

INNER-MODEL REFLECTION PRINCIPLES

NEIL BARTON, ANDRÉS EDUARDO CAICEDO, GUNTER FUCHS,
JOEL DAVID HAMKINS, JONAS REITZ, AND RALF SCHINDLER

ABSTRACT. We introduce and consider the inner-model reflection principle, which asserts that whenever a statement $\varphi(a)$ in the first-order language of set theory is true in the set-theoretic universe V , then it is also true in a proper inner model $W \subsetneq V$. A stronger principle, the ground-model reflection principle, asserts that any such $\varphi(a)$ true in V is also true in some non-trivial ground model of the universe with respect to set forcing. These principles each express a form of width reflection in contrast to the usual height reflection of the Lévy-Montague reflection theorem. They are each equiconsistent with ZFC and indeed Π_2 -conservative over ZFC, being forceable by class forcing while preserving any desired rank-initial segment of the universe. Furthermore, the inner-model reflection principle is a consequence of the existence of sufficient large cardinals, and lightface formulations of the reflection principles follow from the maximality principle MP and from the inner-model hypothesis IMH. We also consider some questions concerning the expressibility of the principles.

CONTENTS

1. Introduction	1
2. Forcing inner model reflection	4
3. Large cardinals	5
4. The maximality principle and the inner model hypothesis	11
5. Forcing axioms	13
6. Expressibility of inner model reflection	14
Acknowledgements	15
References	15

1. INTRODUCTION

Every set theorist is familiar with the classical Lévy-Montague reflection principle, which explains how truth in the full set-theoretic universe V reflects down to truth in various rank-initial segments V_θ of the cumulative hierarchy. Thus, the Lévy-Montague reflection principle is a form of height-reflection, in that truth in V is reflected vertically downwards to truth in some V_θ .

2010 *Mathematics Subject Classification*. Primary 03E45; Secondary 03E35, 03E55, 03E65.

Key words and phrases. Inner-model reflection principle, ground-model reflection principle.

This article grew out of an exchange held by some of the authors on the Mathematics.StackExchange site in response to an inquiry posted by the first-named author concerning the nature of width-reflection in comparison to height-reflection [Bar16]. Commentary concerning this paper can be made at <http://jdh.hamkins.org/inner-model-reflection-principles>.

In this brief article, in contrast, we should like to introduce and consider a form of width-reflection, namely, reflection to non-trivial inner models. Specifically, we shall consider the following reflection principles.

Definition 1.1.

- (1) *The inner-model reflection principle asserts that if a statement $\varphi(a)$ in the first-order language of set theory is true in the set-theoretic universe V , then there is a proper inner model W , a transitive class model of ZF containing all ordinals and with $a \in W \subsetneq V$, in which $\varphi(a)$ is true.*
- (2) *The ground-model reflection principle asserts that if $\varphi(a)$ is true in V , then there is a non-trivial ground model $W \subsetneq V$ with $a \in W$ and $W \models \varphi(a)$.*
- (3) *Variations of the principles arise by insisting on inner models of a particular type, such as ground models for a particular type of forcing, or by restricting the class of parameters or formulas that enter into the schema.*
- (4) *The lightface forms of the principles, in particular, make their assertion only for sentences, so that if σ is a sentence true in V , then σ is true in some proper inner model or ground W , respectively.*

In item (1), we could equivalently insist that the inner model W satisfies ZFC, simply by reflecting the conjunction $\varphi(a) \wedge \text{AC}$ instead of merely $\varphi(a)$. For the rest of this article, therefore, our inner models will satisfy ZFC. In item (2), a *ground model* or simply a *ground* of the universe V is a transitive inner model $W \models \text{ZFC}$ over which the universe V is obtained by set forcing, so that $V = W[G]$ for some forcing notion $\mathbb{P} \in W$ and some W -generic filter $G \subseteq \mathbb{P}$. The universe V , of course, is a ground of itself by trivial forcing, but the ground-model reflection principle seeks grounds $W \subsetneq V$ that are properly contained in V .

Aside from the usual height-reflection principles, the closest relatives of the principles we are considering are the inner model hypotheses proposed by Friedman [Fri06]. These meta-principles assert that certain sentences obtainable in outer models are already satisfied in inner models. Although inner model hypotheses were one of the original motivators for asking in [Bar16] the question that prompted this paper, they end up behaving rather differently, with two immediate salient differences between them and inner-model reflection principles being that

- (1) the inner-model reflection principles make no reference to outer models, and
- (2) the inner models asserted to exist in inner model hypotheses are not necessarily proper.

As we will see, these principles also differ in their expressibility. However, the lightface version of the inner-model reflection principle is obtainable from the standard inner model hypothesis, see §4.

The full inner-model reflection principle is expressible in the second-order language of set theory, such as in Gödel-Bernays GBC set theory, as a schema:

$$\forall a [\varphi(a) \rightarrow \exists W \subsetneq V \varphi^W(a)],$$

with a separate statement for each formula φ and where the quantifier $\exists W$ ranges over the inner models of ZFC.

(We add a clarification at the request of a referee. In GBC, choice is usually formulated as a global principle asserting that there is a class function selecting an element from every non-empty set. In contrast, it is stated in ZFC as the existence

of such set functions for each set-sized collection of non-empty sets. Nevertheless, this difference plays no role here, and even the weakening of GBC that uses the ZFC version of choice suffices to express the inner-model reflection principle as indicated above.)

The ground-model reflection principle, in contrast, is expressible as a schema in the first-order language of set theory. To see this, consider first the ground-model enumeration theorem of Fuchs-Hamkins-Reitz [FHR15, theorem 12], which asserts that there is a definable class $W \subseteq V \times V$ for which

- (i) every section $W_r = \{x \mid (r, x) \in W\}$ is a ground of V by set forcing; and
- (ii) every ground arises as such a section W_r .

Thus, the collection of ground models $\{W_r \mid r \in V\}$ is uniformly definable and we may quantify over the grounds in a first-order manner by quantifying over the indices r used to define them. In light of the enumeration theorem, the ground-model reflection principle is expressed in the first-order language of set theory as the following schema:

$$\forall a [\varphi(a) \rightarrow \exists r W_r \subsetneq V \wedge \varphi^{W_r}(a)].$$

We may therefore undertake an analysis of the ground-model reflection principle in a purely first-order formulation of set theory, such as in ZFC .

Clearly, the ground-model reflection principle strengthens the inner-model reflection principle, since ground models are inner models. Both principles are obviously false under the axiom of constructibility $V = L$, since L has no non-trivial inner models. Similarly, the ground-model reflection principle is refuted by the ground axiom, which asserts that there are no non-trivial grounds, see Hamkins [Ham05] and Reitz [Rei06, Rei07]. In particular, the ground axiom holds in many of the canonical inner models of large cardinal assumptions, such as the Dodd-Jensen core model K^{DJ} , the model $V = L[\mu]$, and also the Jensen-Steel core model K , provided that there is no inner model with a Woodin cardinal. The reason is that these inner models are definable in a way that is generically absolute, see Jensen-Steel [JS13] and Mitchell [Mit12] (one uses the hypothesis of no inner models with a Woodin cardinal in the case of K), and so they have no non-trivial ground models.

In the rest of this paper we verify and discuss some other properties of these principles, and how they can be obtained. §2 provides a discussion of how the principles can be forced. In particular we show in theorems 2.1 and 2.2 that both the lightface and boldface versions of the ground-model reflection principle are obtainable from models of ZFC using forcing constructions. In §3 we explain how the principles interact with large cardinal axioms, proving in theorems 3.1 and 3.3 that the inner-model reflection principle is a consequence of sufficient large cardinals. We also discuss how these principles behave in many canonical inner models. In particular, as we shall explain in corollary 3.4 and following remarks, under the right large cardinal assumption, the core model satisfies the inner-model reflection principle, but not the ground-model reflection principle. In contrast to this, in theorem 3.8 we show that fine-structural inner models of sufficiently strong large cardinal assumptions satisfy the ground-model reflection principle. This is complemented by theorem 3.9, which shows that the assumption of theorem 3.8 is optimal. In §4 we show how the maximality principle of Stavi-Väänänen [SV02] and Hamkins [Ham03] implies the lightface ground-model reflection principle, and how the inner model hypothesis of [Fri06] implies the lightface inner-model reflection principle.

Next, §5 considers the relationship with forcing axioms. We point out that while the bounded proper forcing axiom is consistent with the failure of inner-model reflection, the ground-model reflection principle is consistent with several strong forcing axioms. Finally, in §6 we discuss limitations concerning the expressibility of the inner-model reflection principle, and conclude with an open question.

2. FORCING INNER MODEL REFLECTION

Let us begin by showing that we may easily force instances of the lightface reflection principles as follows.

Theorem 2.1. *In the forcing extension $V[c]$ arising by forcing to add a Cohen real, the lightface ground-model reflection principle holds, and indeed, the ground-model reflection principle holds for assertions with arbitrary parameters from V .*

Proof. Suppose that an assertion $\varphi(a)$ is true in $V[c]$, where $a \in V$. By the homogeneity of the forcing, it follows that $\varphi(a)$ is forced by every condition, and so it will be true in $V[d]$, where d is the Cohen real obtained by retaining every other digit of c and using them to form a new real number. Since $V[d]$ is a proper inner model and indeed a ground of $V[c]$, it follows that the ground-model reflection principle holds in $V[c]$ for first-order assertions having parameters in V . \square

We could have allowed any parameter $a \in V[c]$ for which $V[a] \subsetneq V[c]$, since the quotient forcing is again that of adding a Cohen real. We cannot necessarily allow c itself as a parameter, since $L[c]$ satisfies the statement $V = L[c]$, using c as a parameter, but no proper inner model of $L[c]$ satisfies this statement. Meanwhile, the proof of theorem 2.2 will show that over some models, one can allow even c as a parameter for the ground-model reflection principle in $V[c]$.

Cohen forcing is hardly unique with the property mentioned in theorem 2.1, since essentially the same argument works with many other kinds of forcing, such as random forcing or Cohen forcing at higher cardinals. Indeed, let us now push the idea a little harder with class forcing so as to achieve the full principle, with arbitrary parameters, including all the new parameters of the forcing extension.

Theorem 2.2. *Every model of ZFC has a class-forcing extension satisfying the ground-model reflection principle, with arbitrary parameters from the extension.*

Proof. By preliminary forcing if necessary, we may assume that GCH holds. Let \mathbb{P} be the proper-class Easton-support product of the forcing posets $\text{Add}(\delta, 1)$ that add a Cohen subset to every cardinal δ in some proper class of regular cardinals. Suppose that $G \subseteq \mathbb{P}$ is V -generic, and consider the extension $V[G]$, which is a model of ZFC. Suppose that $V[G] \models \varphi(a)$ for some first-order assertion φ and set a . So there is some condition $p \in \mathbb{P}$ forcing $\varphi(\dot{a})$ for some name \dot{a} with $\dot{a}_G = a$. Let δ be a stage of forcing that is larger than the support of p and any condition in the name \dot{a} , and let G_0 be just like G on all coordinates other than δ , except that on coordinate δ itself, we take only every other digit of the generic subset of δ that was added, re-indexed so as to make a generic subset of δ . Thus, the filter $G_0 \subseteq \mathbb{P}$ is V -generic for this forcing, $p \in G_0$ and $V[G] = V[G_0][g]$, where g consists of the information on the complementary digits of the subset of δ . So $V[G_0]$ is a proper inner model of ZFC, and since $p \in G_0$ and $\dot{a}_{G_0} = \dot{a}_G$, it follows that $V[G_0] \models \varphi(a)$, fulfilling the desired instance of ground-model reflection. \square

Corollary 2.3. *The inner-model and ground-model reflection principles are each conservative over ZFC for Π_2 assertions about sets. In other words, any Σ_2 assertion that is consistent with ZFC is also consistent with ZFC plus the ground-model reflection principle.*

Proof. The point is that the previous argument, by starting the forcing sufficiently high, shows that any given model of ZFC can be extended to a model of ground-model reflection while preserving any particular V_θ and therefore the truth of any particular Σ_2 assertion. Thus, any Σ_2 assertion that is consistent with ZFC is also consistent with ground-model reflection and hence also with inner-model reflection. By contraposition, any Π_2 assertion that is provable from ZFC or GBC plus the inner-model or ground-model reflection principles is provable in ZFC alone. \square

The Lévy-Montague reflection principle produces for every natural number n in the metatheory a proper class club $C^{(n)}$ of cardinals θ , the Σ_n -correct cardinals, for which $V_\theta \prec_{\Sigma_n} V$. But this kind of reflection can never hold for inner models:

Observation 2.4. *If W is an inner model of ZF and $W \prec_{\Sigma_1} V$, then $W = V$.*

Proof. Assume that W is a transitive class model of ZF containing all ordinals and that $W \prec_{\Sigma_1} V$. If $W \neq V$, then there is some set a in V that is not in W . Let θ be above the rank of a and let $u = (V_\theta)^W$. So V models “there is a set of rank less than θ that is not in u .” This is a Σ_1 assertion about θ and u , witnessed by a rank function into θ . But it is not true in W , by the choice of u . So it must be that $W = V$. \square

So the situation of width reflection is somewhat different in character from that of height reflection, where we have $H_\kappa \prec_{\Sigma_1} V$ for every uncountable cardinal κ and more generally $V_\theta \prec_{\Sigma_n} V$ for all cardinals θ in the class club $C^{(n)}$. Observation 2.4 shows that there is no analogue of this for width reflection.

3. LARGE CARDINALS

Next, we point out that the inner-model reflection principle is an outright consequence of sufficient large cardinals. We present several such hypotheses that suffice, in decreasing order of magnitude. The theorem below discusses measurability.

Theorem 3.1.

- (1) *If there is a measurable cardinal, then the lightface inner-model reflection principle holds.*
- (2) *Indeed, if κ is measurable, then the inner-model reflection principle holds for assertions with parameters in V_κ .*
- (3) *Consequently, if there is a proper class of measurable cardinals, then the full inner-model reflection principle holds for arbitrary parameters.*

Proof. The theorem is easy to prove. Suppose that $\varphi(a)$ is true in V , where the parameter a is in V_κ for some measurable cardinal κ . Let $j : V \rightarrow M$ be an ultrapower embedding by a measure on κ , with critical point κ , into a transitive class M , which must be a proper inner model, definable from the measure. Since $a \in V_\kappa$, below the critical point, it follows that $j(a) = a$ and consequently $M \models \varphi(a)$ by the elementarity of j . So $\varphi(a)$ is true in a proper inner model, thereby witnessing this instance of the inner-model reflection principle. \square

Using this, we can separate the inner-model reflection principle from the ground-model reflection principle. They do not coincide.

Corollary 3.2. *If ZFC is consistent with a proper class of measurable cardinals, then there is a model of ZFC in which the inner-model reflection principle holds, but the ground-model reflection principle fails.*

Proof. One can prove this as in corollary 3.4 below, using the fact that the core model K has no non-trivial grounds; but let us give a forcing proof. If there are a proper class of measurable cardinals in V , then there is a class-forcing extension $V[G]$ preserving them, in which every set is coded into the GCH pattern. This idea was a central theme of [Rei06, Rei07]; but let us sketch the details. After forcing GCH, if necessary, we perform a lottery iteration, which at the successor of every measurable cardinal either forces a violation of GCH or performs trivial forcing. Generically, every set is coded into the resulting pattern. The standard lifting arguments, such as those in Hamkins [Ham00, Ham], show that all measurable cardinals are preserved, and so by theorem 3.1 it follows that $V[G]$ satisfies the inner-model reflection principle. Meanwhile, because every set in $V[G]$ is coded into the GCH pattern, by placing sets into much larger sets it follows that every set is coded unboundedly often. Since set forcing preserves the GCH pattern above the size of the forcing, every ground model has this coding. So $V[G]$ can have no non-trivial grounds and consequently does not satisfy the ground-model reflection principle. \square

We can improve the large cardinal hypothesis of the preceding results by using the work of Vickers-Welch [VW01].

Theorem 3.3. *If Ord is Ramsey, then the inner-model reflection principle holds.*

Proof. Following [VW01, definition 2.2; see also definition 1.1], we say that Ord is Ramsey if and only if there is an unbounded class $I \subseteq \text{Ord}$ of good indiscernibles for $\langle V, \in \rangle$. This is a second-order assertion in GBC. One can arrange set models of this theory with first-order part V_κ , if κ is a Ramsey cardinal, and so the hypothesis “Ord is Ramsey” is strictly weaker in consistency strength than ZFC plus the existence of a Ramsey cardinal, which in turn is strictly weaker in consistency strength than ZFC plus the existence of a measurable cardinal.

Meanwhile, the argument of [VW01, theorem 2.3] shows how to construct from the class I a transitive class M with a non-trivial elementary embedding $j: M \rightarrow V$, where the critical point of j can be arranged so as to be any desired element of I . Note that $M \subsetneq V$ in light of the Kunen inconsistency. If $\varphi(a)$ holds in V , then there is such a $j: M \rightarrow V$ with critical point above the rank of a and therefore with $a \in M$ and $j(a) = a$. It follows by elementarity that $M \models \varphi(a)$, thereby fulfilling the inner-model reflection principle. \square

In particular, any statement that is compatible with Ord being Ramsey is also compatible with the inner-model reflection principle, which makes a contrast to corollary 2.3. One can use theorem 3.3 to weaken the hypothesis of corollary 3.2 as follows, where we now use the core model rather than forcing. Note that if “Ord is Ramsey” holds, then the hypothesis of corollary 3.4 holds in an inner model.

Corollary 3.4. *If the core model K exists and satisfies “Ord is Ramsey”, then K satisfies the inner-model reflection principle, but not the ground-model reflection principle.*

Proof. If $K \models \text{Ord is Ramsey}$, then it satisfies the inner-model reflection principle by the previous theorem. And since K is definable in a way that is generically absolute, it has no non-trivial grounds and therefore cannot satisfy the ground-model reflection principle. \square

We thank Philip Welch for allowing us to include the following observation showing that the assumption above can be further weakened. A large cardinal hypothesis in the region of long unfoldable cardinals suffices, see Welch [Wel04, definition 1.3]. For the present purpose, consider the following weakening of long unfoldability:

Definition 3.5. *A cardinal κ is LU if and only if there is a definable proper inner model M with $V_\kappa \prec M$.*

For simplicity, we stated the definition in GBC, using the language of proper classes. Welch’s definition appears more involved since it is being formalized in ZFC; the appropriate version here would be the result of taking $S = \emptyset$ in [Wel04, definition 1.3], see [Wel04, fact 1.1].

Theorem 3.6 (Welch). *If the class of LU cardinals is Mahlo in V , then the inner-model reflection principle holds.*

“Mahlo” is understood here to mean “definably Mahlo”, that is, the LU cardinals form a class stationary with respect to all definable class clubs, as in [Wel04, remark following theorem 1.4].

Proof. Suppose $\varphi(a)$ holds in V . By the Lévy-Montague reflection principle, there is a club C of ordinals α such that $V_\alpha \models \varphi(a)$. Since the class of LU cardinals is Mahlo, there is such a cardinal κ with $a \in V_\kappa$ and $V_\kappa \models \varphi(a)$. Letting M be as in definition 3.5 for κ , we have $M \models \varphi(a)$. \square

Similarly, corollary 3.4 holds replacing the assertion that K satisfies “Ord is Ramsey” with the weaker claim that K models “the LU cardinals are Mahlo”.

As shown in [Wel04, theorem 1.4], the assertion that the class of LU cardinals is Mahlo is strictly weaker in consistency strength than ω^2 - Π_1^1 -Determinacy, an assumption significantly weaker than “Ord is Ramsey”.

Remark 3.7. As suggested by item (3) of definition 1.1, there are natural hierarchies of inner-model reflection principles obtained by restricting the complexity of the statements φ under consideration. In particular, let Π_n -IMR be the Π_n -inner-model reflection principle, the version where the φ are restricted to be Π_n . In private communication, Welch remarked that this is a proper hierarchy in the following strong sense: (if it exists) there is in K an equivalence for each n between Π_n -IMR and a (somewhat technical) statement concerning the length of the mouse order and admissibility. In turn, this can be used to show that there is a strictly increasing sequence of inner models $(W_n)_{n < \omega}$, with W_n the least $L[E]$ model for which Π_n -INR holds (but Π_{n+1} -INR fails).

Since the core model can never satisfy the ground-model reflection principle, it is thus natural to ask whether any fine-structural extender model can. The following two theorems provide the precise consistency strength of the existence of such a model. The first direction builds on the methods of Fuchs-Schindler [FS16].

Theorem 3.8. *If $L[E]$ is a minimal iterable extender model with a proper class of Woodin cardinals, then it satisfies the ground-model reflection principle.*

Proof. The argument uses some techniques from [FS16], adapted to the present context. One key ingredient is the “ δ generator version” of Woodin’s extender algebra, see Schindler-Steel [SS09, lemma 1.3]. If \mathcal{M} is a normally $(\omega, \kappa^+ + 1)$ -iterable premouse and δ is a Woodin cardinal in \mathcal{M} , then this forcing notion $\mathbb{P} = \mathbb{P}^{\mathcal{M}|\delta}$ has the property that for every subset $A \subseteq \kappa^+$ there is a normal non-dropping iteration tree on \mathcal{M} with a last model \mathcal{N} such that if π is the iteration embedding, then $A \cap \pi(\delta)$ is $\pi(\mathbb{P})$ -generic over \mathcal{N} .

The other key ingredient is the \mathcal{P} -construction of [SS09]. If \mathcal{M} is a premouse, δ is a cutpoint of \mathcal{M} (that is, δ is not overlapped by any extender on the \mathcal{M} sequence), $\bar{\mathcal{P}}$ is a premouse of height $\delta + \omega$, δ is a Woodin cardinal in $\bar{\mathcal{P}}$, $\bar{\mathcal{P}}|\delta$ is definable in $\mathcal{M}|\delta$, and $\bar{\mathcal{P}}[G] = \mathcal{M}|\delta$ for some G which is generic over $\bar{\mathcal{P}}$ for its version of the δ generator version of Woodin’s extender algebra, then $\mathcal{P}(\mathcal{M}, \bar{\mathcal{P}}, \delta)$ is the result of the maximal \mathcal{P} -construction over $\bar{\mathcal{P}}$ with respect to \mathcal{M} , above δ .

Essentially, this construction appends to the extender sequence of $\bar{\mathcal{P}}$ the restrictions of the extenders on the extender sequence of \mathcal{M} that are indexed beyond δ , as long as this results in a structure in which δ is still a Woodin cardinal. We will define $\mathcal{P}(\mathcal{M}, \bar{\mathcal{P}}, \delta)$ also in the case that δ is not a cutpoint of \mathcal{M} , by letting α be least such that $\alpha \geq \delta$, $E_\alpha^{\mathcal{M}} \neq \emptyset$ and $\kappa = \text{crit}(E_\alpha^{\mathcal{M}}) \leq \delta$, letting $\zeta \leq \text{ht}(\mathcal{M})$ be maximal such that $\alpha \leq \zeta$ and $\kappa^{+\mathcal{M}|\alpha} = \kappa^{+\mathcal{M}|\zeta}$, and setting

$$\mathcal{P}(\mathcal{M}, \bar{\mathcal{P}}, \delta) = \mathcal{P}(\text{ult}_n(\mathcal{M}|\zeta, E_\alpha^{\mathcal{M}}), \bar{\mathcal{P}}, \delta),$$

where n is least such that $\rho_{n+1}(\mathcal{M}|\zeta) \leq \kappa$, if such an n exists, and $n = 0$ otherwise. See the discussion after the statement of [SS09, lemma 3.21] for details.

An iteration tree \mathcal{T} on an extender model W which is definable in $L[E]$ is said to be guided by \mathcal{P} -constructions in $L[E]$ if and only if the branches in \mathcal{T} at limit stages are determined by \mathcal{Q} -structures that are pullbacks of \mathcal{Q} -structures obtained in $L[E]$ by maximal \mathcal{P} -constructions, see the discussion after [FS16, definition 3.22] for details.

We modify [FS16, definition 3.25] to say that an extender model W definable in $L[E]$ is *minimal* for $L[E]$ if and only if it has a proper class of Woodin cardinals, and for every δ that is Woodin in W , whenever $\mathcal{T} \in L[E]$ is a normal iteration tree on W which is based on $W|\delta$ and is guided by \mathcal{P} -constructions in $L[E]$ and uses only extenders indexed above the supremum of the Woodin cardinals of W below δ , the following holds true:

- if \mathcal{T} has limit length, then \mathcal{T} lives strictly below δ if and only if

$$\mathcal{P}(L[E], \mathcal{M}(\mathcal{T}) + \omega, \delta(\mathcal{T})),$$

if defined, is not a proper class, and

- if \mathcal{T} has successor length and $[0, \infty]_{\mathcal{T}}$ does not drop, then

$$\mathcal{P}(L[E], \mathcal{M}_\infty^{\mathcal{T}}|i_{0,\infty}^{\mathcal{T}}(\delta) + \omega, i_{0,\infty}^{\mathcal{T}}(\delta)),$$

if defined, is a proper class.

Here, as in [FS16], given a sound premouse \mathcal{R} , by $\mathcal{R} + \omega$ we denote the premouse end-extending \mathcal{R} and obtained from \mathcal{R} by constructing over it one step further.

It follows that $L[E]$ itself is minimal in this sense. For example, if δ, \mathcal{T} are as above, where the length of \mathcal{T} is a limit ordinal and \mathcal{T} lives strictly below δ , then if $\mathcal{P} = \mathcal{P}(L[E], \mathcal{M}(\mathcal{T}) + \omega, \delta(\mathcal{T}))$ were a proper class, it would be an iterable extender model with a proper class of Woodin cardinals that lies below $L[E]$ in the canonical pre-well-ordering of iterable extender models. The point is that \mathcal{T}

would be according to the iteration strategy of $L[E]$, and so there would be a cofinal well-founded branch such that $\mathcal{M}_b^{\mathcal{T}}$ is iterable. But since \mathcal{T} lives strictly below δ , it would follow that $\pi_{0,b}^{\mathcal{T}}(\delta) > \delta(\mathcal{T})$, and so $\delta(\mathcal{T})$ would not be Woodin in $\mathcal{M}_b^{\mathcal{T}}$, but it is Woodin in \mathcal{P} , which implies that \mathcal{P} is below $L[E]$; see the proof of [FS16, lemma 3.23], which shows that there is no $L[E]$ -based sequence of length 2, in the terminology introduced there. If the length of \mathcal{T} is a limit ordinal and \mathcal{T} does not live strictly below δ , then one can argue as in the proof of [FS16, lemma 3.26] to show that the relevant \mathcal{P} -construction yields a proper class model, and similarly in the case that \mathcal{T} has successor length.

Now, to show that $L[E]$ satisfies the ground-model reflection principle, let $\varphi(a)$ be a statement true in $L[E]$. By assumption, there is a δ larger than the rank of a that is Woodin in $L[E]$. Let $\eta > \delta$ be a cutpoint of the extender sequence E . Let \tilde{E} code E as a class of ordinals in some natural way. Form an iteration tree on $L[E]$ as follows (the construction is much as in the proof of [FS16, lemma 3.29]): first, hit η many times some total extender on the E -sequence with critical point greater than the rank of a and indexed below δ but above every Woodin cardinal less than δ . After that, at successor stages, choose the least total extender in the current model with an index greater than the supremum of the Woodin cardinals below the current image δ' of δ that violates an axiom of the extender algebra with respect to $\tilde{E} \cap \delta'$. Since such extenders suffice to witness the Woodinness of δ , one can work with the version of the extender algebra with this added restriction. If there is no such extender, or if a limit stage λ is reached such that

$$\mathcal{P} = \mathcal{P}(L[E], \mathcal{M}(\mathcal{T} \upharpoonright \lambda) + \omega, \delta(\mathcal{M}(\mathcal{T} \upharpoonright \lambda)))$$

is a proper class, then the construction is complete. Otherwise, as in the proof of [FS16, lemma 3.29], it follows that \mathcal{P} can serve as a \mathcal{Q} -structure, and the branch for $\mathcal{T} \upharpoonright \lambda$ given by the iteration strategy in V can be found inside $L[E]$, allowing us to extend the iteration tree in this case. Further, as in that proof, it follows that this process terminates at a limit stage $\lambda = \eta^{+L[E]} = \delta(\mathcal{T})$, and

$$\mathcal{P} = \mathcal{P}(L[E], \mathcal{M}(\mathcal{T}) + \omega, \delta(\mathcal{T}))$$

is a proper class, and hence a proper ground of $L[E]$. The tree \mathcal{T} does not have a cofinal well-founded branch inside $L[E]$, but the model $\mathcal{M}(\mathcal{T})$ can be formed within $L[E]$.

Since $L[E]$ is iterable in V , it follows that \mathcal{T} has a cofinal well-founded branch b in V such that $\mathcal{M}_b^{\mathcal{T}}$ is iterable in V . Since η is a cutpoint of E , it follows that it is a cutpoint of \mathcal{P} , and moreover \mathcal{P} is iterable above $\eta^{+L[E]}$. The coiteration of \mathcal{P} and $\mathcal{M}_b^{\mathcal{T}}$ has to result in a common (proper class) iterate \mathcal{Q} . Let $\pi: L[E] \rightarrow \mathcal{M}_b^{\mathcal{T}}$, $\sigma: \mathcal{M}_b^{\mathcal{T}} \rightarrow \mathcal{Q}$ and $\tau: \mathcal{P} \rightarrow \mathcal{Q}$ be the iteration embeddings. Note that a is not moved by π . Since the coiteration between $\mathcal{M}_b^{\mathcal{T}}$ and \mathcal{P} is above $\eta^{+L[E]}$, a is not moved by σ or τ either. Hence, we get:

$$L[E] \models \varphi(a) \iff \mathcal{Q} \models \varphi(\sigma(\pi(a))) \iff \mathcal{Q} \models \varphi(\tau(a)) \iff \mathcal{P} \models \varphi(a).$$

Thus, \mathcal{P} is a non-trivial ground of $L[E]$ which reflects the truth of $\varphi(a)$, as desired. \square

It was noted earlier that the ground axiom, stating that there is no proper ground, implies the failure of the ground-model reflection principle. There is a natural way of relativizing the ground axiom to an arbitrary set a , thus weakening

it: let us say that the ground axiom holds *relative to a* if and only if there is no proper ground containing a . Thus, the usual ground axiom is the ground axiom relative to \emptyset . Clearly, if there is an a such that the ground axiom holds relative to a , then ground-model reflection fails, since there is then no non-trivial ground reflecting the statement “ $a = a$ ”. The following theorem shows that a sufficiently iterable $L[E]$ model that is below a proper class of Woodin cardinals satisfies the ground axiom relative to some set, and hence fails to satisfy the ground-model reflection principle.

Theorem 3.9. *Let $L[E]$ be an extender model that is fully iterable in every set generic extension of V and that is below a proper class of Woodin cardinals, in the sense that its Woodin cardinals are bounded and no initial segment of $L[E]$ is a sharp for an inner model with a proper class of Woodin cardinals. Under these assumptions, $L[E]$ satisfies the ground axiom relative to some set $x \in L[E]$.*

Proof. Let $L[E]$ be as described, and let α be the supremum of its Woodin cardinals. By [SS09, theorem 0.2], there is a $\beta \geq \alpha$ which is a cutpoint of $L[E]$ and such that $L[E]$ has an iteration strategy for iteration trees on $L[E]$ that only use extenders whose critical points are above β . By our assumption that $L[E]$ is fully iterable in every set-generic forcing extension of V , the argument of the proof of [SS09, theorem 0.2] generalizes to show that there is a β as above such that every set-generic forcing extension of $L[E]$ has an iteration strategy for iteration trees on $L[E]$ which only use extenders with critical points above β . We claim that $L[E]$ satisfies the ground axiom relative to $x = L[E] \restriction \beta$.

To see this, let W be a ground of $L[E]$ with $x \in W$. We must show that W is the trivial ground, that is, that $W = L[E]$. Let $\mathbb{P} \in W$ be a notion of forcing, and let $g \subseteq \mathbb{P}$ be generic over W such that $W[g] = L[E]$. Let γ be the cardinality of \mathbb{P} in $L[E]$, and let h be $\text{Col}(\omega, \gamma)$ -generic over V . By the absorption property of the collapse, there is then an $h' \subseteq \text{Col}(\omega, \gamma)$ generic over $L[E]$ such that $L[E][h] = W[g][h] = W[h']$. It follows that $L[E]$ is a definable class in $W[h']$ and is fully iterable there with respect to iteration trees that live above β . In fact, E (and hence $L[E]$) is definable in $L[E][h] = W[h']$ by a formula φ using a parameter $z \in L[E]$. This is because $L[E]$ is a ground of $W[h']$, and is hence definable in $W[h']$ using a parameter from $L[E]$ (namely the power set $\mathcal{P}(\text{Col}(\omega, \gamma))^{L[E]}$; see the discussion of the ground-model reflection principle in the introduction, and [FHR15, theorem 5]).

Further, E is definable inside $L[E]$, using the argument of the proof of Schlutzenberg [Sch14, theorem 4.3]—since we only have self-iterability above β , we start the inductive definition with $E \restriction \beta$, and get a definition of E using $E \restriction \beta$ as a parameter. Combining these two parameters $\mathcal{P}(\text{Col}(\omega, \gamma))^{L[E]}$ and $E \restriction \beta$ we obtain the parameter z indicated above.

Now, $z = \tau^{h'}$ for some $\text{Col}(\omega, \gamma)$ -name $\tau \in W$. Let $p \in h'$ be a condition that forces over W that $\varphi(-, \tau)$ defines a universal extender model that is fully iterable above the cutpoint β and that agrees with $L[E]$ up to β . Inside $W[h']$, for any $q \in \text{Col}(\omega, \gamma)$ extending p , let h'_q be the finite variant of h' compatible with q , that is, viewing h' as a function from ω to γ , $h'_q = q \cup h' \restriction (\omega \setminus \text{dom}(q))$. Still working inside $W[h']$, for q as above, let $L[E_q]$ be the inner model defined by $\varphi(-, \tau^{h'_q})$.

All these models $L[E_q]$ are universal, iterable above β , coincide up to β , and β is a cutpoint for each of them. Thus, they can all be simultaneously coiterated

inside $W[h']$, yielding a common iterate, which we denote $L[F]$. This is a definable extender model in $W[h']$. But note that $L[F]$ is definable in W from p and τ , since its definition does not depend on h' , but only on the collection of all finite variants of h' . In more detail, view F as a class of ordinals (this is no problem, as $L[F]$ has a canonical well-order). We claim that F is definable in W as the class of ordinals ξ such that p forces “ $\xi \in F$ ” (note that this statement uses the parameters p and τ). To see this, let h'' be $\text{Col}(\omega, \gamma)$ -generic over W with $p \in h''$, let ξ be given, and suppose that $\xi \in F = F^{W[h']}$, say. If $\xi \notin F^{W[h'']}$, then there is a condition $q \in h''$ that forces “ $\xi \notin F$ ”. Since $p \in h''$, we may assume that $q \leq p$. But then, since $q \in h'_q$, it follows that $\xi \notin F^{W[h'_q]}$, which is absurd since $W[h'_q] = W[h']$, as h'_q is just a finite variant of h' .

One can now argue as in the proof of Sargsyan-Schindler [SS18, theorem 2.14]: in $W[h']$, the iteration embedding $j: L[E] \rightarrow L[F]$ is a definable class, and since $W[h']$ is a set-generic forcing extension of W , it follows that there is a thick class Γ definable in W and consisting of ordinal fixed points of j such that $L[E]$ is isomorphic to the hull in $L[F]$ of $\Gamma \cup L[E] \upharpoonright \beta$. Since Γ and $L[F]$ are definable in W , so is this hull. Hence, $L[E]$, the Mostowski collapse of this hull, is a class in W , and therefore $L[E] \subseteq W$. Thus, $W = L[E]$, as was to be shown. \square

Thus, the consistency strength of the ground-model reflection principle, which is the same as that of ZFC, increases dramatically if it is relativized to a sufficiently iterable $L[E]$ model. This is in line with our earlier results that the core model K cannot satisfy this principle, and supports the view that ground-model reflection is in a sense an “anti-canonical inner model” statement. There are several well-known instances of this phenomenon: for example, the statement (A) that every projective set of reals is Lebesgue-measurable has consistency strength an inaccessible cardinal, but does not hold in a canonical inner model at a low level, where there are easily definable well-orderings of the reals. In fact, if (A) holds in an iterable $L[E]$ model M , then for every $n < \omega$ and every real x in M , $M_n^\sharp(x)$ must be in M . Another example is the statement (B) that there is no definable well-ordering of the reals in $L(\mathbb{R})$. The consistency strength of (B) is just that of ZFC, but by a result of Steel, the least iterable $L[E]$ model that satisfies (B) must be above M_ω (the least $L[E]$ model with infinitely many Woodin cardinals) in the mouse order, see the remark following the statement of [SS09, theorem 0.2].

4. THE MAXIMALITY PRINCIPLE AND THE INNER MODEL HYPOTHESIS

Consider next the *maximality principle* of [SV02] and [Ham03], which asserts that whenever a statement is forceably necessary, which is to say that it is forceable in such a way that it remains true in all further extensions, then it is already true. This is expressible in modal terms by the schema $\diamond \Box \varphi \rightarrow \varphi$, the principal axiom of the modal theory S5, where the modal operators are interpreted so that $\diamond \psi$ means that ψ is true in some set-forcing extension and $\Box \psi$ means that ψ is true in all set-forcing extensions.

Theorem 4.1. *The maximality principle implies the lightface ground-model reflection principle.*

Proof. Suppose that a sentence σ is true in V . Consider the statement, “ σ is true in some non-trivial ground.” In light of the ground-model enumeration theorem, this supplementary statement is expressible in the first-order language of set theory.

Furthermore, it becomes true in any non-trivial forcing extension $V[G]$, since V is a non-trivial ground of $V[G]$, and the statement remains true in any further forcing extension. Thus, the supplementary statement is forceably necessary in V , and therefore by the maximality principle it must already be true in V . So there must be a non-trivial ground model $W \subsetneq V$ in which σ is true. \square

The same argument works with the various other versions of the maximality principle, such as $\text{MP}_\Gamma(X)$, where only forcing notions in a class Γ are considered and statements with parameters from X . The same argument as in theorem 4.1 shows that $\text{MP}_\Gamma(X)$ implies the Γ -ground model reflection principle with parameters from X .

A similar argument can be made from the *inner-model hypothesis* IMH, which is the schema of assertions made for each sentence σ , that if there is an outer model with an inner model of σ , then there is already an inner model of σ without first moving to the outer model. This principle also can be described in modal vocabulary as the schema of assertions $\diamond\diamond\sigma \rightarrow \diamond\sigma$, where the up-modality \diamond refers here to possibility in outer models and the down-modality \diamond refers to possibility in inner models. See [Fri06] for details about IMH; the axiom is naturally formalized in a multiverse context of possible outer models, although Antos-Barton-Friedman [ABF] shows that modified versions of the axiom are expressible in the second-order language of set theory in models of a variant of $\text{GBC} + \Sigma_1^1$ -comprehension, without direct reference to outer models. In formulating IMH, one may equivalently insist on proper inner models.

Theorem 4.2. *The inner-model hypothesis implies the lightface inner-model reflection principle.*

Proof. If σ is true in V then, in any non-trivial extension of V , there is a proper inner model in which σ holds, namely V itself. So if the inner-model hypothesis holds, then there must already be such an inner model of V , and so the lightface inner-model reflection principle holds. \square

Let us now consider the downward-directed version of the maximality principle studied in Hamkins-Löwe [HL13], which can be viewed itself as a kind of reflection principle. Namely, let us say that the *ground-model maximality principle* holds if and only if any statement σ that holds in some ground model and all grounds of that ground, is true in V . This is expressible as $\diamond\Box\sigma \rightarrow \sigma$, where \diamond and \Box are the modal operators of “true in some ground model” and “true in all ground models,” respectively.¹

By considering $\neg\sigma$ and the contrapositive, the ground-model maximality principle is easily seen to be equivalent to the following assertion: if a sentence σ is true, then every ground model has itself a ground in which σ is true. This formulation of the maximality principle reveals it to be a particularly strong reflection principle, when combined with the assertion that indeed there are non-trivial grounds.

¹In light of the downward orientation of this axiom, however, the ‘maximality’ terminology may be distracting, as any deeper ground, for example, will also satisfy the ground-model maximality principle. What is being maximized here is not the model, but the collection of truths that are downward-necessary. The principle is related to S5 for grounds, as in [HL13], since the axiom $\diamond\Box\varphi \rightarrow \varphi$ is the defining axiom of S5 over S4; but the ground-model maximality principle is not identical to the validity of S5, as the principle is (trivially) true under the ground axiom, whereas the modal logic of grounds in this situation strictly exceeds S5.

Theorem 4.3. *If $V_\delta \prec V$, then V has a ground model W that satisfies the ground-model maximality principle for assertions allowing parameters of rank less than δ in W .*

Proof. The proof uses the recent result of Usuba [Usu17] showing that the strong downward-directed grounds hypothesis (DDG) holds; that is, for any set I , there is a ground W contained in $\bigcap_{r \in I} W_r$, where W_r denotes the r^{th} ground as in the statement of the ground-model enumeration theorem.

Assume $V_\delta \prec V$. By the strong DDG, there is a ground W with $W \subseteq W_r$ for all $r \in V_\delta$. Suppose that $W \models \diamond \square \varphi(a)$ for some $a \in V_\delta \cap W$. So W has a ground model U such that $U \models \square \varphi(a)$; that is, U satisfies $\varphi(a)$ and so does every ground of U . Since U is a ground of W , which is a ground of V , it follows that U is a ground of V and consequently $U = W_r$ for some r by the ground-model enumeration theorem. The least rank of an r whose corresponding ground W_r has the properties of U that we have mentioned is definable in V using a as a parameter, and consequently there is such an r already in V_δ . In this case, $W \subseteq U = W_r$ by the assumption on W . Since $W \subseteq U \subseteq V$, where W is a ground of V , it follows from the intermediate-model theorem (see [FHR15, fact 11] or Jech [Jec03, corollary 15.43]) that W is also a ground of U where $\square \varphi(a)$ holds, and so $\varphi(a)$ is true in W , as desired. \square

The previous argument, using the strong DDG, provides a more direct method than [HL13, theorem 8] of finding a model of the ground-model maximality principle.

Corollary 4.4. *If $V_\delta \prec V$ and V has no bedrock, that is, no minimal ground, then there is a ground model W satisfying the ground-model maximality principle and the ground-model reflection principle for assertions using parameters of rank less than δ .*

Proof. Assume $V_\delta \prec V$ and V has no bedrock. Let W be the ground model identified in theorem 4.3, which satisfies the ground-model maximality principle for assertions with parameters of rank less than δ in W . Since there is no minimal ground, it follows that W also has no minimal ground. By the ground-model maximality principle, every statement $\varphi(a)$ true in W with $a \in V_\delta \cap W$ is true densely often in the grounds of W , that is, true in some deeper ground of any given ground of W . Since there are such proper grounds of W , it follows that any such statement $\varphi(a)$ is true in some proper ground of W , and so W satisfies the ground-model reflection principle for these assertions. \square

5. FORCING AXIOMS

Let us consider next the question of whether strong forcing axioms might settle the inner-model or ground-model reflection principles. Work of Caicedo-Veličković [CV06, corollary 2] shows that it is relatively consistent with the bounded proper forcing axiom BPFA that the universe is a minimal model of BPFA, that is, that the universe satisfies BPFA, but has no proper inner model satisfying BPFA. In this situation, of course, the inner-model reflection principle must fail, and so it seems natural to inquire whether PFA or MM might outright refute the ground-model reflection principle. The next theorem shows, however, that this is not the case.

Theorem 5.1. *The proper forcing axiom PFA, as well as Martin’s Maximum MM, if consistent, are consistent with the ground-model reflection principle, as well as with its failure.*

Proof. We use the fact that both PFA and MM are necessarily indestructible by $<\omega_2$ -directed closed forcing, see Larson [Lar00, theorem 4.3]. Let \mathbb{P} be the Easton-support class product used in the proof of theorem 2.2, but with non-trivial forcing factors only at stages ω_2 and above. This forcing is consequently ω_2 -directed closed, and therefore preserves PFA and MM, if these forcing axioms should hold in the ground. The proof of theorem 2.2 then shows that the extension satisfies the ground-model reflection principle, as desired.

Meanwhile, if (after suitable preparatory forcing coding sets into the GCH pattern) one should use an Easton-support iteration, rather than a product, then again the forcing is ω_2 -directed closed, and the main result of Hamkins-Reitz-Woodin [HRW08] shows that the extension satisfies the ground axiom: it has no non-trivial ground models for set forcing. Thus, the ground-model reflection principle fails in this extension. \square

6. EXPRESSIBILITY OF INNER MODEL REFLECTION

Lastly, let us consider the question of whether the inner-model reflection principle might be expressible in the first-order language of set theory, or whether, as we expect, it is a fundamentally second-order assertion. The ground-model reflection principle, as we have pointed out, is expressible as a schema in the first-order language of set theory. But the same does not seem to be true for the inner-model reflection principle, in light of the quantification over inner models. How can we prove that indeed there is no first-order means of expressing the principle?

As a step towards this, let us first show that the existence of an inner model satisfying a given sentence σ is not necessarily first-order expressible.

Denote by (m) the following consistency assumption, expressed in the language of proper classes:

$$V = L + \text{“there is a truth-predicate } \text{Tr}^L \text{ for truth in } L\text{.”}$$

The m stands for *mild*. Note that Tr^L is a non-definable class, which satisfies the Tarskian recursion for the definition of satisfaction for first-order truth. The consistency strength of (m) is strictly smaller than that of Kelly-Morse set theory, which is itself strictly weaker than ZFC plus the existence of an inaccessible cardinal, since KM implies the existence of such a truth predicate. So (m) is indeed a mild consistency assumption. See Gitman-Hamkins-Holy-Schlicht-Williams [GHH⁺] for further discussion of truth predicates and the strength of this hypothesis in the hierarchy between GBC and KM.

Theorem 6.1. *Assuming (m) , it is consistent that there are models $\langle M, \in, S_0 \rangle$ and $\langle M, \in, S_1 \rangle$ of GBC set theory with the same first-order part M and a particular first-order sentence σ , such that S_1 has a proper inner model of M satisfying σ , but S_0 does not.*

Proof. In L , let \mathbb{P} be the Easton-support product that adds a Cohen subset via $\text{Add}(\delta, 1)$ for every regular cardinal δ . Suppose $G \subseteq \mathbb{P}$ is L -generic, and consider $L[G]$. Let S_0 consist of the classes that are definable in the structure $\langle L[G], \in, G \rangle$, and let S_1 be the classes definable in $\langle L[G], \in, G, \text{Tr}^L \rangle$, where we also add the truth

predicate. (A similar argument is made in Hamkins-Reitz [HR].) Note that in light of the definability of the forcing relation for this forcing, it follows that S_1 includes a truth predicate for the extension $L[G]$.

Both models $\langle L[G], \in, S_0 \rangle$ and $\langle L[G], \in, S_1 \rangle$ satisfy GBC. Inside the latter model, let T be a class of regular cardinals that codes the information of the truth predicate Tr^L in some canonical and sufficiently absolute manner, such as by including $\aleph_{\alpha+1}^L$ in T exactly when α codes a formula-parameter pair $\varphi[\vec{a}]$ that is declared true by Tr^L . Let G_T be the restriction of the generic filter G to include the Cohen sets only on the cardinals in T . It is not difficult to see that in $L[G_T]$, the cardinals that have L -generic Cohen subsets of a regular cardinal δ are precisely the cardinals in T . Therefore, the model $L[G_T]$ satisfies the assertion that “the class of regular cardinals δ for which there is an L -generic Cohen subset of δ codes a truth-predicate for truth in L .” So S_1 has an inner model, namely $L[G_T]$, that satisfies this statement. But S_0 can have no such inner model, since there can be no truth predicate for L definable in $\langle L[G], \in, G \rangle$, as in this case we could use the definable forcing relation and thereby define a truth predicate for the full structure $\langle L[G], \in, G \rangle$ itself, contrary to Tarski’s theorem. \square

In order to show that the inner-model reflection principle is not first-order expressible, however, one would need much more than this. It would suffice to exhibit a positive instance of the following:

Question 6.2. Are there two models of GBC with the same first-order part, such that one of them is a model of the inner-model reflection principle and the other is not?

In particular, for this to happen we would at the very least need to strengthen theorem 6.1 by producing a model $\langle M, \in, S \rangle \models \text{GBC}$ and a sentence σ that is true in M and also true in some inner model $W \subsetneq M$ in S , but which is not true in any proper inner model of M that is first-order definable in M allowing set parameters. We are unsure how to arrange even this much.

Acknowledgements. The authors would like to thank Philip Welch for his interest in this project, helpful comments and suggestions. Thanks are also due to the anonymous referee for their valuable feedback. Neil Barton is very grateful for the generous support of the FWF (Austrian Science Fund) through Project P 28420 (*The Hyperuniverse Programme*). Gunter Fuchs was supported in part by PSC-CUNY grant 60630-00 48. Ralf Schindler gratefully acknowledges support by the DFG grant SCHI 484/8-1, “Die Geologie Innerer Modelle”.

REFERENCES

- [ABF] Carolin Antos, Neil Barton, and Sy-David Friedman. Universism and extensions of V . <https://arXiv.org/abs/1708.05751>.
- [Bar16] Neil Barton. What is the consistency strength of width reflection? Mathematics Stack Exchange question, 2016. <https://math.stackexchange.com/q/1912624> (version: 2016-09-03).
- [CV06] Andrés Eduardo Caicedo and Boban Veličković. The bounded proper forcing axiom and well orderings of the reals. *Math. Res. Lett.*, 13(2-3):393–408, 2006.
- [FHR15] Gunter Fuchs, Joel David Hamkins, and Jonas Reitz. Set-theoretic geology. *Annals of Pure and Applied Logic*, 166(4):464–501, 2015.
- [Fri06] Sy-David Friedman. Internal consistency and the inner model hypothesis. *Bull. Symbolic Logic*, 12(4):591–600, 2006.

- [FS16] Gunter Fuchs and Ralf Schindler. Inner model theoretic geology. *Journal of Symbolic Logic*, 81(3):972–996, 2016.
- [GHH⁺] Victoria Gitman, Joel David Hamkins, Peter Holy, Philipp Schlicht, and Kameryn Williams. The exact strength of the class forcing theorem. <https://arXiv.org/abs/1707.03700>.
- [Ham] Joel David Hamkins. *Forcing and large cardinals*. Book manuscript in preparation.
- [Ham00] Joel David Hamkins. The lottery preparation. *Ann. Pure Appl. Logic*, 101(2-3):103–146, 2000.
- [Ham03] Joel David Hamkins. A simple maximality principle. *J. Symbolic Logic*, 68(2):527–550, 2003.
- [Ham05] Joel David Hamkins. The Ground Axiom. *Mathematisches Forschungsinstitut Oberwolfach Report*, 55:3160–3162, 2005.
- [HL13] Joel David Hamkins and Benedikt Löwe. Moving up and down in the generic multiverse. *Logic and its Applications, ICLA 2013 LNCS*, 7750:139–147, 2013.
- [HR] Joel David Hamkins and Jonas Reitz. The set-theoretic universe V is not necessarily a class-forcing extension of HOD. <https://arxiv.org/abs/1709.06062>.
- [HRW08] Joel David Hamkins, Jonas Reitz, and W. Hugh Woodin. The ground axiom is consistent with $V \neq \text{HOD}$. *Proc. Amer. Math. Soc.*, 136(8):2943–2949, 2008.
- [Jec03] Thomas Jech. *Set Theory*. Springer Monographs in Mathematics, 3rd edition, 2003.
- [JS13] Ronald Jensen and John Steel. K without the measurable. *J. Symbolic Logic*, 78(3):708–734, 2013.
- [Lar00] Paul Larson. Separating stationary reflection principles. *J. Symbolic Logic*, 65(1):247–258, 2000.
- [Mit12] William J. Mitchell. Inner models for large cardinals. In *Sets and extensions in the twentieth century*, volume 6 of *Handb. Hist. Log.*, pages 415–456. Elsevier/North-Holland, Amsterdam, 2012.
- [Rei06] Jonas Reitz. *The Ground Axiom*. PhD thesis, The Graduate Center of the City University of New York, September 2006.
- [Rei07] Jonas Reitz. The ground axiom. *J. Symbolic Logic*, 72(4):1299–1317, 2007.
- [Sch14] Farmer Schlutzenberg. The definability of \mathbb{E} in self-iterable mice. <https://arxiv.org/abs/1412.0085>.
- [SS18] Grigor Sargsyan and Ralf Schindler. Varsovian models I. *J. Symb. Log.*, 83(2):496–528, 2018.
- [SS09] Ralf Schindler and John Steel. The self-iterability of $L[E]$. *Journal of Symbolic Logic*, 74(3):751–779, 2009.
- [SV02] Jonathan Stavi and Jouko Väänänen. Reflection principles for the continuum. In *Logic and algebra*, volume 302 of *Contemp. Math.*, pages 59–84. Amer. Math. Soc., Providence, RI, 2002.
- [Usu17] Toshimichi Usuba. The downward directed grounds hypothesis and very large cardinals. *J. Math. Log.*, 17(2), 1750009, 24 pp., 2017.
- [VW01] J. Vickers and P. D. Welch. On elementary embeddings from an inner model to the universe. *J. Symbolic Logic*, 66(3):1090–1116, 2001.
- [Wel04] P. D. Welch. On unfoldable cardinals, ω -closed cardinals, and the beginning of the inner model hierarchy. *Arch. Math. Logic*, 43(4):443–458, 2004.

(Neil Barton) KURT GÖDEL RESEARCH CENTER FOR MATHEMATICAL LOGIC, WÄHRINGER STRASSE, 25, 1090, VIENNA, AUSTRIA

E-mail address: neil.barton@univie.ac.at

URL: <https://neilbarton.net>

(Andrés Eduardo Caicedo) MATHEMATICAL REVIEWS, 416 FOURTH STREET, ANN ARBOR, MI 48103-4820, USA

E-mail address: aec@ams.org

URL: <http://www-personal.umich.edu/~caicedo>

(Gunter Fuchs) MATHEMATICS, 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, USA

E-mail address: gunter.fuchs@csi.cuny.edu

URL: <http://www.math.csi.cuny.edu/~fuchs>

(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, USA

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

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

(Jonas Reitz) NEW YORK CITY COLLEGE OF TECHNOLOGY OF THE CITY UNIVERSITY OF NEW YORK, MATHEMATICS, 300 JAY STREET, BROOKLYN, NY 11201, USA

E-mail address: jreitz@citytech.cuny.edu

(Ralf Schindler) INSTITUT FÜR MATHEMATISCHE LOGIK UND GRUNDLAGENFORSCHUNG, FACHBEREICH MATHEMATIK UND INFORMATIK, UNIVERSITÄT MÜNSTER, EINSTEINSTRASSE 62, 48149 MÜNSTER, GERMANY

E-mail address: rds@wwu.de