

LARGE CARDINAL AXIOMS FROM TAMENESS IN AECs

WILL BONEY AND SPENCER UNGER

ABSTRACT. We show that various tameness assertions about abstract elementary classes imply the existence of large cardinals under mild cardinal arithmetic assumptions. For instance, we show that if κ is an uncountable cardinal such that $\mu^\omega < \kappa$ for every $\mu < \kappa$ and every AEC with Löwenheim-Skolem number less than κ is $< \kappa$ -tame, then κ is almost strongly compact. This is done by isolating a class of AECs that exhibits tameness exactly when sufficiently complete ultrafilters exist.

1. INTRODUCTION

The birth of modern model theory is often said to be Morley's proof [12] of what was then called the Loś Conjecture. This is now called Morley's Categoricity Theorem. It is only natural that this same question be an important test question when studying nonelementary model theory. In one of the most popular contexts for this study, Abstract Elementary Classes, this question is known as Shelah's Categoricity Conjecture.

While still open, there are many partial results towards this conjecture that add various model-theoretic and set-theoretic assumptions. The most relevant for this discussion is the first author's [5, Theorem 7.5], which shows that if there are class many strongly compact cardinals, then any Abstract Elementary Class (AEC) that is categorical in *some* high enough successor cardinal is categorical in every high enough cardinal. One of the central concepts in the proof is the notion of tameness, which says roughly that if two types differ, then they differ over some small subset of their domain. Types here do not have the syntactic form familiar from first order logic, since AECs lack syntax. Instead, a semantic version of type (called Galois or orbital type) is introduced as, roughly, the orbit of elements under automorphisms of a sufficiently homogeneous (or monster) model fixing the domain. In practice tameness has two cardinal parameters, the size of the domain of the types and the cardinal measuring how small the subset of the domain must be.

A key instance of the advances in the first author's work is the following, which with a little more work allows the application of previous results of Shelah [14] and Grossberg and VanDieren [7] to obtain a version of Shelah's categoricity conjecture.

Fact 1.1 ([5].4.5). *If \mathbb{K} is an AEC with $LS(\mathbb{K}) < \kappa$ and κ is strongly compact, then \mathbb{K} is $< \kappa$ -tame.*

Sections 5 and 6 of [5] give similar theorems for measurable and weakly compact cardinals. The main theorems of this paper give converses to these results under mild cardinal arithmetic assumptions. We state the following theorem as a sample application of our methods. We prove below that by strengthening the tameness hypothesis we can drop the "almost" from the conclusion of the theorem.

Theorem. *Let κ be uncountable such that $\mu^\omega < \kappa$ for every $\mu < \kappa$.*

- (1) *If $\kappa^{<\kappa} = \kappa$ and every AEC with Löwenheim-Skolem number less than κ is $(< \kappa, \kappa)$ -tame, then κ is almost weakly compact.*
- (2) *If every AEC with Löwenheim-Skolem number less than κ is κ -local, then κ is almost measurable.*
- (3) *If every AEC with Löwenheim-Skolem number less than κ is $< \kappa$ -tame, then κ is almost strongly compact.*

The first step in this direction is Shelah [13], where the measurable version appears as Theorem 1.3. The example constructed in this paper is a generalization of Shelah's. Note that Shelah's proof is essentially correct, but requires minor correction (see Remark 4.6 for a discussion).

The proof of the main theorems all follow the same plan, which we outline here. First, Section 2 codes large cardinals into a combinatorial statement $\#(\mathcal{D}, \mathcal{F})$ (see Definition 2.4). Then Section 3 defines two structures H_1 and H_2 such that corresponding small substructures of them are isomorphic, but H_1 and H_2 are only isomorphic if the relevant $\#(\mathcal{D}, \mathcal{F})$ holds. Finally, Section 4 defines an AEC \mathbb{K}_σ that contains H_1 and H_2 and codes their isomorphism (and the isomorphism of their substructures) into equality of Galois types. This forms the connection between large cardinals and equality of Galois types.

Our work has an immediate application to category theory. Makkai and Paré proved a theorem [11] about the accessibility of powerful images from the assumption of class many strongly compact cardinals and Lieberman and Rosicky [9] later applied this to AECs to give an alternate proof of Fact 1.1 above. Brooke-Taylor and Rosicky [6] have recently weakened the hypotheses of Makkai and Paré's result to almost strongly compact and our result completes the circle and shows that the conclusion of Makkai and Paré's result is actually a large cardinal statement in disguise. See Corollary 4.14 and the surrounding discussion.

Turning back to model theory, this shows that any attempt to prove that all AECs (even with the extra assumption of amalgamation) are eventually tame as a strategy to prove Shelah's Categoricity Conjecture will fail in ZFC. However, the AECs constructed in this paper are unstable and don't really fit into the picture of classification theory so far or the categoricity conjecture. This leaves open the possibility that eventual tameness can be proven in ZFC from *model-theoretic* assumptions, such as stability or categoricity. Partial work towards this goal has already been done by Shelah [14], which derives a variant of tameness from categoricity; see [2, Theorem 11.15] for an exposition. A related question of Grossberg asks if amalgamation can be derived from categoricity.

In Section 5, we prove that even without our cardinal arithmetic assumption we can derive large cardinal strength from tameness assertions. Roughly speaking we show that if κ carries the tameness property corresponding to weak compactness, then κ is weakly compact in L .

The reader is advised to have some background in both set theory and model theory. The set-theoretic background is in large cardinals for which we recommend Kanamori's book [8]. For the model-theoretic background, see a standard reference on AECs such as Baldwin's book [2]. We would like to thank John Baldwin, Andrew Brooke-Taylor, and the anonymous referee for helpful comments on this paper.

2. LARGE CARDINALS

We begin by recalling some relevant large cardinal definitions. These are slight tweaks on standard definitions in the spirit of \aleph_1 -strongly compact cardinals (see for example [1]). The basic framework is to take a large cardinal property that has κ being large if there is a κ -complete object of some type and parameterizing the completeness by some δ . Then κ is “almost large” if a δ -complete object exists for all $\delta < \kappa$ rather than at κ .

Definition 2.1. *Let κ be an uncountable cardinal.*

- (1) (a) κ is δ -weakly compact if for every field $\mathcal{A} \subset \mathcal{P}(\kappa)$ of size κ there is a nonprincipal δ -complete uniform filter measuring each set in \mathcal{A} .
- (b) κ is almost weakly compact if it is δ -weakly compact for all $\delta < \kappa$.
- (c) κ is weakly compact if it is κ -weakly compact.
- (2) (a) κ is δ -measurable if there is a uniform, δ -complete ultrafilter on κ .
- (b) κ is almost measurable if it is δ -measurable for all $\delta < \kappa$.
- (c) κ is measurable if it is κ -measurable.
- (3) (a) κ is (δ, λ) -strongly compact for $\delta \leq \kappa \leq \lambda$ if there is a δ -complete, fine ultrafilter on $\mathcal{P}_\kappa \lambda$.
- (b) κ is (δ, ∞) -strongly compact if it is (δ, λ) -strongly compact for all $\lambda \geq \kappa$.
- (c) κ is λ -strongly compact if it is (κ, λ) -strongly compact.
- (d) κ is almost strongly compact if it is (δ, ∞) -strongly compact for all $\delta < \kappa$.
- (e) κ is strongly compact if it is (κ, ∞) -strongly compact.

We note that the notions of δ -weakly compact and δ -measurable are not standard. From the definitions, it can be seen that being almost measurable implies being a limit of measurables. For almost weak and almost strong compactness, the relation is not so clear. For instance, the following seems open.

Question 2.2. *Is “there exists a proper class of almost strongly compact cardinals” equiconsistent with “there exists a proper class of strongly compact cardinals?”*

For the section, we fix an upward directed partial ordering $(\mathcal{D}, \triangleleft)$ with \triangleleft strict. The intended applications are (κ, \in) and $(\mathcal{P}_\kappa \lambda, \subset)$.

Definition 2.3. *For $d \in \mathcal{D}$ we define $\lceil d \rceil = \{d' \in \mathcal{D} \mid d' \triangleleft d\}$ and $\lfloor d \rfloor = \{d' \in \mathcal{D} \mid d \triangleleft d'\}$.*

We also fix a collection \mathcal{F} of functions each of which has domain \mathcal{D} . For $f_1, f_2 \in \mathcal{F}$ we set $f_1 \leq f_2$ if and only if there is an $e : \text{ran}(f_2) \rightarrow \text{ran}(f_1)$ such that $f_1 = e \circ f_2$. Note the witnessing e is unique. Obviously for each $f \in \mathcal{F}$, the set $\{f^{-1}\{i\} \mid i \in \text{ran}(f)\}$ partitions \mathcal{D} . So $f_1 \leq f_2$ is equivalent to saying that the partition from f_2 refines the partition from f_1 . We require that \mathcal{F} is upward directed under \leq .

We are interested in elements that appear cofinally often as values of f so we define

$$\text{ran}^*(f) = \bigcap_{d \in \mathcal{D}} \text{ran}(f \upharpoonright \lfloor d \rfloor).$$

With this notation in mind we formulate the following principle, which is implicit in [13].

Definition 2.4. Suppose \mathcal{F} is a directed family of functions with domain \mathcal{D} . Let $\#(\mathcal{D}, \mathcal{F})$ be the assertion that there are $f^* \in \mathcal{F}$ and a collection $\{u_f \subseteq \text{ran}^*(f) \mid f \in \mathcal{F} \wedge f \geq f^*\}$ of nonempty finite sets such that if e witnesses that $f \geq f^*$, $e \upharpoonright u_f : u_f \rightarrow u_{f^*}$ is a bijection.

Note that $e \upharpoonright u_f$ is unique, since there is a unique e witnessing $f \geq f^*$.

This principle allows us to define a filter on \mathcal{D} . Assume that $\#(\mathcal{D}, \mathcal{F})$ holds; this assumption is active until Corollary 2.10. Then we can choose $i_f \in u_f$ such that $e(i_f) = \min u_{f^*}$ where e witnesses $f^* \leq f$. Then we define $U \subseteq P(\mathcal{D})$ by $A \in U$ if and only if there are $d \in \mathcal{D}$ and $f \in \mathcal{F}$ with $f \geq f^*$ such that $f^{-1}\{i_f\} \cap [d] \subseteq A$. Note that U depends on the many parameters we have defined so far: $\mathcal{D}, \mathcal{F}, \{u_f\}$, and i_f . Also, the choice of i_{f^*} as the minimum of u_{f^*} was arbitrary, any element would have done. Indeed, different elements generate different ultrafilters.

Remark 2.5. The formulation of $\#(\mathcal{D}, \mathcal{F})$ given above is chosen because it is the easiest to work with in general. However, there is an alternate formulation in terms of the partitions of \mathcal{D} generated by the functions of \mathcal{F} that can make the definition of the filter more clear. In that language, $\#(\mathcal{D}, \mathcal{F})$ holds if and only if there is a special partition \mathcal{P}^* such that any finer partition \mathcal{P} has a distinguished piece $X_{\mathcal{P}}$ that is chosen in a coherent way: if \mathcal{Q} is finer than \mathcal{P} , then $X_{\mathcal{Q}} \subseteq X_{\mathcal{P}}$. We also require that for $d \in \mathcal{D}$ and \mathcal{P} finer than \mathcal{P}^* , we have $X_{\mathcal{P}} \cap [d] \neq \emptyset$.

Then we can define the filter as follows, given $A \subset \mathcal{D}$, we form a partition \mathcal{P}_A that is finer than both \mathcal{P}^* and $\{A, \mathcal{D} - A\}$. Then we set $A \in U$ if and only if the distinguished piece $X_{\mathcal{P}_A}$ is a subset of A rather than $\mathcal{D} - A$.

The choice of $i_f \in u_f$ corresponds to choices of different distinguished pieces, showing that there are $|u_{f^*}|$ -many filters with the desired property.

Claim 2.6. U is a proper filter and for all $d \in \mathcal{D}$, $[d] \in U$.

Proof. It is not hard to see that $f^{-1}\{i_f\} \cap [d]$ is nonempty for all d and f , so $\emptyset \notin U$ provided that it forms a filter. The fact that $[d] \in U$ for all d is immediate from the definition.

To see that U is a filter, let $A, B \in U$ witnessed by f_1, d_1 and f_2, d_2 respectively. Let $d_3 \in \mathcal{D}$ be above d_1 and d_2 and $f \geq f_1, f_2$. This is possible since both \mathcal{D} and \mathcal{F} are directed. It follows that $f^{-1}\{i_f\} \cap [d_3] \subseteq A \cap B$. \square

We would like to generate highly complete filters. To do so we use the following ad-hoc definition, which is essentially a closure property of the set of functions \mathcal{F} .

Definition 2.7. We say that \mathcal{F} is τ -replete if for every $\mu < \tau$ and sequence $\langle B_\epsilon \mid \epsilon < \mu \rangle$ of subsets of \mathcal{D} such that for each ϵ there is a function f_ϵ such that $B_\epsilon = f_\epsilon^{-1}\{i_\epsilon\}$ for some i_ϵ , there is a function $f \in \mathcal{F}$ and $\{i_\alpha : \alpha < \mu\} \subset \text{ran} f$ such that $f^{-1}\{i_0\} = \bigcap_{\epsilon < \mu} B_\epsilon$ and for all $\alpha < \mu$ and $d \in \mathcal{D}$, $f(d) = i_{\alpha+1}$ if and only if $d \notin \bigcap_{\epsilon < \mu} B_\epsilon$ and α is least such that $d \notin B_\alpha$.

Claim 2.8. If $(\mathcal{D}, \triangleleft)$ is τ -directed and (\mathcal{F}, \leq) is τ -replete, then U is τ -complete.

By τ -directed we mean that sets of size less than τ have an upperbound.

Proof. Let A_ϵ for $\epsilon < \mu$ be elements of U where $\mu < \tau$. By the definition of U , for each $\epsilon < \mu$ we have f_ϵ and d_ϵ so that $f_\epsilon^{-1}\{i_{f_\epsilon}\} \cap [d_\epsilon] \subseteq A_\epsilon$. Let $B_\epsilon = f_\epsilon^{-1}\{i_{f_\epsilon}\}$ for $\epsilon < \mu$ and use the τ -repleteness of \mathcal{F} to find f . Using the directedness of \mathcal{F} we can

find $\hat{f} \geq f, f^*$ (recall f^* is given by $\#(\mathcal{D}, \mathcal{F})$). Using the τ -directedness of \mathcal{D} , let d be above each d_ϵ for $\epsilon < \mu$.

Let e witness that $f \leq \hat{f}$, ie $f = e \circ \hat{f}$. We want to show that $e(i_{\hat{f}}) = 0$, since then $\hat{f}^{-1}\{i_{\hat{f}}\} \cap [d] \subseteq \bigcap_{\epsilon < \mu} (B_\epsilon \cap [d_\epsilon]) \subseteq \bigcap_{\epsilon < \mu} A_\epsilon$.

Suppose that $e(i_{\hat{f}}) = \epsilon$ is not zero. Then $\hat{f}^{-1}\{i_{\hat{f}}\} \subseteq \mathcal{D} - B_\epsilon$ by the definition of f . This contradicts that U is filter containing all the sets $[d]$ for $d \in \mathcal{D}$. \square

Claim 2.9. *If $A \subseteq \mathcal{D}$ and there is an f in \mathcal{F} such that $A = f^{-1}X$ for some $X \subseteq \text{ran}(f)$, then U measures A .*

Proof. Let f and X witness the hypotheses of the claim. Since \mathcal{F} is directed, we can find $\hat{f} \in \mathcal{F}$ such that $f, f^* \leq \hat{f}$. Let e be such that $f = e \circ \hat{f}$. Now it is not hard to see that if $e(i_{\hat{f}}) \in X$, then $\hat{f}^{-1}\{i_{\hat{f}}\} \subseteq A$ and if $e(i_{\hat{f}}) \notin X$, then $\hat{f}^{-1}\{i_{\hat{f}}\} \subseteq \mathcal{D} - A$. In the first instance we have $A \in U$ and in the second we have $\mathcal{D} - A \in U$. \square

If \mathcal{F} satisfies the hypothesis of the previous claim for A , then we say that \mathcal{F} has a characteristic function for A .

We can now reformulate many large cardinal notions that are witnessed by the existence of measures. Our first corollary is an equivalent formulation of weak compactness.

Corollary 2.10. *Let κ be a regular cardinal. κ is weakly compact if and only if for all fields \mathcal{A} of subsets of κ with $|\mathcal{A}| = \kappa$, $\#(\kappa, \mathcal{F})$ holds for some set of functions \mathcal{F} on κ which is directed, κ -replete and contains characteristic functions for all elements of \mathcal{A} .*

Proof. Assume that κ is weakly compact. Let \mathcal{A} be a field of subsets of κ with $|\mathcal{A}| = \kappa$. We can assume that \mathcal{A} is closed under intersections of size less than κ . Using the weak compactness of κ , we fix a κ -complete \mathcal{A} -ultrafilter U .

Let $\mathcal{F}_{\mathcal{A}}$ be the collection of functions $f : \kappa \rightarrow \kappa$ such that $\text{ran}(f) \subseteq \alpha < \kappa$ for some α and for all $\beta \in \text{ran}(f)$, $f^{-1}\{\beta\} \in \mathcal{A}$. It is not difficult to show that $\mathcal{F}_{\mathcal{A}}$ is \leq -directed and it is κ -replete since U is κ -complete.

For each $f \in \mathcal{F}_{\mathcal{A}}$ let $u_f = \{i_f\}$ where i_f is the unique element of $\text{ran} f$ such that $f^{-1}\{i_f\} \in U$. If we take f^* to be the constantly zero function, then it is straightforward to see that f^* and $\{u_f \mid f \in \mathcal{F}_{\mathcal{A}}\}$ satisfy $\#(\mathcal{D}, \mathcal{F}_{\mathcal{A}})$.

For the reverse direction, for each field \mathcal{A} we apply Claims 2.6, 2.8 and 2.9 to see that the filter U generated by $\#(\mathcal{D}, \mathcal{F})$ is a nonprincipal κ -complete \mathcal{A} -ultrafilter. \square

Remark 2.11. *A similar proof characterizes σ^+ -weak compactness where we just replace κ -replete with σ^+ -replete and consider functions with codomain σ .*

We also have characterizations of σ^+ -measurable, (δ, λ) -strongly compact and λ -strongly compact.

Corollary 2.12. *κ is σ^+ -measurable if and only if $\#(\kappa, \mathcal{F})$ holds for some set \mathcal{F} of functions from κ to σ such that \mathcal{F} is directed, σ^+ -replete and has characteristic functions for all subsets of κ .*

Corollary 2.13. *Let $\kappa \leq \lambda$ be cardinals. κ is λ -strongly compact if and only if $\#(\mathcal{P}_\kappa(\lambda), \mathcal{F})$ holds for some set \mathcal{F} of functions with domain $\mathcal{P}_\kappa(\lambda)$ and range bounded in κ , such that \mathcal{F} is directed, κ -replete and has characteristic functions for all subsets of $\mathcal{P}_\kappa(\lambda)$.*

Remark 2.14. *The previous corollary can be modified to give a natural characterization of “ κ is (δ, λ) -strongly compact”.*

The proofs of these corollaries are all similar to the proof of Corollary 2.10 and will be omitted.

3. MODEL CONSTRUCTIONS

In this section we describe a family of constructions of models which take \mathcal{D} and \mathcal{F} from the previous section as parameters. We will also use a countable closure hypothesis on \mathcal{F} , but we delay this specification until Lemma 3.9.

We define the languages and structures that are the key objects in this section and the next.

Definition 3.1. *Fix a \mathcal{D} and \mathcal{F} as in the previous section.*

- (1) Set $X := \cup\{\text{ran}f : f \in \mathcal{F}\}$, $\sigma := |X|$ and $(G, +) := ([X]^{<\omega}, \Delta)$.
- (2) \mathcal{L}_σ^- is the language with two sorts A and I ; functions $\pi : A \rightarrow I$ and $F_c : A \rightarrow A$; and relations $P, D_c \subset A$ and $E', E, R \subset A^2$ as c ranges over G .
- (3) \mathcal{L}_σ is the language \mathcal{L}_σ^- with an additional sort J and a function $Q : A \rightarrow J$.

The sorts here are disjoint. Δ is the symmetric difference on finite subsets of X . Note that $(G, +)$ is the free group of order 2 on $\sigma = |X|$ many generators, so the definitions above only depend on σ up to renaming. Given an \mathcal{L}_σ^- -structure H , we will often expand it trivially to a \mathcal{L}_σ -structure M by putting a single point in Q . Also, we allow structures with empty sorts for A and J .

For this section, we focus on \mathcal{L}_σ^- . For $\ell = 1, 2$, we will build $H_{\ell, \mathcal{D}}$ as the colimit of the $\subset_{\mathcal{L}_\sigma^-}$ -directed system $\langle H_{\ell, d} \mid d \in \mathcal{D} \rangle$.

We focus first on $H_{1, \mathcal{D}}$. For $d \in \mathcal{D}$, $H_{1, d}$ is the substructure of $H_{1, \mathcal{D}}$ with universe $A_d = \mathcal{F} \times [d] \times G$ and $I_d = \mathcal{F} \times [d]$. For each of the functions and relations below, we replace ‘ \mathcal{D} ’ with ‘ d ’ to denote the restriction to $H_{1, d}$, for example $\pi_d = \pi_{\mathcal{D}} \upharpoonright A_d$.

Definition 3.2. *$H_{1, \mathcal{D}}$ is the \mathcal{L}_σ^- -structure with universe $A_{\mathcal{D}} = \mathcal{F} \times \mathcal{D} \times G$ and $I_{\mathcal{D}} = \mathcal{F} \times \mathcal{D}$ with the following functions and relations:*

- $\pi_{\mathcal{D}}$ is the natural projection from $A_{\mathcal{D}}$ to $I_{\mathcal{D}}$, and $E'_{\mathcal{D}}$ is the derived equivalence relation;
- $E_{\mathcal{D}}$ refines $E'_{\mathcal{D}}$ and is given by $(f, d, u)E_{\mathcal{D}}(f', d', u')$ iff $(f, d) = (f', d')$ and there are $d_0, \dots, d_{2n-1} \in [d]$ such that $u\Delta\{f(d_0)\}\Delta\dots\Delta\{f(d_{2n-1})\} = u'$;
- $P_{\mathcal{D}}$ is the unary parity predicate and holds at (f, d, u) iff $|u \cap \text{ran}(f \upharpoonright [d])|$ is odd;
- For $v \in G$, $D_v^{\mathcal{D}}$ is a unary difference predicate and holds at (f, d, u) iff $u - \text{ran}(f \upharpoonright [d]) \subset v$ and $v \cap \text{ran}f \upharpoonright [d] = \emptyset$;
- $F_c^{\mathcal{D}}$ describe the transitive action of G on each $E'_{\mathcal{D}}$ -class given by $F_c^{\mathcal{D}}(f, d, u) = (f, d, u\Delta c)$; and
- $(f, d, u)R_{\mathcal{D}}(f', d', u')$ if and only if $f \leq f'$ and if e witnesses this, then $u = \{i \in X \mid \exists^{odd} j \in u'.e(j) = i\}$.

The use of E' is redundant given π , but makes the discussion of its equivalence classes easier. We also have that E is redundant.

Claim 3.3. *For all $d^* \in \mathcal{D}$ and $(f, d, u), (f', d', u') \in H_{1, d^*}$, we have $(f, d, u)E_{d^*}(f', d', u')$ if and only if the following hold:*

- $(f, d, u)E'_{d^*}(f', d', u')$;
- $P_{d^*}(f, d, u)$ if and only if $P_{d^*}(f', d', u')$; and
- for all $v \in G$, $D_v^{d^*}(f, d, u)$ if and only if $D_v^{d^*}(f', d', u')$.

In particular, within a particular E' -class, E -equivalence is determined by the quantifier-free type of the singletons from A .

Proof. Clearly $(f, d, u)E(f', d', u')$ if and only if

- $(f, d, u)E'(f', d', u')$ (hence $f = f'$ and $d = d'$);
- $u\Delta u' \subset \text{ran}(f)$ is even; and
- $u - \text{ran}(f \upharpoonright [d]) = u' - \text{ran}(f' \upharpoonright [d'])$

It is not hard to see that this is equivalent to the list from the claim. \square

Definition 3.4. For $d \in \mathcal{D}$, set g_d to be the permutation on $A_d \cup I_d$ of order two given by $g_d(f, d', u) = (f, d', u + \{f(d)\})$ on A_d and the identity on I_d .

Note that g_d is defined on (f, d', u) only if $d' \triangleleft d$. If $\neg(d' \triangleleft d)$, we could define $g_d(f, d', u)$, but it would not be a member of A_d . This is simply a bijection on the underlying set of $H_{1,d}$, but we can describe its interaction with the \mathcal{L}_σ^- structure as well.

Claim 3.5. Given $d_1, d_2 \in [d]$, $g_{d_1} \circ g_{d_2}$ is an \mathcal{L}_σ^- -automorphism of $H_{1,d}$.

Proof. Most of this is clear from the definition of $H_{1,d}$. For a permutation f of M , we say that some f preserves a predicate U iff U holds of x iff it holds of $f(x)$ (in M) and it flips U iff U holds of x iff it fails to hold at $f(x)$. It is easy to see that g_{d_ℓ} preserves each difference predicate and flips each parity predicate. We show that g_{d_1} already preserves $R_{\mathcal{D}}$ on any pair for which the function is defined. Suppose $(f, d, u)R_{\mathcal{D}}(f', d', u')$ and let $f = e \circ f'$. Since $u = \{i \in X \mid \exists^{\text{odd}} j \in u'.e(j) = i\}$, we have that

$$\begin{aligned} u\Delta\{f(d_1)\} &= \{i \in X \mid \exists^{\text{odd}} j \in u'.e(j) = i\} \Delta \{e \circ f'(d_1)\} \\ &= \{i \in X \mid \exists^{\text{odd}} j \in u'\Delta\{f'(d_1)\}.e(j) = i\}. \end{aligned}$$

The second equality holds because $u'\Delta\{f'(d_1)\}$ changes the number of preimages of $f(d_1)$ by 1 (when compared to u'). It follows that $(f, \alpha, u\Delta\{f(d_1)\})R_{\mathcal{D}}(f', \alpha', u'\Delta\{f'(d_1)\})$.

Finally, we note that E_d is preserved because the same elements of \mathcal{D} witnessing E_d -relatedness in $H_{1,d}$ will witness E_d relatedness of the $g_{d_1} \circ g_{d_2}$ -images. Since each predicate is flipped or preserved, the composition $g_{d_1} \circ g_{d_2}$ preserves each predicate and is an \mathcal{L}_σ^- -isomorphism. \square

Claim 3.6. If $\langle (f, d, u_1), (f, d, u_2) \rangle$ has the same quantifier-free type in H_{1,d^*} as $\langle (f, d, v_1), (f, d, v_2) \rangle$ for any $d \triangleleft d^*$, then $u_1\Delta u_2 = v_1\Delta v_2$.

Proof. Note that “ $F_{u_1\Delta u_2}(x) = y$ ” is in the quantifier free type of the first and $F_c^{d^*}(g, d, u) = (g, d, v)$ if and only if $c = u\Delta v$. \square

The $H_{2,d}$'s are built as the g_d -images of the $H_{1,d}$.

Definition 3.7. For $d \in \mathcal{D}$, set $H_{2,d}$ to be the \mathcal{L}_σ^- -structure with universe $A_d \cup I_d$ defined so that g_d is an \mathcal{L}_σ^- -isomorphism from $H_{1,d}$. Set $H_{2,\mathcal{D}}$ be the direct union of the sequence $\langle H_{2,d} \mid d \in \mathcal{D} \rangle$.

This definition is justified because g_d is a bijection from $A_d \cup I_d$ to itself and, by Claim 3.5, if $d \triangleleft d' \in \mathcal{D}$, then $H_{2,d} \subseteq H_{2,d'}$; here, \subseteq refers to the substructure relation. Note that $H_{2,\mathcal{D}}$ has the same universe as $H_{1,\mathcal{D}}$.

Tameness type assumptions about our class of models will give an \mathcal{L}_σ^- -isomorphism h from $H_{1,\mathcal{D}}$ to $H_{2,\mathcal{D}}$ with an additional property. We call this additional property “respecting π ”.

Definition 3.8. *We say that $h : H_{1,d} \rightarrow H_{2,d}$ respects π iff for all $a \in A_{\mathcal{D}}$, $\pi_{\mathcal{D}}(a) = \pi_{\mathcal{D}}(h(a))$.*

In other words, we require $h \upharpoonright I_{\mathcal{D}}$ to be the identity.

From such an isomorphism we will prove $\#(\mathcal{D}, \mathcal{F})$. Note that setting $h_0(f, d, u) = (f, d, u \Delta \{f(d)\})$ is a $\mathcal{L}_\sigma^- - \{R\}$ -isomorphism respecting π . This h_0 does not preserve R because different d give different $f(d)$. This could be remedied by picking a “generic” or “average” value to play the role of $f(d)$, and we could use an ultrafilter to find such a value. Since we can derive an ultrafilter from the existence of such an isomorphism, this argues that this average construction is essentially the only way to construct such an isomorphism.

We say that \mathcal{F} is countably closed if for any \leq -increasing sequence $\langle f_n \mid n < \omega \rangle$ from \mathcal{F} , there is $f \in \mathcal{F}$ such that $f \geq f_n$ for all $n < \omega$.

Lemma 3.9. *Suppose that \mathcal{F} is countably closed. If there is an \mathcal{L}_σ^- -isomorphism h from $H_{1,\mathcal{D}}$ to $H_{2,\mathcal{D}}$ respecting π , then $\#(\mathcal{D}, \mathcal{F})$ holds.*

For the remainder of the section, we assume that \mathcal{F} is countably closed and that there is an h as in the lemma to derive $\#$.

Claim 3.10. *The following are true of h :*

- (1) *If we let $u_{f,d}$ be the unique element of G such that $h(f, d, \emptyset) = (f, d, u_{f,d})$, then $u_{f,d}$ doesn't depend on d . Hence we denote the common value by u_f .*
- (2) *For all $d \in \mathcal{D}$, $f \in \mathcal{F}$ and $u \in G$, $h(f, d, u) = (f, d, u \Delta u_f)$.*
- (3) *If $f \leq f'$, then $|u_f| \leq |u_{f'}|$.*
- (4) *For all $f \in \mathcal{F}$, $u_f \neq \emptyset$.*
- (5) *For all $f \in \mathcal{F}$, $u_f \subseteq \text{ran}^* f$.*

Proof. For (1), applying h

$$H_{1,\mathcal{D}} \models (f, d, \emptyset)R(f, d', \emptyset) \rightarrow H_{2,\mathcal{D}} \models (f, d, u_{f,d})R(f, d', u_{f,d'}).$$

Recall from the proof of Claim 3.5 that g_d preserves R and note that id is the witness that $f \leq f$. Applying the definition of R , we have

$$u_{f,d} = \{i \in X \mid \exists^{\text{odd}} j \in u_{f,d'} \cdot j = i\} = u_{f,d'}$$

For (2) we apply Claim 3.6 to $\langle (f, d, \emptyset), (f, d, u) \rangle$ and $\langle (f, d, u_f), (f, d, v) \rangle$ where $h(f, d, u) = (f, d, v)$.

For (3), we let e be any function such that $f = e \circ f'$. Then $(f, d, \emptyset)R_{\mathcal{D}}(f', d', \emptyset)$ implies $H_{2,\mathcal{D}} \models “(f, d, u_f)R(f', d, u_{f'})”$ implies $(f, d, u_f)R_{\mathcal{D}}(f', d, u_{f'})$ because g_d preserves R . So $u_f \subseteq e “u_{f'}$ and $|u_f| \leq |u_{f'}|$.

For (4), note that $H_{2,d}$ interprets the parity predicate to mean “ $|u|$ is even” since g_d flips P and that $H_{2,\mathcal{D}} \models \neg P(f, d, u_f)$, since $H_{1,\mathcal{D}} \models \neg P(f, d, \emptyset)$. Thus $|u_f|$ is odd and can't be empty.

For (5), for all $d \in \mathcal{D}$, $H_{2,\mathcal{D}} \models D_\emptyset(f, d, u_f)$, since $H_{1,\mathcal{D}} \models D_\emptyset(f, d, \emptyset)$. Moreover g_d preserves this predicate. So $u_f \subseteq \text{ran}(f \upharpoonright [d])$. \square

We are ready to produce the f^* for $\#$. It is here that we use for the first (and only) time the countable closure of the space of functions \mathcal{F} under the order \leq .

Claim 3.11. *There is $f^* \in \mathcal{F}$ such that $|u_f| = |u_{f^*}|$ for all $f \geq f^*$ from \mathcal{F} . Moreover, if e witnesses $f^* \leq f$, then $e \upharpoonright u_f$ is a bijection from u_f to u_{f^*} .*

Proof. The moreover part follows from the first part because every member of u_{f^*} is in the image of u_f under an e witnessing $f^* \leq f$ by the proof of Claim 3.10 part (3).

Suppose there is no such f^* . Then there is a \leq -increasing sequence $\langle f_n \in \mathcal{F} \mid n < \omega \rangle$ such that $|u_{f_n}| < |u_{f_{n+1}}|$ for all $n < \omega$. By our assumption that \mathcal{F} is countably closed we can find $f^* \geq f_n$ for all $n < \omega$, but then $|u_{f^*}|$ is a natural number above infinitely many natural numbers, a contradiction. \square

So we have derived $\#(\mathcal{D}, \mathcal{F})$ as witnessed by f^* and the finite sets u_f . This finishes the construction of our sequence of models. We will need further work to show that these models can be thought of as elements of some AEC and that tameness assumptions about that AEC give the hypothesis of Lemma 3.9.

4. ABSTRACT ELEMENTARY CLASSES

The goal of this section is twofold. First, we put the algebraic constructions of Section 3 into the context of AECs. Second, we put the necessary pieces together to conclude large cardinal principles from global tameness and locality axioms.

The AEC is designed to precisely take in the algebraic examples constructed in Section 3 with a single twist. Recall \mathcal{L}_σ from Definition 3.1. The extra predicate Q is an index for copies of the algebraic construction, as in Baldwin and Shelah [3]. This allows us to turn the ‘‘incompactness’’ results about the existence of isomorphisms above into the desired nonlocality results for Galois types.

We define an AEC parameterized by σ with very minimal structure. In applications we require $\sigma^\omega = \sigma$. The strong substructure relation is as weak as possible, leaving open the question of whether restricting to the case of stronger strong substructure relations carries the same large cardinal implications.

Definition 4.1. *We define $\mathbb{K} = \mathbb{K}_\sigma$ to be the collection of \mathcal{L}_σ -structures (recall Definition 3.1) given by $M \in \mathbb{K}$ if and only if M is an \mathcal{L}_σ -structure satisfying:*

- (1) $\{F_c \mid c \in G\}$ is an 1-transitive action of G on $\pi^{-1}\{i\}$ for every $i \in I$.
- (2) E' and E are equivalence relations on A and $aE'b$ if and only if $\pi(a) = \pi(b)$.
- (3) For all $i \in I$ and $j \in J$, there is an $a \in A$ such that $\pi(a) = i$ and $Q(a) = j$.

We let $\prec_{\mathbb{K}}$ be the \mathcal{L}_σ -substructure relation.

Note that this is an AEC with $LS(\mathbb{K}) = \sigma$; in fact, \mathbb{K}_σ is the class of models of an $L_{\sigma^+, \omega}$ -sentence. The main difference between this definition and Shelah [13, \boxtimes_2 in Proof of Theorem 1.3] is that we have encoded the entire group G into the language rather than adding a separate sort for it (see Remark 4.6). We note that the structures we call $H_{\ell, d}$ and $M_{\ell, d}$ are called $M_{\ell, \alpha}$ and $M_{\ell, \alpha}^+$ respectively in [13].

As mentioned in the last section, any \mathcal{L}_σ^- -structure can be trivially expanded to an \mathcal{L}_σ -structure by putting a single point in J and fixing Q to be the constant function with this value on A . For $\ell = 1, 2$, let $M_{\ell, d}$ be this expansion of $H_{\ell, d}$ and name the single element of $J^{H_{\ell, d}}$ as i_ℓ . We similarly expand $H_{\ell, \mathcal{D}}$ to $M_{\ell, \mathcal{D}}$.

We also define $M_{0, d}$ to be the \mathcal{L}_σ -structure with A and J empty and $I = I_d$; note that $M_{0, d} \in \mathbb{K}_\sigma$. We similarly define $M_{0, \mathcal{D}}$.

Note that $M_{0,d} \subseteq M_{\ell,d}$ for $\ell = 1, 2$. Thus we can define $p_d := gtp(i_1/M_{0,d}; M_{1,d})$ and $q_d := gtp(i_2/M_{0,d}; M_{2,d})$; $p_{\mathcal{D}}$ and $q_{\mathcal{D}}$ are defined similarly. The connection between this AEC and the previous work is the following proposition.

Claim 4.2.

- (1) For all $d \in \mathcal{D}$, $p_d = q_d$.
- (2) There is an isomorphism as in Lemma 3.9 if and only if $p_{\mathcal{D}} = q_{\mathcal{D}}$.

To prove the forward direction of (2), we need the notion of *admitting intersection* coming from [3, Definition 1.2] in the AEC case.

Definition 4.3. \mathbb{K} admits intersections if and only if for all $X \subseteq M \in \mathbb{K}$, $cl_M(X) \prec_{\mathbb{K}} M$, where $cl_M(X)$ is the substructure of M with universe $\cap\{N : X \subseteq N \prec_{\mathbb{K}} M\}$.

The key consequence of closure under intersection is that it simplifies checking if two types are equal.

Fact 4.4 ([3].1.3). Suppose \mathbb{K} admits intersections. Then $gtp(a_1/M_0; M_1) = gtp(a_2/M_0; M_2)$ if and only if there is $h : cl_{M_1}(M_0 a_1) \cong_{M_0} cl_{M_2}(M_0 a_2)$ with $h(a_1) = a_2$.

Claim 4.5. \mathbb{K}_{σ} admits intersections.

Proof. We define the closure on each of the predicates. Then $cl_M(X)$ will be the substructure with the union of the $cl_M^i(X)$ as the universe.

- $cl_M^1(X) = (X \cap J) \cup \{j \in J^M : \exists a \in X \cap A^M. Q^M(a) = j\}$;
- $cl_M^2(X) = (X \cap I) \cup \{i \in I^M : \exists a \in X \cap A^M. \pi^M(a) = i\}$;
- $cl_M^3(X) = \{a \in A^M : \exists (i, j) \in cl_M^1(X) \times cl_M^2(X). Q^M(a) = j \text{ and } \pi^M(a) = i\}$.

It is routine to verify that cl_M satisfies Definition 4.3. □

Remark 4.6. The AEC as constructed in [13] is not closed under intersections. Shelah does not require that the entire group G be included in every model. This means that if there is an empty E' -equivalence class of A that must be filled (due to it projecting into I and J), then any proper subgroups $G' < G$ allows a choice of orbits to fill the equivalence class. This choice is incompatible with closure under intersection. However, an argument similar to Claim 4.7 still shows it has amalgamation.

Proof of Claim 4.2. We begin by showing (1) and the forward direction of (2). The type equality comes from the fact that if f is an \mathcal{L}_{σ}^- -isomorphism from $H_{1,d}$ to $H_{2,d}$ that respects π , then $f^* := f \cup \{(i_1, i_2)\}$ is a \mathcal{L}_{σ} -isomorphism from $M_{1,d}$ to $M_{2,d}$ that fixes $M_{0,d}$ and sends i_1 to i_2 . Since each g_d respects π , this witnesses $p_d = q_d$. The same argument gives the forward direction of (2).

For the other direction, suppose $p_{\mathcal{D}} = q_{\mathcal{D}}$. It is easy to compute that $cl_{M_{\ell,\mathcal{D}}}(M_{0,\mathcal{D}} i_{\ell}) = M_{\ell,\mathcal{D}}$. So by Fact 4.4 we have an isomorphism $h : M_{1,\mathcal{D}} \cong_{M_{0,\mathcal{D}}} M_{2,\mathcal{D}}$. This restricts to an isomorphism from $H_{1,\mathcal{D}}$ to $H_{2,\mathcal{D}}$ that respects π as in Lemma 3.9 and so $\#(\mathcal{D}, \mathcal{F})$ follows. □

Much of the work on AECs takes place under the assumption of amalgamation. Although not necessary for this proof, we also point out that \mathbb{K}_{σ} has amalgamation. This means the use of the construction [3, Definition 4.5] in [13] is unnecessary.

Claim 4.7. \mathbb{K}_{σ} has amalgamation.

Proof. Suppose that $M_0 \subset M_1, M_2 \in K$ and without loss of generality $M_1 \cap M_2 = M_0$. We will define the amalgam to essentially be the disjoint union of M_1 and M_2 over M_0 written in the standard way with A equal to $I \times G \times J$. Thus, we define the universe of M^* as follows:

- $J^* = J_1 \cup J_2$;
- $I^* = I_1 \cup I_2$; and
- $A^* = I^* \times G \times J^*$.

For each $(i, j) \in I_\ell \times J_\ell$ and $\ell = 1, 2$, pick some $x_{i,j}^\ell \in A^{M_\ell}$ such that

- $\pi^{M_\ell}(x_{i,j}^\ell) = i$;
- $Q^{M_\ell}(x_{i,j}^\ell) = j$; and
- if $(i, j) \in I_0 \times J_0$, then $x_{i,j}^1 = x_{i,j}^2$.

The $x_{i,j}^\ell$ serves as the “zero” to define the action of G on A^* . We define $f_\ell : M_\ell \rightarrow M^*$ as the identity on J_ℓ, I_ℓ , and G and, given $y \in A^{M_\ell}$,

$$f_\ell(y) = (i, g, j) \iff F_g^{M_\ell}(x_{i,j}^\ell) = y; \pi^{M_\ell}(y) = i; \text{ and } Q^{M_\ell}(y) = j$$

We put the \mathcal{L}_σ -structure on A^* only as required by M_1 and M_2 . For instance, E^* holds of (i, g, j) and (i', g', j') iff they are images of f_ℓ and their preimages are E^{M_ℓ} related. Then $M^* \in \mathbb{K}_\sigma$ and is the amalgam. \square

We are now able to put the pieces together and generate several equivalences between global tameness principles for AECs and large cardinal axioms.

Definition 4.8 ([2], Chapter 11). *Let \mathcal{K} be an AEC and $\kappa \leq \lambda$.*

- (1) \mathcal{K} is $(< \kappa, \lambda)$ -tame if for every $M \in \mathcal{K}_\lambda$ and $p \neq q \in gS(M)$, there is a $M_0 \prec_{\mathcal{K}} M$ of size $< \kappa$ such that $p \upharpoonright M_0 \neq q \upharpoonright M_0$.
- (2) \mathcal{K} is $< \kappa$ -tame if it is $(< \kappa, \mu)$ -tame for all $\mu \geq \kappa$.
- (3) \mathcal{K} is eventually tame if it is $< \kappa$ -tame for some $\kappa > \text{LS}(\mathcal{K})$.
- (4) \mathcal{K} is κ -local if for every $M \in \mathcal{K}$, $p \neq q \in gS(M)$, and resolution $\langle M_i \in \mathcal{K} \mid i < \kappa \rangle$ of M , there is $i_0 < \kappa$ such that $p \upharpoonright M_{i_0} \neq q \upharpoonright M_{i_0}$.

Theorem 4.9. *Let $\sigma < \kappa$ be infinite cardinals with $\sigma^\omega = \sigma$.*

- (1) If $\kappa^\sigma = \kappa$ and every AEC \mathbb{K} with $\text{LS}(\mathbb{K}) = \sigma$ is $(< \kappa, \kappa)$ -tame, then κ is σ^+ -weakly compact.
- (2) If every AEC \mathbb{K} with $\text{LS}(\mathbb{K}) = \sigma$ is κ -local, then κ is σ^+ -measurable.
- (3) If every AEC \mathbb{K} with $\text{LS}(\mathbb{K}) = \sigma$ is $(< \kappa, \sigma^{(\lambda^{< \kappa})})$ -tame, then κ is (σ^+, λ) -strongly compact.

So we have the following corollary.

Corollary 4.10. *Let κ be an infinite cardinal such that $\mu^\omega < \kappa$ for all $\mu < \kappa$.*

- (1) if $\kappa^{< \kappa} = \kappa$ and every AEC \mathbb{K} with $\text{LS}(\mathbb{K}) < \kappa$ is $(< \kappa, \kappa)$ -tame, then κ is almost weakly compact.
- (2) if every AEC \mathbb{K} with $\text{LS}(\mathbb{K}) < \kappa$ is $< \kappa$ -tame, then κ is almost strongly compact.

Proof of Theorem 4.9. We start with the proof of part (1). We wish to apply our AEC construction together with Remark 2.11. Construct \mathbb{K}_σ as in Definition 4.1; by assumption, this is $(< \kappa, \kappa)$ -tame. Let \mathcal{A} be a field of subsets of κ with $|\mathcal{A}| = \kappa$. Let θ be a big regular cardinal and take $X \prec H_\theta$ of size κ with $\mathcal{A} \in X$. By our cardinal arithmetic assumption we can take X to be closed under σ -sequences. Let

\mathcal{F} be the collection of functions in X with domain κ and range a bounded subset of κ and $(\mathcal{D}, \triangleleft) = (\kappa, \in)$. A straightforward argument using the fact that X is closed under σ -sequences shows that \mathcal{F} is σ^+ -replete. It is also clear that \mathcal{F} has a characteristic function for each $A \in \mathcal{A}$. Build the Galois types $\{p_\alpha, q_\alpha \mid \alpha < \kappa\}$ corresponding to this system. By the $(< \kappa, \kappa)$ -tameness of \mathbb{K}_σ , $p_\kappa = q_\kappa$. By Claim 4.2 and Lemma 3.9, we have that $\#(\mathcal{D}, \mathcal{F})$ holds. Note that Lemma 3.9 requires that \mathcal{F} is countably closed and this follows from the fact that $\sigma^\omega = \sigma$. So by Remark 2.11 κ is σ^+ -weakly compact.

Part (2) is essentially Shelah's theorem from [13], but with the required corrections. We let $\mathcal{F} = {}^\kappa\sigma$ and $(\mathcal{D}, \triangleleft) = (\kappa, \in)$. By our locality assumption, Claim 4.2 and Lemma 3.9, we have $\#(\mathcal{D}, \mathcal{F})$. Hence by Corollary 2.12, we have that κ is σ^+ -measurable.

For part (3) we let $(\mathcal{D}, \triangleleft) = (\mathcal{P}_\kappa(\lambda), \subset)$ and $\mathcal{F} = {}^{\mathcal{D}}\sigma$. By our tameness assumption, Claim 4.2 and Lemma 3.9, we have $\#(\mathcal{D}, \mathcal{F})$. Hence by Remark 2.14, κ is (σ^+, λ) -strongly compact. \square

By strengthening the hypotheses a little we can remove the 'almost' from the above theorem. To do so we need a definition that generalizes [5, Definition 2.10].

Definition 4.11. $(\mathbb{K}, \prec_{\mathbb{K}})$ is quasi-essentially below κ if and only if $\text{LS}(\mathbb{K}) < \kappa$ or there is a theory T in $L_{\kappa, \omega}$ such that $\mathbb{K} = \text{Mod } T$ and $\prec_{\mathbb{K}}$ is implied by $\prec_{L_{\kappa, \omega}}$.

We have introduced quasi-essentially below instead of just essentially below from [5], because although the class of models in \mathbb{K}_σ are axiomatizable in $L_{\sigma^+, \omega}$, the strong substructure relation is even weaker than first-order elementary.

Theorem 4.12. Let κ be an infinite cardinal with $\kappa^{<\kappa} = \kappa$ and for every $\mu < \kappa$, $\mu^\omega < \kappa$. If every AEC \mathbb{K} which is quasi-essentially below κ is $(< \kappa, \kappa)$ -tame, then κ is weakly compact.

Proof. The proof follows the proof of Theorem 4.9.(1), and we point out the differences. The additional cardinal arithmetic implies that $X \prec H_\theta$ can be taken to closed under $< \kappa$ -sequences. Then set \mathcal{F} to be the collection of functions in X with domain κ and range bounded in κ . Note that \mathcal{F} is countably closed since κ has uncountable cofinality and ${}^\kappa\sigma$ is countably closed provided that $\sigma^\omega = \sigma$. As before, we code this into \mathbb{K}_κ to conclude that $\#(\mathcal{D}, \mathcal{F})$ holds. Moreover \mathcal{F} is κ -replete by the closure of X , hence by Corollary 2.10, κ is weakly compact. \square

Theorem 4.13. Let κ be a cardinal with $\text{cf}(\kappa) > \omega$ and for all $\mu < \kappa$, $\mu^\omega < \kappa$. If every AEC \mathbb{K} that is quasi-essentially below κ is $< \kappa$ -tame, then κ is strongly compact.

This theorem has some level by level information. In particular $(< \kappa, \sup_{\alpha < \kappa} \alpha^{(\lambda^{<\kappa})})$ -tameness will give that κ is λ -strongly compact.

Proof. The proof is similar to the other proofs above. We take $(\mathcal{D}, \triangleleft) = (\mathcal{P}_\kappa(\lambda), \subset)$ and \mathcal{F} to be the set of functions with domain \mathcal{D} and range bounded in κ . A similar argument to the one in the previous theorem shows that \mathcal{F} is countably complete. Our tameness assumption gives $\#(\mathcal{D}, \mathcal{F})$ and hence κ is λ -strongly compact by Corollary 2.13. \square

It is important to note that the converses of the main theorems (and corollaries) from this section are true. The bulk of the work is already done in the first

author's paper [5]. In some cases the converses of the stated results are stronger than theorems appearing in the literature. With some small adjustment the proofs in the literature already give these stronger claims. We make a brief list of the improvements required.

- (1) Loś' Theorem for AECs [5, Theorem 4.3] holds for AECs with Löwenheim-Skolem number σ and any σ^+ -complete ultrafilter. This allows us to prove for example that every AEC \mathbb{K} with $\text{LS}(\mathbb{K}) < \kappa$ is $< \kappa$ -tame from κ is almost strongly compact.
- (2) Loś' Theorem for AECs does not require an ultrafilter only that the filter measures enough sets. This allows us to prove every AEC \mathbb{K} with $\text{LS}(\mathbb{K}) < \kappa$ is $(< \kappa, \kappa)$ -tame from the assumption that κ is almost weakly compact. In particular we can build everything into a transitive model of set theory of size κ and the weak compactness assumption gives a filter measuring all subsets of κ in the model. This is enough to complete the proof.
- (3) Loś' Theorem for AECs applies to the class of AECs which are quasi-essentially below κ . This allows us to prove that every AEC \mathbb{K} which is quasi-essentially below κ is $< \kappa$ -tame from κ is strongly compact. It could be that such AECs have no models of size less than κ , but existing arguments are enough to give tameness over sets¹ rather than models.

We collect a few remarks on our construction:

- (1) We do not know if the cardinal arithmetic assumptions are necessary in the main theorems of this section. For example, if we assume that κ is weakly compact and we force to add κ^+ many subsets to some σ^+ where $\sigma < \kappa$, then κ remains σ^+ -weakly compact in the extension. It follows that every AEC with Löwenheim-Skolem number σ is $(< \kappa, \kappa)$ -tame in the extension. We do not know if κ satisfies any stronger tameness properties in the extension.
- (2) Under our mild cardinal arithmetic assumptions, the global full tameness and type shortness and compactness results from [5] follow from the global tameness for 1-types, as this tameness is already enough to imply the necessary large cardinals.

We conclude this section with an application to category theory. There has been recent activity in exploring the connection between AECs and accessible categories; see Lieberman [10], Beke and Rosicky [4], and Lieberman and Rosicky [9]. In [9, Theorem 5.2], the authors apply a result of Makkai and Pare to derive a global version of Fact 1.1 from class many strong compacts. Here we show that this application is in fact equivalent to the whole result. Note that in the global version we do not need any cardinal arithmetic assumptions.

Corollary 4.14. *The following are equivalent:*

- (1) *The powerful image of any accessible functor is accessible.*
- (2) *Every AEC is eventually tame.*
- (3) *There are class many almost strongly compact cardinals.*

¹Typically, Galois types are defined so that the domains are always models. The same definition works for defining Galois types over arbitrary sets. However, many model-theoretic arguments ([14, Claim 3.3] on local character of non-splitting from stability is an early example) only work for Galois types over models, explaining their prevalence. The set-theoretic nature of the arguments from large cardinals, on the other hand, mean that they carry through with little change.

Fix and infinite cardinal κ with $\mu^\omega < \kappa$ for all $\mu < \kappa$. The following are equivalent:

- (1) The powerful image of a $< \kappa$ -accessible functor is $< \kappa$ -accessible.
- (2) Every AEC with $LS(K) < \kappa$ is $< \kappa$ -tame.
- (3) κ is almost strongly compact.

In saying that every AEC is eventually tame, we allow AECs with no models of size κ or larger to be trivially tame by saying they are $< \kappa$ -tame. For the category theoretic notions in this corollary, see [11]. In particular, given a functor $F : \mathcal{K} \rightarrow \mathcal{L}$, the powerful image of F is the subcategory of \mathcal{L} whose objects are Fx for $x \in \mathcal{K}$ and whose arrows are any arrow from \mathcal{L} between these objects.

Proof. In the first set of equivalences, the first implies the second by [9, Theorem 5.2]. The third implies the first by Brooke-Taylor and Rosicky [6, Corollary 3.5], which is a modification of Makkai and Pare's original [11, 5.5.1].

To see the second implies the third, for all σ , we know that there is some κ_σ such that \mathbb{K}_σ is $< \kappa_\sigma$ -tame. Let \mathbf{S} be the set of all limit points of the map that takes σ to $\sigma^\omega + \kappa_\sigma$. Clearly, \mathbf{S} is class sized.

We claim that each $\kappa \in \mathbf{S}$ is almost strongly compact. First note that $\sigma^\omega < \kappa$ for all $\sigma < \kappa$. Let $\sigma < \kappa \leq \lambda$. Then $\mathbb{K}_{\sigma^\omega}$ is $< \kappa$ -tame, so it's $(< \kappa, (\sigma^\omega)^{\lambda^{< \kappa}})$ -tame. The proof of Theorem 4.9.(3) only involves this AEC, so it implies that κ is $((\sigma^\omega)^+, \lambda)$ -strongly compact. Of course, this means that it is (σ, λ) -strongly compact. Since σ and λ were arbitrary, κ is almost strongly compact, as desired.

The second set of equivalences is just the parameterized version of the first one, and follows by the parameterized versions of the relevant results. The cardinal arithmetic is only needed for (2) implies (3). \square

5. THE CONSISTENCY STRENGTH OF $(< \kappa, \kappa)$ -TAMENESS

We have already remarked that we do not know if the cardinal arithmetic assumptions are necessary in the theorems of Section 4. In this section, we show that in the absence of cardinal arithmetic assumptions, the degree of tameness we associate to weak compactness has the expected *consistency strength*.

Theorem 5.1. *Let κ be a regular cardinal greater than \aleph_1 . If every AEC \mathbb{K} which is quasi-essentially below κ is $(< \kappa, \kappa)$ -tame, then κ is weakly compact in L .*

Proof. We may assume that $0^\#$ does not exist, since otherwise every uncountable cardinal is weakly compact in L (see [8, Theorems 9.17.(b) and 9.14.(b)]). Let \mathcal{A} in L be a collection of κ many subsets of κ which is closed under complements and intersections of size less than κ . Choose an ordinal $\beta < (\kappa^+)^L$ with $\mathcal{A} \in L_\beta$ and such that L models $[L_\beta \cap \mathcal{P}(\kappa)]^{< \kappa} \subseteq L_\beta$. Let \mathcal{F} be the collection of functions in L_β from κ to κ whose ranges are bounded in κ .

We claim that \mathcal{F} is countably closed in V under the ordering on functions defined in Section 2. Suppose that X is a countable subset of \mathcal{F} . By the covering lemma and our assumption that $0^\#$ doesn't exist, there is a set $Y \in L$ of size \aleph_1 with $X \subseteq Y$. By the choice of L_β , $Y \in L_\beta$ and hence it has an upperbound in L_β .

By our tameness assumption, Claim 4.2 and Lemma 3.9, we have $\#(\kappa, \mathcal{F})$ from which we can derive a filter U on κ . From the way we chose \mathcal{F} , U measures all subsets of κ in L_β and is κ -complete with respect to sequences in L_β . We are now ready to give a standard argument that U restricted to \mathcal{A} is in L .

Let $j : L_\beta \rightarrow L_\gamma \simeq \text{Ult}(L_\beta, U)$ be the elementary embedding derived from the ultrapower by U . Standard arguments show that the critical point of j is κ . Let $\langle A_\alpha \mid \alpha < \kappa \rangle$ be an enumeration of \mathcal{A} in L_β . The sequence $\langle j(A_\alpha) \mid \alpha < \kappa \rangle$ is in L , since it is just $j(\langle A_\alpha \mid \alpha < \kappa \rangle) \upharpoonright \kappa$. So the set $\bar{U} = \{A_\alpha \mid \kappa \in j(A_\alpha)\}$ is in L . It is easy to see that $\bar{U} \subseteq U$ and hence is the κ -complete \mathcal{A} -ultrafilter that we require. \square

We expect that a similar result can be proved for locality and almost measurability with an appropriate inner model in place of L . Of course, extending this result to the almost strongly compact case would require major advances in inner model theory.

REFERENCES

1. Joan Bagaria and Menachem Magidor, *Group radicals and strongly compact cardinals*, Trans. Amer. Math. Soc. **366** (2014), no. 4, 1857–1877. MR 3152715
2. John T. Baldwin, *Categoricity*, University Lecture Series, vol. 50, American Mathematical Society, Providence, RI, 2009.
3. John T Baldwin and Saharon Shelah, *Examples of non-locality*, The Journal of Symbolic Logic **73** (2008), no. 03, 765–782.
4. T. Beke and J. Rosick, *Abstract elementary classes and accessible categories*, Annals of Pure and Applied Logic **163** (2012), no. 12, 2008 – 2017.
5. Will Boney, *Tameness from large cardinal axioms*, J. Symb. Log. **79** (2014), no. 4, 1092–1119.
6. Andrew Brooke-Taylor and Jiri Rosický, *Accessible images revisited*, preprint.
7. Rami Grossberg and Monica VanDieren, *Categoricity from one successor cardinal in tame abstract elementary classes*, J. Math. Log. **6** (2006), no. 2, 181–201. MR 2317426 (2008k:03070)
8. Akihiro Kanamori, *The higher infinite*, second ed., Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003, Large cardinals in set theory from their beginnings.
9. Michael Lieberman and Jiri Rosický, *Classification theory for accessible categories*, preprint.
10. Michael J. Lieberman, *Category-theoretic aspects of abstract elementary classes*, Annals of Pure and Applied Logic **162** (2011), no. 11, 903 – 915.
11. Michael Makkai and Robert Paré, *Accessible categories: the foundations of categorical model theory*, Contemporary Mathematics, vol. 104, American Mathematical Society, Providence, RI, 1989.
12. Michael Morley, *Categoricity in power*, Trans. Amer. Math. Soc **114** (1965), no. 2, 514–538.
13. Saharon Shelah, *Maximal failures of sequence locality in aec*, Preprint.
14. ———, *Categoricity for abstract classes with amalgamation*, Ann. Pure Appl. Logic **98** (1999), no. 1-3, 261–294. MR 1696853 (2000m:03087)

E-mail address: `wboney@math.harvard.edu`

MATHEMATICS DEPARTMENT, HARVARD UNIVERSITY, CAMBRIDGE, MA, USA

E-mail address: `sunger@math.ucla.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA-LOS ANGELES, LOS ANGELES, CA, USA