notions from accessible categories (most of which are drawn from [14]) into the context of AECs. This exercise bears immediate fruit in Sections 6 and 7, as simple category-theoretic manipulations yield a pair of novel results: respectively, a structure theorem for categorical AECs, and a more technical partial stability spectrum result for weakly tame AECs satisfying the notion of total transcendence outlined in [11].

2 Abstract Elementary Classes

We begin with a very brief introduction to AECs, Galois types, and a few relevant properties thereof. Readers interested in further details may wish to consult [2] or [5]. To begin:

Definition 2.1. Let L be a 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 both $\mathcal{K}$ and $\prec_{\mathcal{K}}$ are closed under isomorphism, and satisfy 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 (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 $\bigcup_{\alpha<\delta}M_\alpha\prec_{\mathcal{K}}M$.

A3 (Coherence) If $M_0,M_1\prec_{\mathcal{K}}M$ in $\mathcal{K}$, and $M_0\subseteq_L M_1$, then $M_0\prec_{\mathcal{K}}M_1$.

A4 (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, in which case $\mathrm{LS}(\mathcal{K})$ is, naturally, $\aleph_0+|L(T)|$.

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 $f:M\hookrightarrow_{\mathcal{K}}N$.

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

1. We say that an AEC $\mathcal{K}$ 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 AEC $\mathcal{K}$ 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 $g_1 \circ f_1 = g_2 \circ f_2$.

Notice that both properties hold in elementary classes (provided the associated first-order theory is complete), 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. We will state clearly when AP and JEP are assumed—this will not occur wholesale until Section 7.

It is not immediately clear what we might embrace as a suitable notion of type in AECs, given that 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 [16]. Although Galois types can be defined in very general AECs, they take a particularly simple form in those with amalgamation and joint embedding. In any AEC $\mathcal{K}$ of this form, we may fix a monster model $\mathfrak{C} \in \mathcal{K}$, and consider all models $M \in \mathcal{K}$ as strong submodels of $\mathfrak{C}$. In this case, Galois types have the following characterization:

Definition 2.3. Let $M \in \mathcal{K}$, and $a \in \mathfrak{C}$. The *Galois type of a over M*, denoted $\text{ga-tp}(a/M)$, is the orbit of a in $\mathfrak{C}$ under $\text{Aut}_M(\mathfrak{C})$, the group of automorphisms of $\mathfrak{C}$ that fix M. We denote by $\text{ga-S}(M)$ the set of all Galois types over M.

In case $\mathcal{K}$ is an elementary class with $\prec_{\mathcal{K}}$ as elementary submodel, the Galois types over M correspond to the complete first-order types over M:

$$\text{ga-tp}(a/M) = \text{ga-tp}(b/M) \text{ if and only if } \text{tp}(a/M) = \text{tp}(b/M)$$

In general, however, Galois types and syntactic types do not match up, even in cases when the logic underlying the AEC is clear (say, $\mathcal{K} = \text{Mod}(\psi)$, with $\psi \in L_{\omega_1,\omega}$). A few basic definitions and notations:

Definition 2.4. 1. We say that $\mathcal{K}$ is *λ-Galois stable* if for every $M \in \mathcal{K}_\lambda$, $|\text{ga-S}(M)| = \lambda$.

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

3. 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 $\text{ga-tp}(a/M) \restriction N = p$. Equivalently, M realizes p if the orbit in $\mathfrak{C}$ corresponding to p meets M.

4. 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. It bears mentioning that Galois saturation is closely related to a notion of homogeneity peculiar to AECs:

Definition 2.5. 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 will have occasion to use the following fact (Theorem 8.14 in [2]):

Proposition 2.6. *Let $\mathcal{K}$ be an AEC satisfying AP and JEP. For $\lambda > LS(\mathcal{K})$, a model $M \in \mathcal{K}$ is λ-Galois-saturated if and only if it is λ-model homogeneous.*

Initial attempts at establishing a classification theory for AECs have focused on classes satisfying a variety of broad structural conditions: for example, excellence (as described in [6]) and the existence of good or semi-good frames (as considered in [15] and [9]), which echo classical notions of simplicity, stability, and superstability, respectively. We here concern ourselves primarily with AECs satisfying the property known as tameness, which says, roughly speaking, that types are determined by their restrictions to small submodels of their domains, a condition reminiscent of the locality properties of syntactic types.

We present two measures of tameness:

Definition 2.7. We say that $\mathcal{K}$ is *χ-tame* if for every $M \in \mathcal{K}$, if p and p' are distinct types over M, then there is an $N \in \mathrm{Sub}_{\leq\chi}(M)$, such that $p \restriction N \neq p' \restriction N$. We say that $\mathcal{K}$ is *weakly χ-tame* if the condition above holds for saturated $M \in \mathcal{K}$.

A certain amount of tameness is necessary for essentially all existing results on the classification theory of AECs, which field remains, it must be said, a work very much still in progress. In regard to questions of eventual categoricity, results are typically measured against Shelah's Categoricity Conjecture: if an AEC $\mathcal{K}$ is categorical in one cardinal $\mu \geq \mathrm{Hanf}(\mathcal{K})$, the Hanf number of the class (see [2] for a detailed treatment of this notion), $\mathcal{K}$ is categorical in every $\kappa \geq \mathrm{Hanf}(\mathcal{K})$. Approximations of this result hold in tame AECs, the most promising of which, a result of [8], implies categoricity in every cardinal $\kappa \geq \mathrm{H}_2$, the second Hanf number—related to, but substantially larger than, $\mathrm{Hanf}(\mathcal{K})$—given categoricity in a single successor cardinal $\mu^+ \geq \mathrm{H}_2$. It is important to note that the computation of the Hanf number requires the reintroduction of syntax via Shelah's Presentation Theorem, and the proof the the eventual categoricity result mentioned above depends on syntax and a resort to a classical, and decidedly ensembliste, toolkit: indiscernibles, EM-models, and so on. The natural question: if we retain our category-theoretic perspective on AECs, is it still possible to prove interesting theorems about categorical AECs? Section 6 represents a partial answer (in the affirmative) and suggests that there are results which, although readily apparent from the aforementioned perspective, would be otherwise unobtainable.

Stability spectrum results are even patchier. For tame AECs, Grossberg and VanDieren have proven in [7], using splitting and the techniques mentioned above, that stability in a cardinal λ implies stability in any κ such that $\kappa^\lambda = \kappa$. In [3], the machinery of splitting in involved again, this time to prove stability transfer from a cardinal λ to λ^+, a result that carries over to the weakly tame context. In [11], the author introduces a few new tools for the analysis of Galois stability in AECs: a family of Morley-like ranks RM^λ, indexed by cardinals $\lambda \geq \mathrm{LS}(\mathcal{K})$, and related notions of λ-total transcendence (where an AEC $\mathcal{K}$ is λ-totally transcendental if RM^λ is ordinal-valued on all types associated with the class). Since λ-total transcendence allows one to bound the number of types over structures in $\mathcal{K}_{>\lambda}$ and since, at least in tame AECs, λ-total transcendence follows from λ-stability provided that $\lambda^{\aleph_0} > \lambda$, this allows one to prove a number of upward stability transfer theorems. In particular, one can prove a generalization of another result of [3] which implies, for example, that for any tame $\aleph_0$-stable AEC $\mathcal{K}$, if $\mathcal{K}$ is stable in a sequence of cardinals cofinal in a cardinal κ, $cf(\kappa) > \aleph_0$, then it is κ-stable as well.

For weakly tame AECs, there is nothing more than the transfer of stability from a cardinal to its successor, mentioned above, and the following result of [11]:

Theorem 2.8. *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.*

It is remarkable that the existence of saturated extensions of the sort required in the proposition above can be guaranteed by the purely category-theoretic condition known as weak λ-stability. It is a still more remarkable fact that weak stability occurs in many cardinalities in any accessible category hence also, as we will see, in any AEC. The result of this reasoning, a partial stability transfer result for weakly tame AECs, is examined in Section 7.

3 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. Moreover, the Downward Löwenheim-Skolem Property for AECs guarantees that any structure $M \in \mathcal{K}$ can be obtained as the directed union of its submodels of cardinality at most $\mathrm{LS}(\mathcal{K})$, meaning that an AEC $\mathcal{K}$ is, in fact, generated from the set of all such small models, $\mathcal{K}_{\mathrm{LS}(\mathcal{K})}$. Although accessible categories—the category theorists' preferred generalization of classes of structures, both elementary and nonelementary (see [1] and [13])—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 [4], and has subsequently been treated in a more accessible fashion in [1] and [13], the latter being a particularly good source of concrete examples. We begin with the simplest and most mathematically natural case:

Definition 3.1. 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 $D : (I, \leq) \to \mathbf{C}$ 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.

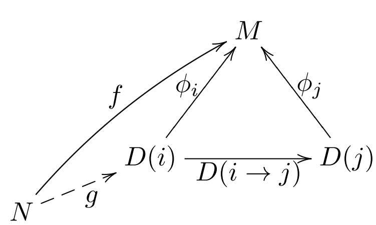

Moreover, this factorization must be essentially unique, in the sense that for any two such, say g and g' from N to $D(i)$ with $f = \phi_i \circ g = \phi_i \circ g'$, there is a $j \geq i$ in I such that $D(i \to j) \circ g = D(i \to j) \circ g'$.

Examples:

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

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

3. 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.
   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 3.2. A category $\mathbf{C}$ is *finitely accessible* if

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

- $\mathbf{C}$ is closed under directed colimits.

Finitely accessible categories abound in mainstream mathematics: the category **Grp** or, indeed, any category of finitary algebraic varieties; $\mathbf{Rel}(\mathbf{\Sigma})$, under the conditions described above; **Set**, the category of sets; and **Pos**, the category of posets and monotone functions.

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

Definition 3.3. 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 3.4. 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 3.1 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 3.5. 1. Let λ be an infinite regular cardinal. A category $\mathbf{C}$ is λ-*accessible* if

   - $\mathbf{C}$ is closed under λ-directed colimits

- **C** contains only a set of λ-presentable objects up to isomorphism, and every object in **C** is a λ-directed colimit of λ-presentables.

2. We say that a category **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. Invoking Theorem 2.11 in [1]:

Theorem 3.6. *For regular cardinals $\lambda < \mu$, the following are equivalent:*

1. *Each λ-accessible category is μ-accessible.*
2. *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 3.7. For regular cardinals λ and μ, we say that λ is sharply less than μ, denoted $\lambda \lhd \mu$, if they satisfy the equivalent conditions of the theorem above.

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

1. $\omega \lhd \mu$ 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 $\lambda \in L$. This will play an important role in the partial stability spectrum result in Section 7.

4 AECs as Accessible Categories

Given an AEC $\mathcal{K}$, 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 $\mathcal{K}$. The first step in our analysis of the connections between AECs and accessible categories involves proving the following:

Theorem 4.1. *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 4.2. *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 4.3. *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

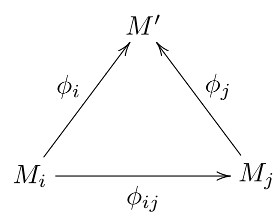

Consider an embedding $f : M \hookrightarrow_{\mathcal{K}} M'$. The image $f[M]$ is a strong submodel of M, and is of cardinality $\nu < \mu$. Since $\mathcal{K}$ is a concrete category, the submodels $\phi_i[M_i]$ of M' cover M', meaning that for each $m \in f[M]$ we may choose a $\phi_{i_m}[M_{i_m}]$ 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 4.4. *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 above:

Lemma 4.5. *$\mathcal{K}$ is closed under directed colimits.*

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

Lemma 4.6. *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 4.1 above with the following result of [14]:

Proposition 4.7. *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 such 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 [10], 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 capturing 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 [10] under the heading of accessibility, thereby making clear the connection between AECs and the existing body of work on accessible categories. This perspective clarifies, for example, that the abstract notion of size laid out in the aforementioned piece corresponds to the well-established notion of presentability.

We introduce two definitions, the second of which is drawn from [10]:

Definition 4.8. 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

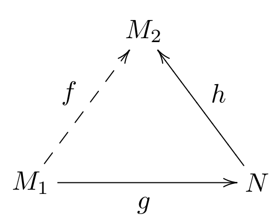

with h and g (hence also their domains and codomains) in $\mathbf{C}$ and with f in $\mathbf{B}$, then in fact f is in $\mathbf{C}$.

Purely from Theorem 4.1 and the axioms for AECs,

Proposition 4.9. *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 4.10. *The class* $\mathbf{C}$ *is an AEC.*

Proof: As in the proof of Lemma 4 in [10], 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 4.3 above, of cardinality at most $\mu = |A| + \lambda$. One can see, then, that λ will do as LS($\mathbf{C}$). □

The amalgamation and joint embedding properties for AECs are purely diagrammatic, and coincide exactly with the analogues for accessible categories included in [14]. If we add them to the axioms in Proposition 4.9, we obtain an axiomatization of AECs with the AP and JEP. On the other hand, if we replace $L-$**Struct** with a particular category of metric L-structures (as in [10]), our axioms describe the abstract metric classes in the signature L.

5 Model Theory and Category Theory: Correspondences

We turn 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 originally defined in [14]. The latter piece is, of course, concerned with accessible categories with directed colimits—almost AECs, as we now know.

We begin with the easiest correspondence. Accompanying our notion of size for objects in accessible categories—presentability—is a natural notion of categoricity:

Definition 5.1. A category $\mathbf{C}$ is λ-*categorical* if it contains, up to isomorphism, a unique object N which is λ^+-presentable, but not μ-presentable for any $\mu < \lambda^+$. $\mathbf{C}$ is said to be *strongly λ-categorical* if it contains, up to isomorphism, a unique λ^+-presentable object.

From Lemma 4.3, we have:

Corollary 5.2. *For an 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).*

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 5.3. 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 [14]:

Definition 5.4. Let λ be a regular cardinal. An object M in a category $\mathbf{C}$ 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:

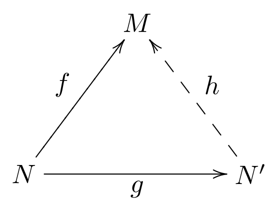

This looks more like a homogeneity condition, although it matches up nicely with the classical notion of λ-saturation in elementary classes. For AECs, we have:

Proposition 5.5. *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 4.3, 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:

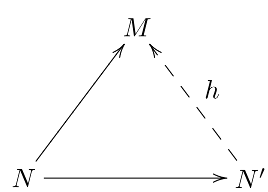

with $N \hookrightarrow_{\mathcal{K}} N'$ and $N \hookrightarrow_{\mathcal{K}} M$ the inclusions. This says precisely that h is an embedding of N' into M fixing N. So M is λ-model homogeneous.
($\Leftarrow$) Using one application of each of the facts in Remark 5.3, one can see that it suffices to consider diagrams of strong inclusions

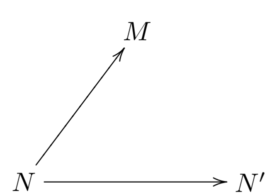

with N and N' both λ-presentable. Then $N, N' \in \mathcal{K}_{<\lambda}$ and λ-model homogeneity of M guarantees the existence of a strong embedding $h : N' \hookrightarrow_{\mathcal{K}} M$ fixing N, and therefore making the relevant diagram commute. Hence M is λ-saturated. $\square$

Recalling that λ-model homogeneity implies λ-Galois-saturation in any AEC, and that the converse holds in AECs with amalgamation (see Proposition 2.6), we get

Corollary 5.6. *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.*

Definition 5.7. 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

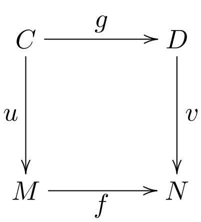

in which C and D are λ-presentable, there is a morphism $h : D \to M$ such that $h \circ g = u$.

In elementary classes, one can show 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 5.8. *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

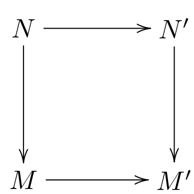

By λ-purity of the bottom inclusion, there is a strong embedding $h : N' \hookrightarrow_{\mathcal{K}} M$ that makes the upper triangle of

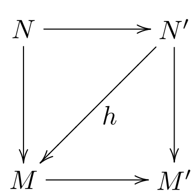

commute. This commutativity condition simply means that $h : N' \hookrightarrow_{\mathcal{K}} M$ fixes N, so we are done.
($\Leftarrow$) By Remark 5.3 and a little diagram-wrangling, it suffices to consider diagrams of the form

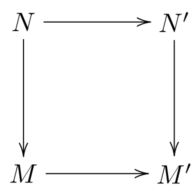

with the bottom and vertical morphisms strong inclusions. In that case, of course, the upper map must be an inclusion as well (and strong, by coherence), so we may as well take all maps to be strong inclusions.

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

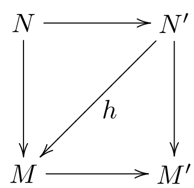

commute. Hence the inclusion $M \hookrightarrow_{\mathcal{K}} M'$ is λ-pure. $\square$

Proposition 5.9. *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 5.8, 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 considered λ-purity of inclusions. As a first step in generalizing to arbitrary strong embeddings, the fact that compositions (pre- or post-) of isomorphisms with λ-pure maps are λ-pure implies:

Proposition 5.10. *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.*

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

Corollary 5.11. *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'.*

The converse of the second statement holds if $\mathcal{K}$ has the AP. If $\mathcal{K}$ has the AP and the codomain is known to be λ-model homogeneous (or, equivalently, λ-Galois-saturated), we have:

Proposition 5.12. *Let $\mathcal{K}$ be an AEC with amalgamation. For $\lambda > LS(\mathcal{K})$, if M' is λ-model homogeneous (λ-Galois-saturated), a strong embedding $M \hookrightarrow_{\mathcal{K}} M'$ is λ-pure if and only if M is λ-model homogeneous (λ-Galois-saturated).*

In particular,

Corollary 5.13. *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).*

We turn now to the category-theoretic property most indispensable for our purposes: weak λ-stability.

Definition 5.14. Let λ be a regular cardinal. A category $\mathbf{C}$ is said to be *weakly λ-stable* if for any λ^+-presentable M and morphism $f : M \to M'$, f factors as $M \xrightarrow{g} \bar{M} \xrightarrow{h} M'$, where $\bar{M}$ is λ^+-presentable, and h is λ-pure.

One can show that in the elementary case, weak λ-stability of $\mathbf{Elem}(T)$ follows from λ-stability of the first order theory T (see the discussion following Definition 2.31 in [12]). One would hope that λ-Galois stability of an AEC would imply weak λ-stability of the associated category, but it is not clear that this is the case. It is still possible to give a reasonably model-theoretic condition sufficient to guarantee weak λ-stability, however. Leaving the details to [12], the result is the following:

Theorem 5.15. *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.*

More interesting for our purposes is the converse, roughly speaking: the way in which weak λ-stability controls the proliferation of Galois types over models in AECs.

Proposition 5.16. *Let $\mathcal{K}$ be an AEC with amalgamation. For any $\lambda \geq LS(\mathcal{K})$, if $\mathcal{K}$ is weakly λ-stable, then any $M \in \mathcal{K}_\lambda$ has a saturated extension $\bar{M} \in \mathcal{K}_\lambda$.*

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 (meaning that $\bar{M} \in \mathcal{K}_{\leq\lambda}$, 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 5.13, the claim follows. $\Box$

The condition in the consequent of Proposition 5.16—the existence of saturated extensions of size λ—is precisely the condition from which we are able, by Theorem 2.8, to conclude λ-stability in a weakly χ-tame and μ-totally transcendental AEC with the AP and JEP, provided $\lambda \geq \chi$ and $\mathrm{cf}(\lambda) > \mu$. That is, weak λ-stability actually implies full λ-Galois stability in this context. This fact lies at the heart of the spectrum result in Section 7.

6 A Structure Theorem For Categorical AECs

In [14], Jiří 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. We work things out only in the special case, and point interested readers to [14] for the general result (although, in fact, 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 4.1), and strongly λ-categorical in the sense of Definition 5.1. 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 6.1. 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 6.2. 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 6.3. 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 6.4. *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: 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 are natural transformations $\phi_{ij} : \mathrm{Hom}_{\mathcal{K}'}(-, C) \to \mathrm{Hom}_{\mathcal{K}'}(-, C)$ for $i \leq j$ in I. By the Yoneda Lemma, the functor F is full and faithful, meaning that that this diagram arises (morphisms and all) from an I-indexed diagram in $\mathcal{K}'$.

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 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 6.5. *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 $\mathrm{Hom}_{\mathcal{K}'}(C, C)$ (that is, 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 M^{op} acting 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** as above. $\square$

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 verified.

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 a radically different context in which to consider questions originally posed in relation to AECs. 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. 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 hope that this translation provides a simplification not merely in appearance, but in the sense of providing genuine traction in addressing such problems.

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.

7 Implications for Galois Stability

We now return to the subject broached after Proposition 5.16: 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 a result of [14], 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 in due course. For reference, though, the result in question is:

Proposition 7.1. *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 7.2. *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 4.1, $\mathcal{K}$ is λ-accessible. The result then follows directly from the proposition above. □

We are now finally in a position to apply Theorem 2.8:

Theorem 7.3. *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 7.2. We show that the conditions of Theorem 2.8 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 5.16, 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 $\mu^{<\mu}$. $\Box$

If $\mathcal{K}_{<\lambda}$ contains only a single isomorphism class, say with representative C, $(\mathcal{A}_{<\lambda})^{mor}$ is simply $\mathrm{Hom}_{\mathcal{K}}(C, C)$. This leads to a clearer picture in the following special case:

Proposition 7.4. *If $\mathcal{K}$ is weakly $\aleph_0$-tame, $\aleph_0$-categorical, and $\aleph_0$-t.t., then for any regular μ with $\aleph_1 \lhd \mu$ and $\mu > |Hom_{\mathcal{K}}(C, C)|$, $\mathcal{K}$ is $\mu^{<\mu}$-stable.*

- Worst case: $|\mathrm{Hom}_{\mathcal{K}}(C, C)| = 2^{\aleph_0}$. We have $(2^{\aleph_1})^+ > |\mathrm{Hom}_{\mathcal{K}}(C, C)|$ and sharply greater than $\aleph_1$, hence also stability in $[(2^{\aleph_1})^+]^{2^{\aleph_1}}$. Similarly, we have stability in $[(2^{\aleph_1})^{+(n+1)}]^{(2^{\aleph_1})^{+n}}$ for $n < \omega$. We also have stability in $[(2^{\aleph_k})^{+(n+1)}]^{(2^{\aleph_k})^{+n}}$ for $1 \leq k \leq \omega$ and $n < \omega$, among other cardinals. Under GCH, this gives stability in all κ with $\aleph_{k+2} \leq \kappa < \aleph_\omega$.

- Better: $|\mathrm{Hom}_{\mathcal{K}}(C, C)| = \aleph_k$ with $0 \leq k < \omega$. Then we have stability in $\aleph_{k+1}^{\aleph_k}$, $\aleph_{k+2}^{\aleph_{k+1}}$, $\aleph_{k+3}^{\aleph_{k+2}}$, and, more generally, $\aleph_{n+1}^{\aleph_n}$ for $k \leq n < \omega$. Naturally, we also have stability in the cardinals listed in the worst case scenario above. Under GCH, this gives stability in all κ with $\aleph_{k+1} \leq \kappa < \aleph_\omega$.

One would hope that total transcendence could be replaced by a more straightforward assumption of stability, thereby transforming the above result into a pure upward transfer theorem like those of [3], [7], and [11]. Unfortunately, the proof of the inference from stability to total transcendence hinges on full tameness of the AEC—weak tameness does not suffice. It is to be hoped that a more general argument can be found.

Regardless, we have a partial stability spectrum result (of sorts) for weakly tame AECs and, moreover, the only such result that is not limited to local transfer of the kind covered in [3]. 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 $(\mathcal{A}_{<\lambda})^{mor}$. As in Section 6, this reduction of broad structural questions to ones involving only the smallest models emerges as a central feature—and central virtue—of AECs as seen through the lens of accessible category theory.

References

[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.* Number 50 in University Lecture Series. American Mathematical Society, 2009. Available at `http://www.math.uic.edu/~jbaldwin/pub/AEClec.pdf`.

[3] 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.

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

[5] 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.

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

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

[8] 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.

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

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

[11] Michael Lieberman. Rank functions and partial stabililty spectra for AECs. Submitted. Draft at `http://arxiv.org/abs/1001.0624v1.`

[12] Michael Lieberman. *Topological and category-theoretic aspects of abstract elementary classes*. PhD thesis, University of Michigan, 2009.

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

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

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

[16] 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.