MODEL THEORETIC CHARACTERIZATIONS OF LARGE CARDINALS REVISITED

WILL BONEY, STAMATIS DIMOPOULOS, VICTORIA GITMAN, AND MENACHEM MAGIDOR

Abstract. In [Bon20], model theoretic characterizations of several established large cardinal notions were given. We continue this work, by establishing such characterizations for Woodin cardinals (and variants), various virtual large cardinals, and subtle cardinals.

1. Introduction

The compactness of strong logics and set theory have been intertwined since Tarski [Tar62] defined (weakly and strongly) compact cardinals in terms of the properties of the infinitary logic $\mathbb{L}_{\kappa,\omega}$. This established a strong connection between abstract model theory and the theory of large cardinals, which has also become apparent by the recent breakthroughs in the theory of Abstract Elementary Classes–a purely model theoretic framework–where certain important results depend on the existence of large cardinal axioms (e.g., [Bon14, BU17, SV]).

The interaction between the two fields is also strengthened by the first author's article [Bon20], which establishes new characterizations of established large cardinal notions, expressed in model theoretic terms as compactness properties. This paper is a sequel to [Bon20] and characterizes more large cardinals this way, namely Woodin, various virtual large cardinals and subtle cardinals. Notably, this has led us to generalise or define new concepts in abstract model theory, that may be useful outside the scope of the current exposition.

One of the main philosophical open questions about the large cardinal hierarchy is to explain the fact that it appears to be linear. Hence, apart from the intrinsic interest, we believe that the framework of compactness principles that we invoke offers a new insight into this problem.

The structure of the paper is as follows. In Section 2 we fix our notation and terminology and recall definitions and known results from abstract model theory and large cardinals. In Section 3 we give a model-theoretic characterisation of Woodin cardinals by introducing a notion of Henkin models for arbitrary abstract logics. In Section 4 we characterise various virtual large cardinals, by introducing the notion of a pseudo-model for a theory. Finally, in Section 5 we characterise (a class version of) subtle cardinals as a natural weakening of Vopěnka's principle by showing that if Ord is subtle, then every abstract logic has a stationary class of weak compactness cardinals.

The conversations leading to this paper began at the conference "Accessible categories and their connections" organized by Andrew Brooke-Taylor, and we would like to thank Brooke-Taylor for organizing a fascinating meeting.

2. Logics and large cardinals

2.1. **Abstract logics.** We begin by fixing a notion of abstract logic. The following is a sublist of standard properties, e.g., that appear in [CK12, Definition 2.5.1], and we have omitted the clauses that do not factor into our analysis (e.g., the closure and quantifier properties). From the definition of a language, it might appear that we have restricted ourselves to single-sorted,

Date: October 9, 2020.

first-order structures. However, many-sorts, higher-order relations, etc. can be coded into this framework.

Definition 2.1.

(1) A *language* τ is a collection of function and relation symbols that come with a finite number as an arity, as well as constant symbols. Formally, this means τ is an ordered quadruple $(\mathfrak{F}, \mathfrak{R}, \mathfrak{C}, n)$ where $\mathfrak{F}$, $\mathfrak{R}$, and $\mathfrak{C}$ are disjoint sets and $n : \mathfrak{F} \cup \mathfrak{R} \to \omega$ is the arity function.

(2) Given a language τ, Str τ is the collection of all τ-structures M, which consist of $\langle |M|, F^M, R^M, c^M\rangle_{F\in\mathfrak{F}, R\in\mathfrak{R}, c\in\mathfrak{C}}$, where $|M|$ is a set (called the *universe* or *underlying set* of M); $F^M : |M|^{n(F)} \to |M|$, $R^M \subseteq |M|^{n(F)}$, and $c^M \in |M|$. We often do not notationally distinguish between M and $|M|$.

(3) A *morphism* f between two languages $\tau = (\mathfrak{F}, \mathfrak{R}, \mathfrak{C}, n)$ and $\rho = (\mathfrak{F}', \mathfrak{R}', \mathfrak{C}', n')$ is an injective function $f : \mathfrak{F} \cup \mathfrak{R} \cup \mathfrak{C} \to \mathfrak{F}' \cup \mathfrak{R}' \cup \mathfrak{C}'$ that preserves the partition and maintains the arity. A *renaming* is a bijective morphism. Note a renaming $f : \tau \to \rho$ induces a bijection f^* from Str τ to Str ρ that fixes the underlying sets.

(4) A (abstract) logic is a pair $(\mathcal{L}, \vDash_{\mathcal{L}})$ satisfying the following conditions.
   (a) $\mathcal{L}$ is a (class) map from languages and we call $\mathcal{L}(\tau)$ the set of τ-sentences.
   (b) $\vDash_{\mathcal{L}} \subseteq \bigcup_{\text{languages } \tau} \text{Str } \tau \times \mathcal{L}(\tau)$ is the satisfaction relation.
   (c) (monotonicity) If $\tau \subseteq \rho$, then $\mathcal{L}(\tau) \subseteq \mathcal{L}(\rho)$.
   (d) (expansion) If $\phi \in \mathcal{L}(\tau)$, $\rho \supseteq \tau$, and M is a ρ-structure, then $M \vDash_{\mathcal{L}} \phi$ if and only if the reduct of M to τ, $M \upharpoonright \tau \vDash_{\mathcal{L}} \phi$.
   (e) (isomorphism) If $M \cong N$, then $M \vDash_{\mathcal{L}} \phi$ if and only if $N \vDash_{\mathcal{L}} \phi$.
   (f) (renaming) Every renaming $f : \tau \to \rho$ induces a unique bijection $f_* : \mathcal{L}(\tau) \to \mathcal{L}(\rho)$ such that, for any τ-structure M and $\phi \in \mathcal{L}(\tau)$, we have

$$M \vDash_{\mathcal{L}} \phi \text{ if and only if } f^*(M) \vDash_{\mathcal{L}} f_*(\phi).$$

   We often refer to a logic as just $\mathcal{L}$ and drop the subscript from satisfaction–simply writing $\vDash$–when the context makes it clear.

(5) The *occurrence number* of a logic $\mathcal{L}$–written $o(\mathcal{L})$–is the minimal cardinal κ such that, for every $\phi \in \mathcal{L}(\tau)$, there is $\tau_0 \in \mathcal{P}_\kappa \tau$ such that $\phi \in \mathcal{L}(\tau_0)$ ($\mathcal{P}_\kappa A$ here is the collection of all subsets of A of size less than κ).

Note that abstract logics are defined only for sentences, there is no incorporation of free variables, although these can be tacitly handled by adding and properly interpreting constants. So using this work around, we can, in fact, assume that free variables are available. Also, note that we are requiring abstract logics to have an occurrence number.

It will be useful for later to note that the unique bijections f_* associated to renamings f respect both composition and restriction. More precisely if $f : \tau \to \sigma$ and $g : \sigma \to \rho$ are renamings, then clearly $g \circ f$ is a renaming and, by uniqueness, it must be case that $g_* \circ f_* = (g \circ f)_*$. Also, if $\sigma \subseteq \tau$ are languages and $f : \tau \to \rho$ is a renaming, then clearly $f \upharpoonright \sigma : \sigma \to f\,"\,\sigma$ is a renaming and, again by uniqueness, $(f \upharpoonright \sigma)_* = f_* \upharpoonright \sigma$.

By definition, all our languages τ are set-sized and there can be only set-many sentences $\mathcal{L}(\tau)$ in a fixed language. The intuition behind most of the properties of an abstract logic is clear. The occurrence number captures our intuition that there should be a bound on the number of elements of a language that a single assertion can reference. For instance, first-order logic $\mathbb{L}_{\omega,\omega}$ has occurrence number ω because no single assertion can mention more than finitely much of the language and infinitary logics $\mathbb{L}_{\kappa,\omega}$ (see below for definition) have occurrence number κ.

We will often consider unions of logics. If $\mathcal{L}_0$ and $\mathcal{L}_1$ are logics, then $\mathcal{L}_0 \cup \mathcal{L}_1$ is the natural union of them, with sentences identified if they are satisfied by the same models.[1]

An $\mathcal{L}$-theory is $<\kappa$-satisfiable when every $<\kappa$-sized subset of it has a model. A cardinal κ is a *strong compactness cardinal* of an abstract logic $\mathcal{L}$ if and only if every $<\kappa$-satisfiable $\mathcal{L}$-theory is satisfiable. A cardinal κ is a *weak compactness cardinal* of an abstract logic $\mathcal{L}$ if and only if every $<\kappa$-satisfiable $\mathcal{L}$-theory of size κ is satisfiable. For example, ω is the strong compactness cardinal of first-order logic, a weakly compact cardinal κ is a weak compactness cardinal of the infinitary logic $\mathbb{L}_{\kappa,\kappa}$, and a strongly compact cardinal κ is a strong compactness cardinal of $\mathbb{L}_{\kappa,\kappa}$. Makowski showed that every abstract logic has a strong compactness cardinal if and only if Vopěnka's principle holds [Mak85, Theorem 2] (see Section 4 for more details).

Next, we will give an overview of some specific logics that come up in the article and their key properties. To distinguish abstract logics from a specific logic, we use $\mathcal{L}$ to denote abstract logics and $\mathbb{L}$ (with some decoration) to denote specific logics.

The logic $\mathbb{L}(Q^{WF})$ is first-order logic augmented by the quantifier Q^{WF} that takes in two variables so that $Q^{WF}xy\varphi(x,y)$ is true if and only if $\varphi(x,y)$ defines a well-founded relation: there is no sequence $\langle x_n \mid n<\omega\rangle$ such that $\varphi(x_{n+1},x_n)$ holds for all $n<\omega$.

The logic $\mathbb{L}_{\kappa,\mu}$ extends first-order logic by closing the rules of formula formation under conjunctions (and disjunctions) of $<\kappa$-many formulas that are jointly in $<\mu$-many free variables; and under existential (and universal) quantification of $<\mu$-many variables. $\mathbb{L}_{\omega,\omega}$ is just first-order logic, and if κ or μ are uncountable, we refer to $\mathbb{L}_{\kappa,\mu}$ as an infinitary logic. Note that $\mathbb{L}(Q^{WF}) \subseteq \mathbb{L}_{\omega_1,\omega_1}$ since the quantifier Q^{WF} is expressible in this logic. As we already alluded to above, compactness and other properties of infinitary logics are connected to the existence of large cardinals.

Second-order logic $\mathbb{L}^2$ extends first-order logic by allowing quantification over all relations on the universe (from the overarching universe V of sets). In structures where coding is available, such as arithmetic or set theory, this reduces to quantification over all subsets of the universe. We will often make use of the fact that in $\mathbb{L}^2(\{\in\})$, there is an assertion, known as Magidor's Φ, encoding the well-foundedness of $\in$ and that the model is isomorphic to some V_β; Magidor's Φ is used in proofs in [Mag71] and is explicitly discussed at [Bon20, Fact 2.1]. We can also extend the second-order logic $\mathbb{L}^2$ by allowing infintary conjunctions and quantification, resulting in logics $\mathbb{L}^2_{\kappa,\mu}$.

A particularly powerful logic extending second-order logic is *sort logic* $\mathbb{L}^s$, which was introduced by Väänänen (see [Vää14] for a precise definition and properties). It is a multi-sorted logic, where the predicate quantifiers are allowed to range over all relations on the universe of that sort, and which allows additional sort quantifiers $\exists^\sim$ and $\forall^\sim$ which range over all sets in V (not just relations on a given sort) searching for additional universes with desired predicates on them. A sort quantifier $\exists^\sim X$ answers the question about whether the model can be expanded to include universes with finitely many new sorts satisfying the predicate X. For example, in sort logic we can ask inside a group whether there is a field of which it is the multiplicative group by asking for a new sort with a predicate $+$ on it whose elements are the same as of the original sort of the multiplicative group and which interact according to field rules. Since sort logic extends second-order logic, we can use Magidor's sentence Φ to pick out the structures $(V_\alpha,\in)$, and using the sort quantifiers we can even express whether V_α is Σ_n-elementary in V.

Proposition 2.2. *In sort logic $\mathbb{L}^s$, we can express that a universe of a given sort is (isomorphic to) a V_α and that $V_\alpha \prec_{\Sigma_n} V$. This assertion has complexity Σ_n for the sort quantifiers.*

[1] WB: Is there a place this identification helps us? It seems unnatural abstractly; for instance, we can have distinct sentences in a single logic that are modeled by the same class of models. But I get that we don't really want/need two copies of $\mathbb{L}_{\omega,\omega}$ in $\mathbb{L}_{\kappa,\omega} \cup \mathbb{L}(Q^{WF})$. Maybe just replace it with a note that we don't double up in the obvious cases?

Proof. We use Magidor's Φ to express that the universe is V_α. Next, we argue, by induction on complexity, that for every sentence $\phi(x)$ in the language of set theory, there is a corresponding sentence ϕ^* in the language of sort logic such that V_α reflects V with respect to $\phi(x)$ if and only if V_α satisfies ϕ^* in sort logic. For the base case, observe that V_α already reflects V with respect to all Δ_0-assertions. For a Σ_1-assertion $\phi(x) := \exists y\, \psi(y,x)$ with $\psi(y,x)$ being Δ_0, we let ϕ^* informally be the formula $\forall a \in V_\alpha\, V_\alpha \models \phi(a) \leftrightarrow \exists V_\delta\, (\exists y \in V_\delta\, V_\delta \models \psi(y,a))$. We say here "informally" because in actuality we would have to say that there exists a predicate satisfying Magidor's Φ and an embedding of V_α into the universe of this predicate so that we have $\psi(y,a')$ holds where a' is where a is mapped by the isomorphism, etc. For Π_1-formulas, we replace the $\exists V_\delta\, \exists y$ quantifiers by $\forall V_\delta\, \forall y$ and for formulas of complexity $n+1$, we use the translation for formulas of complexity n to quantify only over V_δ's that are sufficiently elementary in V. Note that for Σ_1-formulas $\phi(x)$, the formula ϕ^* has Σ_1-complexity for sort quantifiers because the only sort quantifier is "$\exists V_\delta$". □

Because defining a satisfaction relation for $\mathbb{L}^s$ runs into definability of truth issues, we limit our analysis to logics $\mathbb{L}^{s,\Sigma_n}$ where we are only allowed to use Σ_n-formulas with sort quantifiers.

As a curiosity, observe that, for example, the logic $\mathbb{L}_{\mathrm{Ord},\omega}$ is not an abstract logic under our criteria because, in particular, there is a proper class of sentences for a given language. This logic has several other undesirable properties as well. It does not have an occurrence number and it can never have a strong compactness cardinal. We will call such logics *pseudo-abstract logics.*

2.2. **Large cardinals.** We collect here several of the large cardinal notions that we use. Occasionally, we defer a definition to later if it is tailored to a specific situation.

First, we have several variants of Woodin cardinals. The notion of externally definable Woodin cardinals is new. The definition of Woodin for strong compactness (due to Dimopoulos) is phrased in the equivalent form of [Dim19, Proposition 3.3].

Definition 2.3.

(1) A cardinal δ is *Woodin* if it is inaccessible and for all $A \subseteq V_\delta$, there is $\kappa < \delta$ such that for all $\kappa < \alpha < \delta$, there is an elementary embedding $j : V \to \mathcal{M}$ with
 (a) crit $j = \kappa$,
 (b) $j(\kappa) > \alpha$,
 (c) $V_\alpha \subseteq \mathcal{M}$, and
 (d) $j(A) \cap V_\alpha = A \cap V_\alpha$.

(2) A cardinal δ is *externally definable Woodin* if it satisfies the definition of a Woodin cardinal when restricting the A's to be externally definable sets. Explicitly, this means that δ is inaccessible and for every formula $\phi(x,a)$ with $a \in V_\delta$, there is $\kappa < \delta$ such that $a \in V_\kappa$ and for all $\kappa < \alpha < \delta$, there is an elementary embedding $j : V \to \mathcal{M}$ with
 (a) crit $j = \kappa$,
 (b) $j(\kappa) > \alpha$,
 (c) $V_\alpha \subseteq \mathcal{M}$, and
 (d) $\phi(\mathcal{M},a) \cap V_\alpha = \phi(V,a) \cap V_\alpha$.
 Here, A has been replaced with the externally definable $\phi(V,a) \cap V_\delta \subseteq V_\delta$.

(3) A cardinal δ is *Woodin for strong compactness* if for every $A \subseteq V_\delta$ there is $\kappa < \delta$ which is $<\delta$-strongly compact for A.

Let us explain more about Woodin for strong compactness cardinals introduced in [Dim19]. Recall first that a cardinal κ is *λ-strong* if there is an elementary embedding $j : V \to \mathcal{M}$ with crit $j = \kappa$, $j(\kappa) > \lambda$, and $V_\lambda \subseteq \mathcal{M}$. We say that κ is *λ-strong for a set A* if additionally $j(A) \cap V_\lambda = A \cap V_\lambda$. With this definition, we can rephrase that a cardinal δ is Woodin if and only if it is inaccessible and for every $A \subseteq V_\delta$ there is $\kappa < \delta$ which is $<\delta$-strong for A. Recall next that a cardinal κ is *λ-strongly compact* if there is an elementary embedding $j : V \to \mathcal{M}$ with crit $j = \kappa$, $j(\kappa) > \lambda$, and there exists a set $s \in \mathcal{M}$ with $|s|^M < j(\kappa)$ such that $j\,"\,\lambda \subseteq s$. We can similarly extend the standard definition of λ-strongly compact cardinals to assert λ-strong compactness for a set A. Now we can consider what happens when we replace the $<\delta$-strongness for A requirement in the definition of Woodin cardinals with $<\delta$-strong compactness, calling the resulting notion *Woodin for strong compactness* cardinals. The second author showed in [Dim19] that every Woodin for strong compactness cardinal δ is already inaccessible and that for every $A \subseteq V_\delta$, there is a $\kappa < \delta$ simultaneously witnessing both $<\delta$-strong compactness and $<\delta$-strongness for A.

The next class of large cardinals we will consider are the recently introduced virtual large cardinals. Given a set-theoretic property P characterized by the existence of elementary embeddings between (set) first-order structures, we say that P holds *virtually* if embeddings characterizing P between structures from V exist in set-forcing extensions of V. We will make repeated use of the following absoluteness lemma for such embeddings of countable structures.

Lemma 2.4. *Suppose M is a countable first-order structure and $j : M \to N$ is an elementary embedding. If W is a transitive (set or class) model of (some sufficiently large fragment of)* ZFC *such that M is countable in W and $N \in W$, then for any finite subset of M, W has some elementary embedding $j^* : M \to N$, which agrees with j on that subset. Moreover, if both M and N are transitive $\in$-structures and j has a critical point, we can additionally assume that crit $(j^*) =$ crit (j).*

In particular, it follows from the lemma that if any forcing extension $V[G]$ has an elementary embedding $j : M \to N$ for some $M, N \in V$, then such an embedding must already exist in every $\mathrm{Coll}(\omega, M)$-extension (see [GS18] for proofs).

Since most large cardinals can be characterized by the existence of elementary embeddings between (set) models of set-theory (in the case of class embeddings $j : V \to \mathcal{M}$ we chop off the universe at an appropriate rank initial segment), they are natural candidates for virtualization. The study of virtual large cardinals was initiated by Schindler when he introduced the notion of a remarkable cardinal and showed that it is equiconsistent with the assertion that the theory of $L(\mathbb{R})$ cannot be changed by proper forcing [Sch00]. He later observed that remarkable cardinals have an equivalent characterization as virtually supercompact cardinals. Other virtual large cardinals were subsequently studied in [GS18]. Unlike their philosophical cousins, the generic large cardinals, virtual large cardinals are actual large cardinals (they are ineffable and more), but they sit much lower in the hierarchy than their original counterparts. They are consistent with L and are in the neighborhood of an ω-Erdos cardinal. In the definitions of virtual large cardinals given below, we will abbreviate the statement that an elementary embedding exists in some forcing extension by saying that there is a “virtual elementary embedding”.

Definition 2.5.

(1) A cardinal κ is *virtually measurable* if for every $\alpha > \kappa$, there is a transitive $\mathcal{M}$ and a virtual elemetary embedding $j : V_\alpha \to \mathcal{M}$ with crit $j = \kappa$.
(2) A cardinal κ is *virtually supercompact* (remarkable) if for every $\lambda > \kappa$, there is $\alpha > \lambda$ and a transitive $\mathcal{M}$ with $\mathcal{M}^\lambda \subseteq \mathcal{M}$ such that there is a virtual elementary embedding $j : V_\alpha \to \mathcal{M}$ with crit $j = \kappa$ and $j(\kappa) > \lambda$.
(3) A cardinal κ is *virtually extendible* if for every $\alpha > \kappa$, there is a virtual elementary embedding $j : V_\alpha \to V_\beta$ with crit $j = \kappa$ and $j(\kappa) > \alpha$. A cardinal κ is *weakly virtually extendible* if we omit the assumption that $j(\kappa) > \alpha$.

Let $C^{(n)}$ be the class club of cardinals α such that $V_\alpha \prec_{\Sigma_n} V$.

(4) A cardinal κ is *virtually $C^{(n)}$-extendible* if for every $\alpha > \kappa$ in $C^{(n)}$, there is $\beta \in C^{(n)}$ and a virtual elementary embedding $j : V_\alpha \to V_\beta$ with crit $j = \kappa$ and $j(\kappa) > \alpha$.[2] A cardinal κ is *weakly virtually $C^{(n)}$-extendible* if we omit the assumption that $j(\kappa) > \alpha$.

Although, for the most part the hierarchy of virtual large cardinals mirrors the hierarchy of their original counterparts, but much lower down, there are anomalies that appear to arise from the following two circumstances. The first is that Kunen's inconsistency does not hold for virtual large cardinals, so that we can have virtual elementary embeddings $j : V_\alpha \to V_\alpha$ with α much larger than the supremum of the critical sequence of j. This accounts for the split in the definition of virtually extendible cardinals into the weak and strong forms given above. In the case of the actual extendible cardinals we can argue using Kunen's inconsistency that the assumption $j(\kappa) > \alpha$ is superfluous, which gives that the weak extendible and extendible cardinals are equivalent. But the equivalence fails in a surprising way in the virtual context, where it is consistent that there are weakly virtually extendible cardinals that are not virtually extendible, and moreover the consistency strength of the existence of such a notion is higher than that of the existence of a virtually extendible cardinal [GH19]. The second circumstance is that the more robust virtual large cardinals arise from large cardinals that have characterizations in terms of the existence of virtual embeddings between rank initial segments V_α, such as rank-into-rank, supercompact, and extendible cardinals. Virtual versions of large cardinal notions lacking such characterizations do not fit properly into the hierarchy. For example, the following is not difficult to see.

Theorem 2.6 ([GS18])**.** *A cardinal κ is virtually supercompact if and only if for every $\lambda > \kappa$, there is $\alpha > \lambda$ and a transitive $\mathcal{M}$ with $V_\lambda \subseteq \mathcal{M}$ such that there is a virtual elementary embedding $j : V_\alpha \to \mathcal{M}$ with crit $j = \kappa$ and $j(\kappa) > \lambda$.*

It follows that the notions of virtually strong and virtually supercompact cardinals coincide, and one should note here that strong cardinals do not have a characterization in terms of existence of embeddings between rank initial segments. This phenomena is also exhibited by the virtually measurable cardinals introduced and studied by Nielsen and Welch [NW19], who showed that virtually measurable cardinals and virtually supercompacts are equiconsistent (in particular, the first virtually measurable in L must be virtually supercompact).

Finally, Section 5 uses the hypothesis that Ord is subtle. Recall that a regular cardinal κ is *subtle* if for every club $C \subseteq \kappa$ and sequence $\langle A_\alpha \mid \alpha < \kappa \rangle$ with $A_\alpha \subseteq \alpha$, there are $\alpha < \beta$ in C such that $A_\alpha = A_\beta \cap \alpha$.

Definition 2.7. We say that Ord *is subtle* if for every class club $C \subseteq$ Ord and every class sequence $\langle A_\alpha \subseteq \alpha \mid \alpha \in C \rangle$ there are $\alpha < \beta \in C$ such that $A_\alpha = A_\beta \cap \alpha$.

The definition requires us to specify precisely what a class is in a universe of set theory. There are a number of approaches to this and we will be intentionally vague about which approach we take in this article. We can be working in first-order logic in the theory ZFC and specify that classes are definable collections. We can be working in any one of the numerous second-order set theories where classes are second-order objects, such as Gödel-Bernays set theory GBC or the much stronger Kelley-Morse set theory KM. We can also assume that our universe is the V_κ of a much larger ZFC-model in which κ is subtle and the classes in this case are the $V_{\kappa+1}$ of this model. In the last case, V_κ together with the classes given by $V_{\kappa+1}$ satisfy Kelley-Morse (and more). Our arguments will require the class axiom *global choice* which asserts that there is a class well-ordering of the universe of sets. Both GBC and Kelley-Morse include global choice,

[2] Bagaria's original definition of $C^{(n)}$-extendibility required only that $V_{j(\kappa)} \prec_{\Sigma_n} V$, but the third author and Hamkins showed in [GH19] that the two definitions are equivalent.

but a universe of set theory need not have a definable well-ordering of all sets, so global choice can fail for definable classes.

3. Woodin cardinals and abstract Henkin structures

In this section, we will give compactness characterizations for Woodin cardinals and their variants. To motivate our characterizations, we start by recalling the compactness characterization of strong cardinals.

Recall that a *standard* model (M, P) of second-order logic has the second-order part P consisting of all subsets of M, so that the second-order quantifiers range over all subsets of the domain. A *Henkin* model (M, P) has the second-order part $P \subseteq \mathcal{P}(M)$ that is a possibly proper sub-collection of the subsets of M, so that the second-order quantifiers cannot access all subsets of the domain.

Theorem 3.1 ([Bon20, Theorem 4.7])**.** *A cardinal κ is λ-strong if and only if the following holds. Given a $\mathbb{L}^2_{\kappa,\omega}(Q^{WF})$ theory T that can be written as an increasing union $T = \cup_{\eta<\kappa} T_\eta$, if every T_η has a (standard) model, then T has a Henkin model whose universe is an ordinal and whose second-order part has all subsets of rank $< \lambda$.*

Requiring that the universe of the Henkin model is an ordinal insures that the condition of having all small subsets of the universe is non-vacuously satisfied because a Henkin model with an arbitrary universe can simply have no subsets of rank $< \lambda$.

Due to the similarity between Woodin and strong cardinals, the compactness characterization of Woodin cardinals will also use Henkin models. However, to accommodate the "$A \subseteq V_\delta$" parameter, we will need a much more specialized notion of a Henkin structure for abstract logics. In a Henkin model we have a nonstandard conception of evaluating the truth of sentences (and the data to carry out this evaluation) because we do not have access to all true subsets, but the collection of sentences remain the same. We encapsulate this idea in the following definition of Henkin structures for abstract logics.

Definition 3.2. Fix a logic $(\mathcal{L}, \vDash_{\mathcal{L}})$ and a language τ.

1. A *Henkin τ-structure for $\mathcal{L}$* is a model of set theory $\hat{M} = (\mathcal{M}, \mathrm{E}, \vDash^*, M)$, with an additional binary relation $\vDash^*$ and a distinguished element M, satisfying the following properties:
   (a) $\hat{M} \vDash \mathrm{ZFC}^*$ (a large finite fragment of ZFC).
   (b) $\tau \in \mathcal{M}$, $\mathcal{L}(\tau) \subseteq \mathcal{M}$.
   (c) For Σ_n-formulas $\varphi(x)$ (where n is taken to be very large), we have $\varphi(\vDash_{\mathcal{L}})$ if and only if $\mathcal{M} \vDash \varphi(\vDash^*)$.
2. Given an $\mathcal{L}(\tau)$-theory T, we say that a Henkin structure $\hat{M} = (M, \mathrm{E}, \mathcal{M}, \vDash^*)$ *Henkin-models* T if for every $\varphi \in T$, $\mathcal{M}$ has a partial $\tau \supseteq \tau_\varphi$-structure on M such that $M \vDash^* \varphi$ and the assignment of the partial τ_φ-structures is coherent in the sense that if $\tau_\varphi \subseteq \tau_\psi$, then the structure assignments extend correspondingly.

We think of $\hat{M}$ as a model of set theory that is correct about the sentences in $\mathcal{L}(\tau)$ but which has a nonstandard satisfaction for them given by the relation $\vDash^*$. The statement "$\tau \in \mathcal{M}$" needs further clarifying. Here we intend that not only is τ an element of the universe $\mathcal{M}$ of $\hat{M}$, but also $\mathcal{M}$ sees τ as a language with its functions, relations, and constants. In practice, we will be working with Henkin models for which E is $\in$ and $(\mathcal{M}, \in)$ is transitive, so these issues will not arise. For our purposes, it appears to be too strong to require that a transitive $\in$-structure $\hat{M}$ has a total τ-structure on M which witnesses that M satisfies T according to $\vDash^*$. Thus, we require only that sufficiently large pieces of τ can be interpreted over M making it satisfy $\phi \in T$

and the interpretation is coherent. It would be easy to achieve a full interpretation of τ if we were willing to give up transitivity, but this appears to us to be unnatural.

Note that, given a Henking τ-structure $\hat{M} = (M, \mathrm{E}, \mathcal{M}, \vDash^*)$ that Henkin models a theory T, we do get in V a total τ-structure on M that is the union of the coherent interpretations from $\mathcal{M}$.

It should be noted that our set-up does not generalize the usual notion of Henkin models for second-order logic, which we described above, because it cannot accommodate models with arbitrary second-order parts, but instead restricts it only to models with a "sensible" second-order part, namely one that arises from a model of set theory that reflects fundamental properties of second-order logic. This is a natural class of strong Henkin models for second-order logic. This restricted notion of a Henkin model for second-order logic, which lies somewhere between the standard and Henkin models, appears to be interesting in its own right.

The motivation behind the clauses of Definition 3.2 will become apparent after the proof of Theorem 3.6, but we will try to give some explanation here. Henkin models came up in the proof of Theorem 3.1 through elementary embeddings of the form $j : V \to \mathcal{M}$ with $V_\lambda \subseteq M$ for some cardinal λ. More precisely, for a second-order theory T, the proof gives a second-order model of $j\,``\,T$ in the sense of $\mathcal{M}$ and such a model can interpret the second-order variables correctly only if they refer to subsets of $V_\lambda$. It is crucial that the relation $j(\vDash_2)$ is the same as $\vDash_2$ for structures in $\mathcal{M}$. Also, the renaming of $T$ to $j\,``\,T$ is not necessarily in $\mathcal{M}$.

Our proof will work with abstract logics $\mathcal{L}$, so we will once again use elementary embeddings $j : V \to \mathcal{M}$ and obtain models of the desired theory in $\mathcal{M}$. Here, we face the challenge that the satisfaction relation $\vDash_\mathcal{L}$ may not be the same as $j(\vDash_\mathcal{L})$ and the latter is essentially the $\vDash^*$ relation of the definition. Also, we once again find a model for $j\,``\,T$ instead of $T$, but since $j\,``\,T$ may not be in $\mathcal{M}$, $\mathcal{M}$ may not have the right τ-structure for M obtained via the renaming $j : \tau \to j\,``\,\tau$, although it will have coherent partial τ-structures resulting from this renaming.

In order to ensure some similarity between $\mathcal{L}$ and $j(\mathcal{L})$, we add Clause (1c) which guarantees that any Σ_2-definable aspect of $\mathcal{L}$ is mirrored by $\mathcal{M}$. For instance, $\mathcal{M}$ will mirror any absolute properties of $\mathcal{L}$ (such as expressing the statements of $\mathbb{L}_{\mathrm{Ord},\omega}$) and statements about well-foundedness.

The following are a list of useful properties of Henkin structures. Often, whether or not a Henkin structure correctly computes satisfaction for a logic will depend on these properties. For instance, a Henkin structure $\hat{M}$ will correctly verify the well-foundedness quantifier Q^{WF} when $\mathcal{M}$ is itself well-founded. Note that we emphasize well-foundedness rather than transitivity to make the notion closed under isomorphism.

Definition 3.3. Given a Henkin τ-structure $\hat{M} = (\mathcal{M}, \mathrm{E} \vDash^*, M)$ for a logic $(\mathcal{L}, \vDash_\mathcal{L})$ and language τ, we say:

(1) $\hat{M}$ is *well-founded* whenever $(\mathcal{M}, \mathrm{E})$ is well-founded.
(2) $\hat{M}$ is *transitive* whenever E is $\in$ and $(\mathcal{M}, \in)$ is transitive.
(3) $\hat{M}$ is *full up to* λ whenever $V_\lambda \subseteq \mathcal{M}$ (if $\mathcal{M}$ is not well-founded, this means for every $x \in V_\lambda$, there is $y \in \mathcal{M}$ such that E is well-founded below y and the transitive collapse of (y, E) is x).
(4) Suppose that $\hat{M}$ is transitive and full up to λ. We say $\hat{M}$ is *$\mathcal{L}$-correct up to* λ if $(\vDash^*) \restriction V_\lambda \times V_\lambda = (\vDash) \restriction V_\lambda \times V_\lambda$.
(5) Suppose that $\hat{M}$ is transitive and $A \subseteq \mathcal{M}$. We say that $\hat{M}$ is *n-correct for A* if
    - for every Σ_n-formula $\phi(x)$ and $a \in A$, $\mathcal{M} \vDash \phi(a)$ if and only if $\phi(a)$ holds,
    - for every Σ_n-formula $\varphi(x)$, $\varphi(\vDash_\mathcal{L})$ if and only if $\hat{M} \vDash \varphi(\vDash^*)$.

We will now be able to characterize Woodin cardinals in terms of chain compactness for sufficiently correct Henkin models of arbitrary logics. Before we give the proof, we are going to need a notion of closure points for abstract logics, which is captured by the following result.

Proposition 3.4. *Let $(\mathcal{L}, \models_{\mathcal{L}})$ be an abstract logic. There is a closed unbounded class of cardinals α such that*

(1) $\mathcal{L} \restriction V_\alpha : V_\alpha \to V_\alpha$,
(2) $o(\mathcal{L}) < \alpha$.

The proof is a standard closure argument.

Definition 3.5. For an abstract logic $(\mathcal{L}, \models_{\mathcal{L}})$, a cardinal α that satisfies the properties of Proposition 3.4 is called a *closure point of $\mathcal{L}$*.

Theorem 3.6. *The following are equivalent for a cardinal δ.*

(1) *δ is Woodin.*
(2) *For every logic $(\mathcal{L}, \models_{\mathcal{L}})$ with closure point δ, there is $\kappa < \delta$ such that given any $\kappa < \lambda < \delta$ any theory $T \subseteq \mathcal{L} \cup \mathbb{L}_{\kappa,\omega}(\tau)$ for a language $\tau \in V_\lambda$, if T can be written as an increasing union $T = \bigcup_{\eta<\kappa} T_\eta$ of satisfiable theories, then T has a transitive Henkin τ-structure that is full up to λ and $\mathcal{L}$-correct up to λ.*

Proof. For the forward direction, fix a logic $\mathcal{L}$ with a closure point at δ. By definition the occurrence number of $\mathcal{L}$, $o(\mathcal{L}) < \delta$. For this proof, we need to consider another class related to $\mathcal{L}$. Let R be the class of quadruples (τ, σ, f, f_*) such that τ and σ are languages, f is a renaming from σ to τ and $f_* : \mathcal{L}(\sigma) \to \mathcal{L}(\tau)$ is the associated bijection. Using the fact that δ is Woodin, let κ be a $<\delta$-strong for $\mathcal{L} \restriction V_\delta$, $\models_{\mathcal{L}} \restriction V_\delta$, and $R \restriction V_\delta$ cardinal above $o(\mathcal{L})$.

Fix a language $\tau \in V_\delta$ and a $\mathcal{L}' = \mathcal{L} \cup \mathbb{L}_{\kappa,\omega}(\tau)$-theory $T = \bigcup_{\eta<\kappa} T_\eta$ as in the second clause. Note that, since a Woodin cardinal δ is Mahlo, δ is also a closure point of $\mathbb{L}_{\kappa,\omega}$. Thus δ is a closure point of $\mathcal{L}'$, and so the requirement that $\tau \in V_\delta$ implies that $T \in V_\delta$ as well. Fix $\lambda < \delta$ which is above the rank of τ and $\mathcal{L}(\tau)$. Let $j : V \to \mathcal{N}$ be an elementary embedding with

- crit $(j) = \kappa$, $j(\kappa) > \lambda$,
- $V_\lambda \subseteq \mathcal{N}$,
- $\mathcal{L} \restriction V_\lambda = j(\mathcal{L} \restriction V_\lambda) \restriction V_\lambda$
- $(\models_{\mathcal{L}}) \restriction V_\lambda \times V_\lambda = j((\models_{\mathcal{L}}) \restriction V_\lambda \times V_\lambda) \restriction V_\lambda \times V_\lambda$.
- $R \restriction V_\lambda^4 = j(R) \restriction V_\lambda^4$.

Consider the sequence $j(\langle T_\eta \mid \eta < \kappa \rangle) := \langle T^*_\eta \mid \eta < j(\kappa) \rangle$. By elementarity, $\mathcal{N}$ thinks that T^*_κ is a satisfiable $j(\mathcal{L}')(j(\tau))$-theory, and, we have that $j \,"\, T \subseteq T^*_\kappa$.

Let $M \in \mathcal{N}$ be a $j(\mathcal{L}')$-model of T^*_κ in the sense of $\mathcal{N}$. Let $\mathcal{M} = V_\theta^{\mathcal{N}}$ so that θ is a closure point of $j(\mathcal{L}')$ and that $V_\theta^{\mathcal{N}} \prec_{\Sigma_m} \mathcal{N}$ for a very large m. Define $\models^*$ to be the satisfaction relation $\models_{j(\mathcal{L}')}$ from $\mathcal{N}$ restricted to $V_\theta^{\mathcal{N}}$. We claim that $\hat{M} = (\mathcal{M}, \in, \models^*, M)$ is a Henkin τ-structure for T. It is easy to see that it satisfies clauses 1(a), 1(b) and 1(c) of Definition 3.2. It remains to verify clause (2).

Fix $\phi \in T$ and let τ_ϕ be the fragment of $\mathcal{L}(\tau)$ that occurs in ϕ. Note that τ_ϕ has size less than κ by our assumption that $o(\mathcal{L}) < \kappa$. We will argue that M is a τ_ϕ-structure in $\mathcal{M}$ via the renaming $f : \tau_\phi \to j(\tau_\phi) = j \,"\, \tau_\phi$ defined by $f = j \restriction \tau_\phi$. Note that the renaming f is an element of $\mathcal{N}$, although the entire map $j : \tau \to j \,"\, \tau$ might not be in $\mathcal{N}$ (since the embedding is not a supercompactness embedding).[3] It suffices then to argue that $f_*(\phi) = j(\phi)$ in $\mathcal{N}$ since by elementarity, we know that in $\mathcal{N}$, $M \models_{j(\mathcal{L})} j(\phi)$. The argument that $f_*(\phi) = j(\phi)$

[3]Note that technically the bijection $j : \tau \to j \,"\, \tau$ is not a renaming in our sense because it is not an element of $\mathcal{N}$, however it does have the required renaming properties because it is the union of a coherent collection of renamings that are elements of $\mathcal{N}$.

is straightforward for sentences from $\mathbb{L}_{\kappa,\omega}(\tau)$, but is not obvious for sentences from $\mathcal{L}(\tau)$ because we do not know how the satisfaction for $\mathcal{L}$ works. So let us argue as follows.

By our assumptions, there is a language $\sigma \in V_\kappa$ with $\mathcal{L}(\sigma) \in V_\kappa$ and a renaming $g : \sigma \to \tau_\phi$. Let $g_*(\bar{\phi}) = \phi$. By elementarity, $j(g) : \sigma \to j(\tau_\phi)$ is a renaming from σ to $j(\tau_\phi) = j\text{ " }\tau_\phi$ in $\mathcal{M}$. It is easy to check that $j(g) = f \circ g$. Now, by elementarity, we have

$$(j(g))_*(\bar{\phi}) = j(g_*)(\bar{\phi}) = j(g_*(\bar{\phi})) = j(\phi).$$

By our observations from Section 2 we have, $j(g)_* = (f \circ g)_* = f_* \circ g_*$. The g_* on the right of the last equation is the bijection between $j(\mathcal{L})(\sigma)$ and $j(\mathcal{L})(\tau_\phi)$ in $\mathcal{M}$ associated with g. A priori there is no reason to believe that the two maps which we both denoted by g_* are the same, but we will argue that they are by our assumption that $R \restriction V_\lambda^4 = j(R) \restriction V_\lambda^4$ since g and g_* are elements of V_λ. Thus, we have $f_*(\phi) = f_*(g_*(\bar{\phi})) = j(g)_*(\bar{\phi}) = j(\phi)$.

Since $V_\lambda \subseteq M$ we have fullness up to λ and by the properties of j, we have correctness up to λ as well.

For the converse direction, suppose the second clause is true and fix any set $A \subseteq V_\delta$. We define the associated logic $\mathbb{L}^A$ as follows. We extend $\mathbb{L}^2$ by adding, for a binary relation E, a single new formula $\phi_{A,E}(x)$, which holds whenever E is well-founded on the transitive closure of x, $\text{tc}_E\, x$, and $\text{tc}_E\, x$ is isomorphic to $\text{tc}_\in a$ for some $a \in A$.

It is clear that $o(\mathbb{L}^A) = \omega$ and δ is a closure point of $\mathbb{L}^A$, so there is a cardinal κ for $\mathbb{L}^A$ as in the hypothesis, which we want to show is $<\delta$-strong for A.

Fix $\alpha > \kappa$ and fix $\lambda \gg \alpha$ below δ. Let τ be the language consisting of a binary relation $\in$ and constants $\{c_x \mid x \in V_{\alpha+1}\} \cup \{c\}$. Let T be the $\mathbb{L}^A \cup \mathbb{L}_{\kappa,\omega}(\tau)$-theory T consisting of the following sentences.

(1) $\text{ED}_{\mathbb{L}_{\kappa,\omega}}(V_{\alpha+1}, \in, c_x)_{x \in V_{\alpha+1}}$,
(2) $\{c_i < c < c_\kappa \mid i < \kappa\}$,
(3) Magidor's Φ
(4) $\forall x\, (x \in c_{A \cap V_\alpha} \to \phi_{A,\in}(x))$,
(5) $\forall x\, (\phi_{A,\in}(x) \wedge x \in c_{V_\alpha} \to x \in c_{A \cap V_\alpha})$.

We can write T as an increasing κ-union of satisfiable theories by filtrating the sentences in (2). Let $\hat{M} = (\mathcal{M}, \mathrm{E}, \models^*, M)$ be a transitive Henkin τ-structure for T that is full up to λ and $\mathbb{L}^A$-correct up to λ.

Note once again that even though $\mathcal{M}$ has only partial interpretations of τ on M, we get a total interpretation by unioning up the coherent interpretations from $\mathcal{M}$ in V.

With the standard set theoretic arguments, it suffices to produce an elementary embedding $j : V_{\alpha+1} \to M$ with crit $(j) = \kappa$, $V_\alpha \subseteq M$, $j(\kappa) > \alpha$ and $A \cap V_\alpha = j(A \cap V_\alpha) \cap V_\alpha$.

By the first clause in the definition of T, there is an elementary embedding $j : V_{\alpha+1} \to M$ (because first-order elementarity is absolute). The second clause guarantees that j has a critical point $\leq \kappa$ and since every ordinal $\alpha < \kappa$ is definable in the logic $\mathbb{L}_{\kappa,\omega}$, the critical point must be exactly κ. Note that the formulas used to ensure this are in $\mathbb{L}_{\kappa,\omega}$, and thus reflected correctly in the transitive structure $\hat{M}$. Since $\hat{M}$ is transitive, it must be correct about M being well-founded. So we can assume without loss that M is transitive. Since $\hat{M}$ is full up to λ, it is full up to α and we have $V_\alpha \subseteq \mathcal{M}$. But then since $\models^*$ in $\mathcal{M}$ satisfies the basic properties of $\mathbb{L}^2$, and it believes that $M \models^* \Phi$, it follows that $V_\alpha \subseteq M$.

Next, we will argue that for $a \in V_\alpha$, $M \models^* \phi_{A,\in}(a)$ if and only if $a \in A$. In V, the logic $\mathbb{L}^A$ has the property that if N is a well-founded model in $\mathbb{L}^A$ and $\bar{N}$ is any transitive first-order submodel of N such that for some transitive (from the perspective of N) set B, $B \subseteq \bar{N}$, then $\bar{N}$ and N must agree on the formulas $\phi_{A,\in}(a)$ for $a \in B$. Thus, $\mathcal{M}$ satisfies the same property of the relation $\models^*$. It follows, for our particular case, that V_α and M agree on $\phi_{A,\in}(a)$ for $a \in V_\alpha$.

But now the point is that V_α is small enough that the correctness of $\models^*$ up to λ guarantees that $\phi_{A,\in}(a)$ holds true if and only if $a \in A$. Thus, M must also be correct with respect to the formula $\phi_{A,\in}(a)$ for $a \in V_\alpha$. Since M satisfies $\forall x\, (x \in c_{A\cap V_\alpha} \to \phi_{A,\in}(x))$, it follows that for $a \in V_\alpha$ and $a \in j(A \cap V_\alpha)$, we have $a \in A$. Thus, $j(A) \cap V_\alpha \subseteq A \cap V_\alpha$. Since M satisfies $\forall x\, (\phi_{A,\in}(x) \wedge x \in c_{V_\alpha} \to x \in c_{A\cap V_\alpha})$, it follows that for $a \in V_\alpha$ and $a \in A$, we have $a \in j(A \cap V_\alpha)$. Thus, $A \cap V_\alpha \subseteq j(A) \cap V_\alpha$, which gives the desired equality.

□

This sort of characterization also suggests a new hierarchy of Woodin cardinals based on the logics that they endow non-standard chain compactness to. For instance, using Väänänen's sort logics, we get definable notions of Woodin cardinals.

Theorem 3.7. *The following are equivalent for a cardinal δ.*

1. *δ is externally definable Woodin.*
2. *Fix $n < \omega$ and $a \in V_\delta$. There is a cardinal $\kappa < \delta$ above the rank of a such that for every $\kappa < \lambda < \delta$ and theory $T \subseteq \mathbb{L}^{s,\Sigma_n}_{\kappa,\omega}(\tau)$ for a language $\tau \in V_\lambda$, if T can be written as an increasing union $T = \bigcup_{\eta<\kappa} T_\eta$ of satisfiable theories, then T has an n-correct for V_λ Henkin τ-structure that is full up to λ.*

Proof. Suppose δ is externally definable Woodin. Fix $n < \omega$. Let

$$A = \{(\ulcorner\varphi(x)\urcorner, b) \mid \varphi(x) \text{ is } \Sigma_n \text{ and } V \models \varphi(b)\},$$

and note that A is definable because Σ_n-truth is definable. Fix a language τ and let $\kappa > \text{rank }(\tau)$ be as in the definition of externally definable Woodin cardinals for the definable class A. Fix a $\mathbb{L}^{s,\Sigma_n}_{\kappa,\omega}(\tau)$-theory T as in (2). Pick a limit $\lambda > \kappa$ below δ such that $T \in V_\lambda$. Let $j : V \to \mathcal{N}$ be an elementary embedding with crit $j = \kappa$, $j(\kappa) > \lambda$, $V_\lambda \subseteq \mathcal{N}$, and $A \cap V_\lambda = j(A) \cap V_\lambda$.

Now consider the sequence $j(\langle T_\eta \mid \eta < \kappa\rangle) = \langle T^*_\eta \mid \eta < j(\kappa)\rangle$. By elementarity, it is an increasing sequence and hence, T^*_κ is a theory containing $j\,"\,T$. So, by elementarity, $\mathcal{N}$ thinks that T^*_κ has a model M in $\mathbb{L}^{s,\Sigma_n}_{\kappa,\omega}$ (technically it would have to be $\mathbb{L}^{s,\Sigma_n}_{j(\kappa),\omega}$, but we can always prune to get rid of the extra assertions). Let $\mathcal{M} = V^{\mathcal{N}}_\theta \prec_{\Sigma_m} \mathcal{N}$ for some large enough θ and $m > n$, and let $\models^*$ be the satisfaction $\models_{\mathbb{L}^{s,\Sigma_n}_{\kappa,\omega}}$ of $\mathcal{N}$ restricted to $V^{\mathcal{N}}_\theta$. We claim that $\hat{M} = (V^{\mathcal{N}}_\theta, \in, \models^*, M)$ is an n-correct for V_λ Henkin τ-structure for T.

The argument that $\hat{M}$ is a Henkin τ-structure for T is even easier than in the proof of Theorem 3.6 because we know exactly what the formulas in $\mathbb{L}^{s,\Sigma_n}_{\kappa,\omega}$ look like. So it remains to check n-correctness for V_λ. Since A is definable and j is elementary, we have

$$j(A) = \{(\ulcorner\varphi(x)\urcorner, b) \mid \varphi(x) \text{ is } \Sigma_n \text{ and } \mathcal{N} \models \varphi(b)\}$$

Also, since λ is limit, we have that

$$A \cap V_\lambda = \{(\ulcorner\varphi(x)\urcorner, b) \mid b \in V_\lambda,\ \varphi(x) \text{ is } \Sigma_n \text{ and } V \models \varphi(b)\}$$

and similarly for $j(A) \cap V_\lambda$. By our assumptions on j, $A \cap V_\lambda = j(A) \cap V_\lambda$. But this means precisely that for $a \in V_\lambda$, we have $V \models \varphi(a)$ if and only if $\mathcal{N} \models \varphi(a)$ if and only if $V^{\mathcal{N}}_\theta \models \varphi(a)$.

For the converse direction, we fix a Σ_n-formula $\varphi(x, y)$ and a parameter $a \in V_\delta$. Let $\kappa < \delta$ be as in the statement of (2) for $n + 1$ and a. Fix $\kappa < \alpha < \delta$. We adapt the corresponding argument from Theorem 3.6. We replace the assertion $\phi_{A,\in}(x)$ of that proof with the formula $\phi_{\varphi,E}(x)$ asserting that the transitive closure of x is isomorphic to the transitive closure of an element b such that $\phi(b, a)$. To express $\phi_{\varphi,E}$ in the logic $\mathbb{L}^{s,\Sigma_{n+1}}_{\kappa,\omega}$ we use that $\varphi(x, y)$ is Σ_n and the parameter a is definable using an infinitary atomic formula since it has rank below κ. So we can let T be the theory as in that proof, but expressed using the sort logic $\mathbb{L}^{s,\Sigma_{n+1}}_{\kappa,\omega}$. By assumption T must have an n-correct for V_λ Henkin τ-structure $\hat{M} = (\mathcal{M}, \in, \models^*, M)$ that is full up to λ for $\lambda \gg \alpha$. As in that proof we get an elementary embedding $j : V_{\alpha+1} \to M$ with crit $j = \kappa$. So it remains to verify that $M \models^* \phi_{\varphi,\in}(a)$ for $a \in V_\alpha$ if and only if $\varphi(a)$ holds in V. Since by $n+1$-correcteness, $\mathcal{M}$ is correct about the Σ_{n+1} properties of $\models^*$, it recognizes that $\models^*$ is satisfaction for $\mathbb{L}^{s,\Sigma_n}_{\kappa,\omega}$ (this is where we need $n+1$ instead of n). Thus, $\mathcal{M}$ believes that $M \models^* \phi_{\varphi,\in}(a)$ if and only if $\varphi(a)$ holds. Again, by n-correctness, $\mathcal{M}$ is correct about $\varphi(a)$ for $a \in V_\alpha$, and thus, so is M.

Finally, the formula $\varphi(x, y)$ is Σ_n, so the n-correctness of $\hat{M}$ and the its fullness up to α guarantees that

$$\varphi(\mathcal{M}, a) \cap V_\alpha = \varphi(V, a) \cap V_\alpha.$$

□

Next, we observe that if we relax clause (2) of Theorem 3.4 to just $<\kappa$-satisfiable theories, then we get a characterisation of Woodin for strong compactness cardinals.

Theorem 3.8. *The following are equivalent for a cardinal δ.*

1. *δ is Woodin for strong compactness.*
2. *For every logic $(\mathcal{L}, \models_{\mathcal{L}})$ with occurrence number $o(\mathcal{L}) < \delta$,there is $\kappa < \delta$ such that given any $\kappa < \lambda < \delta$ and theory $T \subseteq \mathcal{L} \cup \mathbb{L}_{\kappa,\omega}(\tau)$ with $\tau \in V_\lambda$, if T is $<\kappa$-satisfiable, then T has a transitive Henkin τ-structure that is full up to λ and $\mathcal{L}$-correct up to λ.*

Proof. For (1) implies (2), the only difference from Theorem 3.6, is that the embedding $j : V \to \mathcal{N}$ that we use can be assumed to satisfy the strong compactness λ-covering property. Thus, there is a set $s \in \mathcal{N}$ such that $j " \lambda \subseteq s$ and $|s|^{\mathcal{N}} < j(\kappa)$. It follows that there is a set $t \in \mathcal{N}$ such that $j " T \subseteq t$ and $|t|^{\mathcal{N}} < j(\kappa)$. Then, $t \cap j(T)$ is a $j(\mathcal{L})$-theory containing $j " T$ and we use the fact that by elementarity, $j(T)$ is $<j(\kappa)$-satisfiable in the sense of $\mathcal{N}$.

For (2) implies (1), if we fix $A \subseteq V_\delta$, the same proof as for Theorem 3.6 can be used to show that there is $\kappa < \delta$ which is $<\delta$-strong for A. We fix $\kappa < \alpha \ll \lambda < \delta$. We augment our language there by a constant s and augment our theory T to contain statements $\{c_\xi \in s \mid \xi < \lambda\}$ and the statement $|s| < c_\kappa$. The theory T is clearly $<\kappa$-satisfiable since we can always satisfy $<\kappa$-many statements $c_\xi \in s$ together with $|s| < c_\kappa$. Clearly the embedding $j : V_{\alpha+1} \to \mathcal{M}$ is then a λ-strong compactness embedding as witnessed by the interpretation of s.

□

We could ask about pushing these results higher, mixing characterizations from [Mag71, Ben78, Bon20] with our new notion of abstract Henkin structures to find model-theoretic characterizations for Woodin for supercompactness (see [AS07, Apt12]) or prospective notions of Woodin for extendibility. However, Perlmutter [Per15, Theorem 5.10] has shown that Woodin for supercompactness is equivalent to being a Vopěnka cardinal, which has a compactness characterization due to Makowsky [Mak85, Theorem 2], namely, κ is Vopěnka if and only if $V_\kappa \models$ "Every logic has a strong compactness cardinal".

4. Virtual large cardinals

In this section, we will give compactness characterizations for various virtual large cardinals, starting with virtually extendible cardinals, using a new notion of pseudo-models.

Sometimes to prove that a cardinal κ is some virtual large cardinal from a given compactness assumption, we will need to argue that the compactness assumption yields the existence of a *virtual $\mathcal{L}$-embedding* $j : M \to N$ (with $M, N \in V$) for an abstract logic $\mathcal{L}$. Let us explain what we mean by virtual $\mathcal{L}$-embeddings. Having fixed a language τ, we have that both $\mathcal{L}(\tau)$ and $\models_{\mathcal{L}}$ (restricted to $M \cup N$) are sets. Using these two actual sets from V, we can interpret $\mathcal{L}$ (restricted to these models) in a forcing extension of V de re (using the sets from V) and not de dicto (using its definition from V as interpreted by the forcing extension). In this way, a virtual $\mathcal{L}$-elementary embedding $j : M \to N$ has the property that for every formula $\varphi(x)$ in the set $\mathcal{L}(\tau)$ and $a \in M$, M satisfies $\varphi(a)$ according to $\models_{\mathcal{L}}$ if and only if N satisfies $\varphi(f(a))$ according to $\models_{\mathcal{L}}$. We should note here that de dicto interpretations of an abstract logic in a forcing extension, using its V-definition, might yield a logic with entirely different properties. For instance the logic $\mathbb{L}_{\kappa,\kappa}$ for a weakly compact κ defined using the parameter κ in V will turn into the logic $\mathbb{L}_{\omega_1,\omega_1}$ in a forcing extension by $\mathrm{Coll}(\omega, {<}\kappa)$.

The following simple observation about virtual embeddings will be used repeatedly.

Observation 4.1. *Suppose $j : V_\alpha \to M$ is a virtual embedding with crit $j = \kappa$ and V_α has a bijection $f : \kappa \to A$. Then j " $A \in V$.*

Proof. For $a \in A$, given $f(\xi) = a$, we have $j(a) = j(f)^{-1}(\xi)$. Since $M \in V$ by assumption, we have both f and $j(f)$ in V, and therefore we can recover j " A. □

We start with virtually extendible cardinals. Magidor showed that extendible cardinals are the strong compactness cardinals for the second-order infinitary logic $\mathbb{L}^2_{\kappa,\kappa}$ [Mag71].

Theorem 4.2 (Magidor)**.** *Every extendible cardinal is a strong compactness cardinal for the second-order infinitary logic $\mathbb{L}^2_{\kappa,\kappa}$.*

We will adapt Magidor's characterization to virtually extendible cardinals using the new notion of pseudo-models. Given a theory T of a logic $\mathcal{L}(\tau)$ and a sub-language σ of τ, we let $T^\sigma := T \cap \mathcal{L}(\sigma)$. The notion of a "forth system" below is simply the forward direction of back-and-forth systems from model theory, see Definition 4.24.

Definition 4.3. Let T be a $\mathcal{L}(\tau)$-theory, $\mathcal{M}$ be a τ^*-structure, and δ be a cardinal.

(1) A *δ-forth system* $\mathcal{F}$ from τ to τ^* is a collection of renamings $f : \sigma \to \sigma^*$ with $\sigma \in \mathcal{P}_\delta\tau$ and $\sigma^* \in \mathcal{P}_\delta\tau^*$ satisfying the following properties.
  (a) $\emptyset \in \mathcal{F}$.
  (b) If $f \in \mathcal{F}$ and $\tau_0 \in \mathcal{P}_\delta\tau$, then there is $g \in \mathcal{F}$ with $f \subseteq g$ and $\tau_0 \subseteq \mathrm{dom}\, g$.

(2) $\mathcal{M}$ is a *δ-pseudo-model for T* if there is a δ-forth system $\mathcal{F}$ from τ to τ^* such that for every $f \in \mathcal{F}$, $\mathcal{M}$ is a model of f_* " $T^{\mathrm{dom}(f)}$.

We will call ω-forth systems and ω-pseudo-models simply forth systems and pseudo-models respectively.

Definition 4.4. Let $\mathcal{L}$ be a logic. We say that κ is a *strong pseudo-compactness cardinal for $\mathcal{L}$* if every ${<}\kappa$-satisfiable $\mathcal{L}$-theory has a κ^+-pseudo model.

Note that a strong pseudo-compactness cardinal for $\mathcal{L}$ is a weak compactness cardinal for $\mathcal{L}$. We will see below that the converse fails to hold.

Theorem 4.5. *The following are equivalent for a cardinal κ.*

(1) *κ is virtually extendible.*
(2) *Every ${<}\kappa$-satisfiable $\mathbb{L}^2_{\kappa,\kappa}(\tau)$-theory has a pseudo-model.*
(3) *Every ${<}\kappa$-satisfiable $\mathbb{L}^2 \cup \mathbb{L}_{\kappa,\omega}(\tau)$-theory has a pseudo-model.*
(4) *Same as (2) but with pseudo-model replaced by κ^+-pseudo-model.*

In particular, virtually extendible cardinals κ are precisely the strong pseudo-compactness cardinals for $\mathbb{L}^2_{\kappa,\kappa}$.

Proof. Assume the compactness principle in (3) and fix $\alpha > \kappa$. Let τ be the language consisting of a binary relation $\in$ and constants $\{c_x \mid x \in V_\alpha\} \cup \{d_\xi \mid \xi \leq \alpha\} \cup \{c\}$. Let T be the following $\mathbb{L}^2 \cup \mathbb{L}_{\kappa,\omega}(\tau)$-theory:

$$\mathrm{ED}_{\mathbb{L}_{\kappa,\omega}}(V_\alpha, \in, c_x)_{x\in V_\alpha} \cup \{c_\xi < c < c_\kappa \mid \xi < \kappa\} \cup \{\Phi\} \cup \{d_\xi < d_\eta < c_\kappa \mid \xi < \eta \leq \alpha\},$$

where ED stands for elementary diagram and each constant c_x is interpreted as x in V_α, and Φ is Magidor's $\mathbb{L}^2(\{\in\})$-sentence encoding the well-foundedness of $\in$ and that the model is isomorphic to some V_β. Clearly V_α satisfies every piece of T of size less than κ.

By assumption, there is some τ^*-structure $\mathcal{M}$ and a forth system $\mathcal{F}$ witnessing that $\mathcal{M}$ is a pseudo-model for T. We can fix a way of renaming $\in$ and look at the forth system extending this renaming, so that without loss, we assume that τ and τ^* use the same symbol for $\in$. Thus, $\mathcal{M}$ is a model of $\mathrm{Th}_{\mathbb{L}_{\kappa,\omega}}(V_\alpha,\in)$ and Φ. In particular, $\mathcal{M}$ is well-founded and the Mostowski collapse gives $\pi:\mathcal{M}\cong V_\beta$ for some β. So V_β is a pseudo-model for T. The witnessing forth system under the inclusion ordering is a poset $\mathbb{P}$. Forcing with $\mathbb{P}$ yields a bijection[4] $f:\tau\to\tau^*$ in the forcing extension that is a union of the generically chosen collection of renamings from $\mathcal{F}$. The bijection f allows us to build a virtual $\mathbb{L}_{\kappa,\omega}$-elementary embedding $j:V_\alpha\to V_\beta$. We also need to verify that crit $j=\kappa$ and $j(\kappa)>\alpha$. Since every ordinal $\xi<\kappa$ is definable in the logic $\mathbb{L}_{\kappa,\omega}$ in any transitive model of set theory of which it is an element (argue by induction on $\xi<\kappa$), the critical point of j must be κ. The inclusion of the d_ξ constants forces there to be ordinals of order-type α below $j(\kappa)$, so $j(\kappa)>\alpha$. Thus, (3) $\to$ (1).

Now suppose that κ is virtually extendible. Given a $<\kappa$-satisfiable $\mathbb{L}^2_{\kappa,\kappa}(\tau)$-theory T, let $F:\mathcal{P}_\kappa T\to \mathit{Str}\,\tau$ be a map such that $F(s)\models s$. Let α be large enough so V_α contains F and witnesses this property. Using virtual extendibility, in a forcing extension $V[G]$ by $\mathrm{Coll}(\omega,V_\alpha)$, there an elementary embedding $j:V_\alpha\to V_\beta$ with critical point κ and $j(\kappa)>\alpha$. Since the forcing $\mathrm{Coll}(\omega,V_\alpha)$ has size $|V_\alpha|$, we can cover $j\,"\,T$ by a set Y in V of size $<j(\kappa)$. We will argue that the $\tau^*=j(\tau)$-structure $\mathcal{M}=j(F)(Y)$ is a κ^+-pseudo-model of T.

What we would like to say is that $\mathcal{F}$ is the collection of all $f\subseteq j\restriction\tau$ of size κ, but we cannot do this because we do not have access to j in V. However, we can do something relatively close using the forcing relation. Let $\dot{j}$ be a $\mathrm{Coll}(\omega,V_\alpha)$-name that is forced to be an elementary embedding from V_α to V_β with the required properties. Define $\mathcal{F}$ to be the collection of all renamings $f:\sigma\to\sigma^*$ with $\sigma\in\mathcal{P}_{\kappa^+}\tau$ and $\sigma^*\in\mathcal{P}_{\kappa^+}\tau^*$ such that there is a condition $p\in\mathrm{Coll}(\omega,V_\alpha)$ with $p\Vdash\check{f}=\dot{j}\restriction\check{\sigma}:\check{\sigma}\to\dot{j}\,"\,\check{\sigma}$. The system $\mathcal{F}$ is non-empty by Observation 4.1. It has the extension property because whenever a condition $p\Vdash\check{f}=\dot{j}\restriction\check{\sigma}:\check{\sigma}\to j\,"\,\check{\sigma}$ and there is some $\tau_0\in\mathcal{P}_{\kappa^+}\tau$, then there is a condition $q\leq p$ deciding that some g is a renaming between $\sigma\cup\tau_0$ and $j\,"\,(\sigma\cup\tau_0)$, and since $q\leq p$, g must extend f.

It is obvious that (4) $\to$ (3). □

Corollary 4.6. *Every virtually extendible cardinal κ is a weak compactness cardinal for $\mathbb{L}^2_{\kappa,\kappa}$.*

The converse fails to hold. Let us observe here that much weaker large cardinals κ can be weak compactness cardinals for $\mathbb{L}^2_{\kappa,\kappa}$. Hamkins and Johnstone defined that an inaccessible cardinal κ is *strongly uplifting* if for every $A\subseteq\kappa$, there are arbitrarily large regular $\theta>\kappa$ such that $(V_\kappa,\in,A)\prec(V_\theta,\in,\bar{A})$ for some $\bar{A}\subseteq V_\theta$ [HJ17]. Strongly uplifting cardinals are weaker than subtle cardinals in strength. Given a $<\kappa$-satisfiable theory T in $\mathbb{L}^2_{\kappa,\kappa}(\tau)$ of size κ, we can assume without loss that $T\subseteq V_\kappa$. Choose a large enough θ such that $(V_\kappa,\in,T)\prec(V_\theta,\in,\bar{T})$ and V_θ sees that every proper initial segment of T has a model. By elementarity, $(V_\kappa,\in,T)$ must then also satisfy that every proper initial segment of T has a model, but then again by elementarity, we get that $(V_\theta,\in,\bar{T})$ satisfies that every proper initial segment of $\bar{T}$, in particular T, has a model. But, of course, V_θ is correct about it. It is interesting whether the converse holds.

Question 4.7. *Does a weakly compact cardinal κ for the logic $\mathbb{L}^2_{\kappa,\kappa}$ have to be strongly uplifting?*

[4]Although a bijection between languages coming from outside the universe V may not satisfy the properties of a renaming with respect to a particular logic from V, the bijection f does because it is a union of renamings from V.

Magidor showed that the least cardinal κ which is a strong compactness cardinal for $\mathbb{L}^2$ must be extendible [Mag71]. We get an interesting not quite analogous result in the virtual case.

Theorem 4.8.
(1) *Suppose κ is the least strong pseudo-compactness cardinal for $\mathbb{L}^2$. Then either κ is virtually extendible or there is a measurable cardinal below it.*
(2) *Suppose the* GCH *holds. Then either there is a measurable cardinal or every cardinal κ that is a strong pseudo-compactness cardinal for $\mathbb{L}^2$ is virtually extendible.*

Proof. We will only prove (1) because (2) follows from the argument. Suppose κ is the least strong pseudo-compactness cardinal for $\mathbb{L}^2$ and fix a strong limit cardinal $\alpha > \kappa$ of countable cofinality. Let the language τ and theory T be as in the proof of Theorem 4.5, with the only difference that we take the elementary diagram of $(V_\alpha, \in, c_x)_{x \in V_\alpha}$ in $\mathbb{L}^2$ as opposed to $\mathbb{L}^2_{\kappa,\kappa}$. By the compactness assumption, there is $\beta > \kappa$ such that V_β is a κ^+-pseudo-model of T. So there is a virtual elementary embedding $j : V_\alpha \to V_\beta$ with crit $j = \gamma \leq \kappa$ and $j(\kappa) > \alpha$ whose κ-sized pieces are in V. First, suppose that $2^\gamma \leq \kappa$. In this case, because we have κ-sized pieces of j, it follows that γ is measurable. So suppose that $2^\gamma > \kappa$. It follows, by elementarity, that $2^{j(\gamma)} > j(\kappa)$. But since we assumed that α is a strong limit and $j(\kappa) > \alpha$, it must be that $j(\gamma) > \alpha$ as well (note that $j(\gamma)$ is inaccessible and therefore cannot be α).

Assuming that there is no measurable cardinal below κ, by the pigeon-hole principle, there must be some critical point $\gamma \leq \kappa$, which works for unboundedly many ordinals α. So γ is virtually extendible, which means that it cannot be below κ by the leastness assumption. □

We do not know whether an analogous result holds with the existence of a κ^+-pseudo-model replaced with just a pseudo-model. More generally, we do not have any results where the apparently stronger assumption of the existence of a κ^+-pseudo-model is known to be necessary.

Question 4.9. *For some abstract logic $\mathcal{L}$ and language τ, is there an uncountable theory T in $\mathcal{L}(\tau)$ which has a pseudo-model, but not an ω_1-pseudo-model?*

Next, we give a compactness characterization of weakly virtually extendible cardinals via chain compactness for $\mathbb{L}^2_{\kappa,\kappa}$.

Theorem 4.10. *The following are equivalent for a cardinal κ.*
(1) *κ is weakly virtually extendible.*
(2) *Every $\mathbb{L}^2_{\kappa,\kappa}(\tau)$-theory T that can be written as an increasing union $T = \bigcup_{\eta<\kappa} T_\eta$ of satisfiable theories has a pseudo-model.*
(3) *Every $\mathbb{L}^2 \cup \mathbb{L}_{\kappa,\omega}(\tau)$-theory T that can be written as an increasing union $T = \bigcup_{\eta<\kappa} T_\eta$ of satisfiable theories has a pseudo-model.*
(4) *Same as (2) but with pseudo-model replaced by κ^+-pseudo-model.*

Proof. Assume the compactness principle in (3) and fix $\alpha > \kappa$. Let $\bar{\tau}$ be the language from the proof of Theorem 4.5 without the constants d_ξ and let $\bar{T}$ be the theory T without the statements about the d_ξ. We can filtrate $\bar{T}$ by including in $\bar{T}_\eta$ the statements $\{c_\xi < c < c_\kappa \mid \xi < \eta\}$. By hypothesis, some V_β is a pseudo-model for the theory $\bar{T}$, and this yields a virtual elementary embedding $j : V_\alpha \to V_\beta$ with crit $j = \kappa$.

Now suppose that κ is weakly virtually extendible and $T = \bigcup_{\eta<\kappa} T_\eta$ is a $\mathbb{L}^2_{\kappa,\kappa}(\tau)$-theory such that $\vec{T} = \langle T_\eta \mid \eta < \kappa \rangle$ is an increasing sequence of satisfiable theories. Let α be large enough so that V_α witnesses all this and let $j : V_\alpha \to V_\beta$ with crit $j = \kappa$ be a virtual elementary embedding. By elementarity, V_β satisfies that $j(\vec{T})$ is an increasing sequence of theories and for all $\eta < j(\kappa)$, $j(\vec{T})(\eta)$ is satisfiable, so in particular, V_β has a model $\mathcal{N} \models j(T)(\kappa) \supseteq T$. The model $\mathcal{N}$ is the required κ^+-pseudo-model for T.

□

Theorem 4.11. *Suppose κ is the least cardinal such that every $\mathbb{L}^2(\tau)$-theory T that can be written as an increasing union $T = \bigcup_{\eta<\kappa} T_\eta$ of satisfiable theories has a pseudo-model. Then κ is weakly virtually extendible.*

Proof. Suppose κ is least with the compactness property and fix $\alpha > \kappa$. Let the language $\bar{\tau}$ and theory $\bar{T}$ be as in the proof of Theorem 4.10, with the only difference that we take the elementary diagram of $(V_\alpha, \in, c_x)_{x\in V_\alpha}$ in $\mathbb{L}^2$ as opposed to $\mathbb{L}^2_{\kappa,\kappa}$. By the compactness assumption, there is a virtual elementary embedding $j : V_\alpha \to V_\beta$ with crit $j = \gamma \leq \kappa$. By the pingeon-hole principle, there is some $\gamma \leq \kappa$ that works for unboundedly many α. But then γ must be weakly virtually extendible and so $\gamma = \kappa$ by the leastness assumption. □

We will now give compactness characterizations of several other virtual large cardinal notions by reformulating the known compactness properties of the original large cardinals in terms of κ^+-pseudo models. At the same time, we will see that such a translation fails to hold for virtual Vopěnka's principle as a consequence of the splitting of virtual $C^{(n)}$-extendibility into the weak and strong forms.

It is a folklore result that measurable cardinals κ satisfy chain compactness for the logic $\mathbb{L}_{\kappa,\kappa}$ (see for instance, [CK12], Exercise 4.2.6) .

Theorem 4.12. *The following are equivalent.*

1. *κ is virtually measurable.*
2. *Every $\mathbb{L}_{\kappa,\kappa}(\tau)$-theory T that can be written as an increasing union $T = \bigcup_{\eta<\kappa} T_\eta$ of satisfiable theories has a pseudo-model.*
3. *Same as (2) but with pseudo-model replaced by κ^+-pseudo-model.*

Proof. Assume the compactness principle in (2) and fix $\alpha > \kappa$. Let τ be the language consisting of a binary relation $\in$ and constants $\{c_x \mid x \in V_\alpha\} \cup \{c\}$. Let T be the following $\mathbb{L}_{\kappa,\kappa}(\tau)$-theory:

$$\mathrm{ED}_{\mathbb{L}_{\kappa,\kappa}}(V_\alpha, \in, c_x)_{x\in V_\alpha} \cup \{c_\xi < c < c_\kappa \mid \xi < \kappa\},$$

where each constant c_x is interpreted as x. We can filtrate $\bar{T}$ by including in T_η the statements $\{c_\eta < c < c_\kappa \mid \eta < \alpha\}$. By assumption, T has a well-founded pseudo-model $\mathcal{M}$, and hence there is a virtual elementary embedding $j : V_\alpha \to \mathcal{M}$ with crit $j = \kappa$ (since every ordinal below κ is $\mathbb{L}_{\kappa,\kappa}$-definable). Note that $\mathcal{M}$ must be well-founded because well-foundedness is expressible in $\mathbb{L}_{\kappa,\kappa}$, and so this assertion must be contained in elementary diagram of V_α.

In the other direction, suppose that κ is virtually measurable and $T = \bigcup_{\eta<\kappa} T_\eta$ is a $\mathbb{L}_{\kappa,\kappa}(\tau)$-theory such that $\vec{T} = \langle T_\eta \mid \eta < \kappa\rangle$ is an increasing sequence of satisfiable theories. Let α be large enough so that V_α witnesses all this and let $j : V_\alpha \to \mathcal{M}$ with crit $j = \kappa$ be a virtual elementary embedding. By elementarity, $\mathcal{M}$ satisfies that $j(\vec{T})$ is an increasing sequence of theories and for all $\eta < j(\kappa)$, $j(\vec{T})(\eta)$ is satisfiable, so in particular, it has a model $\mathcal{N} \models j(T)(\kappa) \supseteq T$. The model $\mathcal{N}$ is the required κ^+-pseudo-model for T. □

Benda [Ben78] provided a compactness characterization of supercompact cardinals in terms of a variant of chain compactness together with omitting types, which has been extended by the first author to other cardinals [Bon20]. We will need to incorporate omitting types into our pseudo-models framework in order to give a reformulation for virtually supercompact cardinals.

Definition 4.13. We will say that a δ-pseudo-model $\mathcal{M}$ in a language τ^* *omits* an $\mathcal{L}(\tau)$-type $p(x)$ if there is a δ-forth system $\mathcal{F}$ from τ to τ^* such that, for all $f : \sigma \to \sigma^*$ from $\mathcal{F}$, $\mathcal{M}$ models $f^* \text{"} T^\sigma$ and omits $f^* \text{"} p^\sigma$.[5]

[5]The notation p^σ is analogous to T^σ, defined earlier, and denotes the subset of the type $p(x)$ referencing only σ.

We will say that an $\mathcal{L}(\tau)$-theory T is *increasingly filtered by* $\mathcal{P}_\kappa\delta$ if T is the union of a sequence $\vec{T} = \langle T_s \mid s \in \mathcal{P}_\kappa\delta\rangle$ such that whenever $s \subseteq s'$, then $T_s \subseteq T_{s'}$, and we will call $\vec{T}$, an *increasing filtration* of T.

Theorem 4.14. *The following are equivalent for a cardinal* κ.

(1) κ *is virtually supercompact* (*remarkable*).
(2) *For every* $\delta > \kappa$, *whenever* T *is an* $\mathbb{L}_{\kappa,\kappa}(\tau)$ *theory that is increasingly filtered by* $\mathcal{P}_\kappa\delta$ *and* $p^a(x)$ *for* $a \in A$ *is some set of increasingly filtered by* $\mathcal{P}_\kappa\delta$ *types such that every* T_s *has a model omitting all* $p^a_s(x)$, *then there is a pseudo-model of* T *omitting all* $p^a(x)$.
(3) *Same as* (2) *but with pseudo-model replaced by* κ^+*-pseudo-model.*

Proof. Suppose κ is virtually supercompact. Fix an $\mathbb{L}_{\kappa,\kappa}(\tau)$-theory $T = \bigcup_s T_s$ with an increasing filtration $\vec{T} = \langle T_s \mid s \in \mathcal{P}_\kappa\delta\rangle$ and $\mathbb{L}_{\kappa,\kappa}(\tau)$-types $p^a(x) = \bigcup_s p^a_s(x)$ indexed by $a \in A$ with increasing filtrations $\vec{p}^a = \langle p^a_s(x) \mid s \in \mathcal{P}_\kappa\delta\rangle$ satisfying the hypothesis of the theorem. Let λ be a large enough $\beth$-fixed point of cofinality κ^+ so that V_λ sees all this. Choose $\alpha > \lambda$ such that there is a transitive model $\mathcal{N}$ closed under λ-sequences and a virtual elementary embedding $j : V_\alpha \to \mathcal{N}$ with crit $j = \kappa$ and $j(\kappa) > \lambda$. Consider the restriction $j : V_\lambda \to j(V_\lambda)$. Observe that since V_λ is closed under κ-sequences by cofinality considerations, $j(V_\lambda)$ is closed under $j(\kappa)$-sequences in $\mathcal{N}$ by elementarity. Thus, $j(V_\lambda)$ is truly closed under λ-sequences, since $\mathcal{N}$ is. We will not use the embedding j, but move to a $\mathrm{Coll}(\omega, V_\lambda)$-extension $V[G]$ with an elementary embedding $h : V_\lambda \to \bar{\mathcal{N}} = j(V_\lambda)$ with crit $h = \kappa$ and $h(\kappa) > \lambda$.

Since the forcing $\mathrm{Coll}(\omega, \mathrm{V}_\lambda)$ has size $|V_\lambda| = \lambda$ because we chose λ to be a $\beth$-fixed point, we can cover $h\,"\,\lambda$ by a set $s^* \subseteq j(\lambda)$ of size $\lambda < h(\kappa)$ in V. Next, observe crucially that s^* is an element of $\bar{\mathcal{N}}$ by λ-closure. By elementarity, $\bar{\mathcal{N}}$ satisfies that $h(\vec{T})_{s^*}$ has a model $\mathcal{M}$ omitting all $h(p)^a_{s^*}$ for $a \in j(A)$. Since $h\,"\,\delta \subseteq s^*$, we have for every $s \in \mathcal{P}_\kappa\delta$ that $h(s) = h\,"\,s \subseteq s^*$. It follows that $h\,"\,T_s \subseteq h(\vec{T}_s) = h(\vec{T})_{h(s)} \subseteq h(\vec{T})_{s^*}$, and similarly $h\,"\,\vec{p}^a_s \subseteq h(\vec{p}^a_s) \subseteq h(\vec{p})^{h(a)}_{s^*}$. It follows that $\mathcal{M}$ is the required κ^+-pseudo-model.

In the other direction, suppose that we have the compactness assumption and fix a singular $\beth$-fixed point $\lambda > \kappa$ and $\alpha > \lambda$. Let τ be the language consisting of a binary relation $\in$ and constants $\{c_x \mid x \in V_\alpha\} \cup \{d_x \mid x \in V_\lambda\} \cup \{c\}$. Let T be the following $\mathbb{L}_{\kappa,\kappa}(\tau)$-theory:

$$\mathrm{ED}_{\mathbb{L}_{\kappa,\kappa}}(V_\alpha, \in, c_x)_{x\in V_\alpha} \cup \{c_\xi < c < c_\kappa \mid \xi < \kappa\} \cup \{d_b \in d_a \mid b \in a,\ a \in V_\lambda\} \cup \{d_\xi < d_\eta < c_\kappa \mid \xi < \eta < \lambda\},$$

and let $p^a(x)$, for $a \in V_\lambda$, be the following $\mathbb{L}_{\kappa,\kappa}(\tau)$-types:

$$\{x \in d_a\} \cup \{x \neq d_b \mid b \in a\}.$$

Now let us find a filtration for the theory T and the types $p^a(x)$ such that V_α can be made, by correctly interpreting the constants d_a, into a model of T_s omitting all $p^a_s(x)$. Fix a bijection $f : \lambda \to V_\lambda$. Given $s \in \mathcal{P}_\kappa(\lambda)$, let $X_s \prec V_\alpha$ be the elementary substructure of V_α generated by $f\,"\,s \subseteq V_\lambda$. Let $\mathrm{ED}_{\mathbb{L}_{\kappa,\kappa}}(V_\alpha, \in, c_x)_{x\in V_\alpha} \subseteq T_s$, but T_s is only allowed to mention sentences $c_\xi < c < c_\kappa$ if $c_\xi \in X_s$, it is only allowed to mention sentences with constants d_a if $a \in X_s$. Let $p^a_s(x) = \emptyset$ if $a \notin X_s$. Otherwise, suppose $a \in X_s$. In this case, let $p^a_s(x)$ mention only formulas $d_b \in d_a$ for $b \in X_s$. Let $\pi : X_s \to M$ be the Mostowski collapse. To make V_α into a model of T_s omitting $p^a_s(x)$, we will first interpret all c_x as x. For every $b \in X_s \cap V_\lambda$, let d_b be interpreted as $\pi(b)$. This ensures that there is no space to interpret x to satisfy $p^a_s(x)$.

By assumption, T has a well-founded pseudo-model $\mathcal{M}$ omitting all $p^a(x)$, and we can assume without loss that it is transitive. Let $\mathcal{F}$ be the associated forth system between the languages τ and τ^*, and let us force with $\mathcal{F}$ to add a generic injection $F : \tau \to \tau^*$. Define $F_*(\phi(x)) = f_*(\phi(x))$ for some/any $f \in \mathcal{F}$ such that $f \subseteq F$. In the forcing extension, we get an elementary embedding $j : V_\alpha \to \mathcal{M}$ with crit $j = \kappa$, $j(\kappa) > \lambda$ (we cannot have $j(\kappa) = \lambda$ since we chose λ to be singular) and the model $\mathcal{M}$ omits all types $p^{*a}(x) := F_* \,"\, p^a(x)$. It follows that $\mathcal{M}$ has a transitive subset isomorphic to V_λ and so $V_\lambda \subseteq \mathcal{M}$. □

Recall that Vopěnka's principle is the assertion that every proper class of first-order structures in the same language has two structures which elementarily embed. Analogously, *virtual Vopěnka's principle* (also known in the literature as *generic Vopěnka's principle*) asserts that every proper class of first-order structures in the same language has two structures which virtually embed.

Makowsky showed that Vopěnka's principle is equivalent to the assertion that every abstract logic has a strong compactness cardinal [Mak85]. Bagaria showed that Vopěnka's principle is equivalent to the assertion that for every $n < \omega$, there is a $C^{(n)}$-extendible cardinal [Bag12]. The third author and Hamkins showed in [GH19] that virtual Vopěnka's principle is equivalent to the assertion that for every $n < \omega$, there is a proper class of weakly virtually $C^{(n)}$-extendible cardinals, but at the same same time it is consistent that virtual Vopěnka's principle holds and yet there are no virtually $C^{(n)}$-extendible cardinals.

We will reprove Moskowski's theorem and show that one direction generalizes to the case of virtually $C^{(n)}$-extendible cardinals, but the other direction fails to generalize because of the split in the virtual case into the weak and strong forms of $C^{(n)}$-extendibility.

Theorem 4.15 (Makowsky, Bagaria)**.** *For every $n < \omega$, there is a $C^{(n)}$-extendible cardinal if and only if every abstract logic has a strong compactness cardinal.*

We will need the following easy fact about $C^{(n)}$-extendible cardinals.

Proposition 4.16 ([Bag12])**.** *Suppose that for every $n < \omega$, there is a $C^{(n)}$-extendible cardinal. Then for every $n < \omega$, there is a proper class of $C^{(n)}$-extendible cardinals.*

Proof. Suppose towards a contradiction that for some fixed n, the $C^{(n)}$-extendible cardinals are bounded by δ. Let κ be a $C^{(m)}$-extendible cardinal for some $m \gg n$. Obviously $\kappa < \delta$. Fix $\alpha > \delta$ in $C^{(m)}$ and take an elementary embedding $j : V_\alpha \to V_\beta$ with crit $j = \kappa$, $j(\kappa) > \alpha$ and $\beta \in C^{(m)}$. By elementarity, V_α sees that κ is $C^{(n)}$-extendible. It follows that V_β thinks that $j(\kappa)$ is $C^{(n)}$-extendible and it must be correct about this by the level of elemenarity. But $j(\kappa) > \alpha > \delta$, which is the desired contradiction. □

Clearly the argument would work identically for virtually $C^{(n)}$-extendible cardinals as well, but not for weakly virtually $C^{(n)}$-extendible cardinals.

Proof of Theorem 4.15. Fix an abstract logic $\mathcal{L}$ with occurrence number $o(\mathcal{L}) = \delta$. The logic $\mathcal{L}$ and satisfaction $\models_\mathcal{L}$ are defined by some Σ_n-formulas with a parameter a of rank δ_a. Clearly there are only set-many languages of size less than δ modulo a renaming. So let us fix an ordinal $\bar{\delta}$ such that $V_{\bar{\delta}}$ is closed under $\mathcal{L}$ and any language τ of size less than δ has a renaming to a language $\tau' \in V_{\bar{\delta}}$. Let $\kappa > \bar{\delta}, \delta_a$ be $C^{(n)}$-extendible. We will argue that κ is a strong compactness cardinal for $\mathcal{L}$. Let T be a $<\kappa$-satisfiable $\mathcal{L}(\tau)$-theory. Choose an elementary embedding $j : V_\alpha \to V_\beta$ with $\alpha, \beta \in C^{(n)}$, crit $j = \kappa$ and $j(\kappa) > \alpha$ such that V_α sees that T is $<\kappa$-satisfiable. Note that both V_α and V_β are correct about $\mathcal{L}$ and $\models_\mathcal{L}$ because they are Σ_n-elementary in V.

By elementarity V_β satisfies that $j \,"\, T$ has a model M in the logic $\mathcal{L}(j \,"\, \tau)$, and since $V_\beta \prec_{\Sigma_n} V$, it must be correct about this. Let $f : \tau \to j \,"\, \tau$ be the renaming taking elements of τ to their images under j. Under the renaming M is a τ-structure, and it suffices to show that $f_*(\phi) = j(\phi)$ for every $\phi \in T$. Fix $\phi \in T$ and let τ_ϕ be the $<\delta$-sized subset of τ used in ϕ. Let $f^\phi = f \restriction \tau_\phi$. By our assumptions, there is a language $\sigma \in V_\kappa$ with $\mathcal{L}(\sigma) \in V_\kappa$ and a renaming $g : \sigma \to \tau_\phi$. Let $g_*(\bar{\phi}) = \phi$. By elementarity, $j(g) : \sigma \to j(\tau_\phi)$ is a renaming from σ to $j(\tau_\phi) = j \,"\, \tau_\phi$. It is easy to check that $j(g) = f^\phi \circ g$. Now, by elementarity, we have $(j(g))_*(\bar{\phi}) = j(g_*)(\bar{\phi}) = j(g_*(\bar{\phi})) = j(\phi)$. By our observations from Section 2 we have, $j(g)_* = (f^\phi \circ g)_* = f^\phi_* \circ g_*$. Thus, we have $f^\phi_*(\phi) = f^\phi_*(g_*(\bar{\phi})) = j(g)_*(\bar{\phi}) = j(\phi)$. Again, by our observations from Section 2 about restrictions of renamings, we have $f_*(\phi) = j(\phi)$, but in fact it already suffices to know that $f^\phi_*(\phi) = j(\phi)$ since V_β is correct about satisfaction for $\mathcal{M}$.

In the other direction, we will argue that if there is a strong compactness cardinal for the infinitary sort logic $\mathbb{L}^{s,\Sigma_n}$, then there must be a $C^{(n)}$-extendible cardinal. So suppose that γ is a strong compactness cardinal for $\mathbb{L}^{s,\Sigma_n}$. Fix $\alpha > \gamma$ in $C^{(n)}$. We can write the usual theory whose model gives an elementary embedding $j : V_\alpha \to V_\beta$ with crit $j = \kappa_\alpha \leq \gamma$, and using sort logic, by Proposition 2.2, we can express that $\beta \in C^{(n)}$. Since there are only boundedly many $\kappa \leq \gamma$, by the pigeon-hole principle we can choose some κ_{α^*} which works for unboundedly many α, and this $\kappa = \kappa_{\alpha^*}$ must be $C^{(n)}$-extendible. Note that because we are not in the virtual case we do not need to show additionally that $j(\kappa) > \alpha$. □

The proof above gives us the following results for the virtual case.

Theorem 4.17. *If for every $n < \omega$, there is a virtually $C^{(n)}$-extendible cardinal, then every abstract logic has a strong pseudo-compactness cardinal.*

Proof. Following the proof of Theorem 4.15 for the forward direction, it suffices to observe that even though j " T may not be in V_β, we can cover it by a theory $T^* \in V_\beta$ of size less than $j(\kappa)$. □

Indeed, as in previous arguments, we do not need κ^+-pseudo models in the argument; having pseudo-models suffices to obtain the desired virtual elementary embeddings.

Theorem 4.18. *If γ is a strong pseudo-compactness cardinal for the sort logic $\mathbb{L}^{s,\Sigma_n}$ and there are no measurable cardinals below γ, then there is a virtually $C^{(n)}$-extendible cardinal.*

Proof. This follows directly from the proof of Theorem 4.15 combined with the proof of Theorem 4.8 (1). Note that in this argument we do make a use of κ^+-pseudo models. □

Corollary 4.19. *Suppose there are no measurable cardinals. Then the assertion that for every abstract logic there is a strong pseudo-compactness cardinal is equivalent to the assertion that for every $n < \omega$, there is a virtually $C^{(n)}$-extendible cardinal.*

Corollary 4.20. *Virtual Vopěnka's principle is not equivalent to the assertion that every abstract logic has a strong pseudo-compactness cardinal.*

Proof. The model constructed in [GH19] in which virtual Vopěnka's principle holds but there are no virtually $C^{(n)}$-extendible cardinals does not have any measurable cardinals. Therefore in this model it cannot be the case that every abstract logic has a strong pseudo-compactness cardinal. □

Question 4.21. *Is the assertion that for every abstract logic there is a strong pseudo-compactness cardinal equivalent to the assertion that for every $n < \omega$, there is a virtually $C^{(n)}$-extendible cardinal?*

The first author showed in [Bon20] that a cardinal κ is $C^{(n)}$-extendible if and only if κ is a strong compactness cardinal for $\mathbb{L}^{s,\Sigma_n}_{\kappa,\omega}$. This result can be reformulated in the pseudo-compactness framework.

Theorem 4.22. *A cardinal κ is virtually $C^{(n)}$-extendible if and only if κ is a strong pseudo-compactness cardinal for $\mathbb{L}^{s,\Sigma_n}_{\kappa,\omega}$.*

Proof. The forward direction follows from the proof of Theorem 4.15 and the backward direction follows because in the infinitary logic $\mathbb{L}_{\kappa,\omega}$ we can express that the critical point of our embedding is exactly κ. □

We do, however, get a chain compactness characterization of virtual Vopěnka's principle.

Theorem 4.23. *The following are equivalent.*

(1) *Virtual Vopěnka's principle.*
(2) *For every abstract logic $\mathcal{L}$, there is a cardinal κ such that every theory $T = \bigcup_{\eta<\kappa} T_\eta$ in $\mathcal{L}$ that is an increasing union of satisfiable theories has a pseudo-model.*
(3) *Same as* (2) *but with pseudo-model replaced by κ^+-pseudo model.*

Proof. Suppose that virtual Vopěnka's principle holds and so for every $n < \omega$, we have a proper class of weakly virtually $C^{(n)}$-extendible cardinals. The compactness principle in (3) then follows directly by the proof of Theorem 4.15.

Assume the compactness principle in (2). We need to show that given a fixed ordinal β, there is $\gamma > \beta$ which is weakly virtually $C^{(n)}$-extendible. We use the logic $\mathbb{L}^{s,\Sigma_n}_{\delta,\omega}$ for some $\delta > \beta$ to ensure that the critical point of the virtual embedding we obtain is above β and argue as in the proof of Theorem 4.15. □

The above analysis was carried out in ZFC permitting only the definable classes. However, all of it straightforwardly generalizes to the second-order context of GBC by replacing the notion of (virtually) $C^{(n)}$-extendible with the notion of (virtually) A-extendible cardinals for a class A. A cardinal κ is A-*extendible*, for a class A, if for every $\lambda > \kappa$, there is $\theta > \kappa$ such that there is an elementary embedding $j : (V_\lambda, \in, A \cap V_\lambda) \to (V_\theta, \in, A \cap V_\theta)$ with crit $j = \kappa$ and $j(\kappa) > \lambda$. The virtual versions are defined analogously (for more details, see [GH19]). Thus, we get for example, that in GBC, Vopěnka's principle holds if and only if for every class A, there is an A-extendible cardinal, etc.

In Definition 4.3, we introduced the notion of a forth system between languages. Below, we define a forth system between models of the same language (Definition 4.24.(3)). First, we explain the motivation and their connection to back-and-forth systems and Ehrenfeucht-Fraïssé games.

It is a classical result that two structures M and N from V are isomorphic in a forcing extension if and only if the *good* player has a winning strategy in the ω-length Ehrenfeucht-Fraïssé game $\Game_F(M, N)$. The game starts with the *bad* player choosing to play an element $a_0 \in M$ or $b_0 \in N$ and the *good* player has to respond with a move $b_0 \in N$ or $a_0 \in M$ respectively so that the map $f = \{\langle a_0, b_0\rangle\}$ is a finite partial isomorphism between M and N. Given a finite partial isomorphism f between M and N that results from an initial play of $\Game_F(M, N)$, at the next step of the game, the bad player again chooses to play an element out of either M or N and the good player has to respond to extend f to a partial isomoprhism. It is easy to see that the good player having a winning strategy in $\Game_F(M, N)$ is equivalent to the existence of a back-and-forth system. A *back-and-forth system* between two models M and N in the same language τ is a collection $\mathcal{P}$ of finite partial isomorphisms between M and N such that whenever $a \in M$ and $f \in \mathcal{P}$, then there is $g \in \mathcal{P}$ extending f with $a \in \text{dom}(g)$, and conversely whenever $b \in N$, then there is a $g' \in \mathcal{P}$ extending f with $b \in \text{ran}(g')$. Clearly, forcing with a back-and-forth system gives a virtual isomorphism and, in the other direction, a virtual isomorphism suffices to obtain a back-and-forth system via the forcing relation. The existence of a back-and-forth system is also equivalent to the assertion that the models M and N are elementary equivalent in the pseudo-abstract logic $\mathbb{L}_{\text{Ord},\omega}$.

Schindler showed that two structures from V have a virtual embedding if and only if the *good* player has a winning strategy in the ω-length modified Ehrenfeucht-Fraïssé game $\Game_v(M, N)$

[BGS17]. The game starts with the *bad* player playing an element $a_0 \in M$ and the *good* player has to respond with a move $b_0 \in N$ so that the map $f = \{\langle a_0, b_0\rangle\}$ is a finite partial isomorphism between M and N. The good player wins if at each step she can maintain a finite partial isomorphism between M and N. For the game $\Game_v(M, N)$, the corresponding notion to a back-and-forth system is a *forth-system* $\mathcal{P}$, which is a collection of finite partial isomorphisms f between M and N with only the forth extension property. The corresponding pseudo-abstract logic will be a subclass of $\mathbb{L}_{\mathrm{Ord},\omega}$, which we will call *virtual logic*. Indeed, we can define the notion of a virtualization for any abstract logic.

Definition 4.24. Fix a logic $\mathcal{L}$.

(1) Given two structures M and N in the same language, the game $\Game_{v_{\mathcal{L}}}(M, N)$ is defined exactly as the game $\Game_v(M, N)$ with satisfaction given by the logic $\mathcal{L}$ in place of first-order logic.
(2) The logic $\mathcal{L}^v(\tau)$, for a language τ, consists of the formulas given by the following closure rules.
  (a) Every formula of $\mathcal{L}(\tau)$ is a formula of $\mathcal{L}^v(\tau)$,
  (b) If $\phi(x) \in \mathcal{L}^v(\tau)$ (with possibly other free variables), then $\exists x\phi(x) \in \mathcal{L}^v(\tau)$.
  (c) If $\{\phi_i \mid i \in I\}$ is a collection of formulas from $\mathcal{L}^v(\tau)$ jointly in finitely many variables, then so is $\bigwedge_{i\in I} \phi_i$.
(3) An $\mathcal{L}$-forth system $\mathcal{P}$ from a τ-structure M to a τ-structure N is a collection of $\mathcal{L}$-elementary embeddings with the following properties.
  (a) $\emptyset \in \mathcal{P}$.
  (b) If $f \in \mathcal{P}$ and $a \in M$, then there is $g \supseteq f$ in $\mathcal{P}$ such that $a \in \mathrm{dom}\, g$.

Recall, that technically our abstract logic does not support formulas with free variables; instead something like "$\phi(x) \in \mathcal{L}^v(\tau)$" with x free really means "$\phi(c) \in \mathcal{L}^v(\tau \cup \{c\})$", where c is a new constant symbol. However, we use the notation of free variables for better readability.

Theorem 4.25. *For structures M and N in a common language τ, the following are equivalent.*

(1) *The good player has a winning strategy in $\Game_{\mathcal{L}_v}(M, N)$.*
(2) *$N \models Th_{\mathcal{L}^v}(M)$, the collection of all sentences in $\mathcal{L}^v(\tau)$ that M satisfies.*
(3) *There is an $\mathcal{L}$-forth system from M to N.*
(4) *There is a virtual $\mathcal{L}$-elementary embedding $f : M \to N$.*

The new parts of Theorem 4.25 are proved by varying classical proofs involving Ehrenfeucht-Fraïssé games. If M and N satisfy any one of the equivalent conditions of Theorem 4.25, then we shall say that M is a *virtual submodel* of N. This lets us give Löwenheim-Skolem-Tarski style characterizations of virtual large cardinals. For example, a virtualization of classical arguments about supercompact cardinals from [Mag71] show:

Theorem 4.26. *Suppose κ is virtually supercompact and Φ is an $\mathbb{L}^2$-sentence. If M is an $\tau(\Phi)$-structure such that $M \models \Phi$, then M has a virtual submodel of size less than κ that also models Φ.*

Proof. Magidor showed that a cardinal κ is supercompact if and only if for every $\delta > \kappa$, there is $\bar{\delta} < \kappa$ and an elementary embedding $j : V_{\bar{\delta}} \to V_\delta$ with crit $j = \gamma$ and $j(\gamma) = \kappa$. The obvious virtualized reformulation holds for virtually supercompact cardinals as well, and this immediately implies the desired reflection. □

Magidor also showed that the least cardinal with this reflection property is supercompact. We do not expect the result to generalize to the virtual context for the usual reasons.

Question 4.27. *Is the least cardinal κ with the property that for every $\mathbb{L}^2$-sentence Φ, every $\tau(\Phi)$-structure $M \models \Phi$ has a virtual submodel of size less than κ that also satisfies Φ supercompact?*

Finally, using virtual logic we can give an alternative compactness-style characterization of weakly virtually extendible cardinals asserting that whenever $T \subseteq \mathbb{L}^2_{\kappa,\kappa}(\tau^*)$ has size κ and all of its $<\kappa$-sized pieces can be realized in expansions of a single model $\mathcal{M}$ of some sub-language $\tau \subseteq \tau^*$, then T has a model $\mathcal{N}$ also satisfying the virtual theory of $\mathcal{M}$.

Theorem 4.28. *A cardinal κ is weakly virtually extendible if and only if for every theory $T \subseteq \mathbb{L}^2_{\kappa,\kappa}(\tau^*)$ of size κ, if $\mathcal{M}$ is a τ-structure, with $\tau \subseteq \tau^*$, such that for every $T^* \subseteq T$ of size less than κ, there is an expansion of $\mathcal{M}$ to a τ^*-structure making $\mathcal{M} \models T^*$, then there is a τ^*-structure $\mathcal{N}$ satisfying $Th_{(\mathbb{L}^2_{\kappa,\omega})^v(\tau)}(\mathcal{M})$ together with T.*

Proof. Suppose the compactness characterization holds and fix $\alpha > \kappa$. Let τ^* be the language consisting of a binary relation $\in$ and constants $\{c_\xi \mid \xi \leq \kappa\} \cup \{c\}$ and let τ be τ^* without c. Let T be the following $\mathbb{L}^2_{\kappa,\omega}(\tau)$-theory.

$$\mathrm{ED}_{\mathbb{L}_{\kappa,\omega}}(V_\alpha, \in, c_\xi)_{\xi \leq \kappa} \cup \{c_\xi < c < c_\kappa \mid \xi < \kappa\} \cup \{\Phi\}.$$

It is easy to see that if T^* is any sub-theory of T of size less than κ, then we can interpret c by a large enough $\lambda < \kappa$. So by hypothesis, there is a virtual $\mathbb{L}_{\kappa,\omega}(\tau)$-embedding $j : (V_\alpha, \in, c_\xi)_{\xi \leq \kappa} \to (V_\beta, \in, c_\xi)_{\xi \leq \kappa}$. The only question is how each c_ξ is interpreted in V_β. Since V_β satisfies the $\mathbb{L}_{\kappa,\omega}$-theory of $(V_\alpha, \in)$, it knows that for $\xi < \kappa$, that c_ξ must be interpreted as ξ. Also, since V_β is a model T, it knows that $c_\kappa > \kappa$. Therefore κ is the critical point of j.

In the other direction, suppose that κ is weakly virtually extendible. Fix a κ-sized $\mathbb{L}^2_{\kappa,\kappa}(\tau^*)$-theory T and let $\mathcal{M}$ be a τ-structure as in the hypothesis. Since τ^* has size κ, we can assume that both τ^* and T are subsets of V_κ. Fix V_α large enough that it contains the model $\mathcal{M}$ and fix a virtual elementary embedding $j : V_\alpha \to V_\beta$ with critical point κ. So V_β satisfies that for every $T^* \subseteq j(T)$ of size less than $j(\kappa)$ there is an expansion of $j(\mathcal{M})$ to a $j(\tau^*)$-structure such that $j(\mathcal{M}) \models T^*$. In particular, there is an expansion of $j(\mathcal{M})$ to a τ^*-structure satisfying $j\,''\,T = T$. It is not difficult to see that $j : \mathcal{M} \to j(\mathcal{M})$ is a virtual $\mathbb{L}^2_{\kappa,\kappa}$-elementary embedding. Thus, $j(\mathcal{M})$ is a model of T that satisfies the virtual theory $Th_{(\mathbb{L}^2_{\kappa,\omega})^v(\tau)}(\mathcal{M})$. □

5. Vopěnka for weak compactness

Recall once again that Vopěnka's principle holds if and only if every abstract logic has a strong compactness cardinal. The proof can be pushed further to show that every abstract logic has a stationary class of strong compactness cardinals because Vopěnka's principle implies that for every $n < \omega$, the $C^{(n)}$-extendible cardinals are stationary in Ord [Bag12]. Here, we will show that the principle Ord is subtle holds if and only if every abstract logic has a stationary class of weak compactness cardinals.

This section deals with the set-theoretic principle Ord *is subtle*, which is an assertion about properties of classes. So we will start by explaining some frameworks that we can use for our formal setting. In first-order set theory, classes are collections definable (with parameters) over the model. Second-order set theories are formalized in a two-sorted logic with objects for both sets and classes. A standard axiomatization of second-order set theory is Gödel-Bernays set theory GBC. The theory GBC consists of ZFC together with the following axioms for classes: extensionality, replacement (a class function restricted to a set is a set), global choice axiom asserting that is there well-ordering of the universe, and the comprehension scheme for first-order formulas asserting that every first-order formula (with set/class parameters) defines a class. A first-order model of ZFC together with its definable collections satisfies all the axioms of GBC except possibly global choice (it is consistent with ZFC that there is no definable global well-ordering). Since our theorems appear to require global choice, we have to work either in GBC or over a first-order ZFC-model with a definable global well-order. Alternatively, we can assume that we are working with a rank-initial segment V_κ, for an inaccessible cardinal κ and our classes are precisely $V_{\kappa+1}$. The second-order model consisting of the sets V_κ and classes $V_{\kappa+1}$ satisfies GBC with full comprehension for all second-order assertions, which together comprise the much stronger second-order set theory Kelley-Morse.

The proof of the main theorem of this section relies on the following proposition asserting that the ordinals $\alpha < \beta$ given by subtleness can be assumed to be regular cardinals.

Proposition 5.1. *If* Ord *is subtle, then for any class sequence $\langle A_\alpha \subseteq \alpha \mid \alpha \in C\rangle$ indexed by a class club C, there are regular cardinals $\alpha < \beta$ in C such that $A_\alpha = A_\beta \cap \alpha$.*

The proof is an easy modification of the proof of [AHKZ77, Theorem 3.6.3].

Proof. By shrinking C if necessary, we can suppose that C consists only of cardinals. Define an auxiliary sequence $\langle B_\alpha \subseteq \alpha \mid \alpha \in C\rangle$ by:

(1) If α is singular with cofinality μ, then we fix a cofinal sequence $s^{(\alpha)}$ of length μ in α and define
$$B_\alpha = \{(0, \mu), (1, s^{(\alpha)}_\xi) \mid \xi < \mu\}.$$
(2) If α is regular, then
$$B_\alpha = \{(2, \xi) \mid \xi \in A_\alpha\}.$$

Note that we make use here of global choice to pick the cofinal sequences $s^{(\alpha)}$. By applying subtlety to the club C and the sequence $\langle B_\alpha \mid \alpha \in C\rangle$, there are $\alpha < \beta$ from C such that $B_\alpha = B_\beta \cap \alpha$. If β is regular, then α must be regular because B_α consists of pairs of the form $(2, \xi)$. So suppose β is singular. Then α must also be singular and $(0, \mathrm{cf}\,\beta) \in B_\alpha$. So $\mathrm{cf}\,\alpha = \mathrm{cf}\,\beta$. It follows that $\langle \nu \mid (1, \nu) \in B_\alpha\rangle$ is a $\mathrm{cf}\,\beta$-sequence cofinal in α that is end-extended to a $\mathrm{cf}\,\beta$-sequence cofinal in β, but this is impossible. So α and β are both regular and also, $A_\alpha = A_\beta \cap \alpha$ by the definition of the B_α-sequence. □

Theorem 5.2. *The following are equivalent over* GBC.

(1) Ord *is subtle.*
(2) *For every abstract logic $\mathcal{L}$ the collection of all cardinals κ that are weakly compact for $\mathcal{L}$ is a stationary class.*

Proof. First, suppose that Ord is subtle and fix an abstract logic $\mathcal{L}$ with occurrence number $o(\mathcal{L}) = \lambda$. We will assume towards a contradiction that there is a class club $C_\mathcal{L}$ of cardinals such that no element of $C_\mathcal{L}$ is weakly compact for $\mathcal{L}$. Using global choice, we can choose, for every $\alpha \in C_\mathcal{L}$, a language τ_α and an $\mathcal{L}(\tau_\alpha)$-theory T_α of size α that is $<\alpha$-satisfiable, but not satisfiable.

We can assume that the first element of C is very high above λ, so that each τ_α has size at most α. But then we can further assume that each τ_α is the 'maximal' language of size α, having relations R^n_ξ of arity n, functions F^n_ξ of arity n, and constants c_ξ, for $n < \omega$ and $\xi < \alpha$. Thus, in particular, we have that τ_α extends τ_β for $\beta < \alpha$, and that the sequence of the τ_α is continuous. This, in turn, gives that the sequence of the $\mathcal{L}(\tau_\alpha)$ is increasing and continuous at cardinals of cofinality $\geq \lambda$. Let's code the sentences of $\mathcal{L}$ by ordinals and from now on associate elements of $\mathcal{L}$ with their ordinal codes. Now observe that cardinals α such that $\mathcal{L}(\tau_\alpha) \subseteq \alpha$ form an unbounded class D that is closed under sequences of cofinality $\geq \lambda$. It is easy to find α such that $\bigcup_{\xi<\alpha} \mathcal{L}(\tau_\xi) \subseteq \alpha$ and if $\mathrm{cf}\,\alpha \geq \lambda$, then, using $o(\mathcal{L}) = \lambda$ and the continuity of the τ_α sequence, $\bigcup_{\xi<\alpha} \mathcal{L}(\tau_\xi) = \mathcal{L}(\tau_\alpha)$. Let $\bar{D}$ be D together with all its limit points of small cofinality, which is easily seen to be a club. Finally, apply Proposition 5.1 to the club $C \cap \bar{D}$ and the sequence $\langle T'_\alpha \mid \alpha \in C \cap \bar{D}\rangle$, where $T'_\alpha = T_\alpha$ for regular α and $T'_\alpha = T_\alpha \cap \alpha$ otherwise, to obtain regular

$\alpha < \beta$ such that $T_\alpha = T_\beta \cap \alpha$. But this is impossible because we assumed both that T_β is $<\beta$-satisfiable and that T_α is not satisfiable.

In the other direction, suppose that every abstract logic $\mathcal{L}$ has a stationary class of weakly compact cardinals. Fix a class club C of ordinals and a class sequence $A = \langle A_\alpha \subseteq \alpha \mid \alpha \in C\rangle$. We can thin C out to assume that $|V_\alpha| = \alpha$ for every $\alpha \in C$ and that every successor in C has uncountable cofinality. For each $\alpha \in C$, define the associated structure

$$\mathbb{A}_\alpha = \langle V_\alpha, \in, D_\alpha, C \cap \alpha, A_\alpha\rangle,$$

where $D_\alpha = \{(\gamma, \delta) \in \alpha \times \alpha \mid \delta \in A_\gamma\}$ codes all smaller A_γ.

Let τ be the language consisting of one binary relation and three unary relations, so that in particular $\mathbb{A}_\alpha$ are τ-structures. We define a logic $\mathbb{L}^{C,A}$ which extends first-order logic by adding a sentence Ψ such that a τ-structure $\langle M, E, J, K, L\rangle \vDash \Psi$ if and only if there is some $\beta \in C$ such that $\langle M, E, J, K, L\rangle$ is isomorphic to $\mathbb{A}_\beta$. Let τ_α be the language τ extended by adding constants $\{c_x \mid x \in V_\alpha\} \cup \{c\}$. For each $\alpha \in C$, define the $\mathbb{L}^{C,A}(\tau_\alpha)$-theory

$$T_\alpha = \{\Psi\} \cup \mathrm{ED}(\mathbb{A}_\alpha, c_x)_{x \in V_\alpha} \cup \{c \neq c_\beta \mid \beta < \alpha\},$$

where each constant c_x is interpreted as x. The theory T_α has size α and is $<\alpha$-satisfiable (in an expansion of $\mathbb{A}_\alpha$).

By assumption, there is some limit point κ^* of C that is weakly compact for $\mathbb{L}^{C,A}$. So, in particular, T_{κ^*} has some model $\mathcal{M}$. By definition of Ψ, $\mathcal{M} \restriction \tau$ is isomorphic to some $\mathbb{A}_\delta$. Since $\mathbb{A}_\delta$ models $\mathrm{ED}(\mathbb{A}_{\kappa^*}, c_x)_{x \in V_{\kappa^*}}$, this induces an elementary embedding $\pi : \mathbb{A}_{\kappa^*} \to \mathbb{A}_\delta$. The interpretation of c witnesses that π is not onto the ordinals of $\mathbb{A}_\delta$, so let $\eta < \delta$ be the minimum ordinal not hit. Note that $\eta \leq \kappa^*$.

If $\eta = \kappa^*$, then π is the identity and $\mathbb{A}_{\kappa^*} \prec \mathbb{A}_\delta$. Since the sequence A is coded in the language, this means $A_\delta \cap \kappa^* = A_{\kappa^*}$.

If $\eta < \kappa^*$, then crit $\pi = \eta$. We claim that $\eta \in C$. If not, then we can find $\nu, \mu \in C$ so ν is the largest member of C below η and μ is the smallest member of C above it. Since κ^* is a limit member of C, we have

$$\nu < \eta < \mu < \kappa^*.$$

We know that μ is definable from ν (being its successor in C), so π fixes μ since it fixes ν. It follows that $\pi \restriction V_\mu$ is a nontrivial elementary embedding of V_μ into itself, which violates Kunen's theorem (that there is no nontrivial elementary embedding from $V_{\alpha+2}$ into $V_{\alpha+2}$ for any ordinal α), since by our assumption on successor elements of C, μ has uncountable cofinality. Thus, $\eta \in C$, which means $\pi(\eta) \in C$. Applying the elementary embedding to the predicate D_{κ^*}, we have that $D_\eta = D_{\pi(\eta)} \cap \eta$, as desired. □

Question 5.3. *Does Theorem 5.2 hold in the absence of the global choice axiom?*

References

[AHKZ77] Fred Abramson, Leo Harrington, E. Kleinberg, and W. Zwicker, *Flipping properties: a unifying thread in the theory of large cardinals*, Annals of Mathematical Logic **12** (1977), no. 1, 25–58.

[Apt12] Arthur Apter, *Some applications of Sargsyan's equiconsistency method*, Fundamenta Mathematicae **216** (2012), 207–222.

[AS07] Arthur Apter and Grigor Sargsyan, *A reduction in consistency strength for universal indestructibility*, Bulletin of the Polish Academy of Sciences Mathematics **55** (2007), no. 1, 1–6.

[Bag12] Joan Bagaria, *$C^{(n)}$-cardinals*, Arch. Math. Logic **51** (2012), no. 3-4, 213–240. MR 2899689

[Ben78] Miroslav Benda, *Compactness for omitting of types*, Annals of Mathematical Logic **4** (1978), 39–56.

[BGS17] Joan Bagaria, Victoria Gitman, and Ralf Schindler, *Generic Vopěnka's Principle, remarkable cardinals, and the weak Proper Forcing Axiom*, Archive for Mathematical Logic **56** (2017), 1–20.

[Bon14] Will Boney, *Tameness from large cardinal axioms*, Journal of Symbolic Logic **79** (2014), no. 4, 1092–1119.

[Bon20] ______, *Model-theoretic characterizations of large cardinals*, Israel Journal of Mathematics **236** (2020), 133–181.

[BU17] Will Boney and Spencer Unger, *Large cardinal axioms from tameness in abstract elementary classes*, Proceedings of the American Mathematical Society **145** (2017), no. 10, 4517–4532.

[CK12] C. C. Chang and H. Jerome Keisler, *Model theory*, 3rd ed., Dover Publications, 2012.

[Dim19] Stamatis Dimopoulos, *Woodin for strong compactness cardinals*, Journal of Symbolic Logic **84** (2019), no. 1, 301–319.

[GH19] Victoria Gitman and Joel David Hamkins, *A model of the generic Vopěnka principle in which the ordinals are not Mahlo*, Arch. Math. Logic **58** (2019), no. 1-2, 245–265. MR 3902814

[GS18] Victoria Gitman and Ralf Schindler, *Virtual large cardinals*, Ann. Pure Appl. Logic **169** (2018), no. 12, 1317–1334. MR 3860539

[HJ17] Joel David Hamkins and Thomas Johnstone, *Strongly uplifting cardinals and the boldface resurrection axioms*, Archive for Mathematical Logic **56** (2017), no. 7, 1115–1133.

[Mag71] Menachem Magidor, *On the role of supercompact and extendible cardinals in logic*, Israel Journal of Mathematics **10** (1971), 147–157.

[Mak85] Johann Makowsky, *Vopenka's principle and compact logics*, Journal of Symbolic Logic **50** (1985), 42–48.

[NW19] Dan Nielsen and Philip Welch, *Games and Ramsey-like cardinals*, Journal of Symbolic Logic **84** (2019), 408–437.

[Per15] Norman Lewis Perlmutter, *The large cardinals between supercompact and almost-huge*, Archive for Mathematical Logic **54** (2015), no. 3-4, 257–289.

[Sch00] Ralf-Dieter Schindler, *Proper forcing and remarkable cardinals*, Bull. Symbolic Logic **6** (2000), no. 2, 176–184. MR 1765054 (2001h:03096)

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

[Tar62] Alfred Tarski, *Some problems and results relevant to the foundations of set theory*, Logic, Methodology and Philosophy of Science, Proceedings of the 1960 International Congress, Stanford University Press, 1962.

[Vää14] Jouko Väänänen, *Sort logic and foundations of mathematics*, Infinity and truth, Lect. Notes Ser. Inst. Math. Sci. Natl. Univ. Singap., vol. 25, World Sci. Publ., Hackensack, NJ, 2014, pp. 171–186. MR 3205075

(W. Boney) Mathematics Department, Texas State University, San Marcos, TX, USA
E-mail address: `wb1011@txstate.edu`

(S. Dimopoulos) Department of Mathematics and Statistics, University of Cyprus, Nicosia, Cyprus
E-mail address: `stamatiosdimopoulos@gmail.com`

(V. Gitman) Mathematics Department, CUNY Graduate Center, NY, USA
E-mail address: `vgitman@gmail.com`

(M. Magidor) Institute of Mathematics, Hebrew university of Jerusalem, Jerusalem 91904, Israel
E-mail address: `mensara@savion.huji.ac.il`