Gödel's Program

John R. Steel
University of California, Berkeley

June 2013

Reference: Gödel's program, to appear in *Interpreting Gödel*, Juliette Kennedy ed., Cambridge Univ. Press.

Abstract: Set theorists have discovered many mutually incompatible natural theories extending ZFC. It is possible that these incompatibilities will be resolved by interpreting all such theories in a useful common framework theory.

Plan:

I. Introduction.

II. The consistency-strength hierarchy

III. A theory of the concrete.

IV. A boundary.

V. The multiverse language.

VI. A core for the multiverse?

VII. Gödel's program.

I. Introduction

Let LST be the language of set theory, i.e. its syntax coupled with the meaning we currently assign to that syntax. Let ZFC be the axioms of Zermelo-Fraenkel with Choice.

I. Introduction

Let LST be the language of set theory, i.e. its syntax coupled with the meaning we currently assign to that syntax. Let ZFC be the axioms of Zermelo-Fraenkel with Choice.

(1) All mathematical language can be translated into LST.
(2) Not all mathematically interesting statements are decided by ZFC.

LST is semantically complete, but ZFC is proof-theoretically incomplete.

Gödel's program: Decide mathematically interesting questions independent of ZFC in well-justified extensions of ZFC.

Gödel's program: Decide mathematically interesting questions independent of ZFC in well-justified extensions of ZFC.

(1) Gödel was concerned mainly with the Continuum Hypothesis (CH).

Gödel's program: Decide mathematically interesting questions independent of ZFC in well-justified extensions of ZFC.

(1) Gödel was concerned mainly with the Continuum Hypothesis (CH).

(2) How does one justify statements in LST? General philosophical questions concerning meaning, evidence, and belief arise.

Gödel's program: Decide mathematically interesting questions independent of ZFC in well-justified extensions of ZFC.

(1) Gödel was concerned mainly with the Continuum Hypothesis (CH).

(2) How does one justify statements in LST? General philosophical questions concerning meaning, evidence, and belief arise.

(3) For those who believe the truth value of CH is not determined by the meaning we currently assign to the syntax of LST, the Continuum Problem does not disappear. Certainly we don't want to employ a syntax which encourages us to ask psuedo-questions, and the problem then becomes how to flesh out the current meaning, or trim back the current syntax, so that we can stop asking psuedo-questions.

Maximize!

The key methodological maxim that epistemology can contribute to the search for a stronger foundation for mathematics is:

maximize interpretative power.

Our foundational language and theory should enable us to say as much as possible, as efficiently as possible.

Maximize!

The key methodological maxim that epistemology can contribute to the search for a stronger foundation for mathematics is:

maximize interpretative power.

Our foundational language and theory should enable us to say as much as possible, as efficiently as possible.

The idea that set theorists ought to seek a language and theory that maximizes interpretative power seems to carry us a long way. We shall discuss just how far it goes in this talk.

Maximize!

The key methodological maxim that epistemology can contribute to the search for a stronger foundation for mathematics is:

maximize interpretative power.

Our foundational language and theory should enable us to say as much as possible, as efficiently as possible.

The idea that set theorists ought to seek a language and theory that maximizes interpretative power seems to carry us a long way. We shall discuss just how far it goes in this talk.

Language and theory evolve together.

The consistency strength hierarchy

One thing set theorists have understood much better in the years since Gödel is the family of possible extensions of ZFC.

The consistency strength hierarchy

One thing set theorists have understood much better in the years since Gödel is the family of possible extensions of ZFC.

Underlying the great variety of consistent extensions of ZFC, and the corresponding wealth of models of ZFC, there is a good deal more order than might at first be apparent.

The consistency strength hierarchy

One thing set theorists have understood much better in the years since Gödel is the family of possible extensions of ZFC.

Underlying the great variety of consistent extensions of ZFC, and the corresponding wealth of models of ZFC, there is a good deal more order than might at first be apparent.

Definition

Let T and U be axiomatized theories extending ZFC; then $T \leq_{\mathrm{Con}} U$ iff ZFC proves $\mathrm{Con}(U) \Rightarrow \mathrm{Con}(T)$. If $T \leq_{\mathrm{Con}} U$ and $U \leq_{\mathrm{Con}} T$, then we write $T \equiv_{\mathrm{Con}} U$, and say that T and U have the same consistency strength, or are *equiconsistent*.

Large cardinal hypotheses play a very special role in our understanding of the consistency of theories extending ZFC.

Large cardinal hypotheses play a very special role in our understanding of the consistency of theories extending ZFC.

Many natural extensions T of ZFC have been shown to be consistent relative to some large cardinal hypothesis H, via the method of forcing. This method is so powerful that, at the moment, we know of no interesting T extending ZFC which seems unlikely to be provably consistent relative to some large cardinal hypothesis via forcing.

Large cardinal hypotheses play a very special role in our understanding of the consistency of theories extending ZFC.

Many natural extensions T of ZFC have been shown to be consistent relative to some large cardinal hypothesis H, via the method of forcing. This method is so powerful that, at the moment, we know of no interesting T extending ZFC which seems unlikely to be provably consistent relative to some large cardinal hypothesis via forcing.

Thus the extensions of ZFC via large cardinal hypotheses seem to be cofinal in the part of the consistency strength order on extensions of ZFC which we know about.

Large cardinal hypotheses play a very special role in our understanding of the consistency of theories extending ZFC.

Many natural extensions T of ZFC have been shown to be consistent relative to some large cardinal hypothesis H, via the method of forcing. This method is so powerful that, at the moment, we know of no interesting T extending ZFC which seems unlikely to be provably consistent relative to some large cardinal hypothesis via forcing.

Thus the extensions of ZFC via large cardinal hypotheses seem to be cofinal in the part of the consistency strength order on extensions of ZFC which we know about.

These days, the way a set theorist convinces people that T is consistent is to show by forcing that $T \leq_{\mathrm{Con}} H$ for some large cardinal hypothesis H.

We do have pretty good evidence that even quite strong large cardinal hypotheses like the existence of rank-to-rank embeddings are consistent with ZFC.

We do have pretty good evidence that even quite strong large cardinal hypotheses like the existence of rank-to-rank embeddings are consistent with ZFC.

In all cases, the evidence is basically the existence of a coherent theory in which the hypothesis plays a central role, a theory that extends in a natural way the theory we obtain from weaker hypotheses.

We do have pretty good evidence that even quite strong large cardinal hypotheses like the existence of rank-to-rank embeddings are consistent with ZFC.
In all cases, the evidence is basically the existence of a coherent theory in which the hypothesis plays a central role, a theory that extends in a natural way the theory we obtain from weaker hypotheses.

Natural consistency strengths wellordered: If T is a natural extension of ZFC, then there is an extension H axiomatized by large cardinal hypotheses such that $T \equiv_{\mathrm{Con}} H$. Moreover, $\leq_{\mathrm{Con}}$ is a prewellorder of the natural extensions of ZFC. In particular, if T and U are natural extensions of ZFC, then either $T \leq_{\mathrm{Con}} U$ or $U \leq_{\mathrm{Con}} T$.

III. A theory of the concrete

A set theory T is consistent just in case all its Π^0_1 consequences are true.

III. A theory of the concrete

A set theory T is consistent just in case all its Π^0_1 consequences are true.

Remarkably, climbing the consistency strength hierarchy in any natural way seems to decide uniquely not just Π^0_1 sentences, but more complicated sentences about the concrete as well. *Concrete* refers here to natural numbers, real numbers, and certain sets of real numbers.

III. A theory of the concrete

A set theory T is consistent just in case all its Π^0_1 consequences are true.

Remarkably, climbing the consistency strength hierarchy in any natural way seems to decide uniquely not just Π^0_1 sentences, but more complicated sentences about the concrete as well. *Concrete* refers here to natural numbers, real numbers, and certain sets of real numbers.

Definition

Let Γ be a set of sentences in the syntax of LST, and T a theory; then $\Gamma_T = \{\varphi \mid \varphi \in \Gamma \wedge T \vdash \varphi\}$.

A theory of the natural numbers:

Phenomenon: If T and U are natural extensions of ZFC, then

$$\begin{aligned} T \leq_{\mathrm{Con}} U &\Leftrightarrow (\Pi^0_1)_T \subseteq (\Pi^0_1)_U \\ &\Leftrightarrow (\Pi^0_\omega)_T \subseteq (\Pi^0_\omega)_U \end{aligned}$$

A theory of the natural numbers:

Phenomenon: If T and U are natural extensions of ZFC, then

$$\begin{aligned} T \leq_{\mathrm{Con}} U &\Leftrightarrow (\Pi^0_1)_T \subseteq (\Pi^0_1)_U \\ &\Leftrightarrow (\Pi^0_\omega)_T \subseteq (\Pi^0_\omega)_U \end{aligned}$$

Thus the wellordering of natural consistency strengths corresponds to a wellordering by inclusion of theories of the natural numbers. There is no divergence at the arithmetic level, if one climbs the consistency strength hierarchy in any natural way we know of.

A theory of the reals:

Phenomenon: Let T, U be natural theories of consistency

strength at least that of "there are infinitely many Woodin cardinals"; then either $(\Pi^1_\omega)_T \subseteq (\Pi^1_\omega)_U$, or $(\Pi^1_\omega)_U \subseteq (\Pi^1_\omega)_T$.

A theory of the reals:

Phenomenon: Let T, U be natural theories of consistency

strength at least that of "there are infinitely many Woodin cardinals"; then either $(\Pi^1_\omega)_T \subseteq (\Pi^1_\omega)_U$, or $(\Pi^1_\omega)_U \subseteq (\Pi^1_\omega)_T$.

In other words, the second-order arithmetic generated by natural theories is an eventually monotonically increasing function of their consistency strengths.

This "one road upward" phenomenon extends to statements about sets of reals generated by reasonably simple means.

This "one road upward" phenomenon extends to statements about sets of reals generated by reasonably simple means.

There is a partial explanation of the phenomena of non-divergence, eventual monotonicity, and practical completeness in the realm of the concrete, for theories of sufficiently high consistency strength. It lies in the way we obtain independence theorems, by interpreting one theory in another.

This "one road upward" phenomenon extends to statements about sets of reals generated by reasonably simple means.

There is a partial explanation of the phenomena of non-divergence, eventual monotonicity, and practical completeness in the realm of the concrete, for theories of sufficiently high consistency strength. It lies in the way we obtain independence theorems, by interpreting one theory in another.

Our model-producing methods lead to eventual Γ-monotonicity because in order to produce a model for a theory T that is sufficiently strong with respect to Γ, we must produce a Γ-correct model.

IV. The Levy-Solovay boundary

None of our current large cardinal axioms decide CH, because they are preserved by small forcing, whilst CH can both be made true and made false by small forcing. Because CH is provably not generically absolute, it cannot be decided by large cardinal hypotheses that are themselves generically absolute.

IV. The Levy-Solovay boundary

None of our current large cardinal axioms decide CH, because they are preserved by small forcing, whilst CH can both be made true and made false by small forcing. Because CH is provably not generically absolute, it cannot be decided by large cardinal hypotheses that are themselves generically absolute.

Theorem (Levy, Solovay)

Let A be one of the current large cardinal axioms, and suppose $V \models A$; then there are set generic extensions M and N of V which satisfy $A + \mathrm{CH}$ and $A + \neg\mathrm{CH}$ repectively.

IV. The Levy-Solovay boundary

None of our current large cardinal axioms decide CH, because they are preserved by small forcing, whilst CH can both be made true and made false by small forcing. Because CH is provably not generically absolute, it cannot be decided by large cardinal hypotheses that are themselves generically absolute.

Theorem (Levy, Solovay)

Let A be one of the current large cardinal axioms, and suppose $V \models A$; then there are set generic extensions M and N of V which satisfy $A + \mathrm{CH}$ and $A + \neg\mathrm{CH}$ repectively.

CH, is a Σ^2_1 statement. It is the simplest sort of statement large cardinals do not decide. There are many more of them in general set theory.

V. The multiverse language

What does this picture of what is possible suggest as to what we should believe, or give preferred development, as a framework theory?

V. The multiverse language

What does this picture of what is possible suggest as to what we should believe, or give preferred development, as a framework theory?

We have good evidence that the consistency hierarchy is not a mirage, that the theories in it we have identified are indeed consistent. This argues for developing the theories in this hierarchy. All their Π^0_1 consequences are true, and we know of no other way to produce new Π^0_1 truths.

Developing one natural theory develops them all, via the boolean-valued interpretations. At the level of statements about the concrete (including most of what non-set-theorists say), all the natural theories agree.

Developing one natural theory develops them all, via the boolean-valued interpretations. At the level of statements about the concrete (including most of what non-set-theorists say), all the natural theories agree.

This might suggest that we need no further framework: why not simply develop all the natural theories in our hierarchy as tools for generating true statements about the concrete? Let 1000 flowers bloom! This is Hilbertism without the consistency proof, and with perhaps an enlarged class of "real" statements.

Unify!

The problem with this watered-down Hilbertism is that we don't want everyone to have his own private mathematics. We want one framework theory, to be used by all, so that we can use each other's work. It's better for all our flowers to bloom in the same garden. If truly distinct frameworks emerged, the first order of business would be to unify them.

Unify!

The problem with this watered-down Hilbertism is that we don't want everyone to have his own private mathematics. We want one framework theory, to be used by all, so that we can use each other's work. It's better for all our flowers to bloom in the same garden. If truly distinct frameworks emerged, the first order of business would be to unify them.

In fact, the different natural theories we have found in our hierarchy are not independent of one another. Their common theory of the concrete stems from logical relationships that go deeper, and are brought out in our relative consistency proofs. These logical relationships may suggest a unifying framework.

Large cardinals, our source of interpretative power.

The central role of the theories axiomatized by large cardinal hypotheses argues for adding such hypotheses to our framework.

Large cardinals, our source of interpretative power.

The central role of the theories axiomatized by large cardinal hypotheses argues for adding such hypotheses to our framework.

The goal of our framework theory is to *maximize interpretative power*, to provide a language and theory in which all mathematics, of today, and of the future so far as we can anticipate it today, can be developed.

Large cardinals, our source of interpretative power.

The central role of the theories axiomatized by large cardinal hypotheses argues for adding such hypotheses to our framework.

The goal of our framework theory is to *maximize interpretative power*, to provide a language and theory in which all mathematics, of today, and of the future so far as we can anticipate it today, can be developed.

Maximizing interpretative power entails maximizing consistency strength, but it requires more, in that we want to be able to translate other theories/languages into our framework theory/language in such a way that we preserve their meaning. The way we interpret set theories today is to think of them as theories of inner models of generic extensions of models satisfying some large cardinal hypothesis, and this method has had amazing success. We don't seem to lose any meaning this way. It is natural then to build on this approach.

Beyond large cardinals?

Nevertheless, large cardinal hypotheses like our current ones cannot decide CH, and so our theory of the concrete still has many different possible theoretical superstructures, some with CH, some with $\diamondsuit$, some with MM, some with $2^{\aleph_0}$ being real-vaulued measurable, and so on: all the behaviors that can hold in set-generic extensions of V, no matter what large cardinals exist.

Beyond large cardinals?

Nevertheless, large cardinal hypotheses like our current ones cannot decide CH, and so our theory of the concrete still has many different possible theoretical superstructures, some with CH, some with $\diamondsuit$, some with MM, some with $2^{\aleph_0}$ being real-vaulued measurable, and so on: all the behaviors that can hold in set-generic extensions of V, no matter what large cardinals exist.

Before we try to decide whether some such theory is preferable to the others, let us try to find a neutral common ground on which to compare them. We seek a language in which all these theories can be unified, without bias toward any, in a way that exhibits their logical relationships, and shows clearly how they can be used together. That is, we want one neat package they all fit into.

Our neutral common ground.

We describe a *multiverse language*, and an open-ended *multiverse theory*, in an informal way. It is routine to formalize completely.

Multiverse language: usual syntax of set theory, with two sorts, for the *worlds* and for the *sets*.

Axioms of MV*:*

$(1)_\varphi$ φ^W, for every world W. (For each axiom φ of ZFC.)

(2) (a) Every world is a transitive proper class. An object is a set just in case it belongs to some world.
(b) If W is a world and $\mathbb{P} \in W$ is a poset, then there is a world of the form $W[G]$, where G is $\mathbb{P}$-generic over W.
(c) If U is a world, and $U = W[G]$, where G is $\mathbb{P}$-generic over W, then W is a world.
(d) (Amalgamation.) If U and W are worlds, then there are G, H set generic over them such that $W[G] = U[H]$.

The natural way to get a model of MV is as follows.

Let M be a transitive model of ZFC, and let G be M-generic for $\mathrm{Col}(\omega, < \mathsf{OR}^M)$. The worlds of the multiverse M^G are all those W such that

$$W[H] = M[G \restriction \alpha],$$

for some H set generic over W, and some $\alpha \in \mathsf{OR}^M$.

The natural way to get a model of MV is as follows.

Let M be a transitive model of ZFC, and let G be M-generic for $\mathrm{Col}(\omega, < \mathsf{OR}^M)$. The worlds of the multiverse M^G are all those W such that

$$W[H] = M[G \upharpoonright \alpha],$$

for some H set generic over W, and some $\alpha \in \mathsf{OR}^M$.

It follows from a result of Laver and Woodin that the full first order theory of M^G is independent of G, and present in M, uniformly over all M.

The natural way to get a model of MV is as follows.

Let M be a transitive model of ZFC, and let G be M-generic for $\mathrm{Col}(\omega, < \mathsf{OR}^M)$. The worlds of the multiverse M^G are all those W such that

$$W[H] = M[G \restriction \alpha],$$

for some H set generic over W, and some $\alpha \in \mathsf{OR}^M$.

It follows from a result of Laver and Woodin that the full first order theory of M^G is independent of G, and present in M, uniformly over all M.

That is, there is a recursive translation function t such that whenever M is a model of ZFC and G is $\mathrm{Col}(\omega, < \mathsf{OR}^M)$-generic over M, then

$$M^G \models \varphi \Leftrightarrow M \models t(\varphi),$$

for all sentences φ of the multiverse language. $t(\varphi)$ just says "φ is true in some (equivalently all) multiverse(s) obtained from me".

If $\mathcal{W}$ is a model of MV, then for any world $M \in \mathcal{W}$, there is a G such that $\mathcal{W} = M^G$. Thus assuming MV indicates then that we are using the multiverse language as a sublanguage of the standard one, in the way described above. Also, it is clear that if φ is any sentence in the multiverse language, then MV proves

$$\varphi \Leftrightarrow \text{ for all worlds } M,\ t(\varphi)^M \ \Leftrightarrow \text{ for some world } M,\ t(\varphi)^M.$$

Thus everything that can be said in the multiverse language can be said using just one world-quantifier.

One can add large cardinal hypotheses that are preserved by small forcing to MV as follows: given such a large cardinal hypothesis φ, we add "φ^W, for all worlds W" to MV.

One can add large cardinal hypotheses that are preserved by small forcing to MV as follows: given such a large cardinal hypothesis φ, we add "φ^W, for all worlds W" to MV.

By adding large cardinal hypotheses to MV this way, we get as theorems "for all worlds W, φ^W", for any φ in the theory of the concrete they generate.

One can add large cardinal hypotheses that are preserved by small forcing to MV as follows: given such a large cardinal hypothesis φ, we add "φ^W, for all worlds W" to MV.

By adding large cardinal hypotheses to MV this way, we get as theorems "for all worlds W, φ^W", for any φ in the theory of the concrete they generate.

There is no obvious way to state CH in the multiverse language.

Have we lost expressive power?

One can think of the standard language as the multiverse language, together with a constant symbol $\dot{V}$ for a reference universe. Statements like CH are intended as statements about the reference universe. To what extent is this constant symbol meaningful? Does one lose anything by retreating to the superficially less expressive multiverse language? We distinguish three answers to this question:

Have we lost expressive power?

One can think of the standard language as the multiverse language, together with a constant symbol $\dot{V}$ for a reference universe. Statements like CH are intended as statements about the reference universe. To what extent is this constant symbol meaningful? Does one lose anything by retreating to the superficially less expressive multiverse language? We distinguish three answers to this question:

Weak relativist thesis: Every proposition that can be expressed in the standard language LST can be expressed in the multiverse language.

Have we lost expressive power?

One can think of the standard language as the multiverse language, together with a constant symbol $\dot{V}$ for a reference universe. Statements like CH are intended as statements about the reference universe. To what extent is this constant symbol meaningful? Does one lose anything by retreating to the superficially less expressive multiverse language? We distinguish three answers to this question:

Weak relativist thesis: Every proposition that can be expressed in the standard language LST can be expressed in the multiverse language.

Strong absolutist thesis: "$\dot{V}$" makes sense, and that sense is not expressible in the multiverse language.

Finally, perhaps weak relativism and the absolutist's idea of a distinguished reference world can be combined, in that that there is an individual world that is definable in the multiverse language.

An elementary forcing argument shows that if the multiverse has a definable world, then it has a unique definable world, and this world is included in all the others. (An observation due to Woodin.) In this case, we call this unique world included in all others the *core* of the multiverse.

Weak absolutist thesis: There are individual worlds that are definable in the multiverse language; that is, the multiverse has a core.

Why weak relativism?

The strongest evidence for the weak relativist thesis is that the mathematical theory based on large cardinal hypotheses that we have produced to date can be naturally expressed in the multiverse sublanguage.

Perhaps we lose something when we do that, some future mathematics built around an understanding of the symbol $\dot{V}$ that does not involve defining $\dot{V}$ in the multiverse language. But at the moment, it's hard to see what that is.

Why weak relativism?

The strongest evidence for the weak relativist thesis is that the mathematical theory based on large cardinal hypotheses that we have produced to date can be naturally expressed in the multiverse sublanguage.

Perhaps we lose something when we do that, some future mathematics built around an understanding of the symbol $\dot{V}$ that does not involve defining $\dot{V}$ in the multiverse language. But at the moment, it's hard to see what that is.

The weak relativist thesis can be considered as a piece of advice: don't go looking for it.

VI. Does the multiverse have a core?

Whatever one thinks of the semantic completeness of the multiverse language, it does bring the weak absolutist thesis to the fore, as a fundamental question. Because the multiverse language is a sublanguage of the standard one, this is a question for everyone. If the multiverse has a core, then surely it is important, whether it is the denotation of the absolutist's $\dot{V}$ or not! Indeed, if there is an inclusion-least world in the multiverse, why don't we use $\dot{H}$ to denote it , and agree to retire $\dot{V}$ until we need it? The question as to whether the multiverse has a core is an important question for everyone, relativist or absolutist.

VI. Does the multiverse have a core?

Whatever one thinks of the semantic completeness of the multiverse language, it does bring the weak absolutist thesis to the fore, as a fundamental question. Because the multiverse language is a sublanguage of the standard one, this is a question for everyone. If the multiverse has a core, then surely it is important, whether it is the denotation of the absolutist's $\dot{V}$ or not! Indeed, if there is an inclusion-least world in the multiverse, why don't we use $\dot{H}$ to denote it , and agree to retire $\dot{V}$ until we need it? The question as to whether the multiverse has a core is an important question for everyone, relativist or absolutist.

Neither MV nor its extensions by large cardinal hypotheses of the sort we currently understand decides whether there is a core to the multiverse, or the basic theory of this core if it exists. (Hamkins, Reitz, Woodin.) So what we have here is another basic question, like the CH, that large cardinals do not decide. But it is a different question, and its role in our search for a universal framework theory seems crucial.

Woodin's Axiom H.

There is some reason to hope for a positive answer. Hugh Woodin has recently proposed an axiom which

(a) implies the multiverse has a core,

(b) suggests an approach toward developing a detailed, systematic "fine structure theory" for this core, and

(c) may be consistent with all our large cardinal hypotheses.

The new mathematics needed in order to turn (b) and (c) from promise into reality is formidable, but there is some reason for optimism. One can hope that this axiom pins down our multiverse, without restricting it.

V looks like the HOD of a model of AD

Recall that a set is *ordinal definable* (OD) iff it is definable over the universe of sets from ordinal parameters, and is *hereditarily ordinal definable* (HOD) just in case it and all members of its transitive closure are OD. Gödel first isolated HOD in the 1940s. Myhill and Scott showed that if $M \models \mathrm{ZF}$, then $\mathrm{HOD}^M \models \mathrm{ZFC}$.

V looks like the HOD of a model of AD

Recall that a set is *ordinal definable* (OD) iff it is definable over the universe of sets from ordinal parameters, and is *hereditarily ordinal definable* (HOD) just in case it and all members of its transitive closure are OD. Gödel first isolated HOD in the 1940s. Myhill and Scott showed that if $M \models \mathrm{ZF}$, then $\mathrm{HOD}^M \models \mathrm{ZFC}$.

Woodin's axiom says that V looks like HOD^M, for models M of the axiom of determinacy.

Axiom H. For any sentence φ of LST: if if φ is true, then for some $M \models \mathrm{AD}^+ + V = L(P(\mathbb{R}))$ such that $\mathbb{R} \cup \mathrm{OR} \subseteq M$, $(\mathrm{HOD} \cap V_\Theta)^M \models \varphi$.

V looks like the HOD of a model of AD

Recall that a set is *ordinal definable* (OD) iff it is definable over the universe of sets from ordinal parameters, and is *hereditarily ordinal definable* (HOD) just in case it and all members of its transitive closure are OD. Gödel first isolated HOD in the 1940s. Myhill and Scott showed that if $M \models \mathrm{ZF}$, then $\mathrm{HOD}^M \models \mathrm{ZFC}$.

Woodin's axiom says that V looks like HOD^M, for models M of the axiom of determinacy.

Axiom H. For any sentence φ of LST: if if φ is true, then for some $M \models \mathrm{AD}^+ + V = L(P(\mathbb{R}))$ such that $\mathbb{R} \cup \mathrm{OR} \subseteq M$, $(\mathrm{HOD} \cap V_\Theta)^M \models \varphi$.

The schema is stated above in the standard language, but Woodin has shown that it implies that V is the core of its own multiverse. So one could state Axiom H in the multiverse language: the multiverse has a core, and it satisfies Axiom H.

The hope is that Axiom H is consistent with all the large cardinal hypotheses, so that adopting it does not restrict interpretative power. It is known to be consistent with the existence of Woodin cardinals. Whether it is consistent with significantly stronger large cardinal hypotheses is a crucial open problem.

The hope is that Axiom H is consistent with all the large cardinal hypotheses, so that adopting it does not restrict interpretative power. It is known to be consistent with the existence of Woodin cardinals. Whether it is consistent with significantly stronger large cardinal hypotheses is a crucial open problem.

At the same time, one hopes that Axiom H will yield a detailed fine structure theory for V, removing the incompleteness that large cardinal hypotheses by themsevles can never remove. It is known that Axiom H implies the CH, and many instances of the GCH. Whether it implies the full GCH is a crucial open problem.

The hope is that Axiom H is consistent with all the large cardinal hypotheses, so that adopting it does not restrict interpretative power. It is known to be consistent with the existence of Woodin cardinals. Whether it is consistent with significantly stronger large cardinal hypotheses is a crucial open problem.

At the same time, one hopes that Axiom H will yield a detailed fine structure theory for V, removing the incompleteness that large cardinal hypotheses by themsevles can never remove. It is known that Axiom H implies the CH, and many instances of the GCH. Whether it implies the full GCH is a crucial open problem.

Axiom H can be stated in the multiverse language. The strong absolutist who believes that Axiom H is false must still face the question whether the multiverse has a core satisfying Axiom H. If he agrees that it does, then the argument between him and someone who accepts Axiom H as a strong absolutist to have little practical importance.

VII. Gödel's program.

What are our prospects today for reaching Gödel's original goal, deciding the CH?

VII. Gödel's program.

What are our prospects today for reaching Gödel's original goal, deciding the CH?

The work of Gödel, Cohen, and their successors has shown us just how important the metamathematics of set theory is in this endeavor. An understanding of the possible set theories is useful in finding the true one. Our metamathematical work has also shown us just how many other natural questions are in the same boat as the CH, and thereby broadened our focus. Most likely, the Continuum Problem will not be solved without a significant, far-reaching clarification of the notion of set. Our metamathematical work is a necessary prelude to that.

Our current understanding of the possibilities for maximizing interpretative power has led us to to one theory of the concrete, and a family of theoretical superstructures for it, each containing all the large cardinal hypotheses. These different theories are logically related in a way that enables us to use them all together. Whatever the strong absolutist may believe about V, it is surely an important fact about the global structure of V that it has the generic extensions it does have.

Our current understanding of the possibilities for maximizing interpretative power has led us to to one theory of the concrete, and a family of theoretical superstructures for it, each containing all the large cardinal hypotheses. These different theories are logically related in a way that enables us to use them all together. Whatever the strong absolutist may believe about V, it is surely an important fact about the global structure of V that it has the generic extensions it does have.

The logical relationships between the different theories extending ZFC plus large cardinal hypotheses we have discovered are brought out clearly by formalizing them in the multiverse language. This language is a sublanguage of the standard one, and in it we can formalize naturally all the mathematics that set theorists have done. Remaining within this sublanguage has the additional virtue that our attention is directed away from CH, which has no obvious formalization within it, and toward the global question as to whether the multiverse has a core.

Remarkably, we can see now the outlines of a positive answer to this question, a way in which the multiverse may indeed have a core, and this core may admit a detailed fine-structural analysis that resembles that of Gödel's L. There are formidable technical mathematical problems that need to be answered in a certain way to realize this promise: we must show that there are models M of AD^+ such that HOD^M satisfies "there are supercompact cardinals", for example, and we must produce a fine structure theory for HOD^M. Although these are difficult questions, large cardinal hypotheses should settle them.

Remarkably, we can see now the outlines of a positive answer to this question, a way in which the multiverse may indeed have a core, and this core may admit a detailed fine-structural analysis that resembles that of Gödel's L. There are formidable technical mathematical problems that need to be answered in a certain way to realize this promise: we must show that there are models M of AD^+ such that HOD^M satisfies "there are supercompact cardinals", for example, and we must produce a fine structure theory for HOD^M. Although these are difficult questions, large cardinal hypotheses should settle them.

Perhaps the mathematics will turn out some other way. Perhaps the multiverse has no core, but some other, more subtle structure. There are many basic open questions at the foundations of set theory: the extent of generic absoluteness, the existence of iterable structures, the Ω-conjecture, the form of canonical inner models with supercompacts, and the properties of HOD in models of determinacy, to give my own partial list. Our path toward a stronger foundation will be lit by the answers to such questions.