

THE EXACT STRENGTH OF THE CLASS FORCING THEOREM

VICTORIA GITMAN, JOEL DAVID HAMKINS, PETER HOLY, PHILIPP SCHLICHT,
AND KAMERYN WILLIAMS

ABSTRACT. The class forcing theorem, which asserts that every class forcing notion \mathbb{P} admits a forcing relation $\Vdash_{\mathbb{P}}$, that is, a relation satisfying the forcing relation recursion—it follows that statements true in the corresponding forcing extensions are forced and forced statements are true—is equivalent over Gödel-Bernays set theory GBC to the principle of elementary transfinite recursion ETR_{Ord} for class recursions of length Ord . It is also equivalent to the existence of truth predicates for the infinitary languages $\mathcal{L}_{\text{Ord},\omega}(\in, A)$, allowing any class parameter A ; to the existence of truth predicates for the language $\mathcal{L}_{\text{Ord},\text{Ord}}(\in, A)$; to the existence of Ord -iterated truth predicates for first-order set theory $\mathcal{L}_{\omega,\omega}(\in, A)$; to the assertion that every separative class partial order \mathbb{P} has a set-complete class Boolean completion; to a class-join separation principle; and to the principle of determinacy for clopen class games of rank at most $\text{Ord} + 1$. Unlike set forcing, if every class forcing notion \mathbb{P} has a forcing relation merely for atomic formulas, then every such \mathbb{P} has a uniform forcing relation applicable simultaneously to all formulas. Our results situate the class forcing theorem in the rich hierarchy of theories between GBC and Kelley-Morse set theory KM.

1. INTRODUCTION

We shall characterize the exact strength of the class forcing theorem, which asserts that every class forcing notion \mathbb{P} has a corresponding forcing relation $\Vdash_{\mathbb{P}}$, a relation satisfying the forcing relation recursion. When there is such a forcing relation, then statements true in any corresponding forcing extension are forced and forced statements are true in those extensions.

Unlike set forcing, for which one may prove in ZFC that every set forcing notion has corresponding forcing relations, with class forcing it is consistent with Gödel-Bernays set theory GBC that there is a proper class forcing notion lacking a corresponding forcing relation, even merely for the atomic formulas. For certain forcing notions (see [HKL⁺16, Kra17], also theorem 17), the existence of an atomic forcing relation implies $\text{Con}(\text{ZFC})$ and much more, and so the consistency strength of the class forcing theorem strictly exceeds GBC, if this theory is consistent. Nevertheless, the class forcing theorem is provable in stronger theories, such as Kelley-Morse set theory. What is the exact strength of the class forcing theorem?

Our project here is to identify the strength of the class forcing theorem by situating it in the rich hierarchy of theories between GBC and KM, displayed in part in figure 1, with the class forcing theorem highlighted in blue. It turns out

The research of the second author has been supported by grant #69573-00 47 from the CUNY Research Foundation. He is also grateful for the support he has received from the Hausdorff Center at the University of Bonn for several visits there. The third and fourth author were partially supported by DFG-grant LU2020/1-1. Commentary concerning this paper can be made at <http://jdh.hamkins.org/class-forcing-theorem>.

that the class forcing theorem is equivalent over GBC to an attractive collection of several other natural set-theoretic assertions; it is a robust axiomatic principle.

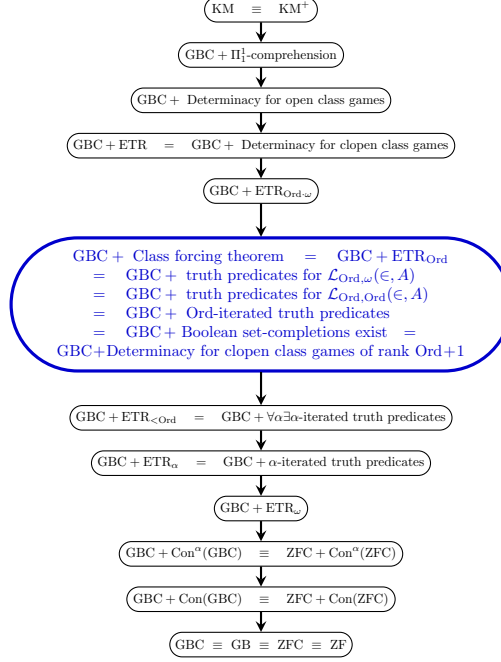


FIGURE 1. A hierarchy of theories between GBC and KM by consistency strength (\equiv means equiconsistent)

The main theorem is naturally part of the emerging subject we call the reverse mathematics of second-order set theory, a higher analogue of the perhaps more familiar reverse mathematics of second-order arithmetic. In this new research area, we are concerned with the hierarchy of second-order set theories between GBC and KM and beyond, analyzing the strength of various assertions in second-order set theory, such as the principle ETR of elementary transfinite recursion, the principle of Π_1^1 -comprehension or the principle of determinacy for clopen class games. We fit these set-theoretic principles into the hierarchy of theories over the base theory GBC . The main theorem of this article does exactly this with the class forcing theorem by finding its exact strength in relation to nearby theories in this hierarchy.

Specifically, extending the analysis of [HKL⁺16, HKS17, Kra17, GH16], we show that the class forcing theorem is equivalent over GBC to the principle of elementary transfinite recursion ETR_{Ord} for transfinite class recursions of length Ord ; to the existence of various kinds of truth predicates and iterated truth-predicates; to the existence of Boolean completions for any separative class partial order; to a class-join separation principle; and to the principle of determinacy for clopen class games of rank at most $\text{Ord} + 1$. In addition, by separating the class forcing theorem from the nearby theories of figure 1, placing it strictly between the theory with ETR_α simultaneously for all ordinals α and the theory $\text{ETR}_{\text{Ord},\omega}$, we locate it finely in the hierarchy of second-order set theories.

Main Theorem. *The following are equivalent over Gödel-Bernays set theory GBC.*

- (1) *The atomic class forcing theorem: every class forcing notion admits forcing relations for atomic formulas*

$$p \Vdash \sigma = \tau \qquad p \Vdash \sigma \in \tau.$$

- (2) *The class forcing theorem scheme: for each first-order formula φ in the forcing language, with finitely many class names $\dot{\Gamma}_i$, there is a forcing relation applicable to this formula and its subformulas*

$$p \Vdash \varphi(\vec{\tau}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m).$$

- (3) *The uniform first-order class forcing theorem: every class forcing notion \mathbb{P} admits a uniform forcing relation*

$$p \Vdash \varphi(\vec{\tau}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$$

applicable to all assertions $\varphi \in \mathcal{L}_{\omega, \omega}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$ in the first-order forcing language with finitely many class names.

- (4) *The uniform infinitary class forcing theorem: every class forcing notion \mathbb{P} admits a uniform forcing relation*

$$p \Vdash \varphi(\vec{\tau}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$$

applicable to all assertions $\varphi \in \mathcal{L}_{\text{Ord}, \text{Ord}}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$ in the infinitary forcing language with finitely many class names.

- (5) *Names for truth predicates: every class forcing notion \mathbb{P} has a class name \dot{T} and a forcing relation for which $\mathbb{1} \Vdash \dot{T}$ is a truth-predicate for the first-order forcing language with finitely many class names $\mathcal{L}_{\omega, \omega}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$.*

- (6) *Boolean completions: Every class forcing notion \mathbb{P} , that is, every separative class partial order, admits a Boolean completion \mathbb{B} , a set-complete class Boolean algebra into which \mathbb{P} densely embeds.*

- (7) *The class-join separation principle plus ETR_{Ord} -foundation.*


- (8) *For every class A , there is a truth predicate for $\mathcal{L}_{\text{Ord}, \omega}(\in, A)$.*

- (9) *For every class A , there is a truth predicate for $\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, A)$.*

- (10) *For every class A , there is an Ord-iterated truth predicate for $\mathcal{L}_{\omega, \omega}(\in, A)$.*

- (11) *The principle of determinacy for clopen class games of rank at most $\text{Ord}+1$.*

- (12) *The principle ETR_{Ord} of elementary transfinite recursion for Ord-length recursions of first-order properties, using any class parameter.*



We shall prove the theorem by establishing the complete cycle of implications pictured above. For convenience, the implication diagram is clickable, with each implication arrow linking to the corresponding theorem where the implication is

proved. In addition, all subsequent mentions of an implication, such as (8) \rightarrow (10), link to the corresponding theorem statements. Precise details for the terms appearing in the theorem statements appear in the various sections below where the corresponding implications are proved. The red implication arrows indicate more substantial or critical implication arguments; blue arrows correspond to the essentially immediate implications; and the dashed green arrow indicates implication (12) \rightarrow (1), which although not needed to complete the cycle, is nevertheless used in the proof of theorem 13, establishing the implication (12) \rightarrow (4).

We should like particularly to emphasize that statement (2) is a scheme over the formulas φ that are finite in the meta-theory, stating as a scheme that for each such formula, there is a forcing relation class that works with that formula and its subformulas. Statements (3) and (4), in contrast, are not schemes, but assert in each case that there is a single uniform relation that works with all formulas simultaneously. In the set-forcing ZFC context, we are used to having the forcing relations $p \Vdash \varphi(\vec{\tau})$ available only as a scheme, a separate relation for each formula φ , and because of Tarski's theorem on the non-definability of truth, it follows that one cannot prove in ZFC that there is a single unified forcing relation $\Vdash_{\mathbb{P}}$ that works with all formulas (although by theorem 12 one can do this in the quantifier-free infinitary case). In the case of class forcing, however, our main theorem shows that if every class forcing notion \mathbb{P} admits forcing relations merely for atomic formulas, then in fact they all have fully uniform forcing relations $\Vdash_{\mathbb{P}}$, and not only for the first-order forcing languages, but for the infinitary forcing languages, as in statement (4).

Let us remark specifically on the role of the global axiom of choice in this analysis. The base theory for the subject is Gödel-Bernays set theory GBC, which includes the global axiom of choice, the assertion that there is a class well-ordering of the universe. But actually, none of the arguments in the proof of the main theorem require the global axiom of choice, as opposed to the ordinary axiom of choice for sets, with the exception of the implications to the clopen determinacy assertion of statement (11). In particular, our arguments show that all the statements in the main theorem except statement (11) are equivalent over the theory GB+AC, which has the axiom of choice only for sets. Meanwhile, the clopen determinacy assertion of statement (11) in the main theorem implies the global axiom of choice, by the folklore result mentioned in [GH16, theorem 4], and so global choice is required if one wants to include statement (11).

For definiteness, in this article we say \mathbb{P} is a *class forcing notion* if \mathbb{P} is a separative class pre-order.

2. ETR_{Ord} IMPLIES THE CLASS FORCING THEOREM SCHEME

In this section, we shall prove the implications (12) \rightarrow (1) and (1) \rightarrow (2) in the main theorem. Let's begin by defining the notions carefully.

Definition 1. The principle of *elementary transfinite recursion* ETR_{Ord} for recursions of length Ord , is the scheme asserting of any first order formula $\varphi(x, X, A)$ with a class parameter A , that there is a class $S \subseteq \text{Ord} \times V$ that is a solution of the following recursion

$$S_\alpha = \{ x \mid \varphi(x, S \restriction \alpha, A) \},$$

where $S_\alpha = \{ x \mid \langle \alpha, x \rangle \in S \}$ denotes the α^{th} slice of S and $S \restriction \alpha = S \cap (\alpha \times V)$ is the part of the solution prior to stage α .

Thus, the axiom asserts that we may undertake transfinite recursive definitions of classes by recursions of length Ord . In general, in GBC one may not necessarily undertake class recursions even merely of length ω , since first-order truth, of course, is defined by the Tarskian recursion on formulas, and this recursion has length merely ω ; but in GBC, if consistent, one cannot prove the existence of a truth-predicate for first-order truth and therefore this recursion may have no solution. So even ETR_ω , which asserts that class recursions of length ω have solutions, is not provable in GBC, if GBC is consistent, and consequently neither is ETR_{Ord} .

Meanwhile, ETR_{Ord} is a consequence of the full principle ETR of elementary transfinite recursion, which allows recursion along any class well-founded relation, including relations much taller than Ord , and this principle is strictly weaker than $\text{GBC} + \Pi_1^1$ -comprehension [Sat14], which is itself strictly weaker than Kelley-Morse set theory KM. Gitman and Hamkins [GH16] proved that ETR is equivalent over GBC to the principle of determinacy for clopen class games.

The idea of [GH16, lemma 7] shows that the principle ETR_{Ord} is equivalently formulated in terms of recursions along arbitrary well-founded class relations of rank Ord .

Definition 2. A class forcing notion \mathbb{P} *admits forcing relations for atomic formulas*, if there are relations

$$p \Vdash \sigma \in \tau, \quad p \Vdash \sigma \subseteq \tau, \quad p \Vdash \sigma = \tau$$

which obey the following recursive properties:

- (a) $p \Vdash \sigma \in \tau$ if and only if there is a dense class of conditions $q \leq p$ for which there is $\langle \rho, r \rangle \in \tau$ with $q \leq r$ and $q \Vdash \sigma = \rho$.
- (b) $p \Vdash \sigma = \tau$ if and only if $p \Vdash \sigma \subseteq \tau$ and $p \Vdash \tau \subseteq \sigma$.
- (c) $p \Vdash \sigma \subseteq \tau$ if and only if whenever $\langle \rho, r \rangle \in \sigma$ and $q' \leq p, r$, there is $q \leq q'$ with $q \Vdash \rho \in \tau$.

Since the formulas are distinguished syntactically, we may unify the three forcing relations into a single relation \Vdash , applied to any atomic assertion. One may take statement (c) as a definition of the relation $p \Vdash \sigma \subseteq \tau$ in terms of the forcing relation for \in , and in this case statements (a) and (b) are expressible as a recursion solely in terms of $p \Vdash \sigma \in \tau$ and $p \Vdash \sigma = \tau$. So the use of \subseteq here is merely a convenience.

Theorem 3. *The principle ETR_{Ord} of elementary transfinite recursion for class recursions of length Ord implies that every class forcing notion \mathbb{P} admits a forcing relation for atomic formulas.*

This will establish implication (12) \rightarrow (1) in the main theorem.

Proof. The main point is that having a forcing relation for atomic formulas, by definition, is to have a solution of the recursion that is expressed by definition 2. Since this is an \in -recursion on the \mathbb{P} -names, we may organize it as a recursion of length Ord , using the natural Ord -ranking of pairs of names $\langle \sigma, \tau \rangle$, which are ordered first by the maximum of their ranks and then lexically by rank. We may now place the forcing relation $p \Vdash \sigma \in \tau$ and $p \Vdash \sigma = \tau$ on the α^{th} slice, when the pair $\langle \sigma, \tau \rangle$ has rank α with respect to that relation; the forcing relation for such names is defined in terms of the forcing relation on preceding pairs of names. So ETR_{Ord} is sufficient to find a solution of the recursion. \square

Now, let us explain how to extend the forcing relation beyond the atomic formulas. A class \mathbb{P} -name $\dot{\Gamma}$ is simply a class of pairs $\langle \rho, p \rangle$, where ρ is a \mathbb{P} -name and $p \in \mathbb{P}$; so, it is a class that is a \mathbb{P} -name. The canonical name for the generic filter is the class $\dot{G} = \{ \langle \check{p}, p \rangle \mid p \in \mathbb{P} \}$. We denote by $\mathcal{L}_{\omega, \omega}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$ the usual first-order forcing language, allowing the names $\sigma \in V^{\mathbb{P}}$ as constants and allowing finitely many class name parameters $\dot{\Gamma}_i$. This notation will mesh with the more general notation we introduce in definition 10 for the various infinitary languages.

Definition 4. A class forcing notion \mathbb{P} admits a forcing relation for a collection of first-order formulas, closed under subformulas, if there is a relation \Vdash obeying the following recursive properties, for the formulas on which it is defined:

- (a) The forcing relation \Vdash is defined on atomic formulas $\sigma = \tau$ and $\sigma \in \tau$ in accordance with definition 2;
- (b) For a class name $\dot{\Gamma}$, we have $p \Vdash \sigma \in \dot{\Gamma}$ if and only if there is a dense class of $q \leq p$ for which there is $\langle \tau, r \rangle \in \dot{\Gamma}$ with $q \leq r$ and $q \Vdash \sigma = \tau$.
- (c) $p \Vdash \varphi \wedge \psi$ if and only if $p \Vdash \varphi$ and $p \Vdash \psi$;
- (d) $p \Vdash \neg \varphi$ if and only if there is no $q \leq p$ with $q \Vdash \varphi$; and
- (e) $p \Vdash \forall x \varphi(x)$ if and only if $p \Vdash \varphi(\tau)$ for every \mathbb{P} -name τ .

What we mean is that in each case, if the left-hand side of the equivalence is defined, then the relations appearing on the right hand side are defined and furthermore, are defined in such a way that fulfills the equivalence. We say that \mathbb{P} admits a forcing relation for a formula φ , if it admits a forcing relation defined on a collection of formulas including all instances of $\varphi(\vec{\tau})$ for any choice of \mathbb{P} -names $\vec{\tau}$.

We should like particularly to emphasize that this definition, as well as the definition of what it means to have a forcing relation for atomic formulas, makes no reference whatsoever to generic filters or to genericity of any kind or to the truth of any formula in any forcing extension. Rather, for a forcing relation to exist means, by definition, precisely that there is a class relation $\Vdash_{\mathbb{P}}$ exhibiting the recursive properties expressed in definitions 2 and 4 (and later, for the infinitary language case, definition 11). These recursive properties are entirely first-order expressible properties of the class relation $\Vdash_{\mathbb{P}}$, and the question of whether $\Vdash_{\mathbb{P}}$ has those properties can be answered entirely in the ground model. The question for a given forcing notion \mathbb{P} is whether or not there is indeed a class relation $\Vdash_{\mathbb{P}}$ exhibiting those recursive properties.

Theorem 5. *If a forcing notion \mathbb{P} admits a forcing relation for atomic formulas, then it admits a forcing relation for any particular first-order formula φ in the forcing language $\mathcal{L}_{\omega, \omega}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$.*

This establishes (1) \rightarrow (2) in the main theorem, which although not necessary for the main cycle of implications, will be used in subsequent arguments.

Proof. This is also proved in [HKL⁺16]. This is a theorem scheme, proved by meta-theoretic induction on the formula φ . Given the forcing relation \Vdash defined on atomic formulas $\sigma = \tau$ and $\sigma \in \tau$, one may proceed simply to define the forcing relation for any given first-order formula in the meta-theory. For any particular formula φ , we may form the finite set of subformulas of φ , plus the atomic formulas, and then simply apply the recursive definitions expressed by the requirements of definition 4. For an actual formula, this recursion takes place in the meta-theory, and so we get the desired forcing relation in finitely many recursive steps. \square

Because the induction in the proof of theorem 5 takes place in the meta-theory, the result applies only to formulas φ that are finite in the meta-theory. We cannot use this theorem directly, for example, to get a forcing relation for nonstandard formulas in a nonstandard model of GBC. Nevertheless, the main theorem shows that the principle ETR_{Ord} implies that indeed we can have a uniform forcing relation applying to such formulas in such a model, as in statement (3) of the main theorem.

Lemma 6. *Suppose that \mathbb{P} is a class forcing notion and \Vdash is a forcing relation defined on the relevant formulas.*

- (1) *If $p \Vdash \varphi$ and $q \leq p$, then $q \Vdash \varphi$.*
- (2) *If it is dense below p to force φ , then $p \Vdash \varphi$.*
- (3) *If $\varphi \rightarrow \psi$ is a logical validity and $p \Vdash \varphi$, then $p \Vdash \psi$.*

Proof. Statements (1) and (2) are each proved for atomic formulas by induction on names, and then easily extended to all formulas by induction on formulas. For example, for the negation case of statement (2), if it is dense below p to force $\neg\varphi$, then for any $q \leq p$ there is $r \leq q$ with $r \Vdash \neg\varphi$, which means by (1) that q cannot force φ , and so $p \Vdash \neg\varphi$, as desired.

Statement (3) is proved by induction on proofs with respect to a standard deduction system. It is easy to verify, for example, that forcing respects *modus ponens*:

$$\text{if } p \Vdash \varphi \text{ and } p \Vdash \varphi \rightarrow \psi, \quad \text{then } p \Vdash \psi.$$

Forcing respects substitution, since by definition, $p \Vdash \forall x\varphi(x)$ just in case $p \Vdash \varphi(\tau)$ for any particular name τ . Also, if $p \Vdash \varphi$ and x is not a variable in φ , then $p \Vdash \forall x\varphi$.

One can prove by induction on names that every condition forces every instance of the atomic equality axioms:

$$\begin{aligned} \sigma &= \sigma \\ \sigma &= \tau \rightarrow \tau = \sigma \\ \sigma &= \tau \rightarrow (\tau = \rho \rightarrow \sigma = \rho) \\ \sigma &= \tau \rightarrow (\rho \in \sigma \leftrightarrow \rho \in \tau). \end{aligned}$$

We leave to the reader the further elementary exercises to prove for every condition p that

$$\begin{aligned} p &\Vdash \forall x(\varphi \rightarrow \psi) \rightarrow (\forall x\varphi \rightarrow \forall x\psi) \\ p &\Vdash \varphi \rightarrow (\psi \rightarrow \varphi) \\ p &\Vdash (\neg\psi \rightarrow \neg\varphi) \rightarrow (\varphi \rightarrow \psi) \\ p &\Vdash [\varphi \rightarrow (\psi \rightarrow \theta)] \rightarrow [(\varphi \rightarrow \psi) \rightarrow (\varphi \rightarrow \theta)] \end{aligned}$$

In each case, one can prove the statement by considering the definition of what it means to be a forcing relation and by using statements (1) and (2) in a density argument. Since we have therefore observed that forcing respects the axioms and rules of a complete deduction system, statement (3) now follows by induction on proofs. \square

The lemma also holds by essentially similar arguments for the infinitary languages we introduce in section 5. For example, the forcing relation respects the infinitary conjunction rule and the infinitary quantifier rules.

3. CONSTRUCTING ACTUAL FORCING EXTENSIONS

Before proceeding to the rest of the main theorem, we should like to clarify some issues concerning the forcing relation and class forcing and the construction of forcing extensions. It has been traditional to highlight what has been called the forcing theorem, which explains how the forcing relation interacts with truth in a forcing extension. What we would like to do here is to explain how that part of the forcing theorem is a consequence of the existence of the forcing relation, which we defined here as a solution of the forcing relation recursion. The central question with regard to a class forcing notion \mathbb{P} , therefore, becomes whether indeed one has such a forcing relation; for if one does, then it will interact with truth in the forcing extension in the desired manner.

In order to show this, let us first review how one constructs a forcing extension for a forcing notion \mathbb{P} . We shall explain how to construct forcing extensions of an arbitrary model M of GBC, without assuming that that model is countable or transitive.

Given a class forcing notion \mathbb{P} inside a model M of GBC, let's say that a filter $G \subseteq \mathbb{P}$ is M -generic, if G meets every dense subclass $D \subseteq \mathbb{P}$ that is in M . Suppose that we have a forcing relation \Vdash available in M for \mathbb{P} . For any such G , we may define the following relations on the \mathbb{P} -names available in M :

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

Equivalently, using the Boolean values described in theorem 14, we have defined

$$\begin{aligned} \sigma =_G \tau &\iff \llbracket \sigma = \tau \rrbracket \in G \\ \sigma \in_G \tau &\iff \llbracket \sigma \in \tau \rrbracket \in G. \end{aligned}$$

It follows easily using lemma 6 that $=_G$ is an equivalence relation and a congruence with respect to \in_G . Let $M[G]$ denote the collection of equivalence classes $[\sigma]_{=_G}$, equipped with the relation induced by \in_G . This amounts to the same quotient procedure one undertakes with the Boolean-valued model approach to forcing and the Boolean ultrapower (discussed in [HS]).

It is well-known that $\langle M[G], \in_G \rangle$ is not necessarily a model of ZFC, since with class forcing one can, for example, collapse all cardinals to ω , which of course contradicts ZFC; indeed, one can even collapse M itself to become countable, as will happen with the forcing \mathbb{F}_A that we consider in section 8. Nevertheless, $\langle M[G], \in_G \rangle$ is a structure of some kind and has its theory, whatever it may be, and so it still makes sense to inquire about which statements in that theory are true in $M[G]$ or forced by a condition, and so on.

If we had equipped the forcing relation with a class name \dot{G} , then we define its extension in $M[G]$ by

$$\Gamma([\sigma]) \iff \exists p \in G \quad p \Vdash \sigma \in \dot{G},$$

which is the same as saying that there is a dense class of conditions $q \leq p$ for which there is $\langle \rho, r \rangle \in \dot{G}$ with $q \leq r$ and $q \Vdash \sigma = \rho$. For example, if we use the canonical name of the generic filter \dot{G} , then we have $G([\sigma])$ just in case there is $p \in G$ with $p \Vdash \sigma \in \dot{G}$. If G is M -generic, this is equivalent to saying that there is some $p, r \in G$ with $p \leq r$ and $p \Vdash \sigma = \check{r}$.

Theorem 7. *Suppose that M is a model of GBC with a class forcing notion \mathbb{P} that admits a forcing relation in M for a first-order formula φ , which is standard in the meta-theory, in a forcing language with finitely many class names $\dot{\Gamma}_i$ and $G \subseteq \mathbb{P}$ is M -generic. Then*

$$\langle M[G], \in_G, \Gamma_0, \dots, \Gamma_n \rangle \models \varphi([\tau]) \quad \text{if and only if} \quad \exists p \in G \ p \Vdash \varphi(\tau).$$

Proof. This is proved by induction in the meta-theory on the formula φ , which is why we need φ to be standard in the meta-theory. The theorem is true for atomic formulas basically by definition of the relations $=_G$ and \in_G and the definitions of the extensions Γ_i in $M[G]$ for any of the class names $\dot{\Gamma}_i$. If the theorem statement is true for φ , then it is also true for $\neg\varphi$, since the class of conditions p with $p \Vdash \varphi$ or $p \Vdash \neg\varphi$ is easily seen to be dense, and so there is such a condition in G . If the theorem statement is true for φ and ψ , then it is also true for the conjunction $\varphi \wedge \psi$, because the class of conditions p that either force both φ and ψ or else force one of the negations $\neg\varphi$ or $\neg\psi$ is dense, and so there is such a $p \in G$, which then gives the theorem for the conjunction. If the theorem is true for φ , then it is true for $\forall x \varphi(x)$, since the class of conditions p that either force $\neg\varphi(\tau)$ for some τ or force $\varphi(\tau)$ for all such names τ , is dense, and so there is such a condition $p \in G$, which then gives the theorem for $\forall x \varphi(x)$, as desired. \square

The argument can just as easily handle assertions φ in the infinitary languages $\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, A)$, introduced in definitions 10 and 11, provided that the formulas are actually well-founded in the meta-theory.

We should like to emphasize that although theorem 7 works with nonstandard models M of GBC, we do still have a standardness assumption for the formula φ , since one proves the theorem by induction on formulas in the meta-theory. We shall later explain how one can move beyond that to a setting accommodating nonstandard formulas φ by using theorem 9 to construct suitable truth predicates in the extension.

4. FORCING RELATION AS NAME FOR A TRUTH PREDICATE

We shall now prove the equivalence (3) \leftrightarrow (5) in the main theorem. Let us begin by making precise what it means to have a truth predicate for first-order truth.

Definition 8. A *truth predicate* for first-order truth (also known as a *satisfaction class*), with a class parameter A , is a class T consisting of pairs $\langle \varphi, \vec{a} \rangle$, where φ is a formula in the first-order language of set theory augmented with a predicate symbol \hat{A} for the class A and \vec{a} is a valuation mapping the free variables of φ to corresponding set parameters, such that the following recursion is satisfied:

- (a) T judges the truth of $\{=, \in, \hat{A}\}$ -atomic statements correctly:

$$\begin{aligned} T(x = y, \langle a, b \rangle) & \quad \text{if and only if} \quad a = b \\ T(x \in y, \langle a, b \rangle) & \quad \text{if and only if} \quad a \in b \\ T(x \in \hat{A}, a) & \quad \text{if and only if} \quad a \in A \end{aligned}$$

- (b) T performs Boolean logic correctly:

$$\begin{aligned} T(\varphi \wedge \psi, \vec{a}) & \quad \text{if and only if} \quad T(\varphi, \vec{a}) \text{ and } T(\psi, \vec{a}) \\ T(\neg\varphi, \vec{a}) & \quad \text{if and only if} \quad \neg T(\varphi, \vec{a}) \end{aligned}$$

(c) T performs quantifier logic correctly:

$$\mathsf{T}(\forall x \varphi, \vec{a}) \quad \text{if and only if} \quad \forall b \mathsf{T}(\varphi, b \frown \vec{a})$$

When a truth predicate exists in a model of GBC, then it is unique, since there cannot be a least formula where the disagreement occurs. Nevertheless, classical results of Krajewski [Kra74] show that there are (necessarily nonstandard) models of ZFC that admit different incompatible truth predicates as ZFC-amenable classes (see discussion and further results in [HY]). Similar reasoning produces models of ZFC with different incompatible forcing relations for a given forcing notion, even for set forcing, each of them ZFC-amenable but not jointly ZFC-amenable.

Let us introduce the notation $\text{op}(\sigma, \tau)$ for the canonical name for the ordered pair of σ and τ ; technically, if we use the Kuratowski definition of the ordered pair, then $\text{op}(\sigma, \tau) = \{\{\langle \sigma, \mathbb{1} \rangle\}, \mathbb{1}\}, \{\langle \sigma, \mathbb{1} \rangle, \langle \tau, \mathbb{1} \rangle\}, \mathbb{1}\}$.

Theorem 9. *For any class forcing notion \mathbb{P} , and any finitely many class name parameters $\dot{\Gamma}_0, \dots, \dot{\Gamma}_m$, the following are equivalent:*

- (i) \mathbb{P} admits a uniform forcing relation for the first-order forcing language $\mathcal{L}_{\omega, \omega}(\in, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$.
- (ii) There is a class name $\dot{\mathsf{T}}$ and a forcing relation defined on the following statement, for which

$$\mathbb{1} \Vdash \dot{\mathsf{T}} \text{ is a truth predicate for } \mathcal{L}_{\omega, \omega}(\in, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m).$$

The main lesson of this argument should be that a uniform forcing relation essentially *is* the name of a truth predicate.

Proof. (i) \rightarrow (ii) Suppose that we have a uniform forcing relation \Vdash for \mathbb{P} in the first-order forcing language $\mathcal{L}_{\omega, \omega}(\in, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$. Let

$$\dot{\mathsf{T}} = \{ \langle \text{op}(\check{\varphi}, \vec{\tau}), p \rangle \mid p \Vdash \varphi(\vec{\tau}) \},$$

which is a class name for the pairs $\langle \varphi, \vec{a} \rangle$ such that $\varphi(\vec{a})$ will be true in the extension. Since we have the atomic forcing relation, it follows by theorem 5 that we may extend \Vdash to cover any particular first-order statement using any other fixed class parameter. In particular, we can talk about whether specific assertions about $\dot{\mathsf{T}}$ are forced, even if $\dot{\mathsf{T}}$ is not one of the $\dot{\Gamma}_i$ (that is, theorem 5 allows us to apply the forcing relation to specific individual assertions about $\dot{\mathsf{T}}$). Using this, we claim that

$$\mathbb{1} \Vdash \dot{\mathsf{T}} \text{ is a first-order truth predicate.}$$

This is a first-order-expressible property of the predicate $\dot{\mathsf{T}}$.

The main point is that the recursive requirements on the forcing relation are exactly what one needs to prove that $\dot{\mathsf{T}}$ is (forced to be) a truth predicate. The atomic case is clear, because a condition p forces (the canonical name of) $\langle x \in y, \langle \sigma, \tau \rangle \rangle$ to be in $\dot{\mathsf{T}}$ if and only if $p \Vdash \sigma \in \tau$, and similarly with $=$ and with any class name parameter $\dot{\Gamma}$. Since $p \Vdash \neg \varphi$ just in case no stronger condition $q \leq p$ forces φ , it follows that the class of conditions q either forcing φ or forcing $\neg \varphi$, but never both, is dense, and so $\mathbb{1}$ forces that $\dot{\mathsf{T}}$ exhibits the negation requirement for a truth predicate. Similarly, since any condition p forces $\bigwedge_i \varphi_i$ just in case it forces each φ_i separately, it follows that $\mathbb{1}$ forces that $\dot{\mathsf{T}}$ obeys the conjunction rule. And since $p \Vdash \forall \vec{x} \varphi(\vec{x})$ just in case $p \Vdash \varphi(\vec{\tau})$ for all particular $\vec{\tau}$, it follows that $\mathbb{1}$ forces that $\dot{\mathsf{T}}$ obeys the quantifier requirement.

(ii \rightarrow i) Conversely, suppose that \mathbb{P} admits a class \mathbb{P} -name \dot{T} and a forcing relation \Vdash for which

$$\mathbb{1} \Vdash \dot{T} \text{ is a } \mathcal{L}_{\omega, \omega}(\in, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)\text{-truth predicate.}$$

We take it as part of this assertion, following definition 4, that the forcing relation \Vdash is defined on all atomic formulas. For assertions φ in the forcing language, let us define a new relation \Vdash as follows:

$$p \Vdash \varphi(\vec{\tau}) \iff p \Vdash \text{op}(\check{\varphi}, \vec{\tau}) \in \dot{T}.$$

That is, we say that an assertion $\varphi(\tau)$ is forced via \Vdash by a condition p , if p forces that it is true according to the truth predicate \dot{T} . On the right hand side of this definition, we use the extension of the atomic forcing relation \Vdash as in theorem 5 to any particular first-order expression in the forcing language.

First, let's notice that the two forcing relations agree on atomic formulas. Specifically, because of what it means to be a truth predicate, this relation gets the right answer for atomic formulas, since we have that $p \Vdash \sigma = \tau$ if and only if $p \Vdash \text{op}((x = y), \text{op}(\sigma, \tau)) \in \dot{T}$, which holds if and only if $p \Vdash \sigma = \tau$; and similarly, $p \Vdash \sigma \in \tau$ if and only if $p \Vdash \text{op}((x \in y), \text{op}(\sigma, \tau)) \in \dot{T}$, which holds if and only if $p \Vdash \sigma \in \tau$. So \Vdash and \Vdash agree on atomic assertions (and therefore we didn't really need a new forcing symbol).

We have to verify that this relation that we have just defined satisfies the required recursive properties to be a forcing relation. We have already observed that it agrees with the forcing relation for atomic assertions. Since $\mathbb{1}$ forces that \dot{T} is a truth predicate, it follows that for any particular assertion $\varphi(\vec{\tau})$, there are densely many conditions q either forcing that $\varphi(\vec{\tau})$ is true according to \dot{T} or that $\neg\varphi(\vec{\tau})$ is true according to \dot{T} , but never both. From this, it follows that $p \Vdash \neg\varphi(\vec{\tau})$ just in case no stronger condition forces $\varphi(\vec{\tau})$. Similarly, the truth-predicate requirements on conjunctions lead to the fact that $p \Vdash \varphi \wedge \psi$ if and only if $p \Vdash \varphi$ and $p \Vdash \psi$. And $p \Vdash \forall x \varphi(x)$ just in case $p \Vdash \varphi(\tau)$ for all \mathbb{P} -names τ . So we've got a uniform forcing relation, as desired. \square

Theorem 9 establishes the equivalence (3) \leftrightarrow (5) in the main theorem on a case-by-case basis for any class forcing notion \mathbb{P} . A similar argument shows that to have a forcing relation \Vdash defined on a fragment of the forcing language, such as the collection of subformulas of a given formula or the formulas of a certain complexity, then $\mathbb{1}$ forces that the corresponding predicate \dot{T} is a truth predicate on the same language fragment. In the case of the infinitary languages introduced in section 5, an analogue of theorem 9 provides a name for a truth predicate for the language consisting of all ground-model assertions in that infinitary language.

5. THE INFINITARY LANGUAGES

Let us now explain how to extend the forcing relation concept to the case of various infinitary languages. These languages are unified as instances of the following general definition.

Definition 10. Assume $\kappa, \lambda \leq \text{Ord}$ are infinite cardinals or Ord itself.

- (a) The language $\mathcal{L}_{\kappa, \lambda}(\in, \hat{A})$ has a signature consisting of the binary relation \in , a unary predicate symbol \hat{A} and sufficiently many (e.g. $\lambda^{<\kappa}$ many) variable symbols x_i . The formulas of the language are obtained from the

atomic formulas $x = y$, $x \in y$ and $x \in \hat{A}$ by closing under negation $\neg\varphi$, under conjunctions $\bigwedge_{j \in J} \varphi_j$ of size $|J| < \kappa$, provided that there are fewer than $\max(\omega, \lambda)$ many free variables collectively in the φ_j , and under quantification $\forall \vec{x} \varphi(\vec{x})$ by quantifier blocks $\vec{x} = \langle x_i \mid i \in I \rangle$ of size $|I| < \lambda$.

- (b) The forcing language $\mathcal{L}_{\kappa, \lambda}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$, for any class forcing notion \mathbb{P} , augments the previous language with all the various \mathbb{P} -names as constant symbols and predicate symbols $\dot{\Gamma}_i$ for finitely many class names, closing under conjunctions of size less than κ and quantifier blocks of size less than λ .

This general definition has natural special cases, with which we shall be concerned:

- The usual first-order language of set theory $\mathcal{L}_{\omega, \omega}(\in)$;
- the usual first-order forcing language $\mathcal{L}_{\omega, \omega}(\in, V^{\mathbb{P}}, \dot{G})$;
- the quantifier-free infinitary forcing language $\mathcal{L}_{\text{Ord}, 0}(\in, V^{\mathbb{P}}, \dot{G})$;
- the partial infinitary language $\mathcal{L}_{\text{Ord}, \omega}(\in, \hat{A})$;
- the full infinitary language $\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, \hat{A})$.
- the full infinitary forcing language $\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$.

Every formula in each of these languages is determined by its well-founded parse tree, which specifies at each node how the formula was constructed from its various subformulas, and one may undertake inductive proofs on formulas by means of these parse trees. Indeed, let us simply identify every formula φ with its parse tree.

Extending definitions 2 and 4, we now define what it means to have uniform forcing relations for these languages.

Definition 11. If \mathbb{P} is a class forcing notion and $\mathcal{L}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$ is one of the forcing languages mentioned above, then we say that \mathbb{P} admits a uniform forcing relation for this language, if there is a relation \Vdash obeying the following recursive properties:

- (a) The forcing relation \Vdash is defined on all atomic formulas $\sigma = \tau$, $\sigma \in \tau$, and $\sigma \in \dot{\Gamma}_i$ and fulfills the requirements expressed by definitions 2 and 4 for these atomic formulas;
- (b) for conjunctions in the language, $p \Vdash \bigwedge_{j \in J} \varphi_j$ if and only if $p \Vdash \varphi_j$ for each j ;
- (c) $p \Vdash \neg\varphi$ if and only if there is no $q \leq p$ with $q \Vdash \varphi$;
- (d) if the language allows quantification, then $p \Vdash \forall x \varphi(x)$ if and only if $p \Vdash \varphi(\tau)$ for every \mathbb{P} -name τ ; and
- (e) if the language allows infinitary blocks of quantifiers, then $p \Vdash \forall \vec{x} \varphi(\vec{x})$, where $\vec{x} = \langle x_i \mid i \in I \rangle$, if and only if $p \Vdash \varphi(\vec{\tau})$ for all sequences of \mathbb{P} -names $\vec{\tau} = \langle \tau_i \mid i \in I \rangle$.

If we regard disjunction $\bigvee_j \varphi_j$ as an abbreviation for $\neg \bigwedge_j \neg \varphi_j$, then the definition requires that $p \Vdash \bigvee_j \varphi_j$ if and only if there are densely many $q \leq p$ for each of which there is some j such that $q \Vdash \varphi_j$, allowing different j for different q . Similarly, if we regard $\exists x \varphi(x)$ as an abbreviation for $\neg \forall x \neg \varphi(x)$, then the definition requires that $p \Vdash \exists x \varphi(x)$ if and only if there are densely many $q \leq p$ for each of which there is some τ such that $q \Vdash \varphi(\tau)$.

It turns out that every class forcing relation for atomic formulas can be extended to a uniform forcing relation on the quantifier-free infinitary forcing language. Indeed, the proof of theorem 12 shows that the expressive power of the forcing language $\mathcal{L}_{\text{Ord},0}(\in, V^{\mathbb{P}}, \dot{G})$ does not actually exceed the expressive power of the atomic equality assertions $\dot{a} = \dot{b}$, and this is why having a forcing relation for merely the atomic assertions suffices to provide a forcing relation for $\mathcal{L}_{\text{Ord},0}(\in, V^{\mathbb{P}}, \dot{G})$. This theorem will be used in section 6

Theorem 12. *If a class forcing notion \mathbb{P} admits a forcing relation for atomic formulas, then it admits a uniform forcing relation in the quantifier-free infinitary forcing language $\mathcal{L}_{\text{Ord},0}(\in, V^{\mathbb{P}}, \dot{G})$.*

Proof. This argument adapts the main ideas of [HKL⁺16, lemma 5.2, theorem 5.5], generalizing from the context there of a countable transitive model of set theory M to our development here of class forcing as an internal GBC matter. The difference is that in the countable-transitive-model case, one is able to define a forcing relation externally by reference to what is true in the various extensions $M[G]$, and then use that relation in inductive arguments; but here, we must define a suitable class relation internally to GBC and then prove that it fulfills the forcing-relation recursion, even in uncountable or nonstandard models and with nonstandard formulas, whose truth conditions are not necessarily sensible in the meta-theory.

To begin, we assume that \mathbb{P} is a class forcing notion for which we have a forcing relation \Vdash for the atomic formulas. We aim to define a forcing relation $p \Vdash \varphi$ for the sentences φ in the quantifier-free infinitary forcing language $\mathcal{L}_{\text{Ord},0}(\in, V^{\mathbb{P}}, \dot{G})$. In order to do so, we shall recursively assign to each sentence φ in the quantifier-free infinitary forcing language $\mathcal{L}_{\text{Ord},0}(\in, V^{\mathbb{P}}, \dot{G})$ an atomic formula of the form $\dot{a}_\varphi = \dot{b}_\varphi$ and then define the desired forcing relation \Vdash as follows:

$$p \Vdash \varphi \quad \longleftrightarrow \quad p \Vdash \dot{a}_\varphi = \dot{b}_\varphi. \quad (*)$$

Basically, we aim to find suitable atomic equality assertions $\dot{a}_\varphi = \dot{b}_\varphi$ that track the truth and forcing conditions for any given quantifier-free infinitary assertion φ , and we shall then prove that the relation \Vdash defined by $(*)$ is indeed a forcing relation.

The assignment $\varphi \mapsto \dot{a}_\varphi = \dot{b}_\varphi$ will be the end result of a certain transfinite syntactic translation process, which gradually reduces a given sentence φ in $\mathcal{L}_{\text{Ord},0}(\in, V^{\mathbb{P}}, \dot{G})$ to sentences appearing earlier in the (well-founded) translation hierarchy that is implicit in our construction, eventually reaching the atomic equalities $\dot{a}_\varphi = \dot{b}_\varphi$ as irreducible terminal nodes. At each step, the translation process will respect the forcing-relation requirements of definition 11.

For the first step of the translation, we systematically eliminate use of the \dot{G} predicate, except for the check names, by applying the following transformation whenever σ is not itself a check name:

$$\sigma \in \dot{G} \quad \mapsto \quad \bigvee_{p \in \mathbb{P} \cap V_{\text{rank}(\sigma)+1}} (\check{p} \in \dot{G} \wedge \sigma = \check{p})$$

Since it is not difficult to see by induction on names that $\mathbb{1} \Vdash \sigma \in \check{V}_{\text{rank}(\sigma)+1}$ for any name σ , the idea of this transformation is that the only way σ can name a condition in \dot{G} is if it is naming one of the conditions $p \in \mathbb{P} \cap V_{\text{rank}(\sigma)+1}$. It follows that $\mathbb{1}$ forces the equivalence of the assertion $\sigma \in \dot{G}$ with its translation at the right. The result of the transformation is a formula for which the only occurrences of the \dot{G} predicate are with check names of the form $\check{p} \in \dot{G}$.

Next, we systematically apply the infinitary de Morgan laws to push all remaining negations to the bottom of the parse tree, so that they appear, if at all, immediately in front of atomic formulas.

$$\begin{aligned} \neg \bigvee_{i \in I} \varphi_i &\mapsto \bigwedge_{i \in I} \neg \varphi_i \\ \neg \bigwedge_{i \in I} \varphi_i &\mapsto \bigvee_{i \in I} \neg \varphi_i \end{aligned}$$

These transformations are logical validities and therefore, if the forcing relation is defined on the formulas at the right, then it can be legitimately extended to the formulas at the left.

After this, we eliminate most of the remaining negations by systematically applying the following reductions.

$$\begin{aligned} \sigma \neq \tau &\mapsto \sigma \not\subseteq \tau \vee \tau \not\subseteq \sigma \\ \sigma \not\subseteq \tau &\mapsto \bigvee_{\langle \rho, r \rangle \in \sigma} (\check{r} \in \dot{G} \wedge \rho \notin \tau) \\ \sigma \notin \tau &\mapsto \bigwedge_{\langle \rho, r \rangle \in \tau} (\check{r} \notin \dot{G} \vee \sigma \neq \rho) \end{aligned}$$

What we intend here is that the reductions are applied iteratively, until they can no longer be applied; an easy inductive argument on names shows that the reduction process eventually terminates. It is easy to see in each case that $\mathbb{1}$ forces the equivalence of each of these negated atomic formulas with its translation.

The translation process mentioned so far reduces any given sentence to a positive infinitary Boolean combination (using iterated infinitary conjunction and disjunction) of formulas of the form $\sigma = \tau$, $\sigma \in \tau$, $\check{p} \in \dot{G}$ and $\check{p} \notin \dot{G}$. In order to eliminate all but the atomic equalities, we now apply the following reductions

$$\begin{aligned} \sigma \in \tau &\mapsto \tau = \tau \cup \{ \langle \sigma, \mathbb{1} \rangle \} \\ \check{q} \in \dot{G} &\mapsto \{ \langle \emptyset, q \rangle \} = \{ \langle \emptyset, \mathbb{1} \rangle \} \\ \check{q} \notin \dot{G} &\mapsto \{ \langle \emptyset, q \rangle \} = \emptyset, \end{aligned}$$

whose equivalences are in each case forced by $\mathbb{1}$.

Thus, we have transformed every sentence of the quantifier-free infinitary forcing language $\mathcal{L}_{\text{Ord},0}(\in, V^{\mathbb{P}}, \dot{G})$ to a positive Boolean combination of atomic equalities $\sigma = \tau$. We shall now systematically apply further reductions to eliminate the need for conjunctions and disjunctions, and thereby reduce every infinitary sentence φ to a single atomic equality $\dot{a}_\varphi = \dot{b}_\varphi$.

If φ_i has already been mapped to the atomic equality $\dot{a}_i = \dot{b}_i$, then we eliminate the infinitary conjunction by the transformation

$$\bigwedge_{i \in I} \varphi_i \mapsto \{ \langle \text{op}(\check{i}, \dot{a}_i), \mathbb{1} \rangle \mid i \in I \} = \{ \langle \text{op}(\check{i}, \dot{b}_i), \mathbb{1} \rangle \mid i \in I \}.$$

It is easy to see that a condition p forces the atomic equality $\dot{a} = \dot{b}$ at the right if and only if $p \Vdash \dot{a}_i = \dot{b}_i$ for all $i \in I$.

Infinitary disjunctions are a little more troublesome, but we can eliminate them by means of the following transformation.

$$\bigvee_{i \in I} \varphi_i \quad \mapsto \quad \dot{a} = \dot{b},$$

where

$$\begin{aligned} \dot{u} &= \{ \langle \text{op}(\check{i}, \dot{a}_i), \mathbb{1} \rangle \mid i \in I \} \cup \{ \langle \text{op}(\check{i}, \dot{b}_i), \mathbb{1} \rangle \mid i \in I \} \\ \dot{u}_j &= \{ \langle \text{op}(\check{i}, \dot{a}_i), \mathbb{1} \rangle \mid i \in I \} \cup \{ \langle \text{op}(\check{i}, \dot{b}_i), \mathbb{1} \rangle \mid i \in I, i \neq j \} \\ \dot{a} &= \{ \langle u_j, \mathbb{1} \rangle \mid j \in I \}, \\ \dot{b} &= \{ \langle u_j, \mathbb{1} \rangle \mid j \in I \} \cup \{ \langle \dot{u}, \mathbb{1} \rangle \}. \end{aligned}$$

The idea is that the names \dot{u}_j each remove what might be a redundancy from \dot{u} , and so the names \dot{a} and \dot{b} will name the same set just in case there is such a redundancy. Note that if $p \Vdash \dot{a}_j = \dot{b}_j$ for some particular j , then $p \Vdash \dot{u}_j = \dot{u}$, because of the corresponding redundancy in \dot{u} , and consequently $p \Vdash \dot{a} = \dot{b}$. Conversely, if $p \Vdash \dot{a} = \dot{b}$, then $p \Vdash \dot{u} \in \dot{a}$, and so there are densely many $q \leq p$ for which there is some j such that $q \Vdash \dot{u} = \dot{u}_j$ and consequently $q \Vdash \dot{a}_j = \dot{b}_j$. In short, $p \Vdash \dot{a} = \dot{b}$ if and only if there are densely many $q \leq p$ with $q \Vdash \dot{a}_i = \dot{b}_i$ for some particular i . Thus, our transformation respects the desired forcing relation requirement for disjunctions.

We now complete the proof of the theorem. We have described a translation of any sentence φ in the quantifier-free infinitary forcing language $\mathcal{L}_{\text{Ord},0}(\in, V^{\mathbb{P}}, \dot{G})$ to a corresponding atomic equality $\dot{a}_\varphi = \dot{b}_\varphi$. This translation process was an ordinary set-like recursion on formulas, which can be undertaken in GBC without the need for any ETR-like assumption. We defined the forcing relation by $(*)$ above, namely, $p \Vdash \varphi$ if and only if $p \Vdash \dot{a}_\varphi = \dot{b}_\varphi$. This relation \Vdash , we claim, obeys the recursive requirements of definition 11. This is proved by induction on the translation order. If \Vdash is a forcing relation on all formulas appearing before φ in the translation process, then because as we have observed, each step of the translation process respects the requirements of the forcing relation, it follows that \Vdash is also a forcing relation on the sentence φ . So we have defined a uniform forcing relation on $\mathcal{L}_{\text{Ord},0}(\in, V^{\mathbb{P}}, \dot{G})$, as desired. \square

The statement in the conclusion of theorem 12, if stated for all class forcing notions—that is, the assertion that every class forcing notion \mathbb{P} admits a uniform forcing relation for the quantifier-free infinitary forcing language—is equivalent to all the other statements made in the main theorem, because the previous theorem shows that it is implied on a case-by-case basis by statement (1) and conversely it clearly also implies statement (1). So we could actually have added this statement to the main theorem as yet another equivalent assertion.

The method of the previous theorem extends to the case of limited-complexity infinitary assertions, which allow finitely many quantifiers at the top level of the parse tree. That is, if one has a forcing relation for atomic formulas, then one can have a forcing relation for any formula having finitely many quantifiers at the front, alternating in any desired pattern, followed by a quantifier-free infinitary assertion.

One first gets the forcing relation for the quantifier-free infinitary language as above, and then applies induction in the meta-theory as in theorem 5.

Let us now prove the implication $(12) \rightarrow (4)$. Meanwhile, the implications $(4) \rightarrow (3) \rightarrow (2) \rightarrow (1)$ are essentially immediate.

Theorem 13. *The principle ETR_{Ord} of elementary transfinite recursion for class recursions of length Ord implies that every class forcing notion \mathbb{P} admits a uniform forcing relation for assertions in the forcing language $\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$, allowing any fixed class names $\dot{\Gamma}_i$.*

Proof. For any class forcing notion \mathbb{P} , we get the forcing relation for atomic formulas by theorem 3. And now the point is that the requirements stipulated for the rest of the forcing relation by definition 11 amount exactly to a recursion on formulas. That is, to have a forcing relation is to have the solution of a certain recursion, the recursion expressed by definition 11. Since every formula has a well-founded parse tree, which has some ordinal rank, we may organize this recursive definition of the uniform forcing relation as a class recursion of length Ord , by recursing on the rank of the parse tree of the formula. Thus, ETR_{Ord} provides a solution of this recursion, which is the desired uniform infinitary forcing relation. \square

We should like to call attention to the contrast between theorems 12 and 13. In the case of theorem 12, we constructed a uniform forcing relation for quantifier-free infinitary assertions in the forcing language $\mathcal{L}_{\text{Ord}, 0}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$ on a case-by-case basis for the forcing \mathbb{P} , without any need for ETR_{Ord} . This was an ordinary set-like recursion on formulas. In theorem 13, however, the recursion is no longer set-like, because for the quantifier case, the relation $p \Vdash \forall x \varphi(x)$ reduces to a proper class of smaller instances $p \Vdash \varphi(\tau)$, even when φ is merely first-order. Thus, this is no longer an ordinary recursion on sets, but a class recursion of length Ord , which can be undertaken using ETR_{Ord} .

In the ZFC context, set theorists have grown accustomed to having separate forcing relations for each formula φ , since it is not possible ever to have a uniform forcing relation as a definable class, as this would lead quickly to a definable truth predicate, contrary to Tarski's theorem. Nevertheless, the main theorem shows that if every class forcing notion \mathbb{P} has its atomic forcing relations, then ETR_{Ord} holds, and therefore every class forcing relation \mathbb{P} has a fully uniform forcing relation, even in the full infinitary forcing language $\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, V^{\mathbb{P}}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$. And therefore also we get the accompanying truth predicates that these relations provide. This does not violate Tarski's theorem, because while the forcing relations exist as classes, they are not first-order definable classes.

For that reason, the ‘definability lemma’ terminology in the literature, used to refer to the assertion that the class forcing relations exist, is somewhat misleading, because this terminology should not be interpreted as asserting that the forcing relations are actually (first-order) definable classes. Indeed, certain definable class forcing notions, such as those used in section 8, cannot have first-order definable forcing relations, even in the case of the forcing relations for atomic formulas only, although those relations can exist as GBC classes. It is true, however, that when the forcing relations exist, then they are the unique class relations satisfying the forcing relation recursion, and therefore they are always first-order implicitly definable, in the sense of [HL16], and therefore also second-order definable relations.

6. BOOLEAN COMPLETIONS

In this section we prove the equivalence (1) \leftrightarrow (6). This argument basically follows [HKL⁺16, theorem 5.5], generalizing it from the context of countable transitive models of set theory to the general case of arbitrary models of GBC, using our recursion conception of what it means in GBC to have a forcing relation. The equivalence holds on a case-by-case basis for each class forcing notion \mathbb{P} separately.

Theorem 14. *For any class forcing notion \mathbb{P} , the following are equivalent:*

- (i) \mathbb{P} admits a forcing relation for atomic formulas.
- (ii) \mathbb{P} admits a uniform forcing relation for quantifier-free infinitary formulas in the forcing language $\mathcal{L}_{\text{Ord},0}(\in, \dot{G})$.
- (iii) \mathbb{P} admits a Boolean completion, a set-complete class Boolean algebra \mathbb{B} into which \mathbb{P} densely embeds.

Proof. The implication (i \rightarrow ii) is provided by theorem 12.

For (ii \rightarrow iii), assume that \mathbb{P} has a uniform forcing relation for quantifier-free infinitary formulas in the forcing language $\mathcal{L}_{\text{Ord},0}(\in, \dot{G})$. For φ, ψ in this language, define $\varphi \approx \psi$ just in case $\mathbb{1} \Vdash \varphi \leftrightarrow \psi$. This is a class equivalence relation, and the quotient $\mathbb{B} = \mathcal{L}_{\text{Ord},0}(\in, \dot{G}) / \approx$, inheriting the logical structure of the language itself, is easily seen to be a set-complete Boolean algebra, as in the Lindenbaum algebra, where $\neg[\varphi]_{\approx} = [\neg\varphi]_{\approx}$ and $\bigwedge_{i \in I} [\varphi_i]_{\approx} = [\bigwedge_{i \in I} \varphi_i]_{\approx}$. (We use Scott's trick to represent each equivalence class canonically by the set of its minimal-rank members, in order to avoid the inconvenience that an equivalence class would otherwise be a proper class.) Finally, \mathbb{P} embeds densely into \mathbb{B} by the map $p \mapsto [\check{p} \in \dot{G}]_{\approx}$.

Lastly, for (iii \rightarrow i), suppose that a class forcing notion \mathbb{P} has a Boolean completion \mathbb{B} , a set-complete class Boolean algebra \mathbb{B} with a dense embedding $i : \mathbb{P} \rightarrow \mathbb{B}$. We shall prove that there is a forcing relation for atomic formulas. To do so, we define the following Boolean values, by recursion on names:

$$\begin{aligned} \llbracket \sigma \in \tau \rrbracket &= \bigvee_{\langle \rho, r \rangle \in \tau} (\llbracket \sigma = \rho \rrbracket \wedge i(r)) \\ \llbracket \sigma = \tau \rrbracket &= \llbracket \sigma \subseteq \tau \rrbracket \wedge \llbracket \tau \subseteq \sigma \rrbracket \\ \llbracket \sigma \subseteq \tau \rrbracket &= \bigwedge_{\langle \rho, r \rangle \in \sigma} (\neg i(r) \vee \llbracket \rho \in \tau \rrbracket) \end{aligned}$$

This is a straightforward recursion on names, which we may undertake in GBC, without need for any ETR assumption. We now define the corresponding forcing relation $p \Vdash \varphi$ if and only if $i(p) \leq \llbracket \varphi \rrbracket$, for atomic φ .

It remains to check that this relation is indeed a forcing relation for atomic truth, that is, that it fulfills the recursive requirements of definition 2. This is an exercise in the usual Boolean-valued reasoning. For example, if $p \Vdash \sigma \in \tau$, then

$$i(p) \leq \llbracket \sigma \in \tau \rrbracket = \bigvee_{\langle \rho, r \rangle \in \tau} (\llbracket \sigma = \rho \rrbracket \wedge i(r)),$$

and so every $p' \leq p$ has $i(p')$ compatible with $\llbracket \sigma = \rho \rrbracket \wedge i(r)$ for some $\langle \rho, r \rangle \in \tau$, which means there is some $q \leq p'$ with $q \leq r$ and $i(q) \leq \llbracket \sigma = \rho \rrbracket$. Thus, there is a dense class of $q \leq p$ with $q \Vdash \sigma = \rho$ for some $\langle \rho, r \rangle \in \tau$ with $q \leq r$, as desired. The other properties are similar and left for the reader. \square

7. ETR_{Ord} IMPLIES TRUTH PREDICATE FOR $\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, A)$

In this section, we shall prove the implication $(12) \rightarrow (9)$. The implication $(9) \rightarrow (8)$ is essentially immediate. We begin by providing the natural generalization of the truth predicate concept of definition 8 to the infinitary context as follows.

Definition 15. If $\kappa, \lambda \leq \text{Ord}$ are cardinals or Ord itself and A is a fixed class parameter, then a *truth predicate* for the $\mathcal{L}_{\kappa, \lambda}(\in, \hat{A})$ language of set theory with a predicate for A is a class T consisting of pairs $\langle \varphi, \vec{a} \rangle$, where φ is a formula in that language and \vec{a} is a valuation mapping the free variables of φ to corresponding set parameters, such that the following recursion is satisfied:

- (a) T judges the truth of $\{=, \in, \hat{A}\}$ -atomic statements correctly:

$$T(x = y, \langle a, b \rangle) \quad \text{if and only if} \quad a = b$$

$$T(x \in y, \langle a, b \rangle) \quad \text{if and only if} \quad a \in b$$

$$T(x \in \hat{A}, a) \quad \text{if and only if} \quad a \in A$$

- (b) T performs Boolean logic correctly:

$$T\left(\bigwedge_{i \in I} \varphi_i, \vec{a}\right) \quad \text{if and only if} \quad T(\varphi_i, \vec{a}) \text{ for all } i \in I$$

$$T(\neg \varphi, \vec{a}) \quad \text{if and only if} \quad \neg T(\varphi, \vec{a})$$

- (c) T performs quantifier logic correctly:

$$T(\forall \vec{x} \varphi, \vec{a}) \text{ if and only if } \forall \vec{b} \, T(\varphi, \vec{b} \smallfrown \vec{a}),$$

where $\vec{b} \smallfrown \vec{a}$ is the valuation extending \vec{a} by mapping the variables of \vec{x} to the objects listed by \vec{b} .

This generalizes definition 8 from the case of the first-order language $\mathcal{L}_{\omega, \omega}(\in, A)$ to the infinitary languages, such as $\mathcal{L}_{\text{Ord}, \omega}(\in, A)$ and $\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, A)$. Let us subsume a few minor syntactical details, such as the fact that in (b) when referring to $T(\varphi_i, \vec{a})$, we should restrict the valuation \vec{a} to the free variables of φ_i , if these were fewer than in $\bigwedge_i \varphi_i$.

Theorem 16. *Assume the principle of elementary transfinite recursion ETR_{Ord} for recursions of length Ord . Then there is a truth predicate for $\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, A)$, with any class parameter A .*

Proof. The point is that the existence of a such a truth predicate for this infinitary logic is an elementary transfinite recursion of length Ord , defined by recursion on formulas. The formulas each have an ordinal rank coming from the rank of their parse trees, and we may define the truth of such a formula by reference to the truth of its constituent pieces. \square

In light of the results of section 13, we should not expect to get Ord -iterated truth predicates for the infinitary logic $\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, A)$, or even an ω -iterated truth predicate for $\mathcal{L}_{\text{Ord}, \omega}(\in, A)$, since those recursions have length Ord^2 or $\text{Ord} \cdot \omega$, and ETR_{Ord} is not able to prove that such recursions have a solution. Indeed, it follows from the separation results of that section that the existence of an ω -iterated truth predicate for $\mathcal{L}_{\text{Ord}, \omega}(\in)$ is strictly stronger in consistency strength than ETR_{Ord} .

8. FORCING THEOREM IMPLIES TRUTH PREDICATE FOR $\mathcal{L}_{\text{Ord}, \omega}$

In this section, we shall prove the implication (1) \rightarrow (8) in the main theorem. In order to do so, let us define a particular class forcing notion \mathbb{F}_A , from whose atomic forcing relation we shall be able to extract a truth predicate. This is an adaptation and generalization of the forcing defined in [HKL⁺16, definition 2.4], and the argument here is based on the analysis of [HKL⁺16, section 7]. Let A be any proper class parameter. Let $\text{Coll}(\omega, V)$ be the class partial order having as conditions all finite injective partial functions $f : \omega \rightarrow V$. This is the usual forcing to add a bijection from ω to V . To form \mathbb{F}_A , we shall augment the forcing $\text{Coll}(\omega, V)$ with additional conditions

$$\mathbb{F}_A = \text{Coll}(\omega, V) \sqcup \{e_{n,m} \mid n, m \in \omega\} \sqcup \{a_n \mid n \in \omega\}$$

where for $f \in \text{Coll}(\omega, V)$ we define

$$\begin{aligned} f \leq e_{n,m} &\iff f(n) \in f(m), \text{ and} \\ f \leq a_n &\iff f(n) \in A. \end{aligned}$$

In other words, the condition $e_{n,m}$ is by definition precisely the supremum of the conditions f with $f(n) \in f(m)$, and a_n is by definition the supremum of the conditions f for which $f(n) \in A$. Let us also take $\mathbb{1} = \emptyset \in \text{Coll}(\omega, V)$ to be above all the new conditions $e_{n,m}$ and a_n .

In short, we define the forcing \mathbb{F}_A to be basically the collapse forcing $\text{Coll}(\omega, V)$, but augmented with some additional conditions that are the suprema of certain useful classes of conditions. In particular, $\text{Coll}(\omega, V)$ is a dense subclass of \mathbb{F}_A , and consequently a generic filter $G \subseteq \mathbb{F}_A$ will be fully determined by $G \cap \text{Coll}(\omega, V)$, which will be generic for that forcing. Nevertheless, class forcing differs from set forcing in that a dense subclass of the forcing does not necessarily give rise to the same forcing extensions (see [HKS17, section 5]). The reason is that the extra conditions, such as $e_{n,m}$ and a_n above, allow one to form names in the larger forcing that are not equivalent to any name in the smaller forcing, even though it is dense. Basically, we augmented the forcing with those extra conditions precisely so that we could use those conditions to form \mathbb{F}_A -names for objects that we could not have been able to name in the $\text{Coll}(\omega, V)$ forcing alone.

Let us illustrate with a few examples. Define the name

$$\dot{\varepsilon} = \{ \langle \text{op}(\check{n}, \check{m}), e_{n,m} \rangle \mid n, m \in \omega \}.$$

Notice that $\dot{\varepsilon}$ is a set-sized name—actually, it is countable—even though it seems to carry information about what will ultimately happen with a proper class of conditions f . The reason it can do so is precisely because of the supremum conditions $e_{n,m}$ that we added to \mathbb{F}_A .

Similarly, we may define the name

$$\dot{A} = \{ \langle \check{n}, a_n \rangle \mid n \in \omega \}.$$

This is not a name for the class A itself, but rather it is the name for the copy of A on ω that will be induced by the generic bijection. That is, \dot{A} is the name for the collection of $n \in \omega$ that will correspond to an element of A by the generic bijection. Again, this is a set-sized name, actually countable, which we can form precisely because of the extra conditions a_n that we had added to the forcing.

Lastly, without using the augmented part of \mathbb{F}_A , let us define for each set a the name

$$\dot{n}_a = \{ \langle \check{k}, (n \mapsto a) \rangle \mid k < n \in \omega \},$$

where $(n \mapsto a) = \{ \langle n, a \rangle \} \in \text{Coll}(\omega, V)$ is the finite partial function defined on only the point n and mapping it to a . Thus, \dot{n}_a is the name of the set of numbers k that are less than the natural number n that will get mapped to a by the generic bijection. In other words, since every natural number is the set of smaller numbers, \dot{n}_a is the name of the number n that will get mapped to a .

The following theorem is a generalization of [HKL⁺16, theorem 7.3] to the internal GBC treatment of class forcing and also to allow a class parameter.

Theorem 17. *If the forcing \mathbb{F}_A , for a fixed class parameter A , admits forcing relations for atomic formulas, then there is a truth predicate for $\mathcal{L}_{\text{Ord}, \omega}(\in, A)$.*

Proof. Fix the class A and suppose that the forcing \mathbb{F}_A admits forcing relations for atomic formulas. It follows by theorem 12 that we have a uniform forcing relation for quantifier-free infinitary assertions in the forcing language $\mathcal{L}_{\text{Ord}, 0}(\in, V^{\mathbb{F}_A})$.

Our proof will make use of the following purely syntactical translation $\varphi \mapsto \varphi^*$, where $\varphi \in \mathcal{L}_{\text{Ord}, \omega}(\in, A)$ and $\varphi^* \in \mathcal{L}_{\text{Ord}, 0}(\in, V^{\mathbb{F}_A})$. The translation is defined by recursion on φ as follows:

$$\begin{aligned} (x \in y)^* &= x \dot{\in} y \\ (x = y)^* &= x = y \\ (x \in A)^* &= x \in \dot{A} \\ (\varphi \wedge \psi)^* &= \varphi^* \wedge \psi^* \\ (\neg \varphi)^* &= \neg \varphi^* \\ \left(\bigwedge_i \varphi_i \right)^* &= \bigwedge_i \varphi_i^* \\ (\forall x \varphi)^* &= \bigwedge_{m \in \omega} \varphi^*(\check{m}). \end{aligned}$$

The interesting cases are $(x \in y)^*$ and $(\forall x \varphi)^*$. The idea is that the translation is transforming truth assertions about the structure $\langle V, \in, A \rangle$ to the corresponding truth assertions about the structure $\langle \omega, \dot{\in}, \dot{A} \rangle$, which intuitively will be made isomorphic by the generic bijection $a \mapsto (\dot{n}_a)_G$, where G is generic for \mathbb{F}_A . The point is that \in assertions in V will correspond to $\dot{\in}$ relations in the latter structure, and universal assertions $\forall x \varphi$ in V will correspond to countable conjunctions in the latter structure, since every object will be placed at some n by the generic bijection, mapping a to $(\dot{n}_a)_G$. Intuitively, the translation aims at establishing the equivalence

$$\langle V, \in, A \rangle \models \varphi(a) \iff V[G] \models (\langle \omega, \dot{\in}_G, \dot{A}_G \rangle \models \varphi^*((\dot{n}_a)_G)).$$

Formalizing this equivalence at the outset, however, is somewhat problematic, because the left-hand side is not expressible unless we already have the desired truth predicate, and the right-hand side makes truth assertions in $V[G]$, which is not a model of any decent set theory, as we have collapsed V to become countable, and so it isn't clear to what extent $V[G]$ has a truth predicate for the structure $\langle \omega, \dot{\in}_G, \dot{A}_G \rangle$. One might hope alternatively for the right hand side to work directly with the structure $\langle \omega, \dot{\in}_G, \dot{A}_G \rangle$, using an externally defined truth predicate; but this will not work

properly with non-standard models, since there will be nonstandard formulas φ^* whose truth in that structure will not have a clear meaning in the meta-theory (although this approach does work when the original model is transitive).

Nevertheless, without formalizing the equivalence displayed above, we may instead take it as inspiring an idea that we are able to formalize, specifically, the definition in GBC of a certain predicate $\text{Tr}(\varphi, \vec{a})$, which we shall prove is a truth predicate. The idea is similar to that of theorem 9. Namely, if φ is an infinitary formula in the language $\mathcal{L}_{\text{Ord}, \omega}(\in, A)$ and $\vec{a} = \langle a_0, \dots, a_k \rangle$ is a valuation of the (finitely many) free variables of φ , then we define

$$\text{Tr}(\varphi, \vec{a}) \iff \mathbb{1} \Vdash_{\mathbb{F}_A} \varphi^*(\dot{n}_{a_0}, \dots, \dot{n}_{a_k}).$$

The definition uses the forcing relation in \mathbb{F}_A only in the case of the formulas φ^* , which are quantifier-free infinitary assertions in the forcing language $\mathcal{L}_{\text{Ord}, 0}(\in, V^{\mathbb{F}_A})$, and we have a uniform forcing relation for that language by theorem 12. Let us prove that this is indeed a truth predicate.

Lemma 17.1. *For any formula $\varphi \in \mathcal{L}_{\text{Ord}, \omega}(\in, A)$, any set a and any condition p ,*

$$p \Vdash \varphi^*(\dot{n}_{a_0}, \dots, \dot{n}_{a_k}) \quad \text{if and only if} \quad \mathbb{1} \Vdash \varphi^*(\dot{n}_{a_0}, \dots, \dot{n}_{a_k}).$$

Proof. We proceed by induction on the formula φ . It is important here that the only names appearing as parameters in φ^* have the form \dot{n}_a for some a ; for example, the lemma is not true if one allows parameters of the form \dot{n} , since if $p(n)$ is defined, then p will force things about \dot{n} that other conditions will not. For the atomic case, it is easy to verify that

$$\begin{aligned} p \Vdash \dot{n}_a = \dot{n}_b & \quad \text{if and only if} \quad a = b \\ p \Vdash \dot{n}_a \dot{\in} \dot{n}_b & \quad \text{if and only if} \quad a \in b \end{aligned}$$

because we can simply extend p to a condition that decides the exact values of \dot{n}_a and \dot{n}_b . It follows that if $p \Vdash \dot{n}_a = \dot{n}_b$ or $p \Vdash \dot{n}_a \dot{\in} \dot{n}_b$, then $\mathbb{1}$ also forces that statement, which verifies the atomic case of the lemma. Next, for conjunctions, suppose that $p \Vdash (\bigwedge_i \varphi_i(\dot{n}_a))^*$, which by the definition of the \star -translation means $p \Vdash \varphi_i^*(\dot{n}_a)$ for every i (for notational simplicity, we consider just one parameter \dot{n}_a). By the induction hypothesis, this means $\mathbb{1} \Vdash \varphi_i^*(\dot{n}_a)$ for every i and so $\mathbb{1} \Vdash (\bigwedge_i \varphi_i(\dot{n}_a))^*$, as desired. For negation, if $p \Vdash \neg \varphi^*(\dot{n}_a)$, it means there is no $q \leq p$ with $q \Vdash \varphi^*(\dot{n}_a)$. But now, in fact there is no q at all with $q \Vdash \varphi^*(\dot{n}_a)$, for if there were such a q , then by induction we would have $\mathbb{1} \Vdash \varphi^*(\dot{n}_a)$, contrary to our assumption that $p \Vdash \neg \varphi^*(\dot{n}_a)$.

Finally, we consider the quantifier case. Since universal quantifiers get \star -translated to conjunctions over the natural numbers, let us suppose that $p \Vdash \bigwedge_m \varphi^*(\dot{m}, \dot{n}_a)$. If $\mathbb{1}$ does not force this conjunction, then there is some condition forcing the negation, and by strengthening further we find a condition q forcing $\neg \varphi^*(\dot{m}, \dot{n}_a)$ for some particular m . Notice that $\varphi^*(\dot{m}, \dot{n}_a)$ does not fall under the induction assumption, since the name \dot{m} is not of the form \dot{n}_a . Nevertheless, by strengthening further if necessary, we may assume $q(m)$ is defined and thus $q \Vdash \dot{m} = \dot{n}_b$ for some b . It follows by lemma 6 that $q \Vdash \neg \varphi^*(\dot{n}_b, \dot{n}_a)$. By the induction hypothesis (and the negation case), it follows now that $\mathbb{1} \Vdash \neg \varphi^*(\dot{n}_b, \dot{n}_a)$. By strengthening p to a condition p' forcing $\dot{n}_b = \dot{m}$ for some m , this contradicts our assumption that $p \Vdash \bigwedge_m \varphi^*(\dot{m}, \dot{n}_a)$. \square

We now proceed to prove that the class Tr defined before the lemma is a truth predicate for the language $\mathcal{L}_{\text{Ord},\omega}(\in, A)$. We need simply to verify the recursive requirements of definition 15. For the atomic case, the equivalences

$$\begin{aligned} a = b & \quad \text{if and only if} \quad \mathbb{1} \Vdash \dot{n}_a = \dot{n}_b \\ a \in b & \quad \text{if and only if} \quad \mathbb{1} \Vdash \dot{n}_a \dot{\in} \dot{n}_b \\ a \in A & \quad \text{if and only if} \quad \mathbb{1} \Vdash \dot{n}_a \in \dot{A}, \end{aligned}$$

follow essentially by design from the definitions of \dot{n}_a , $\dot{\in}$ and \dot{A} .

Next, let's check that our definition performs Boolean logic correctly. For simplicity, allow us to consider the case of just one parameter a rather than \vec{a} . In the case of conjunctions, we have $\text{Tr}(\bigwedge_i \varphi_i, a)$ just in case $\mathbb{1} \Vdash \bigwedge_i \varphi_i^*(\dot{n}_a)$, which holds if and only if $\mathbb{1} \Vdash \varphi_i^*(\dot{n}_a)$ for each i , which by induction is equivalent to $\text{Tr}(\varphi_i, a)$ for all i , as desired.

Consider next negation. If $\text{Tr}(\varphi, a)$ holds, then $\mathbb{1} \Vdash \varphi^*(\dot{n}_a)$, and so it is not the case that $\mathbb{1} \Vdash \neg \varphi^*(\dot{n}_a)$ and so $\text{Tr}(\neg \varphi, a)$ fails. Conversely, if $\text{Tr}(\varphi, a)$ does not hold, then $\mathbb{1} \nVdash \varphi^*(\dot{n}_a)$. It follows by the lemma that no condition p can force $\varphi^*(\dot{n}_a)$ and consequently $\mathbb{1} \Vdash \neg \varphi^*(\dot{n}_a)$, which implies that $\text{Tr}(\neg \varphi, a)$ holds, thereby fulfilling the desired negation requirement.

Finally, consider the quantifier case. Suppose that $\text{Tr}(\forall x \varphi(x), a)$ holds. By definition, this means $\mathbb{1} \Vdash (\forall x \varphi(x))^*(\dot{n}_a)$, and by definition of the \star -translation, this means $\mathbb{1} \Vdash \bigwedge_{m \in \omega} \varphi^*(\dot{m}, \dot{n}_a)$, which is equivalent to $\mathbb{1} \Vdash \varphi^*(\dot{m}, \dot{n}_a)$ for every m . For any set b , there is a dense class of conditions $q \leq \mathbb{1}$ with $q \Vdash \dot{n}_b = \dot{m}$ for some m and consequently $q \Vdash \varphi^*(\dot{n}_b, \dot{n}_a)$. Since the q are dense, this implies $\mathbb{1} \Vdash \varphi^*(\dot{n}_b, \dot{n}_a)$, and so we have $\text{Tr}(\varphi, b \cap a)$ for every set b , as required. Conversely, if $\text{Tr}(\varphi, b \cap a)$ for every set b , then $\mathbb{1} \Vdash \varphi^*(\dot{n}_b, \dot{n}_a)$ for every b , and from this it follows that $\mathbb{1} \Vdash \varphi^*(\dot{m}, \dot{n}_a)$ for any particular m , since it is dense to force $\dot{m} = \dot{n}_b$ for some b . Thus, $\mathbb{1} \Vdash \bigwedge_m \varphi^*(\dot{m}, \dot{n}_a)$ and consequently $\text{Tr}(\forall x \varphi, a)$, as desired. \square

9. TRUTH PREDICATE FOR $\mathcal{L}_{\text{Ord},\omega}(\in, A)$ IMPLIES ORD-ITERATED TRUTH PREDICATE FOR $\mathcal{L}_{\omega,\omega}(\in, A)$

In this section, we prove the implication (8) \rightarrow (10). To begin, let us define what it means to have an Ord-iterated truth predicate. The idea is to have a truth predicate that applies not only to statements in the language of set theory, but to statements about (earlier stages of) truth in that language. So the Ord-iterated truth predicate will make truth assertions in ordinal stages.

Definition 18. An Ord-iterated truth predicate for first-order truth, with a class parameter A , is a class Tr consisting of triples $\langle \beta, \varphi, \vec{a} \rangle$, where β is an ordinal, φ is a formula in the first-order language of set theory augmented with a predicate for A and also with a trinary predicate symbol $\hat{\text{Tr}}$ to be used for (iterated) truth assertions, and \vec{a} is a valuation mapping the free variables of φ to corresponding parameters, such that the following recursion is satisfied:

(a) Tr judges the truth of $\{=, \in, \hat{A}\}$ -atomic statements correctly:

$$\begin{aligned} \text{Tr}(\beta, x = y, \langle a, b \rangle) & \quad \text{if and only if} \quad a = b \\ \text{Tr}(\beta, x \in y, \langle a, b \rangle) & \quad \text{if and only if} \quad a \in b \\ \text{Tr}(\beta, x \in \hat{A}, a) & \quad \text{if and only if} \quad a \in A \end{aligned}$$

- (b) Tr judges atomic assertions of the truth predicate self-coherently:

$$\begin{aligned} & \text{Tr}(\beta, \hat{\text{Tr}}(x, y, z), \langle \alpha, \varphi, \vec{a} \rangle) \\ & \text{if and only if } \alpha < \beta \text{ and } \text{Tr}(\alpha, \varphi, \vec{a}) \end{aligned}$$

- (c) Tr performs Boolean logic correctly:

$$\begin{aligned} & \text{Tr}(\beta, \varphi \wedge \psi, \vec{a}) \quad \text{if and only if} \quad \text{Tr}(\beta, \varphi, \vec{a}) \text{ and } \text{Tr}(\beta, \psi, \vec{a}) \\ & \text{Tr}(\beta, \neg \varphi, \vec{a}) \quad \text{if and only if} \quad \neg \text{Tr}(\beta, \varphi, \vec{a}) \end{aligned}$$

- (d) Tr performs quantifier logic correctly:

$$\text{Tr}(\beta, \forall x \varphi, \vec{a}) \quad \text{if and only if} \quad \forall b \text{Tr}(\beta, \varphi, b \hat{\sim} \vec{a})$$

Note the crucial clause (b), which insists that atomic truth assertions made at stage β can refer only to earlier stages of truth, $\alpha < \beta$. When the formula φ has no free variables, then to improve readability we shall write $\text{Tr}(\beta, \varphi)$ in place of $\text{Tr}(\beta, \varphi, \langle \rangle)$, since in this case the valuation is empty.

When $\text{Tr}(\beta, \varphi, \vec{a})$, then we shall say that $\varphi[\vec{a}]$ is declared true by the predicate at stage β . In this way, we can extract from the uniform truth predicate the sequence of individual truth predicates $\text{Tr}_\beta(\varphi, \vec{a})$ for the truth assertions made at each stage β . These predicates perform as one would want for the non-uniform manner of iterated truth, where one makes truth assertions always only at a particular stage of truth, with a separate predicate symbol for each stage. We should like to emphasize, however, that the uniform truth predicate is stronger than this, because it allows formulas to quantify over the stages of truth. For example, the uniform iterated truth predicate allows one to express various liar-type sentences, which have an interesting nature with respect to the iterated truth predicate.

Consider the following instance, a sentence σ whose truth value will systematically and endlessly alternate between true and false at successive ordinal stages. Specifically, the sentence σ expresses that there is no immediately preceding stage at which σ is true. To formalize this in the language of the iterated truth predicate, let \top be any tautologically valid sentence, such as $\forall x x = x$, and notice that in light of requirement (ii) above, the assertion $\hat{\text{Tr}}(\alpha, \top)$ is judged true at stage β if and only if $\alpha < \beta$; we can use this feature to quantify in effect over the ‘earlier’ stages of truth. By the Gödel-Carnap fixed-point theorem, which asserts of every formula $\varphi(x)$ that there is a sentence σ such that our theory proves $\sigma \leftrightarrow \varphi(\sigma)$, it follows that there is a sentence σ for which:

$$\sigma \longleftrightarrow \neg \exists \alpha [\text{Tr}(\alpha, \sigma) \wedge \text{Tr}(\alpha, \top) \wedge \neg \text{Tr}(\alpha + 1, \top)],$$

and furthermore, this equivalence is valid at every stage of truth. Said plainly, σ asserts that it isn’t the case that σ is true at a stage which is the largest ‘previous’ stage. At stage 0, there is no largest previous stage, and so indeed σ is true at stage 0. Thus, it will become false at stage 1, and therefore true again at stage 2. Whenever σ is true at stage β , then it will become false at stage $\beta + 1$ and true again at stage $\beta + 2$. The sentence σ will be true at limit ordinal stages, since there is no largest previous stage. So σ flips between true and false for the iterated truth predicate, being true at every even ordinal stage and false at every odd ordinal stage, even though it is the very same sentence being considered each time.

Similar constructions yield sentences σ that are true exactly at the stages in a class A or exactly at certain other stages in a way that is convenient. Such kind of sentences are not generally possible in the weaker language having only

truth predicates Tr_β for each stage of truth separately, without the possibility of quantifying over the index β , since in that language a sentence φ would refer to only finitely many of those predicates, and therefore the truth or falsity of the sentence would stabilize at stages of truth beyond those explicitly mentioned in φ .

Theorem 19. *For any class A , if there is a truth predicate for the infinitary language $\mathcal{L}_{\text{Ord},\omega}(\in, A)$, then there is an Ord-iterated truth predicate in the first-order language $\mathcal{L}_{\omega,\omega}(\in, \hat{\text{Tr}}, A)$.*

This will establish implication (8) \rightarrow (10) in the main theorem.

Proof. Suppose that $T(\psi, \vec{a})$ is a (non-iterated) truth predicate for the infinitary language $\mathcal{L}_{\text{Ord},\omega}(\in, A)$, where A is a fixed class parameter. Note that in light of lemma 24, we don't actually need the parameters \vec{a} , since every set is definable, and so the entire semantic content of the truth predicate is actually contained in its sentences. Nevertheless, we shall carry on with the parameters \vec{a} .

Using the truth predicate T , we shall define another truth predicate, an Ord-iterated truth predicate $\text{Tr}(\beta, \varphi, \vec{a})$ for first-order assertions φ in the language $\mathcal{L}_{\omega,\omega}(\in, \hat{\text{Tr}}, A)$. In order to do so, we first define a certain syntactic translation

$$(\beta, \varphi) \mapsto \varphi_\beta^*,$$

where β is an ordinal and φ is a formula in the first-order language with an iterated truth predicate $\mathcal{L}_{\omega,\omega}(\in, \hat{\text{Tr}}, A)$. The resulting formula φ_β^* is an assertion in $\mathcal{L}_{\text{Ord},\omega}(\in, A)$, without any truth predicate. This translation is defined by induction on β and φ , as follows:

- (a) Atomic formulas not mentioning the truth predicate are not changed by the translation:

$$\begin{aligned} (x \in y)_\beta^* &= (x \in y), \\ (x = y)_\beta^* &= (x = y), \\ (x \in \hat{A})_\beta^* &= (x \in \hat{A}). \end{aligned}$$

- (b) $\hat{\text{Tr}}(x, y, z)_\beta^*$ is the assertion

$$\bigvee_{\substack{\xi < \beta \\ \psi \in \mathcal{L}_{\omega,\omega}(\in, \hat{\text{Tr}}, \hat{A})}} \left["x = \xi" \wedge "y = \psi" \wedge \exists \vec{a} \text{ valuation}_\psi(z, \vec{a}) \wedge \psi_\xi^*(\vec{a}) \right].$$

- (c) $(\varphi \wedge \psi)_\beta^* = \varphi_\beta^* \wedge \psi_\beta^*$.

- (d) $(\forall x \varphi)_\beta^* = \forall x \varphi_\beta^*$.

The key translation occurs in step (b), replacing atomic instances of the truth predicate with certain infinitary formulas. The formula $\text{valuation}_\psi(z, \vec{a})$ asserts that z is a valuation mapping the variables that happen to be free in ψ to the objects in \vec{a} . Basically, the stage β translation of the atomic truth assertion $\hat{\text{Tr}}(x, y, z)$ is the assertion that (i) x is some stage ξ less than β ; (ii) y is some formula ψ ; and (iii) z is a valuation of the free variables of that formula to objects \vec{a} for which $\psi_\xi^*(\vec{a})$ holds. In particular, in statement (b) we used the expressions " $x = \xi$ " and " $y = \psi$ ", and by this we mean the formulas $\theta_\xi(x)$ and $\theta_\psi(y)$ provided by lemma 24, which as we mentioned does not involve using ξ or ψ (or any code for ψ) as a parameter. We should also like to emphasize that we do not need ETR_{Ord} in order to make this recursive definition of the translation, since it is an ordinary transfinite recursion of length Ord on sets, not a class recursion.

We now define our proposed iterated truth predicate $\text{Tr}(\beta, \varphi, \vec{a})$ to hold if and only if $T(\varphi_\beta^*, \vec{a})$. We claim that this relation fulfills the requirements to be an iterated truth predicate.

Because T is a truth predicate for assertions in $\mathcal{L}_{\text{Ord}, \omega}(\in, A)$, it follows easily that our iterated truth predicate Tr works correctly on $\{\in, =, \hat{A}\}$ -atomic formulas and on Boolean combinations and quantifiers. The only difficult part is to verify that we have judged the atomic truth assertions themselves in a self-coherent manner. What we need to prove is that $\text{Tr}(\beta, \hat{\text{Tr}}(x, y, z), \langle \alpha, \varphi, \vec{a} \rangle)$ holds if and only if $\alpha < \beta$ and $\text{Tr}(\alpha, \varphi, \vec{a})$.

To see this, notice that $\text{Tr}(\beta, \hat{\text{Tr}}(x, y, z), \langle \alpha, \varphi, \vec{a} \rangle)$ holds, by the definition of Tr , just in case $T(\hat{\text{Tr}}(x, y, z)_\beta^*, \langle \alpha, \varphi, \vec{a} \rangle)$ holds, which is equivalent to

$$T\left(\bigvee_{\substack{\xi < \beta \\ \psi \in \mathcal{L}_{\omega, \omega}(\in, \hat{\text{Tr}}, \hat{A})}} ("x = \xi" \wedge "y = \psi" \wedge \psi_\xi^*), \langle \alpha, \varphi, \vec{a} \rangle\right).$$

Since T is a truth predicate, we may unwrap the meaning of this disjunction to see that this holds if and only if $\alpha < \beta$ and $T(\varphi_\alpha^*, \vec{a})$, since the disjuncts can be realized only with $\alpha = \xi$ and $\varphi = \psi$ themselves. By the definition of Tr , this is equivalent to $\text{Tr}(\alpha, \varphi, \vec{a})$. So we have verified that Tr is an iterated truth predicate for first-order truth with the class parameter A , as desired. \square

One can extract from this argument a corresponding result for having merely a κ -iterated truth predicate, given a truth predicate for the $\mathcal{L}_{\kappa, \omega}$ language of set theory.

Theorem 20. *For any class A , if there is a truth predicate for the infinitary language $\mathcal{L}_{\kappa, \omega}(\in, \hat{A})$, where κ is any uncountable cardinal, then there is an κ -iterated truth predicate in the first-order language $\mathcal{L}_{\omega, \omega}(\in, \hat{\text{Tr}}, \hat{A})$.*

Proof. The point is that to define $\text{Tr}(\beta, \varphi, \vec{a})$ in the previous theorem, we took a disjunction of size $\omega \cdot \beta$, which will still be less than κ for $\beta < \kappa$, if κ is an uncountable cardinal. So exactly the same translation and definition of Tr works up to κ . \square

10. ITERATED TRUTH PREDICATES IMPLY THE ELEMENTARY TRANSFINITE RECURSION PRINCIPLES

In this section, we shall prove implication (10) \rightarrow (12) in the main theorem, and indeed, we shall prove (10) \leftrightarrow (12). To do so, we shall undertake a refinement of the following theorem of Fujimoto [Fuj12] (see also Gitman and Hamkins [GH16]).

Theorem 21. *The principle of elementary transfinite recursion ETR is equivalent over GBC to the existence of iterated truth predicates along any well-founded class relation.*

The forward implication of this is straightforward, as the truth predicate itself is defined as the solution to a certain recursion. The real content of the theorem is the converse, that from any sufficiently iterated truth predicate one can extract a solution of a given recursion. For the full ETR, one must consider recursions of length exceeding Ord , as well as iterated truth predicates of length strictly longer than Ord , for any well-founded class relation.

What we want to prove here is that the equivalence goes through when one restricts the recursions and iterations to length Ord , or indeed, to length Γ for any well-ordered infinite class Γ .

Theorem 22. *For any class well-order Γ , with $\omega^\omega \leq \Gamma$, the principle of elementary transfinite recursion ETR_Γ for recursions of length Γ is equivalent over GBC to the existence of Γ -iterated truth predicates, allowing any class parameter in each case.*

This will establish the equivalence $(10) \leftrightarrow (12)$ in the main theorem, if we consider the case $\Gamma = \text{Ord}$. The expression ω^ω refers to the countable ordinal arising via ordinal exponentiation of ω with itself.

Proof. We follow the proof of [GH16, theorem 8]. Let us emphasize that both statements in the theorem make their assertions universally for all class parameters.

The forward implication is basically straightforward, since the iterated truth predicate itself is defined by a transfinite recursion of length $\omega \cdot \Gamma$, by recursion on formulas. Namely, if we have the predicate $\text{Tr} \upharpoonright \alpha$ up to stage $\alpha \in \Gamma$, then we can define Tr at stage α by a length ω -recursion on formulas, by reference to the partial solution $\text{Tr} \upharpoonright \alpha$. So the entire recursion has length $\omega \cdot \Gamma$. Note that ETR_Γ implies $\text{ETR}_{\Gamma+\Gamma}$, since one need only perform the recursion up to stage Γ , and then define a new recursion for the rest of the way, using the partial solution as a new parameter. And since $\omega^\omega \leq \Gamma$, it follows that $\omega \cdot \Gamma < \Gamma + \Gamma$ (use that $\Gamma = \omega^\omega \cdot \Lambda + \alpha$ for some Λ and some $\alpha < \omega^\omega$, and observe $\omega \cdot \Gamma = \omega \cdot (\omega^\omega \cdot \Lambda + \alpha) = \omega^\omega \cdot \Lambda + \omega \cdot \alpha = \Gamma + \omega \cdot \alpha < \Gamma + \Gamma$). Thus, we can get the Γ -iterated truth predicate from ETR_Γ , as desired.

Conversely, suppose that for class parameter A , we have an iterated truth predicate Tr of length Γ for first-order truth relative to parameter A . Now suppose that we have an instance of ETR_Γ , iterating a formula $\varphi(x, \alpha, S, A)$, where we seek a solution S up to Γ , a class $S \subseteq \Gamma \times V$ for which $S_\alpha = \{x \mid \varphi(x, \alpha, S \upharpoonright \alpha, A)\}$ for every $\alpha < \Gamma$. We claim that using the truth predicate as a class parameter, we may define such a solution S . To do this, we claim first that there is a formula $\bar{\varphi}$ such that if one extracts from Tr the class defined by $\bar{\varphi}$, namely, $S = \{\langle \alpha, x \rangle \mid \text{Tr}(\alpha, \bar{\varphi}, x)\}$, then S is a solution to the recursion of φ along Γ . The formula $\bar{\varphi}$ should simply be chosen so that $\langle V, \in, A, \text{Tr} \upharpoonright \alpha \rangle \models \bar{\varphi}(x, \alpha)$ if and only if $\langle V, \in, A, S \upharpoonright \alpha \rangle \models \varphi(x, \alpha)$, where S is defined as just mentioned using $\bar{\varphi}$. Such a formula $\bar{\varphi}$ exists by the Gödel-Carnap fixed-point lemma: for any e , let $\psi(e, x, \alpha)$ be the assertion $\langle V, \in, A, \{\langle \alpha, x \rangle \mid \text{Tr}(\alpha, e, x)\} \rangle \models \varphi(x, \alpha)$, and then by the usual fixed-point trick find a formula $\bar{\varphi}(x, \alpha)$, for which $\langle V, \in, A, \text{Tr} \upharpoonright \alpha \rangle \models \psi(\bar{\varphi}, x, \alpha) \leftrightarrow \bar{\varphi}(x, \alpha)$. It follows that the class S iteratively defined from Tr by $\bar{\varphi}$ satisfies φ at each step and therefore is a solution to the recursion of φ up to Γ , as desired. \square

An essentially similar idea works with infinitary formulas, provided that $\text{Ord} \cdot \Gamma \leq \Gamma + \Gamma$, which is to say that Γ is at least Ord^ω , with ordinal exponentiation.

Theorem 23. *For any infinite class well-order Γ , of order type at least Ord^ω , the principle $\text{ETR}_\Gamma(\mathcal{L}_{\text{Ord}, \text{Ord}}(\in, A))$ of elementary transfinite recursion for recursions of length Γ in the infinitary language with a class parameter A is equivalent over GBC to the existence of Γ -iterated truth predicates for that language.*

Proof. The forward implication is proved by observing that the iterated truth predicate is precisely a solution of a recursion. The converse implication is proved as in the previous theorem by defining a solution of the recursion by reference to the iterated truth predicate. \square

11. CLASS-JOIN SEPARATION

In this section, we shall prove the implications $(12) \rightarrow (7) \rightarrow (8)$ in the main theorem. To begin, we define the principle of ETR_{Ord} -*foundation* to be the assertion that every instance of elementary transfinite recursion of length Ord either has a solution, or else fails at some least stage $\alpha \leq \text{Ord}$. That is, either there is a solution or there is $\alpha \leq \text{Ord}$ such that for every $\beta < \alpha$ there is a solution of the recursion of length β , but there is no solution of length α . Since the partial solutions are unique when they exist, perhaps some readers expect that one could simply combine those earlier solutions into one uniform solution; but such an argument would appeal to a class-replacement principle that is not provable in GBC. (Consider the difficulty, for example, of combining the various Σ_n -truth predicates into a uniform truth predicate, if the model is ω -standard and has only definable classes.) Meanwhile, the principle of ETR_{Ord} -foundation is a consequence of the principle of Π_1^1 -*foundation*, which asserts that every Π_1^1 -definable class of ordinals has a least element, and both of these principles are true in any well-founded model of GBC. That is, in transitive models, we get it for free.

Let us now define a separation-like principle that we call the *class-join separation principle*. This is the assertion that for any class Φ of $\mathcal{L}_{\text{Ord},\omega}(\in, A)$ -formulas in finitely many free variables, where A is any class parameter, if every formula $\varphi \in \Phi$ admits a truth predicate T_φ , then $\{\vec{a} \mid \exists \varphi \in \Phi \ T_\varphi(\varphi, \vec{a})\}$ exists as a class. The idea is that this class is essentially what we would want to mean by the class $\{\vec{a} \mid \bigvee_{\varphi \in \Phi} \varphi(\vec{a})\}$, defined by a class-sized join. The principle asserts that such class-sized joins can be used to define classes, even when one lacks a uniform truth predicate and only has truth predicates for each formula φ individually. Indeed, if there is a uniform truth predicate application to all $\varphi \in \Phi$, then the instance of the class-join separation principle would follow from the ordinary separation axiom of GBC simply by using that predicate as a parameter. The principle, instead, is about unifying a class-indexed collection of separate truth predicates.

The dual principle, the *class-meet separation principle*, asserts of every Φ as above that $\{\vec{a} \mid \forall \varphi \in \Phi \ T_\varphi(\varphi, \vec{a})\}$ exists. In other words, the principle allows us to use the class-sized conjunction $\bigwedge_{\varphi \in \Phi} \varphi(\vec{x})$ to define a class, provided that we have truth predicates for each individual $\varphi \in \Phi$.

It will be convenient to have the following folklore lemma, showing that every set is definable by a suitable infinitary formula.

Lemma 24. *For every set a , there is a formula $\theta_a(x)$ in the infinitary language $\mathcal{L}_{\text{Ord},\omega}(\in)$ that defines a . More specifically, there is a truth predicate for the class of these formulas θ_a and their subformulas, and with respect to this truth predicate, $\theta_a(x)$ is true only when $x = a$.*

In particular, in any transitive model of set theory, θ_a defines a , meaning that $M \models \theta_a[b]$ just in case $b = a$.

Proof. We define the formulas θ_a by the following \in -recursion:

$$\theta_a(x) = \forall z [z \in x \leftrightarrow \bigvee_{u \in a} \theta_u(z)].$$

The formula asserts that the elements of x are precisely the objects satisfying the definition of an element of a . In GBC we can define a truth predicate for the class of $\theta_a(x)$ and their subformulas simply by extending the class $\{\langle \theta_a(x), a \rangle \mid a \in V\}$

in the natural way to the subformulas of the θ_a , which are these, joins of these and the formulas $z \in x \leftrightarrow \bigvee_{u \in a} \theta_u(z)$. Basically, since we know that we want $\theta_a(x)$ to define $x = a$, we can use that to build the truth predicate. One can verify inductively that this is indeed a truth predicate for these formulas. \square

Subsequently, we shall write simply “ $x = a$ ” for the formula $\theta_a(x)$, with the understanding that a is not appearing here as a parameter, but rather the hereditary \in -structure of a appears essentially in the parse tree of the formula itself.

Theorem 25. *The following are equivalent over GBC:*

- (i) *The principle ETR_{Ord} .*
- (ii) *For every class A , there is a truth predicate for $\mathcal{L}_{\text{Ord}, \omega}(\in, A)$.*
- (iii) *The class-join separation principle plus ETR_{Ord} -foundation.*
- (iv) *The class-meet separation principle plus ETR_{Ord} -foundation.*

Proof. (i) \leftrightarrow (ii) This was established by the results of sections 7 and 10.

(ii) \rightarrow (iii) The principle of ETR_{Ord} -foundation is an immediate consequence of ETR_{Ord} . And once we have a uniform $\mathcal{L}_{\text{Ord}, \omega}(\in, A)$ -truth predicate T , any instance of class-meet separation with predicate A reduces to an ordinary instance of separation relative to the class predicate T .

(iii) \leftrightarrow (iv) This follows easily from the de Morgan law.

(iv) \rightarrow (ii) Assume the class-join separation principle plus ETR_{Ord} -foundation. Consider the recursive definition of a truth predicate for $\mathcal{L}_{\text{Ord}, \omega}(\in, A)$. This is an Ord-recursion, defined by recursion on the rank of the parse trees of the formulas. If the recursion has a solution, then we have the desired truth predicate. If it does not, then by ETR_{Ord} -foundation, the recursion fails at some least rank $\alpha \leq \text{Ord}$. So for every $\beta < \alpha$, we have a uniform truth predicate T_β for the class \mathcal{L}^β consisting of formulas whose parse tree has rank less than β , but there is no such uniform truth predicate covering all formulas of rank less than α . Since we can easily extend a truth predicate for formulas of rank β to rank $\beta + 1$, it follows that α must be a limit ordinal or Ord itself, and so the class of formulas of rank less than α is the union of those of rank less than some $\beta < \alpha$. Let Φ be the class of formulas in \mathcal{L}^α with free variable x . Using the class-join separation principle, we may define a predicate T_α as follows:

$$T_\alpha = \{ (y, z) \mid \bigvee_{\psi \in \Phi} (“y = \psi” \wedge \psi(z)) \}.$$

What we mean by the join is the assertion $\exists \psi \in \Phi T_\psi (“y = \psi” \wedge \psi(z))$, where “ $y = \psi$ ” is the formula $\theta_\psi(y)$ of lemma 24 and where T_ψ is a truth predicate for the conjunction “ $y = \psi” \wedge \psi(z)$. Indeed we have such a truth predicate, because we can simply combine the truth predicate of ψ arising from the fact that it has rank less than α with the truth predicates provided by lemma 24.

Thus, we have included (φ, a) in T_α just in case φ is some formula ψ of rank less than α for which $\psi(z)$ is true. It is now easy to verify that this is indeed a uniform truth predicate for all formulas of rank less than α , contrary to the assumption that there was no such solution at stage α . \square

At bottom, the argument is that ETR_{Ord} -foundation tells you that if the recursion fails, then it does so at a particular stage, and class-join separation allows you to unify the earlier-stage truth predicates into a uniform predicate, showing that that stage was not a failure after all. So the recursion must succeed.

12. CLOPEN DETERMINACY FOR GAMES OF RANK AT MOST $\text{Ord} + 1$

We shall now prove the implications $(12) \rightarrow (11) \rightarrow (8)$, using arguments that amount to a refinement of corresponding results of Gitman and Hamkins in [GH16].

A clopen class game is played on a well-founded class tree $T \subseteq V^{<\omega}$, whose terminal nodes are labeled as a win for one or the other player. The game starts at the root node, which we place at the top, so the play proceeds downward. The players take turns, each subsequently selecting a child node of the current node and then continuing in turn from this node. The game ends when a terminal node is reached—by well-foundedness this must happen at some finite stage—and the winner is determined by the label of that node.

A well-founded class tree T admits a ranking function with a well-ordered class relation Γ , if there is a labeling of the nodes of the tree with elements of that relation, in such a way that every child node has lower rank than its parent. A tree has rank at most $\text{Ord} + 1$, therefore, if we can label the root of the tree with Ord and all other nodes of the tree with ordinals, in such a way that these ordinals descend as one moves down in the tree. A ranking is *continuous*, if it obeys the recursive property that the rank of any node p is exactly the supremum of $\text{rank}(q) + 1$ in Γ for every child q of p . The principle ETR_Γ implies that every tree with a Γ ranking function has a continuous Γ ranking function, but it isn't clear whether one can prove this without an appeal to some fragment of ETR.

Meanwhile, we claim that the question of whether or not a well-founded class tree T has rank at most $\text{Ord} + 1$ is actually a first-order-expressible property of the tree, and furthermore in GBC such trees always admit first-order-definable continuous ranking functions, without requiring any appeal to a fragment of ETR. To see this, notice that if we consider $T \cap V_\theta$, which is a well-founded set-sized tree, then it has an continuous ranking function, since ZFC proves that every well-founded set relation has a continuous ranking function. As θ increases, the rank of any fixed node in $T \cap V_\theta$ never decreases. If every non-root node in T has the property that its ranks in these approximation trees $T \cap V_\theta$ eventually stabilize for large enough θ , then in fact those limit values form an acceptable continuous ranking of the whole tree T , if we should place label Ord or a suitable ordinal on the root node. And conversely, if we are able to rank the whole tree, then those ranks also serve as ranks in the trees $T \cap V_\theta$. So a tree has rank at most $\text{Ord} + 1$ if and only if every non-root node has an eventually stabilizing ordinal rank in $T \cap V_\theta$, which is a first-order property about the tree. Since the assignment to the nodes of the corresponding limit rank value is first-order definable, we thereby achieve a definable continuous ranking of such trees, as we claimed.

The existence of rankings in a game tree is connected with the requirement in some games that a particular player must count down in a well-order during play. This counting-down feature in a game is often a convenient way to ensure that the game is clopen, since the player cannot count down forever, and the outcome of the game is known when the clock runs out. If a game tree has one player counting down in Γ during play, then we can rank the tree with elements of $2 \cdot \Gamma + 1$, since we label the root node with $2 \cdot \Gamma$, and then whenever the count-down player has just announced α , we label that node with $2 \cdot \alpha$, and if it is the other players turn, we label with $2 \cdot \alpha + 1$. Thus, using the fact that $2 \cdot \text{Ord} = \text{Ord}$, a game where one of the players must count down in the ordinals has rank at most $\text{Ord} + 1$.

Definition 26. The principle of *determinacy* for clopen class games of rank at most Γ , a class well-order, is the assertion that for every clopen class game with a game tree of rank at most Γ , one of the players has a winning strategy.

We emphasize that we are referring here to the ordinal rank of the game tree, which is not the same as the game values that would arise in the open determinacy analysis of the game. For example, the game tree can have a very high rank, even if the first player has a winning move on the first move, which would make the game value very low. Meanwhile, the game value for a clopen game, when it exists (it is defined by a class recursion and in GBC there needn't be a solution of that recursion) is bounded by the rank of the tree.

A winning strategy for a player is a function on the game tree, which selects of every parent node a child of it, in such a way that every play of the game that conforms with those choices on that player's moves, leads to a terminal node that is a win for that player.

Theorem 27. *The principle of elementary transfinite recursion ETR_{Ord} for Ord-length recursions implies the principle of determinacy for clopen class games of rank at most $\text{Ord} + 1$. More generally, for any class well-order Γ , the principle ETR_Γ for recursions of length Γ implies determinacy for clopen class games of rank at most $\Gamma + 1$.*

This will establish implication (12) \rightarrow (11) in the main theorem.

Proof. Using the back-propagation method, due originally to Zermelo in his proof of the fundamental theorem of finite games, we shall label every node in the game tree T as a win either for player I or for player II, and these designations will provide a winning strategy for whoever gets their label on the root node, the strategy being: stay on the nodes with your label. To begin, assume ETR_Γ for a class well-order Γ , and suppose T is a well-founded game tree T of rank at most $\Gamma + 1$. Thus, the root node in the tree gets rank Γ , but all other nodes will have a rank below Γ . Using ETR_Γ , we apply the back-propagation method to label the nodes of the tree with player I or player II, by recursion on rank. The terminal nodes, with rank 0, are already labeled for us. If all children of a node t are labeled, then if it is player I's turn to play and there is a child node labeled I, then we place label I on t , and otherwise II; similarly, if it is player II's turn to play, and there is a child node labeled II, then we place label II on t , and otherwise I. By the principle ETR_Γ , this labels all the non-root nodes of the tree. We may now place the corresponding label on the root, labeled following the same back-propagation rule. From this labeling, we can get a winning strategy: whoever has their label on the root node can always stay on their own labels, and thereby win the game. \square

Next, we establish the implication (11) \rightarrow (8) in the main theorem.

Theorem 28. *The principle of determinacy for clopen class games of rank at most $\text{Ord} + 1$ implies the existence of a truth predicate for $\mathcal{L}_{\text{Ord},\omega}(\in, A)$ for any class parameter A .*

Proof. We follow the main ideas of Gitman and Hamkins [GH16, theorem 9], using a natural infinitary analogue of the truth-telling game, where the interrogator counts down in the ordinals. Specifically, consider the truth-telling game for assertions in the logic $\mathcal{L}_{\text{Ord},\omega}(\in, A)$. There are two players, the *interrogator* and the *truth-teller*. At each move, the interrogator issues a challenge in the form of a set Φ of

infinitary formulas $\varphi \in \mathcal{L}_{\text{Ord},\omega}(\in, A)$ and a valuation \vec{a} of their free variables. The interrogator must also state an ordinal α , which will strictly descend during play; we call this the count-down clock. The truth-teller replies to the inquiry by stating of each formula $\varphi \in \Phi$ whether it is true or false at \vec{a} (not necessarily truthfully). If the truth-teller happens to declare an existential formula $\exists x \psi(x)$ to be true, then she is also obligated to provide a witness b and declare that $(\psi, b \smallfrown \vec{a})$ is true. The truth-teller loses, if she should ever explicitly violate the Tarskian recursion. The interrogator loses if the clock runs out (which must happen eventually, in finitely many moves).

This game has rank at most $\text{Ord} + 1$, because of the ordinal count-down clock. Now, we simply argue as in [GH16] that if the truth-teller has a winning strategy, then the truth assertions made by that strategy will be independent of the play, for all plays in which there remains sufficient time on the clock. This can be proved by induction on formulas. It is clearly true for atomic formulas. And if it is true for a formula, it will be true for the negation, since any violation of this can be transformed into a violation of the Tarski recursion. If it is true for a set of formulas, then it will be true for the conjunction, just by taking the supremum of the stabilizing clock values plus one. And the quantifier case is also easy to handle.

Thus, these plays provide a truth predicate for our infinitary language, as desired.

Finally, we need to argue that the interrogator can have no winning strategy. If σ is any strategy for the interrogator, then find an ordinal θ such that V_θ is closed under that strategy, and have the truth-teller play in accordance with truth in V_θ . This will never violate the Tarski recursion and therefore it will defeat σ .

So if clopen determinacy holds for this game, then there is a truth predicate for $\mathcal{L}_{\text{Ord},\omega}(\in, A)$ truth, as desired. \square

We would like to remark that Gitman and Hamkins have pointed out that there is a flaw in their published proof of [GH16, theorem 9] for the implication of clopen determinacy to ETR. While that implication is indeed correct as they state, nevertheless it does not suffice for the interrogator to count down merely in the natural numbers, as they had initially claimed. Rather, the interrogator should count down in the order $\omega \cdot \Gamma$, where Γ is the length of the iteration, and one can prove by induction in this case that the truth assertions made by the truth-teller about the solution up to any stage α are invariant of the play, provided that the count-down clock is at least $\omega \cdot (\Gamma \upharpoonright \alpha)$. Gitman and Hamkins plan to release an updated version of their paper addressing this matter.

13. SEPARATING THE THEOREM FROM OTHER SECOND-ORDER THEORIES

In order to situate the class forcing theorem more precisely in the hierarchy of theories between GBC and KM, let us prove a few theorems that separate ETR_{Ord} from other similar principles in the vicinity.

Theorem 29. *The theory $\text{GBC} + \text{ETR}_{\text{Ord},\omega}$ is strictly stronger in consistency strength than $\text{GBC} + \text{ETR}_{\text{Ord}}$.*

Proof. Assume $\text{GBC} + \text{ETR}_{\text{Ord},\omega}$, and fix a class global well-order. By theorem 22, it follows that we may form an $(\text{Ord} \cdot \omega)$ -iterated truth predicate Tr for first-order truth relative to the fixed global well-order (in the language with a predicate for that order). Using this predicate, consider the GBC model having the same sets as V , but having as classes only those classes that are definable from a proper initial

segment $\text{Tr} \upharpoonright (\text{Ord} \cdot n)$ for some $n < \omega$. This is a GBC model, and furthermore, it satisfies ETR_{Ord} , because if a class A is definable from $\text{Tr} \upharpoonright \text{Ord} \cdot n$, then we have an Ord -iterated truth predicate definable from $\text{Tr} \upharpoonright (\text{Ord} \cdot (n+1))$. So by statement (10) in the main theorem, we have ETR_{Ord} in this model. But this model cannot have its own $(\text{Ord} \cdot \omega)$ -iterated truth predicate, because no such predicate is definable from a proper initial segment of itself. Thus, this is a model of ETR_{Ord} without $\text{ETR}_{\text{Ord} \cdot \omega}$. But furthermore, since the entire collection of classes of our constructed model is coded by a single class in our original model, it follows that $\text{ETR}_{\text{Ord} \cdot \omega}$ implies the consistency of $\text{GBC} + \text{ETR}_{\text{Ord}}$, as desired. \square

Theorem 30. *The theory $\text{GBC} + \text{ETR}_{\text{Ord}}$ has a strictly stronger consistency strength than the theory $\text{GBC} + \text{ETR}_{<\text{Ord}}$, which asserts ETR_α for every ordinal α , provided this latter theory is consistent.*

Proof. We may use a similar argument for this. Let Tr be an Ord -iterated truth predicate for first-order truth, with a fixed class global well-order parameter. Consider the GBC model arising from the sets in V together with any class that is definable from a proper initial segment of the truth predicate Tr_α . This latter GBC model satisfies $\text{GBC} + \text{ETR}_{<\text{Ord}}$ by theorem 22, since one has an α -iterated truth predicate with respect to any class that arises. \square

One may undertake similar arguments to separate many other levels of the ETR_Γ hierarchy. For example, ETR_ω is not provable in GBC , since one can use it to construct a truth predicate for first-order truth. But ETR_ω does not establish ETR_{ω^2} , since ETR_{ω^2} is enough to construct an ω -iterated truth predicate, and then one can take the classes definable from a proper initial segment of it. This will be a GBC model that satisfies ETR_ω , but not ETR_{ω^2} .

14. FINAL REMARKS

We have now proved all the implications necessary to establish the main theorem. Let us make a few final remarks on the topic of the class forcing theorem.

First, we should like to call attention to the fact that only some of the implications in the main theorem hold on a case-by-case basis for the various class forcing notions \mathbb{P} and class parameter A . For example, if \mathbb{P} admits a forcing relation for atomic formulas, then by theorems 5 and 12 we get a scheme of forcing relations for first-order assertions and a uniform forcing relation for the quantifier-free infinitary forcing language $\mathcal{L}_{\text{Ord},0}(\in, V^\mathbb{P}, \dot{\Gamma}_0, \dots, \dot{\Gamma}_m)$, and by theorem 14 we get a Boolean completion of \mathbb{P} , but we do not generally get a uniform forcing relation for the first-order language of set theory or for the stronger infinitary languages. To see this, consider a very nice forcing notion, such as the Easton forcing to control the GCH pattern or even trivial forcing, in a model of GBC having only its definable classes. There is a definable forcing relation for atomic formulas in this model, but there can be no uniform forcing relation for first-order assertions, since from such a relation we could define a truth predicate for the ground model, which is impossible as the model has only definable classes. To get the uniform forcing relation for a class \mathbb{P} , what one needs is an instance of ETR_{Ord} relative to the forcing relation as a class parameter. In order to complete the cycle of implications in the main theorem, therefore, one applies the statements with stronger and more robust class parameters.

Second, an observant reader might have noticed that we established the equivalences of the main theorem mainly by appealing to an instance of class forcing, the forcing \mathbb{F}_A , which probably one would rarely want to perform, since this forcing is highly destructive, collapsing the entire universe to become countable, and not just making all sets countable, but rather making V itself countable. Perhaps the reader may wonder if the strength of the class forcing theorem is ETR_{Ord} simply because one is including these strange notions of forcing, and that with a more well-behaved collection of class forcing, the principle might be weaker.

In a sense, the objection is correct. To see this, consider the *pretame* class forcing notions, a prominent collection of well-behaved class forcing notions (see [HKS17]). When forcing over the countable transitive models of GB, it turns out that pretameness is equivalent to the preservation of the axioms of GB^- , that is, GB without the power set axiom. Furthermore, Maurice Stanley has proved that the class forcing theorem holds outright for all pretame class forcing (see [HKS17] for the case of forcing over a countable transitive model of set theory; the argument works generally in GBC and in fact already in GB^-). The class forcing theorem for pretame forcing, consequently, has no extra consistency strength beyond the base theory.

Meanwhile, if one attempts to move beyond pretame class forcing, then we claim that one is immediately again in the realm of the analysis of this article, with highly destructive forcing notions such as \mathbb{F}_A . The reason is that [HKS17, lemmas 2.6, 2.7] shows that if a class forcing notion \mathbb{P} is not pretame, then there is some cardinal δ and a class \mathbb{P} -name \dot{F} forced to be a surjection from δ to V . In other words, any class forcing notion that is not pretame necessarily collapses the entire universe V to a cardinal. Furthermore, making use of the class \mathbb{P} -name \dot{F} , an easy reworking of the arguments of section 8 shows that for every proper class A , the forcing \mathbb{P} is dense in a notion of class forcing \mathbb{Q}_A that can be used in place of the forcing \mathbb{F}_A in the argument of theorem 17. So if we consider any natural collection of class forcing notions that goes strictly beyond pretame forcing and which includes a forcing notion whenever it includes a dense subclass of that forcing, then the class forcing theorem for this collection will already be as strong as the class forcing theorem for the collection of all class forcing notions.

Let us provide a rough sketch of the reworking idea we mentioned. Suppose that \mathbb{P} is a non-pretame notion of class forcing and that \dot{F} is a \mathbb{P} -name for a surjection from a cardinal δ to Ord . Given a proper class A , we construct the forcing \mathbb{Q}_A from \mathbb{P} in the same manner that we had constructed \mathbb{F}_A from $\text{Coll}(\omega, V)$ in section 8. Namely, we add the conditions $e_{n,m}$ as before, as upper bounds for the conditions forcing $\dot{F}(\check{n}) \in \dot{F}(\check{m})$, but we do this now for all $n, m < \delta$, and we similarly add the conditions a_n for $n < \delta$ as upper bounds for the conditions forcing $\dot{F}(\check{n}) \in A$. In addition, we add a condition $(n \mapsto a)$, for every set a and $n < \delta$, which will be the supremum of all conditions forcing that \check{n} is the least ordinal with $\dot{F}(\check{n}) = \check{a}$. Since \dot{F} is only a surjection and not necessarily a bijection, we define \dot{n}_a so that it is the name of the *least* ordinal $n < \delta$ that will get mapped to a . Under this setup, we may carry out the analogue of the proof of theorem 17, replacing \mathbb{F}_A by \mathbb{Q}_A and ω by δ , respectively. One minor change consists of isolating a least, rather than an arbitrary, ordinal $n < \delta$ for which q forces $\neg\varphi^*(\check{n}, \dot{n}_a)$ in the quantifier case in the proof of lemma 17.1; this ensures that we can choose q and find some set b so that q forces $\check{n} = \dot{n}_b$, as in the corresponding step in the proof of lemma 17.1.

In summary, the class forcing theorem for pretame forcing is provable in GBC, but for any sufficiently robust collection of class forcing notions going beyond the pretame forcing, the class forcing theorem will have the full strength of ETR_{Ord} and all the other statements of the main theorem.

REFERENCES

- [Fuj12] Kentaro Fujimoto. Classes and truths in set theory. *Annals of Pure & Applied Logic*, 163(11):1484–1523, 11 2012.
- [GH16] Victoria Gitman and Joel David Hamkins. Open determinacy for class games. In Andrés E. Caicedo, James Cummings, Peter Koellner, and Paul Larson, editors, *Foundations of Mathematics, Logic at Harvard, Essays in Honor of Hugh Woodin's 60th Birthday*, Contemporary Mathematics. American Mathematical Society, 2016. also available as Newton Institute preprint ni15064.
- [HKL⁺16] Peter Holy, Regula Krapf, Philipp Lücke, Ana Njgomir, and Philipp Schlicht. Class forcing, the forcing theorem and Boolean completions. *The Journal of Symbolic Logic*, 81(4):15001530, 2016.
- [HKS17] Peter Holy, Regula Krapf, and Philipp Schlicht. Characterizations of pretameness and the Ord-cc. 2017. Accepted for Annals of Pure and Applied Logic.
- [HL16] Joel David Hamkins and Cole Leahy. Algebraicity and implicit definability in set theory. *Notre Dame J. Formal Logic*, 57(3):431–439, 2016.
- [HS] Joel David Hamkins and Daniel Seabold. Well-founded Boolean ultrapowers as large cardinal embeddings. pages 1–40. preprint <http://jdh.hamkins.org/boolean-ultrapowers/>.
- [HY] Joel David Hamkins and Ruizhi Yang. Satisfaction is not absolute. *to appear in the Review of Symbolic Logic*, pages 1–34.
- [Kra74] S. Krajewski. Mutually inconsistent satisfaction classes. *Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys.*, 22:983–987, 1974.
- [Kra17] Regula Krapf. *Class forcing and second-order arithmetic*. PhD thesis, University of Bonn, 2017.
- [Sat14] Kentaro Sato. Relative predicativity and dependent recursion in second-order set theory and higher-order theories. *J. Symb. Log.*, 79(3):712–732, 2014.

VICTORIA GITMAN, THE CITY UNIVERSITY OF NEW YORK, CUNY GRADUATE CENTER, MATHEMATICS PROGRAM, 365 FIFTH AVENUE, NEW YORK, NY 10016, USA

E-mail address: vgitman@nylogic.org

URL: <http://boolesrings.org/victoriagitman>

JOEL DAVID HAMKINS, MATHEMATICS, PHILOSOPHY, COMPUTER SCIENCE, THE GRADUATE CENTER OF THE CITY UNIVERSITY OF NEW YORK, 365 FIFTH AVENUE, NEW YORK, NY 10016 & MATHEMATICS, COLLEGE OF STATEN ISLAND OF CUNY, STATEN ISLAND, NY 10314

E-mail address: jhamkins@gc.cuny.edu

URL: <http://jdh.hamkins.org>

PETER HOLY, MATH. INSTITUT, UNIVERSITÄT BONN, ENDENICHER ALLEE 60, 53115 BONN, GERMANY

E-mail address: pholy@math.uni-bonn.de

PHILIPP SCHLICHT, MATH. INSTITUT, UNIVERSITÄT BONN, ENDENICHER ALLEE 60, 53115 BONN, GERMANY

E-mail address: schlicht@math.uni-bonn.de

KAMERYN WILLIAMS, THE CITY UNIVERSITY OF NEW YORK, CUNY GRADUATE CENTER, MATHEMATICS PROGRAM, 365 FIFTH AVENUE, NEW YORK, NY 10016, USA

E-mail address: kwilliams4@gradcenter.cuny.edu

URL: <http://kamerynblog.wordpress.org>