Chapter 1

The Transfinite Universe

W. Hugh Woodin

Professor of Mathematics
Department of Mathematics
University of California, Berkeley
Berkeley, CA USA

The 20$^{\text{th}}$ century choice for the axioms [1] of Set Theory are the Zermelo-Frankel axioms together with the Axiom of Choice, these are the ZFC axioms. This particular choice has led to a 21$^{\text{th}}$ century problem:

> **The** ZFC **Delemma:** *Many of the fundamental questions of Set Theory are formally unsolvable from the* ZFC *axioms.*

Perhaps the most famous example is given by the problem of the Continuum Hypothesis: Suppose X is an infinite set of real numbers, must it be the case that either X is countable or that the set X has cardinality equal to the cardinality of the set of all real numbers?

One interpretation of this development is:

> **Skeptic's Attack:** The Continuum Hypothesis is neither true nor false because the entire conception of the universe of sets is a complete fiction. Further, all the theorems of Set Theory are merely finitistic truths, a reflection of the mathematician and not of any genuine mathematical "reality".

Here and in what follows, the "Skeptic" simply refers to the meta-mathematical position which denies any genuine meaning to a conception of uncountable sets. The counter-view is that of the "Set Theorist":

[1] This paper is dedicated to the memory of Paul J. Cohen.

> **The Set Theorist's Response:** The development of Set Theory, after Cohen, has led to the realization that formally unsolvable problems have *degrees of unsolvability* which can be calibrated by *large cardinal axioms.*

Elaborating further, as a consequence of this calibration, it has been discovered that in many cases very different lines of investigation have led to problems whose degree of unsolvability is the same. Thus the hierarchy of large cardinal axioms emerges an intrinsic, fundamental conception within Set Theory. To illustrate this I discuss two examples.

An excellent reference for both the historical development and the background material for much of what I will discuss is the book by Kanamori [5]. The present paper is not intended to be a survey: my intent is to discuss some very recent results which I think have the potential to be relevant to the concept of the Universe of Sets. The danger of course is that this invariably involves speculation and this is compounded whenever such speculation is based on research in progress (as in my manuscript [18]).

1.1 The examples

The first example, infinitary combinatorics. A natural class of objects for study are the subsets of ω_1 which is the least uncountable ordinal. Recall that ω_1 is the set of all *countable ordinals* and so the collection of all subsets of ω_1 is exactly the collection of all sets of countable ordinals. I shall be concerned with two varieties of subsets of ω_1 which I define below.

Definition 1 (1) A set $C \subseteq \omega_1$ is *closed* if, for all $\alpha < \omega_1$, if $C \cap \alpha$ is cofinal in α then $\alpha \in C$.

(2) A set $S \subseteq \omega_1$ is *stationary* if $S \cap C \neq \emptyset$ for all closed, cofinal sets $C \subseteq \omega_1$.□

The sets, $S \subseteq \omega_1$, which are stationary and co-stationary are in many respects the simplest manifestation of the Axiom of Choice. For example one can show, without appealing to the Axiom of Choice, that if there exists a wellordering of the real numbers, then such a set S must exist. The converse is not true as the existence of such a set S does not imply the existence of a wellordering of the real numbers. Recall that a wellordering of the real numbers is total order of the real numbers relative to which every nonempty set of real numbers has a least element.

Therefore, it is natural to ask how complicated the structure of the stationary subsets of ω_1 (modulo non-stationary subsets of ω_1) is or even if there can exist a small generating family for these sets. Consider the following Stationary Basis Hypothesis:

> (SBH) *There exists* ω_1 *many stationary sets,* $\langle S_\alpha : \alpha < \omega_1 \rangle$*, such that for* every *stationary set* $S \subseteq \omega_1$*, there exists* $\alpha < \omega_1$ *such that* $S_\alpha \subseteq S$ *modulo a nonstationary set.*

The assertion that $S_\alpha \subseteq S$ modulo a nonstationary set is simply the assertion that the set,

$$S_\alpha \backslash S = \{\beta < \omega_1 \mid \beta \in S_\alpha \text{ and } \beta \notin S\},$$

is not stationary. Such a sequence $\langle S_\alpha : \alpha < \omega_1 \rangle$ of stationary subsets of ω_1 would give in a natural sense, a *basis* for the stationary subsets of ω_1 which is of cardinality $\aleph_1$.

There is a remarkable theorem of Shelah, [14]:

Theorem 2 (Shelah) *The hypothesis* SBH *implies that* CH *is false.* □

This theorem in conjunction with the subsequent consistency results of [17], suggests the following intriguing conjecture: *The hypothesis* SBH *implies that* $2^{\aleph_0} = \aleph_2$.

The second example, infinite games. Suppose $A \subseteq \mathcal{P}(\mathbb{N})$ where $\mathcal{P}(\mathbb{N})$ denotes the set of all sets $\sigma \subseteq \mathbb{N}$ and $\mathbb{N} = \{1, 2, \ldots, k, \ldots\}$ is the set of all natural numbers.

Associated to the set A is an infinite game involving two players, Player I and Player II. The players alternate declaring at stage k whether $k \in \sigma$ or $k \notin \sigma$ as follows:

> Stage 1: Player I declares $1 \in \sigma$ or declares $1 \notin \sigma$;
>
> Stage 2: Player II declares $2 \in \sigma$ or declares $2 \notin \sigma$;
>
> Stage 3: Player I declares $3 \in \sigma$ or declares $3 \notin \sigma$; ...

After infinitely many stages a set $\sigma \subseteq \mathbb{N}$ is specified. Player I wins this run of the game if $\sigma \in A$; otherwise Player II wins. (Note: Player I has control of which odd numbers are in σ, and Player II has control of which even numbers are in σ.).

A *strategy* is simply a function which provides moves for the players given just the current state of the game. More formally a strategy is a function

$$\tau : [\mathbb{N}]^{<\omega} \times \mathbb{N} \to \{0, 1\},$$

where $[\mathbb{N}]^{<\omega}$ denotes the set of all finite subsets of $\mathbb{N}$. At each stage k of the game the relevant player can choose to follow τ by declaring "$k \in \sigma$" if $\tau(a, k) = 1$

and declaring "$k \notin \sigma$" if $\tau(a, k) = 0$, where

$$a = \{i < k \mid \text{“}i \in \sigma\text{” was declared at stage } i\text{”}\}\,.$$

The strategy τ is a *winning strategy* for Player I if by following the strategy at each stage k where it is Player I's turn to play (i.e., for all odd k), Player I wins the game *no matter how Player II plays*. Similarly τ is a *winning strategy* for Player II if by following the strategy at each stage k where it is Player II's turn to play (i.e., for all even k), Player II wins the game *no matter how Player I plays*. The game is *determined* if there is a *winning strategy* for one of the players. Clearly it is impossible for there to be winning strategies for *both* players.

It is easy to specify sets $A \subseteq \mathcal{P}(\mathbb{N})$ for which the corresponding game is determined. For example if $A = \mathcal{P}(\mathbb{N})$ then any strategy is a winning strategy for Player I. On the other hand, if A is countable, then one can fairly easily show that there exists a strategy which is a winning strategy for Player II.

The problem of specifying a set $A \subseteq \mathcal{P}(\mathbb{N})$ for which the corresponding game is *not* determined, turns out to be quite a bit more difficult. The *Axiom of Determinacy*, AD, is the axiom which asserts that for all sets $A \subseteq \mathcal{P}(\mathbb{N})$, the game given by A, as described above, *is* determined. This axiom was first proposed by Mycielski and Steinhaus, [10], and contradicts the Axiom of Choice. So the problem here is whether the Axiom of Choice is necessary to construct a set $A \subseteq \mathcal{P}(\mathbb{N})$ for which the corresponding game is not determined. Clearly if the Axiom of Choice is necessary, then the existence of such set A is quite a subtle fact.

Three problems and three formal theories. I now add a third problem to the list and specify formally a list of three problems. As indicated, the first and third problems are within ZFC and the second problem is within just the theory ZF (this is the theory given by the axioms ZFC without the Axiom of Choice).

Problem 1: (ZFC) *Does* SBH *hold?*

Problem 2: (ZF) *Does* AD *hold?*

Problem 3: (ZFC) *Do there exist infinitely many Woodin cardinals?*

I shall not give the formal definition of a Woodin cardinal here as it is a large cardinal notion whose definition is a bit technical; see [5] for one definition.

The first problem, **Problem 1**, is *formally unsolvable* if assuming the axioms ZFC one can neither prove or refute the SBH. Similarly, **Problem 2** is formally unsolvable if assuming the axioms ZF, one can neither prove or refute AD. Finally, **Problem 3** is formally unsolvable if assuming the axioms ZFC one can neither prove or refute the existence of infinitely many Woodin cardinals. In each case the assertion of formal unsolvability is simply a statement of Number Theory. The remarkable fact is that *these three assertions of Number Theory are equivalent*, and this is a theorem of Number Theory from the classical (Peano) axioms for Number Theory. Thus, two completely different lines of investigation have resulted in problems whose degree of formal unsolvability is the same, and this is exactly calibrated by a large cardinal axiom.

Assuming that the axioms ZFC are formally consistent, then, for the three problems indicated here, the *only possible* formal solutions are as follows: "No" for the first and third problems, and "yes" for the second problem. Therefore, it is more natural to rephrase these assertions of formal unsolvability as assertions that particular theories are formally consistent. I have implicitly defined three theories, and the assertions of unsolvability discussed above correspond to the assertions that these theories are each formally consistent.

Theory 1: ZFC + SBH.

Theory 2: ZF + AD.

Theory 3: ZFC + "There exist infinitely many Woodin cardinals".

The following theorem is the theorem which implies that the degree of unsolvability of the three problems that I have listed is the same; see [5] for a discussion of this theorem.

Theorem 3 *The three theories,* **Theory 1**, **Theory 2**, *and* **Theory 3**, *are equiconsistent.* □

1.2 A prediction and a challenge for the Skeptic

Are the three theories I have defined really formally consistent? The claim that they are consistent is a prediction which can be refuted by finite evidence (a formal contradiction). Just knowing the first two theories are equiconsistent does not justify this prediction at all. I claim:

> *It is through the calibration by a large cardinal axiom* **in conjunction with** *our understanding of the hierarchy of such axioms as* **true axioms about the universe of sets** *that this prediction is justified.*

As a consequence of my belief in this claim, I make a prediction:

> *In the next 10,000 years there will be no discovery of an inconsistency in these theories.*

This is a specific and unambiguous prediction about *the physical universe*. Further it is a prediction which does *not* arise by a reduction to a previously held truth (as for example is the case for the prediction that no counterexample to Fermat's Last Theorem will be discovered).

This is a genuinely new prediction which I make based on the development of Set Theory over the last 50 years and which I make based on my belief that the conception of the transfinite universe of sets is meaningful. Finally, I make this prediction independently of all speculation of what computational devices might be developed in the next $10,000$ years which increase the effectiveness of research in Mathematics—it is a prediction based on Mathematics and not on consideration of the Mathematician.

Now the Skeptic might object that this prediction is not interesting or natural because the formal theories are not interesting or natural. But such objections are not allowed in Physics: the ultimate physical theory should explain *all* (physical) aspects of the physical universe, not just those which we regard as natural. How can we apply a lessor standard for the ultimate mathematical theory? In fact, I make the stronger prediction:

> *There will be no discovery* ***ever*** *of an inconsistency in these theories.*

One can arguably claim that if this stronger prediction is true, then it is a physical law.

> **Skeptic's Retreat:** OK, I accept the challenge noting that I only have to explain the predictions of formal consistency given by the large cardinal axioms. The formal theory of Set Theory as given by the axioms ZFC is so "incomplete" that: *Any large cardinal axiom, in the natural formulation of such axioms, is either consistent with the axioms of Set Theory, or there is an elementary proof that the axiom cannot hold.*

We shall see this is a very shrewd counter-attack, *even* framed within the specific context of the current list of large cardinal axioms where it is a much more plausible position. I shall need to review some elementary concepts from Set Theory. This is necessary to specify the basic template for large cardinal axioms.

1.3 The cumulative hierarchy of sets.

As is customary in modern Set Theory, V denotes the Universe of Sets. The purpose of this notation is to facilitate the (mathematical) discussion of Set

Theory, it does not presuppose any meaning to the concept of the Universe of Sets.

The *ordinals* calibrate V through the definition of the cumulative hierarchy of sets, [19]. The relevant definition is given below.

Definition 4 Define for each ordinal α a set V_α by induction on α.

(1) $V_0 = \emptyset$.

(2) $V_{\alpha+1} = \mathcal{P}(V_\alpha) = \{X \mid X \subseteq V_\alpha\}$.

1. If β is a limit ordinal, then $V_\alpha = \bigcup\{V_\beta \mid \beta < \alpha\}$. □

It is a consequence of the ZF axioms that for every set a there must exist an ordinal α such that $a \in V_\alpha$.

A set N is *transitive* if every element of N is a subset of N. Transitive sets are fragments of V which are analogous to initial segments. For each ordinal α the set V_α is a transitive set.

Every ordinal is a transitive set; in fact, the ordinals are precisely those transitive sets X with the property that for all $a, b \in X$, if $a \neq b$ then either $a \in b$ or $b \in a$. Thus if X is an ordinal and if $Y \in X$, then necessarily Y is an ordinal. This defines a natural order on the ordinals. If α and β are ordinals then $\alpha < \beta$ if $\alpha \in \beta$. Thus every ordinal is simply the set of all ordinals which are smaller than the given ordinal, relative to this order.

The simplest (proper) class is the class of all ordinals. This class is a transitive class and more generally a class $M \subseteq V$ is a transitive class if every element of M is a subset of M. The basic template for large cardinal axioms is as follows.

There is a transitive class M and an elementary embedding

$$j : V \to M$$

which is not the identity.

With the exception of the definition of a *Reinhardt cardinal* which I shall come to below, one can always assume that the classes, M and j, are classes which are logically definable from parameters by formulas of a fixed bounded level of complexity (Σ_2-formulas). Moreover the assertion that j is an elementary embedding, is the assertion:

For all formulas $\phi(x)$ and for all sets a, $V \models \phi[a]$ if and only if $M \models \phi[j(a)]$;

and this is equivalent to the assertion:

For all formulas $\phi(x)$, for all ordinals α, and for all sets $a \in V_\alpha$, $V_\alpha \models \phi[a]$ if and only if $j(V_\alpha) \models \phi[j(a)]$.

Therefore, this template makes no essential use of the notion of a class. It is simply for convenience that I refer to classes (and this is the usual practice in Set Theory).

Suppose that M is a transitive class and that $j : V \to M$ is an elementary embedding which not the identity. Suppose that $j(\alpha) = \alpha$ for all ordinals α. Then one can show by transfinite induction that for all ordinals α, the embedding j is the identity on V_α. Therefore, since j is not the identity, there must exist an ordinal α such that $j(\alpha) \neq \alpha$ and since j is order-preserving on the ordinals this is equivalent to the requirement that $\alpha < j(\alpha)$. The least such ordinal is called the *critical point* of j and it can be shown that this must be a cardinal. The critical point of j is the large cardinal specified and the existence of the transitive class M and the elementary embedding j are the witnesses for this.

A cardinal κ is a *measurable cardinal* if there exists a transitive class M and an elementary embedding $j : V \to M$ such that κ is the critical point of j.

It is by requiring M to be *closer* to V that one can define large cardinal axioms far beyond the axiom, "There is a measurable cardinal". In general the closer one requires M to be to V, the stronger the large cardinal axiom. In [12], the natural maximum axiom was proposed ($M = V$). The associated large cardinal axiom is that of a *Reinhardt cardinal*.

Definition 5 A cardinal κ is a *Reinhardt cardinal* if there is an elementary embedding, $j : V \to V$ such that κ is the critical point of j. □

The definition of a Reinhardt cardinal makes essential use of classes, but the following variation does not and is more useful for this discussion. The definition requires a logical notion. Suppose that α and β are ordinals such that $\alpha < \beta$. Then we write $V_\alpha \prec V_\beta$ to mean that for all formulas $\phi(x)$ and for all $a \in V_\alpha$, $V_\alpha \models \phi[a]$ if and only if $V_\beta \models \phi[a]$.

Definition 6 A cardinal κ is a *weak Reinhardt cardinal* if there exist $\gamma > \lambda > \kappa$ such that

(1) $V_\kappa \prec V_\lambda \prec V_\gamma$,

(2) there exists an elementary embedding, $j : V_{\lambda+2} \to V_{\lambda+2}$ such that κ is the critical point of j. □

The definition of a weak Reinhardt cardinal only involves sets. The relationship between Reinhardt cardinals and weak Reinhardt cardinals is unclear;

but, given the original motivation for the definition of a Reinhardt cardinal, one would conjecture that at least in terms of consistency strength, Reinhardt cardinals are stronger than weak Reinhardt cardinals. Hence my choice of terminology. In any case, the concept of a weak Reinhardt cardinal is better suited to illustrate the key points I am trying to make.

The following theorem is an immediate corollary of the fundamental inconsistency results of Kunen [7].

Theorem 7 (Kunen) *There are no weak Reinhardt cardinals.* □

The proof is elementary so this does not refute the Skeptic's Retreat. But Kunen's proof makes essential use of the Axiom of Choice. The problem is open without this assumption, and this is not just an issue for weak Reinhardt cardinals, which is just a notion of large cardinal defined in this paper. There really is no known interesting example of a strengthening of the definition of a Reinhardt cardinal that yields a large cardinal axiom which can be refuted without using the Axiom of Choice. The difficulty is that without the Axiom of Choice it is extraordinarily difficult to prove anything about sets.

Kunen's proof leaves open the possibility that the following large cardinal axiom might be consistent with the Axiom of Choice. This therefore is essentially the strongest large cardinal axiom not known to be refuted by the Axiom of Choice; see [5] for more on this as well as for the actual statement of Kunen's theorem.

Definition 8 A cardinal κ is an *ω-huge cardinal* if there exists $\lambda > \kappa$ and an elementary embedding $j : V_{\lambda+1} \to V_{\lambda+1}$ such that κ is the critical point of j. □

One could strengthen this axiom still further by requiring in addition that for some $\gamma > \lambda$, we have $V_\kappa \prec V_\lambda \prec V_\gamma$ and so match in form the definition of a weak Reinhardt cardinal, the only modification being that $\lambda + 2$ is replaced by $\lambda + 1$. This change would not affect *any* of the claims below concerning ω-huge cardinals.

The issue of whether the existence of a weak Reinhardt cardinal is consistent with the axioms ZF is an important issue for the Set Theorist because by the results of [18], the theory

$$\text{ZF} + \text{“There is a weak Reinhardt cardinal”}$$

proves the formal consistency of the theory

$$\text{ZFC} + \text{“There is a proper class of } \omega\text{-huge cardinals”}.$$

This number-theoretic statement is a theorem of Number Theory. But, as indicated above, the notion of an ω-huge cardinal is essentially the strongest large cardinal notion which is not known to be refuted by the Axiom of Choice.

Therefore the number theoretic assertion that the theory

ZF + "There is a weak Reinhardt cardinal"

is consistent is a *stronger* assertion than the number theoretic assertion that the theory

ZFC + "There is a proper class of ω-huge cardinals"

is consistent. More precisely, the former assertion implies, *but is not implied by*, the latter assertion; unless of course the theory

ZFC + "There is a proper class of ω-huge cardinals"

is formally inconsistent. This raises an interesting question:

> *How could the Set Theorist ever be able to argue for the prediction that the existence of weak Reinhardt cardinals is consistent with axioms of Set Theory without the Axiom of Choice?*

Moreover this *one* prediction implies *all* the predictions (of formal consistency) the Set Theorist can currently make based on the *entire* large cardinal hierarchy as presently conceived (in the context of a universe of sets which satisfies the Axiom of Choice). My point is that by appealing to the Skeptic's Retreat, one could reasonably claim that the theory

ZF + "There is a weak Reinhardt cardinal",

is formally consistent—and in making this *single* claim one would subsume *all* the claims of consistency that the Set Theorist can make based on our current understanding of the Universe of Sets (without abandoning the Axiom of Choice).

The only tools currently available seem powerless to resolve this issue. Reinterpreting the number theoretic statement that the theory,

ZF + "There is a weak Reinhardt cardinal",

is formally consistent, in a way that allows the Set Theorist to argue for the truth of this statement, seems equally hopeless. Finally, unlike the axiom AD there is no candidate presently known for a fragment of V for which the existence of weak Reinhardt cardinals is the correct (or even a possible) axiom.

1.4 Probing the Universe of Sets: the Inner Model Program

The *Inner Model Program* is the detailed study of large cardinal axioms. The first construction of an inner model is due to Gödel, [3] and [4]. This construction founded the Inner Model Program, and the transitive class constructed is denoted by L. This is the *minimum* possible universe of sets containing all ordinals.

If X is a transitive set, then $\mathrm{Def}(X)$ denotes the set of all $A \subseteq X$ such that A is logically definable in the structure $(X, \in)$ from parameters. The definition of L is simply given by replacing the operation $\mathcal{P}(X)$ in the definition of $V_{\alpha+1}$ by the operation $\mathrm{Def}(X)$. More precisely:

Definition 9 (1) Define L_α by induction on the ordinal α:

(a) $L_0 = \emptyset$ and $L_{\alpha+1} = \mathrm{Def}(L_\alpha)$;

(b) If α is a limit ordinal then $L_\alpha = \bigcup\{L_\beta \mid \beta < \alpha\}$.

(2) L is the class of all sets a such that $a \in L_\alpha$ for some ordinal α. □

It is perhaps important to note that while there must exist a proper class of ordinals α such that

$$L_\alpha = L \cap V_\alpha,$$

this is not true for all ordinals α.

The question of whether $V = L$ is an important one for Set Theory. The answer has profound implications for the conception of the Universe of Sets.

Theorem 10 (Scott, [13]) *Suppose there is a measurable cardinal. Then* $V \neq L$. □

The *Axiom of Constructibility* is the axiom which asserts $V = L$; more precisely this is the axiom which asserts that for each set a there exists an ordinal α such that $a \in L_\alpha$. Scott's Theorem provided the first indication that the Axiom of Consructibility is independent of the ZFC axioms. At the time there was no compelling reason to believe that the existence of a measurable cardinal was consistent with the ZFC axioms, so one could not make the claim that Scott's theorem established the formal independence of the Axiom of Consructibility from the ZFC axioms. Of course it is an immediate corollary of Cohen's results that the Axiom of Consructibility is formally independent of the ZFC axioms. The modern significance of Scott's theorem is more profound: Scott's theorem establishes that the Axiom of Consructibility is *false*. This claim (that $V \neq L$) is not universally accepted, but in my view no one has come up with a credible argument against this claim.

The Inner Model Program seeks generalizations of L for the large cardinal axioms; in brief it seeks generalizations of the Axiom of Consructibility which are compatible with large cardinal axioms (such as the axioms for measurable cardinals and beyond). It has been a very successful program and its successes have led to the realization that the large cardinal hierarchy is a very "robust" notion. The results which have been obtained provide some of our deepest glimpses into the Universe of Sets, and its successes have led to a "meta-prediction":

> **A Set Theorist's Cosmological Principle:** *The large cardinal axioms for which there is an inner model theory are consistent; the corresponding predictions of unsolvability are true* ***because the axioms are true****.*

Despite the the rather formidable merits as indicated above, there is a fundamental difficulty with the prospect of using the Inner Model Program to counter the Skeptic's Retreat. The problem is in the basic methodology of the Inner Model Program, but to explain this I must give a (brief) description of the (technical) template for inner models.

The inner models which are the goal and focus of the Inner Model Program are defined layer by layer working up through the hierarchy of large cardinal axioms, which in turn is naturally revealed by the construction of these inner models. Each layer provides the foundation for the next, and L is the first layer. Roughly (and in practice) in constructing the inner model for a specific large cardinal axiom, one obtains an exhaustive analysis of all weaker large cardinal axioms. There can be surprises here in that seemingly different notions of large cardinals can coincide in the inner model. Finally as one ascends through the hierarchy of large cardinal axioms, the construction generally becomes more and more difficult.

1.5 The building blocks for inner models: Extenders

Suppose that M is a transitive class and that $j : V \to M$ is an elementary embedding with critical point κ. As with the basic template for large cardinal axioms I discussed above, here and below one can restrict to the classes which are definable classes (by Σ_2-formulas) etc., so that no essential use of classes is involved.

It is immediate from the definitions that for all ordinals γ, $j(\gamma)$ is an ordinal and moreover $j(\gamma) \geq \gamma$. Suppose that $\kappa < \gamma < j(\kappa)$ and that $\mathcal{P}(\gamma) \subseteq M$ where $\mathcal{P}(\gamma) = \{A \mid A \subseteq \gamma\}$. The function:

$$E(A) = j(A) \cap \gamma$$

with domain $\mathcal{P}(\gamma)$ is the *extender* E of length γ defined from j. Note that since $\gamma > \kappa$, necessarily E is not the identity function. Extenders are nontrivial *fragments* of the elementary embedding j. (The concept of an extender is due to Mitchell.) The definition that I have given is really that of a *strong* extender because of the assumption that $\mathcal{P}(\gamma) \subseteq M$. This I do for expository reasons. In the case that $\gamma < j(\kappa)$, which is the present case, one could drop this requirement without affecting much of the discussion.

The importance of the concept of an extender is the following. Suppose that E is an extender of length γ derived from an elementary embedding $j : V \to M$ and that N is a transitive class such that $N \models \text{ZFC}$. Suppose that $E \cap N \in N$ and that $\gamma = \kappa + 1$ where κ is the critical point of j. Then there exists a transitive class $M_E \subseteq N$ and an elementary embedding,

$$j_E : N \to M_E$$

such that $E \cap N$ is the extender of length γ derived from j_E. The point here of course is that the assumption is only that $E \cap N \in N$ as opposed to the much stronger assumption that $E \in N$. Both M_E and j_E can be chosen to be definable classes of N (by Σ_2-formulas) using just the parameter, $E \cap N$.

Without the assumption that $\gamma = \kappa + 1$, which is a very special case, these claims still hold provided that one drops the requirement $\mathcal{P}(\gamma) \subseteq M$ in the definition of an extender that I have given.

These remarks suggest that one should seek, as generalizations of L, transitive classes N such that N contains *enough* extenders of the form, $E \cap N$, for some extender $E \in V$ to witness that the targeted large cardinal axiom holds in N. One can then regard such transitive classes as *refinements* of V which are constructed to "preserve" certain extenders of V. The complication is in specifying just which extenders are to be preserved.

For each set A one can naturally define a class $L[A]$ which is L *relativized* to the set A, as follows:

Definition 11 (1) Define $L_\alpha[A]$ by induction on the ordinal α:

(a) $L_0 = \emptyset$ and $L_{\alpha+1}[A] = \text{Def}(X) \cap \mathcal{P}(L_\alpha[A])$ where $X = L_\alpha[A] \cup \{L_\alpha[A] \cap A\}$;

(b) If α is a limit ordinal then $L_\alpha[A] = \bigcup\{L_\beta[A] \mid \beta < \alpha\}$.

(2) $L[A]$ is the class of all sets a such that $a \in L_\alpha[A]$ for some ordinal α. $\square$

If F is a function, then $L[F]$ is defined to be $L[A]$ where $A = F$. Thus, if the domain of F is disjoint from L then $L[F] = L$.

Constructing from a single extender E yields $L[E]$, which is a true generalization of L and solves the inner model problem for the large cardinal axiom: "There is a measurable cardinal". This claim follows from the results and methods of [6] and is illustrated in part by the theorem below.

There is a feature of inner models for large cardinals which is implicit in this example:

> For a specific large cardinal axiom there is in general no unique inner model for that axiom, but rather a family of inner models. But all these inner models are equivalent in a natural (but technical) sense.

One illustration of this is given by the following theorem, which is a modern formulation of the fundamental results of Kunen [6] on the inner model problem for one measurable cardinal. For each extender E, let κ_E denote the least ordinal α such that $E(\alpha) \neq \alpha$. This coincides with critical point of the elementary embedding, $j : V \to M$ from which E is derived.

Theorem 12 *Suppose that E and F are extenders.*

(1) *If $\kappa_E = \kappa_F$, then $L[E] = L[F]$.*

(2) *If $\kappa_E < \kappa_F$, then $L[F] \subset L[E]$ and there is an elementary embedding,*

$$j : L[E] \to L[F].$$ □

For the generalizations of $L[E]$ which one must consider to solve the inner model problem for large cardinals beyond the level of measurable cardinals, this ambiguity is much more subtle and lies at the core of the difficulty in even defining the inner models.

By the theorem above one cannot use a single extender to build an inner model for essentially any large cardinal axiom beyond the level of a single measurable cardinal. For example suppose that E is an extender. Combining elements of Gödel's basic analysis of L, generalized to an analysis of $L[E]$ with Theorem 12, it follows that the inner model $L[E]$ will fail satisfy the large cardinal axiom "There are two measurable cardinals". There is an obvious remedy: to reach inner models for stronger large cardinal axioms one should use sequences,

$$\tilde{E} = \langle E_\alpha : \alpha < \theta \rangle,$$

where each E_α is the extender derived from some elementary embedding as above. The complication is in how to actually define the sequence; in fact, one must ultimately allow the sequence to contain *partial* extenders which creates still further complications.

A partial extender of length γ is the extender of length γ derived from a Σ_0-elementary embedding $j : N \to M$ where N and M are transitive sets which are only assumed to be closed under the Gödel operations, and, instead of requiring $\mathcal{P}(\gamma) \subseteq M$, one requires that

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

The requirement that j be a Σ_0-elementary embedding is the requirement that for a very restricted collection of formulas, $\phi(x_0)$, and for all $a \in N$ $N \models \phi[a]$ if and only if $M \models \phi[j(a)]$. The relevant formulas are the Σ_0-formulas.

The difficulty mentioned above is in determining exactly when such partial extenders are acceptable. In fact, things get so complicated that, unlike the situation with measurable cardinals, one can only define the inner model by simultaneously developing the detailed analysis of the inner model in an elaborate induction.

The current state of the art is found in the inner models defined by Mitchell and Steel, [9]. The definition of these inner models is the culmination of a nearly 20 year program of developing the theory of inner models. The Mitchell-Steel inner models can accommodate large cardinals up to the level of *superstrong cardinals*, but existence has only been proved—from the relevant large cardinal axioms—at the level of a Woodin cardinal which is a limit Woodin cardinals. In this program of establishing existence of Mitchell-Steel inner models, the best results to date are due to I. Neeman, [11].

The distinction between developing the theory of the inner models and proving existence of the inner models is perhaps a confusing one at first glance. The precise explanation requires details of the Mitchell-Steel Theory which are beyond the scope of the present discussion. Roughly, the Mitchell-Steel Theory reduces the problem of the existence of the generalization of L for the large cardinal axiom under consideration to a specific combinatorial hypothesis, provided that the large cardinal axiom is at the level of a superstrong cardinal or below. This combinatorial (iteration) hypothesis can be specified without any reference to the Mitchell-Steel Theory and, more generally, without reference to inner model theory at all. There is of course the possibility that this is symptomatic of a far more serious problem and that by answering one of the test questions of the Inner Model Program negatively one can prove that the Inner Model Program as presently conceived fails for some large cardinal axiom below the level of a superstrong cardinal.

1.6 The Inner Model Program, the Core Model Program and the Skeptic's Retreat

As I have claimed, there is a fundamental problem with appealing to the Inner Model Program to counter the Skeptic's Retreat. The precise nature of the problem is subtle so I shall begin by describing what might seem to be a plausible version of the problem. I then briefly will try to describe the actual problem. This will involve the *Core Model Program* which is a variant of the Inner Model

Program.

Suppose (for example) that a hypothetical large cardinal axiom "Ω" provides a counterexample to the Skeptic's Retreat, and this is accomplished by the Inner Model Program.

> *To use the Inner Model Program to refute the existence of an "Ω-cardinal" one first must be able to successfully construct the inner models for all smaller large cardinals, and this hierarchy would be fully revealed by the construction.*

Perhaps this could happen, but it can only happen *once*. This is the problem. Having refuted the existence of an "Ω-cardinal", how could one then refute the existence of any *smaller* large cardinals, for one would have solved the inner model problem for these smaller large cardinals. This would **refute** the Set Theorist's Cosmological Principle. So the fundamental problem is:

> *The Inner Model Program seems inherently unable, by virtue of its inductive nature, to provide a framework for an evolving understanding of the boundary between the possible and the impossible (large cardinal axioms).*

Upon close inspection it is perhaps not entirely convincing that there is a problem here. Arguably there is the potential for a problem, but the specific details of how the Inner Model Program might succeed in countering the Skeptic's Retreat are clearly critical in determining whether or not there really is a problem.

Though idea that the Inner Model Program could ever yield an inconsistency result has always seemed unlikely, there is another way that the Inner Model Program might succeed in establishing inconsistency results in a manner that refutes the Skeptic's Retreat. The *Core Model Program* can be described as follows:

Suppose that $L[\tilde{E}]$ is an inner model as constructed by the Inner Model Program. In general would one not expect that the inner model $L[\tilde{E}]$ to contain even all the real numbers; for example if the Continuum Hypothesis is false in V, then necessarily there are real numbers which are not in $L[\tilde{E}]$. Therefore, *every* extender on the sequence $\tilde{E}$ when restricted to $L[\tilde{E}]$ is necessarily a partial extender in V. This suggests that one might attempt to construct the inner model $L[\tilde{E}]$ without using extenders at all, just partial extenders.

While this might seem reasonable, there is of course a problem. If there are no extenders in V, then there are no measurable cardinals in V; and so one cannot in general expect to be able to build an inner model in which there are measurable cardinals. But suppose one assumes that there is no proper transitive class N in which a particular large cardinal axiom holds. Then a reasonable conjecture is that there is an inner model of the form $L[\tilde{E}]$ which is "close" to V.

One measure of the closeness of an inner model N to V is a *weak covering principle*. This requires a definition. A cardinal γ is *singular* if there exists a cofinal set $X \subseteq \gamma$ such that $|X| < \gamma$, and γ^+ refers to the least cardinal κ such that $\kappa > \gamma$. Suppose that N is a (proper) transitive class and that $N \models \mathrm{ZFC}$. Then *weak covering* holds for V relative to the inner model N if for all uncountable singular cardinals γ, if $\gamma = |V_\gamma|$, then

$$(\gamma^+)^N = \gamma^+.$$

Allowing that case that $\tilde{E} = \emptyset$, so that $L[\tilde{E}] = L$, this becomes a very interesting problem. The program to solve this family of problems is the Core Model Program. Both the Inner Model Program and the Core Model Program seek to construct exactly the same form of an inner model: the only difference is in the assumptions from which one starts. The Inner Model Program starts with the assumption that a particular large cardinal axiom holds in V, whereas the Core Model Program starts with the assumption that a particular large cardinal axiom *does not* hold in any transitive class $N \subseteq V$. It is customary to refer to the transitive classes constructed by the methods of the Core Model Program as *core models*.

The Core Model Program was inspired by Jensen's Covering Lemma and began with the results of Dodd and Jensen [2]. The strongest results to date are are primarily due to Steel, who extended the Core Model Program to the level of Woodin cardinals in [15]. As with the Inner Model Program, the solutions provided by the Core Model Program increase in complexity as the associated large cardinal axiom is strengthened.

The Core Model Program has been quite successful, and out of it have come a number of deep combinatorial theorems. For example the methods and constructions of Core Model Program play an essential role in the proof of Theorem 3. One might suspect that utility of the Core Model Program is limited for proving the kinds of theorems that an inconsistency result would require, because of the requisite hypothesis that there be no proper transitive class N in which a specific large cardinal axiom holds. But despite this requirement, the Core Model Program has yielded some surprising theorems of Set Theory.

Recall that the *Generalized Continuum Hypothesis* (GCH) is the assertion that for all infinite cardinals, γ, $2^\gamma = \gamma^+$ where 2^γ is the cardinality of $\mathcal{P}(\gamma) = \{X \mid X \subseteq \gamma\}$. The following theorem is an example of a theorem proved by the methods of the Core Model Program.

Theorem 13 ([16]) *Suppose that there exists a countable set A such that A is a set of ordinals and $V = L[A]$. Then the* GCH *holds.* □

On general grounds, to prove the theorem it suffices from the hypothesis of the theorem to just prove that the Continuum Hypothesis holds. The special case where $A \subseteq \omega$ follows from Gödel's analysis of L, generalized to the analysis of $L[A]$. For the general case where one does not assume $A \subseteq \omega$, there is no elementary proof of the Continuum Hypothesis known.

By a theorem of Jensen [1], for *any* sentence ϕ, if the sentence is consistent with the axioms ZFC, then the existence of a proper class N within which the sentence holds is consistent with the hypothesis of Theorem 13. If in addition the sentence is consistent with the axioms ZFC + GCH, then one can even require that the transitive class N be close to $L[A]$. For example one can require that N and $L[A]$ have the same cardinals. While this additional consistency assumption may seem like a very restrictive assumption, at least for the current generation of large cardinal axioms, it is not.

Therefore, it is perhaps not unreasonable that the Core Model Program might yield that some proposed large cardinal axiom is inconsistent and in doing so, refute the Skeptic's Retreat. But to accomplish this, the Core Model Program would seem to have to produce an "ultimate" core model corresponding to the ultimate inner model. But, if this is an inner model of the form $L[\tilde{E}]$, for some sequence of (partial) extenders, as is the case for essentially all core models which have been constructed to date, then the nature of the extenders on the sequence $\tilde{E}$ should reveal the entire large cardinal hierarchy—and we again are in a situation where further progress looks unlikely.

Thus, it would seem that the Skeptic's Retreat is in fact a powerful counterattack. But there is something wrong here, some fundamental misconception. The answer lies in understanding large cardinal axioms which are much stronger than those within reach of the Mitchell-Steel hierarchy of inner models. Ironically, one outcome of my proposed resolution to this misconception is that for these large cardinal axioms, the Set Theorist's Cosmological Principle is either false or useless. I shall discuss these ramifications after Theorem 22.

1.7 Supercompact cardinals and beyond

We begin with a definition.

Definition 14 (Solovay) A cardinal κ is a *supercompact cardinal* if for each ordinal α there exists a transitive class M and an elementary embedding $j : V \to M$ such that

(1) κ is the critical point of j and $j(\kappa) > \alpha$;

(2) M contains all functions, $f : \alpha \to M$. □

Still stronger are *extendible cardinals*, *huge cardinals*, and *n-huge cardinals* where $n < \omega$. These I shall not define here. As I have already indicated, the strongest large cardinal axioms not known to be inconsistent with the Axiom of Choice are the family of axioms asserting the existence of ω-huge cardinals. These axioms have seemed so far beyond any conceivable inner model theory that they simply are not understood.

The possibilities for an inner model theory at the level of supercompact cardinals and beyond has been essentially a complete mystery until recently. The reason lies in the nature of extenders. Suppose that E is an extender of length γ derived from an elementary embedding $j : V \to M$ with critical point κ and such that $\mathcal{P}(\gamma) \subseteq M$. If $\gamma \leq j(\kappa)$ then E is a *short* extender, otherwise E is a *long* extender. Up to this point I have only considered short extenders. The properties of long extenders can be quite subtle and it is for this reason that I impose the requirement $\mathcal{P}(\gamma) \subseteq M$. Even with this requirement many subtleties remain. For example, Theorem 12, which I implicitly stated for short extenders, is *false* for long extenders. One can prove the following variation provided the extenders are not *too* long. For expository purposes let me define an extender E to be a *suitable* extender if E is the extender of length γ derived from an elementary embedding $j : V \to M$ such that $\mathcal{P}(\gamma) \subseteq M$ and such that $\gamma < j(\alpha)$ for some $\alpha < j(\kappa)$, where κ is the critical point of j. Suitable extenders can be long extenders but they cannot be too long.

Theorem 15 *Suppose that E and F are suitable extenders. Then*

$$\mathbb{R} \cap L[E] \subseteq \mathbb{R} \cap L[F] \quad \textit{or} \quad \mathbb{R} \cap L[F] \subseteq \mathbb{R} \cap L[E]. \qquad \square$$

Without the restriction to suitable extenders, it is not known if this theorem holds.

The following lemma of Magidor reformulated in terms of suitable extenders gives a useful reformulation of supercompactness, [18]. The statement of the lemma involves the following notation which I previously defined for short extenders. Suppose E is an extender of length γ derived from an elementary embedding $j : V \to M$. Then κ_E is the critical point of j. Suppose that $\alpha < \gamma$. Then by the definition of E, and since $\alpha = \{\beta \mid \beta < \alpha\} \subseteq \gamma$, we find that α is in the domain of E and either $E(\alpha) = \gamma$ or $E(\alpha) < \gamma$. Moreover the following hold:

1. $E(\alpha) < \gamma$ if and only if $j(\alpha) < \gamma$ and $E(\alpha) = j(\alpha)$,

2. $E(\alpha) = \gamma$ if and only if $j(\alpha) \geq \gamma$.

Thus, κ_E is simply the least α such that $E(\alpha) \neq \alpha$.

Lemma 16 *Suppose that δ is a cardinal. Then the following are equivalent:*

(1) *δ is supercompact;*

(2) *For each ordinal $\gamma > \delta$, there exists a suitable extender E of length γ such that $E(\kappa_E) = \delta$.* □

My convention in what follows is that a class $\mathcal{E}$ of extenders *witnesses that δ is a supercompact cardinal* if, for each $\gamma > \delta$, there exists an extender $E \in \mathcal{E}$ such that

1. E has length γ, and
2. $E(\kappa_E) = \delta$ and for some $\alpha < \delta$, $E(\alpha) = \gamma$.

Note that condition (2) implies that E is a suitable extender.

The Mitchell-Steel inner models are constructed from sequences of short extenders. But to build inner models at the level of supercompact cardinals and beyond one must have long extenders on the sequence, and this creates serious obstacles if the these extenders are "too" long. In fact, Steel has isolated a specific obstacle which becomes severe at the level of one supercompact cardinal with a measurable cardinal above.

But by some fairly recent theorems from [18] something completely unexpected and remarkable happens. Suppose that N is a transitive class, for some cardinal δ,

$$N \models \text{"}\delta \text{ is a supercompact cardinal"},$$

and that this is witnessed by class of all $E \cap N$ such that $E \cap N \in N$ and such that E is an extender. Then the transitive class N is close to V and N inherits essentially all large cardinals from V.

For example, suppose that for each n there is a proper class of n-huge cardinals in V. Then in N, for each n, there is a proper class of n-huge cardinals. The amazing thing is that this must happen *no matter how N is constructed.* This would seem to undermine my earlier claim that inner models should be constructed as refinements of V which preserve enough extenders from V to witness that the targeted large cardinal axiom holds in the inner model. It does not, and the reason is that by simply requiring that $E \cap N \in N$ for enough suitable extenders from V to witness that the large cardinal axiom, "There is a supercompact cardinal", holds in N, one (and this is the surprise) necessarily must have $E \cap N \in N$ for a *much* larger class of extenders of V. So the *principle* that there are enough extenders of N which are of the form $E \cap N$ for some extender $E \in V$, to witness the targeted large cardinal axiom holds in N, *is preserved*. The *change*, in the case that N is constructed from a sequence of extenders, is that these extenders do not have to be on the sequence from which N is constructed. In particular, in the case that the sequence of extenders from which N is constructed contains only suitable extenders, large cardinal axioms can be witnessed to hold in N by the "phantom" extenders (these are extenders of N which are not on the sequence) which *cannot* be witnessed to hold by *any* extender on the sequence.

As a consequence of this, one can completely avoid the cited obstacles because:

> *One does not need to have the kinds of long extenders on the sequence which give rise to the obstacles.*

Specifically, one can restrict consideration to extender sequences of just suitable extenders and this is a paradigm shift in the whole conception of inner models.

The analysis yields still more. Suppose that there is a positive solution (in ZFC) to the inner model problem for just one supercompact cardinal. Note that this seems at the edge of feasibility without encountering the serious obstacles raised by long extenders. More precisely, suppose that, if κ is a supercompact cardinal, then (provably) there is a definable sequence

$$\tilde{E} = \langle E_\alpha : \alpha \in \mathrm{Ord} \rangle$$

of (partial) extenders such that $L[\tilde{E}]$ is an L-like inner model in which κ is supercompact, and, that this is witnessed by the extenders on the sequence (which are the restriction of true extenders to the associated inner model $L[\tilde{E}]$). Then as a corollary one would would obtain a proof of the following conjecture:

Conjecture (ZF) *There are no weak Reinhardt cardinals.* □

It is possible to isolate a specific conjecture which must be true if there is a positive solution to the inner problem for one supercompact cardinal, as described above, and which itself suffices for this inconsistency result. To explain this further I must give one last definition.

This is the definition of the class, HOD, which originates in remarks of Gödel at the Princeton Bicentennial Conference in December, 1946. The first detailed reference appears to be [8] (see the review of [8] by G. Kreisel).

Definition 17 (ZF) (1) For each ordinal α, let HOD_α be the set of all sets a such that there exists a set $A \subseteq \alpha$ such that

(a) A is definable in V_α from ordinal parameters;

(b) $a \in L_\alpha[A]$.

(2) HOD is the class of all sets a such that $a \in \mathrm{HOD}_\alpha$ for some α. □

The definition of HOD_α combines features of the definition of L_α and features of the definition of V_α. I caution that, just as is the case for L_α, in general we have

$$\mathrm{HOD}_\alpha \neq \mathrm{HOD} \cap V_\alpha,$$

though for a proper class of ordinals α it is true that $\mathrm{HOD}_\alpha = \mathrm{HOD} \cap V_\alpha$.

The class HOD is quite interesting for a number of reasons one of which is illustrated by the following observation of Gödel which as indicated is stated within just the theory ZF, in other words without assuming the Axiom of Choice.

Theorem 18 (ZF) *HOD $\models$ ZFC.* □

This theorem gives a completely different approach to showing that if ZF is formally consistent then so is ZFC.

One difficulty with HOD is that the definition of HOD is not absolute; for example, in general HOD is not even the same as defined *within* HOD. As a consequence almost any set theoretic question one might naturally ask about HOD is formally unsolvable. Two immediate such questions are whether $V = \mathrm{HOD}$ and, more simply, whether HOD contains all the real numbers. Both of these questions are formally unsolvable but are of evident importance because they specifically address the complexity of the Axiom of Choice. If $V = \mathrm{HOD}$ then there is no mystery as to why the Axiom of Choice holds, but of course one is left with the problem of explaining why $V = \mathrm{HOD}$. If $V = L$ then it is easy to verify that $V = \mathrm{HOD}$. Further the inner models of Mitchell-Steel (in the situations where existence can be proved) can always be constructed to be contained in HOD even though the axiom $V = \mathrm{HOD}$ can *fail* in a Mitchell-Steel inner model.

I now present a key conjecture. This conjecture involves the notion that an uncountable cardinal γ is a *regular* cardinal. This is the property that for all $X \subseteq \gamma$, if $|X| < \gamma$, then X is bounded in γ. Alternatively, referring to notions already defined, an uncountable cardinal γ is a regular cardinal if it is not a singular cardinal.

The HOD **Conjecture**: (ZFC) *Suppose that κ is a supercompact cardinal. Then there exists a regular cardinal $\gamma > \kappa$ which is not a measurable cardinal in* HOD. □

If there is a supercompact cardinal in V then the HOD Conjecture implies that HOD is "close" to V. Assuming that a slightly stronger large cardinal hypothesis holds in V (and that the HOD Conjecture also holds in V), then one actually obtains that HOD is quite close to V. As evidence for this latter claim, I note the following theorem. Part (3) of this theorem is a very strong closure condition for HOD and part (2) follows directly from part (3). The analogous closure condition holds for *any* transitive (proper) class $N \models \mathrm{ZFC}$ such that for some cardinal δ,

$$N \models \text{“}\delta \text{ is a supercompact cardinal”},$$

and such that this is witnessed by class of all restrictions $E \cap N$, where E is an extender and $E \cap N \in N$. Referring back to the discussion on the distinction between the Core Model Program and the Inner Model Program, this closure condition establishes that in the context of the existence of just one supercompact cardinal, there is *no difference* between the primary objectives of these two programs.

Theorem 19 (ZFC) *Suppose that there is an extendible cardinal and that the* HOD Conjecture *holds. Then the following hold:*

(1) *There exists an ordinal α such that for all cardinals $\gamma > \alpha$, if γ is a sinular cardinal then $\gamma^+ = (\gamma^+)^{\mathrm{HOD}}$;*

(2) *suppose for each n there is a proper class of n-huge cardinals. Then for each n,*

$$\mathrm{HOD} \models \text{“There is a proper class of } n\text{-huge cardinals.”};$$

(3) *there exists an ordinal α such that for all $\gamma > \alpha$, if*

$$j : \mathrm{HOD} \cap V_{\gamma+1} \to \mathrm{HOD} \cap V_{j(\gamma)+1}$$

is an elementary embedding with critical point above α, then $j \in \mathrm{HOD}$. □

If the HOD Conjecture is provable in ZFC then there is a striking corollary in ZF:

Theorem 20 (ZF) *Suppose that λ is a limit of supercompact cardinals and that there is an extendible cardinal below λ. Then there is no elementary embedding,*

$$j : V_{\lambda+2} \to V_{\lambda+2},$$

which is not the identity. □

This corollary would prove the conjecture above (that there are no weak Reinhardt cardinals) and would give an inconsistency result which more closely matches the version of Theorem 7 that Kunen actually proved (assuming the Axiom of Choice).

Theorem 21 (Kunen, [7]) *Suppose that λ is an ordinal. Then there is no elementary embedding $j : V_{\lambda+2} \to V_{\lambda+2}$ which is not the identity.* □

The connection between the HOD Conjecture and the inner model problem for one supercompact cardinal is illustrated by the next theorem. Arguably, statement (3) of this theorem would follow from any reasonable solution to the inner model problem for one supercompact cardinal.

Theorem 22 (ZFC) *Suppose that there is an extendible cardinal. Then the following are equivalent;*

(1) *The* HOD *Conjecture holds;*

(2) *there is a cardinal δ such that*

$$\mathrm{HOD} \models \text{“}\delta \text{ is a supercompact cardinal”}$$

and this is witnessed by the class of all $E \cap \mathrm{HOD}$ such that E is an extender and $E \cap \mathrm{HOD} \in \mathrm{HOD}$.

(3) *there exists a class $N \subseteq \mathrm{HOD}$ and there exists a cardinal δ such that*

$$N \models \mathrm{ZFC} + \text{“}\delta \text{ is a supercompact cardinal”}$$

and this is witnessed by the class of all $E \cap N$ such that E is an extender and $E \cap N \in N$. □

These developments come with a price. For the large cardinal axioms stronger than the axiom which asserts the existence of one supercompact cardinal, the Set Theorist's Cosmological Principle must either be abandoned or revised. The reason is that the solution to the inner model problem for the specific axiom "There exists one supercompact cardinal" necessarily will solve the inner model problem (as currently defined) for essentially *all* of the known large axioms up to and including the axiom "There is an ω-huge cardinal".

We therefore face a very simple dichotomy of possibilities: the inner model problem for the axiom "There exists one supercompact cardinal" is solvable or it is not. The other possibility—this is the possibility that this solvability question is itself unsolvable—is *not* an option based on any reasonable notion of mathematical truth. The simple reason is that if the solvability question is itself formally unsolvable then the inner model problem for the axiom "There exists one supercompact cardinal" is *not* solvable. The situation here is exactly like the situation for number (but not all) of the prominent open questions of modern mathematics. For example, if the "*Riemann Hypothesis*" is formally unsolvable (and there is absolutely no evidence for this), then the Riemann Hypothesis is *true*.

Whatever the outcome to this dichotomy of possibilities, one outcome seems certain: The Set Theorist's Cosmological Principle cannot be applied to argue for the truth of (any) large cardinal axioms beyond the axiom "There exists one supercompact cardinal". Of course, it could be that solution to the Inner Model Problem for the axiom "There exists one supercompact cardinal" involves the construction of an inner model which is not of the form $L[\tilde{E}]$ where $\tilde{E}$ is a sequence of partial extenders etc. But this alone would not suffice to resolve the issue raised by the preceding theorem. The reason is that statement (3) of the theorem makes no assumption that the inner model N is constructed from an extender sequence. Further, the necessity of the closeness of N to V (which is the only issue here) does not require that $N \subseteq \text{HOD}$ but only requires that for some cardinal δ,

$$N \models \text{ZFC} + \text{“}\delta \text{ is a supercompact cardinal”}.$$

And this is witnessed by the class of all $E \cap N$ such that E is an extender and $E \cap N \in N$. A solution which solves the inner model problem for the axiom "There exists one supercompact cardinal" and yet involves only the construction of inner models for which this fails, would be completely unlike all the current solutions to the inner model problem for the various large cardinal axioms where a solution exists.

One can correctly speculate that the difficulty is in the requirement that for the inner model N the relevant large cardinal axiom is witnessed to hold by extenders of N which are of the form $E \cap N$ for some extender $E \in V$. But extenders are the witnesses for large axioms and therefore any genuine construction of an inner model N should arguably satisfy this requirement. Moreover, to avoid "trivial" solutions, one has to require that the associated inner models satisfy some form of being close to V. More precisely, one has to require that the large cardinals of the inner models N which constitute the solution have some form of ancestry in the large cardinals of V.

There is a silver lining to this dark cloud. Suppose that the Inner Model Problem at the level of one supercompact cardinal can be solved and that the solution *does* involve defining inner models which are of the form $L[\tilde{E}]$ where $\tilde{E}$ is a sequence of (partial) extenders. Then it is possible to analyze the relationship of the inner models $L[\tilde{E}]$ to V without knowing how the corresponding extender sequences $\tilde{E}$ are actually constructed. The results to date in [18] have greatly clarified the axioms which assert the existence of ω-huge cardinals and revealed a new hierarchy of such axioms. The emerging structure theory for these axioms could well develop to the point where it serves as a surrogate for the existence of an inner model theory in a revised version of the Set Theorist's Cosmological Principle. On general grounds one can argue that if these axioms are consistent, then $L[\tilde{E}]$ must provide a structure theory for these axioms; because, if the axioms can hold in V, then they can hold in $L[\tilde{E}]$ (by the closeness of $L[\tilde{E}]$ to V). Therefore, the revision of the Set Theorist's Cosmological Principle that is required may actually not be so severe. The revision would be in what constitutes an inner model theory for those large cardinal axioms beyond the large cardinal axiom that asserts the existence of one supercompact cardinal. My point is that simply requiring that there be a generalization of the Axiom of Consructibility which is compatible with the large cardinal axiom may (for evident reasons) not be sufficient for these axioms. In addition, one may have to impose much stronger requirements, perhaps more in line with the often quoted speculation of Gödel (1947):

> *There might exist axioms so abundant in their verifiable consequences, shedding so much light upon a whole discipline, and furnishing such powerful methods for solving given problems (and even solving them, as far as possible, in a constructivistic way) that quite irrespective of their intrinsic necessity they would have to be assumed at least in the same sense as any established physical theory.*

To summarize the point I am attempting to make:

> *The foundational basis for asserting that large cardinal axioms beyond the level of one supercompact cardinal are "true" might lie in the structural consequences for $L[\tilde{E}]$ that their existence implies. Moreover, this claim of truth may require (and reinforce) some version of the claim that $V = L[\tilde{E}]$.*

This speculation is grounded in a number of preliminary results, [18].

1.8 Summary

There is now a body of mathematical evidence that if there is a supercompact cardinal then there is transcendent version, say L^{Ω}, of Gödel's inner model L: in brief, there is an ultimate L. This development *if realized* will yield a much deeper understanding of the large cardinal axioms:

1. Identifying much more precisely the transition for large cardinal axioms from the possible to the impossible; and,

2. providing a framework for a continuing evolution in the understanding of this transition.

The analysis will reveal some very subtle theorems about the nature of sets which in turn will eliminate essentially all the large cardinal axioms known to contradict the Axiom of Choice. How then could one account for the new predictions of consistency (and formal unsolvability) which will arise? Certainly not by invoking the Skeptic's Retreat.

Finally, we know Gödel rejected the axiom $V = L$. The current view rejects this axiom primarily because it is a *limiting* axiom which denies large cardinal axioms. This particular argument would *not apply* to the axiom $V = L^{\Omega}$. Further, ***assuming*** the analysis can be carried out to construct L^{Ω}:

> **There is no known candidate for a sentence which is independent from the axiom $V = L^{\Omega}$ and which is not a consequence of some large cardinal axiom.**

But all large cardinal axioms are merely axioms for the "height" of L^{Ω}, since no (known) large cardinal axiom can transcend L^{Ω}.

As this point is a key point of the thesis of this paper, I shall discuss it a bit further. It is well known that large cardinal axioms yield new theorems of number theory; more precisely, assuming the large cardinal axioms to be true one can infer as true specific statements of Number Theory *which arguably cannot otherwise be proved*. The foundational issue raised by this is that the large cardinals do not exist in the universe of number theory and yet their existence generates new truths of that universe. How then can the number theorist account for these truths? This of course is presented in the previous discussion as the debate between the Skeptic and the Set Theorist.

Why do not similar issues arise for the universe given by L^{Ω}? Because the large cardinals can (and therefore do) exist in this universe. This is the key new feature that L^{Ω} would possess which sets it apart from all the current known generalizations of L. As a consequence of this feature, L^{Ω} would provide an unambiguous conception of the transfinite universe, giving an example of an axiom which achieves this goal and which is arguably compatible with all large cardinal axioms where **there is no example currently known.**

Would this alone be sufficient to argue that the axiom $V = L^{\Omega}$ is the axiom for V? No more than one could argue that the axiom $V = L$ is the axiom for V should all the large cardinal axioms which imply that $V \neq L$ turn out to be inconsistent (which will not happen). But, the successful construction of L^{Ω}

would provide substantial evidence that there is a single axiom for V which yields a conception of the Universe of Sets which is in fact (for the reasons articulated above) *more* unambiguous then our present conception of number theory. Moreover, the successful construction of L^{Ω} would provide a starting point for discovering that axiom.

This development would be a significant milestone in our understanding of the Transfinite Universe. I make this claim completely independently of any speculation that there are number theoretic problems which are "orthogonal" to all large cardinal axioms, such as is the case for the problem of the Continuum Hypothesis.

Bibliography

[1] A. Beller, R. B. Jensen, and P. Welch. *Coding the universe*, volume No. 47 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge-New York, 1982.

[2] A. Dodd and R. Jensen. The core model. *Ann. Math. Logic*, 20(1):43–75, 1981.

[3] Kurt Gödel. Consistency-Proof for the Generalized Continuum-Hypothesis. *Proc. Nat. Acad. Sci. U.S.A*, 25:220–224, 1938.

[4] Kurt Gödel. *The Consistency of the Continuum Hypothesis*. Annals of Mathematics Studies, no. 3. Princeton University Press, Princeton, N. J., 1940.

[5] Akihiro Kanamori. *The higher infinite*. Perspectives in Mathematical Logic. Springer-Verlag, Berlin, 1994. Large cardinals in set theory from their beginnings.

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

[7] Kenneth Kunen. Elementary embeddings and infinitary combinatorics. *J. Symbolic Logic*, 36:407–413, 1971.

[8] Azriel Lévy. Definability in axiomatic set theory. I. In *Logic, Methodology and Philos. Sci. (Proc. 1964 Internat. Congr.)*, pages 127–151. North-Holland, Amsterdam, 1965.

[9] William J. Mitchell and John R. Steel. *Fine structure and iteration trees*. Springer-Verlag, Berlin, 1994.

[10] Jan Mycielski and H. Steinhaus. A mathematical axiom contradicting the axiom of choice. *Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys.*, 10:1–3, 1962.

[11] Itay Neeman. Inner models in the region of a Woodin limit of Woodin cardinals. *Ann. Pure Appl. Logic*, 116(1-3):67–155, 2002.

[12] William N. Reinhardt. Ackermann's set theory equals ZF. *Ann. Math. Logic*, 2(2):189–249, 1970.

[13] Dana Scott. Measurable cardinals and constructible sets. *Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys.*, 9:521–524, 1961.

[14] S. Shelah. *Around classification theory of models*, volume 1182 of *Lecture Notes in Mathematics*. Springer–Verlag, 1986.

[15] John R. Steel. *The core model iterability problem.* Springer-Verlag, Berlin, 1996.

[16] W. Hugh Woodin. The universe constructed from a sequence of ordinals. *Arch. Math. Logic*, 35(5-6):371–383, 1996.

[17] W. Hugh Woodin. *The Axiom of Determinacy, forcing axioms, and the nonstationary ideal.* Walter de Gruyter & Co., Berlin, 1999.

[18] W. Hugh Woodin. Suitable extender sequences. *Preprint*, pages 1–676, 2009.

[19] Ernst Zermelo. Über Grenzzahlen und Mengenbereiche: Neue Untersuchungenuberdie Grundlagen der Mengenlehre. *Fundamenta Mathematicae*, 16:29–47, 1930.