HANF NUMBERS VIA ACCESSIBLE IMAGES

M. LIEBERMAN AND J. ROSICKÝ

Department of Mathematics and Statistics, Masaryk University, Czech Republic
e-mail address: {lieberman,rosicky}@math.muni.cz

Dedicated to Jiří Adámek on the occasion of his 70th birthday.

Abstract. We present several new model-theoretic applications of the fact that, under the assumption that there exists a proper class of almost strongly compact cardinals, the powerful image of any accessible functor is accessible. In particular, we generalize to the context of accessible categories the recent Hanf number computations of Baldwin and Boney, namely that in an abstract elementary class (AEC) if the joint embedding and amalgamation properties hold for models of size up to a sufficiently large cardinal, then they hold for models of arbitrary size. Moreover, we prove that, under the above-mentioned large cardinal assumption, every metric AEC is strongly d-tame, strengthening a result of Boney and Zambrano and pointing the way to further generalizations.

1. Introduction

The connections of accessible categories to problems in abstract model theory, and particularly to those involving abstract elementary classes (AECs) and metric abstract elementary classes (mAECs), have by now been the subject of considerable investigation, e.g. [14], [4], [15], [6], and [16]. Accessible categories serve as a common generalization, certainly—AECs, mAECs, and further generalizations including μ-CAECs and μ-AECs, all amount to accessible categories with additional structure—and offer a perspective in which certain powerful tools for analysing the structure of such classes become evident. As in [15], all of the machinery at work in this paper can be traced back, ultimately, to Theorem 5.5.1 in [17]: assuming arbitrarily large strongly compact cardinals, given any accessible functor $F : \mathcal{K} \to \mathcal{L}$, the powerful image of F (that is, the closure of its image under subobjects in $\mathcal{L}$) is an accessible category.

The dependence on set theory, while perhaps surprising, is easily motivated. While the category of abelian groups, **Ab**, is finitely accessible, the full subcategory of free abelian groups, **FrAb**, which is the powerful image of the accessible free abelian group functor $F : \textbf{Set} \to \textbf{Ab}$, need not be. In particular, it is not accessible assuming $V = L$; if there are arbitrarily large strongly compact cardinals, on the other hand, it is accessible ([11]). The aforementioned theorem of [17], which generalizes this fact, admits a further generalization in [9]: if $F : \mathcal{K} \to \mathcal{L}$ is a λ-accessible functor that preserves $\mu_{\mathcal{L}}$-presentable objects—we

recall the definition of $\mu_{\mathcal{L}}$ in Section 3 below—and κ is a $\mu_{\mathcal{L}}$-strongly compact cardinal, the powerful image of F is $\mu_{\mathcal{L}}$-accessible.

The applications of the accessibility of powerful images in abstract model theory, both past (the proof in [15], via [17], of Boney's Theorem on the tameness of AECs under the assumption of arbitrarily large strongly compact cardinals in [5]—subsequently completed in [7] 4.14 to a grand equivalence of tameness, accessibility of powerful images, and arbitrarily large almost strongly compact cardinals) and present (the arguments to follow, which lean on the theorem of [9] above) are united by a single, very simple idea. In particular, many of the essential properties of abstract classes of structures $\mathcal{K}$—including the amalgamation and joint embedding properties (abbreviated as AP and JEP, respectively), and the equality of, or distance between, Galois types—reduce to questions about whether diagrams of a particular shape $\mathcal{A}$ can be extended to diagrams of shape $\mathcal{A}'$. In the case of the joint embedding and amalgamation properties, and to an extent their disjoint versions (DJEP, DAP), $\mathcal{A}$ and $\mathcal{A}'$ are simple finite categories and the diagrams live entirely in $\mathcal{K}$—which we now take to be a general accessible category, meaning that we can identify the respective categories of diagrams with the functor categories $\mathcal{K}^{\mathcal{A}}$ and $\mathcal{K}^{\mathcal{A}'}$. This sets up the first, and perhaps clearest, application of the accessibility of powerful images.

The functor whose image we will consider is simply the forgetful functor

$$F : \mathcal{K}^{\mathcal{A}'} \to \mathcal{K}^{\mathcal{A}}$$

which forgets the details of the extension. Clearly, the category of extendable diagrams of shape $\mathcal{A}$ corresponds to its (powerful) image. When this functor is accessible, as it will be in the cases considered below, we know that the image is κ-accessible for any sufficiently compact cardinal κ. If we know that any $\mathcal{A}$-diagram of κ-presentable objects is extendible (corresponding to, e.g. the $< \kappa$-JEP), the latter fact—in particular, the closure of the image under κ-directed colimits—is precisely what we need to infer that every $\mathcal{A}$-diagram is extendible (corresponding to, e.g. the full JEP). In Sections 3 and 4, we will use this template to generalize Theorem 1.1 of [3], proving that if one of the JEP, AP, DJEP, or DAP holds in an accessible category $\mathcal{K}$ for objects of presentability rank up to a sufficiently large cardinal κ—a Hanf number, of sorts—it must hold for objects of arbitrary presentability rank.

In Section 5, we turn to the question of d-tameness of Galois types in mAECs. Here again, the problem reduces to the accessibility of a particular category of completable diagrams, and is a straightforward (albeit decidedly more technical) generalization of the corresponding category-theoretic proof for the discrete case, namely conventional AECs ([15] 5.2). The discrete case is instructive: in an AEC $\mathcal{K}$, a Galois type over $M \in \mathcal{K}$ is an equivalence class of pairs $(f : M \to N, a \in UN)$, where UN denotes the underlying set of N, and where pairs $(f_i : M \to N_i, a_i \in UN_i)$ for $i = 1, 2$ are equivalent if there is an amalgam in which a_0 and a_1 are identified. This is, essentially, a pointed version of the picture for the AP, with the category of equivalent pairs as the image of the obvious forgetful functor. This (powerful) image will be κ-accessible for sufficiently compact κ, meaning that the category of equivalent pairs is determined by the subcategory of equivalent pairs over models of size no greater than κ. In the terminology of AECs, this means that $\mathcal{K}$ is κ-tame—see [15] for details.

In an mAEC $\mathcal{K}$, the objects have not just underlying sets, but underlying complete metric spaces. Provided $\mathcal{K}$ is reasonably well-behaved, we can speak meaningfully of the distance between Galois types $(f_i : M \to N_i, a_i \in UN_i)$ for $i = 1, 2$, here defined not in

an ambient monster model, as is customary, but as the infimum of the distances between the images of a_0 and a_1 in all possible amalgams of $N_0 \leftarrow M \rightarrow N_1$ in $\mathcal{K}$. For each $\epsilon > 0$, we then have a forgetful functor G_ϵ whose (powerful) image is the category of pairs within distance ϵ of one another. By choosing sufficiently compact κ, we may assume that this image is κ-accessible, and can thereby force a very strong version of metric tameness (i.e. d-tameness). Namely, we show that for any $\epsilon > 0$, a pair of types over an arbitrary model $M \in \mathcal{K}$ are at distance at least ϵ from one another only if there is a submodel N of M of density character at most κ such that their restrictions to N are already at distance at least ϵ. This strengthens, at least superficially, the tameness portion of Theorem 3.11 in [8], which is concerned with the standard $\epsilon - \delta$ definition of d-tameness; we prove that one may take $\delta = \epsilon$. More significantly, this approach points a way to further generalizations of tameness phenomena.

We conclude with a discussion of possible extensions of these results. In ongoing work with Zambrano, the authors are investigating natural generalizations from metric structures to quantale-valued structures, with the enticing prospect of meaningful applications both in fuzzy, probabilistic contexts, and in the more rarefied realm of model theory over sheaves.

We thank the anonymous referees for their insightful comments on this paper, which have led to noteworthy improvements in the presentation.

2. Preliminaries

We here restate a few basic definitions and results concerning accessible categories and their connections with AECs and mAECs; exhaustive treatments of the former topic can be found in [17] and [1], while the current state of knowledge on the latter is summarized in [16]. Additional background on AECs, and the broader context of abstract model theory, can be found in, e.g. [2].

For λ a regular cardinal, a λ-*accessible category* is a category $\mathcal{K}$ with λ-directed colimits in which every object can be expressed as a λ-directed colimit of λ-presentable objects, of which $\mathcal{K}$ contains only set many, up to isomorphism. Here an object N in $\mathcal{K}$ is λ-*presentable* if the hom-functor $\mathcal{K}(N, -) : \mathcal{K} \rightarrow \mathbf{Set}$ preserves λ-directed colimits; the *presentability rank* of an object N in $\mathcal{K}$, denoted $\pi(N)$, is the least λ such that N is λ-presentable. Presentability rank serves as a notion of size in an abstract category $\mathcal{K}$, and connects nicely to the natural notions of size in familiar model-theoretic contexts.

Examples 2.1. (1) Let $\mathcal{K}$ be an AEC with Löwenheim-Skolem number $LS(\mathcal{K})$, and let $U : \mathcal{K} \rightarrow \mathbf{Set}$ be the underlying set functor. Then $\mathcal{K}$ is $LS(\mathcal{K})^+$-accessible with directed colimits. Moreover, for any regular $\kappa > LS(\mathcal{K})$, a model $N \in \mathcal{K}$ is κ-presentable if and only if $|UN| < \kappa$ ([14] 4.1, 4.3). In particular, $\pi(N) = |UN|^+$.

(2) Let $\mathcal{K}$ be an mAEC with Löwenheim-Skolem number $LS^d(\mathcal{K})$, and let $U : \mathcal{K} \rightarrow \mathbf{Met}$ be the functor assigning to each model its underlying complete metric space: recall that $\mathbf{Met}$ is the category of complete metric spaces and isometric embeddings. Here the relevant notion of size for models $M \in \mathcal{K}$ is not the cardinality of UM, but its density character, which we denote by $\mathrm{dc}(UM)$—this is the reason for the superscript in $LS^d(\mathcal{K})$. By [16] 3.1, $\mathcal{K}$ is $LS^d(\mathcal{K})^+$-accessible with directed colimits. Moreover, for regular $\kappa > LS^d(\mathcal{K})$, a model $N \in \mathcal{K}$ is κ-presentable if and only if $\mathrm{dc}(UN) < \kappa$. In particular, $\pi(N) = \mathrm{dc}(UN)^+$.

In the examples above, and, indeed, in any accessible category with directed colimits, the categories $\mathcal{K}$ are not merely accessible in $LS(\mathcal{K})$ or $LS^d(\mathcal{K})$, but in all larger regular

cardinals, i.e. they are *well-accessible* ([4] 4.1). This is not true in general: for $\mu > \lambda$, a λ-accessible category is μ-accessible if and only if μ is *sharply greater than* λ:

Definition 2.2. For regular cardinals λ and μ, we say that μ is *sharply greater than* λ, denoted $\mu \rhd \lambda$, if $\mu > \lambda$ and for any $\kappa < \mu$ the set $[\kappa]^{<\lambda}$ consisting of subsets of κ of size less than λ, ordered by subset inclusion, contains a cofinal set of size less than μ.

A number of equivalent characterizations of this relation can be given, e.g. in [17] 2.3. We note that, barring the assumption of GCH, this condition on μ and λ is weaker than the cardinal inequality $\mu^\lambda = \mu$. The latter condition is not a bad intuition to keep around, however: following [16] 4.11, if $\mu > 2^\lambda$, then $\mu^+ \rhd \lambda^+$ if and only if $\mu^\lambda = \mu$.

Definition 2.3. Let $F : \mathcal{K} \to \mathcal{L}$ be a functor.

(1) We say that F is λ-*accessible* if it preserves λ-directed colimits and the categories $\mathcal{K}$ and $\mathcal{L}$ are both λ-accessible.
(2) The *powerful image* of F is the smallest subcategory of $\mathcal{L}$ that contains the image of F and is closed under subobjects.

Finally, we recall the definitions of the large cardinals with which we will be working:

Definition 2.4. Given cardinals $\kappa \geq \mu$, we say that κ is μ-*strongly compact* if any κ-complete filter can be extended to a μ-complete ultrafilter. A cardinal κ is *almost strongly compact* if it is μ-strongly compact for all $\mu < \kappa$.

We note that the assumption of μ-strongly compact cardinals for arbitrary μ, preferred in [9], is equivalent to the assumption, preferred in [7], of a proper class of almost strongly compact cardinals (see [9], Proposition 2.4). We note, too, that whereas [9] derives the accessibility of powerful images from these equivalent large cardinal assumptions, [7] shows that this assumption is optimal: the powerful image of any accessible functor is accessible if and only if there is a proper class of almost strongly compact cardinals.

The key result in the argument that follows is Theorem 3.4 in [9], which guarantees that for any well-behaved λ-accessible functor $F : \mathcal{K} \to \mathcal{L}$, its powerful image is κ-accessible for $\mu_\mathcal{L}$-strongly compact κ, where $\mu_\mathcal{L}$ is a cardinal that depends on both λ and $\mathcal{L}$. We are now in a position to describe this $\mu_\mathcal{L}$, following the notation of [9] 3.1.

Let $\mathbf{Pres}_\lambda(\mathcal{L})$ denote a skeleton of the full subcategory of λ-presentable objects in $\mathcal{L}$; that is, a subcategory containing a single object for each isomorphism class of λ-presentable objects. Let $\beta = |\mathbf{Pres}_\lambda(\mathcal{L})|$. Let $\gamma_\mathcal{L}$ be the smallest cardinal such that $\gamma_\mathcal{L} \geq \beta$ and $\gamma_\mathcal{L} \rhd \lambda$. We note that this $\gamma_\mathcal{L}$ will not be too large: if $\lambda < \beta$,

$$\gamma_\mathcal{L} \leq (2^\beta)^+,$$

and clearly $\gamma_\mathcal{L} = \lambda$ otherwise ([1] 2.13(3)). In case $\mathcal{K}$ is well λ-accessible, we may simply take $\gamma_\mathcal{L} = \max(\beta, \lambda)$, noting the parallel between this cardinal and the corresponding bound for AECs used in [3], i.e. $I(\mathcal{K}, \lambda) + \lambda$. Finally, we define $\mu_\mathcal{L} = (\gamma_\mathcal{L}^{<\gamma_\mathcal{L}})^+$. By design, $\mu_\mathcal{L} \rhd \lambda$ ([1] 2.13(5)). This notation is slightly ambiguous, given that $\mathcal{L}$ will be λ-accessible for many λ, but the choice of λ will usually be clear. In the one place where there is a risk of confusion—Theorems 4.4 and 4.5—we specify the cardinal parameter in the subscript, e.g. $\mu_{\mathcal{L},\lambda}$.

For the sake of completeness, we state the theorem of [9] on which we will depend in the sequel:

Theorem 2.5 ([9] 3.4). *Let λ be a regular cardinal and $\mathcal{L}$ an accessible category such that there exists a $\mu_\mathcal{L}$-strongly compact cardinal κ. Then the powerful image of any λ-accessible functor to $\mathcal{L}$ that preserves $\mu_\mathcal{L}$-presentable objects is accessible.*

We need a slight refinement, however. In the proof of this theorem in [9], the only obstacle to the powerful image being κ-accessible is that it may not be the case that $\kappa \unrhd \lambda$. If, as in the theorem above, one were not concerned with the index of accessibility of the powerful image, one could simply pass to $\kappa' > \kappa$ (hence also $\mu_\mathcal{L}$-strongly compact) such that $\kappa' \unrhd \lambda$. As we are concerned with said index, however, we must build in the assumption:

Theorem 2.6. *Let λ be a regular cardinal and $\mathcal{L}$ an accessible category such that there exists a $\mu_\mathcal{L}$-strongly compact cardinal κ. Suppose, moreover, that $\kappa \unrhd \lambda$. Then the powerful image of any λ-accessible functor to $\mathcal{L}$ that preserves $\mu_\mathcal{L}$-presentable objects is κ-accessible.*

We note that if the λ-accessible category $\mathcal{L}$ is well-λ-accessible (that is, if it is accessible in every $\kappa > \lambda$), we may dispense with the sharp inequality above. This is significant because, as already noted, both AECs and mAECs fall under this heading.

Corollary 2.7. *Let λ be a regular cardinal and $\mathcal{L}$ a well-accessible category such that there exists a $\mu_\mathcal{L}$-strongly compact cardinal κ. Then the powerful image of any λ-accessible functor to $\mathcal{L}$ that preserves $\mu_\mathcal{L}$-presentable objects is κ-accessible.*

3. Structural Properties Via Accessible Images

As promised in the introduction, we begin by considering the amalgamation and joint embedding properties, setting aside the disjoint versions for the moment. Through a suitable rewriting of each one in terms of functor categories, we will move the problems to a context where we can apply Theorem 2.6. We thereby obtain Hanf numbers (in fact, the same Hanf number) for each property.

Definition 3.1. Let $\mathcal{K}$ be an accessible category.
(1) We say that $\mathcal{K}$ has the $< \kappa$-*joint embedding property* (or $< \kappa$-*JEP*) if any pair of κ-presentable objects M_0, M_1 extends to a cospan $M_0 \to N \leftarrow M_1$. We say that $\mathcal{K}$ has the *joint embedding property* (or *JEP*) if the same is true for pairs of objects of arbitrary presentability.
(2) We say that $\mathcal{K}$ has the $< \kappa$-*amalgamation property* (or $< \kappa$-*AP*) if any span $N_0 \xleftarrow{f_0} M \xrightarrow{f_1} N_1$ in which the objects are κ-presentable extends to a commutative square

$$
\begin{array}{ccc}
N_0 & \xrightarrow{g_0} & N \\
{\scriptstyle f_0}\big\uparrow & & \big\uparrow{\scriptstyle g_1} \\
M & \xrightarrow{f_1} & N_1
\end{array}
$$

We say that $\mathcal{K}$ has the *amalgamation property* (or *AP*) if the same is true for spans of objects of arbitrary size.

Naturally, these notions reduce to the familiar ones if $\mathcal{K}$ is an AEC or mAEC.

Definition 3.2. We now define the categories and functors we will use in the analysis of amalgamation and joint embedding.
(1) (JEP): The category of pairs M_0, M_1 in $\mathcal{K}$ can be identified with the category of functors from $\bullet \; \bullet$, the discrete category on two objects; that is, with the functor category $\mathcal{K}^{\bullet\bullet}$. We may identify the category of cospans $M_0 \to N \leftarrow M_1$ in $\mathcal{K}$ with the category of

functors from the three-object category $\bullet \to \bullet \leftarrow \bullet$ to $\mathcal{K}$, i.e. $\mathcal{K}^{\bullet \to \bullet \leftarrow \bullet}$. For convenience, we will refer to the latter category as $\mathcal{K}^{cosp}$. Consider the forgetful functor

$$F_J : \mathcal{K}^{cosp} \to \mathcal{K}^{\bullet \bullet}$$

that takes a cospan—a joint embedding diagram—and forgets everything but the leftmost and rightmost objects.

Clearly, the image of F_J is precisely the collection of jointly embeddable pairs in $\mathcal{K}$. We also note that the image of F_J is closed under subobjects, and is therefore powerful.

(2) (AP): Similarly the category of spans can be identified with the functor category $\mathcal{K}^{\bullet \leftarrow \bullet \to \bullet}$, which we denote by $\mathcal{K}^{sp}$. Moreover, the category of commutative squares is morally the same as the category of functors into $\mathcal{K}$ from the category with four objects and morphisms arranged in a commutative square—we denote this functor category by $\mathcal{K}^{csq}$. Again, the forgetful functor

$$F_A : \mathcal{K}^{csq} \to \mathcal{K}^{sp}$$

picks out precisely the amalgamable spans. Once again, the image of F_A is itself powerful.

Remark 3.3. (1) Notice that $\mathcal{K}$ satisfies the $< \kappa$-JEP if and only if the powerful image of F_J contains all pairs of κ-presentable objects, i.e. $\mathbf{Pres}_\kappa(\mathcal{K})^{\bullet \bullet}$. By contrast, $\mathcal{K}$ satisfies the JEP if and only if F_J is an essential surjection; that is, the image is all of $\mathcal{K}^{\bullet \bullet}$.

(2) Similarly, $\mathcal{K}$ satisfies the $< \kappa$-AP if and only if the image of F_A contains $\mathbf{Pres}_\kappa(\mathcal{K})^{sp}$. Furthermore, $\mathcal{K}$ satisfies the AP if and only if F_A is an essential surjection.

By [1] 2.39, whenever $\mathcal{K}$ is accessible and $\mathcal{A}$ is a small category, $\mathcal{K}^{\mathcal{A}}$ is accessible as well, although the index of accessibility may increase. As the categories $\mathcal{A}$ invoked here are very simple, things are decidedly less complicated:

Lemma 3.4. *If $\mathcal{K}$ is λ-accessible, so are $\mathcal{K}^{\bullet \bullet}$, $\mathcal{K}^{sp}$, $\mathcal{K}^{cosp}$, and $\mathcal{K}^{csq}$. In each case, a diagram is λ-presentable if and only if each of the objects in the diagram is λ-presentable. Moreover, if $\mathcal{K}$ is well-λ-accessible, $\mathcal{K}^{\bullet \bullet}$, $\mathcal{K}^{sp}$, $\mathcal{K}^{cosp}$, and $\mathcal{K}^{csq}$ are well-λ-accessible.*

Proof. We note that in any functor category $\mathcal{K}^{\mathcal{A}}$, colimits are computed componentwise, meaning that if $\mathcal{K}$ has all colimits of a particular shape, so does $\mathcal{K}^{\mathcal{A}}$. In particular, if $\mathcal{K}$ is λ-accessible, it has all λ-directed colimits, and hence so does $\mathcal{K}^{\mathcal{A}}$. So in proving λ-accessibility of our functor categories, we need only show that each contains a set of λ-presentable objects, and that any object can be realized as a λ-directed colimit of such objects.

In the case of $\mathcal{K}^{\bullet \bullet}$, this is clear: for any regular cardinal μ, an object in $\mathcal{K}^{\bullet \bullet}$ is λ-presentable if and only if the pair of elements M_0 and M_1 in $\mathcal{K}$ picked out by the two objects of $\bullet \bullet$ are both λ-presentable in $\mathcal{K}$. Thus there is clearly a set of λ-presentable objects, up to isomorphism. By λ-accessiblity of $\mathcal{K}$, we can express each of M_0 and M_1 as colimits of λ-directed diagrams of λ-presentable $\mathcal{K}$-objects, say $M_0 = \mathrm{colim}_{i \in I} M_0^i$ and $M_1 = \mathrm{colim}_{j \in J} M_1^j$. Then the diagram of pairs (M_0^i, M_1^j) indexed by $I \times J$ is λ-directed, and clearly has colimit (M_0, M_1). So we are done.

For the remaining categories, we reduce the problem to the fact that if $\mathcal{K}$ is λ-accessible, its category of morphisms $\mathcal{K}^{\to}$ is also λ-accessible, and the λ-presentable morphisms are precisely those with λ-presentable domain and codomain (see [1] Ex. 2c). Indeed, since $\mathcal{K}^{csq} = (\mathcal{K}^{\to})^{\to}$, we immediately get the result for the category of commutative squares.

Let $\mathcal{K}_0$ be the category obtained by formally adjoining an initial object 0 to $\mathcal{K}$. In particular, $\mathcal{K}_0$ is obtained from $\mathcal{K}$ by adding the new object 0 and a unique morphism $0 \to M$ for each M in $\mathcal{K}$. There are no morphisms $M \to 0$ for M in $\mathcal{K}$. Then $\mathcal{K}_0$ is λ-accessible

whenever $\mathcal{K}$ is λ-accessible, and has precisely the same λ-presentable objects, with the addition of the finitely presentable object 0. The result for $\mathcal{K}^{cosp}$ follows from identifying the cospans $M_0 \to N \leftarrow M_1$ in $\mathcal{K}$ with the subcategory of $\mathcal{K}_0^{csq}$ consisting of squares

$$
\begin{array}{ccc}
M_0 & \longrightarrow & N \\
\uparrow & & \uparrow \\
0 & \longrightarrow & M_1
\end{array}
$$

Working in $\mathcal{K}_0^{csq}$, it is easy to verify that any cospan is a λ-directed colimit of cospans of λ-presentable objects, and that such cospans of λ-presentable objects are λ-presentable. It follows, of course, that $\mathcal{K}^{cosp}$ is λ-accessible.

Analogously, we can represent spans $N_0 \leftarrow M \to N_1$ as squares

$$
\begin{array}{ccc}
N_0 & \longrightarrow & 1 \\
\uparrow & & \uparrow \\
M & \longrightarrow & N_1
\end{array}
$$

in the category $\mathcal{K}_1$ obtained by formally adjoining a terminal object 1 to $\mathcal{K}$. This yields the result for $\mathcal{K}^{sp}$.

The moreover clause follows in a similar fashion. $\qquad \square$

From this perspective, the proofs of the existence of a Hanf number for JEP and AP in accessible categories are essentially trivial, modulo Theorem 2.6.

Theorem 3.5. *Let $\mathcal{K}$ be a λ-accessible category, and let κ be a $\mu_\mathcal{K}$-strongly compact cardinal with $\kappa \trianglerighteq \lambda$. If $\mathcal{K}$ has the $< \kappa$-JEP, then it has the JEP.*

Proof. Consider the functor $F_J : \mathcal{K}^{cosp} \to \mathcal{K}^{\bullet\bullet}$. Each of $\mathcal{K}^{\bullet\bullet}$ and $\mathcal{K}^{cosp}$ are λ-accessible, by Lemma 3.4, and clearly F_J preserves λ-directed colimits. That is, F_J is λ-accessible.

Notice that, since $\mathcal{K}$ and $\mathcal{K}^{\bullet\bullet}$ are both λ-accessible and have subcategories of λ-presentables of the same size, $\mu_\mathcal{K} = \mu_{\mathcal{K}^{\bullet\bullet}}$. As the $\mu_\mathcal{K}$-presentable objects in $\mathcal{K}^{cosp}$ are precisely the cospans whose objects are $\mu_\mathcal{K}$-presentable, they are certainly preserved by F_J.

Hence by Theorem 2.6, the powerful image of F_J is κ-accessible. As noted in Definition 3.2(1), this powerful image consists of precisely the jointly embeddable pairs of objects in $\mathcal{K}$. Let (M_0, M_1) be an arbitrary pair of objects in $\mathcal{K}$, i.e. an arbitrary object of $\mathcal{K}^{\bullet\bullet}$. Because $\kappa \trianglerighteq \lambda$, $\mathcal{K}^{\bullet\bullet}$ is κ-accessible, meaning that (M_0, M_1) in $\mathcal{K}^{\bullet\bullet}$ can be realized as a κ-directed colimit of pairs of κ-presentable objects; that is,

$$(M_0, M_1) = \mathrm{colim}_{i \in I}(M_0^i, M_1^i)$$

with I a κ-directed set.

Since $\mathcal{K}$ satisfies the $(< \kappa)$-JEP, by assumption, Remark 3.3(1) implies that each (M_0^i, M_1^i) lies in the powerful image of F_J. As the powerful image of F_J is κ-accessible, it is closed under κ-directed colimits, and therefore contains (M_0, M_1). In particular, M_0 and M_1 admit a joint embedding. So $\mathcal{K}$ satisfies the JEP. $\qquad \square$

Theorem 3.6. *Let $\mathcal{K}$ be a λ-accessible category, and let κ be a $\mu_\mathcal{K}$-strongly compact cardinal with $\kappa \trianglerighteq \lambda$. If $\mathcal{K}$ has the $< \kappa$-AP, then it has the AP.*

Proof. The theorem follows by the same argument as Theorem 3.5, with $\mathcal{K}^{sp}$ and $\mathcal{K}^{csq}$ in place of $\mathcal{K}^{\bullet\bullet}$ and $\mathcal{K}^{cosp}$, respectively. $\qquad \square$

We have simpler statements in case $\mathcal{K}$ is well-accessible:

Theorem 3.7. *Let $\mathcal{K}$ be a well λ-accessible category, and let κ be a $\mu_{\mathcal{K}}$-strongly compact cardinal. If $\mathcal{K}$ has the $< \kappa$-JEP ($< \kappa$-AP), then it satisfies the JEP (AP).*

Proof. This follows by Corollary 2.7 and the moreover clause of Lemma 3.4. ☐

Remark 3.8. We note that Theorem 2.0.5 in [3], which covers the case in which $\mathcal{K}$ is an AEC, assumes that κ is a strongly compact cardinal greater than $LS(\mathcal{K})$. In particular, κ must be strongly inaccessible. In this case $\mathcal{K}$ is $LS(\mathcal{K})^+$-accessible, and the cardinals $\lambda = LS(\mathcal{K})^+$, $\beta = |\mathbf{Pres}_\lambda(\mathcal{K})|$, $\gamma_{\mathcal{K}} = \max(\lambda, \beta)$, and $\mu_{\mathcal{K}} = (\gamma_{\mathcal{K}}^{<\gamma_{\mathcal{K}}})^+$ are all less than κ. It follows that κ is $\mu_{\mathcal{K}}$-strongly compact. Thus the specialization of Theorem 3.7 to the AEC case involves, *a priori*, a weaker large cardinal assumption: we note, however, that, per [7] 2.2, it remains open whether this assumption is genuinely (consistently) weaker.

4. Disjoint Analogues

Capturing the disjoint versions of amalgamation and joint embedding requires more care, insofar as we must frame the notion of disjointness in an abstract category. The disjoint JEP for an AEC $\mathcal{K}$, for example, asserts that any models M_0 and M_1 can be completed to a cospan $M_0 \xrightarrow{g_0} N \xleftarrow{g_1} M_1$ where the ranges of the embeddings are disjoint, i.e. the following diagram is a pullback in **Set**:

$$
\begin{array}{ccc}
UM_0 & \xrightarrow{U(g_0)} & UN \\
\uparrow & & \uparrow{\scriptstyle U(g_1)} \\
\emptyset & \longrightarrow & UM_1
\end{array}
$$

Here the essential role is played by the empty set, which is a *strictly initial object* in **Set**: it is initial, and any morphism $M \to \emptyset$ is an isomorphism. The difficulty is that we cannot expect $\mathcal{K}$ itself to have a strict initial object (or, indeed, pullbacks), so we cannot necessarily pull this diagram back to one in $\mathcal{K}$, considered as an abstract category.

We skirt this difficulty by broadening the notion of disjointness to incorporate categories that may lack a strict initial object: A cospan $M_0 \to N \leftarrow M_1$ in a category $\mathcal{K}$ will be called *disjoint* if it cannot be completed to a commutative square or

$$
\begin{array}{ccc}
M_0 & \longrightarrow & N \\
\uparrow & & \uparrow \\
\emptyset & \longrightarrow & M_1
\end{array}
$$

is the only completion, where $\emptyset$ is strictly initial in $\mathcal{K}$. Note that this square, if it exists, is a pullback square in $\mathcal{K}$.

The category of disjoint cospans in $\mathcal{K}$ will be denoted $\mathcal{K}^{dcosp}$, while $\mathcal{K}^{psq}$ will be the category of pullback squares in $\mathcal{K}$.

Lemma 4.1. *Let $\mathcal{K}$ be a λ-accessible category. Then the category $\mathcal{K}^{dcosp}$ is λ-accessible and $\mathcal{K}^{psq}$ is λ^+-accessible.*

Proof. Any morphism in $\mathcal{K}^{cosp}$ to a disjoint cospan has disjoint domain. Thus the result for $\mathcal{K}^{dcosp}$ follows from that for $\mathcal{K}^{cosp}$ (see 3.4).

Following [17] 6.2.1(i), $\mathcal{K}^{psq}$ is given by the pseudopullback

$$\begin{array}{ccc} \mathcal{K}^{cosp} & \xrightarrow{\ H\ } & \mathbf{Set}^{\mathcal{A}^{\mathrm{op}}} \\ \big\uparrow & & \big\uparrow{\scriptstyle E} \\ \mathcal{K}^{psq} & \longrightarrow & \mathcal{K} \end{array}$$

where E is the canonical functor $\mathcal{K} \to \mathbf{Set}^{\mathcal{A}^{\mathrm{op}}}$ given by $EM = \mathcal{K}(-, M) \restriction \mathcal{A}^{\mathrm{op}}$, where $\mathcal{A}$ denotes $\mathbf{Pres}_\lambda(\mathcal{K})$ (see [1] 1.25), and H is the composition

$$\mathcal{K}^{cosp} \to (\mathbf{Set}^{\mathcal{A}^{\mathrm{op}}})^{cosp} \to \mathbf{Set}^{\mathcal{A}^{\mathrm{op}}}$$

where the first functor is E^{sp} and the second functor sends a cospan in the presheaf category $\mathbf{Set}^{\mathcal{A}^{\mathrm{op}}}$ to its pullback—note that, while $\mathcal{K}$ may not have pullbacks, $\mathbf{Set}^{\mathcal{A}^{\mathrm{op}}}$ always will. Since E and H preserve λ-directed colimits and λ-presentable objects, $\mathcal{K}^{psq}$ is λ^+-accessible (see [18] 2.2). □

Definition 4.2. A category $\mathcal{K}$ has the *disjoint embedding property* (or *DJEP*) if any pair of objects M_0 and M_1 can be completed to a disjoint cospan $M_0 \to N \leftarrow M_1$.

$\mathcal{K}$ has the *disjoint amalgamation property* (or *DAP*) if any span $N_0 \leftarrow M \to N_1$ can be completed to a pullback square

$$\begin{array}{ccc} N_0 & \longrightarrow & N \\ \big\uparrow & & \big\uparrow \\ M & \longrightarrow & N_1 \end{array}$$

The meanings of $< \kappa\text{-}DJEP$ and $< \kappa\text{-}DAP$ are clear.

Remark 4.3. Those familiar with the AEC versions of these properties will recognize our formulation of the DAP as a straightforward generalization. The DJEP is as formulated here is slightly more delicate: in fact, even in case $\mathcal{K}$ is an AEC, it yields a notion still more general than the *nearly disjoint JEP*, in which any two models can be embedded in such a way that they overlap only on the submodel generated from the constant symbols. It suffices for our purposes here, however. That is, if an AEC $\mathcal{K}$ satisfies the $< \kappa$-DJEP for any κ, then it cannot contain any strictly initial object other than the empty model. If it contains no strictly initial object, we formally adjoin the empty model to serve this purpose, a step which does not change the accessibility of $\mathcal{K}$. In either case, we have the $< \kappa$-DJEP in the sense of Definition 4.2. When we show below that $\mathcal{K}$ therefore satisfies the global DJEP, again in the sense of Definition 4.2, we can be certain that this translates into the familiar DJEP for $\mathcal{K}$ as an AEC, simply because we still have the empty model as the strictly initial object.

Theorem 4.4. *Let $\mathcal{K}$ be a λ-accessible category. If κ is a $\mu_{\mathcal{K},\lambda}$-strongly compact cardinal with $\kappa \trianglerighteq \lambda$ and $\mathcal{K}$ has the $< \kappa$-DJEP, then it satisfies the DJEP. If κ is $\mu_{\mathcal{K},\lambda^+}$-strongly compact with $\kappa \trianglerighteq \lambda^+$ and $\mathcal{K}$ has the $< \kappa$-DAP, then it satisfies the DAP.*

Proof. For DJEP we use the λ-accessible functor

$$\mathcal{K}^{dcosp} \to \mathcal{K}^{\bullet\bullet}$$

in the same way as in 3.5. For DAP we use the functor

$$\mathcal{K}^{psq} \to \mathcal{K}^{sp}$$

from the proof of 4.1. Since this functor is λ^+-accessible (see [10] 3.1), we proceed as in 3.6. $\square$

In the case of DAP, $\mu_\mathcal{K}$ is taken for λ^+ and not for λ. Again, we have simpler statements for well-accessible categories.
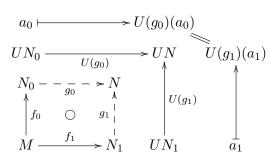
Theorem 4.5. *Let $\mathcal{K}$ be a well λ-accessible category, and let κ be $\mu_{\mathcal{K},\lambda}$-strongly compact. If $\mathcal{K}$ has the $< \kappa$-DJEP, then it satisfies the DJEP. Similarly, if κ is $\mu_{\mathcal{K},\lambda^+}$-strongly compact and $\mathcal{K}$ has the $< \kappa$-DAP, it has the AP.*

Remark 4.6. As noted in Remark 3.8, the specialization of Theorem 4.5 to AECs involves an *a priori* weaker large cardinal hypothesis than [3] 2.0.5.
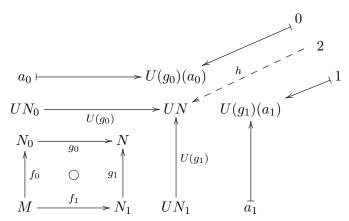
5. METRIC TAMENESS

We turn now to the most technical and involved of the results to be presented here, the *d*-tameness of metric AECs under the assumption of sufficiently strongly compact cardinals. We will again provide the necessary basic definitions; extensive background and motivation for this question, and for the broader study of mAECs, can be found in, e.g. [12] and [20].

As suggested in the introduction, the method to be used here is a generalization of that employed in [15] to establish the discrete version of this result, namely the tameness of AECs under the assumption of arbitrarily large strongly compact cardinals. We dwell briefly on that case here, largely as an opportunity to recast it in a form better suited to the needed generalization. In the discrete context, a Galois type over a model M is an equivalence class of pairs (f, a), where $f : M \to N$ is a $\mathcal{K}$-embedding and $a \in UN$. Here two pairs (f_0, a_0) and (f_1, a_1) are said to be equivalent if there is an amalgam of f_0 and f_1 in which a_0 and a_1 are identified, i.e. there are $g_0 : N_0 \to N$ and $g_1 : N_1 \to N$ such that the following diagram commutes:

This is an equivalence relation provided $\mathcal{K}$ has amalgamation. Equivalently, two pairs (f_0, a_0) and (f_1, a_1) represent distinct types just in case for any amalgam $g_0 : N_0 \to N$ and $g_1 : N_1 \to N$, there is a monomorphism h from the two point set $2 = \{0, 1\}$ to UN such $h(0) = U(g_0)(a_0)$ and $h(1) = U(g_1)(a_1)$, as in the diagram below.

$$
\begin{array}{c}
\\
a_0 \longmapsto U(g_0)(a_0) \quad\quad\quad 0 \\
\\
UN_0 \xrightarrow{\ U(g_0)\ } UN \quad U(g_1)(a_1) \quad 2,\ 1 \\
N_0 \xrightarrow{\ g_0\ } N \\
f_0 \quad \bigcirc \quad g_1 \qquad U(g_1) \\
M \xrightarrow{\ f_1\ } N_1 \quad UN_1 \quad a_1
\end{array}
$$

Although the argument of [15] is built around the former characterization of equivalence, the latter is of greater significance for the discussion here: we measure distinctness of types in AECs by monomorphisms into amalgams from 2, the natural source for such probes in **Set**. What if, instead, we consider the metric case, where structures have underlying complete metric spaces, rather than sets? We will want to measure not inequivalence of pairs, but the distance between them. Hence the natural objects with which to probe in **Met** are the two point metric spaces of diameter ϵ, 2_ϵ, for $\epsilon \in \mathbb{R}_{>0}$.

There is a great deal to be made precise. We recall first that for all the minor differences between metric and discrete AECs—chiefly that the operative notion of size is density character rather than cardinality and that directed colimits need not be concrete—the fundamental definition of Galois types is precisely the same: as equivalence classes of pairs $(f_0 : M \to N_0, a_0)$, $(f_1 : M \to N_1, a_1)$. One can, however, speak not merely of the equality or inequality of types, as in the discrete case, but of the distance between them. This distance is typically defined as the Hausdorff distance between the corresponding orbits in a monster model $\mathfrak{C}$: the assumption of AP, JEP, and the existence of arbitrarily large models guarantees that $\mathfrak{C}$ exists; the *perturbation property* ([12]), also known as the *continuity of types property* ([20]), guarantees that this distance is a metric. We make precisely the same assumptions here:
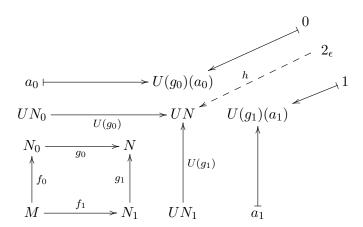
Assumption 5.1. *From now on, we assume that our mAEC $\mathcal{K}$ has the AP and JEP, has arbitrarily large models, and satisfies the perturbation property.*

We need an alternative, albeit equivalent, definition of the distance between types, however, for reasons which will soon become clear.

Definition 5.2. Let $\mathcal{K}$ be an mAEC (satisfying the conditions of Assumption 5.1). Given types $p_0 = \text{ga-tp}(f_0, a_0)$ and $p_1 = \text{ga-tp}(f_1, a_1)$ over $M \in \mathcal{K}$, $f_i : M \to N_i$, we define the *distance between p_0 and p_1*, denoted $d(p_0, p_1)$, to be the infimum of the distances between the images of a_0 and a_1 in all amalgams of f_0 and f_1. That is,

$$d(p_0, p_1) = \inf\{d_N(U(g_0)(a_0), U(g_1)(a_1)) \mid g_i : N_i \to N \text{ and } g_0 f_0 = g_1 f_1\}$$

Notice that $d_N(U(g_0)(a_0), U(g_1)(a_1)) = \epsilon$ if and only if there is an isometry $h : 2_\epsilon \to N$, where $2_\epsilon = \{0, 1\}$ is the two point space of diameter ϵ, such that $h(0) = U(g_0)(a_0)$ and $h(1) = U(g_1)(a_1)$. This is witnessed by a diagram of the following shape:

Definition 5.3. For each $\epsilon > 0$, we define $\mathcal{L}_\epsilon$ to be the category of (pointed) commutative squares in **Met** equipped with a specified such isometry h. Let $\mathcal{L}$ denote the category of pairs (f_0, a_0) and (f_1, a_1) or, in the notation of [15], (f_0, f_1, a_0, a_1). Finally, we denote by $G_\epsilon : \mathcal{L}_\epsilon \to \mathcal{L}$ the obvious forgetful functor.

In the proof of d-tameness that follows, we will use the eventual κ-accessibility of the powerful images of these functors G_ϵ to show that if two types over a model are separated by distance at least ϵ, then this fact is witnessed by restrictions to a model of density character κ, in the sense that the restrictions are also at distance at least ϵ. That is, we show that any mAEC is strongly d-tame:

Definition 5.4. We say that an mAEC $\mathcal{K}$ is κ-d-tame if for every $\epsilon > 0$ there is a $\delta > 0$ such that for any $M \in \mathcal{K}$ and Galois types p, q over M, if $d(p, q) \geq \epsilon$ then there exists $N \preceq_\mathcal{K} M$ with $\mathrm{dc}(N) \leq \kappa$ such that $d(p \restriction N, q \restriction N) \geq \delta$. We say that $\mathcal{K}$ is *strongly κ-d-tame* if the above holds with $\delta = \epsilon$.

We say that $\mathcal{K}$ is *(strongly) d-tame* if it is (strongly) κ-d-tame for some κ.

Theorem 5.5. *Suppose that for every cardinal μ there exists a μ-strongly compact cardinal. Then every mAEC is strongly d-tame.*

Proof. Let $\mathcal{K}$ be an mAEC, and let $U : \mathcal{K} \to$ **Met** be the forgetful functor to the category of complete metric spaces and isometric embeddings. Let $\mathcal{L}$, $\mathcal{L}_\epsilon$, and $G_\epsilon : \mathcal{L}_\epsilon \to \mathcal{L}$ be as described in Definition 5.3. Then both $\mathcal{L}, \mathcal{L}_\varepsilon$ are accessible and the forgetful functor $G_\varepsilon : \mathcal{L}_\varepsilon \to \mathcal{L}$ is accessible as well. In fact, $\mathcal{L}$ is the category of commutative squares

$$
\begin{array}{ccc}
N_0 & \xrightarrow{g_0} & N \\
{\scriptstyle f_0}\big\uparrow & & \big\uparrow{\scriptstyle g_1} \\
M & \xrightarrow{f_1} & N_1
\end{array}
$$

equipped with elements $a_0 \in UN_0$ and $a_1 \in UN_1$. Thus its accessibility follows from that of $\mathcal{K}^{csq}$ and from the preservation of $\aleph_1$-directed colimits by the forgetful functor $U : \mathcal{K} \to$ **Met**. Similarly for $\mathcal{L}_\epsilon$ where one adds the condition that $d_N(g_0(a_0), g_1(a_1)) = \epsilon$. Clearly, if the quadruple (f_0, f_1, a_0, a_1) belongs to the image of G_ε then the distance of the corresponding Galois types is $\leq \varepsilon$. Thus the distance $d(p_0, p_1)$ of Galois types $p_0 = (f_0, a_0)$ and $p_1 = (f_1, a_1)$

is equal to the infimum of the set

$$\{\epsilon > 0 \,|\, (f_0, f_1, a_0, a_1) \in G_\varepsilon(\mathcal{L}_\varepsilon)\}$$

It is easy to see that the full image of G_ε is a sieve: for a morphism $(u, v, w) : (f_0', f_1', a_0', a_1') \to (f_0, f_1, a_0, a_1)$, i.e.
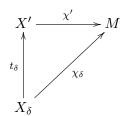
$$
\begin{array}{ccccc}
UN_0' \ni a_0' & N_0' \xleftarrow{f_0'} M \xrightarrow{f_1'} N_1' & a_1' \in UN_1' \\
U(u) \downarrow & u \downarrow \quad v \downarrow \quad w \downarrow \quad U(w) \downarrow \\
UN_0 \ni a_0 & N_0 \xleftarrow{f_0} M \xrightarrow{f_1} N_1 & a_1 \in UN_1
\end{array}
$$

with $(f_0, f_1, a_0, a_1) \in G_\varepsilon(\mathcal{L}_\varepsilon)$, we have $(f_0', f_1', a_0', a_1') \in G_\varepsilon(\mathcal{L}_\varepsilon)$ as well. Thus the full image of G_ε is powerful. Let λ be a regular cardinal such that all functors G_ε are λ-accessible and preserve λ-presentable objects. Let $\mu_\mathcal{L}$ be the cardinal from defined in Section 2—just before Theorem 2.6—and let κ be a $\mu_\mathcal{L}$-strongly compact cardinal. Then, following Theorem 2.6, the full image of any G_ε is κ-accessible.

Assume that the distance $d(p_0\chi, p_1\chi) < \varepsilon$ for each morphism $\chi : X \to M$ in $\mathcal{K}$ with X κ-presentable, where we write $p_i\chi$ as an abbreviation for the type represented by $(f_i\chi, a_i)$. We claim that there are cofinally many $\chi : X \to M$ giving the same distance $d(p_0\chi, p_1\chi) = \delta$ for some $\delta < \varepsilon$. In fact, assuming the contrary, for any $\delta < \varepsilon$ there is $\chi_\delta : X_\delta \to M$ with no factorization $\chi_\delta = \chi' t_\delta$ such that $d(p_0\chi', p_1\chi') = \delta$. There are at most $2^{\aleph_0}$ many witnesses χ_δ, and given that M is a κ-directed colimit of the $\chi : X \to M$, there must therefore be a map $\chi' : X' \to M$ through which all of the χ_δ factor. That is, for every $\delta < \varepsilon$, there is $t_\delta : X_\delta \to X'$ such that the following diagram commutes:

$$
\begin{array}{ccc}
X' & \xrightarrow{\chi'} & M \\
t_\delta \uparrow & \nearrow \chi_\delta & \\
X_\delta & &
\end{array}
$$

Note that $d(p_0\chi', p_1\chi') = \delta'$ for some $\delta' < \varepsilon$. Of course, $\chi_{\delta'}$ must factor through χ', which is a contradiction.

Let $\mathcal{D}$ denote the resulting cofinal set. For each $\chi \in \mathcal{D}$, there is a sequence of $\delta_n > \delta$ converging to δ such that $(p_0\chi, p_1\chi)$ belongs to the image of G_{δ_n} for each n. Since $(2^{\aleph_0})^{\aleph_0} = 2^{\aleph_0} < \kappa$, there is a cofinal subset of $\mathcal{D}$ realizing the same sequence δ_n. Otherwise, for any sequence $\delta_* = (\delta_n)_n$ there is χ_{δ_*} with no factorization $\chi_{\delta_*} = \chi' t$ such that $(p_0\chi', p_1\chi')$ does not belong to the image of G_{δ_n} for each n. Since all of the χ_{δ_*} factorize through some χ' we get a contradiction, as above. Thus (p_0, p_1) belongs to the image of G_{δ_n} for each n, by accessibility of the images. Hence $d(p_0, p_1) \leq \delta$. Consequently, $d(p_0, p_1) < \varepsilon$. $\square$

Noting that a strongly d-tame $\mathcal{K}$ is d-tame in the sense of [21], we have:

Corollary 5.6. *Suppose that for every cardinal μ there exists a μ-strongly compact cardinal. Then every mAEC is d-tame.*

This last result forms half of Theorem 3.11 in [8], although it appears that their proof should entail strong d-tameness as well. A more substantive difference is that, while they

rely on a delicate argument involving metric ultraproducts, the argument above more clearly illustrates the common underlying structure of the discrete and continuous problems, and points the way to further generalizations.

6. Extensions

We note two clear directions in which the results of this paper can be extended. First, the focus on joint embedding and amalgamation in Sections 3 and 4 is due entirely to practical considerations: these properties, amounting to the existence of coproduct-like cones over pairs and pushout-like cones over spans, are the ones of greatest interest to model theorists. The form of argument via accessible images is vastly more general, though. For example, while model theorists work almost exclusively with monomorphisms, meaning that coequalizer-like (or quotient-like) cones over pairs of parallel arrows are rarely of interest, these make perfect sense in a general accessible category. We might say, for example, that an accessible category has the coequalizing property (CP) if every parallel pair of arrows $f, g : M \to N$ is coequalized by a map $h : N \to N'$, i.e. $hf = hg$. The existence of such an h can again be formulated in terms of the image of a forgetful functor $F_C : \mathcal{K}^{\mathcal{A}'} \to \mathcal{K}^{\mathcal{A}}$, with $\mathcal{A}' = \bullet \rightrightarrows \bullet \to \bullet$ (with both compositions equal) and $\mathcal{A} = \bullet \rightrightarrows \bullet$. As the image is again closed under subobjects, i.e. powerful, we obtain a Hanf number for the CP, just as in Theorem 3.5 and 3.6. There are is an abundance of possible analogues, in fact, limited only by this last stipulation: the image of the forgetful functor, i.e. the class of completable diagrams, must be closed under subobjects.

We note also that ongoing work with Zambrano suggests that the results of Section 5 can be extended well beyond the case of mAECs. Of particular interest are classes of structures with underlying Q-sets, with Q a suitably cocontinuous quantale, in roughly the sense of [13]. This would mean immediate tameness results across a wide array of fields, including fuzzy semantics, model theory over probability spaces, and model theory in categories of sheaves over sufficiently well-behaved spaces.

References

[1] J. Adámek and J. Rosický. *Locally Presentable and Accessible Categories*. Cambridge University Press, 1994.

[2] J. Baldwin. *Categoricity*. AMS, 2009.

[3] J. Baldwin and W. Boney. Hanf numbers and presentation theorems in AECs. To appear in *Beyond First Order Model Theory*, ed. J. Iovino, CRC Press. arXiv:1511.02935v2.

[4] T. Beke and J. Rosický. Abstract elementary classes and accessible categories. *Ann. Pure Appl. Log.*, 163:2008-2017, 2012.

[5] W. Boney. Tameness from large cardinal axioms. *J. Symb. Log.*, 79(4):1092-1119, 2014.

[6] W. Boney, R. Grossberg, M. Lieberman, J. Rosický, and S. Vasey. μ-abstract elementary classes and other generalizations. *J. Pure Appl. Alg.*, 220(9):3048-3066, 2016.

[7] W. Boney and S. Unger. Large cardinal axioms from tameness in AECs. To appear in *Proc. AMS*. arXiv:1509.01191v3.

[8] W. Boney and P. Zambrano. Around the set-theoretical consistency of d-tameness of metric abstract elementary classes. arXiv:1509.01191v1.

[9] A. Brooke-Taylor and J. Rosický. Accessible images revisited. *Proc. AMS*, 145(3):1317-1327.

[10] B. Chorny and J. Rosický. Locally class-presentable and class-accessible categories. *J. Pure Appl. Alg.*, 216: 2113-2125, 2012.

[11] P. Eklof and A. Mekkler. *Almost Free Modules: Set-theoretic methods*. North Holland, 1990.

[12] A. Hirvonen and T. Hyttinen. Categoricity in homogeneous complete metric spaces. *Arch. Math. Log.*, 48:269-322, 2009.

[13] R. Flagg. Quantales and continuity spaces. *Alg. Univ.*, 37:257-276, 97.

[14] M. Lieberman. Category theoretic aspects of abstract elementary classes. *Ann. Pure Appl. Log.*, 162:903-915, 2011.

[15] M. Lieberman and J. Rosický. Classification theory for accessible categories. *J. Symb. Log.*, 81(1):151-165, 2016.

[16] M. Lieberman and J. Rosický. Metric abstract elementary classes as accessible categories. To appear in the *J. Symb. Log.*. arXiv:1504.02660v5.

[17] M. Makkai and R. Paré. *Accessible Categories: The Foundations of Categorical Model Theory*. AMS, 1989.

[18] G. Raptis and J. Rosický. The accessibility rank of weak equivalences. *Theory Appl. Categ.*, 30:687-703, 2015.

[19] J. Rosický. Accessible categories, saturation and categoricity. *J. Symb. Log.*, 62:891-901, 1997.

[20] A. Villaveces and P. Zambrano. Limit models in metric abstract elementary classes: the categorical case. *Math. Log. Quart.*, 62(4-5):319-334, 2016.

[21] P. Zambrano. Around superstability in metric abstract elementary classes. PhD thesis.