An Outline of Inner Model Theory

John R. Steel

October 26, 2006

Contents

I. An Outline of Inner Model Theory

John R. Steel

1. Introduction

This article is an exposition of the theory of canonical inner models for large cardinal hypotheses, or *extender models*. We hope to convey the most important ideas and methods of this theory without sinking into the morass of fine-structural detail surrounding them. The resulting outline should be accessible to anyone familiar with the theory of iterated ultrapowers and $L[\mu]$ contained in Kunen's paper [14], and with the fine structure theory for L contained in Jensen's paper [11].

We shall present basic inner model theory in what is roughly the greatest generality in which it is currently known. This means that the theory we shall outline applies to extender models which may satisfy large cardinal hypotheses as strong as "There is a Woodin cardinal which is a limit of Woodin cardinals". Indeed, granted the iterability conjecture 6.5, the theory applies to extender models satisfying "There is a superstrong cardinal". Measuring the scope of the theory descriptive-set-theoretically, we can say that it applies to any extender model containing only reals which are ordinal definable over $L(\mathbb{R})$, and in fact to extender models containing somewhat more complicated reals. One can obtain a deeper analysis of a smaller class of inner models by restricting to models satisfying at most "There is a strong cardinal" (and therefore having only Δ^1_3 reals). The basic theory of this smaller class of models is significantly simpler, especially with regard to the structure of the iterated ultrapowers it uses. One can find expositions of this special case in the papers [19] and [20], and in the book [49].

Our outline of basic inner model theory occupies sections two through six of this paper. In sections seven and eight we present an application of this theory in descriptive set theory: we show that the model $\mathrm{HOD}^{L(\mathbb{R})}$ of all sets hereditarily ordinal definable in $L(\mathbb{R})$ is (essentially) an extender model.

The reader can find in [15] an exposition of basic inner model theory which is similar to this one, but somewhat less detailed. That paper then turns toward applications of inner model theory in the realm of consistency-strength lower bounds, an important area driving much of the evolution of the subject which we shall, nevertheless, avoid here. There is a more thorough and modern exposition of this area in [31]. We shall also abstain here from any extended discussion of the history of inner model theory. The reader can find philosophical/historical essays on the subject in the introductory sections of [18] and [15], and in [10], [24], [45], and in the chapter notes of [49].

2. Premice

The models we consider will be of the form $L[\vec{E}]$, where $\vec{E}$ is a coherent sequence of extenders. This framework seems quite general; indeed, it is plausible that there are models of the $L[\vec{E}]$ form for all the known large cardinal hypotheses. The framework is due, for the most part, to W.J. Mitchell ([21], [22]).

2.1. Extenders

An *extender* is a system of ultrafilters which fit together in such a way that they generate a single elementary embedding. The concept was originally introduced by Mitchell ([22]), and then simplified to its present form by Jensen.

2.1 Definition. Let $\kappa < \lambda$ and suppose that M is transitive and rudimentarily closed. We call E a (κ, λ) *extender over* M iff there is a nontrivial Σ_0-elementary embedding $j\colon M \to N$, with N transitive and rudimentarily closed, such that $\kappa = \text{crit}(j), \lambda < j(\kappa)$, and

$$E = \{\, (a, x) \mid a \in [\lambda]^{<\omega} \wedge x \subseteq [\kappa]^{|a|} \wedge x \in M \wedge a \in j(x)\}.$$

We say in this case that E is *derived from* j, and write $\kappa = \text{crit}(E), \lambda = \text{lh}(E)$.

If the requirement that N be transitive is weakened to $\lambda \subseteq \text{wfp}(N)$, where $\text{wfp}(N)$ is the wellfounded part of N, then we call E a (κ, λ) *pre-extender over* M. For the most part, this weakening is important only in the sort of details we intend to suppress.

If E is a (κ, λ) pre-extender over M and $a \in [\lambda]^{<\omega}$, then setting $E_a = \{\, x \mid (a, x) \in E\}$, we have that E_a is an M, κ complete nonprincipal ultrafilter on the field of sets $P([\kappa]^{|a|}) \cap M$. Thus we can form the ultrapower $\text{Ult}(M, E_a)$. The fact that all the E_a's come from the same embedding implies that there is a natural direct limit of the $\mathrm{Ult}(M, E_a)$'s, and we call this direct limit $\mathrm{Ult}(M, E)$. We can present $\mathrm{Ult}(M, E)$ more concretely as follows.

Let E be a (κ, λ) pre-extender over M. Let us identify finite sets of ordinals with their increasing enumerations. Let $a, c \in [\lambda]^{<\omega}$ with $a \subseteq c$, and let s be the increasing enumeration of $\{\, i \mid c(i) \in a\}$. For $x \subseteq [\kappa]^{|a|}$, we set

$$x_{ac} = \{\, u \in [\kappa]^{|c|} \mid u \circ s \in x\}.$$

If we think of x as a $|a|$-ary predicate on κ, then x_{ac} is just the result of blowing it up to a $|c|$-ary predicate by adding dummy variables at spots corresponding to ordinals in $c \setminus a$. It is easy to see that

$$x \in E_a \Leftrightarrow x_{ac} \in E_c.$$

That this is true of all x, a, c is a property of E known as *compatibility*. Notice that it really is a property of E alone; M only enters in through $P(\kappa) \cap M$, and E determines $P(\kappa) \cap M$. Similarly, if f is a function with domain $[\kappa]^{|a|}$, then f_{ac} is the function with domain $[\kappa]^{|c|}$ given by $f_{ac}(u) = f(u \circ s)$, which comes from f by adding the appropriate dummy variables. It is easy to see that E has the following property, known as *normality* : if $a \in [\lambda]^{<\omega}$, $i < |a|$, $f \in M$ is a function with $\mathrm{dom}(f) = [\kappa]^{|a|}$, and

$$\text{for } E_a \text{ a.e. } u,\ f(u) \in u(i)\ ,$$

then

$$\exists \xi < a(i)(f_{a, a\cup\{\xi\}}(v) = v(j) \text{ for } E_{a\cup\{\xi\}} \text{ a.e. } v),$$

where j is such that ξ is the j^{th} element of $a \cup \{\xi\}$. (Just take $\xi = j(f)(a)$, where E is derived from j.) Again, normality is a property of E alone.

Suppose M is transitive and rudimentarily closed, and that $E = \langle E_a \mid a \in [\lambda]^{<\omega}\rangle$ is a family of M, κ complete ultrafilters, E_a on $[\kappa]^{|a|}$, having the compatibility and normality properties. We construct $\mathrm{Ult}(M, E)$ as follows. Suppose $a, b \in [\kappa]^{<\omega}$ and f, g are functions in M with domains $[\kappa]^{|a|}$ and $[\kappa]^{|b|}$; then we put

$$\langle a, f\rangle \sim \langle b, g\rangle \text{ iff for } E_{a\cup b} \text{ a.e. } u\ (f_{a,a\cup b}(u) = g_{b,a\cup b}(u)).$$

(Here and in the future we use the "almost every" quantifier: given a filter F, we say $\phi(u)$ holds for F a.e. u iff $\{u \mid \phi(u)\} \in F$.) It is easy to check that $\sim$ is an equivalence relation; we use $[a, f]^M_E$ to denote the equivalence class of $\langle a, f\rangle$, and omit the subscript and superscript when context permits. Let

$$[a, f]\ \tilde{\in}\ [b, g] \text{ iff for } E_{a\cup b} \text{ a.e. } u\ (f_{a,a\cup b}(u) \in g_{b,a\cup b}(u)).$$

Then $\mathrm{Ult}(M, E)$ is the structure consisting of the set of all $[a, f]$ together with $\tilde{\in}$. We shall identify the wellfounded part of $\mathrm{Ult}(M, E)$ with its transitive isomorph, so that $\tilde{\in} = \in$ on the wellfounded part.

Suppose also that M satisfies the Axiom of Choice, as will indeed be the case in our applications. We then have Łoś's theorem for Σ_0 formulae, in that if φ is Σ_0 and $c = \bigcup_{i=1}^n a_i$, then

$$\mathrm{Ult}(M, E) \models \varphi[[a_1, f_1], \ldots, [a_n, f_n]]$$

if and only if

$$\text{for } E_c \text{ a.e. } u \ (M \models \varphi[(f_1)_{a_1 c}(u), \ldots, (f_n)_{a_n c}(u)]).$$

(The full Łoś theorem may fail, as M may not satisfy enough ZFC.) It follows that the canonical embedding

$$i_E^M : M \to \mathrm{Ult}(M, E)$$

is Σ_1-elementary, where i_E^M is given by $i_E^M(x) = [\{0\}, c_x]$, with $c_x(\alpha) = x$ for all α.

We have $[a, \mathrm{id}] = a$ for all $a \in [\lambda]^{<\omega}$ by an easy induction using the normality of E. From this and Łoś's theorem we get

$$x \in E_a \Leftrightarrow a \in i_E^M(x),$$

for all a, x, and

$$[a, f] = i_E^M(f)(a),$$

for all a, f. The first of these facts implies that E is the (κ, λ) pre-extender over M derived from i_E^M. Thus compatibility and normality are equivalent to pre-extenderhood; moreover, if E is a (κ, λ) pre-extender over Q, then E is also a (κ, λ) pre-extender over any transitive, rudimentarily closed M such that $P(\kappa) \cap M = P(\kappa) \cap Q$. It is definitely not the case, however, that the wellfoundedness of $\mathrm{Ult}(Q, E)$ implies the wellfoundedness of $\mathrm{Ult}(M, E)$.

If E is derived from $j : M \to N$, then there is a natural embedding $k : \mathrm{Ult}(M, E) \to N$ given by $k([a, f]) = j(f)(a)$, and the diagram

$$\begin{array}{ccc} M & \xrightarrow{\ j\ } & N \\ & {\scriptstyle i_E}\searrow & \uparrow{\scriptstyle k} \\ & & \mathrm{Ult}(M, E) \end{array}$$

commutes. It is easy to see that $k \restriction \lambda = \mathrm{id}$.

If E is a (κ, λ) pre-extender over M and $\xi \leq \lambda$, then we set $E \restriction \xi = \{(a, x) \in E \mid a \subseteq \xi\}$. There is a natural embedding σ from $\mathrm{Ult}(M, E \restriction \xi)$ into $\mathrm{Ult}(M, E)$ given by: $\sigma([a, f]^M_{E\restriction\xi}) = [a, f]^M_E$. We call ξ a *generator* of E just in case $\xi = \mathrm{crit}(\sigma)$; that is, $\xi \neq [a, f]^M_E$ for all $f \in M$ and $a \subseteq \xi$. The idea is that in this case $E \restriction \xi + 1$ has more information than $E \restriction \xi$, in that it determines a "bigger" ultrapower. The smallest generator of E is κ. All other generators are $> \kappa^{+M}$.

2.2 Definition. If E is a (κ, λ) pre-extender over M, then $\nu(E) = \sup(\kappa^{+M} \cup \{\xi + 1 \mid \xi \text{ is a generator of } E\})$. We call $\nu(E)$ the *support* of E.

The (κ, λ) extender derived from j can capture significantly more of the strength of j than the normal measure (that is, $(\kappa, \kappa + 1)$ extender) derived from j. For example, if $|V_\alpha^N|^N \leq \lambda$, then the existence of the factor map k implies that $V_\alpha^N = V_\alpha^{\mathrm{Ult}(M,E)}$. So if there is an embedding $j\colon V \to N$ such that $V_{\mathrm{crit}(j)+2} \subseteq N$, then there is an extender whose ultrapower gives rise to such an embedding. Indeed, if we remove the requirement that $\lambda < j(\kappa)$ from the definition of "extender", the results just discussed still go through, and we see that any embedding can be fully captured by such a generalized extender. We have included the restriction $\lambda < j(\kappa)$ in 3.1 only because nothing we shall prove here requires these "long" extenders, and it simplifies the exposition.

2.2. Fine Extender Sequences

Our models will be constructed from *coherent sequences* of extenders. Roughly speaking, this means that each E_α is either trivial (i.e. $E_\alpha = \emptyset$), or is an extender over $L[\vec{E} \upharpoonright \alpha]$ satisfying certain conditions. The extenders in a coherent sequence must appear in order of increasing strength, in that $\beta < \alpha$ implies $i_{E_\alpha}(\vec{E})_\beta = \vec{E}_\beta$. There can be no gaps, in that $i_{E_\alpha}(\vec{E})_\alpha = \emptyset$. These two conditions constitute *coherence*, a key idea which goes back to [21]. There are further conditions on the extender sequences we consider which insure that if $E_\alpha \neq \emptyset$, then α is completely determined by the embedding coded in E_α; this prevents us from coding random information into our model via the indexing of its extenders. There are different ways of handling the details here, all of which lead to the same class of models in the end. We shall adopt the indexing scheme of [26].

We shall use the Jensen J-hierarchy to stratify our models. If A is any set or class,

$$L[A] = \bigcup_{\alpha \in \mathrm{On}} J_\alpha^A,$$

where $J_0^A = \emptyset$, $J_\lambda^A = \bigcup_{\alpha<\lambda} J_\alpha^A$ for λ limit, and

$$J_{\alpha+1}^A = \mathrm{rud}^A(J_\alpha^A),$$

the closure of $J_\alpha^A \cup \{J_\alpha^A\}$ under rudimentary functions and the function $x \mapsto A \cap x$. If $\vec{E}$ is a sequence, then we shall abuse notation slightly by writing $J_\alpha^{\vec{E}}$ for J_α^A, where $A = \{(\beta, z) \mid z \in E_\beta\}$. In the case of interest to us, each E_α is either $\emptyset$ or a pre-extender over $J_\alpha^{\vec{E}}$ of length α, and $E_\alpha = \emptyset$ if α is a successor ordinal. It follows then that $J_\alpha^A = J_\alpha^{\vec{E}\upharpoonright\alpha}$ and $E_\alpha \subseteq J_\alpha^{\vec{E}\upharpoonright\alpha}$;

from this we get that for all $X \subseteq J_\alpha^{\vec{E}}$,

$$X \in J_{\alpha+1}^{\vec{E}} \text{ iff } X \text{ is definable over } (J_\alpha^{\vec{E}}, \in, \vec{E} \restriction \alpha, E_\alpha),$$

where the definition of X may use parameters from $J_\alpha^{\vec{E}}$. (See [38, 1.4].) Although we are officially using the J-hierarchy, we might have used Gödel's L-hierarchy instead, and the reader who prefers can change the J's to L's in what follows. (The advantages of using the J hierarchy show up in details we shall suppress.)

There is one important point here: in our setup, if $E_\alpha \neq \emptyset$, then E_α is an extender *over* $J_\alpha^{\vec{E}}$; it only measures the subsets of its critical point constructed before stage α. There may or may not be subsets of $\mathrm{crit}(E_\alpha)$ constructed in $L[\vec{E}]$ after stage α; if there are, then E_α does not measure them, and so fails to be an extender over all of $L[\vec{E}]$. The idea of adding such "partial" extenders to our sequences $\vec{E}$ is due to S. Baldwin and W. Mitchell. It leads to a stratification of core models much simpler than the sort studied previously. In particular, the hierarchies we shall study are (strongly) acceptable in the sense of [6].

2.3 Definition. A set A is *acceptable at α* iff
$\forall \beta < \alpha \forall \kappa ((P(\kappa) \cap (J_{\beta+1}^A \setminus J_\beta^A) \neq \emptyset) \to J_{\beta+1}^A \models |J_\beta^A| \leq \kappa)$.

Notice that if A is acceptable at α and $J_\alpha^A \models$ "κ^+ exists", then $J_\alpha^A \models$ "$P(\kappa)$ exists and $P(\kappa) \subseteq J_{\kappa^+}^A$". It follows that GCH is true in J_α^A.

It is a basic fact in the fine structure of L that $\emptyset$ is acceptable at all α. On the other hand, if μ is a normal measure on κ, then μ is not acceptable at $\kappa+2$, since there are subsets of ω in $J_{\kappa+2}^\mu \setminus J_{\kappa+1}^\mu$ (such as $0^\sharp$), while κ is not countable in $J_{\kappa+2}^\mu$ (or anywhere else).

Suppose that E is a pre-extender over M, and that $M \models \kappa^+$ exists, where $\kappa = \mathrm{crit}(E)$. Let $\nu = \nu(E)$ and $\eta = (\nu^+)^{\mathrm{Ult}(M,E)}$ be in the wellfounded part of $\mathrm{Ult}(M, E)$. We shall use the ordinal η to index E in extender sequences. Let E^* be the (κ, η) pre-extender derived from E. It is easy to check that $\nu = \nu(E^*)$ and $E \restriction \nu = E^* \restriction \nu$, so that E and E^* are equivalent. For a minor technical reason, it is E^* which we shall index at η. We call E^* the *trivial completion* of E.

We shall need the following very technical concept. Let E be an extender over M. We say that E is of *type Z* iff $\nu(E) = \lambda+1$ for some limit ordinal λ such that (a) $\lambda = \nu(E \restriction \lambda)$, and (b) $(\lambda^+)^{\mathrm{Ult}(M,E)} = (\lambda^+)^{\mathrm{Ult}(M,E \restriction \lambda)}$. Notice that our indexing convention would require that the trivial completions E^* and $(E \restriction \lambda)^*$ be indexed at the same place, if E is type Z. We resolve this conflict by giving $(E \restriction \lambda)^*$ preference, and therefore putting no type Z extenders on our sequences.

We are ready for one of the most important definitions in this article.

2.4 Definition. A *fine extender sequence* is a sequence $\vec{E}$ such that for each $\alpha \in \mathrm{dom}(\vec{E})$, $\vec{E}$ is acceptable at α, and either $\vec{E}_\alpha = \emptyset$, or E_α is a (κ, α) pre-extender over $J^{\vec{E}}_\alpha$ for some κ such that $J^{\vec{E}}_\alpha \models \kappa^+$ exists, and:

1. E_α is the trivial completion of $E_\alpha \restriction \nu(E_\alpha)$, and E_α is not of type Z,
2. (Coherence) $i(\vec{E} \restriction \kappa) \restriction \alpha = \vec{E} \restriction \alpha$ and $i(\vec{E} \restriction \kappa)_\alpha = \emptyset$, where $i\colon J^{\vec{E}}_\alpha \to \mathrm{Ult}(J^{\vec{E}}_\alpha, E_\alpha)$ is the canonical embedding, and
3. (Closure under initial segment) for any η such that $(\kappa^+)^{J^{\vec{E}}_\alpha} \leq \eta < \nu(E_\alpha)$, $\eta = \nu(E_\alpha \restriction \eta)$, and $E_\alpha \restriction \eta$ is not of type Z, one of the following holds:
   (a) there is $\gamma < \alpha$ such that E_γ is the trivial completion of $E_\alpha \restriction \eta$, or
   (b) $E_\eta \neq \emptyset$, and letting $j\colon J^{\vec{E}}_\eta \to \mathrm{Ult}(J^{\vec{E}}_\eta, E_\eta)$ be the canonical embedding and $\mu = \mathrm{crit}(j)$, there is a $\gamma < \alpha$ such that $j(\vec{E} \restriction \mu)_\gamma$ is the trivial completion of $E_\alpha \restriction \eta$.

2.5 Remarks. Let $\vec{E}$ be a fine extender sequence, $E_\alpha \neq \emptyset$, and let $i\colon J^{\vec{E}}_\alpha \to \mathrm{Ult}(J^{\vec{E}}_\alpha, E_\alpha)$ be the canonical embedding.

1. Although $\mathrm{Ult}(J^{\vec{E}}_\alpha, E_\alpha)$ may be illfounded, it must be that $\alpha + 1$ is contained in the wellfounded part of the ultrapower, and this is enough to make sense of the conditions in 2.4. Also, $\vec{E} \restriction \beta \in J^{\vec{E}}_\alpha$ for all $\beta < \alpha$, and it is natural then to set $i(\vec{E} \restriction \alpha) = \bigcup_{\beta<\alpha} i(\vec{E} \restriction \beta)$.
2. Let $\nu = \nu(E_\alpha)$. By coherence, $J^{i(\vec{E}\restriction\alpha)}_\alpha = J^{\vec{E}}_\alpha$. Since $\alpha = \nu^+$ in $\mathrm{Ult}(J^{\vec{E}}_\alpha, E_\alpha)$, and since $i(\vec{E} \restriction \alpha)$ is acceptable at all $\beta < \sup_{\gamma<\alpha} i(\gamma)$ by Łoś's theorem (acceptability being a Π_1 property of $\vec{E} \restriction \alpha$ whenever α is a limit), there are no cardinals $> \nu$ in $J^{\vec{E}}_\alpha$. The ordinal ν itself may be a successor ordinal. It is not hard to show that if ν is a limit ordinal, then ν is a cardinal in both $J^{\vec{E}}_\alpha$ and $\mathrm{Ult}(J^{\vec{E}}_\alpha, E_\alpha)$.
3. Let $\kappa = \mathrm{crit}(E_\alpha)$. By clause 1 of 2.4, there is a map of $(P(\kappa) \cap J^{\vec{E}}_\nu) \times [\nu]^{<\omega}$ onto α, the map being in $J^{\vec{E}}_{\alpha+1}$. Thus α is not a cardinal in $J^{\vec{E}}_{\alpha+1}$.
4. For the fine sequences $\vec{E}$ we construct, E_α is an extender over $L[\vec{E} \restriction \alpha]$, and $\alpha = \nu(E_\alpha)^+$ in both $L[\vec{E} \restriction \alpha]$ and $\mathrm{Ult}(L[\vec{E} \restriction \alpha], E_\alpha)$. This in fact follows from the clauses of 2.4 if we can iterate from $J^{\vec{E}}_\alpha$ via E_α and its images On times.

Definition 2.4 diverges slightly from the definition of "good extender sequence" in [26, section 1]. The latter definition is wrong, in that the extender sequences constructed in section 11 of [26] and section 6 of the present paper do not satisfy it. This was recently shown by Martin Zeman. The problem lies in the initial segment condition of [26], which does not contain the proviso in clause 3 of 2.4 that $E_\alpha \restriction \eta$ is not of type Z. Zeman showed that on any reasonably rich sequence of the sort constructed in [26] or section 6 of this paper, there must be extenders E such that for some $\eta < \nu(E)$, $\eta = \nu(E \restriction \eta)$ and $E \restriction \eta$ is of type Z.[1] Our indexing scheme implies that the conclusion of clause 3 of 2.4 must then fail for one of $E \restriction \eta$ and $E \restriction (\eta - 1)$. R.D. Schindler and W.H. Woodin independently found the correct axiomatization of the properties of the extender sequences constructed in [26] and here: one simply adds that type Z extenders do not occur on the sequence, and weakens the initial segment condition to take this into account.[2]

It might be hoped that alternative 3(b) of 2.4 could be dropped, but we suspect that if $L[\vec{E}]$ is to have a Woodin cardinal, or even many strong cardinals, then one cannot demand this stronger form of the initial segment condition. The initial segment condition in 2.4 is crucial in the proof that the comparison process terminates. We need some form of it as an axiom on our extender sequences in order to get a decent theory going.

Following a suggestion of S. Friedman, R. Jensen has investigated an indexing of extenders different from the sort described in 2.4 (cf. [49]). In this framework, the extender E is indexed at the cardinal successor of $i_E(\text{crit}(E))$ in its ultrapower. For any fine extender sequence $\vec{E}$ there is a Friedman-Jensen sequence $\vec{F}$ such that $L[\vec{E}] = L[\vec{F}]$, and vice-versa, so both approaches lead to the same class of models. The Friedman-Jensen hierarchy grows more slowly than the one we are using, in that certain extenders are put on a Friedman-Jensen sequence which only appear on ultrapowers of its translation to a fine extender sequence. In particular, one can drop the counterpart of clause 3(b) of 2.4 in the Friedman-Jensen approach.

2.6 Definition. A *potential premouse*(or ppm) is a structure of the form $(J_\alpha^{\vec{E}}, \in, \vec{E} \restriction \alpha, E_\alpha)$, where $\vec{E}$ is a fine extender sequence. We use $\mathcal{J}_\alpha^{\vec{E}}$ to denote this structure.

2.7 Definition. Let $\mathcal{M} = \mathcal{J}_\alpha^{\vec{E}}$ be a ppm. We say $\mathcal{M}$ is *active* if $E_\alpha \neq \emptyset$, and passive otherwise. If $\mathcal{M}$ is active, then letting $\nu = \nu(E_\alpha)$ and $\kappa = \text{crit}(E_\alpha)$, we say $\mathcal{M}$ is *type* I if $\nu = (\kappa^+)^{\mathcal{M}}$, $\mathcal{M}$ is *type* II if ν is a successor ordinal, and $\mathcal{M}$ is *type* III if ν is a limit ordinal $> (\kappa^+)^{\mathcal{M}}$.

[1] See [37], which also corrects some further errors in [26] and [32].

[2] The "proof" in [26] of the stronger initial segment condition goes wrong in the proof of theorem 10.1, where on p. 98, in the "$\eta = \gamma$" case, the authors ignore the possibility that G might be of type Z. Schindler found this error. What the argument of [26] does prove is the weaker initial segment condition of 2.4.

The distinctions among potential premice introduced in 2.7 are mostly important in the sort of details we shall suppress, but we need them in order to make certain definitions formally correct.

2.3. The Levy Hierarchy, Cores, and Soundness

Although it is possible to avoid fine structure theory entirely in the proofs of basic facts about smaller core models (for example, in the proof that $L[\mu] \models \mathsf{GCH}$), there is little one can show about larger core models (such as the minimal model satisfying "there is a Woodin cardinal") without fine structure theory.[3] It seems that one must marshall all one's forces in good order in order to advance; indeed, the very definition of the models requires fine structural notions. Therefore, in order to be able even to state precise definitions and theorems, we must lay out some of the fine structure theory of definability over potential premice.

We shall simplify matters by concentrating on the representative special case of Σ_1 definability, and indicating only briefly the appropriate notions at higher levels of the Levy hierarchy. In those few places where fine structural details crop up in proofs we give in later sections, the reader will lose little by considering only the special case $\Sigma_{n+1} = \Sigma_1$. The reader should see [38] for an excellent full account of the fine structural underpinnings of the theory we present here.[4]

The subsets of $J_\alpha^{\vec{E}}$ belonging to $J_{\alpha+1}^{\vec{E}}$ are precisely those first-order definable over the ppm $\mathcal{J}_\alpha^{\vec{E}}$, but unfortunately, this structure is not amenable if $E_\alpha \neq \emptyset$.

2.8 Definition. A structure $(M, \in, A_1, A_2, \ldots)$ is *amenable* iff $\forall x \in M \forall i (A_i \cap x \in M)$.

Since amenability is important in basic ways[5], we need an amenable structure with the same definable subsets as $(J_\alpha^{\vec{E}}, \in, \vec{E} \restriction \alpha, E_\alpha)$; that is, we need an amenable predicate coding E_α. The following lemma is the key.

---

[3] Fine stucture theory begins with Jensen's landmark paper [11]. R. Solovay (unpublished manuscript) extended Jensen's work to $L[\mu]$, and then Dodd and Jensen showed in [6], [7], [8], and [5] just how remarkably fruitful this extension could be. Dodd, Jensen, and Mitchell extended this older fine structure theory to still larger core models (in [23], and unpublished work), but the complexities became unmanageable just past core models with strong cardinals. The Baldwin-Mitchell idea of putting partial extenders on a coherent sequence cut through these difficulties. [26] was the first paper to develop the Baldwin-Mitchell idea.

[4] Jensen has developed a more general fine structure theory, using terminology somewhat different from that used here. See [46] or [49]. We shall not need this extra generality here.

[5] For example, in the proof that satisfaction for Σ_1 formulae is Σ_1, and in the proof of the Łoś's theorem for Σ_0 formulae. See [38, 1.12, 8.4]

2.9 Lemma. *Let $\vec{E}$ be a fine extender sequence, $E_\alpha \neq \emptyset$, $\kappa = \mathrm{crit}(E_\alpha)$, and $\nu = \nu(E_\alpha)$; then for any $\eta < \alpha$ and $\xi < (\kappa^+)^{J^{\vec{E}}_\alpha}$, $E_\alpha \cap ([\eta]^{<\omega} \times J^{\vec{E}}_\xi) \in J^{\vec{E}}_\alpha$. Moreover, if for $\xi < (\kappa^+)^{J^{\vec{E}}_\alpha}$ we set*

$$\gamma_\xi = \textit{ least } \gamma < \alpha \textit{ such that } E_\alpha \cap ([\nu]^{<\omega} \times J^{\vec{E}}_\xi) \in J^{\vec{E}}_\gamma,$$

then

$$\sup(\{\gamma_\xi \mid \xi < (\kappa^+)^{J^{\vec{E}}_\alpha}\}) = \alpha.$$

Proof. Fix $\xi < (\kappa^+)^{J^{\vec{E}}_\alpha}$. Let $\langle A_\beta \mid \beta < \kappa \rangle$ be an enumeration of $\bigcup_{n<\omega}(P([\kappa]^n) \cap J^{\vec{E}}_\xi)$ belonging to $J^{\vec{E}}_\alpha$. Let

$$i\colon J^{\vec{E}}_\alpha \to \mathrm{Ult}(J^{\vec{E}}_\alpha, E_\alpha)$$

be the canonical embedding, and notice that

$$\langle i(A_\beta) \mid \beta < \kappa \rangle \in \mathrm{Ult}(J^{\vec{E}}_\alpha, E_\alpha),$$

since $\langle i(A_\beta) \mid \beta < \kappa \rangle = i(\langle A_\beta \mid \beta < \kappa \rangle) \restriction \kappa$. But

$$E_\alpha \cap ([\eta]^{<\omega} \times J^{\vec{E}}_\xi) = \{(a, A_\beta) \mid a \in [\eta]^{<\omega} \wedge a \in i(A_\beta)\},$$

so $E_\alpha \cap ([\eta]^{<\omega} \times J^{\vec{E}}_\xi) \in \mathrm{Ult}(J^{\vec{E}}_\alpha, E_\alpha)$. Since α is a cardinal in this ultrapower, we have by acceptability that $E_\alpha \cap ([\eta]^{<\omega} \times J^{\vec{E}}_\xi) \in J^{i(\vec{E}\restriction\alpha)}_\alpha$. But $J^{i(\vec{E}\restriction\alpha)}_\alpha = J^{\vec{E}}_\alpha$ by coherence, so we are done with the first part of the lemma.

In order to show the γ_ξ are cofinal in α, it suffices to show that whenever $A \subseteq \nu$ and $A \in \mathrm{Ult}(J^{\vec{E}}_\alpha, E_\alpha)$, then there is a ξ such that $A \in J^{\vec{E}}_{\gamma_\xi+1}$. So fix such an A, and let $A = [a, f]$, where $a \subseteq \nu$ and, without loss of generality, $f \in J^{\vec{E}}_\alpha$ and $f\colon J^{\vec{E}}_\kappa \to J^{\vec{E}}_\kappa$. By acceptability, we have $\xi < (\kappa^+)^{J^{\vec{E}}_\alpha}$ such that $f \in J^{\vec{E}}_\xi$. Now for $\eta < \nu$, $\eta \in A \Leftrightarrow$ for $(E_\alpha)_{a\cup\{\eta\}}$ a.e. u, $\mathrm{id}_{\{\eta\},a\cup\{\eta\}}(u) \in f(u)$, and the set to be measured in answering this question about η is in $J^{\vec{E}}_\xi$. Thus A can be computed from $E_\alpha \cap ([\nu]^{<\omega} \times J^{\vec{E}}_\xi)$, so $A \in J^{\vec{E}}_{\gamma_\xi+1}$. ⊣

Given now a fine extender sequence $\vec{E}$ with $E_\alpha \neq \emptyset$, we can code E_α as follows: let E^c_α be the set of quadruples (γ, ξ, a, x) such that

$$(\nu(E_\alpha) < \gamma < \alpha) \wedge (\mathrm{crit}(E_\alpha) < \xi < (\mathrm{crit}(E_\alpha)^+)^{J^{\vec{E}}_\alpha}) \wedge$$

$$(E_\alpha \cap ([\nu(E_\alpha)]^{<\omega} \times J^{\vec{E}}_\xi) \in J^{\vec{E}}_\gamma) \wedge ((a, x) \in (E_\alpha \cap ([\gamma]^{<\omega} \times J^{\vec{E}}_\xi))).$$

It follows from Lemma 2.9 that $(J^{\vec{E}}_\alpha, \in, \vec{E} \restriction \alpha, E^c_\alpha)$ is amenable.

Certain ordinal parameters are important in the description of a ppm. Let $\mathcal{M} = \mathcal{J}_\alpha^{\vec{E}}$. If $\mathcal{M}$ is active, then we set

$$\nu^{\mathcal{M}} = \nu(E_\alpha) \text{ and } \mu^{\mathcal{M}} = \text{crit}(E_\alpha).$$

If $\mathcal{M}$ is passive, set $\nu^{\mathcal{M}} = \mu^{\mathcal{M}} = 0$. If $\mathcal{M}$ is active of type II , then there is a longest non-type-Z proper initial segment F of E_α containing properly less information than E_α itself, and we let $\gamma^{\mathcal{M}}$ determine where F appears on $\vec{E}$ or an ultrapower of $\vec{E}$. More precisely, set

$$F = \begin{cases} (E_\alpha \restriction \nu^{\mathcal{M}} - 1)^* & \text{if } (E_\alpha \restriction \nu^{\mathcal{M}} - 1)^* \text{ is not type Z} \\ (E_\alpha \restriction \nu(E_\alpha \restriction \nu^{\mathcal{M}} - 1) - 1)^* & \text{otherwise.} \end{cases}$$

Then we let

$$\gamma^{\mathcal{M}} = \text{ the unique } \xi \in \text{dom}(\vec{E}) \text{ such that } F = E_\xi,$$

if there is such a ξ.[6] If there is no such ξ, then setting $\eta = \nu(F)$, we have by 3(b) of 2.4 that F is on the extender sequence of $\text{Ult}(J_\eta^{\vec{E}}, E_\eta)$. We then let

$$\gamma^{\mathcal{M}} = (\eta, a, f), \text{ where } F = [a, f]_{E_\eta}^{J_\eta^{\vec{E}}},$$

and (a, f) is least in the order of construction on $J_\eta^{\vec{E}}$ with this property. Finally, if $\mathcal{M}$ is not active type II, then we set $\gamma^{\mathcal{M}} = 0$.

Since we shall put these parameters in all hulls we form, we might as well have names for them in our language.

2.10 Definition. $\mathcal{L}$ is the language of set theory with additional constant symbols $\dot{\mu}, \dot{\nu}, \dot{\gamma}$, and additional unary predicate symbols $\dot{E}$ and $\dot{F}$.

2.11 Definition. Let $\mathcal{M} = \mathcal{J}_\alpha^{\vec{E}}$ be a ppm; then the Σ_0 *code* of $\mathcal{M}$, or $\mathcal{C}_0(\mathcal{M})$, is the $\mathcal{L}$-structure $\mathcal{N}$ given by:

1. if $\mathcal{M}$ is passive, then $\mathcal{N}$ has universe $J_\alpha^{\vec{E}}$, $\dot{E}^{\mathcal{N}} = \vec{E} \restriction \alpha$, $\dot{F}^{\mathcal{N}} = \emptyset$, and $\dot{\mu}^{\mathcal{N}} = \dot{\nu}^{\mathcal{N}} = \dot{\gamma}^{\mathcal{N}} = 0$;

2. if $\mathcal{M}$ is active of types I or II, then $\mathcal{N}$ has universe $J_\alpha^{\vec{E}}$, $\dot{E}^{\mathcal{N}} = \vec{E} \restriction \alpha$, $\dot{F}^{\mathcal{N}} = E_\alpha^*$ (where E_α^* is the amenable coding of E_α), and $\dot{\mu}^{\mathcal{N}} = \mu^{\mathcal{M}}$, $\dot{\nu}^{\mathcal{N}} = \nu^{\mathcal{M}}$, and $\dot{\gamma}^{\mathcal{N}} = \gamma^{\mathcal{M}}$;

3. if $\mathcal{M}$ is active type III, then letting $\nu = \nu(E_\alpha)$, $\mathcal{N}$ has universe $J_\nu^{\vec{E}}$, $\dot{E}^{\mathcal{N}} = \vec{E} \restriction \nu$, $\dot{F}^{\mathcal{N}} = E_\alpha \restriction \nu$, $\dot{\mu}^{\mathcal{N}} = \mu^{\mathcal{M}}$, and $\dot{\nu}^{\mathcal{N}} = \dot{\gamma}^{\mathcal{N}} = 0$.

[6] $\gamma^{\mathcal{M}} =$ lh(F) in this case.

The Σ_0 code $\mathcal{C}_0(\mathcal{M})$ is amenable; this follows from our lemma unless $\mathcal{M}$ is active type III, in which case it follows at once from the initial segment condition of 2.4. The reader may wonder why we treated the type III ppm differently in the definition above, but fortunately, the answer lies in fine structural details we shall avoid here.[7] The reader will lose nothing of importance if he pretends that all active premice are of type II. Notice that $\mathcal{M}$ is indeed coded into $\mathcal{C}_0(\mathcal{M})$; this is obvious unless $\mathcal{M}$ is active type III, and in that case we can recover $\mathcal{M}$ by forming $\mathrm{Ult}(\mathcal{C}_0(\mathcal{M}), \dot{F}^{\mathcal{C}_0(\mathcal{M})})$, then adding the trivial completion of $\dot{F}^{\mathcal{C}_0(\mathcal{M})}$ to its sequence at the proper place. There is little harm in identifying $\mathcal{M}$ with $\mathcal{C}_0(\mathcal{M})$.

We can now define the Σ_1 projectum, first standard parameter, and first core of a ppm $\mathcal{M}$.

2.12 Definition. Let $\mathcal{M}$ be a ppm; then the Σ_1 *projectum* of $\mathcal{M}$, or $\rho_1(\mathcal{M})$, is the least ordinal α such that for some boldface $\boldsymbol{\Sigma}_{\mathbf{1}}^{\mathcal{C}_0(\mathcal{M})}$ set $A \subseteq \alpha$, $A \notin \mathcal{C}_0(\mathcal{M})$. (Thus $\rho_1(\mathcal{M}) \leq \mathrm{On} \cap \mathcal{C}_0(\mathcal{M})$.)

Notice that the new set A may not be (lightface) Σ_1 definable. Since there is a $\Sigma_1^{\mathcal{C}_0(\mathcal{M})}$ map from the class of finite sets of ordinals onto $\mathcal{C}_0(\mathcal{M})$, we can take the parameter from which A is defined to be a finite set of ordinals. We standardize the parameter by minimizing it in a certain wellorder.

2.13 Definition. A *parameter* is a finite (perhaps empty) sequence $\langle \alpha_0, \ldots, \alpha_n \rangle$ of ordinals such that $\alpha_0 > \ldots > \alpha_n$. If $\mathcal{M}$ is a ppm, then the *first standard parameter of* $\mathcal{M}$, or $p_1(\mathcal{M})$, is the lexicographically least parameter p such that there is a $\Sigma_1^{\mathcal{C}_0(\mathcal{M})}(\{p\})$ set A such that $(A \cap \rho_1(\mathcal{M})) \notin \mathcal{C}_0(\mathcal{M})$.

2.14 Definition.
1. For any $\mathcal{L}$ structure $\mathcal{Q}$ and set $X \subseteq |\mathcal{Q}|$, $\mathcal{H}_1^{\mathcal{Q}}(X)$ is the transitive collapse of the substructure of $\mathcal{Q}$ whose universe consists of all $y \in |\mathcal{Q}|$ such that $\{y\}$ is $\Sigma_1^{\mathcal{Q}}$ definable from parameters in X.
2. For any ppm $\mathcal{M}$, the *first core* of $\mathcal{M}$, or $\mathcal{C}_1(\mathcal{M})$, is defined by: $\mathcal{C}_1(\mathcal{M}) = \mathcal{H}_1^{\mathcal{C}_0(\mathcal{M})}(\rho_1(\mathcal{M}) \cup \{p_1(\mathcal{M})\})$.

It is a routine matter to show that for any ppm $\mathcal{M}$, $\mathcal{C}_1(\mathcal{M})$ is the Σ_0 code of some ppm $\mathcal{N}$. One need only check that being a Σ_0 code can be expressed using Π_2 sentences of $\mathcal{L}$. (See [26, 2.5].)

We introduce two important ways in which the standard parameter $p_1(\mathcal{M})$ can behave well.

2.15 Definition. Let $\mathcal{M}$ be a ppm.

1. We say $p_1(\mathcal{M})$ is *1-universal* iff whenever $A \subseteq \rho_1(\mathcal{M})$ and $A \in \mathcal{C}_0(\mathcal{M})$, then $A \in \mathcal{C}_1(\mathcal{M})$.

[7] See [26, section 3]).

2. Let $p_1(\mathcal{M}) = \langle \alpha_0, \ldots, \alpha_n \rangle$. We say $p_1(\mathcal{M})$ is *1-solid* iff whenever $i \leq n$ and A is $\Sigma_1^{\mathcal{C}_0(\mathcal{M})}(\{\alpha_0, .., \alpha_{i-1}\})$, then $A \cap \alpha_i \in \mathcal{C}_0(\mathcal{M})$.

3. We say $\mathcal{M}$ is *1-solid* just in case $p_1(\mathcal{M})$ is 1-solid and 1-universal.

If $p_1(\mathcal{M})$ is 1-universal, then letting $\mathcal{C}_1(\mathcal{M}) = \mathcal{C}_0(\mathcal{N})$, one has that $\rho_1(\mathcal{N}) = \rho_1(\mathcal{M})$, and $p_1(\mathcal{N})$ is the image of $p_1(\mathcal{M})$ under the transitive collapse.[8] The 1-solidity of $p_1(\mathcal{M})$ is important in showing that $i(p_1(\mathcal{M})) = p_1(\mathcal{Q})$ for certain ultrapower embeddings $i \colon \mathcal{M} \to \mathcal{Q}$.[9]

2.16 Definition. $\mathcal{M}$ is *1-sound* iff $\mathcal{M}$ is 1-solid and $\mathcal{C}_1(\mathcal{M}) = \mathcal{C}_0(\mathcal{M})$.

Let $\mathcal{N}$ be the ppm whose Σ_0 code is $\mathcal{C}_1(\mathcal{M})$. It is easy to see that $\mathcal{C}_1(\mathcal{N}) = \mathcal{C}_1(\mathcal{M})$, so that if $\mathcal{N}$ is 1-solid, then $\mathcal{N}$ is 1-sound.

We should now go on and define the n^{th} projectum $\rho_n(\mathcal{M})$, the n^{th} standard parameter $p_n(\mathcal{M})$, and the n^{th} core $\mathcal{C}_n(\mathcal{M})$, as well as the notions of n-solidity and n-universality for $p_n(\mathcal{M})$ and n-soundness for $\mathcal{M}$, in the case $n > 1$. The definitions run parallel to those in the $n = 1$ case, but there are enough annoying details that we prefer to shirk our duty and refer the conscientious reader to [26, section 2]. (Formally speaking, these objects and notions are defined by induction on n in such a way that $\rho_n(\mathcal{M}), p_n(\mathcal{M})$, etc., only make sense if $\mathcal{M}$ is $(n-1)$-solid.) There is one point worth mentioning here, namely, $\rho_n(\mathcal{M})$, $p_n(\mathcal{M})$, $\mathcal{C}_n(\mathcal{M})$, etc., are defined *from the viewpoint of* $\mathcal{C}_{n-1}(\mathcal{M})$. For example, $\rho_2(\mathcal{M})$ is the least ordinal α such that there is an $r\Sigma_2^{\mathcal{C}_1(\mathcal{M})}$-in-parameters set $A \subseteq \alpha$ such that $A \notin \mathcal{C}_1(\mathcal{M})$.[10] The class of $\Sigma_2^{\mathcal{C}_0(\mathcal{M})}$ definable relations is not relevant at this (or any) point, since random information can be coded into such relations by iterating some $\mathcal{C}_0(\mathcal{N})$ above $\rho_1(\mathcal{N})$.[11]

2.17 Definition. Let $\mathcal{M}$ be a ppm; then $\mathcal{M}$ is *ω-solid* iff $\mathcal{M}$ is n-solid for all $n < \omega$, and $\mathcal{M}$ is *ω-sound* iff $\mathcal{M}$ is n-sound for all $n < \omega$. If $\mathcal{M}$ is ω-solid,

---

[8] Let r be the image of $p_1(\mathcal{M})$ under the collapse. As the collapse is the identity on $\rho_1(\mathcal{M})$, r defines over $\mathcal{C}_0(\mathcal{N})$ a new Σ_1 subset of $\rho_1(\mathcal{M})$, so that $\rho_1(\mathcal{N}) \leq \rho_1(\mathcal{M})$ and $p_1(\mathcal{N}) \leq_{\text{lex}} r$. It is easy to see $\rho_1(\mathcal{N}) \geq \rho_1(\mathcal{M})$. Finally, if $s <_{\text{lex}} r$ and $A \subseteq \rho_1(\mathcal{M})$ is $\Sigma_1^{\mathcal{C}_0(\mathcal{N})}$ definable from s, then $A \in \mathcal{M}$ by the minimality of $p_1(\mathcal{M})$, so $A \in \mathcal{N}$ by the universality of $p_1(\mathcal{M})$. Thus $r \leq_{\text{lex}} p_1(\mathcal{N})$.

[9] For any parameter $s <_{\text{lex}} p_1(\mathcal{M})$, let T_s be the Σ_1 theory in $\mathcal{C}_0(\mathcal{M})$ of parameters from $\rho_1(\mathcal{M}) \cup \{s\}$; then $T_s \in \mathcal{M}$ by the definition of $p_1(\mathcal{M})$. The solidity of $p_1(\mathcal{M})$ is equivalent to the assertion that the map $s \mapsto T_s$ is a member of $\mathcal{M}$.

[10] The $r\Sigma_2$ relations are, roughly speaking, just those which are Σ_1 definable from the function T, where $T(\eta, q) = \Sigma_1$ theory of parameters in $\eta \cup \{q\}$, for $\eta < \rho_1$, and $T(\eta, q) = 0$ if $\eta \geq \rho_1$.

[11] The following example is due to Mitchell. Suppose $\langle \kappa_i | i \in \omega \rangle$ are an increasing sequence of measurable cardinals of $\mathcal{N}$ with $\rho_1(\mathcal{N}) \leq \kappa_0$, and suppose $\mathcal{N}$ is 1-sound and iterable. Let $a \subseteq \omega$ be arbitrary. Let $\mathcal{M}$ result from iterating $\mathcal{N}$ by hitting a normal measure with critical point κ_i iff $i \in a$. Then a is $\Sigma_2^{\mathcal{M}}$ since $i \in a$ iff κ_i is not $\Sigma_1^{\mathcal{M}}$ definable from parameters in $\kappa_i \cup \{p_i(\mathcal{M})\}$.

then we let $\rho_\omega(\mathcal{M})$ be the eventual value of $\rho_n(\mathcal{M})$ and $\mathcal{C}_\omega(\mathcal{M})$ the eventual value of $\mathcal{C}_n(\mathcal{M})$ as $n \to \omega$.

If $n < m$, then $\rho_n(\mathcal{M}) \geq \rho_m(\mathcal{M})$, so there is indeed an eventual value for $\rho_n(\mathcal{M})$, and hence $\mathcal{C}_n(\mathcal{M})$). Clearly, $\mathcal{M}$ is ω-sound iff $\mathcal{C}_0(\mathcal{M}) = \mathcal{C}_\omega(\mathcal{M})$. All levels of the core models we shall construct will be ω-sound. Nevertheless, we must study potential premice which are not ω-sound, since these can be produced from ω-sound potential premice by taking ultrapowers. (See 2.23 below.) However, all proper initial segments of such an ultrapower are ω-sound, so we can restrict ourselves to ppm all of whose proper initial segments are ω-sound.

2.18 Definition. Let $\mathcal{M} = \mathcal{J}_\alpha^{\vec{E}}$ be a ppm, and let $\beta \leq \alpha$; then we write $\mathcal{J}_\beta^{\mathcal{M}}$ for $\mathcal{J}_\beta^{\vec{E}}$, and call $\mathcal{J}_\beta^{\mathcal{M}}$ an *initial segment* of $\mathcal{M}$. We write $\mathcal{N} \trianglelefteq \mathcal{M}$ ($\mathcal{N}$ is an initial segment of $\mathcal{M}$) iff $\exists\beta(\mathcal{N} = \mathcal{J}_\beta^{\mathcal{M}})$, and $\mathcal{N} \triangleleft \mathcal{M}$ ($\mathcal{N}$ is a proper initial segment of $\mathcal{M}$) iff $\exists\beta < \alpha(\mathcal{N} = \mathcal{J}_\beta^{\mathcal{M}})$.

2.19 Definition. A *premouse* is a potential premouse all of whose proper initial segments are ω-sound. A *coded premouse* is a structure of the form $\mathcal{C}_0(\mathcal{M})$, where $\mathcal{M}$ is a premouse.

It is easy to see that $\vec{E}$ is an extender sequence with domain α such that all proper initial segments of $\mathcal{J}_\alpha^{\vec{E}}$ are ω-sound, then $\vec{E}$ is acceptable at α. Indeed, soundness is simply a refinement of acceptability, in that we demand that whenever a new subset of κ appears in $J_{\tau+1}^{\vec{E}} - J_\tau^{\vec{E}}$, the surjection $f \in J_{\tau+1}^{\vec{E}}$ from κ onto $J_\tau^{\vec{E}}$ required by acceptability must actually be definable over $\mathcal{J}_\tau^{\vec{E}}$ at the same quantifier level that the new subset was. The acceptability of the fine extender sequences we shall construct will come from soundness in this way.

Perhaps the first substantial theorem in the fine structural analysis of L is Jensen's result that if $E_\beta = \emptyset$ for all $\beta \leq \alpha$, then $\mathcal{J}_\alpha^{\vec{E}}$ is ω-sound ([11]). If μ is a normal ultrafilter on κ, then $(J_{\kappa+1}^{\mu}, \in, \mu)$ is not 1-sound (in the naturally adapted meaning of the term). It is because we have followed the Baldwin-Mitchell approach in putting partial extenders on $\vec{E}$ that we have the very useful L-like fact that all levels of $L[\vec{E}]$ are ω-sound.

2.4. Fine structure and ultrapowers

If $\mathcal{M}$ is a premouse and E is an extender over $\mathcal{C}_0(\mathcal{M})$, then we can form $\mathrm{Ult}(\mathcal{C}_0(\mathcal{M}), E)$. One can show without too much difficulty that this structure is the Σ_0 code of a premouse. The key here is that the canonical embedding i into the ultrapower is not just Σ_1 elementary, but *cofinal*, in that both $i``(\mathrm{On} \cap \mathcal{C}_0(\mathcal{M}))$ is cofinal in $\mathrm{On} \cap \mathrm{Ult}(\mathcal{C}_0(\mathcal{M}), E)$, and $i``(\dot{\mu}^+)^{\mathcal{C}_0(\mathcal{M})}$ is cofinal in $i((\dot{\mu}^+)^{\mathcal{C}_0(\mathcal{M})})$. The second condition is of course only interesting if $\mathcal{M}$ is active.[12] If $\mathrm{crit}(E) < \rho_n(\mathcal{M})$, where $1 \leq n \leq \omega$, one can form a stronger ultrapower of $\mathcal{M}$, one for which Łoś's theorem holds for $r\Sigma_n$ formulae. Roughly speaking, instead of using only functions $f \in \mathcal{C}_0(\mathcal{M})$, one uses all functions f which are $r\Sigma_n$ definable from parameters over $\mathcal{C}_0(\mathcal{M})$. (See [26, section 4] and [38] for details, and generally for the $r\Sigma_n$ hierarchy.) Since $\mathrm{crit}(E) < \rho_n(\mathcal{M})$, E measures enough sets that the construction makes sense, and Łoś's theorem holds for $r\Sigma_n$ formulae. We call this stronger ultrapower $\mathrm{Ult}_n(\mathcal{C}_0(\mathcal{M}), E)$, and sometimes call the earlier ultrapower $\mathrm{Ult}_0(\mathcal{C}_0(\mathcal{M}), E)$.

We shall only form $\mathrm{Ult}_n(\mathcal{C}_0(\mathcal{M}), E)$ in the case that $\mathcal{M}$ is n-sound. In this case, all of $\mathcal{C}_0(\mathcal{M})$ can be coded by the $r\Sigma_n$ theory of $\rho_n(\mathcal{M}) \cup \{p_n(\mathcal{M})\}$, which we can regard as a subset A_n of $\rho_n(\mathcal{M})$. The structure $(J^{\mathcal{M}}_{\rho_n}, A_n)$ is amenable. If one decodes $\mathrm{Ult}_0((J^{\mathcal{M}}_{\rho_n(\mathcal{M})}, A_n), E)$ in the natural way, one gets $\mathrm{Ult}_n(\mathcal{C}_0(\mathcal{M}), E)$. This is how Σ_n ultrapowers were treated by Dodd and Jensen ([6]), and the reader can find an exposition of their method in [38, §8]. The equivalence of the two approaches in the case that $\mathcal{M}$ is n-sound is proved in [26, §2].

We wish to record some basic facts concerning the elementarity of the canonical embedding associated to a Σ_n ultrapower. As a notational convenience, for any ppm $\mathcal{M}$ we let $\rho_0(\mathcal{M}) = \mathrm{On} \cap \mathcal{C}_0(\mathcal{M})$ and $p_0(\mathcal{M}) = \emptyset$, and we say $\mathcal{M}$ is 0-sound. Again, the concept of being $r\Sigma_n$ is treated in [26] and [38].

2.20 Definition. Let $\pi \colon \mathcal{C}_0(\mathcal{M}) \to \mathcal{C}_0(\mathcal{N})$, and let $n < \omega$. We call π an *n-embedding* iff

1. $\mathcal{M}$ and $\mathcal{N}$ are n-sound,
2. π is $r\Sigma_{n+1}$-elementary,
3. $\pi(p_i(\mathcal{M})) = p_i(\mathcal{N})$ for all $i \leq n$, and
4. $\pi(\rho_i(\mathcal{M})) = \rho_i(\mathcal{N})$ for all $i < n$, and $\sup \pi``\rho_n(\mathcal{M}) = \rho_n(\mathcal{N})$.

We call π an *ω-embedding* iff π is fully elementary. Such an embedding preserves all projecta and standard parameters.

2.21 Lemma. *For any $n \leq \omega$, the canonical embedding associated to a Σ_n-ultrapower is an n-embedding.*

We must also consider the behavior of $\rho_{n+1}(\mathcal{M})$ and $p_{n+1}(\mathcal{M})$ in Σ_n ultrapowers. Here we must impose an additional condition on the extender used to form the ultrapower.

[12] This is why we defined $\mathcal{C}_0(\mathcal{M})$ as we did in the case $\mathcal{M}$ is of type III. Had we defined it as in the type II case, the fact that i might not be continuous at $\nu^{\mathcal{M}}$ might lead to a failure of the initial segment condition for $\mathrm{Ult}(\mathcal{C}_0(\mathcal{M}), E)$. Having said this, we ask the reader to once again forget the type III case, and go back to identifying $\mathcal{C}_0(\mathcal{M})$ with $\mathcal{M}$.

2.22 Definition. Let E be a (κ, λ) extender over $\mathcal{C}_0(\mathcal{M})$; then we say E *is close to* $\mathcal{C}_0(\mathcal{M})$ (or to $\mathcal{M}$ itself) iff for every $a \in [\lambda]^{<\omega}$

1. E_a is Σ_1 definable over $\mathcal{C}_0(\mathcal{M})$ from parameters, and
2. if $\mathcal{A} \in \mathcal{C}_0(\mathcal{M})$ and $\mathcal{C}_0(\mathcal{M}) \models |\mathcal{A}| \leq \kappa$, then $E_a \cap \mathcal{A} \in \mathcal{C}_0(\mathcal{M})$.

2.23 Lemma. *Let $\mathcal{M}$ be a premouse, and E a (κ, λ) extender over $\mathcal{C}_0(\mathcal{M})$ which is close to $\mathcal{C}_0(\mathcal{M})$, with $\kappa < \rho_n(\mathcal{M})$ where $n \leq \omega$. Let $\mathcal{N}$ be such that $\mathcal{C}_0(\mathcal{N}) = \mathrm{Ult}_n(\mathcal{C}_0(\mathcal{M}), E)$. Then*

$$P(\kappa) \cap \mathcal{M} = P(\kappa) \cap \mathcal{N}.$$

If in addition $n < \omega$, $\mathcal{M}$ is n-sound and $n+1$-solid, and $\rho_{n+1}(\mathcal{M}) \leq \kappa$, then the canonical embedding $\pi\colon \mathcal{C}_0(\mathcal{M}) \to \mathcal{C}_0(\mathcal{N})$ satisfies

$$\rho_{n+1}(\mathcal{M}) = \rho_{n+1}(\mathcal{N}) \text{ and } \pi(p_{n+1}(\mathcal{M})) = p_{n+1}(\mathcal{N}),$$

so that

$$\mathcal{C}_{n+1}(\mathcal{M}) = \mathcal{C}_{n+1}(\mathcal{N}),$$

and $\pi \restriction \mathcal{C}_{n+1}(\mathcal{M})$ is (an isomorphic copy of) the collapse map from $\mathcal{C}_{n+1}(\mathcal{N})$ to $\mathcal{C}_n(\mathcal{N})$. In particular, $\mathcal{N}$ is n-sound but not $(n+1)$-sound.

We omit the proof of 2.23, which the reader can find in [26, 4.5,4.6]. See also [38, 8.10]. It is a reasonable exercise to prove the lemma in the case $n = 0$. Here the only tricky part is showing that $\pi(p_1(\mathcal{M})) = p_1(\mathcal{N})$. At that point one uses heavily the solidity of $p_1(\mathcal{M})$. The prewellordering property for $\Sigma_1^{\mathcal{C}_0(\mathcal{M})}$ relations is also used.[13]

Let $\mathcal{M}$ be a premouse, and E an extender over $\mathcal{C}_0(\mathcal{M})$ with $\mathrm{crit}(E) < \rho_n(\mathcal{M})$; then by $\mathrm{Ult}_n(\mathcal{M}, E)$ we shall mean the unique premouse $\mathcal{N}$ such that $\mathcal{C}_0(\mathcal{N}) = \mathrm{Ult}_n(\mathcal{C}_0(\mathcal{M}), E)$.[14]

3. Iteration Trees and Comparison

The key to Kunen's theory of $L[U]$ is the method of *iterated ultrapowers*. Given a structure $\mathcal{M}_0 = \langle L_\zeta[U], \in, U\rangle$ with appropriate ultrafilter U, one can form ultrapowers by U and its images under the canonical embeddings repeatedly, taking direct limits at limit ordinals. One obtains thereby

[13]Let $p_1(\mathcal{M}) = \langle \alpha_0, \ldots, \alpha_k\rangle$, and let T be a universal $\Sigma_1^{\mathcal{M}}(\{\alpha_0, .., \alpha_{i-1}\})$ subset of α_i. Let $\leq$ be the prewellorder of T given by the stages at which Σ_1 formulae are verified. Then the universal $\Sigma_1^{\mathcal{N}}(\{\pi(\alpha_0), \ldots, \pi(\alpha_{i-1})\})$ subset of $\pi(\alpha_i)$ is an initial segment of $\pi(T)$ under $\pi(\leq)$, and is therefore in $\mathcal{N}$. Thus $\pi(p_1(\mathcal{M}))$ is solid, and from this we easily see that $\pi(p_1(\mathcal{M})) = p_1(\mathcal{N})$.

[14]This gives us two definitions of $\mathrm{Ult}_0(\mathcal{M}, E)$, but they clearly agree with one another except possibly when $\mathcal{M}$ is active type III. In that case, we are now discarding the earlier definition.

structures $\mathcal{M}_\alpha$ and embeddings $i_{\alpha,\beta}\colon \mathcal{M}_\alpha \to \mathcal{M}_\beta$ for $\alpha < \beta$. We call the structures $\mathcal{M}_\alpha$ *iterates* of $\mathcal{M}_0$, and say that $\mathcal{M}_0$ is *iterable* just in case all its iterates are wellfounded. Kunen's key *comparison lemma* states that if $\mathcal{M}_0$ and $\mathcal{N}_0$ are two iterable structures of this form, then there are iterates $\mathcal{M}_\alpha$ and $\mathcal{N}_\alpha$ such that one of the two is an initial segment of the other.[15]

One can form iterated ultrapowers of an arbitrary premouse $\mathcal{M}_0$ similarly. In this case, the $\mathcal{M}_\alpha$–sequence may have more than one extender, and we are allowed to choose any one of them to continue. If E_α is the extender chosen, then we take $\mathcal{M}_{\alpha+1}$ to be $\mathrm{Ult}(\mathcal{M}_\alpha, E_\alpha)$.[16] At limit stages we form direct limits and continue. We call any such sequence $\langle(\mathcal{M}_\alpha, E_\alpha)\colon \alpha < \beta\rangle$ a *linear iteration* of $\mathcal{M}_0$, and the structures $\mathcal{M}_\alpha$ in it *linear iterates* of $\mathcal{M}_0$. We say $\mathcal{M}_0$ is *linearly iterable* just in case all its linear iterates are wellfounded.[17]

Given linearly iterable premice $\mathcal{M}_0$ and $\mathcal{N}_0$, there is a natural way to try to compare the two via linear iteration. Having reached $\mathcal{M}_\alpha$ and $\mathcal{N}_\alpha$, and supposing neither is an initial segment of the other (as otherwise our work is finished), we pick extenders E and F representing the least disagreement between $\mathcal{M}_\alpha$ and $\mathcal{N}_\alpha$, and use these to form $\mathcal{M}_{\alpha+1}$ and $\mathcal{N}_{\alpha+1}$.

If the extenders of the coherent sequence of $\mathcal{M}_0$ do not overlap one another too much, and similarly for $\mathcal{N}_0$, then this process must terminate with all disagreements between some $\mathcal{M}_\alpha$ and $\mathcal{N}_\alpha$ eliminated, so that one is an initial segment of the other. This is the key to core model theory at the level of strong cardinals. At bottom, the reason this comparison process must terminate is the following: if E and F are the extenders used at a typical stage α, then there will be a finite set a of generators and sets $\overline{X}$ and $\tilde{X}$ such that $X = i_{\eta,\alpha}(\overline{X}) = j_{\xi,\alpha}(\tilde{X})$, and X is measured differently by E_a and F_a.[18] But then $a \in i_{\alpha,\alpha+1}(X) \Leftrightarrow a \notin j_{\alpha,\alpha+1}(X)$, so $i_{\eta,\alpha+1}(\overline{X}) \neq j_{\xi,\alpha+1}(\tilde{X})$, and the images of $\overline{X}$ and $\tilde{X}$ do not participate in a disagreement at stage $\alpha+1$ the way they did at stage α. *If* all future extenders used in either iteration have critical point above $\sup(a)$, then $i_{\eta,\beta}(\overline{X}) \neq j_{\xi,\beta}(\tilde{X})$ for all β,

[15] This means that there is a filter F such that $\mathcal{M}_\alpha$ and $\mathcal{N}_\alpha$ are of the form

$$\langle L_\xi[F], \in, F\rangle$$

and

$$\langle L_\eta[F], \in, F\rangle$$

for some ξ and η. (Here and elsewhere we identify wellfounded, extensional structures with their transitive isomorphs.) In fact, in this simple case we can take α to be $\sup(|\mathcal{M}_0|, |\mathcal{N}_0|)^+$ and F to be the club filter on α.

[16] This must be qualified, since if E_α does not measure all subsets of its critical point in $\mathcal{M}_\alpha$, then $\mathrm{Ult}(\mathcal{M}_\alpha, E_\alpha)$ makes no sense. In this case we take the "largest" E_α–ultrapower of an initial segment of $\mathcal{M}_\alpha$ we can in order to form $\mathcal{M}_{\alpha+1}$. See below.

[17] In which case we identify these iterates with the premice to which they are isomorphic.

Linear iterability should be taken to include the condition that no linear iteration of $\mathcal{M}_0$ drops to proper initial segments infinitely often.

[18] We use i for the embeddings in the $\mathcal{M}$–iteration, and j in the $\mathcal{N}$–iteration.

so the images of $\overline{X}$ and $\tilde{X}$ never again participate in a disagreement, and we have made real progress at stage α. A simple reflection argument shows that if we never "move generators" in one of our iterations,[19] then eventually all disagreements are removed.[20] The lack of overlaps in the sequences of mice below a strong cardinal means that this process of iterating away the least disagreement does not move generators, and hence terminates in a successful comparison.

However, beyond a strong cardinal this linear comparison process definitely will lead to moving generators. There are tricks for making do with linear iterations a bit beyond strong cardinals, but the right solution is to give up linearity. If the extender E_α from the $\mathcal{M}_\alpha$–sequence we want to use has critical point less than $\nu(E_\beta)$ for some $\beta < \alpha$, then we apply E_α not to $\mathcal{M}_\alpha$, but to $\mathcal{M}_\beta$, for the least such β: *i.e.*, we set $\mathcal{M}_{\alpha+1} = \mathrm{Ult}(\mathcal{M}_\beta, E_\alpha)$, where β is least such that $\mathrm{crit}(E_\alpha) < \nu(E_\beta)$.[21] We have an embedding $i_{\beta,\alpha+1} \colon \mathcal{M}_\beta \to \mathcal{M}_{\alpha+1}$. Thus this new iteration process gives rise to a *tree* of models, with embeddings along each branch of the tree. Along each branch the generators of the extenders used are not moved by later embeddings, and this is good enough to show that if a comparison process involving the formation of such "iteration trees" goes on long enough, it must eventually succeed.

What one needs to keep the construction of an iteration tree going past some limit ordinal λ is a branch of the tree which has been visited cofinally often before λ and is such that the direct limit of the premice along the branch is wellfounded. Thus the iterability we need for comparison amounts to the existence of some method for choosing such branches. We can formalize this as the existence of a winning strategy in a certain game. In giving the details of the necessary definitions, it is more convenient to introduce this "iteration game" first. We turn to this now.

3.1. Iteration trees

Let $\mathcal{M}$ be an k-sound premouse, and let θ be an ordinal; we shall define the *iteration game* $\mathcal{G}_k(\mathcal{M}, \theta)$.

3.1 Definition. A *tree order on* α (for α an ordinal) is a strict partial order T of α with least element 0 such that for all $\gamma < \alpha$

1. $\beta T \gamma \Rightarrow \beta < \gamma$,
2. $\{\beta \mid \beta T \gamma\}$ is wellordered by T,

[19] That is, if $\nu(E) \leq \mathrm{crit}(E')$ whenever E is used before E' in the $\mathcal{M}$ iteration, and similarly on the $\mathcal{N}$–side.

[20] More precisely, there must be a stage $\alpha < \sup(|\mathcal{M}_0|, |\mathcal{N}_0|)^+$ at which $\mathcal{M}_\alpha$ is an initial segment of $\mathcal{N}_\alpha$, or *vice versa*.

[21] Again, if E_α fails to measure all sets in $\mathcal{M}_\beta$, we take the ultrapower of the longest possible initial segment of $\mathcal{M}_\beta$.

3. γ is a successor ordinal $\Leftrightarrow$ γ is a T-successor, and

4. γ is a limit ordinal $\Rightarrow \{\beta \mid \beta T \gamma\}$ is $\in$-cofinal in γ.

3.2 Definition. It T is a tree order then

$$[\beta, \gamma]_T = \{\eta \mid \eta = \beta \vee \beta T \eta T \gamma \vee \eta = \gamma\},$$

and similarly for $(\beta, \gamma]_T$, $[\beta, \gamma)_T$, and $(\beta, \gamma)_T$. Also, if γ is a successor ordinal, we let $\mathrm{pred}_T(\gamma)$ be the unique ordinal $\eta T \gamma$ such that $(\eta, \gamma)_T = \emptyset$.

3.3 Definition. Premice $\mathcal{M}$ and $\mathcal{N}$ *agree below* γ iff $\mathcal{J}_\beta^{\mathcal{M}} = \mathcal{J}_\beta^{\mathcal{N}}$ for all $\beta < \gamma$

We now describe a typical run of $\mathcal{G}_k(\mathcal{M}, \theta)$. As play proceeds the players determine

- a tree order T on θ,
- premice $\mathcal{M}_\alpha$ for $\alpha < \theta$, with $\mathcal{M}_0 = \mathcal{M}$,
- an extender F_α from the $\mathcal{M}_\alpha$ sequence, for $\alpha < \theta$, and
- a set $D \subseteq \theta$, and embeddings $i_{\alpha,\beta} \colon \mathcal{C}_0(\mathcal{M}_\alpha) \to \mathcal{C}_0(\mathcal{M}_\beta)$ defined whenever $\alpha T \beta$ and $D \cap (\alpha, \beta]_T = \emptyset$.

The rules of the game guarantee the following agreement among the premice produced:

- $\alpha \leq \beta \Longrightarrow \mathcal{M}_\alpha$ agrees with $\mathcal{M}_\beta$ below $\mathrm{lh}(F_\alpha)$,
- $\alpha < \beta \Longrightarrow \mathrm{lh}(F_\alpha)$ is a cardinal of $\mathcal{M}_\beta$.

Notice that the last condition implies that if $\alpha < \beta$, then $\mathcal{M}_\alpha$ does not agree with $\mathcal{M}_\beta$ below $\mathrm{lh}(F_\alpha) + 1$. This is because from F_α one can easily compute a map from $\nu(F_\alpha)$ onto $\mathrm{lh}(F_\alpha)$.

The game is played as follows. Suppose first we are at move $\alpha + 1$, and have already defined F_ξ for $\xi < \alpha$, $\mathcal{M}_\xi$ for $\xi \leq \alpha$, and T and D on $\alpha + 1$. (The first move is move 1, and in this case all we need is $\mathcal{M} = \mathcal{M}_0$ to get going.) At move $\alpha + 1$, I must pick an extender F_α from the $\mathcal{M}_\alpha$ sequence such that $\mathrm{lh}(F_\xi) < \mathrm{lh}(F_\alpha)$ for all $\xi < \alpha$. (If he does not, the game is over and he loses.) Now let $\beta \leq \alpha$ be least such that $\mathrm{crit}(F_\alpha) < \nu(F_\beta)$. Let

$$\mathcal{M}^*_{\alpha+1} := \mathcal{J}_\gamma^{\mathcal{M}_\beta}, \text{ where } \gamma \text{ is the largest } \eta \text{ such that } F_\alpha \text{ is a pre-extender over } \mathcal{J}_\eta^{\mathcal{M}_\beta}.$$

Our agreement hypotheses imply that γ exists, $\text{lh}(F_\beta) \le \gamma$, and F_α is a pre-extender over $\mathcal{C}_0(\mathcal{J}_\gamma^{\mathcal{M}_\beta})$. [Proof: this is clear if $\beta = \alpha$, so let $\beta < \alpha$. Let $\kappa = \text{crit}(F_\alpha)$. Since $\text{lh}(F_\beta) < \text{lh}(F_\alpha)$ and $\text{lh}(F_\beta)$ is a cardinal of $\mathcal{M}_\alpha$,

$$P(\kappa) \cap \mathcal{J}^{\mathcal{M}_\beta}_{\text{lh}(F_\beta)} = P(\kappa) \cap \mathcal{M}_\alpha = P(\kappa) \cap \mathcal{J}^{\mathcal{M}_\alpha}_{\text{lh}(F_\alpha)}.$$

Thus F_α is a pre-extender over $\mathcal{J}^{\mathcal{M}_\beta}_{\text{lh}(F_\beta)}$, so γ exists and $\text{lh}(F_\beta) \le \gamma$. The last statement needs proof only in the case $\mathcal{J}_\gamma^{\mathcal{M}_\beta}$ is of type III. In this case, $\nu := \nu(\mathcal{J}_\gamma^{\mathcal{M}_\beta})$ is the largest cardinal of $\mathcal{J}_\gamma^{\mathcal{M}_\beta}$. Thus if $\text{lh}(F_\beta) < \gamma$, then $\text{lh}(F_\beta) \le \nu$, so that $\kappa < \nu$, as desired. If $\text{lh}(F_\beta) = \gamma$, then $\nu = \nu(F_\beta)$, so once again $\kappa < \nu$, as desired.] We put

$$\alpha + 1 \in D \Leftrightarrow \mathcal{M}^*_{\alpha+1} \text{ is a proper initial segment of } \mathcal{M}_\beta.$$

Let $n \le \omega$ be largest such that: (i) $\text{crit}(F_\alpha) < \rho_n(\mathcal{M}^*_{\alpha+1})$ and (ii) if $D \cap [0, \alpha+1]_T = \emptyset$, then $n \le k$. Set

$$\mathcal{M}_{\alpha+1} := \text{Ult}_n(\mathcal{M}^*_{\alpha+1}, F_\alpha),$$

if this ultrapower is wellfounded. (If the ultrapower is not wellfounded, then the game is over and II has lost.) Finally, we let $\beta T(\alpha+1)$, and if $\alpha + 1 \notin D$, then $i_{\beta,\alpha+1} \colon \mathcal{C}_0(\mathcal{M}_\beta) \to \mathcal{C}_0(\mathcal{M}_{\alpha+1})$ is the canonical ultrapower embedding, and $i_{\gamma,\alpha+1} = i_{\beta,\alpha+1} \circ i_{\gamma,\beta}$ whenever $\gamma T \beta$ and $D \cap (\gamma, \beta]_T = \emptyset$. If $\alpha + 1 \in D$, then we leave $i_{\beta,\alpha+1}$ undefined.

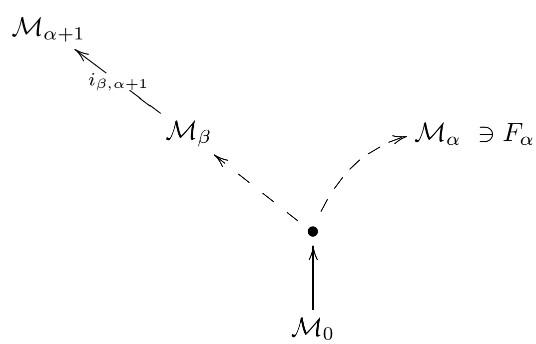

We must verify the agreement hypothesis we have carried along. For this, it suffices by induction to show that $\mathcal{M}_\alpha$ and $\mathcal{M}_{\alpha+1}$ have the necessary agreement. Let $\kappa = \text{crit}(F_\alpha)$, and let $i \colon \mathcal{M}^*_{\alpha+1} \to \mathcal{M}_{\alpha+1}$, $j \colon \mathcal{M}^*_{\alpha+1} \to \text{Ult}_0(\mathcal{M}^*_{\alpha+1}, F_\alpha) := \mathcal{P}$, and $h \colon \mathcal{J}^{\mathcal{M}_\alpha}_{\text{lh}(F_\alpha)} \to \text{Ult}_0(\mathcal{J}^{\mathcal{M}_\alpha}_{\text{lh}(F_\alpha)}, F_\alpha) := \mathcal{Q}$ be the canonical embeddings. We have just shown, in effect, that $\mathcal{M}^*_{\alpha+1}$ and $\mathcal{J}^{\mathcal{M}_\alpha}_{\text{lh}(F_\alpha)}$ agree below their common value λ for κ^+. It follows at once that $\mathcal{P}$ and $\mathcal{Q}$ agree below $j(\lambda) = h(\lambda)$. But $\mathcal{P}$ agrees below $i(\lambda) = j(\lambda)$ with $\mathcal{M}_{\alpha+1}$

because $\kappa < \rho_n(\mathcal{M}^*_{\alpha+1})$ (so that the $r\Sigma_n^{\mathcal{M}^*_{\alpha+1}}$ functions from κ to itself are all in $\mathcal{M}^*_{\alpha+1}$). Finally, $\mathcal{Q}$ agrees with $\mathcal{M}_\alpha$ below $\mathrm{lh}(F_\alpha)$, which is a cardinal of $\mathcal{Q}$, from the definition of fine extender sequences. Since $\mathrm{lh}(F_\alpha) < h(\lambda)$ we have the required agreement.

At a limit move λ, II picks a branch b of the tree T on λ determined by the play thus far. The branch b must be *cofinal* (i.e. $\in$-cofinal in λ), and *wellfounded*; otherwise II loses. (We say b is wellfounded iff $D \cap b$ is bounded below λ, and the direct limit of the $\mathcal{C}_0(\mathcal{M}_\beta)$ for $\beta \in (b \setminus \sup(D \cap \beta))$ under the embeddings $i_{\alpha,\beta}$ along b is wellfounded.) If II picks such a b, we set

$$\mathcal{M}_\lambda := \ \mathrm{dirlim}_{\alpha \in b} \mathcal{M}_\alpha,$$

where we understand the direct limit here to be the premouse whose Σ_0 code is the direct limit of the $\mathcal{C}_0(\mathcal{M}_\alpha)$, for $\alpha \in b$ sufficiently large. We put $\alpha T \lambda$ for all $\alpha \in b$, and let $i_{\alpha,\lambda}$ be the canonical embedding into the direct limit for $\alpha \in b \setminus \sup(D \cap b)$.

This completes the rules of play for $\mathcal{G}_k(\mathcal{M}, \theta)$. If no one has lost after θ moves, then II wins.

3.4 Definition. A *k-maximal iteration tree on* $\mathcal{M}$ is a partial play of $\mathcal{G}_k(\mathcal{M}, \theta)$ in which neither player has yet lost.

We shall use calligraphic letters (e.g. $\mathcal{T}$) for iteration trees, and the corresponding roman letters (e.g. T) for their associated tree orders. ($\mathcal{T}$ is an *iteration tree* if it is a k-maximal iteration tree for some $k \leq \omega$.) We use $\mathcal{M}^\mathcal{T}_\alpha$ for the α^{th} premouse of $\mathcal{T}$, $E^\mathcal{T}_\alpha$ for the α^{th} extender used in $\mathcal{T}$, and $i^\mathcal{T}_{\alpha,\beta}$ for the canonical embeddings. (So $E^\mathcal{T}_\alpha$ is on the sequence of $\mathcal{M}^\mathcal{T}_\alpha$.) We use $D^\mathcal{T}$ for the set of all $\alpha+1$ such that $\mathcal{M}^{*\mathcal{T}}_{\alpha+1} \neq \mathcal{M}^\mathcal{T}_{\mathrm{pred}_T(\alpha+1)}$. In order to avoid a forest of superscripts, we shall often say "$\mathcal{T}$ is an iteration tree with models $\mathcal{N}_\alpha$, extenders F_α, and emdeddings $j_{\alpha,\beta}$" when $\mathcal{N}_\alpha = \mathcal{M}^\mathcal{T}_\alpha$, $F_\alpha = E^\mathcal{T}_\alpha$, and $j_{\alpha,\beta} = i^\mathcal{T}_{\alpha,\beta}$. We will then write $\mathcal{N}^*_{\alpha+1}$ for $\mathcal{M}^{*\mathcal{T}}_{\alpha+1}$, and so forth. In general, we drop superscripts keeping track of an iteration tree whenever it seems like a good idea.

The *length* $\mathrm{lh}(\mathcal{T})$ of an iteration tree $\mathcal{T}$ is the domain of the associated tree order, so that $\mathrm{lh}(\mathcal{T}) = \alpha + 1$ iff $\mathcal{T}$ has last model $\mathcal{M}^\mathcal{T}_\alpha$.

In the course of describing $\mathcal{G}_k(\mathcal{M}, \theta)$ we proved the following lemma.

3.5 Lemma. *Let $\mathcal{T}$ be an iteration tree with models $\mathcal{M}_\alpha$ and extenders E_α, and let $\alpha < \beta < \mathrm{lh}(\mathcal{T})$; then*

1. *$\mathcal{M}_\alpha$ and $\mathcal{M}_\beta$ agree below* $\mathrm{lh}(E_\alpha)$*, and*

2. $\mathrm{lh}(E_\alpha)$ *is a cardinal of $\mathcal{M}_\beta$, so that $\mathcal{M}_\alpha$ and $\mathcal{M}_\beta$ do not agree below* $\mathrm{lh}(E_\alpha) + 1$.

Here is another elementary fact:

3.6 Lemma. *Let $\mathcal{T}$ be an iteration tree, and let $\alpha + 1 < \mathrm{lh}(\mathcal{T})$; then E_α is close to $\mathcal{M}^*_{\alpha+1}$.*

The proof is a straightforward induction (see [26, 6.1.5]). This lemma puts the elementarity lemma 2.23 at our disposal, and we can then describe the elementarity of the embeddings along the branches of an iteration tree as follows.

3.7 Definition. If $\mathcal{T}$ is an iteration tree with models $\mathcal{M}_\alpha$ and extenders E_α, and $\alpha + 1 < \mathrm{lh}(\mathcal{T})$, then $\deg^{\mathcal{T}}(\alpha+1)$ is the largest $n \leq \omega$ such that $\mathcal{M}_{\alpha+1} = \mathrm{Ult}_n(\mathcal{M}^*_{\alpha+1}, E_\alpha)$. Also, we use $i^{*\mathcal{T}}_{\alpha+1}$ for the canonical embedding from $\mathcal{M}^*_{\alpha+1}$ into this ultrapower.

3.8 Theorem. *Let $\mathcal{T}$ be a k-maximal iteration tree on a k-sound premouse, with models $\mathcal{M}_\alpha$ and embeddings $i_{\alpha,\beta}$, and let $(\alpha+1)T\beta$ and $D^{\mathcal{T}} \cap (\alpha + 1, \beta]_T = \emptyset$; then*

1. $\deg^{\mathcal{T}}(\alpha+1) \geq \deg^{\mathcal{T}}(\xi+1)$ *for all* $\xi + 1 \in (\alpha+1, \beta]_T$, *and*
2. *if* $\deg^{\mathcal{T}}(\alpha+1) = \deg^{\mathcal{T}}(\xi+1) = n$ *for all* $\xi + 1 \in (\alpha+1, \beta]_T$, *then*

$$i_{\alpha+1,\beta} \circ i^*_{\alpha+1} \text{ is an } n\text{-embedding;}$$

moreover if $D^{\mathcal{T}} \cap [0, \alpha] \neq \emptyset$ or $n < k$, then

$$\rho_{n+1}(\mathcal{M}^*_{\alpha+1}) = \rho_{n+1}(\mathcal{M}_\beta) < \mathrm{crit}(i_{\alpha+1,\beta} \circ i^*_{\alpha+1}),$$

$$i_{\alpha+1,\beta} \circ i^*_{\alpha+1}(p_{n+1}(\mathcal{M}^*_{\alpha+1})) = p_{n+1}(\mathcal{M}_\beta),$$

and

$$\mathcal{C}_{n+1}(\mathcal{M}^*_{\alpha+1}) = \mathcal{C}_{n+1}(\mathcal{M}_\beta).$$

We omit the proof (see [26, 4.7]), which proceeds by induction on β, using the proof (not just the statement) of 2.23. Because of 3.8, we can for limit λ set $\deg^{\mathcal{T}}(\lambda) =$ eventual value of $\deg^{\mathcal{T}}(\alpha+1)$, for $(\alpha+1)T\lambda$ sufficiently large. When we are considering $\mathcal{T}$ as a play in $\mathcal{G}_k(\mathcal{M}, \theta)$, we set also $\deg^{\mathcal{T}}(0) = k$.[22] We then have that for any $\alpha < \mathrm{lh}(\mathcal{T})$, $\deg^{\mathcal{T}}(\alpha)$ is the largest $n \leq \omega$ such that $\mathcal{M}_\alpha$ is n-sound and $n \leq \deg^{\mathcal{T}}(0)$ if $D \cap [0, \alpha+1]_T = \emptyset$. If $\mathcal{M}^*_{\alpha+1}$ is $n+1$ sound, where $n+1 \leq \deg^{\mathcal{T}}(0)$ if $D \cap [0, \alpha+1]_T = \emptyset$, and $D \cap (\alpha+1, \beta]_T = \emptyset$ and $\deg^{\mathcal{T}}(\alpha+1) = \deg^{\mathcal{T}}(\beta) = n$, then by 3.8 the branch embedding $i_{\alpha+1,\beta} \circ i^*_{\alpha+1}$ is just the uncollapse map from $\mathcal{C}_{n+1}(\mathcal{M}_\beta)$ to $\mathcal{C}_n(\mathcal{M}_\beta)$.

3.9 Definition. A *(k, θ)-iteration strategy for $\mathcal{M}$* is a winning strategy for II in $\mathcal{G}_k(\mathcal{M}, \theta)$. We say *$\mathcal{M}$ is (k, θ)-iterable* iff there is such a strategy.

[22] It is an awkward feature of our terminology that an iteration tree may be a play of $\mathcal{G}_k(\mathcal{M}, \theta)$ for more than one k.

The iteration trees we have introduced have some special properties. If one drops the restriction on I in $\mathcal{G}_k(\mathcal{M}, \theta)$ that he pick extenders of increasing lengths, and allow him to apply the extender chosen to any initial segment of any earlier model over which it is an extender, one obtains a stronger notion of iterability which is perhaps more natural. We shall need an approximation to this stronger notion later.

It is customary to call an iterable premouse a *mouse*, and we shall follow this custom in informal discussion. We shall make no formal definition of "mouse", however, as it is not clear what sort of iterability one should demand. The definition above captures only one variety of iterability. The question of iterability and its applications is of central importance and, at the same time, not very well understood. For this reason, we prefer to spell out in each instance how much iterability we can prove, or how much we need for a given purpose.

3.2. The comparison process

The most important use of iterability lies in the *comparison process* for mice. There are certainly mice $\mathcal{M}$ and $\mathcal{N}$ such that neither is an initial segment of the other, but if $\mathcal{M}$ and $\mathcal{N}$ are sufficiently iterable, then one can form iteration trees on $\mathcal{M}$ and $\mathcal{N}$ with last models $\mathcal{P}$ and $\mathcal{Q}$ respectively such that $\mathcal{P}$ is an initial segment of $\mathcal{Q}$ or vice-versa. Moreover, one can arrange that if, say, $\mathcal{P}$ is an initial segment of $\mathcal{Q}$, then the branch of the tree on $\mathcal{M}$ leading to $\mathcal{P}$ does not drop, and thus gives rise to an elementary embedding from $\mathcal{M}$ to $\mathcal{P}$. Intuitively, this means that $\mathcal{M}$ has been compared with $\mathcal{N}$, and found to be no stronger.

3.10 Definition. A branch b of the iteration tree $\mathcal{T}$ *drops (in model or degree)* iff $D^{\mathcal{T}} \cap b \neq \emptyset$ or $\deg^{\mathcal{T}}(b) < deg^{\mathcal{T}}(0)$.

If b does not drop in model, then $i_{0,b}$ exists, and if in addition b does not drop in degree, then $i_{0,b}$ is a $\deg^{\mathcal{T}}(0)$-embedding. We shall also speak of "partial branches" of the form $[0, \alpha]_T$ dropping (in model or degree), with the obvious meaning. Again, if there is no such dropping, then $i_{0,\alpha}$ exists and is a $\deg^{\mathcal{T}}(0)$-embedding.

3.11 Theorem (The Comparison Lemma). *Let $\mathcal{M}$ and $\mathcal{N}$ be k-sound premice of size $\leq \theta$, and suppose Σ and Γ are $(k, \theta^+ + 1)$ iteration strategies for $\mathcal{M}$ and $\mathcal{N}$ respectively; then there are iteration trees $\mathcal{T}$ and $\mathcal{U}$ played according to Σ and Γ respectively, and having last models $\mathcal{M}_\alpha^{\mathcal{T}}$ and $\mathcal{M}_\eta^{\mathcal{U}}$, such that either*

- $[0, \alpha]_T$ *does not drop in model or degree, and $\mathcal{M}_\alpha^{\mathcal{T}}$ is an initial segment of $\mathcal{M}_\eta^{\mathcal{U}}$, or*

- $[0, \eta]_U$ *does not drop in model or degree, and* $\mathcal{M}^{\mathcal{U}}_\eta$ *is an initial segment of* $\mathcal{M}^{\mathcal{T}}_\alpha$.

Proof. We build $\mathcal{T}$ and $\mathcal{U}$ by an inductive process known as "iterating away the least disagreement". Before step $\alpha+1$ of the construction we have initial segments $\mathcal{T}_\alpha$ and $\mathcal{U}_\alpha$ of the trees we shall eventually construct, and these have last models $\mathcal{P}$ and $\mathcal{Q}$ respectively. ($\mathcal{T}_0$ and $\mathcal{U}_0$ are one-model trees with last models $\mathcal{P} = \mathcal{M}$ and $\mathcal{Q} = \mathcal{N}$.) If one of $\mathcal{P}$ and $\mathcal{Q}$ is an initial segment of the other, then the construction of $\mathcal{T}$ and $\mathcal{U}$ is finished. Otherwise, let

$$\lambda = \text{ least } \gamma \text{ such that } \mathcal{J}^{\mathcal{P}}_\gamma \neq \mathcal{J}^{\mathcal{Q}}_\gamma.$$

This means that the predicates $\dot{F}^{\mathcal{J}^{\mathcal{P}}_\lambda}$ and $\dot{F}^{\mathcal{J}^{\mathcal{Q}}_\lambda}$ are different. If $\dot{F}^{\mathcal{J}^{\mathcal{P}}_\lambda} \neq \emptyset$, then letting $\mathrm{lh}(\mathcal{T}_\alpha) = \beta + 1$, we set

$$E^{\mathcal{T}_{\alpha+1}}_\beta := \text{ pre-extender coded by } \dot{F}^{\mathcal{J}^{\mathcal{P}}_\lambda}$$

and let $\mathcal{T}_{\alpha+1}$ be the unique one-model extension of $\mathcal{T}_\alpha$ determined by this and the rules of $\mathcal{G}_k(\mathcal{M}, \theta^+ + 1)$. If $\dot{F}^{\mathcal{J}^{\mathcal{P}}_\lambda} = \emptyset$, then we just let $\mathcal{T}_{\alpha+1} = \mathcal{T}_\alpha$. Similarly, if $\dot{F}^{\mathcal{J}^{\mathcal{Q}}_\lambda} \neq \emptyset$, then letting $\mathrm{lh}(\mathcal{U}_\alpha) = \eta + 1$, we set

$$E^{\mathcal{U}_{\alpha+1}}_\eta := \text{ pre-extender coded by } \dot{F}^{\mathcal{J}^{\mathcal{Q}}_\lambda}$$

and let $\mathcal{U}_{\alpha+1}$ be the one model extension of $\mathcal{U}_\alpha$ thereby determined; otherwise we let $\mathcal{U}_{\alpha+1} = \mathcal{U}_\alpha$. Notice that in any case, the last models of $\mathcal{T}_{\alpha+1}$ and $\mathcal{U}_{\alpha+1}$ agree below $\lambda + 1$. This means that future extenders used in the two trees will have length $> \lambda$, so that player I is not losing one of the iteration games by failing to play extenders increasing in length.

At limit steps λ in our construction, we set $\mathcal{T}_\lambda = \bigcup_{\alpha<\lambda} \mathcal{T}_\alpha$ if this tree has a last model, that is, if $\mathcal{T}_\alpha$ is eventually constant as $\alpha \to \lambda$. Otherwise we let $\mathcal{T}_\lambda$ be the one-model extension of $\bigcup_{\alpha<\lambda} \mathcal{T}_\alpha$ determined by the cofinal, wellfounded branch of this tree chosen by Σ. We define $\mathcal{U}_\lambda$ in parallel fashion.

The main thing we need to prove is that the inductive process just described stops at some step $\alpha < \theta^+$.

Claim. There is an $\alpha < \theta^+$ such that the last model of $\mathcal{T}_\alpha$ is an initial segment of the last model of $\mathcal{U}_\alpha$, or vice-versa.

Proof. If not, then we have trees $\mathcal{T} = \mathcal{T}_{\theta^+}$ and $\mathcal{U} = \mathcal{U}_{\theta^+}$. It is easy to see that, since $\mathcal{M}$ and $\mathcal{N}$ have size $\leq \theta$, both $\mathcal{T}$ and $\mathcal{U}$ have length $\theta^+ + 1$.

Let us say that extenders E and F are *compatible* iff for some η, E is the trivial completion of $F \restriction \eta$ or F is the trivial completion of $E \restriction \eta$. (This implies that the extenders have the same critical point, and measure the same subsets of that critical point.)

Subclaim. For any $\alpha, \beta < \theta^+$, $E^{\mathcal{T}}_\alpha$ is incompatible with $E^{\mathcal{U}}_\beta$.

Proof. Let $E = E^{\mathcal{T}}_\alpha$, $F = E^{\mathcal{U}}_\beta$, and suppose E is the trivial completion of $F \restriction \eta$, for some η. Let ξ be such that E is the extender used to go from $\mathcal{T}_\xi$ to $\mathcal{T}_{\xi+1}$, and let γ be such that F is used to go from $\mathcal{U}_\gamma$ to $\mathcal{U}_{\gamma+1}$. Since $\mathrm{lh}(E) \le \mathrm{lh}(F)$, we have $\xi \le \gamma$. But if $\xi = \gamma$, then E and F are used at the same stage in our process, so $\mathrm{lh}(E) = \mathrm{lh}(F)$, so $E = F$, contrary to the fact that we were iterating away disagreements. Thus $\xi < \gamma$, and hence $\mathrm{lh}(E) < \mathrm{lh}(F)$. Now let $\mathcal{P}$ and $\mathcal{Q}$ be the last models of $\mathcal{T}_\gamma$ and $\mathcal{U}_\gamma$ respectively. By 3.5, $\mathrm{lh}(E)$ is a cardinal of $\mathcal{P}$, and since $\mathcal{P}$ agrees with $\mathcal{Q}$ below $\mathrm{lh}(F)$, this means $\mathrm{lh}(E)$ is a cardinal of $\mathcal{J}^{\mathcal{Q}}_{\mathrm{lh}(F)}$. On the other hand, the initial segment condition of 2.4 implies (in both its cases) that $E \in \mathcal{J}^{\mathcal{Q}}_{\mathrm{lh}(F)}$. Since E collapses its length in an easily computable way, this is a contradiction. ⊣

We now use a reflection argument to produce compatible extenders used on the branches $[0, \theta^+]_T$ and $[0, \theta^+]_U$, the desired contradiction. Let $X \prec V_\eta$ for some large η, with $\mathcal{T}, \mathcal{U} \in X$, $|X| = \theta$, and $X \cap \theta^+$ transitive. Let H be the transitive collapse of X, $\pi \colon H \to V_\eta$ the collapse map, and $\alpha = \mathrm{crit}(\pi) = X \cap \theta^+$. (Note $\theta < \alpha$.) Let $\bar{\mathcal{T}} = \pi^{-1}(\mathcal{T})$ and $\bar{\mathcal{U}} = \pi^{-1}(\mathcal{U})$.

Since $\mathcal{M}$ and $\mathcal{N}$ have size $\le \theta$, $\bar{\mathcal{T}}$ and $\bar{\mathcal{U}}$ are trees on $\mathcal{M}$ and $\mathcal{N}$ respectively. Similarly, $\bar{\mathcal{T}} \restriction \alpha = \mathcal{T} \restriction \alpha$ and $\bar{\mathcal{U}} \restriction \alpha = \mathcal{U} \restriction \alpha$. Also, $[0, \alpha]_{\bar{T}} = [0, \theta^+]_T \cap \alpha$ and $[0, \alpha]_{\bar{U}} = [0, \theta^+]_U \cap \alpha$. Since $[0, \alpha]_{\bar{T}}$ has limit order type, and any branch of an iteration tree must be closed below its sup (by clauses 3 and 4 of 3.1), we have $\alpha \in [0, \theta^+]_T$, and thus $[0, \alpha]_{\bar{T}} = [0, \alpha]_T$. Similarly $\alpha \in [0, \theta^+]_U$ and $[0, \alpha]_{\bar{U}} = [0, \alpha]_U$. Since the direct limit construction is absolute to H, these facts imply that $\bar{\mathcal{T}} = \mathcal{T} \restriction (\alpha + 1)$ and $\bar{\mathcal{U}} = \mathcal{U} \restriction (\alpha + 1)$.

We can find $\gamma \in [0, \alpha]_T$ such that $D^{\mathcal{T}} \cap [0, \alpha]_T \subseteq \gamma$, and using π we see that $D^{\mathcal{T}} \cap [0, \theta^+]_T \subseteq \gamma$. This means that $i^{\mathcal{T}}_{\alpha,\theta^+}$ is defined. In fact, if $x \in \mathcal{C}_0(\mathcal{M}^{\mathcal{T}}_\alpha)$, then letting

$$x = i^{\mathcal{T}}_{\gamma,\alpha}(\bar{x}) = i^{\bar{\mathcal{T}}}_{\gamma,\alpha}(\bar{x}),$$

we have

$$\pi(x) = i^{\mathcal{T}}_{\gamma,\theta^+}(\bar{x}) = i^{\mathcal{T}}_{\alpha,\theta^+}(i^{\mathcal{T}}_{\gamma,\alpha}(\bar{x})) = i^{\mathcal{T}}_{\alpha,\theta^+}(x).$$

In other words

$$i^{\mathcal{T}}_{\alpha,\theta^+} = \pi \restriction \mathcal{C}_0(\mathcal{M}^{\mathcal{T}}_\alpha).$$

Similarly, we get

$$i^{\mathcal{U}}_{\alpha,\theta^+} = \pi \restriction \mathcal{C}_0(\mathcal{M}^{\mathcal{U}}_\alpha).$$

Thus $i^{\mathcal{T}}_{\alpha,\theta^+}$ and $i^{\mathcal{U}}_{\alpha,\theta^+}$ agree wherever both are defined. Notice that they are defined on the same subsets of α, since

$$P(\alpha)^{\mathcal{M}^{\mathcal{T}}_\alpha} = P(\alpha)^{\mathcal{M}^{\mathcal{T}}_{\theta^+}} = P(\alpha)^{\mathcal{M}^{\mathcal{U}}_{\theta^+}} = P(\alpha)^{\mathcal{M}^{\mathcal{U}}_\alpha}.$$

Here the first and third identities hold because $\mathrm{crit}(i^{\mathcal{T}}_{\alpha,\theta+}) = \mathrm{crit}(i^{\mathcal{U}}_{\alpha,\theta+}) = \alpha$, and the second holds because $\mathcal{M}^{\mathcal{T}}_{\theta+}$ agrees with $\mathcal{M}^{\mathcal{U}}_{\theta+}$ below θ^+.

Now let $\xi+1 \in [0,\theta^+]_T$ be such that $\mathrm{pred}_T(\xi+1) = \alpha$, and $\gamma+1 \in [0,\theta^+]_U$ be such that $\mathrm{pred}_U(\gamma+1) = \alpha$. Let $\nu = \inf(\nu(E^{\mathcal{T}}_\xi), \nu(E^{\mathcal{U}}_\gamma))$. Then for any $a \in [\nu]^{<\omega}$ and $B \in (\mathcal{C}_0(\mathcal{M}^{\mathcal{T}}_\alpha) \cap \mathcal{C}_0(\mathcal{M}^{\mathcal{U}}_\alpha))$,

$$\begin{aligned} B \in (E^{\mathcal{T}}_\xi)_a \quad &\iff \quad a \in i^{\mathcal{T}}_{\alpha,\xi+1}(B) \\ &\iff \quad a \in i^{\mathcal{T}}_{\alpha,\theta+}(B) \\ &\iff \quad a \in i^{\mathcal{U}}_{\alpha,\theta+}(B) \\ &\iff \quad a \in i^{\mathcal{U}}_{\alpha,\gamma+1}(B) \\ &\iff \quad B \in (E^{\mathcal{U}}_\gamma)_a. \end{aligned}$$

The first and last equivalences displayed come from the relationship of an extender to its embedding, and the middle equivalence comes from the agreement between $i^{\mathcal{T}}_{\alpha,\theta+}$ and $i^{\mathcal{U}}_{\alpha,\theta+}$ our reflection argument produced. The second and fourth equivalences come from the fact that $\nu(E^{\mathcal{T}}_\xi) \leq \mathrm{crit}(i^{\mathcal{T}}_{\xi+1,\theta+})$ and $\nu(E^{\mathcal{U}}_\gamma) \leq \mathrm{crit}(i^{\mathcal{U}}_{\gamma+1,\theta+})$. This is because *generators are not moved* along the branches of an iteration tree: if e.g. $(\xi+1)T(\eta+1)$, then $E^{\mathcal{T}}_\eta$ has been applied to a model with index $> \xi$, so $\nu(E^{\mathcal{T}}_\xi) \leq \mathrm{crit}(E^{\mathcal{T}}_\eta)$.

This completes the proof of the claim. $\dashv$

Now let α be as in the claim, and set $\mathcal{T} = \mathcal{T}_\alpha$, $\mathcal{U} = \mathcal{U}_\alpha$, $\beta+1 = \mathrm{lh}(\mathcal{T})$, and $\gamma+1 = \mathrm{lh}(\mathcal{U})$. In order to complete our proof, we must show that we have not dropped in model or degree in a way which would make our comparison meaningless. Now if $\mathcal{M}^{\mathcal{T}}_\beta$ is a proper initial segment of $\mathcal{M}^{\mathcal{U}}_\gamma$, then $\mathcal{M}^{\mathcal{T}}_\beta$ is ω-sound, and hence by the remarks following 3.8 there can have been no dropping in model or degree along $[0,\beta]_T$, so that $i^{\mathcal{T}}_{0,\beta}$ exists and is a k-embedding, as desired. Similarly, if $\mathcal{M}^{\mathcal{U}}_\gamma$ is a proper initial segment of $\mathcal{M}^{\mathcal{T}}_\beta$, then $i^{\mathcal{U}}_{0,\gamma}$ exists and is a k-embedding. Thus we may assume $\mathcal{M}^{\mathcal{T}}_\beta = \mathcal{M}^{\mathcal{U}}_\gamma$. If $D^{\mathcal{T}} \cap [0,\beta]_T = \emptyset$ and $\deg^{\mathcal{T}}(\beta) = k$, then we are done, so let us assume otherwise. Similarly, we may assume that $D^{\mathcal{U}} \cap [0,\gamma]_U \neq \emptyset$ or $\deg^{\mathcal{U}}(\gamma) < k$. It follows from these assumptions that $\deg^{\mathcal{T}}(\beta) = \deg^{\mathcal{U}}(\gamma) = n$, where n is largest such that $\mathcal{M}^{\mathcal{T}}_\beta = \mathcal{M}^{\mathcal{U}}_\gamma$ is n-sound. (See 3.8.) But then, from 3.8 and the remarks following it, we see that there are $\xi+1 \in [0,\beta]_T$ and $\eta+1 \in [0,\gamma]_U$ such that

$$\begin{aligned} i^{\mathcal{T}}_{\xi+1,\beta} \circ i^{*\mathcal{T}}_{\xi+1} \quad &= \quad \text{uncollapse map from } \mathcal{C}_{n+1}(\mathcal{M}^{\mathcal{T}}_\beta) \text{ to } \mathcal{C}_n(\mathcal{M}^{\mathcal{T}}_\beta) \\ &= \quad \text{uncollapse map from } \mathcal{C}_{n+1}(\mathcal{M}^{\mathcal{U}}_\gamma) \text{ to } \mathcal{C}_n(\mathcal{M}^{\mathcal{U}}_\gamma) \\ &= \quad i^{\mathcal{U}}_{\eta+1,\gamma} \circ i^{*\mathcal{U}}_{\eta+1} \end{aligned}$$

Because generators are not moved along the branches of an iteration tree, we get as in the proof of the claim that the extender $E_\xi^{\mathcal{T}}$ giving rise to $i_{\xi+1}^{*\mathcal{T}}$ is compatible with the extender $E_\eta^{\mathcal{U}}$ giving rise to $i_{\eta+1}^{*\mathcal{U}}$. This contradicts the subclaim, and thereby completes the proof of the comparison theorem. $\dashv$

We note that the conclusion of the comparison lemma can be strengthened a bit in the case that one is comparing ω-sound mice using ω-maximal trees, which is the case of greatest interest. In this case, if $\mathcal{T}$ drops in model or degree along the branch leading to its last model, then $\mathcal{U}$ does not, and the last model of $\mathcal{U}$ is a proper initial segment of the last model of $\mathcal{T}$. This follows at once from 3.11 and the observation that the last model of $\mathcal{T}$ cannot be ω-sound in this case.

We can draw some simple corollaries concerning the definability of the reals belonging to mice.

3.12 Corollary. *Let $\mathcal{M}$ and $\mathcal{N}$ be ω-sound (ω, ω_1+1)-iterable premice such that $\rho_\omega(\mathcal{M}) = \rho_\omega(\mathcal{N}) = \omega$; then $\mathcal{M}$ is an initial segment of $\mathcal{N}$, or vice-versa.*

Proof. Since $\mathcal{M}$ and $\mathcal{N}$ are ω-sound and project to ω, they are countable, and so we have enough iterability to compare them. Let $\mathcal{T}$ on $\mathcal{M}$ and $\mathcal{U}$ on $\mathcal{N}$ be as in the conclusion of the comparison lemma 3.11, with last models $\mathcal{M}_\alpha$ and $\mathcal{N}_\eta$ respectively, and suppose without loss of generality that $\mathcal{M}_\alpha \trianglelefteq \mathcal{N}_\eta$ and $[0, \alpha]$ does not drop in model or degree. Since $\rho_\omega(\mathcal{M}) = \omega$, there are no extenders over $\mathcal{M}$ with critical point $< \rho_\omega(\mathcal{M})$, and therefore $\alpha > 0$ implies that $[0, \alpha]$ must drop in model or degree. So $\alpha = 0$. If $\eta = 0$ we are done, so assume $\eta > 0$. Since $\rho_\omega(\mathcal{N}) = \omega$, this implies $\mathcal{N}_\eta$ is not ω-sound. Thus $\mathcal{M}$ is a *proper* initial segment of $\mathcal{N}_\eta$, and $\mathcal{M}$ is countable in $\mathcal{N}_\eta$ because $\rho_\omega(\mathcal{M}) = \omega$. It is easy to see that this implies that $\mathcal{M}$ is an initial segment of $\mathcal{N}$, as desired. (One cannot gain reals by iterating, although one can lose them along some branch that drops.) $\dashv$

3.13 Corollary. *If $\mathcal{M}$ and $\mathcal{N}$ are (ω, ω_1+1)-iterable premice, then the $\mathcal{M}$-constructibility order on $\mathbb{R}\cap\mathcal{M}$ is an initial segment of the $\mathcal{N}$-constructibility order on $\mathbb{R} \cap \mathcal{N}$, or vice-versa.*

Proof. If $x \in \mathbb{R} \cap (\mathcal{J}_{\alpha+1}^{\mathcal{M}} \setminus \mathcal{J}_\alpha^{\mathcal{M}})$, then $\rho_\omega(\mathcal{J}_\alpha^{\mathcal{M}}) = \omega$. This observation and Corollary 3.12 easily yield the desired conclusion. $\dashv$

3.14 Corollary. *If $x \in \mathbb{R} \cap \mathcal{M}$ for some (ω, ω_1+1)-iterable premouse $\mathcal{M}$, then x is ordinal definable, and in fact x is Δ_2^2-definable from some countable ordinal.*

Proof. Say x is the α^{th} real in the $\mathcal{M}$-constructibility order. By 3.13 we know that the formula "v is the α^{th} real of some (ω, ω_1+1)-iterable premouse" characterizes x uniquely, so x is definable from α. In fact, by simply counting quantifiers one sees that (ω, ω_1+1)-iterability is Σ^2_3-definable, so x is Δ^2_3-definable from α. To see that x is Δ^2_2- definable, one uses the following equivalence:

$$\begin{aligned} y = x \quad \Leftrightarrow \quad & \exists \mathcal{M} \exists \Sigma (\mathcal{M} \text{ is a countable premouse and} \\ & \Sigma \text{ is an } (\omega, \omega_1)\text{-iteration strategy for } \mathcal{M} \text{ and} \\ & \forall \mathcal{N} \forall \Gamma (\text{ if } \mathcal{N} \text{ is a countable premouse which} \\ & \text{ has an } \alpha^{th} \text{ real } z \neq y, \text{ and} \\ & \Gamma \text{ is an } \omega_1\text{-iteration strategy for } \mathcal{N}, \text{ then} \\ & \text{ if } (\mathcal{T}, \mathcal{U}) \text{ is the } (\Sigma, \Gamma)\text{-coiteration of } \mathcal{M} \text{ with } \mathcal{N}, \\ & \text{ then } \mathcal{U} \text{ has no cofinal branch}) \end{aligned}$$

Here by the (Σ, Γ)-coiteration we mean the pair of iteration trees determined by Σ and Γ through the process of iterating away the least disagreement, as in 3.11. Since an ω_1 iteration strategy is essentially a set of reals, and the property of being an ω_1-iteration strategy is expressible using only real quantifiers, the formula displayed above is Σ^2_2, and hence x is Δ^2_2 in α. ⊣

We shall refine the proof of 3.14 later, and thereby obtain sharper upper bounds on the complexity of the reals in certain small mice. The refinement involves producing a logically simpler condition equivalent to (ω_1+1)-iterability in the case of these small mice.

4. The Dodd-Jensen Lemma

The Dodd-Jensen Lemma on the minimality of iteration maps is a fundamental, often-used tool in inner model theory.

4.1. The copying construction

Given a k-embedding $\pi \colon \mathcal{M} \to \mathcal{N}$ and a k-maximal iteration tree $\mathcal{T}$ on $\mathcal{M}$ with models $\mathcal{M}_\alpha$, we can lift $\mathcal{T}$ to a k-maximal iteration tree $\pi\mathcal{T}$ on $\mathcal{N}$ with models $\mathcal{N}_\alpha$. In fact, we need slightly less elementarity for π in order to construct $\pi\mathcal{T}$.

4.1 Definition. Let $\pi \colon \mathcal{C}_0(\mathcal{M}) \to \mathcal{C}_0(\mathcal{N})$ and let $k < \omega$. We call π a *weak k-embedding* iff

1. $\mathcal{M}$ and $\mathcal{N}$ are k-sound,
2. π is $r\Sigma_k$-elementary, and $r\Sigma_{k+1}$-elementary on parameters from some set X cofinal in $\rho_k(\mathcal{M})$,

3. $\pi(p_i(\mathcal{M})) = p_i(\mathcal{N})$, for all $i \leq k$, and

4. $\pi(\rho_i(\mathcal{M})) = \rho_i(\mathcal{N})$ for all $i < k$, and $\sup \pi``\rho_k(\mathcal{M}) \leq \rho_k(\mathcal{N})$.

A *weak ω-embedding* is just an ω-embedding, that is, a fully elementary map.

We shall construct $\pi\mathcal{T}$ by induction; at stage α we define its α^{th} model $\mathcal{N}_\alpha$, together with an embedding π_α from $\mathcal{C}_0(\mathcal{M}_\alpha)$ to $\mathcal{C}_0(\mathcal{N}_\alpha)$, as in the following figure:

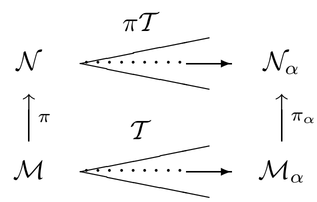

The next lemma describes the successor steps of this construction.

4.2 Lemma (Shift Lemma). *Let $\bar{\mathcal{M}}$ and $\bar{\mathcal{N}}$ be premice, let $\bar{\kappa} = \mathrm{crit}(\dot{F}^{\bar{\mathcal{N}}})$, and let*

$$\psi\colon \mathcal{C}_0(\bar{\mathcal{N}}) \to \mathcal{C}_0(\mathcal{N})$$

be a weak 0*-embedding, and*

$$\pi\colon \mathcal{C}_0(\bar{\mathcal{M}}) \to \mathcal{C}_0(\mathcal{M})$$

be a weak n-embedding. Suppose that $\bar{\mathcal{M}}$ and $\bar{\mathcal{N}}$ agree below $(\bar{\kappa}^+)^{\bar{\mathcal{M}}}$ and $(\bar{\kappa}^+)^{\bar{\mathcal{M}}} \leq (\bar{\kappa}^+)^{\bar{\mathcal{N}}}$, while $\mathcal{M}$ and $\mathcal{N}$ agree below $(\kappa^+)^{\mathcal{M}}$ and $(\kappa^+)^{\mathcal{M}} \leq (\kappa^+)^{\mathcal{N}}$, where $\kappa = \psi(\bar{\kappa})$. Suppose also

$$\pi \restriction (\bar{\kappa}^+)^{\bar{\mathcal{M}}} = \psi \restriction (\bar{\kappa}^+)^{\bar{\mathcal{N}}}.$$

Let $\bar{\kappa} < \rho_n(\bar{\mathcal{M}})$, so that $\mathrm{Ult}_n(\mathcal{C}_0(\bar{\mathcal{M}}), \dot{F}^{\bar{\mathcal{N}}})$ and $\mathrm{Ult}_n(\mathcal{C}_0(\mathcal{M}), \dot{F}^{\mathcal{N}})$ make sense, and suppose the latter ultrapower is wellfounded. Then the former ultrapower is wellfounded; moreover, there is a unique embedding $\sigma\colon \mathrm{Ult}_n(\mathcal{C}_0(\bar{\mathcal{M}}), \dot{F}^{\bar{\mathcal{N}}}) \to \mathrm{Ult}_n(\mathcal{C}_0(\mathcal{M}), \dot{F}^{\mathcal{N}})$ satisfying the conditions:

1. *σ is a weak n-embedding,*

2. *$\mathrm{Ult}_n(\mathcal{C}_0(\bar{\mathcal{M}}), \dot{F}^{\bar{\mathcal{N}}})$ agrees with $\bar{\mathcal{N}}$ below $\rho_0(\bar{\mathcal{N}})$, and $\mathrm{Ult}_n(\mathcal{C}_0(\mathcal{M}), \dot{F}^{\mathcal{N}})$ agrees with $\mathcal{N}$ below $\rho_0(\mathcal{N})$,*

3. *$\sigma \restriction (\rho_0(\bar{\mathcal{N}})) = \psi \restriction (\rho_0(\bar{\mathcal{N}}))$,*

4. *the diagram*

$$\begin{array}{ccc} \mathrm{Ult}_n(\mathcal{C}_0(\bar{\mathcal{M}}), \dot{F}^{\bar{\mathcal{N}}}) & \xrightarrow{\sigma} & \mathrm{Ult}_n(\mathcal{C}_0(\mathcal{M}), \dot{F}^{\mathcal{N}}) \\ {\scriptstyle i}\uparrow & & \uparrow{\scriptstyle j} \\ \mathcal{C}_0(\bar{\mathcal{M}}) & \xrightarrow[\pi]{} & \mathcal{C}_0(\mathcal{M}) \end{array}$$

commutes, where i and j are the canonical ultrapower embeddings.

The proof of the lemma is straightforward, so we omit it. In the representative special case $n = 0$, the desired map σ is defined by

$$\sigma([a, f]^{\bar{\mathcal{M}}}_{\dot{F}^{\bar{\mathcal{N}}}}) = [\psi(a), \pi(f)]^{\mathcal{M}}_{\dot{F}^{\mathcal{N}}}.$$

This is of course how it must be defined if we are to have conditions (3) and (4).

Now let $\pi\colon \mathcal{C}_0(\mathcal{M}) \to \mathcal{C}_0(\mathcal{N})$ be a weak k-embedding, and let $\mathcal{T}$ be a k-maximal iteration tree on $\mathcal{M}$. We define the models of a k-maximal copied tree $\pi\mathcal{T}$ on $\mathcal{N}$ by induction. In order to avoid some fine structural details, we shall assume first that no model on $\mathcal{T}$ is a type III premouse. In that case, $\pi\mathcal{T}$ will be a tree with the same order and drop structure as $\mathcal{T}$, and we shall have embeddings

$$\pi_\alpha\colon \mathcal{C}_0(\mathcal{M}_\alpha) \to \mathcal{C}_0(\mathcal{N}_\alpha).$$

We shall have $\deg^{\mathcal{T}}(\alpha) \le \deg^{\pi\mathcal{T}}(\alpha)$, with perhaps strict inequality being forced on us by the desire that $\pi\mathcal{T}$ be k-maximal. We use E_β and $i_{\beta,\alpha}$ for the extenders and embeddings of $\mathcal{T}$, and F_β and $j_{\beta,\alpha}$ for the extenders and embeddings of $\pi\mathcal{T}$, and we maintain inductively:

- π_α is a weak $\deg^{\mathcal{T}}(\alpha)$-embedding,
- if $\beta < \alpha$ and E_β is the last extender of the initial segment $\mathcal{P}$ of $\mathcal{M}_\beta$, then $\pi_\beta \restriction \rho_0(\mathcal{P}) = \pi_\alpha \restriction \rho_0(\mathcal{P})$, and
- if $\beta T \alpha$ and $(\beta, \alpha]_T \cap D = \emptyset$, then

$$\begin{array}{ccc} \mathcal{C}_0(\mathcal{N}_\beta) & \xrightarrow{j_{\beta,\alpha}} & \mathcal{C}_0(\mathcal{N}_\alpha) \\ {\scriptstyle \pi_\beta}\uparrow & & \uparrow{\scriptstyle \pi_\alpha} \\ \mathcal{C}_0(\mathcal{M}_\beta) & \xrightarrow[i_{\beta,\alpha}]{} & \mathcal{C}_0(\mathcal{M}_\alpha) \end{array}$$

  commutes.

We define $\mathcal{N}_{\alpha+1}$ and $\pi_{\alpha+1}$ by applying the Shift Lemma. Following the notation of the Shift lemma, we take $\bar{\mathcal{N}}$ to be the initial segment of $\mathcal{M}_\alpha$ whose last extender is E_α, and $\mathcal{N}$ to be $\pi_\alpha(\bar{\mathcal{N}})$ if $\bar{\mathcal{N}}$ is a proper initial segment of $\mathcal{M}_\alpha$, and $\mathcal{N} = \mathcal{N}_\alpha$ otherwise. (Because we have assumed $\mathcal{M}_\alpha$ is not of type III, $\mathcal{M}_\alpha$ is contained in the domain of π_α.) We take ψ to be the embedding with domain $\mathcal{C}_0(\bar{\mathcal{N}})$ induced by π_α. We let $F_\alpha = \dot{F}^{\mathcal{N}}$. Following further the Shift Lemma notation, $\bar{\mathcal{M}}$ is the initial segment $\mathcal{M}^*_{\alpha+1}$ of $\mathcal{M}_{\text{pred}_T(\alpha+1)}$ to which E_α is applied, and $\pi\colon \mathcal{C}_0(\bar{\mathcal{M}}) \to \mathcal{C}_0(\mathcal{M})$ is the map induced by π_β, for $\beta = \text{pred}_T(\alpha+1)$.) Let $n = \deg^{\mathcal{T}}(\alpha+1)$, and let $m = \deg^{\pi\mathcal{T}}(\alpha+1)$ be the degree dictated by F_α and our requirement that $\pi\mathcal{T}$ be k-maximal. One can check $n \leq m$. If the ultrapower $\text{Ult}_m(\mathcal{C}_0(\mathcal{N}), F_\alpha)$ giving rise to $\mathcal{N}_{\alpha+1}$ is illfounded, as may very well happen, then we stop the construction of $\pi\mathcal{T}$. Otherwise, let $\pi_{\alpha+1} = \tau \circ \sigma$, where σ is given by the Shift Lemma, and $\tau\colon \text{Ult}_n(\mathcal{C}_0(\mathcal{N}), F_\alpha) \to \text{Ult}_m(\mathcal{C}_0(\mathcal{N}), F_\alpha)$ is the natural map. It is easy to verify the induction hypotheses, and so we can continue.

At limit steps $\lambda < \text{lh}(\mathcal{T})$ we let $\mathcal{N}_\lambda$ be the direct limit over all $\alpha \in [0, \lambda)_T$, α sufficiently large, of the $\mathcal{N}_\alpha$, provided that this limit is wellfounded. We let π_λ be the embedding given by our induction hypothesis (3): $\pi_\lambda(i_{\alpha,\lambda}(x)) = j_{\alpha,\lambda}(\pi_\alpha(x))$. It is easy to verify the induction hypotheses. If the direct limit is illfounded, as may very well happen, we stop the construction of $\pi\mathcal{T}$.

Suppose now α is such that $\mathcal{M}_\alpha$ is type III. Letting $\bar{\mathcal{N}}$ be the initial segment of $\mathcal{M}_\alpha$ whose last extender is E_α, it is possible then that π_α does not act on $\mathcal{N}$, because the domain of π_α is only the squashed structure $\mathcal{C}_0(\mathcal{M}_\alpha)$. In the next paragraph, we include an outline of how to deal with this case, as a service to the scrupulous reader. We advise the unscrupulous reader to skip it.[23]

Let α be least such that $\mathcal{M}_\alpha$ is type III and let $\beta = \text{pred}_T(\alpha+1)$. If $\bar{\mathcal{N}} = \mathcal{M}_\alpha$, then we can just take F_α to be the last extender of $\mathcal{N}_\alpha$, and everything works out. The problem comes when $\bar{\mathcal{N}}$ is a proper initial segment of $\mathcal{M}_\alpha$, but not in the domain of π_α. But notice then that "un-squashing" upstairs gives $\psi\colon \text{Ult}(\mathcal{C}_0(\mathcal{M}_\alpha), \bar{F}) \to \text{Ult}(\mathcal{C}_0(\mathcal{N}_\alpha), F)$ which extends π_α, where $\bar{F}$ and F are the last extenders of $\mathcal{M}_\alpha$ and $\mathcal{N}_\alpha$ respectively. Let $\mathcal{N} = \psi(\bar{\mathcal{N}})$. The problem is that $\mathcal{N}$ may not be an initial segment of $\mathcal{N}_\alpha$. So we extend $\pi\mathcal{T}$ by two steps: first apply F to the appropriate initial segment of the appropriate model (as dictated by maximality), forming $\mathcal{N}_{\alpha+1} = \text{Ult}(\mathcal{Q}, F)$. It is easy to see that $\mathcal{N}$ is a proper initial segment of $\mathcal{P}$. We then take the last extender from $\mathcal{N}$ and apply it to the appropriate initial segment of $\mathcal{N}_\beta$ to get $\mathcal{N}_{\alpha+2}$. We have $\pi_{\alpha+1}\colon \mathcal{M}_{\alpha+1} \to \mathcal{N}_{\alpha+2}$ given by $\pi_{\alpha+1}([a,f]) = [\psi(a), \pi_\beta(f)]$. Again, everything works out. Thus in general, one step forward in $\mathcal{T}$ may correspond to two steps forward in $\pi\mathcal{T}$, and our copy maps π_γ map $\mathcal{M}^{\mathcal{T}}_\gamma$ to $\mathcal{N}^{\pi\mathcal{T}}_{\tau(\gamma)}$, where $\gamma < \tau(\gamma)$ is possible.

[23]We ignored this problem in [26]. Farmer Schlutzenberg found that error, and its repair.

This completes the definition of $\pi\mathcal{T}$.

4.3 Remark. A *near k-embedding*, is a weak k-embedding which is fully $r\Sigma_{k+1}$- elementary. If π_0 is a *near* k-embedding, then all π_α) are near $\deg^{\mathcal{T}}$-embeddings, and moreover $\deg^{\mathcal{T}}(\alpha) = \deg^{\pi\mathcal{T}}(\alpha)$. See [34, 1.3]. There is an error in [26], where it is claimed that one can copy under weak embeddings, while maintaining both $\deg^{\mathcal{T}}(\alpha) = \deg^{\pi\mathcal{T}}(\alpha)$ and that $\pi\mathcal{T}$ is maximal.[24] See [34] for more on how various degrees of elementarity are propagated in the copying construction.

The Dodd-Jensen lemma applies only to mice with a slightly stronger iterability property than the one we have introduced. In order to describe this property, we introduce an elaboration of the iteration game $\mathcal{G}_k(\mathcal{M}, \theta)$; a run of the new game is a linear composition of appropriately maximal iteration trees, rather than just a single such tree.

Let θ be an ordinal. In $\mathcal{G}_k(\mathcal{M}, \alpha, \theta)$, there are α rounds, the β^{th} being played as follows: Let $\mathcal{Q}$ be the last model in the linear composition produced before round β; that is, let $\mathcal{Q} = \mathcal{M}$ if $\beta = 0$, $\mathcal{Q}$ be the last model of the tree played during round $\beta - 1$ if $\beta > 0$ is a successor, and $\mathcal{Q}$ be the direct limit along the unique cofinal branch in the linear composition of trees produced before β, if β is a limit ordinal. (I wins if this branch is illfounded.) We let q, the degree of $\mathcal{Q}$, be k if $\beta = 0$, the degree of $\mathcal{Q}$ as a model of the tree played during round $\beta - 1$ (see 3.7) if $\beta > 0$ is a successor, and the eventual value of the degrees of previous rounds if β is a limit ordinal. I begins round β by choosing an initial segment $\mathcal{P}$ of $\mathcal{Q}$, and an $i \leq \omega$ such that if $\mathcal{P} = \mathcal{Q}$ then $i \leq q$, where q is the degree of $\mathcal{Q}$. The rest of round β is a run of $\mathcal{G}_i(\mathcal{P}, \theta)$,[25] except that we allow I to exit to round $\beta + 1$ before all θ moves have been played, and we require him to do so, on pain of losing, if θ is limit ordinal. (So if I has not lost, then when round β ends there will be in any case a last model to serve as $\mathcal{Q}$ for round $\beta+1$.) II wins $\mathcal{G}_k(\mathcal{M}, \alpha, \theta)$ just in case he does not lose any of the component games and, for $\beta \leq \alpha$ a limit ordinal, the unique cofinal branch in the composition of trees previously produced is wellfounded. A play of this game in which II has not yet lost is called a *k-bounded* iteration tree on $\mathcal{M}$. Notice that any winning strategy Γ for II in $\mathcal{G}_k(\mathcal{M}, \alpha, \theta)$ determines a winning strategy Σ for II in $\mathcal{G}_k(\mathcal{M}, \theta)$ in an obvious way: Σ calls for II to play as if he were using Γ in the first round of $\mathcal{G}_k(\mathcal{M}, \alpha, \theta)$, and I had not dropped to begin that round.

4.4 Definition. Let $\mathcal{M}$ be a k-sound premouse, where $k \leq \omega$; then a *(k, α, θ)-iteration strategy* for $\mathcal{M}$ is a winning strategy for II in $\mathcal{G}_k(\mathcal{M}, \alpha, \theta)$, and $\mathcal{M}$ is *(k, α, θ)-iterable* just in case there is such a strategy.

[24] Schlutzenberg also found this error, and its repair.
[25] So by our earlier conventions, i is the degree of $\mathcal{P}$.

The copying construction enables us to pull back iteration strategies for $\mathcal{N}$ to iteration strategies for premice embedded in $\mathcal{N}$.

4.5 Definition. Let $\pi\colon \mathcal{M} \to \mathcal{N}$ be a weak k-embedding, and Σ a strategy for II in $\mathcal{G}_k(\mathcal{N}, \theta)$, or in $\mathcal{G}_i(\mathcal{P}, \alpha, \theta)$ for some $\mathcal{P}$ such that $\mathcal{N}$ is an initial segment of $\mathcal{P}$ and i such that $i \leq k$ if $\mathcal{N} = \mathcal{P}$; then the *pullback of* Σ *under* π is the strategy Σ^π in the corresponding game on $\mathcal{M}$ such that for any k-bounded $\mathcal{T}$ on $\mathcal{M}$,

$$\mathcal{T} \ \textit{is by}\ \Sigma^\pi \Longleftrightarrow \pi\mathcal{T} \ \textit{is by}\ \Sigma.$$

Clearly, if Σ is a winning strategy for II an iteration game on $\mathcal{N}$, and $\pi\colon \mathcal{M} \to \mathcal{N}$ is sufficiently elementary, then Σ^π is a winning strategy for II in the corresponding game on $\mathcal{M}$. Thus

4.6 Theorem. *Suppose $\mathcal{N}$ is (k, θ)-iterable (respectively, (k, α, θ)-iterable), and there is a weak k-embedding from $\mathcal{M}$ into $\mathcal{N}$; then $\mathcal{M}$ is (k, θ)-iterable (respectively, (k, α, θ)-iterable.)*

4.2. The Dodd-Jensen Lemma

The following definition enables us to state an abstract form of the Dodd-Jensen Lemma.

4.7 Definition. Let Σ be a (k, λ, θ)-iteration strategy for $\mathcal{M}$, where λ is additively closed, and let $\mathcal{T}$ be an iteration tree played according to Σ; then we say $\mathcal{T}$ *is* (k, λ, θ)-*unambiguous* iff whenever $\alpha < \mathrm{lh}(\mathcal{T})$ is a limit ordinal, then $[0, \alpha]_T$ is the unique cofinal branch b of $\mathcal{T} \restriction \alpha$ such that $\mathcal{M}_b^{\mathcal{T}}$ is $(deg(b), \lambda, \theta)$-iterable.

So the unambiguous trees are just those which are played according to *every* (k, λ, θ)-iteration strategy for $\mathcal{M}$.

4.8 Theorem (The Dodd-Jensen Lemma). *Let λ be additively closed, let Σ be a (k, λ, θ)-iteration strategy for $\mathcal{M}$, and let $\mathcal{T}$ be an unambiguous iteration tree of length $\alpha + 1$ played according to Σ. Suppose* $\deg^{\mathcal{T}}(\alpha) = k$, *and $\pi\colon \mathcal{M} \to \mathcal{N}$ is a weak k-embedding, where $\mathcal{N}$ is an initial segment of $\mathcal{M}_\alpha^{\mathcal{T}}$; then*

1. *$\mathcal{N} = \mathcal{M}_\alpha^{\mathcal{T}}$,*
2. *$[0, \alpha]_T$ does not drop (in model or degree), and*
3. *for all $x \in \mathcal{M}$, $i_{0,\alpha}^{\mathcal{T}}(x) \leq_L \pi(x)$, where $\leq_L$ is the order of construction.*

Proof. Assume first toward contradiction that $\mathcal{N}$ is a proper initial segment of $\mathcal{M}_\alpha^{\mathcal{T}}$. We shall construct a run r of $\mathcal{G}_k(\mathcal{M}, \lambda, \theta)$ which is a loss for Σ. The run r is divided into ω blocks, each consisting of a number of rounds of $\mathcal{G}_k(\mathcal{M}, \lambda, \theta)$ equal to the number of rounds in $\mathcal{T}$. We shall use $\mathcal{T}_n$ for the iteration tree played in the n^{th} block of r, and $\mathcal{M}_n$ for the base model of $\mathcal{T}_n$. Thus $\mathcal{M}_{n+1}$ is the model player I drops to at the beginning of the first round in block $n+1$ of r; we have I drop to the degree k at the beginning of this round. We shall arrange that $\mathcal{M}_{n+1}$ is a proper initial segment of the last model of $\mathcal{T}_n$, so that the unique cofinal branch of the composition of the $\mathcal{T}_n$'s is illfounded, and r is indeed a loss for Σ. As an auxiliary we define maps $\pi_n\colon \mathcal{M}_n \to \mathcal{M}_{n+1}$ as we proceed.

Set $\mathcal{M}_0 = \mathcal{M}$, $\mathcal{T}_0 = \mathcal{T}$, $\mathcal{M}_1 = \mathcal{N}$, and $\pi_0 = \pi$.

Now suppose $\mathcal{M}_n$, $\mathcal{T}_n$, $\mathcal{M}_{n+1}$, and π_n are given. Set $\mathcal{T}_{n+1} = \pi_n \mathcal{T}_n$. We shall check shortly that $\mathcal{T}_{n+1}$ is played according to Σ, so that $\mathrm{lh}(\mathcal{T}_{n+1}) = \mathrm{lh}(\mathcal{T}_n)$, and we have from the copying construction an embedding σ from the last model of $\mathcal{T}_n$ to the last model of $\mathcal{T}_{n+1}$. Now $\mathcal{M}_{n+1} \in \mathrm{dom}(\sigma)$, so we can set $\mathcal{M}_{n+2} = \sigma(\mathcal{M}_{n+1})$ and $\pi_{n+1} = \sigma \restriction \mathcal{M}_{n+1}$. This completes the construction of r, and thereby gives the desired contradiction.

We now show that $\mathcal{T}_{n+1}$ is a play according to Σ. Let us call a position u which is according to Σ *transitional* if $u = (s, (\mathcal{P}, i))$ where s represents some number $\beta < \lambda$ of complete rounds of play according to Σ in which I has not lost, and $(\mathcal{P}, i)$ is a way I might legally begin round β. Notice that in this situation, Σ determines an (i, λ, θ)-iteration strategy for $\mathcal{P}$. We call this strategy Σ_u. Now let u and v be the transitional initial segments of r ending with $(\mathcal{M}_n, k)$ and $(\mathcal{M}_{n+1}, k)$ respectively. Let $\psi = \pi_{n-1} \circ \ldots \circ \pi_0$ and $\tau = \pi_n \circ \ldots \circ \pi_0$, so that $\psi\colon \mathcal{M} \to \mathcal{M}_n$ and $\tau\colon \mathcal{M} \to \mathcal{M}_{n+1}$ are weak k embeddings. Since $(\Sigma_u)^\psi$ and $(\Sigma_v)^\tau$ are (k, λ, θ)-iteration strategies for $\mathcal{M}$ and $\mathcal{T}$ is unambiguous, $\mathcal{T}$ is a play by each of $(\Sigma_u)^\psi$ and $(\Sigma_v)^\tau$. Therefore $\psi\mathcal{T}$ and $\tau\mathcal{T}$ are plays according to Σ, and since $\tau\mathcal{T} = \pi_n \circ \psi\mathcal{T} = \pi_n\mathcal{T}_n = \mathcal{T}_{n+1}$, we are done.

The proofs of conclusions 2 and 3 of the Dodd-Jensen lemma are similar. We construct $\mathcal{M}_n$, $\mathcal{T}_n$, and π_n as above, but now we have that $\mathcal{M}_{n+1}$ is the last model of $\mathcal{T}_n$. If the branch of $\mathcal{T}$ from $\mathcal{M}$ to $\mathcal{N} = \mathcal{M}_1$ drops, then the branch of $\mathcal{T}_n$ from $\mathcal{M}_n$ to $\mathcal{M}_{n+1}$ drops for each n, and the unique cofinal branch of the composition of the $\mathcal{T}_n$'s is illfounded. Thus we may assume that the branch of $\mathcal{T}$ from $\mathcal{M}$ to $\mathcal{N}$ does not drop, so that 2 holds. This implies that for all n, the branch of $\mathcal{T}_n$ from $\mathcal{M}_n$ to $\mathcal{M}_{n+1}$ does not drop, so that we have an iteration map $i_n\colon \mathcal{M}_n \to \mathcal{M}_{n+1}$ given by $\mathcal{T}_n$. Assume that conclusion 3 fails, and fix $x_0 \in \mathcal{M}_0$ such that $\pi_0(x_0) <_L i_0(x_0)$. For any $n \geq 0$, define x_{n+1} by: $x_{n+1} = \pi_n(x_n)$. It is easy to check that $x_{n+1} <_L i_n(x_n)$ for all n. (This is true for $n = 0$ by hypothesis. But if $x_{n+1} <_L i_n(x_n)$, then

$$x_{n+2} = \pi_{n+1}(x_{n+1}) <_L \pi_{n+1}(i_n(x_n)) = i_{n+1}(\pi_n(x_n)) = i_{n+1}(x_{n+1}),$$

because $\pi_{n+1} \circ i_n = i_{n+1} \circ \pi_n$ by the commutativity of the copy maps.) Thus again, the unique cofinal branch of the composition of the $\mathcal{T}_n$'s is illfounded, and we have a loss for Σ.

The following diagram illustrates the proof we have given for conclusion 3.

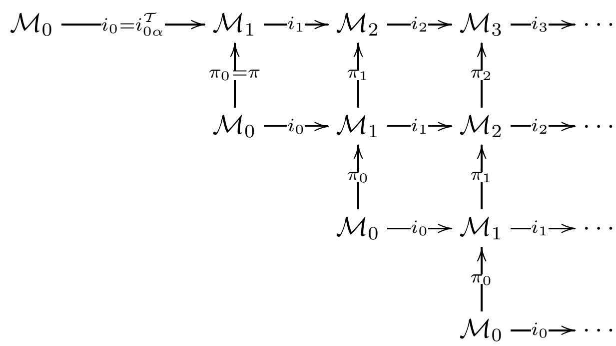

$\dashv$

4.3. The Weak Dodd-Jensen Property

Unfortunately, there are important contexts in which one wants to use the Dodd-Jensen Lemma, but in which one does not know that the given iteration strategy is unambiguous. One such context is the proof of the key fine structural fact that the standard parameters of a sufficiently iterable mouse are solid and universal. (We shall prove this in the next section.) Fortunately, one can construct from any iteration strategy for a countable mouse another iteration strategy which satisfies a weak version of the Dodd-Jensen Lemma, and this weak version suffices for the proof of solidity and universality. Since the construction is simple and natural, we shall give it here.

The notions and results in this subsection come from [30].

Let $\mathcal{M}$ and $\mathcal{P}$ be premice; then we say that $\mathcal{P}$ is $(\mathcal{M}, k)$-*large* just in case there is a near k-embedding from $\mathcal{M}$ to an initial segment of $\mathcal{P}$. (A near k-embedding is a weak k-embedding which is $r\Sigma_{k+1}$ elementary. See [34, 1.2,1.3], where it is shown that the copying construction gives rise to such embeddings. We could make do with weak k-embeddings here, but it would be a bit awkward at one point.) Let $\vec{e} = \langle e_i \mid i < \omega \rangle$ enumerate the universe of a countable premouse $\mathcal{M}$, and $\pi \colon \mathcal{M} \to \mathcal{P}$ be a near k-embedding; then we say π is $(k, \vec{e})$-*minimal* iff whenever σ is a near k-embedding from $\mathcal{M}$ to an initial segment $\mathcal{N}$ of $\mathcal{P}$, then $\mathcal{N} = \mathcal{P}$ and either $\sigma = \pi$, or $\sigma(e_i) >_L \pi(e_i)$ where i is least such that $\sigma(e_i) \neq \pi(e_i)$. Notice that if $\mathcal{P}$ is $(\mathcal{M}, k)$-large but no proper initial segment of $\mathcal{P}$ is $(\mathcal{M}, k)$-large, then there is a $(k, \vec{e})$-minimal embedding from $\mathcal{M}$ to $\mathcal{P}$. This embedding is just the leftmost branch through a certain tree.

4.9 Definition. Let Σ be a (k, α, θ)-iteration strategy for a countable premouse $\mathcal{M}$, and let $\vec{e} = \langle e_i \mid i < \omega \rangle$ enumerate the universe of $\mathcal{M}$ in order type ω; then we say Σ has the *weak Dodd-Jensen property* (relative to $\vec{e}$) iff whenever $\mathcal{T}$ is an iteration tree on $\mathcal{M}$ played according to Σ, and $\beta < lh(\mathcal{T})$ is such that $\mathcal{M}^{\mathcal{T}}_{\beta}$ is $(\mathcal{M}, k)$-large, then $i^{\mathcal{T}}_{0,\beta}$ exists and is $(k, \vec{e})$-minimal.

4.10 Theorem (The Weak Dodd-Jensen Lemma). *Suppose $\mathcal{M}$ is (k, ω_1, θ)-iterable, and that $\vec{e}$ enumerates the universe of $\mathcal{M}$ in order type ω; then there is a (k, ω_1, θ)-iteration strategy for $\mathcal{M}$ which has the weak Dodd-Jensen property relative to $\vec{e}$.*

Proof. Let Σ be any (k, ω_1, θ)-iteration strategy for $\mathcal{M}$. We shall construct a transitional position $u = (r, (\mathcal{P}, k))$ of Σ and a $(k, \vec{e})$-minimal embedding $\pi\colon \mathcal{M} \to \mathcal{P}$ such that π is strongly $(k, \vec{e})$ minimal, in the sense that whenever $\mathcal{R}$ is an $(\mathcal{M}, k)$-large Σ_u-iterate of $\mathcal{P}$, then there is no dropping in the iteration from $\mathcal{P}$ to $\mathcal{R}$, and if $i\colon \mathcal{P} \to \mathcal{R}$ is the iteration map, then $i \circ \pi$ is $(k, \vec{e})$ minimal. It is then easy to see that the π-pullback of Σ_u has the weak Dodd-Jensen property.

Let us call a pair $(r, \mathcal{Q})$ *suitable* if $(r, (\mathcal{Q}, k))$ is transitional, and $\mathcal{Q}$ is $(\mathcal{M}, k)$-large but no proper initial segment of $\mathcal{Q}$ is $(\mathcal{M}, k)$-large. In order to obtain the desired u and π, we define by induction on $n < \omega$ suitable pairs $(r_n, \mathcal{P}_n)$. We maintain inductively that r_{n+1} extends $(r_n, (\mathcal{P}_n, k))$. We begin by letting r_0 be the empty position, and $\mathcal{P}_0 = \mathcal{M}$. Now suppose r_n and $\mathcal{P}_n$ have been defined.

Case 1. There is a suitable $(s, \mathcal{Q})$ such that s extends $(r_n, (\mathcal{P}_n, k))$ and the branch $\mathcal{P}_n$-to-$\mathcal{Q}$ in the iteration given by s has a drop.

In this case, we simply let $(r_{n+1}, \mathcal{P}_{n+1})$ be any such $(s, \mathcal{Q})$.

Case 2. Otherwise.

Let $\tau\colon \mathcal{M} \to \mathcal{P}_n$ be $(k, \vec{e})$-minimal.

Subcase 2a. There is a suitable $(s, \mathcal{Q})$ such that s extends $(r_n, (\mathcal{P}_n, k))$, and letting $i\colon \mathcal{P}_n \to \mathcal{Q}$ be the iteration map given by s, $i \circ \tau$ is not $(k, \vec{e})$-minimal.

In this case, let $m < \omega$ be least such that for some such s, $\mathcal{Q}$, and i we have, letting $\sigma\colon \mathcal{M} \to \mathcal{Q}$ be $(k, \vec{e})$-minimal, that $\sigma(e_m) \neq i \circ \tau(e_m)$ (and thus $\sigma(e_m) <_L i \circ \tau(e_m)$). We then let $(r_{n+1}, \mathcal{P}_{n+1})$ be a suitable pair $(s, \mathcal{Q})$ witnessing this property of m.

Subcase 2b. Otherwise.

In this case τ is strongly $(k, \vec{e})$-minimal in the sense advertised earlier, so we set $u = (r_n, (\mathcal{P}_n, k))$ and $\pi = \tau$, and stop the construction.

Now suppose the construction never stops. Notice that case 1 can only apply finitely often, since otherwise we get an iteration tree played according to Σ whose unique cofinal branch has infinitely many drops. Suppose then that case 2 applies at all $n \geq n_0$, so that for all $n_0 \leq n \leq m$ we have a k-embedding $i_{n,m}\colon \mathcal{P}_n \to \mathcal{P}_m$ given by r_m. For $n \geq n_0$, let $\pi_n\colon \mathcal{M} \to \mathcal{P}_n$

be $(k, \vec{e})$-minimal; then if $n_0 \leq n < m$, π_m is "to the left" of $i_{n,m} \circ \pi_n$. It follows that for any j, $i_{n,m}(\pi_n(e_j)) = \pi_m(e_j)$ for all sufficiently large n, m (by induction on j). Let

$$r = \bigcup_{n<\omega} r_n, \qquad \mathcal{P} = \lim_{n\to\infty} \mathcal{P}_n, \qquad u = (r, (\mathcal{P}, k)),$$

let $i_{n,\infty} : \mathcal{P}_n \to \mathcal{P}$ be the direct limit map (a k-embedding), and define $\pi\colon \mathcal{M} \to \mathcal{P}$ by

$$\pi(e_j) = \text{ eventual value of } \ i_{n,\infty}(\pi_n(e_j)), \text{ as } n \to \infty.$$

We claim that u and π are as advertised earlier.

Clearly π is a near k-embedding, and so $\mathcal{P}$ is $(\mathcal{M}, k)$-large. No proper initial segment $\mathcal{R}$ of $\mathcal{P}$ can be $(\mathcal{M}, k)$-large, as then $(u, \mathcal{R})$ could serve as the $(s, \mathcal{Q})$ witnessing the occurrence of case 1 at a stage $n > n_0$. Similarly, π is $(k, \vec{e})$-minimal. For if σ is a near k-embedding of $\mathcal{M}$ into $\mathcal{P}$ which is to the left of π, then take m_0 to be the least j such that $\sigma(e_j) \neq \pi(e_j)$, and let $l < \omega$ be so large that $n_0 < l$ and $\pi(e_j) = i_{l,\infty}(\pi_l(e_j))$ for all $j \leq m_0$ (and so $m > m_0$, where m is as in case 2a at stage l). Then r,$\mathcal{P}$, and σ could serve as the s,$\mathcal{Q}$, and σ witnessing $m \leq m_0$ at stage l, contradiction. Finally, let $\mathcal{R}$ be any $(\mathcal{M}, k)$-large iterate of $\mathcal{P}$ via Σ_u. Clearly, there is a transitional position $(v, (\mathcal{R}, k))$ such that v extends u. We can argue as above that there is no dropping in the iteration tree given by v from $\mathcal{P}$ to $\mathcal{R}$, and that if $i\colon \mathcal{P} \to \mathcal{R}$ is the iteration map, then $i \circ \pi$ is $(k, \vec{e})$-minimal. Thus u and π are as advertised.

We leave to the reader the easy verification that $(\Sigma_u)^\pi$ has the weak Dodd-Jensen property. ⊣

The weak Dodd-Jensen property isolates a unique iteration strategy, modulo the enumeration $\vec{e}$. Since the main ideas in the proof of this fact are used very often in inner model theory, we give it here.

4.11 Theorem. *Let $\vec{e}$ enumerate the universe of the k-sound premouse $\mathcal{M}$ in order type ω; then there is at most one $(k, \omega_1 + 1)$-iteration strategy for $\mathcal{M}$ which has the weak Dodd-Jensen property relative to $\vec{e}$.*

Proof. Suppose Σ and Γ are distinct such strategies. We can find a k-maximal iteration tree $\mathcal{T}$ on $\mathcal{M}$ such that $\mathcal{T}$ has limit length $\lambda < \omega_1$, $\mathcal{T}$ is played according to both Σ and Γ, and $\Sigma(\mathcal{T}) \neq \Gamma(\mathcal{T})$. Let $\mathcal{U}^*$ and $\mathcal{V}^*$ be the iteration trees of length $\lambda + 1$ extending $\mathcal{T}$ produced by Σ and Γ respectively. We now proceed as if we had produced $\mathcal{U}^*$ and $\mathcal{V}^*$ on the two sides of a coiteration, and continue "iterating the least disagreement". We thereby extend $\mathcal{U}^*$ and $\mathcal{V}^*$ to k-maximal trees $\mathcal{U}$ and $\mathcal{V}$, played according to Σ and Γ respectively, in such a way that the last model of one is an initial segment of the last model of the other. We may as well assume that $\mathcal{M}^{\mathcal{U}}_\alpha$

is an initial segment of $\mathcal{M}^{\mathcal{V}}_\beta$. As in the comparison lemma 3.11, one of the two trees does not drop along the branch leading to its last model, so we can assume that $D^{\mathcal{U}} \cap [0,\alpha]_U = \emptyset$ and $\deg^{\mathcal{U}}(\alpha) = k$, and hence $i^{\mathcal{U}}_{0,\alpha}$ exists and is a k-embedding.

It follows that $\mathcal{M}^{\mathcal{V}}_\beta$ is $(\mathcal{M},k)$-large. Since Γ has the weak Dodd-Jensen property relative to $\vec{e}$, $i^{\mathcal{V}}_{0,\beta}$ exists and is $(k,\vec{e})$-minimal. This implies that no proper initial segment of $\mathcal{M}^{\mathcal{V}}_\beta$ is $(\mathcal{M},k)$-large, so $\mathcal{M}^{\mathcal{U}}_\alpha = \mathcal{M}^{\mathcal{V}}_\beta$. Because Σ also has the weak Dodd-Jensen property relative to $\vec{e}$, $i^{\mathcal{U}}_{0,\alpha}$ is also $(k,\vec{e})$-minimal. It follows that $i^{\mathcal{U}}_{0,\alpha} = i^{\mathcal{V}}_{0,\beta}$.

Notice that since $\Sigma(\mathcal{T}) \neq \Gamma(\mathcal{T})$, $[0,\alpha]_U \cap [0,\beta]_V$ is bounded in λ. As branches in an iteration tree are closed below their sups, we have a largest ordinal γ such that $\gamma \in [0,\alpha]_U \cap [0,\beta]_V \cap \lambda$. Let $\nu = \sup\{\nu(E^{\mathcal{T}}_\xi) \mid \xi T \gamma\}$. Every member of $\mathcal{M}^{\mathcal{T}}_\gamma$ is of the form $i^{\mathcal{T}}_{0,\gamma}(f)(a)$, for some $f \in \mathcal{M}$ and $a \in [\nu]^{<\omega}$. (We take $k = 0$ for notational simplicity; otherwise we have f $r\Sigma_k$ over $\mathcal{M}$.) Since $i^{\mathcal{U}}_{\gamma,\alpha}$ and $i^{\mathcal{V}}_{\gamma,\beta}$ have critical point at least ν, this representation of $\mathcal{M}^{\mathcal{T}}_\gamma$ and the fact that $i^{\mathcal{U}}_{0,\alpha} = i^{\mathcal{V}}_{0,\beta}$ yield that $i^{\mathcal{U}}_{\gamma,\alpha} = i^{\mathcal{V}}_{\gamma,\beta}$.

Let $\xi + 1 \in (\gamma,\alpha]_U$ be such that U-pred$(\xi+1) = \gamma$. Let $\sigma + 1 \in (\gamma,\beta]_V$ be such that V-pred$(\sigma+1) = \gamma$. Since $i^{\mathcal{U}}_{\gamma,\alpha} = i^{\mathcal{V}}_{\gamma,\beta}$, the extenders $E^{\mathcal{U}}_\xi$ and $E^{\mathcal{V}}_\sigma$ are compatible, that is, they agree up to the inf of the sups of their generators. If $\xi < \lambda$ or $\sigma < \lambda$, this is impossible as no extender used in an iteration tree is compatible with any extender used later in the same tree. (If $\alpha < \beta$ and E_α is compatible with E_β, then $E_\alpha \in \mathcal{M}_\beta$ by the initial segment condition. This implies that $\mathrm{lh}(E_\alpha)$ is not a cardinal in $\mathcal{M}_\beta$, contrary to 3.5.) If $\lambda \leq \xi$ and $\lambda \leq \sigma$, this is impossible as no two extenders used in a coiteration are compatible. (This was a subclaim in the proof of 3.11.) This contradiction completes the proof. ⊣

5. Solidity and Condensation

In this section we shall sketch the proofs of two theorems which are central in the fine structural analysis of definability over mice. These results are much deeper than the fine structural results of section 2. Their proofs involve comparison arguments, and hence require an iterability hypothesis. The proofs also use the weak Dodd-Jensen property, and they illustrate a very useful technique for insuring that in certain comparisons, the critical point of the embedding from the first to the last model in one of the trees is not too small.

Our first theorem is a condensation result.

5.1 Theorem. *Let $\mathcal{M}$ be ω-sound and $(\omega,\omega_1,\omega_1+1)$-iterable. Suppose $\pi\colon \mathcal{H} \to \mathcal{M}$ is fully elementary, and crit$(\pi)=\rho^{\mathcal{H}}_\omega$; then either*

1. *$\mathcal{H}$ is a proper initial segment of $\mathcal{M}$, or*

2. *there is an extender E on the $\mathcal{M}$-sequence such that lh(E)= $\rho_\omega^{\mathcal{H}}$, and $\mathcal{H}$ is a proper initial segment of* $\mathrm{Ult}_0(\mathcal{M}, E)$.

5.2 Remarks. The complexities in the statement of 5.1 are necessary.

1. The hypothesis that $\mathrm{crit}(\pi) = \rho_\omega^{\mathcal{H}}$ is necessary in 5.1. For notice that $\mathrm{crit}(\pi) > \rho_\omega^{\mathcal{H}}$ is impossible since otherwise we would have $\rho_\omega^{\mathcal{H}} = \rho_\omega^{\mathcal{M}}$, and since $\mathcal{M}$ is ω-sound, this would imply that $\mathrm{crit}(\pi)$ is definable over $\mathcal{M}$ from points in the range of π. On the other hand, $\mathrm{crit}(\pi) < \rho_\omega^{\mathcal{H}}$ can occur while the conclusions of 5.1 fail: for example, let $\mathcal{M} = \mathrm{Ult}_\omega(\mathcal{H}, E)$, where E is on the $\mathcal{H}$-sequence and $\mathrm{crit}(E) < \rho_\omega^{\mathcal{H}}$, and let π be the canonical embedding.

2. The alternatives in the conclusion of 5.1 are mutually exclusive, since in the second case the extender E is on the $\mathcal{M}$-sequence, but not on the $\mathcal{H}$-sequence. The following example shows that the second alternative can occur. Suppose $\mathcal{P}$ is an active, ω-sound mouse, and F is the last extender on the $\mathcal{P}$-sequence. Let $\kappa = \mathrm{crit}(F)$, and suppose $F \restriction \alpha$ is on the $\mathcal{P}$-sequence, where $\alpha > (\kappa^+)^{\mathcal{P}}$. (Under weak large cardinal hypotheses, there is such a $\mathcal{P}$.) Let

$$\sigma \colon \mathrm{Ult}_0(\mathcal{P}, F \restriction \alpha) \to \mathrm{Ult}_0(\mathcal{P}, F)$$

be the natural embedding. Since α is a cardinal in $\mathrm{Ult}_0(\mathcal{P}, F \restriction \alpha)$ by clause 1 of 2.4, and not a cardinal in $\mathrm{Ult}_0(\mathcal{P}, F)$ because $F \restriction \alpha$ is in this model and collapses α, we have that $\alpha = \mathrm{crit}(\sigma)$. Let

$$\mathcal{H} = \mathcal{J}_{\alpha+1}^{\mathrm{Ult}_0(\mathcal{P}, F\restriction\alpha)}$$

and

$$\mathcal{M} = \sigma(\mathcal{H}), \pi = \sigma \restriction \mathcal{H}.$$

Clearly $\alpha = \mathrm{crit}(\pi) = \rho_\omega^{\mathcal{H}}$, π is fully elementary, and $\mathcal{H}$ is not an initial segment of $\mathcal{M}$.

Proof of 5.1. Let $\mathcal{H}$ and $\mathcal{M}$ constitute a counterexample. Let $X \prec V_\lambda$ for some limit ordinal λ, with X countable and $\mathcal{H}, \mathcal{M} \in X$, and let $\bar{\mathcal{H}}$ and $\bar{\mathcal{M}}$ be the images of $\mathcal{H}$ and $\mathcal{M}$ under the transitive collapse of X. It is easy to see that $\bar{\mathcal{H}}$ and $\bar{\mathcal{M}}$ still constitute a counterexample to 5.1. Thus we may assume without loss of generality that $\mathcal{M}$ is countable. We can therefore fix an enumeration $\vec{e}$ of $\mathcal{M}$ in order type ω, and an $(\omega, \omega_1, \omega_1 + 1)$-iteration strategy Σ for $\mathcal{M}$ having the weak Dodd-Jensen property relative to $\vec{e}$.

The natural plan is to compare $\mathcal{H}$ with $\mathcal{M}$, using Σ to iterate $\mathcal{M}$ and Σ^π to iterate $\mathcal{H}$. Suppose $\mathcal{P}$ is the last model of the tree $\mathcal{T}$ on $\mathcal{H}$ and $\mathcal{Q}$ is the last model of the tree $\mathcal{U}$ on $\mathcal{M}$ in this comparison. We would like to see that $\mathcal{P} = \mathcal{H}$, for then it is clear that $\mathcal{H}$ is an initial segment of $\mathcal{Q}$, and a little further argument, given below, shows that $\mathcal{U}$ uses at most one extender, so that one of the alternatives in the conclusion of 5.1 must hold. Assume then that $\mathcal{P} \neq \mathcal{H}$.

If the branch $\mathcal{H}$-to-$\mathcal{P}$ of $\mathcal{T}$ drops in model or degree, then $\mathcal{M}$-to-$\mathcal{Q}$ does not drop in model or degree, and $\mathcal{Q}$ is a proper initial segment of $\mathcal{P}$. (Here we use that $\mathcal{T}$ and $\mathcal{U}$ are ω-maximal.) But then, letting $j \colon \mathcal{M} \to \mathcal{Q}$ be the iteration map, and $\tau \colon \mathcal{P} \to \mathcal{R}$ be the copy map from $\mathcal{P}$ to the last model of $\pi\mathcal{T}$, we have that $\tau \circ j$ maps $\mathcal{M}$ to a proper initial segment of $\mathcal{R}$, and $\mathcal{R}$ is a Σ-iterate of $\mathcal{M}$. This contradicts the weak Dodd-Jensen property of Σ. Thus $\mathcal{H}$-to-$\mathcal{P}$ does not drop in model or degree, and we have a fully elementary iteration map $i \colon \mathcal{H} \to \mathcal{P}$.

Since the branch $\mathcal{H}$-to-$\mathcal{P}$ does not drop in model or degree, we must have $\text{crit}(i) < \rho_\omega^{\mathcal{H}}$. Let $\rho = \rho_\omega^{\mathcal{H}}$. Since $\text{crit}(\pi) = \rho$, $\mathcal{H}$ and $\mathcal{M}$ agree below ρ, so that all extenders used in their comparison have length at least ρ. Nevertheless, it is possible that the first extender E used along $\mathcal{H}$-to-$\mathcal{P}$ is such that $\text{crit}(E) < \rho \leq \text{lh}(E)$. This possibility ruins our proof, so we must modify the construction of $\mathcal{T}$ so as to avoid it.

We modify the construction so that if E is an extender used in $\mathcal{T}$ and $\text{crit}(E) < \rho$, then E is used in $\mathcal{T}$ to take an ultrapower of $\mathcal{M}$, or rather the longest initial segment of $\mathcal{M}$ containing only subsets of $\text{crit}(E)$ measured by E, instead of being used to take an ultrapower of $\mathcal{H}$, as it would be in a tree on $\mathcal{H}$. This modification is possible because $\mathcal{M}$ and $\mathcal{H}$ agree below ρ. The system $\mathcal{T}$ we form in this way is not an ordinary iteration tree, but rather a "double-rooted" iteration tree whose base is the pair of models $(\mathcal{M}, \mathcal{H})$. We shall use $\mathcal{P}_\alpha$ for the α^{th} model of $\mathcal{T}$, and E_α for the extender taken from the $\mathcal{P}_\alpha$-sequence and used to form $\mathcal{P}_{\alpha+1}$. Let

$$\mathcal{P}_0 = \mathcal{M}, \text{ and } \mathcal{P}_1 = \mathcal{H}.$$

Let $E_0 = \emptyset$, and

$$\nu(E_0) = \rho.$$

For $\alpha \geq 1$, E_α is the extender on the $\mathcal{P}_\alpha$-sequence which participates in its least disagreement with the sequence of the current last model in $\mathcal{U}$. As in an ordinary iteration tree,

$$\text{pred}_T(\alpha + 1) = \text{ least } \beta \text{ such that } \text{crit}(E_\alpha) < \nu(E_\beta),$$

and

$$\mathcal{P}_{\alpha+1} = \text{Ult}_n(\mathcal{P}^*_{\alpha+1}, E_\alpha),$$

where $\mathcal{P}^*_{\alpha+1}$ is the longest initial segment of $\mathcal{P}_\beta$ and n is the largest number $\leq \omega$ such that the ultrapower in question makes sense. (That is, we do so in all but one anomalous case, which we shall explain in the next paragraph.) Our convention on $\nu(E_0)$ and the fact that the $\nu(E_\alpha)$ are increasing then implies that if $\text{crit}(E_\alpha) < \rho$, then $\text{pred}_T(\alpha+1) = 0$, so that E_α is applied in $\mathcal{T}$ to an initial segment of $\mathcal{M}$.

There is one anomalous case here.[26] Suppose $\text{crit}(E_\alpha) := \kappa < \rho$, and let $\mathcal{P}^*_{\alpha+1}$ be the longest initial segment $\mathcal{Q}$ of $\mathcal{M}$ such that $P(\kappa) \cap \mathcal{Q} \subseteq \mathcal{P}_\alpha$. It can happen that $\mathcal{P}^*_{\alpha+1}$ is of type III, with $\nu(\mathcal{P}^*_{\alpha+1}) = \kappa$. (One can show easily then that $\rho = (\kappa^+)^{\mathcal{H}}$, and $\mathcal{P}^*_{\alpha+1} = \mathcal{J}^{\mathcal{M}}_\rho$.) In this case, $\text{Ult}_0(\mathcal{C}_0(\mathcal{P}^*_{\alpha+1}), E_\alpha)$ does not make sense, because $\mathcal{C}_0(\mathcal{P}^*_{\alpha+1})$ has ordinal height $\text{crit}(E_\alpha)$.[27] We must therefore return to our old, naïve meaning for $\text{Ult}_0(\mathcal{P}^*_{\alpha+1}, E_\alpha)$. Let k be the canonical embedding associated to this ultrapower, and let F be the last extender of $\mathcal{P}^*_{\alpha+1}$. Then we set

$$\mathcal{P}_{\alpha+1} = \text{Ult}_\omega(\mathcal{M}, k(F)).$$

Note here that $k(F)$ is indeed a total extender over $\mathcal{M}$ with critical point strictly less than $\rho_\omega(\mathcal{M})$.

Unfortunately, the extender $k(F)$ does not satisfy the initial segment condition, since $F \restriction \kappa$ is an initial segment of it which is not present in $\text{Ult}_0(\mathcal{P}^*_{\alpha+1}, E_\alpha)$. This complicates the comparison argument to follow. We advise the reader who is going through this argument for the first time to simply ignore the anomalous case in the definition of $\mathcal{P}_{\alpha+1}$.

We can lift $\mathcal{T}$ to an ordinary iteration tree on $\mathcal{M}$ as follows. Let

$$\mathcal{R}_0 = \mathcal{R}_1 = \mathcal{M},$$

and let

$$\pi_0 \colon \mathcal{P}_0 \to \mathcal{R}_0 \text{ and } \pi_1 \colon \mathcal{P}_1 \to \mathcal{R}_1$$

be given by: $\pi_0 =$ identity and $\pi_1 = \pi$. Note that π_0 and π_1 agree below $\nu(E_0)$. We can use (π_0, π_1) to lift $\mathcal{T}$ to a double-rooted tree $(\pi_0, \pi_1)\mathcal{T}$ on the pair $(\mathcal{R}_0, \mathcal{R}_1)$ just as we did in the copying construction for ordinary iteration trees. Since $\mathcal{R}_0 = \mathcal{R}_1 = \mathcal{M}$, the tree $(\pi_0, \pi_1)\mathcal{T}$, which we shall call $\mathcal{S}$, is nothing but an ordinary iteration tree on $\mathcal{M}$.[28]

We form $\mathcal{T}$ and $\mathcal{S}$ at limit stages as follows. Suppose the initial segment $\mathcal{S}^*$ of $\mathcal{S}$ built so far is a play by Σ; then we can use Σ to obtain a cofinal wellfounded branch of $\mathcal{S}^*$, and as in the ordinary copying construction, the pullback of this branch is a cofinal wellfounded branch of the initial segment $\mathcal{T}^*$ of $\mathcal{T}$ built so far. We extend $\mathcal{S}^*$ and $\mathcal{T}^*$ by choosing these branches. Thus $\mathcal{S}$ is a play by Σ, and $\mathcal{T}$ is a play by its pullback $\Sigma^{(\pi_0,\pi_1)}$.

Since Σ is an $(\omega, \omega_1, \omega_1+1)$ iteration strategy, this inductive construction of $\mathcal{S}$, $\mathcal{T}$, and $\mathcal{U}$ can last as many as $\omega_1 + 1$ steps. But $\mathcal{H}$ and $\mathcal{M}$ are countable, so as in the proof of the Comparison Lemma 3.11, the comparison

[26] This case was overlooked in [26]. It was discovered by R. Jensen. Our method of dealing with it is due to R. Schindler and M. Zeman; cf. [37].

[27] This problem cannot occur in the construction of an ordinary iteration tree, as we verified in the course of describing the successor steps in an iteration game.

[28] We are ignoring here some complications in the anomalous case.

represented by $\mathcal{T}$ and $\mathcal{U}$ actually terminates successfully at some countable stage. Let $\mathcal{P}$ and $\mathcal{Q}$ be the last models of $\mathcal{T}$ and $\mathcal{U}$ respectively. Let $\mathcal{R}$ be the last model of $\mathcal{S}$, and $\tau\colon \mathcal{P} \to \mathcal{R}$ the copy map. The key claim is:

Claim. $\mathcal{P}$ is above $\mathcal{H}$ in $\mathcal{T}$.

Proof. If not, then $\mathcal{P}$ is above $\mathcal{M}$ in $\mathcal{T}$. Suppose that the branch $\mathcal{M}$-to-$\mathcal{Q}$ of $\mathcal{U}$ drops in model or degree. Since $\mathcal{T}$ and $\mathcal{U}$ are ω-maximal trees on ω-sound mice, we then have that $\mathcal{P}$ is a proper initial segment of $\mathcal{Q}$, and the branch $\mathcal{M}$-to-$\mathcal{P}$ of $\mathcal{T}$ does not drop in model or degree, so that there is a fully elementary iteration map $i\colon \mathcal{M} \to \mathcal{P}$. But then i maps $\mathcal{M}$ to a proper initial segment of a Σ-iterate of $\mathcal{M}$, which contradicts the weak Dodd-Jensen property of Σ. Thus $\mathcal{M}$-to-$\mathcal{Q}$ does not drop, and we have a fully elementary iteration map $j\colon \mathcal{M} \to \mathcal{Q}$ given by $\mathcal{U}$.

Suppose that the branch $\mathcal{M}$-to-$\mathcal{P}$ of $\mathcal{T}$ drops in model or degree. In this case $\mathcal{Q}$ must be a proper initial segment of $\mathcal{P}$. But then $\tau \circ j$ is a fully elementary map from $\mathcal{M}$ to a proper initial segment of $\mathcal{R}$, which is a Σ-iterate of $\mathcal{M}$. This contradicts the weak Dodd-Jensen property of Σ. Thus $\mathcal{M}$-to-$\mathcal{P}$ does not drop, and we have a fully elementary iteration map $i\colon \mathcal{M} \to \mathcal{P}$ given by $\mathcal{T}$.

These arguments also show that $\mathcal{P}$ is not a proper initial segment of $\mathcal{Q}$ and $\mathcal{Q}$ is not a proper initial segment of $\mathcal{P}$, so that $\mathcal{P} = \mathcal{Q}$. We claim that $i = j$ as well. For let x be first in the enumeration $\vec{e}$ of $\mathcal{M}$ such that $i(x) \neq j(x)$. If $i(x) <_L j(x)$, then j is an iteration map produced by Σ which is not $\vec{e}$-minimal, contrary to the weak Dodd-Jensen property of Σ. So $j(x) <_L i(x)$. But now, since $\mathcal{M}$-to-$\mathcal{P}$ did not drop in $\mathcal{T}$, the branch $\mathcal{M}$-to-$\mathcal{R}$ does not drop in the copied tree $\mathcal{S}$, and so we have an iteration map $k\colon \mathcal{M} \to \mathcal{R}$ given by $\mathcal{S}$. The copy maps commute with the tree embeddings, so we have $\tau \circ i = k \circ \pi_0 = k$. But then

$$\tau(j(x)) <_L \tau(i(x)) = k(x),$$

and $\tau \circ j$ witnesses that k is not $\vec{e}$-minimal, contrary to the fact that k is an iteration map produced by Σ. Thus $i = j$.

As in the proofs of 3.11 and 4.11, this implies that the first extenders used along the branches giving rise to i and j are compatible with each other. If these extenders satisfy the initial segment condition, then as in 3.11 and 4.11, that is a contradiction because they participated in disagreements when they were used.

We are left with the possibility that the first extender G used in i comes from our anomalous case. Here $G = k(F)$, where $k\colon \mathcal{J}_\rho^{\mathcal{M}} \to \mathrm{Ult}_0(\mathcal{J}_\rho^{\mathcal{M}}, E_\alpha)$ is the canonical embedding, and F is the last extender of $\mathcal{J}_\rho^{\mathcal{M}}$. We also have $\mathrm{crit}(k) = \nu(F)$, so that $F \restriction \nu(F)$ is an initial segment of G. It is in fact the first initial segment of G which is not in $\mathcal{P}$, and since it is compatible with the first extender used in j (which itself satisfies the initial segment condition), the trivial completion of $F \restriction \nu(F)$ *is* the first extender used in j. One can now show that the second extender used in j is compatible with E_α, and that is a contradiction because both of these extenders satisfy the initial segment condition. To prove the compatibility, one uses that for $A \subseteq \mathrm{crit}(G)$, $i_G(A) = k(i_F(A))$. The reader can find the remaining details in [37]. $\dashv$

So $\mathcal{P}$ is above $\mathcal{H}$ in $\mathcal{T}$. The branch $\mathcal{H}$-to-$\mathcal{P}$ cannot drop in model or degree, since otherwise $\mathcal{Q}$ is a proper initial segment of $\mathcal{P}$ and we have a fully elementary iteration map $j\colon \mathcal{M} \to \mathcal{Q}$, so that $\tau \circ j$ maps $\mathcal{M}$ into a proper initial segment of the Σ-iterate $\mathcal{R}$. Thus we have a fully elementary iteration map $i\colon \mathcal{H} \to \mathcal{P}$ given by $\mathcal{T}$. If i is not the identity, then the rules for $\mathcal{T}$ guarantee $\mathrm{crit}(i) \geq \rho$, so that $\mathcal{H}$-to-$\mathcal{P}$ would have to drop in model or degree at its first step. Therefore i is the identity; that is, $\mathcal{H} = \mathcal{P}$.

$\mathcal{Q}$ cannot be a proper initial segment of $\mathcal{H}$, for otherwise $\mathcal{M}$-to-$\mathcal{Q}$ does not drop, and letting j be the iteration map, $\tau \circ j$ maps $\mathcal{M}$ to a proper initial segment of itself. It cannot be that $\mathcal{H} = \mathcal{Q}$, for if so, then $\mathcal{M}$-to-$\mathcal{Q}$ does not drop, and letting j be the iteration map, $\rho_\omega^{\mathcal{H}} < \rho_\omega^{\mathcal{M}} \leq j(\rho_\omega^{\mathcal{M}}) = \rho_\omega^{\mathcal{Q}}$. Thus $\mathcal{H}$ is a proper initial segment of $\mathcal{Q}$.

We can now complete the proof of 5.1. Suppose that $\mathcal{H}$ is not an initial segment of $\mathcal{M}$, so that $\mathcal{U}$ uses at least one extender $E_0^{\mathcal{U}}$. Now $\rho \leq \mathrm{lh}(E_0^{\mathcal{U}})$ because $\mathcal{H}$ and $\mathcal{M}$ agree below ρ, while $\mathrm{lh}(E_0^{\mathcal{U}}) \leq \mathrm{On}^{\mathcal{H}}$ because $\mathcal{H}$ is not an initial segment of $\mathcal{M}$. But $\mathrm{lh}(E_0^{\mathcal{U}})$ is a cardinal of $\mathcal{Q}$, and $\mathcal{H}$ is a proper initial segment of $\mathcal{Q}$, so that $|\mathrm{On}^{\mathcal{H}}| \leq \rho_\omega^{\mathcal{H}}$ in $\mathcal{Q}$. It follows that $\mathrm{lh}(E_0^{\mathcal{U}}) = \rho$. Similarly, if $E_1^{\mathcal{U}}$ exists, then we must have $\mathrm{On}^{\mathcal{H}} < \mathrm{lh}(E_1^{\mathcal{U}})$, so in fact $E_1^{\mathcal{U}}$ does not exist. This means that $\mathcal{Q} = \mathrm{Ult}_k(\mathcal{M}, E_0^{\mathcal{U}})$ for some k. We can take $k = 0$ because $\mathrm{Ult}_0(\mathcal{M}, E_0^{\mathcal{U}})$ and $\mathrm{Ult}_k(\mathcal{M}, E_0^{\mathcal{U}})$ agree to their common value for ρ^+ and beyond. $\dashv$

One can prove a version of 5.1 in which $\rho_\omega^{\mathcal{H}}$ is replaced by $\rho_n^{\mathcal{H}}$, for some $n < \omega$. See [26, section 8].

The technique by which 5.1 is proved is useful in many circumstances. One wants to compare two mice $\mathcal{H}$ and $\mathcal{M}$ in such a way that the iteration map on the $\mathcal{H}$ side has critical point at least ρ. An ordinary comparison might not have this property, but one finds models (such as $\mathcal{M}$ itself in the proof above) which agree with $\mathcal{H}$ to various extents below ρ, yet in some sense carry more information than $\mathcal{H}$. One then forms a many-rooted iteration tree on $\mathcal{H}$ "backed up" by these other models, and argues that the final model on this tree lies above the root $\mathcal{H}$. One can view the proof of 4.11 in this light.[29] Another important application of the technique lies in

[29] In 4.11 one wanted to compare the last models of $\mathcal{U}^*$ and $\mathcal{V}^*$, but for the proof it was important to back them up with the earlier models of $\mathcal{T}$. Many-rooted iteration trees are also important in the inductive definition of K ([44, section 6]), and in the proof of weak covering for K ([25]).

the proof of the following central fine-structural result concerning the good behavior of the standard parameter.

5.3 Theorem. *Let $k < \omega$, and let $\mathcal{M}$ be a k-sound, $(k, \omega_1, \omega_1 + 1)$-iterable premouse; then $\mathcal{C}_{k+1}(\mathcal{M})$ exists, and agrees with $\mathcal{M}$ below γ, for all γ of $\mathcal{M}$-cardinality $\rho_{k+1}(\mathcal{M})$.*

Sketch of proof. We assume $k = 0$ for notational simplicity, and because only in that case have we given full definitions anyway. Let $r = p_1(\mathcal{M})$ be the first standard parameter of $\mathcal{M}$; we must show that r is 1-solid and 1-universal, so that $\mathcal{C}_1(\mathcal{M})$ exists, and that $\mathcal{C}_1(\mathcal{M})$ agrees with $\mathcal{M}$ as claimed. These properties of r and $\mathcal{M}$ are expressed by sentences in the first order theory of $\mathcal{M}$, so if they fail, they fail in some countable fully elementary submodel of $\mathcal{M}$. Any countable elementary submodel of $\mathcal{M}$ inherits its $(0, \omega_1, \omega_1 + 1)$-iterability. Thus we may assume without loss of generality that $\mathcal{M}$ is countable.

We shall assume that r is solid, and briefly sketch the proof that r is universal and $\mathcal{C}_1(\mathcal{M})$ agrees with $\mathcal{M}$ below the cardinal successor of $\rho_1(\mathcal{M})$ in $\mathcal{M}$. So let $\rho = \rho_1(\mathcal{M})$, and let

$$\mathcal{H} = \mathcal{H}_1^{\mathcal{M}}(\rho \cup \{r\}).$$

We wish to show that $P(\rho) \cap \mathcal{M} \subseteq \mathcal{H}$, and for this the natural strategy is to compare $\mathcal{H}$ with $\mathcal{M}$. If the critical point of the embedding i from $\mathcal{H}$ to the last model $\mathcal{P}$ on the $\mathcal{H}$ side is at least ρ, then the $\Sigma_1^{\mathcal{H}}$ set $A \subseteq \rho$ which is not in $\mathcal{M}$ (witnessing that $\rho = \rho_1^{\mathcal{M}}$) is also $\Sigma_1^{\mathcal{P}}$. Since A is not in the last model $\mathcal{Q}$ on the $\mathcal{M}$ side, $\mathcal{Q}$ is an initial segment of $\mathcal{P}$, and one can then argue that

$$P(\rho)^{\mathcal{M}} = P(\rho)^{\mathcal{Q}} \subseteq P(\rho)^{\mathcal{P}} = P(\rho)^{\mathcal{H}},$$

as desired. In order to insure that $\text{crit}(i) \geq \rho$, we once again form a double-rooted tree on the pair $(\mathcal{M}, \mathcal{H})$ on the $\mathcal{H}$ side of our comparison, going back to $\mathcal{M}$ whenever we use an extender with critical point $< \rho$.

Let $r = \langle \alpha_0, \ldots, \alpha_n \rangle$, where the ordinals α_i are listed in decreasing order. Let $\vec{e}$ be an enumeration of the universe of $\mathcal{M}$ such that $e_i = \alpha_i$ for all $i \leq n$. Let Σ be a $(0, \omega_1, \omega_1 + 1)$ iteration strategy for $\mathcal{M}$ having the weak Dodd-Jensen property relative to $\vec{e}$. Let $\pi_0 =$ identity and $\pi_1 \colon \mathcal{H} \to \mathcal{M}$ be the collapse embedding. We form the double-rooted tree $\mathcal{T}$ on $(\mathcal{M}, \mathcal{H})$ using the pullback $\Sigma^{(\pi_0, \pi_1)}$ of Σ to choose branches at limit stages, and iterating the least disagreement with the last model of the tree $\mathcal{U}$ on $\mathcal{M}$ at successor stages. Let $\mathcal{P}$ and $\mathcal{Q}$ be the last models of $\mathcal{T}$ and $\mathcal{U}$.

As in the proof of 5.1, the weak Dodd-Jensen property of Σ implies that $\mathcal{P}$ is above $\mathcal{H}$, and not above $\mathcal{M}$, and that $\mathcal{H}$-to-$\mathcal{P}$ does not drop, and that $\mathcal{Q}$ is not a proper initial segment of $\mathcal{P}$. Thus we have a 0-embedding $i \colon \mathcal{H} \to \mathcal{P}$ given by $\mathcal{T}$. Since $\text{crit}(i) \geq \rho$, A is $\Sigma_1^{\mathcal{P}}$, and since $A \notin \mathcal{Q}$, $\mathcal{P}$ is not a proper initial segment of $\mathcal{Q}$. Thus $\mathcal{P} = \mathcal{Q}$. We also get that $\mathcal{M}$-to-$\mathcal{Q}$ does not drop, so that $\mathcal{U}$ gives us an embedding $j\colon \mathcal{M} \to \mathcal{Q}$.

Let $\bar{\alpha}_e = \pi_1^{-1}(\alpha_e)$ be the image of α_e under collapse, for $e \le n$. One can show by induction on e that

$$i(\bar{\alpha}_e) = j(\alpha_e),$$

using the solidity of $j(r)$ to show $i(\bar{\alpha}_e) \ge j(\alpha_e)$, and using the weak Dodd-Jensen property for the copied tree $(\pi_0), \pi_1)\mathcal{T}$ to show $i(\bar{\alpha}_e) \le j(\alpha_e)$. (This is where we use the fact that $e_i = \alpha_i$ for all $i \le n$.)

It follows that $\mathrm{crit}(j) \ge \rho$. For otherwise, letting $\kappa = \mathrm{crit}(j)$, and S be the Σ_1 theory in $\mathcal{M}$ of parameters from $\kappa \cup \{r\}$, then $S \in \mathcal{M}$. But then $j(S) \in \mathcal{Q}$, and from $j(S)$ one can compute the Σ_1 theory in $\mathcal{Q}$ of parameters from $j(\kappa) \cup \{j(r)\}$. (This is like the proof of 2.23 which we hinted at earlier.) Now $\rho < j(\kappa)$, $\mathcal{P} = \mathcal{Q}$, and $i(\bar{r}) = j(r)$, so this means the Σ_1 theory of $\rho \cup \{i(\bar{r})\}$ is in $\mathcal{P}$. This implies $A \in \mathcal{P}$, a contradiction.

Since i and j have critical point above ρ, $P(\rho)^{\mathcal{H}} = P(\rho)^{\mathcal{P}} = P(\rho)^{\mathcal{Q}} = P(\rho)^{\mathcal{M}}$, as desired. Also, $\mathcal{H} = \mathcal{C}_{k+1}(\mathcal{M})$ agrees with $\mathcal{P}$, hence $\mathcal{Q}$, hence $\mathcal{M}$, below any γ of $\mathcal{M}$-cardinality ρ, as desired. ⊣

One can use fine-structural condensation results such as 5.1 to show that iterable mice satisfy many of the useful combinatorial principles which Jensen has shown are true in L. For example

5.4 Theorem. *Let $\mathcal{M}$ be an $(\omega, \omega_1, \omega_1 + 1)$-iterable premouse satisfying the axioms of* ZF*, except perhaps Powerset; then the following are true in $\mathcal{M}$:*

1. *for all uncountable regular κ, $\Diamond_\kappa$,*
2. *for all uncountable regular κ, ($\Diamond_\kappa^+ \Leftrightarrow \kappa$ is not ineffable.)*
3. *for all infinite cardinals κ, $\Box_\kappa$.*

Part (1) of 5.4 follows immediately from 5.1 and Jensen's argument for L. Part (2) is due to E. Schimmerling ([32]). Part (3) is work of Schimmerling and M. Zeman ([36]), building on the earlier work of Jensen, Solovay, Welch, Wylie, and Schimmerling. (See [11], [47], [48], [32], and [33].)

It follows immediately from 5.3 that if $\mathcal{M}$ is sufficiently iterable, then $\mathcal{C}_\omega(\mathcal{M})$ exists. We shall use this heavily in the construction of an iterable model, all of whose levels are ω-sound. We turn to that construction now.

6. Background-Certified Fine Extender Sequences

We have been studying mice in the abstract, but we have yet to produce any! In this section we shall describe a certain family of mouse constructions which we call, for obscure reasons, K^c-*constructions*. Such constructions are sufficiently cautious about adding extenders to the model that one gets an iterable model in the end,[30] yet can be sufficiently daring that they can capture the large cardinal strength present in the universe.[31]

6.1. K^c constructions

The natural idea is to construct a fine extender sequence $\vec{E}$ by induction. Given $\vec{E} \restriction \alpha$, we set $E_\alpha = \emptyset$ unless there is a certified[32] extender F such that $(\vec{E} \restriction \alpha)^\frown F$ is still a fine extender sequence; if there is such an F we may either set $E_\alpha = F$ or set $E_\alpha = \emptyset$. Here "certified" means roughly that F is the restriction to $J_\alpha^{\vec{E}\restriction\alpha}$ of a "background extender" F^* which measures a broader collection of subsets of its critical point than does F, and whose ultrapower agrees with V a bit past $\nu(F)$. This background-certificate demand is necessary in order to insure that the premice we are constructing are iterable. Unfortunately, the background certificate demand conflicts with the demand that all levels of the model we are constructing be ω-sound.[33] K^c constructions deal with this conflict by continually replacing the premouse $\mathcal{N}_\alpha$ currently approximating the model being built by its core $\mathcal{C}_\omega(\mathcal{N}_\alpha)$. Taking cores insures soundness, while the background extenders one can resurrect by going back into the history of the construction insure iterability.

This last claim must be qualified. We do not have a general proof of iterability for the premice $\mathcal{N}_\alpha$ produced in K^c constructions. At the moment, in order to prove that such a premouse is appropriately iterable, we need to make an additional "smallness" assumption. One assumption that suffices, and which we shall spell out in more detail shortly, is that no initial segment of $\mathcal{N}_\alpha$ satisfies "there is an extender E on my sequence such that $\nu(E)$ is a Woodin cardinal". We shall call this property of $\mathcal{N}_\alpha$ *tameness*. Iterability is essential from the very beginning, for our proof that $\mathcal{C}_\omega(\mathcal{N}_\alpha)$ exists involves comparison arguments, and hence relies on the iterability of $\mathcal{N}_\alpha$. Thus, for all we know, a K^c construction might simply break down by reaching a non-tame premouse $\mathcal{N}_\alpha$ such that $\mathcal{C}_\omega(\mathcal{N}_\alpha)$ does not exist.

The following definitions describe our background certificate condition. They come from [44, section 1].

6.1 Definition. Let $\mathcal{M}$ be an active premouse, F the extender coded by $\dot{F}^{\mathcal{M}}$ (i.e. its last extender), $\kappa = \mathrm{crit}(F)$, and $\nu = \nu(F)$. Let $\mathcal{A} \subseteq \bigcup_{n<\omega} P([\kappa]^n)^{\mathcal{M}}$; then an *$\mathcal{A}$-certificate for $\mathcal{M}$* is a pair (N, G) such that

[30] This is something between a conjecture and a theorem; see below.

[31] Again, there are qualifications to come.

[32] Whence the "c" in K^c.

[33] Part of the requirement on F^* is that it be countably complete, and so $\mathrm{crit}(F^*)$ must be uncountable; on the other hand, if α is least so that $E_\alpha \neq \emptyset$, then $(J_\alpha^{\vec{E}\restriction\alpha}, \in, \vec{E} \restriction \alpha, E_\alpha)$ has Σ_1 projectum ω, so that $\mathrm{crit}(E_\alpha)$ must be countable if this structure is even 1-sound.

1. N is a transitive, power admissible set, $V_\kappa \cup \mathcal{A} \subseteq N$, N is closed under ω-sequences, and G is an extender over N,

2. $F \cap ([\nu]^{<\omega} \times \mathcal{A}) = G \cap ([\nu]^{<\omega} \times \mathcal{A})$,

3. $V_{\nu+1} \subseteq \text{Ult}(N, G)$, and

4. $\forall\gamma(\omega\gamma < \text{On}^{\mathcal{M}} \Rightarrow \mathcal{J}_\gamma^{\mathcal{M}} = \mathcal{J}_\gamma^{i(\mathcal{J}_\kappa^{\mathcal{M}})})$, where $i = i_G^N$ is the canonical embedding from N to $\text{Ult}(N, G)$.

6.2 Definition. Let $\mathcal{M}$ be an active premouse, and κ the critical point of its last extender. We say $\mathcal{M}$ *is countably certified* iff for every countable $\mathcal{A} \subseteq \bigcup_{n<\omega} P([\kappa]^n)^{\mathcal{M}}$, there is an $\mathcal{A}$-certificate for $\mathcal{M}$.

In the situation described in definition 6.1, we shall typically have $|N| = \kappa$, so that $\text{On}^N < \text{lh}(G)$. We are therefore not thinking of (N, G) as a structure to be iterated; N simply provides a reasonably large collection of sets to be measured by G. The conditions $V_\kappa \subseteq N$ and $V_{\nu+1} \subseteq \text{Ult}(N, G)$ are crucial (although the former can be weakened in a useful way; cf. [35, 2.1]). Power admissibility is simply a convenient fragment of ZFC; it can probably be weakened substantially.

6.3 Definition. A *K^c-construction* is a sequence $\langle \mathcal{N}_\alpha \mid \alpha < \theta \rangle$ of premice such that

1. $\mathcal{N}_0 = (V_\omega, \in, \emptyset, \emptyset)$;

2. if $\alpha + 1 < \theta$, then $\mathcal{N}_\alpha$ is ω-solid, and letting $\mathcal{M}$ be the unique ω-sound premouse such that $\mathcal{C}_\omega(\mathcal{N}_\alpha) = \mathcal{C}_\omega(\mathcal{M})$, either

   (a) $\mathcal{M}$ is passive, and $\mathcal{N}_{\alpha+1}$ is a countably certified premouse of the form $(|\mathcal{M}|, \in, \dot{E}^{\mathcal{M}}, F)$, for some F, or

   (b) letting $\omega\gamma = \text{On}^{\mathcal{M}}$ and $\vec{E} = \dot{E}^{\mathcal{M}} \oplus \dot{F}^{\mathcal{M}}$, we have that $\mathcal{N}_{\alpha+1} = (J_{\gamma+1}^{\vec{E}}, \in, \vec{E}, \emptyset)$;

3. if $\lambda < \theta$ is a limit ordinal, then $\mathcal{N}_\lambda$ is the unique passive premouse $\mathcal{P}$ such that for all β, $\omega\beta < \text{On}^{\mathcal{P}}$ iff $\mathcal{J}_\beta^{\mathcal{N}_\alpha}$ is defined and eventually constant as $\alpha \to \lambda$, and for all β such that $\omega\beta < \text{On}^{\mathcal{P}}$, $\mathcal{J}_\beta^{\mathcal{P}} =$ eventual value of $\mathcal{J}_\beta^{\mathcal{N}_\alpha}$, as $\alpha \to \lambda$.

So at successor steps in a K^c-construction one replaces the previous model with its ω^{th} core, and then either adds a countably certified extender to the resulting extender sequence or takes one step in its constructible closure. At limit steps one forms the natural "lim inf" of the previous premice.

Because we replace $\mathcal{N}_\alpha$ by its core at each step in a K^c–construction, the models of the construction may not grow by end–extension, and we need a little argument to show, for example, that a construction of proper class length converges to a premouse of proper class size. Our Theorem 5.3 on the agreement of $\mathcal{N}$ with $\mathcal{C}_\omega(\mathcal{N})$ is the key here.

6.4 Theorem. *Let κ be an uncountable regular cardinal or $\kappa = \mathrm{On}$, and let $\langle \mathcal{N}_\alpha \mid \alpha < \kappa \rangle$ be a K^c–construction; then there is a unique premouse $\mathcal{N}_\kappa$ of ordinal height κ such that $\langle \mathcal{N}_\alpha \mid \alpha \leq \kappa \rangle$ is a K^c–construction.*

Proof. For any limit ordinal κ and K^c–construction $\langle \mathcal{N}_\alpha \mid \alpha < \kappa \rangle$, there is a unique premouse $\mathcal{N}_\kappa$ satisfying the limit ordinal clause of Definition 6.3. We need only show that $\mathcal{N}_\kappa$ has ordinal height κ in the case κ is an uncountable cardinal or $\kappa = \mathrm{On}$. It is clear that $|\mathcal{N}_\alpha| < \kappa$ for all $\alpha < \kappa$, so $\mathcal{N}_\kappa$ has ordinal height $\leq \kappa$.

For $\nu < \kappa$, let
$$\vartheta_\nu = \inf\{\rho_\omega(\mathcal{N}_\alpha) \mid \nu \leq \alpha < \kappa\}.$$
So $\vartheta_0 = \omega$, and the ϑ's are nondecreasing. By Theorem 5.3, $\mathcal{N}_\nu$ agrees with all later $\mathcal{N}_\alpha$ below ϑ_ν, so if $\kappa = \sup(\{\vartheta_\nu \mid \nu < \kappa\})$, we are done. Since κ is regular, the alternative is that the ϑ's are eventually constant; say $\vartheta_\nu = \rho$ for all ν such that $\eta \leq \nu < \kappa$. Now notice that if $\eta \leq \nu < \kappa$ and $\rho_\omega(\mathcal{N}_\nu) = \rho$, then $\mathcal{C}_\omega(\mathcal{N}_\nu)$ is a proper initial segment of $\mathcal{N}_{\nu+1}$.[34] Moreover, $\mathcal{C}_\omega(\mathcal{N}_\nu)$ has cardinality ρ in $\mathcal{N}_{\nu+1}$ by soundness. It follows from Theorem 5.3 that $\mathcal{C}_\omega(\mathcal{N}_\nu)$ is an initial segment of $\mathcal{N}_\alpha$, for all $\alpha \geq \nu$. Since there are cofinally many $\nu < \kappa$ such that $\rho = \rho_\omega(\mathcal{N}_\nu)$, we again get that $\mathcal{N}_\kappa$ has height κ. ⊣

It is not hard to see that the ϑ_ν defined in the proof above are just the infinite cardinals of $\mathcal{N}_\kappa$.

6.2. The iterability of K^c

It is clear by now that we have gotten nowhere unless we can prove that the premice we have constructed are sufficiently iterable. Here we encounter the central open problem of inner model theory. We formulate one instance of it as a conjecture:

6.5 Conjecture. Suppose $\mathcal{N}$ is a premouse occurring in a K^c construction, that $k \leq \omega$, and that $\mathcal{M}$ is a countable premouse such that there is a weak k-embedding from $\mathcal{M}$ into $\mathcal{C}_k(\mathcal{N})$; then $\mathcal{M}$ is $(k, \omega_1, \omega_1 + 1)$-iterable.

A proof of this conjecture would yield at once the basics of inner model theory at the level of models with superstrong cardinals.[35] At present we can prove the conjecture only for certain small mice.

[34] Assume the last extender predicate of $\mathcal{N}_\nu$ is empty here, as it obviously is for cofinally many such ν.

[35] New problems arise between superstrong and supercompact cardinals.

In general, iterability proofs break up into an *existence* proof and a *uniqueness* proof for "sufficiently good" branches in iteration trees on the premice under consideration. The existence proof itself breaks into two parts, a direct existence argument in the countable case and a reflection argument in the uncountable case.

The direct existence argument applies to countable iteration trees on countable elementary submodels of the premice under consideration, and proceeds by using something like the countable completeness of the extenders involved in the iteration to transform an ill-behaved iteration into an infinite descending $\in$-chain. When coupled with the uniqueness proof, this shows that any countable elementary submodel of a premouse under consideration has an ω_1-iteration strategy, namely, the strategy of choosing the unique cofinal "sufficiently good" branch.[36]

The reflection argument extends this method of iterating to the uncountable: given an iteration tree $\mathcal{T}$ on $\mathcal{M}$, we go to $V[G]$ where G is $\mathrm{Col}(\omega, \kappa)$-generic over V and κ is large enough that $\mathcal{M}$ and $\mathcal{T}$ have become countable, and find a sufficiently good branch there. This branch is unique, and hence by the homogeneity of the collapse it is in V. In order to execute this argument one needs a certain level of absoluteness between V and $V[G]$. Once one gets past mice with Woodin cardinals, "sufficiently good" can no longer be taken simply to mean "wellfounded", and in fact "sufficiently good" is no longer a Σ^1_2 notion at all. Because of this, the generic absoluteness required by our reflection argument needs large cardinal/mouse existence principles that go beyond ZFC.[37]

The conjecture above overlaps slightly with the uncountable case because it is $(\omega_1 + 1)$-iterability, rather than ω_1-iterability, which is at stake. One needs $(\omega_1 + 1)$-iterability to guarantee the comparability of countable mice; the reflection argument that shows coiterations terminate requires a wellfounded branch of length ω_1. Nevertheless, we believe that the conjecture is provable in ZFC.[38]

At present, the strongest partial results on conjecture 6.5 are those of [1], which show that it holds for levels $\mathcal{N}$ of K^c which are of limited complexity, in that they do not have too many extenders overlapping local Woodin cardinals. In this paper we shall consider only premice having no extenders overlapping local Woodin cardinals. We call these special premice "tame". We shall outline a proof of 6.5 for the tame levels of K^c. Our direct existence

[36] Of course a sufficiently good branch must be wellfounded, but in general more is required, for we want to be able to find cofinal wellfounded branches later in the iteration game as well.

[37] For example, if it is consistent that there is a Woodin cardinal, then it is consistent that there is a premouse $\mathcal{N}$ occurring on a K^c-construction which is not θ-iterable for some θ.

[38] We suspect that if κ is strictly less than the infimum of the critical points of the background extenders, then the κ-iterability of the size κ elementary submodels of premice in a K^c-construction is provable in ZFC.

argument in the countable case seems perfectly general, but our uniqueness results are less definitive, and it is here that we resort to the tameness assumption. We begin by stating the existence theorem in the countable case.

We say that a branch b of an iteration tree $\mathcal{T}$ is *maximal* iff b has limit order type but is not continued in $\mathcal{T}$. Such a b must be $\in$-cofinal in some $\lambda \leq \mathrm{lh}(\mathcal{T})$, but different from $[0, \lambda]_T$ if $\lambda < \mathrm{lh}(\mathcal{T})$. Notice that any cofinal branch of $\mathcal{T}$ is maximal; the converse fails in general. Finally, a *putative* iteration tree is just like an ordinary iteration tree, except that we allow the last model, if there is one, to be illfounded.

6.6 Theorem (Branch Existence Theorem). *Let $\pi\colon \mathcal{M} \to \mathcal{C}_k(\mathcal{N}_\alpha)$ be a weak k-embedding, where $\mathcal{M}$ is countable and $\langle \mathcal{N}_\beta \mid \beta < \theta \rangle$ is a K^c construction. Let $\mathcal{T}$ be a countable, k-maximal, putative iteration tree on $\mathcal{M}$; then either*

1. *there is a maximal branch b of $\mathcal{T}$ such that, letting $l = \deg^{\mathcal{T}}(b)$,*

   (a) *$D^{\mathcal{T}} \cap b = \emptyset$, and there is a weak l-embedding $\sigma\colon \mathcal{M}_b^{\mathcal{T}} \to \mathcal{C}_l(\mathcal{N}_\alpha)$ such that*

   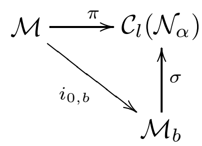

   commutes, or

   (b) *$D^{\mathcal{T}} \cap b \neq \emptyset$, and there is a $\beta < \alpha$ and weak l-embedding $\sigma\colon \mathcal{M}_b^{\mathcal{T}} \to \mathcal{C}_l(\mathcal{N}_\beta)$, or*

2. *$\mathcal{T}$ has a last model $\mathcal{M}_\gamma^{\mathcal{T}}$ such that, letting $l = \deg^{\mathcal{T}}(\gamma)$,*

   (a) *$D^{\mathcal{T}} \cap [0, \gamma]_T = \emptyset$, and there is a weak l-embedding $\sigma\colon \mathcal{M}_\gamma^{\mathcal{T}} \to \mathcal{C}_l(\mathcal{N}_\alpha)$ such that*

   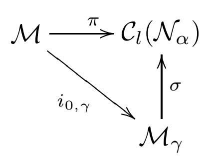

   commutes,or

   (b) *$D^{\mathcal{T}} \cap [0, \gamma]_T \neq \emptyset$, and there is a $\beta < \alpha$ and weak l-embedding $\sigma\colon \mathcal{M}_\gamma^{\mathcal{T}} \to \mathcal{C}_l(\mathcal{N}_\beta)$.*

We shall not attempt to prove this theorem here. The reader can find a proof in [44, sections 2 and 9]. The theorem in the form stated here evolved from earlier results of [18] and [26].

If b is a branch satisfying clause (1) of the conclusion of the Branch Existence Theorem, then we say b (or $\mathcal{M}_b^{\mathcal{T}}$) is *$\pi$-realizable*, and call the map σ described in clause (1) a *π-realization* of b (or $\mathcal{M}_b^{\mathcal{T}}$). Similarly, if γ satisfies clause (2) of the conclusion, then we say γ (or $\mathcal{M}_\gamma^{\mathcal{T}}$) is π-realizable, and call the associated map σ a π-realization.

Given $\mathcal{M}$ and π as in the hypotheses of the Branch Existence Theorem, it is natural to attempt to iterate $\mathcal{M}$ using the following strategy: given $\mathcal{T}$ on $\mathcal{M}$ of countable limit length, pick the unique cofinal π-realizable branch of $\mathcal{T}$ with which to continue. Clause (2) in the conclusion of the Branch Existence Theorem guarantees that this strategy cannot break down at any countable successor stage. Clause (1) guarantees that if this strategy breaks down at some countable limit stage, then there are distinct cofinal π-realizable branches at that stage, since the uniqueness of the branches chosen at earlier stages implies that any maximal π-realizable branch of $\mathcal{T}$ must be cofinal. However, if we ever reach a stage at which our tree has distinct cofinal π-realizable branches (this is possible for some $\mathcal{M}$ and π; see [18, section 5]), our troubles start. The best we can do, it seems, is to choose one such branch b and a π-realization σ of $\mathcal{M}_b^{\mathcal{T}}$. If our opponent in the iteration game is kind enough to continue by playing extenders which can be interpreted as forming a tree on $\mathcal{M}_b^{\mathcal{T}}$, then we can choose unique σ-realizable branches to continue, until we get distinct such branches and must pick one, realize it, and continue, etc. However, we are done for if our opponent applies an extender to a model from $\mathcal{T}$ (that is, a model with index $< \sup(b)$). Nothing in the Branch Existence Theorem even guarantees that the associated ultrapower will be wellfounded.[39]

Clearly, we need a uniqueness theorem to accompany our existence theorem. What we can show, roughly speaking, is that at a non-uniqueness stage in the process just described we pass a local Woodin cardinal.

6.7 Definition. Let $\kappa < \delta$ and $A \subseteq V_\delta$; then we say κ is *A-reflecting in* δ iff for all $\nu < \delta$ there is an extender E over V such that $\mathrm{crit}(E) = \kappa$, $i_E(\kappa) > \nu$, and $i_E(A) \cap V_\nu = A \cap V_\nu$.

6.8 Definition. A cardinal δ is a *Woodin cardinal* iff for all $A \subseteq \delta$ there is a $\kappa < \delta$ which is A-reflecting in δ.

It is perhaps no surprise to the reader that Woodin cardinals were discovered by W.H. Woodin. Woodin was inspired by the results of [9], and by earlier work of S. Shelah reducing the large cardinal hypotheses employed

[39] We have described here how the Branch Existence Theorem yields a winning strategy for II in a game that requires less of him, the *weak iteration game.* We shall introduce this game formally in the next section.

in [9]. The definition of Woodinness given above is different from Woodin's original one, but equivalent to it by an argument essentially due to Mitchell. (See [21, Theorem 4.1].) Mitchell's argument can also be used to show that if δ is Woodin, then δ is witnessed to be Woodin by extenders in V_δ.[40] It follows that the Woodinness of δ can be expressed by a Π_1 sentence about $(V_{\delta+1}, \in)$, so that the least Woodin cardinal is not weakly compact. It is easy to see that all Woodin cardinals are Mahlo.

The (local) Woodin cardinal we get from an iteration tree $\mathcal{T}$ having distinct good branches is the supremum of the lengths of the extenders used in $\mathcal{T}$.

6.9 Definition. Let $\mathcal{T}$ be a k-maximal iteration tree on $\mathcal{M}$ such that $\mathrm{lh}(\mathcal{T})$ is a limit ordinal; then we set

$$\delta(\mathcal{T}) = \sup\{\mathrm{lh}(E_\alpha^{\mathcal{T}}) \mid \alpha < \mathrm{lh}(\mathcal{T})\},$$

and

$$\begin{aligned} \mathcal{M}(\mathcal{T}) \quad = \quad & \text{unique passive } \mathcal{P} \text{ such that } \mathrm{On}^{\mathcal{P}} = \delta(\mathcal{T}) \text{ and} \\ & \forall \alpha < \delta(\mathcal{T})(\mathcal{M}(\mathcal{T}) \text{ agrees with } \mathcal{M}_\alpha^{\mathcal{T}} \text{ below } \mathrm{lh}(E_\alpha^{\mathcal{T}})). \end{aligned}$$

It is clear that if b is a cofinal branch of $\mathcal{T}$ such that $\delta(\mathcal{T}) \in \mathcal{M}_b^{\mathcal{T}}$, then $\delta(\mathcal{T})$ is a limit cardinal of $\mathcal{M}_b^{\mathcal{T}}$.

The main result connecting Woodin cardinals with the uniqueness of cofinal wellfounded branches in iteration trees is the following theorem of [18].

6.10 Theorem (Branch Uniqueness Theorem). *Let b and c be distinct cofinal branches of the k-maximal iteration tree $\mathcal{T}$, let $\delta = \delta(\mathcal{T})$, and suppose $A \subseteq \delta$ is such that $\delta, A \in \mathrm{wfp}(\mathcal{M}_b^{\mathcal{T}}) \cap \mathrm{wfp}(\mathcal{M}_c^{\mathcal{T}})$; then*

$$\mathcal{M}_b^{\mathcal{T}} \models \exists \kappa < \delta(\kappa \textit{ is } A\textit{-reflecting in } \delta).$$

Sketch of Proof. The extenders used on b and c have an overlapping pattern pictured in Figure I.1:

To see this, pick any successor ordinal

$$\alpha_0 + 1 \in b \setminus c,$$

and then let

$$\beta_n + 1 = \min\{\gamma \in c : \gamma > \alpha_n + 1\}$$

and

$$\alpha_{n+1} + 1 = \min\{\eta \in b : \eta > \beta_n + 1\},$$

[40]This observation is due to Woodin.

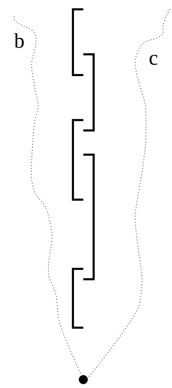

Figure I.1: The overlapping pattern of two distinct well–founded branches

for all $n < \omega$. Now for any n, the T-predecessor of $\beta_n + 1$ is on c and $\leq \alpha_n + 1$, hence $\leq \alpha_n$, so by the rules of the iteration game

$$\mathrm{crit}(F_{\beta_n}) < \nu(F_{\alpha_n}).$$

Similarly, for any n

$$\mathrm{crit}(F_{\alpha_{n+1}}) < \nu(F_{\beta_n}).$$

Now extenders used along the same branch of an iteration tree do not overlap (*i.e.*, if E is used before F, then $\nu(E) \leq \mathrm{crit}(F)$), so we have

$$\begin{aligned}\mathrm{crit}(F_{\beta_n}) < \nu(F_{\alpha_n}) &\leq \mathrm{crit}(F_{\alpha_{n+1}}) < \nu(F_{\beta_n}) \\ &\leq \mathrm{crit}(F_{\beta_{n+1}}) < \nu(F_{\alpha_{n+1}}) \leq \mathrm{crit}(F_{\alpha_{n+2}}),\end{aligned}$$

which is the overlapping pattern pictured.

Now $\sup(\{\alpha_n : n < \omega\}) = \sup(\{\beta_n : n < \omega\})$, and since branches of iteration trees are closed below their suprema in the order topology on On, the common supremum of the α_n and β_n is λ. Let us assume α_0 was chosen large enough that letting

$$\xi = \mathrm{pred}_T(\beta_0 + 1) \text{ and } \eta = \mathrm{pred}_T(\alpha_1 + 1),$$

we have

$$A = i_{\xi,c}(A^*) = i_{\eta,b}(A^{**})$$

for some A^* and A^{**}. Let

$$\kappa = \mathrm{crit}(F_{\beta_0}) = \mathrm{crit}(i_{\xi,c});$$

we shall show that κ is A-reflecting in δ in the model $\mathcal{M}_b$.

Let $E_0 = F_{\beta_0} \restriction \mathrm{crit}(F_{\alpha_1})$. Because of the overlapping pattern, E_0 is a proper initial segment of F_{β_0}, and by initial segment condition on premice and the agreement of the models of an iteration tree, $E_0 \in \mathcal{M}_b$. Moreover, if $j\colon \mathcal{M}_b \to \mathrm{Ult}(\mathcal{M}_b, E_0)$ is the canonical embedding, then because A and A^* agree below κ, $j(A)$ and $i_{\xi,c}(A^*)$ agree below $\mathrm{crit}(F_{\alpha_1})$. That is, $j(A)$ agrees with A below $\mathrm{crit}(F_{\alpha_1})$, and hence E_0 witnesses that κ is A-reflecting up to $\mathrm{crit}(F_{\alpha_1})$ in $\mathcal{M}_b$.

To get A-reflection all the way up to δ, we set

$$E_{2n} = F_{\beta_n} \restriction \mathrm{crit}(F_{\alpha_{n+1}}) \text{ and } E_{2n+1} = F_{\alpha_{n+1}} \restriction \mathrm{crit}(F_{\beta_{n+1}}),$$

for all n. Each of the E_n is in $\mathcal{M}_b$ for the same reason E_0 is in $\mathcal{M}_b$. Therefore the extender E which represents the embedding coming from "composing" the ultrapowers by the E_i for $0 \le i \le 2n$, is in $\mathcal{M}_b$. The argument above generalizes easily to show that E witnesses that κ is A-reflecting up to $\mathrm{crit}(F_{\alpha_{n+1}})$. Since $\mathrm{crit}(F_{\alpha_{n+1}}) \to \delta$ as $n \to \omega$, κ is A-reflecting in δ in the model $\mathcal{M}_b$. ⊣

We shall need a fine-structural refinement of 6.10. For this, we have to look closely at the first level of $\mathcal{M}_b^{\mathcal{T}}$ at which $\delta(\mathcal{T})$ is seen not to be Woodin, if there is one.

6.11 Definition. Let $\mathcal{T}$ be a k-maximal iteration tree on $\mathcal{M}$ of limit length, and let b be a cofinal wellfounded branch of $\mathcal{T}$. Let γ be the least ordinal, if there is one, such that either

$$\omega\gamma < \mathrm{On}^{\mathcal{M}_b} \text{ and } \mathcal{J}_{\gamma+1}^{\mathcal{M}_b} \models \delta(\mathcal{T}) \text{ is not Woodin,}$$

or

$$\omega\gamma = \mathrm{On}^{\mathcal{M}_b} \text{ and } \rho_{n+1}(\mathcal{J}_{\gamma}^{\mathcal{M}_b}) < \delta(\mathcal{T})$$

for some $n < \omega$ such that $n+1 \le k$ if $D^{\mathcal{T}} \cap b = \emptyset$. We set

$$\mathcal{Q}(b, \mathcal{T}) := \mathcal{J}_{\gamma}^{\mathcal{M}_b}$$

if there is such a γ, and let $\mathcal{Q}(b, \mathcal{T})$ be undefined otherwise.

Notice that if $\mathcal{Q}(b, \mathcal{T})$ exists and $\delta(\mathcal{T}) \in \mathcal{Q}(b, \mathcal{T})$, then $\mathcal{Q}(b, \mathcal{T})$ is just the longest initial segment $\mathcal{Q}$ of $\mathcal{M}_b^{\mathcal{T}}$ such that $\mathcal{Q} \models \delta(\mathcal{T})$ is Woodin. There is a failure of $\delta(\mathcal{T})$ to be Woodin definable over $\mathcal{Q}(b, \mathcal{T})$.[41] Notice also that if b drops in either model or degree, then $\rho_n(\mathcal{M}_b^{\mathcal{T}}) < \delta(\mathcal{T})$ for some appropriate n, and therefore $\mathcal{Q}(b, \mathcal{T})$ exists.[42]

[41] The case $\rho_{n+1}(\mathcal{Q}(b, \mathcal{T})) < \delta(\mathcal{T})$ represents a failure of $\delta(\mathcal{T})$ to be a cardinal at all.

[42] Because $\mathcal{T}$ is maximal, b only drops when some extender used on b has critical point above a projectum of the model to which it is applied. At the last drop, this projectum is preserved as a projectum of $\mathcal{M}_b^{\mathcal{T}}$.

Suppose $\mathcal{Q}(b,\mathcal{T}) = \mathcal{Q}(c,\mathcal{T})$ (so both exist), and $\mathcal{Q}(b,\mathcal{T})$ is a proper initial segment of $\mathcal{M}_b^{\mathcal{T}}$ and $\mathcal{M}_c^{\mathcal{T}}$. Since $\mathcal{Q}(b,\mathcal{T})$ codes up a failure of Woodinness, 6.10 implies $b = c$. The following is a fine-structural strengthening of this fact.

6.12 Theorem. *Let $\mathcal{T}$ be k-maximal, and let b and c be distinct cofinal wellfounded branches of $\mathcal{T}$ such that $\mathcal{Q}(b,\mathcal{T})$ and $\mathcal{Q}(c,\mathcal{T})$ exist; then neither is an initial segment of the other.*

Proof. If one is an initial segment of the other, then since they are minimal with respect to the same first-order property, $\mathcal{Q}(b,\mathcal{T}) = \mathcal{Q}(c,\mathcal{T})$. Since this property involves a failure of $\delta(\mathcal{T})$ to be Woodin, $\mathcal{Q}(b,\mathcal{T}) \notin \mathcal{M}_b$ and $\mathcal{Q}(c,\mathcal{T}) \notin \mathcal{M}_c$ by 6.10. Thus $\mathcal{M}_b = \mathcal{Q}(b,\mathcal{T}) = \mathcal{Q}(c,\mathcal{T}) = \mathcal{M}_c$.

It follows that $\mathcal{Q}(b,\mathcal{T})$ and $\mathcal{Q}(c,\mathcal{T})$ are defined by the second clause of 6.11. If we let n be least such that $\rho_{n+1}(\mathcal{M}_b) < \delta(\mathcal{T})$, then there are $\eta \in b$ and $\xi \in c$ such that

$$\mathcal{M}^*_\eta = \mathcal{C}_{n+1}(\mathcal{M}_b) = \mathcal{C}_{n+1}(\mathcal{M}_c) = \mathcal{M}^*_\xi,$$

and $i_{\eta,b} \circ i^*_\eta$ and $i_{\xi,c} \circ i^*_\xi$ exist, and are n-embeddings with critical point at least $\rho_{n+1}(\mathcal{M}^*_\eta)$. But then, as in the fine structure argument at the end of the proof of the Comparison Lemma 3.11,

$$i_{\eta,b} \circ i^*_\eta = i_{\xi,c} \circ i^*_\xi,$$

since each is the core embedding from $\mathcal{C}_{n+1}(\mathcal{M}_b) = \mathcal{C}_{n+1}(\mathcal{M}_c)$ to $\mathcal{M}_b = \mathcal{M}_c$. Thus the extender applied to $\mathcal{M}^*_\eta$ in b is compatible with the extender applied to $\mathcal{M}^*_\xi$ in c, so that $\eta = \xi$.

Let α be the largest ordinal in $b \cap c$, so that $\alpha > \eta$ by the argument above. As usual, let us assume $n = 0$ to simplify matters a bit; the general case is essentially the same. Letting $\nu = \sup\{\nu(E_\beta) \mid \beta T \alpha\}$, we then have

$$\mathcal{M}_\alpha = \{i_{\eta,\alpha} \circ i^*_\eta(f)(a) \mid f \in \mathcal{M}^*_\eta \text{ and } a \in [\nu]^{<\omega}\}.$$

Since $i_{\alpha,b}$ and $i_{\alpha,c}$ are the identity on ν and agree on the range of $i_{\eta,\alpha} \circ i^*_\eta$, we have $i_{\alpha,b} = i_{\alpha,c}$. But this means the extender applied to $\mathcal{M}_\alpha$ in b is compatible with the extender applied to $\mathcal{M}_\alpha$ in c, so that α is not the largest element of $b \cap c$, a contradiction. ⊣

6.13 Definition. We say η is a *cutpoint* of $\mathcal{M}$ iff for all extenders E on the $\mathcal{M}$-sequence, if $\text{crit}(E) < \eta$ then $\text{lh}(E) < \eta$.

6.14 Corollary. *Let $\mathcal{T}$ be k-maximal; then there is at most one cofinal, wellfounded branch b of $\mathcal{T}$ such that*

- *$\mathcal{Q}(b,\mathcal{T})$ exists,*
- *$\delta(\mathcal{T})$ is a cutpoint of $\mathcal{Q}(b,\mathcal{T})$, and*

- *$\mathcal{Q}(b, \mathcal{T})$ is $\delta(\mathcal{T})^+ + 1$-iterable.*

Proof. Suppose b and c are distinct such branches. $\mathcal{Q}(b, \mathcal{T})$ and $\mathcal{Q}(c, \mathcal{T})$ have cardinality $\delta(\mathcal{T})$, so they are sufficiently iterable that their coiteration terminates successfully. Since $\delta(\mathcal{T})$ is a cutpoint of each model, and the two models agree below $\delta(\mathcal{T})$, all extenders used in this coiteration have critical point above $\delta(\mathcal{T})$. Also, each model is $\delta(\mathcal{T})$-sound and projects to $\delta(\mathcal{T})$, in the sense that there is an $n < \omega$ such that $\rho_{n+1}(\mathcal{Q}(b, \mathcal{T})) \leq \delta(\mathcal{T})$ and $\mathcal{Q}(b, \mathcal{T}) = \mathcal{H}_{n+1}^{\mathcal{Q}(b,\mathcal{T})}(\delta(\mathcal{T}) \cup p_{n+1}(\mathcal{Q}(b, \mathcal{T})))$, and similarly for $\mathcal{Q}(c, \mathcal{T})$. Just as in the proof of 3.12, this means that the side which comes out shorter does not move at all in the comparison, so that $\mathcal{Q}(b, \mathcal{T})$ is an initial segment of $\mathcal{Q}(c, \mathcal{T})$ or vice-versa. This contradicts 6.12. $\dashv$

Notice that all we needed in this argument was that $\mathcal{Q}(b, \mathcal{T})$ and $\mathcal{Q}(c, \mathcal{T})$ be iterable enough that we can compare them successfully. We can think of the structure $\mathcal{Q}(b, \mathcal{T})$ as a *branch oracle*, in that the fact that it is sufficiently iterable to be compared with other $\mathcal{Q}$-structures identifies b as the good branch of $\mathcal{T}$, the one any iteration strategy ought to choose. The sufficient-iterability-for-comparison of $\mathcal{Q}(b, \mathcal{T})$ only identifies b as the good branch, however, when $\delta(\mathcal{T})$ is a cutpoint of $\mathcal{Q}(b, \mathcal{T})$. This leads us to restrict our attention to mice all of whose Woodin cardinals are cutpoints.

6.15 Definition. A premouse $\mathcal{M}$ is *tame* iff whenever E is an extender on the $\mathcal{M}$-sequence, and $\lambda = \mathrm{lh}(E)$, then

$$\mathcal{J}_\lambda^{\mathcal{M}} \models \forall\delta \geq \mathrm{crit}(E)(\delta \text{ is not Woodin}).$$

In other words, tame mice cannot have extenders overlapping local Woodin cardinals. It is clear from the definition that any initial segment of a tame mouse is tame. Tame mice can satisfy large cardinal hypotheses as strong as "There is a strong cardinal which is a limit of Woodin cardinals". No tame mouse can satisfy "There is a Woodin cardinal which is a limit of Woodin cardinals".

The iterability conjecture above becomes a theorem when it is restricted to tame premice.

6.16 Theorem. *Let $\mathcal{N}$ be a tame premouse occurring on a K^c construction, let $k \leq \omega$, and let $\mathcal{M}$ be countable and such that there is a weak k-embedding from $\mathcal{M}$ to $\mathcal{C}_k(\mathcal{N})$; then $\mathcal{M}$ is $(k, \omega_1, \omega_1 + 1)$-iterable.*

We shall not prove this theorem here, but in the next section we shall prove a fairly representative special case of it.

6.3. Large cardinals in K^c

The iterability conjectures and theorems above show that K^c-constructions are sufficiently conservative about putting extenders on their sequences. We need also to know that they can be sufficiently liberal.

6.17 Definition. A K^c-construction $\langle \mathcal{N}_\alpha \mid \alpha < \theta \rangle$ is *maximal* iff $\mathcal{N}_{\alpha+1}$ is defined by case (2)(a) of definition 6.3 whenever possible; that is, a new extender is added to the current sequence whenever there is one meeting all the requirements of (2)(a) in 6.3.

One evidence of liberality is that large cardinal hypotheses true in V must also hold in K^c. Here is one such theorem.

6.18 Theorem. *Let δ be Woodin; then either*

- *there is a maximal K^c-construction $\langle \mathcal{N}_\alpha \mid \alpha < \theta + 1 \rangle$ such that $\mathcal{N}_\theta$ is not tame, or*
- *there is a maximal K^c-construction of length* $\mathrm{On}+1$*, and for any such construction $\langle \mathcal{N}_\alpha \mid \alpha \leq \mathrm{On} \rangle$,*

$$\mathcal{N}_{\mathrm{On}} \models \delta \textit{ is Woodin.}$$

Sketch of proof. If no maximal K^c-construction reaches a non-tame premouse, then by 6.16 and 5.3, every premouse occurring in a K^c-construction is ω-solid, and hence there are maximal K^c-constructions of length $\mathrm{On}+1$.

Let $\langle \mathcal{N}_\alpha \mid \alpha \leq \mathrm{On} \rangle$ be such a construction, and let $\mathcal{N}_{\mathrm{On}} = (L[\vec{E}], \in, \vec{E})$. Let $A \subseteq \delta$ and $A \in L[\vec{E}]$; we must find a $\kappa < \delta$ which is satisfied by $L[\vec{E}]$ to be A-reflecting in δ.

Since δ is Woodin in V, we can find a $\kappa < \delta$ which is $(A, \vec{E} \restriction \delta)$-reflecting in δ. Now if F is an extender over V which witnesses this reflection up to η, where $\kappa < \eta < \delta$ and η is, say, inaccessible, then we can show that for any $\xi < \eta$,

$$G_\xi := F \restriction \xi \cap L[\vec{E}] \in L[\vec{E}].$$

This is enough, for the extenders G_ξ witness that κ is A-reflecting in δ up to ξ in $L[\vec{E}]$.

To show that $G_\xi \in L[\vec{E}]$, we show by induction on ξ that if G_ξ is not of type Z, then the trivial completion of G_ξ is either on the sequence $\vec{E}$ or on an ultrapower of it, as in the initial segment condition in the definition of fine extender sequences. It is easy to see that G_ξ satisfies the requirements for being added to $\vec{E}$: coherence comes from the fact that F witnesses

$\vec{E} \upharpoonright \delta$-reflection[43], the initial segment condition comes from our induction hypothesis, and F provides the necessary background certificates. However, there are some problems. First, there is a *timing* problem: the above shows that G_ξ could be added to the $L[\vec{E}]$ sequence somewhere, but we need to find an actual stage $\mathcal{N}_\alpha$ of the construction at which it can be added. Second, there is a *uniqueness of the next extender* problem: we need to conclude from the fact that G_ξ could be added to produce $\mathcal{N}_\alpha+1$ that it *was* added to produce $\mathcal{N}_\alpha + 1$. For these arguments, we refer the reader to [26, Theorem 11.4]. $\dashv$

We note that the proof of 6.18 would have gone through if we had been even more conservative and required in 6.3 that our background extenders measure all sets in V. This requirement simplifies the iterability proof for the resulting model, as it allows us to lift trees on it to trees on V.[44] It is important in some contexts, however, to allow partial background extenders. For example, in proving relative consistency results in which the theory assumed consistent does not imply the existence of measurable cardinals, we must construct core models satisfying large cardinal hypotheses without assuming there are any extenders which are total over V. What assures us that maximal K^c-constructions are sufficiently liberal in that situation is the following.

6.19 Theorem. *Suppose μ is a normal measure on the measurable cardinal Ω, and that no K^c construction reaches a non-tame premouse. Let $\langle \mathcal{N}_\alpha \mid \alpha \leq \Omega \rangle$ be a maximal K^c-construction; then for μ-almost every $\alpha < \Omega$, $(\alpha^+)^{\mathcal{N}_\Omega} = \alpha^+$.*

This is essentially Theorem 1.4 of [44]. That is in turn an extension of earlier work of Jensen and Mitchell which in effect proved 6.19 under the hypothesis that no K^c-construction reaches the sharp for an inner model with a strong cardinal.[45]

Our focus for the rest of this paper will be on applications of core model theory in descriptive set theory, and so for simplicity we shall generally assume that there are Woodin cardinals in V. Therefore it will be 6.18 rather than 6.19 which is important for us. The reader should see [15] for an introductory article which turns at this point toward relative consistency results, results which make use of 6.19 rather than 6.18.

---

[43] This is not actually as obvious as it might seem at first, because the G_ξ ultrapower of $L[\vec{E}]$ only obviously agrees with the F ultrapower (and hence $L[\vec{E}]$) out to $\nu(G_\xi)$, rather than to the successor of $\nu(G_\xi)$ in the G_ξ ultrapower, as required by coherence. The stronger agreement can be proved using the condensation theorem 5.1, applied to the natural embedding of the G_ξ ultrapower into the F ultrapower.

[44] This is the iterability proof given in [26, section 12]. Of course, it only applies to tame mice; that is, it only proves a version of 6.16.

[45] Jensen and Mitchell did not require the measurable cardinal. ("μ-almost every" is replaced by "stationary many".) We suspect that the measurable cardinal is not needed in 6.19, but how to make do without it is an open question.

7. The reals of M_ω

We shall show that the reals in the minimal iterable proper class model satisfying "there are ω Woodin cardinals" are precisely those reals which are ordinal definable over $L(\mathbb{R})$. Of course, in order to do this we must assume that there is such a model. It will simplify matters if we assume something a bit stronger, namely, that there are ω Woodin cardinals with a measurable cardinal above them all (in V). We shall do so throughout the rest of this article, sometimes without explicitly mentioning the assumption. One useful consequence of our assumption is $\mathsf{AD}^{L(\mathbb{R})}$, the axiom of determinacy restricted to sets of reals in $L(\mathbb{R})$.[46]

7.1 Definition. A premouse $\mathcal{M}$ is *ω-small* iff whenever κ is the critical point of an extender on the $\mathcal{M}$-sequence, then

$$\mathcal{J}_\kappa^{\mathcal{M}} \not\models \text{There are } \omega \text{ Woodin cardinals.}$$

An ω-small mouse can satisfy "There are ω Woodin cardinals", but it cannot satisfy any significantly stronger large cardinal hypotheses.

7.2 Theorem. *If there are ω Woodin cardinals with a measurable cardinal above them all, then there is a $(\omega, \omega_1, \omega_1+1)$-iterable premouse which is not ω-small.*

Sketch of proof. Any nontame mouse is not ω-small, so we may assume our maximal K^c construction reaches only tame mice. Let $j\colon V \to M$ witness the measurability of some κ below which there are ω Woodin cardinals. By 6.18, the Woodin cardinals of M are Woodin in $j(K^c)$, and hence there are ω Woodin cardinals of $j(K^c)$ below κ. Now for any $A \subseteq V_{\kappa+1}$ of cardinality κ, the fragment $E_j \cap (A \times [j(\kappa)]^{<\omega})$ of the extender determined by j is in M. These fragments provide sufficient background certificates to show that there is an extender on the K^c sequence whose critical point is above all the Woodin cardinals of $j(K^c)$ which are below κ. Thus our maximal K^c construction reaches an $\mathcal{N}_\alpha$ which is not ω-small. By 6.16, any countable elementary submodel of $\mathcal{C}_\omega(\mathcal{N}_\alpha)$ witnesses the truth of the theorem. ⊣

7.3 Definition. $M_\omega^\sharp$ is the unique sound, $(\omega, \omega_1, \omega_1+1)$-iterable mouse which is not ω-small, but all of whose proper initial segments are ω-small.

It is easy to see that $\rho_1(M_\omega^\sharp) = \omega$, so that $M_\omega^\sharp$ is countable, and in fact every $x \in M_\omega^\sharp$ is Σ_1 definable over $M_\omega^\sharp$.[47] The uniqueness of $M_\omega^\sharp$ follows

[46] This is a result of Woodin, building on the work of [9] and [17]. See [29] for a proof.

[47] Suppose $\mathcal{M}$ is sufficiently iterable, not ω-small, and has only ω-small proper initial segments. The Σ_1 hull $\mathcal{H} := \mathcal{H}_1^{\mathcal{M}}(\emptyset)$ of $\mathcal{M}$ is sufficiently iterable that it can be compared with $\mathcal{J}_\alpha^{\mathcal{M}}$, for any $\alpha < \omega_1^{\mathcal{M}}$. Since $\mathcal{J}_\alpha^{\mathcal{M}}$ is ω-small, $\mathcal{H}$ must iterate past it, and it follows that for $\gamma = \omega_1^{\mathcal{M}}$, $\mathcal{J}_\gamma^{\mathcal{M}}$ is an initial segment of $\mathcal{H}$. Since we can easily compute a counting of $\mathcal{J}_\gamma^{\mathcal{H}}$ from the Σ_1 theory of $\mathcal{M}$, this theory is not a member of $\mathcal{M}$. Thus if $\mathcal{M}$ is 1-sound, $\mathcal{M} = \mathcal{H}$.

from 3.12. It is also clear that $M_\omega^\sharp$ is active; that is, it has a nonempty last extender predicate. We let M_ω be the proper class model left behind when the last extender of $M_\omega^\sharp$ is iterated out of the universe.

7.4 Definition. $M_\omega = \mathcal{J}^{\mathcal{P}}_{\mathrm{On}}$, where $\mathcal{P}$ is the On^{th} iterate of $M_\omega^\sharp$ by the last extender on its sequence.

It is clear that M_ω is an ω-small proper class model with ω Woodin cardinals, and that the Woodin cardinals of M_ω are countable in V. Their supremum is the supremum of the lengths of the extenders on the M_ω-sequence. The iterability of $M_\omega^\sharp$ easily implies that M_ω is $(\omega, \omega_1, \omega_1+1)$-iterable.

We shall show that the reals of M_ω are precisely the reals which are ordinal definable in $L(\mathbb{R})$.[48] We begin by showing that every real in M_ω is $\mathrm{OD}^{L(\mathbb{R})}$. Following the proof of 3.14, we see that for this it is enough to show that if $\alpha = \omega_1^{M_\omega}$, then $L(\mathbb{R})$ satisfies "$\mathcal{J}_\alpha^{\mathcal{M}_\omega}$ is ω_1+1-iterable".[49]

7.1. Iteration strategies in $L(\mathbb{R})$

Our task is complicated by the fact that M_ω is not itself (ω, ω_1+1)-iterable in $L(\mathbb{R})$, as we shall show later. We must drop to slightly smaller mice in order to find iteration strategies in $L(\mathbb{R})$.

7.5 Definition. A premouse $\mathcal{P}$ is *properly small* iff

- $\mathcal{P}$ is ω-small,
- $\mathcal{P} \models$ There are no Woodin cardinals, and
- $\mathcal{P} \models$ There is a largest cardinal $+ZF^-$.

Here ZF^- is ZF without the powerset axiom. It is clear that if α is a successor cardinal of M_ω below its least Woodin cardinal, then $\mathcal{J}_\alpha^{M_\omega}$ is properly small. In particular, this is true when $\alpha = \omega_1^{M_\omega}$.

7.6 Lemma. *Let $\mathcal{T}$ be an ω-maximal iteration tree of limit length on a properly small premouse, and let b be a cofinal wellfounded branch of $\mathcal{M}_b^{\mathcal{T}}$; then $\mathcal{Q}(b, \mathcal{T})$ exists.*

Proof. We have already observed that if b drops in model or degree, then $\rho_{n+1}(M_b) < \delta(\mathcal{T})$ for some n, so that $\mathcal{Q}(b, \mathcal{T})$ exists. Let $\mathcal{M} = \mathcal{M}_0^{\mathcal{T}}$. The

[48]Of course, M_ω and $M_\omega^\sharp$ have the same reals as members. $M_\omega^\sharp$ *is* (coded by) the simplest canonical real which is not $\mathrm{OD}^{L(\mathbb{R})}$; it is definable over $L(\mathbb{R} \cup \{\mathbb{R}^\sharp\})$ in a simple way.

[49]We are regarding this as a statement about the *parameter* $\mathcal{J}_\alpha^{M_\omega}$, which is in $L(\mathbb{R})$ because it is hereditarily countable. $L(\mathbb{R})$ need not believe that $\mathcal{J}_\alpha^{M_\omega}$ is obtained by implementing the definition of M_ω we gave in V.

requirement that $\mathcal{M}$ satisfy ZF^- insures that $\rho_\omega(\mathcal{M}) = \mathrm{On}^{\mathcal{M}}$, so that any iteration map along a non-dropping branch of an ω-maximal tree on $\mathcal{M}$ is fully elementary. The requirement that there are no Woodin cardinals in $\mathcal{M}$ then implies that there are none in $\mathcal{M}_b$, so that if $\delta(\mathcal{T}) < \mathrm{On}^{\mathcal{M}_b}$ then $\mathcal{Q}(b, \mathcal{T})$ exists. But we must have $\delta(\mathcal{T}) < \mathrm{On}^{\mathcal{M}_b}$, since if $\delta(\mathcal{T}) = \mathrm{On}^{\mathcal{M}_b}$, then as $\mathrm{lh}(E_\alpha^{\mathcal{T}})$ is a cardinal of $\mathcal{M}_b$ for all $\alpha < \mathrm{lh}(\mathcal{T})$, there is no largest cardinal of $\mathcal{M}_b$. ⊣

This lemma will, together with 6.12, guarantee that there is at most one iteration strategy for a properly small $\mathcal{M}$, and ultimately the $L(\mathbb{R})$-definability of this strategy when it exists.

It is useful to introduce yet another iteration game, one which requires less of player II than $\mathcal{G}_k(\mathcal{M}, \lambda, \theta))$. We call this new game the *weak iteration game*. Suppose $\mathcal{M}$ is a k-sound premouse; then the weak iteration game $\mathcal{W}_k(\mathcal{M}, \omega)$ is played in ω rounds as follows:

I	$\mathcal{T}_0$		$\mathcal{P}_1, i_1, \mathcal{T}_1$		$\mathcal{P}_2, i_2, \mathcal{T}_2$		$\dots$	
II		b_0		b_1		b_2		$\dots$

Here I begins by playing a countable, k-maximal, putative iteration tree $\mathcal{T}_0$ on $\mathcal{M}$, after which II plays b_0, which may be either "accept" or a maximal wellfounded branch of $\mathcal{T}_0$, with the proviso that II cannot accept unless $\mathcal{T}_0$ has a last model, and this model is wellfounded. Let $\mathcal{Q}_1$ be this last model, if II accepts, and let $\mathcal{Q}_1 = \mathcal{M}_{b_0}^{\mathcal{T}_0}$ otherwise. Let k_1 be the degree of $\mathcal{Q}_1$. Play now goes into the next round as it did in $\mathcal{G}_k(\mathcal{M}, \lambda, \theta)$: I picks an initial segment $\mathcal{P}_1$ of $\mathcal{Q}_1$, and an $i_1 \leq \omega$ such that $i_1 \leq k_1$ if $\mathcal{P}_1 = \mathcal{Q}_1$, together with a countable, i_1-maximal, putative iteration tree on $\mathcal{P}_1$. Then II either accepts or plays a maximal wellfounded branch of $\mathcal{T}_1$, with the proviso that he can only accept if $\mathcal{T}_1$ has a last, wellfounded model. Etc.

If no one breaks any of these rules along the way, then we say II wins this run of $\mathcal{W}_k(\mathcal{M}, \omega)$ iff for all sufficiently large i, $\mathcal{P}_i = \mathcal{Q}_i$, the branch of $\mathcal{T}_i$ from $\mathcal{P}_i$ to $\mathcal{Q}_{i+1}$ does not drop, and the direct limit of the $\mathcal{P}_i$ under the iteration maps given by the $\mathcal{T}_i$ is wellfounded.

7.7 Definition. A *weak (k, ω)-iteration strategy* for $\mathcal{M}$ is a winning strategy for II in $\mathcal{W}_k(\mathcal{M}, \omega)$, and we say $\mathcal{M}$ is *weakly (k, ω)-iterable* (or *$\Game^{\mathbb{R}}\Pi_1^1$-iterable*) just in case there is such a strategy.

It is an immediate consequence of the Branch Existence Theorem 6.6 that every countable elementary submodel $\mathcal{M}$ of $\mathcal{C}_k(\mathcal{N}_\alpha)$, where $\mathcal{N}_\alpha$ occurs in a K^c-construction, is weakly (k, ω)-iterable. In fact, such mice are weakly (k, ω_1)-iterable, in the obvious sense.[50] Weak (k, ω_1)-iteration strategies

[50] In $\mathcal{W}_k(\mathcal{M}, \omega_1)$, player I must play at limit $\lambda < \omega_1$ a tree $\mathcal{T}_\lambda$ on the direct limit of the models $\mathcal{P}_\eta$ for $\eta < \lambda$. Player II must insure that this direct limit is wellfounded.

suffice for the comparison of tame mice, and this fact is what lies behind our iterability theorem 6.16 for tame mice.[51]

If $\mathcal{M}$ is countable, and coded by the real x, then the weak iteration game $\mathcal{W}_k(\mathcal{M},\omega)$ is (can be coded as) a game of length ω on $\mathbb{R}$ with $\Pi^1_1(x)$ payoff. Thus the set of reals coding weakly iterable premice is $\Game^{\mathbb{R}}\Pi^1_1$, which explains the alternate terminology. By [16], $\Game^{\mathbb{R}}\Pi^1_1$ statements are absolute between V and $L(\mathbb{R})$, so we have:

7.8 Theorem. *Let $\mathcal{M}$ be countable and weakly (k,ω)-iterable; then $L(\mathbb{R}) \models$ $\mathcal{M}$ is weakly (k,ω)-iterable.*

It is also shown in [16] that $\Game^{\mathbb{R}}\Pi^1_1 = \Sigma_1^{L(\mathbb{R})}$, that is, that definitions in each form can be translated into the other.[52] We shall do our definability calculations below with Σ_1 formulae interpreted in $L(\mathbb{R})$. It is important here that we allow such formulae to contain a name $\dot{\mathbb{R}}$ for $\mathbb{R}$, so that quantification over $\mathbb{R}$ counts as bounded quantification. (Without this provision, we would have $\Sigma_1^{L(\mathbb{R})} = \Sigma^1_2$.) The sets whose definability we are calculating are generally subsets of HC, the class of hereditarily countable sets. Notice here that a set $A \subseteq HC$ is $\Sigma_1^{L(\mathbb{R})}$ iff the set A^* of all reals coding (in some natural system) a member of A is $\Sigma_1^{L(\mathbb{R})}$. So we have:

7.9 Lemma. *The set of countable, weakly (k,ω)-iterable premice is $\Sigma_1^{L(\mathbb{R})}$.*

If we restrict our attention to properly small premice, weak (k,ω)-iterability suffices for comparison.

7.10 Theorem. *Assume $\mathsf{AD}^{L(\mathbb{R})}$, and let $\mathcal{M}$ be countable, properly small, and weakly (k,ω)-iterable; then*

$$L(\mathbb{R}) \models \mathcal{M} \textit{ is } (k,\omega_1+1)\textit{-iterable.}$$

Proof. We first note

7.11 Lemma. *In $L(\mathbb{R})$, every iteration tree of length ω_1 on a countable premouse has a cofinal, wellfounded branch.*

Proof. Let $\mathcal{T}$ be such a tree. Let j be the embedding coming from the club ultrafilter on ω_1. Now $\mathcal{T}$ can be coded by a subset of ω_1, so $\mathcal{T} \in L[\mathcal{T}]$. As $L[\mathcal{T}]$ is wellordered, $j \restriction L[\mathcal{T}]$ is elementary from $L[\mathcal{T}]$ to $L[j(\mathcal{T})]$. Thus $j(\mathcal{T})$ is an iteration tree of length $j(\omega_1) > \omega_1$, so that $j(\mathcal{T}) \restriction \omega_1$ has a cofinal, wellfounded branch. But $j(\mathcal{T}) \restriction \omega_1 = \mathcal{T}$. $\dashv$

[51] See [41, Theorem 1.1] for the comparison proof. The proof of our unique strategies result 4.11 is the other main ingredient in the proof of 6.16.

[52] We only need here that $\Game^{\mathbb{R}}\Pi^1_1 \subseteq \Sigma_1^{L(\mathbb{R})}$, and this is trivial.

Because of this, it is enough to show that $\mathcal{M}$ is (k, ω_1)-iterable in $L(\mathbb{R})$. We claim that the following is a (k, ω_1)-iteration strategy for $\mathcal{M}$: given that you have reached $\mathcal{T}$ of countable limit length, pick the unique cofinal branch b of $\mathcal{T}$ such that $\mathcal{Q}(b, \mathcal{T})$ is weakly $(\deg^{\mathcal{T}}(b), \omega)$-iterable. Let us call this putative iteration strategy Γ.

Let $\mathcal{T}$ be played according to Γ, and of minimal length such that Γ breaks down at $\mathcal{T}$, either because $\mathcal{T}$ has limit length and there is no such unique branch to serve as $\Gamma(\mathcal{T})$, or because $\mathcal{T}$ has a last, illfounded model. Let Σ be a weak (k, ω)-iteration strategy for $\mathcal{M}$. If $\mathcal{T}$ has a last, illfounded model, then Σ cannot accept $\mathcal{T}$ as I's first move, so $\Sigma(\mathcal{T}) = b$ is a maximal branch of $\mathcal{T}$. Clearly, $\mathcal{Q}(b, \mathcal{T})$ is weakly $(\deg^{\mathcal{T}}(b), \omega)$-iterable, as witnessed by Σ. Letting $\lambda = \sup(b)$, we have from the definition of Γ that $b = \Gamma(\mathcal{T} \restriction \lambda)$, so $b = [0, \lambda]_T$, contrary to the maximality of b. Thus $\mathcal{T}$ has limit length. The argument just given shows that $b := \Sigma(\mathcal{T})$ is a cofinal branch of $\mathcal{T}$, and that $\mathcal{Q}(b, \mathcal{T})$ is weakly $(\deg^{\mathcal{T}}(b), \omega)$-iterable. Therefore there must be a second such branch; call it c. By 6.12 and the proof of 6.14, $\mathcal{Q}(b, \mathcal{T})$ and $\mathcal{Q}(c, \mathcal{T})$ cannot be compared. We shall use their weak iterability to compare them.

Let

$$\delta_0 = \sup\{\, \mathrm{lh}(E_\alpha^{\mathcal{T}}) \mid \alpha < \mathrm{lh}(\mathcal{T})\}.$$

Since δ_0 is Woodin in both $\mathcal{Q}(b, \mathcal{T})$ and $\mathcal{Q}(c, \mathcal{T})$, it is a cutpoint of each model. Since $\mathcal{Q}(b, \mathcal{T})$ and $\mathcal{Q}(c, \mathcal{T})$ agree below δ_0, the comparison we are doing uses only extenders with critical point strictly greater than δ_0.

Let $\Sigma_0 = \Sigma$ and Σ_1 be any weak $(\deg(c), \omega)$-iteration strategy for $\mathcal{Q}(c, \mathcal{T})$. Let $\mathcal{T}_0^0 = \mathcal{T}$, $b_0^0 = b$, and $c_0 = c$. We coiterate $\mathcal{Q}(b, \mathcal{T})$ and $\mathcal{Q}(c, \mathcal{T})$ by iterating the least disagreement at successor steps, and choosing the unique cofinal branch with a weakly iterable $\mathcal{Q}$-structure at limit steps. This process is $L(\mathbb{R})$-definable, and must break down at some countable stage, as otherwise by 7.11 and the proof of 3.11 we shall succeed in comparing $\mathcal{Q}(b, \mathcal{T})$ with $\mathcal{Q}(c, \mathcal{T})$. By the argument given above, the weak iterability of $\mathcal{Q}(b, \mathcal{T})$ and $\mathcal{Q}(c, \mathcal{T})$ implies that uniqueness is what breaks down. (It does not literally follow from 7.6 that cofinal branches always have $\mathcal{Q}$-structures, as the models we are comparing may no longer be properly small. But if, say, $\mathcal{Q}(b, \mathcal{T})$ is not properly small, then we have dropped along b getting to it, and this guarantees that in the tree on $\mathcal{Q}(b, \mathcal{T})$ we are now building, cofinal branches always have $\mathcal{Q}$-structures.) Let $\mathcal{T}_1^0$ on $\mathcal{Q}(b, \mathcal{T})$ and $\mathcal{T}_1^1$ on $\mathcal{Q}(c, \mathcal{T})$ be the trees produced by this process. Let

$$\begin{aligned} \delta_1 &= \sup\{\, \mathrm{lh}(E_\alpha^{\mathcal{T}_1^0}) \mid \alpha < \mathrm{lh}(\mathcal{T}_1^0)\} \\ &= \sup\{\, \mathrm{lh}(E_\alpha^{\mathcal{T}_1^1}) \mid \alpha < \mathrm{lh}(\mathcal{T}_1^1)\}. \end{aligned}$$

Let

$$b_1^0 = \Sigma_0(\langle \mathcal{T}_0^0, (\mathcal{Q}(b_0^0, \mathcal{T}_0^0), \mathcal{T}_1^0)\rangle)$$

and

$$b_1^1 = \Sigma_1(\mathcal{T}_1^1)$$

be the cofinal, weakly iterable branches of $\mathcal{T}_1^0$ and $\mathcal{T}_1^1$ chosen by Σ_0 and Σ_1. By hypothesis we have a third branch c_1 of some $\mathcal{T}_1^i$ (it does not matter which) such that $\mathcal{Q}(c_1, \mathcal{T}_1^i)$ is weakly $(\deg(c_1), \omega)$-iterable, say via the strategy Σ_2. It follows that the premice $\mathcal{Q}(b_1^0, \mathcal{T}_1^0)$, $\mathcal{Q}(b_1^1, \mathcal{T}_1^1)$, and $\mathcal{Q}(c_1, \mathcal{T}_1^i)$ cannot be compared.

We attempt to reach a contradiction by simultaneously comparing these three premice. (This means that we form three iteration trees simultaneously, iterating by the shortest extender on the sequence of any of the three last models which is not present on the sequences of both of the other two last models.) Again, we choose unique weakly iterable branches at limit ordinals, and again this process must break down due to non-uniqueness, giving trees $\mathcal{T}_2^0, \mathcal{T}_2^1$, and $\mathcal{T}_2^2$, with cofinal branches b_2^0, b_2^1, and b_2^2 chosen by Σ_0, Σ_1, and Σ_2. (It is because the $\mathcal{T}_2^i$ use only extenders with critical point above δ_1 that we can interpret them as played by the Σ_i.) We also have a new branch c_2 of some $\mathcal{T}_2^i$, and a weak iteration strategy Σ_3 for $\mathcal{Q}(c_2, \mathcal{T}_2^i)$. We let δ_2 be the sup of the lengths of the extenders used in the $\mathcal{T}_2^i$. And so on.

After ω steps in the process we have for each $i < \omega$ a weak iteration strategy Σ_i and a play by Σ_i in which the iteration trees played by I are the $\mathcal{T}_j^i$ for $j \geq i$ and the branches chosen by II are the b_j^i for $j \geq i$. Let $\mathcal{P}_i$ be the direct limit of the $\mathcal{M}_{b_j^i}^{\mathcal{T}_j^i}$. Since each Σ_i is winning, these direct limits are wellfounded. Clearly, all the δ_k are Woodin in each $\mathcal{P}_i$. Since $\mathcal{P}_i$ is ω-small, it has no extenders with index above the sup of the δ_k, and thus $\mathcal{P}_i$ is an initial segment of $\mathcal{P}_n$ or vice-versa, for all i and n. Since all $\mathcal{P}_i$ project below the sup of the δ_k, they must all be the same. Moreover, as in the proof of the Comparison Lemma 3.11, we can show that for no i does the composition of the trees $\mathcal{T}_j^i$ drop in model or degree on the branch leading to $\mathcal{P}_i$. But this means that $\mathcal{P}_0$ and $\mathcal{P}_1$ are the last models of a successful comparison of $\mathcal{Q}(b, \mathcal{T})$ with $\mathcal{Q}(c, \mathcal{T})$, a contradiction. ⊣

We have at once

7.12 Corollary. *Every real in M_ω is ordinal definable in $L(\mathbb{R})$.*

Proof. Let x be the α^{th} real in the order of constructibility of M_ω; then

$$\begin{aligned} y = x \quad \Leftrightarrow \quad & L(\mathbb{R}) \models \exists \mathcal{M}(\mathcal{M} \text{ is countable, properly small,} \\ & (\omega, \omega_1 + 1)\text{-iterable, and } y \text{ is the } \alpha^{th} \text{ real} \\ & \text{in the constructibility order of } \mathcal{M}.) \end{aligned}$$

⊣

The proof of 7.10 gives at once:

7.13 Corollary. *Assume* $\mathsf{AD}^{L(\mathbb{R})}$*, and let* $\mathcal{M}$ *be countable, properly small, and weakly* (k,ω)*-iterable; then in* $L(\mathbb{R})$*,* $\mathcal{M}$ *has a unique* (k,ω_1)*-iteration strategy* Σ*; moreover,* Σ *is* $\Sigma_1^{L(\mathbb{R})}(\{\mathcal{M}\})$ *definable, uniformly in* $\mathcal{M}$*, and* Σ *extends, in* $L(\mathbb{R})$*, to a* (k,ω_1+1)*-iteration strategy for* $\mathcal{M}$*.*

7.2. Correctness and genericity iterations

We shall prove some correctness results for M_ω, and use them to show that every real ordinal definable over $L(\mathbb{R})$ is in M_ω. The key to these results is the following remarkable theorem of W.H. Woodin.

7.14 Theorem. *Let* Σ *be an* (ω,ω_1+1)*-iteration strategy for* $\mathcal{M}$*, and suppose* δ *is a countable ordinal such that* $\mathcal{M} \models \mathsf{ZF}^- + \ \delta$ *is Woodin ; then there is a* $\mathbb{Q} \subseteq V_\delta^{\mathcal{M}}$ *such that*

- $\mathcal{M} \models \mathbb{Q}$ *is a* δ*-c.c. complete Boolean algebra, and*
- *for any real* x*, there is a countable iteration tree* $\mathcal{T}$ *on* $\mathcal{M}$ *played according to* Σ *with last model* $\mathcal{M}_\alpha$ *such that* $i_{0,\alpha}$ *exists and* x *is* $i_{0,\alpha}(\mathbb{Q})$*-generic over* $\mathcal{M}_\alpha$*.*

Proof. Working in $\mathcal{M}$, let $L_{\delta,0}$ be the infinitary language whose formulae are built up by means of conjunctions and disjunctions of size $< \delta$, and negation, from the propositional letters A_n, for $n < \omega$. (So all formulae are quantifier-free.) Any real x, regarded as a subset of ω, gives us an interpretation of $L_{\delta,0}$:

$$x \models A_n \Leftrightarrow n \in x.$$

We can then define $x \models \varphi$, for arbitrary formulae φ, by the obvious induction.

Still working in $\mathcal{M}$, consider the $L_{\delta,0}$ theory S which has the axioms

$$\bigvee_{\alpha<\kappa} \varphi_\alpha \longleftrightarrow \bigvee_{\alpha<\lambda} i_E(\langle \varphi_\xi \mid \xi < \kappa \rangle))$$

whenever E is on the $\mathcal{M}$-sequence, $\mathrm{crit}(E) = \kappa \leq \lambda$, and $\nu(E)$ is an $\mathcal{M}$-cardinal such that $i_E(\langle \varphi_\xi \mid \xi < \kappa \rangle) \restriction \lambda \in \mathcal{J}_{\nu(E)}^{\mathcal{M}}$.

We let $\mathbb{Q}$ be the Lindenbaum algebra of S. That is, we let

$$\varphi \sim \psi \text{ iff } S \vdash \varphi \leftrightarrow \psi,$$

and

$$[\varphi] \leq [\psi] \text{ iff } S \vdash \varphi \rightarrow \psi,$$

and we let

$$\mathbb{Q} := (\{[\varphi] \mid \varphi \in L_{\delta,0}\}, \leq).$$

Here provability in S means provability using the usual finitary rules together with the rule: from φ_α for all $\alpha < \kappa$ (where $\kappa < \delta$) infer $\bigwedge_{\alpha<\kappa} \varphi_\alpha$. Equivalently, $S \vdash \tau$ iff whenever x is a real in $\mathcal{M}[G]$ for some G generic over $\mathcal{M}$ and $x \models S$, then $x \models \tau$. (See [2]. Clearly, if $S \vdash \tau$, then *any* real which satisfies S satisfies τ.)

Claim 1. $\mathbb{Q}$ is δ-c.c. in $\mathcal{M}$.

Proof. We work in $\mathcal{M}$. Let $\langle [\varphi_\alpha] \mid \alpha < \delta \rangle$ be an antichain in $\mathbb{Q}$. Let $\kappa < \delta$ be $\langle \varphi_\alpha \mid \alpha < \delta \rangle$-reflecting. Let ν be a cardinal such that $\varphi_\kappa \in \mathcal{J}_\nu^{\mathcal{M}}$, and let F on the $\mathcal{M}$-sequence witness the reflection of κ at this ν.[53] Let E be the trivial completion of $F \restriction \nu$. We then have

$$i_E(\bigvee_{\alpha<\kappa} \varphi_\alpha) \restriction (\kappa + 1) = \bigvee_{\alpha\leq\kappa} \varphi_\alpha,$$

so that

$$\bigvee_{\alpha<\kappa} \varphi_\alpha \longleftrightarrow \bigvee_{\alpha\leq\kappa} \varphi_\alpha$$

is provable in S. It follows that $[\varphi_\kappa] \leq \bigvee_{\alpha<\kappa} [\varphi_\alpha]$ in $\mathbb{Q}$, a contradiction. ⊣

Claim 2. $\mathbb{Q}$ is a complete Boolean algebra in $\mathcal{M}$.

Proof. $\mathbb{Q}$ is closed under sums of size $< \delta$ since $\bigvee_{\alpha<\kappa} [\varphi_\alpha] = [\bigvee_{\alpha<\kappa} \varphi_\alpha]$. By claim 1, $\mathbb{Q}$ is closed under arbitrary sums. ⊣

Claim 3. If $x \models S$, then setting $G_x := \{[\varphi] \mid x \models \varphi\}$, we have that G_x is $\mathbb{Q}$-generic over $\mathcal{M}$ and $x \in \mathcal{M}[G_x]$.

Proof. Since $x \models S$, G_x is welldefined on equivalence classes: if $S \vdash (\varphi \leftrightarrow \psi)$, then $x \models \varphi$ iff $x \models \psi$. It is also clear that G_x is an ultrafilter on $\mathbb{Q}$. To see that G_x is $\mathcal{M}$-generic, let $\langle [\varphi_\alpha] \mid \alpha < \nu \rangle$ be a maximal antichain. Since $[\bigvee_{\alpha<\nu} \varphi_\alpha] = 1$, we have $S \vdash \bigvee_{\alpha<\nu} \varphi_\alpha$. Since $x \models S$, we have $x \models \varphi_\alpha$ for some α; that is, $[\varphi_\alpha] \in G_x$ for some α. Finally, $n \in x$ iff $A_n \in G_x$, so $x \in \mathcal{M}[G_x]$. ⊣

An arbitrary real x may not satisfy S, but one can iterate $\mathcal{M}$ in such a way that x satisfies some image of S.

Claim 4. For any real x, there is a countable iteration tree $\mathcal{T}$ on $\mathcal{M}$ which is played according to Σ, has last model $\mathcal{M}_\alpha$, and is such that $[0, \alpha]_T$ does not drop and $x \models i_{0,\alpha}(S)$.

Proof. We keep iterating away the first extender which induces an axiom not satisfied by x. More precisely, set $\mathcal{M}_0 = \mathcal{M}$, and now suppose we have

[53] We are using here the fact that the Woodinness of δ in $\mathcal{M}$ is witnessed by extenders on the $\mathcal{M}$-sequence. We might just have added this to the hypotheses of 7.14, but we need not do so because, by [35], it follows from the other hypotheses.

constructed the model $\mathcal{M}_\beta$ of $\mathcal{T}$, where $\beta < \omega_1$. Suppose also $\mathcal{T}$ has not dropped anywhere yet; that is, $D^{\mathcal{T}} = \emptyset$ as of now. If $x \models i_{0,\beta}(S)$ we are done, so suppose not. Let E on the $\mathcal{M}_\beta$-sequence be such that E induces an axiom of $i_{0,\beta}(S)$ which is false of x, and $\mathrm{lh}(E)$ is minimal among all extenders on the $\mathcal{M}_\beta$-sequence with this property. We set $E := E^{\mathcal{T}}_\beta$, and use E according to the rules for ω-maximal iteration trees to extend $\mathcal{T}$ one more step.

We must check here that $\gamma < \beta \Rightarrow \mathrm{lh}(E_\gamma) < \mathrm{lh}(E_\beta)$. But if not, the agreement of models in an ω-maximal iteration tree implies that E_β is on the sequence of $\mathcal{M}_\gamma$, and it is not hard to check that the false axiom of $i_{0,\beta}(S)$ it induces in $\mathcal{M}_\beta$ is also induced by it in $\mathcal{M}_\gamma$. (To see that $\nu(E_\beta)$ is a cardinal of $\mathcal{M}_\gamma$ in this situation, note that since $\nu(E_\gamma)$ is a cardinal of $\mathcal{M}_\gamma$, any cardinal of $\mathcal{M}_\beta$ which is $\leq \nu(E_\gamma)$ is a cardinal of $\mathcal{M}_\gamma$. But $\nu(E_\beta) < \mathrm{lh}(E_\beta) \leq \mathrm{lh}(E_\gamma)$ and there are no cardinals of $\mathcal{M}_\beta$ in the interval $(\nu(E_\gamma), \mathrm{lh}(E_\gamma))$, so $\nu(E_\beta) \leq \nu(E_\gamma)$.)

We must also check that $[0, \beta{+}1]$ does not drop; that is, that E_β measures all subsets of its critical point κ in the model $\mathcal{M}_\gamma$ to which it is applied. This is true because $\kappa < \nu(E_\gamma)$, $\nu(E_\gamma)$ is a cardinal of $\mathcal{M}_\gamma$, and $\mathcal{M}_\beta$ agrees with $\mathcal{M}_\gamma$ below $\nu(E_\gamma)$.

This finishes the successor step in the construction of $\mathcal{T}$. At limit ordinals $\lambda \leq \omega_1$ we use Σ to extend $\mathcal{T}$.

It is enough to show this process terminates at some countable ordinal, so suppose not. We reach a contradiction much as in the proof that the comparison process terminates. As in that argument, let

$$\pi\colon H \to V_\eta$$

be elementary, where H is a countable, transitive set, and V_η and the range of π are large enough to contain everything of interest. Let $\pi(\bar{\mathcal{T}}) = \mathcal{T}$, etc., and let $\alpha = \mathrm{crit}(\pi) = \omega_1^H$. We get as before, setting $\delta^* = i^{\mathcal{T}}_{0,\alpha}(\delta) = i^{\bar{\mathcal{T}}}_{0,\alpha}(\delta)$,

$$V_{\delta^*}^{\mathcal{M}_\alpha^{\bar{\mathcal{T}}}} = V_{\delta^*}^{\mathcal{M}_\alpha^{\mathcal{T}}}$$

and

$$\pi \restriction V_{\delta^*}^{\mathcal{M}_\alpha^{\bar{\mathcal{T}}}} = i^{\mathcal{T}}_{\alpha,\omega_1} \restriction V_{\delta^*}^{\mathcal{M}_\alpha^{\mathcal{T}}}.$$

Now let $\beta + 1$ be the T-successor of α on $[0, \omega_1]_T$. We have $\mathrm{crit}(E_\beta) = \mathrm{crit}(i_{\alpha,\omega_1}) = \alpha$, and we have an axiom

$$\bigvee_{\gamma<\alpha} \varphi_\gamma \longleftrightarrow i_{E_\beta}(\bigvee_{\gamma<\alpha} \varphi_\gamma) \restriction \lambda$$

of $i_{0,\beta}(S)$ induced by E_β and false of x. The falsity of this axiom means that the right hand side is true of x, but the left hand side is not. But now

$\bigvee_{\gamma<\alpha} \varphi_\gamma$ is essentially a subset of α, and therefore is small enough that it is in $\mathcal{M}_\alpha$. Moreover, $\lambda < \nu(E_\beta)$, and since generators are not moved on $\mathcal{T}$

$$i_{E_\beta}(\bigvee_{\gamma<\alpha} \varphi_\gamma) \restriction \lambda = i_{\alpha,\omega_1}(\bigvee_{\gamma<\alpha} \varphi_\gamma) \restriction \lambda = \pi(\bigvee_{\gamma<\alpha} \varphi_\gamma) \restriction \lambda.$$

But $x \in H$ and $\pi(x) = x$. Since $L_{\delta,0}$ satisfaction is sufficiently absolute and $x \not\models \bigvee_{\gamma<\alpha} \varphi_\gamma$, we have $x \not\models \pi(\bigvee_{\gamma<\alpha} \varphi_\gamma)$. This contradicts the fact that x satisfies the initial segment $i_{E_\beta}(\bigvee_{\gamma<\alpha} \varphi_\gamma) \restriction \lambda$ of this disjunction. $\dashv$

If $\mathcal{M}_\alpha$ is as in claim 4, then we can replace $\mathcal{M}$ by $\mathcal{M}_\alpha$ in claims 1, 2, and 3, and we see then that $\mathcal{T}$ and $\mathcal{M}_\alpha$ witness the conclusion of 7.14. $\dashv$

The complete Boolean algebra $\mathbb{Q}$ of 7.14 is known as the *extender algebra*.

We drop for a moment to smaller mice, and use the extender algebra to prove a correctness result for the minimal proper class model with one Woodin cardinal. (This was Woodin's original application of 7.14.) Let us call a premouse $\mathcal{M}$ *1-small* iff whenever κ is the critical point of an extender on the $\mathcal{M}$-sequence, then $\mathcal{J}_\kappa^{\mathcal{M}} \models$ "there are no Woodin cardinals". Let $M_1^\sharp$ be the least mouse which is not 1-small, and M_1 the result of iterating the last extender of $M_1^\sharp$ out of the universe. (Granted that there is a Woodin cardinal with a measurable above it in V, $M_1^\sharp$ exists and is $(\omega, \omega_1 + 1)$-iterable.) Let $\mathbb{Q}$ be the extender algebra of M_1; then for any Σ_3^1 sentence φ, possibly involving real parameters from M_1, we have

$$\varphi \iff M_1 \models \exists p(p \Vdash \varphi).$$

The right-to-left direction comes from the fact that $P(\mathbb{Q}) \cap M_1$ is countable in V, so that any condition is extended by a generic filter in V. For the left-to-right direction: let x witness the outer existential quantifier of φ, and let $\mathcal{M}_\alpha$ be an iterate of M_1 over which x is $i_{0,\alpha}(\mathbb{Q})$-generic. Clearly, $\mathcal{M}_\alpha[G_x] \models \varphi$, so $\mathcal{M}_\alpha \models \exists p(p \Vdash \varphi)$, so by elementarity $M_1 \models \exists p(p \Vdash \varphi)$.

Thus M_1 can compute Σ_3^1 truth by asking what is forced in its extender algebra. (M_1 is not itself Σ_3^1-correct.) This easily implies that every real which is Δ_3^1 in a countable ordinal is in M_1. A careful look at the sort of iterability needed to compare "properly 1-small" mice (like $\mathcal{J}_\alpha^{M_1}$, for $\alpha = \omega_1^{M_1}$) shows every real in M_1 is Δ_3^1 in a countable ordinal, so we have a descriptive-set-theoretic characterization of the reals in M_1.[54]

$M_1^\sharp$ is essentially a real, and from this real we can recursively construct generic objects for the extender algebra of M_1 below any condition. It follows that every nonempty Σ_3^1 set of reals has a member recursive in

[54] This set of reals is known in descriptive set theory as Q_3, and it has many other characterizations. $M_1^\sharp$ is also known from descriptive set theory; it is Turing equivalent to y_0. See [12].

$M_1^\sharp$. We can relativise the $M_1^\sharp$ construction to an arbitrary real x and obtain $M_1^\sharp(x)$; simply throw x into the model at the bottom. We get that any nonempty $\Sigma^1_3(x)$ set of reals has a member recursive in $M_1^\sharp(x)$, and therefore any premouse closed under the function $x \mapsto M_1^\sharp(x)$ is Σ^1_3-correct. In particular, M_ω is Σ^1_3-correct.

If we give M_n and $M_n^\sharp$ the obvious meaning, then we can show that the reals of M_n are precisely those which are Δ^1_{n+2} in a countable ordinal, and that every nonempty Σ^1_{n+2} set of reals has a member recursive in $M_n^\sharp$. (See [43].) Since M_ω is closed under $x \mapsto M_n^\sharp(x)$ for all $n < \omega$, M_ω is projectively correct. The following theorem gives us much more; it says that M_ω can compute $L(\mathbb{R})$ truth in much the same way that M_1 can compute Σ^1_3 truth.

We let $\mathrm{Col}(\omega, X)$ be the Levy collapsing poset of all finite functions from ω into X. Notice that 7.14 implies that if $\mathcal{M}$, Σ, and δ satisfy its hypotheses, then for any real x there is a countable $\mathcal{T}$ played by Σ, with last model $\mathcal{M}_\alpha$, such that x is $\mathrm{Col}(\omega, i_{0,\alpha}(\delta))$-generic over $\mathcal{M}_\alpha$. This is true because $\mathrm{Col}(\omega,\kappa)$ is universal for forcings of size κ. Unlike the extender algebra, $\mathrm{Col}(\omega,\delta)$ is not δ-c.c.; on the other hand, it is homogeneous.

By $\mathrm{Col}(\omega, {<}\lambda)$ we mean the finite support product of all $\mathrm{Col}(\omega,\alpha)$ such that $\alpha < \lambda$. If G is $\mathcal{M}$-generic over $\mathrm{Col}(\omega,{<}\lambda)$, then we set

$$\mathbb{R}^*_G := \bigcup_{\alpha<\lambda} \mathbb{R} \cap \mathcal{M}[G \cap \mathrm{Col}(\omega, {<}\alpha)],$$

and say that $\mathbb{R}^*_G$ is the set of reals of a symmetric collapse of $\mathcal{M}$ below λ.

7.15 Theorem. *Suppose that $\mathcal{M} \models \lambda$ is a limit of Woodin cardinals, where λ is countable in V, and that Σ is an (ω,ω_1+1)-iteration strategy for $\mathcal{M}$. Let H be $\mathrm{Col}(\omega,\mathbb{R})$-generic over V; then in $V[H]$ there is an iteration map $i\colon \mathcal{M} \to \mathcal{N}$ coming from an iteration tree all of whose proper initial segments are played by Σ, and a G which is $\mathrm{Col}(\omega,{<}i(\lambda))$-generic over $\mathcal{N}$, such that*

$$\mathbb{R}^*_G = \mathbb{R}^V.$$

Proof. We shall need the following slight refinement of 7.14.

7.16 Lemma. *Let $\mathcal{M} \models \delta$ is Woodin, where δ is countable in V, and let Σ be an (ω,ω_1+1)-iteration strategy for $\mathcal{M}$. Let $\kappa < \delta$, and let G be $\mathcal{M}$-generic for a poset $\mathbb{P} \in V^{\mathcal{M}}_\kappa$. Then for any $x \subseteq \omega$, there is a countable iteration tree $\mathcal{T}$ played by Σ and having last model $\mathcal{M}_\alpha$ such that*

- *$D^{\mathcal{T}} = \emptyset$ and $\mathrm{crit}(E^{\mathcal{T}}_\beta) > \kappa$ for all β, and*
- *x is in some $\mathrm{Col}(\omega,\delta)$-generic extension of $\mathcal{M}_\alpha[G]$.*

Sketch of proof. In $\mathcal{M}[G]$, δ is still Woodin via the extenders over $\mathcal{M}[G]$ which are "completions" of extenders on the $\mathcal{M}$-sequence with critical point

$> \kappa$. So in $\mathcal{M}[G]$, the version of the extender algebra which uses only these extenders is a δ-c.c. complete Boolean algebra. The iteration $\mathcal{U}$ of $\mathcal{M}[G]$ we need to do to make x generic can be obtained by from an iteration $\mathcal{T}$ of $\mathcal{M}$: $\mathcal{M}^{\mathcal{U}}_\beta = \mathcal{M}^{\mathcal{T}}_\beta[G]$ for all β. We omit further details. ⊣

We can now prove the theorem. Working in $V[H]$, let $\langle x_n \mid n < \omega\rangle$ be an enumeration of $\mathbb{R}^V$. Let $\langle \delta_n \mid n < \omega\rangle$ be an increasing sequence of Woodin cardinals of $\mathcal{M}$ which is cofinal in λ. We shall use 7.16 to successively absorb the x_n into the collapse of some image of δ_n in an iterate of $\mathcal{M}$.

More precisely, working in V we find a countable iteration tree $\mathcal{T}_0$ on $\mathcal{M}$ played by Σ with last model $\mathcal{P}_0$, and a G_0 which is $\mathcal{P}_0$-generic over $\mathrm{Col}(\omega, i_0(\delta_0))$, where $i_0 \colon \mathcal{M} \to \mathcal{P}_0$ is the iteration map, so that $x_0 \in \mathcal{P}_0[G_0]$. We then find an iteration tree $\mathcal{T}_1$ on $\mathcal{P}_0$ such that $\mathcal{T}_0 \oplus \mathcal{T}_1$ is according to Σ, and if $i_1 \colon \mathcal{P}_0 \to \mathcal{P}_1$ is the iteration map, then $\mathrm{crit}(i_1) > i_0(\delta_0)$, and there is a G_1 which is $\mathcal{P}_1[G_0]$-generic over $\mathrm{Col}(\omega, i_1 \circ i_0(\delta_1))$ such that $x_1 \in \mathcal{P}_1[G_0][G_1]$. And so on: given $\mathcal{P}_n$, we use 7.16 in V to obtain an iteration tree $\mathcal{T}_{n+1}$ on $\mathcal{P}_n$ such that $\mathcal{T}_0 \oplus \ldots \oplus \mathcal{T}_{n+1}$ is according to Σ, and if $i_{n+1} \colon \mathcal{P}_n \to \mathcal{P}_{n+1}$ is the iteration map, then $\mathrm{crit}(i_{n+1}) > i_n \circ \ldots \circ i_0(\delta_n)$, and there is a G_{n+1} which is $\mathcal{P}_{n+1}[G_0, \ldots, G_n]$-generic over $\mathrm{Col}(\omega, i_{n+1} \circ \ldots \circ i_0(\delta_{n+1}))$ such that $x_{n+1} \in \mathcal{P}_{n+1}[G_0, \ldots, G_n][G_{n+1}]$.

Let $\mathcal{T} = \bigoplus_n \mathcal{T}_n \oplus b$, where b is the branch of $\bigoplus_n \mathcal{T}_n$ containing the $\mathcal{P}_n$. By construction, b is the unique cofinal branch of $\bigoplus_n \mathcal{T}_n$, and the $\mathcal{T}_n$ constitute a play by Σ. Let $\mathcal{N}$ be the last model of $\mathcal{T}$; clearly $\mathcal{N}$ is just the direct limit of the $\mathcal{P}_n$ under the i_n. A simple absoluteness argument shows that $\mathcal{N}$ is wellfounded: if not, then the tree of attempts to produce a sequence $\langle \mathcal{U}_n \mid n \in \omega\rangle$ which constitutes a play of ω rounds of $\mathcal{G}_\omega(\mathcal{M}, \omega, \omega_1 + 1)$ by Σ, together with a descending chain of ordinals in the direct limit along the unique cofinal branch, would have a branch in V. Let $i \colon \mathcal{M} \to \mathcal{N}$ be the direct limit map. By construction, each G_n is in V, so we have $x_n \in (\mathbb{R} \cap \mathcal{N}[G_0, \ldots, G_{n+1}]) \subseteq \mathbb{R}^V$, and therefore $\bigcup_n (\mathbb{R} \cap \mathcal{N}[G_0, \ldots, G_n]) = \mathbb{R}^V$. It is easy to see that $\bigcup_n (\mathbb{R} \cap \mathcal{N}[G_0, \ldots, G_n])$ is the set of reals of a symmetric collapse of $\mathcal{N}$ below $i(\lambda)$, so we are done. ⊣

7.17 Corollary. *Let $\mathcal{M}$ be a proper class premouse such that $\mathcal{M} \models \lambda$ is a limit of Woodin cardinals, where λ is countable in V, and suppose $\mathcal{M}$ is $(\omega, \omega_1 + 1)$-iterable; then every real which is ordinal definable over $L(\mathbb{R})$ belongs to $\mathcal{M}$.*

Proof. Let $i \colon \mathcal{M} \to \mathcal{N}$ be as in 7.15, and let x be $\mathrm{OD}^{L(\mathbb{R})}$. We have, by the symmetry of $\mathrm{Col}(\omega, {<}i(\lambda))$ and the fact that $L(\mathbb{R})^V$ is realized as some $L(\mathbb{R}^*_G)$, that $x \in \mathcal{N}$. It follows that $x \in \mathcal{M}$. ⊣

The proof of 6.16 shows that if λ is a limit of Woodin cardinals and there is a measurable cardinal above λ, then $M^\sharp_\omega$ exists and is $(\omega, \omega_1, \omega_1 + 1)$-

iterable , not just in V, but in $V^{\mathbb{P}}$, for any poset $\mathbb{P}$ of cardinailty $< \lambda$. So we get at once:

7.18 Corollary. *If there are ω Woodin cardinals with a measurable above them all in V, then $\mathbb{R} \cap M_\omega = \{x \in \mathbb{R} \mid x$ is $\mathrm{OD}^{L(\mathbb{R})}\}$.*

We are in a position now to see that M_ω has no (ω, ω_1)-iteration strategy in $L(\mathbb{R})$. (We assume here that there are in V ω Woodin cardinals with a measurable above them all.) For if there were such a strategy in $L(\mathbb{R})$, then the set of reals which are not in M_ω would be a $\Sigma_1^{L(\mathbb{R})}$ set: $z \notin M_\omega$ iff $L(\mathbb{R}) \models$(there is an (ω, ω_1)-iterable, ω-small premouse $\mathcal{N}$ of ordinal height ω_1 such that for some countable λ, $\mathcal{N} \models \lambda$ is a limit of Woodin cardinals, and such that $z \notin \mathcal{N}$). However, by [16], any nonempty $\Sigma_1^{L(\mathbb{R})}$ set of reals has an $\mathrm{OD}^{L(\mathbb{R})}$ member.[55] So there is an $\mathrm{OD}^{L(\mathbb{R})}$ real not in M_ω, contrary to 7.18.

The proof of 7.15 shows that any sufficiently iterable proper class model with ω Woodin cardinals can compute $L(\mathbb{R})$ truth by consulting its symmetric collapse; in fact

7.19 Theorem. *Let $\mathcal{M}$ be a proper class premouse such that $\mathcal{M} \models \lambda$ is a limit of Woodin cardinals, where λ is countable in V, and suppose $\mathcal{M}$ is $(\omega, \omega_1 + 1)$-iterable. Let $\mathbb{R}^*$ be the set of reals of a symmetric collapse of $\mathcal{M}$ below λ; then in $V^{\mathrm{Col}(\omega,\mathbb{R})}$ there is an elementary $j \colon L(\mathbb{R}^*) \to L(\mathbb{R})^V$.*

Sketch of proof. Let $\langle \delta_n \mid n < \omega \rangle$ be a sequence of Woodin cardinals with limit λ, and let G_n be $\mathrm{Col}(\omega, \delta_n)$-generic over $\mathcal{M}$ and such that

$$\mathbb{R}^* = \bigcup_n (\mathbb{R} \cap \mathcal{M}[G_n]).$$

Working in $V[H]$, where H is $\mathrm{Col}(\omega, \mathbb{R})$-generic over V, the proof of 7.15 gives for each n an iteration map

$$i_n \colon \mathcal{M} \to \mathcal{P}_n, \text{ with } \mathrm{crit}(i_n) > \delta_n,$$

such that $\mathbb{R}^V$ is the set of reals of a symmetric collapse of $\mathcal{P}_n$ below $i_n(\lambda)$. Let

$$\Gamma = \{\alpha \in \mathrm{On} \mid \forall n (i_n(\alpha) = \alpha)\},$$

and

$$X = \{x \mid x \text{ is definable over } L(\mathbb{R}) \text{ from elements of } \mathbb{R}^* \cup \Gamma\}.$$

Since the i_n are iteration maps, Γ is a proper class. Now i_n induces an elementary embedding $i_n^* \colon \mathcal{M}[G_n] \to \mathcal{P}_n[G_n]$, and by the homogeneity of

[55] We shall give a purely inner-model-theoretic proof of this result immediately after 7.20.

the symmetric collapses we get, for all reals $\vec{x}$ in $\mathcal{M}[G_n]$, ordinals $\vec{\alpha}$, and formulae φ,

$$L(\mathbb{R}^*) \models \varphi[\vec{x}, \vec{\alpha}] \Leftrightarrow L(\mathbb{R})^V \models \varphi[\vec{x}, i_n(\vec{\alpha})].$$

It follows easily that

$$\mathbb{R} \cap X = \mathbb{R}^*.$$

Thus it suffices to show that $X \prec L(\mathbb{R})$, for then the inverse of the transitive collapse of X is the desired elementary embedding. So suppose

$$L(\mathbb{R}) \models \exists v \sigma[\vec{y}, \vec{\alpha}],$$

where $\vec{y} \in (\mathbb{R}^*)^{<\omega}$ and $\vec{\alpha} \in \Gamma^{<\omega}$. Pick n such that $\vec{y} \in M[G_n]$. Using the partial elementarity of i_n displayed above, we get

$$L(\mathbb{R}^*) \models \exists v \sigma[\vec{y}, \vec{\alpha}].$$

Since Γ is a proper class, we can take the witness v from $L(\mathbb{R}^*)$ to be definable over $L(\mathbb{R}^*)$ from z and $\vec{\beta}$, where $z \in \mathbb{R}^*$ and $\vec{\beta} \in \Gamma^{<\omega}$. Let $k \geq n$ be such that $z \in M[G_k]$; then the partial elementarity of i_k guarantees that there is a witness v to σ which is $L(\mathbb{R})$-definable from $z, \vec{y}, \vec{\beta}$, and $\vec{\alpha}$. This shows $X \prec L(\mathbb{R})$, as desired. ⊣

Although iterable class models with ω Woodin cardinals can compute $L(\mathbb{R})$ truth, they need not be correct for arbitrary statements about $L(\mathbb{R})$. We do have, however:

7.20 Theorem. *Let $\mathcal{M}$ be a proper class premouse such that $\mathcal{M} \models \eta$ is a limit of Woodin cardinals, for some $\eta < \omega_1^V$. Suppose $\mathcal{M}$ is $(\omega, \omega_1 + 1)$-iterable; then for any real $x \in \mathcal{M}$ and Σ_1 formula φ, containing perhaps a name $\dot{\mathbb{R}}$ for $\mathbb{R}$,*

$$(L(\mathbb{R}) \models \varphi[x]) \Longrightarrow (L(\mathbb{R})^{\mathcal{M}} \models \varphi[x]).$$

Proof. We shall assume $x \in M_\omega$; the argument in general is only slightly more complicated.

Fix an ω-small proper class premouse $\mathcal{N}$ whose extender sequence is an initial segment of that of $\mathcal{M}$, and such that there is a $\lambda \leq \eta$ such that λ is a limit of Woodin cardinals in $\mathcal{N}$. To see that there is such an $\mathcal{N}$, note that either $\mathcal{M}$ is ω-small, in which case we can take $\mathcal{N} = \mathcal{M}$, or $M_\omega^\sharp = \mathcal{J}_\alpha^{\mathcal{M}}$ for some α, in which case we can take $\mathcal{N} = M_\omega$. The iterability of $\mathcal{M}$ implies that of $\mathcal{N}$. From 7.19 we get some α such that $J_\alpha^{\mathcal{N}} \models \mathsf{ZF}^-$ + "there is a λ which is a limit of Woodin cardinals, and $L(\mathbb{R}^*) \models \varphi[x]$, where $\mathbb{R}^*$ is the set of reals of a symmetric collapse below λ". By taking a Skolem hull inside $\mathcal{N}$ and comparing the result with $\mathcal{N}$, we see that if $\bar{\alpha}$ is the least such α, then $\bar{\alpha}$ is countable in $\mathcal{N}$.

We claim that $\mathcal{J}^{\mathcal{N}}_{\bar\alpha}$ is (ω,ω_1+1)-iterable in $\mathcal{M}$. (This is why we dropped from $\mathcal{M}$ to $\mathcal{N}$.) For $\mathcal{Q} := \mathcal{J}^{\mathcal{N}}_{\bar\alpha+1}$ is properly small, and therefore by 7.13, has a (ω,ω_1^V)-iteration strategy Σ which is $\Sigma_1^{L(\mathbb{R})}(\{\mathcal{Q}\})$. By 7.19, and the homogeneity of $\mathrm{Col}(\omega,{<}\eta)$, $V_\eta^{\mathcal{M}}$ is closed under Σ, and $\Sigma \restriction V_\eta^{\mathcal{M}} \in \mathcal{M}$.

We can now run the construction of 7.15 in $\mathcal{M}[H]$, where H is $\mathcal{M}$-generic over $\mathrm{Col}(\omega,\mathbb{R})$. We obtain an iteration map

$$i\colon \mathcal{J}^{\mathcal{N}}_{\bar\alpha} \to \mathcal{P}$$

such that for some $\mathrm{Col}(\omega,{<}i(\bar\lambda))$-generic G over $\mathcal{P}$

$$\mathbb{R}^{\mathcal{M}} = \mathbb{R}^*_G.$$

Thus, for $\xi = \mathrm{On}^{\mathcal{P}}$, $L_\xi(\mathbb{R}^{\mathcal{M}}) \models \varphi[x]$, and hence $L(\mathbb{R}^{\mathcal{M}}) \models \varphi[x]$ since φ is Σ_1. ⊣

One can also prove this theorem using stationary tower forcing. [By 7.15 we have an iteration map $i\colon \mathcal{M} \to \mathcal{P}$ such that for some G which is $\mathrm{Col}(\omega,{<}i(\eta))$-generic over $\mathcal{P}$, $\mathbb{R}^V = \mathbb{R}^*_G$. Via stationary tower forcing over $\mathcal{P}$ one gets, for any α, a $\mathcal{P}$-generic elementary embedding $j\colon \mathcal{P} \to \mathcal{Q}$ with $\mathbb{R}^{\mathcal{Q}} = \mathbb{R}^*_G$ and $\alpha \in \mathrm{wfp}(\mathcal{Q})$. Then any Σ_1 fact true in $L(\mathbb{R})^V$ is true in some such $L(\mathbb{R})^{\mathcal{Q}}$, hence in $L(\mathbb{R})^{\mathcal{P}}$, and hence in $L(\mathbb{R})^{\mathcal{M}}$.] It is often the case that stationary tower forcing and genericity iterations can be made to do the same work.[56]

The argument of 7.20 yields another proof of the standard basis theorem for $\Sigma_1^{L(\mathbb{R})}$: every nonempty $\Sigma_1^{L(\mathbb{R})}$ set of reals has a $\Delta_1^{L(\mathbb{R})}$ member. For if φ defines our set over $L(\mathbb{R})$, then as in 7.20 we get an initial segment $\mathcal{Q}$ of M_ω of height $< \omega_1^{M_\omega}$ such that for some λ, $\mathcal{Q} \models \lambda$ is a limit of Woodin cardinals and it is forced in the symmetric collapse over $\mathcal{Q}$ below λ that $L(\mathbb{R}^*) \models \exists z\varphi(z)$. Working in M_ω, where λ is countable, we can find a generic object G for some $\mathrm{Col}(\omega,\delta)$, where $\delta < \lambda$, such that $\mathcal{Q}[G] \models \exists z(\mathrm{Col}(\omega,{<}\lambda) \Vdash \varphi(\hat z)^{L(\mathbb{R})})$. Picking such a $z \in \mathcal{Q}[G]$, we see from the iterability of $\mathcal{Q}$ in $V^{\mathrm{Col}(\omega,\mathbb{R})}$ that $L(\mathbb{R})^V \models \varphi[z]$. But z is in M_ω, hence z is $\mathrm{OD}^{L(\mathbb{R})}$. If we pick the least such z in the canonical wellorder of the reals of M_ω, we get that z is $\Delta_1^{L(\mathbb{R})}$.

The argument just given is closely related to the proof we gave that every nonempty Σ^1_3 set of reals has a member recursive in $M_1^\sharp$. One can extend the argument so as to show via inner model theory that the pointclass $\Sigma_1^{L(\mathbb{R})}$ has the scale property. (See [16] for the original proof, which used methods involving games and determinacy due to Moschovakis.) In recent unpublished

[56] One can also show using the scale property for $\Sigma_1^{L(\mathbb{R})}$ that if $\mathcal{M}$ is any model of set theory such that $\mathbb{R}^{\mathcal{M}}$ is countable, and every $\mathrm{OD}^{L(\mathbb{R})}(\{\mathbb{R}^{\mathcal{M}}\})$ set $X \subseteq \mathbb{R}^{\mathcal{M}}$ is in $\mathcal{M}$, then the conclusion of 7.20 holds. Combining this with the natural extension of 7.17 to sets of reals, we get yet another proof of 7.20.

work, Itay Neeman has found a general method which uses definability over mice to produce many pointclasses with the scale property. Neeman's work gives a new proof that Σ^1_{2n} and Π^1_{2n+1} have the scale property, for any $n \geq 1$. Neeman's work builds on earlier ideas of Woodin (unpublished, but see [43]), who found a purely inner-model-theoretic proof of the weaker fact that Σ^1_{2n} and Π^1_{2n+1} have the prewellordering property, for all n.

7.21 Corollary. *Suppose there are ω Woodin cardinals with a measurable above them all; then $M_\omega \models \mathbb{R}$ has a $\Delta_1^{L(\mathbb{R})}$ wellorder.*

Proof. By the reflection theorem,

$$x \in \mathrm{OD}^{L(\mathbb{R})} \Leftrightarrow \exists\alpha(x \in \mathrm{OD}^{L_\alpha(\mathbb{R})}).$$

So being $\mathrm{OD}^{L(\mathbb{R})}$ is a $\Sigma_1^{L(\mathbb{R})}$ property. Thus, by 7.18 and 7.20,

$$M_\omega \models \forall x \in \mathbb{R}(x \text{ is } \mathrm{OD}^{L(\mathbb{R})}).$$

The reals can now be wellordered in M_ω via their definitions in $L(\mathbb{R})^{M_\omega}$. ⊣

One can also prove 7.21 by showing that the natural wellorder of $\mathbb{R} \cap M_\omega$ given by the stages of construction is Δ_1 over $L(\mathbb{R})^{M_\omega}$. The proof of this is implicit in the arguments just given.

The author (unpublished) has shown that $M_\omega \models V = \mathrm{HOD}$. The proof builds on that of 7.21, but more is required.[57]

The correctness theorem 7.20 is best possible, in the sense that, if there are ω Woodin cardinals with a measurable cardinal above them all, then the statement "There is a wellorder of the reals" is a Σ_1 statement which is true in $L(\mathbb{R})^{M_\omega}$, but not true in $L(\mathbb{R})$. Another such statement is "Every real is ordinal definable over some $L_\alpha(\mathbb{R})$".

Iterations to make reals generic can be used to prove the generic absoluteness theorems one gets from stationary tower forcing. For example:

7.22 Theorem (Woodin). *Suppose that λ is a limit of Woodin cardinals, and there is a measurable cardinal above λ. Let G be $\mathbb{P}$-generic over V, where $|\mathbb{P}| < \lambda$, and let H be $\mathrm{Col}(\omega, \mathbb{R})^{V[G]}$-generic over $V[G]$; then in $V[G][H]$ there is an elementary*

$$j \colon L(\mathbb{R})^V \to L(\mathbb{R})^{V[G]}.$$

In particular, $L(\mathbb{R})^V$ is elementarily equivalent to $L(\mathbb{R})^{V[G]}$.

[57] One shows that the inductive definition of K from [44] relativises in such a way that one can define over M_ω its extender sequence in each interval between successive Woodin cardinals of M_ω.

Proof. Let $\langle (i_n, \mathcal{P}_n) \mid n < \omega \rangle$ be a genericity iteration of M_ω such that setting $\mathcal{P} =$ dirlim $\mathcal{P}_n$, we have that $\mathbb{R}^V$ can be realized as the reals $\mathbb{R}^*_K$ of a symmetric collapse of $\mathcal{P}$ below the sup of its Woodin cardinals. We get such an iteration in $V[G][H]$ from the proof of 7.15, and we have from that proof that each $\mathcal{P}_n$ is countable in V, and $\mathbb{R}^*_K = \bigcup_n \mathbb{R} \cap \mathcal{P}_n[K_n]$, where K_n is in V and $\mathrm{Col}(\omega, i_n \circ \ldots \circ i_0(\delta_n))$-generic over $\mathcal{P}_n$. (Here δ_n is the n^{th} Woodin cardinal of M_ω.) Applying 7.15 again, we have for each n an iteration map $j_n \colon \mathcal{P}_n \to \mathcal{Q}_n$ such that $\mathrm{crit}(j_n) > i_n \circ \ldots \circ i_0(\delta_n)$ and $\mathbb{R}^{V[G]}$ is the set of reals of a symmetric collapse of $\mathcal{Q}_n$. Note that j_n lifts to an elementary $\hat{j}_n$ from $\mathcal{P}_n[K_n]$ to $\mathcal{Q}_n[K_n]$. From the homogeneity of the two collapses it then follows that for any real $x \in \mathcal{P}_n[K_n]$, formula φ, and ordinal α, $L(\mathbb{R})^V \models \varphi[x, i_{n,\omega}(\alpha)]$ iff $L(\mathbb{R})^{V[G]} \models \varphi[x, j_n(\alpha)]$. As in the proof of 7.19, this means that if we let $X = \{\alpha \mid \forall n (j_n(\alpha) = \alpha = i_{n,\omega}(\alpha))\}$, and let j be the inverse of the transitive collapse of the hull in $L(\mathbb{R})^{V[G]}$ of $X \cup \mathbb{R}^V$, then $j \colon L(\mathbb{R})^V \to L(\mathbb{R})^{V[G]}$ elementarily. ⊣

One can also use genericity iterations to eliminate stationary tower forcing from the proof of $\mathsf{AD}^{L(\mathbb{R})}$, and in fact this can be done in several different ways. See for example [29], [28], and [39].

The connection between correctness of mice and definability of their iteration strategies extends much further. How much further it extends is one of the central open problems of inner model theory.

7.23 Definition. *Mouse capturing* is the following statement: for all $x, y \in \mathbb{R}$, x is ordinal definable from y if and only if for some (ω, ω_1)-iterable y-premouse $\mathcal{M}$, $x \in \mathcal{M}$.

Here a y-premouse is just like an ordinary premouse, except that we put y in at the bottom of its hierarchy. We have shown in this section that the existence of $M_\omega^\sharp$ implies that mouse capturing holds in $L(\mathbb{R})$. Results of Woodin show that $\mathsf{AD}^{L(\mathbb{R})}$ implies that mouse capturing holds in $L(\mathbb{R})$, and in fact, appropriately interpreted, it holds in every $J_\alpha(\mathbb{R})$. (See [13] and [40].) Woodin has also shown that mouse capturing holds in models of determinacy beyond $L(\mathbb{R})$: in any model of AD in which all ω_1-iterable mice are tame (see [40]), and even beyond that, in the minimal model of $\mathsf{AD}_\mathbb{R} + \mathsf{DC}$. This unpublished result is essentially the current frontier in this direction. It too has a local refinement: mouse capturing holds in any reasonably closed Wadge initial segment of the minimal model of $\mathsf{AD}_\mathbb{R} + \mathsf{DC}$. The capturing mice of the minimal model of $\mathsf{AD}_\mathbb{R} + \mathsf{DC}$ can be nontame, but they are below a Woodin limit of Woodin cardinals.

This leads us to the

Mouse Set Conjecture: Assume AD^+, and that there is no ω_1-iteration strategy for a premouse satisfying "there is a superstrong cardinal"; then mouse capturing holds.

AD^+ is a strengthening of AD which holds in all the models of AD we have constructed under large cardinal hypotheses. See for example [13, §8] for a precise definition. We might have stated the mouse set conjecture with AD as its hypothesis, but preferred to separate it from the open technical question as to whether AD implies AD^+.

It might be possible to drop the hypothesis that there is no ω_1-iteration strategy for a premouse satisfying "there is a superstrong cardinal" from the mouse set conjecture. One would presumably then have to enlarge the notion of mouse, so as to accomodate canonical models with supercompacts and more. The hypothesis that there is no ω_1-iteration strategy for a premouse satisfying "there is a superstrong cardinal" is a convenient way to say that we are in the initial segment of AD^+ models in which the capturing mice are premice in the sense of this paper.

The author believes it is unlikely that one can construct ω_1+1-iterable premice satisfying "there is a superstrong cardinal" under any hypothesis, even the hypothesis that there are superstrong cardinals, without proving the mouse set conjecture.

8. $\mathrm{HOD}^{L(\mathbb{R})}$ below Θ

Having characterized the reals in $\mathrm{HOD}^{L(\mathbb{R})}$ in terms of mice, it is natural to look for a similar characterization of the full model $\mathrm{HOD}^{L(\mathbb{R})}$. In this section we shall describe some work of the author ([42]) and W.H. Woodin (unpublished) which provides such a characterization.

The arguments of the last section give more in this direction than we stated there. Let $\mathcal{N}$ be the linear iterate of M_ω obtained by taking ultrapowers by the unique normal measure on the least measurable cardinal, and its images, ω_1^V times. Thus the least measurable cardinal of $\mathcal{N}$ is ω_1^V. One can show by the methods of the last section that $P(\omega_1^V)\cap\mathrm{HOD}^{L(\mathbb{R})} = P(\omega_1^V)\cap\mathcal{N}$. (See [41, section 4].) This clearly suggests that the whole of $\mathrm{HOD}^{L(\mathbb{R})}$ might be an iterate of M_ω. We shall show in this section that that is almost true.

8.1 Definition.

$$\Theta = \sup\{\, \alpha \mid \exists f \in L(\mathbb{R})(f\colon \mathbb{R} \to \alpha \text{ and } f \text{ is surjective})\}.$$

8.2 Definition.

$$\delta_1^2 = \sup\{\, \alpha \mid \exists f(f\colon \mathbb{R} \to \alpha \text{ and } f \text{ is surjective and } \mathbf{\Delta_1^{L(\mathbb{R})}}\}.$$

Standard notation would require that we write $\Theta^{L(\mathbb{R})}$ and $(\delta_1^2)^{L(\mathbb{R})}$ here, but since we shall only interpret the notions in question in $L(\mathbb{R})$, we have chosen to drop the superscripts. Similarly, we shall occasionally write HOD for $\mathrm{HOD}^{L(\mathbb{R})}$ in this section. We have nothing to say about HOD^V here.

We shall show that below δ^2_1, HOD is the direct limit of a certain class $\mathcal{F}$ of countable, iterable mice, under the iteration maps given by the comparison process. (One gets a typical element of $\mathcal{F}$ by iterating M_ω, then cutting the iterate off at a successor cardinal below its bottom Woodin cardinal.) The mice in $\mathcal{F}$ are properly small, so that $L(\mathbb{R})$ knows how to iterate them correctly. They are as "full" as possible, given this smallness condition. Fullness guarantees that in the comparison of two mice in $\mathcal{F}$, neither side drops along the branch leading to the final model, and thus we have iteration maps on both sides. The Dodd-Jensen Lemma guarantees that these maps commute, so that we can indeed form a direct limit. The whole direct limit system is definable over $L(\mathbb{R})$ in a way that insures its direct limit M_∞ is included in $\mathrm{HOD} \cap V_{\delta^2_1}$. On the other hand, we shall see that in the bigger universe $V^{\mathrm{Col}(\omega_1,\mathbb{R})}$ there is an iterate N of $\mathcal{M}_\omega$ such that M_∞ is just N cut off at the least cardinal κ which is β-strong for all β below the bottom Woodin cardinal of N. The correctness properties of N can then be used to show that $\mathrm{HOD} \cap V_{\delta^2_1} \subseteq M_\infty$.

The maps in our direct limit system will come from compositions of iteration trees. In order to make the Dodd-Jensen lemma applicable, we need to take care of some details regarding unique iterability. Let $\mathcal{M}$ be properly small. By $\mathcal{G}^*(\mathcal{M}, \lambda, \theta)$ we mean the variant of the iteration game $\mathcal{G}_\omega(\mathcal{M}, \lambda, \theta)$ in which player I is not allowed to drop at the beginning of a new round. That is, if $\mathcal{Q}$ is the model we get at the end of round α and q is its degree (with $\mathcal{Q} = \mathcal{M}$ and $q = \omega$ if $\alpha = 0$), then round $\alpha + 1$ of $\mathcal{G}^*(\mathcal{M}, \lambda, \theta)$ must be a play of $\mathcal{G}_q(\mathcal{M}, \theta)$. Let us call a play of $\mathcal{G}^*(\mathcal{M}, \lambda, \theta)$ in which II has not yet lost an *almost ω-maximal* iteration tree on $\mathcal{M}$; such a tree is just a linear composition of appropriately maximal trees, where "appropriately" means that the composition is itself maximal. Our proof of 7.13 gives

8.3 Lemma. *Let $\mathcal{M}$ be countable, properly small, and $\Game^{\mathbb{R}}\Pi^1_1$-iterable; then in $L(\mathbb{R})$, there is a unique winning strategy Σ for $\mathcal{G}^*(\mathcal{M}, \omega_1, \omega_1)$; moreover, Σ is $\Sigma_1^{L(\mathbb{R})}(\{\mathcal{M}\})$ definable, uniformly in $\mathcal{M}$.*

8.4 Definition. Let $\mathcal{M}$ be countable, properly small, and $\Game^{\mathbb{R}}\Pi^1_1$-iterable. An almost ω-maximal iteration tree on $\mathcal{M}$ is *correct* just in case it is played according to the unique winning strategy for II in $\mathcal{G}^*(\mathcal{M}, \omega_1, \omega_1)$. We say that $\mathcal{M}$ *iterates correctly* to $\mathcal{N}$ iff $\mathcal{N}$ is the last model of some correct $\mathcal{T}$ on $\mathcal{M}$ such that the branch $\mathcal{M}$-to-$\mathcal{N}$ of $\mathcal{T}$ has no drops.

From the last lemma we have at once:

8.5 Lemma. *The relations*

$$\{\,(\mathcal{M}, \mathcal{T}) \mid \mathcal{T} \textit{ is a correct tree on } \mathcal{M}\}$$

and

$$\{\,(\mathcal{M}, \mathcal{N}) \mid \mathcal{M} \textit{ iterates correctly to } \mathcal{N}\}$$

on HC are Σ_1 definable over $L(\mathbb{R})$.

There may in fact be more than one iteration tree witnessing that $\mathcal{M}$ iterates correctly to $\mathcal{N}$, but our proof of the Dodd-Jensen lemma, together with the uniqueness lemma 8.3 above, easily implies that all such trees give rise to the same iteration map $\pi\colon \mathcal{M} \to \mathcal{N}$. Because properly small $\mathcal{M}$ satisfy ZF^-, π is fully elementary.

8.6 Definition. A properly small mouse $\mathcal{M}$ is *full* iff whenever $\mathcal{M}$ iterates correctly to $\mathcal{N}$, A is a bounded subset of $\mathrm{On} \cap \mathcal{N}$, and A is ordinal definable over $L(\mathbb{R})$ from the parameter $\mathcal{N}$, then $A \in \mathcal{N}$.

It is clear that fullness is Π_1 definable over $L(\mathbb{R})$.[58] Since the $\mathrm{OD}^{L(\mathbb{R})}(\{\mathcal{N}\})$ sets are captured by mice, we can reformulate fullness in purely inner-model-theoretic terms.

8.7 Definition. We write $\mathcal{N} \trianglelefteq^* \mathcal{P}$ iff $\mathcal{N} = \mathcal{J}_\eta^{\mathcal{P}}$ for some cutpoint η of $\mathcal{P}$. In this case, we also call $\mathcal{N}$ a cutpoint of $\mathcal{P}$.

8.8 Lemma. *The following are equivalent:*

1. *$\mathcal{M}$ is full,*
2. *if $\mathcal{M}$ iterates correctly to $\mathcal{N}$, and $\mathcal{N} \trianglelefteq^* \mathcal{P}$, and $\mathcal{P}$ is $\Game^{\mathbb{R}}\Pi^1_1$-iterable above $\mathrm{On} \cap \mathcal{N}$,[59] then $\rho_\omega(\mathcal{P}) \geq \mathrm{On} \cap \mathcal{N}$.*

Proof. To see that (1)⇒(2), notice that the proof of 7.12 relativises, and thus if $\mathcal{P}$ and $\mathcal{N}$ are as in (2), then $\mathcal{P}$ is $\mathrm{OD}^{L(\mathbb{R})}$ from $\mathcal{N}$ as a parameter.

For the converse, suppose $\mathcal{N}$ is a correct iterate of $\mathcal{M}$, and let A be a bounded subset of $\lambda := \mathrm{On} \cap \mathcal{N}$ which is $\mathrm{OD}^{L(\mathbb{R})}$ from $\mathcal{N}$. We can modify the K^c construction by starting with $\mathcal{N}$ instead of $(V_\omega, \in, \emptyset, \emptyset)$ as our initial structure, and by adding only extenders with critical point strictly greater than λ. All ω-small structures we produce in such a construction are $\Game^{\mathbb{R}}\Pi^1_1$-iterable above λ, and so by (2) no such structure projects strictly below λ. It follows that $\mathcal{N}$ is an initial segment of all structures in the construction; indeed, λ is included in every core we take. Since $\mathcal{N}$ has a largest cardinal, λ is not the critical point of any extender in such a core, so that $\Game^{\mathbb{R}}\Pi^1_1$-iterability above λ is enough for comparison. We therefore get a proper class premouse $M_\omega(\mathcal{N})$ with ω Woodin cardinals which is iterable above λ and has $\mathcal{N}$ as a cutpoint. The proof of 7.18 relativises so as to show that $A \in M_\omega(\mathcal{N})$. But by (2), no level of $M_\omega(\mathcal{N})$ projects strictly below λ, and therefore $A \in \mathcal{N}$. ⊣

[58] Notice that a premouse which is not $\Game^{\mathbb{R}}\Pi^1_1$-iterable is vacuously full, since there are no correct trees on it. Of course, we are only interested in the full mice which *are* $\Game^{\mathbb{R}}\Pi^1_1$-iterable.

[59] This means that II wins the variant of $\mathcal{W}_\omega(\mathcal{N}, \omega)$ in which I is constrained to play only extenders with critical point above $\mathrm{On} \cap \mathcal{N}$.

We can now define our direct limit system. Set

$$\mathcal{F} := \{ \mathcal{M} \mid \mathcal{M} \text{ is properly small, } \Game^{\mathbb{R}}\Pi^1_1\text{-iterable, and full}\},$$

and for $\mathcal{M}, \mathcal{N}$ in $\mathcal{F}$, let

$$\mathcal{M} \prec^* \mathcal{N} \Leftrightarrow \exists \mathcal{P}(\mathcal{M} \text{ iterates correctly to } \mathcal{P} \text{ and } \mathcal{P} \trianglelefteq^* \mathcal{N}).$$

The Dodd-Jensen lemma implies that if $\mathcal{M} \prec^* \mathcal{N}$, then there is a unique $\mathcal{P} \trianglelefteq^* \mathcal{N}$ and a unique fully elementary $\pi\colon \mathcal{M} \to \mathcal{P}$ which is the iteration map given by some play of $\mathcal{G}^*(\mathcal{M}, \omega_1, \omega_1)$ according to the unique winning strategy for II. (There may be more than one such play giving rise to π.) We let

$$\pi_{\mathcal{M},\mathcal{N}} := \text{ unique correct iteration map from } \mathcal{M} \text{ to some } \mathcal{P} \trianglelefteq^* \mathcal{N}.$$

It is clear that $\mathcal{F}, \prec^*$, and the function $(\mathcal{M}, \mathcal{N}) \mapsto \pi_{\mathcal{M},\mathcal{N}}$ are $\mathrm{OD}^{L(\mathbb{R})}$.

8.9 Lemma. *The relation $\prec^*$ is transitive; moreover, if $\mathcal{M} \prec^* \mathcal{N} \prec^* \mathcal{P}$, then $\pi_{\mathcal{M},\mathcal{P}} = \pi_{\mathcal{N},\mathcal{P}} \circ \pi_{\mathcal{M},\mathcal{N}}$.*

Proof. Let $\mathcal{T}$ and $\mathcal{U}$ be correct trees witnessing that $\mathcal{M} \prec^* \mathcal{N}$ and $\mathcal{N} \prec^* \mathcal{P}$ respectively. Let $\mathcal{Q}$ be the last model of $\mathcal{T}$. Since $\mathcal{Q}$ is a cutpoint in $\mathcal{N}$, we can re-arrange $\mathcal{U}$ as an iteration tree $\mathcal{R}$ on $\mathcal{N}$ which uses only extenders from the image of $\mathcal{Q}$, followed by an iteration tree $\mathcal{S}$ on the last model $\mathcal{M}^{\mathcal{R}}_\alpha$ of $\mathcal{R}$ which uses no extenders from $i^{\mathcal{R}}_{0,\alpha}(\mathcal{Q})$. (We leave the details to the reader.) But then $\mathcal{T} \oplus \mathcal{R}$ witnesses that $\mathcal{M} \prec^* \mathcal{P}$. Moreover, the embedding given by $\mathcal{T} \oplus \mathcal{R}$ from $\mathcal{M}$ to its last model is just $i^{\mathcal{R}}_{0,\alpha} \circ \pi_{\mathcal{M},\mathcal{N}}$. Since $i^{\mathcal{R}}_{0,\alpha} = \pi_{\mathcal{N},\mathcal{P}} \restriction \mathcal{Q}$ by construction, the embedding given by $\mathcal{T} \oplus \mathcal{R}$ is $\pi_{\mathcal{N},\mathcal{P}} \circ \pi_{\mathcal{M},\mathcal{N}}$, as desired. $\dashv$

The comparison lemma and fullness imply that $\prec^*$ is directed. For suppose that $\mathcal{M}, \mathcal{N} \in \mathcal{F}$, and let $\mathcal{T}$ and $\mathcal{U}$ be the correct trees on $\mathcal{M}$ and $\mathcal{N}$ constituting their coiteration. Let $\mathcal{P}$ and $\mathcal{Q}$ be their respective last models, and suppose for example that $\mathcal{P} \trianglelefteq^* \mathcal{Q}$. (We can always take one more ultrapower so as to guarantee that $\trianglelefteq^*$, rather than just $\trianglelefteq$, holds between the last models.) From the comparison lemma we get that $\mathcal{M}$-to-$\mathcal{P}$ has no drops, so that $\mathcal{M}$ iterates correctly to $\mathcal{P}$. But $\mathcal{M}$ is full, so $\rho_\omega(\mathcal{Q}) \geq \mathrm{On} \cap \mathcal{P}$. Now if $\mathcal{N}$-to-$\mathcal{Q}$ drops, then letting κ be the extender used at the last drop, we have $\rho_\omega(\mathcal{Q}) \leq \kappa < \mathrm{On} \cap \mathcal{P}$. Thus $\mathcal{N}$-to-$\mathcal{Q}$ has no drops, so that $\mathcal{N}$ iterates correctly to $\mathcal{Q}$, and we have $\mathcal{M} \prec^* \mathcal{Q}$ and $\mathcal{N} \prec^* \mathcal{Q}$.

We wish to show that $\prec^*$ is countably directed, and for this it is most convenient to first relate the system $(\mathcal{F}, \prec^*)$ to a natural system $(\mathcal{F}^+, \prec^+)$ of iterates of M_ω.

8.10 Definition. Let Σ_0 be the unique winning strategy for II in $\mathcal{G}^*(M_\omega, \omega_1, \omega_1 + 1)$.

We can extend definition 8.4 from properly small mice to iterates of M_ω in the natural way. In general, let us say that $\mathcal{M}$ *iterates correctly* to $\mathcal{Q}$, or $\mathcal{Q}$ is a *correct iterate* of $\mathcal{M}$, iff there is a unique winning strategy for II in $\mathcal{G}^*(\mathcal{M}, \omega_1, \omega_1 + 1)$, and $\mathcal{Q}$ is the last model of a countable iteration tree $\mathcal{T}$ on $\mathcal{M}$ played according to this strategy such that the branch $\mathcal{M}$-to-$\mathcal{Q}$ of $\mathcal{T}$ does not drop.

8.11 Definition. We call an iteration tree on a premouse $\mathcal{M}$ which satisfies "there is a Woodin cardinal" δ_0*-bounded* if it uses only extenders from the image of $\mathcal{J}_\delta^{\mathcal{M}}$, where δ is the least Woodin cardinal of $\mathcal{M}$.

Thus a δ_0-bounded tree on $\mathcal{M}$ is just one which can be interpreted as a tree on $\mathcal{J}_\delta^{\mathcal{M}}$, where δ is the least Woodin cardinal of $\mathcal{M}$.

8.12 Definition. We set

$$\mathcal{F}^+ = \{ \mathcal{Q} \mid M_\omega \text{ iterates correctly to } \mathcal{Q} \text{ via a } \delta_0\text{-bounded tree}\},$$

and for $\mathcal{P}, \mathcal{Q} \in \mathcal{F}^+$, put

$$\mathcal{P} \prec^+ \mathcal{Q} \Leftrightarrow \mathcal{P} \text{ iterates correctly via a } \delta_0\text{-bounded tree to } \mathcal{Q}.$$

In this case, we let

$$\pi^+_{\mathcal{P},\mathcal{Q}} := \text{ unique iteration map from } \mathcal{P} \text{ to } \mathcal{Q}.$$

The uniqueness of the iteration map from $\mathcal{P}$ to $\mathcal{Q}$ follows from the Dodd-Jensen lemma.

The pair $(\mathcal{F}^+, \prec^+)$ is not lightface definable over $L(\mathbb{R})$, since from it we can define M_ω. It does happen to be definable over $L(\mathbb{R})$ from M_ω as a parameter, but this is of no use to us now. The function $(\mathcal{P}, \mathcal{Q}) \mapsto \pi^+_{\mathcal{P},\mathcal{Q}}$ does not even belong to $L(\mathbb{R})$. One can regard the system $(\mathcal{F}, \prec^*)$, with its maps, as an $L(\mathbb{R})$-definable approximation to the direct limit system $(\mathcal{F}^+, \prec^+)$, with its maps. We shall spell this out in more detail momentarily, but first we should verify:

8.13 Lemma. *The relation $\prec^+$ is transitive and countably directed; moreover, if $\mathcal{M} \prec^+ \mathcal{N} \prec^+ \mathcal{Q}$, then $\pi^+_{\mathcal{M},\mathcal{Q}} = \pi^+_{\mathcal{N},\mathcal{Q}} \circ \pi^+_{\mathcal{M},\mathcal{N}}$.*

Proof. Transitivity is obvious because we can compose iterations. (The situation here is a little simpler than it was with $\prec^*$.) The commutativity of the maps is clear.

Let $\mathcal{P}_i \in \mathcal{F}^+$ for all $i \in \omega$. Let $\mathcal{Q}_0 = M_\omega$, and given $\mathcal{Q}_i$, let $\mathcal{Q}_{i+1}$ be the last model of the iteration tree $\mathcal{T}_i$ on $\mathcal{Q}_i$ which results from comparing $\mathcal{Q}_i$ with $\mathcal{P}_i$, using their unique iteration strategies in both cases. Let $\mathcal{U}_i$ be the tree on $\mathcal{P}_i$ in this comparison. Clearly, neither $\mathcal{T}_i$ nor $\mathcal{U}_i$ drops along the branch to its last model, so $\mathcal{Q}_{i+1}$ is a correct iterate of both $\mathcal{Q}_i$ and $\mathcal{P}_i$.

Letting $\mathcal{Q}$ be the direct limit of the $\mathcal{Q}_i$, we have that for all i, $\mathcal{Q}$ is a correct iterate of $\mathcal{P}_i$. In order to show $\mathcal{P}_i \prec^+ \mathcal{Q}$ for all i, it is enough to show that all $\mathcal{T}_i$ and $\mathcal{U}_i$ are δ_0-bounded.

Suppose this is true for all $j < i$. Now we can regard M_ω as an initial segment of $M_\omega^\sharp$, and the latter is ω-sound and has Σ_1 projectum ω. The iteration strategy Σ_0 is the restriction to M_ω of a winning strategy in $\mathcal{G}^*(M_\omega^\sharp, \omega_1, \omega_1 + 1)$. Thus $\mathcal{P}_i$ and $\mathcal{Q}_i$ are initial segments of Σ_0-iterates $\mathcal{P}_i^*$ and $\mathcal{Q}_i^*$ of $M_\omega^\sharp$, and since the iterations are δ_0-bounded, each of $\mathcal{P}_i^*$ and $\mathcal{Q}_i^*$ is Σ_1-generated by the ordinals below its bottom Woodin cardinal. Now let $\mathcal{T}$ and $\mathcal{U}$ be the longest δ_0-bounded initial segments of $\mathcal{T}_i$ and $\mathcal{U}_i$, let $\mathcal{R}$ and $\mathcal{S}$ be their last models, and let $\mathcal{R}^*$ and $\mathcal{S}^*$ be the corresponding iterates of $M_\omega^\sharp$. Then $\mathcal{R}^*$ and $\mathcal{S}^*$ agree below their common value δ for the least Woodin cardinal (because this least Woodin is a cutpoint in each, and the last models of $\mathcal{T}_i$ and $\mathcal{U}_i$ so agree). Moreover, each is Σ_1-generated by δ, and they have a common iterate $\mathcal{Q}_{i+1}^*$ obtained from the rest of $\mathcal{T}_i$ and $\mathcal{U}_i$, which is above δ. It follows that $\mathcal{R}^* = \mathcal{S}^*$, so that $\mathcal{R} = \mathcal{S} = \mathcal{Q}_{i+1}$, and $\mathcal{T}_i$ and $\mathcal{U}_i$ are δ_0-bounded. ⊣

We now relate our two direct limit systems.

8.14 Lemma. *1. Let $\mathcal{T}$ be an iteration according to Σ_0 of M_ω with last model $\mathcal{Q}$, and suppose that M_ω-to-$\mathcal{Q}$ does not drop. If η is a successor cardinal of $\mathcal{Q}$ below its bottom Woodin cardinal, then $\mathcal{J}_\eta^{\mathcal{Q}}$ is full, and therefore in $\mathcal{F}$.*

2. Let $\mathcal{P} \in \mathcal{F}$, and let $\mathcal{M}$ be a correct iterate of M_ω; then there is a correct iterate $\mathcal{Q}$ of $\mathcal{M}$, given by a δ_0-bounded iteration tree, such that $\mathcal{P} \prec^* \mathcal{J}_\eta^{\mathcal{Q}}$ for some successor cardinal cutpoint η of $\mathcal{Q}$ below its bottom Woodin cardinal.

3. If $\mathcal{P} \prec^+ \mathcal{Q}$, and $\mathcal{M}$ is a cutpoint of $\mathcal{P}$ at some successor cardinal below its bottom Woodin cardinal, and $\mathcal{N} = \pi^+_{\mathcal{P},\mathcal{Q}}(\mathcal{M})$, then $\mathcal{M} \prec^* \mathcal{N}$, and $\pi_{\mathcal{M},\mathcal{N}} = \pi^+_{\mathcal{P},\mathcal{Q}} \restriction \mathcal{M}$.

Proof. Part 1 follows easily from 8.8: suppose $\mathcal{J}_\eta^{\mathcal{Q}}$ iterates correctly to $\mathcal{N}$, and $\mathcal{N} \trianglelefteq^* \mathcal{P}$, where $\mathcal{P}$ is ω-small and $\Game^{\mathbb{R}}\Pi^1_1$-iterable above $\mathrm{On} \cap \mathcal{N} := \lambda$. We must show $\rho_\omega(\mathcal{P}) \geq \lambda$. Now, since η is a successor cardinal cutpoint of $\mathcal{Q}$, our correct iteration $\mathcal{J}_\eta^{\mathcal{Q}}$-to-$\mathcal{N}$ lifts to an iteration $\mathcal{Q}$-to-$\mathcal{R}$ according to Σ_0; moreover λ is a successor cardinal cutpoint of $\mathcal{R}$. We can now compare $\mathcal{P}$ and $\mathcal{R}$, and the comparison is above λ since it is a cutpoint of each. If $\rho_\omega(\mathcal{P}) < \lambda$, then we must have $\mathcal{P} \trianglelefteq \mathcal{R}$, but this contradicts the fact that λ is a cardinal of $\mathcal{R}$.

For 2, we simply compare $\mathcal{P}$ with $\mathcal{M}$, forming iterations according to the unique $(\omega, \omega_1 + 1)$-iteration strategy on both sides. Since $\mathcal{P}$ is properly small, it must iterate into an initial segment $\mathcal{R}$ of the last model $\mathcal{Q}$ on the

$\mathcal{M}$ side, with no dropping from $\mathcal{P}$ to $\mathcal{R}$. Since $\mathcal{P}$ is full, $\mathcal{M}$-to-$\mathcal{Q}$ does not drop. Since $\mathcal{R}$ is properly small and full, it must have the form described.

For 3, notice that the iteration from $\mathcal{P}$ to $\mathcal{Q}$ can be factored so as to give an iteration from $\mathcal{M}$ to $\mathcal{N}$ because $\mathcal{M}$ is a cutpoint in $\mathcal{P}$. The uniqueness of the iteration strategies gives the rest. ⊣

8.15 Definition. We let M_∞ be the direct limit of $(\mathcal{F}, \prec^*)$ under the $\pi_{\mathcal{M},\mathcal{N}}$, and M_∞^+ be the direct limit of $(\mathcal{F}^+, \prec^+)$ under the $\pi_{\mathcal{M},\mathcal{N}}^+$, transitively collapsed in each case.

Since $\prec^+$ is countably directed, M_∞^+ is wellfounded, so we can regard it as transitive. But 8.14 shows that M_∞ is an initial segment of M_∞^+, so it too is wellfounded. In fact

8.16 Corollary. *Let δ be the least Woodin cardinal of M_∞^+, and let $\kappa < \delta$ be the least cardinal of M_∞^+ which is $< \delta$-strong in M_∞^+; then $M_\infty = \mathcal{J}_\kappa^{M_\infty^+}$.*

Proof. By 8.14, the set of all $\mathcal{M}$ which are cutpoints of some $\mathcal{Q} \in \mathcal{F}^+$ at a successor cardinal below its bottom Woodin cardinal (and hence below the least cardinal strong to its bottom Woodin) are cofinal in $(\mathcal{F}, \prec^*)$; moreover, the π^+ maps act on these $\mathcal{M}$ the same way that the π maps act. Thus M_∞ is the direct limit of all such $\mathcal{M}$ under the π^+ maps. Clearly, this direct limit is M_∞^+ cut at the sup of all its successor cardinal cutpoints below δ. That sup is just κ. ⊣

We shall now show that the ordinal height of M_∞ is δ_1^2.

8.17 Definition. Let $\mathcal{M}$ be a premouse, $\varphi(v)$ a Σ_1 formula, and $x \in \mathbb{R}$. We call $\mathcal{M}$ a *(φ, x)-witness* just in case $\mathcal{M}$ has ω Woodin cardinals with supremum λ, and for some set $\mathbb{R}^*$ of reals of a symmetric collapse below λ over $\mathcal{M}$, we have $x \in \mathbb{R}^*$ and $L_\alpha(\mathbb{R}^*) \models \varphi[x]$, where $\alpha = \mathrm{On} \cap \mathcal{M}$.

8.18 Lemma. *Let φ be Σ_1 and $x \in \mathbb{R}$. The following are equivalent:*

1. *$L(\mathbb{R}) \models \varphi[x]$,*
2. *There is an $(\omega, \omega_1 + 1)$-iterable (φ, x)-witness,*
3. *$\exists \mathcal{M} \in \mathcal{F} \exists \beta (\mathcal{J}_\beta^{\mathcal{M}}$ is a (φ, x)- witness.)*

Proof. For (3) ⇒ (2), notice that $\mathcal{J}_\beta^{\mathcal{M}}$ is $(\omega, \omega_1 + 1)$-iterable, because $\mathcal{M}$ is. For (2) ⇒ (1), we can easily adapt the proofs of 7.19 and 7.15 to mice of set size with ω Woodin cardinals. We get, in some generic extension of V, an iterate of our witness $\mathcal{P}$ which has a symmetric collapse of the form $L_\alpha(\mathbb{R}^V)$ such that $L_\alpha(\mathbb{R}^V) \models \varphi[x]$. Since φ is Σ_1, this implies that $L(\mathbb{R}^V) \models \varphi[x]$.

We now prove (1) ⇒ (3). Let $\mathcal{Q}$ be a correct iterate of M_ω such that x is generic over $\mathcal{Q}$ for the extender algebra at its least Woodin cardinal δ. Now

$\mathcal{Q}$ is a (φ, x)-witness by 7.19, but it is not an initial segment of any $\mathcal{M} \in \mathcal{F}$. We must therefore take some Skolem hulls.

Since φ is Σ_1, we can fix α such that $\mathcal{J}_\alpha^{\mathcal{Q}}$ is a (φ, x)-witness. Let $G_x^{\mathcal{Q}}$ be the generic object on the extender algebra of $\mathcal{Q}$ at δ determined by x. (That is, $[\psi] \in G_x^{\mathcal{Q}}$ iff $x \models \psi$.) We then have some $p \in G_x^{\mathcal{Q}}$ such that

$$\mathcal{J}_\alpha^{\mathcal{Q}} \models \exists\lambda[\lambda \text{ is a limit of Woodins and } p \Vdash (1 \Vdash (L(\mathbb{R}^*) \models \varphi[\check{\dot{x}}]))],$$

where the first forcing is the extender algebra, the second is the symmetric collapse, and $\dot{x}$ is the canonical name for the real determined by the extender algebra generic. This is a Σ_1 fact about p and δ, so we may assume that $\mathcal{J}_\alpha^{\mathcal{Q}}$ is Σ_1-generated by $\delta \cup \{\delta\}$. (The Σ_1 hull of these parameters collapses to an initial segment of $\mathcal{Q}$ by a simple comparison argument. The extender algebra is definable over $\mathcal{J}_\delta^{\mathcal{Q}}$, hence contained in the hull, so that $G_x^{\mathcal{Q}}$ is still generic over the collapse of the hull.)

Now, working in $\mathcal{Q}[x]$, where δ is still a regular cardinal, we can find an η and an elementary submodel $Y \prec \mathcal{J}_\eta^{\mathcal{Q}}[x]$ such that $\delta, \alpha, p, x \in Y$ and $Y \cap \delta \in \delta$. Let N be the transitive collapse of Y, and $\mathcal{P}$ be the image of $\mathcal{J}_\alpha^{\mathcal{Q}}$ under the collapse. Letting $\bar{\delta} = Y \cap \delta$, we have that $\mathcal{P}$ is iterable, Σ_1 projects to $\bar{\delta}$, and agrees with $\mathcal{Q}$ below $\bar{\delta}$. It follows that $\mathcal{P}$ is an initial segment of $\mathcal{J}_{\bar{\delta}}^{\mathcal{Q}}$, by comparison, and therefore $\mathcal{P}$ is an initial segment of some $\mathcal{M} \in \mathcal{F}$. Since the property of being a (φ, x)-witness is first-order over $\mathcal{J}_\eta^{\mathcal{Q}}[x]$, we have that $\mathcal{P}$ is a (φ, x)-witness, as desired. ⊣

8.19 Lemma. $\mathrm{On} \cap M_\infty = \delta_1^2$.

Proof. A direct computation shows that $\mathrm{On} \cap M_\infty \leq \delta_1^2$. For let $\alpha \in \mathrm{On} \cap M_\infty$, and fix $\mathcal{M} \in \mathcal{F}$ so that $\pi_{\mathcal{M},\infty}(\bar{\alpha}) = \alpha$ for some $\bar{\alpha}$. Let $\mathcal{G} := \{\mathcal{P} \mid \mathcal{M}$ iterates correctly to $\mathcal{P}\}$. Then $\mathcal{G} \subseteq \mathcal{F}$, and one can easily check that $\mathcal{G}$ is $\Delta_1^{L(\mathbb{R})}(\{\mathcal{M}\})$. Also, the relation R is $\Delta_1^{L(\mathbb{R})}(\{\mathcal{M}\})$, where

$$R(\langle \mathcal{P}, \bar{\beta}\rangle, \langle \mathcal{Q}, \bar{\gamma}\rangle) \Leftrightarrow (\mathcal{P}, \mathcal{Q} \in \mathcal{G} \wedge \bar{\beta} \in \mathrm{On}^{\mathcal{P}} \wedge \bar{\gamma} \in \mathrm{On}^{\mathcal{Q}} \wedge \pi_{\mathcal{P},\infty}(\bar{\beta}) \leq \pi_{\mathcal{Q},\infty}(\bar{\gamma})).$$

This is because we can check whether $R(\langle \mathcal{P}, \bar{\beta}\rangle, \langle \mathcal{Q}, \bar{\gamma}\rangle)$ by comparing $\mathcal{P}$ with $\mathcal{Q}$, using their unique $\Sigma_1^{L(\mathbb{R})}(\{\mathcal{P}, \mathcal{Q}\})$ iteration strategies. Since every $\beta < \alpha$ is of the form $\pi_{\mathcal{P},\infty}(\bar{\beta})$ for some $\mathcal{P} \in \mathcal{G}$, there is a $\Delta_1^{L(\mathbb{R})}(\{\mathcal{M}\})$ prewellorder of H_{ω_1} of order type at least α. Thus $\alpha \leq \delta_1^2$.

Now suppose $\mathrm{On} \cap M_\infty < \delta_1^2$. Since M_∞ can be coded simply by a subset of $\mathrm{On} \cap M_\infty$, we have by the Coding Lemma ([27, Chapter 7]) that for some real z, M_∞ is coded by a $\Delta_1^{L(\mathbb{R})}(\{z\})$ set of reals. But 8.18 implies that the universal $\Sigma_1^{L(\mathbb{R})}$ set of reals is projective in any set of reals coding M_∞, for we have, for all Σ_1 formulae φ and reals x:

$$L(\mathbb{R}) \models \varphi[x] \Leftrightarrow \exists\mathcal{M}\exists\beta\exists\pi(\mathcal{M} \text{ is a } (\varphi, x) \text{ witness and } \pi\colon \mathcal{M} \to \mathcal{J}_\beta^{M_\infty}).$$

(The left-to-right direction follows at once from (1) $\Rightarrow$ (3) of 8.18, and the right-to-left direction follows from (2) $\Rightarrow$ (1) of 8.18.) This implies that the universal $\Sigma_1^{L(\mathbb{R})}$ set of reals is $\Delta_1^{L(\mathbb{R})}(\{z\})$, a contradiction. $\dashv$

8.20 Theorem. $\mathrm{HOD} \cap V_{\delta_1^2} = M_\infty \cap V_{\delta_1^2}$.

Proof. We have shown that $\mathcal{F}$,$\prec^*$, and the function $(\mathcal{M}, \mathcal{N}) \mapsto \pi_{\mathcal{M},\mathcal{N}}$ are definable over $L(\mathbb{R})$. It follows that $M_\infty \in \mathrm{HOD}$. It is enough, then, to show that every bounded subset A of δ_1^2 which is $\mathrm{OD}^{L(\mathbb{R})}$ is in M_∞. (Note here that δ_1^2 is strongly inaccessible in HOD, by work of H.Friedman and Y. Moschovakis.) So fix such an A. By the reflection theorem, we can fix a Σ_1 formula $\varphi(v_0, v_1)$ and an ordinal $\beta < \delta_1^2$ such that $A \subseteq \beta$, and for all $\alpha < \beta$

$$\alpha \in A \Leftrightarrow L(\mathbb{R}) \models \varphi[\alpha, \beta].$$

Since $M_\infty = \mathcal{J}_{\delta_1^2}^{M_\infty^+}$, and δ_1^2 is a cardinal of M_∞^+ by 8.16, it will be enough to show that $A \in M_\infty^+$. Let λ be the sup of the Woodin cardinals of M_∞^+. By asking what is true in its own symmetric collapse below λ, M_∞^+ will be able to answer membership questions about A. More precisely, let $\bar{\varphi}(u)$ be the Σ_1 formula:

"$u \in \mathbb{R}$ codes $(\mathcal{N}, \gamma, \delta)$ where $\mathcal{N} \in \mathcal{F}$ and $\varphi(\pi_{\mathcal{N},\infty}(\gamma), \pi_{\mathcal{N},\infty}(\delta))$".

Let η be a successor cardinal of M_∞ above β, and for each $\alpha < \beta$ let τ_α be a term for a real in the symmetric collapse below λ over M_∞^+ such that for all generic objects H for this collapse

$$\tau_\alpha^H \text{ codes } (\mathcal{J}_\eta^{M_\infty}, \alpha, \beta).$$

The map $\alpha \mapsto \tau_\alpha$, if chosen naturally, is definable over M_∞^+ from η and β. We claim that for all $\alpha < \beta$,

$$\alpha \in A \Leftrightarrow M_\infty^+ \models (1 \Vdash \bar{\varphi}(\tau_\alpha)^{L(\mathbb{R}^*)}).$$

It clearly suffices to prove this claim.

Fix $\alpha < \beta$. By 8.14, we can find $\mathcal{Q} \in \mathcal{F}^+$ and ordinals $\bar{\eta}, \bar{\beta}$, and $\bar{\alpha}$ in $\mathcal{Q}$ such that

$$\pi_{\mathcal{Q},\infty}^+(\langle \bar{\eta}, \bar{\beta}, \bar{\alpha} \rangle) = \langle \eta, \beta, \alpha \rangle.$$

Let $\bar{\tau}_{\bar{\alpha}}$ be definable over $\mathcal{Q}$ from $\bar{\eta}$, $\bar{\beta}$, and $\bar{\alpha}$ the way τ_α was from η,β, and α over M_∞^+, so that for any H generic over $\mathcal{Q}$ for the symmetric collapse below the sup $\bar{\lambda}$ of its Woodin cardinals, $\bar{\tau}_{\bar{\alpha}}^H$ is a real coding $(\mathcal{J}_{\bar{\eta}}^{\mathcal{Q}}, \bar{\alpha}, \bar{\beta})$. We have

$$\begin{aligned}
\alpha \in A \quad &\Leftrightarrow \quad L(\mathbb{R}) \models \varphi[\alpha, \beta] \\
&\Leftrightarrow \quad \forall H (H \text{ is } \mathrm{Col}(\omega, {<}\bar{\lambda}), \mathcal{Q}\text{- generic } \Rightarrow L(\mathbb{R}_H^*) \models \bar{\varphi}(\bar{\tau}_{\bar{\alpha}}^H)) \\
&\Leftrightarrow \quad \mathcal{Q} \models (1 \Vdash \bar{\varphi}(\bar{\tau}_{\bar{\alpha}})^{L(\mathbb{R}^*)}) \\
&\Leftrightarrow \quad M_\infty^+ \models (1 \Vdash \bar{\varphi}(\tau_\alpha)^{L(\mathbb{R}^*)}).
\end{aligned}$$

The second equivalence above follows from the correctness of $L(\mathbb{R}^*_H)$ and the fact that $\pi_{\mathcal{M},\infty}(\langle\bar{\alpha},\bar{\beta}\rangle) = \langle\alpha,\beta\rangle$, for $\mathcal{M} = \mathcal{J}^{\mathcal{Q}}_{\bar{\eta}}$; this is true because the π and π^+ maps agree.

The displayed equivalences contain our claim. This completes the proof. $\dashv$

A different proof of 8.20 is sketched in [42]. One shows that in $L[M_\infty]$ there is a tree T on $\omega \times \delta^2_1$ projecting to the universal $\Sigma_1^{L(\mathbb{R})}$ set of reals, and that this tree is enough like the tree of a $\Sigma_1^{L(\mathbb{R})}$ scale that, by arguments of Martin, Becker, and Kechris ([4]), $\mathrm{HOD} \cap V_{\delta^2_1} \subseteq L[T]$. The tree T attempts to verify $\varphi(x)$ by building a (φ, x)-witness and embedding it into M_∞. In this version of the proof, the Dodd-Jensen lemma corresponds nicely to the lower semi-continuity of a certain semi-scale.

Assuming sufficient determinacy, and given a pointclass Γ which resembles Π^1_1 in a certain technical sense, Moschovakis has defined a submodel of HOD corresponding to Γ-definability which he calls H_Γ. See [27, 8G]. Becker and Kechris show in [4] that $H_\Gamma = L[T]$, whenever T is the tree of a Γ-scale on a universal Γ set. The argument of the last paragraph actually shows that $L[M_\infty] = H_\Gamma$, where $\Gamma = \Sigma_1^{L(\mathbb{R})}$. The argument generalizes to many other Γ, with M_∞ replaced by a direct limit of mice whose iteration strategies and degree of correctness match Γ appropriately. This gives

8.21 Theorem. *Assume $\mathsf{AD}^{L(\mathbb{R})}$, and let Γ be either Π^1_n for n odd, or the pointclass $\Sigma_1^{L(\mathbb{R})}$; then H_Γ is an extender model.*

The theorem probably holds for all Γ resembling Π^1_1, but this has not been fully proved.

One immediate consequence of 8.20 is

8.22 Corollary. HOD $\models$ GCH .

Proof. By 8.20, the GCH holds in HOD at all $\alpha < \delta^2_1$. But Woodin (unpublished) has shown that δ^2_1 is $<\Theta$-strong in HOD, and thus the GCH holds in HOD at all $\alpha < \Theta$. Since HOD $= L(P)$ for some $P \subseteq \Theta$,[60], the GCH holds in HOD at all α. $\dashv$

We emphasize that HOD $=$ HOD$^{L(\mathbb{R})}$ in the statement of 8.22, and that $\mathsf{AD}^{L(\mathbb{R})}$ is a tacit hypothesis there.[61] Whether $\mathsf{AD}^{L(\mathbb{R})}$ implies that the GCH holds in HOD was open for some time, and various partial results were

[60] This is another result of Woodin; P is a version of the Vopenka algebra which can add $\mathbb{R}$ to HOD.

[61] The proof we have given used a bit more, namely, that $M^\sharp_\omega$ exists and is (ω, ω_1+1)-iterable in $V^{\mathrm{Col}(\omega,\mathbb{R})}$. The proof can be made to work under the weaker hypothesis $\mathsf{AD}^{L(\mathbb{R})}$, however. The key is to prove the existence of mouse-witnesses, as stated in 8.18, assuming only $\mathsf{AD}^{L(\mathbb{R})}$. This is a result of Woodin. The method behind the original proof is described in [13]; there is another proof using the core model induction method.

obtained using the methods of "neo-classical" descriptive set theory, such as games and scales.[62] Our proof of 8.20 is evidence of what inner model theory can contribute to this mix. One gets not just GCH, of course, but the other consequences of fine structure theory, such as $\Diamond$ and $\Box$.

It is natural to ask whether the full $\mathrm{HOD}^{L(\mathbb{R})}$ is a core model. Building on the proof of 8.20, W.H. Woodin has shown that this is essentially, but not literally, the case. We shall state Woodin's results, although it is beyond the scope of this article to prove them. The first is

8.23 Theorem (Woodin). $M_\infty^+ \subseteq \mathrm{HOD}$*; moreover, the least Woodin cardinal of* M_∞^+ *is* Θ*, and* $V_\Theta \cap \mathrm{HOD} = V_\Theta \cap M_\infty^+$.

Since the full HOD is of the form $L(P)$ for some $P \subseteq \Theta$, M_∞^+ is not far from the full HOD. What is missing can be represented in inner-model-theoretic terms. Let X be the class of all δ_0-bounded iteration trees on M_∞^+ which belong to M_∞^+ and are satisfied to have cardinality strictly less than the sup of the Woodin cardinals in M_∞^+. There is a unique iteration strategy for M_∞^+; let us call it Σ.[63] Let $\Sigma^* := \{ (\mathcal{T}, \alpha) \mid \mathcal{T} \in X$ and $\mathcal{T}$ is according to Σ and $\mathrm{lh}(\mathcal{T})$ is a limit ordinal, and $\alpha \in \Sigma(\mathcal{T})\}$. We then have

8.24 Theorem (Woodin). $\mathrm{HOD} = M_\infty^+[\Sigma^*]$.

Woodin has obtained results on HOD^M for M a model of AD larger than $L(\mathbb{R})$; for example, the Mouse Set Conjecture implies that $\mathrm{HOD}^M \restriction \Theta_0^M$ is an extender model. (Here Θ_0 is the supremum of the lengths of prewellorders of $\mathbb{R}$ which are ordinal definable from a real. If $V = L(\mathbb{R})$, then $\Theta_0 = \Theta$.) Woodin has also obtained an analysis of the full HOD^M analogous to that in 8.24. See [13, §8] for something on these results, on local forms of 8.24, and on open questions in the area.

We conclude with some applications of these results on HOD.

8.25 Lemma. *Let* $\kappa < \Theta$ *and suppose* $\mathrm{HOD} \models \kappa$ *is regular; then exactly one of the following holds:*

1. $\mathrm{HOD} \models \kappa$ *is measurable,*
2. $\mathrm{cf}^{L(\mathbb{R})}(\kappa) = \omega$.

Proof. Let $\mathcal{Q} \in \mathcal{F}^+$ and $\bar{\kappa} \in \mathcal{Q}$ be such that $\pi^+_{\mathcal{Q},\infty}(\bar{\kappa}) = \kappa$. Thus $\mathcal{Q} \models \bar{\kappa}$ is regular.

Suppose first that $\bar{\kappa}$ is not measurable in $\mathcal{Q}$. Now since $\pi^+_{\mathcal{Q},\infty}$ is essentially an iteration map, it is continuous at all regular, non-measurable cardinals

[62] For example, Becker ([3]) showed that GCH holds in HOD at all $\alpha < \omega_1^V$.

[63] Granted ω Woodins plus a measurable above in V, Σ_0 prolongs uniquely to trees in $V^{\mathrm{Col}(\omega,\mathbb{R})}$.

of $\mathcal{Q}$. (In $V^{\mathrm{Col}(\omega,\mathbb{R})}$ we can find a $\prec^+$-increasing ω sequence starting with $\mathcal{Q}$ and cofinal in $\prec^+$. The map $\pi^+_{\mathcal{Q},\infty}$ is just the iteration map coming from composing iteration trees witnessing the $\prec^+$ relations along this sequence. So $\pi^+_{\mathcal{Q},\infty}$ is an iteration map in $V^{\mathrm{Col}(\omega,\mathbb{R})}$, which is good enough.) In particular, $\pi^+_{\mathcal{Q},\infty}$"$\bar{\kappa}$ is cofinal in κ. Since $\bar{\kappa}$ is below the least Woodin cardinal of $\mathcal{Q}$ by 8.23, and hence countable, $\mathrm{cf}^V(\kappa) = \omega$. But clearly, V and $L(\mathbb{R})$ have the same ω-sequences of ordinals $< \mu$, whenever $\mu < \Theta$. Thus $\mathrm{cf}^{L(\mathbb{R})}(\kappa) = \omega$. Note also that we have in this case that κ is not measurable in HOD.

Suppose next that $\bar{\kappa}$ is measurable in $\mathcal{Q}$. It is clear then that κ is measurable in HOD, and we need only show that $\mathrm{cf}^V(\kappa) > \omega$. Let X be a countable subset of κ. By the countable directedness of $\prec^+$, we can find an $\mathcal{R} \in \mathcal{F}^+$ such that $\mathcal{Q} \prec^+ \mathcal{R}$ and $X \subseteq \mathrm{dom}(\pi^+_{\mathcal{R},\infty})$. Let $\hat{\kappa} = \pi^+_{\mathcal{Q},\mathcal{R}}(\bar{\kappa})$, and let $\mathcal{S}$ be the ultrapower of $\mathcal{R}$ by some normal measure on $\hat{\kappa}$. Then $\mathcal{R} \prec^+ \mathcal{S}$, and it is easy to see that $X \subseteq \pi^+_{\mathcal{S},\infty}(\hat{\kappa}) < \kappa$, so that X is bounded in κ, as desired. ⊣

We remark that the restriction of 8.25 to ordinals $\kappa < \delta^2_1$ requires only 8.20, rather than the full 8.23.

It follows from 8.25 that all successor cardinals of HOD below Θ have cofinality ω in $L(\mathbb{R})$, or equivalently, V. This is also true if we replace HOD by HOD_x, the sets hereditarily ordinal definable over $L(\mathbb{R})$ from x, for x a real. This is because our results relativise routinely to arbitrary reals x; we simply extend the notion of mouse by requiring that x be put in $\mathcal{J}^{\vec{E}}_0(x)$. The relativisation of our dichotomy 8.25 gives the following result, known as the "boldface GCH" for $L(\mathbb{R})$.

8.26 Theorem. *Assume* AD *and* $V = L(\mathbb{R})$*; then for any* $\kappa < \Theta$*, every wellordered family of subsets of* κ *has cardinality at most* κ*.*

Proof. If not, we have some $A \subseteq \kappa^+$ which codes up a sequence of κ^+ distinct subsets of κ. Since $V = L(\mathbb{R})$, we can find a real x such that $A \in \mathrm{HOD}_x$. We have just observed that $(\kappa^+)^{\mathrm{HOD}_x} < \kappa^+$, by the relativisation of our dichotomy 8.25 to x. But then A witnesses that GCH fails in HOD_x, contrary to the relativised version of 8.22. ⊣

Although we have quoted 8.22 in our proof of 8.26, we really only need 8.20. This is because "the boldface GCH fails at κ" is a $\Sigma^{L(\mathbb{R})}_1$ assertion about κ. Since $L_{\delta^2_1}(\mathbb{R})$ is a Σ_1 elementary substructure of $L(\mathbb{R})$, if the boldface GCH fails at some κ, it fails at some $\kappa < \delta^2_1$. But we can use 8.20 in the proof of 8.26 to see that this is not the case.

Finally, if $\kappa < \Theta$ is regular in $L(\mathbb{R})$, then by our dichotomy result, κ is measurable in HOD, and in fact, κ is measurable in HOD_x for all reals x. We can put the order zero measures on κ from the various HOD_x together, and we obtain:

8.27 Theorem. *Assume* AD *and* $\mathsf{V} = \mathsf{L}(\mathbb{R})$*; then for any regular* $\kappa < \Theta$*, the* ω*-closed unbounded filter on* κ *is a* κ*-complete, normal ultrafilter on* κ*. Thus all regular cardinals below* Θ *are measurable.*

Proof. For any real x, let μ_x be the order zero measure on κ of HOD_x, that is, the unique measure giving the set of measurable cardinals measure zero. There is such a measure by 8.25; it is unique because HOD_x is a core model. It will be enough to show that there is an ω-closed, unbounded set C which generates μ_x, in the sense that for all $A \subseteq \kappa$ such that $A \in \mathrm{HOD}_x$,

$$A \in \mu_x \Rightarrow \exists \alpha < \kappa (C \setminus \alpha \subseteq A).$$

For this implies that the union over x of the μ_x is just the ω-closed unbounded filter on κ. Since every $A \subseteq \kappa$ is in some HOD_x, this union is an ultrafilter. Since every $f \colon \kappa \to \kappa$ is in some HOD_x, that ultrafilter is normal, and hence κ-complete.

We now construct the desired generating set for μ_x. Let us assume $x = 0$, so that we can use our earlier notation for the direct limit system giving $\mathrm{HOD}_x = \mathrm{HOD}$; the general case is only notationally different. Fix $\mathcal{Q} \in \mathcal{F}^+$ such that $\kappa \in \mathrm{ran}(\pi^+_{\mathcal{Q},\infty})$. Let

$$C := \{ \alpha \mid \mathrm{cf}(\alpha) = \omega \text{ and } \mathrm{Hull}^{M^+_\infty}(\alpha \cup \mathrm{ran}(\pi^+_{\mathcal{Q},\infty})) \cap \kappa \subseteq \alpha \},$$

where the hull in question is the "uncollapsed" set of all points definable over M^+_∞ from parameters in $\mathrm{ran}(\pi^+_{\mathcal{Q},\infty})$ and ordinals $< \alpha$. Clearly, C is ω-closed and unbounded in κ. To see that C works, fix $A \in \mu_x = \mu_0$.

For any $\mathcal{S}$ such that $\mathcal{Q} \prec^+ \mathcal{S}$, let

$$\kappa(\mathcal{S}) := \text{ unique } \nu \in \mathcal{S} \text{ such that } \pi^+_{\mathcal{S},\infty}(\nu) = \kappa.$$

Fix $\mathcal{R}$ such that $\mathcal{Q} \prec^+ \mathcal{R}$ and $A \in \mathrm{ran}(\pi^+_{\mathcal{R},\infty})$, and for $\mathcal{S}$ such that $\mathcal{R} \prec^+ \mathcal{S}$ put

$$A(\mathcal{S}) := \text{ unique } B \in \mathcal{S} \text{ such that } \pi^+_{\mathcal{S},\infty}(B) = A.$$

We shall show that

$$C \setminus (\sup(\mathrm{ran}(\pi^+_{\mathcal{R},\infty}) \cap \kappa)) \subseteq A,$$

which will then finish the proof.

We need the following general fact about iterated ultrapower constructions.

Claim 1. If $g \in \mathcal{R}$ and $g \colon [\kappa(\mathcal{R})]^{<\omega} \to \kappa(\mathcal{R})$, then there is a function $f \in \mathcal{Q}$ such that $g = \pi^+_{\mathcal{Q},\mathcal{R}}(f)(b)$ for some finite $b \subseteq \kappa(\mathcal{R})$.

Proof. Let $\mathcal{T}$ be an iteration tree on $\mathcal{Q}$ with last model $\mathcal{R}$. One can show by an easy induction that if $\mathcal{R}^*$ is on the branch of $\mathcal{T}$ leading to $\mathcal{R}$, then the claim holds with $\mathcal{R}^*$ replacing $\mathcal{R}$. $\dashv$

Because our mice do not reach superstrong cardinals, we also have

Claim 2. If $\mathcal{M}$ is a premouse, E is on the $\mathcal{M}$-sequence , $\mathrm{crit}(E) = \kappa$, and $i\colon \mathcal{M} \to \mathrm{Ult}_0(\mathcal{M}, E)$ is the canonical embedding, then $i(\kappa) = \sup\{\, i(f)(\kappa) \mid f\colon \kappa \to \kappa \wedge f \in \mathcal{M}\}$.

Proof. Let λ be the sup in question. Clearly, $\lambda \leq i(\kappa)$, so suppose $\lambda < i(\kappa)$ toward contradiction. Let $\nu = \nu(E)$.

Suppose $\nu \leq \lambda$. Let $a \subseteq \nu$ and g be such that $\lambda = i(g)(a)$. Let h be such that $a \subseteq i(h)(\kappa)$. Now define $f\colon \kappa \to \kappa$ by

$$f(\alpha) := \sup\{\, g(u) \mid u \in [h(\alpha)]^{|a|}\}.$$

Then clearly, $\lambda \leq i(f)(\kappa)$, a contradiction. Therefore $\lambda < \nu$.

Arguing as in the last paragraph, we get that $i(g)(a) < \lambda$ for all finite $a \subseteq \lambda$ and $g\colon [\kappa]^{|a|} \to \kappa$. This means that $\lambda = j(\kappa)$, where $j\colon \mathcal{M} \to \mathrm{Ult}_0(\mathcal{M}, E \restriction \lambda)$ is the canonical embedding. But the initial segment condition on premice implies that the trivial completion E^* of $E \restriction \lambda$ is on the sequence of some premouse. Since $i_{E^*}(\kappa) < \mathrm{lh}(E^*)$, we do not allow such "long extenders" in a fine extender sequence, so this is a contradiction. $\dashv$

Now fix any $\alpha \in C \setminus (\sup(\mathrm{ran}(\pi^+_{\mathcal{R},\infty}) \cap \kappa))$. Fix any $\mathcal{B}^* \in \mathcal{F}^+$ such that $\alpha \in \mathrm{ran}(\pi^+_{\mathcal{B}^*,\infty})$, and let $\mathcal{T}$ be the ω-maximal iteration tree on $\mathcal{R}$ which results from the coiteration of $\mathcal{B}^*$ with $\mathcal{R}$, using Σ_0 on both sides, and let $\mathcal{B}$ be the last model of $\mathcal{T}$. Since neither side drops, $\mathcal{B} \in \mathcal{F}^+$ and $\alpha \in \mathrm{ran}(\pi^+_{\mathcal{B},\infty})$; say

$$\alpha = \pi^+_{\mathcal{B},\infty}(\bar{\alpha}).$$

It will be enough to show that $\bar{\alpha} \in A(\mathcal{B})$.

Let us look closely at the tree $\mathcal{T}$ leading from $\mathcal{R}$ to $\mathcal{B}$. We use $\mathcal{M}_\xi, E_\xi$, and $i_{\xi,\gamma}$ for the models, extenders, and embeddings of $\mathcal{T}$. Let $\mathcal{B} = \mathcal{M}_\eta$. Now $i_{0,\eta}(\kappa(\mathcal{R})) = \kappa(\mathcal{B}) > \bar{\alpha}$, so we can set

$$\xi := \text{ least } \nu \in [0, \eta]_T \text{ such that } i_{0,\nu}(\kappa(\mathcal{R})) > \bar{\alpha}.$$

Note here that $\kappa(\mathcal{R}) \leq \bar{\alpha}$, so that $\xi > 0$; this is because if $\gamma < \kappa(\mathcal{R})$, then $\pi^+_{\mathcal{R},\infty}(\gamma) < \alpha$, so $\pi^+_{\mathcal{R},\mathcal{B}}(\gamma) < \bar{\alpha}$, so $\gamma < \bar{\alpha}$.

Let $(\nu + 1)T\xi$; we claim that $\mathrm{lh}(E_\nu) < \bar{\alpha}$. For letting $\beta = \mathrm{pred}_T(\nu + 1)$, we have $\mathrm{crit}(E_\nu) = \mathrm{crit}(i_{\beta,\xi})$ because $\mathcal{T}$ is ω-maximal, and $\mathrm{crit}(i_{\beta,\xi}) \leq \kappa(\mathcal{M}_\beta)$ by the minimality of ξ. But then $\mathrm{lh}(E_\nu) < i_{0,\nu+1}(\kappa(\mathcal{R})) \leq \bar{\alpha}$ by the minimality of ξ.

It follows that ξ is a successor ordinal. For otherwise, since $\bar{\alpha} < i_{0,\xi}(\kappa(\mathcal{R}))$, we get that $\bar{\alpha} = i_{0,\xi}(g)(a)$ for some $a \subseteq \mathrm{crit}(i_{\xi,\eta}) \cap \bar{\alpha}$ finite and $g\colon [\kappa(\mathcal{R})]^{|a|} \to$

$\kappa(\mathcal{R})$. (We get $a \subseteq \text{crit}(i_{\xi,\eta})$ because $\mathcal{T}$ is ω-maximal, and $a \subseteq \bar{\alpha}$ from the preceding paragraph and the assumption that ξ is a limit ordinal.) But by our first claim, we have $g = \pi^+_{\mathcal{Q},\mathcal{R}}(f)(b)$ for some $f \in \mathcal{Q}$ and $b \subseteq \kappa(\mathcal{R})$. We then have that

$$i_{\xi,\eta}(\bar{\alpha}) = i_{\xi,\eta}(i_{0,\xi}(g)(a)) = i_{0,\eta}(g)(a) = \pi^+_{\mathcal{Q},\mathcal{B}}(f)(\pi^+_{\mathcal{R},\mathcal{B}}(b))(a).$$

Since $\bar{\alpha} \leq i_{\xi,\eta}(\bar{\alpha})$, we can apply $\pi^+_{\mathcal{B},\infty}$ to the identity above and obtain

$$\alpha \leq \pi^+_{\mathcal{Q},\infty}(f)(\pi^+_{\mathcal{R},\infty}(b))(\pi^+_{\mathcal{B},\infty}(a)).$$

Now $\pi^+_{\mathcal{R},\infty}(b) \subseteq \alpha$ because we chose α as large as we did, and $\pi^+_{\mathcal{B},\infty}(a) \subseteq \alpha$ because $a \subseteq \bar{\alpha}$. Thus the ordinal named on the right side of the line just displayed witnesses that $\alpha \notin C$. This is a contradiction, and hence ξ is a successor ordinal.

Let $\xi = \gamma + 1$, $E = E_\gamma$, and $\beta = \text{pred}_T(\xi)$. If $\nu(E) \leq \bar{\alpha}$, then we get the same contradiction we got in the last paragraph, so we have $\nu(E) > \bar{\alpha}$. By the minimality of ξ, $\text{crit}(E) \leq \kappa(\mathcal{M}_\beta)$. We claim that $\text{crit}(E) = \bar{\alpha}$. This is true because otherwise Claim 2 gives some $h\colon \kappa(\mathcal{M}_\beta) \to \kappa(\mathcal{M}_\beta)$ such that $\bar{\alpha} < i_{\beta,\xi}(h)(c)$, where $c = \{\text{crit}(E)\} \subseteq \bar{\alpha}$. One can then proceed to a contradiction as in the last paragraph: represent h as $i_{0,\beta}(g)(d)$ where $d \subseteq \text{crit}(E)$, so that $\bar{\alpha} = i_{0,\xi}(g)(a)$, where $a := c \cup d \subseteq \text{crit}(i_{\xi,\eta}) \cap \bar{\alpha}$. Then let f, b be such that $\pi^+_{\mathcal{Q},\mathcal{R}}(f) = g$ and $b \subseteq \kappa(\mathcal{R})$, etc.

Since $\kappa(\mathcal{M}_\beta) \leq \bar{\alpha}$ by the minimality of ξ, we have $\kappa(\mathcal{M}_\beta) = \text{crit}(E) = \bar{\alpha}$. Now $\bar{\alpha}$ cannot be measurable in $\mathcal{M}_\xi = \text{Ult}(\mathcal{M}_\beta, E)$, since then $\alpha = \pi^+_{\mathcal{B},\infty}(\bar{\alpha}) = \pi^+_{\mathcal{B},\infty}(i_{\xi,\eta}(\bar{\alpha}))$ is measurable in HOD. Since $\text{cf}(\alpha) = \omega$, our dichotomy 8.25 rules this out. It follows that E is the order zero measure on $\kappa(\mathcal{M}_\gamma)$, and since using the order zero measure cannot move generators, that $\beta = \gamma$. We have then that $A(\mathcal{M}_\beta) \in E_a$, for $a = \{\kappa(\mathcal{M}_\beta)\}$, so $\bar{\alpha} = \kappa(\mathcal{M}_\beta) \in A(\mathcal{M}_\xi)$, so $\bar{\alpha} \in A(\mathcal{B})$, so $\alpha \in A$, as desired. ⊣

We remark that, once again, the negation of 8.27 is a Σ_1 statement about $L(\mathbb{R})$ by the Coding Lemma, so that if 8.27 fails, it fails below δ^2_1. Therefore, we really needed only 8.20 for its proof. It is also worth noting that 8.26 and 8.27 make no mention of mice, or even HOD, in their statements.

Bibliography

[1] Alessandro Andretta, Itay Neeman, and John R. Steel. The domestic levels of K^c are iterable. *Israel Journal of Mathematics*, 125:157–201, 2001.

[2] Jon Barwise. *Admissible sets and structures*. Springer-Verlag, Berlin, 1975. An approach to definability theory, Perspectives in Mathematical Logic.

[3] Howard S. Becker. Thin collections of sets of projective ordinals and analogs of L. *Annals of Mathematical Logic*, 19(3):205–241, 1980.

[4] Howard S. Becker and Alexander S. Kechris. Sets of ordinals constructible from trees and the third Victoria Delfino problem. In *Axiomatic set theory (Boulder, Colo., 1983)*, pages 13–29. American Mathematical Society, Providence, R.I., 1984.

[5] Anthony J. Dodd. *The Core Model*. Cambridge University Press, Cambridge, 1982.

[6] Anthony J. Dodd and Ronald B. Jensen. The core model. *Annals of Mathematical Logic*, 20(1):43–75, 1981.

[7] Anthony J. Dodd and Ronald B. Jensen. The covering lemma for K. *Annals of Mathematical Logic*, 22(1):1–30, 1982.

[8] Anthony J. Dodd and Ronald B. Jensen. The covering lemma for $L[U]$. *Annals of Mathematical Logic*, 22(2):127–135, 1982.

[9] Matthew Foreman, Menachem Magidor, and Saharon Shelah. Martin's maximum, saturated ideals, and non-regular ultrafilters. Part I. *Annals of Mathematics*, 127:1–47, 1988.

[10] Ronald Jensen. Inner models and large cardinals. *The Bulletin of Symbolic Logic*, 1(4):393–407, 1995.

[11] Ronald B. Jensen. The fine structure of the constructible hierarchy. *Annals of Mathematical Logic*, 4:229–308, 1972.

[12] Alexander S. Kechris, Donald A. Martin, and Robert M. Solovay. Introduction to Q-theory. In *Cabal seminar 79–81*, pages 199–282. Springer, Berlin, 1983.

[13] Peter Koellner and W. Hugh Woodin. Large cardinals from determinacy. In this Handbook.

[14] Kenneth Kunen. Some applications of iterated ultrapowers in set theory. *Annals of Mathematical Logic*, 1:179–227, 1970.

[15] Benedikt Löwe and John R. Steel. An introduction to core model theory. In *Sets and Proofs, Logic Colloquium 1997, Volume 1*, London Mathematical Society Lecture Notes 258. Cambridge University Press, Cambridge, 1999.

[16] Donald A. Martin and John R. Steel. The extent of scales in $L(\mathbb{R})$. In *Cabal seminar 79–81*, pages 86–96. Springer, Berlin, 1983.

[17] Donald A. Martin and John R. Steel. A proof of projective determinacy. *Journal of the American Mathematical Society*, 2(1):71–125, 1989.

[18] Donald A. Martin and John R. Steel. Iteration trees. *Journal of the American Mathematical Society*, 7(1):1–73, 1994.

[19] William J. Mitchell. Beginning inner model theory. In this Handbook.

[20] William J. Mitchell. The covering lemma. In this Handbook.

[21] William J. Mitchell. Sets constructible from sequences of ultrafilters. *The Journal of Symbolic Logic*, 39:57–66, 1974.

[22] William J. Mitchell. Hypermeasurable cardinals. In *Logic Colloquium '78 (Mons, 1978)*, pages 303–316. North-Holland, Amsterdam, 1979.

[23] William J. Mitchell. The core model for sequences of measures. I. *Mathematical Proceedings of the Cambridge Philosophical Society*, 95(2):229–260, 1984.

[24] William J. Mitchell. The core model up to a Woodin cardinal. In *Logic, methodology and philosophy of science, IX (Uppsala, 1991)*, pages 157–175. North-Holland, Amsterdam, 1994.

[25] William J. Mitchell, Ernest Schimmerling, and John R. Steel. The covering lemma up to a Woodin cardinal. *Annals of Pure and Applied Logic*, 84(2):219–255, 1997.

[26] William J. Mitchell and John R. Steel. *Fine structure and iteration trees*. Springer-Verlag, Berlin, 1994.

[27] Yiannis N. Moschovakis. *Descriptive set theory.* North-Holland Publishing Co., Amsterdam, 1980.

[28] Itay Neeman. Determinacy in $L(\mathbb{R})$. In this Handbook.

[29] Itay Neeman. Optimal proofs of determinacy. *The Bulletin of Symbolic Logic*, 1(3):327–339, 1995.

[30] Itay Neeman and John R. Steel. A weak Dodd-Jensen lemma. *The Journal of Symbolic Logic*, 64(3):1285–1294, 1999.

[31] Ernest Schimmerling. A core model toolbox and guide. In this Handbook.

[32] Ernest Schimmerling. Combinatorial principles in the core model for one Woodin cardinal. *Annals of Pure and Applied Logic*, 74:153–201, 1995.

[33] Ernest Schimmerling. A finite family weak square principle. *The Journal of Symbolic Logic*, 64:1087–1100, 1999.

[34] Ernest Schimmerling and John R. Steel. Fine structure for tame inner models. *The Journal of Symbolic Logic*, 61(2):621–639, 1996.

[35] Ernest Schimmerling and John R. Steel. The maximality of the core model. *Transactions of the American Mathematical Society*, 351:3119–3141, 1999.

[36] Ernest Schimmerling and Martin Zeman. Square in core models. *The Bulletin of Symbolic Logic*, 7:305–314, 2001.

[37] Ralf D. Schindler, John R. Steel, and Martin Zeman. Deconstructing inner model theory. *The Journal of Symbolic Logic*, 67:721–736, 2002.

[38] Ralf D. Schindler and Martin Zeman. Fine structure. In this Handbook.

[39] John R. Steel. A stationary-tower-free proof of the derived model theorem. *To appear.*

[40] John R. Steel. A theorem of Woodin on mouse sets. *To appear.*

[41] John R. Steel. Inner models with many Woodin cardinals. *Annals of Pure and Applied Logic*, 65(2):185–209, 1993.

[42] John R. Steel. $\mathrm{HOD}^{L(\mathbb{R})}$ is a core model below Θ. *The Bulletin of Symbolic Logic*, 1(1):75–84, 1995.

[43] John R. Steel. Projectively well-ordered inner models. *Annals of Pure and Applied Logic*, 74(1):77–104, 1995.

[44] John R. Steel. *The core model iterability problem.* Springer-Verlag, Berlin, 1996.

[45] John R. Steel. Mathematics needs new axioms. *The Bulletin of Symbolic Logic*, 6:422–433, 2000.

[46] Philip D. Welch. Σ^* fine structure. In this Handbook.

[47] Philip D. Welch. Combinatorial principles in the core model. 1979. Ph.D. dissertation, Oxford.

[48] Dorshka Wylie. Condensation and square in a higher core model. 1990. Ph.D. dissertation, MIT.

[49] Martin Zeman. *Inner models and large cardinals.* de Gruyter Series in Logic and Applications #5, Berlin, 2002.

Index