

THE STRUCTURE OF MODELS OF SECOND-ORDER SET THEORIES

by

KAMERYN J WILLIAMS

A dissertation submitted to the Graduate Faculty in Mathematics in partial fulfillment of the requirements
for the degree of Doctor of Philosophy, The City University of New York.

2018

THE STRUCTURE OF MODELS OF SECOND-ORDER SET THEORIES

by

KAMERYN J WILLIAMS

This manuscript has been read and accepted by the Graduate Faculty in Mathematics in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

Professor Joel David Hamkins

 Date

 Chair of Examining Committee

Professor Ara Basmajian

 Date

 Executive Officer

Professor Joel David Hamkins
Professor Arthur Apter
Professor Gunter Fuchs
 Supervisory Committee

Abstract
 THE STRUCTURE OF MODELS OF SECOND-ORDER SET THEORIES
 by
 KAMERYN J WILLIAMS

Advisor: Professor Joel David Hamkins

This dissertation is a contribution to the project of second-order set theory, which has seen a revival in recent years. The approach is to understand second-order set theory by studying the structure of models of second-order set theories. The main results are the following, organized by chapter. First, I investigate the poset of T -realizations of a fixed countable model of ZFC , where T is a reasonable second-order set theory such as GBC or KM , showing that it has a rich structure. In particular, every countable partial order embeds into this structure. Moreover, we can arrange so that these embeddings preserve the existence/nonexistence of upper bounds, at least for finite partial orders. Second I generalize some constructions of Marek and Mostowski from KM to weaker theories. They showed that every model of KM plus the Class Collection schema “unrolls” to a model of ZFC^- with a largest cardinal. I calculate the theories of the unrolling for a variety of second-order set theories, going as weak as $GBC + ETR$. I also show that being T -realizable goes down to submodels for a broad selection of second-order set theories T . Third, I show that there is a hierarchy of transfinite recursion principles ranging in strength from GBC to KM . This hierarchy is ordered first by the complexity of the properties allowed in the recursions and second by the allowed heights of the recursions. Fourth, I investigate the question of which second-order set theories have least models. I show that strong theories—such as KM or $\Pi_1^1\text{-CA}$ —do not have least transitive models while weaker theories—from GBC to $GBC + ETR_{\text{Ord}}$ —do have least transitive models.

Dedication

For Margaret and Martin.

Acknowledgments

First and foremost I must thank Joel. You have been a fantastic advisor throughout my studies and there is no doubt that my growth as a mathematician owes much to your support and guidance. Thank you for the many fruitful conversations, for thoroughly reading through this text, and for the myriad helpful suggestions along the way.

Thank you Arthur and Gunter for serving on my dissertation committee and for reading and giving me comments on this document. Thank you Roman for being on my committee in spirit and for organizing the models of arithmetic group at CUNY. Thank you Kaethe, Miha, Corey, Alex, Ryan, and Eoin. It was great to have other students interested in set theory and to have a student-run seminar in the subject for all of my years at CUNY. Thank you Vika for many conversations and for your hard work in making things run smoothly for the set theory group. And thank you to all the others who make CUNY's logic community so vibrant and such a productive place to be.

Thank you Herbie. You are my best friend and I love you. Thank you for moving across the country with me and for supporting and putting up with me through the joys and trials of graduate school. And finally, thank you mom, dad, Chris, Kalen, Christian, and Lexi.

Contents

Contents	vii
List of Figures	viii
0 Introduction	1
1 A first look at models of second-order set theories	3
1.1 Dramatis Personae	3
1.1.1 Second-order set theories	4
1.1.2 Doing without Powerset	8
1.1.3 Models of second-order set theory	8
1.2 Verifying the axioms	10
1.3 Preserving the axioms	12
1.4 Some basic constructions	16
1.5 GBC-realizations of a countable model	21
2 Many constructions: unrollings, cutting offs, and second-order L	31
2.1 The unrolling construction	33
2.1.1 Unrolling in GBC + ETR	42
2.1.2 How hard is it to unroll?	45
2.2 The cutting off construction	46
2.3 The constructible universe in the classes	49
2.4 The smallest heights of transitive models and β -models	54
3 Truth and transfinite recursion	57
3.1 Truth and iterated truth	59
3.2 Iterated truth and constructibility	62
3.3 Separating levels of ETR	64
3.4 A detour through the finite realm	67
3.5 Non-elementary transfinite recursions	72
4 Least models	77
4.1 β -models and well-founded classes	80
4.2 Admissible covers and nonstandard compactness arguments	81
4.3 Strong theories	86
4.4 Medium theories	91
4.5 Weak theories	96
4.6 Coda: the analogy to second-order arithmetic	97

List of Figures

1.1	Some classes of models of second-order set theory.	9
2.1	A membership code representing the set $a = \{0, 2\}$	34
2.2	To construct a membership code for $\bigcup X$ we remove all edges ending at the top node and add edges (n_Z, t_X) for all $n_Z \dashv\circ n_Y \dashv\circ t_X$	37
2.3	A partial picture of W . The nodes x, y are from elts A with $x < y$. Below x and y the membership code W looks like A	38
2.4	A partial picture of F^*	39
2.5	The least heights of transitive and β -models of fragments of KM. The ordinals in the picture are ordered by \ll	55
3.1	Transfinite recursion, from GBC to KM. Ordered by consistency strength.	58
3.2	The tree of inductive iterated full satisfaction classes over a model of arithmetic. The node at depth m has continuum many extensions to leaf nodes of depth k	72
4.1	An illustration of Friedman's theorem. The dashed line represents that N is ill-founded at A	82
4.2	A picture of the desired generalization of Friedman's theorem to the ill-founded realm.	84
4.3	Friedman's theorem gives $N \models S$ with $V_\kappa^N = M = V_\kappa^W$ and N ill-founded at A	87
4.4	The (GBC + ETR)-realizations for M form a disjoint collection of cones, each one with a basis element at the bottom.	93
4.5	Arithmetic versus set theory.	97

Chapter 0

Introduction

Eine Vielheit kann nämlich so beschaffen sein, daß die Annahme eines “Zusammenseins” *aller* ihrer Elemente auf einen Widerspruch führt, so daß es unmöglich ist, die Vielheit als eine Einheit, als “ein fertiges Ding” aufzufassen. Solche Vielheiten nenne ich *absolut unendliche* oder *inconsistente Vielheiten*.

Georg Cantor

The distinction between set and class can be traced back to Cantor. He distinguished sets from those multiplicities he termed *absolutely infinite* or *inconsistent*. As the name suggests, inconsistent multiplicities are those which lead to contradiction if taken as a set. For example, if we assume that the collection of ordinals is a set then we can derive the Burali–Forti paradox.

From a more modern perspective based upon the iterative conception of set we have a clear distinction between sets and classes. A collection is a set if it appears at some stage α in the cumulative hierarchy, while a (proper) class consists of elements unbounded in rank. So sets can be elements of other collections, while classes can never be elements.

In ordinary set theoretic practice in the early twenty-first century classes are treated as mere syntactic sugar; for example, $x \in \text{Ord}$ is an abbreviation for the formula expressing that x is a transitive set linearly ordered by \in . But set theory can also be formalized with classes as actual objects, rather than relegating them to a metatheoretic role.

The first axiomatization of second-order set theory—set theory with both sets and classes, sometimes called class theory—is due to von Neumann [Neu25]. His system survives into the modern day as GBC, allowing only predicative definitions of classes.¹ But GBC is not the only well-studied axiomatization of second-order set theory. The other major axiomatization KM, which allows impredicative comprehension, was independently proposed by multiple logicians—among them Morse, Quine, and Tarski. See the appendix to [Kel75] for a popularization of this system.

In the decades following von Neumann’s axiomatization, second-order axiomatizations of set theory saw significant use among set theorists. Perhaps most notably, Gödel’s original presentation [Göd38] of his relative consistency proof for the axiom of choice was in terms of von Neumann’s system. But over time the use of second-order systems waned, with the first-order system ZFC becoming the de facto standard.

¹But note that the contemporary GBC is quite different from von Neumann’s original axiomatization. Most strikingly, von Neumann’s system did not use sets and classes but rather what he termed I-objects and II-objects. I-objects are the sets whereas II-objects are not classes but rather functions (possibly class-sized), with I-II-objects being the set-sized functions. The modern formulation in terms of sets and classes, originally due to Bernays, is much more convenient to work with.

In recent years, however, second-order set theory has enjoyed a revived interest. Several mathematicians independently arrived at second-order set theory as the natural arena in which to pursue certain projects. It has seen use in work on the foundations of class forcing [Ant15; HKLNS16; HKSb; HKSa; GHHSW17], hyperclass forcing [AF], formalizing the inner model hypothesis [ABF], determinacy for class games [GH17; Hac16], and in truth theoretic work [Fuj12].

One fact that has emerged is that GBC and KM are not the only interesting second-order set theories. For some applications, KM is not quite strong enough so we need to extend to a stronger system. And for other applications GBC is too weak while KM is overkill, so we want to study natural intermediate theories.

This dissertation is a contribution towards this project of second-order set theory. Rather than apply the tools of second-order set theory to some domain, the aim is to study second-order set theories themselves. My approach will be model theoretic, aiming to understand these theories by understanding their models. A better knowledge of the foundations of second-order set theory will then facilitate applications thereof.

The layout of this dissertation is as follows.

Chapter 1 begins with the main theories we will consider, as well as several important classes of models. This is followed by a discussion of how to check whether a collection of classes gives a model of such and such theory, and when the axioms are preserved by forcing. To conclude the chapter, I show that the partial order consisting of GBC -realizations of a fixed countable model of ZFC has a rich structure. Much of the material in this chapter is already known, but I include it for the sake of giving a complete presentation.

Chapter 2 is dedicated to three constructions, which were originally studied in the context of models of KM by Marek and Mostowski [Mar73; MM75]. The first of these constructions, which I call the unrolling construction, takes a model of KM (plus Class Collection) and gives a model of ZFC^- with a largest cardinal, which is inaccessible. The second construction, I call it the cutting-off construction, takes a model of ZFC^- with a largest cardinal and gives a model of second-order set theory. Together, these two constructions show that KM (plus Class Collection) is bi-interpretable with a first-order set theory without powerset. The third construction is a version of Gödel's constructible universe in the classes. Given a model of KM this gives a smaller model of KM plus Class Collection with the same ordinals.

I investigate these constructions over a weaker base theory than KM (plus Class Collection), generalizing Marek and Mostowski's results to weaker theories. In particular, this shows that for the second-order set theories T in which we are interested that being T -realizable is closed under taking inner models. I close the chapter with an application of the constructions, showing that the least height of a transitive model of $\text{GBC} + \Pi_k^1\text{-CA}$ is less than the least height of a β -model of $\text{GBC} + \Pi_k^1\text{-CA}$. This generalizes an analogous result—due to Marek and Mostowski, as the reader may have guessed—about transitive and β -models of KM .

Chapter 3 is about transfinite recursion principles in second-order set theory. The main result is that there is a hierarchy of theories, ranging in strength from GBC to KM , given by transfinite recursion principles. This hierarchy is ordered first by the complexity of the properties we can do recursion for and second by the lengths of recursion that can be done.

Chapter 4 investigates the phenomenon of minimal models of second-order set theories. The main result is that strong second-order set theories—e.g. KM or $\text{GBC} + \Pi_1^1\text{-CA}$ —do not have least transitive models whereas weaker second-order set theories—e.g. GBC or $\text{GBC} + \text{ETR}_{\text{Ord}}$ —do. Indeed, the results of that chapter show that no countable model of ZFC can have a least KM -realization (and similarly for other strong theories). Left open is the question of whether $\text{GBC} + \text{ETR}$ has a least transitive model. I show that there is a basis of minimal $(\text{GBC} + \text{ETR})$ -realizations for any $(\text{GBC} + \text{ETR})$ -realizable model M so that if $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}$ then \mathcal{X} sits above precisely one of these basis realizations. I also show that the second-order set theories considered in this dissertation have least β -models.

Some of the material in this dissertation also appeared in a paper of mine [Wil17] which, at time of writing, is under review. In particular, that paper contains most of the results of chapter 4, parts of chapter 3, and a little bit from the end of chapter 1.

Chapter 1

A first look at models of second-order set theories

Wesentlich aber ist, da auch "zu große" Mengen Gegenstand dieser Mengenlehre sind, nämlich diejenigen II. Dinge, die keine I.II. Dinge sind. Anstatt sie gänzlich zu verbieten, werden sie nur für unfähig erklärt Argumente zu sein (sie sind keine I. Dinge!). Zum Vermeiden der Antinomien reicht das aus und ihre Existenz ist für gewisse Schluweisen notwendig.

John von Neumann

The purpose of this chapter is to introduce the important objects of study for this dissertation and lay out some basic properties thereof.

I begin by introducing the major second-order set theories under study, followed by some important classes of models for those theories. Next comes a discussion of means for checking whether a structure satisfies the axioms of these theories. In particular, I look at when these axioms are preserved by class forcing. This transitions into some well-known constructions for producing models of weak second-order set theories.

I end the chapter with an in-depth look at the collection of GBC-realizations for a fixed countable model of ZFC. Some of those results can be generalized to stronger theories, and I discuss to what extent this can be done. But a fuller look at the topic for stronger theories is delayed until chapter 4, after we have built up more tools.

Most of the work in this chapter is not new. I have strived to indicate clearly where each theorem originates. At times, however, I have resorted to labeling a result as folklore when its origin is lost in the literature to me.

1.1 Dramatis Personae

In this section I present the main players in the drama. I introduce the main second-order set theories of interest and the important classes of models of these theories. First, let me set up some framework.

The reader may know that there are two main approaches to formalizing second-order set theory. The first is to use a one-sorted theory, where the only objects are classes and sets are those classes which are elements of another class. I will not take that approach. Instead, I will take the two-sorted approach, where

there are two types of objects: sets and classes. The domain of a model (M, \mathcal{X}) has two parts, where M is the *first-order part of the model*, with elements of M being the *sets*, and \mathcal{X} is the *second-order part of the model*, with elements of \mathcal{X} being the *classes*. I will suppress writing the membership relation for the model, referring simply to (M, \mathcal{X}) . In case we need to refer to the membership relation I will write $\in^{(M, \mathcal{X})}$.

In this dissertation we will often be interested in different models which have the same sets but different classes. This easily fits to the two-sorted approach, where we can easily talk about models (M, \mathcal{X}) and (M, \mathcal{Y}) with the same first-order part. In contrast, the one-sorted approach is awkward here.

I will use \mathcal{L}_\in to refer to the language of set theory, whether second-order or first-order. The only non-logical symbols in the language are for membership. Extensions of this language by adding new symbols will be denoted e.g. $\mathcal{L}_\in(A)$.

It should be emphasized that the “second-order” in second-order set theory refers to the use of classes, not to the logic. The theories we consider here will all be formalized in first-order logic.¹ Consider the analogous situation of second-order arithmetic, where two-sorted theories of arithmetic—with numbers and sets of numbers as objects—are formulated in first-order logic.

In the formal language for these theories I will distinguish between variables for sets and variables for classes by using lowercase letters for the former and uppercase letters for the latter. For instance, the formula $\forall X \exists y y \in X$ gives the (false) proposition that every class has some set as a member. Say that a formula in the language of set theory is *first-order* if it has no class quantifiers (though classes may appear as variables). Those formulae which do have class quantifiers are called *second-order*. The first-order formulae are stratified according to the usual Lévy hierarchy, which will be denoted by Σ_n^0 and Π_n^0 . There is also a stratification of the second-order formulae. A formula is Σ_1^1 if it is of the form $\exists X \varphi(X)$ or is Π_1^1 if it is of the form $\forall X \varphi(X)$, where φ is first-order. This extends upward to Σ_n^1 and Π_n^1 in the obvious manner. It will sometimes be convenient to have a name for the first-order and second-order formulae. I will use Σ_ω^0 , Π_ω^0 , Σ_0^1 , or Π_0^1 for the first-order formulae and Σ_ω^1 or Π_ω^1 for the second-order formulae.

All second-order set theories I consider will include Class Extensionality as an axiom. It will often be convenient to assume that our models are of the form (M, \mathcal{X}) where $\mathcal{X} \subseteq \mathcal{P}(M)$ and the set-membership relation is the true \in . However, one must be a little careful here. It may be that some of the elements of M are subsets of M . If M is transitive and its membership relation is $\in \upharpoonright M$, then this is no problem as in this case the set-set and set-class membership relations will cohere for any $\mathcal{X} \subseteq \mathcal{P}(M)$ we choose. But it could be that the membership relation of M is some weird thing. Nevertheless, it is always true that our models (M, \mathcal{X}) are isomorphic to a model of the form (M, \mathcal{X}) where $\mathcal{X} \subseteq \mathcal{P}(M)$ and the set-membership relation is the true \in . We first find $M \cong \bar{M}$ so that $M \cap \mathcal{P}(M) = \emptyset$. Next, using that all our models will satisfy Extensionality for classes, we can realize the classes of (the isomorphic copy of) our model as literal subsets of the model.²

1.1.1 Second-order set theories

The theories can be roughly grouped into three groups: weak, strong, and medium. I will present them in that order.

Definition 1.1. Gödel–Bernays set theory with Global Choice GBC is axiomatized with the following.³

- ZFC for sets.
- Extensionality for classes.
- Class Replacement—if F is a class function and a is a set then $F''a$ is a set.

¹One could also think of them as formalized in second-order logic with Henkin semantics, but officially we will use first-order logic.

²There is a small point here which needs to be addressed. Under this approach the sets and classes are disjoint. So, for example, if we are being really careful we must distinguish between the set of finite ordinals and the class of finite ordinals, as they are different objects. Nevertheless, it will follow from the axioms we use that every set has the same elements as some class. In practice I will not always be careful to distinguish a set from the class it is co-extensive with.

³In the literature one also sees this axiom system called NBG .

- Global Choice—there is a class bijection $\text{Ord} \rightarrow V$.
- Elementary Comprehension—if $\varphi(x)$ is a first-order formula, possibly with set or class parameters, then $\{x : \varphi(x)\}$ is a class.

Dropping Global Choice from the axiomatization gives the theory **GBC**. The **c** reminds one that while Global Choice is lacking, there are still choice functions for sets. Though it will not be used in this dissertation, **GB** is used to refer to **ZF** plus Class Extensionality, Class Replacement, and Elementary Comprehension.

This axiomatization is not parsimonious. In particular, it has infinitely many axioms whereas **GBC** is known to be finitely axiomatizable. An advantage of this axiomatization is that it makes immediately apparent the distinction between classes and sets and how **GBC** relates to **ZFC**. It is obvious from this axiomatization that we can obtain a model of **GBC** by taking a model of **ZFC** (perhaps we require more from the model⁴) and adding certain classes. Later in this chapter we will see how to verify whether a collection of classes for a model of **ZFC** gives a model of **GBC**.

Note that by Elementary Comprehension for every set x there is a class X which has the same elements. But not all classes are co-extensive with sets. For instance, By Elementary Comprehension there is a class of all sets. But there can be no such set, by a well-known argument of Russell's. A class which is not co-extensive with a set is called a *proper class*.

Also note that **GBC** proves Separation for classes, i.e. that $A \cap b$ is a set for every class A and every set b . To see this, let F be the class function which is the identity on A and sends every set not in A to some designated element, say \emptyset . By Class Replacement $a = F''b$ is a set. Then either a or $a \setminus \{\emptyset\}$ will be $A \cap b$, depending upon whether $b \subseteq A$ and $\emptyset \in A \cap b$.

Next we look at much stronger theories. The difference in axiomatization may appear slight—allowing impredicative definitions in Comprehension—but the effects are profound.

Definition 1.2. Kelley–Morse set theory **KM** is axiomatized with the axioms of **GBC** plus the full Comprehension schema. Instances of this schema assert that $\{x : \varphi(x)\}$ is a set for any formula φ , possibly with class quantifiers and set or class parameters.⁵

We can strengthen **KM** by adding the Class Collection schema. For some purposes, **KM** is not quite enough and we need the extra strength of this schema.

Definition 1.3. The theory **KMCC** is obtained from **KM** by adding the Class Collection axiom schema.⁶ Informally, this schema asserts that if for every set there is a class satisfying some property, then there is a coded hyperclass⁷ consisting of witnesses for each set. Formally, let $\varphi(x, Y)$ be a formula, possibly with parameters. The instance of Class Collection for φ asserts

$$[\forall x \exists Y \varphi(x, Y)] \Rightarrow [\exists C \forall x \exists y \varphi(x, (C)_y)]$$

where $(C)_y = \{z : (y, z) \in C\}$ is the y -th slice of C .

⁴Though we will see later in this chapter that we do not need to require more, at least for countable models. Any countable model of **ZFC** can be expanded to a model of **GBC**. Consequently, **GBC** is conservative over **ZFC**.

⁵In the literature **KM** has many other names—I have seen **MK** for Morse–Kelley (e.g. [AF]), **MT** for Morse–Tarski (e.g. [Chu80]), **MKT** for Morse–Kelley–Tarski (e.g. [Chu81]), and **QM** for Quine–Morse (e.g. [Die83]). If Monty Python did sketches about set theory instead of breakfast [Pyt70] no doubt we would also have Morse–Kelley–Morse–Morse–Tarski–Morse, or **MKMMTM**.

⁶Continuing a theme of previous footnotes, both **KMCC** and Class Collection have different names in the literature. Antos and Friedman [AF] call them **MK*** and Class Bounding while Gitman and Hamkins [GH] call them **KM+** and Class Choice. I myself previously have used **KM+** [Wil17], but in this dissertation I will consider second-order set theories formulated without the axiom of Powerset. Following the standard of referring to **ZFC**–Powerset as **ZFC**[−] I will call these theories **GBC**[−], **KMCC**[−], and so forth. Using **KM+** would lead to the infelicitous **(KM+)**[−]. So **KMCC** it is.

⁷A *hyperclass* is a collection of classes. A hyperclass \mathcal{A} is *coded* if there is a class C so that $\mathcal{A} = \{(C)_x : x \in V\}$ where $(C)_x = \{y : (x, y) \in C\}$ is the x -th slice of C . Officially, of course, hyperclasses are not objects in the models and any talk of such is a paraphrase, similar to the usage of classes in first-order set theory.

Observe that under Global Choice, Class Collection is equivalent to the schema with instances

$$[\forall x \exists Y \varphi(x, Y)] \Rightarrow [\exists C \forall x \varphi(x, (C)_x)],$$

that is where x is the the index of the slice in C witnessing the property for x . This version of the schema has the flavor of a choice principle, hence it sometimes being called Class Choice.

The set theorist who does not work with second-order set theories may wonder why we would want to work with something even stronger than KM. To her I have two responses. First, KM behaves badly with some constructions. For instance, set theorists like to take ultrapowers of the universe using some measure. In order for Łoś's theorem to be satisfied for the full second-order language, we need Class Collection. Gitman and Hamkins showed that KM alone does not suffice [GH]. Second, the natural models of KM are actually models of KMCC. If κ is inaccessible then $(V_\kappa, V_{\kappa+1})$ is a model of KMCC. This may not satisfy the skeptic who is worried about a jump in consistency strength, but we will see in chapter 2 that the skeptic need not worry, as KMCC does not exceed KM in consistency strength.

It is immediate that KM is stronger than GBC. Indeed, KM proves the existence of Σ_k^1 truth predicates for every (standard) k . Therefore, KM proves Con(ZFC) so once we see that GBC and ZFC are equiconsistent we will see that the separation is also in terms of consistency strength. However, there is a significant gap between the two theories. We can weaken Comprehension to get intermediate theories, though the following are still grouped among the strong theories.

Definition 1.4. Let k be a (standard) natural number. The Π_k^1 -Comprehension Schema Π_k^1 -CA is the restriction of the Comprehension schema to Π_k^1 -formulae. Note that, over GBC^- , this is equivalent to restricting Comprehension to Σ_k^1 -formulae.

Recall that Π_ω^1 refers to the second-order formulae, of any complexity. It will sometimes be convenient to use Π_ω^1 -CA or Π_ω^1 -Comprehension as a synonym for the full second-order Comprehension schema.

Observe that Π_0^1 -CA is Elementary Comprehension. So we are really only interested in the case where $k > 0$.

We get that $\text{GBC} + \Pi_1^1$ -CA proves Con(GBC) and $\text{GBC} + \Pi_k^1$ -CA proves Con($\text{GBC} + \Pi_n^1$ -CA) for $n < k$, as the Σ_n^1 truth predicate can be defined via a Σ_{n+1}^1 -formula. So there is a hierarchy of theories between GBC and KM, increasing in consistency strength.

It is also useful to consider fragments of Class Collection.

Definition 1.5. Let k be a (standard) natural number. The Σ_k^1 -Class Collection axiom schema, denoted by Σ_k^1 -CC, is the restriction of the Class Collection schema to Σ_k^1 -formulae. Elementary Class Collection ECC is another name for Σ_0^1 -CC.

In chapter 2 we will see that $\text{GBC} + \Pi_k^1$ -CA + Σ_k^1 -CC does not exceed $\text{GBC} + \Pi_k^1$ -CA in consistency strength. See corollary 2.48.

Observation 1.6. Over GBC^- we have that Σ_k^1 -CC implies Π_k^1 -Comprehension.

Proof. Let $\varphi(x)$ be a Σ_k^1 -formula, possibly with (suppressed) parameters. Apply the instance of Class Collection to the formula

$$(\varphi(x) \wedge Y = \{x\}) \vee (\neg\varphi(x) \wedge Y = \emptyset)$$

to get a class C so that $(C)_x = \{x\}$ if $\varphi(x)$ and $(C)_x = \emptyset$ otherwise. Then, $\{x : \varphi(x)\} = \{x : (C)_x \neq \emptyset\} \in \mathcal{X}$, as desired. \square

On the other hand, Gitman and Hamkins [GH] produced a model of KM which does not satisfy even Σ_0^1 -Class Collection.

Between the weak GBC and the strong Π_k^1 -CA we have the medium theories.

Definition 1.7. We define the **E**lementary **T**ransfinite **R**ecursion schema ETR. This schema asserts that recursions of first-order properties along well-founded relations have solutions. Formally, let $\varphi(x, Y, A)$ be a

first-order formula, possibly with a class parameter A and let R be a well-founded class relation. Denote by $<_R$ the transitive closure of R . The instance of ETR for φ and R asserts that there is a class $S \subseteq \text{dom } R \times V$ which satisfies

$$(S)_r = \{x : \varphi(x, S \upharpoonright r, A)\}$$

for all $r \in \text{dom } R$. Here, $(S)_r = \{x : (r, x) \in S\}$ denotes the r -th slice of S and

$$S \upharpoonright r = S \cap (\{r' \in \text{dom } R : r' <_R r\} \times V)$$

is the partial solution below r .

One example of an elementary recursion is the Tarskian definition of a (first-order) truth predicate. Thus, $\text{GBC} + \text{ETR}$ proves $\text{Con}(\text{ZFC})$ and thereby exceeds GBC in consistency strength. On the other hand, $\text{GBC} + \Pi_1^1\text{-CA}$ proves $\text{Con}(\text{GBC} + \text{ETR})$ (see [Sat14]) so ETR sits below the strong second-order set theories.

It is equivalent, over GBC , to formulate ETR for recursions over well-founded relations, well-founded partial orders, or well-founded tree orders. See [GH17, lemma 7].

We get fragments of ETR by restricting the length of recursions.

Definition 1.8. Let Γ be a class well-order. Then ETR_Γ is the Elementary Transfinite Recursion schema restricted to well-orders of length $\leq \Gamma$.

There is a subtlety here. Namely, the issue is whether ETR_Γ can be expressed as a theory in the language \mathcal{L}_\in of set theory. If Γ is a definable well-order, say $\Gamma = \omega$ or $\Gamma = \text{Ord}$, then this can be done in the obvious way. Different models of set theory may disagree on what Ord is, but it is sensible to ask whether they satisfy ETR_{Ord} .

But we will also be interested in the case where Γ is a specific well-order, possibly undefinable. To be more precise, consider a model (M, \mathcal{X}) of second-order set theory with $\Gamma \in \mathcal{X}$ a well-order. We can then ask whether $(M, \mathcal{X}) \models \text{ETR}_\Gamma$. This may not be expressible as a theory in the language \mathcal{L}_\in of set theory, but because $\Gamma \in \mathcal{X}$ we can use it as a parameter to define ETR_Γ in the expanded language $\mathcal{L}_\in(\Gamma)$.

It will be clear from context which of the two meanings is had in mind, so I will refer to both as simply ETR_Γ .

Let me give an example to illustrate where the distinction matters. Take countable $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}_\gamma$ where $\gamma = \omega_1^M$. Assume that $(M, \mathcal{X}) \not\models \text{ETR}_{\gamma \cdot \omega}$. We will see in chapter 3 that this assumption can be made without loss—if (M, \mathcal{X}) does satisfy $\text{ETR}_{\gamma \cdot \omega}$ then we can throw out classes to get $\vec{\mathcal{X}}$ so that $(M, \vec{\mathcal{X}})$ satisfies $\text{GBC} + \text{ETR}_\gamma$ but does not satisfy $\text{ETR}_{\gamma \cdot \omega}$. Let $g \subseteq M$ be generic over (M, \mathcal{X}) for the forcing to collapse ω_1 to be countable. It is not difficult to check that $(M, \mathcal{X})[g] \models \text{GBC} + \text{ETR}_\gamma$ but $(M, \mathcal{X})[g] \not\models \text{ETR}_{\gamma \cdot \omega}$. So $(M, \mathcal{X})[g]$ will not be a model of the \mathcal{L}_\in -theory ETR_{ω_1} , even though it is a (set) forcing extension of a model of ETR_{ω_1} .

Another issue with expressing ETR_Γ as an \mathcal{L}_\in -theory is that we can have definitions for a well-order Γ which are highly non-absolute. For instance, suppose Γ is defined as “if $V = L$ then $\Gamma = \text{Ord}$ and otherwise $\Gamma = \omega_1$ ”. Then there is a model of $\text{GBC} + \text{ETR}_\Gamma$ whose L is not a model of $\text{GBC} + \text{ETR}_\Gamma$. On the other hand, as we will see in chapter 3, $\text{GBC} + \text{ETR}_\Gamma$ as an $\mathcal{L}_\in(\Gamma)$ -theory does go down to inner models.⁸

The reader who is familiar with reverse mathematics may see an analogy to second-order arithmetic. Namely, three of these theories line up with the strongest three of the “big five” subsystems of second-order arithmetic: GBC is analogous to ACA_0 , $\text{GBC} + \text{ETR}$ is analogous to ATR_0 and $\text{GBC} + \Pi_1^1\text{-CA}$ is analogous to the subsystem of second-order arithmetic which is referred to as $\Pi_1^1\text{-CA}_0$. At the highest level, KM is analogous to Z_2 , full second-order arithmetic. This analogy can be useful to keep in mind. However, the reader should beware that results from arithmetic do not always generalize to set theory. For example, Simpson proved that there is no smallest β -model of ATR_0 —see [Sim09] for a proof. But there is a smallest β -model of $\text{GBC} + \text{ETR}$, as we will see in chapter 4. (For the reader who does not know what a β -model is, we will get to that later in this section).

⁸There is a technical caveat here. Namely, Γ must be sufficiently nice over the inner model to avoid pathologies such as $\Gamma \subseteq L$ which codes 0^\sharp . See theorem 3.16 for details.

We will briefly return to this analogy at the end of chapter 4, after we have seen enough theorems about models of second-order set theories to satisfactorily explore its limits.

1.1.2 Doing without Powerset

All of the second-order set theories considered so far include that the sets satisfy ZFC. But we can ask for less out of the first-order part. We get variants on all of the above theories by dropping the requirement that the first-order part satisfy Powerset.

Definition 1.9. The first-order set theory ZFC^- is axiomatized by Extensionality, Pairing, Union, Infinity, Foundation, Choice, Separation, and Collection.

The reader should be warned that in the absence of Powerset that Collection is stronger than Replacement [Zar96] and thus we do not want to use Replacement to axiomatize ZFC^- , as the resulting theory is badly behaved; see [GHJ16] for some discussion of how this theory misbehaves.

We get “minus versions” of all the above-defined second-order set theories by dropping the requirement that the first-order part satisfy Powerset. I will give one definition in full and leave it to the reader to fill in the pattern for the others. The way to think of them is that, for example, KM^- is $KM - \text{Powerset}$. (But one should keep in mind the Collection versus Replacement issue.)

Definition 1.10. The second-order set theory GBC^- is axiomatized by the following.

- ZFC^- for sets.
- Extensionality for classes.
- Class Replacement.
- Global Choice, in the form “there is a bijection $\text{Ord} \rightarrow V$ ”.
- Elementary Comprehension.

In the absence of Powerset, the various equivalent forms of Global Choice are no longer equivalent. (See section 1.2 for a proof.) The strongest is the assertion that there is a bijection from Ord to V , or equivalently, that there is a global well-order of ordertype Ord . I adopt this strongest version as the official form of Global Choice for GBC^- , though at times we could get away with less.

Models of, say $KMCC^-$ are not hard to come by. Indeed, $KMCC^-$ is much weaker than ZFC in consistency strength. In a model of ZFC if κ is a regular uncountable cardinal then $(H_\kappa, \mathcal{P}(H_\kappa)) \models KMCC^-$. A special case of particular interest is that of the hereditarily countable sets: $(H_{\omega_1}, \mathcal{P}(H_{\omega_1}))$ is a model of $KMCC^- +$ every set is countable.

1.1.3 Models of second-order set theory

An important theme of this work is the following: given a fixed model M of first-order set theory what can be said about possible second-order parts that can be put on M to make a model of some second-order set theory? It will be convenient to have a name for these possible collections of classes.

Definition 1.11. Let M be a model of first-order set theory and T be some second-order set theory. A T -realization for M is a set $\mathcal{X} \subseteq \mathcal{P}(M)$ so that $(M, \mathcal{X}) \models T$. If M has a T realization then we say M is T -realizable.

Many properties of first-order models can also be had by second-order models, via the exact same definition. For instance, (M, \mathcal{X}) is ω -standard (synonymously, is an ω -model) if ω^M is well-founded. One important property is transitivity.




Figure 1.1: Some classes of models of second-order set theory.

Definition 1.12. A model (M, \mathcal{X}) of second-order set theory is *transitive* if its membership relations are the true \in . This is equivalent to requiring that M is transitive, due to our convention of only considering models so that $\mathcal{X} \subseteq \mathcal{P}(M)$.

It is well-known that transitive models of ZFC are correct about well-foundedness: if transitive $M \models \text{ZFC}$ thinks that $R \in M$ is a well-founded relation then R really is well-founded. (Indeed, the same is true for much weaker theories, e.g. ZFC^- .) This does not hold for transitive models of second-order set theories. While they will be correct about whether set relations are well-founded they can be wrong about whether a class relation is well-founded.⁹ Models of second-order set theory which are correct about which of their class relations are well-founded are of special interest.

Definition 1.13. A model (M, \mathcal{X}) of second-order set theory is a *β-model* if its membership relations are well-founded and it is correct about well-foundedness. That is, if $R \in \mathcal{X}$ is a relation which (M, \mathcal{X}) thinks is well-founded then R really is well-founded. (The inverse direction, that if R is well-founded then (M, \mathcal{X}) thinks R is well-founded is always true for well-founded models by downward absoluteness.)

Observe that every β -model is isomorphic to a transitive model, so we can usually assume without loss that a β -model is transitive.

The following observation shows that the distinction between β -model and transitive model does not arise a certain class of models, which includes many natural models considered by set theorists.

Observation 1.14. *Suppose $V_\alpha \models \text{ZFC}^-$ is transitive with $\text{cof } \alpha > \omega$. Then (V_α, \mathcal{X}) , equipped with the true membership relation, is a β -model for any $\mathcal{X} \subseteq \mathcal{P}(V_\alpha)$.*

Proof. Suppose $R \in \mathcal{X}$ is ill-founded (i.e. from the perspective of V). But witnesses to ill-foundedness are countable sequences and V_α is closed under countable sequences. So (V_α, \mathcal{X}) thinks that R is ill-founded. Since R was arbitrary, (V_α, \mathcal{X}) is correct about well-foundedness. \square

We will also be interested in various ways models may be contained within each other. The basic notion is that of a submodel. This is a familiar concept, but let me give a definition specialized to the context of second-order set theory.

Definition 1.15. Let (M, \mathcal{X}) and (N, \mathcal{Y}) be models of second-order set theory. Say that (M, \mathcal{X}) is a *submodel* of (N, \mathcal{Y}) , written $(M, \mathcal{X}) \subseteq (N, \mathcal{Y})$, if $M \subseteq N$, $\mathcal{X} \subseteq \mathcal{Y}$, and the membership relation for (M, \mathcal{X}) is the restriction of $\in^{(N, \mathcal{Y})}$ to (M, \mathcal{X}) .

⁹For the reader who was previously unaware of this folklore result, in chapter 4 we will construct, as a by-product toward other goals, transitive models which are wrong about well-foundedness.

This definition can be strengthened in various ways.

Definition 1.16. Let (M, \mathcal{X}) and (N, \mathcal{Y}) be models of second-order set theory.

- Say that (M, \mathcal{X}) is a *V-submodel* of (N, \mathcal{Y}) if $M = N$ and $\mathcal{X} \subseteq \mathcal{Y}$. The name is because in this case $V^{(M, \mathcal{X})} = V^{(N, \mathcal{Y})}$.
- Say that (M, \mathcal{X}) is an *Ord-submodel* of (N, \mathcal{Y}) if $(M, \mathcal{X}) \subseteq (N, \mathcal{Y})$ and $\text{Ord}^M = \text{Ord}^N$. The name is because in this case $\text{Ord}^{(M, \mathcal{X})} = \text{Ord}^{(N, \mathcal{Y})}$.
- Say that (M, \mathcal{X}) is an *inner model* of (N, \mathcal{Y}) if (M, \mathcal{X}) is an Ord-submodel of (N, \mathcal{Y}) and \mathcal{X} is definable over (N, \mathcal{Y}) , possibly via a second-order formula using parameters. Note that it is automatic that $M \in \mathcal{Y}$, since $M \in \mathcal{X} \subseteq \mathcal{Y}$.
- Let (M, \mathcal{X}) be a V-submodel of (M, \mathcal{Y}) . Say that (M, \mathcal{X}) is a *coded V-submodel* of (M, \mathcal{Y}) if \mathcal{X} is coded in \mathcal{Y} . That is, there is a single class $C \in \mathcal{Y}$ so that $\mathcal{X} = \{(C)_x : x \in M\}$.

The reader may find examples to be helpful. Suppose κ is inaccessible and that $L_\kappa \neq V_\kappa$. Then $(V_\kappa, \text{Def}(V_\kappa)) \models \text{GBc}$ is a V-submodel of $(V_\kappa, \mathcal{P}(V_\kappa)) \models \text{KMCC}$ and $(L_\kappa, \text{Def}(L_\kappa)) \models \text{GBC}$ is an Ord-submodel of $(V_\kappa, \mathcal{P}(V_\kappa))$. In fact, they are both inner models because $\mathcal{P}(V_\kappa)$ contains truth predicates for all $A \subseteq V_\kappa$ and can thus uniformly pick out which classes are in $\text{Def}(A)$.

1.2 Verifying the axioms

Many times in this dissertation we will find ourselves in the following situation. We have some first-order model M of set theory and some $\mathcal{X} \subseteq \mathcal{P}(M)$ a collection of classes over M . We will want to be able to say something about the theory of (M, \mathcal{X}) . In this section I present some basic tools one can use in this situation, focusing here on the axioms of **GBC**.

Let us begin with the most obvious of observations.

Observation 1.17. *A second-order model of set theory (M, \mathcal{X}) with $\mathcal{X} \subseteq \mathcal{P}(M)$ and the true \in for its set-class membership relation always satisfies Extensionality for classes. \square*

To see that $(M, \mathcal{X}) \models \text{GBc}$ satisfies Global Choice one just has to see that \mathcal{X} contains a bijection $\text{Ord} \rightarrow V$. As with the ordinary axiom of choice there are several equivalent forms.

Fact 1.18. *Let $(M, \mathcal{X}) \models \text{GBc}$ be a second-order model of set theory. The following are equivalent.*

1. \mathcal{X} contains a bijection $\text{Ord} \rightarrow V$.
2. \mathcal{X} contains a global choice function, that is a class function F whose domain is the class of nonempty sets so that $F(x) \in x$ for all x .
3. \mathcal{X} contains a global well-order, that is a well-order whose domain is the entire universe of sets.
4. \mathcal{X} contains a global well-order of ordertype Ord .

Proof. Both $(1 \Rightarrow 2)$ and $(4 \Rightarrow 1)$ are obvious. That leaves only $(2 \Rightarrow 3)$ and $(3 \Rightarrow 4)$ to check.

$(2 \Rightarrow 3)$ Using the global choice function F we can well-order the V_α 's in a coherent fashion. Suppose we have already defined a well-order $<_\alpha$ of V_α . Then define a well-order $<_{\alpha+1}$ of $V_{\alpha+1}$ which extends $<_\alpha$ by using the global choice function: $<_{\alpha+1}$ is simply $F(w)$, where w is the set of well-orders of $V_{\alpha+1}$ which extend $<_\alpha$. Then $\bigcup_{\alpha < \alpha} <_\alpha$ is a global well-order (of ordertype Ord , in fact).

$(3 \Rightarrow 4)$ Let $<^*$ be a global well-order. We define a new global well-order $<^\dagger$ as:

$$x <^\dagger y \quad \text{iff} \quad \text{rank } x < \text{rank } y \text{ or } (\text{rank } x = \text{rank } y \text{ and } x <^* y).$$

Then $<^\dagger$ has ordertype Ord . \square

As was mentioned in subsection 1.1.2, these are not all equivalent in the absence of Powerset. We officially adopted the strongest form of Global Choice for the powerset-free context, whose equivalent forms are “there is a bijection $\text{Ord} \rightarrow V$ ” and “there is a global well-order of ordertype Ord ”. Let us quickly see that the other forms of Global Choice are weaker in this context.

Fact 1.19.

1. Over KMCC^- – Global Choice,¹⁰ the existence of a global well-order does not imply the existence of a global well-order of ordertype Ord .
2. (Shapiro [Sha91, theorem 5.4]) Assume $\text{AD}_{\mathbb{R}}$, asserting the determinacy of every two-player, perfect information game where the two players play reals, is consistent with ZF .¹¹ Then over KMCC^- – Global Choice, the existence of a global choice function does not imply the existence of a global well-order.¹²

Proof Sketch. (1) Force if necessary to get the continuum to have size \aleph_2 . Then one can check that $(H_{\omega_1}, \mathcal{P}(H_{\omega_1}))$ is a model of all the axioms of KMCC^- except Global Choice. Easily, it has a global well-order (externally seen to have ordertype ω_2) but has no global well-order of ordertype $\omega_1 = \text{Ord}^{H_{\omega_1}}$.

(2) Assume $\text{ZF} + \text{AD}_{\mathbb{R}}$. Consider the model $(M, \mathcal{X}) = (H_{\omega_1}, \mathcal{P}(H_{\omega_1}))$. One can check it satisfies all the axioms of KMCC^- – except Global Choice. Observe that \mathcal{X} cannot contain a global well-order, as that would imply there is a well-order of \mathbb{R} in the ambient universe, contradicting $\text{AD}_{\mathbb{R}}$. But \mathcal{X} does have a global choice function, which arises from a winning strategy for the following game: Player I plays a real which codes a nonempty hereditarily countable set. Player II responds by playing a real, then the game ends. Player II wins if her real codes an element of the set coded by player I’s real and otherwise player I wins. This game is determined by $\text{AD}_{\mathbb{R}}$ and it is clear that player I could not possibly have a winning strategy. So player II has a winning strategy from which we can extract a global choice function for M . \square

Note that (2) requires far from the full strength of $\text{AD}_{\mathbb{R}}$, only needing that one can have clopen determinacy for games played with reals while not having a well-order of the reals. The reader who wishes to know the exact strength needed for (2) is welcome to do that work herself.

Next let us see how to check whether our classes satisfy Elementary Comprehension.

Definition 1.20. Let M be a model of set theory with $A \subseteq M$. Then $\text{Def}(M; A)$ is the collection of classes of M definable from A , possibly with set parameters. Formally,

$$X \in \text{Def}(M; A) \Leftrightarrow X = \{x \in M : (M, A) \models \varphi(x, p)\} \text{ for some first-order } \varphi \text{ and } p \in M.$$

More generally, let \mathcal{X} be a collection of classes from M and A_i , for i in some index set I , be classes of M . Then $\text{Def}(M; \mathcal{X}, A_i : i \in I)$ is the collection of classes of M definable from (finitely many) classes from $\mathcal{X} \cup \{A_i : i \in I\}$.¹³

I will write $\text{Def}(M)$ to refer to $\text{Def}(M; \emptyset)$.

Observation 1.21. Let (M, \mathcal{X}) be a second-order model of set theory. Then (M, \mathcal{X}) satisfies Elementary Comprehension if and only if \mathcal{X} is closed under first-order definability—i.e. for any $A_0, A_1, \dots, A_n \in \mathcal{X}$ we have $\text{Def}(M; A_0, A_1, \dots, A_n) \subseteq \mathcal{X}$.

Proof. (\Rightarrow) If $B \in \text{Def}(M; A_0, A_1, \dots, A_n)$ then B was defined from the A_i ’s by some first-order formula. But then Elementary Comprehension yields that B is a class.

¹⁰Note that this theory includes the well-ordering theorem for sets, which is necessary to have a global well-order at all.

¹¹See [Gaß94] for a proof which does not need this consistency assumption.

¹²The context for Shapiro’s result here is second-order logic, as is the context for Gaßner’s paper cited in the previous footnote. This is a reformulation of his result into the context of second-order set theory.

¹³Two remarks are in order. First, this definition is ambiguous, as a set can be both a subset of M and also a collection of subsets of M . This happens if, for instance, M is transitive and $\mathcal{X} \subseteq M$. But this will not arise in practice and we will sacrifice complete unambiguity in favor of readable notation. Second, while we could generalize the definition further and allow multiple collections \mathcal{X}_i of classes of M , this is not needed for this work. We will be interested in $\text{Def}(M; \mathcal{X}, A)$ when (M, \mathcal{X}) is a model of second-order set theory.

(\Leftarrow) Fix class A_0, A_1, \dots, A_n and a formula $\varphi(x, A_0, A_1, \dots, A_n)$. Then

$$\{x : \varphi(x, A_0, A_1, \dots, A_n)\}$$

is a class because it is definable from the A_i 's. \square

In general Class Replacement can be tricky to check. Nevertheless, there are some circumstances where it is trivial.

Observation 1.22. *Suppose $(M, \mathcal{X}) \models \text{GBC}^-$. Let (N, \mathcal{Y}) be an Ord-submodel of (M, \mathcal{X}) which satisfies all the axioms of GBC^- except possibly Class Replacement. Then in fact (N, \mathcal{Y}) satisfies Class Replacement.*

Proof. Suppose $F \in \mathcal{Y}$ is a class function and $a \in N$ is a set so that $F''a \notin N$. By Global Choice in (N, \mathcal{Y}) , there is a bijection in \mathcal{Y} between $F''a \in \mathcal{Y}$ and Ord^N . So the failure of Class Replacement for F and a gives a map in \mathcal{Y} from $\alpha \in \text{Ord}^N$ to Ord^N . But this same map must be in \mathcal{X} , contradicting that (M, \mathcal{X}) satisfies Class Replacement. \square

This argument does not need Global Choice. It is enough that \mathcal{Y} contains a \subseteq -increasing sequence $\langle n_\alpha : \alpha \in \text{Ord}^N \rangle$ of sets from N so that $\bigcup_\alpha n_\alpha = N$. This allows (N, \mathcal{Y}) to define a ranking function relative to this n_α -hierarchy, and from that get a map from α to Ord^N . In particular, the argument goes through if N satisfies Powerset, since then it has the V_α -hierarchy.

Observation 1.23. *Suppose $(M, \mathcal{X}) \models \text{GBC}^-$. Let (N, \mathcal{Y}) be an Ord-submodel of (M, \mathcal{X}) satisfying all the axioms of GBC except possibly Class Replacement. Then in fact (N, \mathcal{Y}) satisfies Class Replacement.* \square

Together these observations give us the tools to check whether $\mathcal{X} \subseteq \mathcal{P}(M)$ is a GBC -realization (or GBC^- -realization) for M . We do not have such nice tools for stronger theories. Nevertheless, something can be said. For the medium theories, in chapter 3 we will see that Elementary Transfinite Recursion is equivalent to the existence of certain classes, namely iterated truth predicates. A similar result will hold for ETR_Γ .

For stronger forms of Comprehension, it is true that (M, \mathcal{X}) satisfying Π_k^1 -Comprehension is equivalent to \mathcal{X} being closed under Π_k^1 -definability. But this means, of course, definability over (M, \mathcal{X}) . So in practice this characterization is not useful and we want other tools.

One specific situation of interest is when our model arises as a forcing extension of a structure already satisfying a strong form of Comprehension. In this case we can say something about whether our model also satisfies Comprehension. More generally, we can ask about the preservation of the axioms under class forcing, to which we now turn.

1.3 Preserving the axioms

There is more than one approach to formalize class forcing. I will take the following. We work over a model (M, \mathcal{X}) . A *forcing notion* $\mathbb{P} \in \mathcal{X}$ is a separative partial order with a maximum element $\mathbf{1}$. If $p \leq q$ we say that p is *stronger than* q . Two conditions p and q are *compatible*, denoted $p \parallel q$, if there is $r \leq p, q$. Otherwise, p and q are *incompatible*, denoted $p \perp q$.

Given \mathbb{P} we can define the collection of \mathbb{P} -names. These are sets or classes whose elements are of the form (τ, p) where τ is a \mathbb{P} -name and $p \in \mathbb{P}$. This *prima facie* circular definition is actually a recursion on ranks. The convention here will be to use capital letters such as Σ for proper class \mathbb{P} -names and lowercase letters such as σ for set \mathbb{P} -names. In case I want to refer to either, I will use lowercase letters.

The forcing relation \Vdash is defined recursively, via the following schema.

Definition 1.24. Let $p \in \mathbb{P}$ and σ, τ, \dots be \mathbb{P} -names. Unless otherwise indicated, they may be either set names or class names. A forcing relation $\Vdash = \Vdash_{\mathbb{P}}$ for \mathbb{P} is a relation between $p \in \mathbb{P}$ and formulae in the forcing language which satisfies the following recursive schema on its domain.

- $p \Vdash \sigma \in \tau$ if and only if there are densely many $q \leq p$ so that there is $(\rho, r) \in \tau$ with $q \leq r$ and $q \Vdash \sigma = \rho$;
- $p \Vdash \sigma \subseteq \tau$ if and only if for all $(\rho, r) \in \sigma$ and all $q \leq p, r$ we have $q \Vdash \rho \in \tau$;
- $p \Vdash \sigma = \tau$ if and only if $p \Vdash \sigma \subseteq \tau$ and $p \Vdash \tau \subseteq \sigma$;
- $p \Vdash \varphi \wedge \psi$ if and only if $p \Vdash \varphi$ and $p \Vdash \psi$;
- $p \Vdash \neg\varphi$ if and only if no $q \leq p$ forces φ ;
- $p \Vdash \forall x\varphi(x)$ if and only if $p \Vdash \varphi(\sigma)$ for every set \mathbb{P} -name σ ; and
- $p \Vdash \forall X\varphi(X)$ if and only if $p \Vdash \varphi(\dot{A})$ for every class \mathbb{P} -name \dot{A} .

If Φ is a collection of formulae say that \mathbb{P} *admits a forcing relation for* Φ (or, synonymously, $\Vdash_{\mathbb{P}}$ *exists for* Φ) if there is a class \Vdash which satisfies the above schema which covers all $\varphi \in \Phi$. For φ a formula, \mathbb{P} *admits a forcing relation for* φ if \mathbb{P} admits a forcing relation for the collection of all instances of subformulae of φ . Note that if $\Vdash_{\mathbb{P}}$ exists for the atomic formulae then, by an induction in the metatheory, $\Vdash_{\mathbb{P}}$ exists for all φ .

Observe that each step in this recursion, except the last, is done in a first-order way. As such, it is immediate that $\text{GBC} + \text{ETR}$ proves the forcing relation $\Vdash_{\mathbb{P}}$ exists for first-order formulae. Then $\Vdash_{\mathbb{P}}$ restricted to subformulae of a second-order formula is a definable hyperclass, via an induction in the metatheory.

Indeed, the existence of $\Vdash_{\mathbb{P}}$ for every \mathbb{P} is equivalent, over GBC , to a fragment of ETR .

Theorem 1.25 ([GHHSW17]). *Over GBC the following are equivalent.*

- ETR_{Ord} , *Elementary Transfinite Recursion restricted to well-orders of length $\leq \text{Ord}$.*
- *The class forcing theorem, asserting that for every class forcing \mathbb{P} and every formula φ in the forcing language for \mathbb{P} admits a forcing relation $\Vdash_{\mathbb{P}}$ for subformulae of φ .*
- *The uniform first-order class forcing theorem, asserting that every class forcing \mathbb{P} admits a forcing relation $\Vdash_{\mathbb{P}}$ for all first-order formulae in the forcing language.*
- *The atomic class forcing theorem, asserting that every class forcing \mathbb{P} admits a forcing relation for atomic formulae.*

Nevertheless, GBC^- suffices to prove that $\Vdash_{\mathbb{P}}$ exists for a nice collection of class forcings.

Definition 1.26 (S. Friedman [Fri00]). Let \mathbb{P} be a forcing notion.

- $D \subseteq \mathbb{P}$ is *predense below* p (or *predense $\leq p$*) if for every $q \leq p$ is compatible with an element of D .
- \mathbb{P} is *pretame* if given any set-indexed sequence $\langle D_i : i \in a \rangle$ of dense subclasses of \mathbb{P} and any $p \in \mathbb{P}$ there is $q \leq p$ and a sequence $\langle d_i : i \in a \rangle$ of predense $\leq q$ subsets of \mathbb{P} with $d_i \subseteq D_i$ for all $i \in a$.
- (D, D') is a *predense below p partition* (or a *predense $\leq p$ partition*) if $D \cup D'$ is predense $\leq p$ and $p \in D$ and $p' \in D'$ implies $p \perp p'$.
- Two sequences $\mathcal{D} = \langle (D_i, D'_i) : i \in a \rangle$ and $\mathcal{E} = \langle (E_i, E'_i) : i \in a \rangle$ of predense $\leq p$ partitions are *equivalent below q* (or *equivalent $\leq q$*) if for each i the collection of $r \in \mathbb{P}$ so that r meets D_i if and only if r meets E_i is dense below q .¹⁴
- \mathbb{P} is *tame* if it is pretame and additionally for every $p \in \mathbb{P}$ and every set a there is $q \leq p$ and ordinal α so that if $\mathcal{D} = \langle (D_i, D'_i) : i \in a \rangle$ is a sequence of predense $\leq q$ partitions then

$$\{r \in \mathbb{P} : \mathcal{D} \text{ is equivalent } \leq r \text{ to some sequence } \mathcal{E} \in V_\alpha \text{ of predense } \leq q \text{ partitions}\}$$

is dense below q .

¹⁴To clarify, r *meets* predense D means that $r \leq q$ for some $q \in D$.

Theorem 1.27 (S. Friedman, Stanley). GBC^- proves that $\Vdash_{\mathbb{P}}$ exists for φ for every pretame forcing \mathbb{P} and every formula φ in the forcing language.¹⁵

Besides allowing for the forcing language to be definable in a weak theory, the other role of pretameness is in the preservation of the axioms. First, let us make precise how forcing extensions are built. The case for transitive models is well-known, but we can also work over non-transitive models. We work in $(M, \mathcal{X}) \models \text{GBC}^-$ and assume that for our forcing notion $\mathbb{P} \in \mathcal{X}$ we have that for every first-order formula φ of the forcing language the forcing relation $\Vdash_{\mathbb{P}}$ restricted to instances of subformulae of φ exists as a class in \mathcal{X} . In particular, $\Vdash_{\mathbb{P}}$ for atomic formulae is a single class in \mathcal{X} . Given a generic $G \subseteq \mathbb{P}$ and this class we can define the forcing extension as follows.

Work externally to (M, \mathcal{X}) , as is necessary if we have a generic. Define the following relations on set \mathbb{P} -names:

$$\begin{aligned} \sigma \in_G \tau &\iff \exists p \in G p \Vdash \sigma \in \tau \\ \sigma =_G \tau &\iff \exists p \in G p \Vdash \sigma = \tau \end{aligned}$$

and the similar relations \in_G between set \mathbb{P} -names and class \mathbb{P} -names and $=_G$ between class \mathbb{P} -names. It can be straightforwardly checked that $=_G$ is an equivalence relation¹⁶ and \in_G is a congruence modulo $=_G$, meaning that if $\sigma \in_G \tau =_G \rho$ then $\sigma \in_G \rho$. Given a \mathbb{P} -name σ let $[\sigma]_G$ denote the equivalence class of σ modulo $=_G$. Let $[M]_G$ denote the collection of equivalence classes of set \mathbb{P} -names and $[\mathcal{X}]_G$ denote the collection of equivalence classes of class \mathbb{P} -names.

Definition 1.28. Let (M, \mathcal{X}) , \mathbb{P} , and G be as above. Then the *forcing extension* of (M, \mathcal{X}) , denoted $(M, \mathcal{X})[G]$, is the structure $([M]_G, [\mathcal{X}]_G)$ with membership relation \in_G . To refer to the sets of the extension I will use $M[G]$ and to refer to the classes of the extension I will use $\mathcal{X}[G]$.

Proposition 1.29 (The truth lemma). *Consider a model (M, \mathcal{X}) of set theory with forcing notion $\mathbb{P} \in \mathcal{X}$ and generic $G \subseteq \mathbb{P}$. Suppose that for each φ in the forcing language for \mathbb{P} that \mathbb{P} admits a forcing relation for φ in \mathcal{X} . Let $\varphi(x_0, \dots, x_n, Y_0, \dots, Y_m)$ be a formula, τ_0, \dots, τ_n be set \mathbb{P} -names, and $\Sigma_0, \dots, \Sigma_m$ be class \mathbb{P} -names. Then,*

$$\begin{aligned} (M, \mathcal{X})[G] \models \varphi([\tau_0]_G, \dots, [\tau_n]_G, [\Sigma_0]_G, \dots, [\Sigma_m]_G) \\ \text{if and only if} \\ \exists p \in G p \Vdash \varphi(\tau_0, \dots, \tau_n, \Sigma_0, \dots, \Sigma_m). \end{aligned}$$

Proof sketch. By induction on formulae. See [GHHSW17] for more detail. \square

In particular, if (M, \mathcal{X}) is countable then we can always find a generic $G \subseteq \mathbb{P}$. Externally to (M, \mathcal{X}) line up the countably many dense subclasses of \mathbb{P} in \mathcal{X} in ordertype ω and then inductively meet each of them. As such, for countable models the only possible impediment to having forcing extensions is having forcing relations. But so long as we only look at pretame forcings this is no impediment.

Let us turn now to the preservation of the axioms under class forcings. First, let us see the importance of pretameness to this question.

Theorem 1.30 ([Sta84] for $(1 \iff 2)$, [HKSa] for the rest). *Consider $(M, \mathcal{X}) \models \text{GBC}^-$ and let $\mathbb{P} \in \mathcal{X}$ be a forcing notion.¹⁷ The following are equivalent.*

1. \mathbb{P} is pretame.

¹⁵Friedman [Fri00] proved the theorem with the assumption of Powerset, while Stanley [Sta84] independently gave a proof which did not need that assumption.

¹⁶To be clear, both the set-set and class-class relations $=_G$ are equivalence relations on, respectively, the set \mathbb{P} -names and the class \mathbb{P} -names.

¹⁷Holy, Krapf, and Schlicht formulate their result in terms of countable *transitive* models, but it is not hard to see that their result holds more generally.

2. \mathbb{P} preserves GBc^- .
3. \mathbb{P} preserves Collection.
4. \mathbb{P} preserves Replacement.
5. \mathbb{P} preserves Separation and \mathbb{P} satisfies the class forcing theorem.

Tameness has a similar importance for preserving GBc .

Theorem 1.31 (Friedman [Fri00]). *Consider $(M, \mathcal{X}) \models \text{GBc}$ and let $\mathbb{P} \in \mathcal{X}$ be a forcing notion. The following are equivalent.*

1. \mathbb{P} is tame;
2. \mathbb{P} is pretame and preserves Powerset; and
3. \mathbb{P} preserves GBc .

I will not prove these results in full. But as a warmup towards showing that pretame forcings preserve strong forms of Comprehension, let us see that pretame forcings preserve Elementary Comprehension.

Proposition 1.32. *Let $G \subseteq \mathbb{P}$ be generic for $(M, \mathcal{X}) \models \text{GBC}$ with $\mathbb{P} \in \mathcal{X}$ admitting a forcing relation for each φ in the forcing language. Then $(M, \mathcal{X})[G]$ satisfies Elementary Comprehension.*

In particular, pretame forcings always satisfy the forcing theorem so this proposition works for all pretame forcings over some model.

Proof. Consider an instance of Elementary Comprehension. That is, we have a first-order formula $\varphi(x, P)$ with possible parameter P and we want to see that the class $\{x : (M, \mathcal{X})[G] \models \varphi(x, P)\} \in \mathcal{X}[G]$. Towards this end, let $\dot{P} \in \mathcal{X}$ be a name for P . Now consider the name $\dot{B} = \{(\sigma, p) : p \Vdash \varphi(\sigma, \dot{P})\}$. Then, because \mathbb{P} satisfies the forcing theorem, $\dot{B} \in \mathcal{X}$. So

$$\dot{B}^G = \{\sigma^G : \sigma \in M \wedge \exists p \in G \ p \Vdash \varphi(\sigma, \dot{P})\} = \{x \in M[G] : (M, \mathcal{X})[G] \models \varphi(x, P)\}$$

is in \mathcal{X} . □

In fact, the same argument applies higher up.

Corollary 1.33. *Class forcing preserves Π_k^1 -Comprehension, for $1 \leq k \leq \omega$.¹⁸ That is, if $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA}$ and $G \subseteq \mathbb{P} \in \mathcal{X}$ is generic over (M, \mathcal{X}) then $(M, \mathcal{X})[G] \models \Pi_k^1\text{-CA}$.*

Proof. Because $(M, \mathcal{X}) \models \Pi_k^1\text{-CA}$ it in particular satisfies ETR_{Ord} , so it satisfies the uniform first-order class forcing theorem. Now run the same argument as before, but use that $\Vdash_{\mathbb{P}}$ restricted to subformulae of a Σ_n^1 -formula φ , for $1 \leq n < \omega$, is Σ_n^1 -definable from $\Vdash_{\mathbb{P}}$ restricted to first-order formulae. □

We are also interested in the preservation of (fragments) of Class Collection.

Theorem 1.34. *Let $G \subseteq \mathbb{P}$ be generic for $(M, \mathcal{X}) \models \text{KMCC}$ with $\mathbb{P} \in \mathcal{X}$ a pretame forcing. Then $(M, \mathcal{X})[G]$ satisfies Class Collection.*

Proof. Suppose that $(M, \mathcal{X})[G] \models \forall \alpha \exists Y \varphi(\alpha, Y, A)$, for some class A . We want to find a class $C \in \mathcal{X}[G]$ so that $(M, \mathcal{X})[G] \models \forall \alpha \varphi(x, (C)_x, A)$. Take $p \in G$ forcing $\forall \alpha \exists Y \varphi(x, Y, \dot{A})$, where \dot{A} is a \mathbb{P} -name for A . Fix an ordinal α . Then $p \Vdash \exists Y \varphi(\dot{\alpha}, Y, \dot{A})$.

I claim there is a class name \dot{Y}_α so that $p \Vdash \varphi(\alpha, \dot{Y}_\alpha, \dot{A})$. To see this: let D be the dense below p class of conditions $q \leq p$ so that $q \Vdash \varphi(\alpha, \dot{Y}_q, \dot{A})$ for some class name \dot{Y}_q . This D is in \mathcal{X} by an instance of Comprehension. Now construct a maximal antichain $N \subseteq D$ by recursion using a bijection $b : \text{Ord} \rightarrow D$.

¹⁸Recall that Π_ω^1 -Comprehension is another name for the full Comprehension schema.

First, put $b(0)$ into N . Continuing upward, we include $b(\xi)$ in N if and only if $b(\xi)$ is incompatible with all the $b(\zeta)$ for $\zeta < \xi$ we have already committed to being in N . We can now use this antichain N to build the desired \dot{Y}_α via a mixing argument. Namely,

$$\dot{Y}_\alpha = \{(\sigma, r) : \exists q, q' \in N \exists p (\sigma, p) \in \dot{Y}_q \text{ and } r \leq q', p\}.$$

Note that this definition uses Class Collection to pick the \dot{Y}_q corresponding to each $q \in N$. Now given these names \dot{Y}_α for each α we can put them together to get a name for a class C so that the α th slice of C is the interpretation of \dot{Y}_α . \square

Corollary 1.35. *Let $G \subseteq \mathbb{P}$ be generic for $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$, for $k \geq 1$, with $\mathbb{P} \in \mathcal{X}$ a pretame forcing. Then $(M, \mathcal{X})[G]$ satisfies $\Sigma_k^1\text{-Class Collection}$.*

Proof. In the above argument we used Comprehension to get D and Class Collection to choose the \dot{Y}_q . This instance of Comprehension used $p \Vdash \varphi$, which is Σ_k^1 -definable for a Σ_k^1 formula φ . Similarly, the instance of Class Collection is also Σ_k^1 . So both go through in this context. \square

I end this section with an open question. We have seen that strong second-order set theories are preserved by tame forcing, as is GBC. What about intermediate theories?

Question 1.36. *Is ETR preserved by tame forcing?*

1.4 Some basic constructions

In this section I survey some basic constructions for models of second-order set theories. I will focus on models of the weak theories, as the constructions for stronger theories are not so basic.

Observation 1.37. *Let $M \models \text{ZFC}$ and $\text{Def}(M)$ consist of the definable, possibly with parameters, subsets of M . Then $(M, \text{Def}(M)) \models \text{GBC}$. If additionally, $\text{Def}(M)$ contains a global well-order of M then $(M, \text{Def}(M)) \models \text{GBC}$. Similarly, if $M \models \text{ZFC}^-$ then $(M, \text{Def}(M)) \models \text{GBC}^-$, or GBC^- in case M has a definable global well-order of ordertype Ord .*

Of course, M may not have a definable global well-order. So in general it requires more to get a GBC-realization (or GBC^- -realization) for M .

Definition 1.38. Let M be a model of first-order set theory. Say that $A \subseteq M$ is *amenable to M* if $A \cap x \in M$ for all $x \in M$.

Definition 1.39. Let T be a second-order set theory and M be a model of first-order set theory. Say that $A \subseteq M$ is *T -amenable to M* if there is a T -realization \mathcal{X} for M with $A \in \mathcal{X}$.

A special case of interest is when T is GBC or GBC^- .

Observation 1.40. *A is GBC^- -amenable to M if and only if (M, A) satisfies the Separation and Replacement schemata for formulae in the expanded language.*

Proof. (\Rightarrow) Because GBC^- includes Elementary Comprehension and Class Replacement.

(\Leftarrow) Then $\text{Def}(M; A)$ is a GBC^- -realization for M . \square

It is obvious that if G is a GBC -amenable global well-order of M then $(M, \text{Def}(M; G)) \models \text{GBC}$. Accordingly, to show that M is GBC -realizable we want to find such a global well-order.

Theorem 1.41 (Folklore). *Suppose $M \models \text{ZFC}$ is countable. Then M is GBC -realizable. In general, if $(M, \mathcal{X}) \models \text{GBC}$ is countable then there is $\mathcal{Y} \supseteq \mathcal{X}$ a GBC -realization for M .*

The tool used in this proof will be Cohen-forcing to add a generic subclass of Ord .

Definition 1.42. The forcing to add a Cohen-generic subclass of Ord , denoted $\text{Add}(\text{Ord}, 1)$, consists of all set-sized partial functions from Ord to 2 , ordered by reverse inclusion.

Lemma 1.43. *Over GBC , the forcing $\text{Add}(\text{Ord}, 1)$ is tame.*

Proof. First let us see that $\text{Add}(\text{Ord}, 1)$ is pretame. Because $\text{Add}(\text{Ord}, 1)$ is $< \kappa$ -distributive for every κ ,¹⁹ below any set-sized collection of open dense subclasses of $\text{Add}(\text{Ord}, 1)$ we can find a single dense subclass. That is, if $\langle D_i : i \in a \rangle$ is a set-indexed collection of open dense subclasses then we can find a dense subclass D so that any generic which meets D must meet all the D_i . As such, we may assume without loss that that we are dealing with a single dense subclass. That is, the setup is that we have a dense class D and a condition p . We want to find a condition $q \leq p$ and a set $d \subseteq D$ which is predense $\leq q$. But this is trivial: take $q \leq p$ which meets D and let $d = \{q\}$.

Finally, $\text{Add}(\text{Ord}, 1)$ preserves Powerset because forcing with it does not add any new sets, which in turn is because it is $< \kappa$ -closed for every cardinal κ . \square

This appeared in the proof of the above lemma, but it is important enough to be stated on its own.

Observation 1.44. *Over GBC , forcing with $\text{Add}(\text{Ord}, 1)$ does not add new sets.* \square

We can now prove theorem 1.41.

Proof of theorem 1.41. It suffices to prove the more general result, since we get the other result by considering $\mathcal{X} = \text{Def}(M)$ to consist of the (first-order) definable classes.

We obtain \mathcal{Y} by forcing over (M, \mathcal{X}) with $\text{Add}(\text{Ord}, 1)$. That is, let $C \subseteq \text{Ord}^M$ be generic over (M, \mathcal{X}) for $\text{Add}(\text{Ord}, 1)$. Then $\mathcal{Y} = \mathcal{X}[C]$. So $(M, \mathcal{Y}) \models \text{GBC}$ and in particular the class C is GBC -amenable to M . We want to see that we can define a global well-order from C . To do this we use that every set of ordinals in M is coded into C . This is because sets can be coded as sets of ordinals—to code x take an isomorphic copy of $\in \upharpoonright \text{TC}(\{x\})$ as a set of pairs of ordinals, which can be coded as a set of ordinals via a pairing function—and so by density every set is coded into C . Thus, we can define a global well-order $<_C$ as $x <_C y$ if the first place x is coded into C comes before the first place y is coded into C . \square

Corollary 1.45 (Folklore). *GBC is conservative over ZFC . That is, if φ is a first-order sentence in the language of set theory then $\text{GBC} \vdash \varphi$ if and only if $\text{ZFC} \vdash \varphi$.*

Proof. Suppose otherwise. Then there is a countable $M \models \text{ZFC}$ which satisfies $\neg\varphi$ while every model of GBC satisfies φ . But M has a GBC -realization \mathcal{X} and $(M, \mathcal{X}) \models \text{GBC} + \neg\varphi$, a contradiction. \square

We can also get a version of theorem 1.41 that applies to models of ZFC^- , but we need a little more from our model. In the ZFC context we knew that $\text{Add}(\text{Ord}, 1)$ did not add sets because it was $< \kappa$ -closed for every κ . Let me quickly sketch the argument so we know what we would like to generalize.

Fix a set name \dot{a} . Without loss of generality we may assume that \dot{a} gives a set of ordinals. That is, suppose $C \subseteq \text{Add}(\text{Ord}, 1)$ is generic and consider $p \in C$ so that $p \Vdash \dot{a} \subseteq \check{\kappa}$. Let us now see that the ground model can interpret \dot{a} . Start with $p_0 = p$. Given p_α for $\alpha < \kappa$ extend p_α to $p_{\alpha+1}$ which decides whether $\check{\alpha} \in \dot{a}$. And at limits use $< \kappa$ -closure to continue the construction. And because $\text{Add}(\text{Ord}, 1)$ is $< \kappa^+$ -closed we get p_κ below all the p_α 's for $\alpha < \kappa$. Moreover, we may make the choices so that p_α is always in C , since the classes we are meeting are dense. So p_κ , which is in the ground model, contains all the information that the generic uses to interpret \dot{a} . So \dot{a}^C is in the ground model.

The same argument will work for a model of ZFC^- without a largest cardinal. But if the model does have a largest cardinal, that will not work. For concreteness, suppose we are working over a model of ZFC^- plus “every set is countable”. To show that $\text{Add}(\text{Ord}, 1)$ we want to be able to make countably many choices according to some definable procedure and have that those choices cohere. This would give us the desired p_ω which has all the information needed to interpret the name \dot{a} . That is, we want ω -Dependent Choice for definable procedures. In general, if our model has larger cardinals, then we want this but with ω replaced with the largest cardinal in the model.

¹⁹ $\text{Add}(\text{Ord}, 1)$ is $< \kappa$ -distributive because it is $< \kappa$ -closed. Checking this is an easy exercise.

Definition 1.46. Work in the context of ZFC^- and let κ be a cardinal. The principle of *Definable κ -Dependent Choice* asserts the following: if T is a definable tree of sequences of length $< \kappa$ so that for all $\alpha < \kappa$ each node in T of length α has a successor in T , then T has a branch.

Remark 1.47. Observe that the branch is a set. If it were a definable class, then because κ is a set and the branch has length κ , Replacement would imply that the branch is a set.

Definable Dependent Choice is not a theorem of ZFC^- . S. Friedman and Gitman [FG17] produced a model of ZFC^- + “every set is countable” where Definable ω -Dependent Choice fails.

If our model of ZFC^- with a largest cardinal κ satisfies definable κ -dependent choice then $\text{Add}(\text{Ord}, 1)$ over that model will be pretame and not add any new sets. So in that case we can force with $\text{Add}(\text{Ord}, 1)$ to add a global well-order without adding any new sets. Indeed, the two are equivalent.

Proposition 1.48. *Let $M \models \text{ZFC}_1^-$ have a largest cardinal κ . Then the following are equivalent:*

1. M satisfies Definable κ -Dependent Choice; and
2. M admits a GBC^- -amenable global well-order.

Proof sketch. We sketched $(1 \Rightarrow 2)$ above. For $(2 \Rightarrow 1)$, suppose $G : \text{Ord}^M \rightarrow V^M$ is GBC^- -amenable to M . Work in $(M, \mathcal{X}) \models \text{GBC}^-$ with $G \in \mathcal{X}$. We can use G to make choices along κ -trees. Since \mathcal{X} contains every definable class this yields that M satisfies Definable κ -Dependent Choice. \square

For a related result, Gitman, Hamkins, and Johnstone [GHJ16] showed that, over ZFC^- , Definable ω -Dependent Choice is equivalent to the Reflection schema, i.e. the schema asserting that for every formula $\varphi(x, a)$ and set a that there is a transitive set $t \ni a$ so that $\varphi(x, a)$ reflects to t .

Finally, let me observe that we cannot get a version of corollary 1.45 that works for ZFC^- .

Corollary 1.49. *The theory GBC^- is not conservative over ZFC^- .*

Proof. As we just saw, GBC^- proves the Definable κ -Dependent Choice schema. But Friedman and Gitman showed that this schema is not a theorem of ZFC^- . \square

Having investigated what happens without powerset, let us now turn to the uncountable. We saw in theorem 1.41 that every countable model of ZFC is GBC -realizable. This does not generalize to the uncountable. Let us see why.

Definition 1.50. A model $M \models \text{ZFC}$ is called *rather classless* if every amenable $X \subseteq M$ is definable.

Necessarily, every rather classless model has uncountable cofinality. If $X \subseteq \text{Ord}^M$ has ordertype ω and is cofinal then it is amenable, because its intersection with an initial segment of Ord^M is finite, but not definable or danger of contradicting Replacement.

Theorem 1.51 (Keisler [Kei74], Shelah [She78]). *Any countable model of ZFC has an elementary rank extension²⁰ to a rather classless model.²¹*

As a consequence if $M \models \text{ZFC}$ does not have a definable global well-order then any rather classless elementary rank extension of M will not have a GBC -realization. Keisler’s theorem also applies to ZFC^- , so there are models of ZFC^- which fail to have a GBC^- realization.

One might hope that theorem 1.41 could be generalized to stronger second-order set theories. Of course, this could not work for all countable models of ZFC , as these stronger theories are not conservative over ZFC . But one might hope that any countable model of a strong enough theory is, say, KM -realizable.

One’s hopes are in vain.

²⁰If $M \subseteq N$ are models of ZFC then say N is a *rank extension* of M if every new set has a higher rank. For example, if $\kappa < \lambda$ are inaccessible, then V_λ is a rank extension of V_κ .

²¹Keisler showed this theorem under the assumption of \diamond and the assumption of \diamond was later eliminated by Shelah.

Proposition 1.52 (Folklore). *Consider $M \models \text{ZFC}$ so that the truth predicate for M is GBC -amenable. Then there is a club of ordinals $\alpha \in \text{Ord}^M$ so that $V_\alpha^M \prec M$.*

Before moving to the proof, let us recall the Tarskian definition of a truth predicate.

Definition 1.53. Let M be a model of first-order set theory. A *truth predicate* or *satisfaction class* for M is a class $\text{Tr} \subseteq M$ of pairs (φ, \bar{a}) satisfying the following recursive requirements.

1. If φ is atomic then $(\varphi, \bar{a}) \in \text{Tr}$ if and only if $\varphi(\bar{a})$ gives a true fact about M . That is, Tr should declare $a \in b$ to be true if and only if $M \models a \in b$ and declare $a = b$ to be true if and only if $M \models a = b$.
2. $(\varphi \vee \psi, \bar{a})$ is in Tr if and only if (φ, \bar{a}) or (ψ, \bar{a}) are in Tr .
3. $(\neg\varphi, \bar{a})$ is in Tr if and only if (φ, \bar{a}) is not in Tr .
4. $(\exists x \varphi(x), \bar{a})$ is in Tr if and only if there is $b \in M$ so that $(\varphi, b \frown \bar{a})$ is in Tr .

We are interested in adding Tr as a class to M , so in the case that M is ω -nonstandard let us explicitly require that Tr measures the ‘truth’ of every nonstandard formula.²²

Observe that it is a first-order property of a class whether it is a truth predicate, so it does not depend upon what classes are in the model. Also observe that GBC proves the truth predicate is unique. If two different classes both satisfy the definition then there must be a minimal place where they disagree on the truth of (φ, \bar{a}) . But they agree on every previous stage so they must agree on the truth of (φ, \bar{a}) , a contradiction.

Proof of proposition 1.52. This follows from an instance of the Montague reflection principle. Let Tr be the truth predicate for M . We can use Tr as a parameter in the formula we are reflecting precisely because it is GBC -amenable. Namely, reflect to find a club of α so that $(V_\alpha^M; \text{Tr} \cap V_\alpha^M) \prec_{\Sigma_1} (M; \text{Tr})$. Then $V_\alpha^M \prec M$ because by elementarity $(V_\alpha^M; \text{Tr} \cap V_\alpha^M) \models (\varphi, \bar{a}) \in \text{Tr} \cap V_\alpha^M$ if and only if $(M; \text{Tr}) \models (\varphi, \bar{a}) \in \text{Tr}$ if and only if $M \models \varphi(\bar{a})$.²³ \square

Corollary 1.54. *Let $T \supseteq \text{GBC}^-$ be a second-order set theory which proves the existence of the truth predicate for the first-order part. Then no first-order theory characterizes which countable models are T -realizable.*

In particular this works for $T = \text{KM}$, $T = \text{GBC} + \text{ETR}$, or even $T = \text{GBC} + \text{ETR}_\omega$.

Proof sketch. Let S be any consistent first-order set theory. Then, by standard results about nonstandard models, there is $M \models S$ which is ω -nonstandard but not recursively saturated.²⁴ It follows that M does not admit an amenable truth predicate, so it cannot be T -realizable. \square

The reader may find this nonstandard trick to be unsatisfactory. But we get a version of corollary 1.54 for ω -models or even transitive models.

Corollary 1.55. *Let $T \supseteq \text{GBC}^-$ be a second-order set theory which proves the existence of the truth predicate for the first-order part. Suppose that T has an ω -model. Then no first-order set theory characterizes which ω -standard models are T -realizable. Moreover, suppose T has a transitive model. Then no first-order set theory characterizes which transitive models are T -realizable.*

²²In the literature such a class Tr is called a *full satisfaction class*. See also chapter 3 for a discussion of truth predicates over ω -nonstandard models.

²³There is a minor subtlety. Namely, what happens if M is ω -nonstandard? Then Tr makes assertions about the ‘truth’ of nonstandard formulae, and for such formulae φ it does not make sense externally to ask whether $M \models \varphi$. But this is not an issue because Tr must be correct about the truth of standard formulae, as can be checked by an easy induction external to the model.

²⁴For a definition of recursive saturation, see chapter 3.

For trivial reasons, we need the assumption that T has an ω -model (or transitive model for the moreover). If T has no ω -models then it is easy to get a first-order set theory characterizing which ω -models are T -realizable—take your favorite inconsistent theory.

Proof. Suppose S is some first-order set theory. If $T + S$ has no ω -model (or no transitive model, for that case) then S cannot characterize which ω -models (or transitive models) are T -realizable. So assume that $T + S$ has an ω -model (M, \mathcal{X}) . (Or, for the transitive case, assume (M, \mathcal{X}) is transitive. We will find an elementary model of M which is not T -realizable, establishing that satisfying S cannot ensure a model is T -realizable.

The truth predicate for M is in \mathcal{X} so it must be GBC-amenable. Now take the least $\alpha \in \text{Ord}^M$ so that $V_\alpha^M \prec M$. Then $N \models S$ and is ω -standard (or transitive, if M is transitive). I claim that $N = V_\alpha^M$ is not T -realizable. Otherwise, by proposition 1.52 there is a club of ordinals $\beta \in \text{Ord}^N$ so that $V_\beta^N \prec N$. But $V_\beta^N = V_\beta^M$, because M is a rank extension of N . So then $V_\beta^M \prec V_\alpha^M \prec M$, contradicting the leastness of α . \square

Remark 1.56. This argument uses essentially that M (and hence also N) is an ω -model. Suppose $(M, \mathcal{X}) \models T$, where T is as in the corollary, is countable and ω -nonstandard. Let $N = V_{\alpha_0}^M$, where $\alpha_0 \in M$ is least in the club of ordinals α_0 so that $V_{\alpha_0}^M \prec M$, where elementarity here is according to the truth predicate in M . In particular, N is also countable and ω -nonstandard and has the same theory and standard system²⁵ as M does. So by a back-and-forth argument we can show that in fact N and M are isomorphic. So because M is T -realizable, so must N be T -realizable.

To finish off this section, let us see that see that T -amenability for different choices of T can give different notions. We saw that, for countable models, being GBC-amenable is equivalent to being GBC-amenable. But for some T , being T -amenable to M is stronger than being GBC-amenable.

Theorem 1.57. *Let $M \models \text{ZFC}$ be a countable ω -model. Then there is $A \subseteq M$ which is GBC-amenable to M but is not $(\text{GBC} + \text{ETR}_\omega)$ -amenable to M . That is, there is a GBC-realization \mathcal{X} for M with $A \in \mathcal{X}$ but no $(\text{GBC} + \text{ETR}_\omega)$ -realization \mathcal{Y} for M can have $A \in \mathcal{Y}$.*

Proof. I will show a stronger fact, from which this theorem will immediately follow. Namely, I will show that there is $G \subseteq M$ which is GBC-amenable but no GBC-realization for M can contain both G and the truth predicate for M .²⁶ This G will be a carefully constructed Cohen-generic subclass of Ord . I claim that there is a sequence $\langle D_i : i \in \text{Ord} \rangle$ of definable dense subclasses of $\text{Add}(\text{Ord}, 1)$ so that (1) the sequence is definable from the truth predicate and (2) meeting every D_i is sufficient to guarantee that a filter is generic over $(M, \text{Def}(M))$. It is obvious that the truth predicate can define a sequence of all the definable subclasses of $\text{Add}(\text{Ord}, 1)$. Namely, from the truth predicate can be defined the sequence indexed by (φ, a) where the (φ, a) -th dense subclass is the one consisting of all conditions p so that $\varphi(p, a)$ holds (or trivial if $\varphi(x, a)$ does not define a dense subclass). But this sequence is not of the correct ordertype. However, we can use that $\text{Add}(\text{Ord}, 1)$ is $< \kappa$ -distributive for every cardinal κ to get the ordertype to be Ord . Namely let D_i be below all the open dense subclasses which are definable from parameters in V_i . Then $\langle D_i : i \in \text{Ord} \rangle$ is as desired.

We will now use this sequence to define our G . Externally to the model fix an ω -sequence cofinal in Ord^M . Think of this sequence as an Ord^M -length binary sequence $\langle b_i : i \in \text{Ord}^M \rangle$, consisting mostly of zeros with ones showing up rarely. This sequence is amenable to M , since its initial segments have only finitely many ones. On the other hand, it is not GBC-amenable since from this sequence we can define a cofinal map $\omega \rightarrow \text{Ord}$, contradicting an instance of Replacement.

Build G in Ord^M many steps, with partial piece g_i at the i th step. We start with $g_0 = \emptyset$. Given g_i , let $g'_i = g_i \hat{\ } \langle b_i \rangle$. Then get g_{i+1} by extending g'_i to meet D_i , where we do this in the minimal possible length. (If there is more than one way to meet D_i with minimal length, then pick arbitrarily.) And if i is limit

²⁵The *standard system* of an ω -nonstandard model M is the collection of all reals coded in M , i.e. the set of all $x \subseteq \omega$ so that there is $y \in \omega^M$ so that $y \cap \omega = x$.

²⁶This is why I require M to be an ω -model. If M is ω -nonstandard then there will be many different classes which satisfy the definition of being a truth predicate. See chapter 3 for more detail.

then $g_i = \bigcup_{j < i} g_j$. For each $i \in \text{Ord}$, we have that $g_i \in M$ because $\langle b_i : i \in \text{Ord} \rangle$ is amenable to M and g_i can be defined from an initial segment of this sequence. Finally, set $G = \bigcup_{i \in \text{Ord}} g_i$. Then G is generic over $(M, \text{Def}(M))$, since it meets every D_i . But hidden within G is this bad sequence $\langle b_i \rangle$. It is well hidden, but with the truth predicate we can snoop it out.

Suppose towards a contradiction that $(M, \mathcal{X}) \models \text{GBC}$ contains both G and the truth predicate for M . Then $\langle D_i : i \in \text{Ord} \rangle \in \mathcal{X}$. We will see that $\langle b_i \rangle \in \mathcal{X}$ by inductively determining each b_i . First, b_0 is the first bit of G . We then know the minimal length we have to extend $\langle b_0 \rangle$ to meet D_0 . We can use this to recover g_1 and then discover b_1 as the first bit in G after g_1 . We then repeat this process, using $\langle D_i \rangle$ to recover g_2 then get b_2 , and so on. So we can define $\langle b_i \rangle$ from G and the truth predicate, so $\langle b_i \rangle \in \mathcal{X}$. But then (M, \mathcal{X}) cannot satisfy Replacement, a contradiction. \square

1.5 GBC-realizations of a countable model

In this section we look at the structure of the GBC-realizations for a fixed model. At the end of the section I will discuss the extent to which the results generalize for theories stronger than GBC.

Definition 1.58. Let M be a model of first-order set theory and let T be a second-order set theory. Set $T\text{-}\mathfrak{Rc}(M) = \{\mathcal{X} \subseteq \mathcal{P}(M) : \mathcal{X} \text{ is a } T\text{-realization for } M\}$. Then $T\text{-}\mathfrak{Rc}(M)$ is a partial order under \subseteq .

Of course, M may fail to be T -realizable and thus $T\text{-}\mathfrak{Rc}(M)$ may be empty. But theorem 1.41 implies that if M is countable then $\text{GBC-}\mathfrak{Rc}(M)$ is not empty. If we move to the uncountable, however, then $\text{GBC-}\mathfrak{Rc}(M)$ may be anemic. If M is rather classless then $|\text{GBC-}\mathfrak{Rc}(M)| \leq 1$; if such M has a definable global well-order then it has a single GBC-realization—namely its definable classes—otherwise it will have no GBC-realization at all.

But if $M \models \text{ZFC}$ is countable then it will have continuum many GBC-realizations. To prove in theorem 1.41 that countable models of ZFC have GBC realizations we added a Cohen-generic subclass of Ord . But there are continuum many different generic Cohen subclasses of Ord^M for countable M . This gives continuum many different GBC-realizations for M .

Proposition 1.59. *Let $M \models \text{ZFC}$ be countable. Then there are continuum many different subclasses of Ord^M which are Cohen-generic over $(M, \text{Def}(M))$.*

Proof. Consider the full binary tree $T = {}^{<\omega}2$ of finite binary sequences. We will construct a family of Cohen-generic subclasses of Ord^M for each branch through T . This construction is done along T . Order, in ordertype ω , the dense subclasses D_i of $\text{Add}(\text{Ord}, 1)^M$ which are in $\text{Def}(M)$. Start with $p_\emptyset = \emptyset$. Assume that we have already defined p_s for s a node in T . Then, to get $p_{s \smallfrown i}$ extend $p_s \smallfrown i$ to a condition in D_ℓ , where ℓ is the length of s . Then if B is a branch through T we have that $C_B = \bigcup_{s \in B} p_s$ is Cohen-generic over $(M, \text{Def}(M))$, as it met every dense class. And if $B \neq B'$ are distinct branches this is because there is a node $s \in T$ so that $s \smallfrown 0 \in B$ and $s \smallfrown 1 \in B'$. So $C_B \neq C_{B'}$ because they extend p_s in incompatible ways. \square

An intriguing question is whether there is anything between these extremes of continuum many GBC-realizations and ≤ 1 GBC-realization.

Question 1.60. *Is there (necessarily uncountable) $M \models \text{ZFC}$ so that $|\text{GBC-}\mathfrak{Rc}(M)| = 2$? What about $|\text{GBC-}\mathfrak{Rc}(M)| = n$ for finite n ? What about $|\text{GBC-}\mathfrak{Rc}(M)| = \omega$? In general, what cardinals κ are the cardinality of $\text{GBC-}\mathfrak{Rc}(M)$ for some M ?*

The rest of this section will be confined to looking at countable models of ZFC as there it can be shown that $\text{GBC-}\mathfrak{Rc}(M)$ has a rich structure.

Let us begin with some basic properties.

Theorem 1.61. *Let $M \models \text{ZFC}$ be countable. Then $\text{GBC-}\mathfrak{Rc}(M)$ satisfies the following properties.*

1. *If $\{\mathcal{X}_i : i \in I\} \subseteq \text{GBC-}\mathfrak{Rc}(M)$ has a lower bound in $\text{GBC-}\mathfrak{Rc}(M)$ it has a greatest lower bound.*

2. If $\{\mathcal{X}_i : i \in I\} \subseteq \text{GBC-}\mathfrak{Rc}(M)$ has an upper bound in $\text{GBC-}\mathfrak{Rc}(M)$ it has a least upper bound.
3. If M does not have a definable global well-order, then there are pairs of elements of $\text{GBC-}\mathfrak{Rc}(M)$ without a lower bound. On the other hand, if M does have a definable global well-order then any $\{\mathcal{X}_i : i \in I\} \subseteq \text{GBC-}\mathfrak{Rc}(M)$ has a lower bound.
4. There are pairs of elements of $\text{GBC-}\mathfrak{Rc}(M)$ without an upper bound.
5. If $\langle \mathcal{X}_i : i \in I \rangle$ is an increasing chain then it has a supremum.
6. There are maximal elements of $\text{GBC-}\mathfrak{Rc}(M)$. That is, there is $\mathcal{X} \in \text{GBC-}\mathfrak{Rc}(M)$ so that there is no $\mathcal{Y} \in \text{GBC-}\mathfrak{Rc}(M)$ with $\mathcal{Y} \supsetneq \mathcal{X}$.

Proof. (1) The infimum is $\mathcal{Y} = \bigcap_{i \in I} \mathcal{X}_i$. We want to see that $(M, \mathcal{Y}) \models \text{GBC}$. Extensionality is free, as is Replacement since \mathcal{Y} is contained inside a GBC -realization. To see that \mathcal{Y} satisfies Global Choice, fix any \mathcal{Z} a lower bound for the \mathcal{X}_i . Then $\mathcal{Y} \supseteq \mathcal{Z}$ so the global well-order in \mathcal{Z} is in \mathcal{Y} . Finally, we want to see that \mathcal{Y} satisfies Elementary Comprehension. But each \mathcal{X}_i is closed under first-order definability, so their intersection must also be closed under first-order definability.

(2) The supremum is the intersection of all the upper bounds. It follows from (1) that this gives an element of $\text{GBC-}\mathfrak{Rc}(M)$.

(3) Suppose M does not have a definable global well-order. Take C, D subclasses of Ord^M which are mutually Cohen-generic over $(M, \text{Def}(M))$. Then $\text{Def}(M; C)$ and $\text{Def}(M; D)$ are GBC -realizations for M . But their intersection is $\text{Def}(M)$, by mutual genericity, which does not have a global well-order.

For the other case, suppose M does have a definable global well-order. Then $\text{Def}(M)$ is a lower bound for a subset of $\text{GBC-}\mathfrak{Rc}(M)$.

(4) We will construct C, D Cohen-generic over $(M, \text{Def}(M))$ so that no GBC -realization for M can contain both C and D . This will establish that $\text{Def}(M; C)$ and $\text{Def}(M; D)$ are elements of $\text{GBC-}\mathfrak{Rc}(M)$ without an upper bound.

Fix $B : \text{Ord}^M \rightarrow 2$ so that $\{i : B(i) = 1\}$ has ordertype ω and is cofinal in Ord^M . Such B exists because M is countable. But no GBC^- -realization for M can contain this bad class B because B reveals that Ord^M has countable cofinality. We will construct C and D so that together they code B . This is a construction in Ord^M many steps, defining c_i and d_i for $i \in \text{Ord}^M$. Order the dense subclasses of $\text{Add}(\text{Ord}, 1)$ as $\langle D_i : i \in \text{Ord}^M \rangle$.

- Set $c_0 = d_0 = \emptyset$.
- Given c_i and d_i set $c'_{i+1} = c_i \hat{\ } \langle 0 \dots 0 \rangle \hat{\ } \langle 1, B(i) \rangle$, where the sequence of 0's has length chosen so that $c_i \hat{\ } \langle 0 \dots 0 \rangle$ has the same length as d_i . Then get c_{i+1} by extending c'_{i+1} to meet D_i . Next, let $d'_{i+1} = d_i \hat{\ } \langle 0 \dots 0 \rangle \hat{\ } \langle 1 \rangle$, where the sequence of 0's has length chosen so that $d_i \hat{\ } \langle 0 \dots 0 \rangle$ has the same length as c'_{i+1} . Then get d_{i+1} by extending d'_{i+1} to meet D_i .
- If i is limit then $c_i = \bigcup_{j < i} c_j$ and $d_i = \bigcup_{j < i} d_j$.

Finally, set $C = \bigcup_i c_i$ and $D = \bigcup_i d_i$. By construction C and D are Cohen-generic over M . Suppose towards a contradiction that \mathcal{X} is a GBC^- -realization for M with $C, D \in \mathcal{X}$. Let us see that $B \in \mathcal{X}$, a contradiction. Namely, \mathcal{X} can inductively recover $B(i)$, c_i , and d_i from C and D . First, $c_0 = d_0 = \emptyset$. Now given c_i we find $B(i)$ by looking at the bit in C after the first 1 after the block of 0s in C starting after the end of c_i . We also get d_i by using that block of 0s to tell us how far in D we need to go to get d_i . Next, looking at the block of 0s in D starting after the end of d_i tells us how long c_i was extended in C to get c_{i+1} . We then continue this process, getting $B(i+1)$, d_{i+1} , c_{i+2} , and so on.

(5) The supremum is $\mathcal{Y} = \bigcup_{i \in I} \mathcal{X}_i$. We need to see that \mathcal{Y} is a GBC -realization for M . We know for free that (M, \mathcal{Y}) satisfies Extensionality. It satisfies Global Choice because each \mathcal{X}_i contains a global well-order. To see that it satisfies Replacement, suppose $F \in \mathcal{Y}$ witnesses a failure of Class Replacement. But then $F \in \mathcal{X}_i$ for some i so (M, \mathcal{X}_i) fails to satisfy Class Replacement, a contradiction. To check Elementary

Comprehension we want to see that \mathcal{Y} is closed under first-order definability. Towards this fix $\bar{A} \in \mathcal{Y}$. Then $\bar{A} \in \mathcal{X}_i$ for some i . This uses that the \mathcal{X}_i are linearly ordered by \subseteq . So if B is definable from \bar{A} then $B \in \mathcal{X}_i \subseteq \mathcal{Y}$, as desired.

(6) Combine (5) and Zorn's lemma. \square

Let me remark that (3) of theorem 1.61 gives us a criterion in terms of the theory of M for when $\text{GBC-}\mathfrak{Rc}(M)$ has a least element.

Corollary 1.62. *Let $M \models \text{ZFC}$ be countable. Then, whether $\text{GBC-}\mathfrak{Rc}(M)$ has a least element is recognizable from the theory of M .*

Proof. We saw that $\text{GBC-}\mathfrak{Rc}(M)$ has a least element if and only if M has a definable global well order. This happens if and only if $M \models \exists x V = \text{HOD}(\{x\})$. \square

Theorem 1.61.(4) previously appeared as (a special case of) lemma 3.1 of [Mos76]. Mostowski moreover embeds the full binary tree of height ω_1 into $\text{GBC-}\mathfrak{Rc}(M)$ for countable M , thereby concluding that $|\text{GBC-}\mathfrak{Rc}(M)| \geq 2^{\omega_1}$.

What other orders can be embedded into $\text{GBC-}\mathfrak{Rc}(M)$? We will get to this question in time—see theorems 1.76 and 1.77. But first let us consider some local properties of $\text{GBC-}\mathfrak{Rc}(M)$. It will be useful to single out those $\mathcal{X} \in \text{GBC-}\mathfrak{Rc}(M)$ which are countable.

Definition 1.63. Let $M \models \text{ZFC}$ be a countable model of set theory. Set

$$\text{GBC-}\mathfrak{Rc}_{<\omega_1}(M) = \{\mathcal{X} \in \text{GBC-}\mathfrak{Rc}(M) : \mathcal{X} \text{ is countable}\}.$$

With this definition in hand, we can now see that what was proved in (4) of the previous theorem was really the following.

Corollary 1.64. *For any $\mathcal{X} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ there are $\mathcal{Y}, \mathcal{Z} \supseteq \mathcal{X}$ in $\text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ so that \mathcal{Y} and \mathcal{Z} lack an upper bound.*

Proof. Carry out the same argument, but over (M, \mathcal{X}) instead of $(M, \text{Def}(M))$. Generics can be found because \mathcal{X} is countable. \square

Cohen forcing holds the key to establishing other local properties of $\text{GBC-}\mathfrak{Rc}(M)$.

Theorem 1.65. *Every $\mathcal{X} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ has a dense extension. That is, there is $\mathcal{Y} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ so that there is a dense linear order in $\text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ between \mathcal{X} and \mathcal{Y} . Consequently, the real line with its usual order embeds into $\text{GBC-}\mathfrak{Rc}(M)$ between \mathcal{X} and \mathcal{Y} .*

Proof. Let C be generic over (M, \mathcal{X}) for the forcing to add a Cohen-generic subclass of Ord . Then $\mathcal{Y} = \mathcal{X}[C]$ is a GBC- realization for M containing both \mathcal{X} and C . Now take $X \subseteq \text{Ord}^M$ in \mathcal{X} which is unbounded. It follows from the homogeneity of Cohen-forcing that

$$C_X = \{\alpha : \text{the } \alpha\text{th element of } X \text{ is in } C\}$$

is Cohen-generic. (If you think of C as an Ord -length binary sequence then C_X is the bits which appear in X , in order.) Clearly, $C_X \in \mathcal{Y}$. On the other hand, if $\text{Ord}^M \setminus X$ is unbounded then $C \notin \mathcal{X}[C_X]$. The theorem now follows from the fact that there is a dense linear order of subclasses of Ord^M so that both they and their complements are unbounded. Namely, fix your favorite bijection b between ω and the rationals. For a rational q , put $\alpha = \omega \cdot \alpha_0 + n$ into X_q if and only if $b(n) < q$. Then each X_q and its complement is unbounded and $X_q \subseteq X_{q'}$ if and only if $q < q'$.

Namely, let k be a positive integer and $0 \leq n < 2^k$. Set $X(k, n)$ to be those ordinals which are equivalent to n modulo 2^k . For example, $X(0, 1)$ is the class of even ordinals and $X(1, 1)$ is the class of odd ordinals.

For the consequently, let e be an embedding of the rationals into $\text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ between \mathcal{X} and \mathcal{Y} . For r a real set $\mathcal{Z}_r = \bigcup_{q < r} e(q)$, which is a GBC- realization for M by theorem 1.61.(5). Then $r \mapsto \mathcal{Z}_r$ gives an embedding of the real line into $\text{GBC-}\mathfrak{Rc}(M)$ between \mathcal{X} and \mathcal{Y} . \square

We saw in theorem 1.61 that $\text{GBC-}\mathfrak{Rc}(M)$ has a least element if and only if M has a definable well-order. This result can be improved.

Theorem 1.66. *Let M be a countable model of ZFC and consider $\mathcal{X} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$. Then there is $\mathcal{Y} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ so that:*

- *If M has a definable global well-order then the only lower bound of \mathcal{X} and \mathcal{Y} in $\text{GBC-}\mathfrak{Rc}(M)$ is $\text{Def}(M)$.*
- *If M does not have a definable global well-order then \mathcal{X} and \mathcal{Y} have no lower bound in $\text{GBC-}\mathfrak{Rc}(M)$.*

Proof. Let H be (M, \mathcal{X}) -generic for $\text{Add}(\text{Ord}, 1)$. Then $\text{Def}(M; H) \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$, because meeting every dense class in \mathcal{X} implies meeting every dense class in $\text{Def}(M)$. Set $\mathcal{Y} = \text{Def}(M; H)$. Let us see that \mathcal{Y} is as desired.

Consider $A \in \mathcal{X}$ and assume that $A \in \mathcal{Y} = \text{Def}(M; H)$. Then, there is some first-order formula φ , possibly with set parameters and H as a class parameter but with no other class parameters so that $(M, \mathcal{X}[H]) \models \forall x \in A \Leftrightarrow \varphi(x, H)$. By the forcing theorem there is some $p \in H$ so that, in (M, \mathcal{X}) , we have $p \Vdash \forall x \in \check{A} \Leftrightarrow \varphi(x, \check{H})$. So for all $x \in M$, we have $x \in A$ if and only if $(M; \mathcal{X}) \models "p \Vdash \varphi(\check{x}, \check{H})"$. But this formula does not depend upon G , so the same is true in $(M; \text{Def}(M))$. Therefore, $A \in \text{Def}(M)$. So we have seen that if $A \in \mathcal{X} \cap \mathcal{Y}$ then $A \in \text{Def}(M)$, from which the conclusion of the theorem immediately follows. \square

Next we look at the opposite phenomenon from theorem 1.65. Namely, every $\mathcal{X} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ extends to some $\mathcal{Y} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ with nothing in between them. To do so, I will make use of a variant of Sacks forcing for adding new classes of ordinals.

Generalizing Sacks forcing to add generic classes of ordinals has been considered before. Kossak and Schmerl [KS06] considered what they call perfect generics, an adaptation of Sacks forcing to models of arithmetic. A set theoretic variant of their idea to add a subclass of Ord was considered by Hamkins, Linetsky, and Reitz [HLR13]. In the arithmetic case, perfect generics are used to produce minimally undefinable inductive sets over a model M of arithmetic, i.e. inductive $G \subseteq M$ so that for $A \in \text{Def}(M; G)$ either $A \in \text{Def}(M)$ or $G \in \text{Def}(M; A)$. A similar construction works in set theory to produce minimal extensions of countable models of GBC.

Theorem 1.67. *Take $\mathcal{X} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$. Then there is a GBC-realization $\mathcal{Y} \supseteq \mathcal{X}$ for M which is minimal above \mathcal{X} , in the sense that if $\mathcal{Z} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ with $\mathcal{X} \subseteq \mathcal{Z} \subseteq \mathcal{Y}$ then either $\mathcal{Z} = \mathcal{X}$ or $\mathcal{Z} = \mathcal{Y}$.*

Let me sketch the main idea before giving a proof. The desired \mathcal{Y} will be $\text{Def}(M; \mathcal{X}, G)$ where G is a specially chosen subclass of Ord^M . We need to ensure two things. First, we need that $\text{Def}(M; \mathcal{X}, G)$ is a GBC-realization for M , so we need to ensure that we satisfy the Separation and Replacement schemata with G and parameters from \mathcal{X} . Second, we need to ensure that every new class codes G , so that \mathcal{Y} is minimal above \mathcal{X} .

To achieve the first of these we will define G to be the intersection of a certain \subseteq -descending sequence $\langle \mathbb{P}_n : n \in \omega \rangle$ of Ord -height perfect binary trees. To ensure that G is a branch we will have that \mathbb{P}_n only splits above α_n , where $\langle \alpha_n : n \in \omega \rangle$ is a fixed sequence cofinal in Ord^M . That adjoining G gives a GBC-realization will be due to genericity properties of G . We can think of each of the \mathbb{P}_n as a forcing notion. While G will not be fully generic over any \mathbb{P}_n , it will be increasingly generic over each one. This will suffice.

Definition 1.68. Let $\mathbb{Q} \subseteq \mathbb{P}$ be Ord -height perfect binary trees. Say that \mathbb{Q} has the Σ_n -branch genericity property over \mathbb{P} if every branch through \mathbb{Q} (not necessarily in the model) is Σ_n -generic over \mathbb{P} , meaning that the branch meets every Σ_n -definable dense subclass of \mathbb{P} .

We will need a refinement of the forcing theorem for partial generics. Let me state what we need here, sans proof. Any Ord -height perfect binary tree \mathbb{P} , considered as a forcing notion, is tame. So we have that the forcing relation is definable. The refinement we need is that $\Vdash_{\mathbb{P}}$ restricted to the Σ_k -formulae is Σ_j definable for some $j > k$. This is proved via the usual proof of the definability lemma, but with careful bookkeeping of quantifiers. More generally, if we want $\Vdash_{\mathbb{P}}$ restricted to the Σ_k^0 -formulae in some class parameter P , then this is Σ_j^0 definable from that same parameter.

The other half of the forcing theorem is the truth lemma, which asserts that if $G \subseteq \mathbb{P}$ is generic then $(M, \mathcal{X})[G] \models \varphi$ if and only if there is $p \in G$ so that $p \Vdash \varphi$. The refinement is that this works in a restricted fashion for partial generics. That is, for each k there is $j > k$ so that if G is Σ_k -generic then for any Σ_k -formula φ we have $(M, \mathcal{X})[G] \models \varphi$ if and only if there is $p \in G$ so that $p \Vdash \varphi$. Again, this refinement comes from considering the usual proof of the truth lemma, but paying close attention to the quantifiers involved.

We may assume that the $j > k$ for both lemmata is the same, simply by taking the maximum of the two. To have a name for this j , call it *the Σ_k -genericity witness*.

Lemma 1.69. *Work with $(M, \mathcal{X}) \models \text{GBC}$. Suppose $\langle \mathbb{P}_n : n \in \omega \rangle$ is a descending sequence of Ord-height perfect binary trees so that \mathbb{P}_{n+1} has the Σ_n -branch genericity property over \mathbb{P}_n . Let $G = \bigcup \bigcap_n \mathbb{P}_n$. Then $(M, \text{Def}(M; \mathcal{X}, G)) \models \text{GBC}$.*

Proof. Suppose towards a contradiction that $F \in \text{Def}(M; \mathcal{X}, G)$ witnesses a failure of Replacement. Then F is defined from G and parameters from \mathcal{X} via a Σ_n formula. Now let $j > n$ be the Σ_n -genericity witness. So there is some $p \in G$ so that $p \Vdash_{\mathbb{P}_j}$ “ F witnesses a failure of Replacement”. Now take $H \in \mathcal{X}$ which is Σ_j -generic over \mathbb{P}_j ; such exists because Σ_j -truth is definable. Carry out the definition of F but use H instead of G . We get that $F \in \text{Def}(M; \mathcal{X}, H)$ witnesses a failure of Replacement, because H is sufficiently generic. But $\text{Def}(M; \mathcal{X}, H) = \mathcal{X}$ because $H \in \mathcal{X}$. So (M, \mathcal{X}) is not a model of GBC, a contradiction. \square

Before seeing how to ensure that \mathcal{Y} is minimal above \mathcal{X} , let us see that we can always arrange such a sequence $\langle \mathbb{P}_n \rangle$ of perfect binary trees.

Lemma 1.70. *Let $(M, \mathcal{X}) \models \text{GBC}$ be countable. Then there is $\langle \mathbb{P}_n \rangle$ of Ord-height perfect binary trees from \mathcal{X} so that $\bigcup \bigcap_n \mathbb{P}_n$ is a class of ordinals and \mathbb{P}_{n+1} has the Σ_n -branch genericity property over \mathbb{P}_n .*

This sequence will not be (coded) in \mathcal{X} , though each tree in the sequence will be in \mathcal{X} .

Proof. Fix $\langle \alpha_n : n \in \omega \rangle$ cofinal in Ord^M . Start with $\mathbb{P}_0 = {}^{<\text{Ord}}2$ the full binary tree. Assume we have already found $\mathbb{P}_n \in \mathcal{X}$. We define \mathbb{P}_{n+1} by defining a certain embedding $g : {}^{<\text{Ord}}2 \rightarrow \mathbb{P}_n$. Closing $\text{ran } g$ downward in \mathbb{P}_n will give \mathbb{P}_{n+1} . We define g by a set-like recursion of height Ord. Fix in advance a global well-order. Let $g(0)$ be the first node according to this well-order which has length $\geq \alpha_n$. This will ensure that the intersection of the \mathbb{P}_n gives a branch. At limit stages, take unions. If we have already defined $g(s)$, then let t be the least, according to the global well-order, splitting node in \mathbb{P}_n which extends $g(s)$ and decides the $\text{len } s$ -th instance (according to the global well-order) of the universal Σ_n -formula. Then set $g(s \smallfrown i) = t \smallfrown i$.

It is clear that $\mathbb{P}_{n+1} \in \mathcal{X}$, because we defined it from parameters from \mathcal{X} . It is also clear that \mathbb{P}_{n+1} is a perfect tree. It has the Σ_n -branch genericity property over \mathbb{P}_n because any Σ_n -formula is decided by a long enough node in \mathbb{P}_{n+1} . \square

It remains to see how to ensure the minimality of the extension. This is encapsulated by the following lemma, which is a set theoretic counterpart to a result from section 6.5 of [KS06].

Lemma 1.71 (Minimality lemma). *Let $\varphi(x)$ be a formula in the forcing language and $\mathbb{P} \in \mathcal{X}$ be a perfect subtree of the full binary tree. Then there is $\mathbb{Q} \subseteq \mathbb{P}$ in \mathcal{X} so that one of the two cases holds:*

1. *There is an ordinal α so that for all ordinals ξ we have that all $p \in \mathbb{Q}$ of length greater than α decide $\varphi(\check{\xi})$ (in \mathbb{P}) the same.*
2. *For every ordinal α there is $\beta > \alpha$ so that if $p, q \in \mathbb{Q}$ both have length β and $p \upharpoonright \alpha = q \upharpoonright \alpha$ then there is an ordinal ξ so that p and q decide $\varphi(\check{\xi})$ differently (in \mathbb{P}).*

Proof. Fix k so that φ is a Σ_k formula. Take $\mathbb{P}' \subseteq \mathbb{P}$ a k -deciding subtree for \mathbb{P} . We may assume that there is a function $f : \mathbb{B} \rightarrow \mathbb{P}'$ which embeds the full binary tree onto the splitting nodes of \mathbb{P}' and that $f(s)$ decides $\varphi(\text{len } s)$. There are two cases. The first is that there is some $s \in \mathbb{B}$ so that for every $t, t' \succ_{\mathbb{B}} s$ if $\text{len } t = \text{len } t'$ then $f(t)$ and $f(t')$ decide $\varphi(\check{\xi})$ the same for all ordinals ξ . In this case, set $\mathbb{Q} = \mathbb{P}' \upharpoonright f(s)$ and get the first conclusion in the lemma.

The second case is that this does not happen for any $s \in \mathbb{B}$. In this case, we can inductively define a $g : \mathbb{B} \rightarrow \mathbb{P}'$ as follows:

- Set $g(0) = f(0)$.
- Set $g(s \smallfrown 0) = p_0$ and $g(s \smallfrown 1) = p_1$, where p_0, p_1 are least (according to a fixed global well-order) so that $\text{len } p_0 = \text{len } p_1$ and there is an ordinal ξ so that p_0 and p_1 decide $\varphi(\xi)$ differently. Such p_0 and p_1 always exist, as otherwise we would be in the previous case.
- At limit stages take unions.

Set $\mathbb{Q} = \{p \in \mathbb{P}' : \exists s \in \mathbb{B} p \leq_{\mathbb{P}'} g(s)\}$. This yields the second conclusion in the lemma. \square

Observe that we used global choice in an essential manner here. There are possibly many choices for p_0 and p_1 in the successor stage of the construction of g . In order to guarantee that $g \in \mathcal{X}$ and hence that $\mathbb{Q} \in \mathcal{X}$, we need to uniquely specify a choice.

Proof of theorem 1.67. Work with countable $(M, \mathcal{X}) \models \text{GBC}$. Fix a cofinal sequence $\langle \alpha_n \rangle$ of ordinals and an enumeration $\langle \varphi_n(x) \rangle$ of formulae in the forcing language. We construct a descending sequence of perfect trees

$${}^{<\text{Ord}}_2 = \mathbb{Q}_0 \supseteq \mathbb{P}_0 \supseteq \mathbb{Q}_1 \supseteq \mathbb{P}_1 \supseteq \dots \supseteq \mathbb{Q}_n \supseteq \mathbb{P}_n \supseteq \dots$$

so that \mathbb{P}_n has the Σ_n -branch genericity property over \mathbb{Q}_n and does not split below α_n and $\mathbb{Q}_{n+1} \subseteq \mathbb{P}_n$ is as in the previous lemma for φ_n . Set $G = \bigcup(\bigcap_n \mathbb{P}_n)$ and $\mathcal{Y} = \text{Def}(M; \mathcal{X}, G)$. Then $\mathcal{Y} \supseteq \mathcal{X}$ is a GBC-realization for M .

Now suppose \mathcal{Z} is a GBC-realization for M with $\mathcal{X} \subseteq \mathcal{Z} \subseteq \mathcal{Y}$. We want to see that either $\mathcal{Z} = \mathcal{X}$ or $\mathcal{Z} = \mathcal{Y}$. It is enough to see that if $A \in \mathcal{Y}$ then either G is definable from A and parameters in \mathcal{X} or else $A \in \mathcal{X}$. Without loss of generality we may assume that A is a class of ordinals. Then it was defined by some formula φ_n in our enumeration.

Consider $\mathbb{Q}_{n+1} \subseteq \mathbb{P}_n$. If the first case from the minimality lemma holds, then $A \in \mathcal{X}$ because $\xi \in A$ if and only if for every $p \in \mathbb{Q}_{n+1}$ the length of p being sufficiently long implies that $p \Vdash_{\mathbb{P}_n} \varphi_n(\xi)$. If the second case of the previous lemma holds, then we can define G from A . In this case, $p \in \bigcap_n \mathbb{P}_n$ if and only if for every ordinal α there is $q \succ_{\mathbb{Q}_{n+1}} p$ of length greater than α so that $q \Vdash_{\mathbb{P}_n} \varphi_n(\xi) \Leftrightarrow \xi \in A$ for all ordinals ξ . From a definition of $\bigcap_n \mathbb{P}_n$ can easily be produced a definition for G . \square

As remarked earlier, Global Choice was used essentially in the proof of lemma 1.71. Proving this lemma without Global Choice would yield a construction for minimal but not least GBC-realizations. Namely, start with a countable $M \models \text{ZFC}$ with no definable global well-order. Let $\mathcal{X} = \text{Def}(M)$. Then (M, \mathcal{X}) is a model of GBC minus Global Choice. Applying the theorem to (M, \mathcal{X}) would yield a GBC-realization \mathcal{Y} for M which is minimal above \mathcal{X} . But since any GBC-realization must contain \mathcal{X} , this would give that \mathcal{Y} is a minimal GBC-realization for M .

Thus, the problem of constructing a minimal but not least GBC-realization can be reduced down to the problem of proving the minimality lemma without using choice. A similar question can be asked for ordinary Sacks forcing.

Question 1.72. *Is choice needed to prove the minimality lemma for Sacks forcing? That is, is it consistent that there are $M \models \text{ZF} + \neg \text{AC}$, $s \subseteq \omega^M$ Sacks-generic over M , and $A \in M[s]$ so that $M \subsetneq M[A] \subsetneq M[s]$?*

This technique can also be applied to study principal models of GBC.

Definition 1.73. Say that (M, \mathcal{X}) is a *principal* model if there is $A \in \mathcal{X}$ so that $\mathcal{X} = \text{Def}(M; A)$. Let $\text{GBC-}\mathfrak{Rc}_{\text{Pr}}(M)$ denote the collection of principal GBC-realizations for M . Note that $\text{GBC-}\mathfrak{Rc}_{\text{Pr}}(M) \subseteq \text{GBC-}\mathfrak{Rc}(M)$ and if M is countable then $\text{GBC-}\mathfrak{Rc}_{\text{Pr}}(M) \subseteq \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$. Like $\text{GBC-}\mathfrak{Rc}(M)$ and $\text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$, $\text{GBC-}\mathfrak{Rc}_{\text{Pr}}(M)$ is ordered by \subseteq .

Theorem 1.74 (S. Friedman, Kossak–Schmerl). *For M countable, $\text{GBC-}\mathfrak{Rc}_{\text{Pr}}(M)$ is dense in $\text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$. That is, given any countable GBC-realization \mathcal{X} for M there is $\mathcal{Y} \supseteq \mathcal{X}$ which is a principal GBC-realization for M .*

This is theorem 15 of [HLR13]. Hamkins, Linetsky, and Reitz attribute the result independently to Friedman, via private communication, and Kossak and Schmerl for the finite set theory (equivalently, arithmetic) case. I will not give a full argument here, but let me sketch the proof.

Proof sketch. The argument is similar to the proof of theorem 1.67. We want to get a sequence

$$\langle \text{Ord} \rangle_2 = \mathbb{Q}_0 \supseteq \mathbb{P}_0 \supseteq \mathbb{Q}_1 \supseteq \mathbb{P}_1 \supseteq \cdots \supseteq \mathbb{Q}_n \supseteq \mathbb{P}_n \supseteq \cdots$$

so that \mathbb{Q}_{n+1} has the Σ_n -branch genericity property over \mathbb{P}_n . Before, we used the minimality lemma to produce \mathbb{P}_n from \mathbb{Q}_n . Here we need a different lemma.

Lemma 1.75. *Work over $(M, \mathcal{X}) \models \text{GBC}$. Let \mathbb{Q} be an Ord-height perfect binary tree and $A \in \mathcal{X}$ be a class of ordinals. Then, there is $\mathbb{P} \subseteq \mathbb{Q}$ in \mathcal{X} so that from \mathbb{Q} and any branch through \mathbb{P} we can define A .*

Proof sketch. Using A we thin out \mathbb{P} , keeping every other splitting node. This ensures that the tree we get at the end is still perfect. Reaching the $2i$ -th splitting node along a branch, we either go left or right. We go left if $i \in A$ and go right if $i \notin A$. This gives \mathbb{Q} . From \mathbb{P} we know where the $2i$ -th splitting nodes along a branch are. From a branch through \mathbb{Q} we know whether we went left or right to define \mathbb{Q} and thus whether $i \in A$. So we can define A . \square

To get theorem 1.74 we line up the classes in \mathcal{X} in ordertype ω —externally to the model. Then, to define \mathbb{P}_n from \mathbb{Q}_n we apply the lemma to code the n -th class. At the end, the G we get will allow us to define every set in \mathcal{X} , and thus $\text{Def}(M; \mathcal{X}, G) = \text{Def}(M; G) \in \text{GBC-}\mathfrak{Rc}_{\text{Pr}}(M)$. \square

We saw above that the rationals embed into $\text{GBC-}\mathfrak{Rc}(M)$. Indeed, given any $\mathcal{X} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ we can embed rationals into $\text{GBC-}\mathfrak{Rc}(M)$ above \mathcal{X} . This can be generalized to any countable partial order.

Theorem 1.76. *Every countable partial order embeds into $\text{GBC-}\mathfrak{Rc}(M)$, for countable $M \models \text{ZFC}$. More generally, for every finite partial order P and every $\mathcal{X} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ there is an embedding $e : P \rightarrow \text{GBC-}\mathfrak{Rc}(M)$ which maps P above \mathcal{X} —that is, $e(p) \supseteq \mathcal{X}$ for all $p \in P$.*

This theorem is similar to the result that every countable partial order embeds into the Turing-degrees. See below for further discussion.

Proof. Let us first see the special case of an atomic boolean algebra. Let $(B, <^B)$ be a countable atomic boolean algebra, with atoms b_0, b_1, \dots . Let C_0, C_1, \dots be mutually $(M, \text{Def}(M, \mathcal{X}))$ -generic Cohen subclasses of Ord. Then the map $b_i \mapsto \text{Def}(M; \mathcal{X}, C_i)$ induces an embedding of B into $\text{GBC-}\mathfrak{Rc}(M)$; given arbitrary $b \in B$ map b to $\text{Def}(M; \mathcal{X}, C_i : b_i \leq^B b)$.

So we are done once we see that every countable partial order embeds into a countable atomic boolean algebra. Let $(P, <^P)$ be a countable partial order. The desired boolean algebra $(B, <^B)$ will be generated by countably many atoms, with an atom $a(p)$ associated to each $p \in P$. For the embedding e , map p to the unique $b \in B$ so that $a(q) \leq^B b$ if and only if $q \leq^P p$. Then $q \leq^P p$ if and only if $e(q) \leq^B e(p)$. \square

The embedding from this argument destroys a lot of information about the partial order. It may be that $p, q \in P$ have no upper bound. But their image under the embedding will have an upper bound, as we first embed P into a boolean algebra—where all pairs of elements have an upper bound—and then embed that boolean algebra into $\text{GBC-}\mathfrak{Rc}(M)$. But recall from corollary 1.64 that there are $\mathcal{X}, \mathcal{Y} \in \text{GBC-}\mathfrak{Rc}(M)$ without an upper bound.

Can we do better and get embeddings that preserve the non-existence of upper bounds?

Theorem 1.77 (Mostowski [Mos76]). *Let $(M, \mathcal{X}) \models \text{Gbc}$ be countable and F be a finite family of finite sets, closed under subset.²⁷ Then there are (M, \mathcal{X}) -generic Cohen subclasses C_i of Ord for each $i \in \bigcup F$ and an embedding e from F to $\text{GBC-}\mathfrak{Rc}(M)$ satisfying the following.*

²⁷Such F are precisely those orders which are initial segments of a finite boolean algebra of sets.

1. If $f \in F$ then the C_i for $i \in f$ are mutually generic; and
2. If $f \subseteq \bigcup F$ is not in F then the C_i for $i \in f$ do not amalgamate: there is no GBC-realization \mathcal{Y} for M which contains C_i for all $i \in f$.

Before giving the proof let me extract a corollary.

Corollary 1.78. *Let $M \models \text{ZFC}$ be countable.*

1. *Suppose M has a definable global well-order and let P be a finite partial order with a least element $\mathbf{0}^P$. Then there is an embedding e of P into $\text{GBC-}\mathfrak{Rc}(M)$ so that $e(\mathbf{0}) = \text{Def}(M)$ and e preserves the existence/nonexistence of upper bounds and nonzero lower bounds.*
2. *Suppose that M does not have a definable global well-order and let P be a finite partial order without a least element. Then there is an embedding e of P into $\text{GBC-}\mathfrak{Rc}(M)$ which preserves the existence/nonexistence of upper bounds and lower bounds.*
3. *Let $\mathcal{X} \in \text{GBC-}\mathfrak{Rc}_{<\omega_1}(M)$ and let P be a finite partial order with a least element $\mathbf{0}^P$. Then there is an embedding e of P into $\text{GBC-}\mathfrak{Rc}(M)$ above \mathcal{X} so that $e(\mathbf{0}) = \mathcal{X}$ and e preserves existence/nonexistence of upper bounds and nonzero lower bounds. That is, $e(p) \supseteq \mathcal{X}$ for all $p \in P$ and p and p' have a nonzero lower bound if and only if $e(p)$ and $e(p')$ have \mathcal{X} as their greatest lower bound.*

Proof. In each case, we will first embed P into a finite family F of finite sets closed under subset, in such a way that the embedding preserves the existence/non-existence of upper bounds and nonzero lower bounds. We will apply theorem 1.77 using F . For (1) and (2) we will work over $(M, \text{Def}(M))$ while for (3) we will work over (M, \mathcal{X}) . The preservation of upper bounds will be ensured by mutual genericity, which will also ensure the condition about lower bounds. Finally, the preservation of nonexistence of upper bounds follows from non-amalgamability.

It remains to see the embedding of P into F . First embed $(P, <^P)$ into a boolean algebra of sets B by the following e : the boolean algebra is generated by an atom $a(p)$ for each $p \in P$ and we map p to $\{a(q) : q \leq p\}$. Let F be the downward closure of $e''P$ in B . Then $e(p)$ and $e(q)$ have an upper bound in F if and only if there is $r \geq^P p, q$. And two sets $f, g \in F$ have a nonzero lower bound if and only if $f \cap g \neq \emptyset$ which happens if and only if $\{i\} \subseteq f \cap g$ for some $i \in \bigcup F$. So if $f = e(p)$ and $g = e(q)$ then $e(p)$ and $e(q)$ have a nonzero lower bound if and only if they both contain $a(r)$ for some $r \in P$ if and only if $r \leq p$ and $r \leq q$. Thus, we have seen that e is as desired. \square

Proof of theorem 1.77. Without loss of generality $\bigcup F = n \in \omega$. And if $n \in F$ then the result is trivial—merely add n mutually-generic Cohen subclasses of Ord . So assume we are the case where $n \notin F$.

As in the non-amalgamability argument for theorem 1.61.(4), fix a bad sequence $B : \text{Ord}^M \rightarrow 2$ which witnesses that Ord^M is countable. We want to construct the Cohen generics C_i for $i \in n$ so that for $f \subseteq n$ the Cohen generics $\{C_i : i \in f\}$ code B if and only if $f \notin F$. We construct the C_i in Ord^M many stages. Externally to the model, fix an Ord^M -sequence of the dense subclasses in \mathcal{X} of $\text{Add}(\text{Ord}, |f|)$ for some $f \in F$. We can arrange this so that each stage α has a corresponding $f_\alpha \in F$ and for each $f \in F$ each dense subclass of $\text{Add}(\text{Ord}, |f|)$ appears at some stage α with $f = f_\alpha$.

Start with $c_i^0 = \emptyset$ for all $i \in f$. At limit stages, we will simply take unions. All the work is in the successor stage. Suppose we have already built c_i^α for all $i \in f$. Let $D \subseteq \text{Add}(\text{Ord}, |f_\alpha|)$ be the dense class for stage α . We can extend c_i^α to d_i^α for $i \in f$ so that $\prod_{i \in f} d_i^\alpha$ meets D . By padding out with 0s if necessary, we may assume without loss that the d_i^α 's all have the same length, call it γ . Now, for $i \notin f$ extend c_i^α to d_i^α of length γ by adding 0s everywhere new. Finally, set $c_i^{\alpha+1} = d_i^\alpha \cap \langle 1, B(\alpha) \rangle$.

Then $C_i = \bigcup_{\alpha \in \text{Ord}^M} c_i^\alpha$ is Cohen-generic. And it is clear from the construction that $\{C_i : i \in f\}$ is a family of mutually-generic Cohen subclasses of Ord for $f \in F$. It remains to see that if $f \notin F$ then $\{C_i : i \in f\}$ codes B . This is because, $\{C_i : i \in f\}$ can recognize the coding points. They occur just after the rows of all 1s. That is, $B(\alpha) = C_i(\xi_\alpha + 1)$ (any $i \in f$) where ξ_α is the α th ξ so that $C_i(\xi) = 1$ for all $i \in f$. \square

Many questions remain open about the structure of $\text{GBC-}\mathfrak{Rc}(M)$. Let me mention one open-ended question.

Question 1.79. *What can be said about the theory of the structure $(\text{GBC-}\mathfrak{Rc}(M), \subseteq)$? What if we add in predicates for collection of the countable GBC -realizations or the collection of the principal GBC -realizations?*

This project of studying the order structure of $\text{GBC-}\mathfrak{Rc}(M)$ has similarities to two extant projects in mathematical logic. Let me briefly mention them and the connections.

The first and older of the two is the study of the Turing degrees under the order of Turing-reducibility. The structure of this partial order has been well studied and many results about the Turing degrees have counterparts in the context of $\text{GBC-}\mathfrak{Rc}(M)$. For instance, it is known that every countable poset embeds into the Turing degrees.

The reader should be warned, however, that there are differences between the two. Let me illustrate this with the exact pair theorem as an example.

Theorem 1.80 (Spector [Spe56]). *Let $\langle \mathbf{d}_n : n \in \omega \rangle$ be a sequence of Turing degrees so that $\mathbf{d}_n <_T \mathbf{d}_{n+1}$ for all n . Then there are Turing degrees \mathbf{a} and \mathbf{b} which are above each \mathbf{d}_n but if $\mathbf{c} <_T \mathbf{a}, \mathbf{b}$ then $\mathbf{c} <_T \mathbf{d}_n$ for some n .*

The analogous result is not true for GBC -realizations for a fixed countable $M \models \text{ZFC}$. Let $\langle \mathcal{X}_n : n \in \omega \rangle$ be an increasing \subseteq -chain of GBC -realizations for M and let $\mathcal{Y}, \mathcal{Z} \in \text{GBC-}\mathfrak{Rc}(M)$ be upper bounds for the sequence. Then $\bigcup \mathcal{X}_n \in \text{GBC-}\mathfrak{Rc}(M)$ is below both \mathcal{Y} and \mathcal{Z} but above every \mathcal{X}_n . The dis-analogy is because GBC -realizations do not have to be generated by a single class. This suggests that the correct analogy is between GBC -realizations and Turing ideals. But even then, the analogy is not perfect. For example, Turing ideals do not have a non-amalgamability phenomenon, as any collection of Turing ideals are all contained in the Turing ideal $\mathcal{P}(\omega)$.

The second connection is to the generic multiverse. Given a countable transitive²⁸ model M of set theory the generic multiverse of M is the smallest collection of countable transitive models containing M which is closed under (set) forcing extensions and (set) grounds.²⁹ It follows from work by Usuba [Usu17] that the generic multiverse of M can be equivalently defined as the collection of all forcing extensions of grounds of M .

Let \mathcal{M} be the generic multiverse of M . Then \mathcal{M} is partially ordered under inclusion and we can ask about the order-theoretic properties of \mathcal{M} . Many of the arguments about Cohen-generic subclasses of Ord also apply for Cohen-generic subsets of, say, ω which yields similar results for the generic multiverse as it does for the GBC -realizations. See e.g. [Ham16].

Let me conclude by considering to what extent these results generalize from GBC to other theories.

First, we consider what happens if we drop powerset. As mentioned in section 1.2, forcing with $\text{Add}(\text{Ord}, 1)$ may add sets for some models of ZFC^- . So most of the techniques of this section fail badly for those models. However, if $M \models \text{ZFC}^-$ is such that forcing with $\text{Add}(\text{Ord}, 1)$ does not add sets, then the same arguments go through and we get all the same results for $\text{GBC}^- \text{-}\mathfrak{Rc}(M)$.

Next consider stronger theories. First, let us see that $T\text{-}\mathfrak{Rc}_{\text{Pr}}(M)$ may be empty.

Proposition 1.81. *Let $T \supseteq \text{GBC}^-$ prove that for every class A the first-order truth predicate relative to A exists. Then T has no principal models.*

Proof. Suppose otherwise that $(M, \mathcal{X}) \models T$ has that every class is definable from $A \in \mathcal{X}$. But then the truth predicate relative to A is definable from A , contradicting Tarski's theorem on the undefinability of truth. \square

In particular, there are no principal models of theories extending $\text{GBC}^- + \text{ETR}$, or even $\text{GBC}^- + \text{ETR}_\omega$. So for most theories T of interest, $T\text{-}\mathfrak{Rc}_{\text{Pr}}(M)$ is trivial. But we can say something about $T\text{-}\mathfrak{Rc}(M)$ and $T\text{-}\mathfrak{Rc}_{<\omega_1}(M)$.

²⁸One does not need the assumption of transitivity, but let me leave it in to simplify the discussion.

²⁹A *ground* of M is a submodel W so that $M = W[g]$ for some g generic for some forcing notion in W .

The main tool used in this section was Cohen-generic subclasses of Ord . We carefully constructed generics to have certain properties with regard to amalgamability/non-amalgamability, and thereby concluded something about the structure of $\text{GBC-}\mathfrak{Re}(M)$. Some of these results generalize to $T\text{-}\mathfrak{Re}(M)$ and $T\text{-}\mathfrak{Re}_{<\omega_1}(M)$, for T which is preserved by Cohen-forcing.

Theorem 1.82. *Let $M \models \text{ZFC}$ be countable and let T be a second-order set theory which is preserved by forcing with $\text{Add}(\text{Ord}, 1)$ over M .³⁰ Suppose M is T -realizable. Then $T\text{-}\mathfrak{Re}(M)$ satisfies the following.*

1. *For any $\mathcal{X} \in T\text{-}\mathfrak{Re}_{<\omega_1}(M)$ there are $\mathcal{Y}, \mathcal{Z} \supseteq \mathcal{X}$ in $T\text{-}\mathfrak{Re}_{<\omega_1}(M)$ so that \mathcal{Y} and \mathcal{Z} lack an upper bound.*
2. *Every $\mathcal{X} \in T\text{-}\mathfrak{Re}_{<\omega_1}(M)$ has a dense extension. That is, there is $\mathcal{Y} \in T\text{-}\mathfrak{Re}_{<\omega_1}(M)$ so that there is a dense linear order in $T\text{-}\mathfrak{Re}(M)$ between \mathcal{X} and \mathcal{Y} .*
3. *For every countable partial order P and every $\mathcal{X} \in T\text{-}\mathfrak{Re}_{<\omega_1}(M)$ there is an embedding $e : P \rightarrow T\text{-}\mathfrak{Re}(M)$ which maps P above \mathcal{X} .*
4. *Every finite partial order embeds into $T\text{-}\mathfrak{Re}(M)$ in such a way as to preserve the existence/nonexistence of upper bounds and nonzero lower bounds.*

Proof sketch. The same arguments as in the $T = \text{GBC}$ case. □

However, not all arguments generalize. In particular, the argument for when $\text{GBC-}\mathfrak{Re}(M)$ has a least element will not work for theories stronger than GBC . Indeed, in chapter 4 we will see that the result is false for sufficiently strong theories. If $T \supseteq \text{GBC} + \Pi_1^1\text{-CA}$ then $T\text{-}\mathfrak{Re}(M)$ never has a least element, for countable M .

But we are not yet ready to prove this. First we need to know more about the structure of models of strong second-order set theories, which we turn to in chapter 2.

³⁰For example, T could be KM , KMCC , $\text{GBC} + \Pi_k^1\text{-CA}$, or $\text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$.

Chapter 2

Many constructions: unrollings, cutting offs, and second-order L

Second-order logic is set theory in disguise.

W.V.O. Quine

This chapter contains an exposition of several constructions relating to models of strong second-order set theories. These constructions are not new, dating at least as far back as work by Marek and Mostowski in the 1970s [Mar73; MM75]. The ultimate origin of the constructions is not clear to me. Indeed, in the introduction to [Mar73] Marek claims that Jensen, Mostowski, Solovay, and Tharp all independently proved that KM is consistent with $V = L$, which will follow from the third construction considered in this chapter, viz. the second-order constructible universe.¹ More recently, Antos and S. Friedman [AF] independently rediscovered these constructions a few years ago.

Previously these constructions have been done in the context of KM or KMCC. In this chapter I analyze them more finely, generalizing their application from models of KM or KMCC to models of weaker theories.

The first construction I call the unrolling construction. This is essentially the same construction as the one used to code hereditarily countable sets as reals, which has seen wide use within set theory. The point is that the same idea works for sets which are not hereditarily countable. We can code ‘sets’ of rank $> \text{Ord}$ as proper classes. Then, given that our universe satisfies a strong enough second-order set theory these codes can be unrolled to produce a model of first-order set theory. The theory of this unrolled model will depend upon the theory of the ground universe.

The reader who is familiar with reverse mathematics may know that a similar construction is used for theories of second-order arithmetic. This yields that they are bi-interpretable with certain (first-order) set theories. See [Sim09, chapter VII] for a thorough exposition.

The second construction is the cutting off construction. This construction starts with a model of ZFC^- (or weaker theory) with a largest cardinal κ with κ regular and $H_\kappa \models \text{ZFC}^-$. We then get a second-order model by considering $(H_\kappa, \mathcal{P}(H_\kappa))$ in this model, where $\mathcal{P}(H_\kappa)$ is necessarily a (definable) proper class in the model. This will yield a model of GBC^- , with more strength coming from a stronger theory in the ground model. If κ is moreover inaccessible we will get a model of GBC , or more.

The cutting off construction is exactly the inverse of the unrolling construction. Starting with (M, \mathcal{X}) a model of a sufficiently strong second-order set theory, the cutting off of the unrolling of (M, \mathcal{X}) is isomorphic to (M, \mathcal{X}) . In the other direction, start with N a model of a strong enough fragment of ZFC^- with a largest cardinal κ with κ regular and $H_\kappa^N \models \text{ZFC}^-$. Then the unrolling of the cutting off of N is isomorphic to N .

¹Marek does not mention how they proved the result, but it seems likely that it was via a version of this construction.

Together, these two constructions yield that strong enough second-order set theories are bi-interpretable with certain first-order set theories. I summarize these bi-interpretability results below. First, however, we will need names for the first-order set theories we get from unrolling.

Definition 2.1. The following are the first-order set theories theories which are satisfied by the unrolled models arising from a model of second-order set theory. Each includes the basic axioms of set theory—namely Extensionality, Pairing, Union, Foundation, Choice, and Infinity—and the assertion that there is a largest cardinal κ .

- ZFC_I^- consists of the basic axioms plus Separation, Collection, and the assertion that κ is inaccessible. To be clear, since this theory does not include Powerset, by “ κ is inaccessible” is meant that κ is regular and every set in V_κ has a powerset which is also in V_κ . In particular, ZFC_I^- proves that $V_\kappa = H_\kappa$ is a model of ZFC.
- ZFC_R^- consists of the basic axioms plus Separation, Collection, the assertion that κ is regular, and the assertion that H_κ exists. In particular, ZFC_R^- proves that H_κ is a model of ZFC^- .²
- wZFC_I^- consists of the basic axioms plus Separation and the assertion that κ is inaccessible. In particular, wZFC_I^- proves that $V_\kappa = H_\kappa$ is a model of ZFC.
- wZFC_R^- consists of the basic axioms plus Separation, the assertion that κ is regular, and the assertion that H_κ exists. In particular, wZFC_R^- proves that H_κ is a model of ZFC^- .

Let $k < \omega$.

- $\text{ZFC}_I^-(k)$ consists of the basic axioms plus Σ_k -Separation, Σ_k -Collection, and the assertion that κ is inaccessible. In particular, $\text{ZFC}_I^-(k)$ proves that $V_\kappa = H_\kappa$ is a model of ZFC.
- $\text{ZFC}_R^-(k)$ consists of the basic axioms plus Σ_k -Separation, Σ_k -Collection, the assertion that κ is regular, and the assertion that H_κ exists. In particular, $\text{ZFC}_R^-(k)$ proves that H_κ is a model of ZFC^- .
- $\text{wZFC}_I^-(k)$ consists of the basic axioms plus Σ_k -Separation and the assertion that κ is inaccessible. In particular, $\text{wZFC}_I^-(k)$ proves that $V_\kappa = H_\kappa$ satisfies every axiom of ZFC.
- $\text{wZFC}_R^-(k)$ consists of the basic axioms plus Σ_k -Separation, the assertion that κ is regular, and the assertion that H_κ exists. In particular, $\text{wZFC}_R^-(k)$ proves that H_κ satisfies every axiom of ZFC^- .

Let me explain the mnemonic behind the names of these theories for the benefit of the reader, to whom I apologize for giving eight theories to remember. The subscripts tell you what is being asserted about κ , the largest cardinal. ‘I’ reminds you that κ is inaccessible while ‘R’ tells you κ is merely regular. The w in front stands for weak, wimpy, and why would you ever want to work with a theory which does not have even a fragment of Collection?^{3,4} The parenthetical k tells us what fragment of Separation and Collection—or just Separation in case w is in front—is in the theory.

Theorem 2.2. *The following pairs of theories are bi-interpretable. Below, $k \geq 1$.*

- (Marek [Mar73]) KMCC and ZFC_I^- .
- KMCC^- and ZFC_R^- .
- KM and $\text{wZFC}_I^- + \Sigma_0$ -Transfinite Recursion.⁵

² ZFC_R^- has natural models, for instance H_{ω_2} . In general, if κ is regular then $H_{\kappa^+} \models \text{ZFC}_R^-$.

³For an extensive case study in why one would want Collection, see [Mat01].

⁴So under this naming system wZFC , although not used here, would be Zermelo set theory plus Foundation and Choice. Unfortunately the natural name for this theory, ZFC, is already used to refer to Zermelo set theory plus Foundation, Choice, and Collection.

⁵ Σ_0 -Transfinite Recursion, which means what you think it means, will be formally defined in subsection 2.1.2.

- KM^- and $\text{wZFC}_R^- + \Sigma_0$ -Transfinite Recursion.
- $\text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$ and $\text{ZFC}_I^-(k)$.
- $\text{GBC}^- + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$ and $\text{ZFC}_R^-(k)$.
- $\text{GBC} + \Pi_k^1\text{-CA}$ and $\text{wZFC}_I^-(k) + \Sigma_0$ -Transfinite Recursion.
- $\text{GBC}^- + \Pi_k^1\text{-CA}$ and $\text{wZFC}_R^-(k) + \Sigma_0$ -Transfinite Recursion.
- $\text{GBC} + \text{ETR}$ and $\text{wZFC}_I^-(0) + \Sigma_0$ -Transfinite Recursion.
- $\text{GBC}^- + \text{ETR}$ and $\text{wZFC}_R^-(0) + \Sigma_0$ -Transfinite Recursion.

The third construction I look at in this chapter is the construction of Gödel's constructible universe but extended from the sets to the classes. Given a class well-order Γ we can define L_Γ and then consider the definable hyperclass \mathcal{L} consisting of all classes which appear in some L_Γ . We will consider also the construction of L relative to parameters. The main use to which we will put this construction is in showing how to get models satisfying (a fragment of) Class Collection. The classical result here is that if we start with (M, \mathcal{X}) a model of KM then the \mathcal{L} we build gives a model of KMCC [Mar73]. That is, any model of KM contains an Ord-submodel of KMCC , and it is straightforward to tweak the construction to give a V -submodel. I will generalize this result from models of KM to models of $\text{GBC} + \Pi_k^1\text{-CA}$. This gives the following result.

Theorem 2.3. *Let $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_k^1\text{-CA}$ for $1 \leq k \leq \omega$ and suppose $N \in \mathcal{X}$ is an inner model (of ZFC^-) of M . Then there is $\mathcal{Y} \subseteq \mathcal{X}$ a definable hyperclass so that $(N, \mathcal{Y}) \models \text{GBC}^- + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$. In particular, if $N = M$, this implies that every model of $\text{GBC}^- + \Pi_k^1\text{-CA}$ contains a second-order definable V -submodel of $\text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$.*

Finally, I will end this chapter with an application of these constructions. Marek and Mostowski showed [MM75] that the shortest height of a transitive model of KM is less than the shortest height of a β -model of KM . Moreover, if the former ordinal is τ_ω and the latter is β_ω then $L_{\beta_\omega} \models \tau_\omega$ is countable. I will generalize this result to $\text{GBC} + \Pi_n^1\text{-CA}$, showing that τ_n —the least height of a transitive model of $\text{GBC} + \Pi_n^1\text{-CA}$ —is less than β_n —the least height of a beta model of $\text{GBC} + \Pi_n^1\text{-CA}$. Moreover, $L_{\beta_n} \models \tau_n$ is countable.

The results in this chapter will form the bedrock for much of chapter 3 and chapter 4.

2.1 The unrolling construction

The structure of this section is as follows. We work throughout in a fixed model (M, \mathcal{X}) of some second-order set theory. I will first lay out the basic definitions for what will be the unrolling of (M, \mathcal{X}) . I will then investigate what theory is satisfied by the unrolled structure. This will be set up as a series of propositions, showing that if the ground model satisfies such and such then the unrolled model will satisfy so and so. At the end, these propositions will yield one half of the bi-interpretability theorem above. I will summarize the results before moving on to the next section, about the cutting off construction.

In reading the following the reader would do well to keep in mind what she knows about coding hereditarily countable sets as reals.

Definition 2.4 (Over Gbc^-). Call a class binary relation A a *membership code* if A is the relation for a well-founded, extensional directed graph with a top element. Let t_A denote the top element of A . For any $x \in \text{dom}(A)$, let $A \downarrow x = \{(a, b) \in A : a, b \leq_A x\}$, where \leq_A is the reflexive transitive closure of A , be A restricted below the node x .

To suggest the graph theoretic perspective, I will sometimes use $\dashv\!\!\dashv_A$ as a synonym for A . If the context is clear, I will just write $\dashv\!\!\dashv$. As a particular example, to express that x is an immediate predecessor of the top element of A I will write $x \dashv\!\!\dashv t_A$ rather than $x A t_A$.

Let $\text{elts } A = \{x \in \text{dom } A : x \dashv\!\!\dashv t_A\}$ denote the penultimate level of A . In the unrolled structure, for $x \in \text{elts } A$ we have $A \downarrow x$ is a membership code which represents an element of the set represented by A .




Figure 2.1: A membership code representing the set $a = \{0, 2\}$.

Note that this definition is first-order; it does not require any quantification over classes to check whether A is a membership code. Thus, if \mathcal{X} and \mathcal{Y} are GBc^- -realizations for the same $M \models \text{ZFC}^-$ with $A \in \mathcal{X} \cap \mathcal{Y}$ then \mathcal{X} and \mathcal{Y} agree on whether A is a membership code.

For insight into the definition, consider the following definition.

Definition 2.5. Let x be a set. The *canonical membership code for x* is $E_x = \in \upharpoonright \text{TC}(\{x\})$. If X is a class then the *canonical membership code for x* is $E_X = \in \upharpoonright \text{TC}(\{x\}) \cup \{(x, \star) : x \in X\}$ where \star is a new element.

Once we have defined the unrolling it will be immediate that the canonical membership code for a set (respectively class) represents that set (respectively class) in the unrolling.

Unlike with the coding of hereditarily countable sets as reals, many membership codes represent virtual objects, ‘sets’ which are too high in rank to be a class in the ground universe. For example, any order of ordertype $\text{Ord} + \text{Ord} + 1$ is a membership code representing an ‘ordinal’ of ordertype $\text{Ord} + \text{Ord}$. But of course there is no literal such ordinal in the ground universe.

Because membership codes can be high in rank, Mostowski’s collapse lemma does not apply to them. Indeed, it is precisely those membership codes which are not isomorphic to the restriction of \in to some class that are of interest here. Lacking a way to canonically choose membership codes, we work with all membership codes and will quotient by isomorphism to produce the first-order structure.

Let us check some basic properties about isomorphisms of membership codes. The reader should note that these facts require only a weak background theory to prove. As such, the unrolling construction can be carried out over a model of a rather weak theory to produce some sort of structure. But it will take some strength in order for this unrolled structure to satisfy an appreciable set theory.

First, isomorphisms for membership codes are unique. This is analogous to the rigidity of transitive sets, but applied to membership codes rather than literal sets.

Proposition 2.6 (Over GBc^-). *Let A and B be membership codes. If $\pi, \sigma : A \cong B$ then $\pi = \sigma$.*

Proof. By Elementary Comprehension, the class $X = \{a \in \text{dom } A : \pi(a) \neq \sigma(a)\}$ exists. Suppose towards a contradiction that X is nonempty. Then, we can pick a minimal $x \in X$, as this is expressible via a first-order property. By minimality, $y \dashv\!\circ_A x$ implies $\pi(y) = \sigma(y)$. But because π and σ are isomorphisms and A and B are extensional, $\pi(x)$ is determined by $\{\pi(y) : y \dashv\!\circ_A x\}$ and $\sigma(x)$ is determined by $\{\sigma(y) : y \dashv\!\circ_A x\}$. So $\pi(x) = \sigma(x)$, contradicting the choice of x . So X is nonempty and thus $\pi = \sigma$. \square

Next, nodes in a membership code are determined by the isomorphism type of their class of predecessors.

Proposition 2.7 (Over Gbc^-). *Let A be a membership code and $x, y \in \text{dom } A$. If $A \downarrow x \cong A \downarrow y$ then $x = y$.*

Proof. Let $\sigma : A \downarrow x \cong A \downarrow y$. We want to see that σ is the identity. Set $X = \{a <_A x : \sigma(a) \neq a\}$. Suppose towards a contradiction that X is nonempty. Then, pick $a \in X$ minimal. We have that $\sigma(b) = b$ for all $b \dashv\!\!\dashv_A a$ so it must be that $\sigma(a) = a$. This contradicts the choice of a , so it must be that X is empty and σ is the identity. \square

Third, the uniqueness of isomorphism generalizes in an appropriate way to certain partial isomorphisms.

Definition 2.8. Let A and B be membership codes. A partial function $\pi : A \rightarrow B$ is an *initial partial isomorphism* if its domain is downward-closed in A ,⁶ its range is downward closed in B , and for all $a, a' \in \text{dom } \pi$ we have $a \dashv\!\!\dashv_A a'$ if and only if $\pi(a) \dashv\!\!\dashv_B \pi(a')$.

Proposition 2.9 (Over Gbc^-). *Let $\pi, \sigma : A \rightarrow B$ be initial partial isomorphisms. Then π and σ agree on the intersection of their domains.*

Proof. Let $C = \text{dom } \pi \cap \text{dom } \sigma$. Then C is nonempty, as it must contain the least element of A . Now consider $X = \{x \in C : \pi(x) \neq \sigma(x)\}$. As before if X is nonempty then it has a minimal element x , but then by the properties of isomorphism for membership codes it must be that $\pi(x) = \sigma(x)$, contradicting the non-emptiness of X . \square

Observe that there are always initial partial isomorphisms between membership codes. In particular, the partial map sending the least element of A to the least element of B is an initial partial isomorphism.

If our ground universe has a stronger theory then we can prove that there is a maximum initial partial isomorphism between membership codes.

Lemma 2.10 (Over $\text{Gbc}^- + \text{ETR}$). *Let A and B be membership codes. Then there is a maximum initial partial isomorphism π between A and B . That is, if $\sigma : A \rightarrow B$ is any initial partial isomorphism then $\sigma \subseteq \pi$.*

Proof. This maximum initial partial isomorphism π is constructed via an elementary transfinite recursion on A . Namely, π is constructed via the transfinite recursion to construct an isomorphism between A and B , except that we stop constructing higher when we reach a local failure of isomorphism. Formally, π is defined via the following recursive requirement:

- $\pi(a)$ is the unique $b \in B$ so that for all $a' <_A a$ we have $\pi(a') \dashv\!\!\dashv_B b$ if and only if $a' \dashv\!\!\dashv_A a$, if such b exists and $\pi(a')$ is defined for all $a' \dashv\!\!\dashv_A a$; and
- $\pi(a)$ is undefined, otherwise.

Elementary Transfinite Recursion says that this recursion has a solution π . Manifestly π is an initial partial isomorphism. Let us check that π is the maximum initial partial isomorphism. Take $\sigma : A \rightarrow B$ an initial partial isomorphism. Then σ and π agree on their domain, by proposition 2.9. So the recursion to construct π will work on all of $\text{dom } \sigma$ and restricted to $\text{dom } \sigma$ will give σ . So $\sigma \subseteq \pi$. \square

Let me remark on this proof. One might attempt to more easily prove the existence of maximum initial partial isomorphisms between membership codes by considering the hyperclass of all initial partial isomorphisms and then taking the union of all of them. The issue with this argument is that it makes a hidden appeal to Π_1^1 -Comprehension: we wish to define π by saying that $\pi(a)$ is defined if *there exists* some initial partial isomorphism $\sigma : A \rightarrow B$ with $a \in \text{dom } \sigma$ and that then $\pi(a) = \sigma(a)$. This is a Σ_1^1 assertion.⁷ The more convoluted argument is preferred because it works from a weaker base theory.

Isomorphism will become equality in the unrolled structure. We also must say what will become the membership relation.

⁶That is, if $a \in \text{dom } \pi$ and $a' \dashv\!\!\dashv a$ then $a' \in \text{dom } \pi$

⁷Recall that Π_1^1 -Comprehension is equivalent to Σ_1^1 -Comprehension.

Definition 2.11 (Over GBc^-). Given membership codes A and B say that $A \varepsilon B$ if there is $a \multimap_B t_B$ so that $A \cong B \downarrow a$.

In particular, if $x \multimap_A t_A$ then $(A \downarrow x) \varepsilon A$.

Proposition 2.12 (Over GBc^-). *Isomorphism of membership codes is a congruence with respect to ε . That is, if A and B are membership codes so that $A \varepsilon B$, $A \cong A'$, and $B \cong B'$, then $A' \varepsilon B'$.*

Proof. Let $e : A \hookrightarrow B$ embed A onto $B \downarrow a$ for some $a \multimap_B t_B$. Let $\pi : A' \cong A$ and $\sigma : B \cong B'$. Then $\sigma \circ e \circ \pi : A' \hookrightarrow B$ embeds A onto $B' \downarrow a'$ for some $a' \multimap_{B'} t_{B'}$. \square

The remainder of this section is dedicated to working out just what the theory of the ground universe implies about the theory of the unrolling.

Hereon, let \mathcal{U} denote the hyperclass of all membership codes and let $\mathfrak{U} = (\mathcal{U}/\cong, \varepsilon)$ denote the unrolled structure. Note that \mathcal{U} is a definable hyperclass, via a first-order formula.

Theorem 2.13 (Over $\text{GBc}^- + \text{ETR}$). *The unrolled structure \mathfrak{U} satisfies Extensionality.*

Proof. Fix membership codes A and B . It needs to be shown that $A \cong B$ if and only if $\forall C C \varepsilon A \Leftrightarrow C \varepsilon B$. The forward direction of the implication is immediate. It is the other direction which requires work.

Suppose that for any membership code C we have $C \varepsilon A$ if and only if $C \varepsilon B$. In particular, this is true for C of the form $A \downarrow a$ for $a \multimap_A t_A$ or the form $B \downarrow b$ for $b \multimap_B t_B$. For $a \multimap_A t_A$ let π_a be the embedding which maps $A \downarrow a$ onto $B \downarrow b$ for some b . And in the other direction, for $b \multimap_B t_B$ let σ_b be the initial partial isomorphism which maps $B \downarrow b$ onto $A \downarrow a$ for some $a \multimap_A t_A$. Proposition 2.7 gives that the choice of b is unique and thus π_a is well-defined, and similarly for σ_b . Notice, however, that in the absence of Π_1^- -Comprehension we have no way to uniformly refer to the π_a and the σ_b .

By lemma 2.10 let π be the maximum initial partial isomorphism from A to B . First, note that $\text{dom } \pi$ includes all $a \multimap_A t_A$. This is because $\pi \supseteq \pi_a$ for $a \multimap_A t_A$, by maximality. Notice also that $B \downarrow b \subseteq \text{ran } \pi$ for all $b \multimap_B t_B$ because σ_b^{-1} is an initial partial isomorphism from A to B which maps onto $B \downarrow b$. But then $\text{dom } \pi$ must also include t_A and $\text{ran } \pi$ must include t_B as once we have an initial partial isomorphism defined everywhere but the top elements it is obvious how to extend: send t_A to t_B . So $\text{dom } \pi = A$ and $\text{ran } \pi = B$, so π is a full isomorphism. \square

The unrolled structure satisfies other basic axioms of set theory.

Proposition 2.14 (Over $\text{GBc}^- + \text{ETR}$). *The unrolled structure \mathfrak{U} satisfies Union, Pairing, Infinity, and Foundation.*

Proof. (Union) Given a membership code X we need to produce a membership code which represents $\bigcup X = \{Z : \exists Y Z \varepsilon Y \varepsilon X\}$. We define such a Y by cutting out the penultimate level of X :

$$Y = (X \setminus \{(x, t_X) : x \multimap t_X\}) \cup \{(x, t_X) : \exists x' x \multimap x' \multimap t_X\}.$$

This can be constructed by an instance of Comprehension and is easily verified to represent $\bigcup X$. See figure 2.2.

(Pairing) Let A and B be membership codes. By lemma 2.10 find $\pi : A \rightarrow B$ the maximum initial partial isomorphism from A to B . Now let

$$\begin{aligned} P = & A \setminus (A \upharpoonright \text{dom } \pi) \\ & \cup \{(\pi(a), a') : (a, a') \in A \text{ and } a \in \text{dom } \pi \text{ and } a' \notin \text{dom } \pi\} \\ & \cup B \\ & \cup \{(t'_A, p), (t_B, p)\}, \end{aligned}$$

where p is a new point and $t'_A = t_A$ if $t_A \notin \text{dom } \pi$ and $t'_A = \pi(t_A)$ otherwise. It is easy to see that P is a membership code which represents the unordered pair consisting of A and B .




Figure 2.2: To construct a membership code for $\bigcup X$ we remove all edges ending at the top node and add edges (n_Z, t_X) for all $n_Z \multimap n_Y \multimap t_X$.

(Infinity) There is a membership code for ω . It is straightforward to check that it represents an inductive set in \mathfrak{U} .

(Foundation) At bottom, Foundation holds in \mathfrak{U} because membership codes are well-founded.

More formally, suppose towards a contradiction that Foundation fails in \mathfrak{U} . That is, in the ground universe there is a membership code A (for a nonempty set) so that for every $B \in A$ there is C such that $C \in A$ and $C \in B$. In particular, this holds for B of the form $A \downarrow b$ for $b \multimap_A t_A$. Therefore, we get that for all $b \multimap_A t_A$ there is $c \multimap_A t_A$ so that $c \multimap_A t_b$. But then elts A has no minimal element, contradicting that A is well-founded. \square

We also get that the unrolled structure satisfies Choice, due to having Global Choice in the ground universe.

Proposition 2.15 (Over $\text{GBC}^- + \text{ETR}$). *The unrolled structure \mathfrak{U} satisfies Choice, in the guise of the well-ordering theorem.*⁸

Proof. Fix a membership code A . Appealing to Global Choice we may without loss of generality assume that $\text{dom } A \subseteq \text{Ord}$. We want to find a membership code W which codes a well-order of A . We do this by modifying A as follows. First, throw out t_A and any edges pointing to t_A to get A' . While A' is in general not a membership code we will modify it to get the desired W . For each $x, y \in \text{elts } A$ with $x < y$ (under the ordering from the ordinals) we will add the following nodes to A' : $\{x\}$, $\{x, y\}$ and $\{\{x\}, \{x, y\}\}$. We also add edges to A' in the obvious manner. That is, we add edges from x to $\{x\}$ and $\{x, y\}$, from y to $\{x, y\}$, from $\{x\}$ to $\{\{x\}, \{x, y\}\}$, and from $\{x, y\}$ to $\{\{x\}, \{x, y\}\}$. Finally, add a new node which will be t_W and add edges from each of the $\{\{x\}, \{x, y\}\}$ nodes to t_W . (See figure 2.3 for a picture of W .) Then W represents a well-order of A of ordertype $\leq \text{Ord}$. \square

The astute reader will note that Elementary Transfinite Recursion was used only so far in two places, namely to get Extensionality and Pairing. In both cases it was used via the lemma that there are maximum initial partial isomorphisms between membership codes. One might wonder whether this use of Elementary Transfinite Recursion can be avoided, whether by a different argument or by a different definition of membership code. See subsection 2.1.2 below for some discussion.

⁸In the absence of Powerset, the various forms of Choice are no longer equivalent [Zar96]. That every set can be well-ordered is stronger than the existence of choice functions.




Figure 2.3: A partial picture of W . The nodes x, y are from elts A with $x < y$. Below x and y the membership code W looks like A .

Let Γ be a well-order, possibly class sized. Then $\Gamma + 1 = \{(g, g') : g \Gamma g'\} \cup \{(g, \star) : g \in \text{dom } \Gamma\} \cup \{(\star, \star)\}$, where $\star \notin \text{dom } \Gamma$, is a membership code for an ordinal with ordertype Γ .⁹ So every well-order in the ground model corresponds to an ordinal in the unrolling. Of particular interest is the ordinal in the unrolling whose ordertype is the Ord of the ground universe.

Let κ denote the set in \mathfrak{U} represented by the membership code $\text{Ord} + 1$. Then $\mathfrak{U} \models \kappa$ is regular. If the ground universe moreover satisfies Powerset then $\mathfrak{U} \models \kappa$ is inaccessible. To prove these facts—and later results about \mathfrak{U} —it will be convenient to be able to translate facts about functions in \mathfrak{U} to facts about class functions on the penultimate level of membership codes, and vice versa.

I will slightly abuse notation by writing e.g. $F : A \rightarrow B$ to refer to a membership code for a function F in \mathfrak{U} from the set represented by A to the set represented by B .

Lemma 2.16 (Over $\text{GBc}^- + \text{ETR}$). *There is a correspondence between class functions $F : \text{elts } A \rightarrow \text{elts } B$ and membership codes for functions $F^* : A \rightarrow B$.*

Proof. Let us do the simpler direction first. For the backward direction, suppose that G is a membership code for a function from membership code A to membership code B . That is, if $A' \varepsilon A$ then there is $B' \varepsilon B$ so that (in \mathfrak{U}) G maps A' to B' . In particular, this holds for A' of the form $A \downarrow a$ for $a \dashv\!\!\dashv_A t_A$. Fix such a . Then there is a unique $b_a \dashv\!\!\dashv_B t_B$ so that $G(A \downarrow a) \cong B \downarrow b_a$. Set $G_\star(a) = b_a$. This yields (in the ground universe) a class function $G_\star : \text{elts } A \rightarrow \text{elts } B$.

For the forward direction of the correspondence, suppose that $F : \text{elts } A \rightarrow \text{elts } B$ is a class function. Let P be a membership code for the unordered pair whose elements are A and B . Such exists by proposition 2.14. Taking isomorphic copies if necessary we may assume without loss that $A, B \subseteq P$. We will now modify P to produce the membership code F^* . This is done similar to the argument in proposition 2.15 to construct a membership code for a well-ordering of a set. Namely, throw away t_P, t_A , and t_B to produce P' .¹⁰ This P' is in general not a membership code, but we will modify it to produce F^* . For $a \in \text{elts } A$ add nodes for $\{a\}$, $\{a, F(a)\}$ and $(a, F(a))$ to P' . Then add edges in the obvious way: add an edge from a to $\{a\}$, an edge from a to $\{a, F(a)\}$, an edge from $\{a\}$ to $(a, F(a))$, an edge from $F(a)$ to $\{a, F(a)\}$, and an edge from

⁹Of course, I assume here that $\text{dom } \Gamma \neq V$. But in case that does happen one can take an isomorphic copy of Γ with a smaller domain.

¹⁰Unless either $t_A \dashv\!\!\dashv_P t_B$ or $t_B \dashv\!\!\dashv_P t_A$, in which case keep, respectively, t_A or t_B .

$\{a, F(a)\}$ to $(a, F(a))$. Finally, add a new top element t_{F^*} and edges from each $(a, F(a))$ to t_{F^*} . This gives the desired F^* .




Figure 2.4: A partial picture of F^* .

It is now easy to check that $(F^*)_* = F$ and $(F^*)^* \cong F$. □

The same argument works more generally for relations.

Corollary 2.17 (Over $\text{GBC}^- + \text{ETR}$). *There is a correspondence between class relations $R \subseteq \text{elts } A \times \text{elts } B$ and membership codes for relations R^* between A and B .* □

Proposition 2.18 (Over $\text{GBC}^- + \text{ETR}$). *In \mathfrak{U} , κ is regular and the largest cardinal. Additionally, if the ground universe satisfies Powerset then κ is a strong limit, hence inaccessible.*

Proof. (κ is regular) Let Γ be a membership code for an ordinal shorter than Ord . Then Γ must be set-sized. Suppose towards a contradiction that there is a membership code for a cofinal function $F^* : \Gamma \rightarrow \text{Ord}$. Let $F : \text{elts } \Gamma \rightarrow \text{elts}(\text{Ord} + 1)$ be the corresponding function. I claim F must be cofinal in the order on $\text{elts}(\text{Ord} + 1) = \text{Ord}$. To see this, pick $\alpha \in \text{Ord}$. Then because F^* is cofinal there is $\beta > \alpha$ so that for some $g \in \text{elts } \Gamma$ we have $F^*(\Gamma \downarrow g) = (\text{Ord} + 1) \downarrow \beta$. But then $F(g) = \beta > \alpha$, establishing that F is cofinal. This is impossible, however, since Ord is a class and $\text{elts } \Gamma$ is a set.

(κ is the largest cardinal) Let Γ be a membership code for an ordinal longer than Ord . By Global Choice there is a class bijection $F : \text{elts}(\text{Ord} + 1) \rightarrow \text{elts } \Gamma$. Let F^* be the corresponding membership code for a function $F^* : \text{Ord} + 1 \rightarrow \Gamma$. We want to see that F^* represents a bijection. First, suppose that $x \neq y$ are in $\text{elts}(\text{Ord} + 1)$. Then $F(x) \neq F(y)$ and thus $F^*((\text{Ord} + 1) \downarrow x) \not\cong F^*((\text{Ord} + 1) \downarrow y)$. So f is one-to-one. To see f is onto α , take $a \in \text{elts } A$. because F is a bijection, there is $x \in \text{elts}(\text{Ord} + 1)$ so that $F(x) = a$. Thus $F^*((\text{Ord} + 1) \downarrow x) \cong A \downarrow a$.

(κ is strong limit, if the ground universe satisfies Powerset) Take Γ a membership code for a well-order shorter than Ord . Then Γ is set-sized. Because the ground universe satisfies powerset this means there is a set-sized membership code for 2^Γ and thus 2^Γ represents an ordinal less than κ . More explicitly, by the Mostowski collapse theorem we may assume without loss that $\Gamma = \in \uparrow (\gamma + 1)$ for some ordinal γ . Then

2^Γ is represented by the membership code $\in \uparrow (2^\gamma + 1)$, where 2^γ denotes cardinal exponentiation while the addition is ordinal addition. \square

It requires full second-order Comprehension to get full Separation in the unrolled model. However, it goes level by level.

To prove this we will need to translate between first-order formulae to be interpreted over \mathfrak{U} and second-order formulae to be interpreted in the ground universe. That is, given first-order φ there is a formula φ^* so that $\mathfrak{U} \models \varphi$ if and only if the ground universe satisfies φ^* . This purely syntactic translation is given by the following schema.

- If φ is of the form $x = y$ then φ^* is $X \cong Y$.
- If φ is of the form $x \in y$ then φ^* is $X \varepsilon Y$.
- If φ is of the form $\psi \vee \theta$ then φ^* is $\psi^* \vee \theta^*$, and similarly for conjunctions.
- If φ is of the form $\neg\psi$ then φ^* is $\neg\psi^*$.
- If φ is of the form $\exists x\psi(x)$ then φ^* is $\exists X X$ is a membership code and $\psi^*(X)$, and similarly for unbounded universal quantifiers.
- If φ is of the form $\exists x \in y\psi(x)$ then φ^* is $\exists x \multimap_Y t_Y \psi^*(Y \downarrow x)$, and similarly for bounded universal quantifiers.

Recall that capitals are for second-order variables while lowercase letters are for first-order variables. This translation assumes there is some fixed in advance correspondence between first-order variables and second-order variables so that e.g. x can be replaced by X . The details of this correspondence are completely uninteresting and will be suppressed.

Given a formula $\varphi(\bar{x})$ we translate it into a formula $\varphi^*(\bar{X})$. This is a purely syntactic translation defined via a recursion on the countable set of first-order formulae, which can be carried out in a weak fragment of ZFC^- . The translation is transparent to parameters: to handle formulae augmented with parameters $\varphi(\bar{A})$ simply consider $\varphi^*(\bar{A})$.

First we need to see that this translation is coherent, since we gave different translations for $\exists x \in y \varphi(x)$ and $\exists x (x \in y \wedge \varphi(x))$.

Lemma Schema 2.19 (Over GBc^-). *For all formulae $\varphi(x)$ in the language of first-order set theory, GBc^- proves that if A and B are isomorphic membership codes then $\varphi^*(A) \Leftrightarrow \varphi^*(B)$.*

Proof. This is an easy induction on formulae. \square

Lemma Schema 2.20 (Over GBc^-). *For all first-order formulae $\varphi(x)$ the theory GBc^- proves that $(\exists x x \in y \wedge \varphi(x))^*$ is equivalent to $(\exists x \in y \varphi(x))^*$.*

Proof. Consider a formula φ . Work in GBc^- and consider an arbitrary membership code Y .

(\Rightarrow) Suppose there is a membership code X so that $X \varepsilon Y$ and $\varphi^*(X)$. By the definition of ε there is $x \multimap t_Y$ so that $X \cong Y \downarrow x$. But then $\varphi^*(Y \downarrow x)$ holds because of the previous lemma plus the fact that $\varphi^*(X)$ holds, so we are done.

(\Leftarrow) Suppose there is $x \multimap t_Y$ so that $\varphi^*(Y \downarrow x)$ holds. We want to see there is a membership code X so that $X \varepsilon Y$ and $\varphi^*(X)$. Take $X = Y \downarrow x$. Done. \square

Lemma 2.21 (Over $\text{GBc}^- + \text{ETR}$). *Both $X \cong Y$ and $X \varepsilon Y$ are Δ_1^1 .*

Proof. The definitions of these are both Σ_1^1 . So it remains to see that they are equivalent to Π_1^1 assertions. Let us consider $A \varepsilon B$; similar reasoning works for $A \cong B$. If $A \varepsilon B$ then this must be witnessed by the maximum initial partial isomorphism $\pi: A \rightarrow B$. So to say that $A \not\varepsilon B$ it is equivalent to say that the maximum initial partial isomorphism $\pi: A \rightarrow B$ does not have a range of the form $B \downarrow a$ for some $a \multimap_B t_B$. Thus, $A \not\varepsilon B$ is

Σ_1^1 -expressible, via the formula asserting that there is an initial partial isomorphism $\pi : A \rightarrow B$ so that (1) for some $x \in \text{dom } A \setminus \text{dom } \pi$ there is no $y \in \text{dom } B$ so that $\pi \cup \{(x, y)\}$ is an initial partial isomorphism and (2) for all $a \dashv\!\!\dashv_B t_B$ there is $a' \leq_B a$ so that $a' \notin \text{ran } \pi$. Therefore $A \varepsilon B$ is equivalent to a Π_1^1 -formula. \square

Lemma 2.22 (Over $\text{GBC}^- + \text{ETR}$). *If φ is Σ_k for $k \geq 1$ then φ^* is equivalent to a Σ_k^1 formula.*

Proof. First, recall that being a membership code is a first-order property. So that cannot increase the logical complexity of φ^* .

Lemma 2.21 allows us to absorb the class quantifiers from $A \cong B$ and $A \varepsilon B$ into the innermost quantifier block of the translated formula. First, φ^* is equivalent to a formula in prenex normal form. In particular, all unbounded quantifiers are on the outside and \neg only occurs in front of atomic statements. Suppose that the innermost quantifier block consists of existential quantifiers. Then produce an equivalent Σ_k^1 -formula by replacing all instances of $\neg A \cong B$ and $\neg A \varepsilon B$ with the equivalent Σ_1^1 assertion. Thus, all bounded quantifiers in φ become Σ_1^1 assertions and thereby get absorbed into the innermost quantifier block. Similarly, if the innermost quantifier block consists of universal quantifiers then replace all instances of $A \cong B$ and $A \varepsilon B$ with the equivalent Π_1^1 assertion, thereby absorbing them. \square

Proposition 2.23 (Over $\text{GBC}^- + \text{ETR}$). *If the ground universe satisfies Π_k^1 -CA for $k \geq 1$ then \mathfrak{U} satisfies Σ_k -Separation. Thus, if the ground universe satisfies KM then the unrolled structure satisfies Separation.*

In fact, the backwards implication is also true; see section 2.2.

Proof. Fix a membership code A and a Σ_k formula $\varphi(x)$, possibly with (suppressed) parameters. We want to find a membership code B so that B represents the subset of A formed using φ . Because φ^* is equivalent to a Σ_k^1 formula we can by Π_k^1 -CA form the collection of $x \dashv\!\!\dashv_A t_A$ so that $\varphi^*(A \downarrow x)$ holds. Closing this collection downward in A and adding a top element gives a membership code B representing the subset of A formed using φ . \square

The case for Collection is similar. We need full Class Collection to get Collection in the unrolled model, but it goes level by level.

Proposition 2.24 (Over $\text{GBC}^- + \text{ETR}$). *If the ground universe satisfies Σ_k^1 -Class Collection then the unrolled structure \mathfrak{U} satisfies Σ_k -Collection. Thus if the ground universe satisfies KMCC then the unrolled structure satisfies Collection.*

In fact, the backwards implication is also true; see section 2.2.

Proof. We want to show that \mathfrak{U} satisfies every instance of Σ_k collection. That is, consider a Σ_k -formula φ and suppose there are membership codes A and P so that for all $X \varepsilon A$ there is a membership code Y so that $\varphi^*(X, Y, P)$. We want to find a membership code B which collects together all these witnessing Y .

By Σ_k^1 -Class Collection there is a class C of pairs (x, y) so that for each $x \dashv\!\!\dashv_A t_A$ we have $\varphi^*(A \downarrow x, (C)_{x, P})$. We want to glue together the $(C)_{x, P}$ s to get the desired membership code B . This will be done by means of a certain elementary transfinite recursion.

There are two layers to the recursion. The outer layer is along elts A , according to some fixed enumeration $\langle x_\alpha : \alpha \in \text{Ord} \rangle$. Each stage α of the outer layer yields a partial construction of B , let us call it B_α , along with P_α a partial listing of what will be the penultimate level of B . Stage 0 is merely to take $B_0 = \emptyset$ and $P_0 = \emptyset$. Having done stage α , for stage $\alpha + 1$ we must do the inner layer of the recursion to produce the maximum initial partial isomorphism $\pi : (C)_{x_\alpha} \rightarrow B_\alpha$. We then use it to define $B_{\alpha+1}$ by adding on everything in $(C)_{x_\alpha}$ which we have not yet accounted for. Formally, set

$$\begin{aligned} B_{\alpha+1} = & B_\alpha \\ & \cup (C)_{x_\alpha} \upharpoonright (\text{dom } C \setminus \text{dom } \pi) \\ & \cup \{(\pi(c), c') : (c, c') \in C \text{ and } c \in \text{dom } \pi \text{ and } c' \notin \text{dom } \pi\}. \end{aligned}$$

We also add either $t_{(C)x_\alpha}$ or $\pi(t_{(C)x_\alpha})$ as appropriate to P_α to get $P_{\alpha+1}$. At limit stages λ take unions: set $B_\lambda = \bigcup_{\alpha < \lambda} B_\alpha$ and $P_\lambda = \bigcup_{\alpha < \lambda} P_\alpha$. Finally, after Ord steps we produce B from B_{Ord} by adding a top element connecting each element of P_{Ord} to this new top element. Then B is the desired membership code. \square

The following theorem summarizes what \mathfrak{U} 's theory will be, based upon the theory of the ground universe.

Theorem 2.25. *Let (M, \mathcal{X}) be a model of second-order set theory and let \mathfrak{U} be the unrolled model constructed inside (M, \mathcal{X}) .*

- If $(M, \mathcal{X}) \models \text{KMCC}$ then $\mathfrak{U} \models \text{ZFC}_I^-$.
- If $(M, \mathcal{X}) \models \text{KMCC}^-$ then $\mathfrak{U} \models \text{ZFC}_R^-$.
- If $(M, \mathcal{X}) \models \text{KM}$ then $\mathfrak{U} \models \text{wZFC}_I^-$.
- If $(M, \mathcal{X}) \models \text{KM}^-$ then $\mathfrak{U} \models \text{wZFC}_R^-$.
- If $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$ then $\mathfrak{U} \models \text{ZFC}_I^-(k)$, for $k \geq 1$.
- If $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$ then $\mathfrak{U} \models \text{ZFC}_R^-(k)$, for $k \geq 1$.
- If $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA}$ then $\mathfrak{U} \models \text{wZFC}_I^-(k)$, for $k \geq 1$.
- If $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_k^1\text{-CA}$ then $\mathfrak{U} \models \text{wZFC}_R^-(k)$, for $k \geq 1$.

In section 2.2 we will get the backward results. But first we must address the conspicuous absence of $\text{GBC} + \text{ETR}$ and $\text{GBC}^- + \text{ETR}$ from this last theorem. The following subsection will also establish that for all of the above, the unrolling \mathfrak{U} satisfies Σ_0 -Transfinite Recursion. Of course, Σ_0 -Transfinite Recursion follows from Σ_1 -Collection, so this is only additional content for the theories which lack a fragment of Collection.

2.1.1 Unrolling in $\text{GBC} + \text{ETR}$

Above, to show that \mathfrak{U} satisfies a fragment of Separation and Collection—dependent upon the fragment of Comprehension and Class Collection satisfied in the ground universe—we used a translation $\varphi \mapsto \varphi^*$. Given a first-order formula φ in the language of set theory this translation produced a second-order formula φ^* about membership codes. We used that $A \varepsilon B$ and $A \cong B$ are Δ_1^1 over $\text{GBC}^- + \text{ETR}$, so that the class quantifiers arising from the translation of $a = b$ and $a \in b$ could be absorbed by the block of class quantifiers at the front of a prenex-normal form formula equivalent to φ^* .

This of course will not work if we do not have class quantifiers at the front. These came from translating unbounded quantifiers from φ . But, as we will finish seeing later, $\text{wZFC}_R^-(1)$ —which allows Separation for formulae with only a single unbounded quantifier in front—is bi-interpretable with $\text{GBC}^- + \Pi_1^1\text{-CA}$, which is stronger than $\text{GBC}^- + \text{ETR}$. So we cannot hope to absorb the class quantifiers arising from $A \varepsilon B$ and $A \cong B$. If we want to calculate the theory of unrollings of models of $\text{GBC} + \text{ETR}$ and $\text{GBC}^- + \text{ETR}$ we must use a different translation.

The key idea is that to check the truth of an assertion which only has bounded quantifiers it suffices to look at a single membership code. Before giving the translation let me illustrate this with an example.

We work in $\text{GBC}^- + \text{ETR}$. Suppose we are given a membership code E and we wish to know whether E represents an ordered pair in the unrolling, say according to Kuratowski's definition. Formally, we want to know whether

$$\begin{aligned} \exists A, B \varepsilon E \left(\forall X \varepsilon E (X \cong A \vee X \cong B) \right. \\ \wedge \forall Y, Y' \varepsilon A (Y \cong Y' \wedge Y \varepsilon B \\ \left. \wedge \forall Z, Z' \varepsilon B ([Z \not\cong Y \wedge Z' \not\cong Y] \Rightarrow Z \cong Z')) \right) \end{aligned}$$

is true. Expanding out the definitions of \cong and ε this formula has fourteen class quantifiers, which is too many to handle directly in $\text{GBC} + \text{ETR}$. But we do not have to look at, for example, all membership codes A, B so that $A, B \varepsilon E$. It suffices to look just at all membership codes of the form $E \downarrow x$ for $x \in \text{elts } E$. Similar considerations apply for the other instances of ε or \cong in this formula. So it is equivalent to ask whether

$$\begin{aligned} \exists a, b \text{---}\circ_E t_E \left(\forall x \text{---}\circ_E t_E (x = a \vee x = b) \right. \\ \left. \wedge \forall y, y' \text{---}\circ_E a (y = y' \wedge y \text{---}\circ_E b) \right. \\ \left. \wedge \forall z, z' \text{---}\circ_E b ([z \neq y \wedge z' \neq y] \Rightarrow z = z') \right) \end{aligned}$$

is true. This formula only has set quantifiers.

Definition 2.26 (Over $\text{GBC}^- + \text{ETR}$). Let $\varphi(A_1, \dots, A_n)$ be a Σ_0 -formula in the language of first-order set theory with parameters A_1, \dots, A_n , which are membership codes, and possibly with free variables. Then $\varphi^*(\bar{A})$ is defined as follows. First, let P be the membership code for the ‘set’ $\{A_1, \dots, A_n\}$. This can be constructed by an instance of Elementary Transfinite Recursion, similar to the construction of unordered pairs. For $1 \leq i \leq n$ let $a_i \text{---}\circ_P t_P$ be the unique member of the penultimate level of P so that $P \downarrow a_i \cong A_i$. Then $\varphi^*(\bar{A})$ is defined by the following schema. Here, t and s are either variables x, y, z, \dots or one of the a_i ’s.

- If φ is $t = s$ then φ^* is $t = s$.
- If φ is $t \in s$ then φ^* is $t \text{---}\circ_P s$.
- If φ is $\psi \wedge \theta$ then φ^* is $\psi^* \wedge \theta^*$, and similarly for disjunctions.
- If φ is $\neg\psi$ then φ^* is $\neg\psi^*$.
- If φ is $\exists x \in t \psi(x)$ then φ^* is $\exists x \text{---}\circ_P t \psi^*(x)$, and similarly for bounded universal quantification.

Thus, φ^* is a Σ_ω^0 -formula in the parameter P .

This translation is not purely syntactic, since we needed to know the parameters \bar{A} to construct P . Nevertheless, given $\varphi(\bar{A})$ we can construct $\varphi^*(P)$ by an elementary recursion. So $\text{GBC}^- + \text{ETR}$ lets us carry out the translation.

Proposition 2.27 (Over $\text{GBC}^- + \text{ETR}$). *The unrolled model \mathfrak{U} satisfies Σ_0 -Separation.*

Proof. Fix a membership code A and a Σ_0 -formula $\varphi(x, \bar{B})$ with parameters \bar{B} . By Elementary Comprehension form the class of all $x \text{---}\circ_A t_A$ so that $\varphi^*(x, P)$ holds, where P is constructed from \bar{B} and A as in the translation. Closing this collection downward in A and adding a top element gives a membership code representing the subset of A formed using φ . \square

We get more than Σ_0 -Separation. As a warm-up, let us see that the unrolled model has transitive collapses.

Proposition 2.28 (Over $\text{GBC}^- + \text{ETR}$). *The unrolled model \mathfrak{U} satisfies Mostowski’s collapse lemma that every well-founded extensional binary relation is isomorphic to the restriction of the membership relation to some set.*

I provide only a sketch of an argument, as we will prove a more general statement later.

Proof sketch. Let E be a membership code for a well-founded extensional relation on A . By corollary 2.17 take the corresponding class relation $E_\star \subseteq \text{elts } A \times \text{elts } A$. Then E_\star is itself almost a membership code; all that is missing is that it does not have a top element. Consider the membership code $F = E_\star \cup \{(e, \dagger) : e \in \text{dom } E_\star\}$ where \dagger is a new element. Then in the unrolling E is isomorphic to the membership relation restricted to the set represented by F . \square

More generally, the unrolled model \mathfrak{U} starting from a model of $\text{GBC}^- + \text{ETR}$ will satisfy a transfinite recursion principle.

Definition 2.29. Let Φ be a collection of formulae in the language of first-order set theory. Then Φ -*Transfinite Recursion* is the axiom schema consisting of the following axiom for each $\varphi(x, y, a) \in \Phi$:

Suppose a is a parameter so that $\varphi(x, y, a)$ defines a class function $F : V \rightarrow V$ and D is a set equipped with a well-ordering $<_D$. Then there is a function $s : D \rightarrow V$ so that for all $d \in D$ we have $s(d) = F(s \upharpoonright d)$ where $s \upharpoonright d$ means $s \upharpoonright \{d' \in D : d' <_D d\}$.

Before seeing that the unrolled structure satisfies Σ_0 -Transfinite Recursion let us justify the “more generally” above and check that Σ_0 -Transfinite Recursion (along with the other axioms we already know to be satisfied by the unrolled model) proves Mostowski’s collapse lemma.

Proposition 2.30. *The theory $\text{wZFC}_{\mathbb{R}}^-(0) + \Sigma_0$ -Transfinite Recursion¹¹ proves Mostowski’s collapse lemma.*

Proof. Let e be a binary well-founded, existential relation on a set D . We want to see that there is a function π with domain D satisfying $\pi(d) = \{\pi(d') : d' e d\}$. This is a recursive requirement on π , where the property we want to recursively satisfy is Σ_0 . So it exists by Σ_0 -Transfinite Recursion. \square

Proposition 2.31 (Over $\text{GBC}^- + \text{ETR}$). *The unrolled structure \mathfrak{U} satisfies Σ_0 -Transfinite Recursion.*

Proof. Consider an instance of Σ_0 -Transfinite Recursion. That is, F is a class function $\mathfrak{U} \rightarrow \mathfrak{U}$ which is Σ_0 -definable, possibly using a membership code as a parameter, and D is a membership code equipped with $<_D$ a membership code for a well-ordering of D . We want to see that there is a membership code for the desired s . First, observe that in the ground universe that F is first-order definable (from parameters), by the translation.

We will build the the desired membership code S via an instance of Elementary Transfinite Recursion. The idea is to mimic the recursion to produce s , but in membership codes. This introduces some extra work, since we have to deal with the picky details of how membership codes work.

The iteration proceeds as follows, with an outer layer and an inner layer. The outer layer occurs on elts D according to the well-ordering corresponding to the membership code $<_D$ (see corollary 2.17). Each step d in the outer layer produces a partial construction of S , call it S_d . We start with $S_0 = \emptyset$ and take unions at limit stages. The hard work is done in the successor step, where the inner layer of the transfinite recursion occurs. We start with S_d and want to produce S_{d+1} . By construction, each $d' \in \text{elts } D$ which comes before d in $<_D$ is in S_d . More, there is a corresponding node, call it $f(d')$, which represents $F(S_d \upharpoonright d')$ and then nodes for $\{d'\}$, $\{d', f(d')\}$, and $(d', f(d'))$ above, similar to the constructions in proposition 2.15 and lemma 2.16. In particular, S_d itself may not be a membership code. Modify S_d to produce a membership code U by adding a top node t_U and edges from each $(d', f(d'))$ node in S_d to t_U . Then, we have a membership code $F(U)$ by Elementary Comprehension. Construct by transfinite recursion the maximum initial partial isomorphism between S_d and $F(U)$ and use it to glue a copy of $F(U)$ onto S_d , as in the argument for proposition 2.24. Then, add $d+1$ to S_d along with nodes for $\{d+1\}$, $\{d+1, t_{F(U)}\}$, and $(d+1, t_{F(U)})$ and the corresponding edges to produce S_{d+1} . \square

This completes the last step in the calculation of the theory of the unrolled model starting with a ground universe satisfying $\text{GBC} + \text{ETR}$ or $\text{GBC}^- + \text{ETR}$.

Corollary 2.32. *Let (M, \mathcal{X}) be a model of second-order set theory and let \mathfrak{U} be the unrolled model constructed inside (M, \mathcal{X}) .*

- *If $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$ then $\mathfrak{U} \models \text{wZFC}_{\mathbb{R}}^-(0) + \Sigma_0$ -Transfinite Recursion.*
- *If $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}$ then $\mathfrak{U} \models \text{wZFC}_{\mathbb{I}}^-(0) + \Sigma_0$ -Transfinite Recursion.*

In section 2.2 we will get the backward direction.

¹¹That is, the theory with axioms Extensionality, Pairing, Union, Infinity, Foundation, Choice (in the guise of the well-ordering theorem), Σ_0 -Separation, and Σ_0 -Transfinite Recursion.

2.1.2 How hard is it to unroll?

I wish now to address the question of what is needed for the unrolling to satisfy a reasonable theory. As the reader who has familiarity with coding hereditarily countable sets as reals will know, there is more than one way to do the coding. Which coding one prefers is, to a significant extent, a matter of taste. However, it means I must also address whether my choice of coding affects how strong a theory is needed in the ground universe to carry out the unrolling.

First, let us consider other codings based on well-founded extensional directed graphs. For example, rather than requiring the graph to have a largest element one could work with pointed graphs with a designated point. If (A, p_A) and (B, p_B) are two such pointed graphs then they represent the same ‘set’ if $A \downarrow p_A \cong B \downarrow p_B$, with a corresponding definition of ε .¹² It is not hard to see that such a coding will require the existence of maximum initial partial isomorphisms in the same places where they are required by membership codes according to my definition. So this coding is no easier than the one I use.

The question then is, what is needed to carry out my coding? As remarked in section 2.1, GBC^- suffices to define membership codes, isomorphism between them, and their membership relation ε . So it takes very little just to unroll into some structure. But we do not want to unroll to some arbitrary structure, we want to unroll into a model of (first-order) set theory. This lacks a precise definition, but it should be uncontroversial to say that to be a model of set theory a structure must at least satisfy the basic axioms: Extensionality, Pairing, Union, and so forth. Earlier, we used Elementary Transfinite Recursion to show that the unrolling satisfies Extensionality and Pairing. More specifically, $\text{GBC}^- + \text{ETR}$ proves the existence of maximal initial partial isomorphisms between membership codes, which is what we used. Indeed, the existence of such is necessary.

Proposition 2.33 (Over GBC^-). *If \mathcal{U} satisfies Pairing then if A and B are membership codes there is a maximum initial partial isomorphism between them.*

Proof. Fix A and B . Then there is a membership code P so that $E \varepsilon P$ if and only if $E \cong A$ or $E \cong B$. Let $\pi_A : A \rightarrow P$ and $\pi_B : B \rightarrow P$ be the embeddings into P . Then $\pi_B^{-1} \circ \pi_A \upharpoonright (\text{ran } A \cap \text{ran } B)$ is the maximum initial partial isomorphism from A to B . \square

A special case of this is of interest.

Corollary 2.34 (Over GBC^-). *If the unrolling \mathcal{U} satisfies Pairing then the ground model satisfies the comparability of class well-orders: given class well-orders Γ and Δ either Γ embeds as an initial segment of Δ or vice versa.*

Proof. Let A and B be membership codes for ordinals. That is, A and B are class well-orders of successor length. The initial partial isomorphism π between them gives a comparison map. If neither $\text{dom } \pi = A$ nor $\text{ran } \pi = B$ then we can extend π to a larger initial partial isomorphism, namely by mapping $\min(A \setminus \text{dom } \pi)$ to $\min(B \setminus \text{ran } \pi)$. Note that this definition can be done using Elementary Comprehension, because it is a first-order property to be the minimum of a well-order. So it must be that either $\text{dom } \pi = A$, in which case A has ordertype \leq that of B , or $\text{ran } \pi = B$, in which case B has ordertype \leq that of A .

The general case then follows because Γ and Δ are comparable if and only if $\Gamma + 1$ and $\Delta + 1$ are comparable. \square

What does it take to show that class well-orders are always comparable? It is clear that $\text{GBC}^- + \text{ETR}$ suffices; we want a map which satisfies the recursive requirement that $e(g)$ is the least element of $\text{dom } \delta$ which is above the range of $e \upharpoonright g$. Can we get away with less than Elementary Transfinite Recursion?

In second-order arithmetic, the answer to the analogous question is no. Over RCA_0 , the comparability of well-orders is equivalent to ATR_0 (see [Sim09, chapter V.6]). But this proof uses that in arithmetic, “ X is a well-order” is Π_1^1 -universal. This fails badly in set theory, as in this context “ X is a well-order” is an elementary assertion. So the proof from second-order arithmetic will not generalize.

¹²I leave it to the reader to explicitly write down the definition of ε for this coding.

Question 2.35. *Over GBC^- , what is the strength of the comparability of class well-orders?*

Let us now look at alternative codings which use trees instead of graphs. Perhaps there things may be easier. This is the approach taken by Antos and Friedman [AF], whom one should consult for precise definitions. To illustrate the idea, consider a set-sized well-founded tree $(T, <^T)$. By the Mostowski collapse lemma, there is a map π which maps T to a transitive set a so that $s <^T t$ if and only if $\pi(s) \in \pi(t)$. (This π will not in general be one-to-one.) Rather than representing sets with graphs we can represent them with certain trees. Membership is defined similarly as to my membership codes; say that $S \varepsilon^{\text{tree}} T$ if there is a child node t of the root node of T so that S is isomorphic to the subtree of T below t .

An advantage of this coding is that satisfying Pairing is trivial. To form a code for the unordered pair of S and T just make a tree whose root has two children, which are the roots of S and T . However, Extensionality still gives trouble. Similar to the case with membership codes, we want to conclude that $S \cong T$ knowing that $U \varepsilon^{\text{tree}} S$ if and only if $U \varepsilon^{\text{tree}} T$. This amounts to wanting to glue together partial isomorphisms between S and T to get a full isomorphism. Naively, one wants to define an isomorphism π between S and T by mapping subtrees from a child of the root of S to the isomorphic copy under the root of T and mapping the root of S to the root of T . But this uses an obfuscated appeal to Π_1^1 -Comprehension; cf. the remarks after lemma 2.10. The same idea used to show my membership codes satisfy Extensionality can be used here, namely using Elementary Transfinite Recursion to conclude there is a maximal initial partial isomorphism between the trees and then showing that this must be a full isomorphism. So we look to be in no better condition by coding with trees instead of graphs.

2.2 The cutting off construction

The other direction of the bi-interpretability results is more straightforward. Suppose we have a first-order model of set theory which has a largest cardinal κ and H_κ exists in the model so that $H_\kappa \models \text{ZFC}^-$. We can construct a second-order model whose first-order part is H_κ and whose second-order part is the (proper class of) subsets of H_κ in the model. The stronger the theory satisfied by the first-order model, the stronger the theory that will be satisfied by the second-order model gotten by this cutting off.

Formally, we will define an interpretation of formulae in the language of second-order set theory in the language of first-order set theory with a constant symbol κ for the largest cardinal of the (first-order) model. Of course, κ is definable so this use of a constant symbol is only a convenience. This interpretation $\varphi \mapsto \varphi^I$ is given by the following schema.

- The interpretation of $x \in y$ is $x \in y$ and the interpretation of $x \in Y$ is $x \in Y$ —that is, both membership relations for the second-order model will be the membership relation of the first-order model.
- The interpretation of $x = y$ is $x = y$ and the interpretation of $X = Y$ is $X = Y$. Unlike the unrolling construction we can directly use equality and not have to quotient out by an equivalence relation.
- The interpretation of $\varphi \wedge \psi$ is $\varphi^I \wedge \psi^I$, and similarly for disjunction and negation;
- The interpretation of $\forall x \varphi$ is $\forall x (x \in H_\kappa \Rightarrow \varphi^I)$ and similarly for first-order existential quantification; and
- The interpretation of $\forall X \varphi$ is $\forall x (x \subseteq H_\kappa \Rightarrow \varphi^I)$ and similarly for second-order existential quantification.

It is immediate that if φ is Σ_k^1 (in parameters) then φ^I is Σ_k (in parameters).

As a base theory for this section I will take $\text{wZFC}_R^-(0) + \Sigma_0$ -Transfinite Recursion. Recall that $\text{wZFC}_R^-(0)$ is the set theory axiomatized by Extensionality, Union, Pairing, Infinity, Foundation, Choice, Σ_0 -Separation, plus the assertions that there is a largest cardinal κ , that κ is regular, and that H_κ exists. In particular, $\text{wZFC}_R^-(0) + \Sigma_0$ -Transfinite Recursion proves that $H_\kappa \models \text{ZFC}^-$.

Proposition 2.36. *Work in $\text{wZFC}_R^-(0) + \Sigma_0$ -Transfinite Recursion and let κ be the largest cardinal. Then $H_\kappa \models \text{ZFC}^-$.*

Proof sketch. I will show that H_κ satisfies Separation. The rest is an easy exercise for the reader. Consider $a, p \in H_\kappa$ and fix some formula $\varphi(x, p)$. We want to see that $b = \{x \in a : \varphi(x, p)^{H_\kappa}\} \in H_\kappa$. Let T be the truth predicate for H_κ , which exists by an instance of Σ_0 -transfinite recursion. We can then define b by Σ_0 -Separation, namely to consist of those $x \in a$ for which $(\varphi, x \wedge p) \in T$. \square

We can now see that all the axioms of **GBC** are satisfied by the cut-off model.

Proposition 2.37. *If $\varphi \in \text{GBC}^-$, then $\text{wZFC}_R^-(0) + \Sigma_0$ -Transfinite Recursion $\vdash \varphi^I$. Also, if $\varphi \in \text{GBC}$ then $\text{wZFC}_I^-(0) + \Sigma_0$ -Transfinite Recursion $\vdash \varphi^I$.*

Proof. That (the interpretation of) ZFC^- holds for the first-order part is because $H_\kappa \models \text{ZFC}^-$. If κ is moreover inaccessible then $H_\kappa = V_\kappa$ moreover satisfies Powerset, thus full **ZFC**. Extensionality for classes holds because of Extensionality in the first-order model. Global Choice holds because Choice holds in the ground model: Any well-order of H_κ must have ordertype α for some α of size κ , since κ is the largest cardinal. But then we can use this to get a well-order of ordertype κ , from which we can extract a bijection $\kappa \rightarrow H_\kappa$. Replacement holds because κ is regular. If φ is an instance of Elementary Comprehension, then φ^I is Σ_0 and thus holds by Separation applied to H_κ . \square

To get that the cut off model satisfies Π_k^1 -Comprehension requires Σ_k -Separation from the ground universe.

Proposition 2.38. *If $\varphi \in \Pi_k^1\text{-CA}^-$, for $k \geq 1$, then $\text{wZFC}_R^-(k) + \Sigma_0$ -Transfinite Recursion, i.e. $\text{wZFC}_R^-(0) + \Sigma_0$ -Transfinite Recursion $+ \Sigma_k$ -Separation, proves φ^I . Thus, if $\varphi \in \Pi_k^1\text{-CA}$, for $k \geq 1$, then $\text{wZFC}_I^-(k) + \Sigma_0$ -Transfinite Recursion proves φ^I .*

Proof. Consider φ is an instance of Π_k^1 -Comprehension, i.e. φ asserts that there is a class whose members are precisely those sets satisfying ψ where ψ is Π_k^1 , possibly with parameters. By Σ_k -Comprehension form A the subset of H_κ consisting of those sets which satisfy ψ^I . This A is the desired class in the cut off model. \square

As an immediate corollary we get full Comprehension if the ground universe satisfies $\text{wZFC}_R^- + \Sigma_0$ -Transfinite Recursion.

Corollary 2.39. *If $\varphi \in \text{KM}^-$ then $\text{wZFC}_R^- + \Sigma_0$ -Transfinite Recursion $\vdash \varphi^I$. Thus, if $\varphi \in \text{KM}$ then $\text{wZFC}_I^- + \Sigma_0$ -Transfinite Recursion $\vdash \varphi^I$. \square*

Next we turn to Class Collection.

Proposition 2.40. *If $\varphi \in \Sigma_k^1$ -Class Collection, for $k \geq 1$, then $\text{ZFC}_R^-(k)$ proves φ^I .¹³*

Proof. Let φ be an instance of Σ_k^1 -Class Collection. Then, φ^I asserts that if for every $x \in H_\kappa$ there is $Y \subseteq H_\kappa$ so that $\psi^I(x, Y, P)$, then there is $Z \subseteq H_\kappa$ so that for every $x \in H_\kappa$ we have $\psi^I(x, (Z)_x, P)$. Because ψ is Σ_k^1 we have that ψ^I is Σ_k . Apply Σ_k -Collection to get B so that for all $x \in H_\kappa$ there is $Y \in B$ so that $Y \subseteq H_\kappa$ and $\psi^I(x, Y, P)$. Define

$$Z = \left\{ (x, b) \in V_\kappa \times \bigcup B : b \in Y \text{ where } Y \text{ is least so that } \psi^I(x, Y, P) \right\}.$$

Here, Y is least by a fixed well-order of B . Then Z is manifestly a subset of H_κ and for all $x \in H_\kappa$ we have $\psi^I(x, (Z)_x, P)$. \square

As an immediate corollary, full Collection in the ground model translates to full Class Collection in the cut off model.

¹³Note that $\text{ZFC}_R^-(k)$ includes Σ_0 -Transfinite Recursion.

Corollary 2.41. *If $\varphi \in \text{KMCC}^-$ then $\text{ZFC}_R^- \vdash \varphi^I$. Thus if $\varphi \in \text{KMCC}$ then $\text{ZFC}_I^- \vdash \varphi^I$.* \square

Finally, let us see what we need to get merely $\text{GBC} + \text{ETR}$ or $\text{GBC}^- + \text{ETR}$.

Proposition 2.42. *If $\varphi \in \text{GBC}^- + \text{ETR}$ then $\text{wZFC}_R^-(0) + \Sigma_0$ -Transfinite Recursion proves φ^I .*

Proof. Consider an instance of Elementary Transfinite Recursion. We want to find the subset of H_κ which witnesses that instance holds. This is done in the obvious way by an instance of Σ_0 -Transfinite Recursion. \square

Altogether we can summarize the results so far in this section.

Theorem 2.43. *Let N be an appropriate model for performing the cut off construction and let (M, \mathcal{X}) be the cut off model constructed from N .*

- *If $N \models \text{ZFC}_I^-$ then $(M, \mathcal{X}) \models \text{KMCC}$.*
- *If $N \models \text{ZFC}_R^-$ then $(M, \mathcal{X}) \models \text{KMCC}^-$.*
- *If $N \models \text{wZFC}_I^- + \Sigma_0$ -Transfinite Recursion then $(M, \mathcal{X}) \models \text{KM}$.*
- *If $N \models \text{wZFC}_R^- + \Sigma_0$ -Transfinite Recursion then $(M, \mathcal{X}) \models \text{KM}^-$.*
- *If $N \models \text{ZFC}_I^-(k)$ then $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$, for $k \geq 1$.*
- *If $N \models \text{ZFC}_R^-(k)$ then $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$, for $k \geq 1$.*
- *If $N \models \text{wZFC}_I^-(k) + \Sigma_0$ -Transfinite Recursion then $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA}$, for $k \geq 1$.*
- *If $N \models \text{wZFC}_R^-(k) + \Sigma_0$ -Transfinite Recursion then $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_k^1\text{-CA}$, for $k \geq 1$.*
- *If $N \models \text{wZFC}_I^-(0) + \Sigma_0$ -Transfinite Recursion then $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}$.*
- *If $N \models \text{wZFC}_R^-(0) + \Sigma_0$ -Transfinite Recursion then $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$.*

Let us now see that these interpretations, combined with the ones from the previous section, give bi-interpretability results. We need to see that performing one construction then its inverse brings us back to the original model.

Theorem 2.44. *Let $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$ and let \mathfrak{U} be the unrolled model constructed from (M, \mathcal{X}) . Then, if (M', \mathcal{X}') is the cut off model constructed from \mathfrak{U} we get $(M, \mathcal{X}) \cong (M', \mathcal{X}')$. Moreover, this isomorphism is definable over (M, \mathcal{X}) .*

Proof. Fix a class A . Then in \mathfrak{U} this class A is represented by E_A , the canonical membership code representing A . So the map $A \mapsto E_A$ gives an isomorphism $(M, \mathcal{X}) \cong (M', \mathcal{X}')$. \square

And in the other direction.

Theorem 2.45. *Let $N \models \text{wZFC}_R^-(0) + \Sigma_0$ -Transfinite Recursion and let (M, \mathcal{X}) be the cut off model constructed from N . Then if \mathfrak{U} is the unrolled model constructed from (M, \mathcal{X}) we get $N \cong \mathfrak{U}$. Moreover, this isomorphism is definable over N .*

Proof. Let κ be the largest cardinal in N . Then for any set $a \in N$ there is a binary relation E_a on H_κ so that $(\text{TC}(\{a\}), \in) \cong E_a$. This uses that $\text{TC}(\{a\})$ exists in N , which follows from Σ_0 -Transfinite Recursion. So in the cut off model (M, \mathcal{X}) we have that E_a is a membership code representing the set a in the unrolling. This map $a \mapsto E_a$ gives the isomorphism $N \cong \mathfrak{U}$. \square

Altogether this finishes the proof of theorem 2.2.

Let me remark on the role of Σ_0 -Transfinite Recursion here. The reader may have noticed that Σ_0 -Transfinite Recursion played little role in the arguments for calculating the theory of the cut off model. It was used to conclude the cut off model satisfies ETR given a weak base theory, but it was not used to get (fragments of) Comprehension in the cut off model. However, it plays an important role for the bi-interpretability results.

To illustrate this, take κ an inaccessible cardinal. Then $V_{\kappa+\omega}$ satisfies all the axioms of ZFC except the Collection schema. Let W consist of all subsets of H_{κ^+} which appear in $V_{\kappa+\omega}$. Then $W \models \text{wZFC}_1^-$. So cutting off W gives $(V_\kappa, \mathcal{X}) \models \text{KM}$. But when we unroll (V_κ, \mathcal{X}) we get back more than W . For instance, this unrolling contains an ordinal γ for each well-order in \mathcal{X} , but these will of course often be larger than $\kappa + \omega$. So applying the cutting off construction followed by the unrolling construction starting from W does not produce an isomorphic copy of W .

To finish off this section let us consider a basic but useful fact about unrollings.

Proposition 2.46. *Let $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$ and let \mathfrak{U} be its unrolling. Then \mathfrak{U} is well-founded if and only if (M, \mathcal{X}) is a β -model.*

Proof. (\Rightarrow) If \mathfrak{U} is well-founded then every membership code in (M, \mathcal{X}) is well-founded. But every successor class well-order is a membership code and every relation (M, \mathcal{X}) thinks is well-founded has a ranking function into a class well-order, by ETR. Since the class well-orders are actually well-founded, so must every relation the model thinks is well-founded.

(\Leftarrow) If (M, \mathcal{X}) is a β -model then in particular every membership code in \mathcal{X} is well-founded. \square

2.3 The constructible universe in the classes

In this section I will expoit a construction of the constructible universe where we iterate longer than Ord. If (M, \mathcal{X}) satisfies a sufficiently strong theory then it can carry out this construction. Let \mathcal{L} denote the hyperclass consisting of constructible classes from \mathcal{X} . Then we will see that (L^M, \mathcal{L}) has a nice theory, with the precise details depending upon the theory of (M, \mathcal{X}) . We will also carry out the construction of the constructible relative to a class parameter. The main parameter of interest to us will be (N, G) where $N \in \mathcal{X}$ is an inner model of M (possibly $N = M$) and G is a global well-order of N . Letting $\mathcal{L}(N, G)$ denote the hyperclass of (N, G) -constructible classes we will then get that $(N, \mathcal{L}(N, G))$ satisfies a nice theory. This will imply that for many natural second-order set theories T that being T -realizable is closed under taking inner models.

Theorem 2.47. *Let $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$ and suppose $N \in \mathcal{X}$ is an inner model (of ZFC^- or of ZFC , as appropriate) of M . Then there is a hyperclass $\mathcal{Y} \subseteq \mathcal{X} \cap \mathcal{P}(N)$ second-order definable in \mathcal{X} so that the following.*

- If $(M, \mathcal{X}) \models \text{KM}$ then $(N, \mathcal{Y}) \models \text{KMCC}$.
- If $(M, \mathcal{X}) \models \text{KM}^-$ then $(N, \mathcal{Y}) \models \text{KMCC}^-$.
- If $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA}$ then $(N, \mathcal{Y}) \models \text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$.
- If $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_k^1\text{-CA}$ then $(N, \mathcal{Y}) \models \text{GBC}^- + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$.
- If $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}$ then $(N, \mathcal{Y}) \models \text{GBC} + \text{ETR} + \text{ECC}$.
- If $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$ then $(N, \mathcal{Y}) \models \text{GBC}^- + \text{ETR} + \text{ECC}$.

The last two, that being $(\text{GBC} + \text{ETR})$ -realizable or $(\text{GBC}^- + \text{ETR})$ -realizable is closed under inner models, will be proved in the next chapter. The rest will follow from results in this section.

The special case where $N = M$ yields the following result, which I state separately.

Corollary 2.48. *Let $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$. Then there is a definable hyperclass $\mathcal{Y} \subseteq \mathcal{X}$ so that the following.*

- *If $(M, \mathcal{X}) \models \text{KM}$ then $(M, \mathcal{Y}) \models \text{KMCC}$.*
- *If $(M, \mathcal{X}) \models \text{KM}^-$ then $(M, \mathcal{Y}) \models \text{KMCC}^-$.*
- *If $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA}$ then $(M, \mathcal{Y}) \models \text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$, for $k \geq 1$.*
- *If $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_k^1\text{-CA}$ then $(M, \mathcal{Y}) \models \text{GBC}^- + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$, for $k \geq 1$.*
- *If $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}$ then $(M, \mathcal{Y}) \models \text{GBC} + \text{ETR} + \text{ECC}$.*
- *If $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$ then $(M, \mathcal{Y}) \models \text{GBC}^- + \text{ETR} + \text{ECC}$.*

That is, we can always obtain (a fragment of) Class Collection by moving to a possibly smaller V -submodel.

The strategy in this section will be to define \mathcal{L} somewhat indirectly. I will define \mathfrak{L} to consist of the constructible membership codes. Then, \mathcal{L} will be the classes which are coded in \mathfrak{L} . (Similar remarks apply when constructing relative to a parameter.) This approach follows the usual construction of L . But since we want to iterate the construction beyond Ord we will be too high in rank to use actual sets and will settle for membership codes for sets. This approach also has the advantage that some of the arguments are trivial modifications of arguments from section 2.1 of this chapter. As such, the reader is strongly encouraged to read section 2.1 before reading this section.

Let us move now to the definitions.

Definition 2.49 (Over $\text{GBC}^- + \text{ETR}$). Let Γ be a well-order. Denote by L_Γ the class obtained by iterating the definition of L along Γ . To clarify, L_Γ is the membership code constructed according to the following recursion:

- L_0 is a membership code for \emptyset .
- $L_{\alpha+1}$ is a membership code for $\text{Def}(L_\alpha)$.
- L_λ is a membership code for $\bigcup_{\alpha < \lambda} L_\alpha$, for λ a limit element of Γ .

This definition cries out for clarification. Recall that $\text{GBC}^- + \text{ETR}$ proves the existence of (first-order) truth predicates relative to a class. In particular, if (A, R_0, R_1, \dots) is a class-sized structure then by Elementary Transfinite Recursion can be constructed the truth predicate for (A, R_0, R_1, \dots) ; simply take the truth predicate for $(V, \in, A, R_0, R_1, \dots)$ and throw away the irrelevant facts about the background universe. In particular, this works for membership codes. So if E is a membership code which represents a ‘set’ e we can build the truth predicate for $(\text{dom } E, E)$ and from it extract a membership code D_E for the ‘set’ $\text{Def}(e)$. So ETR lets us construct membership codes that are of the form L_Γ . We will see in the next chapter that weaker principles will not suffice.

Of course, these membership codes are not unique, as there will be many isomorphic copies of L_Γ . But if we fix in advance the details of how we build D_E from E then we have fixed a way to build L_Γ from Γ .¹⁴ So in this sense it is justified to refer to L_Γ as *the* membership code obtained by iterated the definition of L along Γ .

Definition 2.50 (Over $\text{GBC}^- + \text{ETR}$). Say that a membership code E is *constructible* if $E \in L_\Gamma$ for some Γ . Let \mathfrak{L} consist of the membership codes for all constructible classes. We consider \mathfrak{L} as a first-order model by quotienting out by isomorphism and using ε (see definition 2.11) for its membership relation. I will call \mathfrak{L} the *constructible unrolled model*. Let \mathcal{L} consist of all classes coded in \mathfrak{L} .

¹⁴I will not go into the details of the necessary coding, as it is routine and uninteresting.

In particular, each L_Γ is constructible, as $L_\Gamma \varepsilon L_{\Gamma+1}$. Note that this is true even if Γ is very complicated and codes a lot of information. It will nevertheless be isomorphic to $L_{\Gamma+1} \downarrow e$ for some $e \in \text{elts}(L_{\Gamma+1})$.

As in the usual first-order setting we can carry out the construction of L relative to some parameter(s).

Definition 2.51. Let P be a class and Γ a well-order. Denote by the $L_\Gamma(P)$ the membership code obtained by iterating the definition of $L(P)$ along Γ .

Definition 2.52. A membership code E is P -constructible if $E \varepsilon L_\Gamma(P)$ for some Γ .

Similar to before, let $\mathfrak{L}(P)$ consist of the membership codes for all constructible classes. We consider $\mathfrak{L}(P)$ as a first-order model by quotienting out by isomorphism and using ε as its membership relation. I will call $\mathfrak{L}(P)$ the P -constructible unrolled model. Let $\mathcal{L}(P)$ consist of all classes coded in $\mathfrak{L}(P)$.

Hereon work in a fixed $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$, where we will require more from (M, \mathcal{X}) in various propositions. Fix $N \in \mathcal{X}$ an inner model (of ZFC^- or ZFC , as appropriate) of M and fix $G \in \mathcal{X}$ a GBC -amenable global well-order of N . We want to calculate the theory of $(N, \mathcal{L}(N, G))$, based upon how strong the theory of (M, \mathcal{X}) is. Below, κ will denote the ordinal in $\mathfrak{L}(N, G)$ which represents the Ord of N , i.e. κ is represented by the membership code $\text{Ord} + 1$.

But first, let us check that there really is such a G .

Lemma 2.53. Let $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$ and let $N \in \mathcal{X}$ be an inner model of M which satisfies Choice. Then there is a class $G \in \mathcal{X}$ which is a GBC -amenable global well-order of N .

Proof. Because \mathcal{X} contains first-order truth predicates relative to any class it has uniform access to the definable subclasses of a definable forcing notion over N . In particular, this works for the Cohen forcing $\text{Add}(\text{Ord}, 1)^N$ to add a subclass of Ord . Let $C \in \mathcal{X}$ meet every definable dense subclass of $\text{Add}(\text{Ord}, 1)^N$. Then C codes a global well-order of N , which is necessarily GBC -amenable to N . \square

Now let us see that $\mathfrak{L}(N, G)$ satisfies some basic axioms.

Proposition 2.54 (Over $\text{GBC}^- + \text{ETR}$). The (N, G) -constructible unrolled model $\mathfrak{L}(N, G)$ satisfies Extensionality, Pairing, Union, Infinity, Foundation, Choice, and that the cardinal $\kappa = \text{Ord}^N$ is regular.

Proof sketch. This is proved similarly to the analogous results from section 2.1 (namely, theorem 2.13 and propositions 2.14, 2.15, and 2.18). The new content is to check that the arguments never take us outside of the constructible membership codes. Note in particular that $\mathfrak{L}(N, G)$ satisfies Choice precisely because G is coded in $\mathfrak{L}(N, G)$. \square

Note that the previous proposition does not claim that $\mathfrak{L}(N, G) \models \text{“}\kappa \text{ is the largest cardinal”}$. Indeed, this will not in general be true. For a counterexample, consider countable $L_\alpha \models \text{ZFC}$ with $\kappa < \alpha$ so that $L_\alpha \models \kappa$ is inaccessible. Now do a class forcing over L_α to collapse all cardinals $> \kappa$. This produces $W \models \text{ZFC}_1^-$ with $\text{Ord}^M = \alpha$ and κ is the largest cardinal in W . Let $(M, \mathcal{X}) \models \text{KMCC}$ be the model obtained by the cutting off construction applied to W . Then we get that $\mathfrak{L}^{(M, \mathcal{X})}$ is (isomorphic to) L_α . So \mathfrak{L} does not think that κ is the largest cardinal. Observe however, that $L^M = L_\kappa$ is still KMCC -realizable, as $V_{\kappa+1}^{L_\alpha}$ is a KMCC -realization for L^M .

Back to calculating the theory of $\mathfrak{L}(N, G)$, if the ground universe satisfies Powerset then $\mathfrak{L}(N, G)$ will think that κ , the cardinal in the unrolling corresponding to the Ord of the ground universe, is inaccessible.

Corollary 2.55 (Over $\text{GBC} + \text{ETR}$). The unrolled model $\mathfrak{L}(N, G)$ satisfies that κ is inaccessible.

Proof. Unrolling with all membership codes, not just the (N, G) -constructible ones, gave that κ is inaccessible. Using fewer membership codes—and thus possibly having fewer sets in the unrolling—can only make it easier for κ to be inaccessible. \square

Next, let us see that $\mathfrak{L}(N, G)$ satisfies the level of Separation corresponding to the level of Comprehension in the ground universe (M, \mathcal{X}) . Similar to the proof of proposition 2.23 we will need to translate first-order formulae φ into second-order formulae about membership codes. Here, however, we want to confine to talking only about constructible membership codes. This translation $\varphi \mapsto \varphi^{*L}$ is given by the following schema.

- If φ is of the form $x = y$ then φ^{*L} is $X \cong Y$.
- If φ is of the form $x \in y$ then φ^{*L} is $X \varepsilon Y$.
- If φ is of the form $\psi \vee \theta$ then φ^{*L} is $\psi^{*L} \vee \theta^{*L}$, and similarly for conjunctions.
- If φ is of the form $\neg\psi$ then φ^{*L} is $\neg\psi^{*L}$.
- If φ is of the form $\exists x\psi(x)$ then φ^{*L} is $\exists X$ X is a constructible membership code and $\psi^{*L}(X)$, and similarly for unbounded universal quantifiers.
- If φ is of the form $\exists x \in y\psi(x)$ then φ^{*L} is $\exists x \multimap_Y t_Y \psi^{*L}(Y \downarrow x)$, and similarly for bounded universal quantifiers.

Like before, this translation does not increase complexity.

Lemma 2.56. *If φ is Σ_k for $k \geq 1$ then φ^{*L} is equivalent to a Σ_k^1 -formula.*

Proof sketch. First, let us see by induction that $(\varphi^L)^*$, using the $*$ -translation from lemma 2.22, is equivalent to φ^{*L} . The only step to check is the unbounded quantifier step, as that is the only nontrivial step in the definition of the relativization $\varphi \mapsto \varphi^L$. Suppose φ is of the form $\exists x \psi(x)$. Then φ^L is of the form $\exists x \in L \psi(x)^L$. By inductive hypothesis, $(\psi(x)^L)^*$ is equivalent to $\psi(x)^{*L}$. But then we are done, since $(x \in L)^*$ is equivalent to “ X is a constructible membership code”.

So, using lemma 2.22, we are done once we know that φ^L is equivalent to a Σ_k formula when φ is Σ_k . This is also seen by induction. The base case, $k = 0$, is trivial because φ^L is just φ for a Σ_0 -formula φ . Now suppose φ is of the form $\exists x \psi(x)$ where ψ is Π_k . By inductive hypothesis ψ^L is equivalent to a Π_k -formula. So φ^L , which is $\exists x \in L \psi(x)^L$ is equivalent to a Σ_{k+1} -formula, because “ $x \in L$ ” is Σ_1 , as desired. \square

Proposition 2.57 (Over $\text{GBC}^- + \text{ETR}$). *If the ground universe (M, \mathcal{X}) satisfies Π_k^1 -CA then $\mathfrak{L}(N, G)$ satisfies Σ_k -Separation, for $k \geq 1$.*

As an immediate corollary we get the following.

Corollary 2.58 (Over $\text{GBC}^- + \text{ETR}$). *If the ground universe (X, \mathcal{X}) satisfies full Comprehension then $\mathfrak{L}(N, G)$ satisfies Separation.* \square

Proof of proposition 2.57. Fix a constructible membership code A and a Σ_k -formula $\varphi(x)$, possibly with (suppressed) parameters. Because φ^{*L} is equivalent to a Σ_k^1 -formula we may by Π_k^1 -Comprehension form the collection of $x \multimap_A t_A$ so that $\varphi^{*L}(A \downarrow x)$. Closing this collection downward in A and adding a top element gives a membership code B representing the subset of A formed using φ . \square

Altogether, we have seen the following.

Theorem 2.59. *Let $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$, $N \in \mathcal{X}$ be an inner model (of ZFC^- or ZFC , as appropriate) of M , and $G \in \mathcal{X}$ a GBC -amenable global well-order of N .*

- If $(M, \mathcal{X}) \models \text{KM}$ then $\mathfrak{L}(N, G) \models \text{wZFC}_1^-$.
- If $(M, \mathcal{X}) \models \text{KM}^-$ then $\mathfrak{L}(N, G) \models \text{wZFC}_R^-$.
- If $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA}$ then $\mathfrak{L}(N, G) \models \text{wZFC}_1^-(k)$.
- If $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_k^1\text{-CA}$ then $\mathfrak{L}(N, G) \models \text{wZFC}_R^-(k)$.

Combined with theorem 2.43 we immediately get the following.

Theorem 2.60. *Let $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$, $N \in \mathcal{X}$ be an inner model (of ZFC^- or ZFC , as appropriate) of M , and $G \in \mathcal{X}$ a GBC -amenable global well-order of N .*

- If $(M, \mathcal{X}) \models \text{KM}$ then $(N, \mathcal{L}(N, G)) \models \text{KM}$.
- If $(M, \mathcal{X}) \models \text{KM}^-$ then $(N, \mathcal{L}(N, G)) \models \text{KM}^-$.
- If $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA}$ then $(N, \mathcal{L}(N, G)) \models \text{GBC} + \Pi_k^1\text{-CA}$.
- If $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_k^1\text{-CA}$ then $(N, \mathcal{L}(N, G)) \models \text{GBC}^- + \Pi_k^1\text{-CA}$.

This is most of theorem 2.47. We have not yet handled the cases of $\text{GBC} + \text{ETR}$ and $\text{GBC}^- + \text{ETR}$, which will happen in the next chapter. It also remains to see that $(N, \mathcal{L}(N, G))$ satisfies a fragment of Class Collection corresponding to the amount of Comprehension satisfied by the ground universe.

Proposition 2.61. *If the ground universe (M, \mathcal{X}) satisfies Π_k^1 -Comprehension, then $\mathfrak{L}(N, G)$ satisfies Σ_k -Collection.*

As an immediate corollary we get the following.

Corollary 2.62 (Over $\text{GBC}^- + \text{ETR}$). *If the ground universe satisfies full Comprehension, then $\mathfrak{L}(N, G)$ satisfies Collection.* \square

Proof of proposition 2.61. The key step in the argument for this proof is encapsulated by the following lemma.

Lemma 2.62.1 (Over $\text{GBC}^- + \Pi_k^1\text{-CA}$). *The (N, G) -constructible unrolled model $\mathfrak{L}(N, G)$ satisfies Σ_k -reflection with respect to the $L(N, G)$ -hierarchy: For any $a \in \mathfrak{L}(N, G)$ and any Σ_k -formula $\varphi(x, p)$ there is an ordinal γ so that $a \in L_\gamma(N, G)^{\mathfrak{L}(N, G)}$ and for all $x \in L_\gamma(N, G)^{\mathfrak{L}(N, G)}$ we have $L_\gamma(N, G)^{\mathfrak{L}(N, G)} \models \varphi(x, a)$ if and only if $\mathfrak{L}(N, G) \models \varphi(x, a)$. Moreover, such γ are unbounded in the ordinals of $\mathfrak{L}(N, G)$.*

Before proving the lemma, let us see how it lets us prove the proposition. This is the familiar argument that reflection implies collection.

Work in $\mathfrak{L}(N, G)$. Fix a and assume that for every $x \in a$ there is y so that $\varphi(x, y, p)$, where p is some parameter. We want to find a b so that for all $x \in a$ there is $y \in b$ so that $\varphi(x, y, p)$. By the lemma, there is γ so that φ reflects to $L_\gamma(N, G) \ni a$. Hence, for every $x \in a$ there is $y \in L_\gamma(N, G)$ so that $\varphi(x, y, p)$. The desired b is then $L_\gamma(N, G)$.

Proof of lemma 2.62.1. I will give a false proof, indicate the error therein, and then explain how to fix the error, giving a proof that works.

The first way one might attempt to prove this is to work in $\mathfrak{L}(N, G)$ and try to find γ so that $L_\gamma(N, G)$ is closed under witnesses to existential subformulae of φ . We start with γ_0 so that $p \in L_{\gamma_0}(N, G)$. Then, define γ_{n+1} to be the least ordinal η so that $L_\eta(N, G)$ is closed under witnesses to existential subformulae of φ with parameters in $L_\gamma(N, G)$. Such γ_{n+1} exists because the ground universe satisfies Π_k^1 -Comprehension and φ^{*L} is Σ_k^1 . We then define γ to be the supremum of the sequence $\langle \gamma_n \rangle$.

The problem with this argument is that it is circular; Collection says we can collect the γ_n , but that is what we are trying to prove! If we look from the perspective of the KM model, we are appealing to Class Collection to pick out (membership codes for) these γ_n . If we could find a way to pick these γ_n without using Class Collection, then the argument would go through. We would get a sequence of meta-ordinals $\langle \Gamma_n \rangle$ coded in \mathcal{X} and from that be able to construct a membership code for $L_{\sup_n \Gamma_n}(N, G)$.

Fix a relation $C \subseteq \text{Ord}^2$ which is isomorphic to \in^N . Call a membership code *compliant* if $E \upharpoonright \text{Ord} = C$. Let Γ be a meta-ordinal. I claim that in (M, \mathcal{X}) we can uniquely pick an isomorphic copy of Γ . First, note that if $\Gamma \in L_\Delta(N, G)$ then in $L_\Delta(N, G)$ there is a least $\Gamma_\Delta \cong \Gamma$. That is, $\Gamma_\Delta = L_\Delta(N, G) \downarrow g$ for some $g \rightarrow t_{L_\Delta(N, G)}$ where this g is picked according to the L -order for $L_\Delta(N, G)$. The key observation now is that if $L_\Delta(N, G)$ and $L_{\Delta'}(N, G)$ are compliant, then $\Gamma_\Delta = \Gamma_{\Delta'}$. This is because they agree on how they code N but elements of Γ_Δ and $\Gamma_{\Delta'}$ are pairs of elements of the code of N in $L_\Delta(N, G)$ and $L_{\Delta'}(N, G)$, respectively. This yields a way to uniquely pick an isomorphic copy of Γ : pick any Δ so that $\Gamma \in L_\Delta(N, G)$ and then look at Γ_Δ . This is well-defined.

We can use this now to pick the Γ_n for the above argument. Since we can do this for all n , we can code the sequence $\langle \Gamma_n \rangle$ by the class $Z = \{(n, a) : a \in \Gamma_n\}$. From this produce a meta-ordinal Γ of ordertype $\sup_n \Gamma_n$ and build $L_\Gamma(N, G)$. Moving to the first-order model, $L_\Gamma(N, G)$ is a representative of the equivalence class for the desired $L_\gamma(N, G)$.

Finally, note that the moreover is immediate as to ensure that $\gamma > \alpha$ for some ordinal α one can run the argument using $a' = (a, \alpha)$ rather than a . \square

The same argument shows that we can get a fragment of Class Collection starting from ETR, except using the \star -translation of definition 2.26 rather than the $*$ -translation.¹⁵ In fact, the argument is a little easier in this case since we only have to worry about Σ_0 -formulae and if φ is Σ_0 then φ^L is just φ .

Corollary 2.63 (Over $\text{GBC}^- + \text{ETR}$). *The (N, G) -constructible unrolling $\mathfrak{L}(N, G)$ satisfies Elementary Class Collection ECC.* \square

This is the final step needed to prove corollary 2.48, which I restate below for the convenience of the reader who does not want to flip back a few pages.

Theorem. *Let $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$. Then there is a definable hyperclass $\mathcal{Y} \subseteq \mathcal{X}$ so that the following.*

- *If $(M, \mathcal{X}) \models \text{KM}$ then $(M, \mathcal{Y}) \models \text{KMCC}$.*
- *If $(M, \mathcal{X}) \models \text{KM}^-$ then $(M, \mathcal{Y}) \models \text{KMCC}^-$.*
- *If $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA}$ then $(M, \mathcal{Y}) \models \text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$, for $k \geq 1$.*
- *If $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_k^1\text{-CA}$ then $(M, \mathcal{Y}) \models \text{GBC}^- + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$, for $k \geq 1$.*
- *If $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}$ then $(M, \mathcal{Y}) \models \text{GBC} + \text{ETR} + \text{ECC}$.*
- *If $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$ then $(M, \mathcal{Y}) \models \text{GBC}^- + \text{ETR} + \text{ECC}$.*

2.4 The smallest heights of transitive models and β -models

As an application of the results in this chapter I would like to investigate the smallest heights of transitive and β -models of $\text{GBC} + \Pi_k^1\text{-CA}$. This builds upon prior work of Marek and Mostowski.

Theorem 2.64 (Marek and Mostowski [MM75]). *Assume there is a β -model of KM. Let β_ω be the least height of a β -model of KM and let τ_ω be the least height of a transitive model of KM. Then*

- $\tau_\omega < \beta_\omega$; and moreover
- $L_{\beta_\omega} \models \tau_\omega$ is countable.

I will generalize their argument to fragments of KM.

Definition 2.65. Let β_ω be the least height of a β -model of KM and let τ_ω be the least height of a transitive model of KM. For $n \in \omega$ let

- β_n be the least height of a β -model of $\text{GBC} + \Pi_n^1\text{-CA}$; and
- τ_n be the least height of a transitive model of $\text{GBC} + \Pi_n^1\text{-CA}$.

By corollary 2.48 it is equivalent to define β_n as the least height of a β -model of $\text{GBC} + \Pi_n^1\text{-CA} + \Sigma_n^1\text{-CC}$ and τ_n as the least height of a transitive model of $\text{GBC} + \Pi_n^1\text{-CA} + \Sigma_n^1\text{-CC}$.

Definition 2.66. For $\gamma < \delta$ ordinals countable in L say that $\gamma \ll \delta$ if $L_\delta \models \gamma$ is countable.

Theorem 2.67. *Let $m < n \leq \omega$. Assume there is a transitive model of $\text{GBC} + \Pi_n^1\text{-CA}$. Then, $\beta_m \ll \tau_n$.*




Figure 2.5: The least heights of transitive and β -models of fragments of KM. The ordinals in the picture are ordered by \ll .

Theorem 2.68. *Let $1 \leq n \leq \omega$. Assume there is a β model of $\text{GBC} + \Pi_n^1\text{-CA}$. Then, $\tau_n \ll \beta_n$.*

See figure 2.5 for a pictorial representation of these theorems.

Proof of theorem 2.67. Let $m < n \leq \omega$. We need to see that $\beta_m \ll \tau_n$.

Consider $(N, \mathcal{Y}) \models \text{GBC} + \Pi_n^1\text{-CA}$ with $\text{Ord}^N = \tau_n$. Then, by the existence of Σ_m^1 -truth predicates in (N, \mathcal{Y}) we get $C \in \mathcal{Y}$ coding $\bar{\mathcal{X}} \subseteq \mathcal{Y}$ so that $(N, \bar{\mathcal{X}}) \models \text{GBC} + \Pi_m^1\text{-CA}$. Work in (N, \mathcal{Y}) . Apply reflection to this C to get a club of ordinals α with $C_\alpha \subseteq V_\alpha$ coding $\mathcal{X}_\alpha \subseteq \mathcal{P}(V_\alpha)$ so that $(V_\alpha, \mathcal{X}_\alpha) \models \text{GBC} + \Pi_m^1\text{-CA}$. Pick α with uncountable cofinality from this club. Then $N \models (V_\alpha, \mathcal{X}_\alpha)$ is a β -model because any transitive model with uncountable cofinality is a β -model.¹⁶ But N is a transitive model of ZFC so it is correct about well-foundedness, so $(V_\alpha^N, \mathcal{X}_\alpha)$ really is a β -model.

This establishes that $\beta_m < \tau_n$. To further see that $\beta_m \ll \tau_n$ observe that N sees the unrolling \bar{W} of the model of $\text{GBC} + \Pi_m^1\text{-CA} + \Sigma_m^1\text{-CC}$ contained inside $(V_\alpha^N, \mathcal{X}_\alpha)$. But then this \bar{W} has a countable (from the perspective of N) submodel \bar{W} . Cutting off this \bar{W} gives a countable β -model of $\text{GBC} + \Pi_m^1\text{-CA}$. \square

Proof of theorem 2.68. Fix n with $1 \leq n \leq \omega$. We want to see that $\tau_n \ll \beta_n$. Let $H_n(\alpha)$ denote L_ξ where $\xi > \alpha$ is least so that $L_\xi \models \text{ZFC}^-(n)$.¹⁷ And let $\text{Hyp}(\alpha)$ be the least admissible set containing α , i.e. L_ξ where $\xi > \alpha$ is least so that $L_\xi \models \text{KP}$. It is then obvious that $\text{Hyp}(\alpha) \subseteq H_n(\alpha)$. Consider the admissible sets $A = \text{Hyp}(\beta_n)$ and $B = H_n(\beta_n)$. Consider the \mathcal{L}_A -theory¹⁸ T whose axioms consist of

- $\text{ZFC}_1^-(n)$;

¹⁵See subsection 2.1.1 for discussion of why we need the alternate translation.

¹⁶See observation 1.14.

¹⁷Following the naming convention of definition 2.1, $\text{ZFC}^-(n)$ is the theory axiomatized by Extensionality, Pairing, Union, Foundation, Infinity, Choice, Σ_n -Separation, and Σ_n -Collection. So, for instance, $\text{ZFC}^-(0)$ is KP plus Choice plus Infinity.

¹⁸Recall that if A is an admissible set then \mathcal{L}_A is the admissible fragment of $\mathcal{L}_{\text{Ord}, \omega}$ associated with A , i.e. the infinitary language consisting of formulae in A . See [Bar75] for details.

- The infinitary \in -diagram of A ;¹⁹ and
- The assertion that β_n is inaccessible.

Then T is Σ_1 -definable over A via a formula θ . I claim that $B \models T$ and thus T is consistent. First, it is clear that $A \in B$ and $B \models \text{ZFC}^-(n)$. What remains is to see that $B \models \beta_n$ is inaccessible. To see this, let $(L_{\beta_n}, \mathcal{X})$ be a β -model of $\text{GBC} + \Pi_n^1\text{-CA} + \Sigma_n^1\text{-CC}$, whose classes are all constructible. Then $(L_{\beta_n}, \mathcal{X})$ unrolls to a transitive model $L_\xi \models \text{ZFC}_1^-(n)$ with $\beta_n \in L_\xi$ inaccessible and the largest cardinal of L_ξ . By definition of B , we have that $\text{Ord}^B \leq \xi$. So B must agree with L_ξ that β_n is inaccessible and is the largest cardinal. Thus, $B \models T$.²⁰

It looks like we are setting up to apply the Barwise compactness theorem, but we are not yet ready to do so. (Indeed, we will not apply the Barwise compactness theorem directly to T .) First, work in B . By condensation there are countable (i.e. from the perspective of B) ordinals γ, δ so that $j : (L_\gamma, \in, \delta) \rightarrow (A, \in, \beta_n)$ is an elementary embedding. Then, $A \models \gamma$ is countable; this is true in B and the bijection from γ to ω must occur earlier than β_n in the L -hierarchy because $A \models \beta_n$ is inaccessible.

Now let T' be L_γ 's version of T ; formally, $T' = \{\varphi : L_\gamma \models \theta(\varphi)\}$. We can explicitly list the axioms of T' , namely:

- $\text{ZFC}_1^-(n)$;
- The infinitary \in -diagram of L_γ ; and
- The assertion that δ is inaccessible.

Because $A \models \text{Con}(T)$ we get by elementarity that $L_\gamma \models \text{Con}(T')$. Also by elementarity, $L_\gamma = \text{Hyp}(\delta)$.

Finally we apply the Barwise compactness theorem, but within A to the theory T' . This yields $C \in A$ with $C \models T'$ and $C \supseteq_{\text{end}} L_\xi$. So the same is true in V , by absoluteness. Applying the cutting off construction to C we get that L_ξ is $(\text{GBC} + \Pi_n^1\text{-CA} + \Sigma_n^1\text{-CC})$ -realizable. So $\tau_n \leq \delta \ll \beta_n$, as desired. \square

Essentially the same argument works for the version of the theories without Powerset, i.e. KM^- and $\text{GBC}^- + \Pi_k^1\text{-CA}$. Let β_ω^- be the least height of a β -model of KM^- and τ_ω^- be the least height of a transitive model of KM^- . For $k \in \omega$ let β_k^- be the least height of a β -model of $\text{GBC}^- + \Pi_k^1\text{-CA}$ and let τ_k^- be the least height of a transitive model of $\text{GBC}^- + \Pi_k^1\text{-CA}$.

Theorem 2.69. *Assume there is a β -model of KM^- . Let $n < m \leq \omega$. Then the following hold.*

- $\beta_n^- \ll \tau_m^-$; and
- $\tau_n^- \ll \beta_n^-$.

Proof sketch. Just like the proof of theorems 2.67 and 2.68. The difference is that when axiomatizing the theory T use $\text{ZFC}_R^-(n)$ rather than $\text{ZFC}_1^-(n)$. \square

It is quite obvious that $\beta_n^- \ll \tau_m^-$ for any $n, m \leq \omega$. So the above picture could be extended by putting the minus ordinals at the bottom, ordered the same but with superscripts everywhere.

We can ask the same question for other theories. For a second-order set theory T , let $\tau(T)$ be the least height of a transitive model of T and let $\beta(T)$ be the least height of a β -model of T .

Question 2.70. *Do we have $\tau(\text{GBC}) < \beta(\text{GBC})$? Do we have $\tau(\text{GBC} + \text{ETR}) < \beta(\text{GBC} + \text{ETR})$?*

¹⁹That is, the collection of all sentences of the form $\forall x x \in a \Leftrightarrow \bigvee_{b \in a} x = b$. Any structure which satisfies all these sentences will have A as a transitive submodel. That is, if N satisfies all these sentences then $A \subseteq N$ and for every $a \in A$ and $b \in N$ if $N \models b \in a$ then $b \in A$.

²⁰If we think in terms of the $\text{GBC} + \Pi_n^1\text{-CA} + \Sigma_n^1\text{-CC}$ model $(L_{\beta_n}, \mathcal{X})$, then what we have seen is that this structure contains a membership code for $\text{Hyp}(V)$, the smallest admissible 'set' containing its V . See chapter 4 for more on $\text{Hyp}(V)$, where it will play an important role.

Chapter 3

Truth and transfinite recursion

Die alte und berühmte Frage, womit man die Logiker in die Enge zu treiben vermeinte und sie dahin zu bringen suchte, dass sie sich entweder auf einer elenden Dialelle mussten betreffen lassen oder ihre Unwissenheit, mithin die Eitelkeit ihrer ganzen Kunst bekennen sollten, ist diese: *Was ist Wahrheit?*

Immanuel Kant

In this chapter I will explicate the relationship between transfinite recursion and iterated truth predicates. This will be used to separate fragments of ETR and Σ_k^1 -TR. (See definition 3.39 below, and note that Σ_0^1 -Transfinite Recursion is a synonym for Elementary Transfinite Recursion.) The main result of this chapter is the following.

Theorem 3.1. *Consider $(M, \mathcal{X}) \models \text{GBC}$ and $\Gamma \in \mathcal{X}$ a well-order with $\Gamma \geq \omega^\omega$. Fix finite $k \geq 0$. Then, if (M, \mathcal{X}) satisfies the Σ_k^1 -Transfinite Recursion principle for recursions of height $\leq \Gamma \cdot \omega$, there is $\mathcal{Y} \subseteq \mathcal{X}$ coded in \mathcal{X} so that $(M, \mathcal{Y}) \models \text{GBC}$ plus the Σ_k^1 -Transfinite Recursion principle for recursions of height $\leq \Gamma$.*

This theorem gives a separation of fragments of Σ_k^1 -Transfinite Recursion by consistency strength. Combined with the easy facts that Π_{k+1}^1 -Comprehension proves Σ_k^1 -Transfinite Recursion and that Σ_{k+1}^1 -Transfinite Recursion for recursions of finite length proves Π_k^1 -Comprehension, this gives a hierarchy of transfinite recursion principles ranging in strength from GBC to KM. Figure 3.1 gives a visual representation of the hierarchy of the second-order transfinite induction principles, spanning from GBC to KM.

This chapter is organized as follows. First, we consider iterated truth predicates and see that (fragments of) Elementary Transfinite Recursion can be equivalently formulated as asserting the existence of certain iterated truth predicates. Next we look at the connections between iterated truth predicates and the second-order constructible universe. This will show that $\text{GBC} + \text{ETR}$ and $\text{GBC}^- + \text{ETR}$ go down to inner models, completing the proof of theorem 2.47 from chapter 2. We will also see the same for fragments of ETR, though by a different argument. We then settle theorem 3.1 for the case $k = 0$, separating fragments of ETR. We further will see that fragments of ETR can be separated in a strong sense, getting transitive models of ZFC satisfying a strong theory which are $(\text{GBC} + \text{ETR}_\Gamma)$ -realizable but not $(\text{GBC} + \text{ETR}_{\Gamma \cdot \omega})$ -realizable. So the separation here is due to an inherently second-order property of the model, not its (first-order) theory. This is followed by a detour to arithmetic, where we consider an analogous situation for iterated truth predicates over models of arithmetic. We see that the strong separation results for fragments of ETR do not have a counterpart in the arithmetic realm. Finally, we turn to non-elementary recursion and settle the $k \geq 1$ case of theorem 3.1.




Figure 3.1: Transfinite recursion, from GBC to KM. Ordered by consistency strength.

Much of this chapter will make use of the constructions from chapter 2, so the reader is strongly encouraged to read that chapter first.

3.1 Truth and iterated truth

In chapter 1 we discussed truth predicates over models of set theory. I wish to return to this subject for a deeper look. As we will see in this chapter, certain principles of second-order set theory can be characterized in terms of the existence of truth predicates as classes. The best known case of such is the Tarskian truth predicate for first-order formulae, which we saw in chapter 1. But we can also consider truth predicates relative to a class parameter.

Definition 3.2. Let A be a class over a model M of first-order set theory. The *truth predicate for M relative to A* —or, synonymously, the *satisfaction class for M relative to A* —is the class T satisfying the recursive Tarskian definition of truth for formulae in the language $\mathcal{L}_{\in}(A)$, i.e. the language of first-order set theory with a symbol for A . Formally, these recursive requirements are the same as in definition 1.53, with the following addition:

- $(x \in A, a)$ is in T if and only if $a \in A$.

Observation 3.3. Let $(M, \mathcal{X}) \models \text{GBC}^-$. If $T, T' \in \mathcal{X}$ both satisfy the definition of a truth predicate relative to $A \in \mathcal{X}$ then $T = T'$.

Proof. If $T \neq T'$ then they disagree at a minimal stage. But this would contradict the recursive requirement at that stage. \square

As such, we are justified in talking about *the* truth predicate relative to A .¹ I will use $\text{Tr}(A)$ to denote the truth predicate relative to A .

As a first application of this idea, let us see that there is no principal model of $\text{GBC} + \text{ETR}$.

Proposition 3.4. Let (M, \mathcal{X}) be a model of second-order set theory. If for every $A \in \mathcal{X}$ we have $\text{Tr}(A) \in \mathcal{X}$ then (M, \mathcal{X}) is not principal.

Proof. Otherwise, if $\mathcal{X} = \text{Def}(M; P)$ for some $P \in M$ then $\text{Tr}(P)$ is definable from P , contradicting Tarski's theorem on the undefinability of truth. \square

Proposition 3.5. The theory $\text{GBC}^- + \text{ETR}_\omega$ proves that $\text{Tr}(A)$ exists for every class A .

Proof. From A we can define $\text{Tr}(A)$ by means of an elementary transfinite recursion of height ω . See definitions 3.2 and 1.53. \square

Corollary 3.6. No model of any $T \supseteq \text{GBC}^- + \text{ETR}_\omega$ is principal. \square

Proposition 3.5 gives that every model of $\text{GBC}^- + \text{ETR}_\omega$ contains truth predicates. This includes non- ω -models, so it would be helpful to discuss what that means in this context.

The trouble is simple. Consider a ω -nonstandard model M and its externally defined satisfaction relation $S = \{(\varphi, a) : M \models \varphi(a)\}$. This S will only include standard formulae in its domain, so it cannot be amenable to M . So this cannot be the class in \mathcal{X} which (M, \mathcal{X}) thinks is the truth predicate if \mathcal{X} is to be a $(\text{GBC}^- + \text{ETR}_\omega)$ -realization for M . Instead, the class in \mathcal{X} which (M, \mathcal{X}) thinks is the truth predicate must measure the ‘truth’ of all formulae in M , including the nonstandard formulae. In the literature on nonstandard models, such a class is also known as a *full satisfaction class*.

Let us step away for a moment from the context of second-order set theory. Classical results of Krajewski [Kra76] show that full satisfaction classes can be non-unique, in contrast to the uniqueness of truth predicates in the second-order context. If M is a countable ω -nonstandard model of, say, ZFC which admits a full

¹The reader who is familiar with satisfaction classes over non- ω -models may want to object here. She should hold her objection. We will discuss that case shortly.

satisfaction class then it admits continuum many different full satisfaction classes. But the disagreement must be confined to the nonstandard realm. An easy induction shows that any full satisfaction class for M agrees with the truth predicate of M (as seen externally from V) for standard formulae.

This implies that for ω -nonstandard models M of set theory we can have different GBC^- -realizations for M which contain different truth predicates. On the other hand, if M is ω -standard and two GBC^- -realizations for M each contain a truth predicate then their truth predicates are the same.

Let us return now to truth predicates relative to a class parameter. Suppose for the moment that we have a model of, say, GBC satisfying $\forall A \text{Tr}(A)$ exists. In particular this holds when A itself is a truth predicate. So we can talk about truth about truth, truth about truth about truth, and so on. It is easy to see that this gives us n th order truth for every standard n .

But we are set theorists here. Our inclination is to extend things to the transfinite. We would like to do that for truth about truth about... Formally, this is captured by the following definition of an iterated truth predicate. Informally, an iterated truth predicate is a class of triples $(\gamma, \varphi, \bar{a})$ so that $\varphi(\bar{a})$ is true at level γ , where φ is allowed to make reference to level $< \gamma$.

Definition 3.7. Let (M, \mathcal{X}) be a model of second-order set theory with $\Gamma \in \mathcal{X}$ a well-order. An *iterated truth predicate of length Γ* (or, synonymously, a Γ -*iterated truth predicate*²) for M is a class T of triples $(\gamma, \varphi, \bar{a})$ with $\gamma \in \text{dom } \Gamma$ satisfying the following recursive requirements. Here, φ is in $\mathcal{L}_{\in}(\text{T})$, the language of set theory augmented with a trinary predicate T , and \bar{a} is a valuation for φ .

1. $(\gamma, x = y, \bar{a})$ is in T if and only if $a_x = a_y$.³
2. $(\gamma, x \in y, \bar{a})$ is in T if and only if $a_x \in a_y$.
3. $(\gamma, \text{T}(x, y, z), \bar{a})$ is in T if and only if
 - $a_x <_{\Gamma} \gamma$;
 - a_y is an $\mathcal{L}_{\in}(\text{T})$ -formula;
 - a_z is a valuation for a_y ; and
 - (a_x, a_y, a_z) is in T .
4. $(\gamma, \varphi \vee \psi, \bar{a})$ is in T if and only if $(\gamma, \varphi, \bar{a})$ or (γ, ψ, \bar{a}) are in T .
5. $(\gamma, \neg\varphi, \bar{a})$ is in T if and only if $(\gamma, \varphi, \bar{a})$ is not in T .
6. $(\gamma, \exists x \varphi(x), \bar{a})$ is in T if and only if there is $b \in M$ so that $(\gamma, \varphi, b \hat{\wedge} \bar{a})$ is in T .⁴

We can also have iterated truth predicates relative to a class parameter A . This has the same definition, except that the formulae must be in the language $\mathcal{L}_{\in}(\text{T}, A)$ and the following additional criterion must be satisfied:

- $(\gamma, x \in A, a)$ is in T if and only if $a \in A$.

²The reader may now see one advantage of using “truth predicate” in place of “full satisfaction class”, even in the ω -nonstandard case. The latter would then lead to talk of iterated full satisfaction classes, which is a mouthful—as can be attested by anyone who has given a talk about iterated truth predicates where he made the mistake of using the wrong terminology.

³To clarify the notation: if \bar{a} is a valuation with variable x in its domain, then a_x is the value assigned to x .

⁴Here we of course have the implicit requirement that x be free in φ .

To clarify the situation with the valuations: By $b \hat{\wedge} \bar{a}$ I mean the valuation which modifies \bar{a} by assigning x the value b . Note that it could be that \bar{a} already assigns x a value, as there is no requirement that our valuations only assign values to free variables which appear in the formula. (Indeed, clause (4) of this definition will imply that this always happens; consider e.g. the formula φ given by $x = x \vee y = y$. Then $(\gamma, \varphi, \langle a_x, a_y \rangle) \in T$ if and only if $(\gamma, x = x, \langle a_x, a_y \rangle) \in T$ or $(\gamma, y = y, \langle a_x, a_y \rangle) \in T$.) So if \bar{a} does not assign x a value then $b \hat{\wedge} \bar{a}$ extends \bar{a} by assigning b to x . Otherwise, if \bar{a} does assign x a value, then we get $b \hat{\wedge} \bar{a}$ by dropping that assignment from \bar{a} and then adding in the assignment of b to x .

In the sequel I will avoid repeating this footnote, but the reader should keep these issues in mind.

Note that the property of being the Γ -iterated truth predicate relative to A is first-order expressible (in parameters Γ and A). So if (M, \mathcal{X}) and (M, \mathcal{Y}) are GBC_m models which both contain T, Γ, A then they agree on whether T is the Γ -iterated truth predicate relative to A .

The following observation generalizes observation 3.3 and is proved in the same manner. It justifies talk of *the* Γ -iterated truth predicate.

Observation 3.8. *Let $(M, \mathcal{X}) \models \text{GBC}^-$ with $\Gamma \in \mathcal{X}$ a well-order. Suppose $T, T' \in \mathcal{X}$ both satisfy the definition of a Γ -iterated truth predicate. Then $T = T'$. The same fact holds for iterated truth predicates relative to a parameter.* \square

I will use $\text{Tr}_\Gamma(A)$ to refer to the Γ -iterated truth predicate relative to A and Tr_Γ to refer to the parameter-free Γ -iterated truth predicate. Observe that $\text{Tr}_1(A)$ and $\text{Tr}(A)$, the ordinary Tarskian truth predicate relative to A , are inter-definable.

Note that if Γ is (standard) finite, then the existence of $\text{Tr}_\Gamma(A)$ is equivalent to the existence of certain ordinary truth predicates. That is, for standard finite n we have $\text{Tr}_n(A)$ exists if and only if $T_1 = \text{Tr}(A)$, $T_2 = \text{Tr}(T_1)$, \dots , $T_n = \text{Tr}(T_{n-1})$ all exist. So the main interest in iterated truth is when the length is transfinite or nonstandard finite.

Finally, let me address the skeptic. She may worry there is danger in allowing the iterated truth predicate to measure the truth of statements that themselves make reference to the iterated truth predicate, thus allowing the liar's paradox to slip in. This worry is unfounded. Clause (3) in the definition of an iterated truth predicate legislates that it be ramified, with truth at level γ unable to make reference to truth at level $\geq \gamma$. So any vicious circles are avoided.

We are now ready to begin to see why the word “recursion” is in the title of this chapter. Namely, the existence of iterated truth predicates gives an equivalent characterization of Elementary Transfinite Recursion.

Theorem 3.9 (Fujimoto [Fuj12]). *The following are equivalent over GBC .*

1. *The principle of Elementary Transfinite Recursion; and*
2. *For all class well-orders Γ and all classes A the class $\text{Tr}_\Gamma(A)$ exists.*

Proof. (1 \Rightarrow 2) Note that $\text{Tr}_\Gamma(A)$ is definable via an elementary transfinite recursion of rank $\omega \cdot \Gamma$. To define truth at level $g \in \text{dom } \Gamma$ requires to have first defined truth at all levels $< g$ and each level is defined via a recursion of rank ω .

(2 \Rightarrow 1) Take an instance of ETR, iterating $\varphi(x, i, F, A)$ along a well-order Γ . I claim that from $\text{Tr}_\Gamma(A)$ can be defined a solution F to this recursion. Specifically, let $F = \{(i, x) : (i, \psi, x) \in \text{Tr}_\Gamma(A)\}$, where ψ is a formula so that $(M, \in, P, \text{Tr}_{\Gamma \upharpoonright i}(A)) \models \psi(x, i)$ if and only if $(M, \in, P, F \upharpoonright i) \models \varphi(x, i)$. (For the latter, F is defined via ψ as above.) Such ψ exists by an application of the Gödel fixed-point lemma. It then follows that F satisfies φ at each stage and is therefore a solution to the full recursion. \square

This result can be refined to give equivalences for fragments of ETR.

Corollary 3.10. *Let $\Gamma \geq \omega^\omega$ be a well-order. Over GBC the following are equivalent.*

1. *The principle of Elementary Transfinite Recursion for recursions of rank $\leq \Gamma$; and*
2. *For all classes A the class $\text{Tr}_\Gamma(A)$ exists.*

Proof. (1 \Rightarrow 2) Again, $\text{Tr}_\Gamma(A)$ is definable by a recursion of rank $\omega \cdot \Gamma$. While it may be that $\Gamma < \omega \cdot \Gamma$, because $\omega^\omega \leq \Gamma$ it must be that $\omega \cdot \Gamma < \Gamma + \Gamma$. So we can carry out this recursion as we can get solutions to recursions of rank $\Gamma + \Gamma$ by first getting a solution to the first Γ many stages, then doing a second recursion to get the rest.

(2 \Rightarrow 1) The same argument goes through as before. \square

This also implies that $\text{GBC} + \text{ETR}$ and $\text{GBC} + \text{ETR}_\Gamma$ are finitely axiomatizable.

Corollary 3.11. *The theories $\text{GBC} + \text{ETR}$ and $\text{GBC} + \text{ETR}_\Gamma$ are finitely axiomatizable, for $\Gamma \geq \omega^\omega$.*

Proof. Recall the folklore result that GBC is finitely axiomatizable. Hence the only work is to see that Elementary Transfinite Recursion (or Elementary Transfinite Recursion restricted to Γ) is finitely axiomatizable. But as Fujimoto's theorem shows, Elementary Transfinite Recursion is equivalent to $\forall A \forall \Gamma \text{Tr}_\Gamma(A)$ exists. Similarly, Elementary Transfinite Recursion restricted to Γ is equivalent to $\forall A \text{Tr}_\Gamma(A)$ exists. \square

3.2 Iterated truth and constructibility

The purpose of this section is to explicate the connection between iterated truth predicates and constructibility in the second-order realm.

As a warm-up, let us first see that $\text{Tr}(A)$ gives a code for $\text{Def}(M; A)$.

Proposition 3.12. *Suppose $(M, \mathcal{X}) \models \text{GBC}^-$ is an ω -model with $A \in \mathcal{X}$. Then if $\text{Tr}(A)$ is in \mathcal{X} we get that $\text{Def}(M; A)$ is coded in \mathcal{X} . In particular, if \mathcal{X} contains a truth predicate then it contains a coded V -submodel.*

Proof. Because M is ω -standard $\text{Tr}(A) = \{(\varphi, \bar{a}) : (M, A) \models \varphi(\bar{a})\}$. The following class, which is first-order definable from $\text{Tr}(A)$, is a code for $\text{Def}(M; A)$:

$$\{((\varphi, \bar{a}), b) : (\varphi, \bar{a} \wedge b) \in \text{Tr}(A) \text{ and } \varphi(\bar{a}) \text{ has a single free variable}\}. \quad \square$$

For possibly ω -nonstandard models we can do the same coding to get access to what (M, \mathcal{X}) thinks are the definable classes. I will use $\text{Def}^{(M, \mathcal{X})}(M; A)$, for $A \in \mathcal{X}$, to refer to the hyperclass coded by

$$\{((\varphi, \bar{a}), b) : (\varphi, \bar{a} \wedge b) \in \text{Tr}(A) \text{ and } \varphi(\bar{a}) \text{ has a single free variable}\}.$$

Of course, this can only be done if $(M, \mathcal{X}) \models \text{Tr}(A)$ exists. If M is ω -standard then $\text{Def}^{(M, \mathcal{X})}(M; A) = \text{Def}(M; A)$. But if M is ω -nonstandard then we get that $\text{Def}^{(M, \mathcal{X})}(M; A) \supsetneq \text{Def}(M; A)$.

Indeed, over GBC , the existence of coded V -submodels is equivalent to the existence of truth predicates.

Proposition 3.13. *Let $(M, \mathcal{X}) \models \text{GBC}$. Then the following are equivalent.*

1. *For all $A \in \mathcal{X}$ we have $\text{Tr}(A) \in \mathcal{X}$; and*
2. *For all $A \in \mathcal{X}$ there is a coded V -submodel $(M, \mathcal{Y}) \models \text{GBC}$ of (M, \mathcal{X}) with $A \in \mathcal{Y}$.*

Proof. (1 \Rightarrow 2) We saw earlier that we have a coded V -submodel of GBC . Specifically, the model $(M, \text{Def}^{(M, \mathcal{X})}(M; A))$ is coded in \mathcal{X} . We want to see that we can get Global Choice. Because $\text{Def}^{(M, \mathcal{X})}(M; A)$ is coded in \mathcal{X} we have in \mathcal{X} uniform access to the dense subclasses of $\text{Add}(\text{Ord}, 1)$ which appear in $\text{Def}(M; A)$. (We also get dense subclasses which (M, \mathcal{X}) thinks are defined by nonstandard formulae, but the important thing is that we get all the dense subclasses defined by standard formulae.) So inside (M, \mathcal{X}) we can line them up in ordertype Ord to find $C \in \mathcal{X}$ which meets each of them. Then $(M, \text{Def}^{(M, \mathcal{X})}(M; A, C)) \models \text{GBC}$ is coded in \mathcal{X} .

(2 \Rightarrow 1) Fix A and let C be a code for a GBC -realization \mathcal{Y} for M which contains A . Then, by reflection, there are club many ordinals α so that $(V_\alpha^M, \bar{A}, \bar{C}) \prec_{\Sigma_1} (M, A, C)$, where $\bar{A} = A \cap V_\alpha^M$ and $\bar{C} = C \cap V_\alpha^M$. I claim that $(V_\alpha^M, \bar{A}) \prec (M, A)$. We can see this by the Tarski–Vaught test. The setup is we assume by induction that we know Σ_n -elementarity and we would like to show Σ_{n+1} -elementarity. We fix a Σ_{n+1} -formula $\exists x \varphi(x, p, \bar{A})$, where \bar{A} is a predicate symbol for A or \bar{A} as appropriate. Assume there is $x \in M$ so that $\varphi(x, p, \bar{A})$ holds for $p \in V_\alpha^M$. We would like to find such an x' in V_α^M . But $\{x \in M : M \models \varphi(x, p, \bar{A})\}$ is $(C)_y$ for some y . So by Σ_1 -elementarity $W = \{x \in V_\alpha^M : V_\alpha^M \models \varphi(x, p, \bar{A})\}$ is $(\bar{C})_y$ for some y . And by Σ_1 -elementarity it is nonempty. So pick $x' \in W$ and we are done.

This then gives us $\text{Tr}(A)$ as a class. Namely, to decide whether $(\varphi, \bar{a}) \in \text{Tr}(A)$ look at some large enough α from this club so that V_α^M can see all the parameters. Then say that $(\varphi, \bar{a}) \in \text{Tr}(A)$ if and only if $(V_\alpha^M, \bar{A}) \models \varphi(\bar{a})$. Because (what M thinks is) the satisfaction relation for (V_α, \bar{A}) satisfies the

recursive Tarskian requirements, so will $\text{Tr}(A)$. This definition is first-order in the parameters A and C , so by Elementary Comprehension $\text{Tr}(A) \in \mathcal{X}$. \square

This lays bare the basic idea behind this section: iterating the Def operation is essentially the same thing as iterating a truth predicate.

Let us now move to the main result of this section.

Theorem 3.14. *The following are equivalent over GBC^- .*

1. *The principle of elementary transfinite recursion.*
2. *For any class A and any well-order Γ the iterated truth predicate $\text{Tr}_\Gamma(A)$ exists.*
3. *For any class A and any well-order Γ the membership code $L_\Gamma(A)$ exists.⁵*

Proof. We have already seen $(1 \Leftrightarrow 2)$. The new content is $(1 \Rightarrow 3)$ and $(3 \Rightarrow 2)$.

$(1 \Rightarrow 3)$ We saw this back in chapter 2. Briefly: fix A and Γ . Then $L_\Gamma(A)$ is constructed via an elementary recursion of height $\omega \cdot \Gamma$. It takes ω many steps to construct (a membership code for) $\text{Def}(X)$ from X and this must be done Γ many times to get $L_\Gamma(A)$.

$(3 \Rightarrow 2)$ This essentially comes down to the fact that $L_{\gamma+2}$ contains the truth predicate for L_γ . Fix a class A . Work in the (M, A) -constructible unrolling $\mathfrak{L}(M, A)$. Inductively see that $\text{Tr}_\gamma^{(M, A)}$ —the γ -iterated truth predicate for the structure (M, \in^M, A) —appears in $L_{\gamma+\gamma}(M, A)$; if $\text{Tr}_\delta^{(M, A)}$ is in $L_\beta(M, A)$ then $\text{Tr}(\text{Tr}_\delta^{(M, A)})$ is in $L_{\beta+2}(M, A)$, from which we can define $\text{Tr}_{\delta+1}^{(M, A)}$. Now given $\Gamma \in \mathcal{X}$ let $\gamma \in \mathfrak{L}(M, A)$ be isomorphic to Γ . Then we can transform $\text{Tr}_\gamma^{(M, A)}$ into $\text{Tr}_\Gamma(A) \in \mathcal{X}$, as desired. \square

Corollary 3.15. *If $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}$ and $N \in \mathcal{X}$ is an inner model of M of ZFC then N is $(\text{GBC}^- + \text{ETR})$ -realizable. In particular, if $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}$ then N is $(\text{GBC} + \text{ETR})$ -realizable.*

Proof. Fix $G \in \mathcal{X}$ a GBC^- -amenable global well-order of N . Such exists by lemma 2.53. By the theorem, $L_\Gamma(N, G)$ exists in \mathcal{X} for all $\Gamma \in \mathcal{X}$. So we can work in the (N, G) -constructible unrolling $\mathfrak{L}(N, G)$, i.e. the structure consisting of all membership codes E —quotiented out by isomorphism— E so that $E \varepsilon L_\Gamma(N, G)$ for some Γ .⁶ Now let $\mathcal{Y} = \mathcal{L}(N, G)$. That is, \mathcal{Y} consists of the classes of N which are coded by membership codes in $\mathfrak{L}(N, G)$. It is immediate that $(N, \mathcal{Y}) \models \text{GBC}^-$. We want to see that it also satisfies ETR. Fix $A, \Gamma \in \mathcal{Y}$ where $(N, \mathcal{Y}) \models \Gamma$ is a well-order. Then, it must be that (M, \mathcal{X}) agrees that Γ is a well-order. Otherwise, there is an ordinal $\alpha \in M$ so that $M \models \Gamma \upharpoonright F''\alpha$ is ill-founded, where $F \in \mathcal{Y}$ is a bijection between Ord and $\text{dom } \Gamma$, which exists by Global Choice. But by Replacement $\Gamma \upharpoonright F''\alpha$ must be in N . And N is a transitive submodel of M and they are both models of ZFC^- , so they must agree on what sets are well-founded. So $N \models \Gamma \upharpoonright F''\alpha$ is ill-founded, so $(N, \mathcal{Y}) \models \Gamma$ is ill-founded, a contradiction.

Then there is some $\Delta \in \mathcal{X}$ so that A, Γ are coded in $L_\Delta(N, G)$. More formally, E_A and E_Γ , the canonical membership codes for A and Γ , are ε -elements of $L_\Delta(N, G)$. But then $L_\Gamma(A) \varepsilon L_{\Delta+\Gamma+1}(N, G)$. So $L_\Gamma(A) \in \mathcal{Y}$. Since this worked for arbitrary A and Γ , we get by the theorem that $(N, \mathcal{Y}) \models \text{ETR}$, completing the argument. \square

Let us now turn to ETR_Γ . We get that $\text{GBC} + \text{ETR}_\Gamma$ goes down to inner models, but we need a different argument. The trouble is that without satisfying full ETR it is not clear that the unrolling process gives a model of a sensible theory, so how are we to build second-order L ? (Cf. subsection 2.1.2.)

Theorem 3.16. *Let $(M, \mathcal{X}) \models \text{GBC}^-$ and let $N \in \mathcal{X}$ be an inner model of M . Suppose $\Gamma \in \mathcal{X}$ is a well-order $\geq \omega^\omega$ and is a GBC -amenable subclass of N so that $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}_\Gamma$. Then N is $(\text{GBC}^- + \text{ETR}_\Gamma)$ -realizable, via some $\mathcal{Y} \subseteq \mathcal{X}$.*

⁵See chapter 2 for a definition of $L_\Gamma(A)$.

⁶See chapter 2 for a definition of ε the membership relation between membership codes.

Recall that “ $(N, \mathcal{Y}) \models \text{ETR}_\Gamma$ ” only makes sense when $\Gamma \in \mathcal{Y}$ is a well-order. (You can express it as a theory in first-order logic by using a parameter for Γ .) So in order to have N be $(\text{GBC} + \text{ETR}_\Gamma)$ -realizable it must be that Γ can be put into the classes for N . The condition that Γ be GBC^- -amenable for N ensures this can happen. Otherwise, you can run into pathologies. For example, if $N = L^M$ and Γ codes 0^\sharp , then Γ is not GBC -amenable to N .

And of course, if M is a model of Powerset then N must also be a model of Powerset, so this shows that being $(\text{GBC} + \text{ETR}_\Gamma)$ -realizable goes down to inner models.

Proof. Fix $G \in \mathcal{X}$ a bijection $\text{Ord}^M \rightarrow N$ which is GBC^- -amenable to N . (See lemma 2.53.) We define \mathcal{Y} an ETR_Γ -realization for N as a certain subset of \mathcal{X} , built as a union of an ω -chain of GBC^- -realizations for N . Start with $\mathcal{Y}_0 = \text{Def}^{(M, \mathcal{X})}(N; G, \Gamma)$. Then \mathcal{Y}_0 is coded in \mathcal{X} because \mathcal{X} has a truth predicate for N relative to G and Γ . Also, note that $(N, \mathcal{Y}_0) \models \Gamma$ is a well-order; otherwise, (M, \mathcal{X}) would also see the witness that Γ is ill-founded, contradicting that $(M, \mathcal{X}) \models \Gamma$ is a well-order. Observe that $(N, \mathcal{Y}_0) \models \text{GBC}^-$. Now, given \mathcal{Y}_n let \mathcal{Y}_{n+1} be the smallest extension of \mathcal{Y}_n which contains all Γ -iterated truth predicates relative to parameters from \mathcal{Y}_n . Formally,

$$\mathcal{Y}_{n+1} = \bigcup \left\{ \text{Def}^{(M, \mathcal{X})} \left(N; (\text{Tr}_\Gamma(A)^N)^{(M, \mathcal{X})} \right) : A \in \mathcal{Y}_n \right\}.$$

Some remarks are in order. First, because \mathcal{Y}_n is coded in \mathcal{X} , so is \mathcal{Y}_{n+1} . Second, it must be addressed what this iterated truth predicate is. By way of a transfinite induction of height $\omega \cdot \Gamma$, our model (M, \mathcal{X}) can build what it thinks is the Γ -iterated truth predicate for N , relative to a parameter. This is $(\text{Tr}_\Gamma(A)^N)^{(M, \mathcal{X})}$.

Finally, set $\mathcal{Y} = \bigcup_{n \in \omega} \mathcal{Y}_n$. We know that $(N, \mathcal{Y}) \models \text{GBC}^-$, because \mathcal{Y} is the union of an increasing chain of GBC^- -realizations for N . We now want to see that $(N, \mathcal{Y}) \models \text{ETR}_\Gamma$. Fix $A \in \mathcal{Y}$. Then $A \in \mathcal{Y}_n$ for some n . Thus

$$\text{Tr}_\Gamma(A)^{(N, \mathcal{Y})} = (\text{Tr}_\Gamma(A)^N)^{(M, \mathcal{X})} \in \mathcal{Y}_{n+1} \subseteq \mathcal{Y}.$$

So (N, \mathcal{Y}) contains $\text{Tr}_\Gamma(A)$ for all $A \in \mathcal{Y}$, as desired. \square

A similar strategy can be used to show that $\text{GBC} + \text{ETR}$ is closed under inner models, rather than going through second-order L . For full ETR , however, to build \mathcal{Y}_{n+1} we want to include iterated truth predicates of all lengths in \mathcal{Y}_n , not just those of length Γ . We will also see a version of this construction reappear in chapter 4.

3.3 Separating levels of ETR

In this section we will see that the levels of ETR form a hierarchy in consistency strength.

Let us begin with a lemma that $\text{GBC} + \text{ETR}_\Gamma$ proves well-order comparability for Γ . That is, if Δ is any well-order then exactly one of the following holds: $\Delta < \Gamma$, $\Delta = \Gamma$, or $\Delta > \Gamma$.⁷ This is a refinement of the fact that ETR proves that any two class well-orders are comparable.

Lemma 3.17. *The theory $\text{GBC} + \text{ETR}_\Gamma$ proves that Γ is comparable to any class well-order. That is, if Δ is a class well-order then either there is an embedding of Γ onto an initial segment of Δ or else an embedding of Δ onto an initial segment of Γ .*

Proof. Fix Δ . Consider the transfinite recursion which attempts to construct an embedding of Γ onto an initial segment of Δ . That is, this recursion builds such an embedding π according to the following rule: for $g \in \text{dom } \Gamma$ set $\pi(g)$ to be the least element of $\text{dom } \Delta \setminus \text{ran}(\pi \upharpoonright (\Gamma \upharpoonright <_\Gamma g))$, if such exists, otherwise $\pi(g)$ is undefined. There are two cases. If $\pi(g)$ is always defined then we have embedded Γ onto an initial segment of Δ . If $\pi(g)$ is ever undefined at a stage then it will be undefined at every subsequent stage. So we get that π^{-1} embeds Δ onto an initial segment of Γ . \square

It is clear that if $\Delta < \Gamma$ then ETR_Γ implies ETR_Δ . But if Δ and Γ are sufficiently close then they are in fact equivalent. For instance, ETR_Δ is equivalent (over GBC) to $\text{ETR}_{\Delta+\Delta}$ because to carry out a recursion

⁷If Δ and Γ are class well-orders then $\Delta \leq \Gamma$ if there is an embedding of Δ onto an initial segment of Γ .

of height $\Delta + \Delta$ one first carries out a recursion of height Δ , then using the solution of such as a parameter carries out a second recursion of height Δ . In general, ETR_Δ is equivalent to $\text{ETR}_{\Delta \cdot n}$ for any standard $n > 0$.

This does not generalize from n to ω .

Theorem 3.18. *Suppose $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}_\Gamma$ for $\Gamma \in \mathcal{X}$ a class well-order $\geq \omega^\omega$ and let $\Delta \in \mathcal{X}$ be a class well-order such that $\Delta \cdot \omega \leq \Gamma$. Then, there is (M, \mathcal{Y}) a coded V -submodel of (M, \mathcal{X}) so that $(M, \mathcal{Y}) \models \text{GBC}^- + \text{ETR}_\Delta$.*

Proof. Fix $G \in \mathcal{X}$ a bijection $\text{Ord}^M \rightarrow M$. Consider $T_{\Delta \cdot \omega}(G)$ which is in \mathcal{X} because it can be constructed via an elementary recursion of rank $(\omega \cdot \Delta) \cdot \omega \leq \Gamma$. Now let \mathcal{Y} consist of all sets (internally) definable from an initial segment of $T_{\Delta \cdot \omega}(G)$. That is

$$\mathcal{Y} = \{Y : Y \in \text{Def}^{(M, \mathcal{X})}(M; T_\Upsilon(G)) \text{ for some } \Upsilon < \Delta \cdot \omega\}.$$

Then \mathcal{Y} is coded in \mathcal{X} by using $T_{\Delta \cdot \omega}(G)$.

It is immediate that (M, \mathcal{Y}) satisfies Extensionality, Replacement, and Global Choice. It satisfies First-order Comprehension because \mathcal{Y} is an increasing union of GBC^- -realizations, namely the $\text{Def}(M; T_\Upsilon(G))$ for $\Upsilon < \Delta \cdot \omega$. Finally, it satisfies ETR_Δ because if $A \in \mathcal{Y}$ then $A \in \text{Def}(M; T_\Upsilon(G))$ for some $\Upsilon < \Delta \cdot \omega$ and thus $\text{Tr}_\Delta(A) \in \text{Def}(M; T_{\Upsilon+\Delta}(G)) \subseteq \mathcal{Y}$. \square

This establishes theorem 3.1 for the $\text{GBC} + \text{ETR}$ case.

As an immediate corollary we get that levels of ETR can be separated by consistency strength.

Corollary 3.19. *Let $(M, \mathcal{X}) \models \text{GBC}$ and suppose $\Gamma \in \mathcal{X}$ is a well-order $\geq \omega^\omega$. Then if $(M, \mathcal{X}) \models \text{ETR}_{\Gamma \cdot \omega}$ we have that $(M, \mathcal{X}) \models \text{Con}(\text{GBC} + \text{ETR}_\Gamma)$.* \square

There is also a version of this corollary for GBC^- .

To turn this into a statement about theories in \mathcal{L}_\in , i.e. the language of set theory without any names for distinguished well-orders, we need that Γ is definable. Moreover, in order for (M, \mathcal{X}) to agree with its V -submodels as to what Γ is we need that Γ is defined by a first-order formula (without parameters). So we can say that, for instance, $\text{GBC} + \text{ETR} \vdash \text{Con}(\text{GBC} + \text{ETR}_{\text{Ord}})$, where both theories are in \mathcal{L}_\in . See the discussion in section 4.4 for further details.

The proof for theorem 3.18 also separates ETR_Γ and $\text{ETR}_{<\Gamma}$ for Γ closed under addition.

Definition 3.20. Let $(M, \mathcal{X}) \models \text{GBC}^-$ and suppose $\Gamma \in \mathcal{X}$ is a well-order. Then $(M, \mathcal{X}) \models \text{ETR}_{<\Gamma}$ if it satisfies ETR_Δ for all $\Delta < \Gamma$.

Theorem 3.21. *Suppose $(M, \mathcal{X}) \models \text{GBC}^- + \text{ETR}_\Gamma$ for $\Gamma \in \mathcal{X}$ a well-order $\geq \omega^\omega$ so that $\Delta + \Delta < \Gamma$ for all $\Delta < \Gamma$. Then, (M, \mathcal{X}) has a coded V -submodel $(M, \mathcal{Y}) \models \text{ETR}_{<\Gamma}$.*

Proof sketch. Similar to the proof of theorem 3.18, but set

$$\mathcal{Y} = \{Y : Y \in \text{Def}^{(M, \mathcal{X})}(M; T_\Upsilon(G)) \text{ for some } \Upsilon < \Gamma\}$$

where $G \in \mathcal{X}$ is some bijection $\text{Ord}^M \rightarrow M$. Then $(M, \mathcal{Y}) \models \text{ETR}_{<\Gamma}$. \square

Confining one's attention to transitive models (or, more broadly, ω -standard models) this is the end of the story. For any standard n we have that ETR_Γ is equivalent to $\text{ETR}_{\Gamma \cdot n}$ so ETR_Γ is equivalent to $\text{ETR}_{<\Gamma \cdot \omega}$ for models with the correct ω . But if a model has ill-founded ω then there is a gap. Can an intermediate theory be found in this gap?

The answer is yes.

Fix $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}_{\Gamma \cdot \omega}$ an ω -nonstandard model where $\Gamma \in \mathcal{X}$ so that $(M, \mathcal{X}) \models \Gamma$ is well-founded. Given $\mathcal{Y} \subseteq \mathcal{X}$ define the Γ -recursion cut for \mathcal{Y} to be $I_\Gamma(\mathcal{Y}) = \{e \in \omega^M : (M, \mathcal{Y}) \models \text{GBC} + \text{ETR}_{\Gamma \cdot e}\}$. Note that $I_\Gamma(\mathcal{Y})$ must be closed under addition, as being closed under multiplication by standard n is equivalent to being closed under addition. This is the only restriction on what $I_\Gamma(\mathcal{Y})$ can be.

Theorem 3.22. *Let $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}_{\Gamma \cdot \omega}$ be ω -nonstandard where $\Gamma \in \mathcal{X}$ is well-founded according to (M, \mathcal{X}) . Let $I \subseteq \omega^M$ be a cut closed under addition. Then there is $\mathcal{Y} \subseteq \mathcal{X}$ so that $I_{\Gamma}(\mathcal{Y}) = I$.*

Observe that, unlike for theorem 3.18, \mathcal{Y} cannot be coded in \mathcal{X} as then (M, \mathcal{X}) could define I , which is impossible.

Proof. Fix $G \in \mathcal{X}$ a global well-order of M . Set

$$\mathcal{Y} = \bigcup_{e \in I} \text{Def}(M; (\text{Tr}_{\Gamma \cdot e}(G))^{(M, \mathcal{X})}).$$

A comment is in order. Because M is ω -nonstandard it in general can admit multiple incompatible full satisfaction classes. So $(\Gamma \cdot e)$ -iterated full satisfaction classes will not be unique. Nevertheless, (M, \mathcal{X}) will have at most one Δ -iterated truth predicate relative to a given parameter, because iterated truth predicates are unique in a fixed second-order model. So $(\text{Tr}_{\Gamma \cdot e}(G))^{(M, \mathcal{X})}$ is well-defined.

Note also that each $\text{Def}(M; (\text{Tr}_{\Gamma \cdot e}(G))^{(M, \mathcal{X})})$ is a **GBC**-realization for M . Thus, because \mathcal{Y} is the increasing union of these **GBC**-realizations it too must be a **GBC**-realization. It remains only to check that $(M, \mathcal{Y}) \models \text{ETR}_{\Gamma \cdot e}$ if and only if $e \in I$. The backward direction of this implication is immediate from the definition of \mathcal{Y} .

For the forward direction, take $a > I$. Suppose towards a contradiction that $(M, \mathcal{Y}) \models \text{ETR}_{\Gamma \cdot a}$. Then \mathcal{Y} has (what it thinks is) $\text{Tr}_{\Gamma \cdot a}(G)$. But then $\text{Tr}_{\Gamma \cdot a}(G)$ is definable from $\text{Tr}_{\Gamma \cdot e}(G)$ for some $e \in I$. In particular, this means that $\text{Tr}_{\Gamma \cdot e+1}(G)$ is definable from $\text{Tr}_{\Gamma \cdot e}(G)$, contradicting Tarski's theorem on the undefinability of truth. \square

So while the separation between ETR_{Γ} and $\text{ETR}_{\Gamma \cdot \omega}$ is optimal for ω -standard models, for ω -nonstandard models there are always intermediate levels of **ETR**.

Let us return now to transitive models. Earlier when we separated ETR_{Γ} from $\text{ETR}_{\Gamma \cdot \omega}$ we did so via the second-order part of the model. Starting from $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}_{\Gamma \cdot \omega}$ we found $\mathcal{Y} \subseteq \mathcal{X}$ so that (M, \mathcal{Y}) satisfies ETR_{Γ} but not $\text{ETR}_{\Gamma \cdot \omega}$. So the separation is entirely due to which classes we allow in each model.

Can we separate ETR_{Γ} and $\text{ETR}_{\Gamma \cdot \omega}$ via the first-order part of a model? Can we do so with a transitive model? That is, can we find transitive $M \models \text{ZFC}$ which is $(\text{GBC} + \text{ETR}_{\Gamma})$ -realizable but not $(\text{GBC} + \text{ETR}_{\Gamma \cdot \omega})$ -realizable?

Yes we can, if Γ is a set, rather than a proper class.

Theorem 3.23. *Let γ be an ordinal $\geq \omega^\omega$ given by a first-order definition. That is, there is a first-order formula $\varphi(x)$ without parameters so that **ZFC** proves $\varphi(x)$ has a unique witness γ and this witness is an ordinal $\geq \omega^\omega$. Suppose there is a transitive model of $\text{ETR}_{\gamma \cdot \omega}$. Then there is a transitive model of **ZFC** which is $(\text{GBC} + \text{ETR}_{\gamma})$ -realizable but not $(\text{GBC} + \text{ETR}_{\gamma \cdot \omega})$ -realizable.*

Proof. Take $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}_{\gamma \cdot \omega}$ transitive with a definable global well-order. That such (M, \mathcal{X}) exists is a consequence of theorem 3.15. Let $\vec{T} = \langle (\text{Tr}_{\delta})^M : \delta < \gamma^M \cdot \omega \rangle$ be the sequence of δ -iterated truth predicates for M for $\delta < \gamma \cdot \omega$. Then $\vec{T} \in \mathcal{X}$ because \mathcal{X} contains $(\gamma \cdot \omega)$ -iterated truth predicates. Now let $(N, \gamma^N, \vec{S}) \prec (M, \gamma^{(M, \mathcal{X})}, \vec{T})$ be the Skolem hull of the empty set, using the global well-order of M to pick witnesses. Then, $(N, \gamma^N, \vec{S}) \models \vec{S}$ consists of δ -iterated truth predicates for $\delta < \gamma^N$. I claim that N is $(\text{GBC} + \text{ETR}_{\gamma})$ -realizable but not $(\text{GBC} + \text{ETR}_{\gamma \cdot \omega})$ -realizable.

For the former, let $\mathcal{Y} = \text{Def}(N; \gamma^N, \vec{S})$ consist of the subsets of N which are definable (with set parameters) over the structure $(N; \gamma^N, \vec{S})$. It is immediate that (N, \mathcal{Y}) satisfies Extensionality, Global Choice, and First-Order Comprehension. Suppose towards a contradiction that (N, \mathcal{Y}) does not satisfy Replacement. Then there is $F \in \mathcal{Y}$, defined from γ^N and \vec{S} via a formula φ possibly with parameters, and a set $a \in N$ so that $F''a \notin N$. But then by elementary in (M, γ^M, \vec{T}) the formula φ defines a class function G and there is a set $b \in M$ so that $G''b \notin M$. But then (M, \mathcal{X}) fails to satisfy Replacement, contradicting that it is a model of **GBC**. Altogether we get that \mathcal{Y} is a **GBC**-realization for N .

Next let us see that $(N, \mathcal{Y}) \models \text{ETR}_\gamma$. But this is immediate; given any class $A \in \mathcal{Y}$ there is $\delta < \gamma^N$ so that A is definable from Tr_δ so $\text{Tr}_\gamma(A)$ is definable from $\text{Tr}_{\delta+\gamma}$ and hence $\text{Tr}_\gamma(A) \in \mathcal{Y}$.

Finally, N cannot be $(\text{GBC} + \text{ETR}_{\gamma \cdot \omega})$ -realizable because if \mathcal{Z} were an $(\text{GBC} + \text{ETR}_{\gamma \cdot \omega})$ -realization for N then \mathcal{Z} would contain $\text{Tr}_{\gamma \cdot \omega}$ but then \mathcal{Z} would see that N is a Skolem hull, hence countable. Note that this uses that N is an ω -model, so that there is only one subset of N which can satisfy the definition of a $(\gamma^N \cdot \omega)$ -iterated truth predicate, namely the externally constructed one. But no model of GBC thinks its first-order part is countable, so the existence of such \mathcal{Z} is impossible. \square

Observe that we can ensure that $\gamma^N = \gamma^M$ by requiring N to be the Skolem hull of, say, V_γ^M rather than the Skolem hull of the empty set.

This theorem gives a strong separation for sufficiently weak fragments of ETR . The model N we constructed cannot be made into a model of $\text{ETR}_{\gamma \cdot \omega}$ not because it fails to have a compatible (first-order) theory, but rather due to inherently second-order properties of the model. In the next section we will see that this phenomenon depends essentially upon the transfinite; it does not occur for models of finite set theory, equivalently models of arithmetic.

3.4 A detour through the finite realm

While my analysis has mainly been confined to models of set theory, analogous results are possible for models of arithmetic. I wish to take a brief detour from the infinite world to consider the applications of these ideas to the finite world.

Let me recall some standard facts about satisfaction classes for nonstandard models of arithmetic. First, we will need a few definitions.

Definition 3.24. A structure M is *resplendent* if it realizes any consistent Σ_1^1 -formula. That is, if \hat{X} is a new predicate symbol, \bar{a} are elements of M , and $\varphi(\hat{X}, \bar{a})$ is consistent with $\text{Th}(M, \bar{a})$ then there is $X \subseteq M$ so that $(M, X) \models \varphi(\hat{X}, \bar{a})$.

Further say that M is *chronically resplendent* if X may be chosen so that (M, X) is resplendent.

Definition 3.25. A structure M is *recursively saturated* if it realizes any consistent computable type. That is, if $p(x, \bar{a})$ is a consistent type so that the set of formulae $\varphi \in p$ form a computable subset of ω then there is $t \in M$ so that $M \models \varphi(t, \bar{a})$ for all $\varphi \in p$.

I am primarily interested in structures which allow an appreciable amount of coding, such as models of arithmetic or models of set theory. For such structures we can write down a formula $\varphi(\hat{X})$ which asserts that \hat{X} is a truth predicate. Every countable recursively saturated model admits a full satisfaction class—a theorem of Kotlarski, Krajewski, and Lachlan [KSL81].⁸ So because every completion of PA has a countable recursively saturated model we get that resplendent models admit full satisfaction classes. But note that in general these models will not satisfy induction in the expanded language with a predicate for the full satisfaction class, as induction in the expanded language allows one to prove $\text{Con}(\text{PA})$.

Theorem 3.26 (Lachlan [Lac81]). *If $M \models \text{PA}$ admits a full satisfaction class then M is recursively saturated.*

For countable models all these notions are equivalent, but separations can happen in the uncountable. Kaufmann [Kau77] produced recursively saturated, rather classless⁹ models. Combined with a theorem of Smith's [Smi89] that models with a full satisfaction class must have undefinable classes, this gives recursively saturated models which do not admit a full satisfaction class.

Theorem 3.27 (Barwise–Schlipf [BS76]). *If M is countable and recursively saturated then M is chronically resplendent.*

⁸But see [EV15] for a more elegant proof.

⁹ $M \models \text{PA}$ is rather classless if every class of M is definable, where $A \subseteq M$ is a class if $A \cap [0, x]^M$ is definable for every $x \in M$. Cf. definition 1.50.

As an immediate consequence we get that countable resplendent models are chronically resplendent. It is open whether this is true in general.

Question 3.28. *Does resplendency always imply chronic resplendency?*

Let us turn now to iterated full satisfaction classes¹⁰ over models of arithmetic. Essentially the same argument that resplendent models admit full satisfaction classes yields that resplendent models admit iterated full satisfaction classes. But that is not the only way to get models with iterated full satisfaction classes. Another way goes through models of second-order arithmetic.

Recall that ATR_0 is the theory of second-order arithmetic axiomatized by the following: PA for the first-order part; Extensionality for sets; comprehension for arithmetical (i.e. first-order) properties; and arithmetical transfinite recursion, asserting that inductions of arithmetical properties along a well-founded relation have solutions. That is, ATR_0 is the arithmetic counterpart to $\text{GBC} + \text{ETR}$. Many arguments from $\text{GBC} + \text{ETR}$ generalize to ATR_0 and vice versa. One example is the construction of iterated truth predicates. Let me give a definition specialized to this context, for the sake of clarity.

Definition 3.29. Let $M \models \text{PA}$ and $n \in M$. An *iterated full satisfaction class of length n* (or, synonymously, an *n -iterated full satisfaction class* for M is a set $T \subseteq M$ of triples¹¹ (i, φ, \bar{a}) with $i < n$ satisfying the following recursive requirements. Here, φ is in $\mathcal{L}_{\text{PA}}(\text{T})$, the language of arithmetic augmented with a trinary predicate T , and \bar{a} is a valuation for φ .

1. For atomic φ : (i, φ, \bar{a}) is in T if and only if φ given the values from \bar{a} is true.¹²
2. $(i, x \in y, \bar{a})$ is in T if and only if $a_x \in a_y$.
3. $(i, \text{T}(x, y, z), \bar{a})$ is in T if and only if
 - $a_x < i$;
 - a_y is an $\mathcal{L}_{\in}(\text{T})$ -formula;
 - a_z is a valuation for a_y ; and
 - (a_x, a_y, a_z) is in T .
4. $(i, \varphi \vee \psi, \bar{a})$ is in T if and only if (i, φ, \bar{a}) or (i, ψ, \bar{a}) are in T .
5. $(i, \neg\varphi, \bar{a})$ is in T if and only if (i, φ, \bar{a}) is not in T .
6. $(i, \exists x \varphi(x), \bar{a})$ is in T if and only if there is $b \in M$ so that $(i, \varphi, b \hat{\ } \bar{a})$ is in T .¹³

We can also consider iterated full satisfaction classes of length ω (in the sense of the model) by allowing i to be any element of the model, not just those $< n$.

In this section I will only consider iterated full satisfaction classes of length $\leq \omega$. One could consider longer lengths, but the general theory requires additional care. In particular, “ Γ is a well-order” is a second-order assertion in the arithmetical case, so there are subtleties in formulating $(M, \Gamma, S) \models$ “ S is an iterated full satisfaction class of length Γ ” for $\Gamma, S \subseteq M$.

¹⁰In this section and this section alone I will talk about iterated full satisfaction classes instead of iterated truth predicates. We are not working in a second-order context so the uniqueness of iterated truth predicates within a second-order model does not apply here. I wish to emphasize this change of perspective with a change of language. The exception to this choice is when I talk about standard models, in which case the only possible choice of a full satisfaction class is the truth predicate for the model, as seen externally.

¹¹Recall that PA allows for the coding of sequences as single numbers, so we can think of T as a subset of M , rather than a subset of M^3 .

¹²For example, $2 + 3 < 5 + 1$ is declared true at every level $< n$, because $2 + 3$ really is less than $5 + 1$.

¹³Here we of course have the implicit requirement that x be free in φ .

Observe that each step in the recursive requirements above is first-order. So ATR_0 proves the existence of iterated full satisfaction classes of all lengths $\leq \omega$. Moreover, these will be inductive¹⁴ iterated full satisfaction classes, because any set in a model of ATR_0 must be inductive.

Observe that not every resplendent $M \models \text{PA}$ admits inductive iterated full satisfaction classes. This is a consequence of the following observation.

Observation 3.30. *If $M \models \text{PA}$ admits an inductive iterated full satisfaction class of length n , possibly nonstandard, then $M \models \text{Con}^n(\text{PA})$. So if M admits an inductive iterated full satisfaction class of length ω (in the sense of M), then $M \models \forall n \text{Con}^n(\text{PA})$.*

Proof. An inductive iterated full satisfaction class of length $k + 1$ allows one to get a completion of $\text{PA} + \text{Con}^k(\text{PA})$. But if the model can see a completion of T then it must think $\text{Con}(T)$. \square

In the previous section we considered transitive models of set theory which admitted GBC-amenable iterated truth predicates of length η but not ones of longer lengths. Moreover, we could find such models satisfying any extension T of ZFC which has a transitive model and is consistent with the existence of iterated truth predicates of length η . See theorem 3.23. In particular, T could be the \mathcal{L}_\in -reduct of $\text{GBC} +$ there is an iterated truth predicate of length ζ , for ζ larger than η . Can we have the same phenomenon for models of arithmetic? That is, for any $T \supseteq \text{PA}$ can we find a model of T which admits an inductive full satisfaction class of length η but not an inductive iterated full satisfaction class of length $\eta + 1$?

For the standard model, there is only one choice for a full satisfaction class, namely the truth predicate for the model. But this will not work because \mathbb{N} admits iterated truth predicates of any length and any class over \mathbb{N} is inductive. So we must look at nonstandard models.

Observation 3.31. *Let $M \models \text{PA}$ be countable and admit an inductive full satisfaction class. Then M admits an inductive full satisfaction class of any length compatible with the theory of M . That is, if “ S is an inductive iterated full satisfaction class of length n ” is consistent with $\text{Th}(M)$ then there is $S \subseteq M$ which is an inductive n -iterated full satisfaction class for M . (And similarly for length ω .)*

Proof. Because M admits a full satisfaction class it must be recursively saturated. But since M is countable and recursively saturated it is resplendent. So if “ S is an inductive full satisfaction class of length n ” is consistent with $\text{Th}(M)$ then there is $S \subseteq M$ realizing that theory. (And similarly for length ω .) \square

This observation tells us that for resplendent models the only way to rule out having an inductive iterated full satisfaction class of length n is by the theory of the model. This is unlike the case for transitive models of set theory, where we can find a model of any reasonable T so that the model does not admit a GBC-amenable iterated truth predicate of length n .

If we want to find an inductive iterated full satisfaction class over $M \models \text{PA}$ which cannot be extended to a longer one then there are two ways to do such. One would be if (M, T) is not recursively saturated for T an inductive iterated full satisfaction class. Then (M, T) does not admit a full satisfaction class, inductive or otherwise. But this is a boring case. The more interesting case is when (M, T) is recursively saturated (which, for countable M , is equivalent to being resplendent). In this case the only potential obstacle is $\text{Th}(M, T)$.

As a warm-up let us see that different full satisfaction classes can give different consequences in the extended model.

Proposition 3.32. *There is $M \models \text{PA}$ with $U, I \subseteq M$ inductive full satisfaction classes so that $\text{Th}(M, U) \neq \text{Th}(M, I)$.*

I learned of this argument from [HY14], where they attribute the argument to Schmerl.

¹⁴A set $X \subseteq M \models \text{PA}$ is *inductive* if (M, X) satisfies the induction schema in the expanded language. Equivalently, $X \subseteq M$ is inductive if there is an ACA_0 -realization for M which contains X . Compare to GBC-amenable, definition 1.39.

Proof. Consider the $\mathcal{L}_{\text{PA}}(S)$ theory $T = \text{TA} + \text{“}S \text{ is an inductive full satisfaction class, where TA} = \text{Th}(\mathbb{N}) \text{ is true arithmetic. By a well-known fact from computability theory } T \text{ is Turing-equivalent to } 0^{(\omega)}, \text{ the } \omega\text{-th iterate of the Turing jump of the empty set. However, } \text{Th}(\mathbb{N}, \text{TA}), \text{ which extends } T, \text{ is Turing-equivalent to } 0^{(\omega+\omega)}. \text{ So if } T \text{ were complete then we would get } 0^{(\omega)} \equiv_T 0^{(\omega+\omega)}, \text{ which is of course impossible. So } T \text{ must be incomplete. Now take } T_I \text{ and } T_U \text{ incompatible extensions of } T. \text{ That is, } T_I \cup T_U \text{ is inconsistent while each is individually consistent. Notice, though, that } T_I \text{ and } T_U \text{ have the same } \mathcal{L}_{\text{PA}} \text{ consequences, namely TA. So if } M \models \text{TA} \text{ is resplendent then there are } U, I \subseteq M \text{ so that } (M, U) \models T_U \text{ and } (M, I) \models T_I. \text{ These are the desired } M, U, \text{ and } I. \square$

Note that the restricting U and I to ω both yield TA. So U and I agree on standard truth, but not nonstandard truth. On the other hand, for truth about truth disagreement happens on the standard level.

Definition 3.33. For $n \leq \omega$ let ITR^n be the $\mathcal{L}_{\text{PA}}(T)$ theory asserting PA plus that T is an inductive full satisfaction class of length n .¹⁵ Let itr^n be the \mathcal{L}_{PA} consequences of ITR^n .

With this definition in mind, the above observation can be phrased as: if $M \models \text{itr}^n$ is countable and recursively saturated then M can be extended to a model of ITR^n .

Definition 3.34. Let $m < n \leq \omega$. Let itr_m^n be the reduct of ITR^n to language for a structure with an m -iterated full satisfaction class. More formally, we can use an $\mathcal{L}_{\text{PA}}(T)$ -formula to define an m -iterated full satisfaction class from an n -iterated full satisfaction class by restricting to the first m levels of the iterated full satisfaction class. Then itr_m^n is what ITR^n proves about this reduct.

Theorem 3.35. *Let $M \models \text{itr}^2$ be nonstandard, countable, and recursively saturated. Then there are $S, S' \subseteq M$ so that the following hold:*

- (M, S) and (M, S') are recursively saturated;
- S and S' are inductive full satisfaction classes;
- S can be extended to an inductive iterated full satisfaction class of length 2; and
- S' cannot be extended in this way.

Proof. It is easy to find S . Just take an inductive iterated full satisfaction class of length 2 over M and restrict it to get S . We can ensure (M, S) is recursively saturated because M is chronically resplendent. So the work is in getting S' . This reduces to the following claim.

Claim 3.36. The theory itr_1^2 is independent over $\text{Th}(M) + \text{ITR}^1$.

Given the claim, find S' by chronic resplendency to get a class over M so that the expansion satisfies $\text{Th}(M) + \text{ITR}^1$ but not itr_1^2 . So to finish the proof let us prove the claim. The basic idea is that having an inductive full satisfaction class allows a model to get a handle of the theory of its arithmetic reduct, enabling a diagonalization trick.

Consider the sentence $\text{Con}(\text{ITR}^1 + \text{Tr})$, where Tr is a name for the full satisfaction class. This sentence can be expressed in the language of ITR^1 , because ITR^1 is computably axiomatizable and the full satisfaction class gives access to Tr. Note that for any standard φ in the language of arithmetic and any $(N, S) \models \text{Th}(M) + \text{ITR}^1$ we get that $\varphi \in S$ if and only if $\varphi \in \text{Th}(M)$. Clearly, $\text{ITR}^2 \vdash \varphi$ so $\text{itr}_1^2 \vdash \varphi$. Let us see that $\text{Th}(M) + \text{ITR}^1$ does not prove $\text{Con}(\text{ITR}^1 + \text{Tr})$, which will then yield the claim. By the Gödel fixed-point lemma there is a sentence ψ so that ITR^1 proves

$$\psi \Leftrightarrow \underbrace{\forall x \Pr_{\text{ITR}^1 + \text{Tr}}(\psi, x) \Rightarrow \exists y < x \Pr_{\text{ITR}^1 + \text{Tr}}(\neg\psi, y)}_{=\rho(\psi)}$$

¹⁵The reader should think **iterated truth** for ITR.

where $\text{Pr}_{\text{ITR}^1 + \text{Tr}}(\theta, x)$ asserts x codes a proof of θ from the axioms of $\text{ITR}^1 + \text{Tr}$. It is immediate that ITR^1 proves $\rho(\psi) \Rightarrow \neg \text{Con}(\text{ITR}^1 + \text{Tr})$. So ITR^1 proves $\text{Con}(\text{ITR}^1 + \text{Tr}) \Rightarrow \neg\psi$. Now suppose towards a contradiction that $\text{Th}(M) + \text{ITR}^1$ proves $\neg\psi$. Then there is a standard natural number which codes this proof. Now work in a model of $\text{Th}(M) + \text{ITR}^1$. Then, this model thinks that $\text{ITR}^1 + \text{Tr}$ proves $\neg\rho(\psi)$, which is equivalent to

$$\exists x \text{Pr}_{\text{ITR}^1 + \text{Tr}}(\psi, x) \wedge \forall y < x \neg \text{Pr}_{\text{ITR}^1 + \text{Tr}}(\neg\psi, x).$$

There are two cases to consider, the first being the case that there is a witnessing x which is standard. Then the proof coded by x could only use formulae from the standard part of Tr , which is $\text{Th}(M)$. Thus we would get a standard proof of ψ from $\text{Th}(M) + \text{ITR}^1$, which would be a contradiction. The second case is then that all witnessing x are nonstandard. But then $\forall y < x \neg \text{Pr}_{\text{ITR}^1 + \text{Tr}}(\neg\psi, x)$ cannot be a theorem of $\text{ITR}^1 + \text{Tr}$, as there is a standard y which codes a proof of $\neg\psi$ from $\text{Th}(M) + \text{ITR}^1$, and so our model will think y codes a proof of $\neg\psi$ from $\text{ITR}^1 + \text{Tr}$. In either case we get a contradiction, so our original assumption that $\text{Th}(M) + \text{ITR}^1$ proves $\neg\psi$ must be false, and so it cannot prove $\text{Con}(\text{ITR}^1 + \text{Tr})$. This completes the proof of the claim, which completes the proof of theorem. \square

Krajewski's methods give that there are continuum many such S and S' .

There is nothing special about about 1 and 2 in theorem 3.35. We can get the same result for m -iterated full satisfaction classes and n -iterated full satisfaction classes for $m \leq n$.

Theorem 3.37. *Let $M \models \text{PA}$ be countable and recursively saturated and $m \leq n \in M$. Assume that $M \models \text{itr}^n$. Then there is $S \subseteq M$ so that*

- (M, S) is recursively saturated;
- S is an inductive m -iterated full satisfaction class for M ; and
- S cannot be extended to an inductive $(m + 1)$ -iterated full satisfaction class.

Theorem 3.35 is a direct consequence of this result: apply it with $m = n = 2$ and with $m = 1$ and $n = 2$ to get, respectively, S and S' .

Proof sketch. If $n = m$ this is easy. For $n < m$, use a similar Gödel–Rosser trick to show that itr_m^n is independent over $\text{Th}(M) + \text{ITR}^m$. \square

Again, Krajewski's work implies that there are continuum many such S .

Corollary 3.38. *Let $M \models \text{PA}$ be countable and recursively saturated and $k \leq m < n \in M$. Assume that $M \models \text{itr}^n$. Then there is $S \subseteq M$ so that*

- (M, S) is recursively saturated;
- S is an inductive k -iterated full satisfaction class for M ; and
- S can be extended to an inductive m -iterated full satisfaction class but no further.

Proof. Apply the theorem, then restrict the m -iterated full satisfaction class to get a k -iterated full satisfaction class. \square

Given countable and recursively saturated $M \models \text{PA}$ we can form a tree consisting of the inductive iterated full satisfaction classes over M . Namely, for $m < n \in M$ an m -iterated full satisfaction class S_m is before an n -iterated full satisfaction class S_n in the tree if S_m is the restriction of S_n to the first m levels. The results in this section tell us that this tree has lots of branches of all possible lengths.

For concreteness, suppose $M \models \text{itr}^\omega$, i.e. the \mathcal{L}_{PA} -consequences of $\text{PA} + \text{“}T$ is an inductive ω -iterated full satisfaction class”. (Here, ω is in the sense of the model M .) The tree of inductive iterated full satisfaction classes for M —see figure 3.2—has continuum many branches, coming from the continuum many inductive ω -iterated full satisfaction classes. The above corollary implies that for any $m < n \in M$ this tree has




Figure 3.2: The tree of inductive iterated full satisfaction classes over a model of arithmetic. The node at depth m has continuum many extensions to leaf nodes of depth k .

continuum many nodes of depth m which extend to a node of depth n , but no further. Moreover, if a node s of depth m has extensions to a node of depth n , then for any k between m and n we have that s has continuum many extensions to a leaf node of depth k .

3.5 Non-elementary transfinite recursions

We now return to the infinite. We have investigated transfinite recursion for elementary properties. What about second-order properties?

Definition 3.39. Let k be a natural number. We define the Σ_k^1 -transfinite Recursion schema Σ_k^1 -TR. Instances of Σ_k^1 -Transfinite Recursion take the following form: let $\varphi(x, Y, P)$ be a Σ_k^1 -formula, possibly with a class parameter P , and R be a well-founded class relation with transitive closure $<_R$. The instance of recursion for φ and R asserts that there is a class $S \subseteq \text{dom } R \times V$ which satisfies

$$(S)_r = \{x : \varphi(x, S \upharpoonright r, A)\}$$

for all $r \in \text{dom } R$. Here, $(S)_r = \{x : (r, x) \in S\}$ denotes the r -th slice of S and

$$S \upharpoonright r = S \cap (\{r' \in \text{dom } R : r' <_R r\} \times V)$$

is the partial solution below r .

If Γ is a well-order then Σ_k^1 -TR $_\Gamma$ is the schema obtained by restricting Σ_k^1 -TR to only ask for solutions to recursions of height $\leq \Gamma$.

Note that Σ_0^1 -TR is another name for ETR.

Before wading into a finer analysis, let us put the Σ_k^1 -TR in the context of the more familiar Π_k^1 -CA. First, an upper bound.

Proposition 3.40 (Over GBc^-). Π_{k+1}^1 -CA proves Σ_k^1 -TR.

Proof. Let R be a well-founded relation and φ be Σ_k^1 , possibly with parameters. By Comprehension form the class $\{r \in \text{dom } R : \text{there is no partial solution of } \varphi \text{ up to } r\}$. This uses Π_{k+1}^1 -Comprehension because it asserts there is no class satisfying the Σ_k^1 -property of being a partial solution to a Σ_k^1 -recursion. We want to see this class is empty, so assume otherwise. Then it has a minimal member r . That is, for all $r' R r$ we have a partial solution up for φ up to r' . But then there is a partial solution up to r , a contradiction. \square

We also get a lower bound.

Observation 3.41 (Over GBC^-). Π_k^1 -CA is equivalent to Σ_k^1 -TR₁.

Proof. Simply observe that a Σ_k^1 recursion of length 1 is just asking for a single Σ_k^1 -definable class. \square

So, because Σ_k^1 -TR clearly proves Σ_k^1 -TR₁ we have that Σ_k^1 -TR is between Π_k^1 -CA and Π_{k+1}^1 -CA. Later—specifically theorem 3.50 and a special case of theorem 3.1—we will see that this separation is by consistency strength.

This observation can also be used to show that KM can be alternatively axiomatized by a transfinite recursion principle.

Proposition 3.42. Over GBC^- , the following are equivalent.

1. The full second-order Comprehension schema; and
2. Σ_ω^1 -Transfinite Recursion, the schema asserting that every transfinite recursion of a second-order property has a solution.

Proof. (1 \Rightarrow 2) Use Π_{k+1}^1 -Comprehension to get an instance of Σ_k^1 -Transfinite Recursion. (2 \Rightarrow 1) Use Σ_{k+1}^1 -Transfinite Recursion to get an instance of Π_k^1 -Comprehension. \square

In the sequel we will stratify Σ_k^1 -TR by the length of recursions, separating them by consistency strength similar to the previous analysis of ETR. Before doing so, however, let us see why the arguments for ETR do not carry over immediately.

To separate $\text{ETR}_{\Gamma.\omega}$ from ETR_Γ we started with a model of $\text{ETR}_{\Gamma.\omega}$ and looked at the coded V -submodel generated from initial segments of $\text{Tr}_{\Gamma.\omega}(G)$, where G was a global well-order. This submodel then satisfied ETR_Γ because iterated first-order truth predicates from the original model continued to be iterated first-order truth predicates in the submodel. This is because the two models have the same first-order part, so they agree on well-orders and they agree on first-order truth. We then inductively get that they agree on iterated truth. And since fragments of Elementary Transfinite Recursion are equivalent to the existence of certain iterated truth predicates, we get ETR_Γ in the submodel.

But that strategy cannot work for the non-elementary transfinite recursion principles Σ_k^1 -TR for $k \geq 1$. The problem is simple: a Σ_k^1 -truth predicate must be a hyperclass. But this hyperclass contains $(X = X, A)$ for all classes A . If we had a class coding this truth predicate we could thus from it define a class coding the hyperclass of all classes. But that is impossible, as seen by an easy diagonalization argument. Because of this, there is no hope for having an iterated Σ_k^1 -truth predicate. So we must take a different strategy.

Instead, we will work directly with solutions to Σ_k^1 -transfinite recursions. Here we run into a potential obstacle. Unlike with elementary transfinite recursions which depend only upon the sets (and possibly finitely many class parameters), determining whether a class is a solution to a Σ_k^1 -transfinite recursion requires quantifying over all the classes. So if we restrict the classes to a V -submodel it may be that the smaller model disagrees about what is a solution to the recursion. We overcome this obstacle by ensuring that our V -submodel is sufficiently elementary in the larger model.

It will be convenient to work in the unrolling, so let us see what the theory of the unrolling is. First, recall a special case of definition 2.29.

Definition 3.43. Let $k \in \omega$. Then the Σ_k -Transfinite Recursion principle is the axiom schema consisting of the following axiom for each Σ_k -formula $\varphi(x, y, a)$:

Suppose a is a parameter so that $\varphi(x, y, a)$ defines a class function $F : V \rightarrow V$ and δ is an ordinal. Then there is a function $s : \delta \rightarrow V$ so that for all $i \in \delta$ we have $s(i) = F(s \upharpoonright i)$.

We can refine this to the Σ_k -Transfinite Recursion $\leq \gamma$ principle, which restricts the ordinals δ allowed to only those $\delta \leq \gamma$.

Theorem 3.44. *Let $(M, \mathcal{X}) \models \text{GBC} + \Pi_k^1\text{-CA}$ be a model of second-order set theory, with $k \geq 1$, and let \mathfrak{U} be the unrolled model obtained from (M, \mathcal{X}) . Then, if $(M, \mathcal{X}) \models \Sigma_k^1\text{-TR}$ we have $\mathfrak{U} \models \text{wZFC}_1^-(k) + \Sigma_k\text{-Transfinite Recursion}$.*

Proof. We have already seen in chapter 2 that the unrolling of a model of $\text{GBC} + \Pi_k^1\text{-CA}$ must satisfy $\text{wZFC}_1^-(k)$. The new content is that \mathfrak{U} satisfies Σ_k -Transfinite Recursion. This is done similar to the argument for the proof of proposition 2.31.

Consider an instance of Σ_k -transfinite recursion. That is, F is a class function $\mathfrak{U} \rightarrow \mathfrak{U}$ which is Σ_k -definable, possibly using a membership code as a parameter, and D is a membership code equipped with $<_D$ a membership code for a well-ordering of D . We want to see that there is a membership code for the desired s . First, observe that in the ground universe that F is Σ_k^1 -definable; cf. lemma 2.22.

We will build the the desired membership code S via an instance of Σ_k^1 -Transfinite Recursion. The idea is to mimic the recursion to produce s , but in membership codes. This introduces some extra work, since we have to deal with the picky details of membership codes.

The iteration proceeds as follows, with an outer layer and an inner layer. The outer layer occurs on elts D according to the well-ordering corresponding to the membership code $<_D$ (see corollary 2.17). Each step d in the outer layer produces a partial construction of S , call it S_d . We start with $S_0 = \emptyset$ and take unions at limit stages. The hard work is done in the successor step, where the inner layer of the transfinite recursion occurs. We start with S_d and want to produce S_{d+1} . By construction, each $d' \in \text{elts } D$ which comes before d in $<_D$ is in S_d . More, there is a corresponding node, call it $f(d')$, which represents $F(S_d \upharpoonright d')$ and then nodes for $\{d'\}$, $\{d', f(d')\}$, and $(d', f(d'))$ above, similar to the constructions in proposition 2.15 and lemma 2.16. In particular, S_d itself may not be a membership code. Modify S_d to produce a membership code U by adding a top node t_U and edges from each $(d', f(d'))$ node in S_d to t_U . Then, we have a membership code $F(U)$ by Elementary Comprehension. In the proof for proposition 2.31 at this point we build the maximum initial partial isomorphism between S_d and $F(U)$ by elementary transfinite recursion. Here we can do it in a single step, as asserting the existence of such is a Σ_1^1 -assertion. Given this partial isomorphism we can glue a copy of $F(U)$ onto S_d , as in the argument for proposition 2.24. Then, add $d + 1$ to S_d along with nodes for $\{d + 1\}$, $\{d + 1, t_{F(U)}\}$, and $(d + 1, t_{F(U)})$ and the corresponding edges to produce S_{d+1} . \square

Corollary 3.45. *Let $(M, \mathcal{X}) \models \text{GBC} + \Sigma_k^1\text{-TR}_\Gamma$ where $k \geq 1$ and Γ is a well-order, possibly class-sized. Then the unrolled model \mathfrak{U} satisfies $\text{wZFC}_1^-(k)$ plus the Σ_k -Transfinite Recursion $\leq \Gamma$ principle.*

Proof. In the case for ETR, we could not control the height of the recursion to produce the desired membership code S because we did not know how long a recursion was needed to construct partial isomorphisms between membership codes. So although the outer layer of our recursion had Γ many steps—where Γ is the height of the first-order recursion we were mimicking—each of those steps consisted of an inner layer of recursion which could be very long. But for $\Sigma_k^1\text{-TR}$ those inner steps are of finite, bounded length. We produce the partial isomorphism in a single step, because the existence of such a function follows from an instance of Π_1^1 -Comprehension, and then immediately use to define the next approximation to S . So this recursion is $n \cdot \Gamma$ for some standard finite $n \geq 2$. Conclude that $n \cdot \Gamma \leq \Gamma \cdot n$ by some simple arithmetic of well-orders¹⁶ and observe that $\Sigma_k^1\text{-TR}_\Gamma$ is equivalent to $\Sigma_k^1\text{-TR}_{\Gamma \cdot n}$. So $\Sigma_k^1\text{-TR}_\Gamma$ suffices to prove there is a solution. \square

And in the other direction.

Proposition 3.46. *Suppose N is a model of $\text{wZFC}_1^-(k)$ plus the Σ_k -Transfinite Recursion $\leq \gamma$ principle. Let κ be the largest model of N and (M, \mathcal{X}) be the cut off model obtained from N , i.e. $M = V_\kappa^N$ and \mathcal{X} is the (definable) class in N consisting of all subsets of M . Then, $(M, \mathcal{X}) \models \text{GBC} + \Sigma_k^1\text{-TR}_\Gamma$, where $\Gamma \in \mathcal{X}$ is such that $N \models \Gamma \cong \gamma$.*

Consequently, if N further satisfies the full Σ_k -Transfinite Recursion principle, then $(M, \mathcal{X}) \models \Sigma_k^1\text{-TR}$.

¹⁶If Γ is finite then this is obvious and indeed equality holds. So consider the case where $\Gamma \geq \omega$. Take Δ and finite m so that $\Gamma = \omega \cdot \Delta + m$. Then, $n \cdot \Gamma = \omega \cdot \Delta + nm = \Gamma + (n - 1)m$. This is less than $\Gamma \cdot 2 \leq \Gamma \cdot n$ because $\Gamma \geq \omega$.

Proof. Consider an instance of Σ_k -Transfinite Recursion for a recursion along Γ , possibly using a parameter from \mathcal{X} . We want to find the subset of V_κ^N which witnesses that this recursion has a solution. This is done in the obvious way in N by means of an instance of Σ_k -Transfinite Recursion of height γ , where γ is the ordinal isomorphic to Γ . \square

Now let us see that we may assume a fragment of Class Collection without loss.

Theorem 3.47. *Let $(M, \mathcal{X}) \models \text{GBC} + \Sigma_k^1\text{-TR}_\Gamma$ for $\Gamma \in \mathcal{X}$ and $k \geq 1$. Then there is $\mathcal{Y} \subseteq \mathcal{X}$ second-order definable over (M, \mathcal{X}) so that $(M, \mathcal{Y}) \models \text{GBC} + \Sigma_k^1\text{-TR}_\Gamma + \Sigma_k^1\text{-Class Collection}$.*

Proof. Fix $G \in \mathcal{X}$ a global well-order of M . Set $\mathcal{Y} = \mathcal{L}(M, G)$, the hyperclass of (M, G) -constructible classes. (See section 2.3 for a definition.) Corollary 2.48 tells us that $(M, \mathcal{Y}) \models \text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$. It remains only to see that $(M, \mathcal{Y}) \models \Sigma_k^1\text{-Transfinite Recursion}$ for recursions of height $\leq \Gamma$. By proposition 3.46 it suffices to show that $\mathfrak{L}(M, G)$, the (M, G) -constructible unrolling of (M, \mathcal{X}) , satisfies the Σ_k -Transfinite Recursion $\leq \gamma$ principle, where $\gamma \in \mathfrak{L}(M, G)$ is the ordinal isomorphic to Γ .

Work in the unrolling \mathfrak{U} of (M, \mathcal{X}) . Then, $\mathfrak{L}(M, G) = L(M, G)^\mathfrak{U}$. Let $\varphi(x, y, a)$ be a Σ_k -formula with $a \in \mathfrak{L}(M, G)$ a parameter. Then, $\varphi(x, y, a)^{L(M, G)}$ is Σ_k (in parameters). So applying an instance of Σ_k -Transfinite Recursion $\leq \gamma$ in \mathfrak{U} gives a solution to the recursion of $\varphi(x, y, a)^{L(M, G)}$ along γ . We want to see that this solution is in $\mathfrak{L}(M, G)$. Using Σ_k -Collection inside $\mathfrak{L}(M, G)$, we can for each $i \in \gamma$ find α_i so that $L_{\alpha_i}(M, G)$ sees a partial solution up to i for this recursion. So we have the sequence $\langle \alpha_i : i \in \gamma \rangle \in \mathfrak{L}(M, G)$ and thus can get the entire solution in $\mathfrak{L}(M, G)$, as desired. \square

Because this \mathcal{Y} is a definable hyperclass in (M, \mathcal{X}) any consistency assumptions witnessed in (M, \mathcal{Y}) are also visible to (M, \mathcal{X}) . Of particular interest to us is the following consequence thereof: If $\Sigma_k^1\text{-TR}_\Gamma + \Sigma_k^1\text{-Class Collection}$ proves there is a coded V -submodel of $\Sigma_k^1\text{-TR}_\Delta$ then so does $\Sigma_k^1\text{-TR}_\Gamma$.

Theorem 3.48. *Suppose $(M, \mathcal{X}) \models \text{GBC} + \Sigma_k^1\text{-TR}_{\Gamma \cdot \omega} + \Sigma_k^1\text{-Class Collection}$, for $k \geq 1$ and $\Gamma \in \mathcal{X}$, has that every class is (M, G) -constructible for some fixed global well-order G . Then there is $\mathcal{Y} \subseteq \mathcal{X}$ a coded V -submodel of (M, \mathcal{X}) so that $(M, \mathcal{Y}) \models \text{GBC} + \Sigma_k^1\text{-TR}_\Gamma$.*

By the above remarks, this immediately yields the following corollary.

Corollary 3.49. *Suppose $(M, \mathcal{X}) \models \Sigma_k^1\text{-TR}_{\Gamma \cdot \omega}$ for $k \geq 1$ and $\Gamma \in \mathcal{X}$. Then there is $\mathcal{Y} \subseteq \mathcal{X}$ a coded V -submodel of (M, \mathcal{X}) so that $(M, \mathcal{Y}) \models \Sigma_k^1\text{-TR}_\Gamma$. \square*

Proof of theorem 3.48. Unroll (M, \mathcal{X}) to $\mathfrak{U} \models \text{ZFC}_\Gamma^-(k)$ plus Σ_k -Transfinite Recursion $\leq \gamma$ plus $V = L(M, G)$ for some $G \in \mathcal{X}$ a global well-order.

Recall lemma 2.62.1, which asserted that our unrolled model satisfies Σ_k -reflection along the $L_\alpha(M, G)$ -hierarchy. Using the instance of this for the universal Σ_k -formula gives that there are club many δ so that $L_\delta(M, G) \prec_{\Sigma_n} L(M, G)$. Therefore, we are done if we can show that there are club many ordinals α so that $L_\alpha(M, G)$ is closed under solutions to Σ_k -transfinite recursions of height $\leq \gamma$, where γ is the ordinal isomorphic to Γ . This is because if v is in both of these clubs then $L_v(M, G)$ will satisfy $\text{ZFC}_\Gamma^-(k)$ plus Σ_k -Transfinite Recursion $\leq \gamma$. So if $\mathcal{Y} \subseteq \mathcal{X}$ is the hyperclass consisting of classes which appear in $L_v(M, G)$, then \mathcal{Y} is coded in \mathcal{X} , since $L_v(M, G)$ is a set in the unrolling.

Work inside the unrolling \mathfrak{U} . Fix an arbitrary ordinal α_0 . We will use an instance of Σ_k -Transfinite Recursion $\leq \gamma \cdot \omega$ to find $\alpha > \alpha_0$ as in the above paragraph. It is obvious that the class of such α is closed, so this will suffice to show that it is club. We build α by means of a Σ_k -transfinite recursion of height $\gamma \cdot \omega$. The outer layer of this recursion, which has ω many steps, builds a sequence $\langle \alpha_n : n \in \omega \rangle$ so that $L_{\alpha_{n+1}}(M, G)$ has solutions for Σ_k -transfinite recursions of height γ with parameters from $L_{\alpha_n}(M, G)$. The inner layer, building α_{n+1} from α_n has height γ , so that the whole recursion has height $\gamma \cdot \omega$.

For the inner recursion, we build a grid of ordinals with width $\omega \times L_{\alpha_n}(M, G)$ and height γ . Each column of this grid is indexed by (φ, a) where $\varphi(x, y, a)$ is some formula and $a \in L_{\alpha_n}(M, G)$.¹⁷ We then use the

¹⁷Since we have a column for each (n, a) , even when n is not (the Gödel code of) a Σ_k -formula of appropriate arity, for those bad n just have the ordinals in the column all be α_n .

universal Σ_k -formula to simultaneously build these columns upward. Namely, in the (φ, a) -th column at row i , the ordinal $\xi_{(\varphi, a)}^i$ we put in is an ordinal $\xi \geq \sup_{j < i} \xi_{(\varphi, a)}^j$ so that $L_\xi(M, G)$ has the length i partial solution to the recursion given by $\varphi(x, y, a)$. This a single step because recognizing such a ξ is a Σ_k property. So we fill out the grid in γ many steps and then set $\alpha_{n+1} = \sup \xi_{(\varphi, a)}^i$ to be the supremum of the ordinals in the grid.

Then at $\alpha = \sup_n \alpha_n$ we have caught our tail and have that $L_\alpha(M, G)$ is closed under solutions to Σ_k -transfinite recursions of height $\leq \gamma$. And since we could get such $\alpha > \alpha_0$ for arbitrary α_0 , there is a club of such α .

Finally, note that both of the class clubs we are looking at are Σ_k -definable. To say that $L_\delta(M, G)$ reflects the universal Σ_k -formula is a Σ_k property. And to say that $L_\alpha(M, G)$ is closed under solutions to γ -length Σ_k -transfinite recursions is Σ_k , since it is a Σ_k to check whether something is a solution. So $\text{ZFC}^-(k)$ suffices to prove that these class clubs have nonempty intersection. \square

This completes the proof of theorem 3.1.

Finally, let us see by a similar argument that Π_{k+1}^1 -CA and Σ_k^1 -TR can be separated by consistency strength.

Theorem 3.50 (Over GBC^-). *Let k be a natural number. Suppose Π_{k+1}^1 -CA holds. Then there is a coded V -submodel of $\text{GBC}^- + \Sigma_k^1$ -TR.*

Corollary 3.51. $\text{GBC} + \Pi_{k+1}^1$ -CA proves $\text{Con}(\text{GBC} + \Sigma_k^1\text{-TR})$ and $\text{GBC}^- + \Pi_{k+1}^1$ -CA proves $\text{Con}(\text{GBC}^- + \Sigma_k^1\text{-TR})$.

Proof of theorem 3.50. Work over $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_{k+1}^1$ -CA. Fix a global well-order G and consider the (M, G) -constructible unrolling $W = \mathfrak{L}(M, G)$. (See section 2.3 for further details of this construction.) Recall that W satisfies Σ_{k+1} -Separation and Σ_{k+1} -Collection. Work in W .

Let $C \subseteq W$ be the definable club class of ordinals α so that $L_\alpha(M, G) \prec_{\Sigma_k} W$. The existence of such C follows from a reflection argument using the $L(M, G)$ -hierarchy. Let α_0 be least $> \text{Ord}^M$ in this club C . Given α_n , pick $\alpha_{n+1} > \alpha_n$ from C which is closed under solutions to Σ_k -transfinite recursions with parameters and lengths in $L_{\alpha_n}(M, G)$. Such solutions exist because Σ_{k+1} -Collection implies the existence of solutions to transfinite recursions of Σ_k properties. And being the solution to such a recursion is a Σ_k -expressible property, so α_{n+1} exists by an instance of Σ_k -Collection. Set $\alpha = \sup_n \alpha_n$, which exists by yet another instance of Collection. Then $L_\alpha(M, G) \models \Sigma_k$ -Transfinite Recursion. Let \mathcal{Y} be the cutting off of $L_\alpha(M, G)$, i.e. the definable class over W consisting of all subsets of M which are in $L_\alpha(M, G)$. Then $(M, \mathcal{Y}) \models \text{GBC}^- + \Sigma_k^1$ -TR, as desired. \square

Chapter 4

Least models

Es liegt nämlich nahe, das Axiom [der Beschränktheit] in vermeintlich präziserer Form so zu fassen, daß unter allen möglichen Realisierungen des Axiomensystems—wobei isomorphe als nicht verschieden zu betrachten wären—der “Durchschnitt”, d.h. der kleinste gemeinsame Teilbereich, gewählt werden soll. Sofern man dieser Fassung nicht überhaupt einen scharfen Sinn abstreiten will, so ist es jedenfalls möglich, daß die dem Umfang nach verschiedenen möglichen Realisierungen des Axiomensystems nicht einen kleinsten gemeinsamen Teilbereich aufweisen, in dem gleichfalls sämtliche Axiome befriedigt würden.

Abraham Fraenkel

One desideratum for early axiomatizers of set theory was categoricity, similar to the categoricity results about \mathbb{N} and \mathbb{R} . Fraenkel [Fra22] and [Fra28, pp. 355–356] wanted a “Beschränktheitsaxiom” which would state, essentially, that the only objects that exist are those which are guaranteed to exist by the other axioms. We know now that there can be no such axiom.¹ But we can transmute this question of axioms into a model theoretic question. At first approximation, what we would like to know is: What are the objects that must be in every model of T ? Do they form a model of T ? A positive answer would give a partial realization of Fraenkel’s desire. While we cannot write down an axiom (or a set of axioms) which uniquely picks out this structure, we would know that if we restrict to the bare minimum possible we still get a model of T .

As stated, this admits a trivial answer. By the nonstandardness phenomenon, the only objects in every model of set theory are those appearing in V_n for some *standard* n . These form V_ω , which of course lacks any infinite sets. So in this naive form, the question is not interesting. But we can refine it to a more interesting question by restricting which models we look at. A natural restriction is to only look at transitive models.

¹This must be qualified. The Löwenheim–Skolem theorem implies that there can be no axiomatization of set theory *in first-order logic* which admits a unique model. But in different logics we can have categoricity. For instance, it follows from work of Zermelo [Zer30] that second-order ZFC—i.e., ZFC but with Separation and Collection formulated as single axioms in second-order logic—plus “there are no inaccessible cardinals” has a unique model, namely V_κ where κ is the least inaccessible. (It must be noted, however, that this theory has a very ad hoc feel.)

But the set theories considered in this dissertation are all formulated in first-order logic, ruling out any Beschränktheitsaxiom. It would go too far astray to give a defense here of why we would want to restrict to first-order logic, but let me mention [Vää01]. See also the epigraph to chapter 2.

They hold a special place in set theoretic practice and many set theorists believe we have a determinate notion of well-foundedness and can thereby pick out the transitive models. So the question becomes: is the intersection of all the transitive models of T itself a model of T ? Equivalently, is there a least transitive model of T ?

Before moving to the main topic—models of second-order set theories—let me quickly review what is known for models of first-order set theory. Shepherdson [She53] and, independently, Cohen [Coh63] proved that there is a least transitive model of ZFC.² This model is L_α where α is the least ordinal so that there is a transitive model of ZFC of height α . Their argument, which uses that ZFC is absolute to L , generalizes to stronger theories.³ In particular, it generalizes to theories extending ZFC by asserting the existence of “small” large cardinals. Formally, say that a first-order set theory T is *absolute to L* if $M \models T$ implies that $L^M \models T$. Then if there is a transitive model of T which is absolute to L there is a least transitive model of T . So there is a least transitive model of, for example, ZFC plus there is a proper class of Mahlo cardinals.

But this phenomenon does not extend too far up the large cardinal hierarchy. It fails for large cardinals which give elementary embeddings of the universe into an inner model.

Proposition 4.1. *Let $T \supseteq$ ZFC be a theory which proves there is a measurable cardinal. Then there is not a least transitive model of T .*

Proof. Suppose otherwise that N is the least transitive model of T . Let $M \subseteq N$ be the inner model obtained from taking an ultrapower of M using a measure on a measurable cardinal in N . By leastness, $M = N$, a contradiction. \square

On the other hand, we can recover something of this phenomenon for measurable cardinals and beyond. Results from inner model theory show that if an ordinal κ is measurable in some model then there is a least model in which κ is measurable. And this has been extended higher up the large cardinal hierarchy, although it remains open in many cases, most notably for κ supercompact.

The lesson to be had is that for strong enough first-order set theories, we do not have least transitive models. However, if we restrict the models we look at in some further (non-first-order expressible) way, then we do get least models. In this chapter we will see that there is a similar phenomenon for second-order set theories, except the reason and the ‘fix’ to get leastness are different.

Definition 4.2. Let T be a second-order set theory. The *least transitive model of T* —if it exists—is the unique transitive $(M, \mathcal{X}) \models T$ so that (M, \mathcal{X}) is a submodel of any transitive model of T . The *least β -model of T* —if it exists—is the unique transitive β -model $(M, \mathcal{X}) \models T$ which is a submodel of any transitive β -model of T .⁴

The main theorem of this chapter answers which second-order set theories have least transitive models for a broad class of theories. In short, strong theories do not have least transitive models while weaker theories do.

Theorem 4.3.

- *There is not a least transitive model of KM nor of KMCC.*
- *For $k \geq 1$ there is not a least transitive model of $\text{GBC} + \Pi_k^1\text{-CA}$ nor of $\text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$.*
- *There is a least transitive model of $\text{GBC} + \text{ETR}_\Gamma$, for $\omega^\omega \leq \Gamma \leq \text{Ord}$.*

²Of course, their proofs require a consistency assumption, namely that there is some transitive model of ZFC at all.

³Or more precisely, Shepherdson’s argument generalizes to stronger theories. The essence of Shepherdson’s argument is the same as the standard contemporary argument that there is a least transitive model of ZFC. But Cohen uses a different argument which goes through what he calls “strongly constructible” sets, a strengthening of constructability, which I do not see how to generalize to get results about stronger theories.

⁴In chapter 1, we did not require β -models to be transitive. But every β -model is isomorphic to a transitive model so the extra requirement here is harmless. If one prefers to drop it, then one would need to tweak the definition so that the least β -model of T embeds into every β -model of T , rather than being a literal submodel.

- (Shepherdson [She53]) There is a least transitive model of GBC.

There is some redundancy in the statement of this theorem. In chapter 2 we saw that any model of KM contains a V -submodel of KMCC, so there is a least transitive model of KM if and only if there is a least transitive model of KMCC. (And similar remarks apply for $\text{GBC} + \Pi_k^1\text{-CA}$ versus $\text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$.)

The argument for the negative part of theorem 4.3, that strong second-order set theories do not have least transitive models, actually shows something stronger. Namely, it shows that a given fixed first-order part does not have a least KM-realization (or $(\text{GBC} + \Pi_k^1\text{-CA})$ -realization or . . .).

Definition 4.4. Let M be a model of first-order set theory and T be some second-order set theory. The *least T -realization for M* —if it exists—is the T -realization \mathcal{X} for M so that for any T -realization \mathcal{Y} for M we have $\mathcal{X} \subseteq \mathcal{Y}$.

Theorem 4.5. No countable $M \models \text{ZFC}$ has a least KM-realization. Moreover, if $k \geq 1$ then no countable $M \models \text{ZFC}$ has a least $(\text{GBC} + \Pi_k^1\text{-CA})$ -realization.

The proof of this uses little about the theory of M itself. All the important work takes place above M in the unrolling. As such, a version of this theorem goes through for KM^- and $\text{GBC}^- + \Pi_k^1\text{-CA}$. It also yields that least models cannot be recovered by moving to stronger theories.

Theorem 4.6. No computably axiomatizable extension of KM (in \mathcal{L}_\in) has a least transitive model. More generally, no computably axiomatizable extension of $\text{GBC}^- + \Pi_k^1\text{-CA}$, for $k \geq 1$, has a least transitive model.⁵

Before discussing β -models, where we get positive results even for strong theories, let me highlight the conspicuous absence of $\text{GBC} + \text{ETR}$ in theorem 4.3.

Question 4.7. Is there a least transitive model of $\text{GBC} + \text{ETR}$?

While this question remains open, something can be said about the structure of $(\text{GBC} + \text{ETR})$ -realizations for a model M .

Theorem 4.8. Let $M \models \text{ZFC}$ be $(\text{GBC} + \text{ETR})$ -realizable. Then M has a basis of minimal $(\text{GBC} + \text{ETR})$ -realizations, where amalgamable $(\text{GBC} + \text{ETR})$ -realizations⁶ sit above the same basis element. That is, there is a set $\{\mathcal{B}\}$ of $(\text{GBC} + \text{ETR})$ -realizations for M satisfying the following.

1. Elements of the basis are pairwise non-amalgamable;
2. If \mathcal{Y} is any $(\text{GBC} + \text{ETR})$ -realization for M then there is a unique basis element \mathcal{B} so that $\mathcal{Y} \supseteq \mathcal{B}$; and
3. If \mathcal{X} and \mathcal{Y} are amalgamable $(\text{GBC} + \text{ETR})$ -realizations for M then they sit above the same \mathcal{B} .

And we get the same result for $\text{GBC} + \text{ETR}$ if M has a definable global well-order.

Let us now turn to β -models. Weak theories have least transitive models while strong theories do not. For first-order set theories, we could get least models for strong theories by requiring extra from our models, namely by nailing down which ordinals are to be the large cardinals of the model. For second-order set theories we recover leastness by requiring that the model be correct about well-foundedness.

Theorem 4.9.

- (Folklore) There is a least β -model of KM.
- For $k \geq 1$ there is a least β -model of $\text{GBC} + \Pi_k^1\text{-CA}$.

⁵Indeed, this is true for more than just computably axiomatizable extensions. What we get is that if $T \supseteq \text{GBC}^- + \Pi_1^1\text{-CA}$ is an element of every transitive model of T then T cannot have a least transitive model. In particular, no arithmetical T or even hyperarithmetical $T \supseteq \text{GBC}^- + \Pi_1^1\text{-CA}$ can have a least transitive model.

⁶Two T -realizations \mathcal{X} and \mathcal{Y} for M are *amalgamable* if there is a GBC^- -realization \mathcal{Z} for M so that \mathcal{X} and \mathcal{Y} are both subsets of \mathcal{Z} .

- There is a least β -model of $\text{GBC} + \text{ETR}$.
- For any $\Gamma \geq \omega^\omega$ there is a least β -model of $\text{GBC} + \text{ETR}_\Gamma$.
- (Folklore) There is a least β -model of GBC .

Of course, the above results that there is a least model of such and such theory have consistency requirements, namely that the theory has any model (of appropriate type) at all! For the sake of readability I have suppressed mentioning the exact consistency assumptions here. Every result in this chapter follows from $\text{ZFC} +$ “there is an inaccessible cardinal”. For the reader who wants something more precise, see the statements of the theorems in the body of this chapter.

This chapter is organized as follows. First I collect some observations about β -models and well-founded classes that will be used later in the chapter. I then review Barwise’s notion of the admissible cover, which will be used to prove the results for models of strong theories. This sets us up to finally get to the results about the existence and non-existence of different kinds of least models. I have organized this chapter by the strength of the theories, as the methods used vary. I first give the results about strong theories. Next comes results for theories of medium strength. Finally, I give results about weak theories, or rather collect some results from the literature and from chapter 1. The chapter concludes with a coda on the analogy between second-order set theory and second-order arithmetic.

The section on models of strong theories will make essential use of the results from chapter 2. As such, the reader is strongly encouraged to look at that chapter before reading section 4.3.

4.1 β -models and well-founded classes

In this section I collect several observations about β -models and well-founded classes. They will be used many places in this chapter, usually without explicit citation.

Observation 4.10. *Suppose (M, \mathcal{X}) is a β -model of GBC^- and $(M, \mathcal{Y}) \models \text{GBC}^-$ is a V -submodel of (M, \mathcal{X}) . Then (M, \mathcal{Y}) is a β -model.*

Proof. Suppose towards a contradiction that (M, \mathcal{Y}) is not a β -model. Then there is $R \in \mathcal{Y}$ so that (M, \mathcal{Y}) thinks that R is ill-founded but in V we can see that R is well-founded. Because (M, \mathcal{Y}) thinks R is ill-founded there is a set $r \in M$ which witnesses this; R being ill-founded means there is some infinite descending sequence in R , but this sequence is countable and hence must be a set. But then (M, \mathcal{X}) thinks R is ill-founded since it also sees that r witnesses the ill-foundedness of R . But then, because (M, \mathcal{X}) is a β -model, R must really be ill-founded, a contradiction. \square

Note that we did not need near the full strength of GBC^- for this argument to go through. All we need is that the theory of the models be strong enough to verify that a class being ill-founded is witnessed by a set.

We can strengthen this observation.

Observation 4.11. *Suppose $(N, \mathcal{Y}) \models \text{GBC}^-$ is an Ord -submodel of a β -model $(M, \mathcal{X}) \models \text{GBC}^-$. Then (N, \mathcal{Y}) is a β -model.*

Proof. Take arbitrary ill-founded $R \in \mathcal{Y}$. Then, because (M, \mathcal{X}) is a β -model, there is a set $r \in M$ which witnesses the ill-foundedness of R . This r might not be in N , but it has some rank α in M . Now consider $r' = R \cap V_{\alpha+1}^N \in N$. This r' is well-founded if and only if R is, since the ill-foundedness in R occurs by rank α . But M thinks r' is ill-founded and well-foundedness is absolute between transitive models of ZFC^- , so N thinks r' is ill-founded. So N correctly thinks that R is ill-founded and, as R was arbitrary, N is correct about well-foundedness. \square

The following observation generalizes observation 4.10 in a different direction to tell us to what extent non- β -models with the same first-order part must agree on what is well-founded. In short, they must agree as much as possible.

Observation 4.12. *Suppose \mathcal{X} and \mathcal{Y} are GBC^- -realizations for M . If $R \in \mathcal{X} \cap \mathcal{Y}$ is a class relation then (M, \mathcal{X}) and (M, \mathcal{Y}) agree on whether R is well-founded.*

Proof. Because the two models have the same first-order part. \square

The following observation also appeared in the proof of corollary 3.15, but for the sake of clarity I reproduce it with a proof here.

Observation 4.13. *Suppose (N, \mathcal{Y}) is an Ord-submodel of (M, \mathcal{X}) , where both are models of GBC^- . Suppose $(N, \mathcal{Y}) \models R$ is well-founded. Then $(M, \mathcal{X}) \models R$ is well-founded.*

Proof. Suppose otherwise. Then there is an ordinal $\alpha \in M$ so that $M \models R \upharpoonright F''\alpha$ is ill-founded, where $F \in \mathcal{Y}$ is a bijection between Ord and $\text{dom } R$, which exists by Global Choice. But by Replacement $R \upharpoonright F''\alpha$ must be in N . And N is a transitive submodel of M and they are both models of ZFC^- , so they must agree on what sets are well-founded. So $N \models R \upharpoonright F''\alpha$ is ill-founded, so $(N, \mathcal{Y}) \models R$ is ill-founded, a contradiction. \square

4.2 Admissible covers and nonstandard compactness arguments

To get theorems 4.3 and 4.5 will require more than the elementary tools of admissible set theory. I will review the necessary material in this section.

Speaking roughly, the strategy to prove that strong second-order set theories do not have least transitive models will be to work with their unrollings. Given an unrolling W of a model (M, \mathcal{X}) , we want to produce a new model N so that N and W agree up to their largest cardinal, but we introduced ill-foundedness in N above that. Then, cutting off N we get (M, \mathcal{Y}) . But the ill-foundedness will mean that \mathcal{Y} cannot be contained in \mathcal{X} , so (M, \mathcal{X}) cannot be the least transitive model of our theory.

It has previously been studied by H. Friedman where we can introduce ill-foundedness in a model of set theory. He proved the following very general theorem.

First though, let me set up some terminology. If N is a model of set theory with A as a transitive submodel⁷ then say that N is *ill-founded at A* if there is an omega sequence of ordinals in $N \setminus A$ which is co-initial in $\text{Ord}^N \setminus \text{Ord}^A$.⁸ If A is an admissible set then \mathcal{L}_A denotes the associated admissible fragment of $\mathcal{L}_{\text{Ord}, \omega}$, i.e. the infinitary language consisting of the formulae in A . If A is countable then the Barwise compactness theorem applies to \mathcal{L}_A -theories.

Theorem 4.14 ([Fri73, theorem 2.2]). *Let A be a countable admissible set and let $T \subseteq A$ be an \mathcal{L}_A theory which is Σ_1 -definable in A . If there is a model of T which contains A then there is an ill-founded model N of T so that the following hold.*

- N contains A as a transitive submodel and N is ill-founded at A . In particular:
- $\text{wfp}(N) \supseteq A$;⁹ and
- $\text{Ord}^{\text{wfp}(N)} = \text{Ord}^A$.

Figure 4.1 illustrates the theorem.

Observe that in general we cannot get $\text{wfp}(N) = A$. For a counterexample, suppose that α is a countable ordinal so that $L_\alpha \models \text{ZFC}^- + \text{“every set is countable”}$. Then by Friedman’s theorem there is $N \supseteq_{\text{end}} L_\alpha$ satisfying $\text{KP} + V = L + \text{“every set is countable”}$ so that N is ill-founded at L_α . Consider an ordinal $\gamma \in N$ which is not in $\text{wfp}(N)$. Then, N thinks γ is countable so there is $G \subseteq \omega^2$ so that $(\omega, G) \cong (\gamma, \in)$. In particular, G is in $\text{wfp}(N)$. I claim that this G cannot be in L_α . To see this, note that L_α satisfies enough

⁷ A is a *transitive submodel* of N if $A \subseteq N$ and for all $a \in A$ and all $b \in N$ if $N \models b \in a$ then $b \in A$.

⁸This is equivalent to asking that A be *topless* in N , meaning that there is no smallest ordinal in N above the ordinals of A .

⁹ $\text{wfp}(N)$ is the well-founded part of N , that is the subset of N consisting of elements a so that the membership relation of N below a is well-founded.




Figure 4.1: An illustration of Friedman’s theorem. The dashed line represents that N is ill-founded at A .

Replacement that it can compute $\text{wfp}(G)$ from G . But $\text{wfp}(G) \cong \alpha$. So if $G \in L_\alpha$ then by Mostowski’s collapse lemma we must have $\alpha \in L_\alpha$, a contradiction.

The reader may find it useful to see a simple application of this theorem before moving on.

Proposition 4.15. *Consider Zermelo set theory Z , the first-order set theory axiomatized by Extensionality, Pairing, Union, Infinity, Powerset, and the axiom schema of Separation. There are transitive models of Z which are wrong about well-foundedness. That is, there is $M \models Z$ transitive with $R \in M$ so that $M \models R$ is well-founded but externally we can see that R is ill-founded.*

Since we are concerned with transitive models, they will automatically satisfy Foundation. As well, we can easily arrange so that M also satisfies Choice. So this cute little proposition can be seen as yet more evidence for the importance of Collection in the axioms for set theory. If we include no fragment of Collection, then we do not get the incredibly useful fact that well-foundedness is absolute for transitive models.

Proof. Let $A = L_{\omega_1^{\text{CK}}}$, where ω_1^{CK} is the Church–Kleene ordinal, i.e. the least admissible ordinal $> \omega$. It is easy to see that there are countable models of $\text{KP} + “V_{\omega+\omega}$ exists” which contain A . So by Friedman’s theorem there is $N \models \text{KP} + “V_{\omega+\omega}$ exists” which is ill-founded at A . Take $\gamma \in N$ a countable ordinal which is not in the well-founded part. Then, there is $G \subseteq \omega^2$ in N so that $(\omega, G) \cong (\gamma, \in)$.

Now let $M = V_{\omega+\omega}^N$. It is well-known that $V_{\omega+\omega} \models Z$. Moreover, KP is enough to verify this so that indeed $M \models Z$. Observe also that M is transitive, because all of its elements have rank $< \omega + \omega < \omega_1^{\text{CK}}$. But M is not correct about well-foundedness, because M thinks G is well-founded. \square

While Friedman’s theorem is fantastic, it is not quite general enough for my purposes. To illustrate the difficulty, start with a countable β -model $(M, \mathcal{X}) \models \text{KMCC}$. Consider its unrolling W . Then, $\text{Hyp}(M) \in W$ (see proposition 4.22 below). So by Friedman’s theorem we can find $N \models \text{ZFC}_\Gamma^-$ which is ill-founded at $\text{Hyp}(M)$ and so that M is a rank-initial segment of N . Then, N thinks that $\text{Hyp}(M)$ exists. However, $\text{Hyp}(M)^N$ cannot be in the well-founded part of N by construction. So Friedman’s theorem cannot be used to produce a model of ZFC_Γ^- which is ill-founded at $\text{Hyp}(M)^N$.

The problem is that Friedman’s theorem only applies to well-founded models of KP , whereas I need to be able to handle ill-founded models. Fortunately, Barwise developed machinery for compactness arguments over an ill-founded domain.

The important notion here is that of the admissible cover of a model of set theory, which I review here. For further details see the appendix to [Bar75]. Briefly, the admissible cover of $U \models \text{KP}$ is a certain admissible structure with U as its urelements. The admissible cover of U allows us to then apply the tools of admissible set theory to U , even though U itself is ill-founded.

First we must discuss set theory with urelements. An *urelement* is an object that is neither a set nor a class but can be an element of sets.¹⁰ In contemporary set theory we usually formulate things just in

¹⁰The reader (or maybe just me) may find it amusing to think of urelements as the opposite of classes. A class can have elements but cannot be an element whereas an urelement cannot have elements but can be an element. A set, of course, can both have elements and be an element.

terms of pure sets, i.e. those sets with no urelements in their transitive closures. This is harmless, as any structures of interest can be simulated in the pure sets. For instance, we do not need the natural numbers as urelements since we can instead work with the finite ordinals. But for this context we will want models with urelements. These structures, of course, can be simulated with pure sets. So despite the use of set theories with urelements, all of the below can be formalized in ordinary ZFC.

Formally, the theory we will work with is KPU, **K**ripke–**P**latek set theory with **u**relements. I give an axiomatization here, both for the benefit of the reader unfamiliar with KP—that is, KPU sans urelements—and to highlight the fact that KPU does not prove there is a set of all urelements.

Definition 4.16. The theory KPU is a first-order set theory formulated with urelements. That is, KPU is a two-sorted theory whose objects are *sets* and *urelements*. In addition to basic axioms asserting that sets and urelements are distinct, nothing is a member of an urelement, etc., KPU is axiomatized by the following axioms: Extensionality for sets, Pairing, Union, Δ_0 -Separation, Δ_0 -Collection, and Foundation. Here, Foundation is the schema whose instances are of the form

$$\exists x \varphi(x) \Rightarrow (\exists x \varphi(x) \wedge \forall y \in x \neg \varphi(y))$$

for each \mathcal{L}_ε -formula φ .¹¹

I will write (M, U) for the model of KPU with sets M and urelements U , suppressing writing the set-set and urelement-set membership relations. Often I will give a single name to this structure, usually in the fraktur font, such \mathfrak{C} or \mathfrak{R} . We will also consider structures with additional functions and relations, which I will denote by $(M, U; R_0, R_1, \dots)$. These structures will satisfy the schemata of Δ_0 -Separation and Δ_0 -Collection in the expanded language.

For an example of a model of KPU, let us consider a model with the reals as its urelements. That is, consider $(\mathbb{R}, +, \times, <)$ the set of real numbers with its arithmetic operations. Let M consist of all hereditarily finite sets above the reals. Formally, let $M_0 = \emptyset$ and M_{n+1} be the set of all finite subsets of $\mathbb{R} \cup M_n$. Then $M = \bigcup_{n \in \omega} M_n$. The reader can easily check that $\mathfrak{R} = (M, \mathbb{R}; +, \times, <)$ is a model of KPU. But notice that M does not contain \mathbb{R} , as all sets in M are finite. So \mathfrak{R} gives an example of a model of KPU without a set of all urelements. It also gives a model of KPU which does not satisfy the axiom of Infinity.

Let us consider another example of a model of KPU, this one more relevant to the present discussion. Let U be an ω -nonstandard first-order model of set theory with membership relation E . We will treat U as the urelements for a model of KPU. Similar to the previous example we can define the hereditarily finite sets above U . If M is the collection of such then $\mathfrak{U} = (M, U; E)$ is a model of KPU.

We would like to use a structure like this to mimic internal talk in U with talk of honest-to- V well-founded sets. But this \mathfrak{U} does not have enough sets to do so. For $u \in U$ let $u_E = \{v \in U : u E v\}$ be the set of what U thinks are the elements of u . Because M only consists of finite sets the only $u \in U$ for which $u_E \in M$ are those which really are finite. In particular, $(\omega^U)_E \notin M$ and any $\alpha \in \omega^U \setminus \omega$ will have $\alpha_E \notin M$. So this \mathfrak{U} cannot directly talk about all the ‘sets’ in U . To do so, we need more.

Definition 4.17. Let U be a (possibly ill-founded) model of first-order set theory with membership relation E . Then $\mathfrak{M} = (M, U; E, F) \models \text{KPU}$, a model of KPU with U as urelements, *covers* U if F is a function from U to M so that $F(u) = u_E$.

Any \mathfrak{M} which covers U can mimic \mathcal{L}_U talk. One useful fact about \mathcal{L}_A for admissible A is that every element of A is definable by a single \mathcal{L}_A -formula. Namely, $x = a$ is defined by the formula

$$\forall y y \in x \Leftrightarrow \bigvee_{b \in a} y = b$$

where “ $y = b$ ” is an abbreviation for the formula defining b . Because A is well-founded this recursive definition unwraps into a single \mathcal{L}_A -formula. A similar idea allows us to define elements of U by $\mathcal{L}_{\mathfrak{M}}$ -formulae: define

¹¹The models of KPU we will consider will all be well-founded, so they will automatically satisfy the strongest form of Foundation.




Figure 4.2: A picture of the desired generalization of Friedman's theorem to the ill-founded realm.

$x = a$ for $a \in U$ by

$$\forall y \text{ an urelement } y E x \Leftrightarrow \bigvee_{b \in F(a)} y = b.$$

More, F lets us translate bounded quantification for U to bounded quantification for \mathfrak{M} : replace $\exists x E y$ with $\exists x \in F(y)$ and similarly for universal quantification. So corresponding to each (of what U thinks is a) Δ_0 \mathcal{L}_U -formula is a Δ_0 $\mathcal{L}_{\mathfrak{M}}$ -formula, and similarly for Σ_1 or Π_1 formulae.

Barwise proved that there is a smallest admissible structure which covers U . This structure, the *admissible cover of U* , is the intersection of all admissible structures which cover U and enjoys many nice properties. I summarize them here.

Theorem 4.18 (Barwise [Bar75, appendix]). *Let $U \models \text{KP}$ be a possibly ill-founded model of set theory with membership relation E and let $\mathfrak{C} = (C, U; E, F)$ be the admissible cover of U .*

- *If U is countable then \mathfrak{C} is countable.*
- *The pure sets of \mathfrak{C} are isomorphic to the well-founded part of U .*
- *For any $A \subseteq U$ we have $A \in \mathfrak{C}$ if and only if there is $a \in U$ so that $A = a_E$.*
- *The infinitary \in -diagram of U ,¹² considered as a set of $\mathcal{L}_{\mathfrak{C}}$ -sentences, is Σ_1 -definable over \mathfrak{C} .*

With this notion in hand we are now ready to generalize Friedman's theorem to the ill-founded. We want, when starting with a possibly ill-founded model A to produce the picture in figure 4.2, a variation of the picture in figure 4.1 for Friedman's theorem. That is, given a theory T satisfying an appropriate consistency assumption, we want $N \supseteq A$ a model of T which is ill-founded at A .

Theorem 4.19. *Let $(A, E^A) \models \text{KP}$ be countable and $\mathfrak{C} = \text{Cov}_A$. Suppose that T is an $L_{\mathfrak{C}}$ theory which is Σ_1 -definable over \mathfrak{C} . If there is a model of T which contains A then there is $(N, E^N) \models T$ so that:*

- *A is a transitive submodel of N ;*
- *Ord^A is a proper initial segment of Ord^N ;*
- *There is an ω -sequence coinital in $\text{Ord}^N \setminus \text{Ord}^A$.*

Proof. Friedman's proof can be adapted to this context, using the technology of the admissible cover.

Extend T , if necessary, to include the infinitary \in -diagram of A . This extension is consistent as there is a model of T containing A . The goal is now to construct a further extension T' in a language with countably many new constants c_n so that the following conditions hold:

¹²That is, the collection of all sentences of the form $\forall x x \in a \Leftrightarrow \bigwedge_{b \in a} x = b$.

1. Each $\varphi \in T'$ is consistent;
2. For $\varphi \in \mathcal{L}_{\mathfrak{C}}$, either $\varphi \in T'$ or “ $\neg\varphi$ ” $\in T'$;
3. For $\bigwedge \Phi \in \mathcal{L}_{\mathfrak{C}}$ if $\Phi \subseteq T'$ then $\bigwedge \Phi \in T'$;
4. For “ $\forall x\varphi(x)$ ” $\in \mathcal{L}_{\mathfrak{C}}$ if for each c_n we have $\varphi(c_n) \in T'$ then “ $\forall x\varphi(x)$ ” $\in T'$;
5. For each $a \in A$, “ $a \in c_0$ ” $\in T'$; and
6. If “ $a \in c_n$ ” $\in T'$ for each $a \in A$ then there is $m > n$ so that “ $c_m \in c_n$ ” $\in T'$ and “ $a \in c_m$ ” $\in T'$ for each $a \in A$.

Conditions (1–3) ensure there is a model of T' . Condition (5) forces any model of T' to contain new ordinals. Conditions (4–6) force that the model of T' is ill-founded above A . To see this note, that if $\beta < \gamma$ are ordinals above Ord^A but below $\inf\{\text{rank}(c_n) : c_n \text{ above } A\}$ then for every c_n we get “ $\text{rank}(c_n) \geq \beta \Rightarrow \text{rank}(c_n) > \gamma$ ” $\in T'$. So by condition (4) we can conclude $\beta > \gamma$, a contradiction.

T' is constructed from T in ω many stages. We continually add new formulae to ensure properties (1–6) hold at the end. Fix an enumeration $\langle \varphi_n \rangle$ of the $\mathcal{L}_{\mathfrak{C}}(c_n : n \in \omega)$ -sentences¹³ so that c_n first appears after φ_n and before the first appearance of c_{n+1} .

- Define $T'_0 = T \cup \{a \in c_0 : a \in A\}$. This theory is consistent by Barwise compactness. This ensures property (5). Set $m_0 = 1$ and $m_{-1} = 0$.
- For $n \geq 0$, define T'_{3n+1} to be $T'_{3n} \cup \{\psi\}$, where ψ is chosen from φ_n and $\neg\varphi_n$ so as to be consistent with T'_{3n} . This step will ensure properties (1) and (2).
- For $n \geq 0$, define T'_{3n+2} as follows, according to which of three cases we fall into.
 - If φ_n is of the form $\bigwedge \Phi$ and $\neg\varphi_n$ is in T'_{3n+1} , then take $T'_{3n+2} = T'_{3n+1} \cup \{\neg\varphi\}$ where $\varphi \in \Phi$ and this is consistent. This step ensures property (3).
 - If φ_n is of the form $\forall x\psi(x)$ and $\neg\varphi_n$ is in T'_{3n+1} , then take $T'_{3n+2} = T'_{3n+1} \cup \{\neg\psi(c_m)\}$, where m is the index of the least unused c_m . This step ensures property (4).
 - Otherwise, just take $T'_{3n+2} = T'_{3n+1}$.
- For $n \geq 0$, define T'_{3n+3} as follows, according to which of two cases we fall into.
 - If there is $a \in A$ so that $T'_{3n+2} \cup \{a \notin c_{m_n}\}$ is consistent, take this to be T'_{3n+3} . Set $m_{n+1} = m_n$.
 - Otherwise, the theory $T'_{3n+2} \cup \{a \in c_{m_n} : a \in M\}$ is consistent. By Barwise compactness, so is the theory $T'_{3n+2} \cup \{a \in c_{m_n} : a \in A\} \cup \{c_{m_n} \in c_{m_{n-1}}\}$. Take this to be T'_{3n+3} and set m_{n+1} to be the index of the least unused c_m .
- Set $T' = \bigcup_n T'_n$.

By the construction for T'_{3n+3} , for every $n \geq -1$ we have that $c_{m_{n+1}} \in c_{m_n}$ and $a \in c_{m_n}$, for all $a \in A$ are in T' . This gives property (6). \square

¹³To be clear, by $\mathcal{L}_{\mathfrak{C}}(c_n : n \in \omega)$ I mean the infinitary language consisting of formulae in \mathfrak{C} in the language of \mathfrak{C} with additional symbols c_n for $n \in \omega$.

4.3 Strong theories

In this section we will get results about least models of strong theories, those of strength $\text{GBC} + \Pi_1^1\text{-CA}$ and above. The main results, that strong theories do not have least transitive models, will be derived from the following master lemma.

Master Lemma 4.20. *Let $T \supseteq \text{GBC}^- + \text{ETR}$ be a second-order set theory which proves the existence of $\text{Hyp}(V)$. Suppose $(M, \mathcal{X}) \models T$ is countable and T is in M . Then there is $\mathcal{Y} \subseteq \mathcal{P}(M)$ so that $(M, \mathcal{Y}) \models T$ but $\mathcal{X} \not\subseteq \mathcal{Y}$.*

Before proving this master lemma I must clarify what it means for a second-order set theory to prove that $\text{Hyp}(V)$ exists. Recall that for a set a that $\text{Hyp}(a)$ is the smallest admissible set h with $a \in h$. Always, $\text{Hyp}(a) = L_\alpha(a)$ where α is the least ordinal ξ so that $L_\xi(a) \models \text{KP}$. Of course, if A is a proper class then there can be no class, admissible or otherwise, with A as an element. So it does not make literal sense to talk of $\text{Hyp}(V)$ inside a model of second-order set theory. But recall from chapter 2 that models of ETR can reach higher than Ord , coding ‘sets’ of high rank by class-sized relations. (Indeed, this is why the master lemma asks that $T \supseteq \text{GBC}^- + \text{ETR}$.) In particular, there are codes for ‘sets’ which look like $L_\Gamma(A)$, for Γ a class well-order and A a class. In $\text{GBC}^- + \text{ETR}$, we can talk about the theory of a coded transitive ‘set’, so it makes sense to ask whether $L_\Gamma(A)$ satisfies KP . If there is a class well-order Γ so that $L_\Gamma(A) \models \text{KP}$ then we say that $\text{Hyp}(A)$ exists. Given such a Γ there is a least initial segment Γ_0 of Γ so that $L_{\Gamma_0}(A) \models \text{KP}$. This is $\text{Hyp}(A)$.

From the perspective of the unrolling, if $(M, \mathcal{X}) \models \text{Hyp}(V)$ exists then the unrolling has a set which (it thinks) is $\text{Hyp}(M)$.

Proof of master lemma 4.20. Because $T \supseteq \text{GBC}^- + \text{ETR}$ we can unroll (M, \mathcal{X}) into W . Taking isomorphic copies if necessary we may assume without loss that $M = H_\kappa^W$ where κ is the largest cardinal in W . And because T proves the existence of $\text{Hyp}(V)$ we get that $\text{Hyp}(M)^W \in W$. In general, $A = \text{Hyp}(M)^W$ may be ill-founded, for example if M is ill-founded. Let \mathfrak{C} be the admissible cover of A . Now consider the $\mathcal{L}_{\mathfrak{C}}$ theory S axiomatized by the following.

- Every theorem T proves about the unrolling;¹⁴
- $M = H_\kappa^W$ is an H_α -initial segment of the universe. That is, this statement asserts that if x is hereditarily of cardinality $< \kappa$ then $x \in M$; and¹⁵
- κ is the largest cardinal.

This S can be expressed as a conjunction of a countable set (in A) of $\mathcal{L}_{\omega, \omega}$ -formulae with two $\mathcal{L}_{\mathfrak{C}}$ -formulae. So it is a single $\mathcal{L}_{\mathfrak{C}}$ -sentence and hence is Σ_1 -definable over \mathfrak{C} . This puts us in a position to apply the generalization of Friedman’s theorem, since A is countable. That is, there is $N \models S$ which is ill-founded at A . Put differently, there is a descending sequence of ordinals in N co-initial in $\text{Ord}^N \setminus \text{Ord}^A$. Let $\mathcal{Y} = \{Y \in N : N \models Y \subseteq M\}$.¹⁶ Then $(M, \mathcal{Y}) \models T$, by construction.

To motivate the following, suppose for a moment that (M, \mathcal{X}) is a β -model. Then Ord^A really is an ordinal. Thus, \mathcal{Y} has no element with ordertype Ord^A because otherwise Ord^A would be in the well-founded part of N , contrary to the construction. So $\mathcal{X} \not\subseteq \mathcal{Y}$, as desired.

But in general (M, \mathcal{X}) need not be a β -model, and it may even be that $\text{Hyp}(V)^{(M, \mathcal{X})}$ is ill-founded. So the above argument cannot work. Nevertheless, it provides the right idea. Fix a membership code $\Upsilon \in \mathcal{X}$ which represents Ord^A in the unrolling. Suppose towards a contradiction that $\mathcal{X} \subseteq \mathcal{Y}$.

In particular this implies that $\Upsilon \in \mathcal{Y}$. Because \mathcal{X} and \mathcal{Y} must agree about which of their common classes are well-founded, (M, \mathcal{Y}) thinks that Υ is a membership code for an ordinal. Let $v \in N$ be that ordinal.

¹⁴If you think of the special case where T is KMCC , then this theory is ZFC_1^- . In general, this theory is in M because $T \in M$ and it is computable from T .

¹⁵This can be expressed as a single $\mathcal{L}_{\mathfrak{C}}$ -sentence because $M \in A$.

¹⁶If one wants to be picky, since we officially only work with models whose second-order part consists of subsets of the first-order part, we actually take an isomorphic copy so that elements of \mathcal{Y} are subsets of M .




Figure 4.3: Friedman's theorem gives $N \models S$ with $V_\kappa^N = M = V_\kappa^W$ and N ill-founded at A .

We also have that every initial segment of Υ is in \mathcal{Y} .¹⁷ In W we have that every initial segment of Υ is isomorphic to an ordinal in A . Because of the assumption that $\mathcal{X} \subseteq \mathcal{Y}$ the same isomorphisms exist in N . So for every ordinal $\gamma \in A$ we have $N \models \gamma < v$.

If $N \models \gamma < v$ then γ is isomorphic to an initial segment of Υ . But W sees that initial segments of Υ represent ordinals in A . So N must see the same and thus $\gamma \in A$. The upshot of all this is that Ord^A is topped in N , namely by v . This contradicts the construction of N . So our assumption that $\mathcal{X} \subseteq \mathcal{Y}$ must be false, completing the argument. We have found a T -realization for M which does not contain \mathcal{X} . \square

Remark 4.21. The attentive reader may worry about what happens if M is ω -nonstandard. She is right to worry! There is a subtlety that must be addressed. Namely, if M is ω -nonstandard then no infinite first-order theory T can be in M ; otherwise, the standard cut would be definable as the supremum of the ranks of elements of T . So as written the master lemma does not apply to ω -nonstandard models.

Nevertheless, there is a version which does apply to ω -nonstandard models. While it does not make sense to ask for T to be an element of M it is sensible to ask that T is coded in M , that is whether there is $t \in V_\omega^M \setminus \omega$ so that $t \cap V_\omega = T$. In particular, this always happens if T is computable; run the Turing machine which enumerates T in the nonstandard model of arithmetic coming from M and cut it off at some nonstandard level to get t .

In the ω -nonstandard case replace the assumption that $T \in M$ with the assumption that T is coded in M . Then the above proof works. In particular, if (strong enough) T is computable then no countable T -realizable M has a least T -realization.

It remains to see that the master lemma yields the nonexistence of least transitive models of strong second-order set theories. Because these theories are all computably axiomatizable it is immediate that they appear as elements of any transitive model. As such, the only thing we need to show is that these theories prove the existence of $\text{Hyp}(V)$. As a warm-up let us prove that KMCC proves $\text{Hyp}(V)$ exists. This follows from a stronger statement.

Proposition 4.22. *Let $\varphi(x, \bar{y})$ be a first-order formula in the language of set theory. Then ZFC^- proves that for every \bar{b} there is an ordinal α so that for all $a \in L_\alpha(\bar{b})$, we have $L_\alpha(\bar{b}) \models \varphi(a, \bar{b})$ if and only if $\varphi(a, \bar{b})^{L(\bar{b})}$.*

Proof. This is the standard argument for reflection along the L -hierarchy. Namely, pick an ordinal α_0 . Given α_n let α_{n+1} be the least ordinal so that $L_{\alpha_{n+1}}(\bar{b})$ is closed under witnesses for existential subformulae of φ with parameters from $L_{\alpha_n}(\bar{b})$. Note that we use Collection to find α_{n+1} , since we need to collect witnesses into a single set. Then, if $\alpha = \sup_n \alpha_n$ we have caught our tail and $L_\alpha(\bar{b}) \models \varphi(a, \bar{b})$ if and only if $\varphi(a, \bar{b})$ is true in $L(\bar{b})$. \square

Corollary 4.23. *KMCC proves that $\text{Hyp}(V)$ exists.*

¹⁷Recall from chapter 2 that a membership code for an ordinal γ is a class well-order of ordertype $\gamma + 1$, so it makes sense to talk about initial segments of Υ .

Proof. Work in the unrolling, which satisfies ZFC_1^- . Then, since the axioms of KP are of bounded complexity, there is $\alpha > \kappa$, where κ is the largest cardinal, so that $L_\alpha(V_\kappa) \models \text{KP}$. The least such α gives $\text{Hyp}(V_\kappa)$. But V_κ is the V of the original KMCC model. So there is a membership code for $\text{Hyp}(V)$. \square

But we need much less than KMCC to get the existence of $\text{Hyp}(V)$.

Lemma 4.24. *The theory $\text{GBC}^- + \Pi_1^1\text{-CA}$ proves the existence of $\text{Hyp}(V)$. Consequently any $T \supseteq \text{GBC}^- + \Pi_1^1\text{-CA}$ proves the existence of $\text{Hyp}(V)$.*

See also [ABF, theorem 64] where the same result appears. (They state the result in terms of $\text{GBC} + \Pi_1^1\text{-CA}$ instead of $\text{GBC}^- + \Pi_1^1\text{-CA}$, but nowhere does their proof use Powerset in the first-order part.)

Proof. Work with $(M, \mathcal{X}) \models \text{GBC}^- + \Pi_1^1\text{-CA}$ and consider $W = \mathfrak{L}(M, G)$ the (M, G) -constructible unrolling of (M, \mathcal{X}) , for some $G \in \mathcal{X}$. Then, by results of section 2.3 we have that $W = \mathfrak{L} \models \text{ZFC}^-(1) + V = L(M, G)$.¹⁸ Moreover, there is $\kappa \in W$ so that $M = (H_\kappa)^W$. We want to find $\alpha \in W$ so that $L_\alpha(M) \models \text{KP}$. This reduces down to proving an instance of reflection along the $L(M)$ -hierarchy, which I give here so the reader can see it can be carried out in the weak theory in which we are currently working.

Every $L_\alpha(M)$ satisfies Σ_0 -Separation, so the work is in getting Σ_0 -Collection. We will see that there are unboundedly many α so that $L_\alpha(M) \models \Sigma_0$ -Collection. Let v be the formula giving truth for Σ_0 -formulae. Then v is Σ_1 . It is convenient here to assume (without loss) that v has four free variables, so that $v(\varphi, x, y, p)$ asserts that $\varphi(x, y, p)$ holds for a Σ_0 -formula φ . To show that $L_\alpha(M) \models \Sigma_0$ -Collection it suffices to prove the instance of Σ_0 -Collection for v .

Fix arbitrary α_0 . By Σ_1 -Collection find α_1 the least ordinal $> \alpha_0$ so that if $x, p \in L_{\alpha_0}(M)$ and φ is a formula then there is $y \in L_{\alpha_1}(M)$ so that $v(\varphi, x, y, p)$. Now repeat the process: given α_n let α_{n+1} be the least ordinal $> \alpha_n$ so that if $x, p \in L_{\alpha_n}(M)$ and φ is a formula then there is $y \in L_{\alpha_{n+1}}(M)$ so that $v(\varphi, x, y, p)$. Finally, set $\alpha = \sup \alpha_n$, again using an instance of Σ_1 -Collection. Then $L_\alpha(M) \models \Sigma_0$ -Collection. Since α_0 was arbitrary, this proves there are unboundedly many such α .

Now take $\alpha > \kappa$ least so that $L_\alpha(M) \models \text{KP}$. Then $L_\alpha(M) = \text{Hyp}(M)$, so $(M, \mathcal{X}) \models \text{Hyp}(V)$ exists, as desired. \square

As a corollary we get the negative part of theorem 4.3.

Corollary 4.25. *There is not a least transitive model of KM, nor of KM^- . For $k \geq 1$ there is not a least transitive model of $\text{GBC} + \Pi_k^1\text{-CA}$, nor of $\text{GBC}^- + \Pi_k^1\text{-CA}$. Moreover, the same holds for any computably axiomatizable extensions of these theories.*

Proof. I will state the proof in terms of KM. The same argument goes through, *mutatis mutandis*, for the other theories.

Suppose for sake of a contradiction that there is a least transitive model of KM. Then it must be some countable $(M, \mathcal{X}) \models \text{KMCC}$. By the master lemma there is a KMCC-realization \mathcal{Y} for M so that $\mathcal{X} \not\subseteq \mathcal{Y}$. So (M, \mathcal{X}) is not actually least, a contradiction. \square

This result holds for more than just computably axiomatizable extensions. What is needed about the extension T is that T is an element of any transitive model of T . In particular, the first-order part of T can be complete. If $\text{Hyp}(V)$ exists then the truth predicate for the first-order part must exist, since $\text{Hyp}(V) \models \text{KP}$ and KP proves the existence of truth predicates for set-sized structures and so any model of a second-order set theory $T \subseteq \text{GBC}^-$ which proves the existence of $\text{Hyp}(V)$ must contain its (first-order) theory as an element, since it can be obtained by restricting the truth predicate to sentences.

We also get that countable models do not have least realizations for strong theories.

Corollary 4.26. *Let $M \models \text{ZFC}^-$ be countable and T an extension of $\text{GBC}^- + \Pi_1^1\text{-CA}$ with $T \in M$.¹⁹ Then M does not have a least T -realization.*

¹⁸Recall that $\text{ZFC}^-(1)$ is axiomatized by the axioms of ZFC^- but with Separation and Collection restricted to Σ_1 -formulae.

¹⁹Or, if M is ω -nonstandard, with T coded in M . Cf. remark 4.21.

Proof. Suppose towards a contradiction that \mathcal{X} is the least T -realization for M . By the master lemma there is $\mathcal{Y} \not\subseteq \mathcal{X}$ a T -realization for M . But this contradicts the leastness of \mathcal{X} . \square

In personal communication, Ali Enayat pointed out to me an alternative argument that countable models do not have least KM-realizations. I reproduce his argument, which goes by way of an old theorem by Barwise, here.

Theorem 4.27. *Let $M \models \text{ZFC}$ be a countable transitive model. Then M does not have a least KM-realization.*

From this it immediately follows that KM does not have a least transitive model.

Proof (Enayat): If M is not KM-realizable, then the conclusion is trivial. So work in the case where M is KM-realizable. Recall the following theorem.

Theorem (Barwise, theorem IV.1.1 of [Bar75]). *Let U be a countable structure. Let $A = \text{Hyp}(U)$. Let T be an \mathcal{L}_A -theory which is Σ_1 -definable over A and which has a model of the form $\mathfrak{B} = (B, U; E, \dots)$, where E is a binary relation.²⁰ Suppose $S \subseteq U$ has the property that for every such model there is $b \in B$ so that $S = b_E = \{x \in \mathfrak{B} : x E b\}$. Then $S \in \text{Hyp}(U)$.*

Suppose we have a model of the form $(B, M; E)$. Consider the theory T asserting that M forms the first-order part of and B forms the second-order part of a model of KM with membership relation E . This is a computable $\mathcal{L}_{\omega, \omega}$ -theory, so in particular it is Σ_1 -definable over $\text{Hyp}(M)$. Because M is KM-realizable there is a model of T of form $\mathfrak{B} = (B, M; E)$. Now suppose that $S \subseteq M$ is in every KM-realization for M . Then if $\mathfrak{B} = (B, M; E) \models T$ we can find $b \in B$ so that $S = b_E$. So by Barwise's theorem we get that $S \in \text{Hyp}(M)$.

This yields the following lemma, which is of independent interest.

Lemma 4.28. *Let $M \models \text{ZFC}$ be countable, transitive, and KM-realizable. Then the intersection of all the KM-realizations for M is $\text{Hyp}(M) \cap \mathcal{P}(M)$.*

Proof. Let \mathcal{X} be the intersection of all the KM-realizations of M . We have seen that $\mathcal{X} \subseteq \text{Hyp}(M) \cap \mathcal{P}(M)$. For the other direction, take $A \in \text{Hyp}(M) \cap \mathcal{P}(M)$. Then there is $\gamma < \text{Ord}^{\text{Hyp}(M)}$ so that $A \in L_\gamma(M)$. To conclude that $A \in \mathcal{X}$ it is enough to know that every KM-realization for M unrolls to a structure which is well-founded up to γ . To see that, take $(M, \mathcal{Y}) \models \text{KM}$. By throwing out classes if necessary, assume without loss that $(M, \mathcal{Y}) \models \text{KMCC}$. Let $W \models \text{ZFC}_1^-$ be the unrolling of (M, \mathcal{Y}) . It follows from a result of H. Friedman [Fri73, theorem 3.1] that W , being well-founded up to at least Ord^M , must have that $\text{Ord}^{\text{wfp}(W)} \geq \text{Ord}^M$. But it could be that $\text{Hyp}(M)$ is taller than the next admissible ordinal above Ord^M , so we need a small argument.

Take $\gamma \in \text{Hyp}(M)$ an ordinal. Then there is $\Gamma \subseteq M$ in $\text{Hyp}(M)$ which is isomorphic to (γ, \in) . Because $\Gamma \in \text{Hyp}(M)$ it is Δ_1^1 -definable over M —see [Bar75, corollary IV.3.4]. So $\Gamma \in \mathcal{Y}$, because (M, \mathcal{Y}) satisfies Δ_1^1 -Comprehension. And since Γ is seen to be well-founded from the external universe, it must be that (M, \mathcal{Y}) agrees that Γ is a well-order. So the unrolling of (M, \mathcal{Y}) , namely W , must contain γ . So $\text{Ord}^{\text{wfp}(W)} > \gamma$, completing the argument. \square

Now suppose that $\mathcal{X} = \mathcal{P}(M) \cap \text{Hyp}(M)$ is KM-realization for M . Then, because \mathcal{X} unrolls to $\text{Hyp}(M)$ and KM proves that $\text{Hyp}(V)$ exists we get that $\text{Hyp}(M) \in \text{Hyp}(M)$, a contradiction. So M cannot have a least KM-realization. \square

The full strength of KM is not needed here. We used two facts about KM: first, that KM proves that $\text{Hyp}(V)$ exists; and second, that KM proves Δ_1^1 -Comprehension. So the same argument goes through for weaker theories theories, in particular any theory extending $\text{GBC}^- + \Pi_1^1\text{-CA}$.

Let me turn now to β -models, in which context we do get least models.

Theorem 4.29 (Folklore). *There is a least β -model of KM, if there is any β -model of KM.*

²⁰Recall from section 4.2 that $(B, U; R_0, R_1, \dots)$ has sets B , urelements U , and additional relations R_0, R_1, \dots

Proof. First, note that it is equivalent to ask for a least β -model of KMCC, by theorem 2.47 from chapter 2. Next, observe that if $(M, \mathcal{X}) \subseteq (N, \mathcal{Y})$ are β -models of KMCC, then their unrollings are transitive and the unrolling of (M, \mathcal{X}) must be contained in the unrolling of (N, \mathcal{Y}) . So we just have to see that there is a least transitive model of ZFC_1^- . Because there is a β -model of KM, there is some transitive model of ZFC_1^- .

Take $M \models \text{ZFC}_1^-$ transitive with largest cardinal κ . Then $L^M \models \text{ZFC}^- + \kappa$ is inaccessible. However, it could be that L^M satisfies Powerset. (Imagine if M were obtained by class forcing over a model of $\text{ZFC} + V = L +$ “there is an inaccessible cardinal” to collapse all cardinals above the first inaccessible.) Nevertheless, there is some ordinal $\alpha \in M$ so that $L_\alpha \models \text{ZFC}_1^-$. So while ZFC_1^- is not absolute to L , every transitive model of ZFC_1^- contains an L_α which is a model of ZFC_1^- .

Therefore, if α is least so that $L_\alpha \models \text{ZFC}_1^-$ then L_α is the least transitive model of ZFC_1^- , as desired. \square

This proof generalizes to a fixed first-order part. First we need a definition.

Definition 4.30. Let T be a second-order set theory and $M \models \text{ZFC}^-$. Then $\mathcal{X} \subseteq \mathcal{P}(M)$ is a β - T -realization for M if $(M, \mathcal{X}) \models T$ is a β -model. If such \mathcal{X} exists, then M is β - T -realizable. The least β - T -realization for M , if it exists, is the unique β - T -realization for M which is contained in every β - T -realization.

Corollary 4.31. Let M be a β -KM-realizable model of set theory with a definable global well-order. Then M has a least β -KM-realization.

Proof. Because M is β -KM-realizable, there is a model W of ZFC_1^- with largest cardinal κ so that $M = V_\kappa^W$. Consider α least so that $L_\alpha(M) \models \text{ZFC}_1^-$ and $M = V_\kappa^{L_\alpha(M)}$. Such exists by an argument as in the proof of the previous theorem, using the fact that M has a definable global well-order to get that $L_\alpha(M)$ satisfies Choice. Then the cutting off (M, \mathcal{X}) for $L_\alpha(M)$ gives the least β -KM-realization for M . \square

Essentially the same argument, using tools from chapter 2, gives that there is a least β -model of $\text{GBC} + \Pi_k^1\text{-CA}$.

Theorem 4.32. Let $k \geq 1$. There is a least β -model of $\text{GBC} + \Pi_k^1\text{-CA}$, if there is any β -model of $\text{GBC} + \Pi_k^1\text{-CA}$.

Proof. Again, this reduces to showing that there is a least transitive model of $\text{ZFC}_1^-(k)$, using that β -models of $\text{GBC} + \Pi_k^1\text{-CA} + \Sigma_k^1\text{-CC}$ unroll to transitive models of $\text{ZFC}_1^-(k)$.

If $M \models \text{ZFC}_1^-(k)$ has largest cardinal κ then $L^M \models \text{ZFC}_1^-(k) + \kappa$ is inaccessible. So there is $\alpha \in M$ so that $L_\alpha \models \text{ZFC}_1^-(k)$ has largest cardinal κ . Thus, if α is least so that $L_\alpha \models \text{ZFC}_1^-(k)$ then L_α is the least transitive model of $\text{ZFC}_1^-(k)$. \square

Corollary 4.33. Fix $k \geq 1$ and let M be a β - $(\text{GBC} + \Pi_k^1\text{-CA})$ -realizable model of set theory with a definable global well-order. Then M has a least β - $(\text{GBC} + \Pi_k^1\text{-CA})$ -realization.

Proof. Essentially the same as the proof of corollary 4.31. \square

This is not the end of the story. Does corollary 4.31 give an exact characterization of when M has a least β -KM-realization? (And in light of corollary 4.33 the same question can be asked about $\text{GBC} + \Pi_k^1\text{-CA}$, for $k \geq 1$, instead of KM.)

Question 4.34. Suppose $M \models \text{ZFC}$ is β -KM-realizable but does not have a definable global well-order. Can we conclude that M does not have a least β -KM-realization?

It is Global Choice that is the possible culprit here. If we drop that from the axioms then we do always get least realizations.

Proposition 4.35. Let $\text{KM}^{-\text{GC}}$ denote KM with Choice for sets but without Global Choice. Suppose that $M \models \text{ZFC}$ is β - $\text{KM}^{-\text{GC}}$ -realizable. Then M has a least β - $\text{KM}^{-\text{GC}}$ -realization.

Proof sketch. Similar to the proof of corollary 4.31. The least β - $\text{KM}^{-\text{GC}}$ -realization for M is the cut off model obtained from $L_\alpha(M)$ where α is least so that $L_\alpha(M) \models \text{ZF}_1^-$ and the largest cardinal of $L_\alpha(M)$ is Ord^M . \square

A similar fact holds for $\text{GBC} + \Pi_k^1\text{-CA}$, i.e. $\text{GBC} + \Pi_k^1\text{-CA}$ but where we drop Global Choice.

To finish off this section, let us see that there are minimal but non-least KM-realizations. That is, there are $M \models \text{ZFC}$ which have a KM-realization \mathcal{X} so that there is no KM-realization \mathcal{Y} for M which is strictly contained inside \mathcal{X} . But by theorem 4.5 \mathcal{X} cannot be least.

Observation 4.36. *Suppose there is a β -model of KM. Then, there are $M \models \text{ZFC}$ which have minimal but non-least KM-realizations.*

Proof. Let countable $M \models \text{ZFC}$ have a definable global well-order and be β -KM-realizable. By corollary 4.31 it has a least β -KM-realization, call it \mathcal{X} . Because any KM-realization $\mathcal{Y} \subseteq \mathcal{X}$ must also give a β -model and thus $\mathcal{Y} = \mathcal{X}$, this \mathcal{X} is a minimal KM-realization. But as we saw earlier in this section, M does not have a least KM-realization. \square

Left open is the question of when, if ever, other countable $M \models \text{ZFC}$ have minimal but non-least KM-realizations.

Question 4.37. *Is there countable $M \models \text{ZFC}$ which does not have a minimal KM-realization?*

The same question can be asked for $\text{GBC} + \Pi_k^1\text{-CA}$ instead of KM.

4.4 Medium theories

I turn now to the theories of medium strength, namely $\text{GBC} + \text{ETR}$ and its variants. The main results of this section are that $\text{GBC} + \text{ETR}$ has a least β -model and that for nice enough choice of Γ that $\text{GBC} + \text{ETR}_\Gamma$ has a least transitive and a least β -model. (See the discussion around the relevant theorems in this section for what “nice enough” means.) The major question left open is whether there is a least transitive model of ETR.

As a starting-off point, let us see that some models have least β -($\text{GBC} + \text{ETR}$)-realizations.

Theorem 4.38. *Let $M \models \text{ZFC}$ be a transitive model with a definable global well-order. Then if M has a β -($\text{GBC} + \text{ETR}$)-realization it has a least β -($\text{GBC} + \text{ETR}$)-realization \mathcal{X} . Moreover, \mathcal{X} is also the least ($\text{GBC} + \text{ETR}$)-realization of M .*

Proof. Fix \mathcal{Y} a β -($\text{GBC} + \text{ETR}$)-realization for M . The strategy is to define \mathcal{X} , which will be the least β -($\text{GBC} + \text{ETR}$)-realization for M contained inside \mathcal{Y} . We will then see that in fact \mathcal{X} is contained inside any ($\text{GBC} + \text{ETR}$)-realization for M . In particular, if we started with a different β -realization we would define the same \mathcal{X} .

We define \mathcal{X} in ω many steps. First, let $\mathcal{X}_0 = \text{Def}(M)$. Then $(M, \mathcal{X}_0) \models \text{GBC}$ because M has a definable global well-order. Now given $\mathcal{X}_n \subseteq \mathcal{Y}$ define \mathcal{X}_{n+1} to consist of all classes in \mathcal{Y} which are definable from $\text{Tr}_\Gamma(A)$ for some $\Gamma, A \in \mathcal{X}_n$. Formally,

$$\mathcal{X}_{n+1} = \bigcup \{ \text{Def}(M; \text{Tr}_\Gamma(A)) : A, \Gamma \in \mathcal{X}_n \text{ and } \Gamma \text{ is a well-order} \}.$$

Because (M, \mathcal{Y}) is a β -model, it is correct about which Γ 's are well-orders. So we could equivalently ask in the definition of \mathcal{X}_{n+1} that $(M, \mathcal{Y}) \models \Gamma$ is a well-order. Then $\mathcal{X}_{n+1} \subseteq \mathcal{Y}$ is a GBC -realization for M . Finally, set $\mathcal{X} = \bigcup_n \mathcal{X}_n$. It is clear that $\mathcal{X} \subseteq \mathcal{Y}$.

Let us check that $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}$. It satisfies GBC because \mathcal{X} is the union of an increasing chain of GBC -realizations for M . To see that it satisfies Elementary Transfinite Recursion, pick $A, \Gamma \in \mathcal{X}$ where $(M, \mathcal{X}) \models \Gamma$ is a well-order. Observe that \mathcal{Y} agrees with \mathcal{X} that Γ is a well-order. Since $A, \Gamma \in \mathcal{X}_n$ for some n this means that $\text{Tr}_\Gamma(A) \in \mathcal{X}_{n+1} \subseteq \mathcal{X}$.

Finally, let us see that \mathcal{X} is contained in any ($\text{GBC} + \text{ETR}$)-realization \mathcal{Z} , which will establish that \mathcal{X} is both the least β -($\text{GBC} + \text{ETR}$)-realization for M and the least ($\text{GBC} + \text{ETR}$)-realization for M . Clearly, $\mathcal{X}_0 = \text{Def}(M)$ is contained inside \mathcal{Z} . We continue upward inductively. Having already seen that $\mathcal{X}_n \subseteq \mathcal{Z}$, consider $\Gamma, A \in \mathcal{X}_n$. Then all three of \mathcal{Y} , \mathcal{X}_n , and \mathcal{Z} must agree on which classes are well-founded and which

classes are iterated truth predicates. And since Γ really is well-founded, externally we can see that there is only one option for what class is the Γ -iterated truth predicate relative to A . So since $(M, \mathcal{Z}) \models \text{ETR}$ we get that $(\text{Tr}_\Gamma(A))^{(M, \mathcal{Z})} = \text{Tr}_\Gamma(A) \in \mathcal{X}_{n+1} \cap \mathcal{Z}$. And since \mathcal{Z} is closed under first-order definability, any class definable from $\text{Tr}_\Gamma(A)$ must be in \mathcal{Z} . So $\mathcal{X}_{n+1} \subseteq \mathcal{Z}$. This holds for all n , so $\mathcal{X} \subseteq \mathcal{Z}$, as desired. \square

The proof did not use Powerset. So we get a version for theories without Powerset. And the only place we used that M has a definable global well-order was to get Global Choice in \mathcal{X} . So we also get a version for theories without Global Choice. The following corollary encapsulates both results.

Corollary 4.39. *Let $M \models \text{ZFC}^-$ be a transitive model. Suppose M is β -($\text{GBC}^- + \text{ETR}$)-realizable. Then M has a least β -($\text{GBC}^- + \text{ETR}$)-realization. If M moreover has a definable global well-order then M has a least β -($\text{GBC}^- + \text{ETR}$)-realization. \square*

Observe that although \mathcal{X}_n is always a coded V -submodel of \mathcal{Y} , in general \mathcal{X} need not be coded in \mathcal{Y} . In particular, this will happen when $\mathcal{Y} = \mathcal{X}$.

We get a version of this result for non- β -models. In this broader context we cannot ensure that different \mathcal{Y} 's will define the same \mathcal{X} . But any V -submodel of (M, \mathcal{Y}) will have to agree with (M, \mathcal{Y}) as to what is a well-order and whether a class is $\text{Tr}_\Gamma(A)$. So a similar argument yields a local leastness result.

Theorem 4.40. *Let $M \models \text{ZFC}$ be ($\text{GBC} + \text{ETR}$)-realizable. Then M has a basis of minimal ($\text{GBC} + \text{ETR}$)-realizations, where amalgamable ($\text{GBC} + \text{ETR}$)-realizations²¹ sit above the same basis element. That is, there is a set $\{\mathcal{B}_i : i \in I\}$ of ($\text{GBC} + \text{ETR}$)-realizations for M satisfying the following.*

1. *Elements of the basis are pairwise non-amalgamable;*
2. *If \mathcal{Y} is any ($\text{GBC} + \text{ETR}$)-realization for M then there is a unique basis element \mathcal{B} so that $\mathcal{Y} \supseteq \mathcal{B}$; and*
3. *If \mathcal{X} and \mathcal{Y} are amalgamable ($\text{GBC} + \text{ETR}$)-realizations for M then they sit above the same \mathcal{B} .*

See figure 4.4 for a picture of the ($\text{GBC} + \text{ETR}$)-realizations for M .

In case M has a definable global well-order we get a basis of minimal ($\text{GBC} + \text{ETR}$)-realizations, since every GBC -realization for M must contain the definable global well-order and thereby satisfy Global Choice.

Proof. Fix \mathcal{Y} a ($\text{GBC} + \text{ETR}$)-realization for M . We define the basis element \mathcal{B} below \mathcal{Y} similar to how we defined \mathcal{X} in the proof of theorem 4.38. Start with $\mathcal{B}_0 = \text{Def}(M)$. Clearly $\mathcal{B}_0 \subseteq \mathcal{Y}$ is a GBC -realization for M . Given $\mathcal{B}_n \subseteq \mathcal{Y}$ a GBC -realization for M we set

$$\mathcal{B}_{n+1} = \bigcup \left\{ \text{Def} \left(M; (\text{Tr}_\Gamma(A))^{(M, \mathcal{Y})} \right) : A, \Gamma \in \mathcal{B}_n \text{ and } (M, \mathcal{Y}) \models \Gamma \text{ is a well-order} \right\}.$$

Some illuminating remarks are in order. First, because $\mathcal{B}_n \subseteq \mathcal{Y}$ and $(M, \mathcal{Y}) \models \text{ETR}$ we get that for $A, \Gamma \in \mathcal{B}_n$ there is a unique class in \mathcal{Y} which (M, \mathcal{Y}) thinks is $\text{Tr}_\Gamma(A)$. So \mathcal{B}_{n+1} is well-defined. And since \mathcal{Y} is closed under first-order definability this moreover shows that $\mathcal{B}_{n+1} \subseteq \mathcal{Y}$.

Next, let me emphasize that Def here is the *external* Def operator. This only makes a difference in case M is an ω -model. In this case, none of the \mathcal{B}_n will be coded in \mathcal{Y} . Nevertheless, we still get that $\mathcal{B}_n \subseteq \mathcal{Y}$, so that \mathcal{B}_{n+1} is well-defined.

Third, let us check that $(M, \mathcal{B}_{n+1}) \models \text{GBC}$. I will be more detailed than in the proof of theorem 4.38 to reassure the reader who is worried things may go wrong in an ω -nonstandard model. Both Class Extensionality and Class Replacement are immediate. To see Elementary Comprehension we want to see that \mathcal{B}_{n+1} is closed under first-order definability. It suffices to check the case where we define a class from two class parameters, so consider $X, Y \in \mathcal{B}_{n+1}$. Then, by construction, there are $\Gamma, \Delta, A, B \in \mathcal{B}_n$ so that X is definable from $(\text{Tr}_\Gamma(A))^{(M, \mathcal{Y})}$ and Y is definable from $(\text{Tr}_\Delta(B))^{(M, \mathcal{Y})}$. Then any class definable from X and Y must be definable from $(\text{Tr}_{\max\{\Gamma, \Delta\}}(A \oplus B))^{(M, \mathcal{Y})} \in \mathcal{B}_n$, where $A \oplus B = A \times \{0\} \cup B \times \{1\}$. So any class definable from X and Y is in \mathcal{B}_{n+1} .

²¹Two T -realizations \mathcal{X} and \mathcal{Y} for M are *amalgamable* if there is a GBC^- -realization \mathcal{Z} for M so that \mathcal{X} and \mathcal{Y} are both subsets of \mathcal{Z} .




Figure 4.4: The $(\text{Gbc} + \text{ETR})$ -realizations for M form a disjoint collection of cones, each one with a basis element at the bottom.

Finally, set $\mathcal{B} = \bigcup_n \mathcal{B}_n$. It is immediate that $\mathcal{B} \subseteq \mathcal{Y}$. Let us see that $(M, \mathcal{B}) \models \text{Gbc} + \text{ETR}$. Fix $\Gamma, A \in \mathcal{B}$ so that $(M, \mathcal{B}) \models \Gamma$ is a well-order. Then $\Gamma, A \in \mathcal{B}_n$ for some n . Because (M, \mathcal{B}_n) is a V -submodel of (M, \mathcal{B}) which in turn is a V -submodel of (M, \mathcal{Y}) they all agree as to whether Γ is a well-order and whether a class is $\text{Tr}_\Gamma(A)$. So

$$(\text{Tr}_\Gamma(A))^{(M, \mathcal{B})} = (\text{Tr}_\Gamma(A))^{(M, \mathcal{Y})} \in \mathcal{B}_{n+1} \subseteq \mathcal{B}.$$

So $(M, \mathcal{B}) \models \text{ETR}$.

The proof will be finished once we see that any \mathcal{X} which is amalgamable with \mathcal{Y} defines the same \mathcal{B} . To see this, take \mathcal{Z} a Gbc -realization for M which contains both \mathcal{X} and \mathcal{Y} . Then (M, \mathcal{Y}) and (M, \mathcal{Z}) must agree whether a class is a well-order and whether a class is $\text{Tr}_\Gamma(A)$ and the same holds for (M, \mathcal{X}) and (M, \mathcal{Z}) , so in fact all three agree. Clearly \mathcal{B}_0 is the same whether defined using \mathcal{Y} or \mathcal{X} . And inductively upward they must agree on \mathcal{B}_n because they agree as to whether a class is $\text{Tr}_\Gamma(A)$. So no matter whether we start with \mathcal{X} or \mathcal{Y} we define the same basis element \mathcal{B} . \square

It follows from theorem 4.38 that if M has a β - $(\text{Gbc} + \text{ETR})$ -realization then there is only one basis element in the poset of $(\text{Gbc} + \text{ETR})$ -realizations for M . Does this hold in general?

Question 4.41. *Is there a $(\text{Gbc} + \text{ETR})$ -realizable model M so that the basis for the $(\text{Gbc} + \text{ETR})$ -realizations for M has more than one element? Or, asked in the negative, is it true that every $(\text{Gbc} + \text{ETR})$ -realizable model has a least $(\text{Gbc} + \text{ETR})$ -realization?*

Let me detour to discuss fragments of ETR . Essentially the same argument as in theorem 4.38 gives that $\text{Gbc} + \text{ETR}_\Gamma$ will have least β -realizations. For example, to show that a β - $(\text{Gbc} + \text{ETR}_\Gamma)$ -realizable model M with a definable global well-order has a least β - $(\text{Gbc} + \text{ETR}_\Gamma)$ -realization we define \mathcal{X} in a similar manner. Set $\mathcal{X}_0 = \text{Def}(M; \Gamma)$. The definition for \mathcal{X}_{n+1} is then

$$\mathcal{X}_{n+1} = \bigcup \{ \text{Def}(M; \text{Tr}_\Gamma(A)) : A \in \mathcal{X}_n \}$$

where \mathcal{Y} is some fixed-in-advance β - $(\text{Gbc} + \text{ETR}_\Gamma)$ -realization for M . Observe that Γ here really is well-founded, because a β -model thinks it is well-founded, so $\text{Tr}_\Gamma(A)$ is externally seen to be unique. Then

$\mathcal{X} = \bigcup_n \mathcal{X}_n$ will be the least β -(GBC+ETR $_\Gamma$)-realization for M . This gives us the following results, analogous to the above results about ETR.

Theorem 4.42. *Suppose $(M, \mathcal{Y}) \models \text{GBC} + \text{ETR}_\Gamma$ is an β -model with a definable global well-order and $\omega^\omega \leq \Gamma \in \mathcal{Y}$. Then M has a least β -(GBC+ETR $_\Gamma$)-realization \mathcal{X} . Moreover, \mathcal{X} is also the least (GBC+ETR $_\Gamma$)-realization for M . \square*

Remark 4.43. The purpose of requiring $\Gamma \geq \omega^\omega$ is that this ensures ETR $_\Gamma$ is equivalent to the existence of Γ -iterated truth predicates relative to any class. The same applies to later results about ETR $_\Gamma$, but I will suppress making this comment every time.

Since we are concerned only with a fixed Γ , we can get a least (GBC + ETR $_\Gamma$)-realization even if our $M \models \text{ZFC}$ has realizations which are wrong about well-foundedness. All that matters is whether they are correct about Γ being well-founded.

Theorem 4.44. *Consider an ω -model $(M, \mathcal{Y}) \models \text{GBC}$ and $\Gamma \in \mathcal{Y}$ so that $\Gamma \geq \omega^\omega$ really is well-founded, as seen externally. Suppose M has a definable global well-order and $(M, \mathcal{Y}) \models \text{ETR}_\Gamma$. Then M has a least (GBC + ETR $_\Gamma$)-realization. \square*

If M is ill-founded then this works for Γ of length in the well-founded part of M .²² If M is transitive then we can go up to Ord^M , and even longer. For transitive M we always get that $\text{Ord}^M + \text{Ord}^M$, $\text{Ord}^M \cdot \text{Ord}^M$, and so on are well-founded.

And like before, for non- β -models we get a local leastness result, even if (M, \mathcal{Y}) is wrong about Γ being well-founded.

Corollary 4.45. *Suppose $(M, \mathcal{Y}) \models \text{GBC} + \text{ETR}_\Gamma$ has a definable global well-order and $(M, \mathcal{Y}) \models \omega^\omega \leq \Gamma$. Then M has an (GBC + ETR $_\Gamma$)-realization which is least below \mathcal{Y} . \square*

We are now ready to see that GBC + ETR has a least β -model.

Theorem 4.46. *There is a least β -model of GBC + ETR, if there is any β -model of GBC + ETR.*

Proof. The least β -model of GBC + ETR will be the least β -(GBC + ETR)-realizable L_α along with its least β -(GBC+ETR)-realization. First though we have to know that if M is β -(GBC+ETR)-realizable then so is L^M . We saw in chapter 3 that M being (GBC + ETR)-realizable implies that L^M is also (GBC + ETR)-realizable. More specifically, if $(M, \mathcal{X}) \models \text{GBC} + \text{ETR}$ then there is $\mathcal{Y} \subseteq \mathcal{X} \cap \mathcal{P}(L^M)$ so that $(L^M, \mathcal{Y}) \models \text{GBC} + \text{ETR}$. If (M, \mathcal{X}) is a β -model then so is (L^M, \mathcal{Y}) , by observation 4.11. Therefore, if there is a β -model of GBC + ETR then L_α is β -(GBC + ETR)-realizable where α is the least height of a β -model of GBC + ETR.

Now let \mathcal{X} be the least β -(GBC + ETR)-realization for L_α , which exists by theorem 4.38. We want to see that (L_α, \mathcal{X}) is contained inside every β -model of GBC + ETR. Fix $(N, \mathcal{Y}) \models \text{GBC} + \text{ETR}$ a β -model. If $\text{Ord}^N = \alpha$, then theorem 4.38 yields that $\mathcal{X} \subseteq \mathcal{Y}$. If $\text{Ord}^N > \alpha$ then $L_\alpha \in N$ and thus N can construct \mathcal{X} as in theorem 4.38 by ordinary transfinite recursion on sets. That is, N starts with $\mathcal{X}_0 = \text{Def}(L_\alpha)$. Then given \mathcal{X}_n add in all the classes definable from $\text{Tr}_\Gamma(A)$ for $\Gamma, A \in \mathcal{X}_n$ to get \mathcal{X}_{n+1} . This makes sense, because each $(L_\alpha, \mathcal{X}_n)$ is a β -model, as can be seen externally from V , and thus in N as N is correct about well-foundedness. Then $\mathcal{X} = \bigcup_n \mathcal{X}_n$ must be in N , as otherwise would imply that N does not satisfy an instance of Replacement. Thus, $(L_\alpha, \mathcal{X}) \subseteq (N, \mathcal{Y})$ \square

Essentially the same argument gives least β -models for GBC + ETR $_\Gamma$. But first a subtlety needs to be cleared up. When dealing with a fixed model with a fixed class well-order Γ it was sensical to ask whether it satisfies ETR $_\Gamma$. However, this will not work if we do not have a fixed model in mind. How are we even to express ETR $_\Gamma$ as an \mathcal{L}_\in -theory?

²²This explains why I restricted the statement of the theorem to ω -models. If M is ω -nonstandard and $(M, \mathcal{Y}) \models \Gamma \geq \omega^\omega$ then Γ must be seen from outside to be ill-founded. Nevertheless, the conclusion of the theorem is still true if Γ is in the well-founded part of ω -nonstandard M , because in such a case Γ is standard finite and ETR $_n$ for standard finite n is equivalent to Elementary Comprehension. So in this case it is just asking for M to have a least GBC-realization, which indeed does happen if M has a definable global well-order.

What we can do is ask that Γ be given some definition which evaluated in a model of **GBC** always gives a well-order. For instance, Γ could be **Ord** or ω_1 . Then, although different models may disagree on what Γ is, ETR_Γ can be expressed as an \mathcal{L}_\in theory. To distinguish this case from when Γ is a literal class in a model and ETR_Γ is expressed as an $\mathcal{L}_\in(\Gamma)$ -theory I will talk of Γ being given by a definition. For example, if I say that Γ is given by a first-order definition I mean that there is a certain first-order \mathcal{L}_\in -formula $\varphi(x, y)$ so that **GBC** proves that the class defined by $\varphi(x, y)$ is a well-order. If $(M, \mathcal{X}) \models \text{GBC}$ then I will write $\Gamma^{(M, \mathcal{X})}$ for the well-order in \mathcal{X} given by applying the definition of Γ inside (M, \mathcal{X}) . In case Γ is given by a first-order definition I will simply write Γ^M , as the evaluation depends only upon M .

Theorem 4.47. *Let $\Gamma \geq \omega^\omega$ be given by a first-order definition without parameters.²³ Assume that Γ is necessarily absolute to L , meaning that if any $(M, \mathcal{X}) \models \text{GBC}$ then $\Gamma^{L^M} = \Gamma^M$. If there is a β -model of **GBC** + ETR_Γ then there is a least β -model of **GBC** + ETR_Γ .*

I will explain after the proof why we need the absoluteness condition on Γ . For now, observe that the condition can be ensured for many ordertypes of interest, e.g. ω^ω , **Ord**, **Ord** + **Ord**, and so on. One can give a definition for a well-order of ordertype, e.g., **Ord** + **Ord** which is not absolute to L . For example: “If 0^\sharp exists then Γ is **Ord** followed by $\text{Ord} \times \{0^\sharp\}$ (with the obvious order) and if 0^\sharp does not exist then Γ is the even ordinals followed by the odd ordinals.” But there is a perfectly good definition of Γ with ordertype **Ord** + **Ord** which is absolute to L .

Proof. Let α be least such that there is a β -(**GBC** + ETR_Γ)-realizable M of height α . By theorem 3.16 we can conclude L_α is (**GBC** + ETR_Γ)-realizable. This uses that Γ is absolute to L , as the result in chapter 3 was for Γ being a fixed well-order, rather than being given by a definition. Absoluteness to L ensures that M and L_α have the same Γ . By theorem 4.42 we have \mathcal{X} the least β -(**GBC** + ETR_Γ)-realization for L_α . This uses that Γ is defined by a first-order formula, so we get the same well-order regardless of what collection of classes we put on L_α .

It remains only to see that (L_α, \mathcal{X}) is the least β -model of **GBC** + ETR_Γ . Take (N, \mathcal{Y}) a β -model of **GBC** + ETR_Γ . There are two cases. First, consider the case $\text{Ord}^N = \alpha$. Then $\Gamma^N = \Gamma^{L_\alpha}$. We saw in theorem 3.16 that there is $\bar{\mathcal{Y}} \subseteq \mathcal{Y}$ so that $(L_\alpha, \bar{\mathcal{Y}}) \subseteq (N, \mathcal{Y})$. So, by the leastness of \mathcal{X} we get $(L_\alpha, \mathcal{X}) \subseteq (L_\alpha, \bar{\mathcal{Y}}) \subseteq (N, \mathcal{Y})$. Second, consider the case $\text{Ord}^N > \alpha$. Then $L_\alpha \in N$ so $\Gamma^{L_\alpha} \in N$ and N can build \mathcal{X} by ordinary transfinite recursion. So $\mathcal{X} \in N$ and thus $(L_\alpha, \mathcal{X}) \subseteq (N, \mathcal{Y})$. \square

Let me now give an example to explain why we want to require Γ to be absolute to L . Consider Γ defined by: “If $V = L$ then $\Gamma = \text{Ord}$ and if $V \neq L$ then $\Gamma = \omega^\omega$.” Let us see that ETR_Γ does not have a least β -model. Take α least so that there is a β -model of **GBC** + $\text{ETR}_{\omega^\omega}$ of height α . Then α must be countable. Let x and y be mutually generic Cohen-reals over L_α . Then $L_\alpha[x]$ and $L_\alpha[y]$ have least β - ETR_Γ -realizations, call them \mathcal{X} and \mathcal{Y} respectively. But there is no β -model of ETR_Γ which is contained in both $(L_\alpha[x], \mathcal{X})$ and $(L_\alpha[y], \mathcal{Y})$. By leastness of α , such a model would have to have height α . But $L_\alpha[x] \cap L_\alpha[y] = L_\alpha$ by mutual genericity, so the first-order part of such a model would have to be L_α . However L_α is not β -(**GBC** + ETR_Γ)-realizable because $\Gamma^{L_\alpha} = \text{Ord}^{L_\alpha}$ and **GBC** + ETR_{Ord} proves there is a set-sized β -model of $\text{ETR}_{\omega^\omega}$.²⁴ So if L_α were β -(**GBC** + ETR_Γ)-realizable that would contradict the leastness of α .

Requiring Γ to be absolute to L rules out definitions like this one.

This argument only needs that the models are correct about their Γ being well-founded. So if Γ is given by a definition which interpreted in any transitive model gives a relation which really is a well-order, then we can get a least transitive model of **GBC** + ETR_Γ .

Theorem 4.48. *Let $\Gamma \geq \omega^\omega$ be given by a first-order definition without parameters. Assume the following.*

1. Γ is absolute to L , meaning that if $(M, \mathcal{X}) \models \text{GBC}$ then $\Gamma^{L^M} = \Gamma^M$; and
2. If $(M, \mathcal{X}) \models \text{GBC}$ is transitive then Γ^M really is well-founded, as seen externally.

²³That is, **GBC** proves that $\Gamma \geq \omega^\omega$.

²⁴This is because **GBC** + ETR_{Ord} proves there is a coded V -submodel of $\text{ETR}_{\omega^\omega}$ and this can be reflected down to get a set-sized β -model of $\text{ETR}_{\omega^\omega}$.

Then, if there is a β -model of $\text{GBC} + \text{ETR}_\Gamma$ there is a least transitive model of $\text{GBC} + \text{ETR}_\Gamma$. \square

As particular cases of interest, this works when $\Gamma = \text{Ord}$ or Γ is a given by a definition for a (set-sized) ordinal.

Left open is how high up this can be pushed. Can it be pushed all the way up to ETR ?

Question 4.49 (Open). *Is there a least transitive model of $\text{GBC} + \text{ETR}$?*

Let me note this question is not immediately settled by master lemma 4.20.

Proposition 4.50. *The theory $\text{GBC} + \text{ETR}$ does not prove that $\text{Hyp}(V)$ exists, assuming that it is consistent.*

Proof. Suppose $(M, \mathcal{X}) \models \text{GBC} + \text{ETR} + \text{Hyp}(V)$ exists. (If there is no such model, then we are already done.) By shrinking down to an inner model if necessary, assume without loss that M has a definable global well-order. Unroll (M, \mathcal{X}) to W . From the results in chapter 2, we get that W satisfies Σ_0 -Transfinite Recursion. Now let $\mathcal{Y} = \{A \in W : W \models A \in \text{Hyp}(M) \text{ and } A \subseteq M\}$ consist of the classes of M which are in $\text{Hyp}(M)^W$. Then \mathcal{Y} is closed under first-order definability, so it satisfies Elementary Comprehension. It satisfies Class Replacement because it is a V -submodel of a model of GBC . Class Extensionality is obvious and Global Choice holds because M has a definable global well-order. Altogether, we have seen that $(M, \mathcal{Y}) \models \text{GBC}$. But it also must satisfy ETR , because $\text{Hyp}(M)^W \models \text{KP}$ and KP proves Σ_0 -Transfinite Recursion. But $(M, \mathcal{Y}) \not\models \text{Hyp}(V)$ exists, by construction. so $\text{GBC} + \text{ETR}$ does not prove that $\text{Hyp}(V)$ exists. \square

Let me mention a couple related questions. The first question was also asked at the end of chapter 2.

Question 4.51. *Let $\tau(\text{GBC} + \text{ETR})$ be the least height of a transitive model of $\text{GBC} + \text{ETR}$ and let $\beta(\text{GBC} + \text{ETR})$ be the least height of a β -model of $\text{GBC} + \text{ETR}$. Can we conclude that $\tau(\text{GBC} + \text{ETR}) < \beta(\text{GBC} + \text{ETR})$?*

If the answer is no, then there is a least transitive model of $\text{GBC} + \text{ETR}$, which would be the same as the least β -model of $\text{GBC} + \text{ETR}$.

Question 4.52. *Can there be $M \models \text{ZFC}$ which is $(\text{GBC} + \text{ETR})$ -realizable with a class well-order $\Gamma \in \text{Def}(M)$ so that there are two different $(\text{GBC} + \text{ETR})$ -realizations \mathcal{X} and \mathcal{Y} for M so that $\text{Tr}_\Gamma^{(M, \mathcal{X})} \neq \text{Tr}_\Gamma^{(M, \mathcal{Y})}$? What if we restrict to countable M ?*

It follows from theorem 4.38 that this cannot happen if M has a β - $(\text{GBC} + \text{ETR})$ -realization.

4.5 Weak theories

The results in this section are either old or appeared in chapter 1. I state them here to round out the presentation in this chapter.

Let us start with a classical result.

Theorem 4.53 (Shepherdson [She53]). *There is a least transitive model of GBC , if there is any transitive model of GBC .*

Proof. Let L_α be the least transitive model of ZFC . Then $(L_\alpha, \text{Def}(L_\alpha))$ is the least transitive model of GBC . \square

A slight modification also gives a least β -model.

Corollary 4.54. *There is a least β -model of GBC , if there is any β -model of GBC .*

Proof. First, observe that if (M, \mathcal{X}) is a β -model of GBC then $(L^M, \text{Def}(L^M))$ is also a β -model of GBC . This is because $(L_\alpha, \text{Def}(L_\alpha))$ is always a model of GBC and because Ord -submodels of β -models are β -models.

Now let α be least such that there is a β -model of GBC of height α . Then $(L_\alpha, \text{Def}(L_\alpha))$ is the least β -model of GBC . \square

As with the case for ETR, it is not clear to me that these are actually two different results.

Question 4.55. *Is the height of the least transitive model of GBC less than the height of the least β -model of GBC? Phrased differently, is the least transitive model of GBC a β -model?*

We can also characterize when a countable model of ZFC has a least GBC-realization.

Theorem 4.56. *Let $M \models \text{ZFC}$ be countable. Then the following are equivalent.*

1. M has a least GBC-realization.
2. M has a definable global well-order.
3. $M \models \exists x V = \text{HOD}(\{x\})$.

Proof. (1 \Leftrightarrow 2) was theorem 1.61.(3) from chapter 1. (2 \Leftrightarrow 3) is a well-known fact.²⁵ □

4.6 Coda: the analogy to second-order arithmetic

Several times throughout the course of this dissertation we have touched upon the analogy between second-order set theory and second-order arithmetic. Now that we are done with the major results I would like to flesh this analogy out more fully. Let us begin by seeing how the theories line up.

Arithmetic	Set theory
Z_2	KM
$\Pi_1^1\text{-CA}_0$	$\Pi_1^1\text{-CA}$
ATR_0	ETR
ACA_0	GBC
WKL_0	
RCA_0	

Figure 4.5: Arithmetic versus set theory.

It's well-known that PA is bi-interpretable with finite set theory $\text{ZFC}^{\neg\infty}$, i.e. ZFC with Infinity replaced with its negation.²⁶ This bi-interpretability carries over for second-order set theories so that e.g. ACA_0 and $\text{GBC}^{\neg\infty}$ are bi-interpretable. So the analogy here is really between the finite and the transfinite.

Let me briefly address the gap in the table. First, WKL_0 . Enayat and Hamkins [EH] proved that (in ZFC) there is a definable Ord-tree whose levels are all set-sized with no definable branch. Consequently, GBC does not prove the analog of König's lemma for Ord-trees instead of ω -trees. So it is not clear how tree properties on ω —such as weak König's lemma or König's lemma—could be generalized to this context. For RCA_0 , different ways of thinking of computability suggests different generalizations to set theory. First, consider the view that computable = Δ_1 . With this view in mind, the set theoretic counterpart to RCA_0 would be based upon Δ_1^0 -Comprehension. Second, take the view that computable means verifiable and refutable by only looking at a bounded segment of the universe. By this view, the set theoretic counterpart to RCA_0 would be based upon Δ_2^0 -Comprehension, as the Δ_2 properties are precisely those which are verifiable and refutable by looking at a rank-initial segment of the universe. Which of these two counterparts is the 'correct' one would depend upon the mathematics one can do with them. If, say, Δ_2^0 -Comprehension allows us to do a lot of interesting mathematics but Δ_1^0 -Comprehension does not, then we would have reason to prefer one

²⁵The reason for using $\text{HOD}(\{x\})$ instead of HOD is that we want to possibly allow parameters for the definition of the global well-order.

²⁶Note that, however, this is sensitive to how Foundation is formulated, as formulations equivalent over infinitary set theory are not equivalent over finite set theory. So finite set theory should be understood as formulated with the right version of Foundation. See [KW07] for a discussion.

over the other. But we cannot make such a call before actually doing that mathematics. For now, it is not clear what the set theoretic counterpart to RCA_0 should be.

That aside, let us return to where there are no gaps. Some results about models of arithmetic have direct generalizations to results about models of set theory. For example, it is well-known that every model of arithmetic has a least ACA_0 -realization. The set theoretic counterpart to this fact is that every model of ZFC with a definable global well-order has a least GBC -realization. We also get counterparts to results about strong theories. That KM has a least β -model but no least transitive model is the transfinite analog of the folklore fact that \mathbf{Z}_2 has a least β -model and H. Friedman’s theorem [Fri73] that \mathbf{Z}_2 has no least ω -model.

But disanalogies appear at the level of $\text{GBC} + \text{ETR}$ versus ATR_0 . We have seen that $\text{GBC} + \text{ETR}$ has a least β -model. On the other hand, ATR_0 has neither a least ω -model nor a least β -model—see [Sim09]. Let us explore this disanalogy further. The main culprit here is the property “ X is a well-order”. In set theory, this is a first-order assertion. Whether X is a well-order is determined by a countable piece of information, so is witnessed by the (non)existence of certain sets, not proper classes. On the other hand, in arithmetic it is a second-order assertion to say that X is a well-order, since infinite sequences are second-order objects in the arithmetic context. Indeed, “ X is a well-order” is Π_1^1 -universal in arithmetic, so which classes are well-ordered is very much caught up in the second-order part of the model.

Another disanalogy concerns the existence of $\text{Hyp}(V)$. It is well-known that $\text{Hyp}(V_\omega) = L_{\omega_1^{\text{CK}}}$, where ω_1^{CK} is the least non-computable ordinal. So ATR_0 suffices to prove that $\text{Hyp}(V)$ is coded²⁷ because ω_1^{CK} is arithmetical and so we can do a transfinite recursion along it to produce the L -hierarchy. Indeed, the intersection of all the β -models of ATR_0 is the collection of hyperarithmetical sets, those reals appearing in $L_{\omega_1^{\text{CK}}}$ —see [Sim09]. (And the same is true if we consider ω -models instead of β -models.) But in the set theoretic context, we have seen that $\text{GBC} + \text{ETR}$ does not suffice to prove that $\text{Hyp}(V)$ is coded.

More disanalogies are known. To pick one last example, in arithmetic Clopen Determinacy and Open Determinacy are both equivalent to ATR_0 , a result originating in Steel’s dissertation [Ste76]. In set theory, Clopen Determinacy (for class games) is equivalent (over GBC) to ETR [GH17]. But Open Determinacy is strictly stronger [Hac16].

These suggest that the emerging field of reverse mathematics of second-order set theory should reveal a landscape with some differences from that of reverse mathematics of arithmetic. As further evidence pointing toward this, in the set theoretic context a fragment of ETR captures natural mathematical principles. Gitman, Hamkins, Holy, Schlicht, and myself showed [GHHSW17] that ETR_{Ord} is equivalent (over GBC) to the class forcing theorem and several other natural statements. On the other hand, the same does not happen for fragments of ATR_0 .

More speculatively, work in second-order set theory might shed some light, even if only privatively, on the project of reverse mathematics in second-order arithmetic. Each of the Big Five subsystems of second-order arithmetic roughly corresponds to a philosophical position. (Cf. chapter I of [Sim09].) For example, ATR_0 corresponds to predicative reductionism, where “predicative” here means “predicative given ω ”. We could instead ask about predicativism given V , or given H_{ω_1} , or some other object. This project has already been taken up, e.g. by Sato [Sat14]. The disanalogies between ETR and ATR_0 suggest that starting from the finite realm has a large impact on predicativism and related projects. If we take the transfinite as our starting point then we end up with a very different theory. On the other hand, where we do see an analogy between second-order arithmetic and second-order set theory it suggests that there the role of the finite versus the transfinite is not so vital.

In short, work in second-order set theory may help clarify what in reverse mathematics (of arithmetic) relies essentially upon the first-order domain consisting of finite objects.

²⁷Remember that in arithmetic, V is V_ω !

Bibliography

- [ABF] Carolin Antos, Neil Barton, and Sy-David Friedman. *Universism and extensions of V* . submitted for publication. arXiv:1708.05751[math.LO].
- [AF] Carolin Antos and Sy-David Friedman. *Hyperclass Forcing in Morse-Kelley Class Theory*. accepted and forthcoming. arXiv:1510.04082[math.LO].
- [Ant15] Carolin Antos. “Class Forcing in Class Theory” (2015). submitted for publication. arXiv:1503.00116[math.LO].
- [Bar75] Jon Barwise. *Admissible sets and structures*. An approach to definability theory, Perspectives in Mathematical Logic. Springer-Verlag, Berlin-New York, 1975. DOI: 10.2307/2271957.
- [BS76] Jon Barwise and John Schlipf. “An introduction to recursively saturated and resplendent models”. *Journal of Symbolic Logic* 41.2 (1976), pp. 531–536.
- [Chu80] Rolando Chuaqui. *Internal and forcing models for the impredicative theory of classes*. Instytut Matematyczny Polskiej Akademi Nauk (Warszawa), 1980.
- [Chu81] Rolando Chuaqui. *Axiomatic Set Theory: Impredicative Theory of Classes*. North-Holland Pub. Co., 1981.
- [Coh63] Paul Cohen. “A minimal model for set theory”. *Bulletin of the American Mathematical Society* 69.4 (1963), pp. 537–540. <http://projecteuclid.org/euclid.bams/1183525369>.
- [Die83] Karl-Heinz Diener. “On constructing infinitary languages $L_{\alpha,\beta}$ without the axiom of choice”. *Mathematical Logic Quarterly* 29.6 (1983), pp. 357–376.
- [EH] Ali Enayat and Joel David Hamkins. “ZFC proves that the class of ordinals is not weakly compact for definable classes” (). manuscript under review. arXiv:1610.02729[math.LO]. <http://jdhamkins.org/ord-is-not-definably-weakly-compact>.
- [EV15] Ali Enayat and Albert Visser. “New constructions of satisfaction classes”. In: *Unifying the Philosophy of Truth. Logic, Epistemology, and the Unity of Science*. Ed. by T. Achourioti, H. Galinon, J. Martínez Fernández, and K. Fujimoto. Springer, Dordrecht, 2015, pp. 321–335.
- [FG17] Sy-David Friedman and Victoria Gitman. “A model of second-order arithmetic satisfying AC but not DC” (2017). submitted for publication.
- [Fra22] A. Fraenkel. “Zu den Grundlagen der Cantor-Zermeloschen Mengenlehre”. *Mathematische Annalen* 86 (1922), pp. 230–237. <http://eudml.org/doc/158946>.
- [Fra28] A. Fraenkel. *Einleitung in die Mengenlehre*. 3rd ed. Die Grundlehren der Mathematischen Wissenschaften 9. Springer-Verlag Berlin Heidelberg, 1928.
- [Fri00] Sy D. Friedman. *Fine Structure and Class Forcing*. Reprint 2011 ed. De Gruyter Series in Logic and Its Applications, 3. De Gruyter, 2000. ISBN: 3110167778,9783110167771.
- [Fri73] Harvey Friedman. “Countable Models of Set Theories”. In: *Cambridge Summer School in Mathematical Logic*. Ed. by A. R. D. Mathias and H. Rogers. New York, Springer-Verlag, 1973, pp. 539–573.

- [Fuj12] Kentaro Fujimoto. “Classes and truths in set theory”. *Annals of Pure & Applied Logic* 163.11 (2012), pp. 1484–1523. ISSN: 0003-4843. DOI: 10.1016/j.apal.2011.12.006.
- [Gaß94] Christine Gaßner. “The axiom of choice in second-order predicate logic”. *Mathematical Logic Quarterly* 40.4 (1994), pp. 533–546.
- [GH] Victoria Gitman and Joel David Hamkins. “Kelley-Morse Set Theory and Choice Principles for Classes” (). Unpublished.
- [GH17] Victoria Gitman and Joel David Hamkins. “Open determinacy for class games”. In: *Foundations of Mathematics, Logic at Harvard, Essays in Honor of Hugh Woodin’s 60th Birthday*. Ed. by Andrés E. Caicedo, James Cummings, Peter Koellner, and Paul Larson. Contemporary Mathematics. American Mathematical Society, 2017. arXiv:1509.01099[math.LO]. <http://jdh.hamkins.org/open-determinacy-for-class-games>.
- [GHHSW17] Victoria Gitman, Joel David Hamkins, Peter Holy, Philipp Schlicht, and Kameryn Williams. “The exact strength of the class forcing theorem” (2017). manuscript under review. arXiv:1707.03700[math.LO]. <http://jdh.hamkins.org/class-forcing-theorem>.
- [GHJ16] Victoria Gitman, Joel David Hamkins, and Thomas A. Johnstone. “What is the theory ZFC without Powerset?” *Mathematical Logic Quarterly* 62.4–5 (2016), pp. 391–406. ISSN: 0942-5616. DOI: 10.1002/malq.201500019. arXiv:1110.2430[math.LO]. <http://jdh.hamkins.org/what-is-the-theory-zfc-without-power-set>.
- [Göd38] Kurt Gödel. “The Consistency of the Axiom of Choice and of the Generalized Continuum-Hypothesis”. *Proceedings of the National Academy of Sciences of the United States of America* 24.12 (1938), pp. 556–557.
- [Hac16] Sherwood Hachtman. “Determinacy separations for class games” (2016). submitted for publication. arXiv:1607.05515[math.LO].
- [Ham16] Joel David Hamkins. “Upward closure and amalgamation in the generic multiverse of a countable model of set theory”. *RIMS Kyôkyûroku* (2016), pp. 17–31. ISSN: 1880-2818. arXiv:1511.01074[math.LO]. <http://jdh.hamkins.org/upward-closure-and-amalgamation-in-the-generic-multiverse>.
- [HKLNS16] Peter Holy, Regula Krapf, Philipp Lücke, Ana Njegomir, and Philipp Schlicht. “Class forcing, the forcing theorem and boolean completions”. *The Journal of Symbolic Logic* 81.4 (2016), pp. 1500–1530. DOI: 10.1017/jsl.2016.4.
- [HKSa] Peter Holy, Regula Krapf, and Philipp Schlicht. “Characterizations of pretameness and the Ord-cc” (). submitted for publication.
- [HKSb] Peter Holy, Regula Krapf, and Philipp Schlicht. “Separation in Class Forcing Extensions” (). submitted for publication.
- [HLR13] Joel David Hamkins, David Linetsky, and Jonas Reitz. “Pointwise definable models of set theory”. *J. Symbolic Logic* 78.1 (2013), pp. 139–156. arXiv:1105.4597[math.LO].
- [HY14] Joel David Hamkins and Ruizhi Yang. “Satisfaction is not absolute”. *to appear in the Review of Symbolic Logic* (2014), pp. 1–34. arXiv:1312.0670[math.LO]. <http://jdh.hamkins.org/satisfaction-is-not-absolute>.
- [Kau77] Matt Kaufmann. “A rather classless model”. *Proceedings of the American Mathematical Society* 62.2 (1977), pp. 330–333.
- [Kei74] H. Jerome Keisler. “Models with tree structures”. In: *Proceedings of the Tarski Symposium (Proc. Sympos. Pure Math., Vol. XXV, Univ. California, Berkeley, Calif., 1971)*. Amer. Math. Soc., Providence, R.I., 1974, pp. 331–348.
- [Kel75] John L. Kelley. *General Topology*. Springer-Verlag New York, 1975.
- [Kra76] S. Krajewski. “Non-standard satisfaction classes”. In: *Set theory and hierarchy theory (Proc. Second Conf., Bierutowice, 1975)*. Springer, Berlin, 1976, 121–144. Lecture Notes in Math., Vol. 537.

- [KS06] Roman Kossak and James H. Schmerl. *The structure of models of Peano arithmetic*. Vol. 50. Oxford Logic Guides. Oxford Science Publications. Oxford: The Clarendon Press Oxford University Press, 2006.
- [KSL81] H. Kotlarski, Krajewski S., and A. H. Lachlan. “Construction of satisfaction classes for non-standard models”. *Canadian Mathematical Bulletin* 24.3 (1981), pp. 283–293.
- [KW07] Richard Kaye and Tin Lok Wong. “On Interpretations of Arithmetic and Set Theory”. *Notre Dame J. Formal Logic* 48.4 (2007), pp. 497–510. <http://dx.doi.org/10.1305/ndjfl/1193667707>.
- [Lac81] A. H. Lachlan. “Full satisfaction classes and recursive saturation”. *Canadian Mathematical Bulletin* 24.3 (1981), pp. 295–297.
- [Mar73] Wiktor Marek. *On the metamathematics of impredicative set theory*. Warszawa, Polska: Panstwowe Wydawn Naukowe, 1973.
- [Mat01] A. R. D. Mathias. “The strength of Mac Lane set theory”. *Ann. Pure Appl. Logic* 110.1-3 (2001), pp. 107–234. ISSN: 0168-0072. [https://doi.org/10.1016/S0168-0072\(00\)00031-2](https://doi.org/10.1016/S0168-0072(00)00031-2).
- [MM75] W. Marek and A. Mostowski. “On extendability of models of ZF set theory to the models of Kelley-Morse theory of classes”. In: *ISILC Logic Conference: Proceedings of the International Summer Institute and Logic Colloquium, Kiel 1974*. Ed. by Gert H. Müller, Arnold Oberschelp, and Klaus Potthoff. Berlin, Heidelberg: Springer Berlin Heidelberg, 1975, pp. 460–542. DOI: 10.1007/BFb0079429.
- [Mos76] Andrzej Mostowski. “A Remark on Models of the Gdel-Bernays Axioms for Set Theory”. In: *Sets and Classes on The Work by Paul Bernays*. Ed. by Gert H. Mller. Vol. 84. Studies in Logic and the Foundations of Mathematics. Elsevier, 1976, pp. 325–340.
- [Neu25] John von Neumann. “Eiene Axiomatisierung der Mengenlehre”. *Journal für die reine und angewandte Mathematik* 154 (1925), pp. 219–240.
- [Pyt70] Monty Python. *Spam*. Sketch appeared in series 2, episode 12 of Monty Python’s Flying Circus. 1970. https://www.youtube.com/watch?v=M_eYSuPKP3Y.
- [Sat14] Kentaro Sato. “Relative predicativity and dependent recursion in second-order set theory and higher-order theories”. *Journal of Symbolic Logic* 79.3 (2014), pp. 712–732.
- [Sha91] Stewart Shapiro. *Foundations without Foundationalism: A Case For Second-Order Logic*. Claredon Press, 1991.
- [She53] J. C. Shepherdson. “Inner Models for Set Theory—Part III”. *The Journal of Symbolic Logic* 18.2 (1953), pp. 145–167. ISSN: 00224812. <http://www.jstor.org/stable/2268947>.
- [She78] Saharon Shelah. “Models with second order properties II. Trees with no undefined branches”. *Annals of Mathematical Logic* 14.1 (1978), pp. 73–87.
- [Sim09] Stephen George Simpson. *Subsystems of second order arithmetic*. Perspectives in Logic. New York: Association for Symbolic Logic, 2009.
- [Smi89] S. T. Smith. “Nonstandard definability”. *Annals of Pure and Applied Logic* 42.1 (1989), pp. 21–43.
- [Spe56] Clifford Spector. “On degrees of recursive unsolvability”. *Annals of Mathematics* 64.3 (1956), pp. 581–592.
- [Sta84] Maurice C. Stanley. “A unique generic real”. PhD thesis. University of California at Berkeley, 1984.
- [Ste76] John R. Steel. “Determinateness and Subsystems of Analysis”. PhD thesis. University of California at Berkeley, 1976.
- [Usu17] Toshimichi Usuba. “The downwarddirectedgrounds hypothesis and very large cardinals”. *Journal of Mathematical Logic* 17.02 (2017), p. 1750009.

- [Vää01] Jouko Väänänen. “Second-order logic and foundations of mathematics”. *The Bulletin of Symbolic Logic* 7.4 (2001), pp. 504–520.
- [Wil17] Kameryn J Williams. “Least models of second-order set theories” (2017). manuscript under review. arXiv:1709.03955[math.LO].
- [Zar96] Andrzej M. Zarach. “Replacement \rightarrow collection”. In: *Gödel '96: Logical foundations of mathematics, computer science and physics—Kurt Gödel's legacy, Brno, Czech Republic, August 1996, proceedings*. Vol. 6. Lecture Notes in Logic. Berlin: Springer-Verlag, 1996, pp. 307–322. <http://projecteuclid.org/euclid.lnl/1235417032>.
- [Zer30] Ernst Zermelo. “Über Grenzzahlen und Mengenbereiche: Neue Untersuchungen über die Grundlagen der Mengenlehre”. *Fundamenta Mathematicae* (1930), pp. 29–47.