17. Beginning Inner Model Theory

William J. Mitchell

Contents

This research was partly supported by NSF grant DPS-9970536.
The author would like to thank John Krueger, Paul Larson, Diego Rojas, and the referee for valuable corrections and suggestions.

M. Foreman, A. Kanamori (eds.), *Handbook of Set Theory*,
DOI 10.1007/978-1-4020-5764-9_18,

This chapter provides an introduction to the basic theory of inner models without fine structure. It assumes that the reader is familiar with Gödel's class L of constructible sets; however Sect. 1 begins by recalling the definition and basic theory of L with an emphasis on the condensation property. This discussion leads to a consideration of relative constructibility—that is, models of the form $L[A]$—and then to $L[U]$, the minimal model with a measurable cardinal. A discussion of $0^{\#}$, and of sharps in general, leads up to a brief description of the Dodd-Jensen core model K^{DJ}, which links the models L and $L[U]$.

Sections 2 and 3 discuss generalizations of the ideas of Sect. 1 to larger cardinals. Section 2 looks at sequences of measurable cardinals and the models $L[\mathcal{U}]$ constructed from such sequences. The use of iterated ultrapowers to compare pairs of models, introduced in Sect. 1 for the model $L[U]$, is extended to these models $L[\mathcal{U}]$. Section 3 introduces the notion of an *extender*, a generalized form of ultrafilter used to express cardinal properties stronger than measurability. Extenders are combined with the ideas of Sect. 2 to obtain models $L[\mathcal{E}]$, constructed from a sequence $\mathcal{E}$ of extenders, which can contain cardinals up to a strong cardinal.

The definition of models for still stronger cardinals requires an understanding of iteration trees and fine structure, which are not covered in this chapter. Section 4 gives a brief survey of such larger cardinals, and the current status of their inner model theory.

The principal goal of research in inner models is to define a *core model* K which can coexist with larger cardinals in the universe V. The construction of the core model is not described in this chapter except for a brief description of K^{DJ} (which is the core model if there is no model with a measurable cardinal) in Sect. 1.2. Because of its centrality, however, the core model itself is mentioned frequently. Briefly, the *core model* K should have two properties: (1) it is like L, and (2) it is close to V. The first property is satisfied by defining it as one of the models $L[\mathcal{U}]$ or $L[\mathcal{E}]$ described in this chapter. For the second property we can ask for some form of a *covering lemma*. In the case when L is the core model—that is, when the only large cardinal properties which hold anywhere are those which hold in L—the second criteria is satisfied by Jensen's covering lemma, which states that every uncountable set x of ordinals in V is contained in a set $y \in L$ of the same cardinality. This also holds of K^{DJ} when it is the core model—that is, when there is no model with a measurable cardinal—but for larger cardinals the core model K can only be expected to satisfy some form of the *weak covering lemma*: that $(\lambda^{+})^{K} = \lambda^{+}$ for every singular cardinal λ.

In the final Sect. 5 there is a further discussion of the core model, but from a somewhat different perspective. This is not an attempt to describe the construction of an existing model, but instead is an attempt to answer the questions "how do we decide that a particular model is 'the core model' " and "how will we recognize a model, newly discovered in the future, as the core model". This attempt is, of course, highly speculative: new discoveries may show that models with the properties we are expecting are impossible or even uninteresting, or a newly discovered model with properties substantially different from what we expect may play a critical role with respect to larger cardinals, which demands that it be recognized as the core model.

Most of the topics related to inner model theory which are not covered in this chapter can be found elsewhere in this Handbook. The core model and covering lemma are introduced in the chapters [24, 33, 38]. An excellent source for further information on large cardinals is Kanamori's book [15, 16]. For more information on L, the standard reference is [3]. The more recent book [42] is an excellent introduction to inner models and core model techniques. In this Handbook, fine structure is covered in the chapters [36, 40].

One other approach to inner models which is not covered in this chapter is the class HOD of hereditarily definable sets and its variants. The model HOD has the serious disadvantage that it is not canonical—for example, it can easily be changed by forcing. However it is frequently used for models in which the axiom of choice fails, where it usually gives more readable proofs than do symmetric models, and has been used in studies of determinacy and of cardinals large enough that the inner models described in this chapter are unknown or poorly understood.

The major goal of this introduction is to establish notation and a certain amount of background for other topics in this Handbook. Where sketches of proofs are given, the intention is not so much to present the proof itself as to introduce techniques which are important to the further development of the theory of inner models and core models.

1. The Constructible Sets

1.1 Definition. Gödel's class L of constructible sets is defined to be $L = \bigcup_{\alpha\in\mathrm{On}} L_\alpha$, where the sets L_α are defined by recursion on α as follows:

1. $L_0 = \varnothing$,
2. $L_{\alpha+1} = \mathrm{def}(L_\alpha, \in)$,
3. $L_\lambda = \bigcup_{\alpha<\lambda} L_\alpha$ if λ is a limit ordinal.

Here $\mathrm{def}(L_\alpha, \in)$ is the set of subsets of L_α which are first-order definable in the structure $(L_\alpha, \in)$, using parameters from L_α.

The most basic property of L is the following:

1.2 Lemma. *There is a Π_2 sentence of set theory, which we denote by "$V = L$", such that the transitive models of the sentence "$V = L$" are exactly the sets L_α and the class L itself. Furthermore, if α is any ordinal then $\langle L_\nu : \nu < \alpha\rangle$ is definable in L_α by a Σ_1 formula.*

The main content of the sentence "$V = L$" is the statement $\forall x \exists \alpha \exists y$ $(y = L_\alpha \ \& \ x \in L_\alpha)$. See Jech [14, Lemma 13.17] for a proof of Lemma 1.2 in the case that α is a limit ordinal, which is sufficient for most uses which do not involve fine structure. The use of fine structure goes beyond Lemma 1.2 by splitting the successor interval between L_α and $L_{\alpha+1}$ into infinitely many levels of definability.

The most important property of the constructible hierarchy follows from Lemma 1.2:

1.3 Lemma (Condensation Lemma).

1. *If $X \prec_1 L_\alpha$ for some ordinal α, then there is an ordinal $\alpha' \leq \alpha$ such that $X \cong L_{\alpha'}$.*

2. *If X is a proper class such that $X \prec_1 L$, then $X \cong L$.*

That is, if X is a Σ_1 elementary substructure of L or of any L_α, then X is isomorphic, via its transitive collapse, to L or some $L_{\alpha'}$.

The simplest application of Lemma 1.3 is Gödel's proof that GCH holds in L.

1.4 Definition. If $\mathcal{M}$ is any structure and X is a subset of the universe of $\mathcal{M}$ then the *Skolem hull* of X in $\mathcal{M}$ is the smallest elementary submodel of $\mathcal{M}$ containing X. We write $\mathcal{H}^{\mathcal{M}}(X)$ for the Skolem hull of X in $\mathcal{M}$.

This definition assumes the existence of such a unique minimal submodel of $\mathcal{M}$ containing X. In all of our applications the model $\mathcal{M}$ will have a definable well-ordering which provides Skolem functions that ensure this. Definition 1.4 can also be used in cases when $\mathcal{M}$ is a well-founded class model of ZF. In this case, provided X contains a proper class of ordinals, the Skolem hull $\mathcal{H}^{\mathcal{M}}(X)$ is equal to $\bigcup_{\alpha \in X} \mathcal{H}^{V_\alpha \cap \mathcal{M}}(X \cap V_\alpha)$, and hence is definable in $\mathcal{M}$.

1.5 Theorem. $L \models \mathrm{GCH}$.

Proof. We work inside L. An easy induction on α shows that $|L_\alpha| = |\alpha|$ for all infinite ordinals α. Hence, to establish $2^\kappa = \kappa^+$ it is enough to show that any set $x \subseteq \kappa$ is a member of L_{κ^+}. To this end, pick τ large enough that $x \in L_\tau$ and set $X = \mathcal{H}^{L_\tau}(\kappa \cup \{x\})$. By Lemma 1.3 there is an ordinal α such that $\pi : (X, \in) \cong L_\alpha$ where π is the transitive collapse map. Then $x = \pi(x) \in L_\alpha$, and $|L_\alpha| = |X| = \kappa$ so $\alpha < \kappa^+$. ⊣

The aim of the inner model theory which we will outline in this chapter is to extend this result to a more general class of models. We will describe (informally, and without attempting a precise definition) a hierarchy satisfying the analog of Lemma 1.3 as a *hierarchy with condensation*.

1.1. Relative Constructibility

Each of the inner models which we will consider is defined as the class of all sets which are *constructible from* some specified set or class A. Two notions of relative constructibility are in general use:

1.6 Definition.

1. If A is a transitive set then $L(A) = \bigcup_{\alpha \in \mathrm{On}} L_\alpha(A)$, where the sets $L_\alpha(A)$ are defined exactly like the hierarchy L_α except that rule 1.1(1) is replaced by $L_0(A) = A$.

2. If A is any set or class, then $L[A] = \bigcup_{\alpha \in \mathrm{On}} L_\alpha[A]$, where the sets $L_\alpha[A]$ are defined exactly like the hierarchy L_α, except that rule 1.1(2) is replaced by
$$L_{\alpha+1}[A] = \mathrm{def}(L_\alpha[A], \in, A),$$
where $\mathrm{def}(L_\alpha[A], \in, A)$ is the set of subsets of $L_\alpha[A]$ first-order definable with parameters from $L_\alpha[A]$, using $A \cap L_\alpha[A]$ as a predicate.

The class $L(A)$ satisfies ZF and contains the set A, and it can be characterized as the smallest such class which contains the ordinals. It need not satisfy the axiom of choice, and indeed it is usually used in cases where the axiom of choice is intended to fail. The most important example is $L(\mathbb{R})$, the smallest model of ZF containing all the reals.[1] This model is heavily used in studies of the axiom of determinacy (AD), where it reconciles that axiom with the axiom of choice in the sense that the axiom "$L(\mathbb{R}) \models \mathrm{AD}$" implies many of the same consequences as the full axiom of determinacy, but is consistent with the axiom of choice in V. We will not consider models of the form $L(A)$ further in this chapter.

If A is a set then the model $L[A]$ always satisfies ZFC. It need not have A itself as a member, but the restriction $A \cap L[A]$ of A to the model $L[A]$ is in $L[A]$. The model $L[A]$ can be characterized as the smallest model M of ZF which contains all the ordinals and has $A \cap M$ as a member. The case when A is a class is similar, provided that replacement holds for formulas with a predicate for A.

In one sense the models $L[A]$ can be fully as complex as any other model of set theory. This is clear in the case that A is a class, since (assuming the axiom of global choice) the universe V can be coded by a class A of ordinals, so that $L[A] = V$. However, a surprising result of Jensen ([2], see [10, Theorem 5.1]) shows that A need not be a proper class: he defines a class generic extension $V[G]$ of the universe V such that $V[G] = L[a]$ for some $a \subseteq \omega$. Thus any class can be contained in a model of the form $L[a]$, with $a \subseteq \omega$.

[1] Strictly speaking Definition 1.6 does not apply to $L(\mathbb{R})$, since $\mathbb{R}$ is not transitive. Taking $L(\mathbb{R})$ to be $L(V_{\omega+1})$ repairs this defect and also gives a more convenient form to the low levels of the $L_\alpha(\mathbb{R})$ hierarchy.

In another sense the models $L[A]$ are quite simple when A is a set—nearly as simple as L itself. This simplicity appears when working above the set A, for submodels M of $L[A]$ such that $A \cap L[A] \in M$. For example, the sentence "$V = L$" can be generalized straightforwardly to a sentence "$V = L[A]$" which satisfies the following generalization of the Condensation Lemma 1.3:

1.7 Lemma. *Suppose that A is a set, and that $X \prec_1 L_\alpha[A]$, where $\alpha \in \mathrm{On} \cup \{\mathrm{On}\}$, and the transitive closure of $A \cap L_\alpha[A]$ is contained in X. Then there is an $\alpha' \leq \alpha$ such that $X \cong L_{\alpha'}[A]$.*

Hence the sets $L_\alpha[A]$, with $\alpha \geq \mathrm{rank}(A)$, also form a hierarchy with condensation, and it follows that all of the basic properties of L, such as GCH, $\diamondsuit_\kappa$ and $\Box_\kappa$, hold in $L[A]$—at least above $\mathrm{rank}(A)$—for the same reason that they hold in L

Lemma 1.7 does not give any information about the set A itself, and it says nothing about how models $L[A]$ and $L[A']$ might be related when $A \neq A'$. If we are to use the techniques of inner model theory to study the set A, then we need a version of Lemma 1.7 which does not assume $A \cap L[A] \in X$. Any such lemma will require some restriction on the class of sets A for which it is valid.

Elementary embeddings (or, rather, sets A encoding elementary embeddings) have proved to be especially fruitful for this purpose. One reason for this fruitfulness is that when A and A' encode different elementary embeddings of $L[A]$ and $L[A']$, respectively, then it is possible, under suitable conditions, to use the embeddings themselves to modify the models so that they can be compared. This gives at least a start on the goal of understanding the relationships between distinct models $L[A]$ and $L[A']$. This idea may be seen in the proof of Theorem 1.9 and in the comparison Lemma 2.8 for sequences of measures.

A second reason for this fruitfulness arises in the consideration of the embeddings $\pi : L_{\bar\alpha}[\bar A] \cong X \prec_1 L_\alpha[A]$ arising from a transitive collapse. If embeddings coded by A are suitably chosen, then the embedding π will be closely related to the embeddings encoded into A. In this case an analog of the Condensation Lemma 1.3 may hold without the restriction $A \cap L[A] \subseteq X$ needed for Lemma 1.7. This phenomenon often occurs, and is heavily used, in the analysis of inner models for large cardinals.

1.2. Measurable Cardinals

The simplest, and oldest, example of a model $L[A]$ in which A encodes an embedding of $L[A]$ is $L[U]$, the minimal inner model for a measurable cardinal.

1.8 Definition. Recall that a cardinal κ is *measurable* if there is an elementary embedding $i : V \to M$, where M is a well-founded class and $\kappa = \mathrm{crit}(i)$. Here $\mathrm{crit}(i)$ is the *critical point* of i, that is, the least ordinal α such that $i(\alpha) > \alpha$.

The *ultrafilter* associated with such an embedding i is the set

$$U = \{x \subseteq \kappa : \kappa \in i(x)\}.$$

This set U is a *κ-complete ultrafilter* on κ, where κ-completeness means that $\bigcap X \in U$ whenever $X \subseteq U$ and $|X| < \kappa$. Indeed, U is *normal*, which is a stronger property: for any function $f : \kappa \to \kappa$, if $\{\nu < \kappa : f(\nu) < \nu\} \in U$ then there is an ordinal $\gamma < \kappa$ such that $\{\nu : f(\nu) = \gamma\} \in U$.

A normal ultrafilter is frequently called a *measure*. The analogy with Lebesgue measure on the real line, from which this terminology is derived, is slightly strained since neither normality nor the property of being two-valued has an analog in the real line; however this usage has a strong historical basis (evidenced by the term "measurable cardinal") and it is useful in a context such as the present chapter, in which non-normal ultrafilters never appear.

In the other direction, from an ultrafilter to an embedding, the *ultrapower* construction gives, for any normal ultrafilter U on κ, an embedding $i^U : V \to M = \mathrm{Ult}(V, U)$ with critical point κ such that U is the ultrafilter associated with i^U. The ultrapower has the property that $M = \{i^U(f)(\kappa) : f \in V \cap {}^{\kappa}V\}$, and as a consequence i^U is minimal among all embeddings related to U in the following sense: Any other embedding $i : V \to N$ with the same associated ultrafilter U can be factored as

$$i : V \xrightarrow{i^U} \mathrm{Ult}(V, U) \xrightarrow{k} N,$$

where the embedding k is defined by $k([f]_U) = k(i^U(f)(\kappa)) = i(f)(\kappa)$.

It is easy to see that if U is a normal ultrafilter on κ, then $U \cap L[U]$ is a normal ultrafilter in $L[U]$. On its face, the model $L[U]$ appears to depend on the choice of the ultrafilter U; however Kunen [18] showed that it depends only on the critical point of U.

The proof, which we outline below, uses iterated ultrapowers. We write $i^U_\alpha : V \to \mathrm{Ult}_\alpha(V, U)$ for the α-fold iterated ultrapower by U, which is defined by setting $\mathrm{Ult}_0(V, U) = V$, $\mathrm{Ult}_{\alpha+1}(V, U) = \mathrm{Ult}(\mathrm{Ult}_\alpha(V, U), i^U_\alpha(U))$, and $\mathrm{Ult}_\alpha(V, U) = \operatorname{dir\,lim}_{\alpha' < \alpha} \mathrm{Ult}_{\alpha'}(V, U)$ if α is a limit ordinal.

We will need the fact that every iterated ultrapower $\mathrm{Ult}_\alpha(L[U], U)$ is well-founded. This is easily proved by induction on α: more generally, let M be any well-founded model containing the ordinals, and suppose that M satisfies that U is a countably complete ultrafilter. A useful observation is that all iterated ultrapowers of M are definable subsets of M, and hence we can work inside M. It is easy to see that $\mathrm{Ult}(M, U)$ is well-founded. For any ordinal α such that $\mathrm{Ult}_\alpha(M, U)$ is well-founded, it then follows, by working inside $\mathrm{Ult}_\alpha(M, U)$, that $\mathrm{Ult}_{\alpha+1}(M, U)$ is also well-founded. Hence the least ordinal α such that $\mathrm{Ult}_\alpha(M, U)$ is ill-founded would be a limit ordinal. Now call an ordinal γ *U-soft in M* if there is an iterated ultrapower i^U_α by U such that the set of ordinals in $\mathrm{Ult}_\alpha(M, U)$ below $i^U_\alpha(\gamma)$ is ill-founded. Let γ be the least U-soft ordinal in M, and let α be least such that i^U_α witnesses that

γ is soft. Now if α' is any ordinal in the interval $0 \leq \alpha' < \alpha$, then $i^U_{\alpha'}(\gamma)$ is, by elementarity, the least $i^U_{\alpha'}(U)$-soft ordinal in $\mathrm{Ult}_{\alpha'}(M, U)$. But this is impossible, since for sufficiently large $\alpha' < \alpha$ there is an ordinal $\xi < i^U_{\alpha'}(\gamma)$ in $\mathrm{Ult}_{\alpha'}(M, U)$ such that $i_{\alpha',\alpha}(\xi)$ is a member of an infinite descending sequence below $i_\alpha(\gamma)$. Then $i^U_{\alpha',\alpha}$ is an iterated ultrapower of $\mathrm{Ult}_{\alpha'}(M, U)$ by $i^U_{\alpha'}(U)$ which witnesses that ξ is $i^U_{\alpha'}(U)$-soft in $\mathrm{Ult}_{\alpha'}(M, U)$.

The proof above can be generalized to any well-founded model M with $\omega_1 \subseteq M$, and to any iterated ultrapower of M by arbitrary measures in M rather than by the single ultrafilter U. We will later see that the situation is much more difficult for iterations involving cardinals beyond a strong cardinal.

1.9 Theorem. *Suppose that U and U' are normal ultrafilters in $L[U]$ and $L[U']$, respectively.*

1. *If* $\mathrm{crit}(U) = \mathrm{crit}(U')$ *then* $U = U'$, *and hence* $L[U] = L[U']$.
2. *If* $\mathrm{crit}(U) < \mathrm{crit}(U')$ *then* $L[U'] = \mathrm{Ult}_\alpha(L[U], U)$ *for some ordinal* α.

1.10 Corollary. *The model $L[U]$ has only the one normal ultrafilter U.*

Sketch of Proof of Theorem 1.9(1). The proof of Theorem 1.9 uses the following two observations about the iterated ultrapower $\mathrm{Ult}_\lambda(L[U], U)$, where $\lambda > \kappa^+$ is a cardinal of uncountable cofinality.

(1) The set $C = \{i^U_\alpha(\kappa) : \alpha < \lambda\}$ is a closed, unbounded set of indiscernibles for $\mathrm{Ult}_\lambda(L[U], U)$ which generates its measure $i^U_\lambda(U)$ in the sense that

$$i^U_\lambda(U) = \{x \subseteq \lambda : \sup(C - x) < i^U_\lambda(\kappa)\}, \tag{17.1}$$

and therefore $\mathrm{Ult}_\lambda(L[U], U) = L[i^U_\lambda(U)] = L[\mathcal{C}_\lambda]$ where $\mathcal{C}_\lambda$ is the filter of closed unbounded subsets of λ. To see that (17.1) holds, let x be any subset of $i_\lambda(\kappa)$ in $\mathrm{Ult}_\lambda(L[U], U)$. Then there is some $\alpha_0 < \lambda$ and $x_{\alpha_0} \in \mathrm{Ult}_{\alpha_0}(L[U], U)$ such that $x = i_{\alpha_0,\lambda}(x_{\alpha_0})$. For ordinals α in the interval $\alpha_0 < \alpha < \lambda$ set $x_\alpha = i_{\alpha_0,\alpha}(x_{\alpha_0})$, so that $x = i_{\alpha,\lambda}(x_\alpha)$. Then $i^U_\alpha(\kappa) \in x \iff i^U_\alpha(\kappa) \in x_{\alpha+1} = i_{\alpha,\alpha+1}(x_\alpha) \iff x_\alpha \in i^U_\alpha(U) \iff x \in i^U_\lambda(U)$.

(2) Let Γ be the class of ordinals $\xi > \lambda$ such that $i^U_\lambda(\xi) = \xi$. Then simple cardinal arithmetic shows that Γ is a proper class, and contains all of its limit points of cofinality greater than λ.

Now suppose that the models $L[U]$ and $L[U']$ are as in the hypothesis of Theorem 1.9(1), with $\kappa = \mathrm{crit}(U) = \mathrm{crit}(U')$. Let $\lambda = (2^\kappa)^+$. By the first observation, $\mathrm{Ult}_\lambda(L[U], U) = \mathrm{Ult}_\lambda(L[U'], U') = L[\mathcal{C}_\lambda]$, with $i^U_\lambda(U) = i^{U'}_\lambda(U') = \mathcal{C}_\lambda \cap L[\mathcal{C}_\lambda]$. By the second observation $\Gamma = \{\xi > \lambda : i^U_\lambda(\xi) = i^{U'}_\lambda(\xi) = \xi\}$ is a proper class.

Let $X = \mathcal{H}^{L[\mathcal{C}_\lambda]}(\kappa \cup \Gamma \cup \{\mathcal{C}_\lambda \cap L[\mathcal{C}_\lambda]\}) \prec L[\mathcal{C}_\lambda]$ be the Skolem hull, and let $\pi : M \cong X$ be its transitive collapse. Then $X \subseteq \mathrm{ran}(i^U_\lambda)$, so $M \prec L[U]$. However $U = \pi^{-1}(\mathcal{C}_\lambda \cap L[\mathcal{C}_\lambda]) \in M$, and the proper class Γ is a subset of M. It follows by Lemma 1.7 that $M = L[U]$. By the same argument $M = L[U']$ and $U' = \pi^{-1}(\mathcal{C}_\lambda \cap L[\mathcal{C}_\lambda])$, so $L[U] = L[U']$ and $U = U'$. $\dashv$

1.11 Theorem (Silver). $L[U] \models \mathrm{GCH}$.

Sketch of Proof. First we recall Gödel's proof of GCH for L. Assume $V = L$, and fix any set $x \in \mathcal{P}^L(\lambda)$. Now pick some $\tau > \lambda$ such that $x \in L_\tau$, and let $\pi : M_x \cong X \prec_1 L_\tau$ where $|X| = \lambda$, $\lambda \cup \{x\} \subseteq X$, and M_x is transitive. Then the Condensation Lemma 1.3 implies that $M_x = L_\alpha$ for some $\alpha < \lambda^+$, so that $x = \pi^{-1}(x) \in L_\alpha$. Thus $\{z \subseteq \lambda : z \leq^L x\} \subseteq L_\alpha$, and since $|L_\alpha| = \lambda$ it follows that no set in $\mathcal{P}(\lambda)$ has more than λ many $<^L$-predecessors. Hence $\mathrm{otp}(\mathcal{P}(\lambda), \leq^L) = \lambda^+$, so $L \models 2^\lambda = \lambda^+$.

Now assume $V = L[U]$, where U is a normal ultrafilter on κ, and fix a cardinal λ. If $\lambda \geq \kappa$ then Gödel's proof for L can be easily adapted to $L[U]$ by substituting Lemma 1.7 for Lemma 1.3. Thus we only need to consider the case $\lambda < \kappa$.

Fix a set $x \subseteq \lambda$, and pick τ such that $x \in L_\tau[U]$. For convenience, also let $L_\tau[U]$ satisfy ZF^-, the axioms of ZF without the Power Set Axiom; this will be true if τ is any successor cardinal. Now let $X \prec L_\tau[U]$ where $x \in X$, $\lambda \subseteq X$, and $|X| = \lambda$. If $M_x \cong X$ is the transitive collapse of X, then $M_x = L_{\alpha_x}[U_x]$ for some $\alpha_x < \lambda^+$ and some filter U_x which is a normal ultrafilter in M_x.

In order to conclude, as in the proof for L, that $\mathrm{otp}\big(\mathcal{P}(\lambda), <^{L[U]}\big) = \lambda^+$, we need to show that $\{z \subseteq \lambda : z \leq^{L[U]} x\} \subseteq M_x$. The fact that $U_x \neq U$ is a complication which is not present in L, and we will use the techniques from the proof of Theorem 1.9 to deal with it. The assumption that $L_\tau[U]$, and hence M_x, satisfies ZF^- makes it is easy to verify that the iterated ultrapower $i^{U_x}_\kappa : L_{\alpha_x}[U_x] \to \mathrm{Ult}_\kappa(L_{\alpha_x}[U_x], U_x)$ can be defined and has all of the required properties: In particular, $\mathrm{Ult}_\kappa(L_{\alpha_x}[U_x], U_x) = L_{\alpha'_x}[i^{U_x}_\kappa(U_x)]$ for some $\alpha'_x < \kappa^+$, and $i^{U_x}_\kappa(U_x) \subseteq U$ since $i^{U_x}_\kappa(U_x)$ is generated by the set $C_x = \{i^{U_x}_\nu(\lambda) : \nu < \kappa\}$, which is in U since it is closed and unbounded. Since $i^{U_x}_\lambda \restriction \mathcal{P}^{M_x}(\lambda)$ is the identity, it follows that $\{z \subseteq \lambda : z <^{L[U]} x\} \subseteq \mathcal{P}^{L_{\alpha'_x}[U_x]}(\lambda) \subseteq M_x$, as desired. ⊣

This proof can be interpreted as showing that $L[U]$ contains a hierarchy with condensation; however this hierarchy has two flaws: (i) the very existence of the model $L[U]$, and hence of this hierarchy, is conditional on the existence of the normal ultrafilter U, and (ii) unlike the structures L_α, the structures M_x do not actually satisfy condensation. That is, the model M_x is not actually an initial segment of $L[U]$, but only a structure which can be compared to an initial segment of $L[U]$ by means of an iterated ultrapower. The first, and more important, of these two flaws was fixed by the Dodd and Jensen [4, 6, 5] with their introduction of the original core model K^{DJ}. They defined a *mouse* to be a structure $M = L_{\alpha^M}[U_M]$ such that (i) M satisfies the sentence "U_M is a normal ultrafilter", (ii) all of the iterated ultrapowers of M are well-founded, and (iii) M satisfies a fine structure condition which implies that there is a $\rho \leq \mathrm{crit}(U_M)$ such that $L_{\alpha^M+1}[U_M] \models \left|\alpha^M\right| = \rho$. The Dodd-Jensen core model K^{DJ} is defined to be $L[\mathcal{M}]$, where $\mathcal{M}$ is the class of all mice. With the emergence of a general concept of "the core model" (see Sect. 5), K^{DJ} came to be seen as the core model below $L[U]$, that is, it is the core model provided that there is no model with a measurable cardinal.

The weakest mouse is equi-constructible with $0^{\#}$, which is described in the next subsection. The model M_x in the proof of Theorem 1.11 is an example of a mouse; however its construction required starting with the model $L[U]$ and it is difficult to prove that such a model exists using any assumption weaker than a measurable cardinal. Dodd and Jensen threw out the assumption $M_x \models \mathrm{ZF}^-$ of Theorem 1.11, and replaced it with clauses (i) and (ii); they then used fine structure to show that iterated ultraproducts of mice can still be defined and have the required properties.

The second flaw, the lack of condensation, is only a minor technical problem at the level of one measurable cardinal but leads to serious difficulties at higher levels. This problem is corrected by the modern presentation of the core model. As adapted to the special case of the Dodd-Jensen core model K^{DJ}, this presentation works as follows: First note that K^{DJ} does satisfy a form of condensation, for if $\pi : L_{\alpha'}[\mathcal{M}'] \cong X \prec_1 L_\alpha[\mathcal{M}]$, then π preserves the property of being a mouse. It follows that $\mathcal{M}'$ is contained in $\mathcal{M}$, and since the Dodd-Jensen mice are well-ordered by relative constructibility it follows that $\mathcal{M}'$ is an initial sequence of the class $\mathcal{M}$. We can extend this to $L[U]$ as follows: each mouse is a model $L_{\alpha_M}[U_M]$. Since U_M and $L_{\alpha_M}[U_M]$ are equi-constructible, K^{DJ} can be equivalently written as $L[\langle U_M : M \in \mathcal{M}\rangle]$ instead of as $L[\mathcal{M}]$, and then $L[U]$ is equal to $L[\mathcal{M}, U] = L[\langle U_M : M \in \mathcal{M}\rangle, U]$. If we let $\mathcal{U}$ be the sequence $\langle U_M : M \in \mathcal{M}\rangle^\frown\langle U\rangle$, then $L[U] = L[\mathcal{U}]$, and the transitive collapse of a substructure $X \prec_1 L_\alpha[\mathcal{U}]$ has the form $L_{\alpha'}[\mathcal{U}']$ where, as in the case of K^{DJ}, the sequence $\mathcal{U}'$ is an initial segment of $\mathcal{U}$. Thus $X \cong L_{\alpha'}[\mathcal{U}{\restriction}\alpha']$, which is an initial segment of $L[\mathcal{U}] = L[U]$.

Notice that this construction has the further advantage of smoothly joining the construction of K^{DJ} with $L[U]$ at one extreme and (taking $\mathcal{U}$ to be empty) L at the other.

1.3. $0^{\#}$, and Sharps in General

This subsection covers the first steps of the core model hierarchy suggested by the proof of Theorem 1.11. They are the first steps historically, since the model $L[0^{\#}]$ was the first canonical inner model to be extensively studied other than L and $L[U]$. They are also the first steps in the sense that they lie at the bottom of the core model hierarchy: $0^{\#}$ is, as we will see later, essentially the same as the first Dodd-Jensen mouse.

Lemma 1.2 implies that if $i : L \to M$, where M is a well-founded class, then $M = L$. As Scott [37] observed, it follows that there are no measurable cardinals in L: otherwise let $U \in L$ be a normal ultrafilter on the least measurable cardinal κ of L. Then $i^U(\kappa) > \kappa$; but this is impossible since $i^U(\kappa)$ is, by the elementarity of the embedding i^U, the least measurable cardinal in $\mathrm{Ult}(L, U) = L$, and that is κ.

Nontrivial embeddings from L into L can, however, exist in V: for example, if U is a normal ultrafilter and $i^U : V \to \mathrm{Ult}(V, U)$ then $i^U \restriction L : L \to L$. Silver's $0^\#$ gives a complete analysis of such embeddings. We say that a class I of ordinals is a class of indiscernibles for a model M if for any formula $\varphi(v_0, \dots, v_{n-1})$ of the language of set theory and any increasing sequences $(c_0, \dots, c_{n-1})$ and $(c'_0, \dots, c'_{n-1})$ of members of I we have $M \models \varphi(c_0, \dots, c_{n-1}) \iff \varphi(c'_0, \dots, c'_{n-1})$.

1.12 Definition. We say that $0^\#$ *exists* if there is closed proper class I of indiscernibles for L. In this case we define $0^\# \subseteq \omega$ to be the set of Gödel numbers of formulas $\varphi(v_0, \dots, v_{n-1})$ such that $L \models \varphi(c_0, \dots, c_{n-1})$ for any increasing sequence $\langle c_0, \dots, c_{n-1} \rangle \in [I]^n$.

Since I is a class of indiscernibles for L, this characterization of the set $0^\#$ does not depend on the choice of the sequence $\vec{c} \in [I]^n$. The fact that I is required to be closed implies that the definition of $0^\#$ does not depend on the choice of the class I. It also implies that the members of I possess the following normality property, which Silver called *remarkability*: if η is any ordinal and $f : \mathrm{On} \to \mathrm{On}$ is any map definable in L from parameters in L_η such that $f(c_0, \dots, c_{n-1}) = \xi < c_0$ for some sequence $\vec{c} = (c_0, \dots, c_{n-1}) \in [I - \eta]^n$, then $f(\vec{d}) = \xi$ for every sequence $\vec{d} \in [I - \eta]^n$.

Silver showed that if $0^\#$ exists then there is a unique maximal class I, the *Silver indiscernibles* such that $L = \mathcal{H}^L(I)$, that is, every set in L is definable in L from parameters in I. This class can be obtained by starting with any remarkable class I' of indiscernibles. Then $\mathcal{H}^L(I') \prec L$ is a proper class and hence is isomorphic to L. If $\pi : \mathcal{H}^L(I') \cong L$ is the transitive collapse map, then $I = \pi``I'$ is a closed class of indiscernibles and $\mathcal{H}^L(I) = L$.

Our Definition 1.12 requires that I be a proper class, but Silver showed that this is not necessary:

1.13 Theorem. *If there is an uncountable set of indiscernibles for L then $0^\#$ exists. Furthermore, there is a Π^1_2 formula ψ such that if a is any subset of ω, then $\psi(a)$ holds if and only if $a = 0^\#$.*

Thus for example, the existence of $0^\#$ is an immediate consequence of the existence of a Ramsey cardinal. The following result shows how $0^\#$ can be used to characterize the elementary embeddings from L into L:

1.14 Theorem (Silver). *Assume that $0^\#$ exists. Then (i) for any strictly increasing map $\pi : I \to I$ there is a unique elementary embedding $i : L \to L$ such that $\pi = i \restriction I$, and (ii) if $i : L \to L$ then $i``I \subseteq I$, and i is determined by $i \restriction I : I \to I$.*

The proof follows easily from the indiscernibility of the members of I and the fact that every constructible set is definable from members of I: if x is the unique set satisfying a formula $\varphi(x, \alpha_0, \dots, \alpha_{n-1})$ for some sequence $(\alpha_0, \dots, \alpha_{n-1}) \in [I]^{<\omega}$, then $i(x)$ must be the unique set x' satisfying $\varphi(x', \pi(\alpha_0), \dots, \pi(\alpha_{n-1}))$.

This leaves open one gap in the use of $0^\#$ to characterize embeddings from L into L: the question of whether the existence of such an embedding implies the existence of $0^\#$. This question was settled by Kunen; the version of the proof which we sketch below is largely due to Silver and is included because it involves ideas which are basic to the proof of the covering lemma:

1.15 Theorem. *If $i : L \to L$ is a nontrivial elementary embedding then $0^\#$ exists.*

Sketch of Proof. Let $\kappa = \mathrm{crit}(i)$. We can assume without loss of generality that i is continuous at every ordinal of cofinality greater than κ: if it is not, then factor the embedding i as $i : L \to X := \{i(f)(\kappa) : f \in L\} \prec L$ and replace i with $i' : L \xrightarrow{i} X \xrightarrow{\pi} L$, where $\pi : X \cong L$ is the transitive collapse.

We will define, for each $\nu \in \mathrm{On}$, a class Γ_ν of ordinals which is unbounded and contains all of its limit points of cofinality greater than κ. If we set $\kappa_\nu = \inf(\Gamma_\nu - \kappa)$ then the class $J = \{\kappa_\nu : \nu \in \mathrm{On}\}$ will be a class of indiscernibles for L, and by Silver's results this implies that $0^\#$ exists.

Set $\Gamma_0 = \mathrm{On} \cap \mathrm{ran}(i)$, and if λ is a limit ordinal then set $\Gamma_\lambda = \bigcap_{\nu<\lambda} \Gamma_\nu$. Now suppose that Γ_ν has been defined, and write $\mathcal{H}^L(X) \prec L$ for the class of sets definable in L from parameters in X. Then $\mathcal{H}^L(\Gamma_\nu) \cong L$ since Γ_ν is a proper class, so consider the map

$$i_\nu : L \cong \mathcal{H}^L(\Gamma_\nu) \prec L.$$

Then $\Gamma_{\nu+1}$ is defined to be the set of ordinals ξ such that $i_\nu(\xi) = \xi$.

Notice that $\Gamma_\nu = \mathrm{On} \cap \mathcal{H}^L(\Gamma_\nu)$, that $\kappa_\nu = \inf(\Gamma_\nu - \kappa) = i_\nu(\kappa)$, and that if $\nu > \nu'$ then $i_{\nu'}(\kappa_\nu) = \kappa_\nu$. Now define, for each pair $\nu' < \nu$ of ordinals, the embedding $i_{\nu',\nu} : L \cong \mathcal{H}^L(\kappa_{\nu'} \cup \Gamma_\nu) \prec L$ to be the inverse of the transitive collapse of $\mathcal{H}^L(\kappa_{\nu'} \cup \Gamma_\nu)$. Thus $i_{\nu',\nu}$ is the identity on $\kappa_{\nu'} \cup \Gamma_{\nu+1}$.

We claim that $i_{\nu',\nu}(\kappa_{\nu'}) = \kappa_\nu$. This claim is equivalent to the statement that $\kappa_\nu \cap \mathcal{H}^L(\kappa_{\nu'} \cup \Gamma_\nu) = \kappa_{\nu'}$, and if it were false then there would be $\vec{\alpha} \in [\Gamma_\nu]^{<\omega}$ and a formula φ such that

$$L \models \exists \eta \in \kappa_\nu - \kappa_{\nu'} \exists \vec{\gamma} \in [\kappa_{\nu'}]^{<\omega} \big(\varphi(\vec{\gamma}, \eta, \vec{\alpha}) \ \& \ \forall \eta' < \eta \, \neg\varphi(\vec{\gamma}, \eta', \vec{\alpha})\big). \quad (17.2)$$

Now the embedding $i_{\nu'} : L \cong \mathcal{H}^L(\Gamma_{\nu'}) \prec L$ is elementary, and $i_{\nu'}(\kappa) = \kappa_{\nu'}$, but $i_{\nu'}(\kappa_\nu) = \kappa_\nu$ and $i_\nu(\vec{\alpha}) = \vec{\alpha}$ since $i_{\nu'} \restriction \Gamma_\nu$ is the identity. Thus formula (17.2) implies that

$$L \models \exists \eta \in \kappa_\nu - \kappa \exists \vec{\gamma} \in [\kappa]^{<\omega} \big(\varphi(\vec{\gamma}, \eta, \vec{\alpha}) \ \& \ \forall \eta' < \eta \, \neg\varphi(\vec{\gamma}, \eta', \vec{\alpha})\big).$$

But this is impossible, since any such ordinal η would be in Γ_ν and $\kappa_\nu = \min(\Gamma_\nu - \kappa)$.

This completes the proof of the claim. Now suppose that $\vec{c}$ and $\vec{c}'$ are two increasing sequences in $[J]^{<\omega}$ which differ only in the ith place; say that, $c'_i = \kappa_{\nu'} < c_i = \kappa_\nu$ while $c_j = c'_j$ for $j \neq i$. Then $i_{\nu',\nu}(\vec{c}') = \vec{c}$, and since $i_{\nu',\nu} : L \to L$ is elementary it follows that $\vec{c}'$ and $\vec{c}'$ satisfy the same formulas over L. But if $\vec{c}'$ and $\vec{c}$ are any two increasing sequences of the same length from $[J]^{<\omega}$, then one can be obtained from the other in a finite sequence of steps in such a way that each step changes only one element of the sequence. Hence J is a class of indiscernibles for L. ⊣

The following result of Silver states that if $0^\#$ exists then the class L of constructible sets is much smaller than V:

1.16 Theorem (Silver). *Assume that $0^\#$ exists, and let I be the class of Silver indiscernibles. Then (i) every uncountable cardinal κ of V is a member of I, and indeed $|I \cap \kappa| = \kappa$, (ii) every Silver indiscernible is weakly compact in L, (iii) $\forall\eta\ \left|\mathcal{P}^L(\eta)\right| = |\eta|$, and (iv) $\forall\eta\ \mathrm{cf}(\eta^{+L}) = \omega$.*

Clause (ii) can be strengthened by replacing "weakly compact" with any large cardinal property which can consistently hold in L. This fact suggests that the existence of $0^\#$ can be viewed as the weakest large cardinal property which cannot consistently be true in L, and further experience has supported this view. Such a statement cannot be proved, or even stated precisely, without a precise definition of "large cardinal property"; however it is true for large cardinals inside the core model, and the covering lemma provides other senses in which $L[0^\#]$ is a minimal extension of L. For example, if M is any class model such that $M \models \lambda^+ \neq (\lambda^+)^L$ for some singular cardinal λ of M, then $L[0^\#]$ is contained in M.

Solovay once suggested that $L[0^\#]$ might be minimal in another sense: that every real $a \in L[0^\#]$ such that $0^\# \notin L[a]$ would be set generic over L. This suggestion was refuted by Jensen [2], who used class forcing to construct a counterexample. A weaker conjecture might be that $0^\#$ is the minimal real which is easily definable; however Friedman [11] has shown that if $0^\#$ exists then there is a set a such that $0 <_L a <_L 0^\#$ and a is a Π^1_2 singleton; furthermore, the set defined by this Π^1_2 formula remains a singleton in any extension with the same ordinals. See [10, Theorem 6.5] for more on this subject.

1.4. Other Sharps

The process used to define $0^\#$ can also be applied to models larger than L. This process is commonly used in two slightly different contexts: in order to define the sharp of a large cardinal property, and in order to define the sharp of a set.

In order to construct the sharp of a large cardinal property, we need to start with a minimal inner model M for the property such that M has a suitable inner model theory. For a measurable cardinal, for example, we could take any model of the form $M = L[U]$ such that U is a normal ultrafilter in M. If J is a closed proper class of indiscernibles for M, then we can define a new real, just as with $0^\#$, to be the set of Gödel numbers of formulas $\varphi(x_0, \ldots, x_{n-1})$ such that $M \models \varphi(c_0, \ldots, c_{n-1})$ for any $(c_0, \ldots, c_{n-1}) \in [J]^n$.

By using the inner model theory for the model M in question, together with Silver's techniques from the theory of $0^{\#}$, it can be shown that this construction yields a unique real even though (as in the case of $L[U]$) the model M may not itself be unique.

This procedure is not limited to properties involving a single cardinal. As we will see shortly, the ideas of $L[U]$ can be extended to a model $M = L[\mathcal{U}]$, having a proper class of measurable cardinals, so that M has an inner model theory similar to that of $L[U]$. The procedure described above, applied to the model M, will then yield a real which is the sharp for a proper class of measurable cardinals.

The sharp construction was first applied to $L[U]$ by Solovay, who gave the name $0^{\dagger}$ to the resulting sharp for a measurable cardinal. This precedent has had the effect of leading to a proliferation of typographical symbols for sharps of various large cardinal properties, the most common of which is $0^{\P}$, used for the sharp of a strong cardinal. The use of these symbols, apparently chosen on a whim and with no relation to the cardinals they are supposed to represent, places an unfortunate and unnecessary burden on the reader's memory. Fortunately the most important example, the sharp for a Woodin cardinal, has escaped the use of such symbols. This sharp is important because a number of applications of Woodin cardinals, particularly to inner model theory, appear to require the sharp for a Woodin cardinal, rather than simply a Woodin cardinal itself. It is commonly denoted by $M_1^{\#}$, where M_1 is the standard symbol for the minimal model with one Woodin cardinal.

It is straightforward to generalize the construction of $0^{\#}$ to obtain the sharp $A^{\#}$ for an arbitrary set A of ordinals: If $A \subseteq \gamma$, and J is a closed, proper class of indiscernibles for $L[A]$, then $A^{\#}$ is the set of pairs $(n, \vec{a})$ where $n = \ulcorner\varphi\urcorner$ is the Gödel number of a formula $\varphi(\vec{v}, z, \vec{u})$, $\vec{a} \in [\gamma]^{<\omega}$, and $L[A] \models \varphi(\vec{c}, A, \vec{a})$ for any $\vec{c} \in [J]^n$. By use of an appropriate coding, we can regard $A^{\#}$ as a subset of γ.

In particular, this construction can be used to iterate the sharp operation: Starting with $0^{\#^1} = 0^{\#}$, we define $0^{\#^{\alpha+1}} = (0^{\#^{\alpha}})^{\#}$. If α is a limit ordinal then $0^{\#^{\alpha}}$ is defined to be a set encoding $\langle 0^{\#^{\gamma}} : \gamma < \alpha\rangle$.

Assuming the existence of a large cardinal (a Ramsey cardinal is much more than enough) it can be shown to be consistent that $0^{\#^{\alpha}}$ exists for all ordinals α. The model $L[\langle 0^{\#^{\alpha}} : \alpha \in \mathrm{On}\rangle]$ forms a hierarchy with condensation, and this hierarchy is an initial segment of the core model hierarchy toward which we are working. This process can easily be continued: the model $M = L[\langle 0^{\#^{\alpha}} : \alpha \in \mathrm{On}\rangle]$ is a minimal model for the large cardinal property "$A^{\#}$ exists for all sets A", and (given a class of indiscernibles for this model) we can define the sharp for this property. This sharp will be a subset of ω and it is the next step $0^{\#^{\mathrm{On}}}$ in the desired hierarchy.

On the one hand there seems to be no obvious bound determining how far the hierarchy obtained through this process can be extended, but on the other hand it is not clear how to generalize the process to give a uniform definition, using indiscernibles, of such a hierarchy. The core model provides such a definition by replacing the indiscernibles by ultrafilters, as suggested by the proof of Theorem 1.11. We will conclude this section by discussing the relationship between ultrafilters and sets of indiscernibles.

1.5. From Sharps to the Core Model

The Dodd-Jensen core model K^{DJ} was briefly described at the end of Sect. 1.3. This subsection will explain how the concept of a mouse, which they invented for this model, generalizes and extends the concept of sharps.

Recall that they defined a *mouse* to be a model $L_{\alpha^M}[U_M]$, similar to the models M_x used in the proof of Theorem 1.11, but with the strong theory ZF^- replaced with a fine structural condition. One of the uses of this fine structural condition was to allow them to define iterated ultrapowers of mice and to show that they have the required properties. Silver, Magidor and others later gave a construction of K^{DJ} without the need for fine structure, but fine structure is still needed to define core models for larger cardinals.

As in the proof of Theorem 1.11, iterated ultrapowers can be used to compare two mice. This comparison process prewellorders the class of mice, and shows that they form, in an appropriate sense, a hierarchy with condensation. The Dodd-Jensen core model is defined to be the model $K = L[\mathcal{M}]$, where $\mathcal{M}$ is the class of all mice; the well-ordering of mice and their condensation properties imply that $L[\mathcal{M}]$ is a model of ZF + GCH.

This model cannot contain a measurable cardinal, but Dodd and Jensen proved a covering lemma for K which asserts that a model $L[U]$ with a measurable cardinal has approximately the same relation to K that $0^{\#}$ has to L.

The Dodd-Jensen core model can be better understood by considering a translation from mice to sharps and vice versa. Suppose first that $M = L_\alpha[U]$ is a mouse. Then the result of the iterated ultrapower $i^U_{\mathrm{On}} : M \to M^* = \mathrm{Ult}_{\mathrm{On}}(M, U)$ is a well-founded model by clause (ii) of the definition of a mouse. The ordinals On^{M^*} of M^* have length greater than On, and On is the measurable cardinal in M^*. If we write i^U_λ for the embedding from $L_\alpha[U]$ to $\mathrm{Ult}_\lambda(L_\alpha[U], U)$, then the class $I = \{i^U_\lambda(\kappa) : \lambda \in \mathrm{On}\}$ is a class of indiscernibles for M^*.

This class I can be used, as described in the last subsection, to define an initial sequence of the class of sharps. This sequence of sharps can be defined by recursion over On^{M^*}, so that the length of On^{M^*} provides an indication of how long a hierarchy of sharps will be generated before the process generates a model for which I is not a class of indiscernibles. Clause (iii) of the definition of a mouse can be used to show that this final model is M^*.

In discussing the other direction, from sharps to mice, we will use Kunen's notion of a M-ultrafilter:

1.17 Definition. A *normal M-ultrafilter* in Kunen's sense is an ultrafilter on $\mathcal{P}(\kappa) \cap M$, for some κ in M, such that

1. If $f : \kappa \to \kappa$ is in M and $\{\nu < \lambda : f(\nu) < \nu\} \in U$, then there is γ such that $\{\nu : f(\nu) = \gamma\} \in U$, and
2. If $x \subseteq \mathcal{P}(\kappa)$ is a member of M, and $|x|^M = \kappa$, then $U \cap x \in M$.

The second condition enables the ultrapower by U to be iterated, even though $U \notin M$. If $i^U : M \to \mathrm{Ult}(M,U)$ then $U_1 = \{[f]_U : f \in ({}^{\kappa}U) \cap M\}$ is an $\mathrm{Ult}(M,U)$-ultrafilter, which can be used as $i^U(U)$.

A member of the sharp hierarchy is a set which encodes the theory of a model M^*, with parameters taken from a class I of indiscernibles for M^*. A mouse will be a model $L_\alpha[U]$, where α is the least ordinal such that U is not a normal ultrafilter in $L_{\alpha+1}[U]$. We could easily get an M^*-ultrafilter U on any limit point of I by setting $U = \{x \subseteq \lambda : \sup((I \cap \lambda) - x) < \lambda\}$, the filter on λ generated by $I \cap \lambda$. A better construction, however, uses the analog of Theorem 1.14 for M^* to get an M^*-ultrafilter on $\lambda = \min(I)$: let $i : I \to I$ be an increasing map such that $i(\lambda) > \lambda$. By Theorem 1.14 the embedding i extends to a map $i^* : M^* \to M^*$ such that $i = i^* \restriction I$, and $U = \{x \subseteq \lambda : \lambda \in i^*(x)\}$ is a normal M^*-ultrafilter.

2. Beyond One Measurable Cardinal

The next step beyond $L[U]$ is to develop an inner model theory for models with many measurable cardinals. This is straightforward so long as all of the measures have different critical points: If $\mathcal{U} = \langle U_\nu : \nu < \lambda \rangle$ is a sequence of measures, with increasing critical points κ_ν, then the model $L[\mathcal{U}]$ has measures $U_\nu \cap L[\mathcal{U}]$, and (as with the model $L[U]$) no other measures. If it is desired to have several measures on the same cardinal then the answer is less obvious: if U_0 and U_1 are two measures on a cardinal κ then $U_0 \cap L[U_0, U_1] = U_1 \cap L[U_0, U_1]$ by Kunen's Theorem 1.10, so the model $L[U_0, U_1]$ has only one normal ultrafilter.

A way to proceed is suggested by the following observation of Kunen:

2.1 Proposition. *Every measurable cardinal κ has a normal ultrafilter U_κ which concentrates on nonmeasurable cardinals.*

Proof. Suppose as an induction hypothesis that for each measurable cardinal $\lambda < \kappa$ there is a normal ultrafilter U_λ concentrating on nonmeasurable cardinals, and let U be a normal ultrafilter on κ. If U concentrates on nonmeasurable cardinals then set $U_\kappa = U$; otherwise

$$U' = [\langle U_\lambda : \lambda < \kappa \rangle]_U = \{x \subseteq \kappa : \{\lambda < \kappa : x \cap \lambda \in U_\lambda\} \in U\}$$

is a second normal ultrafilter on κ which concentrates on nonmeasurable cardinals. In this case set $U_\kappa = U'$. ⊣

Note that in the second case of the proof, the model $L[\langle U_\lambda : \lambda < \kappa \rangle, U', U]$ is a model with at least two normal ultrafilters U' and U on κ. The following partial order captures the relation between U' and U:

2.2 Definition. If U and U' are normal ultrafilters on a cardinal κ then we write $U' \lhd U$ if $U' \in \mathrm{Ult}(V, U)$.

Thus $U' \lhd U$ if and only if there is a function f such that

$$\{\alpha < \kappa : f(\alpha) \text{ is a normal ultrafilter on } \alpha\} \in U$$

and

$$U' = \{x \subseteq \kappa : \{\lambda < \kappa : x \cap \lambda \in f(\lambda)\} \in U\}.$$

The argument which Scott used to prove that there are no measurable cardinals in L proves the follow proposition:

2.3 Proposition. *The ordering $\lhd$ is well-founded.*

Proof. Assuming the contrary, let κ be the least cardinal such that the normal ultrafilters on κ are not well-founded by $\lhd$. Then there is a normal ultrafilter U on κ so that $\{U' : U' \lhd U\}$ is not well-founded by $\lhd$. The normal ultrafilters on κ in $\mathrm{Ult}(V, U)$ are exactly the normal ultrafilters U' on κ in V such that $U' \lhd U$, and the $\lhd$-ordering on these normal ultrafilters in $\mathrm{Ult}(V, U)$ is the same as in V since V and $\mathrm{Ult}(V, U)$ have the same functions from κ to V_κ. Thus the measures on κ in $\mathrm{Ult}(V, U)$ are not well-founded under $\lhd$. Since $\mathrm{Ult}(V, U)$ is well-founded it follows that $\mathrm{Ult}(V, U)$ satisfies that the measures on κ are not well-founded by $\lhd$, but this is impossible since, by the elementarity of the embedding i, there is no cardinal $\lambda < i(\kappa)$ in $\mathrm{Ult}(V, U)$ such that the measures on λ are not well-founded by $\lhd$. $\dashv$

2.4 Definition. The order $o(U)$ of a normal ultrafilter U is its rank in the ordering $\lhd$, that is, $o(U) = \sup\{o(U') + 1 : U' \lhd U\}$. The order of a cardinal κ is $o(\kappa) = \sup\{o(U) + 1 : U \text{ is a normal ultrafilter on } \kappa\}$.

Thus a measure U has order 0 if and only if the set of smaller measurable cardinals is not a member of U. A cardinal κ has order 0 if it is not measurable, and order 1 if it is measurable, but has no measures concentrating on smaller measurable cardinals. Since each measure $U' \lhd U$ is equal to $[f]_U$ for some $f : \kappa \to V_\kappa$, we have the following upper bound:

2.5 Proposition (Solovay). *If U is a normal ultrafilter on a measurable cardinal κ then $o(U) < (2^\kappa)^+$, and hence $o(\kappa) \le (2^\kappa)^+$.*

Under the GCH, it follows that $o(\kappa) \le \kappa^{++}$.

The inner models $L[\mathcal{U}]$ for sequences of measures utilize this ordering $\lhd$. We give here the original presentation of these models as in [25]. This presentation is the simplest way to approach these models, and we will show that it can be generalized with the use of extenders to define inner models with a strong cardinal. However it is not adequate for dealing with cardinals very much larger than this; and in Sect. 3 we will follow up by giving a brief description of the modified presentation now used for the core model and for inner models with larger cardinals.

2.6 Definition. A *coherent sequence of measures* is a function $\mathcal{U}$ such that

1. $\mathrm{dom}(\mathcal{U}) = \{(\kappa, \beta) : \kappa < \mathrm{len}(\mathcal{U}) \text{ and } \beta < o^{\mathcal{U}}(\kappa)\}$, where $\mathrm{len}(\mathcal{U})$ is a cardinal and $o^{\mathcal{U}}$ is a function mapping cardinals $\kappa < \mathrm{len}(\mathcal{U})$ to ordinals.

2. If $(\kappa, \beta) \in \mathrm{dom}(\mathcal{U})$ then $U = \mathcal{U}(\kappa, \beta)$ is a normal ultrafilter on κ.

3. If $U = \mathcal{U}(\kappa, \beta)$ then $o^{i^U(\mathcal{U})}(\kappa) = \beta$ and $i^U(\mathcal{U})(\kappa, \beta') = \mathcal{U}(\kappa, \beta')$ for all $\beta' < \beta$.

The final clause of this definition is the *coherence condition*, which can also be expressed by saying that $(i^{\mathcal{U}(\kappa,\beta)}(\mathcal{U}))\restriction\kappa+1 = \mathcal{U}\restriction(\kappa, \beta)$. Here we write $\mathcal{U}\restriction(\kappa, \beta)$ for the restriction of $\mathcal{U}$ to

$$\{(\kappa', \beta') \in \mathrm{dom}(\mathcal{U}) : \kappa' < \kappa \vee (\kappa' = \kappa \wedge \beta' < \beta)\}$$

and $\mathcal{U}\restriction\lambda$ for $\mathcal{U}\restriction(\lambda, 0)$. Notice that the coherence condition implies that $\mathcal{U}(\kappa, \beta') \lhd \mathcal{U}(\kappa, \beta)$ for all $\beta' < \beta < o^{\mathcal{U}}(\kappa)$, so that $o(\mathcal{U}(\kappa, \beta)) \geq \beta$.

The following theorem is the main result of [25]; it is a generalization of the Corollary 1.10 to Theorem 1.9:

2.7 Theorem. *If $\mathcal{U}$ is a coherent sequence of measures in $L[\mathcal{U}]$, then the only normal ultrafilters in $L[\mathcal{U}]$ are the members of the sequence $\mathcal{U}$.*

It follows from Theorem 2.7 that $o(\mathcal{U}(\kappa, \beta))$ is exactly equal to β in $L[\mathcal{U}]$, and that each cardinal κ has exactly $|o^{\mathcal{U}}(\kappa)|$ many normal ultrafilters in $L[\mathcal{U}]$. Theorem 1.9 itself does not generalize to $L[\mathcal{U}]$: starting in a model with κ^+ measurable cardinals, where κ is measurable, it is possible to construct sequences $\mathcal{U}$ and $\mathcal{U}'$ with the same domain such that $L[\mathcal{U}] \neq L[\mathcal{U}']$.

2.1. The Comparison Process

The main difficulty in generalizing the proofs of Theorem 1.9 and Theorem 1.11 to the models $L[\mathcal{U}]$ is in adapting the iterated ultrapowers used in those proofs. Recall that they used the iterated ultrapower $\mathrm{Ult}_\lambda(L_\alpha[U], U)$, where λ is some larger cardinal, so that the image i^U_λ of U is contained in the closed unbounded filter $\mathcal{C}_\lambda$ on λ. Thus two normal ultrafilters U and U' were compared indirectly, via the filter $\mathcal{C}_\lambda$. In adapting this construction to the models $L[\mathcal{U}]$ we use iterated ultrapowers to compare sequences $\mathcal{U}$ and $\mathcal{U}'$ directly, through a process known as *iterating the least difference*.

2.8 Lemma (Comparison). *Suppose that $L_\alpha[\mathcal{U}]$ and $L_{\alpha'}[\mathcal{U}']$ satisfy* ZF^-, *and that $\mathcal{U}$ and $\mathcal{U}'$ are coherent in $L_\alpha[\mathcal{U}]$ and $L_{\alpha'}[\mathcal{U}']$ respectively. Then for some sequence $\mathcal{W}$ and some ordinals $\bar\alpha$ and $\bar\alpha'$, there are iterated ultrapowers $i : L_\alpha[\mathcal{U}] \to L_{\bar\alpha}[\mathcal{W}\restriction\bar\alpha]$ and $i' : L_{\alpha'}[\mathcal{U}'] \to L_{\bar\alpha'}[\mathcal{W}\restriction\bar\alpha']$.*

Notice that the sequence $\mathcal{W}$ plays the role of the closed unbounded filter in the proof of Theorem 1.9

For simplicity, this statement of Lemma 2.8 assumes that the models being compared are sets; however the process can also be used if one or both of the models is a proper class, that is, if one or both of $L_\alpha[\mathcal{U}]$ or $L_{\alpha'}[\mathcal{U}]$ is replaced by $L[\mathcal{U}]$ or $L[\mathcal{U}']$. In this case one or both of the iterated ultrapowers may have length On, and one (but never both) of the models $L_{\bar\alpha}[\mathcal{W}\restriction\bar\alpha]$ or $L_{\bar\alpha'}[\mathcal{W}\restriction\bar\alpha']$ may have length larger than On. The simplest example of this is when $L[\mathcal{U}] = L[U]$, with a single normal ultrafilter, and $L[\mathcal{U}'] = L$. Then the comparison consists of iterating U, the only ultrapower available, and U must be iterated On many times to move it past L. Thus $L_{\bar\alpha}[\mathcal{W}\restriction\bar\alpha] = \mathrm{Ult}_{\mathrm{On}}(L[U], U)$, which has On as its measurable cardinal and has length greater than On.

Proof of Lemma 2.8. An iterated ultrapower of length θ of a model M is a family of maps $i_{\nu,\nu'} : M = M_\nu \to M_{\nu'}$, commuting in the sense that $i_{\nu,\nu''} = i_{\nu',\nu''} \circ i_{\nu,\nu'}$ whenever $\nu < \nu' < \nu'' < \theta$, which is defined by recursion on ordinals $\nu < \nu' < \theta$ by setting $M_0 = M$ and $M_\nu = \operatorname{dir\,lim}_{\nu'<\nu} M_{\nu'}$ for each limit $\nu \le \theta$, and for successor ordinals $\nu+1 < \theta$ letting either $M_{\nu+1} = M_\nu$ or else $M_{\nu+1} = \mathrm{Ult}(M_\nu, U_\nu)$ for some M_ν-measure U_ν. The iterated ultrapowers used in this proof are *internal*, which means that $U_\nu \in M_\nu$ for all $\nu < \theta$. We write i_ν for $i_{0,\nu}$.

The proof given before Lemma 1.9 that the model $L[U]$ is iterable relied on the fact that every iterated ultrapower of $L[U]$ is internal to $L[U]$ in the stronger sense that every iterated ultrapower of $L[U]$ is definable in $L[U]$. That is not true for the models $L[\mathcal{U}]$ described here, since the choice of which ultrafilters to use in the iteration may be made externally to $L[\mathcal{U}]$; however, every iterated ultrapower of $L[\mathcal{U}]$ can be embedded into an iterated ultrapower which is definable in $L[\mathcal{U}]$, and thus the argument before Lemma 1.9 shows that every iterated ultrapower of $L[\mathcal{U}]$ is well-founded.

We define two iterated ultrapowers: $i_{\nu',\nu} : M_{\nu'} \to M_\nu$ on $L[\mathcal{U}]$ and $i'_{\nu',\nu} : M'_{\nu'} \to M'_\nu$ on $L[\mathcal{U}']$, as follows: Suppose that $i_\nu : M_0 = L_\alpha[\mathcal{U}] \to M_\nu = L_{\alpha_\nu}[\mathcal{U}_\nu]$ and $i'_\nu : M'_0 = L_{\alpha'}[\mathcal{U}'] \to M'_\nu = L_{\alpha'_\nu}[\mathcal{U}'_\nu]$ have already been defined, where $\mathcal{U}_\nu = i_\nu(\mathcal{U})$ and $\mathcal{U}'_\nu = i'_\nu(\mathcal{U}')$. Let $\gamma = \min\{\alpha_\nu, \alpha'_\nu\}$. If $\mathcal{U}'_\nu\restriction\gamma = \mathcal{U}_\nu\restriction\gamma$ then we are finished, since we can take $\bar\alpha = \alpha_\nu$ and $\bar\alpha' = \alpha'_\nu$ and let $\mathcal{W}$ be the longer of the sequences $\mathcal{U}_\nu$ and $\mathcal{U}'_\nu$. In this case we say that the comparison *terminates at stage* ν.

Otherwise we define $M_{\nu+1}$ and $M'_{\nu+1}$ by the process of *iterating the least difference*: Let (κ_ν, β_ν) be the lexicographically least pair of ordinals such that $\kappa_\nu < \gamma$, $\beta_\nu \le \min\{o^{\mathcal{U}_\nu}(\kappa_\nu), o^{\mathcal{U}'_\nu}(\kappa_\nu)\}$, and

$$\mathcal{U}_\nu(\kappa_\nu, \beta_\nu) \ne \mathcal{U}'_\nu(\kappa_\nu, \beta_\nu), \tag{17.3}$$

where the inequality (17.3) may hold either because $o^{\mathcal{U}_\nu}(\kappa_\nu) \ne o^{\mathcal{U}'_\nu}(\kappa_\nu)$ and $\beta_\nu = \min\{o^{\mathcal{U}_\nu}(\kappa_\nu), o^{\mathcal{U}'_\nu}(\kappa_\nu)\}$ (so that only one side of (17.3) is defined) or because there is a set $x_\nu \in M_\nu \cap M'_\nu$ such that $x_\nu \in \mathcal{U}_\nu(\kappa_\nu, \beta_\nu) \not\Longleftrightarrow x_\nu \in \mathcal{U}'_\nu(\kappa_\nu, \beta_\nu)$. If $\beta_\nu = o^{\mathcal{U}_\nu}(\kappa_\nu)$ then set $M_{\nu+1} = M_\nu$; otherwise set $M_{\nu+1} = \mathrm{Ult}(M_\nu, \mathcal{U}(\kappa_\nu, \beta_\nu))$. Similarly, $M'_{\nu+1} = \mathrm{Ult}(M'_\nu, \mathcal{U}'_\nu(\kappa_\nu, \beta_\nu))$ if $\beta_\nu < o^{\mathcal{U}'_\nu}(\kappa_\nu)$, and $M'_{\nu+1} = M'_\nu$ otherwise.

Note that the ordinals κ_ν are strictly increasing: we have $\mathcal{U}_{\nu+1}\restriction\kappa_\nu + 1 = \mathcal{U}_\nu\restriction(\kappa_\nu, \beta_\nu) = \mathcal{U}'_\nu\restriction(\kappa_\nu, \beta_\nu) = \mathcal{U}'_{\nu+1}\restriction\kappa_\nu + 1$, where the outer equalities follow from the coherence condition and the inner equality follows from the minimality of the pair (κ_ν, β_ν). It follows that $i_{\nu+1,\nu'}(\kappa_\nu) = i'_{\nu+1,\nu'}(\kappa_\nu) = \kappa_\nu$ for all $\nu' > \nu$.

In order to complete the proof of the lemma, we need to show that this comparison eventually terminates. The proof relies on the following observation:

2.9 Claim. *Suppose that τ is an infinite regular cardinal and $\langle N_\nu : \nu < \tau\rangle$ is an iterated ultrapower with embeddings $j_{\nu,\nu'} : N_\nu \to N_{\nu'}$. Further suppose that $|N_0| < \tau$, that $S \subseteq \tau$ is stationary, and that $y_\nu \in N_\nu$ for each $\nu \in S$. Then there is a stationary set $S' \subseteq S$ such that $j_{\nu,\nu'}(y_\nu) = y_{\nu'}$ for all $\nu < \nu'$ in S'.*

Proof. For each limit $\nu \in S$ there is an ordinal $\gamma < \nu$ such that $y_\nu \in \operatorname{ran}(j_{\gamma,\nu})$, so by Fodor's Lemma there is a fixed $\gamma_0 < \tau$ such that the set of $\nu \in S$ such that $y_\nu \in \operatorname{ran}(j_{\gamma_0,\nu})$ is stationary. Since $|N_{\gamma_0}| \leq \max\{|N_0|, |\gamma_0|\} < \tau$, there is a fixed $\bar{y} \in N_{\gamma_0}$ such that $S' = \{\nu \in S : y_\nu = j_{\gamma_0,\nu}(\bar{y})\}$ is stationary. Now if $\nu < \nu'$ are in S' then $y_{\nu'} = j_{\gamma_0,\nu'}(\bar{y}) = j_{\nu,\nu'} j_{\gamma_0,\nu}(\bar{y}) = j_{\nu,\nu'}(y_\nu)$. ⊣

Set $\tau = (\max\{\alpha, \alpha'\})^+$, and suppose for the sake of contradiction that the comparison does not terminate in fewer then τ steps. By applying the claim successively to the two iterations of the comparison we get a stationary subset S_0 of τ such that for any two ordinals $\nu < \nu'$ in S_0 we have $i_{\nu,\nu'}(\kappa_\nu) = i'_{\nu,\nu'}(\kappa_\nu) = \kappa_{\nu'}$. It follows that $\beta_\nu < \min\{o^{\mathcal{U}_\nu}(\kappa_\nu), o^{\mathcal{U}'_\nu}(\kappa_\nu)\}$ for all $\nu \in S_0$, for otherwise if we take any $\nu' > \nu$ in S_0, then either $i_{\nu,\nu'}(\kappa_\nu) = \kappa_\nu < \kappa_{\nu'}$ or $i'_{\nu,\nu'}(\kappa_\nu) = \kappa_\nu < \kappa_{\nu'}$.

Thus x_ν is defined for each $\nu \in S_0$, and by applying the claim twice again we get a stationary set $S_1 \subseteq S_0$ such that $i_{\nu,\nu'}(x_\nu) = i'_{\nu,\nu'}(x_\nu) = x_{\nu'}$ for each $\nu < \nu'$ in S_1. But this is impossible, for since $i_{\nu+1,\nu'}(\kappa_\nu) = i'_{\nu+1,\nu'}(\kappa_\nu) = \kappa_\nu$ it follows that

$$\begin{aligned} x_\nu \in \mathcal{U}_\nu(\kappa_\nu, \beta_\nu) &\iff \kappa_\nu \in i_{\nu,\nu+1}(x_\nu) \\ &\iff \kappa_\nu \in i_{\nu,\nu'}(x_\nu) = x_{\nu'} = i'_{\nu,\nu'}(x_\nu) \\ &\iff \kappa_\nu \in i'_{\nu,\nu+1}(x_\nu) \\ &\iff x_\nu \in \mathcal{U}'_\nu(\kappa_\nu, \beta_\nu), \end{aligned}$$

contradicting the choice of x_ν and thus completing the proof of Lemma 2.8. ⊣

As an example of the use of Lemma 2.8, we sketch the proof of Theorem 2.7:

Sketch of Proof. Suppose that Theorem 2.7 is false, so that there is a sequence $\mathcal{U}$ such that $\mathcal{U}$ is coherent in $L[\mathcal{U}]$ and $L[\mathcal{U}]$ contains a normal ultrafilter U which is not a member of the sequence $\mathcal{U}$. We can assume that

Lemma 2.7 does hold for every proper initial segment of the sequence $\mathcal{U}$, that κ and $\beta = o(U)$ are the smallest ordinals such that there is a normal ultrafilter U on κ in $L[\mathcal{U}]$ with $o(U) = \beta$ which is not in the sequence $\mathcal{U}$, and that U is the first such ultrafilter in the order of construction of $L[\mathcal{U}]$. Note that all of these statements can be expressed by sentences in $L[\mathcal{U}]$.

Now apply Lemma 2.8 to the models $L[\mathcal{U}]$ and $\mathrm{Ult}(L[\mathcal{U}], U)$ (with $\alpha = \alpha' = \mathrm{On}$). We must also have $\bar{\alpha} = \bar{\alpha}' = \mathrm{On}$; otherwise if, for example, $\mathrm{On} = \bar{\alpha} < \bar{\alpha}'$, then the lemma would fail in $L_{\bar{\alpha}}[\mathcal{W}\restriction\bar{\alpha}] = L[\mathcal{W}\restriction \mathrm{On}]$, which contradicts the fact that $L_{\bar{\alpha}'}[\mathcal{W}]$ satisfies the sentence stating that Lemma 2.7 does hold for every proper initial segment of $\mathcal{W}$. Thus we have the following diagram:

$$\begin{array}{ccccc} & & \mathrm{Ult}(L[\mathcal{U}], U) & & \\ & \overset{i^U}{\nearrow} & & \overset{i}{\searrow} & \\ L[\mathcal{U}] & & \xrightarrow{\quad i' \quad} & & L[\mathcal{W}] \end{array} \tag{17.4}$$

This diagram obviously commutes on definable members of $L[\mathcal{U}]$, but since the diagram itself is definable in $L[\mathcal{U}]$, the least element of $L[\mathcal{U}]$ for which it failed to commute would be definable. Hence diagram (17.4) commutes.

In particular $i'(\kappa) = i\,i^U(\kappa)$, so $i'(\kappa) > \kappa$. Since $i\restriction\kappa$ and $i'\restriction\kappa$ are the identity it follows that i' begins with an ultrapower by a normal ultrafilter on κ; that is, $\beta_0 = \beta = o(U) < o^{\mathcal{U}}(\kappa)$ and $i' = i'_\theta = i'_{1,\theta}\, i^{\mathcal{U}(\kappa,\beta)}$. But now $U = \mathcal{U}(\kappa, \beta)$, for if x is any subset of κ in $L[\mathcal{U}]$ then $i\,i^U(x) = i'(x) = i'_{1,\theta} i^{\mathcal{U}(\kappa,\beta)}(x)$, so

$$\begin{array}{rclcl} x \in U & \iff & \kappa \in i i^U(x) & \iff & \kappa \in i'_{1,\theta}\, i^{\mathcal{U}(\kappa,\beta)}(x) \\ & \iff & x \in \mathcal{U}(\kappa, \beta). & & \end{array}$$

⊣

Models $L[\mathcal{U}]$ with higher order measures are more difficult to obtain than the model $L[U]$ with one measure. One might try to proceed by analogy with the model $L[U]$, choosing a coherent sequence $\mathcal{U}$ in V and using the model $L[\mathcal{U}]$, but this fails on two counts. In the first place it is not clear that there is a coherent sequence $\mathcal{U}$ in V, for example it is not known whether $o(\kappa) = \omega$ implies that there is a coherent sequence $\mathcal{U}$ of measures in V with $o^{\mathcal{U}}(\kappa) = \omega$. In the second place, if $o(\kappa) > \kappa^+$ then it is not clear that a sequence which is coherent in V need be coherent in $L[\mathcal{U}]$. The first construction of an inner model of $o(\kappa) = \omega$ from the assumption $o^V(\kappa) = \omega$ used the covering lemma; however we outline a proof which avoids this. Call a sequence $\mathcal{U}$ *weakly coherent* if it satisfies conditions 1 and 2 of Definition 2.6, together with the following weakened coherence condition: if $(\kappa, \beta) \in \mathrm{dom}(\mathcal{U})$ and $U = \mathcal{U}(\kappa, \beta)$ then $o^V(U) = \beta$.

We first show that the comparison process can be modified to use sequences which are only weakly coherent. Notice that this proof requires that $\mathcal{U}$ and $\mathcal{W}$ be sequences of measures in V, not just in $L[\mathcal{U}]$ and $L[\mathcal{W}]$. The example described following Theorem 2.7 shows that this hypothesis cannot be eliminated.

2.10 Lemma. *Suppose that $\mathcal{U}$ and $\mathcal{W}$ are weakly coherent sequences of measures in V with the same domain. Then $L[\mathcal{U}] = L[\mathcal{W}]$, and $\mathcal{U}(\kappa,\beta) \cap L[\mathcal{U}] = \mathcal{W}(\kappa,\beta) \cap L[\mathcal{W}]$ for every (κ,β) in their common domain.*

Proof. We compare the model V with itself, using iterated ultrapowers $i_\nu : V = M_0 \to M_\nu$ and $j_\nu : V = N_0 \to N_\nu$. The comparison process is similar to that in Lemma 2.8 except that we simultaneously compare each of the sequences $i_\nu(\mathcal{U})$ and $i_\nu(\mathcal{W})$ in M_ν with each of the sequences $j_\nu(\mathcal{U})$ and $j_\nu(\mathcal{W})$ in N_ν. Thus, condition (17.3) of the proof of Lemma 2.8 is modified as follows: Suppose that M_ν and N_ν have been defined. Define o^{M_ν} and o^{N_ν} by setting $o^{M_\nu}(\kappa) = o^{i_\nu(\mathcal{U})}(\kappa) = o^{i_\nu(\mathcal{W})}(\kappa)$ and $o^{N_\nu}(\kappa) = o^{j_\nu(\mathcal{U})}(\kappa) = o^{j_\nu(\mathcal{W})}(\kappa)$. Now let (κ_ν, β_ν) be the least pair (κ,β) such that one of the following hold:

1. $\beta < \min\{o^{M_\nu}(\kappa), o^{N_\nu}(\kappa)\}$ and there is a set in $M_\nu \cap N_\nu$ on which the four filters $i_\nu(\mathcal{U})(\kappa,\beta)$, $i_\nu(\mathcal{W})(\kappa,\beta)$, $j_\nu(\mathcal{U})(\kappa,\beta)$ and $j_\nu(\mathcal{W})(\kappa,\beta)$ do not all agree.

2. $\beta = \min\{o^{M_\nu}(\kappa), o^{N_\nu}(\kappa)\}$ and $o^{M_\nu}(\kappa) \neq o^{N_\nu}(\kappa)$.

Now proceed with a slightly modified version of the proof of Lemma 2.8. In case 1 pick U_ν to be one of $\{i_\nu(\mathcal{U})(\kappa,\beta), i_\nu(\mathcal{W})(\kappa,\beta)\}$ and U'_ν to be one of $\{j_\nu(\mathcal{U})(\kappa,\beta), j_\nu(\mathcal{W})(\kappa,\beta)\}$ so that $U_\nu \cap M_\nu \cap N_\nu \neq U'_\nu \cap M_\nu \cap N_\nu$, and set $M_{\nu+1} = \mathrm{Ult}(M_\nu, U_\nu)$ and $N_{\nu+1} = \mathrm{Ult}(N_\nu, U'_\nu)$. In case 2 let $M_{\nu+1} = M_\nu$ if $o^{M_\nu}(\kappa) = \beta$ and $M_{\nu+1} = \mathrm{Ult}(M_\nu, i_\nu(\mathcal{U})(\kappa,\beta))$ if $\beta < o^{M_\nu}(\kappa)$, and define $N_{\nu+1}$ similarly.

Unlike the proof of Lemma 2.8, the sequence of ordinals κ_ν need not be strictly increasing; however the sequence is nondecreasing and the fact that $\beta_{\nu+1} < \beta_\nu$ whenever $\kappa_{\nu+1} = \kappa_\nu$ implies that for each ν there is an $n < \omega$ such that $\kappa_{\nu+n} > \kappa_\nu$. This, together with the weak coherence of $\mathcal{U}$ and $\mathcal{W}$, is enough to show that the comparison terminates at some stage θ.

There is a λ such that either $o^{i_\theta(\mathcal{U})} = o^{j_\theta(\mathcal{U})}\restriction\lambda$ or $o^{j_\theta(\mathcal{U})} = o^{i_\theta(\mathcal{U})}\restriction\lambda$; we may assume the former. We will show that $L[i_\theta(\mathcal{U})] = L[i_\theta(\mathcal{W})]$, and since i_θ is an elementary embedding it follows that $L[\mathcal{U}] = L[\mathcal{W}]$, as was to be shown.

The four sequences $i_\theta(\mathcal{U})\restriction\lambda$, $i_\theta(\mathcal{W})\restriction\lambda$, $j_\theta(\mathcal{U})$ and $j_\theta(\mathcal{W})$ agree on sets in $M_\theta \cap N_\theta$, and thus $L[i_\theta(\mathcal{U})] = L[i_\theta(\mathcal{W})]$ will follow if we can show that $L[i_\theta(\mathcal{U})] \subseteq M_\theta \cap N_\theta$. Suppose the contrary, and let α be least such that there is a set in $L_{\alpha+1}[i_\theta(\mathcal{U})]$ which is not in $M_\theta \cap N_\theta$. Now $i_\theta(\mathcal{U})$ and $j_\theta(\mathcal{U})$ agree on all sets in $L_\alpha[i_\theta(\mathcal{U})]$. Thus $L_\alpha[i_\theta(\mathcal{U})] = L_\alpha[j_\theta(\mathcal{U})]$, and the restrictions of $i_\theta(\mathcal{U})$ and $j_\theta(\mathcal{U})$ to this set are equal. However $L_{\alpha+1}[i_\theta(\mathcal{U})]$ is equal to the set of subsets of $L_\alpha[i_\theta(\mathcal{U})]$ definable over $L_\alpha[i_\theta(\mathcal{U})]$ using as a predicate the restriction of $i_\theta(\mathcal{U})$ to $L_\alpha[i_\theta(\mathcal{U})]$, and similarly for $L_{\alpha+1}[j_\theta(\mathcal{U})]$. Hence $L_{\alpha+1}[i_\theta(\mathcal{U})] = L_{\alpha+1}[j_\theta(\mathcal{U})]$, and it follows that $L_{\alpha+1}[i_\theta(\mathcal{U})] \subseteq M_\theta \cap N_\theta$, contradicting the choice of α. This contradiction completes the proof of Lemma 2.10. ⊣

2.11 Corollary. *If $\mathcal{U}$ is weakly coherent in V, then either $\mathcal{U}$ is coherent in $L[\mathcal{U}]$ or there is an inner model of $\exists\kappa(o(\kappa) = \kappa^{++})$.*

Sketch of Proof. Let $\mathcal{U}$ be any weakly coherent sequence which is not coherent in $L[\mathcal{U}]$. Since initial segments of $\mathcal{U}$ are also weakly coherent, we may assume that $\mathcal{U}$ has minimal length, so that $\mathcal{U}{\restriction}(\kappa,\beta)$ is coherent in $L[\mathcal{U}{\restriction}(\kappa,\beta)]$ for all (κ,β) in the domain of $\mathcal{U}$. In particular, if (κ,β) is the least place at which $\mathcal{U}$ is not coherent in $L[\mathcal{U}]$ then $\mathcal{U}{\restriction}(\kappa,\beta)$ is coherent in $L[\mathcal{U}{\restriction}(\kappa,\beta)]$, and it will be sufficient to show that $L[\mathcal{U}{\restriction}(\kappa,\beta)] \models o(\kappa) = \beta = \kappa^{++}$.

To this end, set $U = \mathcal{U}(\kappa,\beta)$ and consider the following triangle:

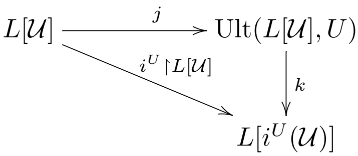

where $i^U : V \to \mathrm{Ult}(V,U)$, j is the ultrapower of $L[\mathcal{U}]$ using functions in $L[\mathcal{U}]$, and k is defined by $k\big(j(f)(\kappa)\big) = i^U(f)(\kappa)$.

We claim that $i^U(\mathcal{U}){\restriction}\kappa+1$ agrees with $\mathcal{U}{\restriction}(\kappa,\beta)$ on all sets in $L[\mathcal{U}]$. To see this, let $\mathcal{U}'$ be the sequence obtained from $\mathcal{U}$ by replacing $\mathcal{U}{\restriction}(\kappa,\beta)$ with $i^U(\mathcal{U}){\restriction}\kappa+1$. Then $\mathcal{U}'$ is weakly coherent, and has the same domain as $\mathcal{U}$, so by Lemma 2.10 it is equal to $\mathcal{U}$ on sets in $L[\mathcal{U}]$.

This implies that k is not the identity on $o^{j(\mathcal{U})}(\kappa)+1$, since otherwise we would have $k(j(\mathcal{U}){\restriction}\kappa+1) = i^U(\mathcal{U}){\restriction}\kappa+1$. Since $L[\mathcal{U}]$ and $L[j(\mathcal{U})]$ have the same subsets of κ, and $i^U(\mathcal{U}){\restriction}\kappa+1$ agrees with $\mathcal{U}$ on these subsets, this would contradict the assumption that $\mathcal{U}$ is not coherent in $L[\mathcal{U}]$ at (κ,β).

Now let $\eta = \mathrm{crit}(k)$. Then $\eta \leq o^{L[j(\mathcal{U})]}(\kappa)$, and since $\beta = k(o^{L[j(\mathcal{U})]}(\kappa))$ it follows that $k(\eta) \leq \beta$. Also $\eta > \kappa$, and η is a cardinal in $L[j(\mathcal{U})]$ and hence in $L[\mathcal{U}]$. But $k(\eta)$ is a cardinal in $L[i^U(\mathcal{U})]$, and hence also in $L[i^U(\mathcal{U}){\restriction}\kappa+1] = L[\mathcal{U}{\restriction}(\kappa,\beta)]$. Thus $\beta \geq k(\eta) \geq \kappa^{++}$ in $L[\mathcal{U}{\restriction}(\kappa,\beta)]$. ⊣

2.2. Indiscernibles from Iterated Ultrapowers

We now look at the use of iterated ultrapowers to generate systems of indiscernibles, and at the relation between these indiscernibles and those added generically by Prikry forcing and its variants. Such forcing is covered extensively in chapter [12].

The simplest case is Prikry forcing [31], which involves only one normal ultrafilter. Let U be a normal ultrafilter on a cardinal κ, and let $i^U_\omega : V \to M_\omega = \mathrm{Ult}_\omega(V,U)$ be the iterated ultrapower of length ω. Then the set $C = \{i^U_n(\kappa) : n < \omega\}$ is a set of indiscernibles over M_ω in the following sense: if x is any subset of $i_\omega(\kappa)$ in M_ω, then there are $n < \omega$ and $x' \subseteq i_n(\kappa)$ in M_n such that $x = i_{n,\omega}(x')$. Then for all $m \geq n$ we have $i_m(\kappa) \in x$ if and only if $x' \in i_n(U)$, which is to say if and only if $x \in i_\omega(U)$. Hence C is almost contained in any set $x \in \mathcal{P}^{M_\omega}(i_\omega(\kappa))$ such that $x \in i_\omega(U)$. By

Mathias's genericity criterion [23], this implies that the sequence C is generic for Prikry forcing over M_ω.

In order to extend this construction to the variants of Prikry forcing discovered by Magidor [21] and Radin [32], let $\mathcal{U} = \langle U_\beta : \beta < \eta\rangle$ be a $\lhd$-increasing sequence of measures on κ, with $o(U_\beta) = \beta$ for $\beta < \eta$ and define an iterated ultrapower $i_\nu : V \to M_\nu$, of length θ, as follows:

As usual, set $M_0 = V$ and set $M_\nu = \operatorname{dir}\lim_{\nu' < \nu} M_{\nu'}$ whenever ν is a limit ordinal. Now suppose that M_ν has been defined. Set $\kappa_\nu = i_\nu(\kappa)$, and let $M_{\nu+1} = \operatorname{Ult}(M_\nu, i_\nu(\mathcal{U})_{\beta_\nu})$ where $\beta_\nu < i_\nu(\eta)$ is the least ordinal β such that $\{\nu' < \nu : i_{\nu',\nu}(\beta_{\nu'}) = \beta\}$ is bounded in ν. If there is no such ordinal β then set $\theta = \nu$ and stop the process.

Assuming $\eta < \kappa^{++}$ and $2^\kappa = \kappa^+$, Fodor's Lemma implies that $\theta < \kappa^{++}$. If $\eta < \kappa$ then a straightforward induction shows that $\theta = \omega^\eta$, and that β_ν is always the least ordinal β such that $\nu = \nu' + \omega^\beta$ for some ordinal $\nu' < \nu$. In particular $\beta_\nu = 0$ if ν is a successor ordinal.

The set $C = \{i_\nu(\kappa) : \nu < \theta\}$ is a closed unbounded subset of $i_\theta(\kappa)$, since the sequence $\langle i_\nu(\kappa) : \nu < \theta\rangle$ is continuous. If $x \in M_\nu$ and $x \subseteq \kappa_\nu$ then $\kappa_\nu \in i_{\nu,\theta}(x) \iff x \in i_\nu(\mathcal{U})_{\beta_\nu} \iff i_{\nu,\theta}(x) \in i_\theta(\mathcal{U})_\beta$, where $\beta = i_{\nu,\theta}(\beta_\nu)$. Thus the sets $C_\beta = \{\kappa_\nu : \nu < \theta \text{ and } i_{\nu,\theta}(\beta_\nu) = \beta\}$ are sets of indiscernibles for the normal ultrafilters $i_\theta(\mathcal{U})_\beta$ on $i_\theta(\kappa)$.

We have already considered the case $n = 1$, when $C = C_0$ is a Prikry sequence for the normal ultrafilter $i(\mathcal{U}_0)$ on $i_\omega(\kappa)$. If $\eta < \kappa$ is an uncountable regular cardinal then $M_\theta[C]$ is a generic extension of the model M_θ by Magidor's generalization [21] of Prikry forcing: the cardinals of $M_\theta[C]$ are the same as those of M_θ, while κ_θ is regular in M_θ and has cofinality η in $M_\theta[C]$. Notice that $\vec{C} = \langle C_\beta : \beta < \eta\rangle \in M_\theta[C]$, since $C_\beta = \{\lambda \in C : o(\lambda) = \beta\}$.

The covering lemma, which is discussed in a separate chapter [24], implies that these results are the best possible in the sense that if there is a cardinal κ which is regular in the core model but is singular of cofinality $\eta > \omega$ in V, then $o(\kappa) \geq \eta$ in the core model. Furthermore, the singularity of κ is witnessed by a set which is similar to the Prikry-Magidor generic set C described above, but which may be more irregular: it satisfies $o(\nu) \geq \limsup\{o(\nu') + 1 : \nu' \in C \cap \nu\}$, while the Prikry-Magidor generic set satisfies the stronger condition $\forall \nu \in C\ o(\nu) = \limsup\{o(\nu') + 1 : \nu' \in C \cap \nu\}$. The case of $\operatorname{cf}(\kappa) = \omega$ can be somewhat more complicated.

If $\kappa < \eta \leq \kappa^{++}$ then the set C obtained from the iterated ultrapower described above is generic for Radin forcing [32], or rather for the variant of Radin forcing described in [26]. It is a closed unbounded subset of $i_\theta(\kappa)$ and it is eventually contained in every member x of the filter $\bigcap i_\theta(\mathcal{U})$ on $\mathcal{P}^{M_\theta}(i_\theta(\kappa))$. If $\eta \geq \kappa^+$ then the sequence $\langle C_\beta : \beta < i_\theta(\eta)\rangle \notin M_\theta[C]$, and the cardinals of $M_\theta[C]$ are the same as those of M_θ. If $\operatorname{cf}(\eta) = \kappa^+$ then κ_θ remains inaccessible in $M_\theta[C]$, and κ can have stronger larger cardinal properties in $M_\theta[C]$ as the ordinal η becomes larger. For the most important example, define β to be a weak repeat point in the sequence $\mathcal{U}$ if for each set $A \in U_\beta$ there is $\beta' < \beta$ such that $A \in U_{\beta'}$. If $\eta = \beta + 1$, where β is a weak repeat point in $\mathcal{U}$, then $i_\theta(\kappa)$ is measurable in $M_\theta[C]$, with a measure on $i_\theta(\kappa)$ in $M_\theta[C]$ which extends the measure $i(U_\beta)$ in M_θ.

If the set C is obtained by Radin forcing or, equivalently, by an iterated ultrapower as described above, then C is eventually contained in any closed unbounded subset of κ which is a member of the ground model M. It can be shown [28] that if this additional condition is imposed, then neither the hypothesis $o(\kappa) \geq \kappa^+$ for preserving the inaccessibility of κ nor the hypothesis of a weak repeat point for preserving measurability can be weakened. If this condition is removed, however, then work of Gitik [13], improved by Mitchell [28], has shown that if $M \models o(\kappa) = \kappa$ then there is a forcing to add a closed, unbounded set $C \subseteq \kappa$ such that every member of C is inaccessible in M, while κ is still measurable in $M[C]$. Gitik also shows that if $\{\nu < \kappa : o(\kappa) > \beta\}$ is stationary in κ for all $\beta < \kappa$ then κ remains inaccessible in κ. Such sets cannot be obtained by iterated ultrapowers alone, without forcing. Both of these results are the best possible.

3. Extender Models

The next step above the hierarchy of measurable cardinals is the hierarchy leading to a strong cardinal:

3.1 Definition. A cardinal κ is λ-*strong* if there is an elementary embedding $j : V \to M$ such that $\kappa = \text{crit}(j)$, $\lambda < j(\kappa)$, and $\mathcal{P}^\lambda(\kappa) \subseteq M$. A cardinal κ is *strong* if it is λ-strong for every ordinal λ.

A cardinal is 1-strong if and only if it is measurable; however an embedding of the form i^U, where U is an ultrafilter on κ, will never witness that a cardinal κ is 2-strong since $U \in \mathcal{P}^2(\kappa) - \text{Ult}(V, U)$. An *extender* is a generalized ultrafilter designed to represent the stronger embeddings needed for strong cardinals. Extenders can be equivalently defined in either of two different ways, as elementary embeddings or as sequences of ultrafilters. We will begin with the simpler of the two:

3.2 Definition. A (κ, λ)-*extender* is an elementary embedding $\pi : M \to N$ where M and N are transitive models of ZF^-, $\kappa = \text{crit}(\pi)$, and $\lambda \leq \pi(\kappa)$.

The model M need not be a model of ZF; indeed we can typically assume that κ is the largest cardinal in M since $\mathcal{P}^M(\kappa)$ is the only part of M which will be used for the ultrapower construction. Extenders are so called because the embedding π can be extended to an embedding on a full model M' of set theory, provided that the subsets of κ in M' are contained in those of M:

3.3 Definition. Suppose that $\pi : M \to N$ is an extender and M' is a model of set theory such that $\mathcal{P}^{M'}(\kappa) \subseteq \mathcal{P}^M(\kappa)$.

If $a, a' \in [\lambda]^{<\omega}$, and f and f' are functions in M' with domains $[\kappa]^{|a|}$ and $[\kappa]^{|a'|}$ respectively, then we say that $(f, a) \sim_\pi (f', a')$ if and only if

$(a,a') \in \pi(\{(v,v') \in [\kappa]^{|a|} \times [\kappa]^{|a'|} : f(v) = f'(v')\})$. We write $[f,a]_\pi$ for the equivalence class $\{(f',a') : (f,a) \sim_\pi (f',a')\}$.

Finally we write $\mathrm{Ult}(M',\pi)$ for the model with universe

$$\{[f,a]_\pi : f \in {}^\kappa M' \cap M' \;\;\&\;\; a \in {}^{<\omega}\lambda\},$$

and with the membership relation $\in_\pi$ defined by $[f,a]_\pi \in_\pi [f',a']_\pi$ if $(a,a') \in \pi(\{(v,v') : f(v) \in f'(v')\}$.

The ultrapower embedding $i^\pi : M' \to \mathrm{Ult}(M',\pi)$ is defined by $i^\pi(x) = [x,\varnothing]_\pi$. Here x is regarded as a constant, that is, a 0-ary function.

We will only be interested in extenders such that $\mathrm{Ult}(M',\pi)$ is well-founded and hence isomorphic to a transitive model, and we will identify $\mathrm{Ult}(M',\pi)$ with the transitive model to which it is isomorphic.

The ordinal λ is called the *length* of the (κ,λ)-extender π, and is written $\mathrm{len}(\pi)$. The embedding π does not actually itself determine the value of λ, since the same embedding π could be used as to represent a (κ,λ') extender for any $\lambda' < \pi(\kappa)$. When necessary, the ordinal λ may be explicitly specified, for example by writing $\mathrm{Ult}(M',\pi,\lambda)$ instead of $\mathrm{Ult}(M',\pi)$ or $[f,a]_{\pi,\lambda}$ instead of $[f,a]_\pi$.

If $\lambda < \lambda'$ then a natural elementary embedding

$$k : \mathrm{Ult}(M',\pi,\lambda) \to \mathrm{Ult}(M',\pi,\lambda')$$

can be defined by setting $k([f,a]_{\pi,\lambda}) = [f,a]_{\pi,\lambda'}$. It can be that $\mathrm{Ult}(M',\pi,\lambda)$ $= \mathrm{Ult}(M',\pi,\lambda')$ and k is the identity, in which case we will say that the (κ,λ)- and (κ,λ')-extenders defined by π are equivalent. This will happen whenever there is, for each $\nu \in \lambda'$, some $a \in [\lambda]^{<\omega}$ and $f \in M$ such that $[f,a]_\pi = [\mathrm{id},\nu]_\pi$. For example, the $(\kappa,\lambda+1)$- and $(\kappa,\lambda+2)$-extenders determined by π will always be equivalent, since if s is the successor function, $s(\nu) = \nu+1$, then $[s,\{\lambda\}]_\pi = [\mathrm{id},\{\lambda+1\}]_\pi$.

Łoś's Theorem for extender ultrapowers is proved in the same way as the Łoś's Theorem for ultrafilters:

3.4 Proposition (Łoś's Theorem). *Suppose that $\varphi(v_0,\dots,v_{n-1})$ is any formula of set theory, and that $a_i \in [\lambda]^{<\omega}$ for $i < n$ and $f_i : [\kappa]^{|a_i|} \to \lambda$. Then*

$$\mathrm{Ult}(M',\pi) \models \varphi([f_0,a_0]_\pi,\dots,[f_{n-1},a_{n-1}]_\pi)$$

if and only if

$$(a_0,\dots,a_{n-1}) \in \pi\big(\{(v_0,\dots,v_{n-1}) : M' \models \varphi(f_0(v_0),\dots,f_{n-1}(v_{n-1}))\}\big).$$

This statement suggests the alternate definition of an extender as a sequence E of ultrafilters:

3.5 Definition. The ultrafilter sequence representing a (κ,λ)-extender π is the sequence $E^\pi = \langle E_a : a \in [\lambda]^{<\omega}\rangle$ of ultrafilters defined by

$$E_a = \{x \subseteq {}^a\kappa : a \in \pi(\{\mathrm{ran}(v) : v \in x\})\}. \tag{17.5}$$

Here we write ran(v) for the sequence $\langle v(a_i) : i < |a|\rangle \in {}^{|a|}\kappa$, where $a = \langle a_i : i < |a|\rangle$. The use of ran$(v)$ instead of v in the right side of (17.5) is necessary because a need not be a member of M. This complication could have been avoided by equivalently defining E_a to be an ultrafilter on subsets of $[\kappa]^{|a|}$ or ${}^{|a|}\kappa$ instead of on ${}^a\kappa$; however the use of ${}^a\kappa$ simplifies some later notation.

3.6 Definition. The ultrapower $\mathrm{Ult}(M', E)$ is defined to be the direct limit of the commuting system of maps

$$\big(\langle \mathrm{Ult}(M', E_a) : a \in {}^{<\omega}\lambda\rangle, \langle \pi_{a,a'} : \mathrm{ran}(a) \subseteq \mathrm{ran}(a')\rangle\big),$$

where $\pi_{a,a'} : \mathrm{Ult}(M', E_a) \to \mathrm{Ult}(M, E_{a'})$ is defined by setting $\pi_{a,a'}([f]_{E_a}) = [v \mapsto f(v{\restriction}a)]_{E_{a'}}$.

It can easily be shown that if π is a (κ, λ)-extender then $\mathrm{Ult}(M, E^\pi) = \mathrm{Ult}(M, \pi, \lambda)$.

In the future we will follow the usual practice of using the ultrafilter representation for extenders. This generally makes for clearer notation, which among other things does not tie down the variables M and N. It also has the advantage of explicitly incorporating the length λ of the extender, but requires additional notation for the shortened extender: if $E = \langle E_a : a \in [\lambda]^{<\omega}\rangle$ is a (κ, λ)-extender and $\lambda' < \lambda$, then we write $E|\lambda'$ for the subsequence $\langle E_a : a \in [\lambda']^{<\omega}\rangle$ of E. Thus $E|\lambda'$ is the (κ, λ')-extender represented by the embedding π^E.

It may happen that $\mathrm{Ult}(V, E|\lambda') = \mathrm{Ult}(V, E)$, in which case we say that the two extenders are equivalent. This will be true whenever there is, for each $\alpha \in \lambda - \lambda'$, a function f and finite set $a \in [\lambda']^{<\omega}$ such that $[f, a]_E = [\mathrm{id}, \{a\}]_E$ or, equivalently, such that $\{v \in {}^{a\cup\{\alpha\}}\kappa : f(v{\restriction}a) = v(\alpha)\} \in E_{a\cup\{\alpha\}}$. As a simple example, by taking f to be the successor function we can see that $E|(\lambda + 1)$ is always equivalent to $E|(\lambda + 2)$.

The notion of countable completeness is somewhat more complicated for extenders than for ultrafilters:

3.7 Definition. An (κ, λ)-extender E is *countably complete* if for each sequence $(a_i : i \in \omega)$ of sets $a_i \in [\lambda]^{<\omega}$ and each sequence $(X_i : i < \omega)$ of sets $X_i \in E_{a_i}$ there is a function $v : \bigcup_i a_i \to \kappa$ such that $v{\restriction}a_i \in X_i$ for each $i < \omega$.

As in the case of ultrafilters, countably complete extenders are important because they ensure well-foundedness of iterated ultrapowers.

3.8 Definition. Suppose that M is a model of set theory and $\mathcal{E}$ is a collection of extenders in M. An iterated ultrapower of M by extenders in $\mathcal{E}$ is a pair of sequences $\langle M_\nu : \nu \leq \theta\rangle$ and $\langle E_\nu : \nu < \theta\rangle$, together with a commuting system of elementary embeddings $i_{\nu,\nu'} : M_\nu \to M_{\nu'}$, such that $M_0 = M$, if ν is a limit ordinal then M_ν is the direct limit of the models $\langle M_{\nu'} : \nu' < \nu\rangle$ under the embeddings $i_{\nu',\nu''}$, and if $\nu < \theta$ then $E_\nu \in i_{0,\nu}(\mathcal{E})$ and $i_{\nu,\nu+1} : M_\nu \to \mathrm{Ult}(M_\nu, E_\nu) = M_{\nu+1}$.

3.9 Lemma. *If $\mathcal{E}$ is a collection of countably complete extenders then any iterated ultrapower using extenders in $\mathcal{E}$ is well-founded.*

Proof. Suppose to the contrary that we have an iterated ultrapower as in Definition 3.8 with M_θ ill-founded. The initial model M_0 could be a proper class, but in that case M_0 can be replaced by an initial segment of M_0 satisfying ZF$^-$ which exhibits the ill-foundedness; thus we can assume that M_0 is a set.

Fix a regular cardinal τ such that the ill-founded iterated ultrapower is a member of $H(\tau)$, the set of sets which are hereditarily of size less than τ, let $X \prec H(\tau)$ be a countable elementary substructure containing the iterated ultrapower, and let $\sigma : P \cong X$ be the inverse of the transitive collapse map. Set $\bar{\theta} = \sigma^{-1}(\theta)$, and set $\bar{E}_\nu = \sigma^{-1}(E_{\sigma(\nu)})$ and $\bar{M}_\nu = \sigma^{-1}(M_{\sigma(\nu)})$ for each $\nu < \bar{\theta}$.

Set $\bar{\mathcal{E}} = \sigma^{-1}(\mathcal{E})$. Then $(\langle \bar{M}_\nu : \nu \leq \bar{\theta}\rangle, \langle \bar{E}_\nu : \nu < \bar{\theta}\rangle)$ is an ill-founded iterated ultrapower of $\bar{M}_0$ of countable length $\bar{\theta}$, using only extenders from $\bar{\mathcal{E}}$.

We will define a commuting sequence of elementary embeddings

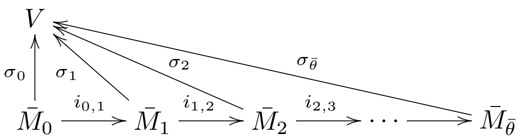

(17.6)

with $\sigma_0 = \sigma \restriction \bar{M}_0$. Thus $\sigma_{\bar{\theta}}$ embeds $\bar{M}_{\bar{\theta}}$ into V, contradicting the assumption that $\bar{M}_{\bar{\theta}}$ is ill-founded and thus completing the proof of the lemma.

The embedding σ_0 has already been defined, and the requirement that the diagram (17.6) commutes determines the choice of σ_α for limit ordinals $\alpha \leq \bar{\theta}$: if $x \in \bar{M}_\alpha$ then $\sigma_\alpha(x) = \sigma_{\alpha'}(i^{-1}_{\alpha,\alpha'}(x))$ where α' is any ordinal less than α such that $x \in i_{\alpha',\alpha}``\bar{M}_{\alpha'}$.

To define $\sigma_{\alpha+1}$, supposing that $\sigma_\alpha : M_\alpha \to H(\tau)$ has been defined, set $\bar{\lambda} = \operatorname{len}(\bar{E}_\alpha)$, and let $\langle (\bar{X}_i, \bar{a}_i) : i < \omega\rangle$ be an enumeration of the set of pairs (X, a) in M_α such that $a \in [\bar{\lambda}]^{<\omega}$ and $X \in (\bar{E}_\alpha)_a$. Then $\sigma_\alpha(\bar{X}_i) \in (\sigma_\alpha(\bar{E}_\alpha))_{\sigma_\alpha(\bar{a}_i)}$, and since $\sigma_\alpha(\bar{E}_\alpha)$ is a member of the collection $\mathcal{E}$ of countably complete extenders there is a function $v : \bigcup_i \sigma_\alpha(\bar{a}_i) \to \sigma_\alpha(\bar{\kappa})$ such that $v \restriction \sigma_\alpha(\bar{a}_i) \in \sigma_\alpha(\bar{X}_i)$ for each $i \in \omega$. Then a straightforward induction shows that the map $\sigma_{\alpha+1} : \bar{M}_{\alpha+1} \to H(\tau)$ defined by setting $\sigma_{\alpha+1}(x) = \sigma_\alpha(f)(v \restriction \sigma_\alpha(a))$ for each $x = [f, a]_{\bar{E}_\alpha} \in \bar{M}_{\alpha+1}$ is an elementary embedding such that $\sigma_\alpha = \sigma_{\alpha+1} \circ \bar{i}_{\alpha,\alpha+1}$. ⊣

This completes the preliminary exposition of extenders, and we now discuss sequences of extenders. The following definition is almost the same as that of a coherent sequence of ultrafilters:

3.10 Definition. A *coherent sequence of nonoverlapping extenders* is a function $\mathcal{E}$ with domain of the form $\{(\kappa, \beta) : \beta < o^{\mathcal{E}}(\kappa)\}$ such that

1. if $o^{\mathcal{E}}(\kappa) > 0$ then $o^{\mathcal{E}}(\lambda) < \kappa$ for every $\lambda < \kappa$,

and if $\beta < o^{\mathcal{E}}(\kappa)$ then

2. $\mathcal{E}(\kappa, \beta)$ is a $(\kappa, \kappa + 1 + \beta)$-extender E, and
3. $i^{\mathcal{E}(\kappa,\beta)}(\mathcal{E})\restriction(\kappa + 1) = \mathcal{E}\restriction(\kappa, \beta)$.

Here $\mathcal{E}\restriction(\kappa, \beta)$ is the restriction of $\mathcal{E}$ to those pairs (κ', β') in its domain which are lexicographically less than (κ, β).

The term *nonoverlapping* refers to clause 1. We will show that nonoverlapping sequences are adequate to construct models with a strong cardinal. It is possible to obtain models with somewhat larger cardinals by weakening clause 1 and modifying the comparison iteration; Baldwin [1] describes a general method of constructing such models. Cardinals very much larger than a strong cardinal, however, require extender sequences $\mathcal{E}$ with overlapping extenders, which greatly complicates the theory of iterated ultrapowers on $L[\mathcal{E}]$, and usually requires the use of iteration trees rather than the linear iterations described in Definition 3.8.

Note that the indexing of the sequences described in Definition 3.10 is the same as that used for sequences of ultrafilters: $\mathcal{E}(\kappa, \beta)$ is the βth extender with critical point κ. This indexing works well for nonoverlapping extenders but fails to be meaningful for sequences with overlapping extenders, where there may be a proper class of extenders with the same critical point κ, and there may be extenders which have critical point κ, but which are stronger than all of the extenders with critical point κ.

All sequences of extenders referred to in this section will be nonoverlapping.

One useful difference between sequences of ultrafilters and sequences of extenders is the fact that the coherence functions for extenders are trivial. The coherence property for a sequence $\mathcal{U}$ of measures depends on the presence, for each $\beta' < \beta < o(\kappa)$, of a function f such that $\beta' = [f]_{\mathcal{U}(\kappa,\beta)}$, or equivalently, such that $\beta' = i^{\mathcal{U}(\kappa,\beta)}(f)(\kappa)$; thus the sequence $\mathcal{U}$ may, for example, be coherent in V but not in $L[\mathcal{U}]$. In the case of a sequence $\mathcal{E}$ of extenders, however, the only coherence function needed is the identity function: if $\beta' < \beta < o^{\mathcal{E}}(\kappa)$ then $\beta' = [\mathrm{id}, \{\beta'\}]_{\mathcal{E}(\kappa,\beta)}$, that is, $\beta' = i^{\mathcal{E}(\kappa,\beta)}(\mathrm{id})(\beta')$. The following proposition, which is not true for sequences of measures, follows immediately:

3.11 Proposition. *If $\mathcal{E}$ is a coherent nonoverlapping sequence of extenders in V and M is an inner model such that the restriction of $\mathcal{E}$ to M is a member of M, then $\mathcal{E}$ is coherent in M.*

In order to define the class $L[\mathcal{E}]$ of sets constructible from $\mathcal{E}$, we can code $\mathcal{E}$ as $\{(\kappa, \beta, a, x) : x \in (\mathcal{E}_{\kappa.\beta})_a\}$. Using this coding, if M is an inner model then $\mathcal{E} \cap M$ is the code for the sequence of restrictions $\langle E_a \cap M : a \in \operatorname{dom} E\rangle$ to M of the extenders E in $\mathcal{E}$.

As we did with sequences of ultrafilters, we need to start with a weaker version of coherence in order to obtain long extender sequences which are coherent in $L[\mathcal{E}]$:

3.12 Definition. A sequence $\mathcal{E}$ of extenders is *weakly coherent* if each extender $E = \mathcal{E}(\kappa, \beta)$ is a $(\kappa, \kappa + 1 + \beta)$-extender such that $o^{i^E(\mathcal{E})}(\kappa) = \beta$.

3.13 Definition. Suppose that N_0 and M_0 are models with countably complete weakly coherent extender sequences $\mathcal{E}_0$ and $\mathcal{F}_0$, respectively. The *comparison iterations* of N_0 and M_0 are defined as follows: Assume $i_\alpha : M_0 \to M_\alpha$ and $j_\alpha : N_0 \to N_\alpha$ have been defined, and let (κ, β) be the least pair such that one of the following holds:

1. $\beta = o^{i_\alpha(\mathcal{E})}(\kappa) < o^{j_\alpha(\mathcal{F})}(\kappa)$.
2. $\beta = o^{j_\alpha(\mathcal{F})}(\kappa) < o^{i_\alpha(\mathcal{E})}(\kappa)$.
3. $\beta < \min\{o^{j_\alpha(\mathcal{F})}(\kappa), o^{i_\alpha(\mathcal{E})}(\kappa)\}$ and there is an $a \in [\kappa + 1 + \beta]^{<\omega}$ and $x \in \mathcal{P}({}^a\kappa) \cap M_\alpha \cap N_\alpha$ such that $x \in (i_\alpha(\mathcal{E})(\kappa, \beta))_a - (j_\alpha(\mathcal{F})(\kappa, \beta))_a$.

If there is no such pair (κ, β) then the sequences $i_\alpha(\mathcal{E})$ and $j_\alpha(\mathcal{F})$ have the same domain and are equal, at least with respect to sets which are in both models. If κ is greater than the length of one of the sequences $i_\alpha(\mathcal{E})$ or $j_\alpha(\mathcal{E})$, that is, if $o^{i_\alpha(\mathcal{E})}(\mu) = 0$ for all $\mu \geq \kappa$ or $o^{j_\alpha(\mathcal{F})}(\mu) = 0$ for all $\mu \geq \kappa$, then one of the sequences is an initial segment of the other (again, at least with respect to sets which are in both models). In either case the process is terminated at this stage.

Otherwise define $i_{\alpha,\alpha+1} : M_\alpha \to M_{\alpha+1}$ to be the ultrapower embedding $i^{i_\alpha(\mathcal{E})(\kappa,\beta)} : M_\alpha \to \mathrm{Ult}(M_\alpha, i_\alpha(\mathcal{E})(\kappa, \beta))$ in cases 2 and 3, and in case 1 define $M_{\alpha+1} = M_\alpha$ and let $i_{\alpha,\alpha+1}$ be the identity. Similarly define $N_{\alpha+1}$ by using the extender $j_\alpha(\mathcal{F})(\kappa, \beta)$ in cases 1 and 3, and set $N_{\alpha+1} = N_\alpha$ in case 2.

The proof that this comparison iteration terminates will use the following proposition, which is proved just like Claim 2.9.

3.14 Proposition. *Suppose that θ is an uncountable regular cardinal, and that we have an iterated extender ultrapower $\langle M_\alpha : \alpha < \theta\rangle$ with iteration embeddings $i_{\alpha',\alpha} : M_{\alpha'} \to M_\alpha$. If X is a set in M_0 such that $|i_\lambda(X)| < \theta$ for each $\lambda < \theta$, and $y_\alpha \in i_{0,\alpha}(X)$ for all $\alpha < \theta$, then for every stationary set $S \subseteq \theta$ there is a stationary set $S' \subseteq S$ such that if $\alpha' < \alpha$ are in S' then $y_\alpha = i_{\alpha',\alpha}(y_{\alpha'})$.*

3.15 Lemma. *If M_0, N_0, $\mathcal{E}$ and $\mathcal{F}$ are as in Definition 3.13, and θ is a regular cardinal such that $\theta \geq \sup\{\, 2^\kappa : o^{\mathcal{E}}(\kappa) > 0$ or $o^{\mathcal{F}}(\kappa) > 0\}$, then the comparison process terminates in fewer than θ steps.*

Proof. Assume the contrary, and at each $\alpha < \theta$ let κ_α and β_α be as in the definition of $M_{\alpha+1}$ and $N_{\alpha+1}$. By applying Proposition 3.14 twice, once to the iterated ultrapower of M_0 and then to that of N_0, we can find a stationary set $S \subseteq \theta$ such that if $\alpha' < \alpha$ are in S then $\kappa_\alpha = i_{\alpha',\alpha}(\kappa_{\alpha'}) = j_{\alpha',\alpha}(\kappa_{\alpha'})$.

Now the sequence $\langle \kappa_\alpha : \alpha < \theta\rangle$ is nondecreasing. Furthermore, whenever $\kappa_{\alpha+1} = \kappa_\alpha$ we have $\beta_{\alpha+1} < \beta_\alpha$, and it follows that for each α there is $k < \omega$

such that $\kappa_\alpha < \kappa_{\alpha+k}$. It follows that $\kappa_{\alpha'} < \kappa_\alpha$ whenever $\alpha' < \alpha$ are limit ordinals.

Now $o^{i_{\alpha+1}(\mathcal{E})}(\kappa_\alpha) = o^{j_{\alpha+1}(\mathcal{F})}(\kappa_\alpha) = \beta_\alpha$ for each $\alpha < \theta$, so case 1 or 2 can only occur at stages α such that $\kappa_{\alpha'} < \kappa_\alpha$ for all $\alpha' < \alpha$. In particular, it never happens that cases 1 and 2 both occur at stages with the same critical point κ_α. For ordinals $\alpha < \alpha'$ in S we have $i_{\alpha,\alpha'}(\kappa_\alpha) = j_{\alpha,\alpha'}(\kappa_\alpha) = \kappa_{\alpha'} > \kappa_\alpha$, so if $\alpha \in S$ and $\alpha^* \geq \alpha$ is the last stage for which $\kappa_{\alpha^*} = \kappa_\alpha$ then case 3 must occur at stage α^*. Finally, let $a_{\alpha^*} \in [\beta_{\alpha^*}]^{<\omega}$ and $x_{\alpha^*} \subseteq [\kappa_\alpha]^{|a_{\alpha^*}|}$ be as in the definition of the comparison at stage α^*. Two more applications of Proposition 3.14 give a stationary $S' \subseteq S$ such that if $\alpha < \gamma$ are in S' then $x_{\gamma^*} = i_{\alpha,\gamma}(x_{\alpha^*}) = j_{\alpha,\gamma}(x_{\alpha^*})$. Set $E_\alpha = i_{0,\alpha}(\mathcal{E})(\kappa_\alpha, \beta_\alpha)$ and $F_\alpha = j_{0,\alpha}(\mathcal{F})(\kappa_\alpha, \beta_\alpha)$. Then we have

$$\begin{aligned} x_\alpha \in (E_{\alpha^*})_{a_{\alpha^*}} \quad &\iff \quad a_{\alpha^*} \in i_{\alpha^*,\alpha^*+1}(x_\alpha) \\ &\iff \quad a_{\alpha^*} \in i_{\alpha^*+1,\gamma} \circ i_{\alpha^*,\alpha^*+1} \circ i_{\alpha,\alpha^*}(x_\alpha) \\ &\qquad\qquad = i_{\alpha,\gamma}(x_\alpha) = x_\gamma, \end{aligned}$$

since $i_{\alpha,\alpha^*}(x_\alpha) \cap [\kappa_\alpha]^{|a_{\alpha^*}|} = x_\alpha$ and $i_{\alpha^*+1,\gamma}(a_{\alpha^*}) = a_\alpha$. Similarly, $x_\alpha \in (F_{\alpha^*})_{a_{\alpha^*}}$ if and only if $a_{\alpha^*} \in j_{\alpha,\gamma}(x_\alpha) = x_\gamma$, and hence $x_{\alpha^*} \in (E_{\alpha^*})_{a_{\alpha^*}}$ if and only if $x_{\alpha^*} \in (F_{\alpha^*})_{a_{\alpha^*}}$. This contradicts the choice of x_{α^*} and hence completes the proof of the lemma. ⊣

The proof of Lemma 3.15 relied crucially on the fact that $i_{\alpha^*+1,\gamma}(a_{\alpha^*}) = j_{\alpha^*+1,\gamma}(a_{\alpha^*}) = a_{\alpha^*}$ for all $\alpha < \gamma$ in S; that is, none of the generators arising from a use of an extender in the iteration is moved by the remainder of the iteration. This problem of *moving generators* is the reason that linear iterations like those used in the proof of Lemma 3.15 are not adequate for comparisons of sequences having overlapping extenders. Thus iteration trees are needed for the analysis of inner models with larger cardinals.

When the comparison process terminates, it is only guaranteed that the sequences match with respect to sets which are in both models, so it is important to observe that this is true of all relevant sets:

3.16 Proposition. *Suppose that the comparison maps $i_\theta : L[\mathcal{E}] \to L[i_\theta(\mathcal{E})]$ and $j_\theta : L[\mathcal{F}] \to L[i_\theta(\mathcal{F})]$ terminate with $i_\theta(\mathcal{E})$ equal to $j_\theta(\mathcal{F})\restriction\eta$ in the sense that $o^{i_\theta(\mathcal{E})} = o^{j_\theta(\mathcal{F})}\restriction\eta$ and*

$$i_\theta(\mathcal{E}) \cap L[i_\theta(\mathcal{E})] \cap L[j_\theta(\mathcal{F})] = (j_\theta(\mathcal{F}))\restriction\eta \cap L[i_\theta(\mathcal{E})] \cap L[j_\theta(\mathcal{F})].$$

Then $L[i_\theta(\mathcal{E})] \subseteq L[j_\theta(\mathcal{F})]$, so that $i_\theta(\mathcal{E}) = (j_\theta(\mathcal{F})\restriction\eta) \cap L[i_\theta(\mathcal{E})]$.

Proof. We prove by induction on α that $L_\alpha[i_\theta(\mathcal{E})] \subseteq L_\alpha[j_\theta(\mathcal{F})]$ for all ordinals α. It is only the successor case that could be problematic: assume as an induction hypothesis that $\alpha < \eta$ and $L_\alpha[i_\theta(\mathcal{E})] \subseteq L_\alpha[j_\theta(\mathcal{F})]$. Notice that it follows that $i_\theta(\mathcal{E}) \cap L_\alpha[i_\theta(\mathcal{E})] = j_\theta(\mathcal{F}) \cap L_\alpha[j_\theta(\mathcal{F})]$ if $\alpha \leq \eta$, and $i_\theta(\mathcal{E}) \cap L_\alpha[i_\theta(\mathcal{E})] = (j_\theta(\mathcal{F})\restriction\eta) \cap L_\alpha[j_\theta(\mathcal{E})]$ if $\alpha > \eta$. In either case both $L_\alpha[i_\theta(\mathcal{E})]$ and $i_\theta(\mathcal{E}) \cap L[i_\theta(\mathcal{E})]$ are definable in $L_\alpha[j_\theta(\mathcal{F})]$, and it follows that $L_{\alpha+1}[i_\theta(\mathcal{E})] \subseteq L_{\alpha+1}[j_\theta(\mathcal{F})]$. ⊣

3.17 Definition. Suppose that $\mathcal{E}$ is a weakly coherent sequence and $\varphi(\mathcal{E})$ is a sentence in the language of set theory. Then $L[\mathcal{E}]$ is said to be φ*-minimal* if $L[\mathcal{E}] \models \varphi(\mathcal{E})$ but there is no proper initial segment $\mathcal{E}' = \mathcal{E}\restriction(\kappa, \beta)$ of $\mathcal{E}$ such that $L[\mathcal{E}'] \models \varphi(\mathcal{E}')$.

3.18 Proposition. *Suppose that $\mathcal{E}$ is weakly coherent and $L[\mathcal{E}]$ is φ-minimal for some formula φ, and suppose that $\pi : L[\mathcal{E}] \to L[\mathcal{E}']$ is an elementary embedding. Then the comparison of $L[\mathcal{E}]$ and $L[\mathcal{E}']$ gives the following diagram:*

$$\begin{array}{ccc} L[\mathcal{E}] & \xrightarrow{\quad j_\theta \quad} & L[j_\theta(\mathcal{E})] \\ & {\scriptstyle\pi}\searrow \qquad \nearrow{\scriptstyle k_\theta} & \\ & L[\mathcal{E}'] & \end{array} \tag{17.7}$$

Furthermore, if π is definable in $L[\mathcal{E}]$ then this diagram commutes.

Proof. If $j_\theta : L[\mathcal{E}] \to L[j_\theta(\mathcal{E})]$ and $k_\theta : L[\mathcal{E}'] \to L[k_\theta(\mathcal{E}')]$ are the two embeddings generated by the comparison process, then Proposition 3.16 implies that one of the two sequences $j_\theta(\mathcal{E})$ and $k_\theta(\mathcal{E}')$ is an initial segment of the other. Since φ-minimality is a first order property, both of the models $L[j_\theta(\mathcal{E})]$ and $L[k_\theta(\mathcal{E}')]$ are φ-minimal; and it follows that neither can be a proper initial segment of the other. Thus $j_\theta(\mathcal{E}) = k_\theta(\mathcal{E}')$.

It follows that the comparison yields the diagram (17.7). To see that the diagram commutes whenever π is definable, suppose the contrary and let x be the least set in the order of construction of $L[\mathcal{E}]$ such that $j_\theta(x) \neq k_\theta \circ \pi(x)$. Since π is definable in $L[\mathcal{E}]$, the set x is also definable, but this is impossible since then $j_\theta(x)$ and $k_\theta \circ \pi(x)$ are both defined in $L[j_\theta(\mathcal{E})]$ by the same formula and hence must be equal. ⊣

3.19 Lemma. *Suppose that $\mathcal{E}$ is a weakly coherent extender sequence and that E is a countably complete $(\kappa, \kappa + 1 + \beta)$-extender in $L[\mathcal{E}]$ such that $o^{i^E(\mathcal{E})}(\kappa) = \beta$. Then $E = \mathcal{E}(\kappa, \beta)$.*

Proof. If this fails then we may assume that $\mathcal{E}$ is φ-minimal for the formula φ asserting that it fails. Pick a counterexample $E \in L[\mathcal{E}]$ with (κ, β) as small as possible and let j_θ and k_θ be the maps arising from the comparison of $L[\mathcal{E}]$ with the model $\mathrm{Ult}(L[\mathcal{E}], E)$. By Proposition 3.18 this gives rise to the following commutative diagram:

$$\begin{array}{ccc} L[\mathcal{E}] & \xrightarrow{\quad j_\theta \quad} & L[j_\theta(\mathcal{E})] \\ & {\scriptstyle i^E}\searrow \qquad \nearrow{\scriptstyle k_\theta} & \\ & L[i^E(\mathcal{E})] & \end{array} \tag{17.8}$$

Now all of the extenders $i^E(\mathcal{E})(\kappa, \beta')$ for $\beta' < o^{i^E(\mathcal{E})}(\kappa) = \beta$ are members of $L[\mathcal{E}]$, and by the minimality of (κ, β) it follows that $i^E(\mathcal{E})\restriction(\kappa+1) = \mathcal{E}\restriction(\kappa, \beta)$.

If $o^{\mathcal{E}}(\kappa) = \beta$ then this would imply $j_\theta(\kappa) = \kappa < i^E(\kappa)$, contradicting the commutativity of diagram (17.8). Hence the comparison starts with case 1, so that $j_{0,1} = i^{E'}$, where $E' = \mathcal{E}(\kappa, \beta)$, and $k_{0,1}$ is the identity. Furthermore, $i^{E'}(\mathcal{E})\restriction\kappa + 1 = \mathcal{E}\restriction(\kappa, \beta) = i^E(\mathcal{E})\restriction\kappa + 1$, so $\kappa_1 > \kappa$. Now suppose that $a \in [\kappa + 1 + \beta]^{<\omega}$ and $x \subseteq [\kappa]^{|a|}$. Then $x \in E_a \iff a \in i^E(x) \iff a \in k_\theta \circ i^E(x)$ and $x \in E'_a \iff a \in i^{E'}(x) = j_{0,1}(x) \iff a \in j_{1,\theta} \circ j_{0,1}(x) = j_\theta(x)$. Since $j_\theta(x) = k_\theta \circ i^E(x)$ it follows that $E = E'$, contrary to the choice of E. ⊣

3.20 Corollary. *If $\mathcal{E}$ is a weakly coherent extender sequence of countably complete extenders, then $\mathcal{E}$ is coherent in $L[\mathcal{E}]$.*

Proof. Suppose to the contrary that $\gamma < o^{\mathcal{E}}(\alpha)$ and $i^{\mathcal{E}(\alpha,\gamma)}(\mathcal{E})\restriction\gamma \neq \mathcal{E}\restriction(\alpha, \gamma)$. Let $\beta < \gamma$ be such that $\mathcal{E}(\alpha, \beta) \neq i^{\mathcal{E}(\alpha,\gamma)}(\mathcal{E})(\alpha, \beta)$, and apply Lemma 3.19 with $E = i^{\mathcal{E}(\alpha,\gamma)}(\mathcal{E})(\alpha, \beta)$. ⊣

It should be noted that the assumption that the extenders in $\mathcal{E}$ are countably complete is used only to assure that any iterated ultrapower using extenders in $\mathcal{E}$ is well-founded.

3.21 Theorem. *If κ is a strong cardinal, then there is a weakly coherent sequence $\mathcal{E}$ of countably complete extenders such that there is a strong cardinal $\kappa' \leq \kappa$ in $L[\mathcal{E}]$.*

Proof. We define the domain $o^{\mathcal{E}}$ of $\mathcal{E}$ and the extenders $\mathcal{E}(\lambda, \beta)$ using recursion on λ with an inner recursion on β. Suppose that $o^{\mathcal{E}}\restriction\lambda$ and $\mathcal{E}\restriction\lambda$ have been defined. If λ is not measurable, or if there is some $\lambda' < \lambda$ such that $o^{\mathcal{E}}(\lambda') \geq \lambda$, then set $o^{\mathcal{E}}(\lambda) = 0$. Otherwise define extenders $\mathcal{E}(\lambda, \beta)$ by recursion on β. Suppose that $\mathcal{E}\restriction(\lambda, \beta)$ has been defined. If there is a countably complete $(\lambda, \lambda + 1 + \beta)$-extender E such that $o^{i^E(\mathcal{E}\restriction\lambda)}(\lambda) = \beta$, then let $\mathcal{E}(\lambda, \beta)$ be any such extender. If there is no such extender E then the inner recursion terminates and $o^{\mathcal{E}}(\lambda)$ is defined to be β.

The sequence $\mathcal{E}$ is coherent in $L[\mathcal{E}]$ by Corollary 3.20. Now a cardinal κ' is strong in $L[\mathcal{E}]$ if and only if $o^{\mathcal{E}}(\kappa') = \mathrm{On}$. The necessity follows from the fact that if $o^{\mathcal{E}}(\kappa') \in \mathrm{On}$ then $\mathcal{E}\restriction\kappa' + 1$ is a set, but there is no extender E on κ' in $L[\mathcal{E}]$ such that $\mathcal{E}\restriction\kappa' + 1 \in \mathrm{Ult}(L[\mathcal{E}], E)$. To see that the condition $o^{\mathcal{E}}(\kappa') = \mathrm{On}$ is sufficient, let X be any set in $L[\mathcal{E}]$ and fix τ so that $X \in L_\tau[\mathcal{E}]$. Now set $E = \mathcal{E}_{\kappa,\tau}$. Then by coherence $\mathcal{E}\restriction(\kappa, \tau) = i^E(\mathcal{E})\restriction\kappa + 1$, so $X \in L_\tau[\mathcal{E}] = L_\tau[\mathcal{E}\restriction(\kappa, \tau)] = L_\tau[i^E(\mathcal{E})] \in \mathrm{Ult}(L[\mathcal{E}], E)$.

To finish the proof we need to show that there is some $\kappa' \leq \kappa$ such that $o^{\mathcal{E}}(\kappa') = \mathrm{On}$. We may suppose that $o^{\mathcal{E}}(\kappa') < \mathrm{On}$ for all $\kappa' < \kappa$. This implies that $o^{\mathcal{E}}(\kappa') < \kappa$ for all $\kappa' < \kappa$: otherwise there is, for all ordinals β, an extender F on κ so that $i(\kappa) > \beta$ and $V_\beta \subseteq \mathrm{Ult}(V, F)$. Then $o^{i^F(\mathcal{E})}(\kappa') = i^F(o^{\mathcal{E}}(\kappa')) > \beta$, but $i^F(\mathcal{E})\restriction\beta = \mathcal{E}\restriction\beta$. Since β was arbitrary, this implies that $o^{\mathcal{E}}(\kappa') = \mathrm{On}$, contrary to assumption.

Now suppose that $\mathcal{E}\restriction(\kappa, \beta)$ has been defined. We must show that there is a countably complete $(\kappa, \kappa + 1 + \beta)$-extender E such that $i^E(o^{\mathcal{E}})(\kappa) = \beta$.

Since κ is strong in V, there is a countably complete extender F on κ such that $\mathcal{E}\restriction(\kappa,\beta) \in \mathrm{Ult}(V,F)$. Now $i^F(\mathcal{E})$ is defined in $\mathrm{Ult}(V,F)$ in the same way as $\mathcal{E}$ is defined in V. Since $\mathcal{E}(\kappa,\gamma) \in \mathrm{Ult}(V,F)$ for each $\gamma < \beta$, and $\mathcal{E}(\kappa,\gamma)$ is a possible choice for $i^F(\mathcal{E})(\kappa,\gamma)$, we must have $o^{i^F(\mathcal{E})}(\kappa) \geq \beta$.

If $o^{i^F(\mathcal{E})} > \beta$ then set $E = i^F(\mathcal{E})(\kappa,\beta)$. Since V and $\mathrm{Ult}(V,F)$ have the same subsets of κ and $E\restriction\kappa = i^F(\mathcal{E})\restriction\kappa$, E is also a countably complete extender on V and satisfies $o^{i^E(\mathcal{E})} = \beta$. Otherwise, if $o^{i^F(\mathcal{E})}(\kappa) = \beta$, let $E = F|(\kappa+1+\beta)$, the $(\kappa,\kappa+1+\beta)$-extender given by the embedding $i^F : V \to \mathrm{Ult}(V,F)$. Since the identity functions serve as coherence functions for extenders, $i^E(o^{\mathcal{E}})(\kappa) = i^F(o^{\mathcal{E}})(\kappa) = \beta$, and hence E is a suitable choice for $\mathcal{E}(\kappa,\beta)$. ⊣

Theorem 3.21 can be generalized to smaller cardinals: If κ is λ-strong in V then there is a sequence $\mathcal{E}$ such that $o^{\mathcal{E}}(\kappa) > (\kappa^{+\lambda})^{L[\mathcal{E}]}$, and this holds if and only if κ is λ-strong in $L[\mathcal{E}]$.

The next result shows that, as was the case for sequences of measures, the sequence $\mathcal{E} \cap L[\mathcal{E}]$ is uniquely determined by its domain, provided that the extenders $\mathcal{E}(\kappa,\beta)$ are countably complete extenders in V, not merely in $L[\mathcal{E}]$.

3.22 Theorem. *Suppose that $\mathcal{E}$ is a weakly coherent sequence of extenders in V, $\beta < o^{\mathcal{E}}(\kappa)$, and F is a countably complete extender of length $\kappa+1+\beta$ such that $o^{i^F(\mathcal{E})}(\kappa) = \beta$. Then $F \cap L[\mathcal{E}] = \mathcal{E}(\kappa,\beta)$.*

Proof. Let $i_\theta : M_0 := L[\mathcal{E},F] \to M_\theta$ and $j_\theta : N_0 := L[\mathcal{E},F] \to N_\theta$ be iterated ultrapowers comparing the model $L[\mathcal{E},F]$ with itself, with the comparison process modified to include F as an alternative to $\mathcal{E}(\kappa,\beta)$. This means that case 3 of Definition 3.13 is modified to allow $M_{\nu+1}$ to be either of $\mathrm{Ult}(M_\nu, i_\nu(\mathcal{E}(\kappa,\beta)))$ or $\mathrm{Ult}(M_\nu, i_\nu(F))$ if the ultrafilter in question differs on a set in $M_\nu \cap N_\nu$ either from $j_\nu(\mathcal{E})(i_\nu(\kappa), i_\nu(\beta))$ or (in the case $i_\nu(\kappa) = j_\nu(\kappa)$ and $i_\nu(\beta) = j_\nu(\beta)$) from $j_\nu(F)$. Similarly, $j_\nu(F)$ is a candidate for use in defining $N_{\nu+1}$.

Lemma 3.15, asserting that the comparison terminates, is still valid for this comparison. Consider the final models $M_\theta = L[i_\theta(\mathcal{E}), i_\theta(F)]$ and $N_\theta = L[j_\theta(\mathcal{E}), j_\theta(F)]$ of this comparison. One of the sequences $i_\theta(\mathcal{E})$ and $j_\theta(\mathcal{E})$ will be an initial segment (possibly proper) of the other; suppose that $i_\theta(\mathcal{E})$ is an initial segment of $j_\theta(\mathcal{E})$. Then we have that $i_\theta(\mathcal{E}(\kappa,\beta))$ and $i_\theta(F)$ agree with $j_\theta(\mathcal{F})(i_\theta(\kappa), i_\theta(\beta))$, and hence with each other, on all sets in $M_\theta \cap N_\theta$. By the elementarity of i_θ there is a set $x \in M_\theta$ on which $i_\theta(F)$ and $i_\theta(\mathcal{E}(\kappa,\beta))$ differ. Let x be the first such set in the order of construction of M_θ, and suppose that $x \in L_{\tau+1}[i_\theta(\mathcal{E}), i_\theta(F)] - L_\tau[i_\theta(\mathcal{E}), i_\theta(F)]$. Then $L_\tau[i_\theta(\mathcal{E}), i_\theta(F)] = L_\tau[i_\theta(\mathcal{E})] = L_\theta[j_\theta(\mathcal{E})]$ so, as in the proof of Lemma 3.19, $x \in L_{\tau+1}[i_\theta(\mathcal{E})] \subseteq N_\theta$. Thus $x \in M_\theta \cap N_\theta$, contradicting the assumption that $i_\theta(F)$ and $i_\theta(\mathcal{E}(\kappa,\beta))$ differ about x. ⊣

3.23 Corollary. *If $\mathcal{E}$ and $\mathcal{E}'$ are weakly coherent sequences of extenders in V with the same domain then $L[\mathcal{E}] = L[\mathcal{E}']$ and $\mathcal{E} \cap L[\mathcal{E}] = \mathcal{E}' \cap L[\mathcal{E}']$.*

It was previously observed that this statement is false, even for ultrapowers of order 0, if the requirement that $\mathcal{E}$ be a sequence of extenders in V is weakened to require only that they be extenders in $L[\mathcal{E}]$.

We conclude this section by showing that the Generalized Continuum Hypothesis holds in $L[\mathcal{E}]$. The same argument shows that other consequences of condensation such as $\diamondsuit_\kappa$ and $\square_\kappa$ also hold in $L[\mathcal{E}]$.

3.24 Theorem. *If $\mathcal{E}$ is a coherent sequence of countably complete extenders in $L[\mathcal{E}]$ then $L[\mathcal{E}] \models \mathrm{GCH}$.*

Proof. The proof of Theorem 3.24 will require the proof of a condensation lemma for $L[\mathcal{E}]$. Let us say that a model M is a *coarse mouse* in $L[\mathcal{E}]$ with projectum ρ if $\pi : M \cong X \prec L_\tau[\mathcal{E}]$ where $L_\tau[\mathcal{E}] \models \mathrm{ZF}^-$ and $X = \mathcal{H}^{L_\tau[\mathcal{E}]}(\{\mathcal{E}\} \cup \rho \cup p)$ for some finite set $p \in L_\tau[\mathcal{E}]$ of parameters. As in the proof of GCH for L, every subset of ρ in $L[\mathcal{E}]$ is in some coarse mouse with projectum ρ, and each such mouse has cardinality $|\rho|$. Hence it will be enough to show that if M and N are coarse mice in $L[\mathcal{E}]$ with the same projectum ρ, then either $\mathcal{P}(\rho) \cap M \subseteq N$ or $\mathcal{P}(\rho) \cap N \subseteq M$.

First, suppose that $o^{\mathcal{E}}(\kappa) \le \rho$ for all $\kappa < \rho$ and let $i_\theta : M \to P$ and $j_\theta : N \to Q$ be the maps arising from the comparison of $M = L_{\tau_0}[\mathcal{F}_0]$ and $N = L_{\tau_1}[\mathcal{F}_1]$. Then $\mathcal{F}_0 \restriction \rho = \mathcal{F}_1 \restriction \rho = \mathcal{E} \restriction \rho$ and hence both $i_\theta \restriction \rho$ and $j_\theta \restriction \rho$ are the identity. Therefore $\mathcal{P}^M(\rho) = \mathcal{P}^P(\rho)$ and $\mathcal{P}^N(\rho) = \mathcal{P}^Q(\rho)$; and since one of P and Q is contained in the other it follows that one of $\mathcal{P}^M(\rho)$ and $\mathcal{P}^N(\rho)$ is contained in the other, as was to be proved.

In particular, the assumption that there is no overlapping in the sequence $\mathcal{E}$ implies that $2^\kappa = \kappa^+$ in $L[\mathcal{E}]$ for any κ such that $o^{\mathcal{E}}(\kappa) > 0$.

Now suppose that there is $\kappa < \rho$ with $o^{\mathcal{E}}(\kappa) > \rho$, and let $M = L_\alpha[\mathcal{F}]$ be any coarse ρ-mouse in $L[\mathcal{E}]$. If we set $\beta = o^{\mathcal{F}}(\kappa)$, then because $L[\mathcal{E}]$ satisfies GCH at κ we have $\mathcal{P}^{L[\mathcal{E}]}(\kappa) \subseteq M$ and hence the extenders $\mathcal{F}(\kappa, \gamma)$ for $\gamma < \beta$ are all extenders in $L[\mathcal{E}]$. It follows by Lemma 3.19 that $\mathcal{F}(\kappa, \gamma) = \mathcal{E}(\kappa, \gamma)$ for all $\gamma < \beta$, and hence $\mathcal{F} \restriction \kappa + 1 = \mathcal{E} \restriction (\kappa, \beta)$.

Now if M and N are two coarse ρ-mice in $L[\mathcal{E}]$ with $\beta^M = \beta^N$, then the same argument as that used for the case when $o^{\mathcal{E}}(\kappa) \le \rho$ for all $\kappa < \rho$ implies that one of $\mathcal{P}^M(\rho)$ and $\mathcal{P}^N(\rho)$ is a subset of the other. Thus, if we hold β fixed then there are at most ρ^+ many subsets of ρ which are in some coarse mouse M for $L[\mathcal{E}]$ with projectum ρ and which have $\beta^M = \beta$. Now $\beta^M < \rho^+$ in $L[\mathcal{E}]$ for any such coarse mouse with projectum ρ, so there can be at most ρ^+-many subsets of ρ in $L[\mathcal{E}]$. ⊣

The natural well-ordering of $\mathcal{P}^{L[\mathcal{E}]}(\rho)$ suggested by this proof is given by setting $x \prec y$ if there is $\beta < o^{\mathcal{E}}(\kappa)$ such that x, but not y, is a member of $\mathrm{Ult}(L[\mathcal{E}], \mathcal{E}(\kappa, \beta))$; and otherwise setting $x \prec y$ if x is less than y in the order of construction either of $\mathrm{Ult}(L([\mathcal{E}], \mathcal{E}(\kappa, \beta))$ where β is least such that $x, y \in \mathrm{Ult}(L([\mathcal{E}], \mathcal{E}(\kappa, \beta))$, or of $L[\mathcal{E}]$ if there is no such β. Note that $i^{\mathcal{E}(\kappa,\beta)}(\rho) < \rho^+$ for any $\beta < \rho^+$, and hence this well-ordering has ordertype ρ^+.

3.1. The Modern Presentation of $L[\mathcal{E}]$

Almost all of the description of $L[\mathcal{U}]$ and $L[\mathcal{E}]$ given so far has followed the original style of [25]; the only exception being the brief description at the end of Sect. 1.2 of the application of the modern presentation to $L[U]$ and K^{DJ}. This presentation was invented in order to accommodate larger cardinals than those considered here, but it has several advantages even for models with smaller cardinals, especially when core model and fine structural techniques are being used.

We will now outline some aspects of this new presentation. There are three major changes.

(1) As was pointed out previously, the method of indexing used in the models of this chapter breaks down beyond a strong cardinal. Instead we index extenders in the sequence with a single ordinal. In the original indexing of these models, the index γ for a extender $E = \mathcal{E}_\gamma$ on the sequence is given by $\gamma = (\nu^+)^{L[\mathcal{E}\restriction\gamma]}$ where ν is the larger of κ^+ and the length of the extender E. This choice of ν ensures that E can easily be coded as a subset of ν.

As part of this indexing, the class coding the sequence $\vec{\mathcal{E}}$ is chosen so that $L_\gamma[\mathcal{E}] = L_\gamma[\mathcal{E}\restriction\gamma]$, while $L_{\gamma+1}[\mathcal{E}]$ is the collection of subsets of $L_\gamma[\mathcal{E}]$ which are definable in the structure $(L_\gamma[\mathcal{E}], \mathcal{E}\restriction\gamma, \mathcal{E}_\gamma)$.

This indexing is still commonly used, but Jensen and others have also worked with indexing schemes using indices as large as $i^E(\kappa^+)$.

(2) More importantly, an extender $E = \mathcal{E}_\gamma$ of the sequence $\mathcal{E}$ does not measure all of the sets in $L[\mathcal{E}]$, but instead only measures the sets in $L_\gamma[\mathcal{E}\restriction\gamma]$, that is, the sets already constructed at the time E appears. This is in contrast to the models of this chapter, in which an extender E is expected to measure sets in $L[\mathcal{E}]$ which require E, and even larger extenders, for their construction. Note that if $\kappa = \mathrm{crit}(\mathcal{E}_\gamma)$ then the choice of $\gamma = (\nu^+)^{L_\gamma[\mathcal{E}\restriction\gamma]}$ implies that $\mathcal{P}(\nu)\cap L[\mathcal{E}\restriction\gamma] \subseteq L_\gamma[\mathcal{E}\restriction\gamma] = L_\gamma[\mathcal{E}]$. Thus $\mathcal{E}_\gamma \subseteq L_\gamma[\mathcal{E}]$, and hence $\mathcal{E}_\gamma$ is a member of $L_{\gamma+1}[\mathcal{E}]$.

An extender $\mathcal{E}_\gamma$ with critical point κ_γ will be a full extender in the final model $L[\mathcal{E}]$ if and only if no new subsets of κ_γ are constructed in $L[\mathcal{E}] - L_\gamma[\mathcal{E}]$. The other extenders, those extenders $\mathcal{E}_\gamma$ for which $\mathcal{P}(\kappa_\gamma) \cap L[\mathcal{E}] \not\subseteq L_\gamma[\mathcal{E}]$, are only partial extenders in $L[\mathcal{E}]$; however (as in the discussion of K^{DJ} at the end of Sect. 1.2) they serve as full extenders inside the mice by which these new subsets of κ_γ are constructed. In fact these mice turn out to be exactly the initial segments $L_\alpha[\mathcal{E}] = L_\alpha[\mathcal{E}\restriction\alpha]$ of the model $L[\mathcal{E}]$.

(3) This use of the partial extenders in mice requires the definition and use of a fine structure which is essentially identical to Jensen's fine structure for L. Fine structure is beyond the purview of this chapter, but one important consequence has already been mentioned in connection with K^{DJ}: whenever $\rho < \alpha$ and there is a set $x \in \mathcal{P}(\rho) \cap L_{\alpha+1}[\mathcal{E}] - L_\alpha[\mathcal{E}]$, then $L_{\alpha+1}[\mathcal{E}] \models |\alpha| \leq \rho$.

This discussion ignores one further difference: the model $L[\mathcal{E}]$ (like all recent fine structural arguments) is defined using Jensen's rudimentary hierarchy $J_\alpha[\mathcal{E}]$ instead of the hierarchy $L_\alpha[\mathcal{E}]$ used in this chapter. This change yields a substantial technical simplification, but makes no conceptual difference.

The main disadvantage of the newer approach is evident. The use of fine structure makes the newer models $L[\mathcal{E}]$ more complex than the models $L[\mathcal{U}]$, and furthermore, the extra complexity cannot be delayed, since the model $L[\mathcal{E}]$ cannot even be defined without it.[2] Thus one would want to have a good understanding of the simpler models described here, as well as of fine structure in the simpler setting of L, before studying the newer extender models.

We list below some of the advantages which justify the extra complexity. It should be noted that for larger cardinals there is no choice: the inner models require the newer style—which was in fact invented in order to make inner models for these cardinals possible. However it turns out that arguments using the newer $L[\mathcal{E}]$ style models are simpler, even though the older style $L[\mathcal{U}]$ could have been used instead. The discussion below indicates some of the reasons for this.

(1) A much stronger condensation property holds for the new fine structural models than for those discussed in this chapter. This point was briefly touched on during the discussion of the model $L[U]$ in Sect. 1.2.

(2) The coherence property is simpler and more robust in the fine structural models. We have already seen this as an advantage of using extenders instead of ultrafilters, and this sometimes gives reason to use extender models even when all extenders used turn out to be equivalent to ultrafilters. This advantage is strengthened in the fine structural models, in which all relevant functions have already been constructed before the extender is added.

(3) The use of partial extenders helps to simplify and strengthen the comparison process. Suppose that the two sequences $\mathcal{E}$ and $\mathcal{E}'$ being compared differ first at an ordinal γ, so that $\mathcal{E}\restriction\gamma = \mathcal{E}'\restriction\gamma$ but $\mathcal{E}_\gamma \neq \mathcal{E}'_\gamma$. Then $\mathcal{E}_\gamma$ measures only the sets in $L_\gamma[\mathcal{E}] = L_\gamma[\mathcal{E}\restriction\gamma] = L_\gamma[\mathcal{E}'\restriction\gamma] = L_\gamma[\mathcal{E}']$, which contains the sets measured by $\mathcal{E}'_\gamma$. Hence there is no need for the maneuver used in Definition 3.13, in which two extenders are deemed to differ for the purposes of defining the iterated ultrapower only if they differ on a set in the intersection $L[\mathcal{E}] \cap L[\mathcal{E}']$ of the two models: If $\mathcal{E}_\gamma$ and $\mathcal{E}_{\gamma'}$ differ at all, then they disagree on a member of their common domain $L_\gamma[\mathcal{E}] = L_\gamma[\mathcal{E}']$.

(4) The development of the core model is greatly simplified in fine structural models, because there is no need to treat mice and ultrafilters separately. Under the old approach, the core model was a structure of the form $K = L[\mathcal{U}, \mathcal{M}]$ where $\mathcal{U}$ is a coherent sequence of measures and $\mathcal{M}$ is the

[2] There do exist inner models for larger cardinals which do not use fine structure. These include the original Martin-Steel model [22] for a Woodin cardinal, the HOD models having Woodin cardinals which Woodin obtained from determinacy hypotheses, and Woodin's recent models for cardinals beyond a supercompact. However all of these models fall badly short of being the L-like models we are looking for: for example, it is still not known whether the Martin-Steel models satisfy GCH.

class of mice over $\mathcal{U}$. In fine structural models the core model K has the form $L[\mathcal{E}]$, and the mice used to construct the model are simply the initial segments $L_\gamma[\mathcal{E}]$ of $L[\mathcal{E}]$, with some of the partial measures of $\mathcal{E}$ being used as full measures in the mouse $L_\gamma[\mathcal{E}]$.

This point becomes more important for core models for larger cardinals. In order for the construction to work properly, the mice must reflect the properties of the full core model, and in particular they must be allowed to recursively contain smaller mice. This seems almost prohibitively complicated when working with a core model in the form $K = L[\mathcal{U}, \mathcal{M}]$, with the measures and the mice treated separately, but it falls out naturally in the fine structural model $K = L[\mathcal{E}]$ where a mouse $M = L_\gamma[\mathcal{E}]$ will contain as smaller mice all its initial segments $L_{\gamma'}[\mathcal{E}]$ for $\gamma' < \gamma$.

(5) The fine structural models come much closer to satisfying the analog of Theorem 1.9 than do the models described in this chapter. To see why this is so, consider an argument like that given for Lemma 3.22, where $\mathcal{E}\restriction\gamma$ has been defined and $E = \mathcal{E}_\gamma$ and F is a second extender which could have been chosen as $\mathcal{E}_\gamma$. In the fine structural model both of these extenders measure the same collection of sets, namely the members of the structure $L_\gamma[\mathcal{E}\restriction\gamma]$. Thus instead of using an iterated ultrapower of the structure $L[\mathcal{E}, F]$, in which both extenders are used in the construction, one can use the *bicephelus* $(L_\gamma, \mathcal{E}\restriction\gamma, E, F)$, in which both extenders are available as predicates but neither is used in the construction. The only extra hypothesis on E and F which is needed, beyond the requirement that each extender individually satisfies the conditions to be $\mathcal{E}_\gamma$, is that they are jointly iterable in the sense that all ultrapowers of this structure are well-founded.

One further point should be noted: the principal disadvantage of the fine structural approach, the need to introduce the extra complexity of fine structure at the very beginning, is not an issue in the development of the core model because the fine structure will be required in any case. Indeed incorporating fine structure into the initial definition of $L[\mathcal{E}]$ allows for a much more natural presentation and development of the core model and its fine structure.

4. Remarks on Larger Cardinals

In this section we briefly list some of the most important large cardinals above measurability, in increasing order of size. The primary focus is on the inner model theory available for these large cardinal properties; more information on some of these inner models can be found in later chapters in this Handbook.

All of the large cardinal properties described here are defined by elementary embeddings. Throughout this section, i is always an elementary embedding and M is a well-founded class.

Strong cardinals

Strong cardinals, together with their inner models, have already been introduced and an inner model has been described. It was also pointed out that such simple models, with comparison defined by linear iterations, are inadequate to handle very much larger cardinals. The line beyond which iteration trees are needed is not sharp. Baldwin [1] uses modified linear iterations to handle cardinals substantially larger than strong cardinals, and Schindler [35] has used nearly linear iterations to define a fine structural core model up to the sharp of a proper class of strong cardinals. In the other direction, a careful analysis shows that fine structural models actually use a simple form of iteration tree even down at the level of a 2-strong cardinal, that is, one with an extender E on κ such that $\mathcal{P}^2(\kappa) \subseteq \mathrm{Ult}(V, E)$.

Because of the need for iteration trees rather than linear iterations, it is much more difficult to obtain iterable models for larger cardinals in this range. Indeed, it is not known[3] whether a core model larger than those constructed by Schindler in [35] can be constructed without an added assumption of some large cardinal strength in the universe. Chapter [33] covers the core model and the covering lemma up to a Woodin cardinal.

Woodin cardinals

A cardinal δ is said to be *Woodin* if for all functions $f : \delta \to \delta$ there is an embedding $i : V \to M$ with critical point $\kappa < \delta$ such that $f``\kappa \subseteq \kappa$ and $V_{i(f)(\kappa)} \subseteq M$.

Woodin cardinals were defined by Woodin in 1984, following work of Foreman, Magidor and Shelah [8], and are the most important large cardinal property for current research in set theory. The most notable result concerning Woodin cardinals is probably the equiconsistency of the axiom of determinacy with the existence of infinitely many Woodin cardinals, due to Woodin, Martin and Steel, which is discussed in chapters [29] and [17].

This, and other consequences of Woodin cardinals, depend largely on two forcing notions which can be used to prove that inner models for Woodin and stronger cardinals must differ in important respects from those for smaller cardinals. The first of these forcing notions is *stationary tower forcing*, which was defined by Woodin using ideas from Foreman, Magidor and Shelah [9, 8]. In one form, this forcing will preserve a Woodin cardinal δ, while making massive changes to the cardinal structure below δ: for example, there is a stationary subset of singular cardinals below δ whose successors are collapsed by the forcing. Hence there cannot be a core model satisfying the weak covering property for (exactly) a Woodin cardinal, although there is one for the sharp of a Woodin cardinal. In addition, stationary tower forcing

[3] It is now known, by a recent unpublished result of Jensen and Steel, that if there is no model with a Woodin cardinal then the core model K can be constructed with no extra large cardinal hypothesis.

can collapse ω_1, and in the process will add new countable mice. This forcing is discussed in the book [20].

The second forcing notion, invented by Woodin, is the remarkable "all sets are generic" forcing: If M is a model with a Woodin cardinal δ, and M is iterable in V, then there is a forcing notion $P \in M$ of size δ such that for any set x in V there is a tree iteration of M, with final model N and embedding $i : M \to N$, such that $N[x]$ is a generic extension of N by the forcing $i(P)$. This forcing can be used to show that the minimal model M for a Woodin cardinal cannot satisfy the sentence asserting that M is iterable, even when M is iterable in the universe V. The implications of this for the core model are discussed further at the end of this chapter.

At present it is not known how to construct core models for cardinals in this range without some large cardinal properties holding in the universe. Jensen has shown that a subtle cardinal, a property weak enough to hold in L, is enough to show prove that the core model exists and satisfies the covering lemma; however it is an open question whether this assumption is needed. Other than this gap, the core model theory through ω many Woodin cardinals is well understood [39]. The strongest current result on existence of iterable inner models is due to Neeman, who has constructed [30] iterable extender models with a Woodin limit of Woodin cardinals. These models, however, are not fine structural, and no core model results are known in this region.

Superstrong Cardinals

A cardinal κ is *superstrong* if there is an embedding $i : V \to M$ with critical point κ such that $V_{i(\kappa)} \subseteq M$.

As was pointed out previously, a superstrong cardinal is at the outer limits of our understanding of inner models. Much of the basic inner model theory is understood up to a superstrong cardinal: for example it is known [34] that $\square_\kappa$ holds in any extender model up through a superstrong cardinal. Indeed they show that $\square_\kappa$ holds in an extender model $L[\mathcal{E}]$ for any cardinal κ short of what Jensen has labeled a *subcompact cardinal*. Jensen has shown that $\square_\kappa$ cannot hold if κ is subcompact. However it is not known, under any large cardinal assumption, that there are any iterable extender models with anything near a superstrong cardinal.

Supercompact Cardinals

A cardinal κ is λ-*supercompact* if there is an embedding $i : V \to M$ with critical point κ such that ${}^\lambda M \subseteq M$, and κ is *supercompact* if κ is λ-supercompact for all cardinals λ.

None of the models described in this chapter give any promise of yielding models with a supercompact cardinal. However Woodin has recently proposed a form of model, using what he calls *suitable extender sequences* which can include supercompact cardinals and even the larger cardinals discussed in the next paragraph, and which he hopes to show have many of the properties enjoyed by the extender models $L[\mathcal{E}]$ which have been discussed in this chapter.

Like these models, Woodin's models have the form $L[\mathcal{E}]$, the class of sets constructible from a sequence of extenders. An important difference is that not all of the extenders witnessing large cardinal properties are members of the sequence $\mathcal{E}$; in fact all of the critical points of extenders on the sequence are below the first supercompact cardinal. It is still not known whether these models have an analog of the comparison process of Lemma 2.8, and no proofs are known for their iterability.

Larger Cardinals

A number of cardinals larger than supercompact have been defined. Some of these have important consequences, notably huge cardinals and variants of these. A cardinal κ is *huge* if there is an elementary embedding $i : V \to M$ with critical point κ such that ${}^{i(\kappa)}M \subseteq M$.

Catalogs of large cardinal properties, such as this one, traditionally end with a nontrivial elementary embedding from V into V, which Kunen proved in [19] to be inconsistent. It is still open whether such an embedding is consistent with ZF without the axiom of choice.

5. What is the Core Model?

This section is not intended to be a description of existing core models, but rather an examination of the term "core model" itself. We will try to determine the meaning of the phrase "the core model", and in particular explain the difference between it and the term "extender model". The structure, construction and properties of known core models is described elsewhere in this chapter and in chapters [24, 33, 38] and [36]. In addition the reader may want to look at [27], which discusses from a relatively non-technical point of view the use of iteration trees and the construction of the Steel core model up to a Woodin cardinal.

Our first approach will be to look at the history of the term "core model", which was introduced by Dodd and Jensen [5, 6] for the model which is variously referred to as the Dodd-Jensen core model, K^{DJ}, or the core model below a measurable cardinal. The history, however, begins earlier—at least as far back as Jensen's discovery of the covering lemma for L, since the Dodd-Jensen core model generalizes this result. The model $L[U]$ also predated K^{DJ}, and although $L[U]$ is not contained in the structure K^{DJ} which Dodd and Jensen referred to as the core model, they proved [7] the covering lemma for $L[U]$ and hence brought this model into the modern pantheon of core models. Their work was extended by Mitchell to include sequences of measures. The core model to this point is described in chapter [24]. The use of extenders as a generalization of normal ultrafilters, and of iteration trees as a generalization of iterated ultrapowers, led to the Steel core model, which is described in chapters [33, 38]. This model is currently at the frontier of the subject.

Of the two terms under consideration, only "extender model" has a precise meaning: an extender model is a model of the form $L[\mathcal{E}]$ where $\mathcal{E}$ is a good sequence of extenders as defined in chapter [38]. Every known core model is an extender model, but this should not be assumed to be true for larger cardinals; indeed it seems unwise to be dogmatic about the properties of as yet unknown core models until we have a better idea of what is possible.

Even keeping this caveat in mind, "the core model" is always singular: there is at most one core model in any given model of set theory, and in particular there is at most one true core model in the true universe of sets.

Some authors have used the term "core model" to mean the same as "extender model". While it is true that every known core model is an extender model, and that generally, or arguably always, an extender model is its own core model, the distinction between the terms is important and should be preserved. The term *extender model* describes the interior structure of the model, while the term *core model* refers to the relation between the model and the class of all sets.

Some illumination on this point can be gained by looking at cases in which we find it useful, in apparent contradiction to the dictum in the last paragraph, to speak of "*a* core model". It is often useful to refer to a model as "a core model" if it is the core model as defined inside some model which is of particular interest, but is not necessarily the universe of all sets. In a related usage, the term "core model" is often used for a model obtained by a particular construction which is known to yield the core model under additional assumptions such as the nonexistence of some large cardinal property. The Dodd-Jensen core model K^{DJ} is an example of both usages: It is characterized by its mode of construction, which is an initial segment of the core model construction in every model for which such a construction is known. It is also characterized by the fact that it (or at least $K^{\mathrm{DJ}} \cap M$) is *the* core model inside *any* model M, so long as M does not have an inner model with a measurable cardinal. Core models for larger cardinals are less clear cut, since the core model for a model M varies with the particular extenders which are members of M, even though the large cardinal strength of the model is held fixed. There is a unique core model, however, for the sharps of such large cardinal properties.

The second approach to understanding the term "core model" is through consideration of the properties of the known core models. These properties fall into two classes. The properties in the first class are those which hold in any extender model: These models are built up from below, in a manner analogous to the construction of L, and as a consequence they satisfy some sort of condensation. They satisfy combinatorial principles such as $\Diamond_\kappa$ and (for cardinals small enough that we currently have a core model) $\Box_\kappa$. They satisfy the generalized continuum hypothesis, they satisfy the global axiom of choice and their well-ordering, both of their reals, and of their full universe, has a logical form which is as simple as possible in any model with the same large cardinal properties.

The other class of properties of the core models are those which might be seen as asserting that the model is close to V. The most important of these is the covering lemma, or at the least some form of the weak covering lemma. A second is rigidity: there is no nontrivial elementary embedding $i : K \to K$. A third is absoluteness: the core model is absolute for a class of sentences which falls just short of including the sentence asserting that there is a set not in that core model.

It is unclear to what extent we should assume that these properties will necessarily hold for larger core models. Even down at the level of a Woodin cardinal, without the sharp of a Woodin cardinal, there is no inner model which satisfies both weak covering and invariance under forcing; and properties which seem close to rigidity fail well below a Woodin cardinal.

A final property bridges these two classes: The core models are uniquely defined by a formula which is absolute under set generic extensions. This formula says on the one hand that the model is built up from below as an extender model $L[\mathcal{E}]$, and on the other that the construction is greedy, including everything appropriate into the sequence $\mathcal{E}$. If we take the first class of properties as evidence of minimality then we could take something like the following as the definition of the core model:

5.1 Definition. The *core model* is the minimal class inner model of ZF which contains all of the large cardinal structure which exists in the universe.

We could modify the statement by requiring ZFC rather than ZF, but it seems better to regard the axiom of choice as a consequence (so far, at least) of minimality.

Although it is labeled a "definition", Definition 5.1 is not intended to be a precise mathematical definition. Neither "minimal" nor "large cardinal structure" have a precise meaning. The phrase "minimal class inner model of ZF" is, perhaps, reasonably clear. We can take "minimal" to mean $\subseteq$-minimal, which works for all known core models—provided a suitable meaning for the term "large cardinal structure" is understood.

The meaning of this term is somewhat more problematic. One important point is that "large cardinal structure" is not the same as "large cardinal properties". The model $L[U]$ is not $\subseteq$-minimal among all models having a measurable cardinal; for example $\mathrm{Ult}(L[U], U)$ is a proper subclass of $L[U]$. However $L[U]$ is the minimal model containing the filter $U \cap L[U]$, and it seems quite clear that the ultrafilter U should be included as part of the large cardinal structure. There are more doubtful cases in which Definition 5.1 may be at least potentially circular: once a particular model K has been anointed as "the core model" there will be a tendency to take the "large cardinal structure" of the universe to be just that structure which is contained in K.

As a case study to illustrate how the line might be drawn, we consider the situation when there is a Woodin cardinal δ, but no sharp for a model with a Woodin cardinal. There is an obvious candidate for the core model in this case, namely the extender model $L[\mathcal{E}]$ given by Steel's core model construction described in chapter [33]. It might be objected that this model is not really obtained by Steel's construction of the core model, but rather as a limiting case of that construction: Steel's construction gives a sequence of models $K_\theta = L_\theta[\mathcal{E}_\theta]$ for measurable cardinals $\theta < \delta$. Each of the models K_θ is unequivocally the core model in V_θ, and the extender sequences $\mathcal{E}_\theta$ of the models K_θ agree so as to yield a combined sequence $\mathcal{E} = \bigcup_\theta \mathcal{E}_\theta$ such that δ is Woodin in $L[\mathcal{E}]$. This objection is a reason for caution, but is irrelevant to the application of Definition 5.1, which deliberately avoids specifying a particular means of construction. A second objection to the model $L[\mathcal{E}]$ is that it is not iterable: Woodin's "all sets are generic" forcing demonstrates that there is an iteration tree of height δ which can be defined in $L[\mathcal{E}]$, but which has no well-founded branch in $L[\mathcal{E}]$. Again this is a reason for caution but is not necessarily fatal: the iterability of the model might well be considered as large cardinal structure, but it is large cardinal structure which does not exist in the universe and thus cannot be expected to exist in the core model. In fact, for example, the existence of a model $L[\mathcal{E}]$ with a Woodin cardinal such that every iteration tree in $L[\mathcal{E}]$ has a well-founded branch in V implies the existence of a class of indiscernibles for $L[E]$.

A more significant question is raised by Woodin's stationary tower forcing, which massively violates the weak covering lemma. The cardinal δ is still Woodin in the generic extension, but it is possible to arrange (and is possibly impossible to avoid) that every sufficiently large successor cardinal below δ is collapsed. This probably should not bother us: we can consider this to be analogous to Prikry forcing at a measurable cardinal, which shows that if there is a measurable cardinal then no core model will satisfy the covering lemma in all generic extensions. It is true that the situation at a measurable cardinal is well understood while that at a Woodin cardinal is quite hazy, but the analogy seems reasonable.

It has been argued that it is not really necessary to give up the weak covering lemma because there is a second candidate for the core model. If we assume that the ground model is $L[\mathcal{E}]$, then an $L[\mathcal{E}]$-generic set G for the stationary tower forcing is essentially an extender which gives an elementary embedding $i^G : L[\mathcal{E}] \to L[i^G(\mathcal{E})]$ with the property that $V_\delta \cap L[\mathcal{E}][G] \subseteq L[i^G(\mathcal{E})]$. In particular, $L[i^G(\mathcal{E})]$ does satisfy the covering lemma in the generic extension $L[\mathcal{E}][G]$, and furthermore, $L[i^G(\mathcal{E})]$ is the model obtained as described above using Steel's construction inside $L[\mathcal{E}][G]$. We could take $L[i^G(\mathcal{E})]$ as the core model, provided that we are willing to give up invariance under forcing. In favor of $L[\mathcal{E}]$, we could assert that $i : L[\mathcal{E}] \to L[i^G(\mathcal{E})]$ should be regarded as analogous to $\mathrm{Ult}(L[U], U) = L[i^U(U)]$, and note that $L[i^U(U)]$ is certainly not the core model. This view is supported by Woodin's [41] extensive and fruitful theory of iterated ultrapowers using generic embeddings such as i^G, but it is weakened by the fact that it throws no light on the failure of the weak covering lemma.

The model $L[\mathcal{E}]$ is certainly the core model according to Definition 5.1, at least inside the ground model $L[\mathcal{E}]$ itself. The question is whether the mice in $L[i^G(\mathcal{E})] - L[\mathcal{E}]$ should be included as part of the large cardinal structure of $L[\mathcal{E}][G]$. For an answer to this question we consider another analogy with $L[U]$: Jensen has proved (see Theorem 3.43 in chapter [24]) that if H is a $L[U]$ generic Levy collapse, then in $L[U][H]$ there is an embedding $i : K^{\mathrm{DJ}} \to K^{\mathrm{DJ}}$ such that $\mathrm{crit}(i)$ is smaller than $\mathrm{crit}(U)$. The embedding is constructed from a model $N = L_\alpha[U^N]$ in $L[U][H] - L[U]$ which is iterable and satisfies ZF^-. In fact N is a mouse; however it is not a mouse in the sense of K^{DJ} because there is no subset of $\mathrm{crit}(U^N)$ in $L_{\alpha+1}[U^N] - L_\alpha[U^N]$. Now U^N is not an ultrafilter in $L[U^N]$, so let $\alpha' > \alpha$ be the least ordinal such that there is a subset of $\mathrm{crit}(U^N)$ in $L_{\alpha'+1}[U^N] - L_{\alpha'}[U^N]$. If $N' = L_{\alpha'}[U^N]$ were iterable then it would be a member of K^{DJ}, and that is not true because U^N measures all sets in K^{DJ} while there is a set in $L_{\alpha'+1}[U^N]$ which U^N does not measure. Thus N' is not iterable; in fact the set in $L_{\alpha'+1}[U^N]$ which is not measured by U^N can be constructed from a sequence of functions in N' which witnesses that $\mathrm{Ult}(N', U^N)$ is not well-founded.

The extra information given by the ordinal $\alpha' > \alpha$ shows that $L_\alpha[U^N]$ is, in an extended sense, not really iterable. Similarly, the information given by the extender sequence $\mathcal{E}$ shows that the supposed mice M which are in $L[i^G(\mathcal{E})]$ but not in $L[\mathcal{E}]$ are not really iterable: if we attempt to compare M with $L[\mathcal{E}]$ then the tree on M has height δ and has no well-founded cofinal branch, as any such branch could be used to construct the sharp for a Woodin cardinal. Thus it seems appropriate to conclude that M is not part of the large cardinal structure of $L[\mathcal{E}][G]$, and hence that $L[\mathcal{E}]$ is the core model in $L[\mathcal{E}][G]$.

Why then does Steel's construction seem to go wrong here? As was suggested earlier, it is not the construction which is in error: If θ is a measurable cardinal below δ then every mouse in the model $K_\theta = L_\theta[i^G(\mathcal{E})\restriction\theta]$ is iterable in $V_\theta^{L[\mathcal{E}][G]}$, and hence K_θ really is the core model in the universe $V_\theta^{L[\mathcal{E}][G]}$. The only error is in assuming that the limit of these local core models will be a core model in $V^{L[\mathcal{E}][G]}$: it is not, because its "mice" are not iterable there.

Bibliography

[1] Stewart Baldwin. Between strong and superstrong. The Journal of Symbolic Logic, 51(3):547–559, 1986.

[2] Aaron Beller, Ronald B. Jensen, and Philip D. Welch. Coding the Universe, volume 47 of London Mathematical Society Lecture Note Series. Cambridge University Press, Cambridge, 1982.

[3] Keith J. Devlin. Constructibility. Springer, Berlin, 1984.

[4] Anthony J. Dodd. Core models. The Journal of Symbolic Logic, 48(1):78–90, 1983.

[5] Anthony J. Dodd and Ronald B. Jensen. The core model. Annals of Mathematical Logic, 20(1):43–75, 1981.

[6] Anthony J. Dodd and Ronald B. Jensen. The covering lemma for K. Annals of Mathematical Logic, 22(1):1–30, 1982.

[7] Anthony J. Dodd and Ronald B. Jensen. The covering lemma for $L[U]$. Annals of Mathematical Logic, 22(2):127–135, 1982.

[8] Matthew Foreman, Menachem Magidor, and Saharon Shelah. Martin's maximum, saturated ideals and nonregular ultrafilters. II. Annals of Mathematics (2), 127(3):521–545, 1988.

[9] Matthew Foreman, Menachem Magidor, and Saharon Shelah. Martin's maximum, saturated ideals, and nonregular ultrafilters. I. Annals of Mathematics (2), 127(1):1–47, 1988.

[10] Sy D. Friedman. Constructibility and class forcing. Chapter 8 in this Handbook. 10.1007/978-1-4020-5764-9_9.

[11] Sy D. Friedman. Provable Π^1_2-singletons. Proceedings of the American Mathematical Society, 123(9):2873–2874, 1995.

[12] Moti Gitik. Prikry-type forcings. Chapter 16 in this Handbook. 10.1007/978-1-4020-5764-9_17.

[13] Moti Gitik. On closed unbounded sets consisting of former regulars. The Journal of Symbolic Logic, 64(1):1–12, 1999.

[14] Thomas J. Jech. Set Theory. Springer Monographs in Mathematics. Springer, Berlin, 2002. Third millennium edition, revised and expanded.

[15] Akihiro Kanamori. The Higher Infinite: Large Cardinals in Set Theory from Their Beginnings. Springer, Berlin, 1994.

[16] Akihiro Kanamori. The Higher Infinite: Large Cardinals in Set Theory from Their Beginnings. Springer, Berlin, 2003. Second edition.

[17] Peter Koellner and W. Hugh Woodin. Large cardinals from determinacy. Chapter 23 in this Handbook. 10.1007/978-1-4020-5764-9_24.

[18] Kenneth Kunen. Some applications of iterated ultrapowers in set theory. Annals of Mathematical Logic, 1:179–227, 1970.

[19] Kenneth Kunen. Elementary embeddings and infinitary combinatorics. The Journal of Symbolic Logic, 36:407–413, 1971.

[20] Paul B. Larson. The Stationary Tower. Notes on a Course by W. Hugh Woodin, volume 32 of University Lecture Series. American Mathematical Society, Providence, 2004.

[21] Menachem Magidor. Changing cofinality of cardinals. Fundamenta Mathematicae, 99(1):61–71, 1978.

[22] Donald A. Martin and John R. Steel. Iteration trees. Journal of the American Mathematical Society, 7(1):1–73, 1994.

[23] Adrian R. D. Mathias. On sequences generic in the sense of Prikry. Journal of the Australian Mathematical Society, 15:409–414, 1973.

[24] William J. Mitchell. The covering lemma. Chapter 18 in this Handbook. 10.1007/978-1-4020-5764-9_19.

[25] William J. Mitchell. Sets constructible from sequences of ultrafilters. The Journal of Symbolic Logic, 39:57–66, 1974.

[26] William J. Mitchell. How weak is a closed unbounded ultrafilter? In Dirk van Dalen, Daniel Lascar, and Timothy J. Smiley, editors, Logic Colloquium '80 (Prague, 1998), pages 209–230. North-Holland, Amsterdam, 1982.

[27] William J. Mitchell. The core model up to a Woodin cardinal. In Dag Prawitz, Brian Skyrms, Dag Westerstahl, editors, Logic, Methodology and Philosophy of Science, IX (Uppsala, 1991), volume 134 of Studies in Logic and the Foundations of Mathematics, pages 157–175. North-Holland, Amsterdam, 1994.

[28] William J. Mitchell. A measurable cardinal with a closed unbounded set of inaccessibles from $o(\kappa) = \kappa$. Transactions of the American Mathematical Society, 353(12):4863–4897, 2001.

[29] Itay Neeman. Determinacy in $L(\mathbb{R})$. Chapter 22 in this Handbook. 10.1007/978-1-4020-5764-9_23.

[30] Itay Neeman. Inner models in the region of a Woodin limit of Woodin cardinals. Annals of Pure and Applied Logic, 116(1–3):67–155, 2002.

[31] Karl Prikry. Changing measurable into accessible cardinals. Dissertationes Mathematicae (Rozprawy Mathematycne), 68:359–378, 1971.

[32] Lon B. Radin. Adding closed cofinal sequences to large cardinals. Annals of Mathematical Logic, 22(3):243–261, 1982.

[33] Ernest Schimmerling. A core model toolbox and guide. Chapter 20 in this Handbook. 10.1007/978-1-4020-5764-9_21.

[34] Ernest Schimmerling and Martin Zeman. Characterization of $\square_\kappa$ in core models. Journal of Mathematical Logic, 4(1):1–72, 2004.

[35] Ralf Schindler. The core model for almost linear iterations. Annals of Pure and Applied Logic, 116(1–3):205–272, 2002.

[36] Ralf Schindler and Martin Zeman. Fine structure. Chapter 9 in this Handbook. 10.1007/978-1-4020-5764-9_10.

[37] Dana S. Scott. Measurable cardinals and constructible sets. Bulletin de l'Académie Polonaise des Sciences, Série des Sciences Mathématiques, Astronomiques et Physiques, 9:521–524, 1955.

[38] John R. Steel. An outline of inner model theory. Chapter 19 in this Handbook. 10.1007/978-1-4020-5764-9_20.

[39] John R. Steel. Core models with more Woodin cardinals. The Journal of Symbolic Logic, 67(3):1197–1226, 2002.

[40] Philip D. Welch. Σ^* fine structure. Chapter 10 this Handbook. 10.1007/978-1-4020-5764-9_11.

[41] W. Hugh Woodin. The Axiom of Determinacy, Forcing Axioms, and the Nonstationary Ideal. de Gruyter, Berlin, 1999.

[42] Martin Zeman. Inner Models and Large Cardinals. de Gruyter, Berlin, 2002.