PART V: HOD AND ITS LOCAL VERSIONS

ORDINAL DEFINABILITY IN MODELS OF DETERMINACY INTRODUCTION TO PART V

JOHN R. STEEL

Five of the papers in this section were written in 1977–1983, and published in the original Cabal Seminar volumes. The other three papers describe work done 1993–2007. ([Ste16B] was written in 1996, and is published for the first time here, while [Nee16] and [SW16] have just been written.) Although the three new papers are separated from the older ones by as much as 30 years, they build on themes first developed in the older papers. All the papers are concerned, in one way or another, with ordinal definability in models of determinacy.

The earlier papers focus on restricted fragments of **HOD**, perhaps because the basic tools needed to analyze the full **HOD** in a determinacy model were not available then. The most important of those tools are definable scales and canonical inner models for large cardinal hypotheses. The advances we have made in constructing definable scales and canonical inner models since then have made it more reasonable to attempt an analysis of the full $\mathbf{HOD}^M$, for M a model of AD^+. Such an analysis is important, because $\mathbf{HOD}^M$ is a model of ZFC that, in a sense, has the same information as M.

A fine-structural analysis of $\mathbf{HOD}^M$ in the case $M = \mathbf{L}(\mathbb{R})$ is presented in [SW16], the last paper in our group. In this case, **HOD** is indeed the *higher-level analog of* **L** envisaged in our earliest papers. This analysis has been carried significantly further in [Sar09][1], but it is not known how to do it in general. Whether **HOD** satisfies the GCH is a central test question, identified in our earliest papers.

We begin by describing some of the history that set the stage for these papers. We then discuss the individual papers, focusing mainly on putting the older papers in a modern context.

§1. Some history.

1.1. Definitions and constructions. Effective descriptive set theory is the marriage of descriptive set theory with recursion theory. Perhaps the first theorems of the subject are Kleene's results:

Partially supported by National Science Foundation Grant No. DMS-0855692.

[1] For M a bit past the minimal model of $\mathsf{AD}_{\mathbb{R}}$ + "Θ is regular".

Ordinal Definability and Recursion Theory: The Cabal Seminar, Volume III
Edited by A. S. Kechris, B. Löwe, J. R. Steel
Lecture Notes in Logic, 43

THEOREM 1.1 (Kleene [Kle55A, Kle55B, Kle55C]).
(1) For any real x, x is Δ^1_1 if and only if x is hyperarithmetic.
(2) Every nonempty Σ^1_1 set of reals has a member that is recursive in $\mathcal{O}$.

Here and below, real is an infinite sequence of natural numbers, considered as an element of the Baire space. We call a set of reals **thin** just in case it has no nonempty perfect subsets. J. Harrison strengthened the first of Kleene's theorems by showing:

THEOREM 1.2 (Harrison [Har66]). Every thin Σ^1_1 set of reals can be enumerated by a hyperarithmetic real.

Such theorems, and others like them, are consequences of a constructive analysis of some level of second order truth. In order to extend them to higher levels of truth, for example Σ^1_2 and beyond, we must use more general constructions, and for this, the basic notions of recursion theory are no longer adequate. We must move upward into *inner model theory*, and what was effective descriptive set theory becomes *descriptive inner model theory*.[2]

The theorems of Kleene and Harrison can be seen in this light, when we recall that a real is hyperarithmetic iff it belongs to $\mathbf{L}_{\omega_1^{\mathrm{CK}}}$. Moving upward, to Σ^1_2 and full $\mathbf{L}$, we have

THEOREM 1.3.
(1) For any real x, x is Δ^1_2 if and only if $x \in L_\alpha$, where α is least such that $\mathbf{L}_\alpha$ is a Σ_1 elementary submodel of $\mathbf{L}$. (Shoenfield [Sho61].)
(2) For any real x, x is Δ^1_2 in a countable ordinal iff $x \in \mathbf{L}$. (Solovay; see [KM72].)
(3) Every nonempty Σ^1_2 set of reals has a Δ^1_2 member. (Shoenfield [Sho61].)
(4) Every thin Σ^1_2 set of reals is contained in $\mathbf{L}$. (Mansfield [Man70], Solovay [Sol66].)

1.2. Definable scales. Underlying these theorems, and pretty much everything else in pure descriptive set theory, is the construction of definable Suslin representations (equivalently, scales). We have

THEOREM 1.4.
(1) A set B of reals is Σ^1_1 iff $B = \mathrm{p}[T]$, for some recursive tree T on $\omega \times \omega$. (Kleene [Kle55C].)
(2) A set B of reals is Σ^1_2 iff $B = \mathrm{p}[T]$, for some tree T on $\omega \times \omega_1^{\mathbf{V}}$ such that T is Σ_1-definable over $\mathbf{L}$. (Shoenfield [Sho61].)

In order fully generalize the theorems of Kleene and Harrison to Σ^1_3 and beyond, we must reach Suslin representations for Σ^1_3 sets of reals, and beyond, via some constructive process. This requires axioms beyond those of ZFC.

[2]The term is due to Grigor Sargsyan.

The game method for obtaining optimally definable Suslin representations makes direct use of determinacy. It was discovered in 1971 by Moschovakis, who used it to extend the Novikoff–Kondô–Addison theorems to the higher levels of the projective hierarchy. (The original paper is [Mos71]; see also [KM78B] and [Mos09, Chapter 6].) As part of this work, Moschovakis introduced the basic notion of a *scale*, which we now describe.

Let T be a tree on $\omega \times \lambda$, and $A = \mathrm{p}[T]$. One can get a "small" subtree of T which still projects to A by considering only ordinals $< \lambda$ which appear in some leftmost branch. The **scale of** T does this, then records the resulting subtree as a sequence of **norms**, *i.e.*, ordinal-valued functions, on A. More precisely, for $x \in A$ and $n < \omega$, put

$$\varphi_n(x) = |\langle \ell_x(0), \dots, \ell_x(n)\rangle|_{\mathrm{lex}},$$

where for $u \in \lambda^{n+1}$, $|u|_{\mathrm{lex}}$ is the ordinal rank of u in the lexicographic order on λ^n. Then

$$\vec{\varphi} = \langle \varphi_n \, : \, n < \omega\rangle$$

is the scale of T. It has the properties:

(a) Suppose that $x_i \in A$ for all $i < \omega$, and $x_i \to x$ as $i \to \infty$, and for all n, $\varphi_n(x_i)$ is eventually constant as $n \to \infty$, then
  (i) (limit property) $x \in A$, and
  (ii) (lower semi-continuity) for all n, $\varphi_n(x) \leq$ the eventual value of $\varphi_n(x_i)$ as $i \to \infty$.

(b) (refinement property) if $x, y \in A$ and $\varphi_n(x) < \varphi_n(y)$, then $\varphi_m(x) < \varphi_m(y)$ for all $m > n$.

A sequence of norms on A with property (a) is called a **scale on** A. Any scale on A can be easily transformed into a scale on A with the refinement property. If $\vec{\varphi}$ is a scale on A, then we define the tree of $\vec{\varphi}$ to be

$$T_{\vec{\varphi}} = \big\{(\langle x(0), \dots, x(n-1)\rangle, \langle \varphi_0(x), \dots, \varphi_{n-1}(x)\rangle) \, : \, n < \omega \text{ and } x \in A\big\}.$$

It is not hard to see that $\mathrm{p}[T_{\vec{\varphi}}] = A$. If $\vec{\varphi}$ has the refinement property, and $\vec{\psi}$ is the scale of $T_{\vec{\varphi}}$, then $\vec{\psi}$ is equivalent to $\vec{\varphi}$, in the sense that for all n, x and y, $\psi_n(x) \leq \psi_n(y)$ iff $\varphi_n(x) \leq \varphi_n(y)$. The reader should see [KM78B, 6B] and [Jac08] for more on the relationship between scales and Suslin representations.

There are at least two benefits to considering the scale of a tree: first, it becomes easier to state and prove optimal definability results, and second, the construction of Suslin representations using comparison games becomes clearer. Concerning definability, we have:

Definition 1.5. Let Γ be a pointclass, and $\vec{\varphi}$ a scale on A, where $A \in \Gamma$; then we call $\vec{\varphi}$ a **Γ-scale on** A just in case the relations

$$R(n, x, y) \Leftrightarrow x \in A \wedge (y \notin A \vee \varphi_n(x) \leq \varphi_n(y)),$$

and

$$S(n, x, y) \Leftrightarrow x \in A \wedge (y \notin A \vee \varphi_n(x) < \varphi(y))$$

are each in Γ. We say Γ **has the scale property** just in case every set in Γ admits a Γ-scale, and write $\mathsf{Scale}(\Gamma)$ in this case.

Moschovakis showed that if Γ is a pointclass which is closed under universal real quantification, has other mild closure properties, and has the scale property, then every Γ relation has a Γ uniformization, and the Γ singletons are a basis for Γ [KM78B, 3A-1]. He also showed

THEOREM 1.6 (Moschovakis 1971). Let $n < \omega$, and assume $\boldsymbol{\Delta}^1_{2n}$-determinacy; then the pointclasses Π^1_{2n+1} and Σ^1_{2n+2} have the scale property.

(See [KM78B, 3B, 3C].) From this, one gets the natural generalization of Novikoff–Kondô–Addison to the higher levels of the projective hierarchy.

Around 1977, Moschovakis showed that the pointclass $\Sigma_1^{\mathbf{L}_\kappa(\mathbb{R})}$ has the scale property, where κ is least such that $\mathbf{L}_\kappa(\mathbb{R}) \models \mathsf{KP}$.[3] In late 1979, Moschovakis showed that sets in the dual pointclass admit definable scales as well, and Martin and Steel very quickly extended his result to:

THEOREM 1.7 (Martin, Steel 1979). Assume $\mathsf{AD}^{\mathbf{L}(\mathbb{R})}$; then

(a) The pointclass $(\Sigma^2_1)^{\mathbf{L}(\mathbb{R})}$ has the scale property.
(b) In $\mathbf{L}(\mathbb{R})$, the sets of reals admitting scales are precisely the $\boldsymbol{\Sigma}^2_1$ sets.

See [Mos78, Mos83, MS83]. Part (b) makes use of an earlier proof by Kechris and Solovay that in $\mathbf{L}(\mathbb{R})$, every set admitting a scale of any complexity is $\boldsymbol{\Sigma}^2_1$.[4] In 1980, Martin [Mar83] found a much subtler limitation on the definability of scales, and Steel [Ste83A] knitted the work just described into a thorough description of the pattern of pointclasses with the scale property in $\mathbf{L}(\mathbb{R})$.

As Solovay observed some time in the 1970s, in $\mathbf{L}(\mathbb{R})$, every nonempty Σ^2_1 collection of sets of reals has a $\boldsymbol{\Delta}^2_1$ member. (Applying Theorem 1.7, we get a lightface Δ^2_1 member.) This fact is often used in the following form:

COROLLARY 1.8 (Solovay, 1970s). Assume $\mathsf{AD} + \mathbf{V}{=}\mathbf{L}(\mathbb{R})$. Let $\mathcal{B} \subseteq \wp(\mathbb{R})$ be Π_1-definable from a real, and suppose that every Suslin-co-Suslin sets of reals is in $\mathcal{B}$; then every set of reals is in $\mathcal{B}$.

Theorem 1.7 lets us calibrate precisely the correctness of $M = \mathbf{HOD}^{\mathbf{L}(\mathbb{R})}$. The key is that if T is the tree of a $\Sigma_1^{\mathbf{L}(\mathbb{R})}$ scale on a universal $\Sigma_1^{\mathbf{L}(\mathbb{R})}$ set of reals, then $T \in M$. Using that, we get

THEOREM 1.9. Assume $\mathsf{AD} + \mathbf{V}{=}\mathbf{L}(\mathbb{R})$; then

[3] Here and below, our Σ_1 formulae are always allowed a name for $\mathbb{R}$ itself.
[4] In $\mathbf{L}(\mathbb{R})$, a set of reals is Σ^2_1 just in case it is Σ_1.

(a) There is an α such that $\mathbf{L}_\alpha(\mathbb{R})^{\mathbf{HOD}} \prec_{\Sigma_1} \mathbf{L}(\mathbb{R})$.
(b) $\mathbf{HOD} \vDash$ "Every real is $\mathbf{OD}^{\mathbf{L}(\mathbb{R})}$".
(c) $\mathbf{HOD} \vDash$ "There is a $\Delta_1^{\mathbf{L}(\mathbb{R})}$ wellorder of $\mathbb{R}$ of order type ω_1".

Thus CH holds in $\mathbf{HOD}^{\mathbf{L}(\mathbb{R})}$.

Much of the work in the 1977–83 papers in this block was done before the late 1979–early 1980 work on the existence of scales in $\mathbf{L}(\mathbb{R})$. In any case, the scale-existence work was quite recent. It was therefore natural for the authors of those papers to focus strongly on projectively definable sets. Without definable scales, it's hard to get very far. The full $\mathbf{HOD}^{\mathbf{L}(\mathbb{R})}$ was not an object of concern in the 1977–83 papers.

1.3. ∞-Borel codes. An ∞-Borel code for a set $A \subseteq \mathbb{R}$ is a triple $\langle \alpha, S, \varphi \rangle$ such that $\alpha \in \mathrm{Ord}$, $S \subseteq \alpha$, and $A = \{x : \mathbf{L}_\alpha[S, x] \vDash \varphi[x]\}$. Any such A is built up from open sets via wellordered unions and complementation, and we can always take S to be a set of ordinals recording the history of how that was done. Every Suslin representation yields an ∞-Borel code, but not conversely.[5] Clearly, if M is a proper class having an ∞-Borel code for A, then for all reals x, $A \cap M[x]$ is definable over $M[x]$ from this ∞-Borel code, uniformly in x. However, in contrast to Suslin representations, it can happen that A is nonempty, but $A \cap M = \varnothing$.

A still weaker notion is **term capturing**. If M is a transitive model of a reasonable fragment of ZFC, $\mathbb{P}$ is a poset in M, and $A \subseteq \mathbb{R}$, then we say M captures A at $\mathbb{P}$ just in case there is a term $\tau \in M$ such that whenever g is M-generic over $\mathbb{P}$, then $\tau_g = A \cap M[g]$. If M has an ∞-Borel code for A, then it captures A at all $\mathbb{P} \in M$. The converse is false: any transitive model M of ZFC closed under sharps captures $\{\langle x, n \rangle : x \in \mathbb{R} \wedge n \in x^\#\}$, but if M is countable, it only has ∞-Borel codes for Borel sets. If M captures A at $\mathbb{P}$, and every real is $\mathbb{P}$-generic over M, then M has an ∞-Borel code for A. (Any $S \subseteq \mathrm{Ord}$ coding the regular open algebra of $\mathbb{P}$ will do.) Amazingly, there are important examples of M and $\mathbb{P}$ such that every real is $\mathbb{P}$-generic over M. One is $M = \mathbf{HOD}$ and $\mathbb{P}$ the Vopěnka algebra, and another is M a fully backgrounded extender model, and $\mathbb{P}$ the extender algebra at a Woodin cardinal of M. (See [Ste10B, Section 7].)

It follows from Theorem 1.7 that in $\mathbf{L}(\mathbb{R})$, every $\underset{\sim}{\boldsymbol{\Sigma}}^2_1$ set of reals is ∞-Borel.[6] Woodin showed in 1981

THEOREM 1.10 (Woodin 1981). Assume AD + $\mathbf{V}$=$\mathbf{L}(\mathbb{R})$; then

(1) Every set of reals is ∞-Borel.
(2) Every $\mathbf{OD}$ set of reals has an ∞-Borel code in $\mathbf{HOD}$.
(3) For all reals x, $\mathbf{HOD}[x] = \mathbf{HOD}_x$; that is, $\mathbf{HOD}$ relativizes by adjunction.

[5] The notion of an ∞-Borel code is introduced and used in [KM78B, Section 10], where it is called an **∞-Boolean** code. We do not know whether it appears anywhere else earlier.

[6] We assume $\mathrm{AD}^{\mathbf{L}(\mathbb{R})}$ here.

With regard to part (3), note that every real is generic over **HOD**, by Vopěnka's theorem. One proof of (2) and (3) goes by showing that **HOD** captures all **OD** sets of reals at the Vopěnka algebra. At about the same time, Kechris and Woodin had shown

THEOREM 1.11 (Kechris, Woodin 1981). Assume $\mathsf{AD} + \mathbf{V}{=}\mathbf{L}(\mathbb{R})$; then

(1) there is a partial order $\mathbb{P} \in \mathbf{HOD}$ and a recursive function t such that

$$\mathbf{V} \models \varphi[x] \Leftrightarrow \mathbf{HOD} \models 1 \Vdash t(\varphi)(\check{x}),$$

for all formulae $\varphi(v)$ and all $x \in \mathbf{HOD}$,

(2) $\mathbf{HOD} = \mathbf{L}[P]$, for some $P \subseteq \Theta$.

Here $\mathbb{P}$ is a variant of the Vopěnka poset. See [Ste08C, Section 3] for the proof of an extension of this theorem. According to the theorem, **HOD** can determine truth in the full **V**, by consulting the forcing relation for $\mathbb{P}$.

Thus by 1982, the degree of correctness of $\mathbf{HOD}^{\mathbf{L}(\mathbb{R})}$, as well as what sets of reals it has ∞-Borel codes for, was known. But again, this work either post-dated the 1977–1983 papers, or was quite recent. Those papers focused lower down, below the full $\mathbf{HOD}^{\mathbf{L}(\mathbb{R})}$.

1.4. Definability theory under determinacy assumptions. Much of the descriptive set theory of the first two levels of the projective hierarchy can be generalized to higher levels, and beyond, without any use of the fine hierarchy provided by inner model theory. Determinacy is enough to go pretty far; one need not have a proof of it from large cardinals, or reach the strategies one is assuming to exist in some construction. For example, Moschovakis' scale-existence results yield at once a basis theorem at the even levels of the projective hierarchy:

THEOREM 1.12 (Moschovakis). Assume $\utilde{\mathbf{\Delta}}^1_{2n}$-determinacy, where $0 \leq n < \omega$; then every Σ^1_{2n+2} set of reals has a Δ^1_{2n+2} member.

The parallel generalization of the Kleene basis theorem for Σ^1_1 to all Σ^1_{2n+1} is more subtle. It was discovered by Martin and Solovay in 1975, and is the beginning of "Q-theory". We shall discuss it when we come to the paper on Q-theory in this block.

There are partial generalizations of the the thin-set theorems for Σ^1_1 and Σ^1_2 as well. At the odd levels, we need Moschovakis' "third periodicity" theorem, on the existence of definable winning strategies for games with definably scaled payoff sets. (See [Mos09, 6E].) With this in hand, we get

THEOREM 1.13. Let $0 \leq n < \omega$, and assume $\utilde{\mathbf{\Delta}}^1_{2n}$-determinacy; then

(1) Every thin Σ^1_{2n+1} set of reals is enumerated by a Δ^1_{2n+1} real.

(2) There is a largest thin Π^1_{2n+1} set, which we call C_{2n+1}.

(3) Letting $\mathrm{C}_{2n+2} = \{x : \exists y \in \mathrm{C}_{2n+1}(x \leq_{\mathrm{T}} y)\}$, we have that every thin Σ^1_{2n+2} set of reals is contained in C_{2n+2}.

(4) If $\underset{\sim}{\mathbf{\Pi}}^1_{2n+1}$-determinacy holds, then C_{2n+2} is the largest countable Σ^1_{2n+2} set of reals.

Item (1) is due to Kechris. Item (2) is due to Kechris, Guaspari, Sacks for $n = 0$, Kechris for $n > 0$. See [Kec75B]. Items (3) and (4) are joint work of Kechris and Moschovakis from the early '70s; see [KM72]. Kechris developed the theory of countable analytical sets in [Kec75A], and [Kec78B]. Combining his work with 1977 work of Harrington and Kechris [HK81], we have

Theorem 1.14 (Kechris, Harrington–Kechris). *Assume PD, let $1 \leq n < \omega$; and let T be the tree of a Π^1_{2n-1}-scale on a universal Π^1_{2n-1} set. For any real x, the following are equivalent:*

(1) $x \in C_{2n}$,
(2) $x \in \mathbf{L}[T]$,
(3) x is Δ^1_{2n} in some ordinal $< \underset{\sim}{\boldsymbol{\delta}}^1_{2n-1}$,
(4) x is Δ^1_{2n} in a countable ordinal.

The equivalence of (1),(3), and (4) is due to Kechris (see [Kec75A]). Moschovakis had shown that if $n = 1$, then $T \in \mathbf{L}$, and the $n = 1$ case of Theorem 1.14 follows from that and Solovay's theorem cited above. (See [KM78B, Section 9].) For $n = 2$, the equivalence of (1) and (2) was first proved by Kechris and Martin in [KM16].

Kechris also showed that C_{2n+1} is prewellordered by Δ^1_{2n+1}-degree, a structural feature suggestive of the master code levels of canonical inner models. Indeed, Guaspari, Kechris, and Sacks showed C_1 consists of those reals having the same Δ^1_1-degree as a master code in $\mathbf{L}$.[7] C_2 is the set of reals in $\mathbf{L}$. Moschovakis showed that for all n, C_{2n+2} is the set of reals in $\mathbf{L}[C_{2n+2}]$. These results suggested that there would be *higher-level analogs of* $\mathbf{L}$, canonical models related to Σ^1_{2n} the way $\mathbf{L}$ is related to Σ^1_2.

Theorem 1.13(1) implies there is no largest countable Σ^1_{2n+1} set, and using Theorem 1.12, it is easy to see that there is no largest countable Π^1_{2n} set. There are largest countable sets in many pointclasses beyond the projective hierarchy, in fact:

Theorem 1.15 (Kechris, Moschovakis). *Suppose Γ is adequate, ω-parametrized, has the scale property, and is closed under $\exists^{\mathbb{R}}$, and suppose all $\underset{\sim}{\boldsymbol{\Gamma}}$ games are determined; then there is a largest countable Γ set of reals.*

The theorem is implicit in the proof of [KM78B, Theorem 11B-2], which proves the existence of C_{2n}. When it exists, the largest countable Γ set is called C_Γ. It is natural to look at the largest pointclass $\Gamma \subseteq \mathbf{L}(\mathbb{R})$ satisfying the

[7] Equivalently, $x \in C_1$ iff for some α, x has the same Δ^1_1 degree as the first order theory (without parameters) of $\mathbf{L}_\alpha$.

hypothesis of Theorem 1.15, namely $\Sigma_1^{\mathbf{L}(\mathbb{R})}$. In this case, the closure properties of $\Sigma_1^{\mathbf{L}(\mathbb{R})}$ make the counterpart of Theorem 1.14 an easy exercise:

Theorem 1.16. Assume $\mathsf{AD}^{\mathbf{L}(\mathbb{R})}$, and let T be the tree of a $\Sigma_1^{\mathbf{L}(\mathbb{R})}$-scale on a universal $\Sigma_1^{\mathbf{L}(\mathbb{R})}$ set. For any real x, the following are equivalent:

(1) $x \in C_\Gamma$, for $\Gamma = \Sigma_1^{\mathbf{L}(\mathbb{R})}$,
(2) $x \in \mathbf{L}[T]$,
(3) x is ordinal definable in $\mathbf{L}(\mathbb{R})$,
(4) x is $\Delta_1^{\mathbf{L}(\mathbb{R})}$ in a countable ordinal.

1.5. Inner model theory. In order to fully generalize the theorems of Kleene, Harrison, and Mansfield–Solovay, one needs to do more than define Suslin representations or ∞-Borel codes for more complicated sets of reals. One must *construct* such representations. The Shoenfield tree is not just definable, it is in $\mathbf{L}$. Definable Suslin representations and ∞-Borel codes let us pack the information in a set of reals into a set of ordinals, but there remains the problem of reaching that set of ordinals via a constructive process, that is, in a canonical inner model. Unfortunately, when the papers in our first block were written, it was not known how to construct canonical inner models significantly more correct than $\mathbf{L}$.

By 1977, the work of Silver, Kunen, and Mitchell had extended the basic theory of $\mathbf{L}$ to canonical inner models with many measurable cardinals. (See [Sil71A, Sil71B, Kun70, Mit74].) Although this work is full of beautiful ideas and powerful tools of permanent value, it does not go very far beyond G'odel's work on $\mathbf{L}$ in descriptive set theoretic terms. The canonical inner models studied in [Sil71A, Kun70, Mit74] fail to be even Σ_3^1 correct. Indeed, the set of reals in the union of these models is a countable Σ_3^1 set, and so if $\underset{\sim}{\mathbf{\Delta}}_2^1$-determinacy holds, it is enumerated by a single Δ_3^1 real. A related fact is that each of the models satisfies "there is a Δ_3^1 wellorder of the reals", and therefore Δ_2^1-determinacy fails to hold in any of them.

In an important 1978 breakthrough, Mitchell discovered the central features of the first order form of canonical inner models, all the way up to inner models with superstrong cardinals. (See [Mit79].) His idea of models constructed from *coherent sequences of extenders* has been the basis of all further work in inner model theory. Mitchell himself used it to construct inner models satisfying "there is a κ that is $(\kappa+2)$-strong", a large cardinal hypothesis significantly beyond the existence of measurables, and seemingly well on the way to the existence of superstrongs, and even supercompacts.

Theorem 1.17 (Mitchell, 1978). Suppose there is a cardinal κ such that κ is $(\kappa+3)$-strong; then there is a model M constructed from a coherent sequence of extenders such that

(1) $M \models \mathsf{ZFC} + \exists\kappa(\kappa \text{ is } \kappa+2\text{-strong})$,