ACCESSIBLE CATEGORIES, SET THEORY, AND MODEL THEORY: AN INVITATION

SEBASTIEN VASEY

Abstract. We give a self-contained introduction to accessible categories and how they shed light on both model- and set-theoretic questions. We survey for example recent developments on the study of presentability ranks, a notion of cardinality localized to a given category, as well as stable independence, a generalization of pushouts and model-theoretic forking that may interest mathematicians at large. We give many examples, including recently discovered connections with homotopy theory and homological algebra. We also discuss concrete versions of accessible categories (such as abstract elementary classes), and how they allow nontrivial "element by element" constructions. We conclude with a new proof of the equivalence between saturated and homogeneous which does not use the coherence axiom of abstract elementary classes.

Contents

Date: January 7, 2020
AMS 2010 Subject Classification: Primary: 18C35. Secondary: 03C45, 03C48, 03C52, 03C55, 03C75, 03E05, 03E55.
Key words and phrases. accessible categories, internal size, stable independence, abstract elementary class.

1. Introduction

Category theory, model theory, and set theory are all foundational branches of mathematics. In this paper, I will attempt to give a taste of the food one obtains when mixing the three of them together. I do not really know how to name this dish but, for the present paper at least and at the risk of scaring researchers from all three fields, I will call it *categorical model theory*. This terminology appears in the title of Makkai and Paré's book [MP89], but unfortunately leaves out the set-theoretic aspect. Let me then reassure the set theorists: set theory is very much part of categorical model theory (see Section 4).

So what is categorical model theory? To start with, a *model* (or *structure*) consists of a set (the "universe") together with relation and functions on that set. For example the set of integers together with addition (seen as a binary function) is a model. Part of model theory studies the *definable subsets* of such structures: sets that can be expressed from the relations and functions using simple set operations such as union, intersection, complement, and projection. This approach has had quite a bit of success (one of the easiest outcomes to describe is the theory of o-minimality [vdD98, PW06]). It *is* moreover possible to study definable sets in a categorical setup (a field named categorical logic, see e.g. [MR77]). Nevertheless, this is *not at all* what we will do here.

To make a loose analogy, categorical model theory relates to studying definable sets of a fixed structure in roughly the same way that thermodynamics relates to quantum physics: category theory looks at *an entire class of structures*, together with morphisms between them. Thus instead of only looking at $(\mathbb{Z}, +)$, we will look at the class of abelian groups. What should be the morphisms in that class? The classical answer is that it should be the group homomorphisms – the maps preserving the additive structure – yielding the well known category **Ab**. There are however other possible choice of morphisms. For example, an abelian group is a structure $(A, +)$ satisfying certain axioms that can be expressed using first-order sentences (including for example $(\forall x \forall y)(x + y = y + x)$). In model-theoretic terminology, it is a model of the theory (i.e. set of sentences) T_{ab} of abelian groups. When working abstractly with a class of models of a given first-order theory T, what should be the "right" notion of morphism? One answer that is well studied in model theory is the notion of an *elementary embedding*: an embedding preserving all formulas, possibly with parameters. The elementary embeddings preserve a lot more than the homomorphisms (in particular, they preserve the definable sets) and yield another category, $\mathbf{Elem}(T)$, whose objects are the models of T and whose morphisms are elementary embeddings. Tarski and Vaught showed that $\mathbf{Elem}(T)$ always has directed colimits (the class $\mathbf{Elem}(T)$ is closed under unions of chains ordered by elementary substructures), while this is not always true for the category $\mathbf{Mod}(T)$ of models of T with homomorphisms.

An elementary embedding of abelian groups is quite difficult to describe, and for a category theorist the category $\mathbf{Elem}(T_{ab})$ of abelian groups with elementary embeddings is quite poorly behaved (all morphisms are monomorphisms, so it lacks a lot of limits and colimits; it does not have any notion of quotient for example). On the other hand, from the set-theoretic point of view the category $\mathbf{Elem}(T_{ab})$ is quite interesting and easy to work with, precisely because all morphisms are

monomorphisms (for many purposes, one can think of the maps as inclusions, and think of this category as a kind of concrete poset, with certain isomorphisms between the elements). Still, there may be other interesting types of monomorphisms (for abelian groups, the injective homomorphisms, or the pure morphisms, are obvious choices). Another problem that comes up with the traditional model-theoretic approach is that interesting classes may (provably) not be models of a first-order theory: consider torsion abelian groups, Banach spaces, Zilber's pseudo-exponential fields [Zil05], etc. Moving to infinitary logics may partly fix this second problem but one is still somewhat tied to notions of elementary embeddings. There are other issues with the traditional "Tarskian" definition of a model as a set with functions and relations on it (see Macintyre's essay [Mac03]). Several natural categories, such as arrow categories and more generally categories of functors, cannot easily be described within the Tarskian frame, for example. Let us then agree to take a categorical approach: categorical model theory will study categories that, in some sense, look or behave like the categories of models studied in classical model theory.

What kind of "classical model-theoretic" behavior are we looking for? Two basic results of model theory are the aforementioned Tarski-Vaught chain theorem (closure under chains of elementary substructures) and the downward Löwenheim-Skolem-Tarski theorem (every structure has a "small" elementary substructure). Category-theoretically, the first can be described by the existence of certain colimits (of chains, or more generally of directed diagrams). The second needs a notion of "smallness", i.e. really a notion of size. In a category, objects don't have a "universe", and even if they do, the cardinality of the universe may not tell us much. Nevertheless, it *is* possible to define a notion of size, the presentability rank, by looking at how an object embeds into sufficiently directed colimits (as a simple example, a finite set contained inside an infinite union will be contained inside a component of the union; this property could serve as a definition of finiteness – thus presentability ranks generalize cardinalities to other categories than the category of sets). After making this idea precise, we arrive at the definition of an *accessible category*, one of the main frameworks of categorical model theory. Roughly, it is a category with all sufficiently directed colimits so that any object is a directed colimits of a fixed set of "small" subobjects (see Definition 2.7 here). For example, a set is a union of its finite subsets, an abelian group is a union of its finitely generated subgroups, etc. Accessible categories were first defined by Lair [Lai81], their theory was created by Makkai-Paré [MP89] and further developed in Adámek-Rosický [AR94]. I invite the reader to consult these references for more on the history. There are *a lot* of examples of accessible categories, including all those discussed before, **Ab**, Banach spaces with contractions (or isometries), $\mathbf{Elem}(T)$ for any first-order theory T, and more generally $\mathbf{Elem}(\phi)$ for ϕ an $\mathbb{L}_{\infty,\infty}$-formula[1] and a suitable notion of elementary embeddings. In passing, let us note that the fact that these examples encompass those studied in continuous model theory may well make accessible categories an interesting framework to reunify continuous with discrete model theory (see e.g. [LR17b]).

Before going further, however, let's address a question a logically-inclined reader may have: what about the (first-order) compactness theorem? Well, you cannot

[1]Roughly, $\mathbb{L}_{\infty,\infty}$ is the logic where we allow infinitary quantifications, conjunctions, and disjunctions. See Section 2.4.

make an omelet without breaking eggs: in view of Lindström's theorem [CK90, 2.5.4], going significantly beyond $\mathbf{Elem}(T)$ will imply losing the compactness theorem. In fact, the situation is worse than this: when passing to a category we "forget" a lot of the logical structures on the objects. For example, we cannot really study definable sets anymore: starting with a category of models with elementary embeddings, we can form a new category by "Morleyizing" – adding a relations for each formula. The Morleyized category is (even concretely) isomorphic to the old one, but model theorists studying quantifier elimination would not want to identify the two categories[2]...

Why in the world, then, would one want to forget the logical structure? We have already argued that many mathematical categories of interest cannot be studied model-theoretically. Even in model theory, it can be helpful to distinguish between *internal* properties (visible at the level of a single structure, such as quantifier elimination) and *external* properties (visible by looking at the structure of the category). More precisely, let us define an external property to be a property that is invariant under equivalence of categories. One example of an interesting external property is the existence of a universal object in a given cardinality. Such properties show up in Shelah-style model theory [She90]. In fact, Shelah has observed [She09a, p. 23] that what he calls *dividing lines* (e.g. model-theoretic stability, simplicity, NIP, etc.) can be characterized by both internal and external properties[3]. Stability, for example, is equivalent to failure of the order property (internal), or to the existence of saturated models in certain cardinals (external). Thus, while dividing lines do have a logical characterization, they are also invariant under equivalence of categories. This is interesting insight suggests it may be possible to still do Shelah-style model theory categorically (and this is indeed the case, as exemplified in the large body of work on classification theory for AECs, see the references in Section 8).

At a broader level, working with accessible categories entails a higher level of generality. This has downsides but also benefits: more categories of interest are accessible, and accessible categories are closed under more operations. This can make the theory easier to develop in some cases. For example, starting from the class of models of a first-order T, one can form its category of λ-saturated models (for some λ). One can also form its class of models omitting some type. One could even just look at the class of models of T of cardinality at least λ. *None* of these examples are classes of models of a first-order but they are still accessible categories. As another example, there is a very general definition of stability, using forking independence, that can be given in any accessible category and makes no mention of logic. This definition specializes to (and is arguably much simpler than) the usual first-order one. See Section 5.

Coming back to the failure of the compactness theorem, another counterpoint is that there are many different types and levels of compactness[4]. Some compactness can be recovered using large cardinals (e.g. compactness for $\mathbb{L}_{\kappa,\kappa}$), other types of

[2]In fact I have seen categorical model theory described as "model theory without logic".

[3]Shelah does not give a precise meaning to "external": the interpretation as "invariant under equivalence of category" is my own suggestion.

[4]In fact, the condition in the definition of an accessible category that every object is a directed colimit of a fixed set of small subobjects can itself be thought of as a weak type of compactness.

compactness are implied by the "complexity" of the class or category under consideration (for example universal classes, see Definition 3.21, do satisfy a weak version of the compactness theorem [ABV19, 3.8]; classes with low descriptive set-theoretic complexity also behave better than the general case [BLS15]). Still other types of compactness are implied by the stability-theoretic properties of the class (e.g. in a first-order stable theory, any long-enough sequences contains an indiscernible subsequence, c.f. [She90, I.2.8] or Theorem B.7, but this is not the case in general without large cardinals [Jec03, 18.18]). This interplay between set-theoretic, model-theoretic, and stability-theoretic compactness is a fascinating aspect of categorical model theory, which is harder to see and study in setups where the compactness theorem applies.

At this point a still skeptical reader may say that, while all these philosophical points are interesting, the setup of accessible categories seems too general or difficult for an interesting theory. There are several answers to this objection. First, category theory itself is very general, but it still is an interesting framework in which to present and understand many different branches of mathematics. There *are* interesting theorems in category theory (see the epilogue of [Rie16]), but of course they often inform and supplement rather than completely supersede results in more specialized branches. This does not make the importance of category theory in doubt. Similarly, I think of accessible category as a framework rather than as an all-encompassing object of study. There are less general frameworks (abstract elementary classes, universal classes, first-order model theory, universal algebra, $\mathbb{L}_{\omega_1,\omega}$, locally presentable categories, Grothendieck abelian categories, etc.) in which one will be able to say more, but may also sometimes be more limited by the lower generality. The "best" framework depends on the question; finding it is a hard part of the mathematician's job.

Second, there *are* nontrivial theorems about accessible categories. Let's just mention a basic one for now: *any* accessible category is equivalent to the category of models of an $\mathbb{L}_{\infty,\infty}$ sentence, with morphisms the homomorphisms (Corollary 3.19). For example, we have already seen that we can Morleyize the models of a first-order theory to obtain a category where the elementary embeddings are simply the injective homomorphisms. Since injective homomorphisms are simply homomorphisms that preserve the non-equality relation, we can add this non-equality relation to the models to get an instance of the theorem. This result gives a surprising correspondence between categorical and "logical" model theory (taken in the broad sense of studying classes of models of $\mathbb{L}_{\infty,\infty}$). Thus studying categories of models of $\mathbb{L}_{\infty,\infty}$ theories is just as hard as studying accessible categories generally. The correspondence also shows that, for the purpose of categorical model theory, the logics with generalized quantifiers described for example in [BFB85] are not necessary: if we study the category of models of a certain sentence expressed in a complicated logic with quantifiers such as "there exists uncountably many", together with a certain logical notion of morphism, then as long as this category is an accessible category, we can find an equivalent category that will be axiomatized simply in $\mathbb{L}_{\infty,\infty}$. See Example 3.20(2), the "toy quasiminimal class", for an instance of this phenomenon.

An intermediate type of setup is abstract model theory, which studies frameworks such as abstract elementary classes (AECs) [She87]. There we are still studying "Tarskian" classes of models, but the notion of embedding is axiomatically specified

rather than defined to be elementarity. Any AEC can naturally be seen as an accessible category, and any class of models of an $\mathbb{L}_{\infty,\infty}$ theory can naturally be seen as an ∞-AEC (Definition 2.12). Thus there is a three way correspondence[5] between categorical model theory (accessible categories), abstract model theory (∞-abstract elementary classes), and logical model theory ($\mathbb{L}_{\infty,\infty}$).

Such a correspondence should be compared with Tarski's presentation theorem [Tar54] (see Theorem 3.22 here): classes closed under substructure, isomorphisms, and unions of chains (abstract model theory) are the same as classes of models of universal theories (logical model theory). Another example of this phenomenon is given by the Birkhoff variety theorem [Bir35, 10]: a class of algebras is a variety (logical model theory) if and only it is closed under products, subobjects, and quotients (categorical model theory), see [AR94, 3.9]. In fact it turns out that many types of accessible categories have natural logical and abstract classes characterizations [AR94, LRV19c]. This is useful, because one can often take advantage of the concreteness of abstract model theory to use certain set-theoretic ideas. "Element by element" constructions in abstract elementary classes are a case in point (Section 6).

Third, accessible categories are already well connected to the rest of mathematics. While Makkai and Paré had a model-theoretic motivation, one of the first use of accessible categories was in algebraic topology: a *model category* is a category endowed with three distinguished classes of morphisms, called fibrations, cofibrations and weak equivalences, satisfying axioms that are properties of the category of topological spaces (with the weak equivalences being the weak homotopy equivalences). In particular, weak equivalences should satisfy a "two out of three" property, and the morphisms should form certain weak factorization systems (e.g. any morphism should factor, in a somewhat canonical way, as a cofibration followed by a fibration that is also a weak equivalence). The reader can think of fibrations as "nice surjections" and cofibrations as "nice inclusions". See [Hes02] for a survey of the use of model categories in algebraic topology, and e.g. [Hov99] for the general theory. It turns out that the most convenient model categories to work with are the locally presentable ones (i.e. those that are bicomplete and accessible). This is because, roughly speaking, in such categories we can often find set-sized families of "small" cofibrations that generate the rest. Such model categories are called *cofibrantly generated*. For example, the category of topological spaces is not accessible (Example 2.5(7)), but the category of *simplicial sets* is. The model category induced on it will, it turns out, be cofibrantly generated, and in some homotopical sense equivalent to the one on topological spaces. An application of the theory of accessible categories to model categories is the existence, assuming Vopěnka's principle, of certain localizations [CSS05, RT03]. See also [Bek00] for more on the connections between accessible categories, logic, and model categories.

Accessible categories have also been used in homological algebra. An only recently solved problem, the *flat cover conjecture*, asks whether every module has a flat cover [Eno81]. This was proven by Bican, El Bashir, and Enochs [BBE01]. Two proofs were given, one using set-theoretic tools of Eklof-Trlifaj [ET01]. This proof was recognized by Rosický [Ros02] to really be an instance of a "small object

[5]Since any ∞-AEC has all morphisms monomorphisms, the precise correspondence should say that accessible categories with all morphisms monos are equivalent to ∞-AECs.

argument": a way to build cofibrantly generated weak factorization systems in any locally presentable category.

I believe these connections are interesting, especially because logic and model theory have not historically threaded too much into algebraic topology and homological algebra. If accessible categories are really tied to model theory, then these results should have also model-theoretic meaning. Recently, joint work with Lieberman and Rosický [LRVc] showed that in fact the notion of a model category (or more generally a weak factorization system) being cofibrantly generated is closely tied to the theory of model-theoretic forking. Essentially, in a locally presentable model category, a weak factorization system is cofibrantly generated exactly when the category obtained by restricting to only to cofibrations is stable, in the sense of having a forking-like independence notion. See Section 5 for an overview of this result.

In conclusion, categorical model theory is a fascinating mix of set theory, category theory, and model theory that sheds light on all these topics and on some other parts of mathematics. This paper is meant to survey some basic results, as well as discuss some topics of current research. I will survey set-theoretic aspects (Section 4), as well as questions on stable independence (Section 5), a recent research development that may be of interest to model theorists, category theorists, and mathematicians at large. Some methods applicable to concrete setups such as AECs, Section 6 will also be discussed. I will end with some open problems (Section 7), as well as a short list of helpful resources to learn more (Section 8). A first appendix gives a streamlined method for handling the bookkeeping in point by point constructions, and the second appendix gives a very short introduction to first-order stability theory.

I will assume that the reader is familiar with very basic logic and model theory [CK90], set theory [HJ99], and category theory [Lan98, AHS04]. Still, an effort has been made to repeat many standard definitions and provide examples and intuitions behind proofs for those that are not necessarily proficient in all three fields.

Acknowledgments. I would like to thank Jiří Rosický and Michael Lieberman for introducing me to accessible categories and taking an early look at this survey. The work on category-theoretic sizes and stable independence presented here comes from our joint collaboration.

I thank Marcos Mazari-Armida and the referee for helpful comments. I also thank Justin Cavitt, Rebecca Coulson, and Rehana Patel, for encouraging me to write Appendix B.

2. Main definitions and examples

2.1. **Set theory.** We assume familiarity with ordinals and cardinals. We think of ordinals as transitive sets ordered by membership, and identify cardinals with the corresponding ordinals. We denote by ω the first infinite ordinal, i.e. the set of natural numbers. For a cardinal λ, we write λ^+ for the successor of λ, the minimal cardinal strictly bigger than λ. Cardinals of the form λ^+ are called *successor cardinals*, and cardinals that are not successors are called *limit*. We write ${}^Y X$ for the set of functions from Y to X and, for an ordinal α, ${}^{<\alpha}X$ for $\bigcup_{\beta<\alpha} {}^\beta X$. For λ

and μ cardinal, $\lambda^{<\mu} = |{}^{<\mu}\lambda|$ and $\lambda^{\mu} = |{}^{\mu}\lambda|$. A *partially ordered set* (or poset for short) is a transitive, reflexive, and antisymmetric relation. A subset I_0 of a poset I is *cofinal* if for all $i \in I$ there is $i_0 \in I_0$ so that $i \leq i_0$. The *cofinality*, $\text{cf}(I)$, of I is the minimal cardinality of a cofinal subset of I. Of course, we identify an ordinal α with the corresponding linear order $(\alpha, \in)$. A cardinal λ is *regular* if it is infinite and $\text{cf}(\lambda) = \lambda$.

2.2. **Category theory.** A category is called *small* if it has only a set of objects, and it is called *large* if it has a proper class of *non-isomorphic* objects. All the categories in this paper will be locally small (i.e. the collection $\text{Hom}(A, B)$ of morphisms from the object A to the object B is always a set, never a proper class). Our categorical notation and conventions will mostly be those of [AHS04] and [AR94]. In particular, a *monomorphism* (or *mono* for short) is a morphism f such that $fg_1 = fg_2$ implies $g_1 = g_2$ for any two morphisms g_1 and g_2. The notion of an *epimorphism* (*epi* for short) is defined dually. A *diagram* in a category $\mathcal{K}$ is a functor $D : I \to \mathcal{K}$, where I is a (here always small) category, called the *index* of the diagram. In this paper, the index I will usually be a partially ordered set, which will be identified with the corresponding category. A *cocone*[6] for a diagram $D : I \to \mathcal{K}$ consist in an object A together with morphism $(D_i \xrightarrow{f_i} A)_{i \in I}$ such that whenever $i \xrightarrow{d} j$ is a morphism of D, we have that $f_i = f_j d$. The cocones for D form a category, $\mathcal{K}_D$, where a morphism from $(D_i \xrightarrow{f_i} A)_{i \in I}$ to $(D_i \xrightarrow{g_i} B)_{i \in I}$ is a $\mathcal{K}$-morphism $A \xrightarrow{h} B$ so that $hf_i = g_i$ for all $i \in I$. A *colimit* for D is an initial object in the category $\mathcal{K}_D$ (recall that an initial object in a category is one so that there is a unique morphism from it to any other object).

2.3. **Presentability and accessible categories.** For λ a cardinal, a partially ordered set is called λ-*directed* if every subset of cardinality strictly less than λ has an upper bound. Note that any non-empty poset is λ-directed for $\lambda \leq 2$ (an empty poset is 0-directed but not 1-directed), 3-directed is equivalent to $\aleph_0$-directed, and for a singular cardinal λ, λ-directed is equivalent to λ^+-directed. Thus we usually assume that λ is a regular cardinal. We will say *directed* instead of $\aleph_0$-directed. For λ regular, an example of a λ-directed poset that the reader can keep in mind is λ itself, seen as a chain (i.e. a linear order) of order type λ. Another typical example is the poset $[A]^{<\lambda}$ of all subsets of a fixed set A which have cardinality strictly less than λ, ordered by containment.

A λ-*directed diagram* is a diagram indexed by a λ-directed poset. A category has λ-*directed colimits* if any λ-directed diagram has a colimit. Note that for $\lambda_1 < \lambda_2$, λ_2-directed implies λ_1-directed, hence having λ_1-directed colimits is *stronger* than having λ_2-directed colimits. The following observation will be used without comments:

Lemma 2.1 (Iwamura's lemma [AR94, 1.7])**.** If a category has colimits of all chains indexed by regular cardinals, then it has directed colimits.

Proof sketch. It is immediate using the definition of cofinality that a category with all colimits of chains indexed by regular cardinals has colimits of all chains indexed

[6]In this paper, we will be more interested in colimits than limits, so we also define cocones rather than cones.

by ordinals. Now any finite directed poset has a maximum, so colimits of such diagrams are trivial. Furthermore, any infinite directed partially ordered set I can be written as $\bigcup_{\alpha<|I|} I_\alpha$, where each I_α has strictly smaller cardinality than I, is directed, and $I_\alpha \subseteq I_\beta$ for $\alpha < \beta$. We can therefore proceed by induction on $|I|$ to show that the category has all colimits indexed by I. □

Remark 2.2. A similar result no longer holds in the uncountable case [AR94, Exercise 1.c(2)]: if λ is uncountable, having colimits indexed by ordinals of cofinality at least λ is *strictly weaker* than having all λ-directed colimits.

Example 2.3. In many concrete algebraic cases, directed colimits exist and are essentially computed by taking unions. For example, the category **Ab** of abelian groups has directed colimits, and they can essentially be computed by taking unions of the resulting system of groups. However in the category of Banach spaces with contractions, colimits of increasing chains are given by taking the *completion* of their union. Of course, some categories don't have directed colimits at all. For example, the category of well-orderings, with morphisms the order-preserving maps, does not have directed colimits. It does however have $\aleph_1$-directed colimits. The category of complete Boolean algebras (with homomorphisms), on the other hand, does not have λ-directed colimits for any λ.

The notion of a λ-presentable object is key to the definition of an accessible category.

Definition 2.4. For λ a regular cardinal, an object A of a category $\mathcal{K}$ is λ-*presentable* if the functor $\mathrm{Hom}(A, -)$ preserves λ-directed colimits. Said more transparently, A is λ-*presentable* if for any λ-directed diagram $D : I \to \mathcal{K}$ with colimit $(D_i \xrightarrow{g_i} \operatorname{colim} D)_{i\in I}$, any map $A \xrightarrow{f} \operatorname{colim} D$ factors essentially uniquely through D: there exists $i \in I$ and $A \xrightarrow{f_i} D_i$ such that $f_i g_i = f$. Essentially uniqueness means that, moreover, if $j \in I$ and $A \xrightarrow{f_j} D_j$ are such that $f_j g_j = f$, then there exists $k \geq i, j$ such that $g_{i,k} f_i = g_{j,k} f_j$ (here, $g_{i,j}$ and $g_{i,k}$ are the diagram maps from i to j and i to k respectively).

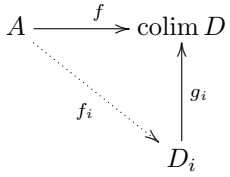

When $\lambda = \aleph_0$, we will say that A is *finitely presentable*. We say that A is *presentable* if it is λ-presentable for some λ. The *presentability rank* of a presentable object A is the least regular cardinal λ such that A is λ-presentable.

Note that a λ-presentable object is also μ-presentable for all regular $\mu > \lambda$.

Example 2.5.

(1) In the category of sets (with morphisms all functions), a set is λ-presentable if and only if it has cardinality strictly less than λ. Indeed, directed colimits are essentially unions, and a set A has cardinality strictly less than λ if and only if whenever $A \subseteq \bigcup_{i<\lambda} A_i$, there exists $i < \lambda$ such that $A \subseteq A_i$. Thus

the presentability rank of an infinite set is the successor of its cardinality. Similarly, in many algebraic categories (in fact in any abstract elementary classes – see Definition 2.12 and Example 3.16(5)), the presentability rank of a big-enough object will be the successor of the cardinality of its universe.

(2) In very algebraic categories, such as **Ab**, an object is λ-presentable if and only if it can be presented using strictly less than λ-many objects and equations [AR94, 3.12]. This is the motivation for the term "presentable".

(3) [LRV19b, 6.2] Consider the category of well-orderings, with morphisms the initial segment embeddings. Note that (as opposed to what would happen if the maps were just order-preserving maps) this category has all directed colimits. An infinite well-order is λ-presentable precisely when it has cofinality strictly less than λ.

(4) Consider the category of all algebras with one ω-ary operation, and morphisms the homomorphisms. Then an object is λ-presentable if and only if it is generated by a set of cardinality strictly less than λ. Here, the free algebra on λ-many generators is λ^+-presentable but has cardinality $\lambda^{\aleph_0}$ (which could be strictly bigger than λ).

(5) Consider the category of complete metric spaces, with isometries as morphisms. For λ regular *uncountable*, a space A is λ-presentable if and only if its it has *density character* strictly less than λ (the density character of a topological space is defined as the least cardinality of a dense subset). Thus the presentability rank is the successor of the density character. This holds more generally in all the "continuous" examples, including Banach and Hilbert spaces [LR17b, 3.1]. Note that in general the density character is different from the cardinality. For example, there are no Hilbert spaces of cardinality λ when $\lambda < \lambda^{\aleph_0}$, but there are no problems finding a Hilbert space generated by an orthonormal basis of cardinality λ. Finite sums from such a basis will give the desired dense subset.

(6) [LRV19b, 6.1][7] Consider the one point metric space $\{0\}$ in the category of complete metric spaces with isometries. This space is $\aleph_1$-presentable (by the preceding discussion) but it is *not* $\aleph_0$-presentable. Indeed, the inclusion of $\{0\}$ into $\{0\} \cup \{\frac{1}{n} \mid 0 < n < \omega\}$ does not factor through $\{\frac{1}{n} \mid 0 < n \leq m\}$, for any $m < \omega$. In fact, the empty metric space is the only $\aleph_0$-presentable object of the category.

(7) [AR94, 1.2(10)] In the category of topological spaces (with morphisms the continuous functions), a discrete spaces (i.e. where every set is open) is λ-presentable exactly when it has cardinality strictly less than λ. On the other hand, a non-discrete space A is *never* presentable. To see this, fix a regular λ and a non-open set $X \subseteq A$. Without loss of generality, $A \cap \lambda = \emptyset$. For $\alpha < \lambda$, let D_α be the space $A \cup \lambda$, where the non-trivial open sets in D_α are of the form $X \cup [\beta, \lambda)$ for $\beta \in [\alpha, \lambda)$. For $\alpha \leq \beta$, the inclusion of D_α into D_β is continuous (D_α has more open sets than D_β), and these inclusions give a λ-directed diagram. The colimit D of this diagram is the indiscrete space on the set $A \cup \lambda$. Because D is indiscrete, the inclusion of A into D is continuous, but it cannot factor through any D_α (because the inverse image of $X \cup [\alpha, \lambda)$, an open set of D_α, is X, which is not open in A).

[7] This example was pointed out to me by Ivan Di Liberti.

In view of the many examples above where the presentability rank was the *successor* of some natural notion of size (see Section 4.1 on what is true in general), it is convenient to have a name for this predecessor. We will call it the internal, or category-theoretic, *size*:

Definition 2.6. The *(internal – or category-theoretic) size* of an object A is the predecessor (if it exists) of its presentability rank.

We can now precisely say what an accessible category is:

Definition 2.7. For a regular cardinal λ, a category $\mathcal{K}$ is λ-*accessible* if:

(1) $\mathcal{K}$ has λ-directed colimits.
(2) (Smallness condition) There is a set S of λ-presentable objects such that every object of $\mathcal{K}$ is a λ-directed colimits of elements of S.

If $\mathcal{K}$ is λ-accessible and bicomplete (i.e. complete and cocomplete: it has all limits and colimits), we say that it is *locally λ-presentable*. When $\lambda = \aleph_0$, we will talk about a finitely accessible or locally finitely presentable category. We say that $\mathcal{K}$ is *accessible* [*locally presentable*] if it is λ-accessible [locally λ-presentable] for some regular cardinal λ.

Remark 2.8. There are three occurrences of the parameter λ in Definition 2.7, and there are no real reasons why these occurrences should all be the same cardinal. Thus we could parameterize the definition further into three cardinals, see [LRV19b, 3.6]. In fact, we can be even more precise and allow singular cardinals in the definition. For example, an $(\aleph_0, \aleph_1, < \aleph_{\omega_1})$-accessible category would be a category with directed colimits where every object is an $\aleph_1$-directed colimits of a fixed set of $(< \aleph_{\omega_1})$-presentable objects, where $(< \aleph_{\omega_1})$-presentable means λ_0-presentable for some regular $\lambda_0 < \aleph_{\omega_1}$. Thus we want each object in the diagram to have presentability rank some $\aleph_\alpha$, $\alpha < \omega_1$, but there may not be a single $\aleph_\alpha$ that bounds the rank of each object. For simplicity, we will not take this approach here but will often discuss, for example, accessible categories with all directed colimits.

Example 2.9.

(1) The category **Set** of sets with functions as morphisms is locally finitely presentable: it is bicomplete and every set is a directed colimit of finite sets.
(2) Any complete lattice is (when seen as a category) a locally presentable category.
(3) The category $\mathbf{Set}_{mono}$ of sets with injective functions as morphisms, is finitely accessible but not bicomplete (for example, it does not have coequalizers; intuitively it is impossible to quotient). There are two general phenomenons at play here: first [LRV19a, 6.2], given any accessible category $\mathcal{K}$, the category $\mathcal{K}_{mono}$ obtained by restricting to its monomorphisms will be accessible as well. Second, if $\mathcal{K}$ is an accessible category where all morphisms are monomorphisms (this is typically the case in model-theoretic setups), then $\mathcal{K}$ cannot be locally presentable, unless it is a complete lattice (see the previous example), so in particular small.
(4) The category **Ab** of a abelian groups is locally finitely presentable. The finitely generated abelian groups are the finitely presentable objects and

any other abelian group is a directed colimit of those. More generally, for a fixed ring R, the category R-**Mod** of R-modules is locally finitely presentable.

(5) The category **Gra** of graphs (symmetric reflexive binary relations) with morphisms the graph homomorphisms is locally finitely presentable (finitely presentable objects are finite graphs).
(6) The category **Ban** of Banach spaces (with morphisms the contractions) is locally $\aleph_1$-presentable, but not locally $\aleph_0$-presentable (for reasons connected to Example 2.5(6)).
(7) The category **Bool** of boolean algebras (with morphisms the boolean algebra homomorphisms) is locally finitely presentable. On the other hand, the category of *complete* boolean algebras is not even accessible: it does not have λ-directed colimits for any λ. For a fixed regular cardinal λ, the category of λ-complete Boolean algebras is however λ-accessible.
(8) The category of well-orderings with morphisms the order-preserving maps is $\aleph_1$-accessible but not finitely accessible (it does not have directed colimits). On the other hand the category of well-orderings with morphisms the initial segment maps is *not* accessible, even though it has directed colimits: for any regular cardinal λ, there is a proper class of non-isomorphic well-orderings of cofinality λ, see Example 2.5(3). We could still look at the full subcategory of well-orderings of order type λ^+ or less. It turns out that this will be a λ^+-accessible category.
(9) The category of all well-founded models of (a sufficiently-big fragment of) ZFC, or of all well-founded models of ZFC + V = L, with morphisms the elementary embeddings, is $\aleph_1$-accessible. A variation of this example is studied in [LRV19b, §6.1].
(10) Any Fraïssé class can be seen as generating a finitely accessible category. In fact, there exists a general categorical theory of Fraïssé constructions [Kub14].
(11) The category of all free abelian groups (with group homomorphisms) is not accessible if $V = L$, but is κ-accessible for κ a strongly compact cardinal [MP89, §5.5]. Thus whether a category is accessible can, in certain cases, be a set-theoretic question. See Section 4.2 for more on this phenomenon.

2.4. **Categories of structures.** In order to get more examples of accessible categories and introduce related frameworks, let us move toward logic. For completeness and because we will work with infinitary languages, we start by repeating the basic definitions. The reader will not lose much by skipping them. More details can be found in [Dic75] or [AR94, Chapter 5]. For κ an infinite cardinal, a *κ-ary vocabulary* (or signature) is a set τ containing[8] relations and functions symbols of arity strictly less than κ. When $\kappa = \aleph_0$, we call τ a finitary vocabulary, and when $\kappa \geq \aleph_1$ an infinitary vocabulary. By default, a *vocabulary* means a κ-ary vocabulary for some κ (so possibly infinitary). More precisely, a κ-ary vocabulary τ is a a set $\langle n_i : i \in I_R, m_j : j \in I_F \rangle$, where I_R and I_F are disjoint index sets and n_i, m_j are cardinals strictly less than κ. For such a vocabulary, a *τ-structure* M consists of a set $A = UM$ (the *universe*), for each $n_i \in \tau$, an n_i-ary relation R_i on A, and for each $m_j \in \tau$ an m_j-ary function f_j on A. A *term* in the vocabulary τ, over a given set

[8]We will work with a single sort for simplicity.

of variables V (disjoint from any other object that we consider, and of cardinality κ – we will never mention V again), is defined inductively as either a variable x, or as $f_j(\rho_\alpha)_{\alpha<m_j}$, for each ρ_α a term. An *atomic formula* is an expression of the form $\top$, $\bot$, $\rho = \rho'$ or $R_i(\rho_\alpha)_{\alpha<n_i}$ for terms ρ ,ρ', ρ_α. For $\lambda \geq \kappa$ an infinite cardinal, an $\mathbb{L}_{\lambda,\kappa}$ *formula* is defined inductively as either an atomic formula, $\bigwedge_{\alpha\in S}\phi_\alpha$, $\bigvee_{\alpha\in S}\phi_\alpha$, $\phi \to \psi$, $\neg\phi$, $(\forall\bar{x})\phi$, $(\exists\bar{x})\phi$, where ϕ_α, ϕ, ψ, are formulas, $|S| < \lambda$, $\bar{x}$ is a sequence of variables of length strictly less than κ, and we require that the formulas obtained from conjunctions and disjunctions still have fewer than κ-many variables. We also define $\mathbb{L}_{\infty,\kappa} = \bigcup_\lambda \mathbb{L}_{\lambda,\kappa}$ and $\mathbb{L}_{\infty,\infty} = \bigcup_\kappa \mathbb{L}_{\infty,\kappa}$ as expected. The *free variables* of a formula are the ones that appear in it and are not bound by any quantifier. We write $\phi(\bar{x})$ for a formula ϕ with free variables among $\bar{x}$. A *sentence* is a formula without free variables, and a *theory* is a set of sentences. For a τ-structure M, $\bar{a}$ a sequence of elements from M, and $\phi(\bar{x})$ a formula ($\bar{x}$ and $\bar{a}$ of the same length), we define what it means for M to satisfy (or be a model of) ϕ, with $\bar{a}$ standing for $\bar{x}$, $M \models \phi(\bar{a})$ for short, as expected. A τ-structure satisfies (or is a model of) a theory if it satisfies all sentences of the theory.

For a vocabulary τ, a *homomorphism* from a τ-structure M to a τ-structure N is a function f from UM to UN that preserves all atomic formulas: if $\phi(\bar{x})$ is atomic, $\bar{a}$ is a sequence in M, and $M \models \phi(\bar{a})$, then $N \models \phi(f(\bar{a}))$. We let $\mathbf{Str}(\tau)$ denote the category of all τ-structures with homomorphisms. For T a theory in $\mathbb{L}_{\infty,\infty}$, we let $\mathbf{Mod}(T)$ denote the full subcategory of $\mathbf{Str}(\tau)$ consisting of all models of T. When ϕ is a sentence, $\mathbf{Mod}(\phi)$ will denote $\mathbf{Mod}(\{\phi\})$. Note that $\mathbf{Mod}(\phi)$ is *not* always an accessible category, essentially because the homomorphisms are not the right notion of embedding when ϕ is too complex. Thus we more generally define, for $\Phi \subseteq \mathbb{L}_{\infty,\infty}$, a Φ-*elementary map from* M *to* N to be a function f from UM to UN that preserves all formulas in Φ. We let $\mathbf{Elem}_\Phi(\tau)$ denote the category of all τ-structures with homomorphism, and $\mathbf{Elem}_\Phi(T)$ denote the full subcategory of $\mathbf{Elem}_\Phi(\tau)$ consisting of models of T. One can check that for λ a regular cardinal and T an $\mathbb{L}_{\lambda,\lambda}$ theory, $\mathbf{Elem}_{\mathbb{L}_{\lambda,\lambda}}(T)$ is λ-accessible [MP89, §3.4].

We will focus on the following simple type of formulas:

Definition 2.10. A formula of $\mathbb{L}_{\infty,\infty}$ is *positive existential* if it can be built from atomic formulas using only conjunctions, disjunctions and existential quantifications. A formula is *basic* if it is of the form $(\forall\bar{x})(\phi \to \psi)$, where ϕ and ψ are positive existential. A *basic* theory is a set of basic sentences.

Note that any theory in $\mathbb{L}_{\infty,\infty}$ can be "Morleyized" to a basic theory, by adding a relation symbol for each formula [MP89, 3.2.8]. In fact, for $\kappa \leq \lambda$ and T a theory in $\mathbb{L}_{\lambda,\kappa}$, $\mathbf{Elem}_{\mathbb{L}_{\lambda,\kappa}}(T)$ is isomorphic to $\mathbf{Mod}(T')$, for some basic $\mathbb{L}_{\lambda,\kappa}$-theory T'. This fact is sometimes called *Chang's presentation theorem* by model theorists and is the reason why we can restrict ourselves to basic formulas without losing generality.

Example 2.11. Let κ and λ be regular cardinals.

(1) [AR94, §5.1] If τ is a κ-ary vocabulary then $\mathbf{Str}(\tau)$ is a locally κ-presentable category.
(2) [AR94, 5.35] For any basic theory T in $\mathbb{L}_{\infty,\lambda}$, $\mathbf{Mod}(T)$ is accessible with λ-directed colimits.

In the second result, existence of λ-directed colimits can be proven directly. The smallness condition in the definition of an accessible category follows, for example, from the infinitary downward Löwenheim-Skolem theorem.

It turns out that Example 2.11 is sharp, in the sense that any accessible category is equivalent to a category of the form $\mathbf{Mod}(T)$, for T a basic theory (Corollary 3.19). Similarly, it is possible to characterize the locally presentable categories as the categories of models of certain theories (called limit). See [AR94, 5.30].

2.5. **∞-abstract elementary classes.** We now introduce ∞-abstract elementary classes (∞-AECs). They are in some sense a compromise: less abstract than accessible categories, but still more abstract than categories of models of basic theories. The definition of an AEC (the case $\mu = \aleph_0$ below) is due to Shelah [She87]. It was generalized to the case of a μ-AEC in [BGL$^+$16]. We introduce essentially the same notion, but to be precise we make μ (and the vocabulary) part of the data. The reader should first read it with the case $\mu = \aleph_0$ in mind.

Definition 2.12. An *∞-abstract elementary class* (∞-AEC for short) $\mathbf{K}$ consists of:

(1) A regular cardinal $\mu = \mu(\mathbf{K})$.
(2) A μ-ary vocabulary $\tau = \tau(\mathbf{K})$.
(3) A class K of τ-structures.
(4) A partial order $\leq_{\mathbf{K}}$ on K.

satisfying the following four axioms:

- Abstract class axiom: K is closed under isomorphisms, $M \leq_{\mathbf{K}} N$ implies that M is a τ-substructure of N (i.e. $UM \subseteq UN$ and the inclusion $M \to N$ is a homomorphism; we write $M \subseteq N$), and $\leq_{\mathbf{K}}$ respects isomorphisms in the sense that if $M, N \in K$, $M \leq_{\mathbf{K}} N$, and $f : N \cong N'$, then $f[M] \leq_{\mathbf{K}} N'$.
- Coherence axiom: if $M_0, M_1, M_2 \in K$, $UM_0 \subseteq UM_1$, $M_1 \leq_{\mathbf{K}} M_2$, and $M_0 \leq_{\mathbf{K}} M_2$, then $M_0 \leq_{\mathbf{K}} M_1$.
- Tarski-Vaught (TV) chain axiom: if $\langle M_i : i \in I \rangle$ is a μ-directed system in $\mathbf{K}$ (i.e. I is a μ-directed poset, all the M_i's are in K, and $i \leq j$ implies $M_i \leq_{\mathbf{K}} M_j$), then letting $M := \bigcup_{i \in I} M_i$ (defined as expected), we have:
  - $M \in K$.
  - $M_i \leq_{\mathbf{K}} M$ for all $i \in I$.
  - If $N \in K$ and $M_i \leq_{\mathbf{K}} N$ for all $i \in I$, then $M \leq_{\mathbf{K}} N$.
- Löwenheim-Skolem-Tarski (LST) smallness axiom: there exists a cardinal $\lambda = \lambda^{<\mu} \geq |\tau| + \mu$ such that for any $M \in K$ and any $A \subseteq UM$, there exists $M_0 \in K$ with $M_0 \leq_{\mathbf{K}} M$, $A \subseteq UM_0$, and $|UM_0| \leq |A|^{<\mu} + \lambda$. We write $\mathrm{LS}(\mathbf{K})$ (the *Löwenheim-Skolem-Tarski number* of $\mathbf{K}$) for the minimal such λ.

Unless K is empty, the vocabulary $\tau(\mathbf{K})$ can always be recovered from K. Thus we usually just write $\mathbf{K} = (K, \leq_{\mathbf{K}})$ and say that $\mathbf{K}$ is a μ-AEC to make the μ associated with it clear. When $\mu = \aleph_0$, we omit it and just say that $\mathbf{K}$ is an AEC. We also will not distinguish between K and $\mathbf{K}$, writing for example $M \in \mathbf{K}$ instead of $M \in K$. Another convention: if we write $M \leq_{\mathbf{K}} N$, we will automatically mean that also $M, N \in K$.

In any ∞-AEC, there is a natural notion of morphism.

Definition 2.13. Let $\mathbf{K}$ be an ∞-AEC. For $M, N \in \mathbf{K}$, a $\mathbf{K}$-*embedding* from M to N is an injective $\tau(\mathbf{K})$-homomorphism f from M to N such that $f[M] \leq_{\mathbf{K}} N$. When $\mathbf{K}$ is clear from context, we write $f : M \to N$ to mean that f is a $\mathbf{K}$-embedding from M to N. We will often identify $\mathbf{K}$ with the category whose objects are the structures in $\mathbf{K}$ and morphisms the $\mathbf{K}$-embeddings.

Remark 2.14. If $\mathbf{K}$ is a μ-AEC, then it has the following properties as a category:

- It is a subcategory of $\mathbf{Str}(\tau(\mathbf{K}))$. Further, it is *isomorphism-closed* in $\mathbf{Str}(\tau(\mathbf{K}))$ (a subcategory $\mathcal{L}$ of a category $\mathcal{K}$ is *isomorphism-closed* – or *replete* – if whenever A is an object of $\mathcal{L}$ and $f : A \to B$ is an isomorphism of $\mathcal{K}$, then both f and B are in $\mathcal{L}$). This is the essential content of the abstract class axiom.
- It is a *concrete* category, as witnessed by the universe functor $U : \mathbf{K} \to \mathbf{Set}$.
- All its morphisms are monomorphisms, and in fact concrete monomorphisms (i.e. they are also monomorphisms in the category of sets – injective functions). More is true: it is noticed in [AHS04, §8] that the "right" notion of monomorphism in many examples ends up being the notion of a concrete embedding [AHS04, 8.6] whose definition mirrors the coherence axiom of AECs. In fact, what the coherence axiom says is exactly that the morphisms of $\mathbf{K}$ are concrete embeddings in the sense of [AHS04, 8.6].
- It has μ-directed colimits (this is the essential content of the chain axiom). In fact these μ-directed colimits are *concrete* in the sense that the functor U preserves them: they are computed the same way as in $\mathbf{Set}$, by taking unions.
- It is an $\text{LS}(\mathbf{K})^+$-accessible category. Indeed, the objects of cardinality at most $\text{LS}(\mathbf{K})$ are all $\text{LS}(\mathbf{K})^+$-presentable, and there is only a set of them up to isomorphism. Moreover, the coherence and smallness axioms together imply that any other object $M \in \mathbf{K}$ can be written as the $\text{LS}(\mathbf{K})^+$-directed union of $\{M_0 \in \mathbf{K} \mid M_0 \leq_{\mathbf{K}} M, |UM_0| \leq \text{LS}(\mathbf{K})\}$ (we think of it as a $\leq_{\mathbf{K}}$-system indexed by itself). Note however that $\mathbf{K}$ need *not* be μ-accessible. In fact it is easy to see that for any regular cardinal λ, the ($\aleph_0$-)AEC of all sets of cardinality at least λ (ordered with subset) is *not* λ-accessible. We will see later (Theorem 3.17) that the smallness axiom is, modulo the other axioms, *equivalent* to the smallness condition in the definition of an accessible category.

In fact, one will not loose much by forgetting about logic and simply thinking of an μ-AEC as a concrete category $(\mathcal{K}, U)$ with concrete μ-directed colimits and where all morphisms are concrete embeddings.

Examples of AECs can be found in [BV17, §3.2], and some more examples of ∞-AECs can be found in [BGL$^+$16, §2]. Some other recently studied examples of interests include the category of flat modules with flat monomorphisms (see Example 5.13(5)), a special case of AECs on class of modules of the form ${}^{\perp}N$ [BET07] and an example of what module theorists call a *Kaplansky classes*, see e.g. [ŠT12]. For now, we note the following general fact:

Example 2.15. The class of models of a basic $\mathbb{L}_{\infty,\lambda}$-theory is a λ-AEC, when ordered with substructure (see Example 2.11). In fact, given any basic $\mathbb{L}_{\infty,\lambda}$-theory T, if the vocabulary contains a binary relation R and T contains the two basic sentences

- $(\forall x \forall y)((R(x,y) \wedge x = y) \rightarrow \bot)$
- $(\forall x \forall y)(x = y \vee R(x,y))$

then all the morphisms of $\mathbf{Mod}(T)$ will be monomorphisms (the two sentences above say that R should be interpreted as the non-equality relation in any model of T), and the category $\mathbf{Mod}(T)$ will essentially be a λ-AEC.

Closing the loop, we will see later that *any* accessible category where all morphisms are monomorphisms is equivalent to an ∞-AEC (Theorem 3.18). This also follows from the already mentioned equivalence between accessible categories and categories of models of $\mathbb{L}_{\infty,\infty}$ sentences.

Two remarks are in order. First, the reader might wonder about the hypothesis that all morphisms are monomorphisms. There is a generalization of the notion of an AEC, called an *abstract elementary category* which removes this restriction [BR12, 5.3]. In fact, let us define a *μ-abstract elementary category* $\mathbf{K}$ to be a subcategory of $\mathbf{Str}(\tau)$, $\tau = \tau(\mathbf{K})$ a μ-ary vocabulary, satisfying all the properties listed in Remark 2.14, except that the morphisms are only required to be initial (in the sense of [AHS04, 8.6]), but no longer required to be concrete monomorphisms. We will not discuss μ-abstract elementary categories further after this section.

Second, and more importantly, the correspondences between $\mathbb{L}_{\infty,\infty}$, accessible categories, and ∞-AECs do not hold cardinal by cardinal: as we have seen in Remark 2.14 it is *not* true that any λ-AEC is λ-accessible (and it is similarly not true that a category of models of a basic $\mathbb{L}_{\infty,\lambda}$ theory is λ-accessible). It *is* however true that a λ-accessible category is a λ-abstract elementary category as well as a category of models of a basic $\mathbb{L}_{\infty,\lambda}$ theory. In fact, we have, for a fixed regular λ, the following hierarchy, where each level has more categories than the next:

(1) λ-accessible categories.
(2) Categories of models of a basic $\mathbb{L}_{\infty,\lambda}$ theory.
(3) λ-abstract elementary categories.
(4) Accessible categories with λ-directed colimits.
(5) ∞-abstract elementary categories = accessible categories = categories of models of a basic $\mathbb{L}_{\infty,\infty}$ sentence.

This is relevant, especially because of the big differences between an $\aleph_1$-AEC and an ($\aleph_0$-)AEC: a lot more can be done in the latter setup (see Section 6).

Example 2.16. Consider the category $\mathbf{Ban}_{mono}$ of Banach spaces with isometries. This is an $\aleph_1$-accessible category with all directed colimits and all morphisms monos. However, these directed colimits are *not* concrete (they are not unions, but completions of unions), so $\mathbf{Ban}_{mono}$ does not seem to obviously be an AEC. One can show [LRVa] that in fact *no* faithful functor from $\mathbf{Ban}_{mono}$ into $\mathbf{Set}$ preserves directed colimits. Thus $\mathbf{Ban}_{mono}$ is indeed *not* equivalent to an AEC. It will however be an $\aleph_1$-AEC.

3. Fundamental results

We start developing the theory of accessible categories from scratch, proving some basic results and ending by sketching the equivalence of the three frameworks described above. Most of the results of this section are well known and appear in [AR94] or [BGL$^+$16].

First, we recall some more category-theoretic terminology. Monomorphisms are in a sense a very weak generalization of the notion of an injection[9]. A much stronger one is given by the following definition: suppose we have $A \xrightarrow{i} B \xrightarrow{r} A$ which compose to the identity ($ri = \mathrm{id}_A$). Then we call i a *section* (or *split monomorphism*) and r a *retraction* (or *split epimorphism*). In such a situation, we say that A is a *retract* of B. It is easily checked that sections and retractions are indeed monomorphisms and epimorphisms respectively. The canonical inclusions $A \to A \oplus B$ and projections $A \oplus B \to A$ in the category of R-modules are good examples of sections and retractions. In case we also have that $ir = \mathrm{id}_B$, then we write $r = i^{-1}$, $i = r^{-1}$, and call them *isomorphisms*. Note that a retraction which is also a mono is an isomorphism [AHS04, 7.36]. Thus when all morphisms are monos (e.g. in an ∞-AEC), any retraction (and thus any section) is an isomorphism.

Our first goal will be to show that any object in an accessible category is presentable. This follows from the following result, which essentially says that a small union (colimit) of small objects is small (a diagram is called λ-*small* if its indexing category has strictly less than λ-many objects):

Theorem 3.1. For λ a regular cardinal, a colimit of a λ-small diagram of λ-presentable objects is λ-presentable. In particular all the objects of an accessible category are presentable.

Proof sketch. Let $D : I \to \mathcal{K}$ be a λ-small diagram consisting of λ-presentable objects, with colimit $(D_i \xrightarrow{d_i} A)_{i \in I}$. Let B be a λ-directed colimit of another diagram $E : J \to \mathcal{K}$, and let $A \xrightarrow{f} B$. For each $i \in I$, the map fd_i factors through some E_{j_i}, $j_i \in J$, by λ-presentability of D_i. Since J is λ-directed and $|I| < \lambda$, there exists $j \in J$ so that $j \geq j_i$ for all $i \in I$. Then the universal property of the colimit implies that f must factor through E_j. The "in particular" part follows from the rest because in a λ-accessible category, any object is (by the smallness condition) a colimit of a (set-sized) diagram consisting of λ-presentable objects. □

We now work toward proving that an accessible category will, for each λ, have only a set (up to isomorphism) of λ-presentable objects. We start with the following technical observations:

Remark 3.2. Work in a category $\mathcal{K}$.

(1) If $A_1 \xrightarrow{i_1} B \xrightarrow{r_1} A_1$ and $A_2 \xrightarrow{i_2} B \xrightarrow{r_2} A_2$ are such that $r_\ell i_\ell = \mathrm{id}_{A_\ell}$ (i.e. they are section/retraction pairs) and $i_1 r_1 = i_2 r_2$, then A_1 and A_2 are isomorphic (as witnessed by $r_2 i_1$ and $r_1 i_2$). Since B has only a set of

[9]For example, in the category of divisible abelian groups, the quotient map $\mathbb{Q} \to \mathbb{Q}/\mathbb{Z}$ turns out to be a monomorphism [AHS04, 7.33(5)].

endomorphisms (i.e. morphisms from and to itself), there is only a set (up to isomorphism) of retracts of any given object.
(2) A retract of a λ-presentable object is λ-presentable (straightforward diagram chase from the definition of λ-presentability).
(3) If a λ-presentable object A is a λ-directed colimit of a diagram $D : I \to \mathcal{K}$, with colimit cocone $(D_i \xrightarrow{d_i} A)_{i \in I}$, then d_i is a retraction for some $i \in I$, so A is a retract of D_i (by λ-presentability, the identity map on A must factor through one of the components of the diagram: $\text{id}_A = d_i f$ for some $i \in I$).

We obtain:

Lemma 3.3. A λ-accessible category has, up to isomorphism, only a set of λ-presentable objects.

Proof. Let S be the set of λ-presentable objects given by the smallness condition. By Remark 3.2, any λ-presentable objects will be a retract of members of S, and there is only a set of such retracts. □

In order to say more, we try to understand when a λ-accessible will be μ-accessible for $\mu > \lambda$. This is *not* a trivial consequence of the definition because a λ-directed poset may not be μ-directed. In fact, as we will see, it is *not* true in general that a λ-accessible category is μ-accessible for all $\mu > \lambda$ (it turns out that the *accessibility spectrum* – the class of cardinals λ such that a category is λ-accessible – is an interesting measure of the complexity of the category, see Section 4.1). Before looking at counterexamples, let us state a positive result. For a regular cardinal μ, an infinite cardinal λ is called *μ-closed*[10] if $\theta^{<\mu} < \lambda$ for all $\theta < \lambda$. Note that any uncountable cardinal is $\aleph_0$-closed and in general for any fixed μ there is a proper class of regular μ-closed cardinal (given any infinite cardinal λ_0, the cardinal $\left(\lambda_0^{<\mu}\right)^+$ is always μ-closed).

Theorem 3.4 (Raising the index of accessibility)**.** Let $\mu \leq \theta \leq \lambda$ be regular cardinals and let $\mathcal{K}$ be a θ-accessible category with μ-directed colimits. If λ is μ-closed, then $\mathcal{K}$ is λ-accessible.

Proof. Given an object A of $\mathcal{K}$, we have to write A as a λ-directed colimit of λ-presentable objects. First, by θ-accessibility we know we can write A as a θ-directed colimit of a diagram $D : I \to \mathcal{K}$ of θ-presentable objects. Now let J be the poset of all μ-directed subsets of I of cardinality strictly less than λ, ordered by containment. For each $I_0 \in J$, we can use that $\mathcal{K}$ has μ-directed colimits to look at the colimit $\text{colim}(D \upharpoonright I_0)$ of the diagram D restricted to I_0. This process induces a new diagram $E : J \to \mathcal{K}$, where $E_{I_0} = \text{colim}(D \upharpoonright I_0)$. Notice that by Theorem 3.1, E consists of λ-presentable objects. Further, because λ is μ-closed, any subset of I of cardinality strictly less than λ will be contained in some member of J. In particular, J is λ-directed and $\text{colim}\, E = \text{colim}\, D = A$. Thus any object is a λ-directed colimits of λ-presentable objects. The argument also shows that each of these λ-presentable object is a colimit of μ-presentable objects indexed by

[10]When λ is regular, this is written $\mu \ll \lambda$ in [Lur09, A.2.6.3]. The notation can be a misleading though, because it is *not* true that $\mu \ll \lambda < \lambda'$ implies $\mu \ll \lambda'$ (take for example $\mu = \aleph_1$, $\lambda = \aleph_2$, $\lambda' = \aleph_{\omega+1}$).

a poset of cardinality strictly less than λ. By Lemma 3.3, there is only a set of μ-presentable objects, hence only a set of such diagrams up to isomorphism, so there is only a set of λ-presentable objects. □

Corollary 3.5. Any accessible category is λ-accessible for a proper class of cardinals λ. Moreover, any θ-accessible category with directed colimits (in particular any finitely accessible or locally θ-presentable category) is λ-accessible for *all* regular cardinals $\lambda > \theta$.

Corollary 3.6. For any regular cardinal λ, an accessible category has only a set, up to isomorphism, of λ-presentable objects.

Proof. Let $\mathcal{K}$ be an accessible category and fix $\theta \geq \lambda$ regular such that $\mathcal{K}$ is θ-accessible. By Lemma 3.3, $\mathcal{K}$ has only a set of θ-presentable objects, hence (because λ-presentable implies θ-presentable) a set of λ-presentable objects. □

Remark 3.7. In the proof of Theorem 3.4, the hypothesis that λ is μ-closed was used to show that for any μ-directed poset I, any subset of I of cardinality strictly less than λ can be completed to a *μ-directed* subset of cardinality strictly less than λ. It more generally suffices to assume that for any $\theta < \lambda$, $\mathrm{cf}([\theta]^{<\mu}) < \lambda$ (recall that $[\theta]^{<\mu}$ is the set of all subsets of θ of cardinality strictly less than μ, ordered by containment). Following [MP89, 2.3.1], we will write $\mu \lhd \lambda$ when this holds (one can check the relation $\lhd$ indeed gives a partial order on the regular cardinals). Since we always have that $\lambda \lhd \lambda^+$, it follows, for example, that any λ-accessible category is also λ^+-accessible [AR94, 2.13(2)]. However, when $\lambda > 2^{<\mu}$, λ is μ-closed if and only if $\mu \lhd \lambda$ (because of the equation $\lambda^{<\mu} = 2^{<\mu} \cdot \mathrm{cf}([\lambda]^{<\mu})$, well known to set theorists, see [LRV19b, 2.5]). Below $2^{<\mu}$, the behavior of $\mathrm{cf}([\lambda]^{<\mu})$ can be somewhat understood through the lens of Shelah's PCF theory [She94, AM10].

Example 3.8 ([AR94, 2.11])**.** For μ a regular uncountable cardinal, let $\mathcal{K}$ be the category of μ-directed posets, with morphisms the order-preserving maps. One can check that $\mathcal{K}$ is μ-accessible. Let $\lambda > \mu$ be a regular cardinal such that there exists a $\theta < \lambda$ with $\mathrm{cf}([\theta]^{<\mu}) \geq \lambda$ (so θ witnesses that $\mu \not\lhd \lambda$; take for example $\theta = \beth_\omega(\mu)$, $\lambda = \theta^+$, see discussion above). Then $\mathcal{K}$ is *not* λ-accessible because the poset $[\theta]^{<\mu}$ is μ-directed, hence an object of $\mathcal{K}$, but cannot be written as a λ-directed colimit of λ-presentable objects. Indeed, suppose it can, and let $D : I \to \mathcal{K}$ be the corresponding diagram. The images of the colimit maps form a collection $(Y_i)_{i \in I}$ of λ-presentable subsets of $[\theta]^{<\mu}$, and the λ-presentable objects are those of cardinality strictly less than λ, so each Y_i has cardinality strictly less than λ. Each $\{\alpha\}$, for $\alpha < \theta$, lies in some Y_i, and because I is λ-directed, there must exist $i^* \in I$ such that each $\{\alpha\}$ lies in Y_{i^*}. Because Y_{i^*} is μ-directed, it must in fact be cofinal in $[\theta]^{<\mu}$, but we know that Y_{i^*} has cardinality strictly less than λ, contradiction.

Different types of categories have different appropriate notions of functors. For accessible categories, an accessible functors play an important role:

Definition 3.9. A functor $F : \mathcal{K} \to \mathcal{L}$ is *λ-accessible* if both $\mathcal{K}$ and $\mathcal{L}$ are λ-accessible categories and F preserves λ-directed colimits. We say that F is *accessible* if it is λ-accessible for some λ.

One reason accessible functors are important is the *adjoint functor theorem*. In general category theory, Freyd's adjoint functor theorem [AHS04, 18.12] tells us

that a functor between complete categories is (left) adjoint if and only if it preserves limits and satisfies a technical "solution set condition" that can be quite difficult to check. The statement simplifies when looking at accessible categories (recall that a bicomplete accessible category is just a locally presentable category). We will not look into this direction much further, so we omit the proof.

Theorem 3.10 (The adjoint functor theorem, [AR94, 1.66])**.** A functor between two locally presentable categories is adjoint if and only if it preserves limits and is accessible.

Remark 3.11. It is also true that any left or right adjoint functor between accessible categories is accessible [AR94, 2.23].

A question we will be interested in is how a given functor interacts with sizes: for a regular cardinal λ, we say that a functor F *preserve λ-presentable objects* if FA is λ-presentable whenever A is λ-presentable. Accessible functors are useful because they preserve certain sizes:

Theorem 3.12 (The uniformization theorem; [AR94, 2.19])**.** For any accessible functor F, there exists a proper class of regular cardinals λ such that F is λ-accessible and preserves λ-presentable objects.

Proof. Let $F : \mathcal{K} \to \mathcal{L}$ be a μ-accessible functor. Up to isomorphism, there is only a set of μ-presentable objects in $\mathcal{K}$, so there exists a regular cardinal $\mu' \geq \mu$ such that FA is μ'-presentable for every μ-presentable objects A. Now let $\lambda \geq \mu'$ be a μ-closed cardinal. Let A be a λ-presentable object in $\mathcal{K}$. By the proof of Theorem 3.4, A can be written as a λ-directed colimit of objects that are each λ-small μ-directed colimits of μ-presentable objects. Since A is λ-presentable, it must be a retract of an object of this diagram (Remark 3.2(3)): hence A is a retract of a λ-small μ-directed colimits of μ-presentable objects. Since by hypothesis F preserves such colimits and any functor preserves retractions, FA is a retract of a λ-small μ-directed colimit of μ'-presentable objects in $\mathcal{L}$. By Theorem 3.1, FA is λ-presentable. □

Remark 3.13. The proof gives more: if $\mu \leq \lambda_0 \leq \lambda_1 \leq \lambda$ are all regular, $F : \mathcal{K} \to \mathcal{L}$, $\mathcal{K}$ and $\mathcal{L}$ are both λ_0-accessible, F preserves μ-directed colimits, λ is μ-closed, and FA is λ_1-presentable whenever A is λ_0-presentable, then F preserves λ-presentable objects. In particular, an accessible functor preserving directed colimits will preserve λ-presentable objects for all high-enough regular λ.

Dually, it is natural to ask when a functor *reflects λ-presentable objects*, i.e. when FA λ-presentable implies A λ-presentable. A sufficient condition is for F to reflect split epimorphisms (if Ff is a split epi – i.e. a retraction – then f is a split epi), see [BR12, 3.6].

Theorem 3.14. If $F : \mathcal{K} \to \mathcal{L}$ is a λ-accessible functor reflecting split epimorphisms, then F reflects λ-presentable objects.

Proof. Assume that FA is λ-presentable. Since $\mathcal{K}$ is λ-accessible, A is a the colimit of a λ-directed diagram $D : I \to \mathcal{K}$ consisting of λ-presentable objects, with colimit cocone $(D_i \xrightarrow{d_i} A)_{i \in I}$. Since F preserves λ-directed colimits, FA is the colimit of

FD, so as FA is λ-presentable, Remark 3.2(3) implies Fd_i is a retraction for some $i \in I$. Because F reflects split epis, d_i is a retraction, so A is a retract of the λ-presentable object D_i, hence is λ-presentable (Remark 3.2(2)). □

Remark 3.15. A functor that reflects isomorphism and whose image contains only monomorphisms will automatically reflect split epimorphisms.

Example 3.16.

(1) Let $\mathcal{K}$ be a locally presentable category. The binary product functor $F : \mathcal{K} \to \mathcal{K}$ sending A to the category-theoretic product $A \times A$ is accessible, because it is adjoint to the diagonal functor. Thus by the uniformization theorem, F preserves λ-presentable objects for a proper class of λ. That is, F does not make the product "too much bigger".
(2) [BR12, 3.2(4)] For μ a regular cardinal, let $F : \mathbf{Set} \to \mathbf{Set}$ send the set X to the set $[X]^{<\mu}$ of subsets of X of cardinality strictly less than μ. This is an accessible functor but, if μ is uncountable and $\lambda < \lambda^{<\mu}$, F will not preserve λ^+-presentable objects.
(3) [BR12, 3.3] Let $F : \mathbf{Grp} \to \mathbf{Ab}$ be the abelianization functor from the category of group to the category of abelian groups. This is a right adjoint functor, so it preserves colimits, and hence by Remark 3.13 preserves λ-presentable objects for all regular cardinals λ. However, if G is a simple group then its abelianization $F(G)$ is the zero group. Thus (since there exists simple groups in all infinite cardinalities) F can make sizes drop. Indeed, one can check that F does not reflect split epimorphisms.
(4) If $\mathcal{K}$ is locally presentable and $U : \mathcal{K} \to \mathbf{Set}$ is an accessible functor preserving limits, then by the adjoint functor theorem, U is left adjoint. If we think of U as a forgetful functor, the right adjoint will be the free functor. Since $\mathcal{K}$ has directed colimits, U will preserve λ-presentable objects for all high-enough regular cardinals λ.
(5) Let $\mathbf{K}$ be a μ-AEC, and let $U : \mathbf{K} \to \mathbf{Set}$ be the forgetful universe functor. By definition of a μ-AEC, U preserves μ-directed colimits, hence is accessible. The abstract class axiom ensures that U reflects isomorphisms. Moreover, any morphism in the image of U must be a monomorphism so by Remark 3.15, U reflects split epis. Finally, it is easy to check that any $\mathrm{LS}(\mathbf{K})^+$-presentable object in $\mathbf{K}$ will have cardinality at most $\mathrm{LS}(\mathbf{K})$ (write the object as the $\mathrm{LS}(\mathbf{K})^+$-directed colimit of its subobject of cardinality at most $\mathrm{LS}(\mathbf{K})$). By the uniformization theorem, if $\lambda > \mathrm{LS}(\mathbf{K})$ is a regular μ-closed cardinal, then F preserves and reflects λ-presentable objects. This means that A is λ-presentable in $\mathbf{K}$ if and only if it its universe has cardinality strictly less than λ. In particular (taking $\mu = \aleph_0$) in an AEC, category-theoretic sizes correspond exactly to cardinalities (above $\mathrm{LS}(\mathbf{K})$). See [LRV19b, §4] for more on such results.
(6) In the $\aleph_1$-AEC $\mathbf{K}$ of Banach spaces (with subspace inclusions), the universe functor U does *not* preserve $\aleph_1$-presentable objects: an $\aleph_1$-presentable Banach space will not have countable cardinality. This is because U does not preserve $\aleph_0$-directed colimits (even though $\mathbf{K}$ *does* have those colimits, they are not concrete: one *cannot* compute them by taking unions).

Let's use the uniformization theorem to better understand the relationship between the smallness axiom in ∞-AEC (Definition 2.12) and the smallness condition in the definition of an accessible category (Definition 2.7).

Theorem 3.17 ([BR12, 5.5],[LRV19a, 7.2])**.** Assume that $\mathbf{K}$ satisfies all the axioms of a μ-AEC (Definition 2.12), except perhaps for the LST smallness axiom. The following are equivalent:

(1) $\mathbf{K}$ is accessible.
(2) $\mathbf{K}$ satisfies the LST smallness axiom, and hence is a μ-AEC.

Proof. We have seen already (Remark 2.14) that (2) implies (1). Assume now that (1) holds: $\mathbf{K}$ is accessible. Then the universe functor $U : \mathbf{K} \to \mathbf{Set}$ is accessible (preservation of μ-directed colimits is just because they are computed the same way in $\mathbf{K}$ and $\mathbf{Set}$ – this is what the chain axiom says). By the uniformization theorem, we can pick a regular cardinal $\theta \geq \mu + |\tau(\mathbf{K})|$ such that U is θ-accessible and preserves θ-presentable objects. We will show that $\text{LS}(\mathbf{K}) \leq \theta^{<\mu}$. Let $M \in \mathbf{K}$, and let $A \subseteq UM$. Set $\lambda := \left((\theta + |A|)^{<\mu}\right)^{+}$. By Remark 3.13, U is λ-accessible and preserves λ-presentable objects. In particular, $\mathbf{K}$ is λ-accessible so we can write M as a λ-directed colimit (union) of λ-presentable objects: $M = \bigcup_{i \in I} M_i$. Now $|A| < \lambda$, so by λ-directedness there must exist $i \in I$ such that $A \subseteq UM_i$. Since M_i is λ-presentable in $\mathbf{K}$, $UM_i = M_i$ must be λ-presentable in $\mathbf{Set}$, hence it must have cardinality strictly less than λ (see Example 2.5(1)), as desired. □

We can now sketch part of the proof that any accessible category can be presented as the class of models of an $\mathbb{L}_{\infty,\infty}$ sentence. Let us make such a statement precise first: recall that two categories $\mathcal{K}$ and $\mathcal{L}$ are *equivalent* if there exists a functor $F : \mathcal{K} \to \mathcal{L}$ that is full (surjective on morphisms), faithful (injective on morphisms), and essentially surjective on objects (any object of $\mathcal{L}$ is isomorphic to an object in the image of F). Essential surjectivity (rather than surjectivity) on objects makes equivalence of categories weaker than *isomorphism* of categories, but it is typically the former notion of "being the same" that is used in category theory[11]. For $\infty \geq \lambda \geq \mu$, let us then define a category to be (λ, μ)-*elementary* if it is equivalent to a category of the form $\mathbf{Mod}(T)$, for T a theory in $\mathbb{L}_{\lambda,\mu}$. We will see that *any* λ-accessible category is (∞, λ)-elementary. The proof proceeds in two steps. First, the category is embedded into a category of structures, and second this category of structures is axiomatized. We will only look at the first step (which is sufficient if one only cares about λ-AECs). The main idea is to represent each object using its Hom functor. The second step will be proven in the very special case of universal classes.

Theorem 3.18 ([BGL$^+$16, 4.8])**.** If $\mathcal{K}$ is a λ-accessible category, then there is a (finitary) vocabulary τ and an embedding[12] $E : \mathcal{K} \to \mathbf{Str}(\tau)$ which is full and preserves λ-directed colimits. In particular, if in addition all the morphisms of $\mathcal{K}$ are monomorphisms, $\mathcal{K}$ is equivalent to a λ-AEC.

[11]For example, the categories of finite dimensional vector spaces over $\mathbb{R}$ and of real matrices (where the objects are natural numbers and the morphisms matrices with the right dimension, with composition defined by matrix multiplication) are equivalent but not isomorphic, see [AHS04, 3.35(2)].

[12]That is, a functor which is faithful and injective on objects.

Proof sketch. Let $\mathcal{K}_\lambda$ be a small full subcategory of $\mathcal{K}$ containing representatives of each λ-presentable object. For each $M \in \mathcal{K}_\lambda$, let S_M be a unary relation symbol and for each morphism f in $\mathcal{K}_\lambda$, let $\underline{f}$ be a binary function symbol. The vocabulary τ will consist of all such S_M and $\underline{f}$. Now map each $M \in \mathcal{K}$ to the following τ-structure EM:

(1) Its universe are the morphisms $g : M_0 \to M$, where $M_0 \in \mathcal{K}_\lambda$.
(2) For each $M_0 \in \mathcal{K}_\lambda$, $S^{EM}_{M_0}$ is the set of morphisms $g : M_0 \to M$.
(3) For each morphism $f : M_0 \to M_1$ of $\mathcal{K}_\lambda$, and each $g : M_1 \to M$, $\underline{f}^{EM}(g) = gf$. When $g \notin S^{EM}_{M_1}$, just let $\underline{f}^{EM}(g) = g$.

Map each morphism $f : M \to N$ to the function $\bar{f} : EM \to EN$ given by $\bar{f}(g) = fg$. Essentially, E is the functor $E(M) = \mathrm{Hom}(-, M)$ from $\mathcal{K}$ to ${}^{\mathcal{K}^{op}_\lambda}\mathbf{Set}$. The Yoneda embedding lemma tells us that $E \restriction \mathcal{K}_\lambda$ is full and faithful [AHS04, 6.20]. The definition of a λ-presentable object also ensures that E preserves λ-directed colimits. By writing any object as a λ-directed colimit of λ-presentable objects, we get that E is full and faithful. To see the "in particular" part, consider the smallest isomorphism-closed subcategory $\mathcal{L}$ of $\mathbf{Str}(\tau)$ that contains $E[\mathcal{K}]$. This category is equivalent to $\mathcal{L}$, satisfies the abstract class axiom, has concrete λ-directed colimits, and (trivially, because the morphisms are homomorphisms) satisfies the coherence axiom. Since $\mathcal{K}$ is accessible and E is full and faithful, $E[\mathcal{K}]$ also is accessible, and hence $\mathcal{L}$ is accessible. Now apply Theorem 3.17. □

Corollary 3.19 ([AR94, 5.35])**.** Any λ-accessible category is (∞, λ)-elementary. In particular, a category is accessible if and only if it is (∞, ∞)-elementary.

Proof. We first apply Theorem 3.18 to reduce the problem to axiomatizing a class of structures, and then axiomatize this class (we will not explain how here). □

Example 3.20.

(1) For λ a regular cardinal, we have seen that the AEC of all sets of cardinality at least λ is not λ-accessible. In fact, this AEC is not even (∞, λ)-elementary but the proof is not trivial, see [Hen19].
(2) [BV19, §4] Consider the following AEC, $\mathbf{K}$, sometimes called the *toy quasi-minimal class*: the vocabulary contains a single binary relation. The objects of $\mathbf{K}$ are the equivalence relations with countably infinite classes. The ordering says that equivalence classes do not grow. It is not difficult to see that $\mathbf{K}$ is finitely accessible: an object is λ-presentable if and only if the relation has fewer than λ-many classes, and every equivalence relation is the directed colimits of its restrictions to finitely-many classes. By Corollary 3.19, $\mathbf{K}$ is (∞, ω)-elementary, hence is equivalent to the category of models of an $\mathbb{L}_{\infty,\omega}$-theory. This latter category is obtained, roughly speaking, by collapsing each equivalence class to a point. However it is well known that $\mathbf{K}$ itself is *not* the category of models of a basic $\mathbb{L}_{\infty,\omega}$ theory. In fact, it is not *finitary* in the sense of Hyttinen-Kesälä [HK06]: roughly, one *cannot* figure out whether a map is a morphism by checking finitely-many points at a time (like we would be able to do for homomorphisms in a finitary vocabulary). This shows that the concept of being a finitary AEC is *not* invariant under equivalence of categories.

(3) An ∞-AEC $\mathbf{K}$ has *intersections* if for any $M \in \mathbf{K}$, and any non-empty collection $\{M_i : i \in I\}$ of $\leq_{\mathbf{K}}$-substructures of M, $\bigcap_{i \in I} M_i$ induces a $\leq_{\mathbf{K}}$-substructure of M. Generalizing the previous example, one can show that any λ-AEC with intersections is a λ-accessible category [LRV19c, 3.3], hence (∞, λ)-elementary. In fact, as the definition makes apparent, λ-AECs with intersections correspond exactly to the λ-accessible categories with wide pullbacks (and all morphisms monos), see [LRV19c, 5.7]. This gives a clear sense in which this class of AECs is natural, and less complex than general AECs. See [Vas17c, §2], [BV19], or [LRV19b, §5] for more on AECs with intersections.

As a consolation prize for not proving the axiomatizability part of Corollary 3.19, let us prove it for a simpler framework: that of universal classes.

Definition 3.21. For μ a regular cardinal, a *μ-universal class* is a μ-AEC $\mathbf{K}$ such that $\leq_{\mathbf{K}}$ is just the $\tau(\mathbf{K})$-substructure relation, and moreover if $M \in \mathbf{K}$ and $M_0 \subseteq M$, then $M_0 \in \mathbf{K}$. Said another way, a μ-universal class is simply a class of structures in a fixed μ-ary vocabulary that is closed under isomorphisms, substructures, and μ-directed unions (we identify the class with the μ-AEC). When $\mu = \aleph_0$, we omit it.

Theorem 3.22 (Tarski's presentation theorem; [Tar54], [LRV19c, 2.2])**.** Let μ be a regular cardinal and let K be a class of structures in a fixed μ-ary vocabulary τ. The following are equivalent:

(1) K is the class of models of a universal $\mathbb{L}_{\infty,\mu}$ theory (i.e. a theory where each sentence is of the form $(\forall \bar{x})\psi$, with ψ quantifier-free).
(2) K is a μ-universal class.

Proof sketch. The implication (1) implies (2) is easy to check. Assume now that (2) holds. Call $M \in K$ *μ-generated* if $M = \text{cl}^M(A)$, for some $A \in [UM]^{<\mu}$ (here, cl^M denotes the closure of A under the functions of M – note that such a closure is always a substructure of M, hence in K by definition of a μ-universal class). Note that μ-generated is equivalent to μ-presentable [LRV19b, 5.7], but this will not be needed. Now by listing all the isomorphism types of the μ-generated models in K, form a quantifier-free formula $\phi(\bar{x})$ of $\mathbb{L}_{\infty,\mu}$ such that $N \models \phi[\bar{a}]$ if and only if $\text{cl}^N(\text{ran}(\bar{a}))$ is in K. The formula $\psi := (\forall \bar{x})\phi(\bar{x})$ axiomatizes K. Indeed, if $N \in K$ then $N \models \psi$. Conversely, if $N \models \psi$ then $\{\text{cl}^N(\text{ran}(\bar{a})) \mid \bar{a} \in {}^{<\mu}UN\}$ is a μ-directed system in K whose union is N, hence N is in K. This proves (1). □

Remark 3.23. There is also a category-theoretic characterization of μ-universal classes: they are the μ-accessible categories with all morphisms monos and all connected limits [LRV19c, 5.9].

Example 3.24.

(1) The class of all vector spaces over a fixed field F is a universal class.
(2) The class of locally finite groups is a universal class (it is a good exercise to try to write down the universal sentence of 3.22(1) for this case).

(3) If T is an $\mathbb{L}_{\infty,\infty}$ theory, we can "Skolemize it" (add functions to pick witnesses of existential sentences) to get a universal class in an expanded vocabulary, whose restriction to the original vocabulary is the class of models of T.
(4) Any μ-universal class is a μ-AEC with intersections (see Example 3.20(3)).
(5) The AEC of algebraically closed fields (with subfield) is not universal: the rationals form a non-algebraically closed subfield of the complex numbers. This AEC does have intersections, however.
(6) The class of all linear orders, or the class of all graphs, is a universal class.

Universal classes are in a sense the simplest type of AECs. In fact, a key result is Shelah's presentation theorem [She09a, I.1.9], which says that, given an AEC $\mathbf{K}$, we can find a universal class[13] $\mathbf{K}^*$ in an expanded vocabulary whose reduct (i.e. restriction to the original vocabulary) to $\mathbf{K}$ is $\mathbf{K}$. In fact, the reduct gives a functor from $\mathbf{K}^*$ to $\mathbf{K}$ which is faithful (injective on morphisms), surjective, and preserves directed colimits. This is the motivation for the following generalization of Shelah's presentation theorem to accessible categories with all morphisms monos.

Theorem 3.25 (Shelah's presentation theorem, categorical version [LRVb, 4.4])**.** If $\mathcal{L}$ is an accessible category with μ-directed colimits and all morphisms monos, then there exists a μ-universal class $\mathbf{K}$ and a faithful essentially surjective functor $H : \mathbf{K} \to \mathcal{L}$ preserving μ-directed colimits.

Proof sketch. First, let $\mathcal{L}^*$ be the category obtained by taking free μ-directed colimits of the μ-presentable objects of $\mathcal{L}$. This is a μ-accessible category [AR94, 2.26], and the natural "projection" functor $F : \mathcal{L}^* \to \mathcal{L}$ is essentially surjective and preserves directed colimits. Moreover, F is also faithful: given two distinct morphisms $f, g : A \to B$ in $\mathcal{L}^*$, separate them on λ-presentable objects by finding $f_0, g_0 : A_0 \to B_0$ distinct with A_0 and B_0 λ-presentable and maps $u : A_0 \to A$, $v : B_0 \to B$ such that $fu = vf_0$, $gu = vg_0$:

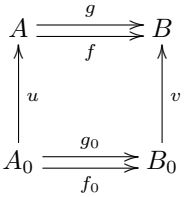

Since F is faithful on λ-presentable objects, $Ff_0 \neq Fg_0$. Moreover, Fv is a mono, so it follows that $Ff \neq Fg$. By Corollary 3.19, $\mathcal{L}^*$ is (∞, μ)-elementary, hence without loss of generality is the category of models of a basic $\mathbb{L}_{\infty,\mu}$ theory T. Skolemizing T (see Example 3.24(3)), one obtains a μ-universal class $\mathbf{K}$ so that the reduct functor $G : \mathbf{K} \to \mathcal{L}^*$ is faithful, essentially surjective, and preserves μ-directed colimits. Let $H := F \circ G$. □

Remark 3.26.

[13]Shelah stated his result in terms of a class of models omitting types, but the proof of Theorem 3.22 shows that any universal class is a type-omitting class: omit the types of μ-generated models outside of K.

(1) Any μ-AEC is an accessible category with μ-directed colimits and all morphisms monos (Remark 2.14), so when $\mu = \aleph_0$ we in particular recover a version of Shelah's original presentation theorem. Note that, when $\mu > \aleph_0$, it is not so easy to imitate Shelah's proof, see [LRVb, §3] for a discussion.
(2) In the case $\mu = \aleph_0$, a classical result of model theory, *Morley's omitting type theorem* tells us in particular (in categorical language, see [MP89, 3.4.1]) that for any (∞, ω)-elementary category $\mathcal{K}$ (hence by Tarski's presentation theorem in particular for a universal class), there is a faithful functor $\mathbf{Lin} \to \mathbf{K}$ preserving directed colimits (where **Lin** is the category of linear orders and order-preserving maps). The functor constructed by the model-theoretic process is called the *Ehrenfeucht-Mostowski (EM) functor*, and its images are called *Ehrenfeucht-Mostowski models*. They can be described as the models generated by order-indiscernibles satisfying a fixed collection of quantifier-free types (an *EM blueprint*). In fact, from any faithful functor $\mathbf{Lin} \to \mathcal{K}$ preserving directed colimits, one can recover such a blueprint [Bonb, 5.5]. Combining Morley's theorem with Tarski's presentation theorem and the categorical version of Shelah's presentation theorem, we obtain that for any large accessible category $\mathcal{L}$ with directed colimits and all morphisms monos there is a faithful functor $\mathbf{Lin} \to \mathcal{L}$ preserving directed colimits. Notice that this works not only for AECs but also for continuous classes (such as Banach spaces).

4. Set-theoretic topics

4.1. **Cardinality vs presentability.** Is the behavior of presentability, the notion of size defined in Definition 2.4, similar to the behavior of cardinality in concrete classes? There are at least two questions one can consider in this direction. First, in the category of sets, a set is λ-presentable if and only if it has cardinality strictly less than λ. In particular, the presentability rank of a set is a successor. This happens in all the examples listed in 2.5. Thus one can ask: *in an arbitrary accessible category, are high-enough presentability ranks always successors*?

Second, the Löwenheim-Skolem theorem for first-order logic says that any theory with an infinite model has models of all high-enough cardinalities. Is there a similar version for accessible categories? Since we do not have the compactness theorem, let us restrict ourselves to *large* categories, and let us consider only eventual behavior. Cardinality here is again not the right notion (consider Example 2.5(5): there are no Hilbert spaces in certain cardinalities), but we can still ask: *does every large accessible category have objects of each high-enough size?* (see Definition 2.6). Following [BR12, 2.4], let us call an accessible category $\mathcal{K}$ *LS-accessible* if there exists a cardinal θ such that $\mathcal{K}$ has an object of size λ for all cardinals $\lambda \geq \theta$. We are then asking whether every large accessible category is LS-accessible.

Both questions turn out to be connected to the *accessibility spectrum*: the set of regular cardinals λ such that a given category is λ-accessible. Regarding the first question, we have:

Theorem 4.1 ([LRVb, 5.3])**.** Let $\mathcal{K}$ be a category and let λ be a weakly inaccessible cardinal[14]. If $\mathcal{K}$ is μ-accessible for unboundedly-many regular $\mu < \lambda$, then no object of $\mathcal{K}$ has presentability rank λ.

Proof idea. We show more generally that every object of $\mathcal{K}$ is a λ-directed colimits of objects of presentability rank strictly less than λ. This will imply the result by Remark 3.2. Let S be the set of regular cardinals $\mu < \lambda$ such that $\mathcal{K}$ is μ-accessible. Given an object A, we know that for each $\mu \in S$, we can write A as a μ-directed colimit of a diagram $D^\mu : I^\mu \to \mathcal{K}$ consisting of μ-presentable objects. One can then put together the D^μ's to obtain a new λ-directed diagram with colimit A and consisting of objects of presentability rank strictly less than λ. □

Corollary 4.2.

(1) [BR12, 4.2] In an accessible category with directed colimits, all high-enough presentability ranks are successors.
(2) [LRV19b, 3.11] If the singular cardinal hypothesis (SCH) holds[15], then in every accessible category, all high-enough presentability ranks are successors.

Proof.

(1) By Theorem 3.4, $\mathcal{K}$ is μ-accessible for *all* high-enough regular μ, so the result is immediate from Theorem 4.1.
(2) Assume that $\mathcal{K}$ is μ-accessible, and let $\lambda > 2^{<\mu}$ be a weakly inaccessible cardinal. Let $\theta_0 < \lambda$ be infinite and let $\theta := \left(\theta_0^{<\mu}\right)^+$. By the SCH hypothesis, $\theta < \lambda$ (see [Jec03, 5.22]) and by Theorem 3.4, $\mathcal{K}$ is θ-accessible. We have shown that $\mathcal{K}$ is accessible in unboundedly-many regular cardinals below λ, hence by Theorem 4.1 $\mathcal{K}$ has no objects of presentability rank λ.

□

Remark 4.3. Since SCH holds above a strongly compact cardinal [Jec03, 20.8], we can replace the SCH assumption by a large cardinal axiom.

The rough idea here is that if an accessible category is of "low-enough" complexity (as measured by its accessibility spectrum), then automatically its presentability ranks will have good behavior. Often additional cardinal arithmetic hypotheses can lower the complexity of the accessibility spectrum, to the point that we can prove results about general accessible categories. Let us now give an example (without proof) of this behavior for the second question above, whether every large accessible category has objects of all high-enough sizes (see Definition 2.6):

Theorem 4.4 ([LRVb, 6.11])**.** Assume SCH. If $\mathcal{K}$ is a large accessible category with all morphisms monos, and λ is a high-enough successor cardinal such that $\mathcal{K}$ is λ-accessible, then $\mathcal{K}$ has an object of presentability rank λ.

Corollary 4.5. Assume SCH.

[14]A cardinal is *weakly inaccessible* if it is regular and limit.

[15]The singular cardinal hypothesis is the statement that $\lambda^{\mathrm{cf}(\lambda)} = \lambda^+ + 2^{\mathrm{cf}(\lambda)}$ for every infinite cardinal λ.

(1) Any large accessible category with directed colimits and all morphisms monos is LS-accessible.
(2) [LRVb, 7.12] Any large accessible category has objects in all sizes of high-enough cofinality.

Proof idea.

(1) Immediate from Theorems 3.4 and 4.4.
(2) This can be obtained from Theorem 4.4 by combining careful use of the proofs of the uniformization theorem and Theorem 3.14, together with the (nontrivial) fact that the inclusion $\mathcal{K}_{mono} \to \mathcal{K}$ is an accessible functor [LRV19a, 6.2].

□

Remark 4.6. Corollary 4.5(1) is due to Lieberman and Rosický, and can be proven without SCH [LR16, 2.7]: let $E : \mathbf{Lin} \to \mathcal{K}$ be faithful and preserving directed colimits (see Remark 3.26(2)). By the uniformization theorem, E preserves λ-presentable objects for all high-enough regular λ, and by Theorem 3.14, E also reflects them (faithful functors reflect epimorphisms, and epimorphisms in **Lin** are isomorphisms). Thus for λ a big-enough cardinal and I a linear order of cardinality λ, $E(I)$ will have size exactly λ.

4.2. **Large cardinals and images of accessible functors.** Consider the functor $F : \mathbf{Set} \to \mathbf{Ab}$ that associates to each set the free abelian group on that set. It is easily checked that F is an accessible functor but, as noticed before (Example 2.9(11)) the question of whether the image of F (i.e. the full subcategory of **Ab** consisting of free abelian groups) is accessible is set-theoretic. One can ask this question generally: when is the image of an accessible functor accessible? The problem of course lies in proving existence of sufficiently directed colimits for this image. For technical reasons, we will close the image under subobjects: in the example of the free abelian group functor, subgroup of free groups are free, so the image is already closed under subobjects the only challenge is to check that a sufficiently directed diagram consisting of free groups has a cocone.

More precisely, define the *powerful image* of an accessible functor $F : \mathcal{K} \to \mathcal{L}$ to be the smallest full subcategory P of $\mathcal{L}$ that contains $F[\mathcal{K}]$ and is closed under subobjects (i.e. if $A \to B$ is a monomorphism and $B \in P$, then $A \in P$). The question becomes: when is the powerful image of an accessible functor accessible? The following result is due to Makkai and Paré:

Theorem 4.7 ([MP89, 5.5.1])**.** If there is a proper class of strongly compact cardinals[16], then the powerful image of any accessible functor is accessible.

Proof idea. Fix a λ-accessible functor F with powerful image P. By the uniformization theorem, we can assume without loss of generality that F preserves λ-presentable objects. Let $\kappa > \lambda$ be strongly compact. We show that P has κ-directed colimits. Using ideas around Corollary 3.19, we can reduce the problem of finding a cocone to the consistency of a certain $\mathbb{L}_{\kappa,\kappa}$-theory. This theory can

[16]A regular cardinal κ is *strongly compact* if every theory in $\mathbb{L}_{\kappa,\kappa}$ with all its subsets of size strictly less than κ consistent (i.e. with a model) is consistent. See for example [Jec03, 20.2].

be shown to have all subsets of size strictly less than κ consistent hence, by the compactness theorem for $\mathbb{L}_{\kappa,\kappa}$, to be consistent. □

Remark 4.8. The large cardinal assumption can be slightly weakened to a proper class of *almost* strongly compact cardinals [BTR17] but this is best possible [BU17]: the powerful image of every accessible functor is accessible if and only if there is a proper class of almost strongly compact cardinals. However, weaker statements can be proven from weaker large cardinal axioms (e.g. measurable or weakly compacts), see [Lie, BL]. We will even mention a ZFC theorem about images of accessible functors (Theorem 4.11).

Questions about the image of accessible functors can be used to study various kinds of compactness. One example is tameness in AECs ([LR16, 5.2]). As a simpler example, we consider the following property:

Definition 4.9. An object A in a category is an *amalgamation base* if any span $C \leftarrow A \rightarrow B$ can be completed (not necessarily canonically) to a commutative square:

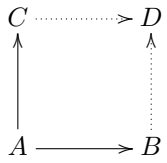

A category has the *amalgamation property* (or *has amalgamation*) if every object is an amalgamation base.

A question one might ask is whether amalgamation up to a certain level (e.g. for all λ-presentable objects for some big λ) implies amalgamation the rest of the way. Large cardinals imply a simple answer (earlier results used model-theoretic techniques, see e.g. [BB17]):

Theorem 4.10 ([LR17a, 3.6])**.** If $\mathcal{K}$ is a λ-accessible category, $\kappa > \lambda$ is strongly compact, and the full subcategory of $\mathcal{K}$ consisting of κ-presentable objects has amalgamation, then $\mathcal{K}$ has amalgamation.

Proof sketch. Let $\mathcal{K}^{\square}$ be the category of commutative squares in $\mathcal{K}$, and let $\mathcal{K}^{sp}$ be the category of spans $B \leftarrow A \rightarrow C$, with the morphisms in each category defined as expected. Consider the functor $F : \mathcal{K}^{\square} \rightarrow \mathcal{K}^{sp}$ that "forgets" the top corner of each square. One can check that this is a λ-accessible functor preserving λ-presentable objects, and moreover its image P is closed under subobjects, hence is equal to its powerful image. This image is, by definition of F, the category of spans that *can* be amalgamated. By Theorem 4.7, P is κ-accessible. Now any span S of $\mathcal{K}^{sp}$ is a κ-directed colimit of κ-presentable objects, and each of these objects can by assumption be amalgamated, hence are in P. Because P has κ-directed colimits, S is also in P. This shows that any span can be amalgamated, hence that $\mathcal{K}$ has amalgamation. □

As a final application, we mention how Shelah's singular compactness theorem [She75] can be restated as a theorem about image of accessible functors. One of the

most well known statement of the singular compactness theorem is that an abelian group of singular cardinality all of whose subgroups of lower cardinality are free is itself free. The proof can be axiomatized to apply to other kinds of objects than groups: modules, well-coloring in graphs, transversals, etc. In [BR16], Beke and Rosický state the following general form:

Theorem 4.11. Let $F : \mathcal{K} \to \mathcal{L}$ be an $\aleph_0$-accessible functor. Assume that F-structures extend along morphisms. Let $A \in \mathcal{L}$ be an object whose size is a singular cardinal. If all subobjects of A of lower size are in the image of F, then A itself is in the image of F.

Here, we say that *F-structure extend along morphisms* if for any $\mathcal{K}$-morphism $g : A \to B$, any object A' of $\mathcal{K}$, and any isomorphism $i : FA' \cong FA$, there exists $f : A' \to B'$ and an isomorphism $j : FB' \cong FB$ such that the following diagram commutes:

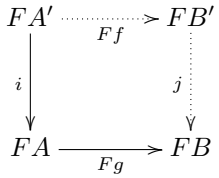

This can be thought of as a generalization of the Steinitz exchange property in vector spaces and fields.

Example 4.12. Let $F : \mathbf{Set}_{mono} \to \mathbf{Ab}$ be the restriction of the free abelian group functor to $\mathbf{Set}_{mono}$. This is an $\aleph_0$-accessible functor and F-structures extend along morphisms (we can rename to the case where g is the inclusion of A into a superset B; then if $i : FA' \cong FA$, we know that both free groups have the same number of generators, and one can add $|B \backslash A|$-many elements to A' to obtain a superset B' so that j extends i to an isomorphism of FB' with FB). Thus we recover Shelah's original application of the singular compactness theorem: if A is an abelian group of singular cardinality, all of whose subobjects of lower cardinality are free, then all these subobjects lie in the image of F, hence by Theorem 4.11 A must also be in this image, i.e. be free.

4.3. **Vopěnka's principle.** is a large cardinal axiom, whose consistency strength is between huge and extendible. A thorough introduction to the category-theoretic implications of Vopěnka's principle is in [AR94, §6]. Students of set theory may be familiar with Vopěnka's principle as the statement that in any proper class of structures in the same vocabulary, there exists an elementary embedding between two distinct members of the class (see [Jec03, p. 380]). Another logical characterization of Vopěnka's principle, due to Stavi, is that every logic has a Löwenheim-Skolem-Tarski number (see [MV11, Theorem 6]). A purely combinatorial characterization of Vopěnka's principle — and the one that Vopěnka first stated — is that there are no rigid proper classes of graphs. Stated in category-theoretic terms, no large full subcategory of the category of graphs is rigid, where a category is *rigid* if all of its morphisms are identities.

In fact, there is a more general category-theoretic formulation of Vopěnka's principle: any locally presentable category can be fully embedded into the category of graphs [AR94, 2.65]. Thus Vopěnka's principle is equivalent to the statement that no locally presentable category has a large rigid full subcategory. Even more strongly (because by Theorem 3.18 any accessible category can be fully embedded into $\mathbf{Str}(\tau)$, a locally presentable category), no large full subcategory of an *accessible* category can be rigid. Thus if C is a proper class of objects from an accessible category, Vopěnka's principle tells us there must be a morphism between two distinct objects of C (the first logical version mentioned above is the special case of the accessible category of τ-structures with elementary embeddings).

The following criteria makes it easy to check that a category is accessible [AR94, 6.9, 6.17]. It can be seen as a category-theoretic version of the fact that every logic has a Löwenheim-Skolem-Tarski number.

Theorem 4.13. Assume Vopěnka's principle. If $\mathcal{L}$ is a full subcategory of an accessible category $\mathcal{K}$ and $\mathcal{L}$ has μ-directed colimits for some μ, then the inclusion $\mathcal{L} \to \mathcal{K}$ preserves λ-directed colimits for some λ and $\mathcal{L}$ is accessible.

This result has been used to prove existence of certain homotopy localizations [CSS05, RT03]. There is also an accessible functor characterization of Vopěnka's principle: every subfunctor of an accessible functor is accessible [AR94, 6.31].

4.4. **Weak diamond and amalgamation bases.** In complicated categories where pushouts are not available (e.g. when all morphisms are monos), the amalgamation property can play a key role. In this subsection, we look at a set-theoretic way to obtain it in a general class of concrete categories. More generally, we will look at *amalgamation bases*: objects A such that any span with base A can be completed to a commutative square (see Definition 4.9). We will study them in AECs (Definition 2.12), although several of the concepts are category-theoretic and the result can be generalized to μ-AECs [BGL$^+$16, 6.12], or other kinds of concrete categories [SV, 5.8].

The key set-theoretic component is the weak diamond, a combinatorial principle introduced by Devlin and Shelah [DS78]. We will use it in the following form (see [DS78, 6.1,7.1]):

Theorem 4.14 (Devlin-Shelah)**.** Let λ be an infinite cardinal and let $\langle f_\eta : \lambda \to \lambda \mid \eta \in {}^\lambda 2\rangle$ be a sequence of functions. If there exists $\theta < \lambda$ such that $2^\theta = 2^{<\lambda} < 2^\lambda$, then ($\lambda$ is regular uncountable and) there exists $\eta \in {}^\lambda 2$ such that the set S_η defined below is stationary[17].

$$S_\eta := \{\delta < \lambda \mid \exists \nu \in {}^\lambda 2 : \eta \restriction \delta = \nu \restriction \delta, \eta \restriction (\delta+1) \neq \nu \restriction (\delta+1), \text{ and } f_\eta \restriction \delta = f_\nu \restriction \delta\}$$

Remark 4.15. Given any fixed cardinal θ, there is a unique cardinal λ such that $2^\theta = 2^{<\lambda} < 2^\lambda$, which can equivalently be described as the minimal cardinal λ such that $2^\theta < 2^\lambda$. Note that $\theta < \lambda \leq 2^\theta$. If θ is finite, $\lambda = \theta^+ = \theta + 1$ but if θ is infinite,

[17] A subset of a regular uncountable cardinal λ is called *stationary* if it intersects every closed unbounded subset of λ. If we think of closed unbounded subsets as having full measure, being stationary means having positive measure.

λ is uncountable and moreover regular (because of the formula $2^\lambda = \left(2^{<\lambda}\right)^{\text{cf}(\lambda)}$, see [Jec03, 5.16(iii)]). If the generalized continuum hypothesis[18] (GCH) holds, then $\lambda = \theta^+$, but in general it could be that $\lambda = \theta^+$ even if GCH fails (e.g. if $\theta = \aleph_0$, $2^{\aleph_0} = \aleph_2$, and $2^{\aleph_1} = \aleph_3$). Nevertheless, it is also consistent with the axioms of set theory that $\lambda > \theta^+$. We will say that the *weak generalized continuum hypothesis (WGCH)* holds if $2^\theta < 2^{\theta^+}$ for all infinite cardinals θ. In this case, the hypothesis of the previous theorem holds exactly when λ is an infinite successor cardinal. Note also that, when λ is regular uncountable, the conclusion of the Devlin-Shelah theorem implies that $2^{\lambda_0} < 2^\lambda$ for all $\lambda_0 < \lambda$. Indeed, given $F : {}^\lambda 2 \to {}^{\lambda_0}2$, we can let $f_\eta(\alpha)$ be $F(\eta)(\alpha)$ if $\alpha < \lambda_0$, or 0 otherwise. Fixing the given $\eta \in {}^\lambda 2$, and δ in S_η bigger than λ_0, we obtain $\nu \neq \eta$ with $F(\nu) = F(\eta)$, so F is not injective [She98, Appendix, 1.B(3)].

The essence of the conclusion of the Devlin-Shelah Theorem (4.14) is that if we think of the f_η's as being indexed by branches of a binary splitting tree of height λ, then there exists two branches (i.e. some η and some ν) that split at some big height $\delta < \lambda$ and where moreover the corresponding functions are equal up to δ. If we think of the f_η's as embedding structures into a common codomain, them being equal up to δ will mean that a certain diagram commutes.

To state the promised application to amalgamation bases in AECs, we first give some terminology: for $\mathbf{K}$ an AEC and λ a cardinal, we write $\mathbf{K}_\lambda$ (resp. $\mathbf{K}_{<\lambda}$) for the class of objects in $\mathbf{K}$ of cardinality λ (resp. strictly less than λ). Of course, we identify it with the corresponding category. An object $N \in \mathbf{K}_\lambda$ is called *universal* if any $M \in \mathbf{K}_\lambda$ embeds into N. We will show that, if λ satisfies the hypotheses of the Devlin-Shelah theorem and $\mathbf{K}$ has a universal model in $\mathbf{K}_\lambda$, then amalgamation bases are cofinal in $\mathbf{K}_{<\lambda}$. The result is due to Shelah [She87]. The proof proceeds by contradiction: if amalgamation bases are not cofinal, we can build a tree of failure, and embed each branch of this tree into the universal model. Applying the weak diamond to this tree will yield enough commutativity to get that the tree has a lot of amalgamation bases.

One reason Theorem 4.16 is interesting is that it turns a one-dimensional property in λ (existence of a universal model) into a two-dimensional property below λ (existence of amalgamation bases). There is in fact a higher-dimensional generalization [SV, 11.16], which is much harder to state (but see Section 5.1).

Theorem 4.16 ([She09a, I.3.8])**.** Let $\mathbf{K}$ be an AEC and let $\lambda > \text{LS}(\mathbf{K})$ be such that there exists $\theta < \lambda$ with $2^\theta = 2^{<\lambda} < 2^\lambda$. If there exists a universal model in $\mathbf{K}_\lambda$, then for any $M \in \mathbf{K}_{<\lambda}$, there exists $N \in \mathbf{K}_{<\lambda}$ such that $M \leq_{\mathbf{K}} N$ and N is an amalgamation base in $\mathbf{K}_{<\lambda}$.

Proof. Suppose not. First recall that λ is regular uncountable (Remark 4.15). Fix $M \in \mathbf{K}_{<\lambda}$ with no amalgamation base extending it in $\mathbf{K}_{<\lambda}$. Because $\mathbf{K}$ is isomorphism-closed in $\mathbf{Str}(\tau(\mathbf{K}))$, we may do some renaming to assume without loss of generality that $UM \subseteq \lambda$. We build $\langle M_\eta \mid \eta \in {}^{<\lambda}2 \rangle$ such that for all $\beta < \lambda$ and all $\eta \in {}^\beta 2$:

(1) $M_{\langle\rangle} = M$.

[18]The statement that $2^\lambda = \lambda^+$ for any infinite cardinal λ.

(2) $M_\eta \in \mathbf{K}_{<\lambda}$.
(3) $UM_\eta \subseteq \lambda$.
(4) $M_{\eta\restriction\alpha} \leq_{\mathbf{K}} M_\eta$ for all $\alpha < \beta$.
(5) If β is limit, $M_\eta = \bigcup_{\alpha<\beta} M_\alpha$.
(6) The span $M_{\eta^\frown 0} \leftarrow M_\eta \rightarrow M_{\eta^\frown 1}$ *cannot* be completed to a commutative square (where the maps are inclusion embeddings).

This is possible: the construction proceeds by transfinite induction on the length of η. The base case is given, at successors we use that M_η cannot be an amalgamation base in $\mathbf{K}_{<\lambda}$ by assumption (and do some renaming to implement (3)). At limit stages, we take unions (and use the axioms of AECs).

This is enough: for each $\eta \in {}^\lambda 2$, define $M_\eta := \bigcup_{\alpha<\lambda} M_{\eta\restriction\alpha}$. The axioms of AECs imply, of course, that $M_{\eta\restriction\alpha} \leq_{\mathbf{K}} M_\eta$ for all $\alpha < \lambda$. Moreover, $M_\eta \in \mathbf{K}_\lambda$. Indeed, $M_{\eta\restriction\alpha} \neq M_{\eta\restriction(\alpha+1)}$, as otherwise we would trivially have been able to amalgamate the span $M_{\eta\restriction\alpha^\frown 0} \leftarrow M_{\eta\restriction\alpha} \rightarrow M_{\eta\restriction\alpha^\frown 1}$. Finally, requirement (3) ensures that $UM_\eta \subseteq \lambda$. Fix a universal model $N \in \mathbf{K}_\lambda$. By doing some renaming again, we can assume without loss of generality that $UN \subseteq \lambda$. For each $\eta \in {}^\lambda 2$, fix a $\mathbf{K}$-embedding $g_\eta : M_\eta \rightarrow N$. Our goal will be to get a contradiction to requirement (6). We are not there yet, because even if $\eta \restriction \delta = \nu \restriction \delta$, the maps g_η and g_ν will not necessarily agree on $M_{\eta\restriction\delta}$. This is where the weak diamond will come in.

Define $f_\eta : \lambda \rightarrow \lambda$ by $f_\eta(\alpha) = g_\eta(\alpha)$ if $\alpha \in UM_\eta$, and $f_\eta(\alpha) = 0$ if $\alpha \notin UM_\eta$. We are now in the setup of the Devlin-Shelah theorem: fix $\eta \in {}^\lambda 2$ such that the set S_η defined there is stationary. Consider the set $C := \{\delta < \lambda \mid UM_{\eta\restriction\delta} \subseteq \delta\}$. The set C is closed (because of requirement (5)) and unbounded. To see the latter, we run a standard "catching your tail" argument[19]: fix $\alpha < \lambda$. We inductively build an increasing sequence of ordinals $\langle \alpha_n : n < \omega \rangle$ as follows: take $\alpha_0 = \alpha$, and given α_n, we know $UM_{\eta\restriction\alpha_n} \subseteq \lambda$ (requirement (3)), and $|UM_{\eta\restriction\alpha_n}| < \lambda$ (requirement (2)), so use regularity of λ to pick $\alpha_{n+1} \in [\alpha, \lambda)$ with $UM_{\eta\restriction\alpha_n} \subseteq \alpha_{n+1}$. At the end, we let $\delta := \sup_{n<\omega} \alpha_n$. The construction, together with requirement (5), implies that $\delta \geq \alpha$ and $\delta \in C$.

Because C is closed unbounded and S_η is stationary, we can pick $\delta \in C \cap S_\eta$. By definition of S_η, there exists $\nu \in {}^\lambda 2$ such that $\eta \restriction \delta = \nu \restriction \delta$, $\eta(\delta) \neq \nu(\delta)$, and $f_\eta \restriction \delta = f_\nu \restriction \delta$. Let $\rho := \eta \restriction \delta$. By definition of C, $f_\eta \restriction M_\rho = f_\nu \restriction M_\rho$. This implies that f_η, f_ν, and N witness the span $M_{\eta\restriction(\delta+1)} \leftarrow M_\rho \rightarrow M_{\nu\restriction(\delta+1)}$ can be completed to a commutative square, contradicting requirement (6). □

If we know weak GCH holds, and moreover we know that $\mathbf{K}$ has a single object (up to isomorphism) in two successive cardinalities[20], the theorem simplifies and we get the amalgamation property locally:

Corollary 4.17. Let $\mathbf{K}$ be an AEC, and let $\mu \geq \mathrm{LS}(\mathbf{K})$ be such that $2^\mu < 2^{\mu^+}$. If $\mathbf{K}$ has a single model (up to isomorphism) in cardinalities μ and μ^+, then $\mathbf{K}_\mu$ has the amalgamation property.

Proof. By its uniqueness, the object of cardinality μ^+ must be universal. Applying Theorem 4.16 with $\lambda = \mu^+$, we get that there exists an amalgamation base in $\mathbf{K}_\mu$,

[19]This is in fact an instance of Theorem A.9 in the appendix, see Example A.10.
[20]Model-theorists call this property *categoricity*, see 5.14.

and this amalgamation base must be isomorphic to any other object of cardinality μ, hence $\mathbf{K}_\mu$ has amalgamation. □

5. Generalized pushouts and stable independence

Many of the "classical" categories of mainstream mathematics are bicomplete: they have all limits and colimits. However, problems will occur if we want to study such categories set-theoretically, and specifically if we want to restrict ourselves to certain classes of monomorphisms. It is clear quotients (coequalizers) will be lost, but one will also lose pushouts: any category with pushouts, all morphisms monos, and an initial object must be thin, hence essentially just a poset [LRV19a, 3.30(3)]. Still, it is natural to ask for approximations to pushouts. The results of this section are a survey of the work of Lieberman, Rosický, and the author on this question [LRV19a, LRVc].

For the purpose of the discussion to follow, let's briefly repeat that for a given a diagram $D : I \to \mathcal{K}$, a *cocone* for that diagram is an object A together with maps $(D_i \xrightarrow{f_i} A)_{i \in I}$ commuting with the diagram. One can look at the category $\mathcal{K}_D$ of cocones for D, where the morphisms are defined as expected. The colimit of a diagram D is then simply an initial object (one that has a unique morphism to every other object) in the category of cocones. In case D is a span $B \xleftarrow{f} A \xrightarrow{g} C$, a cocone is simply an amalgam of this span, and an initial object in the category $\mathcal{K}_D$ of cocones would be a pushout. The amalgamation property (Definition 4.9) simply says that $\mathcal{K}_D$ is non-empty, i.e. that there is *some* cocone, maybe satisfying no universal property whatsoever. A much stronger approximation is the existence of *weak pushouts*: a weak pushout of a given span D is a cocone that is weakly initial in the category of cocones for D: there is a morphism to every other cocone, but that morphism is not required to be unique. In categories where all morphisms are monos, weak pushouts are still too strong of a requirement:

Example 5.1. Consider the category $\mathbf{Set}_{mono}$ of sets with injections. Consider the inclusion of $A = \emptyset$ into $B = \{0, 1\}$ and $C = \{0, 2\}$. Let $D = \{0, 1, 2\}$. Then D, together with the corresponding inclusions, is a cocone/amalgam for the span: $B \leftarrow A \rightarrow C$. On the other hand, consider $D' = \{1, 2, 3, 4\}$ and $f_1 : B \to D'$, $f_2 : C \to D'$ defined by $f_1(0) = 3$, $f_1(1) = 1$, $f_2(0) = 4$, $f_2(2) = 2$. (f_1, f_2) is also an amalgam of $B \leftarrow A \rightarrow C$, but D has no morphisms to D' (in the category of cocones for the span $B \leftarrow A \rightarrow C$), and by cardinality considerations D' has no morphisms to D either. Thus $\mathbf{Set}_{mono}$ does not have weak pushouts.

What is happening in the example is that we had two choices for amalgamating B and C: either sending 0 to the same place, or sending it two different elements. These two choices are then incompatible, in the sense that no amalgam of one type will ever have a morphism into an amalgam of the other type (in the appropriate category of amalgams of a fixed span). In other words, the category of amalgams is not connected. Let us make this explicit, first in complete generality, then for the specific case of amalgams:

Definition 5.2. Two objects A and B are called *comparable*[21] if either $\text{Hom}(A, B) \neq \emptyset$ or $\text{Hom}(B, A) \neq \emptyset$. We say that A and B are *connected* if there exists $A =$

[21]The terminology comes from posets.

$A_0, A_1, \ldots, A_n = B$ such that A_i and A_{i+1} are comparable for all $i < n$. This is an equivalence relation, and the equivalence class is called a *connected component* of the category. The category is called *connected* if any two of its objects are connected. We say that A and B are *jointly connected* if there exists an object C with morphisms $A \to C$, $B \to C$.

Of course, the connected components of a category are exactly the connected component of the undirected graph whose vertices are objects and where there is an edge between A and B exactly when A and B are comparable. Assuming amalgamation, joint connectedness coincides with connectedness:

Lemma 5.3. In a category with the amalgamation property, two objects are connected if and only if they are jointly connected.

Proof. Let $\mathcal{K}$ be a category with the amalgamation property. First note that if two objects A and B are comparable, then they are jointly connected (for example, if $A \xrightarrow{f} B$ then $A \xrightarrow{f} B = C \xleftarrow{\mathrm{id}_B} B$ witness the joint connectivity). Also, if A and B are jointly connected, as witnessed by $A \to C$, $B \to C$, then A and C are comparable, and C and B are comparable, so A and B are connected. We now show that being jointly connected is an equivalence relation. This will be enough because being connected is the smallest equivalence relation extending comparability. So assume that A_0, A_1 are jointly connected, as witnessed by $A_0 \to B \leftarrow A_1$, and A_1, A_2 are jointly connected, as witnessed by $A_1 \to C \leftarrow A_2$. Amalgamate B and C over A_1, forming the diagram below:

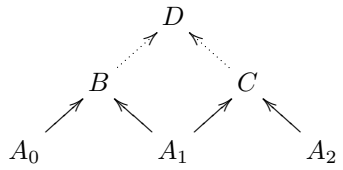

Then D and the expected compositions witness that A_0 and A_2 are jointly connected. □

Definition 5.4. Two amalgams $(B \xrightarrow{f^a} D^a, C \xrightarrow{g^a} D^a)$ and $(B \xrightarrow{f^b} D^b, C \xrightarrow{g^b} D^b)$ of a span $B \leftarrow A \to C$ are *jointly connected* if they are jointly connected in the category of cocones for the appropriate span. Explicitly, there exists D and morphisms into D making the following diagram commute:

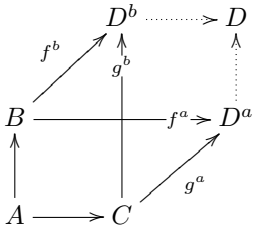

We say that two amalgams D^a and D^b are *connected* if they are connected in the appropriate category of cocones, i.e. there exists a chain of amalgams $D^0, D^1, \ldots, D^n$

(along with respective morphisms) such that $D^0 = D^a$, $D^n = D^b$, and D^i is jointly connected to D^{i+1} for all $i < n$.

Note that if $\mathcal{K}$ has the amalgamation property, then the category of cocones over a diagram in $\mathcal{K}$ also has the amalgamation property, so in this case connectedness already implies joint connectedness. The replacement for pushouts we are looking for will therefore consist in a choice of connected component for each span. Eventually, we may want to look for a weakly initial object for this specific component class (called a *prime* object by model theorists), but this seems to be too strong a requirement to start with: we want to prove it, not assume it.[22] Thus we instead impose a transitivity conditions on the choice of connected component that make the resulting squares into the morphisms of an arrow category (and also holds of pushouts). Still by analogy with pushouts, and for reasons to be discussed later, we require this new category to be accessible and call the result a stable independence notion.

Definition 5.5 ([LRV19a, 3.24])**.** A *stable independence notion* in a category $\mathcal{K}$ is a class of squares (called *independent squares* and marked here with the anchor symbol $\downarrow$) such that:

(1) Independent squares are closed under connectedness: if one amalgam of a given span is independent, then all the connected amalgams are also independent.
(2) Existence: any span can be amalgamated to an independent square.
(3) Uniqueness: any two independent amalgams of the same span are connected.
(4) Transitivity:

$$\begin{array}{ccccc} B & \to & D & \to & F \\ \uparrow & \downarrow & \uparrow & \downarrow & \uparrow \\ A & \to & C & \to & E \end{array} \Rightarrow \begin{array}{ccc} B & \to & F \\ \uparrow & \downarrow & \uparrow \\ A & \to & E \end{array}$$

(5) Symmetry: "swapping the ears" B and C preserves independence.
(6) Accessibility: the arrow category whose objects are morphisms of $\mathcal{K}$ and whose morphisms are independent squares is accessible. We let $\mathcal{K}_\downarrow$ denote this category.

Remark 5.6. The existence property implies the amalgamation property. Hence any two connected amalgams are, in fact, jointly connected. Note also that any two amalgams of a span which contains an isomorphism will be connected, hence (by the existence property) independent squares [LRV19a, 3.12]. In particular, the identity morphism in the arrow category $\mathcal{K}^2$ will indeed be an independent square, hence a morphism of $\mathcal{K}_\downarrow$. The transitivity property ensures that $\mathcal{K}_\downarrow$ is closed under composition, and hence is indeed a category.

Most of the examples below are listed in [LRV19a, 3.31], though of course many are trivial or date back to the beginning of stable first-order theories. See there for more references.

[22]For example, the proofs of existence and uniqueness of differential closure in differentially closed fields (Example 5.7(5)) rely on properties of the independence notion.

Example 5.7.

(1) In an accessible category with weak pushouts, there is a stable independence notion given by all commutative squares. This holds more generally in any accessible category with the amalgamation property where any two amalgams are always connected, for example in accessible categories with a terminal object.
(2) The category $\mathbf{Set}_{mono}$ has a stable independence notion: identifying the morphisms with inclusions, given $A \subseteq B \subseteq D$, $A \subseteq C \subseteq D$, we say that the resulting square A, B, C, D is independent exactly when $B \cap C = A$. That is, the amalgam must be disjoint. In general, the ears of an independent squares "do not interact": independent squares are a notion of "free" amalgam.
(3) The category of vector spaces over a fixed field, with morphisms the injective linear transformations, has a stable independence notion, again given by disjoint amalgamation. This holds more generally for the category of modules over a fixed ring.
(4) The category of all algebraically closed fields of a fixed characteristic (with morphisms field homomorphisms) has a stable independence notion, essentially given by algebraic independence: again identifying the morphisms with inclusions, if we have a square of fields $F_0 \subseteq F_1 \subseteq F_3$, $F_0 \subseteq F_2 \subseteq F_3$, we say it is independent if for any subset $A \subseteq F_1$ and any $b \in F_1$, if b is in the algebraic closure (computed in F_3) of $A \cup F_2$, then it is in the algebraic closure of $A \cup F_0$. Here, it does *not* suffice to require that $F_1 \cap F_2 = F_0$ (because the pregeometry induced by algebraic closure is not modular).
(5) A *differential field* is a field together with an operator D that preserves sums and satisfies Leibnitz' law: $Dfg = gDf + fDg$. A *differentially closed field* is, roughly, a differential field in which every system of linear differential equations that could possibly have a solution has a solution. Model-theoretic methods were used to obtain the first proof that every differential field of characteristic zero has, in some sense, a differential closure. Uniqueness of that differential closure was proven by Shelah, using the fact that the category of differentially closed fields of characteristic zero has a stable independence notion. See [Sac72] for a short overview.
(6) Generalizing the last four examples, let T be a stable first-order theory (T is *stable* if it does *not* have the order property, and T has the *order property* if there exists a formula $\phi(\bar{x}, \bar{y})$, a model M of T, and a sequence $\langle a_i : i < \omega \rangle$ of elements in M such that $M \models \phi[\bar{a}_i, \bar{a}_j]$ if and only if $i < j$, see Definition B.1). Consider the category $\mathcal{K} = \mathbf{Elem}(T)$ of models of T with elementary embeddings. Then $\mathcal{K}$ has a stable independence notion. The proof (originally due to Shelah), is not trivial, see Appendix B for a short exposition. The definition of an independent square mirrors that of fields (the reader should think of the formula ψ in the next sentence as a polynomial): a square $M_0 \preceq M_1 \preceq M_3$, $M_0 \preceq M_2 \preceq M_3$ is *independent* exactly when for any finite sequences $\bar{a}$ from M_1 and $\bar{b}$ from M_2, and any formula $\psi(\bar{x}, \bar{y})$, if $M_3 \models \psi[\bar{a}, \bar{b}]$, then there exists a sequence $\bar{b}_0$ in M_0 so that $M_3 \models \psi[\bar{a}, \bar{b}_0]$. In fact, independent squares are the same as what model theorists call nonforking squares. Thus we have recovered the important model-theoretic notion of forking from simple category-theoretic

considerations! In fact, the original motivation behind the definition of stable independence was to generalize forking.

(7) Conversely, if T is not stable then $\mathbf{Elem}(T)$ does *not* have a stable independence notion [LRV19a, 9.9].

(8) The previous item implies, in particular, that the category of graphs, with morphisms the full subgraph embeddings does *not* have a stable independence notion. Another proof will be given in Remark 5.9. The category **Lin** of linear orders with order-preserving maps similarly does not have a stable independence notion.

(9) The category of graphs with morphisms the (not necessarily full) subgraph embeddings *does* have a stable independence notion: two graphs are independent over a base graph (inside an ambient graph) if all the cross-edges between the two are inside the base graph. See Example 5.13(2).

(10) [LRV19a, 4.8(5)] The category of Hilbert spaces with isometries has a stable independence notion, given by pullback squares. These correspond roughly to squares where everything is "as orthogonal as possible".

(11) We will give later general constructions of a stable independence notion, giving many other nontrivial examples. For example, for a fixed ring, the category of all flat modules with morphisms the flat monomorphisms has a stable independence notion (Example 5.13(5)).

For model theorists, we note that in concrete cases (e.g. inside an ∞-AEC $\mathbf{K}$), the independence notion can be extended to a relation of the form $A \overset{N}{\underset{M}{\downarrow}} B$, where $M \leq_{\mathbf{K}} N$ and $A, B \subseteq UN$, satisfying generalizations of the properties of a stable independence notion. The accessibility axiom can then be shown to be equivalent to the conjunction of the two classical properties of forking: the *witness property* (failure to be independent is witnessed by small subsets of the ears) and the *local character property* (every type is independent over a small set). The proof of this fact is essentially a generalization of the proof of Theorem 3.17, see [LRV19a, 8.14]. We deduce, in particular, that an ∞-AEC has quite a bit of structure when it has a stable independence notion (e.g. it is tame and stable [LRV19a, 8.16]).

One of the most important fact about stable independence (and a strong justification for the definition) is that it leads to a canonical notion: if a category has a stable independence notion, then under very reasonable conditions it can have only one. For first-order forking, this is well known [HH84], but in fact it holds even in the very general categorical setup. The proof for AECs in [BGKV16] was adapted to ∞-AECs with chain bounds in [LRV19a, 9.1], and finally to any category with chain bounds in [LRVc, A.6]. We say a category has *chain bounds* if any ordinal-indexed chain has a cocone (this holds, of course, anytime the category has directed colimits).

Theorem 5.8 (The canonicity theorem, [LRVc, A.6])**.** A category with chain bounds has at most one stable independence notion. In fact, given a stable independence notion $\downarrow$, any other relation satisfying all the axioms of stable independence notion except perhaps accessibility will have to be $\downarrow$.

Proof idea. First, existence of a stable independence notion implies that the category itself is accessible [LRV19a, 3.27]. Let $\overset{1}{\downarrow}$ be a stable independence notion and let $\overset{2}{\downarrow}$ satisfy all the axioms of stable independence, except perhaps accessibility. Given a span $B \leftarrow A \rightarrow C$, it suffices to show that it has an amalgam which is independent in the sense of both $\overset{1}{\downarrow}$ and $\overset{2}{\downarrow}$ (then uniqueness and invariance under connectedness show that $\overset{1}{\downarrow}$ and $\overset{2}{\downarrow}$ must coincide for any amalgam of this span). The idea of the construction is perhaps best described by the following property of a vector space: given any sequence I of vectors and any other vector a, there exists a finite subset I_0 of I such that $(I - I_0) \cup \{a\}$ is independent. Thus the idea is to first complete the span $B \leftarrow A \rightarrow C$ to a $\overset{2}{\downarrow}$-independent square, then build many $\overset{2}{\downarrow}$-independent "copies" of that square. Analogously to the fact mentioned for vector spaces, all except a few of these copies must also be $\overset{1}{\downarrow}$-independent. □

Remark 5.9. The canonicity theorem provides us with a tool to prove the *non-existence* of stable independence notions. Indeed, it suffices for this purpose to find two different notions satisfying all the axioms of stable independence except perhaps for accessibility. For example, the category of graph with full subgraph embeddings does not have a stable independence notion. Indeed, the relations "all cross-edges are contained inside the base" and "all cross-edges outside those in the base are present" satisfy all the axioms of stable independence except for accessibility [LRV19a, 3.31(6)].

We now survey two different general constructions of a stable independence notion. In both cases, the hypotheses are provably optimal in the sense that the existence of stable independence implies them. The first construction works in any ∞-AEC with chain bounds, but uses large cardinals and only obtains an independence notion on a cofinal subclass (consisting of "sufficiently homogeneous" objects).

Theorem 5.10 ([LRV19a, 10.3])**.** Assume Vopěnka's principle. Let $\mathbf{K}$ be an ∞-AEC. The following are equivalent:

(1) $\mathbf{K}$ does not have the order property[23].
(2) $\mathbf{K}$ has a cofinal subclass of "sufficiently homogeneous" objects which has a stable independence notion.

Proof ideas. If $\mathbf{K}$ has the order property, as witnessed by a long sequence $\langle \bar{a}_i : i < \lambda \rangle$ (λ here is a big regular cardinal), we can use the accessibility and uniqueness properties of stable independence to get a subsequence $\langle \bar{a}_i : i \in S \rangle$ with $|S| = \lambda$ that is independent and indiscernible (essentially, this means that the sequence is "very homogeneous" – it looks like a sequence of mutually transcendental elements in an algebraically closed fields). The symmetry axiom can then be use to show that the elements of the subsequence can be permuted without impacting their

[23]This is defined by generalizing the definition given in Example 5.7(6) and in B.1: there is no long sequence that can be ordered by a type in the sense to be given in Section 6. See [LRV19a, 9.7].

properties. In particular, the sequence cannot serve as a witness for the order property, contradiction.

Going from no order property to stable independence, the very rough idea is to imitate the standard construction given Appendix B, but replace $\aleph_0$ by a sufficiently-big strongly compact, and look only at the locally κ-homogeneous models (see [LRV19a, 7.3]). Vopěnka's principle is used to prove that this subclass is an ∞-AEC, and also provides enough compactness to prove the existence property. □

For the second method, we start with a "classical" category $\mathcal{K}$: a locally presentable category. We single out a certain class of morphisms $\mathcal{M}$ (usually some class of nice monomorphisms), and we want to build a stable independence notion on the category $\mathcal{K}_{\mathcal{M}}$ obtained by restricting the morphisms in $\mathcal{K}$ to be those of $\mathcal{M}$. To make this setup more precise, we will have to make some assumptions on $\mathcal{M}$. Of course, we want at minimum $\mathcal{M}$ to contain all isomorphisms and be closed under compositions. It turns out it is very convenient to assume that $\mathcal{M}$ satisfies a coherence property, similar to the coherence axiom of ∞-AECs: for morphisms f, g of $\mathcal{K}$, if $gf \in \mathcal{M}$ and $g \in \mathcal{M}$, then $f \in \mathcal{M}$. We will say that $\mathcal{M}$ is *coherent*. For technical reasons, we will also want $\mathcal{M}$ to be closed under retracts (in the arrow category $\mathcal{K}^2$). This holds under very mild conditions. For example, if $\mathcal{M}$ is closed under compositions, contains all split monos, and is coherent, then it is closed under retracts. Since we want to use pushouts to build our stable independence notion, we will require $\mathcal{M}$ to be *closed under pushouts*: a pushout of a morphism in $\mathcal{M}$ (not necessarily along a morphism in $\mathcal{M}$) should be in $\mathcal{M}$. Finally, to perform iterative constructions, we will need $\mathcal{M}$ to be *closed under transfinite compositions*.

Under all these conditions, we get that $\mathcal{K}_{\mathcal{M}}$ has a stable independence notion *if and only if* it is accessible and $\mathcal{M}$ is cofibrantly generated. Here, $\mathcal{M}$ is *cofibrantly generated* if there is a subset $\mathcal{X}$ of $\mathcal{M}$ (i.e. not a proper class) such that closing $\mathcal{X}$ under retracts, pushouts, and transfinite compositions gives back $\mathcal{M}$. This notion originated in algebraic topology, where the cell complexes are precisely the objects that can be obtained from finitely-many pushouts and composition by starting from inclusions $\partial D^n \to D^n$ of the boundary of the n-ball. As we will see, similar notions have been used in homological algebra as well.

Theorem 5.11 ([LRVc, 4.9])**.** Let $\mathcal{K}$ be a locally presentable category and let $\mathcal{M}$ be a coherent class of morphisms containing all isomorphisms, closed under retracts, transfinite compositions, and pushouts. The following are equivalent:

(1) $\mathcal{K}_{\mathcal{M}}$ has a stable independence notion.
(2) $\mathcal{K}_{\mathcal{M}}$ is accessible and $\mathcal{M}$ is cofibrantly generated.

Proof ideas. Neither direction is easy. First, we need a candidate definition for a stable independence notion. Let us say that a commutative square with morphisms in $\mathcal{M}$ is *effective* if the induced map from the pushout is in $\mathcal{M}$. That is, the outer square in the diagram below is effective if all its maps are in $\mathcal{M}$ and the induced map f from the pushout P is in $\mathcal{M}$.

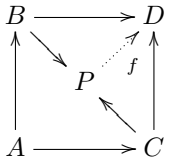

Without any additional hypotheses, it is not too difficult to show that effective squares will satisfy all the axioms of stable independence, except perhaps for accessibility. Let $\mathcal{K}_{\mathcal{M},\downarrow}$ be the category whose objects are morphisms in $\mathcal{M}$ and morphisms are effective squares. We can also show that $\mathcal{K}_{\mathcal{M},\downarrow}$ has directed colimits (and they are computed as in $\mathcal{K}^2$).

- (1) implies (2): Starting from a stable independence notion on $\mathcal{K}_{\mathcal{M}}$, it is not too difficult to show that $\mathcal{K}_{\mathcal{M}}$ must be accessible [LRV19a, 3.27]. Moreover, the stable independence notion must be given by effective squares, by the canonicity theorem (5.8). Let λ be a big-enough regular uncountable cardinal such that all the relevant categories are λ-accessible. For $\mu \geq \lambda$ regular, let $\mathcal{M}_\mu$ denote the class of morphisms in $\mathcal{M}$ with μ-presentable domain and codomain (in $\mathcal{K}$). For a class $\mathcal{H}$ of morphisms, let $\text{cof}(\mathcal{H})$, the *cofibrant closure of* $\mathcal{H}$, denote the closure of $\mathcal{H}$ under retracts, pushouts, and transfinite compositions. We will show that $\mathcal{M} = \text{cof}(\mathcal{M}_\lambda)$. For this, we prove by induction that for all regular cardinals μ, $\mathcal{M}_\mu \subseteq \text{cof}(\mathcal{M}_\lambda)$. For $\mu \leq \lambda$ this is trivial and if μ is limit, $\mathcal{M}_\mu = \bigcup_{\theta<\mu} \mathcal{M}_\theta$ by Corollary 4.2(1). Thus we can assume $\mu = \mu_0^+$. Let $\delta := \text{cf}(\mu_0)$, and fix a morphism $A \xrightarrow{f} B$ in $\mathcal{M}_\mu$. We write f as the directed colimit of a chain $\langle f_i : i < \delta \rangle$ in $\mathcal{K}_{\mathcal{M},\downarrow}$, where each f_i is in $\mathcal{M}_\theta$ for a regular $\theta \leq \mu_0$. Finding such a chain is nontrivial (in accessible categories, objects can be written as directed colimits of directed diagrams of small objects, it is not clear you can find chains like this): assuming for simplicity all morphisms are monos, we use that it is possible in AECs, together with the fact every λ-accessible category $\mathcal{L}$ with directed colimits is a reflexive full subcategory of a finitely accessible category (which is given by taking free directed colimits of the λ-presentable objects of $\mathcal{L}$), and finitely accessible categories with all morphisms monos are AECs (Theorem 3.18).

  Once we have the f_i's, say $A_i \xrightarrow{f_i} B_i$, they are each by the induction hypothesis part of $\text{cof}(\mathcal{M}_\lambda)$. It suffices to use them to generate f. Let $(A_i \xrightarrow{g_i} A, B_i \xrightarrow{h_i} B)_{i<\delta}$ be the colimit maps.

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ {\scriptstyle g_i}\uparrow & \downarrow & \uparrow{\scriptstyle h_i} \\ A_i & \xrightarrow[f_i]{} & B_i \end{array}$$

  Take the pushout P_0 of f_0 along g_0. We get that $f = p_0 \bar{f}_0$, where p_0 is the induced map from the pushout, which is also in $\mathcal{M}$ by assumption. We have managed to generate $\bar{f}_0$, and we now repeat what we did for f but for p_0 instead (write it as the colimit of an increasing chain of morphisms that are above f_1, take the pushout of the first morphism, etc.). In the end, we

will have written f as a δ-length transfinite compositions of pushouts of members of $\text{cof}(\mathcal{M}_\lambda)$, as desired.

- (2) implies (1): As before, fix λ an uncountable regular cardinal so that all relevant categories are λ-accessible and $\mathcal{M} = \text{cof}(\mathcal{M}_\lambda)$. By using what is called "good colimits" [MRV14, B.1], we can in fact show that $\mathcal{M}$ is generated from $\mathcal{M}_\lambda$ by just using pushouts and transfinite compositions (no retracts). Let $\mathcal{H}$ be the class of all morphisms that, in $\mathcal{K}_{\mathcal{M},\downarrow}$ are λ-directed colimits of morphisms of $\mathcal{M}_\lambda$. It suffices to show that $\mathcal{H} = \mathcal{M}$, and for this it suffices to see that $\mathcal{H}$ is closed under pushouts and transfinite compositions. This can readily be done, using the definition of an effective square.

□

Remark 5.12. In Theorem 5.11, if $\mathcal{K}$ is λ-accessible and all morphisms of $\mathcal{M}$ are monos, then $\mathcal{K}_\mathcal{M}$ will be (equivalent to) a λ-AEC [LRVc, 3.10]. This is especially interesting when $\lambda = \aleph_0$ (i.e. $\mathcal{K}$ is locally finitely presentable), in which case we get an AEC.

Compared to Theorem 5.10, Theorem 5.11 does not use large cardinal axioms and more importantly gives a stable independence notion on the whole category. The hypotheses seem to hold often-enough in practice (all the examples below and more are in [LRVc, §6]):

Example 5.13.

(1) Let $\mathcal{K}$ be the category of graphs (reflexive and symmetric binary relations) with homomorphisms. This is locally finitely presentable. Let $\mathcal{M}$ be the full subgraph embeddings. We have seen (Remark 5.9) that $\mathcal{K}_\mathcal{M}$ does not have a stable independence notion. This automatically tells us that $\mathcal{M}$ is *not* cofibrantly generated.

(2) Let $\mathcal{K}$ again be the category of graphs with homomorphisms. This time, let $\mathcal{M}$ be the subgraph embeddings (corresponding to monomorphisms in $\mathcal{K}$ – the full subgraph embeddings correspond to the *regular monomorphisms*: those that are equalizers of some pair of morphism). Then $\mathcal{M}$ is cofibrantly generated by $\emptyset \to 1$ and $1+1 \to 2$, where 1 is a vertex, 2 is an edge, and $1+1$ is an empty graph on two vertices. Therefore $\mathcal{K}_\mathcal{M}$ has a stable independence notion.

(3) An *orthogonal factorization system* in a category $\mathcal{K}$ consist of two classes of morphisms $(\mathcal{M}, \mathcal{N})$ such that both contain all the isomorphisms, both are closed under composition, and every map factors as gf with $f \in \mathcal{M}$ and $g \in \mathcal{N}$ in a way that is unique up to unique isomorphism. For example, in the category of sets (epi, mono) is an orthogonal factorization system. A *weak factorization system* is defined in a somewhat similar way, except the uniqueness condition is considerably relaxed. An example is (mono, epi) in the category of sets. It turns out that (in a cocomplete category) the left part of a weak factorization system is always closed under pushouts, retracts, and transfinite compositions. Thus we can call a weak factorization system $(\mathcal{M}, \mathcal{N})$ *cofibrantly generated* if $\mathcal{M}$ is cofibrantly generated. Thus if $\mathcal{M}$ is the left part of a weak factorization system in a locally presentable

category $\mathcal{K}$, $\mathcal{M}$ is coherent, and $\mathcal{K}_{\mathcal{M}}$ is accessible, then $\mathcal{K}_{\mathcal{M}}$ has a stable independence notion if and only if the weak factorization system is cofibrantly generated. In fact, the *small object argument* tells us that if $\mathcal{X}$ is any set of morphisms in a locally presentable category, then $\text{cof}(\mathcal{X})$ will form the left part of a weak factorization system [Bek00, 1.3]. Thus in the setup of Theorem 5.11, existence of a stable independence notion on $\mathcal{K}_{\mathcal{M}}$ automatically implies that $\mathcal{M}$ forms the left part of a weak factorization system.

(4) In algebraic topology, a *model category* is a bicomplete category together with three classes of morphisms, F (fibrations), C (cofibrations), and W (weak equivalences), such that the weak equivalences satisfy the two out of three property (if two of g, f and gf are weak equivalences, so is the third), and both $(C, F \cap W)$ and $(C \cap W, F)$ form weak factorization systems. A model category is called *cofibrantly generated* when both weak factorization systems are cofibrantly generated. The typical example includes the usual fibrations, cofibrations, and weak homotopy equivalences in the category of topological spaces, or the monomorphisms, Kan fibrations, and weak homotopy equivalences in the category of simplicial sets. There are however model category of a more algebraic flavor, including a model category on chain complexes of modules [Hov99, Chapter 2]. Thus the previous example also describes a two way connection between model categories and stable independence.

(5) Let R be a (associative and unital) ring and let $\mathcal{K}$ be the category of R-modules with homomorphisms. One can check that $\mathcal{K}$ is locally finitely presentable. We want to study the modules that are *flat* (i.e. directed colimits of free modules – this is the easiest equivalent definition for the purpose of this discussion). We will do this through the class $\mathcal{M}$ of *flat monomorphisms*: monomorphisms whose cokernel is flat. In particular, an inclusion $A \subseteq B$ is flat if and only if B/A is flat. One can check that $\mathcal{M}$ contains all isomorphisms, is coherent, and closed under pushouts, retracts, and transfinite compositions. Now the class $\mathcal{K}_{\mathcal{M}}$ is not quite the right category, we really want to study the category $\mathcal{F}_{\mathcal{M}}$, where $\mathcal{F}$ is the full subcategory of $\mathcal{K}$ consisting of flat modules. Note that $\mathcal{F}$ can be described in terms of $\mathcal{M}$: an object A is in $\mathcal{F}$ if and only if the initial morphism $0 \to A$ is in $\mathcal{M}$. We will say that A is a cofibrant object (with respect to $\mathcal{M}$). Before even worrying about stable independence, is $\mathcal{F}_{\mathcal{M}}$ an accessible category? Using Theorem 3.18, one can see that this is equivalent to asking whether it is an AEC. Similar examples were studied by Baldwin-Eklof-Trlifaj [BET07], where it is shown that $\mathcal{F}_{\mathcal{M}}$ is an AEC if and only if $\mathcal{F}$ has refinements: there is a regular cardinal λ such that any flat module A is the colimit of an increasing chain $\langle A_i : i < \delta \rangle$ with $A_0 = 0$ and A_{i+1}/A_i flat and λ-presentable. Earlier, Rosický had essentially shown [Ros02, 4.5] that having refinements is equivalent to $\mathcal{M}$ being cofibrantly generated (even by a subset of $\mathcal{M} \cap \mathcal{F}$)! Thus having refinements is yet another disguise for being cofibrantly generated. Using a slight variation on Theorem 5.11, one can then close the loop to deduce that if $\mathcal{F}$ has refinements then $\mathcal{F}_{\mathcal{M}}$ has a stable independence notion. It turns out that having refinement is closely connected to the (now resolved) *flat cover conjecture* [Eno81] (a dualization

of the fact that every module has an injective envelope). An argument of Bican, El Bashir, and Enochs [BBE01] (which also easily could have been deduced from existing results on accessible categories, see [Ros02, 3.2] or [LRVc, 6.21]) establishes that indeed, $\mathcal{F}$ has refinements, and thus $\mathcal{F}_{\mathcal{M}}$ indeed is an AEC with a stable independence notion.

5.1. **Higher dimensional independence and categoricity.** Recall that stable independence was defined as a class of squares satisfying certain properties, the most important of which was that the arrow category $\mathcal{K}_{\downarrow}$ (whose objects are the arrows of $\mathcal{K}$ and whose morphisms are the independent squares) should be accessible.

Now, given any accessible category, it makes sense to ask whether it has a stable independence notion. Does $\mathcal{K}_{\downarrow}$ have a stable independence notion? If it does (call this stable independence notion $\overset{*}{\perp}$), then what do the objects and arrow look like in the corresponding category $(\mathcal{K}_{\downarrow})_{\downarrow^*}$? Well, the objects are the morphisms of $\overset{*}{\perp}$, i.e. independent squares. The morphisms, in turn, are $\overset{*}{\perp}$-independent "squares" in $\mathcal{K}_{\downarrow}$, but they really are morphisms between two independent squares of $\mathcal{K}$, hence it makes sense to call them independent *cubes*. Continuing in this way, one can define when a category has an n-dimensional stable independence notion, for any $n < \omega$ (the original definition of stable independence is the case $n = 2$).

Very nice, but what are (possibly multidimensional) stable independence notions good for? Very roughly, they are useful to prove the existence and uniqueness of certain objects in a category. As a simple example, let $\mathbf{K}$ be an AEC with $\text{LS}(\mathbf{K}) = \aleph_0$ (e.g. the AEC of abelian groups, ordered with subgroup). Let's assume that there is an object of cardinality $\aleph_0$, and suppose we want to establish that $\mathbf{K}$ has an object of cardinality $\aleph_1$. A simple way would be to first show that for every countable $A \in \mathbf{K}$, there exists a countable $B \in \mathbf{K}$ with $A <_{\mathbf{K}} B$. Equivalently, there is a morphism with domain A that is *not* an isomorphism. If this last property holds, then we can build a strictly increasing chain $\langle A_i : i < \omega_1 \rangle$ of countable objects (taking unions at limits) and the union of this chain will be the desired object of cardinality $\aleph_1$.

One can think of this construction as establishing a 0-dimensional property in $\aleph_1$ (existence of an object) by using a 1-dimensional property in $\aleph_0$ (existence of extensions). There was nothing special about $\aleph_0$ and $\aleph_1$ in this example, we could have replaced them with λ and λ^+ for an arbitrary infinite cardinal λ. In particular, existence in $\aleph_2$ is implied by extensions in $\aleph_1$. Now how do we get existence of extensions in $\aleph_1$? Well, it is a 1-dimensional property, so it seems reasonable that we should look for a 2-dimensional property in $\aleph_0$. Such a property is the *disjoint amalgamation property*: any span can be completed to a pullback square. Note that in many cases, an independent square will be a pullback square [LRV19a, 10.6]. Still, existence is a purely combinatorial property. Stable independence notions become really useful to prove *uniqueness* properties. The simplest uniqueness property is what model theorists call *categoricity* (this is somewhat unfortunate terminology, dating back from before the invention of category theory [Veb04, Łos54]).

Definition 5.14. For an infinite cardinal λ, a category is called λ-*categorical* (or *categorical in* λ) if it has exactly one object of size λ (up to isomorphism).

The following are classical examples of the occurrence of categoricity:

Example 5.15.

(1) The category of abelian groups is not categorical in any infinite cardinal (for a cardinal λ, take λ copies of $\mathbb{Z}$ and λ copies of $\mathbb{Z}/2\mathbb{Z}$: the first is torsion-free the other is not – they are not isomorphic).
(2) The category of all dense linear orders without endpoints is categorical in $\aleph_0$, but not in any uncountable cardinal.
(3) The category of sets is categorical in every infinite cardinal.
(4) The category of vector spaces over $\mathbb{Q}$ is categorical in every uncountable cardinal (but not in $\aleph_0$: consider the spaces of dimension 1 and 2).
(5) The category of all algebraically closed fields of a fixed characteristic is categorical in every uncountable cardinal (but again not in $\aleph_0$).
(6) [LRV19b, 6.3] The category of all Hilbert spaces is categorical in every uncountable size (but not in every uncountable cardinal).

A classical theorem of Morley [Mor65] says that for a countable first-order theory T, if $\mathbf{Elem}(T)$ is categorical in some uncountable cardinal, then it is categorical in all uncountable cardinals. A proof can be sketched as follows: derive the existence of a stable independence notion from categoricity, then use it to build enough of a dimension theory to transfer categoricity. While it seems that having a single object of a given size is a relatively rare occurrence, it seems to be a useful test case as tools developed in proofs of categoricity transfers are often much more general. For example, the category of abelian groups with monomorphisms still has a stable independence notion [LRV19a, 5.3]. For this reason, many generalizations of Morley's theorem have been conjectured. One variation is:

Conjecture 5.16 (Shelah's eventual categoricity conjecture)**.** If an AEC is categorical in a proper class of cardinals, then it is categorical on a tail of cardinals.

Note that we have strengthened the starting assumption to categoricity on a proper class. Often, assuming categoricity in a single "high-enough" cardinal, where "high-enough" has some effective meaning, is enough. Note also that the conjecture is open even for accessible categories (although I suspect it should be wrong in full generality there). The eventual categoricity conjecture for AECs implies the eventual categoricity conjecture for *finitely* accessible categories [BR12, 5.8(2)].

It should be noted that multidimensional stable independence notions were introduced by Shelah to make progress on categoricity questions in $\mathbb{L}_{\omega_1,\omega}$ (in a very different form and a much more special situation than described here), see [She83a, She83b]. Very recently, multidimensional independence was used by Shelah and myself to prove the eventual categoricity conjecture in AECs, assuming the existence of a proper class of strongly compact cardinals [SV].

So how exactly is multidimensional stable independence used? It is a complicated story, that I do not have space to tell here. One idea is that if we have some version of multidimensional stable independence notion just for objects of size λ, then we can lift it up to a multidimensional stable independence notion on the whole category. Another key result is that once we have a multidimensional stable independence notion, we can get a much better approximation to existence of pushouts. In particular, among the independent amalgams of a given span, there will be a

weakly initial one, and it will be unique up to (not necessarily unique) isomorphism. Such an object is called a *prime object* (over the span) by model theorists. In this sense, every span has a "very weak" colimit. This is enough to obtain a notion of generation that helps build a dimension theory.

Finally, let's note that categoricity is a 0-dimensional uniqueness property, and we may want to know higher dimensional versions (even before proving existence of prime objects). For the two-dimensional case, instead of building amalgam that are, in the sense above, minimal, we will want to get amalgams that are maximal, in the sense of being very homogeneous (i.e. with a lot of injectivity). This is given by the notion of a *limit object*, which we survey (in the one dimensional case) in the next section.

6. Element by element constructions in concrete categories

The framework of abstract elementary classes is very useful to perform element by element constructions: we can use the concreteness to build certain objects "point by point". We really will just use that we work in an accessible category with concrete directed colimits and all morphisms embedding (so nothing about the vocabulary will be needed). The reader uncomfortable with logic can think about this more general case.

We first need a definition of a "point" and what it means for two points to be the same. This depends of course on a base set, e.g. in the category of fields, the elements $e^{1/2}$ and $e^{1/4}$ are the same over $\mathbb{Q}$ but not the same over $\mathbb{Q}(e)$. The equivalence class of a point over a given base will be called a *type*[24]. The definition for AECs is due to Shelah, but in the wider setup of concrete categories it was first explored by Lieberman and Rosický [LR16, 4.1]:

Definition 6.1. Let $\mathbf{K} = (\mathcal{K}, U)$ be a concrete category and let A be an object of $\mathcal{K}$. The *category of points over* A, $\mathbf{K}^*_A$, is defined as follows:

- Its objects are pairs (f, a), where $f : A \to B$ and $a \in UB$.
- A morphism from $(A \xrightarrow{f} B, a)$ to $(A \xrightarrow{g} C, b)$ is a $\mathcal{K}$-morphism $B \xrightarrow{h} C$ such that $hf = g$ and $h(a) = b$.

Remark 6.2. This is simply a pointed version of the category of cocones over the diagram consisting of only A. Morphisms between two cocones must respect points. Note that we wrote $h(a)$ instead of the more pedantic $(Uh)(a)$.

Definition 6.3. Let $\mathbf{K} = (\mathcal{K}, U)$ be a concrete category and let A be an object of $\mathcal{K}$. Two points (f, a), (g, b) over A *have the same type* if they are connected in the category of points over A (Definition 6.1). The *type* of a point over A is just its equivalence class under the relation of having the same type (in the category of points over A).

[24]In the AEC literature, the terms "Galois types" or "orbital types" are used. This is to avoid confusion with the model-theoretic syntactic types. However, there is not really a Galois theory and types are not necessarily orbits. Moreover in the first-order case the syntactic types are the same as the orbital types, and we will never refer to syntactic type. Thus we prefer the simpler terminology.

Remark 6.4. If $\mathbf{K}$ has amalgamation, then the category of points over A also has amalgamation. Thus two points $(A \xrightarrow{f} B, a)$ and $(A \xrightarrow{g} C, b)$ have the same type exactly when they are jointly connected in the category of points over A (see Lemma 5.3). Explicitly, this means there exists $B \xrightarrow{h_1} D$ and $C \xrightarrow{h_2} D$ such that the diagram below commutes and moreover $h_1(a) = h_2(b)$:

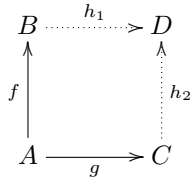

Note that if $(A \xrightarrow{f} B, a)$ is a point in an AEC, then after some "renaming", it has the same type as some point $(A \xrightarrow{i} C, b)$, where i is an inclusion (i.e. $A \leq_{\mathbf{K}} C$). It can sometimes be convenient (to avoid keeping track of morphisms) to just look at points induced by inclusions, and this gives a slightly simpler notation for types. Nevertheless, we will not follow this approach here.

Definition 6.5. Let $\mathbf{K}$ be an AEC and let[25] $M \in \mathbf{K}$.

(1) Let $\mathbf{S}(M)$ be the collection of all types over M.
(2) For $M \leq_{\mathbf{K}} N$, a type $p \in \mathbf{S}(M)$ is *realized in* N if there exists a point $(M \xrightarrow{i} N, b)$ whose type over M is p (here, i is the inclusion).

Remark 6.6. The following are easy exercises from the definitions. Let $\mathbf{K}$ be an AEC, $M \leq_{\mathbf{K}} N$.

(1) If $p \in \mathbf{S}(M)$ is realized in N and $N \leq_{\mathbf{K}} N'$, then p is realized in N'.
(2) If $p \in \mathbf{S}(M)$, then there exists an extension of M in which p is realized. Moreover, by the smallness axiom, we can ensure that extension has size at most $\mathrm{LS}(\mathbf{K}) + |UM|$. In particular, $\mathbf{S}(M)$ is a set.
(3) If $q \in \mathbf{S}(M_0)$, $M_0 \leq_{\mathbf{K}} M$, and $\mathbf{K}$ has the amalgamation property, then there exists an extension of M in which q is realized (and we can similarly ensure the extension has size at most $\mathrm{LS}(\mathbf{K}) + |UM|$).

How can types allow us to construct category-theoretically interesting objects? The following is an interesting definition:

Definition 6.7. Let $\mathbf{K}$ be an AEC, let λ be an infinite cardinal, and let $M \leq_{\mathbf{K}} N$ both be in $\mathbf{K}$. We say that N is λ-*universal over* M if for any given $M \xrightarrow{f} M'$ such that M' has size strictly less than λ, there exists $M' \xrightarrow{g} N$ such that the following diagram commutes (by convention, we will not label inclusions):

[25]We shift the notation to use M and N instead of A and B – this is to emphasize that the objects are structures (models).

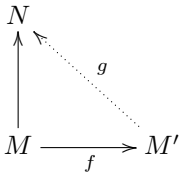

When $\lambda = |UM|^+$, we omit it and say that N is universal over M.

Remark 6.8. By renaming, we can assume without loss of generality that f is also an inclusion. Recall also that "size" in AECs means the same as "cardinality of the universe" (see 3.16(5)).

In an AEC with amalgamation, we can build universal objects categorically by a simple exhaustion argument. However, in typical cases, N above will be much bigger in size than M. For example, if M has size λ and $|\mathbf{S}(M)| > \lambda$, then because a universal object over M must realize all types over M, it must have at least λ^+-many elements. What if, on the other hand, $|\mathbf{S}(M)| = \lambda$? This condition is given a name (recall that we use $\mathbf{K}_\lambda$ to denote the class of objects of size λ):

Definition 6.9. An AEC is said to be λ-*stable* (or *stable in* λ) if $|\mathbf{S}(M)| = \lambda$ for any $M \in \mathbf{K}_\lambda$.

The name "stable" refers to the same kind of model-theoretic stability as "stable independence notion". In fact, in the first-order case, being stable on a proper class of cardinals is equivalent to the existence of a stable independence notion. The reader can think of stability as saying that "most" elements over a given base are transcendentals (i.e. all have the same type). In general, the existence of a stable independence notion in an AEC implies stability in certain cardinals [LRV19a, 8.16]. We will show that, assuming stability in λ, we can build universal extensions of the same cardinality. The result is due to Shelah [She09a, II.1.16] but the proof we give here is new and, as opposed to previous proofs (see [LR16, 6.2]) does *not* use the coherence axiom.

Theorem 6.10. Let $\mathbf{K}$ be an AEC and let $\lambda \geq \text{LS}(\mathbf{K})$. Assume that $\mathbf{K}$ is stable in λ and $\mathbf{K}_\lambda$ has amalgamation. For any $M \in \mathbf{K}_\lambda$, there exists $N \in \mathbf{K}_\lambda$ such that $M \leq_{\mathbf{K}} N$ and N is universal over M.

Proof. We first build $\langle M_i : i \leq \lambda \rangle$ increasing continuous in $\mathbf{K}_\lambda$ (i.e. $i < j$ implies $M_i \leq_{\mathbf{K}} M_j$, all objects are in $\mathbf{K}_\lambda$, and at limit ordinals we take unions) such that $M_0 = M$ and M_{i+1} realizes all types over M_i for all $i < \lambda$. This is possible: the only step to implement is the successor step, so assume that M_i is given and we have to build M_{i+1}. We know that $|\mathbf{S}(M_i)| = \lambda$ by stability, so list the types as $\langle p_j : j < \lambda \rangle$. We build $\langle M_{i,j} : j \leq \lambda \rangle$ increasing continuous in $\mathbf{K}_\lambda$ such that $M_{i,0} = M_i$ and $M_{i,j+1} \in \mathbf{K}_\lambda$ realizes p_j for all $j < \lambda$. This is possible by Remark 6.6 and at the end we can set $M_{i+1} = M_{i,\lambda}$.

We now prove that M_λ is universal over M_0. This is the hard part of the argument. Fix $M_0 \xrightarrow{f_0} M_0'$, with $M_0' \in \mathbf{K}_\lambda$. We want to find $g : M' \to M_\lambda$ with gf_0 fixing M_0. For each $i \leq j \leq \lambda$, we will build $M_i \xrightarrow{f_i} M_i'$ and $M_i' \xrightarrow{h_{ij}} M_j'$ such that for

$i \leq j \leq k \leq \lambda$, $M'_i \in \mathbf{K}_\lambda$, $h_{ik} = h_{jk}h_{ij}$, $h_{ii} = \mathrm{id}_{M'_i}$, and $h_{ij}f_i = f_j \restriction M_i$. That is, the following commutes:

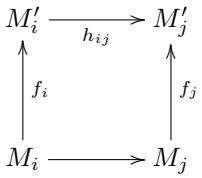

We further require this diagram to be smooth (i.e. at limit ordinals k, f_k is given by the colimit of $(f_i, h_{ij})_{i \leq j < k}$). To go from i to $i+1$, we make sure that the type of a certain point $(M_i \xrightarrow{f_i} M'_i, a_i)$ is realized in M_{i+1}: there is $b_i \in UM_{i+1}$ so that $(M_i \to M_{i+1}, b_i)$ has the same type as (f_i, a_i). This type equality is witnessed by f_{i+1} and $h_{i,i+1}$. In particular, $h_{i,i+1}(a_i) = f_{i+1}(b_i)$, so $h_{i,i+1}(a_i)$ is in the range of f_{i+1}. We make sure to choose the a_i's in such a way that, in the end, we catch our tail: $\{h_{j,\lambda}(a_j) : j < \lambda\} = UM'_\lambda$. This is possible by some bookkeeping (essentially just using that $|\lambda \times \lambda| = \lambda$). We give the details (for a more general setup) in the appendix, see Theorem A.11.

The construction we just described ensures that f_λ is a surjection, hence an isomorphism! Thus $g := f_\lambda^{-1} h_{0\lambda}$ is the desired embedding of M'_0 into M_λ. □

Remark 6.11. A similar argument (see Theorem A.12) establishes the *model-homogeneous = saturated lemma* [She09a, II.1.14]: when $\lambda \geq \mathrm{LS}(\mathbf{K})$, objects that realize all types over every substructure of size λ will in fact be universal over every substructure of size λ.

Once we know we can build universal extensions of the same cardinality, we can consider iterating them, and we get to the notion of a *limit object*:

Definition 6.12. Let $\mathbf{K}$ be an AEC, let $M \leq_{\mathbf{K}} N$, let $\lambda \geq \mathrm{LS}(\mathbf{K})$, and let $\delta < \lambda^+$ be a limit ordinal. We say that N is (λ, δ)-*limit over* M if there exists an increasing continuous chain $\langle M_i : i \leq \delta \rangle$ in $\mathbf{K}_\lambda$ such that $M_0 = M$, $M_\delta = N$, and M_{i+1} is universal over M_i for all $i < \delta$.

Theorem 6.10 establishes that, if $\mathbf{K}$ is stable in λ, there exists limit models. It turns out that limit models have uniqueness properties as well. For example, a back and forth argument establishes that for $\delta_1, \delta_2 < \lambda^+$, if M_ℓ is (λ, δ_ℓ)-limit over M for $\ell = 1, 2$ and $\mathrm{cf}(\delta_1) = \mathrm{cf}(\delta_2)$, then there exists an isomorphism from M_1 to M_2 which fixes M.

What if $\mathrm{cf}(\delta_1) \neq \mathrm{cf}(\delta_2)$? The question of *uniqueness of limit objects* asks whether the above is still true in this case (this can be thought of as a one-dimensional version of categoricity). For model theorists, we note that a positive answer is equivalent to superstability for first-order theories (thus limit objects provide a good way to define superstability categorically). More generally, in a stable theory T, limit models will be isomorphic exactly when their cofinality is at least $\kappa(T)$. In any case, proving uniqueness of limit models requires a good notion of stable independence for elements, allowing one to build very independent objects point by points. This is another (much harder than in Theorem 6.10) example of the use

of "points" in the theory of AECs. See [GVV16] for an exposition of some result in the theory of limit objects in AECs, though stronger theorems have now been found, e.g. [Van16]. A state of the art result from categoricity is [Vas17b, 5.7(2)] (an AEC $\mathbf{K}$ has *no maximal objects* if for any $M \in \mathbf{K}$ there exists $N \in \mathbf{K}$ with $M \leq_{\mathbf{K}} N$ and $M \neq N$):

Theorem 6.13. Let $\mathbf{K}$ be an AEC with amalgamation and no maximal objects. Let $\mu > \text{LS}(\mathbf{K})$. If $\mathbf{K}$ is categorical in μ, then for any $\lambda \in [\text{LS}(\mathbf{K}), \mu)$, any two limit objects of cardinality λ are isomorphic.

The reader may ask what limit models look like in specific categories. This is studied in recent work of Kucera and Mazari-Armida [MA20, KMA] for certain categories of modules.

We end this section by noting that stable independence notions themselves can be built "element by element". For this, one starts with a good notion of stable independence for types (what Shelah calls a good frame), and tries to "paste the points together". See [JS13] for an introduction to the theory of good frames and [Vas17a, Vas18a, Vas19] for recent applications to the categoricity spectrum problem. We state in particular the following result [Vas19]:

Theorem 6.14. Let $\mathbf{K}$ be an AEC with amalgamation. Assume the weak generalized continuum hypothesis: $2^\lambda < 2^{\lambda^+}$ for all infinite cardinals λ. If $\mathbf{K}$ is categorical in some cardinal $\mu \geq \beth_{\left(2^{\text{LS}(\mathbf{K})}\right)^+}$, then $\mathbf{K}$ is categorical in all cardinals $\mu' \geq \beth_{\left(2^{\text{LS}(\mathbf{K})}\right)^+}$, and moreover there is a stable independence notion on the category $\mathbf{K}_{\geq \beth_{\left(2^{\text{LS}(\mathbf{K})}\right)^+}}$.

Note that the statements of the last two theorems are purely category-theoretic, in the sense that their statement does not use concreteness, points, etc (if the reader is worried about whether the definition of an AEC is category-theoretic, they can look at the results for the special case of a finitely accessible category with all morphisms monos, see Theorem 3.17). I am not aware of "purely category-theoretic" proofs of any statements like this, so I suspect that the element by element methods used to study AECs can be useful even if one is interested only in categorical problems.

7. Some known results and open questions

7.1. **Questions on categoricity.** Shelah's eventual categoricity conjecture for AECs (Conjecture 5.16) is still open, in ZFC, but is known to hold from many different types of (quite mild) assumptions. In many cases, we can say more about the "high-enough" bound and even (in (3) below) list exactly what the possibly categoricity spectrums are. For example:

(1) (Large cardinal axioms, [SV]) Assuming there is a proper class of strongly compact cardinals, an AEC categorical in a proper class of cardinals is categorical on a tail of cardinals.
(2) (Large cardinal axioms plus cardinal arithmetic, [SV]) Assuming there is a proper class of *measurable* cardinals and WGCH[26], an AEC categorical in a proper class of cardinals is categorical on a tail of cardinals.

[26]Recall that this means that $2^\lambda < 2^{\lambda^+}$ for all infinite cardinals λ.

(3) (Amalgamation plus cardinal arithmetic [Vas19, 9.7]) Assume WGCH. Let $\mathbf{K}$ be a large AEC with amalgamation (and $\mathbf{K}_{<\mathrm{LS}(\mathbf{K})} = \emptyset$). Exactly one of the following holds:
   (a) $\mathbf{K}$ is not categorical in any cardinal above $\mathrm{LS}(\mathbf{K})$.
   (b) There exists $n \leq m < \omega$ such that $\mathbf{K}$ is categorical in all cardinals in $[\mathrm{LS}(\mathbf{K})^{+n}, \mathrm{LS}(\mathbf{K})^{+m}]$ and no other cardinals.
   (c) There exists $\chi < \beth_{\left(2^{\mathrm{LS}(\mathbf{K})}\right)^+}$ such that $\mathbf{K}$ is categorical in all $\mu \geq \chi$ (and no other cardinals).

   It is known that examples exist of all three types.

(4) (No maximal objects plus strong cardinal arithmetic [Vas19, 10.14]) Assume $\diamondsuit_S$ for every stationary set S (this holds for example if $\mathrm{V} = \mathrm{L}$). An AEC with no maximal objects categorical in a proper class of cardinals is categorical on a tail of cardinals.

(5) (Universal class [Vas17c, Vas17d]) If a universal class $\mathbf{K}$ is categorical in some $\mu \geq \beth_{\beth_{\left(2^{\mathrm{LS}(\mathbf{K})}\right)^+}}$, then it is categorical in all $\mu' \geq \beth_{\beth_{\left(2^{\mathrm{LS}(\mathbf{K})}\right)^+}}$. This holds more generally for multiuniversal classes [ABV19]

Other approximations to categoricity (for example from tameness, a locality property of types that has not been discussed here) are in [She99, GV06b, GV06a]. Note that given any finitely accessible category $\mathcal{K}$, we have that $\mathcal{K}_{mono}$ is an AEC and the embedding $\mathcal{K}_{mono} \to \mathcal{K}$ preserves and reflects presentability ranks, hence categoricity. Thus the results above are, in particular, valid for any finitely accessible category.

Question 7.1.

(1) Is the eventual categoricity conjecture for AECs provable in, or at least consistent with, ZFC?
(2) Is there a counterexample to eventual categoricity for accessible categories? Is it true at least for accessible categories with directed colimits? Can we generalize the proof of eventual categoricity for continuous first-order logic [SU11] to category-theoretic setups?
(3) (Diliberti) Can one give "purely category-theoretic" proofs of categoricity transfers, at least in simple cases (for example, for eventual categoricity in locally finitely presentable categories)?
(4) Is eventual categoricity true for locally presentable categories?

7.2. **Questions on set-theoretic aspects.**

Question 7.2.

(1) Is the presentability rank of every high-enough object in an accessible category always a successor?
(2) What is an example of a large accessible category that is not LS-accessible[27]?
(3) If a category is λ-accessible, for λ big-enough, does it have an object of size λ? (see Theorem 4.4)
(4) Is any large accessible category with directed colimits LS-accessible? (This is asked already in [LR16]).

[27]Such an example would yield to a failure of eventual categoricity: take the coproduct with the category of sets [BR12, 6.3].

(5) What other types of compactness can be expressed using accessible functors (Section 4.2)?
(6) Are there other local methods than Theorem 4.16 that allow us to prove amalgamation in ZFC?

Regarding the second question, we can give a version that does not mention category-theoretic sizes: does there exist a μ-AEC $\mathbf{K}$ such that, for a proper class of cardinals λ with $\lambda < \lambda^{<\mu}$, $\mathbf{K}$ is categorical in $\lambda^{<\mu}$, and $\mathbf{K}$ has no objects of cardinality in $[\lambda, \lambda^{<\mu})$? See [LRV19b, 4.16] for why such a μ-AEC is not LS-accessible.

7.3. **Questions on accessible categories vs AECs.**

Question 7.3.

(1) What is the role of the coherence axiom of AECs?
(2) Is there a short characterization of AECs that is completely category-theoretic, in the sense that it does not refer to concrete functors (as in [LR16]), or embeddings into (variations on) category of structures (as in [BR12, 5.7])?
(3) Is there a natural logic axiomatizing AECs (see [BV19, §4])?

7.4. **Questions on stable independence.**

Question 7.4.

(1) What are more occurrences of stable independence in mainstream mathematics?
(2) How does stable independence interact with accessible functors?
(3) Can one characterize when the uniqueness of limit objects holds in terms of properties of stable independence (say in accessible categories with directed colimits)?
(4) Does stable independence tell us anything interesting about metric classes (Banach spaces, Hilbert spaces, etc.)? See [BYU10] on first-order stability theory for metric classes.
(5) Is there a theory of independence in accessible categories mirroring that of independence in simple unstable first-order theories (see [KP97])? What about other model-theoretic classes of unstable theories?

7.5. **Questions on the model theory of AECs.** The questions below are more technical, and cannot be understood just from the material of this paper. I chose to collect them here for the convenience of the expert reader:

Question 7.5.

(1) If $\mathbf{K}$ is a λ-superstable AEC, does it have λ-symmetry (see [Van16])? More generally, what are the exact relationships between uniqueness of limit objects in λ, λ-symmetry, and λ-superstability?
(2) If $\mathbf{K}$ is a λ-stable AEC and $\mathbf{K}_\lambda$ has amalgamation and no maximal objects, can we prove uniqueness of long-enough limit objects in λ, as in [BV], but without using a continuity property for splitting?

(3) Is the following stability spectrum theorem true? In a large stable LS($\mathbf{K}$)-tame AEC with amalgamation, there exists $\chi < \beth_{(2^{\text{LS}(\mathbf{K})})^+}$ such that for all $\lambda \geq \beth_{(2^{\text{LS}(\mathbf{K})})^+}$, $\mathbf{K}$ is stable in λ if and only if $\lambda = \lambda^{<\chi}$. Approximations are in [Vas18b].

8. Further reading

In addition to all the references cited already, we mention some resources that may help newcomers become acquainted with the field. We repeat again that Makkai-Paré [MP89] and especially Adámek-Rosický [AR94] are the standard textbooks on accessible categories. The category-theoretic singular compactness theorem (Theorem 4.11) appears in [BR16], which has numerous examples and explanations. The relationship between accessible categories and abstract elementary classes is investigated in, for example, [Lie11, BR12, LR16]. The Beke-Rosický paper, specifically, started the abstract study of category-theoretic sizes continued in [LRV19b, LRVb]. The category-theoretic notion of stable independence is introduced in [LRV19a], and a follow-up (establishing the connection with cofibrant generation) is [LRVc].

To become acquainted with abstract elementary classes specifically, two introductions are Grossberg's survey [Gro02] and Baldwin's book [Bal09]. Recently, classes about AECs were given at Harvard University by both Will Boney and myself, and both classes had lecture notes [Bona, Vas] that give an updated take on the subject. The survey about tame AECs [BV17] may also be helpful to get acquainted with the literature. When one starts studying independence for types, Shelah's good frames, a "pointed" and localized version of stable independence, become an unavoidable concept. Currently, the best introduction to good frames is the paper of Jarden and Shelah [JS13]. Finally, it is impossible not to mention Shelah's two volume book [She09a, She09b] which has a very interesting and readable introduction, and is a gold mine of deep (but not always easily readable) results on good frames and AECs generally.

Appendix A. Forcing and construction categories

I give here a general category-theoretic framework for point by point "exhaustion arguments" such as building algebraic closure of fields (or more generally saturated models), or proving a given extension realizing all types many times is universal as in Theorem 6.10. It is also a natural framework in which to understand set-theoretic forcing. To the best of my knowledge, this is new.

Definition A.1. A *construction category* is a triple $\mathbf{K} = (\mathcal{K}, U, U_0)$, where:

(1) $\mathcal{K}$ is a category.
(2) $U : \mathcal{K} \to \mathbf{Set}$ is a faithful functor.
(3) $U_0 : \mathcal{K} \to \mathbf{Set}$ is a faithful[28] subfunctor of U: a faithful functor such that for all morphisms $A \xrightarrow{f} B$, $U_0 A \subseteq UA$ and $U_0 f = (Uf) \restriction U_0 A$.

[28]It seems the faithfulness of U and U_0 is never used.

The idea is that, for an object A, UA gives the elements that could "conceivably" be constructed at some point (e.g. $\mathcal{K}$ could be the category of fields and UA give the polynomials with coefficients from A, see Example A.10), while U_0A gives the elements that have been constructed already (e.g. in the example of the category of fields, U_0A could give the polynomials with coefficients from A that have a root in A). We will be trying to find an object (called *full* below) where everything that can be constructed in some extension has been constructed already. It may help the reader to think of the category $\mathcal{K}$ as a partially ordered set.

Example A.2 (Set-theoretic forcing)**.** Let $\mathbb{P}$ be a partially ordered set (we think of it also as a category). A notion of forcing for $\mathbb{P}$ (e.g. in the sense of [Kei73]) associates to each $p \in \mathbb{P}$ a set of sentences that it "forces" in such a way that if $p \leq q$ then q forces more sentences than p. Setting U_0p to be the formulas that p forces, and Up to be the set of all formulas, we obtain a construction category.

Definition A.3. Let $\mathbf{K}$ be a construction category.

(1) Given an object A and an element $x \in UA$, we say that x is *constructed by stage* A if $x \in U_0A$. We say that x is *constructible from* A if there exists a morphism $A \xrightarrow{f} B$ so that $f(x)$ is constructed by stage B.
(2) A directed diagram $D : I \to \mathbf{K}$ with maps $D_i \xrightarrow{d_{ij}} D_j$ is *full* whenever the following is true: for any $i \in I$ and any $x \in UD_i$, *if* for all $j \geq i$, $d_{ij}(x)$ is constructible from D_j, *then* there exists $j \geq i$ such that $d_{ij}(x)$ is constructed by stage D_j.
(3) For an object A and a set $X \subseteq UA$, we say that A is *full for* X if any $x \in X$ that is constructible from A is constructed by stage A. We say that A is *full* if it is full for UA.

The following are basic properties of the definitions:

Remark A.4. Let $\mathbf{K}$ be a construction category, $A \xrightarrow{f} B$ be a morphism, and $x \in UA$.

(1) If x is constructed by stage A, then x is constructible from A.
(2) If x is constructed by stage A then $f(x)$ is constructed by stage B.
(3) If $f(x)$ is constructible from B, then x is constructible from A.
(4) If A is an amalgamation base (Definition 4.9) and x is constructible from A, then $f(x)$ is constructible from B.
(5) If A is full for X, then B is full for $f[X]$.
(6) An object A is full if and only if the corresponding directed diagram with one object is full.

Lemma A.5. Let $\mathbf{K}$ be a construction category and let $D : I \to \mathcal{K}$ be a directed diagram with maps $D_i \xrightarrow{d_{ij}} D_j$.

(1) If D is full, then for any cocone $(D_i \xrightarrow{f_i} A)_{i \in I}$ for D and any $i \in I$, $f_i^{-1}[U_0A] \cap UD_i \subseteq \bigcup_{j \geq i} d_{ij}^{-1}[U_0D_j]$.
(2) An object A is full if and only if for any morphism $A \xrightarrow{f} B$, $UA \cap f^{-1}[U_0B] = U_0A$.

(3) If D is full and $(D_i \xrightarrow{f_i} A)_{i\in I}$ is a cocone for D, then A is full for $\bigcup_{i\in I} f_i[UD_i]$. In particular, if $UA = \bigcup_{i\in I} f_i[UD_i]$ then A is full.

Proof.

(1) Let $(D_i \xrightarrow{f_i} A)_{i\in I}$ be a cocone for D and let $i \in I$. Let $x \in f_i^{-1}[U_0A] \cap UD_i$. Let $y := f_i(x)$ (so $y \in U_0A$). For any $j \geq i$, f_j (and the fact that $y \in U_0A$) witnesses that $d_{ij}(x)$ is constructible from D_j. By definition of a full diagram, there exists $j \geq i$ such that $d_{ij}(x)$ is constructed by stage D_j, i.e. $d_{ij}(x) \in U_0D_j$. Thus $x \in d_{ij}^{-1}(U_0D_j)$.

(2) If A is full, then the previous part gives that for any morphism $A \xrightarrow{f} B$, $UA \cap f^{-1}[U_0B] \subseteq U_0A$. The reverse inclusion is immediate because U_0 is a subfunctor of U. Conversely, assume that for any morphism $A \xrightarrow{f} B$, $UA \cap f^{-1}[U_0B] = U_0A$. Let $x \in UA$ be constructible from A. This means there exists $A \xrightarrow{f} B$ such that $f(x) \in U_0B$, i.e. $x \in f^{-1}[U_0B]$. By assumption, $x \in U_0A$, so x is constructed by stage A.

(3) Assume that D is full. Let $y \in \bigcup_{i\in I} f_i[UD_i]$ be constructible from A. There exists $i \in I$ such that $y = f_i(x)$ for some $x \in UD_i$. Since y is constructible from A, x is constructible from D_i. In fact, for any $j \geq i$, f_j witnesses that $d_{ij}(x)$ is constructible from D_j. Since D is full, there exists $j \geq i$ so that $d_{ij}(x) \in U_0D_j$. Thus $y = f_jd_{ij}(x) \in f_j[U_0D_j] \subseteq U_0A$, so y is constructed by stage A.

□

Although we will in the end mostly be interested in full objects, it is often helpful (and easier) to first verify that a full *diagram* exists. Its colimit will then usually be the desired full object. In order for full diagrams to exist, objects should not be too big. We will measure size using cofinality of a certain "order of construction".

Definition A.6. Let $\mathbf{K}$ be a construction category and let A be an object.

(1) Define a relation $\leq$ on UA as follows: $x \leq y$ if for all $A \xrightarrow{f} B$, if $f(y)$ is constructed by stage B, then $f(x)$ is constructed by stage B.
(2) Let $\|A\|$ denote the cofinality of $(UA, \leq)$.

Note that $\leq$ is a preorder on UA, and $\|A\| \leq |UA|$. The following will be our main tool to verify existence of full diagrams.

Theorem A.7 (Existence of full diagrams)**.** Let $\mathbf{K}$ be a construction category and let λ be an infinite cardinal. If $\|A\| \leq \lambda$ for all objects A, and (for $\alpha < \lambda$) any α-indexed diagram has a cocone, then $\mathbf{K}$ has a λ-indexed full diagram.

This will be a consequence of the following more general version (used in the proof of Theorem A.11), where we only require diagrams with "small" objects (according to a certain rank) to have an upper bound.

Lemma A.8. Let $\mathbf{K}$ be a construction category, let λ be an infinite cardinal, and let $R : \mathbf{K} \to \lambda$ be given. If:

(1) $\|A\| \leq \lambda$ for all objects A.

(2) For any object A and any $x \in UA$ that is constructible from A, there exists $A \xrightarrow{f} B$ such that $RB \leq RA + 1$ and $f(x)$ is constructed by stage B.
(3) For every $\alpha < \lambda$ and every diagram $D : \alpha \to \mathbf{K}$, if $RD_i < \alpha$ for all $i < \alpha$, then D has a cocone $(D_i \to A)_{i<\alpha}$ with $RA \leq \alpha$.

Then there is a λ-indexed full diagram $D : \lambda \to \mathbf{K}$ with $RD_i \leq i$ for all $i < \lambda$.

Proof of Theorem A.7. Apply Lemma A.8, with $RA = 0$ for all objects A. □

Proof of Lemma A.8. First, we fix a function $F : \lambda \to \lambda \times \lambda$ such that for each pair (α, β) in $\lambda \times \lambda$, there exists unboundedly-many $i < \lambda$ so that $F(i) = (\alpha, \beta)$. This can be done by first partitioning λ into λ-many disjoint pieces of cardinality λ, then bijecting each of these pieces with $\lambda \times \lambda$. Write (α_i, β_i) for $F(i)$.

We inductively build objects $\langle D_i : i < \lambda \rangle$, morphisms $\langle D_i \xrightarrow{d_{ij}} D_j : i \leq j < \lambda \rangle$, and sequences $\langle x_{i,j} : i, j < \lambda \rangle$ such that:

(1) $RD_i \leq i$ for all $i < \lambda$.
(2) $d_{ii} = \mathrm{id}_{D_i}$, $d_{jk} d_{ij} = d_{ik}$ for all $i \leq j \leq k < \lambda$.
(3) For each $i < \lambda$, $\langle x_{i,j}; j < \lambda \rangle$ enumerates the elements of a cofinal set of UD_i (perhaps with repetitions).
(4) For each $i < \lambda$, if $\alpha_i \leq i$ and $d_{\alpha_i,i}(x_{\alpha_i,\beta_i})$ is constructible from D_i, then $d_{\alpha_i,i+1}(x_{\alpha_i,\beta_i})$ is constructed by stage D_{i+1}. Moreover, if there exists an object that is full over D_i, then D_{i+1} is full over D_i.

This is enough: Let $D : \lambda \to \mathcal{K}$ be the diagram with maps d_{ij}. Then D is the desired full diagram. Indeed, let $i < \lambda$, and let $x \in UD_i$ be such that $d_{ij}(x)$ is constructible from D_j for all $j \geq i$. Let $\alpha := i$. Without loss of generality (using cofinality), $x = x_{\alpha,\beta}$ for some $\beta < \lambda$. By construction, there exists $j \geq i$ such that $(\alpha_j, \beta_j) = (\alpha, \beta)$. By assumption, $d_{i,j}(x)$ is constructible from D_j, so by (4), $d_{i,j+1}(x)$ is constructed by stage D_{j+1}, as desired.

This is possible: We proceed by induction on $j < \lambda$. Assume inductively that $\langle D_i : i < j \rangle$, $\langle d_{ii'} : i \leq i' < j \rangle$, and $\langle x_{i,i'} : i < j, i' < \lambda \rangle$ have been constructed. We will build D_j, $\langle d_{ij} : i \leq j \rangle$, and $\langle x_{j,j'} : j' < \lambda \rangle$. First assume that j is limit or zero. By assumption, the diagram $D : j \to \mathcal{K}$ with maps $\langle d_{ii'} : i \leq i' < j \rangle$ has a cocone $(D_i \xrightarrow{d_{ij}} D_j)_{i<j}$, with $RD_j \leq j$. Set $d_{jj} = \mathrm{id}_{D_j}$, and let $\langle x_{j,j'} : j' < \lambda \rangle$ be any enumeration of UD_j. Now assume that j is a successor: $j = i + 1$. If $\alpha_i > i$, or $\alpha_i \leq i$ but $d_{\alpha_i,i}(x_{\alpha_i,\beta_i})$ is not constructible from D_i, then set $B = D_i$, $f = \mathrm{id}_{D_j}$. If $\alpha_i \leq i$ and $d_{\alpha_i,i}(x_{\alpha_i,\beta_i})$ is constructible from D_i, then let $D_i \xrightarrow{f} B$ witness this, with $RB \leq i + 1$. Set $D_i = B$, $d_{i,j} = f$, and let $d_{j,j} = \mathrm{id}_{D_j}$, $d_{i_0,j} = d_{i,j} d_{i_0,i}$ for all $i_0 < i$, and $\langle x_{j,j'} : j' < \lambda \rangle$ be any enumeration of UD_j. □

Even if we are unable to directly build full objects, we can also find a lot of them in continuous-enough chains:

Theorem A.9. Let λ be a regular uncountable cardinal and let $\mathbf{K} = (\mathcal{K}, U, U_0)$ be a construction category where $\mathcal{K}$ is just the ordered set λ and U-images of morphisms are inclusions. If $\|j\| < \lambda$ for all $j < \lambda$, then the set $\{j < \lambda \mid j \text{ is full for } \bigcup_{i<j} Ui\}$ is closed unbounded.

Proof. Let $C := \{j < \lambda \mid j \text{ is full for } \bigcup_{i<j} Ui\}$.

- C is closed: let $j < \lambda$ be a limit ordinal such that unboundedly-many $i < j$ are in C. Let $x \in \bigcup_{i<j} Ui$ be constructible from j. Pick $i' \in C \cap j$ such that $x \in \bigcup_{i<i'} Ui$. Then x is constructible from i' hence, by definition of C, x is constructed by stage i', hence by stage j.
- C is unbounded: let $\alpha < \lambda$. We build $\langle \alpha_n : n < \omega \rangle$ an increasing sequence of ordinals below λ such that for all $n < \omega$:
  (1) $\alpha_0 = \alpha$.
  (2) Any $x \in U\alpha_n$ that is constructible in α_n is constructed by stage α_{n+1}.

  This is possible: given α_n, for each $x \in U\alpha_n$ that is constructible in α_n, there exists a least $i_x < \lambda$ such that x is constructed by stage i_x. Since $\|\alpha_n\| < \lambda$ and λ is regular, there exists $\alpha_{n+1} > \alpha_n$ such that $\alpha_{n+1} \geq i_x$ for all $x \in U\alpha_n$ constructible in α_n.

  This is enough: let $\beta := \sup_{n<\omega} \alpha_n$. Since λ is regular uncountable, $\beta < \lambda$. Moreover, β is full for $\bigcup_{i<\beta} Ui$. Indeed, if $x \in \bigcup_{i<\beta} U_i$ is constructible from β, then there exists $n < \omega$ such that $x \in U\alpha_n$ and x is constructible from α_n, hence constructed by stage α_{n+1}, hence by stage β.

□

Many well known constructions can easily be seen as special cases of the theorems just stated:

Example A.10.

(1) (Zorn's lemma) If $\mathbb{P}$ is a partially ordered set where each chain has an upper bound, then $\mathbb{P}$ has a maximal element. This can be obtained by applying Theorem A.7 to $\lambda = |\mathbb{P}| + \aleph_0$, $\mathbf{K} = (\mathbb{P}, U, U_0)$, where $Up = \mathbb{P}$ for all $p \in \mathbb{P}$ and $U_0 p = \{q \in \mathbb{P} \mid q \leq p\}$. Any upper bound to the full diagram gives the desired maximal element.
(2) (Existence of generics) Let $\mathbb{P}$ be the poset of all finite partial functions from ω to $\{0, 1\}$. Let $\mathbf{K} = (\mathbb{P}, U, U_0)$, where $Us = \omega$, $U_0 s = \text{dom}(s)$. Then a full diagram $D : \omega \to \mathbf{K}$ corresponds to a (total) function $f : \omega \to \{0, 1\}$. This generalizes to the existence of generics, in the sense of set-theoretic forcing [Jec03, 14.4].
(3) (Existence of algebraic closure) Every field F has an algebraic closure: take $\lambda = |F| + \aleph_0$, $\mathbf{K} = (\mathcal{K}, U, U_0)$ where $\mathcal{K}$ is the category of field extensions of F of cardinality at most λ, UA is the set of all polynomials with coefficients from A, and $U_0 A$ is the set of all such polynomials with a root in A. The colimit of the full diagram given by Theorem A.7 is full (Lemma A.5(3)), hence algebraically closed.
(4) (Existence of saturated models) Let $\mathbf{K}^*$ be an AEC with amalgamation and $\lambda > \text{LS}(\mathbf{K}^*)$ be a regular cardinal such that $\mathbf{K}^*$ is stable in λ and $\mathbf{K}^*_\lambda \neq \emptyset$. Let $\mathbf{K} = (\mathcal{K}, U, U_0)$, where $\mathcal{K}$ is the full subcategory of $\mathbf{K}^*$ with objects of cardinality λ, UA is the set of all types over a substructure of A, and $U_0 A$ is the set of all such types that are realized in A. Note that we may well have $|UA| > \lambda$, but by stability we still have that $\|A\| \leq \lambda$ (the types over A form a cofinal set of size λ). The colimit of the full diagram given by

Theorem A.7 is full, hence is a saturated object of $\mathbf{K}^*_\lambda$ (i.e. all types over substructure of cardinality strictly less than λ are realized).

(5) (Disjointness of filtrations on a club) If λ is a regular uncountable cardinal, $A \subseteq B$ are sets of cardinality λ, and $\langle A_i : i < \lambda \rangle$, $\langle B_i : i < \lambda \rangle$ are increasing continuous chains of subsets of cardinality strictly less than λ such that $A = \bigcup_{i<\lambda} A_i$ and $B = \bigcup_{i<\lambda} B_i$, then the set of $i < \lambda$ such that $A \cap B_i = A_i$ is closed unbounded (in particular, if $A = B$ then $A_i = B_i$ on a closed unbounded set). Indeed, let $\mathbf{K} = (\lambda, U, U_0)$, where $Ui = A_i \cup B_i$, $U_0 i = A_i \cap B_i$. By Theorem A.9, the set of i such that i is full is closed unbounded. Now, if i is full and $x \in A \cap B_i$, then $x \in A_j \cap B_j$ for some j, so x is constructible from i, so is in $A_i \cap B_i \subseteq A_i$. Conversely, if $x \in A_i$ then it is constructible from i, hence in $A_i \cap B_i \subseteq A \cap B_i$. Thus $A \cap B_i = A_i$.

(6) Similarly to the previous example, density of reduced towers (in the study of uniqueness models, see e.g. [GVV16, 5.5]) can be seen as describing the existence of a full object in an appropriate construction category.

Let's now give more details on the proof of Theorem 6.10:

Theorem A.11. Let $\mathbf{K}$ be an AEC, let $\lambda \geq \text{LS}(\mathbf{K})$ be such that $\mathbf{K}_\lambda$ has amalgamation, and let $\langle M_i : i \leq \lambda \rangle$ be increasing continuous in $\mathbf{K}_\lambda$ such that M_{i+1} realizes all types over M_i. Then M_λ is universal over M_0.

Proof. Let $\mathbf{K}^* = (\mathcal{K}^*, V, V_0)$ be defined as follows:

- The objects of $\mathcal{K}^*$ are morphisms $M_i \xrightarrow{f} M$ for $M \in \mathbf{K}_\lambda$, $i < \lambda$, such that $f \restriction M_0 = \text{id}_{M_0}$.
- A morphism from $M_i \xrightarrow{f} M$ to $M_j \xrightarrow{g} N$, with $i \leq j < \lambda$, is a map $h : M \to N$ such that the following diagram commutes:

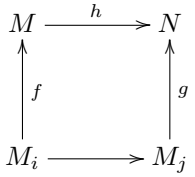

- $V(M_i \to M) = UM$, and the V-image of a morphism h from $M_i \to M$ to $M_j \to N$ is Uh (where U is the universe functor from $\mathbf{K}$ to **Set**)
- $V_0(M_i \xrightarrow{f} M) = Uf[M_i]$, and the V_0-image of a morphism h from $M_i \to M$ to $M_j \to N$ is $Uh \restriction f[M_i]$.

This is easily checked to be a construction category. Now let $M_i \xrightarrow{f} N$ be given, with $i < \lambda$. We show that any $x \in UN$ is constructible from f. Indeed, M_{i+1} realizes all types over M_i, so in particular it realizes the type of $(x, M_i \xrightarrow{f} N)$. Thus there is $x' \in UM_{i+1}$ and a commutative diagram:

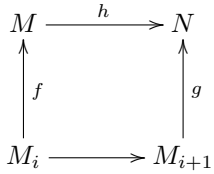

with $g(x') = h(x)$. In particular, $h(x) \in V_0(g)$, so is constructed by stage g.

Assume now that a diagram $D : \lambda \to \mathcal{K}^*$ is full. Let $M_j \xrightarrow{f} N$ be a colimit (in $\mathbf{K}$) of D (so $j \leq \lambda$). From the previous discussion, it is easy to see that f is surjective, hence an isomorphism. This then gives the result: for any $N_0 \in \mathbf{K}_\lambda$, $M_0 \xrightarrow{f_0} N_0$, the subcategory $\mathbf{K}^*_{f_0}$ of objects of $\mathbf{K}^*$ above f_0 satisfies the hypotheses of Lemma A.8 (with $R(M_i \to M) = i$), hence has a full diagram, whose colimit must therefore induce an embedding of N_0 into M_λ. □

We can similarly prove that model-homogeneous is equivalent to saturated (Remark 6.11):

Theorem A.12. Let $\mathbf{K}$ be an AEC, let $\lambda \geq \text{LS}(\mathbf{K})$, and assume that $\mathbf{K}_\lambda$ has amalgamation. If $M \in \mathbf{K}_{\geq\lambda}$ realizes all types over every substructure of size λ, then M is universal over every substructure of size λ.

Proof. Similar to the proof of Theorem A.11: this time the objects of the construction categories are maps $M_0 \to N_0$ with $M_0 \leq_{\mathbf{K}} M$ and M_0, N_0 both of size λ, and the rest of the definition is analogous. □

Appendix B. A very short introduction to first-order stability

For the convenience of the unacquainted reader, I give a quick and self-contained construction of stable independence in the first-order case, and derive two consequences: the equivalence of stability (in terms of counting types) with no order property, and the ability to extract indiscernibles from long-enough sequences. All the material in this appendix is well known but I am not aware of a place where it appears in such compressed form. I assume a very basic knowledge of model theory, but not previous knowledge of stability theory.

Throughout, we fix a complete first-order theory T with only infinite models in a vocabulary τ. For notational simplicity, we assume that $|T| = |\tau| + \aleph_0$. We fix a proper class sized "monster model" $\mathfrak{C}$ for T. This means that $\mathfrak{C}$ is universal and homogeneous, so for convenience we work inside $\mathfrak{C}$. We use the letters $\bar{a}, \bar{b}, \bar{c}$ for (possibly infinite) sequences of elements from $\mathfrak{C}$, $\bar{x}, \bar{y}, \bar{z}$ for (possibly infinite) sequences of variables, A, B, C for subsets of $\mathfrak{C}$, and M, N for elementary substructures of $\mathfrak{C}$. For a sequence $\bar{a}$, $\text{ran}(\bar{a})$ denotes its set of elements (i.e. its range when thought of as a function). We may write $A \cup B$ instead of AB, $A\bar{b}$ instead of $A \cup \text{ran}(\bar{b})$, etc. Formulas are denoted by $\phi(\bar{x}), \psi(\bar{x})$, where $\bar{x}$ is a sequence of variables that contains all free variables from ϕ (but may contain more – ϕ always has finitely-many free variables of course). We write $\models \phi[\bar{a}]$ instead of $\mathfrak{C} \models \phi[\bar{a}]$, which means that ϕ holds in $\mathfrak{C}$ when $\bar{a}$ replaces $\bar{x}$ (we are very casual with arities). As usual, we often do not

distinguish between M and its universe UM. We also abuse notation by writing $\bar{a} \in A$ instead of the more proper $\bar{a} \in {}^{<\infty}A$. We write $\mathbf{tp}(\bar{b}/A)$ (the type of $\bar{b}$ over A) for the set of formulas $\phi(\bar{x}, \bar{a})$, where $\bar{a} \in A$ and $\models \phi(\bar{b}, \bar{a})$. We write $\bar{b}_1 \equiv_A \bar{b}_2$ to mean that $\mathbf{tp}(\bar{b}_1/A) = \mathbf{tp}(\bar{b}_2/A)$. Recall that, by the compactness theorem, this holds if and only if there exists an automorphism f of $\mathfrak{C}$ that fixes A pointwise and sends $\bar{b}_1$ to $\bar{b}_2$. Thus this corresponds to the notion defined in Definition 6.3. We let $\mathbf{S}(A) := \{\mathbf{tp}(b/A) \mid b \in \mathfrak{C}\}$, and more generally $\mathbf{S}^\alpha(A) := \{\mathbf{tp}(\bar{b}/A) \mid \bar{b} \in {}^\alpha\mathfrak{C}\}$, $\mathbf{S}^{<\infty}(A) = \bigcup_\alpha \mathbf{S}^\alpha(A)$. Note that $|\mathbf{S}(A)| \leq 2^{|T|+|A|}$, and that if $A \subseteq B$ then there is a natural surjection of $\mathbf{S}(B)$ into $\mathbf{S}(A)$ (so $|\mathbf{S}(A)| \leq |\mathbf{S}(B)|$). For $p \in \mathbf{S}(B)$ and $A \subseteq B$, we write $p \restriction A$ for the restriction of p to $\mathbf{S}(A)$: the set of formulas from p with parameters in A.

It is time to define the classes "well-behaved" of theories we will work with. As in Definition 6.9, we will say that T is λ-*stable* (or *stable in* λ) if for any A of cardinality λ, $|\mathbf{S}(A)| = \lambda$ (of course, this is exactly the same as stability in λ in the sense of 6.9, in the AEC of models of T ordered by elementary substructure). We say that T is *stable* if it is stable in some cardinal $\lambda \geq |T|$. The following closely related property will play a key role:

Definition B.1. T has the *order property* if there exists a sequence $\langle \bar{a}_i : i < \omega \rangle$ and a formula $\phi(\bar{x}, \bar{y})$ such that $\models \phi[\bar{a}_i, \bar{a}_j]$ if and only if $i < j$.

It turns out that T is stable if and only if it does *not* have the order property. We prove one direction now. The other will be dealt with after we have constructed stable independence.

Theorem B.2. If T has the order property, then T is unstable.

Proof. Fix a cardinal $\lambda \geq |T|$, and fix a linear order I of cardinality λ with strictly more than λ Dedekind cuts (if $\lambda = \aleph_0$, the rationals are such an order; in general take σ minimal such that $\lambda < \lambda^\sigma$ and the set ${}^{<\sigma}\lambda$ ordered lexicographically will do the trick). Fix a formula $\phi(\bar{x}, \bar{y})$ witnessing the order property. Using the compactness theorem, there exists a $\langle \bar{a}_i : i \in I \rangle$ such that for all $i, j \in I$, $\models \phi[\bar{a}_i, \bar{a}_j]$ if and only if $i < j$. Since ϕ has finitely-many free variables we can of course assume wihout loss of generality that the $\bar{a}_i$'s are of finite length. Now each Dedekind cut of I induces a distinct type over $\bigcup_{i \in I} \text{ran}(\bar{a}_i)$, a set of size λ. Thus T is not stable in λ. □

We now define what it means for two sets to be "as independent as possible" over a base. For simplicity, the base is required to be a model. More advanced introductions investigate what happens when the base is an arbitrary set.

Definition B.3. We write $\bar{a} \underset{M}{\downarrow} \bar{b}$, and say that $\bar{a}$ *and* $\bar{b}$ *are independent over* M, if whenever $\bar{c} \in M$ and $\models \phi[\bar{a}, \bar{b}, \bar{c}]$, there exists $\bar{a}' \in M$ such that $\models \phi[\bar{a}', \bar{b}, \bar{c}]$.

Note that if $\bar{a}'$, $\bar{b}'$ have the same range as $\bar{a}$, $\bar{a}'$ respectively, then $\bar{a} \underset{M}{\downarrow} \bar{b}$ if and only if $\bar{a}' \underset{M}{\downarrow} \bar{b}'$. Thus we will also write for example $A \underset{M}{\downarrow} B$ to mean that $\bar{a} \underset{M}{\downarrow} \bar{b}$ for some (equivalently any) enumerations $\bar{a}$, $\bar{b}$ of A and B respectively.

Another way of saying the same thing: if $M \subseteq B$ and $p \in \mathbf{S}^{<\infty}(B)$, we say that p is *free over* M if it is finitely satisfiable over M: any formula $\phi(\bar{x}, \bar{b})$ in p is realized in M. Note that $\bar{a} \underset{M}{\downarrow} B$ if and only if $\mathbf{tp}(\bar{a}/MB)$ is free over M. We will freely go back and forth between these two point of views (types and independence notion $\downarrow$). Which one is easier to work with depends on the specific concepts we are studying.

Theorem B.4 (Properties of independence)**.** Assume that T does not have the order property.

(1) (Invariance) If $A \underset{M}{\downarrow} B$ and f is an automorphism of $\mathfrak{C}$, then $f[A] \underset{f[M]}{\downarrow} f[B]$.
(2) (Normality) If $A \underset{M}{\downarrow} B$, then $AM \underset{M}{\downarrow} BM$.
(3) (Left and right monotonicity) If $A \underset{M}{\downarrow} B$ and $A_0 \subseteq A$, $B_0 \subseteq B$, then $A_0 \underset{M}{\downarrow} B_0$.
(4) (Base monotonicity) If $A \underset{M}{\downarrow} B$ and $M \preceq N \subseteq B$, then $A \underset{N}{\downarrow} B$.
(5) (Finite character) $A \underset{M}{\downarrow} B$ if and only if $A_0 \underset{M}{\downarrow} B_0$ for all finite $A_0 \subseteq A$, $B_0 \subseteq B$.
(6) (Disjointness) If $A \underset{M}{\downarrow} B$, then $A \cap B \subseteq M$.
(7) (Symmetry) $A \underset{M}{\downarrow} B$ if and only if $B \underset{M}{\downarrow} A$.
(8) (Transitivity) If $M_0 \preceq M_1 \preceq M_2$, $A \underset{M_0}{\downarrow} M_1$, and $A \underset{M_1}{\downarrow} M_2$, then $A \underset{M_0}{\downarrow} M_2$.
(9) (Local character) For any A and N, there exists $M \preceq N$ of cardinality at most $|A| + |T|$ such that $A \underset{M}{\downarrow} N$.
(10) (Uniqueness) If $M \subseteq B$, $p, q \in \mathbf{S}^{<\infty}(B)$ are both free over M and $p \restriction M = q \restriction M$, then $p = q$.
(11) (Extension) If $p \in \mathbf{S}^{<\infty}(M)$ and $M \subseteq B$ is a set, there exists $q \in \mathbf{S}^{<\infty}(B)$ that extends p and is free over B.

Proof. Invariance, normality, the monotonicity properties, and finite character are immediate from the definition. To see disjointness, it is enough to see that if $a \underset{M}{\downarrow} a$ then $a \in M$. This follows from the definition applied with the formula $x = y$. Let us prove the other properties:

- Symmetry: Suppose not. Fix $\bar{a}$, $\bar{b}$, and M so that $\bar{a} \underset{M}{\downarrow} \bar{b}$ but $\bar{b} \not\underset{M}{\downarrow} \bar{a}$. Without loss of generality, $\bar{a}$ and $\bar{b}$ are finite and we can pick a formula $\phi(\bar{x}, \bar{y})$ witnessing $\bar{b} \not\underset{M}{\downarrow} \bar{a}$ that has all parameters from M already incorporated in $\bar{a}$: $\models \phi[\bar{a}, \bar{b}]$ but for all $\bar{b}' \in M$, $\models \neg\phi[\bar{a}, \bar{b}']$ (we have swapped the role of $\bar{x}$ and $\bar{y}$ for convenience in the proof that follows).

  We inductively build two sequences $\langle \bar{a}_i : i < \omega \rangle, \langle \bar{b}_i : i < \omega \rangle$ of tuples in M such that $\models \phi[\bar{a}_i, \bar{b}_j]$ if and only if $i \leq j$, and $\models \phi[\bar{a}_i, \bar{b}]$ for all $i < \omega$.

  This is enough: set $\psi(\bar{x}_1, \bar{y}_1, \bar{x}_2, \bar{y}_2)$ to be $\phi(\bar{x}_1, \bar{y}_2) \wedge \bar{x}_1 \bar{y}_1 \neq \bar{x}_2 \bar{y}_2$. Then ψ and the sequence $\langle \bar{a}_i \bar{b}_i : i < \omega \rangle$ witness the order property, contradiction.

  This is possible: Fix $j < \omega$ and assume we are given $\bar{a}_i$ and $\bar{b}_i$ for all $i < j$. By the induction hypothesis, we know that:

$$\models \phi(\bar{a}, \bar{b}) \wedge \bigwedge_{i<j} \phi(\bar{a}_i, \bar{b}) \wedge \bigwedge_{i<j} \neg\phi(\bar{a}, \bar{b}_i)$$

(the last part is from the hypothesis that $\models \neg\phi(\bar{a}, \bar{b}')$ for any $\bar{b}' \in M$). Because $\bar{a} \underset{M}{\downarrow} \bar{b}$, there exists $\bar{a}' \in M$ such that:

$$\models \phi(\bar{a}', \bar{b}) \wedge \bigwedge_{i<j} \phi(\bar{a}_i, \bar{b}) \wedge \bigwedge_{i<j} \neg\phi(\bar{a}', \bar{b}_i)$$

Since $M \preceq \mathfrak{C}$, there exists $\bar{b}' \in M$ such that:

$$\models \phi(\bar{a}', \bar{b}') \wedge \bigwedge_{i<j} \phi(\bar{a}_i, \bar{b}') \wedge \bigwedge_{i<j} \neg\phi(\bar{a}', \bar{b}_i)$$

Set $\bar{a}_j := \bar{a}'$, $\bar{b}_j := \bar{b}'$.

- Transitivity: Using the definition of independence, it is easy to check the "left" version of transitivity: if $M_0 \preceq M_1 \preceq M_2$, $M_2 \underset{M_1}{\downarrow} A$, and $M_1 \underset{M_0}{\downarrow} A$, then $M_2 \underset{M_0}{\downarrow} A$. The "right" version of transitivity then follows from symmetry.
- Local character: This is a downward Löwenheim-Skolem closure argument, that we could do explicitly. Instead, fix A and N, and pick a pair of models $M \preceq M'$ such that $A \subseteq M'$, $|UM'| \leq |A| + |T|$, and $(M, M') \preceq (N, \mathfrak{C})$ (in the vocabulary with an additional predicate for M). From the definition of independence, it follows that $N \underset{M}{\downarrow} M'$, so by monotonicity $N \underset{M}{\downarrow} A$, hence by symmetry $A \underset{M}{\downarrow} N$.
- Uniqueness: This is similar to symmetry: let $\bar{b}$ be an enumeration of $B - M$. Suppose $p = \mathbf{tp}(\bar{a}/M\bar{b})$, $q = \mathbf{tp}(\bar{a}'/M\bar{b})$ both are free over M (so $\bar{a} \underset{M}{\downarrow} \bar{b}$, $\bar{a}' \underset{M}{\downarrow} \bar{b}$), and $p \restriction M = q \restriction M$. We have to see that $p = q$. Without loss of generality, $\bar{b}$, $\bar{a}$, and $\bar{a}'$ are finite. Suppose $p \neq q$, and let $\phi(\bar{x}, \bar{y})$ be such that $\models \phi[\bar{a}, \bar{b}] \wedge \neg\phi[\bar{a}', \bar{b}]$.

  Define sequences $\langle \bar{a}_i : i < \omega \rangle$, $\langle \bar{b}_i : i < \omega \rangle$ in M such that for all $i, j < \omega$:
  (1) $\models \phi[\bar{a}_i, \bar{b}]$.
  (2) $\models \phi[\bar{a}_i, \bar{b}_j]$ if and only if $i \leq j$.
  (3) $\models \neg\phi[\bar{a}, \bar{b}_i]$.

  This is enough: Then $\psi(\bar{x}_1, \bar{y}_1, \bar{x}_2, \bar{y}_2) := \phi(\bar{x}_1, \bar{y}_2) \wedge \bar{x}_1\bar{y}_1 \neq \bar{x}_2\bar{y}_2$ together with the sequence $\langle \bar{a}_i\bar{b}_i : i < \omega \rangle$ witness the order property.

  This is possible: Suppose that $\bar{a}_i, \bar{b}_i$ have been defined for all $i < j$. By the induction hypothesis, we have:

$$\models \bigwedge_{i<j} \phi[\bar{a}_i, \bar{b}] \wedge \bigwedge_{i<j} \neg\phi[\bar{a}, \bar{b}_i] \wedge \phi[\bar{a}, \bar{b}]$$

  Since $\bar{a} \underset{M}{\downarrow} \bar{b}$, there is $\bar{a}'' \in M$ such that:

$$\models \bigwedge_{i<j} \phi[\bar{a}_i, \bar{b}] \wedge \bigwedge_{i<j} \neg\phi[\bar{a}'', \bar{b}_i] \wedge \phi[\bar{a}'', \bar{b}]$$

We also know that $\models \neg\phi[\bar{a}', \bar{b}]$. Combining this with the above and the fact that $\bar{a}' \underset{M}{\downarrow} \bar{b}$, hence by symmetry $\bar{b} \underset{M}{\downarrow} \bar{a}'$, we obtain a $\bar{b}'' \in M$ such that:

$$\models \bigwedge_{i<j} \phi[\bar{a}_i, \bar{b}''] \wedge \bigwedge_{i<j} \neg\phi[\bar{a}'', \bar{b}_i] \wedge \phi[\bar{a}'', \bar{b}''] \wedge \neg\phi[\bar{a}', \bar{b}'']$$

Let $\bar{a}_j := \bar{a}''$, $\bar{b}_j := \bar{b}''$. It is easy to check that this works (for condition (3), we use that $p \restriction M = q \restriction M$ so $\models \neg\phi[\bar{a}', \bar{b}'']$ implies $\models \neg\phi[\bar{a}, \bar{b}'']$).

- Extension: This is the only place where we use the compactness theorem. Consider the set q of formulas $\phi(\bar{x}, \bar{b})$ where $\bar{b} \in B$, $p \cup \{\phi(\bar{x}, \bar{b})\}$ is consistent, and there exists $\bar{a}' \in M$ such that $\models \phi[\bar{a}', \bar{b}]$. Note that q is closed under conjunctions, hence by construction and compactness is consistent. It remains to check that q is complete. Indeed, assume that $\neg\phi(\bar{x}, \bar{b}) \notin q$. There are two cases. If $p \cup \{\neg\phi(\bar{x}, \bar{b})\}$ is inconsistent, then $p \models \phi(\bar{x}, \bar{b})$, so there is $\psi(\bar{x}, \bar{c}) \in p$ such that $\psi(\bar{x}, \bar{c}) \models \phi(\bar{x}, \bar{b})$, with $\bar{c} \in M$. We know that $\mathfrak{C} \models \exists \bar{x} \psi(\bar{x}, \bar{c})$, so $M \models \exists \bar{x} \psi(\bar{x}, \bar{c})$, hence $\phi(\bar{x}, \bar{b})$ is also realized in M. In the second case, $p \cup \{\neg\phi(\bar{x}, \bar{b})\}$ is consistent but $\neg\phi(\bar{x}, \bar{b})$ is not realized in M. In particular, $\models \phi[\bar{a}', \bar{b}]$ for all $\bar{a}' \in M$. As before, for any $\psi(\bar{x}, \bar{c}) \in p$, there exists $\bar{a}' \in M$ so that $M \models \psi[\bar{a}', \bar{c}]$, so $M \models \psi[\bar{a}', \bar{c}] \wedge \phi[\bar{a}', \bar{b}]$, hence $p \cup \{\phi(\bar{x}, \bar{b})\}$ is finitely consistent, hence consistent. This shows that $\phi(\bar{x}, \bar{b}) \in q$, as desired.

□

It is straightforward to check that $\underset{M}{\downarrow}$ as defined here yields a stable independence notion (in the sense of Definition 5.5). The existence of an independence notion as in the theorem implies stability, thus we get:

Theorem B.5. If T does not have the order property (or just the conclusion of Theorem B.4 holds), then T is stable in every infinite cardinal λ with $\lambda = \lambda^{|T|}$. In particular, T is stable if and only if T does not have the order property.

Proof. The "in particular" part will follow from Theorem B.2. Now assume that T does not have the order property, or just that the conclusion of Theorem B.4 holds. Fix an infinite cardinal λ such that $\lambda = \lambda^{|T|}$. Since any set of cardinality λ is contained in a model of cardinality λ, it is enough to count types over models. Fix M of cardinality λ and a sequence $\langle p_i : i < \lambda^+ \rangle$ of types over M. We will show that there exists $i < j$ so that $p_i = p_j$. First, for each $i < \lambda^+$, local character tells us there exists $M_i \preceq M$ of cardinality at most $|T|$ such that p is free over M_i. Note that $|\{M_i : i < \lambda^+\}| \leq \lambda^{|T|} = \lambda$, so by the pigeonhole principle there exists $S \subseteq \lambda^+$ of cardinality λ^+ and $i_0 < \lambda^+$ such that for any $i \in S$, p_i is free over M_{i_0}. Now $|\mathbf{S}(M_{i_0})| \leq 2^{|T|} \leq \lambda^{|T|} = \lambda$, so by the pigeonhole principle again, there exists $i < j$ in S such that $p_i \restriction M_{i_0} = p_j \restriction M_{i_0}$. Since both p_i and p_j are free over M_{i_0}, uniqueness implies that $p_i = p_j$. □

From now on, we will freely use the equivalence between stability and no order property, as well as the conclusion of Theorem B.4 (but not the exact definition of independence).

Lemma B.6. Assume that T is stable. If $\bar{a} \underset{M}{\downarrow} B$ and $\bar{b}_1, \bar{b}_2 \in B$ are such that $\bar{b}_1 \equiv_M \bar{b}_2$, then $\bar{b}_1 \bar{a} \equiv_M \bar{b}_2 \bar{a}$.

Proof. Fix an automorphism f of $\mathfrak{C}$ fixing M and sending $\bar{b}_1$ to $\bar{b}_2$. Let $\bar{a}' := f(\bar{a})$. Then $\bar{b}_1\bar{a} \equiv_M \bar{b}_2\bar{a}'$. On the other hand, $\bar{a} \underset{M}{\downarrow} \bar{b}_1$, so $\bar{a}' \underset{M}{\downarrow} \bar{b}_2$. By monotonicity, also $\bar{a} \underset{M}{\downarrow} \bar{b}_2$, so by uniqueness, $\bar{b}_2\bar{a}' \equiv_M \bar{b}_2\bar{a}$. Combining the two equality of types shows that $\bar{b}_1\bar{a} \equiv_M \bar{b}_2\bar{a}$. □

As a final application, we will show how to extract indiscernibles in stable theories.

Theorem B.7. Assume that $\lambda > |T|$ and T is stable in λ. If $\langle a_i : i < \lambda^+ \rangle$ is a sequence, there exists $I \subseteq \lambda^+$ of cardinality λ^+ such that $\langle a_i : i \in I \rangle$ is indiscernible.

Proof. First, build $\langle M_i : i \leq \lambda^+ \rangle$ an increasing continuous sequence of models of size λ such that $a_i \in M_{i+1}$ for all $i < \lambda^+$. Let $S := \{i < \lambda^+ \mid \text{cf}(i) \geq |T|^+\}$. This is a stationary set. By local character, for each $i \in S$ there exists $j_i < i$ such that $a_i \underset{M_{j_i}}{\downarrow} M_i$. By Fodor's lemma, there exists $S_0 \subseteq S$ stationary and $j^* < \lambda^+$ such that $j_i = j^*$ for all $i \in S_0$. Let's prune a bit more: by stability $|\mathbf{S}(M_{j^*})| = \lambda$. Thus by the pigeonhole principle we can find an unbounded $I \subseteq S$ such that for $i < i'$ in I, $a_i \equiv_{M_{j^*}} a_{i'}$. We prove that $\langle a_i : i \in I \rangle$ is indiscernible over $M := M_{j^*}$.

For this, we show by induction on $n < \omega$ that for any $i_0 < \ldots < i_n$, $i'_0 < \ldots < i'_n$ in I, $a_{i_0} \ldots a_{i_n} \equiv_M a_{i'_0} \ldots a_{i'_n}$. The base case has just been observed. Assume now that $n = m + 1$, set $\bar{b}_1 := a_{i_0} \ldots a_{i_m}$, $\bar{b}_2 := a_{i'_0} \ldots a_{i'_m}$, and assume we know $\bar{b}_1 \equiv_M \bar{b}_2$. Without loss of generality, $i'_n \leq i_n$. By monotonicity, $a_{i_n} \underset{M}{\downarrow} M_{i'_n}$, so by uniqueness $a_{i_n} \equiv_{M_{i'_n}} a_{i'_n}$. In particular, $\bar{b}_2 a_{i_n} \equiv_M \bar{b}_2 a_{i'_n}$. By Lemma B.6, we also have that $\bar{b}_1 a_{i_n} \equiv_M \bar{b}_2 a_{i_n}$. Combining these two type equalities gives that $\bar{b}_1 a_{i_n} \equiv_M \bar{b}_2 a_{i'_n}$, as desired. □

References

[ABV19] Nathanael Ackerman, Will Boney, and Sebastien Vasey, *Categoricity in multiuniversal classes*, Annals of Pure and Applied Logic **170** (2019), no. 11, 102712.

[AHS04] Jiří Adámek, Horst Herrlich, and George E. Strecker, *Abstract and concrete categories*, online edition ed., 2004, Available from `http://katmat.math.uni-bremen.de/acc/`.

[AM10] Uri Abraham and Menachem Magidor, *Cardinal arithmetic*, Handbook of set theory (Matthew Foreman and Akihiro Kanamori, eds.), Springer, 2010, pp. 1149–1227.

[AR94] Jiří Adámek and Jiří Rosický, *Locally presentable and accessible categories*, London Math. Society Lecture Notes, Cambridge University Press, 1994.

[Bal09] John T. Baldwin, *Categoricity*, University Lecture Series, vol. 50, American Mathematical Society, 2009.

[BB17] John T. Baldwin and Will Boney, *Hanf numbers and presentation theorems in AECs*, Beyond first order model theory (José Iovino, ed.), CRC Press, 2017, pp. 327–352.

[BBE01] Ladislav Bican, Robert El Bashir, and Edgar E. Enochs, *All modules have flat covers*, Bulletin of the London Mathematical Society **33** (2001), 385–390.

[Bek00] Tibor Beke, *Sheafifiable homotopy model category*, Mathematical Proceedings of the Cambridge Philosophical society **129** (2000), 447–475.

[BET07] John T. Baldwin, Paul C. Eklof, and Jan Trlifaj, ${}^{\perp}N$ *as an abstract elementary class*, Annals of Pure and Applied Logic **149** (2007), 25–39.

[BFB85] Jon Barwise, Solomon Feferman, and Andreas Baudisch (eds.), *Model-theoretic logics*, Perspectives in mathematical logic, Springer-Verlag, 1985.

[BGKV16] Will Boney, Rami Grossberg, Alexei Kolesnikov, and Sebastien Vasey, *Canonical forking in AECs*, Annals of Pure and Applied Logic **167** (2016), no. 7, 590–613.

[BGL+16] Will Boney, Rami Grossberg, Michael J. Lieberman, Jiří Rosický, and Sebastien Vasey, *μ-Abstract elementary classes and other generalizations*, The Journal of Pure and Applied Algebra **220** (2016), no. 9, 3048–3066.

[Bir35] Garrett Birkhoff, *On the structure of abstract algebras*, Proceedings of the Cambridge Philosophical Society **31** (1935), no. 4, 433–454.

[BL] Will Boney and Michael Lieberman, *Tameness, powerful images, and large cardinals*, Preprint. URL: `https://arxiv.org/abs/1902.10212v4`.

[BLS15] John T. Baldwin, Paul B. Larson, and Saharon Shelah, *Almost Galois ω-stable classes*, The Journal of Symbolic Logic **80** (2015), no. 3, 763–784.

[Bona] Will Boney, *Classification theory for tame abstract elementary classes*, Lecture notes. URL: `http://math.harvard.edu/~wboney/fall17/255-notes.pdf`.

[Bonb] ———, *Erdős-Rado classes*, Preprint. URL: `https://arxiv.org/abs/1810.01513v2`.

[BR12] Tibor Beke and Jiří Rosický, *Abstract elementary classes and accessible categories*, Annals of Pure and Applied Logic **163** (2012), 2008–2017.

[BR16] ———, *Cellular objects and Shelah's singular compactness theorem*, Journal of Pure and Applied Algebra **220** (2016), no. 5, 1813–1836.

[BTR17] Andrew Brooke-Taylor and Jiří Rosický, *Accessible images revisited*, Proceedings of the American Mathematical Society **145** (2017), 1317–1327.

[BU17] Will Boney and Spencer Unger, *Large cardinal axioms from tameness in AECs*, Proceedings of the American Mathematical Society **145** (2017), 4517–4532.

[BV] Will Boney and Monica VanDieren, *Limit models in strictly stable abstract elementary classes*, Preprint. URL: `http://arxiv.org/abs/1508.04717v4`.

[BV17] Will Boney and Sebastien Vasey, *A survey on tame abstract elementary classes*, Beyond first order model theory (José Iovino, ed.), CRC Press, 2017, pp. 353–427.

[BV19] ———, *Structural logic and abstract elementary classes with intersections*, Bulletin of the Polish academy of science (mathematics) **67** (2019), 1–17.

[BYU10] Itay Ben-Yaacov and Alexander Usvyatsov, *Continuous first order logic and local stability*, Transactions of the American Mathematical Society **362** (2010), no. 10, 5213–5259.

[CK90] C.C. Chang and H. Jerome Keisler, *Model theory*, third ed., Studies in logic and the foundations of mathematics, vol. 73, North-Holland, 1990.

[CSS05] Carles Casacuberta, Dirk Scevenels, and Jeffry H. Smith, *Implications of large-cardinal principles in homotopical localization*, Advances in Mathematics **197** (2005), 120–139.

[Dic75] M. A. Dickmann, *Large infinitary languages*, Studies in logic and the foundations of mathematics, vol. 83, North-Holland, 1975.

[DS78] Keith J. Devlin and Saharon Shelah, *A weak version of $\diamondsuit$ which follows from $2^{\aleph_0} < 2^{\aleph_1}$*, Israel Journal of Mathematics **29** (1978), no. 2, 239–247.

[Eno81] Edgar E. Enochs, *Injective and flat covers, envelopes and resolvents*, Israel Journal of Mathematics **39** (1981), no. 3, 189–209.

[ET01] Paul C. Eklof and Jan Trlifaj, *How to make Ext vanish*, Bulletin of the London Mathematical Society **33** (2001), 41–51.

[Gro02] Rami Grossberg, *Classification theory for abstract elementary classes*, Contemporary Mathematics **302** (2002), 165–204.

[GV06a] Rami Grossberg and Monica VanDieren, *Categoricity from one successor cardinal in tame abstract elementary classes*, Journal of Mathematical Logic **6** (2006), no. 2, 181–201.

[GV06b] ———, *Shelah's categoricity conjecture from a successor for tame abstract elementary classes*, The Journal of Symbolic Logic **71** (2006), no. 2, 553–568.

[GVV16] Rami Grossberg, Monica VanDieren, and Andrés Villaveces, *Uniqueness of limit models in classes with amalgamation*, Mathematical Logic Quarterly **62** (2016), 367–382.

[Hen19] Simon Henry, *An abstract elementary class non-axiomatizable in $L_{(\infty,\kappa)}$*, The Journal of Symbolic Logic **84** (2019), no. 3, 1240–1251.

[Hes02] Kathryn Hess, *Model categories in algebraic topology*, Applied Categorical Structures **10** (2002), 195–220.

[HH84] Victor Harnik and Leo Harrington, *Fundamentals of forking*, Annals of Pure and Applied Logic **26** (1984), 245–286.

[HJ99] Karel Hrbacek and Thomas Jech, *Introduction to set theory*, third ed., Marcel Dekker, 1999.

[HK06] Tapani Hyttinen and Meeri Kesälä, *Independence in finitary abstract elementary classes*, Annals of Pure and Applied Logic **143** (2006), 103–138.

[Hov99] Mark Hovey, *Model categories*, Mathematical surveys and monographs, vol. 63, American Mathematical Society, 1999.

[Jec03] Thomas Jech, *Set theory*, 3rd ed., Springer-Verlag, 2003.

[JS13] Adi Jarden and Saharon Shelah, *Non-forking frames in abstract elementary classes*, Annals of Pure and Applied Logic **164** (2013), 135–191.

[Kei73] H. Jerome Keisler, *Forcing and the omitting types theorem*, Studies in model theory (Michael D. Morley and Jon Barwise, eds.), Studies in mathematics, vol. 8, Mathematical Association of America, 1973, pp. 96–133.

[KMA] Thomas G. Kucera and Marcos Mazari-Armida, *On universal modules with pure embeddings*, Preprint. URL: `https://arxiv.org/abs/1903.00414v4`.

[KP97] Byunghan Kim and Anand Pillay, *Simple theories*, Annals of Pure and Applied Logic **88** (1997), 149–164.

[Kub14] Wiesław Kubiś, *Fraïssé sequences: category-theoretic approach to universal homogeneous structures*, Annals of Pure and Applied Logic **165** (2014), 1755–1811.

[Lai81] Christian Lair, *Catégories modelables et catégories esquissables*, Diagrammes **6** (1981), L1–L20.

[Lan98] Saunders Mac Lane, *Categories for the working mathematician*, second ed., Graduate Texts in Mathematics, vol. 5, Springer, 1998.

[Lie] Michael Lieberman, *A category-theoretic characterization of almost measurable cardinals*, Preprint. URL: `https://arxiv.org/abs/1809.05953v4`.

[Lie11] Michael J. Lieberman, *Category-theoretic aspects of abstract elementary classes*, Annals of Pure and Applied Logic **162** (2011), no. 11, 903–915.

[LR16] Michael J. Lieberman and Jiří Rosický, *Classification theory for accessible categories*, The Journal of Symbolic Logic **81** (2016), no. 1, 151–165.

[LR17a] Michael Lieberman and Jiří Rosický, *Hanf numbers via accessible images*, Logical Methods in Computer Science **13** (2017), 1–15.

[LR17b] Michael J. Lieberman and Jiří Rosický, *Metric abstract elementary classes as accessible categories*, The Journal of Symbolic Logic **82** (2017), no. 3, 1022–1040.

[LRVa] Michael Lieberman, Jiří Rosický, and Sebastien Vasey, *Hilbert spaces and C^*-algebras are not finitely concrete*, Preprint. URL: `https://arxiv.org/abs/1908.10200v3`.

[LRVb] Michael J. Lieberman, Jiří Rosický, and Sebastien Vasey, *Sizes and filtrations in accessible categories*, Israel Journal of Mathematics, To appear. URL: `https://arxiv.org/abs/1902.06777v3`.

[LRVc] ———, *Weak factorization systems and stable independence*, Preprint. URL: `https://arxiv.org/abs/1904.05691v2`.

[LRV19a] ———, *Forking independence from the categorical point of view*, Advances in Mathematics **346** (2019), 719–772.

[LRV19b] ———, *Internal sizes in μ-abstract elementary classes*, Journal of Pure and Applied Algebra **223** (2019), no. 10, 4560–4582.

[LRV19c] ———, *Universal abstract elementary classes and locally multipresentable categories*, Proceedings of the American Mathematical Society **147** (2019), no. 3, 1283–1298.

[Lur09] Jacob Lurie, *Higher topos theory*, Annals of Mathematics Studies, no. 170, Princeton University Press, 2009.

[MA20] Marcos Mazari-Armida, *Algebraic description of limit models in classes of abelian groups*, Annals of Pure and Applied Logic **171** (2020), no. 1, 102723.

[Mac03] Angus Macintyre, *Model theory: geometrical and set-theoretic aspects and prospects*, The Bulletin of Symbolic Logic **9** (2003), no. 2, 197–212.

[Mor65] Michael Morley, *Categoricity in power*, Transactions of the American Mathematical Society **114** (1965), 514–538.

[MP89] Michael Makkai and Robert Paré, *Accessible categories: The foundations of categorical model theory*, Contemporary Mathematics, vol. 104, American Mathematical Society, 1989.

[MR77] Michael Makkai and Gonzalo E. Reyes, *First order categorical logic*, Lecture Notes in Mathematics, vol. 611, Springer-Verlag, 1977.

[MRV14] Michael Makkai, Jiří Rosický, and Lukáš Vokřínek, *On a fat small object argument*, Advances in Mathematics **254** (2014), 49–68.

[MV11] Menachem Magidor and Jouko Väänänen, *On löwenheim-skolem-tarski numbers for extensions of first order logic*, Journal of Mathematical Logic **11** (2011), no. 1, 87–113.

[PW06] Jonathan Pila and Alex James Wilkie, *The rational points of a definable set*, Duke Mathematical Journal **133** (2006), no. 3, 591–616.

[Rie16] Emily Riehl, *Category theory in context*, Aurora: Modern Math Originals, Dover, 2016.

[Ros02] Jiří Rosický, *Flat covers and factorizations*, Journal of Algebra **253** (2002), 1–13.

[RT03] Jiří Rosický and Walter Tholen, *Left-determined model categories and universal homotopy theories*, Transactions of the American Mathematical Society **355** (2003), no. 9, 3611–3623.

[Sac72] Gerald E. Sacks, *The differential closure of a differential field*, Bulletin of the American Mathematical Society **78** (1972), no. 5, 629–634.

[She75] Saharon Shelah, *A compactness theorem for singular cardinals, free algebras, Whitehead problem and transversals*, Israel Journal of Mathematics **21** (1975), no. 4, 319–349.

[She83a] ———, *Classification theory for non-elementary classes I: The number of uncountable models of $\psi \in L_{\omega_1,\omega}$. Part A*, Israel Journal of Mathematics **46** (1983), no. 3, 214–240.

[She83b] ———, *Classification theory for non-elementary classes I: The number of uncountable models of $\psi \in L_{\omega_1,\omega}$. Part B*, Israel Journal of Mathematics **46** (1983), no. 4, 241–273.

[She87] ———, *Classification of non elementary classes II. Abstract elementary classes*, Classification Theory (Chicago, IL, 1985) (John T. Baldwin, ed.), Lecture Notes in Mathematics, vol. 1292, Springer-Verlag, 1987, pp. 419–497.

[She90] ———, *Classification theory and the number of non-isomorphic models*, 2nd ed., Studies in logic and the foundations of mathematics, vol. 92, North-Holland, 1990.

[She94] ———, *Cardinal arithmetic*, Oxford Logic Guides, no. 29, Clarendon Press, 1994.

[She98] ———, *Proper and improper forcing*, second edition ed., Perspectives in mathematical logic, Springer-Verlag, 1998.

[She99] ———, *Categoricity for abstract classes with amalgamation*, Annals of Pure and Applied Logic **98** (1999), no. 1, 261–294.

[She09a] ———, *Classification theory for abstract elementary classes*, Studies in Logic: Mathematical logic and foundations, vol. 18, College Publications, 2009.

[She09b] ———, *Classification theory for abstract elementary classes 2*, Studies in Logic: Mathematical logic and foundations, vol. 20, College Publications, 2009.

[ŠT12] Jan Šaroch and Jan Trlifaj, *Kaplansky classes, finite character and $\aleph_1$-projectivity*, Forum Mathematicum **24** (2012), 1091–1109.

[SU11] Saharon Shelah and Alexander Usvyatsov, *Model theoretic stability and categoricity for complete metric spaces*, Israel Journal of Mathematics **182** (2011), 157–198.

[SV] Saharon Shelah and Sebastien Vasey, *Categoricity and multidimensional diagrams*, Preprint. URL: `https://arxiv.org/abs/1805.06291v2`.

[Tar54] Alfred Tarski, *Contributions to the theory of models I*, Indagationes Mathematicae **16** (1954), 572–581.

[Van16] Monica VanDieren, *Superstability and symmetry*, Annals of Pure and Applied Logic **167** (2016), no. 12, 1171–1183.

[Vas] Sebastien Vasey, *Model theory for abstract elementary classes*, Lecture notes. URL: `http://math.harvard.edu/~sebv/aec-spring-2018/aec-lecture-notes_04_26_2018.pdf`.

[Vas17a] ———, *Downward categoricity from a successor inside a good frame*, Annals of Pure and Applied Logic **168** (2017), no. 3, 651–692.

[Vas17b] ———, *Saturation and solvability in abstract elementary classes with amalgamation*, Archive for Mathematical Logic **56** (2017), 671–690.

[Vas17c] ———, *Shelah's eventual categoricity conjecture in universal classes: part I*, Annals of Pure and Applied Logic **168** (2017), no. 9, 1609–1642.

[Vas17d] ———, *Shelah's eventual categoricity conjecture in universal classes: part II*, Selecta Mathematica **23** (2017), no. 2, 1469–1506.

[Vas18a] ———, *Shelah's eventual categoricity conjecture in tame AECs with primes*, Mathematical Logic Quarterly **64** (2018), 25–36.

[Vas18b] ———, *Toward a stability theory of tame abstract elementary classes*, Journal of Mathematical Logic **18** (2018), no. 2, 1850009.

[Vas19] ———, *The categoricity spectrum of large abstract elementary classes with amalgamation*, Selecta Mathematica **25** (2019), no. 5, 65.

[vdD98] Lou van den Dries, *Tame topology and o-minimal structures*, Lecture Note Series, no. 248, Cambridge University Press, 1998.

[Veb04] Oswald Veblen, *A system of axioms for geometry*, Transactions of the American Mathematical Society **5** (1904), no. 3, 343–384.

[Zil05] Boris Zilber, *Pseudo-exponentiation on algebraically closed fields of characteristic zero*, Annals of Pure and Applied Logic **132** (2005), 67–95.

[Łoś54] Jerzy Łoś, *On the categoricity in power of elementary deductive systems and some related problems*, Colloquium Mathematicae **3** (1954), no. 1, 58–62.

E-mail address: sebv@math.harvard.edu

URL: http://math.harvard.edu/~sebv/

Department of Mathematics, Harvard University, Cambridge, Massachusetts, USA