BETWEEN RAMSEY AND MEASURABLE CARDINALS

VICTORIA GITMAN AND PHILIPP SCHLICHT

Abstract. We study several intertwined hierarchies between κ-Ramsey cardinals and measurable cardinals to illuminate the structure of the large cardinal hierarchy in this region. In particular, we study baby versions of measurability introduced by Bovykin and McKenzie and some variants by locating these notions in the large cardinal hierarchy and providing characterisations via filter games. As an application, we determine the theory of the universe up to a measurable cardinal.

Contents

1. Introduction 1
2. Preliminaries 3
3. From amenability to collection 9
4. Baby versions 17
5. The hierarchies 18
6. Games 21
7. From baby to locally measurable 26
8. Indestructibility 27
9. A detour into second-order set theory 28
10. Outlook 31
References 33

1. Introduction

Measurable cardinals and most stronger large cardinals are defined by the existence of elementary embeddings $j : V \to M$ from the universe V into an inner model M with that cardinal as the critical point. Stronger large cardinal axioms impose additional assumptions on the target model M that allow it to capture more and more sets from V. Weakly compact and many other smaller large cardinals are defined via combinatorial properties, often involving existence of large homogeneous sets for colorings. But almost all of them also have elementary embedding characterizations that, like the larger large cardinals, follow prescribed patterns as consistency strength grows. These smaller large cardinals κ are characterized by elementary embeddings of models $\langle M, \in\rangle \models \mathrm{ZFC}^-$ of size κ, most often transitive, with critical point κ, and usually, but not always, with well-founded targets. ZFC^- is the theory ZFC with the powerset axiom removed, the collection scheme in place of the replacement scheme, and the version of the axiom of choice which states that every set can be well-ordered.[1] In many cases, the existence of these embeddings can be equivalently expressed in terms of the existence of certain filters on the subsets of κ contained in M.

Suppose $\langle M, \in\rangle \models \mathrm{ZFC}^-$ is a transitive model of size κ with $\kappa \in M$ (we will give precise definitions and drop the transitivity requirement in the next section). We call such structures

2020 *Mathematics Subject Classification.* (Primary) 03E55; (Secondary) 03E35.

Key words and phrases. Large cardinal, Ramsey cardinal, measurable cardinal, infinite game.

We are grateful to Philip Welch for his kind permission to use an argument of his in the proof of Lemma 3.14.

This research was funded in whole or in part by EPSRC grant number EP/V009001/1 of the second-listed author.

[1]Without the power set axiom, collection and replacement schemes are not equivalent and neither are the various forms of the axiom of choice. See [Zar82] and [GHJ16].

weak κ-models. We will call a filter U on the subsets of κ of a weak κ-model M an *M-ultrafilter* if the structure $\langle M, \in, U\rangle$, together with a predicate for U, satisfies that U is a uniform, normal ultrafilter on κ. What this says is that U is an ultrafilter on the subsets of κ that live in M and if a κ-length sequence of elements of U is an element of M, then its diagonal intersection is in U. Łoś' theorem holds for ultrapowers by an M-ultrafilter, but the ultrapower need not be well-founded. Note that the filter U may not be countably complete for sequences outside of M. The M-ultrafilter U is, in most interesting cases, external to M and even though $\langle M, \in\rangle \models \mathrm{ZFC}^-$, separation and replacement can fail completely in the structure $\langle M, \in, U\rangle$ once we let M know about U. In the same way that making the inner model M closer to V in the characterizations of larger large cardinals increases strength, for the smaller large cardinals, we increase strength by making M be more compatible with U, which amounts to M being more correct about the properties of U or to having more of the ZFC^- axioms in the structure $\langle M, \in, U\rangle$.

Let's look at some examples. A cardinal κ is weakly compact if and only if it is inaccessible and every $A \subseteq \kappa$ is an element of a weak κ-model M for which there is an M-ultrafilter with a well-founded ultrapower. It turns out that it is equivalent to assume the a priori stronger assertion that every $A \subseteq \kappa$ is an element of a weak κ-model for which there is a (externally) countably complete M-ultrafilter U. In particular, the ultrapower is well-founded. A cardinal κ is 1-iterable if every $A \subseteq \kappa$ is an element of a weak κ-model M for which there is an M-ultrafilter U such that $\langle M, \in, U\rangle$ satisfies Σ_0-separation and the ultrapower is well-founded [GW11, Definition 2.11]. The 1-iterable cardinals are stronger than ineffable cardinals, and therefore much stronger than weakly compact cardinals. They are however still compatible with L. Thus, the additional requirement that the structure $\langle M, \in, U\rangle$ satisfies Σ_0-separation pushes up the consistency strength. A cardinal κ is Ramsey if and only if every $A \subseteq \kappa$ is an element of a weak κ-model M for which there is an M-ultrafilter U such that $\langle M, \in, U\rangle$ satisfies Σ_0-separation and which is (externally) countably complete [Mit79, Theorem 3]. Ramsey cardinals sit much higher in the hierarchy than 1-iterable cardinals. Thus, in particular, in the presence of the requirement of Σ_0-separation for the structure $\langle M, \in, U\rangle$ for the M-ultrafilter U, having U be countably complete is much stronger than having it just produce a well-founded ultrapower.

In this article, motivated by the work of Bovykin and McKenzie [BM], we consider a hierarchy of large cardinal notions characterized by the existence, for weak κ-models M, of M-ultrafilters U such that the structure $\langle M, \in, U\rangle$ satisfies fragments up to full ZFC^-. Following Bovykin and McKenzie, we call such cardinals *n-baby measurable* where the fragment is ZFC^-_n (the separation and collection schemes are restricted to Σ_n-formulas). Baby measurable cardinals and some variants were introduced by Bovykin and McKenzie in [BM, Definition 4.2] and used to obtain the following application to the theory NFUM. This theory is a natural strengthening of NFU due to Holmes (see [BM, Section 2.2]) that aims to facilitate mathematics in NFU, the latter being a variant of Quine's *New Foundations* with Urelements introduced by Jensen.

Theorem 1.1 (Bovykin, McKenzie [BM, Section 4]). *The following theories are equiconsistent.*[2]

(1) ZFC *together with the scheme consisting of the assertions for every $n \in \omega$:*

"There exists an n-baby measurable cardinal κ such that $V_\kappa \prec_{\Sigma_n} V$."

(2) NFUM

We collect facts about ultrafilters and large cardinals in Section 2 and study structures consisting of models with ultrafilters satisfying fragments of ZFC^- in Section 3. The new large cardinals notions are introduced in Section 4. We show in Section 5 that increasing levels of collection and separation for the structures $\langle M, \in, U\rangle$ increase the large cardinal strength. For example, adding even Σ_0-replacement to the characterization of 1-iterable cardinals pushes consistency strength well beyond a Ramsey cardinal, and hence beyond L. We then provide a fine analysis of the resulting hierarchies. Surprisingly, the closure of the models does not play a role when working solely with fragments of collection. However, closure conditions do induce a strict hierarchy for the setting of ZFC^-_n with both the collection and separation schemes restricted to Σ_n-formulas. This is shown in Section 6 with the help of filter games resembling those introduced by Holy and the second-listed author [HS18, Section 5]. We thus arrive at

[2] Note that their notion of n-baby measurable cardinal in [BM, Definition 4.2] is slightly different from ours in Definition 4.1 (3), but the theorem holds for our notion as well by Theorem 9.1.

similar patterns of large cardinal notions as the one around strongly Ramsey and α-Ramsey cardinals.

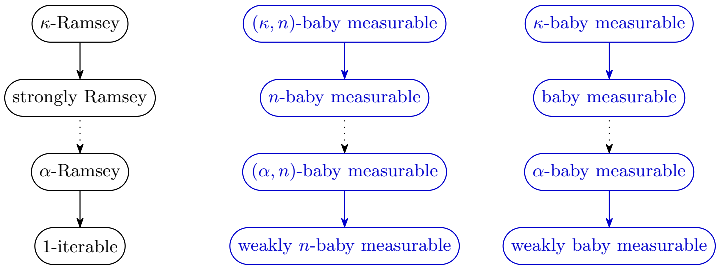

FIGURE 1. Patterns in analogy with α-Ramsey cardinals. Solid arrows denote direct implications, dotted arrows implications in consistency strength.

We show that some of these new large cardinal notions are robust under forcing in Section 8. Finally, we will see that these notions are naturally connected to models of Kelley-Morse set theory in which Ord is measurable and help us to understand the structure of these second-order models in Section 9. In particular, we apply the above methods to show that the theory of structures of the form V_κ^M, where M is a model of ZFC$^-$ or ZFC where κ is measurable, is axiomatizable by the existence of large cardinals similar to the ones above. These results shed light on the large cardinal hierarchy below measurable cardinals and provide a blueprint how to approximate a large cardinal notion from below by a natural hierarchy.

2. Preliminaries

Given a set X, when we say that X is a model of set theory, we will tacitly assume that the membership relation is the actual membership $\in$ restricted to X, so that X is really the model $\langle X, \in\rangle$.

Definition 2.1. Suppose that κ is an inaccessible cardinal.

- A *weak κ-model* is a transitive set $M \models \text{ZFC}^-$ of size κ with $V_\kappa \in M$.
- A *κ-model* is a weak κ-model such that $M^{<\kappa} \subseteq M$.
- A *basic weak κ-model* is a (not necessarily transitive) set $M \models \text{ZFC}^-$ of size κ such that $M \prec_{\Sigma_0} V$ and $V_\kappa \cup \{V_\kappa\} \subseteq M$.
- A *basic κ-model* is a basic weak κ-model such that $M^{<\kappa} \subseteq M$.

In each case, we will say that M is *simple* if κ is the largest cardinal in M.[3]

Note that if κ is a basic weak κ-model, then $\kappa \in M$. To see this, let $\alpha \in M$ be such that

$$M \models \text{“}\alpha = V_\kappa \cap \text{Ord”},$$

and then observe that by Σ_0-elementarity, $\alpha = V_\kappa \cap \text{Ord}$ holds true in V, which means that $\alpha = \kappa$.

Our canonical examples of basic models will be elementary substructures of size κ of some large H_θ for a regular θ.[4] For simple models, the notions of basic weak κ-models and weak κ-models coincide.

Lemma 2.2. *If M is a simple basic weak κ-model, then M is a weak κ-model.*

Proof. We just need to show that M is transitive. Fix $a \in M$. Since M is simple, it thinks that there is a bijection $f : \kappa \to a$. By Σ_0-elementarity, f is really a bijection between κ and a. Thus, since $\kappa \subseteq M$, $a \subseteq M$. □

[3] Note that these definitions differ slightly from those appearing in earlier literature. For instance, in the definition of weak κ-model, it is not usually assumed that κ is inaccessible and $V_\kappa \in M$.

[4] H_θ for a cardinal θ is the collection of all sets whose transitive closure has size less than θ and given that θ is regular, we have $H_\theta \models \text{ZFC}^-$.

Note that even a simple weak κ-model might not be a κ-model. For instance, its height may have countable cofinality.

Given a basic weak κ-model M, we will let $P^M(\kappa)$ denote the collection of all the subsets of κ that are elements of M. Note that in most cases, $P^M(\kappa)$ will be a class, but not an element of M.

Lemma 2.3. *Suppose that M is a basic weak κ-model. If $\bar{M} \in M$ and M thinks that $\bar{M}$ is a basic κ-model, then $\bar{M}$ is a basic κ-model.*

Proof. The model $\bar{M}$ is a basic weak κ-model by the Σ_0-elementarity of M, so it remains to check closure. Let $f : \bar{M} \to V_\kappa$ be a bijection in M, which really must be a bijection by Σ_0-elementarity. Fix a sequence $\vec{a} = \langle a_\xi \mid \xi < \beta \rangle$ such that $\beta < \kappa$ and $a_\xi \in \bar{M}$ for every $\xi < \beta$. Let $b_\xi = f(a_\xi)$ and let $\vec{b} = \langle b_\xi \mid \xi < \beta \rangle$. The sequence $\vec{b} \in V_\kappa$, and hence $\vec{b} \in M$. Thus, $f[\vec{b}] = \vec{a} \in M$ by Σ_0-elementarity, and hence $\vec{a} \in \bar{M}$ since M thinks it is a basic κ-model. □

Definition 2.4. Suppose that M is a basic weak κ-model.

- We will say that $U \subseteq P^M(\kappa)$ is an M-*ultrafilter* if the structure
$$\langle M, \in, U \rangle \models \text{“}U \text{ is a uniform[5] normal ultrafilter on } \kappa.\text{”}$$
- An M-ultrafilter U is *good* if the ultrapower of M by U is well-founded.
- An M-ultrafilter U is *countably complete* if for every sequence $\langle A_n \mid n < \omega \rangle$ with $A_n \in U$ (but the sequence itself not necessarily in M), $\bigcap_{n<\omega} A_n \neq \emptyset$.

If M is a basic weak κ-model and $j : M \to N$ is an elementary embedding with $\text{crit}(j) = \kappa$, then
$$U = \{A \subseteq \kappa \mid A \in M \text{ and } \kappa \in j(A)\}$$
is the M-ultrafilter *derived* from j, and if N is well-founded, then U is good. If j happens to be the ultrapower by an M-ultrafilter U, then the M-ultrafilter derived from j is precisely U. Standard arguments using Łoś' theorem show that if M is a basic weak κ-model and U is a countably complete M-ultrafilter, then U is good. In particular, if a basic weak κ-model M is closed under ω-sequences ($M^\omega \subseteq M$), e.g. if M is a κ-model, then every M-ultrafilter must be countably complete, and hence good. Thus, whenever M is a basic weak κ-model such that $M^\omega \subseteq M$, every M-ultrafilter is automatically good. Even when the ultrapower N of M by U is ill-founded, by our assumptions on U, $V_\kappa \cup \{V_\kappa\} \subseteq N$ and $V_\kappa^M = V_\kappa^N$.[6]

Definition 2.5. Suppose that M is a basic weak κ-model. An M-ultrafilter U is *weakly amenable* if for all $A \in M$ with $|A|^M = \kappa$, $A \cap U \in M$.

Note that for simple models M, a weakly amenable M-ultrafilter U is fully amenable in the sense that for every $A \in M$, we have $A \cap U \in M$. A weakly amenable M-ultrafilter can be iterated to carry out the iterated ultrapower construction for any ordinal length. Recall that in the case of say a measure U on κ (namely if κ is a measurable cardinal), we can define an Ord-length system of iterated ultrapowers of $U = U_0$. We let $j_{01} : V = M_0 \to M_1$ be the ultrapower by U_0 and use j_{01} to obtain the ultrafilter for the next stage of the iteration by defining $U_1 = j_{01}(U_0)$. This gives us a general procedure for obtaining the next stage ultrafilter at the successor stages of the iteration and at limits we take a direct limit of the system of embeddings obtained thus far. By a theorem of Gaifman, all the iterated ultrapowers M_ξ are well-founded [Gai74, Section II, Theorem 5]. We cannot apply the same procedure to obtain successor stage ultrafilters with an M-ultrafilter U because U is (in most interesting cases) not an element of M. However, given weak amenability, we can define, for example, U_1, given that $j : M \to N$ is the ultrapower embedding, by
$$U_1 := \{A = [f]_U \subseteq j(\kappa) \mid \{\alpha < \kappa \mid f(\alpha) \in U\} \in U\}.$$

Weak amenability has an equivalent characterization in terms of the preservation of subsets of κ between the model and its ultrapower. An M-ultrafilter U for a basic weak κ-model M is weakly amenable if and only if M and its ultrapower N have the same subsets of κ. Note that this holds true regardless of whether N is well-founded or not. It is also the case that if

[5] Recall that a filter on a cardinal κ is *uniform* if it contains all tail sets (α, κ) for $\alpha < \kappa$.

[6] We are assuming here that we have collapsed the well-founded part of N.

$j : M \to N$ is an elementary embedding with $\mathrm{crit}(j) = \kappa$ such that M and N have the same subsets of κ, then the M-ultrafilter derived from j is weakly amenable.

Lemma 2.6. *If M is a simple basic weak κ-model, U is a weakly amenable M-ultrafilter, and $\langle N, \bar{\in}\rangle$ is the ultrapower of M by U, then $M = H^N_{\kappa^+}$.*

Proof. Let's assume that we have collapsed the well-founded part of $\langle N, \bar{\in}\rangle$ so that for well-founded sets, $\in = \bar{\in}$. Clearly $M \subseteq H^N_{\kappa^+}$. So suppose that $a \mathrel{\bar{\in}} H^N_{\kappa^+}$ and assume without loss that a is transitive in $\bar{\in}$. Let $f : a \to \kappa$ be a bijection in N, and let

$$A = \{(f(a), f(b)) \mid (a, b) \in \bar{\in}\} \subseteq \kappa \times \kappa.$$

Then $A \in M$ by weak amenability. First, suppose that A is ill-founded. Then M would know this, and hence it would have a descending ω-sequence witnessing the ill-foundedness. But then N would have the sequence as well, which is impossible. Thus, A is well-founded. We can now argue by recursion on rank that the Mostowski collapse of A in M is the Mostowski collapse of A in N, and hence $(a, \bar{\in}) = (a, \in)$. □

The simplest characterization of a large cardinal in terms of embeddings on weak κ-models belongs to weakly compact cardinals.

Theorem 2.7 (folklore)**.** *Suppose that κ is inaccessible. Then the following are equivalent.*

1. *κ is weakly compact.*
2. *Every $A \subseteq \kappa$ is an element of a weak κ-model M for which there is a good M-ultrafilter.*
3. *Every $A \subseteq \kappa$ is an element of a weak κ-model M for which there is a countably complete M-ultrafilter.*
4. *Every weak κ-model has a countably complete M-ultrafilter.*

It is then natural to ask what happens if we addtionally require that the M-ultrafilter is weakly amenable. The answer is that we get a much stronger large cardinal notion. Let's start by introducing the following large cardinal notion.

Definition 2.8. A cardinal κ is 0-*iterable* if every $A \subseteq \kappa$ is an element of a weak κ-model M for which there is a weakly amenable (but not necessarily good) M-ultrafilter.

Recall that a cardinal κ is *weakly ineffable* if for every sequence $\langle A_\alpha \mid \alpha < \kappa\rangle$ with $A_\alpha \subseteq \alpha$, there is a threading set $A \subseteq \kappa$ such that for unboundedly many $\alpha < \kappa$, $A \cap \alpha = A_\alpha$. A cardinal κ is *ineffable* if we can find such a set A that is stationary in κ.

Proposition 2.9. *A* 0*-iterable cardinal is a weakly ineffable limit of ineffable cardinals.*

Proof. Fix a sequence $\vec{A} = \langle A_\alpha \mid \alpha < \kappa\rangle$ with $A_\alpha \subseteq \alpha$, and find a weak κ-model M with $\vec{A} \in M$ for which there is weakly amenable M-ultrafilter U. Let $j : M \to N$ be the possibly ill-founded ultrapower by U. It is easy to check using Łoś' theorem that $A = j(\vec{A})(\kappa)$ has the required property. This argument also shows that κ is ineffable in the ultrapower N. Thus, κ is a limit of ineffable cardinals. □

In particular, requiring that the M-ultrafilter be weakly amenable, but not even requiring that it is good, gives a large cardinal notion much stronger than weakly compact cardinals.

Definition 2.10 ([GW11, Definition 2.11])**.** A cardinal κ is *1-iterable* if every $A \subseteq \kappa$ is an element of a weak κ-model M for which there is a good weakly amenable M-ultrafilter.

Iterating a measure on κ gives rise to Ord-many well-founded iterated ultrapowers. By a theorem of Kunen, the same holds true for iterating a weakly amenable countably complete M-ultrafilter for a weak κ-model M [Kun70]. However, a weakly amenable M-ultrafilter that is not countably complete can give rise to $0 \leq \alpha < \omega_1$ or Ord-many well-founded ultrapowers (the later follows just as in the case of a measure from having ω_1-many). In fact, we can define the hierarchy of α-iterable cardinals (for $1 \leq \alpha \leq \omega_1$), where a cardinal κ is *α-iterable* if every $A \subseteq \kappa$ is an element of a weak κ-model M for which there is a weakly amenable M-ultrafilter with α-many well-founded iterated ultrapowers [GW11]. The 0-iterable cardinals fit naturally at the head of this hierarchy, which is consistent with L for $\alpha < \omega_1$ [GW11], while it is not difficult to see that ω_1-iterable cardinals imply $0^{\#}$.

Theorem 2.11 (Mitchell [Mit79, Theorem 3]). *A cardinal κ is Ramsey if and only if every $A \subseteq \kappa$ is an element of a weak κ-model M for which there is a weakly amenable countably complete M-ultrafilter on κ.*

The ω_1-iterable cardinals are slightly weaker than Ramsey cardinals, which are ω_1-iterable limits of ω_1-iterable cardinals [SW11, Lemma 5.2].

Note that unlike the situation with weakly compact cardinals, asking that the weakly amenable M-ultrafilter be countably complete, and not just good, pushes up the large cardinal strength from 1-iterable cardinals, which are consistent with L, to Ramsey cardinals. For weakly compact cardinals κ, every weak κ-model, and so, in particular, every κ-model, has an M-ultrafilter. If we ask that every $A \subseteq \kappa$ is an element of a κ-model for which there is a weakly amenable M-ultrafilter (which is automatically good), then we get a large cardinal notion stronger than a Ramsey cardinal. Indeed, assuming that we have a weakly amenable M-ultrafilter for every weak κ-model is inconsistent [Git11, Theorem 1.7].

Definition 2.12. Suppose that κ is a cardinal.

- [Git11, Definition 1.4] κ is *strongly Ramsey* if every $A \subseteq \kappa$ is an element of a κ-model M for which there is a weakly amenable M-ultrafilter.
- [Git11, Definition 1.5] κ is *super Ramsey* if every $A \subseteq \kappa$ is an element of a κ-model $M \prec H_{\kappa^+}$ for which there is a weakly amenable M-ultrafilter.

Strongly Ramsey cardinals are limits of Ramsey cardinals, super Ramsey cardinals are limits of strongly Ramsey cardinals, and measurable cardinals are limits of super Ramsey cardinals [Git11, Theorem 1.7]. It is natural to ask given these various notions what will happen (1) if we stratify by closure on the weak κ-model M, weakening, for instance, the κ-model assumption to just countable closure on M, and (2) if we ask for elementarity in a large H_θ. The second question does not make sense as stated because a weak κ-model cannot be elementary in H_θ for any $\theta > \kappa^+$, but this is precisely where the weakening to basic weak κ-models comes into play. Combining both of these ideas, Holy and the second-listed author introduced the α-Ramsey hierarchy below a measurable cardinal [HS18].

Definition 2.13. Suppose that κ is a cardinal.

- [HS18, Definition 5.1] κ is α-Ramsey for a regular $\omega \leq \alpha \leq \kappa$ if for every $A \subseteq \kappa$ and arbitrarily large cardinals θ, there is a $<\alpha$-closed basic weak κ-model $M \prec H_\theta$ with $A \in M$ for which there is a good weakly amenable M-ultrafilter.
- [HS18, Definition 5.7] κ is *faintly ω-Ramsey*[7] if in the definition of an ω-Ramsey cardinal, the M-ultrafilter U is not required to be good.

Note that in the definition of α-Ramsey cardinals, only the ω-Ramsey cardinals need the extra assumption that the M-ultrafilter is good; for the others, the closure on the basic κ-model already implies it. The α-Ramsey cardinals have a natural game theoretic characterization [HS18, Theorem 5.6] arising out of the question of whether given a weak κ-model M, an M-ultrafilter U, and another weak κ-model N extending M, we can always find an N-ultrafilter extending U. The first-listed author showed that this extension property is inconsistent [HS18, Proposition 2.13], but the above characterization can be understood as a game theoretic variant of this property.

Consider the following game $\mathsf{RamseyG}^\theta_\alpha(\kappa)$, for regular cardinals $\omega \leq \alpha \leq \kappa$ and $\theta > \kappa$, of perfect information played by two players the challenger and the judge. The challenger starts the game and plays a basic κ-model $M_0 \prec H_\theta$. The judge responds by playing an M_0-ultrafilter U_0. In the next step, the challenger plays a basic κ-model $M_1 \prec H_\theta$ such that $\langle M_0, U_0\rangle \in M_1$, and the judge responds with an M_1-ultrafilter U_1 extending U_0. More generally, at stage γ, the challenger plays a basic κ-model $M_\gamma \prec H_\theta$ such that $\{\langle M_\xi, U_\xi\rangle \mid \xi < \gamma\} \in M_\gamma$, and the judge responds with an M_γ-ultrafilter U_γ extending $\bigcup_{\xi<\gamma} U_\gamma$. The judge wins the game if she can continue playing for α-many steps (and $U = \bigcup_{\xi<\alpha} U_\xi$ is a good $M = \bigcup_{\xi<\alpha} M_\xi$-ultrafilter). Otherwise, the challenger wins. Since $M = \bigcup_{\xi<\alpha} M_\xi$ is a union of κ-models, it is $<\alpha$-closed, and thus, the additional assumption that $U = \bigcup_{\xi<\alpha} U_\xi$ is good is only necessary for $\alpha = \omega$. Let $\mathsf{faintRamseyG}^\theta_\omega(\kappa)$ be an analogous game to $\mathsf{RamseyG}^\theta_\omega(\kappa)$, but where we don't require the

[7] This property is called the *ω-filter property* in [HS18].

judge to ensure that the union U of her plays is good. It is shown in [HS18, Lemma 3.3] that if either of the players has a winning strategy in the game $\mathsf{RamseyG}^\theta_\alpha(\kappa)$, then the same player has a winning strategy in the game $\mathsf{RamseyG}^\rho_\alpha(\kappa)$ for any other regular cardinal $\rho > \kappa$. An analogous result holds for the game $\mathsf{faintRamseyG}^\theta_\omega(\kappa)$.

Theorem 2.14. *Suppose that κ is a cardinal.*

- [HS18, Theorem 5.6] *For $\omega \leq \alpha < \kappa$, κ is α-Ramsey if and only if the challenger does not have a winning strategy in the game $\mathsf{RamseyG}^\theta_\alpha(\kappa)$ for some/all regular cardinals $\theta > \kappa$.*
- *κ is faintly ω-Ramsey if and only if the challenger does not have a winning strategy in the game $\mathsf{faintRamseyG}^\theta_\omega$ for some/all regular cardinals $\theta > \kappa$.*

The second part of the Theorem, although not explicitly proved in [HS18], is proved analogously to proof of the first part given there.

Nielsen and Welch showed that faintly ω-Ramsey cardinals are precisely the well-known completely ineffable cardinals . The first author showed that the ω-Ramsey cardinals lie between 1-iterable and 2-iterable cardinals in consistency strength. The result was mentioned in [HS18, Section 8] without proof.

Theorem 2.15. *A 2-iterable cardinal κ is a limit of ω-Ramsey cardinals.*

Proof. Suppose that κ is 2-iterable. It follows from the diagram in [Git11, Lemma 4.4] that there is a weak κ-model $M \models \mathrm{ZFC}$ and a good M-ultrafilter U such that for the ultrapower map $j : M \to N$, we have $M = V^N_{j(\kappa)}$. We would like to argue that κ is ω-Ramsey in $V^N_{j(\kappa)}(= M)$. Fix a regular $\theta \in M$ and $A \subseteq \kappa$ in M. We need to produce a basic weak κ-model $m \prec H^M_\theta$ in M with $A \in m$ for which there is a good weakly amenable m-ultrafilter on κ in M. Let $A \in M_0 \prec H^M_\theta$ be any basic weak κ-model in M. Since M_0 has size κ, the restriction $j : M_0 \to j(M_0)$ is in N. Next, let $M_1 \prec H^M_\theta$ be a basic weak κ-model in M such that $M_0, U \cap M_0 \in M_1$. Again, we have $j : M_1 \to j(M_1)$ is in N. In this way, we construct an elementary sequence $\{M_n \mid n < \omega\}$ such that for each $n < \omega$, $M_n, U \cap M_n \in M_{n+1}$, and the restriction $j : M_n \to j(M_n)$ is in N.

Now working in N, we define a tree T whose elements are finite sequences

$$\langle h_0 : m_0 \to k_0, \dots, h_{n-1} : m_{n-1} \to k_{n-1} \rangle,$$

ordered by end-extension, such that for all $i, j < n$:

(1) $m_i \prec H^M_\theta$,
(2) $m_i, u_i \in m_{i+1}$, where u_i is the m_i-ultrafilter derived from h_i.
(3) $m_i \prec m_j$, $k_i \prec k_j$, and $h_i \subseteq h_j$.

Note crucially that the tree T is an element of N by the assumption that $j(\kappa)$ is a set in N. The sequence $\{j_n : M_n \to j(M_n) \mid n < \omega\}$, constructed above, witnesses that the tree T has a branch, and is thus ill-founded in V. By the absoluteness of well-foundedness, T is ill-founded in N as well. Let $h = \bigcup_{i<\omega} h_i : m \to n$, where $m = \bigcup_{i<\omega} m_i$ and $n = \bigcup_{i<\omega} n_i$. Finally, observe that $u = \bigcup_{i<\omega} u_i$ is the m-ultrafilter derived from h, and must be weakly amenable and good by construction. □

It is shown in [HS18, Proposition 5.2] that a κ-Ramsey cardinal κ is a limit of super Ramsey cardinals. Let us say that a cardinal κ is *$<\alpha$-Ramsey* if it is β-Ramsey for all regular $\omega \leq \beta < \alpha$.

Proposition 2.16. *A strongly Ramsey cardinal is a limit of cardinals α that are $<\alpha$-Ramsey.*

Proof. Suppose that κ is strongly Ramsey. Let M be a simple κ-model for which there is a weakly amenable M-ultrafilter U. Let $j : M \to N$ be the ultrapower by U. We will argue that κ is $<\kappa$-Ramsey in $V^N_{j(\kappa)}$. Fix $\alpha < \kappa$. By Theorem 2.14, it suffices to show that the challenger does not have a winning strategy in the game $\mathsf{RamseyG}^{\kappa^+}_\alpha(\kappa)$. Suppose towards a contradiction that, in $V^N_{j(\kappa)}$, the challenger has a winning strategy σ in $\mathsf{RamseyG}^{\kappa^+}_\alpha(\kappa)$. Note that by the weak amenability of U, the moves of the challenger are in M. Consider the following run of the game $\mathsf{RamseyG}^{\kappa^+}_\alpha(\kappa)$. The challenger plays some M_0 according to σ. The judge responds with $U \cap M_0$, which is an element of M by weak amenability. The challenger then plays M_1 according to σ, and the judge again responds with $U \cap M_1$. Suppose that the challenger and the judge continue to play in this manner up to some limit step β. Since M is a κ-model, the run of the game up to β is an element of M, and hence, the challenger can respond to it with some M_α according to σ. Thus, the judge and the challenger can continue playing in this manner for α-many steps. The entire run of the game is also an element of M by closure. But clearly the judge wins this play, contradicting our assumption that σ was a winning strategy for the challenger. Now we can conclude by elementarity that, in V_κ, κ is a limit of cardinals α that are $<\alpha$-Ramsey. Moreover, V_κ is correct about this since to verify that α is β-Ramsey (for some $\beta < \kappa$) we only need to consider games with elementary substructures of $H_{\alpha^+} \subseteq V_\kappa$. □

For future sections, let us consider a variant game $\mathsf{Ramsey}\bar{\mathsf{G}}^\theta_\alpha(\kappa)$, where we ask the judge to play, instead of an M_γ-ultrafilter U_γ extending her moves from the previous stages, a structure $\langle N_\gamma, \in, U_\gamma\rangle$ such that N_γ is a κ-model with $P^{M_\gamma}(\kappa) \subseteq N_\gamma$ and U_γ is an N_γ-ultrafilter, keeping the requirement that the U_γ's must extend. We additionally require the challenger to ensure that the sequence of the judge's previous plays $\{\langle N_\xi, \in, U_\xi\rangle \mid \xi < \gamma\}$ is an element of M_γ. In this variant game, we are giving the judge the extra ability to ensure that certain subsets of κ make it into the final model. Essentially the same argument as for the game $\mathsf{RamseyG}^\theta_\alpha(\kappa)$ shows that the existence of a winning strategy for either player is independent of θ in the game $\mathsf{Ramsey}\bar{\mathsf{G}}^\theta_\alpha(\kappa)$, and moreover κ is α-Ramsey if and only if the challenger doesn't have a winning strategy in the game $\mathsf{Ramsey}\bar{\mathsf{G}}^\theta_\alpha(\kappa)$. It follows that the challenger has a winning strategy in the game $\mathsf{RamseyG}^\theta_\alpha(\kappa)$ if and only if he has a winning strategy in $\mathsf{Ramsey}\bar{\mathsf{G}}^\theta_\alpha(\kappa)$. But, in fact, this is true for the judge as well.

Lemma 2.17. *Either player has a winning strategy in the game* $\mathsf{RamseyG}^\theta_\alpha(\kappa)$ *if and only if they have a winning strategy in the game* $\mathsf{Ramsey}\bar{\mathsf{G}}^\theta_\alpha(\kappa)$.

Proof. We argued above that this is true for the challenger. So suppose that the judge has a winning strategy σ in the game $\mathsf{RamseyG}^\theta_\alpha(\kappa)$. The winning strategy $\bar\sigma$ for the judge in the game $\mathsf{Ramsey}\bar{\mathsf{G}}^\theta_\alpha(\kappa)$ is going to be to play $\langle N_\xi, \in, U_\xi\rangle$ at stage ξ of the game, where U_ξ is the response of the judge according to σ to the moves of the challenger so far (because these could very well be the moves of the challenger in the game $\mathsf{RamseyG}^\theta_\alpha(\kappa)$) and $N_\xi = M_\xi \cap H_{\kappa^+}$. Next, suppose that the judge has a winning strategy $\bar\sigma$ in the game $\mathsf{Ramsey}\bar{\mathsf{G}}^\theta_\alpha(\kappa)$. The winning strategy σ for the judge in the game $\mathsf{RamseyG}^\theta_\alpha(\kappa)$ is going to be to play $U_\xi \cap M_\xi$ at stage ξ of the game, where $\langle N_\xi, \in, U_\xi\rangle$ is the response of the of judge according to $\bar\sigma$ to the moves of the challenger so far (because again these could be the moves of the challenger in the game $\mathsf{Ramsey}\bar{\mathsf{G}}^\theta_\alpha(\kappa)$). Note that if $\{U_\xi \mid \xi < \gamma\} \in M_\gamma$, then the sequence $\{\langle N_\xi, \in, U_\xi\rangle \mid \xi < \gamma\}$ is also in M_γ because each N_ξ is definable from U_ξ since every element of N_ξ is coded by a subset of κ and U_ξ has either the subset or its complement. □

Above the α-Ramsey hierarchy, but still below a measurable cardinal sits a large cardinal notion introduced by Holy and Lücke [HL21].

Definition 2.18 ([HL21, Definition 1.1])**.** A cardinal κ is *locally measurable* if every $A \subseteq \kappa$ is an element of a weak κ-model M which thinks that it has normal ultrafilter on κ.

Standard arguments show that a measurable cardinal is a limit of locally measurable cardinals. It follows immediately from [HL21, Theorem 15.3] that a locally measurable cardinal is a limit of cardinals κ that are κ-Ramsey. To see this, note that the set of $\Delta^\forall_\kappa$-Ramsey cardinals is unbounded in κ by this theorem. Moreover, by the list of large cardinal properties after [HL21, Definition 9.4], a cardinal κ is κ-Ramsey if and only if it is $\mathbf{T}^\forall_\kappa$-Ramsey, where the latter follows from $\Delta^\forall_\kappa$-Ramsey by [HL21, Definition 9.4]. As happens often with these smaller large cardinals, replacing weak κ-models by κ-models in the definition of locally measurable increases consistency strength (the argument is similar to that given in Proposition 5.7). In the situation of Definition 2.18, $P^M(\kappa)$ must be an element of any such weak κ-model M, and so it obviously cannot be simple. While the normal ultrafilter U is not required to be good, Proposition 2.19 shows that this would not add extra strength.

Proposition 2.19. *If κ is locally measurable, then every $A \subseteq \kappa$ is an element of a weak κ-model M such that M contains what it thinks is a normal ultrafilter U on κ with a well-founded ultrapower $N \subseteq M$.*

Proof. Fix $A \subseteq \kappa$. Choose any weak κ-model M such that $A \in M$ and M has what it thinks is a normal ultrafilter U on κ. Let $\delta = (2^\kappa)^M$. We can assume without loss of generality that δ is the largest cardinal of M by replacing M with $H^M_{\delta^+}$ if necessary. Let $\langle N, \bar{\in}\rangle$ be the ultrapower of M by U. The model M believes that the relation $\bar{\in}$ is well-founded in the sense that it doesn't have a sequence of functions $\{f_n : \kappa \to M \mid n < \omega\}$ such that for all $n < \omega$, $\{\xi < \kappa \mid f_{n+1}(\xi) \in f_n(\xi)\} \in U$. Since models of ZFC$^-$ can perform Mostowski collapses, it suffices to argue that the relation $\bar{\in}$ is set-like from the perspective of M, namely for every function $f : \kappa \to M$, there is a set X_f in M such that for every $g : \kappa \to M$ with $[g]_U \bar{\in} [f]_U$, there is some $g' : \kappa \to M$ with $[g']_U = [g]_U$ and $g' \in X_f$. By elementarity, $[c_\delta]_U$, where c_δ is the constant function with value δ, is the largest cardinal of N, and thus, every element of N is bijective with $[c_\delta]_U$. So it actually suffices to argue that M has a set X_{c_δ} as above. For this, observe that $[g]_U \,\bar{\in}\, [c_\delta]_U$ if and only if $[g]_U = [g']_U$ for some $g' : \kappa \to \delta$, and the collection of functions $g' : \kappa \to \delta$ is a set in M because $\delta = 2^\kappa$ is a set in M. □

Thus, as long as 2^κ is the largest cardinal of a weak κ-model M with a normal ultrafilter U on κ, the ultrapower of M by U is contained in M just as in the case of an ultrapower by an actual measure. In fact, it follows from the proof that this will be case whenever M has a largest cardinal γ and γ^κ is a set in M. Note that it is definitely possible for an ultrapower of a weak κ-model M by a normal ultrafilter U in M to have more ordinals than M. Suppose that κ is measurable, U is a normal ultrafilter on kappa, $j : V \to M$ is the ultrapower embedding, and $\lambda > 2^\kappa$ is a cardinal with $\text{cof}(\lambda) = \kappa$. By elementarity, $\text{cof}(j(\lambda)) = j(\kappa)$ in M, which means, since $M^\kappa \subseteq M$, that $j(\lambda) > \lambda^+$, and hence also then $(j(\lambda)^+)^M > \lambda^+$. Thus, $H^M_{(j(\lambda)^+)^M}$, which is the ultrapower of H_{λ^+} by U, has more ordinals than H_{λ^+}. Let N be the Mostowski collapse of an elementary substructure of H_{λ^+} of size κ. Let $\bar{U}$ be the preimage of U under the collapse and let $\bar{\lambda}$ be the preimage of λ. Then, by elementarity, N satisfies that the image of $\bar{\lambda}$ under the ultrapower map cannot be a set in M.

3. From amenability to collection

Suppose that M is a simple weak κ-model and U is an M-ultrafilter. By assumption, M is a model of ZFC$^-$, but separation and replacement can fail very badly in the structure $\langle M, \in, U\rangle$ once the model learns about the ultrafilter. Weak amenability of U is equivalent to having a minimal amount of separation for $\langle M, \in, U\rangle$. Note that for the results in this section, we mostly do not need to assume that the ultrapower of M by U is well-founded, that is, U need not be good.

Lemma 3.1. *Suppose that M is a simple weak κ-model and U is an M-ultrafilter. Then U is weakly amenable if and only if the structure $\langle M, \in, U\rangle$ satisfies Δ_0-separation.*

Proof. Suppose that Δ_0-separation holds in the language with U. Fix $A \in M$ and observe that

$$U \cap A = \{x \in A \mid x \in U\},$$

which indeed requires separation only for atomic formulas. Now suppose that U is weakly amenable to M. Fix $A \in M$ and a Δ_0-formula $\varphi(x, b)$ in the language with U. We need to argue that

$$\{x \in A \mid M \models \varphi(x, b)\} \in M.$$

Now observe that since $\varphi(x, b)$ is a Δ_0-formula, all quantifiers are bounded and therefore, U restricted to $\text{TCl}(b)$, the transitive closure of b, suffices to interpret $\varphi(x, b)$ correctly. But this piece of U is an element of M by weak amenability. □

It follows from Theorem 5.2 below that a structure $\langle M, \in, U\rangle$ with a weakly amenable M-ultrafilter U need not need to satisfy even Δ_0-replacement.

Next, let's observe that our typical structures $\langle M, \in, U\rangle$ have a Δ_1-definable total order that is a strong well-order from the point of view of the structure. Let us say that a total order $\lhd$ on a weak κ-model M is a *strong well-order of M* if for every formula $\varphi(x, a)$ over $\langle M, \in\rangle$, the structure $\langle M, \in, \lhd\rangle$ satisfies that there is a $\lhd$-least b such that $\varphi(b, a)$. Usually, a well-order is defined to have the property that every set has a least element. This property is equivalent to our stronger requirement for a set-like order, but the orders we encounter won't necessarily be set-like. These structures $\langle M, \in, U\rangle$ will also have a Δ_1-definable truth predicate for $\langle M, \in\rangle$.

Lemma 3.2. *Suppose that M is a simple weak κ-model and U is a weakly amenable M-ultrafilter. Then the structure $\langle M, \in, U\rangle$ has a Δ_1-definable strong well-order $\lhd_U$ of M.*

Proof. Let $\langle N, \bar{\in}\rangle$ be the possibly ill-founded ultrapower by U. Since U is weakly amenable, we have $M = H^N_{\kappa^+}$ and $(\kappa^+)^N = \mathrm{Ord}^M$ by Lemma 2.6. Note that every set $a \in M$ can be coded in M by a subset A of κ. The code A codes a subset of $\kappa \times \kappa$, which in turn, codes the membership relation on $\mathrm{TCl}(a)$. While a set a can have many different codes, N has a set $\mathcal{C}$ consisting of a unique code for every set in M by elementarity, since this holds true in M of every H_ν. In fact, N has a membership relation $\mathcal{E}$ for elements of this set as well, making $\langle \mathcal{C}, \mathcal{E}\rangle$ isomorphic to $\langle M, \in\rangle$. Let $[C]_U$ be the equivalence class representing the set $\mathcal{C}$ of N.

Let's consider the complexity of a series of statements in the structure $\langle M, \in, U\rangle$. First observe that any $A \subseteq \kappa$ from M is represented in the ultrapower N by the equivalence class of the function f_A such that $f_A(\xi) = A \cap \xi$. To verify this, it suffices to check that $[f_A]_U \cap \xi = A \cap \xi$ for every $\xi < \kappa$. Since ξ and $A \cap \xi$ are in V_κ, they are represented by the constant functions $[c_\xi]_U$ and $[c_{A\cap\xi}]_U$. So we need to check that

$$[f_A]_U \cap [c_\xi]_U = [c_{A\cap\xi}]_U$$

holds in the ultrapower N. By Łoś' theorem, this holds if and only if

$$\{\alpha < \kappa \mid f_A(\alpha) \cap \xi = A \cap \xi\} \in U,$$

but this is certainly true since it holds on a tail. Let "A is a code" be the assertion that A is an element of $\mathcal{C}$ in the ultrapower N. The complexity of the statement "A is a code" in the structure $\langle M, \in, U\rangle$ is Δ_1. The Σ_1-version says that there exists a function f such that $f(\xi) = A \cap \xi$ and there exists a set $\{\alpha < \kappa \mid f(\xi) \in C(\xi)\}$ and this set is in U, and the Π_1-version says that for every function f such that $f(\xi) = A \cap \xi$ and for every set equal to $\{\alpha < \kappa \mid f(\xi) \in C(\xi)\}$, the set is in U.

Next, let "A is the code for a" be the conjunction of assertions that "A is a code" and that A codes a as explained above. The assertion that A codes a is actually Δ_1 in the structure $\langle M, \in\rangle$ without the ultrafilter. The Σ_1-assertion is that there are sets $\bar{a}$ and π such that $\bar{a} = \mathrm{TCl}(a)$ and π is an isomorphism between $\bar{a}$ and A. Note that $\bar{a} = \mathrm{TCl}(a)$ is Δ_1. The Π_1-assertion is that for every set $\bar{a}$, B, and π, if $\bar{a} = \mathrm{TCl}(a)$ and π is an isomorphism between $\bar{a}$ and B, then $A = B$.

Finally, let $[W]_U$ represent the equivalence class of a well-order $\mathcal{W}$ of $\mathcal{C}$ in the ultrapower N. Given a and b in M, we define that $a \lhd_U b$ whenever there are codes A and B of a and b respectively such that A is below B according to $\mathcal{W}$. The assertion that A is below B according to $\mathcal{W}$ translates to saying that $[f_A]_U$ is below $[f_B]_U$ according to $[W]_U$. With our preliminary analysis this is clearly a Δ_1-assertion in the language with U, using Łoś' theorem. □

If the GCH holds below κ, or just $\{\alpha < \kappa \mid M \models 2^\alpha = \alpha^+\} \in U$, then the ultrapower N has a well-order of M of order-type κ^+, which must then be an actual well-order as κ^{+N} is well-founded. If the M-ultrafilter U is good, and hence N is well-founded, then its well-order of M is an actual well-order. Otherwise, it is possible that N's well-order of M is ill-founded externally.

Recall that a *choice function* for a binary relation R on a set X is a function $F\colon X \to X$ with $(x, F(x)) \in R$ for all $(x, y) \in R$. The next lemma shows that our typical structures $\langle M, \in, U\rangle$ have a Δ_1-definable choice function for every relation R on M that exists in the ultrapower of M by U.

Lemma 3.3. *Suppose that M is a simple weak κ-model and U is an M-ultrafilter such that $\langle M, \in, U\rangle$ is weakly amenable. Then any binary relation R on M that exists in the ultrapower of M by U admits a choice function that is Δ_1-definable in $\langle M, \in, U\rangle$.*

Proof. Let N be the, possibly ill-founded, ultrapower of M by U. Let $R \in N$ be a binary relation on M, and observe that N must have choice functions for R. Let $\bar{F} \in M$ be a function on the class of codes $\mathcal{C}$ corresponding to a choice function on R. Then we can define a choice function F on R by $(x, y) \in F$ whenever A_x is a code for x, A_y is a code for y and $\bar{F}(A_x) = A_y$. Letting $[f]_U$ represent the equivalence class of $\bar{F}$, we use similar arguments as above to show that F is Δ_1 definable over $\langle M, \in, U\rangle$. □

A similar argument shows that the ultrafilter provides a definable truth predicate.

Lemma 3.4. *Suppose that M is a simple weak κ-model and U is a weakly amenable M-ultrafilter. Then $\langle M, \in, U\rangle$ has a Δ_1-definable truth predicate $\mathrm{Tr}_M(\varphi, x)$ for $\langle M, \in\rangle$.*

Proof. Let N be the possibly ill-founded ultrapower of M by U. In the notation of the proof of Proposition 3.2, the model N has a truth predicate $\mathcal{T}$ for the structure $\langle \mathcal{C}, \mathcal{E}\rangle$. Thus, we have $M \models \varphi(a)$ whenever A is the code for a and $\langle \varphi, A\rangle \in \mathcal{T}$. Let $[T]_U$ represent $\mathcal{T}$ in the ultrapower N. Then $\langle \varphi, A\rangle \in \mathcal{T}$ if and only if $\{\alpha < \kappa \mid \langle \varphi, f_A(\alpha)\rangle \in T(\alpha)\} \in U$. □

Lemma 3.5. *Suppose that $M \models \mathrm{ZFC}^-$, U is an amenable predicate on M. Then for every $n < \omega$, $\langle M, \in, U\rangle$ has a Σ_n-definable truth predicate $\mathrm{Tr}^U_n(\varphi, x)$ for Σ_n-formulas in the language with U.*

Proof. Given a Δ_0-formula $\varphi(x, a)$ in the language with U, observe that we only need $U_a := U \cap \mathrm{TCl}(a)$ to evaluate it. By the amenability of U, we have that $U_a \in M$ for every $a \in M$. Thus, a truth predicate for a Δ_0-formula $\varphi(x, a)$ can be defined as usual in a Δ_1 fashion with U_a as a parameter. This means we should require that the set U_a is a part of the witnessing sequence for truth and the extra step in the definition needs to check that $U_a = U \cap \mathrm{TCl}(a)$, but that is a Δ_0-assertion. Once we have truth for Δ_0-formulas, the rest follows by induction on complexity of formulas because we have only finitely many quantifiers. □

We will mainly be interested in properties of models $\langle M, \in, U\rangle$ of the following fragments of ZFC^-:

- ZFC^-_n denotes the theory where the collection and separation schemes are restricted to Σ_n-formulas.
- KP_n denotes the theory where the collection scheme is restricted to Σ_n-formulas and the separation scheme is restricted to Δ_0-formulas.

Note that ZFC^-_0, KP_0 and KP_1 are equivalent. The theory KP_n further implies fragments of separation and recursion. Parts of the next folklore result are stated without proof in [Kra82, Theorem 1.5].

Lemma 3.6. *The theory KP_n implies the schemes of Σ_n-replacement, Δ_n-separation and Σ_n-recursion along the ordinals.*

Proof. The claims for $n = 0$ and $n = 1$ are identical, since $\mathrm{KP}_0 = \mathrm{KP}_1$. Their proofs are easy variants of the following argument.

For $n > 1$, suppose that M is a model of KP_n. We work in M. We have Σ_{n-2}-separation by the induction hypothesis. We first show that Σ_{n-1}-separation holds. Suppose that A is a set and $\varphi(x, y)$ is a Π_{n-2}-formula. We need to show that the set

$$B := \{x \in A \mid \exists y\, \varphi(x, y)\}$$

is in M. Consider the Π_{n-1}-formula $\bar\varphi(x, y) := \varphi(x, y) \vee \forall z\, \neg\varphi(x, z)$ and observe that

$$M \models \forall x \in A\, \exists y\, \bar\varphi(x, y).$$

Thus, by Σ_n-collection, there is a set C such that

$$M \models \forall x \in A\, \exists y \in C\, \bar\varphi(x, y).$$

Thus, also, $M \models \forall x \in A \exists y \in C\, \varphi(x, y)$. It follows that

$$B = \{x \in A \mid \exists y \in C\, \varphi(x, y)\}.$$

The formula $\exists y \in C\, \varphi(x, y)$ is a equivalent to a Π_{n-2}-formula by Σ_{n-2}-collection, and so we can use Σ_{n-2}-separation to conclude that B exists.

For Σ_n-replacement, suppose that A is a set and φ is a Π_{n-1}-formula such that $\exists z\ \varphi(x, y, z)$ defines a function $F\colon A \to M$. By Σ_n-collection, there exists a set B such that for all $x \in A$, there is some $y \in B$ with $\exists z\ \varphi(x, y, z)$. Again by Σ_n-collection, there exists a set $C \in M$ such that for all $x \in A$, there is some $z \in C$ with $\exists y\ \varphi(x, y, z)$. Since y is uniquely determined by x,

$$\mathrm{ran}(F) = \{y \in B \mid \exists z \in C\ \varphi(x, y, z)\}.$$

By Σ_{n-1}-collection, the formula $\exists z \in C\ \varphi(x, y, z)$ is equivalent to a Π_{n-1}-formula, and so $\mathrm{ran}(F)$ is a set by Σ_{n-1}-separation.

For Δ_n-separation, suppose that A is a set and $\varphi(x)$, $\psi(x)$ are Π_n-formulas such that

$$\varphi(x) \leftrightarrow \neg\psi(x)$$

holds for all $x \in A$. We need to show that $B := \{x \in A \mid \varphi(x)\}$ is a set. Assume there exists some $y \in A$ with $\varphi(y)$. Since the function $F\colon A \to A$ defined by letting $F(x) = x$ if $\varphi(x)$ holds and $F(x) = y$ if $\psi(x)$ holds is Σ_n-definable, $B = \operatorname{ran}(F)$ is a set by Σ_n-replacement.

Σ_n-recursion along ordinals now follows just like Σ_1-recursion follows from KP, using Σ_n-replacement for the existence of the recursion at limit stages. In fact, the proof works for recursion along Δ_n-definable strongly well-founded relations. □

Thus, in particular, KP_{n+1} implies ZFC^-_n, which in turn implies KP_n.

The fragments ZFC^-_n and KP_n have their own advantages for different situations. The theory ZFC^-_n implies a fragment of Łoś' theorem for Σ_n-formulas. To state this, we first fix some notation. Suppose that M is a weak κ-model and U is a weakly amenable M-ultrafilter. Let N be the ultrapower of M by U and let

$$W = \{[f]_U \subseteq [c_\kappa]_U \mid \{\alpha < \kappa \mid f(\alpha) \in U\} \in U\}$$

be the N-ultrafilter on $[c_\kappa]_U$ derived from U. We know that Łoś' theorem holds for the ultrapower in the language $\in$. We will say that Łoś' theorem *holds* for an assertion $\varphi(x_1, \ldots, x_n)$ in the language with a predicate for the ultrafilter if for every sequence $\vec{f} = \langle f_1, \ldots, f_n\rangle$ with $f_i : \kappa \to M$, the set

$$A_{\vec{f},\varphi} = \{\alpha < \kappa \mid \langle M, \in, U\rangle \models \varphi(f_1(\alpha), \ldots, f_n(\alpha))\} \in M$$

and $\langle N, \in, W\rangle \models \varphi([f_1]_U, \ldots, [f_n]_U)$ if and only if $A_{\vec{f},\varphi} \in U$.

Lemma 3.7. *Suppose M is a weak κ-model and U is an M-ultrafilter. If $\langle M, \in, U\rangle \models \mathrm{ZFC}^-_n$, then Łoś' theorem holds for Σ_n and Π_n-assertions in the language with a predicate for the ultrafilter.*

Proof. Let's first argue that the extended Łoś' theorem is true for all Δ_0-assertions. It is true for atomic formulas by the definition of W and the case of conjunctions is clear. So let's suppose that the assertion is true for some Δ_0-formula $\varphi(x)$ and argue that it is true for $\neg\varphi(x)$. By our assumption, for every function $f : \kappa \to M$ from M, the set

$$A_{\langle f\rangle} = \{\alpha < \kappa \mid \langle M, \in, U\rangle \models \varphi(f(\alpha))\} \in M.$$

Thus, its complement

$$A_{\langle f\rangle,\neg\varphi} = \{\alpha < \kappa \mid \langle M, \in, U\rangle \models \neg\varphi(f(\alpha))\}$$

is in M as well. Now we have that $A_{\langle f\rangle,\neg\varphi} \in U$ if and only if $A_{\langle f\rangle,\varphi} \notin U$ if and only if $\langle N, \in, W\rangle$ does not satisfy $\varphi([f]_U)$ if and only if $\langle N, \in, W\rangle \models \neg\varphi([f]_U)$. So suppose now that the assertion holds for a Δ_0-formula $\varphi(x, y)$. First, observe that for functions $f : \kappa \to M$ and $g : \kappa \to M$ from M, the set

$$A_{\langle f,g\rangle,\varphi} = \{\alpha < \kappa \mid \langle M, \in, U\rangle \models \exists x \in g(\alpha)\, \varphi(x, f(\alpha))\}$$

is in M by Δ_0-separation because $G = \bigcup_{\alpha<\kappa} g(\alpha)$ is a set in M. Suppose next that

$$\langle N, \in, W\rangle \models \exists x \in [g]_U\, \varphi(x, [f]_U).$$

Then there is $[h]_U$ such that

$$\langle N, \in, W\rangle \models \varphi([h]_U, [f]_U) \wedge [h]_U \in [g]_U$$

and so by our assumption, the set

$$A_{\langle f,g,h\rangle,\varphi} = \{\alpha < \kappa \mid \varphi(h(\alpha), f(\alpha)) \wedge h(\alpha) \in g(\alpha)\} \in U,$$

and hence $A_{\langle f,g\rangle,\varphi} \in U$. In the other direction, suppose that $A_{\langle f,g\rangle,\varphi} \in U$. Let $w_G \in M$ be any well-order of G. Using w_G, we can choose for every $\alpha \in A_{\langle f,g\rangle,\varphi}$ the w_G-least $b \in g(\alpha)$ such that $\langle M, \in, U\rangle \models \varphi(b, f(\alpha))$ and define $h(\alpha) = b$. The definition of h uses Δ_0-separation in the structure $\langle M, \in, U\rangle$.

Next, observe that the negation case holds for any formula provided that we have the inductive assumption for that formula, and so it remains to argue the case of the unbounded existential quantifier. So suppose that the inductive assumption holds for a formula $\varphi(x,y)$ of complexity at most Π_{n-1}. The set

$$A_{\langle f\rangle,\varphi}=\{\alpha<\kappa\mid\langle M,\in,U\rangle\models\exists x\,\varphi(x,f(\alpha))\}$$

is in M since $\langle M,\in,U\rangle$ satisfies Σ_n-separation. If $\langle N,\in,W\rangle\models\exists x\,\varphi(x,[f]_U)$, then as before there is some $[h]_U$ such that $\langle N,\in,W\rangle\models\varphi([h]_U,[f]_U)$, so by the inductive assumption $A_{\langle f,h\rangle,\varphi}\in U$, which in turn implies that $A_{\langle f\rangle,\exists x\varphi}\in U$. In the other direction, supposing that $A_{\langle f\rangle,\varphi}\in U$, we use Σ_n-collection to obtain a set C such that for every $\alpha\in A_{\langle f\rangle,\varphi}$, there is $x\in C$ for which $\langle M,\in,U\rangle\models\varphi(x,f(\alpha))$. Let $w_C\in M$ be a well-order of C, and use w_C to chose for every $\alpha\in A_{\langle f\rangle,\varphi}$, the w_C-least b such that $\langle M,\in,U\rangle\models\varphi(b,f(\alpha))$ and define $h(\alpha)=b$. The definition of h uses Σ_n-separation in the structure $\langle M,\in,U\rangle$. □

The next lemma shows that the fragment KP_n is more natural than ZFC^-_n with respect to forming Σ_n-elementary substructures. Namely, we have that any element of a model of KP_n that is transitive and Σ_n-elementary in the model is itself a model of KP_n. This can fail to hold for models of ZFC_n. For instance, it is easy to see that the Σ_1-Skolem hull is Σ_1-definable in any model $L_\alpha\models\mathrm{ZFC}^-_1$, and therefore must be a set in L_α by collection. Thus, every model $L_\alpha\models\mathrm{ZFC}^-_1$ has a proper Σ_1-elementary substructure, but then the least Σ_1-substructure of L_α cannot be a model of ZFC^-_1. In the complexity calculations that follow, we will repeatedly use the observation that Σ_n-collection implies a normal form theorem for Σ_n-assertions, namely that every assertion with a Σ_n-alternation of unbounded quantifiers is equivalent to a Σ_n-assertion.

Lemma 3.8. *Suppose that $M\models\mathrm{ZFC}^-$ and U is a predicate such that $\langle M,\in,U\rangle\models\mathrm{KP}_n$ for some $n\geq 1$. If $\bar M\in M$ is transitive in M and $\langle\bar M,\in,U\cap\bar M\rangle\prec_{\Sigma_n}\langle M,\in,U\rangle$, then $\langle\bar M,\in,U\cap\bar M\rangle\models\mathrm{KP}_n$.*

Proof. Δ_0-separation in $\langle\bar M,\in,U\cap\bar M\rangle$ can be easily verified as follows. Suppose that $A\in\bar M$ and $\varphi(x)$ is a Δ_0-formula. Then, by Δ_0-separation,

$$\langle M,\in,U\rangle\models\exists z\,(\forall x\in z\,x\in A\wedge\forall x\in A(\varphi(x)\leftrightarrow x\in z)).$$

Since this is Σ_1-assertion, we have that it holds in $\langle\bar M,\in,U\cap\bar M\rangle$ by Σ_1-elementarity.

Next, we verify Σ_n-collection. Suppose that

$$\langle\bar M,\in,U\cap\bar M\rangle\models\forall x\in a\,\exists y\,\varphi(x,y,a),$$

for a Σ_n-formula $\varphi(x,y,a)$. Indeed, we can assume without loss of generality that $\varphi(x,y,a)$ is Π_{n-1}. Since $\bar M$ is transitive in M, Σ_n-elementarity yields

$$\langle M,\in,U\rangle\models\forall x\in a\,\exists y\,\varphi(x,y,a).$$

By Σ_n-collection, $\langle M,\in,U\rangle\models\psi(a)$, where

$$\psi(a):=\exists z\forall x\in a\,\exists y\in z\,\varphi(x,y,a).$$

For $n=1$, the formula $\psi(a)$ is Σ_1, and so $\langle\bar M,\in,U\cap\bar M\rangle\models\psi(a)$ by Σ_1-elementarity, verifying Σ_1-collection. For $n>1$, we can assume inductively that we have already verified Σ_{n-1}-collection in $\langle\bar M,\in,U\cap\bar M\rangle$. Under Σ_{n-1}-collection, the formula $\forall x\in a\,\exists y\in z\,\varphi(x,y,a)$ is equivalent to a Σ_n-formula $\bar\psi(z)$. Since $\langle M,\in,U\rangle$ satisfies the Σ_n-formula $\exists z\,\bar\psi(z)$, by Σ_n-elementarity, $\langle\bar M,\in,U\cap\bar M\rangle\models\exists z\,\bar\psi(z)$. Thus, by Σ_{n-1}-collection, $\langle\bar M,\in,U\cap\bar M\rangle\models\psi(a)$, verifying Σ_n-collection. □

Since Σ_0-elementarity is equivalent to just being a submodel for transitive structures, it is not difficult to find a counterexample to Lemma 3.8 for $n=0$.

The Σ_n-*reflection* scheme states that for every true Σ_n-assertion $\varphi(x,a)$ with parameter a, there is a transitive set M containing a such that for every $x\in M$, $M\models\varphi(x,a)$ if and only if $\varphi(x,a)$. If M is a simple weak κ-model and U is a good M-ultrafilter such that $\langle M,\in,U\rangle\models\mathrm{KP}_{n+1}$, then the structure $\langle M,\in,U\rangle$ satisfies a strong form of Σ_n-reflection.

Lemma 3.9. *Suppose that M is a simple weak κ-model and U is a good M-ultrafilter. If $\langle M,\in,U\rangle\models\mathrm{KP}_{n+1}$, for some $n\geq 1$, then for every $A\in M$ there is a κ-model $\bar M\in M$ such that $A\in M$, $\bar M\prec M$,*

$$\langle\bar M,\in,U\cap\bar M\rangle\prec_{\Sigma_n}\langle M,\in,U\rangle,$$

and $\langle\bar M,\in,U\cap\bar M\rangle\models\mathrm{ZFC}^-_n$.

Proof. Recall that, by Lemma 3.2, the structure $\langle M, \in, U\rangle$ has a Δ_1-definable well-order $\lhd_U$. Since U is good, it follows that $\lhd_U$ is externally a well-order. First, let's argue that for every set $X \in M$, there exists a unique set $X^* \in M$ of $\lhd_U$-least witnesses for Σ_n-assertions $\theta(x, a)$ in the language with the ultrafilter and with $a \in X$. Let C be a collecting set for the assertion

$$\forall a \in X\, \forall \theta \in \omega\, \exists y\, \varphi(a, \theta, y),$$

where $\varphi(a, \theta, y)$ is the assertion

$$\Big(y = \emptyset \wedge \forall z\, \neg \mathrm{Tr}^U_n(\theta, \langle z, a\rangle)\Big) \vee \Big(\mathrm{Tr}^U_n(\theta, \langle y, a\rangle) \wedge \forall z (z \lhd_U y \to \neg \mathrm{Tr}^U_n(\theta, \langle z, a\rangle))\Big).$$

The complexity of the assertion $\varphi(a, \theta, y)$ is clearly at most Σ_{n+1} (it is a disjunction of a Π_n and a Σ_n-formula), and so we can apply Σ_{n+1}-collection. Now using Δ_{n+1}-separation for C by Lemma 3.6, we can obtain the desired unique set X^*.

Next, let's argue that every set $X \in M$ can be put into a κ-model $N \in M$ such that $N \prec M$. For every cardinal $\alpha < \kappa$, we have that every set in H_{α^+} is contained in an elementary submodel of size α that is closed under sequences of length $<\alpha$. Thus, the ultrapower of M by U satisfies that every set in $H_{\kappa^+} = M$ is contained in such an elementary submodel, and all these submodels are in M.

Observe that the assertion that N is a κ-model is Δ_1 for the following reasons. The assertion that N is a weak κ-model is Δ_1 because you just have to say that it is transitive, has κ, V_κ as elements, and satisfies ZFC$^-$ (this last gives the unbounded quantifier). The assertion that the model is closed under $<\kappa$-sequences is clearly Π_1, but in fact it is also Σ_1. The Σ_1-assertion is "there is a bijection $f : V_\kappa \to N$ such that for every $g : \xi \to V_\kappa$ in V_κ with $\xi < \kappa$, there exists $x \in N$ such that $x = f \circ g$." Next, observe that the assertion that N is the $\lhd_U$-least κ-model extending a set X is Π_1. Finally, observe that the assertion that $\langle N, \in\rangle \prec \langle M, \in\rangle$ is Δ_1 using the Δ_1-definable truth predicate $\mathrm{Tr}_M(\varphi, x)$ for $\langle M, \in\rangle$.

Now consider the assertion $\psi(s, \lambda)$, which states that s is a sequence of length λ such that:

(1) $s_0 = \{A\}$
(2) $s_\delta = \bigcup_{\xi<\delta} s_\xi$ for limit ordinals δ,
(3) if $\xi+1$ is an even successor ordinal, then $s_{\xi+1} = s^*_\xi$, the unique closure under existential witnesses for Σ_n-assertions,
(4) if $\xi + 1$ is an odd successor ordinal, then $s_{\xi+1}$ is the $\lhd$-least κ-model such that $s_\xi \subseteq s_{\xi+1} \prec M$.

The recursion defining the s_ξ is Δ_{n+1}, so we can use the Σ_{n+1}-recursion scheme to conclude that there exists a function $f : \kappa \to M$ such that $f(\xi) = s_\xi$. Let $\bar M = \bigcup_{\xi<\kappa} f(\xi)$. By construction, $\langle \bar M, \in, U \cap \bar M\rangle \prec_n \langle M, \in, U\rangle$ with $\bar M \prec M$. Since $\bar M$ is a κ-length union of κ-models in M, it is itself a κ-model in M. The model $\bar M$ really is a κ-model by Lemma 2.3. Note that the model $\langle M, \in, U \cap \bar M\rangle$ satisfies Σ_n-reflection by construction because it is a union of Σ_n-elementary substructures. It will follow by the next Lemma 3.10 that Σ_n-reflection implies ZFC^-_n. □

Again, the result can fail for $n = 0$ because being Σ_0-elementary and transitive is equivalent to being a submodel.

Lemma 3.10. *Suppose that $M \models$ ZFC$^-$ and U is a predicate. If $\langle M, \in, U\rangle$ satisfies Σ_n-reflection for some $n \geq 1$, then it is a model of ZFC^-_n.*

Proof. Concerning Σ_n-collection, suppose that

$$\langle M, \in, U\rangle \models \forall x \in a\, \exists y\, \varphi(x, y, a)$$

for a Σ_n-formula $\varphi(x, y, a)$. Let $a \in \bar M \in M$ such that

$$\langle \bar M, \in, U \cap \bar M\rangle \prec_{\Sigma_n} \langle M, \in, U\rangle.$$

Fix $b \in a$. By assumption, $\langle M, \in, U\rangle \models \exists y\, \varphi(b, y, a)$. By elementarity, we have $\langle \bar M, \in, U \cap \bar M\rangle \models \exists y\, \varphi(b, y, a)$. Thus,

$$\langle \bar M, \in, U \cap \bar M\rangle \models \forall x \in a\, \exists y\, \varphi(x, y, a),$$

and so $\bar M$ is the required collecting set for φ.

For Σ_n-separation, fix $a \in M$ and a Σ_n-formula $\psi(x, b)$ in the language with U. Let $a, b \in \bar N$ such that

$$\langle \bar N, \in, U \cap \bar N\rangle \prec_{\Sigma_n} \langle M, \in, U\rangle.$$

The structure $\langle \bar{N}, \in, U \cap \bar{N}\rangle$ reflects $\langle M, \in, U\rangle$ for the formula $\psi(x, b)$. Therefore

$$\{x \in a \mid \langle M, \in, U\rangle \models \psi(x, b)\} = \{x \in a \mid \langle \bar{N}, \in, U \cap \bar{N}\rangle \models \psi(x, b)\},$$

and the right-hand side set exists by separation in M. □

The hypothesis in Lemma 3.9 cannot be reduced to ZFC^-_n, since the claim fails for an $\in$-minimal model $\langle M, \in, U\rangle$ of ZFC^-_n. For the same reason, the result can fail if $n = 0$. The argument above does not work if the ultrapower by U is ill-founded. In this case, the strong well-order $\lhd_U$ might be ill-founded for formulas in the language with U. Thus, it is not clear whether the above lemma holds in the case where U is not good. If U is not good, we can however take, for any set $A \in M$, a substructure of $\langle M, \in, U\rangle$ with a Δ_1-definable true well-order, namely, $L_\alpha[A, U]$, where α is the height of the model. Indeed, replacing a model $\langle M, \in, U\rangle$ by $\langle L_\alpha[A, U], \in, U \cap L_\alpha[A, U]\rangle$ will prove useful in other ways as well. First though we have to verify that this move preserves the theory.

Suppose that M is a weak κ-model and U is an M-ultrafilter such that

$$\langle M, \in, U\rangle \models \mathrm{KP}_n(\mathrm{ZFC}^-_n),$$

for $n \geq 1$, and $\alpha = \mathrm{Ord}^M$. Let $A \in M$ be a subset of κ. Consider the model $\langle L_\alpha[A, \bar{U}], \in, \bar{U}\rangle$, where $U \cap L_\alpha[A, U]$.

Lemma 3.11. $\langle L_\alpha[A, \bar{U}], \in, \bar{U}\rangle \models \mathrm{KP}_n(\mathrm{ZFC}^-_n)$.

Proof. Observe that the $L[A, U]$ construction (up to α) can be carried out in the structure $\langle M, \in, U\rangle$ by Σ_1-recursion.

First, suppose that $\langle M, \in, U\rangle \models \mathrm{KP}_n$. Σ_0-separation clearly holds in $\langle L_\alpha[A, \bar{U}], \in, \bar{U}\rangle$, so it suffices to verify Σ_n-collection. Suppose that for some Π_{n-1}-formula $\varphi(x, a)$,

$$\langle L_\alpha[A, \bar{U}], \in, \bar{U}\rangle \models \forall x \in a\, \exists y\, \varphi(x, y, a).$$

For every $x \in a$, let α_x be the least ordinal such that $x \in L_{\alpha_x}[A, \bar{U}]$ and let β_x be least ordinal such that $L_{\beta_x}[A, \bar{U}]$ has some y witnessing $\varphi(x, y, a)$. Let $a \in L_\lambda[A, \bar{U}]$. Let $\varphi^*(x, y, a)$ be the formula $\varphi(x, y, a)$ relativized to $L_\alpha[A, \bar{U}]$, that is, for every unbounded quantifier we add the assertion that the variable is in $L[A, U]$. Since the assertion $x \in L[A, \bar{U}]$ is Σ_1 over $\langle M, \in, U\rangle$, the formula $\varphi^*(x, y, a)$ is Π_{n-1}. Next, observe that in $\langle M, \in, U\rangle$, the function $f : \lambda \to \alpha$ defined by $f(\xi) = 0$ if $\xi \neq \alpha_x$ for any $x \in a$ and otherwise $f(\alpha_x) = \beta_x$ is Σ_n-definable. Thus, by Σ_n-collection in $\langle M, \in, U\rangle$, there is some $\beta < \alpha$ such that the range of f is contained in β. It follows that $L_\beta[A, U]$ is a collecting set.

Next, suppose that $\langle M, \in, U\rangle \models \mathrm{ZFC}^-_n$. We already showed that $\langle L_\alpha[A, \bar{U}], \in, \bar{U}\rangle$ satisfies Σ_n-collection. Thus, it suffices to verify Σ_n-separation. Fix a Σ_n-formula $\varphi(x, b) := \exists y\, \psi(y, x, b)$, where $\psi(x, y, b)$ is Π_{n-1}, and a set $C \in L_\alpha[A, \bar{U}]$. The set

$$\bar{C} = \{c \in C \mid \langle L_\alpha[A, \bar{U}], \in, \bar{U}\rangle \models \varphi(c, b)\}$$

exists in M by Σ_n-separation in $\langle M, \in, U\rangle$. For every $c \in \bar{C}$, there is $y \in L_\alpha[A, U]$ such that $\langle L_\alpha[A, U], \in, U \cap L_\alpha[A, U]\rangle \models \psi(y, c, b)$. Thus, by Σ_n-collection in $\langle M, \in, U\rangle$, there is an ordinal β such that $L_\beta[A, U]$ already has all the witnesses y for $c \in \bar{C}$. It follows that we can replace the formula $\varphi(x, b)$ by the formula $\exists y \in L_\beta[A, U]\, \psi(y, x, b)$, in verifying separation for the set C and the formula $\varphi(x, b)$ in $\langle L_\alpha[A, U], \in, U \cap L_\alpha[A, U]\rangle$. Since $\exists y \in L_\beta[A, U]\, \psi(y, x, b)$ is equivalent to a Π_{n-1}-formula by Σ_n-collection in $\langle L_\alpha[A, U], \in, U \cap L_\alpha[A, U]\rangle$, we have separation for it by KP_n, which implies Σ_{n-1}-separation by Lemma 3.6. □

We can ensure that $L_\alpha[A, \bar{U}]$ is a weak κ-model by taking $\bar{A}$, instead of A, that codes A and V_κ. However, if M is simple, it is not necessarily the case that $L_\alpha[A, \bar{U}]$ is simple. In this case, letting $\beta = (\kappa^+)^{L_\alpha[A, \bar{U}]}$, we will be able to replace $L_\alpha[A, \bar{U}]$ by the simple weak κ-model $L_\beta[A, \bar{U}]$.

Proposition 3.12. $\langle L_\beta[A, \bar{U}], \in, \bar{U} \cap L_\beta[A, \bar{U}]\rangle \models \mathrm{KP}_n(\mathrm{ZFC}^-_n)$.

Proof. First, let's argue that $L_\beta[A, \bar{U}] = H_{\kappa^+}^{L_\alpha[A, \bar{U}]}$. Fix a set $B \in L_\alpha[A, \bar{U}]$ whose transitive closure $\mathrm{TCl}(B)$ has size at most κ. Fix some limit ordinal γ such that $\mathrm{TCl}(B) \in L_\gamma[A, \bar{U}]$. In $L_\alpha[A, \bar{U}]$, let

$$\langle X, \in, \bar{U} \cap X\rangle \prec \langle L_\gamma[A, \bar{U}], \in, \bar{U} \cap L_\gamma[A, \bar{U}]\rangle,$$

with $|X| = \kappa$, $\kappa + 1 \subseteq X$, and $A, \mathrm{TCl}(B) \in X$. A Σ_1-recursion suffices to construct X because truth in $\langle L_\gamma[A, \bar{U}], \in, \bar{U} \cap L_\gamma[A, \bar{U}]\rangle$ is Δ_1-definable. Note that $U^* = \bar{U} \cap X$ is a set because X has size κ in $L_\alpha[A, \bar{U}]$. Let $\pi : X \to N$ be the Mostowski collapse. Note that $\pi(\mathrm{TCl}(B)) = \mathrm{TCl}(B)$. Since π fixes subsets of κ, we have that π is an isomorphism between $\langle X, \in, U^*\rangle$ and $\langle N, \in, U^*\rangle$. It follows that $N = L_{\bar{\beta}}[A, \bar{U}]$ for some $\bar{\beta}$ of size κ in $L_\alpha[A, \bar{U}]$. Thus, $B \in L_\beta[A, \bar{U}]$.

If $\beta = \alpha$, then we are done. So we can assume that $\beta < \alpha$, and in particular, that $L_\beta[A, \bar{U}]$ is a set in $L_\alpha[A, \bar{U}]$. Since $L_\beta[A, \bar{U}] = H_{\kappa^+}^{L_\alpha[A, \bar{U}]}$, $\langle L_\beta[A, \bar{U}], \in, \bar{U} \cap L_\beta[A, \bar{U}]\rangle$ has the same amount of comprehension as $\langle L_\alpha[A, \bar{U}], \in, \bar{U}\rangle$. Next, we will argue that $\langle L_\beta[A, \bar{U}], \in, \bar{U} \cap L_\beta[A, U]\rangle$ satisfies full collection. Suppose that collection fails in $\langle L_\beta[A, \bar{U}], \in, \bar{U} \cap L_\beta[A, \bar{U}]\rangle$. We can assume without loss of generality that it fails for a formula $\forall \gamma \in \kappa\, \exists y\, \psi(\gamma, y)$. We can use the failure to obtain an injection from κ into β that is definable over $\langle L_\beta[A, \bar{U}], \in, \bar{U} \cap L_\beta[A, \bar{U}]\rangle$. We define the injection by mapping ξ to the least η_ξ such that $\exists y \in L_{\eta_\xi}[A, \bar{U}]\, \psi(\xi, y)$ holds. Now, we use Σ_1-collection in $\langle L_\alpha[A, \bar{U}], \in, \bar{U}\rangle$ to construct a set consisting of bijections $f_\xi : \kappa \to \eta_\xi$. But this set contradicts that $\beta = (\kappa^+)^{L_\alpha[A, \bar{U}]}$. □

The following lemmas will prove useful in later arguments.

Lemma 3.13. *If $\langle L_\alpha[A], A\rangle \models \mathrm{KP}_{n-1}$ for some set A and $n \geq 1$, then it has Σ_n-definable Skolem functions for Σ_n-formulas.*

Proof. Suppose that $\varphi(x, y)$ is the formula $\exists z\ \psi(x, y, z)$, where $\psi(x, y, z)$ is a Π_{n-1}-formula. Let γ be least such that $L_\gamma[A]$ contains witnesses y and z for $\psi(x, y, z)$. Then let y be the $<_{L[A]}$-least in $L_\gamma[A]$ such that there is a witness $z \in L_\gamma[A]$ for which $\psi(x, y, z)$ holds. Since $\langle L_\alpha[A], A\rangle \models \mathrm{KP}$, we can define the $L[A]$-hierarchy. Since $\langle L_\alpha[A], A\rangle \models \mathrm{KP}_{n-1}$, the statement that $\psi(x, y, z)$ fails for all $<_{L[A]}$-smaller z is a Σ_{n-1}-statement. □

As a consequence of having Σ_n-definable Skolem functions for Σ_n-formulas, we have that the Σ_n-Skolem hull (over any collection of parameters) taken using these functions is Σ_n-elementary in $\langle L_\alpha[A], A\rangle$.

The following proof is based on an argument of Philip Welch which shows that the least admissible ordinal $\alpha > \omega_1$ has uncountable cofinality.

Lemma 3.14. *Suppose that M is a simple weak κ-model and U is an M-ultrafilter such that $M = L_\alpha[A, U]$ for some $A \subseteq \kappa$ and $\langle M, \in, U\rangle \models \mathrm{KP}_n$. Then $\langle M, \in, U\rangle$ has a transitive Σ_n-elementary substructure $\langle \bar{M}, \in, U \cap \bar{M}\rangle$ such that $\bar{M}$ is a κ-model.*

Proof. We can assume without loss of generality that A codes V_κ. Let $\langle \bar{M}, \in, A, U \cap \bar{M}\rangle$ be the Σ_n-Skolem hull of $\kappa + 1$, using the Σ_n-definable Skolem functions (from Lemma 3.13) in $\langle L_\alpha[A, U], \in, A, U\rangle$. By Lemma 3.13, we have

$$\langle \bar{M}, \in, A, U \cap \bar{M}\rangle \prec_{\Sigma_n} \langle L_\alpha[A, U], \in, A, U\rangle.$$

Since every set in $L_\alpha[A, U]$ has size κ, it follows that $\bar{M}$ is transitive. Thus, by Lemma 3.8,

$$\langle \bar{M}, \in, A, U \cap \bar{M}\rangle \models \mathrm{KP}_n.$$

By Σ_n-elementarity, we have $\bar{M} = L_{\bar{\alpha}}[A, U]$ for some $\bar{\alpha} \leq \alpha$ and the structure

$$\langle L_{\bar{\alpha}}[A, U], \in, U \cap L_{\bar{\alpha}}[A, U]\rangle$$

has the additional property that every element is Σ_n-definable using parameters from $\kappa + 1$.

We claim that $\bar{M}^{<\kappa} \subseteq \bar{M}$. To see this, suppose that $\vec{x} = \{x_i \mid i < \xi\}$ is a sequence of elements of $\bar{M}$ for some $\xi < \kappa$. Let each x_i be definable in $\langle L_{\bar{\alpha}}[A, U], \in, A, U \cap L_{\bar{\alpha}}[A, U]\rangle$ by the Σ_n-formula φ_i using parameters $\nu_i < \kappa$ and κ. Since $V_\kappa \in \bar{M}$, the sequences $\{\nu_i \mid i < \xi\}$ and $\{\varphi_i \mid i < \xi\}$ are both in $\bar{M}$. Thus, using the Σ_n-definable Σ_n-truth predicate $\mathrm{Tr}_n^U(x)$ (which exists by Lemma 3.5), the sequence $\{x_i \mid i < \xi\}$ is Σ_n-definable over $\langle L_{\bar{\alpha}}[A, U], \in, A, U \cap L_{\bar{\alpha}}[A, U]\rangle$. By Σ_n-collection, there must be some $\beta < \bar{\alpha}$ such that all x_i are in $L_\beta[A, U]$. At this point, we would be done if we had Σ_n-separation, but we have only Δ_n-separation available. Thus, we need to do some more work to reduce the complexity of the formula defining the sequence $\{x_i \mid i < \xi\}$.

For each $i < \xi$, let $\varphi_i(x, \nu_i, \kappa) := \exists y\, \psi_i(x, \nu_i, \kappa, y)$, where ψ is a Π_{n-1}-formula. Next, we use Σ_n-collection on the formula

$$\forall i < \xi\, \exists y\, \mathrm{Tr}_n^U(\langle \exists y\, \psi_i(x, z, w, y), x_i, \nu_i, \kappa)$$

to obtain a set Y containing, for each $i < \xi$, a witness y_i such that $\psi_i(x_i, \nu_i, \kappa, y_i)$ holds. Using the set Y, we can now reduce the complexity of the definition of the sequence to a Π_{n-1}-formula. □

Thus, in particular, we will be able to replace a model $\langle M, \in, U\rangle \models \mathrm{KP}_n$ with a κ-model with the same properties without an increase in consistency strength.

The next result will allow us to separate large cardinal notions defined using ZFC^-_n and KP_n in Section 5.

Lemma 3.15. *Suppose that M is a simple weak κ-model and U is an M-ultrafilter such that $\langle M, \in, U\rangle \models \mathrm{ZFC}^-_n$. Then for every set $A \in M$, there is a weak κ-model $N \in M$ with $A \in N$ such that $\langle N, \in, U \cap N\rangle \models \mathrm{KP}_n$.*

Proof. By shrinking the model, if necessary, we can assume without loss of generality that $M = L_\alpha[A, U]$. By Lemma 3.14, $\langle M, \in, U\rangle$ has a transitive Σ_n-elementary substructure $\langle \bar{M}, \in, U \cap \bar{M}\rangle$ such that $\bar{M}$ is a κ-model. If $\bar{M} \in M$, then we are done. Otherwise, $\bar{M} = M$. Observe that for any set $B \in M$, the set

$$T_B = \{\langle \varphi, b\rangle \mid b \in B \text{ and } \mathrm{Tr}^U_n(\varphi, b)\}$$

is in M by Σ_n-separation. Also, using Σ_n-collection and the set T_B, M has a set B^* such that for every Σ_n-formula $\exists y\, \psi(x, y)$ and $b \in B$ such that $\langle M, \in, U\rangle \models \exists y\, \psi(y, b)$, there is $y_b \in B^*$ such that $\langle M, \in, U\rangle \models \psi(b, y_b)$. Let $M_0 \in M$ be any weak κ-model with $A \in M_0$. Let $M_1 = \mathrm{TCl}(M^*_0)$, and more generally, given M_n, let $M_{n+1} = \mathrm{TCl}(M^*_n)$. Let $N = \bigcup_{n<\omega} M_n$, and observe that $N \in M$ by closure. By construction N is transitive and

$$\langle N, \in, U \cap N\rangle \prec_{\Sigma_n} \langle M, \in, U\rangle,$$

and so $\langle N, \in, U \cap N\rangle \models \mathrm{KP}_n$ by Lemma 3.8. □

4. Baby versions

We now define n-baby measurable cardinals by slightly simplifying the notions of Bovykin and McKenzie [BM] and we further introduce some related large cardinal notions.

Definition 4.1. A cardinal κ is:

1. *faintly n-baby measurable* if every $A \subseteq \kappa$ is an element of a weak κ-model M for which there is an M-ultrafilter such that $\langle M, \in, U\rangle \models \mathrm{ZFC}^-_n$.
2. *weakly n-baby measurable* if (1) holds and in addition, U is good.
3. *n-baby measurable* if (2) holds and in addition, M is a κ-model.
4. *$[n]$-baby measurable* if (3) holds but with KP_n instead of ZFC^-_n.

Notions (1)-(3) have variants where ZFC^-_n is replaced by ZFC^-. We then omit n from the notation. We will see from Lemma 5.1 below that the concepts of faintly and weakly $[n]$-baby measurable cardinals are equivalent to (4).

We can further assume that all weak κ-models involved in the definitions above are simple. While the M-ultrafilter in the definition of (n-)faintly baby measurable cardinals need not be good, the notion will still turn out to be quite strong.

The faintly (n-)baby measurable cardinals are analogues of completely ineffable cardinals, but with stronger M-ultrafilters. The weakly (n-)baby measurable cardinals are analogues of 1-iterable cardinals, but with stronger M-ultrafilters. The (n-)baby measurable cardinals are analogues of strongly Ramsey cardinals, but with stronger M-ultrafilters.

We will vary the closure condition to study the gap between weakly n-baby measurable and n-baby measurable cardinals. For this purpose, it is useful to discard the requirement that M be transitive and require that M is elementary in some H_θ as in [HS18].

Definition 4.2. Suppose that κ is a cardinal and $\omega \leq \alpha \leq \kappa$ is a regular cardinal.[8] κ is called *(α, n)-baby measurable* if for every $A \subseteq \kappa$ and arbitrarily large θ there is a $<\alpha$-closed basic weak κ-model $M \prec H_\theta$ with $A \in M$ and a (good) M-ultrafilter U such that $\langle H^M_{\kappa^+}, \in, U\rangle \models \mathrm{ZFC}^-_n$.

κ is called *faintly (ω, n)-baby measurable* if the M-ultrafilter is not required to be good.

κ is called *α-baby measurable* if we replace ZFC^-_n by ZFC^- in the above definition.

[8] We use the notation α for cardinals following [HS18] with the motivation that these notions can be characterissed by games of length α and they could thus, in principle, be generalized to ordinals.

Note that we only need to explicitly state that the M-ultrafilter U is good in the case that $\alpha = \omega$. Further variants such as $(<\alpha, n)$-*baby measurable* are used below with the obvious meaning.

The (α, n)-baby measurable (α-baby measurable) cardinals are analogues of the α-Ramsey cardinals, but with stronger M-ultrafilters.

The reflective (α, n)-baby measurable (α-baby measurable) cardinals strengthen this notion by requiring strong Σ_n-reflection (elementary reflection) in the structure $\langle H_{\kappa^+}, \in, U\rangle$.

Definition 4.3. Suppose that κ is a cardinal and $\omega \leq \alpha \leq \kappa$ is a regular cardinal. κ is called *reflective* (α, n)-*baby measurable* if for every $A \subseteq \kappa$ and arbitrarily large θ there is a $<\alpha$-closed basic weak κ-model model $M \prec H_\theta$ with $A \in M$ and a (good) M-ultrafilter U such that for every $B \subseteq \kappa$ in M, there is a κ-model $\bar{M} \in M$ with $B \in \bar{M}$ such that $\langle \bar{M}, \in, U \cap \bar{M}\rangle \prec_{\Sigma_n} \langle H^M_{\kappa^+}, \in, U\rangle$ and $\langle \bar{M}, \in, U \cap \bar{M}\rangle \models \mathrm{ZFC}^-_n$.

κ is called *faintly reflective* (ω, n)-*baby measurable* if the M-ultrafilter is not required to be good.

κ is called *reflective* α-*baby measurable* if we replace Σ_n-elementarity by full elementarity in the above definition.

Note that we get $\langle H^M_{\kappa^+}, \in, U\rangle \models \mathrm{ZFC}^-_n$ for free by Lemma 3.10. These notions will have analogous game theoretic definitions similar to that of α-Ramsey cardinals. We define these games in Section 6.

5. The hierarchies

In this section, we will show where the various (non-game related) notions defined above fit into the large cardinal hierarchy.

First, we show that there are no faint or weak versions of the $[n]$-baby measurable cardinals because this is one of those rare instances where the closure on the model (and hence the well-foundedness of the ultrapower) comes for free.

Lemma 5.1. *Suppose that M is a weak κ-model and U is an M-ultrafilter such that $\langle M, \in, U\rangle \models \mathrm{KP}_n$ for some $n \geq 0$. Then for every $A \subseteq \kappa$ in M, there is a κ-model $\bar{M} \subseteq M$ with $A \in \bar{M}$ such that $\langle \bar{M}, \in U \cap \bar{M}\rangle \models \mathrm{KP}_n$.*

Proof. By a sequence of lemmas at the end of Section 2, we can assume without loss of generality that $M = L_\alpha[A, U]$ with $\alpha = (\kappa^+)^{L_\alpha[A,U]}$. Thus, by Lemma 3.14, there is a κ-model $\bar{M} \subseteq M$ with $A \in \bar{M}$ such that $\langle \bar{M}, \in, U \cap \bar{M}\rangle \models \mathrm{KP}_n$. □

Since the theories ZFC^-_0 and KP_0, and KP_1 are all the same, the following large cardinals are all equivalent:

(1) [0]-baby measurable cardinals
(2) [1]-baby measurable cardinals
(3) faintly 0-baby measurable cardinals
(4) weakly 0-baby measurable cardinals
(5) 0-baby measurable cardinals

Next, we show that the faintly 0-baby measurable cardinals, which are like 0-iterable cardinals with just the additional assumption of Δ_0-collection, are stronger than Ramsey cardinals, and hence cannot exist in L.

Theorem 5.2. *If κ is faintly 0-baby measurable, then κ is a strongly Ramsey limit of strongly Ramsey cardinals.*

Proof. Suppose that M is a simple weak κ-model and U is an M-ultrafilter with $\langle M, \in, U\rangle \models \mathrm{KP}$. Let N be the, not necessarily well-founded, ultrapower of M with U.

We first show that κ is strongly Ramsey in M. Fix any $A \subseteq \kappa$ in M. We construct a continuous increasing sequence $\langle M_\alpha \mid \alpha \leq \kappa\rangle$ with $A \in M_0$ by a Σ_1-recursion in $\langle M, \in, U\rangle$. Note that the Σ_1-recursion scheme holds in $\langle M, \in, U\rangle$ by Lemma 3.6. Let $M_0 \in M$ be arbitrary with $A \in M_0$. For even $\alpha < \kappa$, let $M_{\alpha+1} := M_\alpha \cup \{\{U \cap M_\alpha\}\}$. For odd $\alpha < \kappa$, let $M_{\alpha+1} \prec M$ be a κ-model from the perspective of M, with $M_\alpha \in M_{\alpha+1}$, given by a Δ_1-definable choice function for the relation R that consists of all pairs $(x, y) \in M$, where $x \in y$ and $(y, \in)$ is an elementary substructure of $\langle M, \in\rangle$ that is closed under sequences of length less than κ. Such a choice function exists by Lemma 3.3 since the relation R exists in N. For all limits $\alpha \leq \kappa$, let $M_\alpha := \bigcup_{\bar{\alpha}<\alpha} M_{\bar{\alpha}}$. In M, we have that M_κ is a κ-model with $A \in M_\kappa$ and $U \cap M_\kappa$ is a weakly amenable M_κ-ultrafilter. Since A was arbitrary, κ is strongly Ramsey in M. By Proposition 2.3, M_κ is really a κ-model, and so, since M was arbitrary, κ is strongly Ramsey. Moreover, κ is strongly Ramsey in N, and hence it is a limit of strongly Ramsey cardinals by elementarity. □

Thus, in particular, the existence of faintly 0-baby measurable cardinals already implies that $0^\#$ exists. In fact, faintly 0-baby measurable cardinals have higher consistency strength.

Theorem 5.3. *If there is a faintly* 0*-baby measurable cardinal, then there is a model of* ZFC *with a κ-Ramsey cardinal κ.*

Proof. Suppose that $\langle M, \in, U\rangle \models \text{KP}$ where M is a simple weak κ-model and U is an M-ultrafilter, and let N be the (not necessarily well-founded) ultrapower of M by U. It suffices to show that κ is κ-Ramsey in $V^N_{j(\kappa)}$. Otherwise the challenger has a winning strategy σ in the game $\mathsf{RamseyG}^{\kappa^+}_\kappa(\kappa)$ in $V^N_{j(\kappa)}$ (equivalently in N). The element σ of N is represented by some equivalence class $[S]_U$. Since the ultrapower is κ-powerset preserving, σ is a map from sequences of elements of M to elements of M and all proper initial segments of runs are in M. By coding, we can assume that elements of σ are coding subsets of κ (elements of $\mathcal{C}$ as defined in the proof of Lemma 3.2), which code pairs (s, y), where s is a sequence of plays by the judge and y is the response to the last move. In M, we will use U to play against σ and build a run of the game $\mathsf{RamseyG}^{\kappa^+}_\kappa(\kappa)$ won by the judge, contradicting that σ was a winning strategy. In detail, the judge responds to the model M_α played by the challenger in round α by playing $U \cap M_\alpha$. The challenger plays according to σ. The assertion that the challenger's response to x is y according to σ translates to asking whether the set $\{\alpha < \kappa \mid f_A(\xi) \in S(\xi)\}$ is an element of U, where A is the coding subset for pair (s, y) with s being the sequence of the judge's moves so far ending in x, and $f_A(\xi) = A \cap \xi$ for every $\xi < \kappa$. The above recursion is thus Δ_1, so has a solution M. □

Note that the least 0-baby measurable cardinal κ cannot be super Ramsey because being 0-baby measurable is a property of H_{κ^+} and would therefore be reflected down if κ was super Ramsey.

Next, we show that surprisingly the hierarchies of the variants of n-baby measurable cardinals are intertwined both in the case of ZFC^-_n and that of KP_n.

Theorem 5.4. *For any $n \geq 1$, a faintly n-baby measurable κ is a limit of $[n]$-baby measurable cardinals.*

Proof. Fix $A \subseteq \kappa$ and choose a weak κ-model M, with $A \in M$, for which there is an M-ultrafilter U such that $\langle M, \in, U\rangle \models \text{ZFC}^-_n$. By the sequence of lemmas at the end of Section 2, we can assume without loss of generality that $M = L_\alpha[A, U]$ and $\alpha = (\kappa^+)^{L_\alpha[A,U]}$. By Lemma 3.15, there is a weak κ-model $N \in M$ with $A \in N$ such that $\langle N, \in, U \cap N\rangle \models \text{KP}_n$. Since U is an M-ultrafilter and N is a set in M, $U \cap N$ is countably complete, and hence good. It follows κ is $[n]$-baby measurable.

An analogous argument also shows that κ is $[n]$-baby measurable in the, not necessarily well-founded, ultrapower of M by U. Thus, κ is a limit of $[n]$-baby measurable cardinals. □

Theorem 5.5. *For any $n \geq 1$, a faintly $[n+1]$-baby measurable κ is an n-baby measurable limit of n-baby measurable cardinals.*

Proof. Fix $A \subseteq \kappa$ and choose a weak κ-model M, with $A \in M$, for which there is an M-ultrafilter U such that $\langle M, \in, U\rangle \models \text{KP}_{n+1}$. We can assume without loss of generality that $M = L_\alpha[A, U]$ and thus has a Δ_1-definable true well-order. Thus, the proof of Lemma 3.9 shows that there is a κ-model $\bar{M} \in M$ such that $\langle \bar{M}, \in, U \cap \bar{M}\rangle \models \text{ZFC}^-_n$. Since the subset A was arbitrary, we have verified that κ is n-baby measurable. Since κ is n-baby measurable in the ultrapower N, we also have that κ is a limit of n-baby measurable cardinals. □

Theorem 5.6. *For any $n \geq 1$, a weakly n-baby measurable cardinal κ is a limit of faintly n-baby measurable cardinals.*

Proof. Suppose that $n \geq 1$, M is a weak κ-model, and U is a good M-ultrafilter such that $\langle M, \in, U\rangle \models \mathrm{ZFC}^-_n$. We claim that κ is faintly n-baby measurable in N, and hence a limit of faintly n-baby measurable cardinals.

Recall the structure $\langle M, \in, U\rangle$ has a Σ_n-definable truth predicate $\mathrm{Tr}^U_n(x)$ for Σ_n-formulas (by Lemma 3.5). For the future, we fix a canonical Σ_n-formula $\exists y\, \theta(x, y)$ defining $\mathrm{Tr}^U_n(x)$. Observe that for any set $B \in M$, the set

$$T_B = \{\langle \varphi, b\rangle \mid b \in B \text{ and } \mathrm{Tr}^U_n(\langle \varphi, b\rangle)\}$$

is in M by Σ_n-separation. Also, using Σ_n-collection and the set T_B, M has a set B^* such that for every Σ_n-formula $\exists y\, \psi(x, y)$ and $b \in B$, there is $y_b \in B^*$ such that $\langle M, \in, U\rangle \models \psi(b, y_b)$.

We say that a formula is in Σ_n- or Π_n-*normal form* if it has a block of at most n alternating quantifiers preceding a Σ_0-formula. We can assume that the Σ_n-formula $\exists y\, \theta(x, y)$ defining $\mathrm{Tr}^U_n(x)$ is in normal form. We denote the standard way of converting a negation $\neg\varphi$ of a Σ_n-formula φ in normal form to a logically equivalent Π_n-formula in normal form by $\neg^*\varphi$. We denote the standard way of converting a conjunction $\varphi \wedge \psi$ of two Σ_n- or Π_n-formulas φ and ψ into a logically equivalent Σ_n- or Π_n-formula in normal form by $\varphi \wedge^* \psi$.

Given a structure $\bar{M}$, let us call a set $T \subseteq \bar{M}$ a Σ_n-*pseudo truth predicate* if it consists of pairs $\langle \varphi, \bar{a}\rangle$, where φ is a Σ_n- or Π_n-formula in normal form and $\bar{a} \in \bar{M}$ is a finite tuple, and T satisfies the following conditions. We will simplify the notation by writing $\varphi(\bar{a}) \in T$ instead of $\langle \varphi, \bar{a}\rangle \in T$.

(1) For every Σ_0-formula $\varphi(\bar{x})$ and finite tuple $\bar{a} \in \bar{M}$, $\varphi(\bar{a}) \in T$ if and only if $\bar{M} \models \varphi(\bar{a})$.
(2) For every Σ_n-formula $\varphi(\bar{x})$ and finite tuple $\bar{a} \in \bar{M}$, either $\varphi(\bar{a}) \in T$ or $\neg^*\varphi \in T$.
(3) If $\forall x\, \psi(x, \bar{a}) \in T$ for some finite tuple $\bar{a} \in \bar{M}$, then $\psi(b, \bar{a}) \in T$ for every $b \in \bar{M}$.

Note that what separates T from an actual truth predicate is the omission of the existential quantifier condition that witnesses for existential statements are provided.

Take any $A \subseteq \kappa$ in M. We construct a tree $\mathcal{T}$ of height ω in N whose branches union up to produce structures $\langle \bar{M}, \in, \bar{U}\rangle \models \mathrm{ZFC}^-_n$, where $\bar{M}$ a weak κ-model with $A \in \bar{M}$ and $\bar{U}$ is an $\bar{M}$-ultrafilter on κ. We will then proceed to show that $\mathcal{T}$ has a branch in V and hence also in N by the absoluteness of well-foundedness.

A node on level m of $\mathcal{T}$ is going to be a sequence $\{\langle M_i, \in, U_i, T_i\rangle \mid i < m\}$ satisfying the following conditions:

(1) $A \in M_0$.
(2) Each M_i is a weak κ-model, U_i is an M_i-ultrafilter, and T_i is a Σ_n-pseudo truth predicate for the structure $\langle M_i, \in, U_i\rangle$.
(3) For $i < j < m$, $M_i \subseteq M_j$, $U_i \subseteq U_j$, and $T_i \subseteq T_j$.
(4) For each $i < m$, $U_i \cap M_i \in M_{i+1}$.
(5) For each $j < m - 1$, Σ_n-formula $\exists y\, \psi(x, y) \in T_j$, and $a \in M_j$, there is some $b \in M_{j+1}$ such that $\psi(a, b) \in T_{j+1}$.
(6) For each $j < m - 1$, M_{j+1} contains elements $X_j \subseteq M_j$ and $Y_j \subseteq M_{j+1}$ with the following properties:
    (a) For each $x \in X_j$, $\exists y\, [\theta(x, y) \wedge^* (y \in Y_j)]$ is in T_{j+1}.
    (b) For each $x \in M_j \setminus X_j$, $\forall y\, \neg^*\theta(x, y)$ is in T_{j+1}.

Suppose that $\{\langle M_i, \in, U_i\rangle \mid i < \omega\}$ is a branch of $\mathcal{T}$. Let $\bar{M} = \bigcup_{i<\omega} M_i$, $\bar{U} = \bigcup_{i<\omega} U_i$ and $\bar{T} = \bigcup_{i<\omega} T_i$. Let's argue that $\bar{T}$ is an actual Σ_n-truth predicate for $\langle \bar{M}, \in, \bar{U}\rangle$. We argue by induction on complexity of formulas. Suppose that $\varphi(x)$ is a Σ_0-formula. We have that $\langle \bar{M}, \in, \bar{U}\rangle \models \varphi(a)$ if and only if $\langle M_i, \in, U_i\rangle \models \varphi(a)$, where $a \in M_i$, by absoluteness if and only if $\varphi(a) \in T_i \subseteq \bar{T}$. Next, consider a formula $\exists y\, \psi(x, y)$ such that the hypothesis holds for $\psi(x, y)$ by the inductive assumption. Suppose first that $\langle \bar{M}, \in, \bar{U}\rangle \models \exists y\, \psi(a, y)$. Then $\langle \bar{M}, \in, \bar{U}\rangle \models \psi(b, a)$ for some b. By the inductive assumption $\psi(a, b)$ is in $\bar{T}$ and hence in T_j for some j. Thus, $\forall y\, \neg^*\psi(a, y)$ cannot be in T_j. It follows that $\exists y\, \psi(a, y) \in T_j \subseteq \bar{T}$. Finally, suppose that $\exists y\, \psi(a, y)$ is in $\bar{T}$ and hence in some T_j. Then there is $b \in M_{j+1}$ such that $\psi(a, b)$ is in $T_j \subseteq T$. Thus, by our inductive assumption, $\langle \bar{M}, \in, \bar{U}\rangle \models \psi(a, b)$.

We now show that $\langle \bar{M}, \in, \bar{U}\rangle \models \mathrm{ZFC}^-_n$. Since $\bar{U}$ is amenable to $\bar{M}$ by (4), $\exists y\, \theta(\langle \varphi, a\rangle, y)$ is equivalent to $\exists x\, \varphi(x, a)$ for any Π_{n-1}-formula $\varphi(x, y)$ and any $a \in \bar{M}$. We start with Σ_n-separation. Fix a Σ_n-formula $\exists z\, \psi(x, w, z)$ and sets $a, b \in \bar{M}$. We need to show that the set $X = \{x \in a \mid \exists z\, \psi(x, b, z)\}$ is in $\bar{M}$. Let $a, b \in M_j$. The set $X_j \in M_{j+1}$ consists of all $x \in M_j$

for which $\exists y\,\theta(x,y) \in T_{j+1}$. Since $\bar{T}$ is a truth predicate for $\bar{M}$, the set X_j consists of all $x \in M_j$ for which $\exists y\,\theta(x,y)$ holds in $\bar{M}$. We can use separation in $\langle M_{j+1}, \in\rangle$ to obtain the set $\{x \in a \mid \langle \exists z\,\psi(x,b,z), a\rangle \in X_j\}$, which is precisely the required set X. Next, let's argue that Σ_n-collection holds in $\langle \bar{M}, \in, \bar{U}\rangle$. Fix a Σ_n-formula $\exists y\,\rho(x,y)$ and a set $a \in \bar{M}$ such that $\langle \bar{M}, \in, \bar{U}\rangle \models \forall x \in a\,\exists y\,\rho(x,y)$. We need to find a set $b \in \bar{M}$ such that for every $x \in a$, there is $y_x \in b$ such that $\langle \bar{M}, \in, \bar{U}\rangle \models \rho(x, y_x)$. Let $a \in M_j$. The set $Y_j \in M_{j+1}$ consists of all witnesses y for $x \in M_j$ for the formula $\theta(x,y)$. Since $\exists y\,\theta(x,y)$ was the canonical Σ_n-definition of Σ_n-truth, we have that $\theta(b,c)$ holds if and only if $b = \langle d, \exists z\,\psi(x,z)\rangle$ and $\psi(d,c)$ holds. Thus, in particular, the set Y_j already contains all witnesses for Σ_n-formulas with parameters in M_j true in $\bar{M}$.

It remains to argue that the tree $\mathcal{T}$ has a branch in V. Let M_0 be any weak κ-model in M with $A \in M_0$, let $U_0 = M_0 \cap U$, and let $T_0 = T_{M_0}$, where T_{M_0} denotes the restriction of the Σ_n-truth predicate for $\langle M, \in, U\rangle$ to formulas with parameters in M_0 constructed earlier using Σ_n-separation. Let $M_0^* \in M$ contain witnesses for every Σ_n-formula with parameters in M_0. The set X_0 exists by Σ_n-separation applied to the formula $\exists y\,\theta(x,y)$ and M_0. The set Y_0 is then the collecting set of witnesses for the set X_0 and the formula $\exists y\,\theta(x,y)$. So let M_1 be any weak κ-model in M such that $M_0^* \subseteq M_1$ and $X_0, Y_0 \in M_1$. Moreover, let $U_1 = U \cap M_1$ and $T_1 = T_{M_1}$. We define the structures $\langle M_n, \in, U_n, T_n\rangle$ for $n \geq 1$ analogously. □

Next, we show that weakly n-baby measurable cardinals are weaker than n-baby measurable cardinals.

Theorem 5.7. *An n-baby measurable cardinal is a limit of weakly n-baby measurable cardinals.*

Proof. Suppose that M is a κ-model such that $\langle M, \in, U\rangle \models \mathrm{ZFC}_n^-$. We show that κ is weakly n-baby measurable in the ultrapower N of M by U, and hence a limit of weakly n-baby measurable cardinals. As we already argued in the proof of Theorem 5.6, ZFC_n^- implies that for every set B, there is a set B^* containing for every Σ_n-formula $\varphi(x) := \exists y \psi(x,y)$ and $b \in B$, a witness y_b such that $\psi(b, y_b)$ holds in $\langle M, \in, U\rangle$.

Fix $A \subseteq \kappa$ in M. Let $M_0 \in M$ be any weak κ-model with $A_0 \in M_0$. Let M_0^*, as above, contain witnesses for all Σ_n-formulas with parameters from M_0 that are true over $\langle M, \in, U\rangle$. Let $M_1 \in M$ be any weak κ-model with $M_0^* \subseteq M_1$. Given M_n, we define M_{n+1} analogously. Let $\bar{M}_0 = \bigcup_{n<\omega} M_n$. Then clearly $\langle \bar{M}_0, \in, U \cap \bar{M}\rangle \prec_{\Sigma_n} \langle M, \in, U\rangle$. By closure, the sequence $\{M_n \mid n < \omega\}$ is in M, and hence so is $\bar{M}_0$. Next, we repeat the process, starting with $\bar{M}_0$ instead of A, to build a model $\bar{M}_1 \in M$ such that

$$\langle \bar{M}_0, \in, U \cap \bar{M}_0\rangle \prec_{\Sigma_n} \langle \bar{M}_1, \in, U \cap \bar{M}_1\rangle \prec_{\Sigma_n} \langle M, \in, U\rangle.$$

Continuing in this manner, we define a Σ_n-elementary chain of models $\langle \bar{M}_n, \in, U \cap \bar{M}_n\rangle$, and observe that the sequence must be in M by closure. Let $\bar{M} = \bigcup_{n<\omega} \bar{M}_n$. Then, by construction, $\langle \bar{M}, \in, U \cap \bar{M}\rangle$ satisfies Σ_n-reflection, and so satisfies ZFC_n^- by Lemma 3.10. □

Note that the above proof only required the model M to be closed under countable sequences.

6. Games

The distinction between weakly n-baby measurable and n-baby measurable cardinals (see Proposition 5.7) suggests that one can obtain different large cardinal notions by varying the closure of the models. We will show that in fact closure properties induce a hierarchy between these two notions using games analogous to the ones for α-Ramsey cardinals from [HS18]. We now describe games with perfect information associated to the (α, n)-baby measurable cardinals and their variants. All games have two players, the challenger and the judge. Suppose that κ, α and θ are regular cardinals with $\omega \leq \alpha \leq \kappa < \theta$.

Definition 6.1. The game $\mathsf{G}_\alpha^{\theta,n}(\kappa)$ proceeds for α-many steps. The challenger starts the game and plays a basic κ-model $M_0 \prec H_\theta$. The judge responds by playing a structure $\langle N_0, \in, U_0\rangle$, where $P^{M_0}(\kappa) \subseteq N_0$ is a κ-model and U_0 is an N_0-ultrafilter. At stage γ, the challenger plays a basic κ-model $M_\gamma \prec H_\theta$ such that

$$\{\langle N_\xi, \in, U_\xi\rangle \mid \xi < \gamma\} \in M_\gamma,$$

and the judge responds with a structure $\langle N_\gamma, \in, U_\gamma\rangle$ such that $P^{M_\gamma}(\kappa) \subseteq N_\gamma$ is a κ-model and U_γ is an N_γ-ultrafilter extending $\bigcup_{\xi<\gamma} U_\gamma$. After α-many steps, let $M = \bigcup_{\xi<\alpha} M_\xi$ and $U = \bigcup_{\xi<\alpha} U_\xi$. The judge wins the game if she was able to play for α-many steps such that at the end U is a (good) M-ultrafilter and $\langle H^M_{\kappa^+}, \in, U\rangle \models \mathrm{ZFC}^-_n$. Otherwise, the challenger wins.

Let $\mathsf{faintG}^{\theta,n}_\omega(\kappa)$ be the analogous game where we do not require the M-ultrafilter U to be good.

Note that this game is played as the game $\mathsf{Ramsey}\bar{\mathsf{G}}^\theta_\alpha(\kappa)$, but, while in that game in order to win, the judge needed to ensure that U is a (good) M-ultrafilter, here the judge also needs to ensure that $\langle H^M_{\kappa^+}, \in, U\rangle \models \mathrm{ZFC}^-_n$. Note also that the M-ultrafilter U is automatically good by closure for uncountable α.

Let's argue that $H^M_{\kappa^+} = \bigcup_{\xi<\alpha} N_\xi$ is the union of the judge's moves. If $A \in H^M_{\kappa^+}$, then A can be coded by a subset of κ in some M_ξ, and hence $A \in N_\xi$. If $A \in N_\xi$, then $A \in H_{\kappa^+}$, and M knows this because $M \prec H_\theta$.

We will characterize (n,α)-baby measurable cardinals by the statement that the challenger does not have a winning strategy in the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$ and faintly (ω,n)-baby measurable cardinals by the statement that the challenger does not have a winning strategy in the game $\mathsf{faintG}^{\theta,n}_\omega(\kappa)$.

The next game strengthens the winning requirement that $\langle H^M_{\kappa^+}, \in, U\rangle \models \mathrm{ZFC}^-_n$ to the requirement that Σ_n-reflection holds.

Definition 6.2. The game $\mathsf{RG}^{\theta,n}_\alpha(\kappa)$ is defined just like the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$, except that the judge has to ensure that (1) $\langle N_\gamma, \in, U_\gamma\rangle \models \mathrm{ZFC}^-_n$ and (2) $\langle N_\xi, \in, U_\xi\rangle \prec_{\Sigma_n} \langle N_\gamma, \in, U_\gamma\rangle$ for all $\xi < \gamma$ in move γ. The game $\mathsf{faintRG}^{\theta,n}_\omega(\kappa)$ is defined analogously.

As before, observe that $H^M_{\kappa^+} = \bigcup_{\xi<\alpha} N_\xi$. Moreover, requiring elementarity between the moves ensures that $\langle N_\xi, \in, U_\xi\rangle \prec_{\Sigma_n} \langle H^M_{\kappa^+}, \in, U\rangle$ for all $\xi < \gamma$. Note that this already implies $(H^M_{\kappa^+}, \in, U)$ is a model of ZFC^-_n by Lemma 3.10.

We will characterise reflective (n,α)-baby measurable cardinals by the statement that the challenger does not have a winning strategy in $\mathsf{RG}^{\theta,n}_\alpha(\kappa)$ and faintly reflective (ω,n)-baby measurable cardinals by the statement that the challenger does not have a winning strategy in $\mathsf{faintRG}^{\theta,n}_\alpha(\kappa)$.

Suppose that $\omega \le \alpha \le \kappa$ is a regular cardinal. Let's argue that in the definition of the (reflective) (α,n)-baby measurable cardinals, we can strengthen the assumption that every subset of κ can be put into the required basic κ-model to show that, in fact, we can put every set into such a model.

Proposition 6.3. *If κ is (α,n)-baby measurable, then for every set A and all sufficiently large cardinals θ, there exists a basic weak κ-model $M \prec H_\theta$ with $A \in M$ for which there is a (good) M-ultrafilter U such that $\langle H^M_{\kappa^+}, \in, U\rangle \models \mathrm{ZFC}^-_n$. The same strengthening applies to faintly (ω,n)-baby measurable cardinals, reflective (α,n)-baby measurable cardinals, and faintly reflective (ω,n)-baby measurable cardinals.*

Proof. Suppose that κ is (α,n)-baby measurable and fix a set A. Suppose towards a contradiction that there is a cardinal θ with $A \in H_\theta$ for which there exists no basic weak κ-model as required containing A and elementary in H_θ. Choose any large enough H_ν (1) which sees that there is such a counterexample (θ, A) and (2) for which there is some basic κ-model $M \prec H_\nu$ satisfying the requirements of (α,n)-baby measurability for ν, namely that there is a (good) M-ultrafilter U such that $\langle H^M_{\kappa^+}, \in, U\rangle \models \mathrm{ZFC}^-_n$. By elementarity, M has some counterexample $(\bar\theta, \bar A)$. Let $\bar M = M \cap H_{\bar\theta}$. Then we have $\bar M \prec H_{\bar\theta}$, $\bar A \in \bar M$, and $H^M_{\kappa^+} = H^{\bar M}_{\kappa^+}$. Thus, U is a (good) $\bar M$-ultrafilter such that $\langle H^{\bar M}_{\kappa^+}, \in, U\rangle \models \mathrm{ZFC}^-_n$, but this contradicts that $(\bar\theta, \bar A)$ was a counterexample. The proof obviously generalizes to the other large cardinal notions. □

Indeed, the above proof shows that in the definition of the (reflective) (α,n)-baby measurable cardinals it suffices to assume that for sufficiently large θ, there is a single basic κ-model M satisfying the requirements.

Lemma 6.4. *The existence of a winning strategy for either player in the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$ is independent of $\theta > \kappa$. An analogous result holds for the games $\mathsf{faintG}^{\theta,n}_\omega(\kappa)$, $\mathsf{RG}^{\theta,n}_\alpha(\kappa)$, and $\mathsf{faintG}^{\theta,n}_\omega(\kappa)$.*

Proof. We will prove the result for the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$. The proof for the other games is nearly identical. Fix cardinals $\theta,\rho>\kappa$. Let's argue that if either player has a winning strategy in the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$, then they have a winning strategy in the game $\mathsf{G}^{\rho,n}_\alpha(\kappa)$.

Suppose that the challenger has a winning strategy σ in the game $\mathsf{G}^{\rho,n}_\alpha(\kappa)$. Let τ be the following strategy for the challenger in the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$. Let $M_0\prec H_\rho$ be the first move of the challenger in σ. The first move of the challenger in τ is going to be a basic κ-model $\bar M_0\prec H_\theta$ such that $P^{M_0}(\kappa)\subseteq \bar M_0$. At stage γ in the game, τ needs to respond to the moves $\{\langle N_\xi,\in,U_\xi\rangle\mid \xi<\gamma\}$ of the judge. Note that as long as we continue to choose $\bar M_\xi$ for $\xi<\gamma$ such that $P^{M_\xi}(\kappa)\subseteq\bar M_\xi$, where M_ξ is the move given by σ, then any move $\langle N_\xi,\in,U\rangle$ of the judge in the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$ will also be a valid move in the game $\mathsf{G}^{\rho,n}_\alpha(\kappa)$. So if M_γ is the response of σ to the moves $\{\langle N_\xi,\in,U_\xi\rangle\mid\xi<\gamma\}$ of the judge in the game $\mathsf{G}^{\rho,n}_\alpha(\kappa)$, then τ will tell the challenger to respond with $\bar M_\gamma\prec H_\theta$ such that $P^{M_\gamma}(\kappa)\subseteq\bar M_\gamma$. Suppose the judge can win a run of the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$. The run of the game gives models $\langle\bar M_\xi,\in,U_\xi\rangle$ for $\xi<\alpha$. Let $\bar M=\bigcup_{\xi<\kappa}\bar M_\xi$ and $U=\bigcup_{\xi<\kappa}U_\xi$. Since the judge wins, we have $\langle H^{\bar M}_{\kappa^+},\in,U\rangle\models \mathrm{ZFC}^-_n$. By construction of τ, the models $\bar M_\xi$ were chosen based on the moves M_ξ dictated by σ. But now it follows that the judge would win against the moves M_ξ by playing $\langle N_\xi,\in,U_\xi\rangle$ because $H^M_{\kappa^+}=H^{\bar M}_{\kappa^+}=\bigcup_{\xi<\kappa}N_\xi$.

Next, suppose that the judge has a winning strategy σ in the game $\mathsf{G}^{\rho,n}_\alpha(\kappa)$. Let τ be the following strategy for the judge in the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$. Let $M_0\prec H_\theta$ be the first move of the challenger. Let $\bar M_0\prec H_\rho$ be such that $P^{M_0}(\kappa)\subseteq\bar M_0$. Now let τ respond with the structure $\langle N_0,\in,U_0\rangle$ that is the response of σ to $\bar M_0$. Given a play $\{M_\xi\mid\xi<\gamma\}$ by the challenger, τ will tell the judge to respond with the response of σ to the play $\langle\bar M_\xi\mid\xi<\gamma\}$ such that $\bar M_\xi\prec H_\rho$ and $P^{M_\xi}(\kappa)\subseteq\bar M_\xi$. Now observe that if the challenger wins a play against τ with the moves $\{M_\xi\mid\xi<\alpha\}$, then the challenger would also win with the moves $\{\bar M_\xi\mid\xi<\alpha\}$ against σ because $H^M_{\kappa^+}=H^{\bar M}_{\kappa^+}=\bigcup_{\xi<\kappa}N_\xi$ for $M=\bigcup_{\xi<\alpha}M_\xi$ and $\bar M=\bigcup_{\xi<\alpha}\bar M_\xi$. □

Theorem 6.5. *A cardinal κ is (α,n)-baby measurable if and only if the challenger doesn't have a winning strategy in the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$ for some/all cardinals $\theta>\kappa$. A cardinal κ is reflective (α,n)-baby measurable if and only if the challenger doesn't have a winning strategy in the game $\mathsf{RG}^{\theta,n}_\alpha(\kappa)$ for some/all cardinals $\theta>\kappa$. An analogous result holds for the faint notions and games.*

Proof. Again, we will only prove the result about the (α,n)-baby measurable cardinals because the other proof is nearly identical.

Suppose that the challenger doesn't have a winning strategy in the game $\mathsf{G}^{\theta,n}_\alpha(\kappa)$ for some fixed regular cardinal $\theta>\kappa$. Fix $A\subseteq\kappa$. In particular, starting with a basic κ-model $M_0\prec H_\theta$ with $A\in M_0$ is not a winning strategy for the challenger, and so the judge wins some run of the game, where the challenger starts with such an $M_0\prec H_\theta$. Let M the union of the challenger's moves in this run of the game and U be the union of the ultrafilters played by the judge. Since the game was played for α-many steps, the model in each step was closed under $<\kappa$-sequences and α is a regular cardinal, the union model M is closed under $<\alpha$-sequences. Finally, since the judge wins, we have $\langle H^M_{\kappa^+},\in,U\rangle\models\mathrm{ZFC}^-_n$.

In the other direction, suppose that κ is (α,n)-baby measurable. Suppose towards a contradiction that the challenger has a winning strategy σ in the game $\mathsf{G}^{\kappa^+,n}_\alpha(\kappa)$. It is not hard to see that $\sigma\in H_\theta$ for $\theta=(2^\kappa)^+$. So fix some basic κ-model $M\prec H_\theta$ closed under $<\alpha$-sequences with $\sigma\in M$ for which there is an M-ultrafilter U such that $\langle H^M_{\kappa^+},\in,U\rangle\models\mathrm{ZFC}^-_n$. We will use M and U to play against σ and win, thereby showing that it couldn't have been a winning strategy. First suppose that α is uncountable.

Let M_0 be the first move of the challenger according to σ, and observe that, by elementarity, we have $M_0\in M$. Let $N_0\in M$ be a simple κ-model with $M_0\in N_0$ such that every Σ_n-statement that holds in $H^M_{\kappa^+}$ for an element of M_0 has a witness in N_0. Such a model exists since ZFC^-_n holds in $\langle H^M_{\kappa^+},\in,U\rangle$. Let the judge play $\langle N_0,\in,U\cap N_0\rangle$. Since $\langle N_0,\in,U\cap N_0\rangle\in M$, the response of the challenger according to σ must be in M as well. We continue letting the judge play $\langle N_\xi,\in,N_\xi\cap U\rangle\in M$ for each $\xi<\alpha$ in this fashion. Since M is closed under $<\alpha$-sequences, M will always have the sequence of the judge's moves at each step $\gamma<\alpha$ of the game. Thus, the judge can continue to play for α-many steps. Let $N=\bigcup_{\xi<\alpha}N_\xi$. We have

$$\langle N_\omega,\in,U\cap N_\omega\rangle\prec_{\Sigma_n}\cdots\prec_{\Sigma_n}\langle N_\lambda,\in,U\cap N_\lambda\rangle\prec_{\Sigma_n}\cdots\prec_{\Sigma_n}\langle N,\in,U\cap N\rangle\prec_{\Sigma_n}\langle H^M_{\kappa^+},\in,U\rangle$$

for limit ordinals λ. Thus, $\langle N, \in, U \cap N\rangle$ satisfies Σ_n-reflection, and therefore $\langle N, \in, U \cap N\rangle \models \mathrm{ZFC}^-_n$ by Lemma 3.10. As we already observed, N is precisely the H_{κ^+} of $M = \bigcup_{\xi<\alpha} M_\xi$, the union of the moves of the challenger. Thus, we have shown that the judge can win against σ, contradicting that σ was a winning strategy.

Finally, suppose that $\alpha = \omega$, which means that the model M doesn't have any closure. In this case, the judge responds to the moves of the challenger by choosing the models with the partial truth predicates as in the proof of Theorem 5.6 to ensure that the limit model satisfies ZFC^-_n. □

It follows from the proof that $\theta = (2^\kappa)^+$ suffices in the definition of (α, n)-baby measurable cardinals. Thus, κ is (α, n)-baby measurable if and only if every $A \in H_{(2^\kappa)^+}$ is an element of a $<\alpha$-closed imperfect weak κ-model $M \prec H_{(2^\kappa)^+}$ for which there is an M-ultrafilter such that $\langle H^M_{\kappa^+}, \in, U\rangle \models \mathrm{ZFC}^-_n$. An analogous result holds for reflective (α, n)-baby measurable cardinals and the faint notions.

Next, we show that the (reflective) (α, n)-baby measurable cardinals form a hierarchy.

Proposition 6.6. *Suppose that $\alpha < \beta \leq \kappa$ are uncountable regular cardinals. Then every (β, n)-baby measurable cardinal is a limit of cardinals $\nu > \alpha$ that are (α, n)-baby measurable. An analogous result holds for reflective (β, n)-baby measurable cardinals.*

Proof. As above, we will only prove the result about the (β, n)-baby measurable cardinals because the other proof is nearly identical.

Suppose that κ is (β, n)-baby measurable and fix a $<\beta$-closed basic weak κ-model M for which there is an M-ultrafilter U with $\langle H^M_{\kappa^+}, \in, U\rangle \models \mathrm{ZFC}^-_n$. Note that we do not use $M \prec H_\theta$ for some θ. Let N be the ultrapower of M by U. We will argue that κ is (α, n)-baby measurable in $V^N_{j(\kappa)}$. Otherwise in $V^N_{j(\kappa)}$, the challenger has a winning strategy in the game $\mathsf{G}^{\kappa^+,n}_\alpha(\kappa)$ by Lemma 6.4. We will use U to play against σ and argue that the resulting run of the game is in N. Using the same argument as in the proof of Theorem 6.5, it suffices to observe that M has all the required sequences by $<\beta$-closure. Finally, Lemma 6.4 shows that V_κ must be correct about a cardinal being (α, n)-game baby measurable. □

Theorem 6.7. *For $n \geq 1$, an n-baby measurable cardinal is a limit of cardinals ν that are $(<\nu, n)$-baby measurable. Moreover, an (κ, n)-baby measurable cardinal κ is a limit of n-baby measurable cardinals.*

Proof. The first part follow from the proof of Proposition 6.6.

For the second part, observe that κ is at n-least baby measurable. Now, fix a κ-model $M \prec H_{\kappa^+}$ for which there is an M-ultrafilter such that U such that $\langle M, \in, U\rangle \models \mathrm{ZFC}^-_n$, and let N be the ultrapower of M by U. We will argue that κ is n-baby measurable in N. Since n-baby measurability is verifiable in H_{κ^+}, M satisfies that κ is n-baby measurable by elementarity, and thus, so does N. □

Since being weakly n-baby measurable is a property of H_{κ^+}, it is easy to see that an (ω, n)-baby measurable cardinal is a weakly n-baby measurable limit of weakly n-baby measurable cardinals.

Theorem 6.8. *A weakly n-baby measurable is a limit of faintly (ω, n)-baby measurable cardinals.*

Proof. As usual, we suppose that σ is a winning strategy of the challenger in the ultrapower of M by U and use U to play against the strategy. We use the construction from the proof of Theorem 5.6, but modify the tree so that the even stages in the sequences are as above and the odd ones are moves of the challenger. □

Theorem 6.9. *For $n \geq 1$ and $\alpha < \kappa$, a $[n+1]$-baby measurable cardinal κ is a limit of reflective (α, n)-baby measurable cardinals.*

Proof. Fix $A \subseteq \kappa$ and choose a weak κ-model M, with $A \in M$, for which there is an M-ultrafilter U such that $\langle M, \in, U\rangle \models \mathrm{KP}_{n+1}$. By moving to $L[A, U]$ as constructed in M, we can assume without loss of generality that $\langle M, \in, U\rangle$ has a Δ_1-definable bijection $F : \mathrm{Ord}^M \to M$. Let N denote the, not necessarily well-founded, ultrapower of M by U. Fix $\alpha < \kappa$. We will argue that κ is reflective (α, n)-baby measurable in $V^N_{j(\kappa)}$. By Theorem 6.5, we only need to show that the challenger doesn't have a winning strategy in the game $\mathsf{RG}^{\kappa^+,n}_\kappa$. So suppose that $\sigma \in V^N_{j(\kappa)}$ is a winning strategy for the challenger in $\mathsf{RG}^{\kappa^+,n}_\kappa$. Note that via coding we can view σ as a map from $\mathcal{C}$ to $\mathcal{C}$, the collection of codes for elements of M, defined in the proof of Lemma 3.2.

As in the proof of Theorem 6.5, we will use U to play against σ. Since $\langle M, \in, U\rangle \models \mathrm{KP}_{n+1}$, by Lemma 3.9, every set $B \in M$ is contained in a κ-model $M_B \in M$ such that

$$(M_B, \in, U \cap M_B) \prec_{\Sigma_n} \langle M, \in, U\rangle$$

and $(M_B, \in, U \cap M_B) \models \mathrm{ZFC}^-_n$. Observe that, for a set B, the assertion that $\bar{M}$ is the F-least κ-model of this form is Σ_{n+1}.

Now using a Σ_{n+1}-recursion of length α in the structure $\langle M, \in, U\rangle$, we can construct a sequence of models $\{N_\xi \mid \xi < \alpha\} \in M$ as follows. Let M_0 be the first move of the challenger according to σ. Let N_0 be the F-least κ-model of the form M_{M_0}. Given that we have constructed the sequence $\{\langle N_\xi, \in, U \cap N_\xi\rangle \mid \xi < \beta\}$ for some $\beta < \alpha$ as the judge's moves, let M_β be the response of σ to this sequence for the challenger and let N_β be the F-least κ-model of the form M_{M_β}. To ask whether M_β is a response of the challenger to the sequence $\{\langle N_\xi, \in, U \cap N_\xi\rangle \mid \xi < \beta\}$, we just need to find a code in $\mathcal{C}$ as a subset of κ for the sequence, and then decode the subset of κ which is the response of σ. We use functions representing $\mathcal{C}$ and σ in the ultrapower together with U to determine membership in these sets. Clearly, the sequence $\{\langle N_\xi, \in, U \cap N_\xi\rangle \mid \xi < \alpha\} \in M \subseteq N$ is winning for the judge, which contradicts that σ was a winning strategy for the challenger in $V^N_{j(\kappa)}$. □

Next, we show that the hierarchy of reflective (α, n)-baby measurable cardinals for $n \geq 1$ sits on top of the (α, n)-baby measurable cardinals. Note that for $n = 0$, the (α, n)-baby measurable and the reflective (α, n)-baby measurable are the same large cardinal notion because Σ_0-elementarity is equivalent to being a submodel for transitive structures.

Theorem 6.10. *For $n \geq 1$, a faintly reflective (ω, n)-baby measurable cardinal is a limit of cardinals α that are (α, n)-baby measurable.*

Proof. Suppose that κ is (ω, n)-baby measurable. Fix a simple weak κ-model M for which there is an M-ultrafilter U such that for every $A \in M$ there is a κ-model $\bar{M} \in M$ such that $\langle \bar{M}, \in, U \cap \bar{M}\rangle \prec_{\Sigma_n} \langle M, \in, U\rangle$ and $\langle \bar{M}, \in, U \cap \bar{M}\rangle \models \mathrm{ZFC}^-_n$. Suppose that κ is not (κ, n)-baby measurable in the ultrapower N of M by U. This means that in N, the challenger has a winning strategy σ in the game $\mathsf{G}^{\kappa^+,n}_\kappa(\kappa)$. As in Lemma 3.2, the structure $\langle M, \in, U\rangle$ can check membership in σ using some function s representing it in the ultrapower and the checking procedure is Σ_1. Therefore the models $\langle \bar{M}, \in, U \cap \bar{M}\rangle$, where $s \in \bar{M}$, are going to be correct about membership in σ as well.

So fix some such $\langle \bar{M}, \in, \bar{M} \cap U\rangle$, which we will use to play against σ. The first move M_0 of the challenger must in $\bar{M}$ by Σ_1-elementarity. Let $\langle N_0, \in U \cap N_0\rangle$ be the response of the judge, where N_0 is a κ-model in $\bar{M}$ that has witnesses for all Σ_n-assertions true in $\langle \bar{M}, \in, U \cap \bar{M}\rangle$ with parameters in M_0 (this exists since $\langle \bar{M}, \in, U \cap \bar{M}\rangle \models \mathrm{ZFC}^-_n$). Since $\langle N_0, \in, U \cap N_0\rangle \in \bar{M}$, the next move of the challenger according to σ will also be in $\bar{M}$ by Σ_1-elementarity, and so we can choose $N_1 \in \bar{M}$ analogously. At limits λ, the sequence

$$\{\langle N_\xi, \in, U \cap N_\xi\rangle \mid \xi < \lambda\}$$

will be in $\bar{M}$ by closure, allowing us to analogously choose the next move N_λ. The κ-length sequence $\{\langle N_\xi, \in, U \cap N_\xi\rangle \mid \xi < \kappa\}$ may not be an element of $\bar{M}$, but since $\langle \bar{M}, \in, U \cap \bar{M}\rangle$ is an element of M, M sees the entire construction and therefore has the κ-length sequence, which is clearly winning for the judge. Thus, we have reached the desired contradiction showing that κ is (κ, n)-baby measurable in N. □

Finally, it is not difficult to see that for $n \geq 1$, a reflective (ω, n)-baby measurable cardinal is a limit of faintly reflective (ω, n)-baby measurable cardinals by constructing a tree of partial plays of the judge against the strategy and arguing that the tree has a branch.

7. From baby to locally measurable

The pattern of large cardinal notions around baby measurables is similar to the one around n-baby measurables, often with analogous and in some cases simpler proofs. Weakly baby measurable cardinals are above faintly baby measurable cardinals for a similar reason as in Proposition 5.6.

Proposition 7.1. *A weakly baby measurable cardinal κ is a limit of faintly baby measurable cardinals.*

Proof. Let M be a simple weak κ-model for which there is a good M-ultrafilter U such that $\langle M, \in, U\rangle \models \mathrm{ZFC}^-$. Let N be the ultrapower of M by U. Working in N, we argue that κ is faintly baby measurable. Fix a subset A of κ and consider the tree T whose elements on level n are sequences $\langle\langle M_0, \in, U_0\rangle, \ldots, \langle M_{n-1}, \in, U_{n-1}\rangle\rangle$ such that

- $A \in M_0$,
- $\langle M_i, \in, U_i\rangle \models \mathrm{ZFC}^-_i$ for all $i < n$,
- $\langle M_i, \in, U_i\rangle \prec_{\Sigma_i} \langle M_j, \in, U_j\rangle$ for all $i < j < n$.

The tree is ill-founded in V as witnessed by a chain of elementary substructures

$$\langle M_i, \in, U \cap M_i\rangle \prec_{\Sigma_i} \langle M, \in, U\rangle$$

built using Lemma 3.9. Thus, N has a branch through T as well. The union of models on the branch gives a structure $\langle \bar{M}, \in, W\rangle \models \mathrm{ZFC}^-$ such that $\bar{M}$ is a weak κ-model and W is a $\bar{M}$-ultrafilter. Thus, κ is faintly baby measurable in N, and hence κ is a limit of these cardinals by elementarity. □

Above these notions, it is easy to see that an ω_1-baby measurable cardinal κ is a limit of weakly baby measurable cardinals. The interval between weakly baby measurable and baby measurable cardinals can be studied by replacing ZFC^-_n with ZFC^- and Σ_n-elementarity with full elementarity in the definitions of the games $\mathsf{G}^{\theta,n}_\alpha(\kappa)$ and $\mathsf{RG}^{\theta,n}_\alpha(\kappa)$ above, resulting in games $\mathsf{G}^{\theta}_\alpha(\kappa)$ and $\mathsf{RG}^{\theta}_\alpha(\kappa)$ respectively. Just like in Lemma 6.4, one can now show that the existence of a winning strategy in these games for either player is independent of $\theta \geq \kappa^+$. As in Theorem 6.5, one can then show that a cardinal κ is α-baby measurable if and only if the challenger does not have a winning strategy in $\mathsf{G}^{\theta}_\alpha(\kappa)$, and the analogous characterisation holds for reflective α-baby measurable cardinals and $\mathsf{RG}^{\theta}_\alpha(\kappa)$. The proofs are virtually the same except that in the construction of a winning run for the judge, in step $\lambda + n$ for limits λ one adds witnesses for Σ_n-truths only. The α-baby measurables and the reflective α-baby measurables form strict hierarchies as in Proposition 6.6. Moreover, the former hierarchy sits strictly below the latter, since a faintly reflective ω-baby measurable cardinal is a limit of cardinals α that are α-baby measurable as in Proposition 6.10. At the top of these hierarchies, a baby measurable cardinal is a limit of cardinals ν that are $<\nu$-baby measurable as in Proposition 6.7, similarly to how a strongly Ramsey cardinal κ is a limit of cardinals ν that are $<\nu$-Ramsey, and a κ-baby measurable cardinal κ is a limit of baby measurable cardinals, similarly to how a κ-Ramsey cardinal κ is a limit of strongly Ramsey cardinals.

Locally measurable cardinals (see Definition 2.18) are above all these notions.

Proposition 7.2. *A locally measurable cardinal κ is baby measurable and a limit of cardinals ν that are reflective ν-baby measurable.*

Proof. Let M be a weak κ-model which thinks it has a normal ultrafilter U on κ. Let N be the ultrapower of M by U. If N had a strategy σ for the challenger in the game $\mathsf{RG}^{\kappa^+}_\kappa$, then M would see the strategy via the function representing it in the ultrapower and would be able to use U to play against it.

It remains to show that κ is baby measurable. Fix $A \subseteq \kappa$ and a weak κ-model M, with $A \in M$, having what it thinks is a normal ultrafilter U on κ. Clearly, M can build what it thinks is a κ-model $\bar{M}$ such that $\langle \bar{M}, \in, U \cap \bar{M}\rangle \prec \langle H^M_{\kappa^+}, \in, U\rangle$, with $A \in \bar{M}$, because $H^M_{\kappa^+}$ is a set in M. But then $\bar{M}$ is actually a κ-model. □

8. Indestructibility

In this section, we provide a few basic indestructibility results for faintly baby measurable, weakly baby measurable, and baby measurable cardinals. More specifically, we show that these large cardinals κ are indestructible by small forcing and can be made indestructible by the forcing $\mathrm{Add}(\kappa,1)$ adding a Cohen subset to κ.

The indestructibility arguments will use properties of class forcing over models of second-order set theory. A model of second-order set theory is a triple $\mathbf{M}=\langle M,\in,\mathcal{C}\rangle$, where M consists of the sets of the model and $\mathcal{C}$ consists of the classes. The second-order theory GBc^- consists of the axioms ZFC^- for sets, the extensionality axiom for classes, the class replacement axiom asserting that every class function restricted to a set is a set, and the first-order comprehension scheme asserting that every first-order formula defines a class. The theory GBC^- is the theory GBc^- together with the assertion that there is a class well-order of sets of order-type Ord. Given a class partial order $\mathbb{P}\in\mathcal{C}$, we say that a filter $G\subseteq\mathbb{P}$ is $\mathbf{M}$-*generic* for $\mathbb{P}$ if it meets every dense subclass of $\mathbb{P}$ from $\mathcal{C}$. Given an $\mathbf{M}$-generic filter G, the forcing extension of $\mathbf{M}$ by G is the structure $\langle M[G],\in,\mathcal{C}[G]\rangle$, where $M[G]$ consists of the interpretation of $\mathbb{P}$-names by G and $\mathcal{C}[G]$ consists of the interpretation of class $\mathbb{P}$-names by G, where a *class* $\mathbb{P}$-*name* is any class whose elements are pairs of the form $\langle\sigma,p\rangle$ with $p\in\mathbb{P}$ and σ a $\mathbb{P}$-name. If M is ill-founded, we can still form the generic extension by taking instead of interpretations of names, the structure whose elements are equivalence classes of (class) $\mathbb{P}$-names moded out by the filter G. Although class forcing may not always preserve replacement to the forcing extension, pretame partial orders preserve the theories GBc^- and GBC^-. See [HKS18] for the definition and results on pretameness, and see [AGng] for details on models of second-order set theory and class forcing.

Proposition 8.1. *Faintly baby measurable, weakly baby measurable, and baby measurable cardinals are indestructible by small forcing.*

Proof. Suppose that κ is weakly baby measurable, $\mathbb{P}\in V_\kappa$ is a forcing notion, and $g\subseteq\mathbb{P}$ is V-generic. Fix $A\subseteq\kappa$ in $V[g]$ and let $\dot A$ be a nice $\mathbb{P}$-name such that $\dot A_g=A$. Since $\mathbb{P}\in V_\kappa$, $\dot A$ is a subset of V_κ as well and hence we can put it into a simple weak κ-model M for which there is a good M-ultrafilter U such that $\langle M,\in,U\rangle\models\mathrm{ZFC}^-$. Let $\mathcal{C}$ be the classes of M generated by U, so that the second-order structure $\langle M,\in,\mathcal{C}\rangle\models\mathrm{GBc}^-$. Since $\mathbb{P}$ is a set forcing in M, it is, in particular, trivially pretame. Since pretame forcing preserves GBc^-, we have that the second-order structure $\langle M[g],\in,\mathcal{C}[g]\rangle\models\mathrm{GBc}^-$ as well. Let W the $M[g]$-ultrafilter generated by U in $M[g]$. Clearly $W\in\mathcal{C}[g]$, and so it follows that $\langle M[g],\in,W\rangle\models\mathrm{ZFC}^-$. Also, clearly W is good because we can lift the ultrapower embedding $j:M\to N$ to $j:M[g]\to N[g]$ and the $M[g]$-ultrafilter generated from the lift is precisely W. Since $A\in M[g]$, the structure $\langle M[g],\in,W\rangle$ witnesses weak baby measurability for A.

The case of faintly baby measurable cardinals is even easier because it suffices to note that W is definable in $\langle M[g],\in,\mathcal{C}[g]\rangle$.

For the case of baby measurable cardinals it suffices to observe that a forcing extension $M[g]$ of a κ-model M by $\mathbb{P}$ is again a κ-model in $V[g]$. □

Theorem 8.2. *Faintly baby measurable cardinals, weakly baby measurable cardinals, and baby measurable cardinals κ can be made indestructible by the forcing* $\mathrm{Add}(\kappa,1)$.

Proof. Suppose that κ is weakly baby measurable. Let $\mathbb{P}_\kappa$ be the κ-length Easton support iteration forcing with $\mathrm{Add}(\alpha,1)$ at regular cardinal stages, and let $G*g$ be V-generic for $\mathbb{P}_\kappa*\mathrm{Add}(\kappa,1)$. With some slight renamings, we can assume that the poset $\mathbb{P}_\kappa*\mathrm{Add}(\kappa,1)$ is a subset of V_κ. Every subset $A\subseteq\kappa$ in $V[G*g]$ has a nice $\mathbb{P}_\kappa*\mathrm{Add}(\kappa,1)$-name $\dot A$ in V, which can therefore be put into a weak κ-model M for which there is an M-ultrafilter U such that $\mathbf{M}:=\langle M,\in,U\rangle\models\mathrm{ZFC}^-$. By moving to $L[\dot A,U]$ of M, we can assume without loss of generality that M has a definable class bijection $F:\mathrm{Ord}^M\to M$. Let $\mathcal{C}$ be the classes of M generated by U. Since the forcing $\mathbb{P}_\kappa*\mathrm{Add}(\kappa,1)$ is a set and hence pretame, we have that the second-order structure

$$\langle M[G*g],\in,\mathcal{C}[G*g]\rangle\models\mathrm{GBC}^-.$$

Let Ult be the (not collapsed) ultrapower of M by U that is definable in $\langle M,\in,U\rangle$ and let $\Psi:M\to\mathrm{Ult}$ be the ultrapower map, which is also definable there. Note that $\langle M,\in,U\rangle$ can pick out a unique element of each equivalence class using the global well-order function F. The model $\mathbf{M}[G*g] = \langle M[G*g], \in, \mathcal{C}[G*g]\rangle$ has the classes M and Ult. Using $G*g$, we can define the model $\text{Ult}[G*g]$ inside $\langle M[G*g], \in, \mathcal{C}[G*g]\rangle$. We can think of elements of $\text{Ult}[G*g]$ as equivalence classes $[\tau]$ of $\mathbb{P}_\kappa * \text{Add}(\kappa,1)$-names from Ult, where we have that τ is equivalent to σ whenever there is $p \in G*g$ such that $p \Vdash \tau = \sigma$. The entire construction is definable in $\mathbf{M}[G][g]$. We will lift the ultrapower embedding Ψ of M by U to an ultrapower embedding of $M[G*g]$ inside the structure $\langle M[G*g], \in, \mathcal{C}[G*g]\rangle$.

First, we lift Ψ to $M[G]$. Using the standard lifting criterion for lifting elementary embeddings to a forcing extension, we need to build an Ult-generic filter H for the poset $\Psi(\mathbb{P}_\kappa)$ with $\Psi[G] \subseteq H$. The poset $\Psi(\mathbb{P}_\kappa)$ factors as $\mathbb{P}_\kappa * \text{Add}(\kappa,1) * \mathbb{P}_{\text{tail}}$ (we will associate the initial segment of the ultrapower that is isomorphic to M with M itself to simplify notation). We use $G*g$ for the initial segment of the forcing, thereby trivially satisfying the requirement that $\Psi[G] = G$ will be contained in the filter we end up building. So it remains to find an $\text{Ult}[G*g]$-generic filter G_{tail} for the tail forcing $\mathbb{P}_{\text{tail}}$, which is ${<}\kappa^+$-closed there. The model $\langle M[G*g], \in, \mathcal{C}[G*g]\rangle$ has a class bijection $F' : \text{Ord}^M \to M[G*g]$ constructed from F. Using F', in $\langle M[G*g], \in, \mathcal{C}[G*g]\rangle$, we can enumerate all the dense subsets of $\Psi(\mathbb{P}_\kappa)$ in $\text{Ult}[G*g]$ in a class sequence of length Ord^M. The length of every initial segment of this enumeration is some ordinal $\alpha \in M$. Note next that every sequence of elements of Ult of order type α must be an element of Ult because it is an ultrapower. The closure also transfers to the pairs $M[G]$ and $\text{Ult}[G]$, as well as $M[G*g]$ and $\text{Ult}[G*g]$ (for details of closure arguments, see [GJ22, Section 3]). Thus, as we diagonalize against the sequence of dense sets inside $\langle M[G*g], \in, \mathcal{C}[G*g]\rangle$, every initial segment of our choices of elements from the dense sets is going to be a sequence in $\text{Ult}[G*g]$ and therefore since $\Psi(\mathbb{P}_\kappa)$ is ${<}\kappa^+$-closed in $\text{Ult}[G*g]$, we can find an element below the diagonalization sequence. Thus, we can continue using replacement in our structure to define the generic filter. Note that the generic filter G_{tail} we construct in this manner is a class in $\mathbf{M}[G*g]$.

Once we have the generic filter G_{tail}, we can repeat the process to define an $\text{Ult}[G*g][G_{\text{tail}}]$-generic filter for the image $\Psi(\text{Add}(\kappa,1)) = \text{Add}(\Psi(\kappa),1)$ of $\text{Add}(\kappa,1)$ under the lifted ultrapower map Ψ. Thus, we can lift the ultrapower embedding Ψ to $M[G*g]$ and use the ultrapower map Ψ to define the $M[G*g]$-ultrafilter W extending U. Since W was defined inside the structure $\langle M[G*g], \in, \mathcal{C}[G*g]\rangle$, we have that $\langle M[G*g], \in, W\rangle \models \text{ZFC}^-$. The lifted ultrapower map verifies that W is good.

The case of faintly baby measurable cardinals proceeds identically to above because we never used the well-foundedness of the ultrapower in our construction.

For the case of baby measurable cardinals it suffices to observe that a forcing extension $M[G][g]$ of a κ-model M by $\mathbb{P}_\kappa * \text{Add}(\kappa,1)$ is again a κ-model in $V[G][g]$ (for details, see [GJ22, Section 3]). □

Corollary 8.3. *A faintly baby measurable, weakly baby measurable or baby measurable cardinal κ can be made indestructible by the forcing* $\text{Add}(\kappa,\theta)$ *for any cardinal θ.*

Proof. Suppose that $G \subseteq \text{Add}(\kappa,\theta)$ is V-generic. Observe that any $A \subseteq \kappa \in V[G]$ has an $\text{Add}(\kappa,\theta)$-name $\dot{A}$ in V that uses at most κ-many coordinates in θ. Thus, we can view $\dot{A}$ as an $\text{Add}(\kappa,1)$-name. Hence, it suffices to show that these cardinals can be made indestructible by $\text{Add}(\kappa,1)$. □

It follows that we can make the GCH fail at these cardinals. Also, a simple version of the above lifting argument can be used to show that the GCH forcing that adds Cohen subsets at all successor cardinals preserves these cardinals.

We believe that the indestructibility results should work for the level by level n-versions of the baby measurable cardinals as well (at least for some reasonably large n). This would rely on set forcing preserving an appropriate version of the class theory GBC^-_n extending ZFC^-_n (or KP^-_n) and checking that the complexity of the constructions (ultrapower and generic filters) never beyond Σ_n.

9. A detour into second-order set theory

Models of ZFC^- with a largest cardinal κ that is inaccessible are bi-interpretable with models of the second-order set theory Kelley-Morse strengthened by a choice principle for classes. Let ZFC^-_I be the theory asserting that ZFC^- holds, that there is a largest cardinal

κ, and that κ is inaccessible ($P(\alpha)$ exists and $2^\alpha < \kappa$ for all $\alpha < \kappa$). The assumption that κ is inaccessible implies, in particular, that V_κ exists, and that $V_\kappa \models \text{ZFC}$. Recall that Kelley-Morse (KM) is a second-order set theory whose axioms consist of GBC together with the full comprehension scheme asserting for every second-order formula that it defines a class. We can further strengthen KM by adding various very useful choice principles for classes. Let the *choice scheme* (CC) be the scheme which asserts, for every second-order formula $\varphi(x, X, A)$, that if for every set x, there is a class X witnessing φ, then there is a single class Y collecting witnesses for every set x on its slices $Y_x = \{y \mid \langle x, y\rangle \in Y\}$.

Marek showed that the theory Kelley-Morse together with the choice scheme (KM + CC) is bi-interpretable with ZFC^-_I [Mar73, Section 2]. Given a model $\mathbf{V} = \langle V, \in, \mathcal{C}\rangle \models \text{KM} + \text{CC}$, we obtain the corresponding model $M_{\mathbf{V}} \models \text{ZFC}^-_I$ by taking all the well-founded extensional relations in $\mathcal{C}$, modulo isomorphism, with the natural membership relation that results from viewing these relations as transitive sets. We get that $V_\kappa^{M_{\mathbf{V}}} \cong V$ and $P(V_\kappa)^{M_{\mathbf{V}}} \cong \mathcal{C}$. In the other direction, given any model $M \models \text{ZFC}^-_I$, we obtain the corresponding model

$$\mathbf{V} = \langle V_\kappa^M, \in, P(V_\kappa)^M\rangle \models \text{KM} + \text{CC}.$$

Moreover, what gives us bi-interpretability is that the ZFC^-_I-model $M_{\mathbf{V}}$ corresponding to $\mathbf{V}$ is precisely M.

Let ZFC^-_U be the theory in the language with an additional unary predicate U consisting of ZFC^-_I in the extended language together with the assertion that U is a uniform normal ultrafilter on the largest cardinal κ.

On the second-order side, let KM_U be the theory in the language of second-order set theory with an additional unary predicate U on classes consisting of KM in the extended language together with the assertion that U is a uniform normal ultrafilter on Ord. Let

$$\mathbf{V} = \langle V, \in, \mathcal{C}, U\rangle \models \text{KM}_U.$$

Consider the ultrapower structure $\langle \text{Ult}, \text{E}\rangle$ consisting of the equivalence classes of class functions $F : \text{Ord} \to V$ from $\mathcal{C}$ modulo U. It is not difficult to see that Łoś' Theorem holds for the structures $\langle V, \in\rangle$ and $\langle \text{Ult}, \text{E}\rangle$: $\langle \text{Ult}, \text{E}\rangle \models \varphi([F])$ if and only if $\{\alpha \mid \langle V, \in\rangle \models \varphi(F(\alpha)\} \in U$. It follows that the universe V is isomorphic to a rank-initial segment Ult_κ of Ult consisting of equivalence classes of constant functions $C_a : \text{Ord} \to V$ such that $C_a(\alpha) = a$ for all α, and that $\text{Ult}_{\kappa+1}$ consists precisely of the classes $\mathcal{C}$ via this isomorphism. Since by Łoś' Theorem, $\langle \text{Ult}, \text{E}\rangle \models \text{ZFC}$, it has a well-ordering of $\text{Ult}_{\kappa+1}$. Now, essentially by the argument of Lemma 3.2, we can conclude that $\mathbf{V}$ has a definable well-ordering of its classes. The existence of a definable well-ordering of classes is much stronger than the choice scheme, which clearly follows from it.

Similar arguments as above show that KM_U and ZFC^-_U are bi-interpretable. Therefore we should view ZFC^-_U as essentially being a strong second-order set theory asserting that Ord is measurable.

From the bi-interpretability of these theories, we also know the first-order consequences in models $\mathbf{V} = \langle V, \in, \mathcal{C}, U\rangle$ of KM_U. Using the proof of Theorem 5.3 and Σ_n class reflection mentioned below, it follows that there is a proper class of cardinals κ that are κ-Ramsey. Our results further shed light on the global structure of $\mathbf{V}$. We say that Ord is *ineffably Ramsey* if every class function $F : [\text{Ord}]^{<\omega} \to 2$ has a stationary homogeneous class. By Σ_n *class reflection*, we mean the assertions for each $n \in \omega$ that for every class $A \in \mathcal{C}$, there is a collection $\bar{\mathcal{C}} \subseteq \mathcal{C}$ coded by a single class such that $A \in \bar{\mathcal{C}}$ and $\langle V, \in, \bar{\mathcal{C}}, U \cap \bar{\mathcal{C}}\rangle \prec_{\Sigma_n} \langle V, \in, \mathcal{C}, U\rangle$. Note that the results cited in the following claims can be applied in this situation since the proofs of Lemma 2.6 and Lemmas 3.2, 3.4 and 3.9 only use internal properties of the structures $\langle M, \in, U\rangle$ and therefore, these results hold for all models $\langle M, \in, U\rangle$ of the respective theory such that U is an M-ultrafilter from the viewpoint of $\langle M, \in, U\rangle$.

- $\mathbf{V}$ has a definable global well-order of $\mathcal{C}$ by Lemma 3.2.
- $\mathbf{V}$ has a truth predicate for $\langle V, \in, \mathcal{C}\rangle$ by Lemma 3.4.
- Ord is ineffably Ramsey in $\mathbf{V}$. (This can be seen by carrying out the argument from [Git11, Lemma 3.6] internally to $\mathbf{V}$.)
- $\mathbf{V}$ satisfies Σ_n class reflection for each $n \in \omega$ by Lemma 3.9.

We next provide a version of Bovykin's and McKenzie's Theorem 1.1 that uses our variants of their n-baby measurable cardinals and the above results about them.

Theorem 9.1. *The following theories are equiconsistent:*

1. ZFC^-_U.
2. ZFC *together with the scheme consisting of the assertions for all* $n \in \omega$*:*
   "There exists an n*-baby measurable cardinal* κ *with* $V_\kappa \prec_{\Sigma_n} V$*."*
3. ZFC *together with the scheme consisting of the assertions for all* $n \in \omega$*:*
   "There exists an n*-baby measurable cardinal."*

Moreover, (2) captures precisely the theory of models M *of (1) restricted to* V_κ^M*, where* κ *is the largest cardinal of* M*.*

Proof. To show that the consistency of (1) implies that of (2), suppose that $\langle M, \in, U\rangle \models \mathrm{ZFC}^-_U$ with a largest cardinal κ (although we use $\in$ for the membership relation here, we don't assume that it is the actual membership relation). We show that (2) holds in V_κ^M. Fix $n \in \omega$. The proof of Lemma 3.9 shows that every $A \subseteq \kappa$ in M is an element of $\bar{M}$, which M thinks is a κ-model, such that $\langle \bar{M}, \in \bar{M} \cap U\rangle \prec_{\Sigma_n} \langle M, \in, U\rangle$. Thus, the ultrapower N of M by U satisfies that κ is n-baby measurable. Let's argue that $V_\kappa \prec V_{j(\kappa)}$ in N, where j is the ultrapower map. Note that M and N have the same natural numbers (possibly nonstandard), so they agree about formulas. Given a (possibly nonstandard) formula $\varphi(x)$, we also have that M and N agree on whether $V_\kappa \models \varphi(a)$. Moreover, if $M \models$ "$V_\kappa \models \varphi(a)$", then $N \models$ "$V_{j(\kappa)} \models \varphi(j(a))$" by elementarity. It follows, by Łoś' Theorem, that there is some $\alpha < \kappa$ such that M satisfies that $V_\alpha \prec V_\kappa$ and α is n-baby measurable. Thus, in particular, we actually have that $V_\alpha \prec_{\Sigma_n} V_\kappa$ because M will be correct about satisfaction for standard formulas. Also, V_κ clearly agrees with M that α is n-baby measurable.

(2) clearly implies (3).

To show that the consistency of (3) implies that of (1), suppose that N is a model (3). Suppose towards a contradiction that there is no model of (1), meaning that for some $n < \omega$, the fragment of ZFC^-_U mentioning only instances of collection and separation for formulas of complexity at most Σ_n is inconsistent. Fix such an $n < \omega$. The model N has a model $\langle M, \in, U\rangle$ satisfying what it thinks is all instances of collection and separation for formulas of complexity at most Σ_n, and it must be correct about this for standard formulas. Thus, we have reached a contradiction by producing such a model.

It remains to show the "moreover" part. So suppose that N is a model of (3). We need to argue that the theory ZFC^-_U together with the assertions that $V_\kappa \models \varphi$ for every $\varphi \in \mathrm{Th}(N)$ is consistent. If this were not the case, then there would be some finite fragment T of ZFC^-_U and some φ such that $N \models \varphi$, but there is no model $\langle M, \in, U\rangle \models T$ such that $V_\kappa^M \models \varphi$. Choose $n < \omega$ bounding the complexity of φ and all assertions in T. Let κ be an n-baby measurable cardinal in N such that $V_\kappa^N \prec_{\Sigma_n} N$. Then N has a model $\langle M, \in, U\rangle$ witnessing that κ is n-baby measurable. It follows that $\langle M, \in, U\rangle \models T$ and also $V_\kappa^M = V_\kappa^N \models \varphi$ by Σ_n-elementarity, contradicting our assumption that this theory is inconsistent. □

The above argument works as well for measurable cardinals. Let ZFC_n denote ZFC with the replacement (equivalently, collection) and separation schemes restricted to Σ_n-formulas. We call a cardinal κ *n-junior measurable* if every $A \subseteq \kappa$ is an element of a κ-model M of ZFC_n with a normal ultrafilter $U \in M$.

Remark 9.2. *Let* T *denote* ZFC *together with the existence of a measurable cardinal. Let* S *denote* ZFC *with the scheme consisting of the following sentences for all* $n \in \omega$*:*

"There exists an n*-junior measurable cardinal* κ *with* $V_\kappa \prec_n V$*".*

As in Theorem 9.1, S *captures precisely the consequences in* V_κ *of measurable cardinals* κ*. Since the argument only uses* Π^1_2*-indescribability, a similar claim holds for all other large cardinal notions that imply* Π^1_2*-indescribability, for instance for the notions of strong and supercompact cardinals.*

In particular, there is a hierarchy of natural large cardinal notions, the n-junior measurable cardinals, that reaches up all the way to measurable cardinals. This answers a question of Daniel Isaacson asked at the first-listed author's talk in the Oxford Set Theory Seminar in May 2020.

10. Outlook

We provided a fine analysis of large cardinal notions in the interval between Ramsey and measurable cardinals defined by expanding the amount of collection and separation available in the relevant models. The diagram in Figure 10 below provides an overview of relationships between the large cardinal notions. The patterns around α-Ramsey, (α, n)-baby measurable and α-baby measurable cardinals are enclosed by solid boxes. Repeating steps in a hierarchy that depends on $\omega \leq \alpha < \kappa$ or $1 \leq n < \omega$ are enclosed by dashed boxes. For example, the large dashed box encloses a pattern that repeats for each $1 \leq n < \omega$. An $[n+1]$-baby measurable cardinal is a limit of reflective (α, n)-baby measurable cardinals by Theorem 6.9 and such cardinals are $[n]$-baby measurable. Note that the notions in this box collapse for $n = 0$ by Section 5. The range of α for α-Ramsey cardinals is meant to be $\omega_1 \leq \alpha < \kappa$.

The differences between the properties of KP_n and ZFC^-_n studied in Section 3 entail that closure properties are not relevant for the large cardinal notions defined via KP_n by Lemma 5.1, while they induce to a strict hierarchy for the ones defined via ZFC^-_n by Theorem 5.7 that is studied via the games $\mathsf{G}^{\theta,n}_\alpha(\kappa)$ in Section 5.

We expect a similar pattern as the one around κ-baby measurable cardinals to recur at reflective κ-baby measurable cardinals. It is natural to ask whether reflective (α, n)-baby measurable cardinals are precisely the cardinals κ such that the challenger does not have a winning strategy for the game of length $\kappa \cdot \alpha$. Some issues are left open for the strong variant of the game $\mathsf{RG}^{\theta,n}_\alpha(\kappa)$ in Definition 6.2 where we ask that $M_\gamma \in N_\gamma$ for all $\gamma < \alpha$. We have a similar characterisation as in Theorem 6.5 for winning strategies for the challenger by modifying the reflection in Definition 4.3 to all sets $B \in M$ instead of just subsets of κ, but it is open whether the existence of a winning strategy for the challenger in this game is independent of θ as in Lemma 6.4.

Theorem 9.1 provides a bridge between large cardinals in set theory and class theory. In particular, the theory S in (3) is interpretable in $T := \mathrm{ZFC}^-_U$ and T is conservative over S. Since in familiar examples of conservative extensions such as Gödel-Bernays class theory GBC and ZFC, any model of ZFC can be extending to one of GBC by adding a second-order part, we ask if the same holds here.

Problem 10.1. Is every model of S the restriction to V_κ of a model of T, where κ is the largest cardinal? In other words, is the function from $\mathrm{Mod}(T)$ to $\mathrm{Mod}(S)$ induced by the interpretation surjective?

The results in Section 9 show how to approximate some large cardinal notions from below. For instance, we studied the precise consequences in V^M_κ for a measurable cardinal κ in a model M of ZFC^- in Theorem 9.1. We further ask whether n-baby measurable cardinals can be replaced by (κ, n)-baby measurable cardinals in this theorem. Regarding a finer version of the theorem, let $\mathrm{ZFC}^-_{U,n}$ denote the variant of ZFC^-_U where the collection and separation schemes are restricted to Σ_n-formulas for some $n \geq 1$ and κ denotes the largest cardinal. Are the consequences $\mathrm{ZFC}_{U,n}$ in V_κ axiomatizable by large cardinal properties?

The same problem is of interest for smaller large cardinals.

Problem 10.2. Is the theory of models of the form V^M_κ axiomatizable by large cardinal properties, where M is a model of ZFC and κ is a weakly compact cardinal in M?

The approximation of measurable cardinals from below in Remark 9.2 suggests to ask whether there is a connection with Bagaria's characterization of the existence of measurable cardinals [Bag23, Section 5.2].

The following variant of the above notions for countable models may be connected with properties of sets of reals and the determinacy of infinite games.

Problem 10.3. Consider the statement that every real is contained in a countable transitive model M of ZFC^- with the largest cardinal κ and an M-ultrafilter U on $\mathcal{P}(\kappa)^M$ such that $(M, \in, U)$ is an ω_1-iterable model of KP_n, where $n \geq 1$. Is this equivalent to the determinacy of a natural class of projective sets?

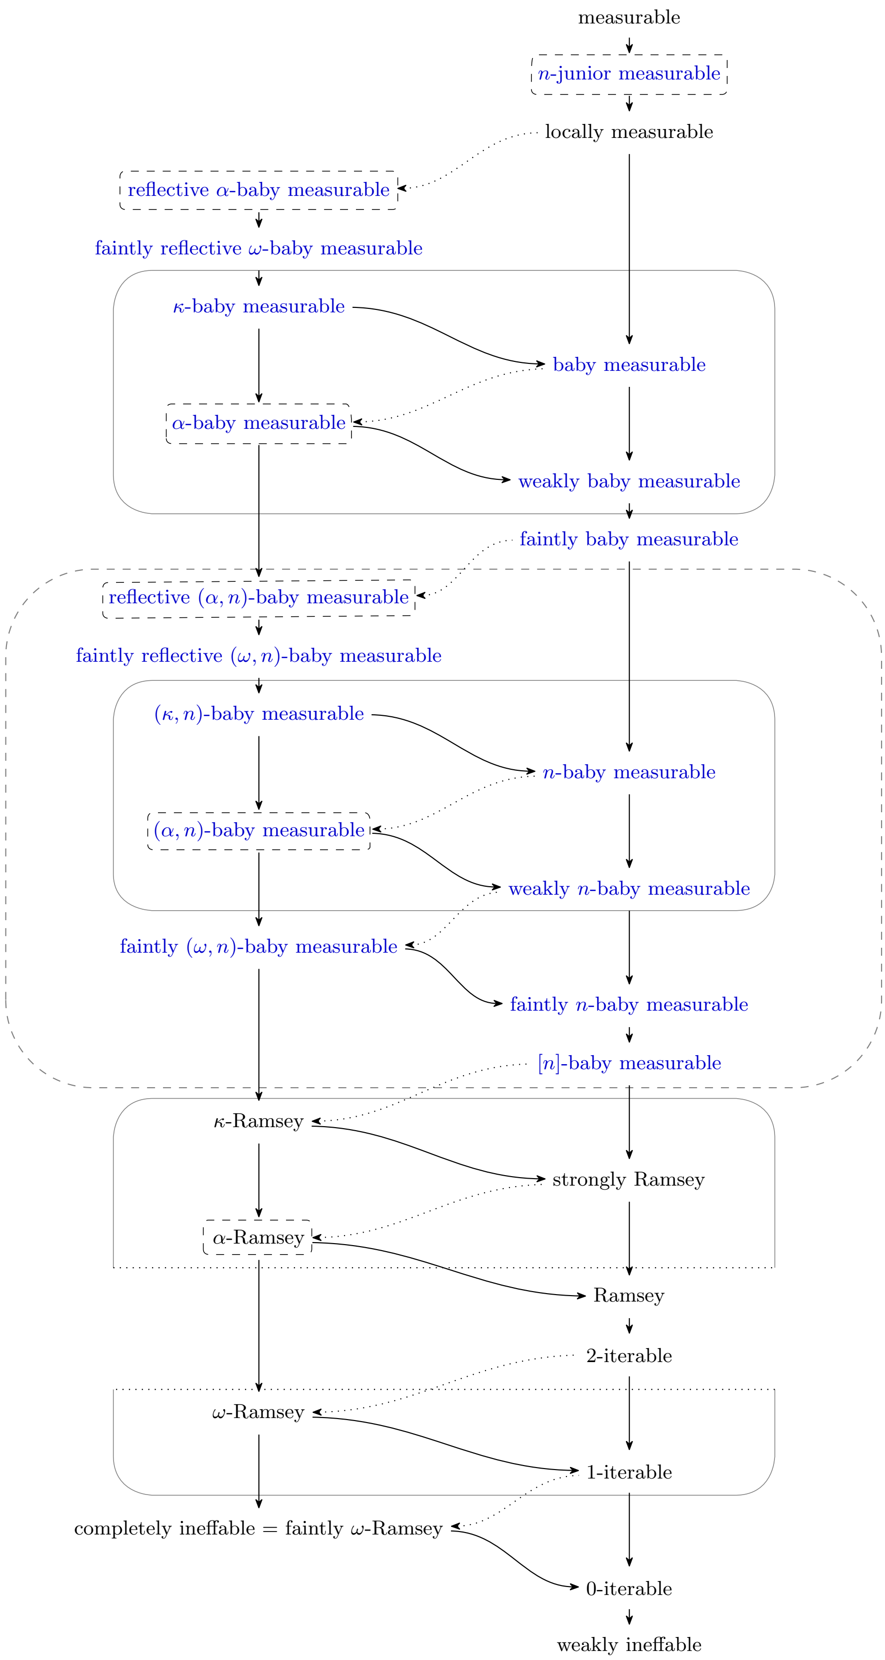

FIGURE 2. Implications between large cardinal notions. Solid arrows denote direct implications, dotted arrows implications in consistency strength.

References

[AGng] Carolin Antos and Victoria Gitman. Modern class forcing. In *Research Trends in Contemporary Logic*. College Publications, Forthcoming.

[Bag23] Joan Bagaria. Large cardinals as principles of structural reflection. *Bull. Symb. Log.*, 29(1):19–70, 2023.

[BM] Andrey Bovykin and Zachiri McKenzie. Ramsey-like cardinals that characterize the exact consistency strength of NFUM. *Manuscript*.

[Gai74] Haim Gaifman. Elementary embeddings of models of set-theory and certain subtheories. In *Axiomatic set theory (Proc. Sympos. Pure Math., Vol. XIII, Part II, Univ. California, Los Angeles, Calif., 1967)*, pages 33–101. Amer. Math. Soc., Providence R.I., 1974.

[GHJ16] Victoria Gitman, Joel David Hamkins, and Thomas A. Johnstone. What is the theory ZFC without power set? *Math. Log. Q.*, 62(4-5):391–406, 2016.

[Git11] Victoria Gitman. Ramsey-like cardinals. *J. Symb. Log.*, 76(2):519–540, 2011.

[GJ22] Victoria Gitman and Thomas A. Johnstone. Indestructibility properties of Ramsey and Ramsey-like cardinals. *Ann. Pure Appl. Logic*, 173(6):Paper No. 103106, 30, 2022.

[GW11] Victoria Gitman and Philip D. Welch. Ramsey-like cardinals II. *J. Symb. Log.*, 76(2):541–560, 2011.

[HKS18] Peter Holy, Regula Krapf, and Philipp Schlicht. Characterizations of pretameness and the Ord-cc. *Ann. Pure Appl. Log.*, 169(8):775–802, 2018.

[HL21] Peter Holy and Philipp Lücke. Small models, large cardinals, and induced ideals. *Ann. Pure Appl. Log.*, 172(2):102889, 2021.

[HS18] Peter Holy and Philipp Schlicht. A hierarchy of Ramsey-like cardinals. *Fund. Math.*, 242(1):49–74, 2018.

[Kra82] Evangelos Kranakis. Reflection and partition properties of admissible ordinals. *Ann. Math. Log.*, 22(3):213–242, 1982.

[Kun70] Kenneth Kunen. Some applications of iterated ultrapowers in set theory. *Ann. Math. Log.*, 1:179–227, 1970.

[Mar73] Wiktor Marek. On the metamathematics of impredicative set theory. *Diss. Math.*, 98:1–41, 1973.

[Mit79] William Mitchell. Ramsey cardinals and constructibility. *J. Symb. Log.*, 44(2):260–266, 1979.

[SW11] Ian Sharpe and Philip Welch. Greatly Erdős cardinals with some generalizations to the Chang and Ramsey properties. *Ann. Pure Appl. Log.*, 162(11):863–902, 2011.

[Zar82] Andrzej Zarach. Unions of ZF^--models which are themselves ZF^--models. In *Logic Colloquium '80 (Prague, 1980)*, volume 108 of *Stud. Logic Foundations Math.*, pages 315–342. North-Holland, Amsterdam, 1982.

The City University of New York, CUNY Graduate Center, Mathematics Program, 365 Fifth Avenue, New York, NY 10016

Email address: vgitman@gmail.com

URL: https://victoriagitman.github.io/

School of Mathematics, University of Bristol, Fry Building, Woodland Road, Bristol, BS8 1UG, UK

Email address: philipp.schlicht@bristol.ac.uk