Tightness in second-order arithmetic

Kameryn J. Williams

Sam Houston State University

Models of Peano Arithmetic Seminar, CUNY
2022 Oct 18

Joint work with Alfredo Roque Freire

PA has some nice properties

As befitting an important foundational theory, PA enjoys some nice properties.

PA has some nice properties

As befitting an important foundational theory, PA enjoys some nice properties.

- PA isn't finitely axiomatizable;
- For each formula $\varphi(x)$, PA proves $\varphi(x)$ admits a definable Skolem function;
- (Mostowski) For each finite $T \subseteq \mathrm{PA}$, PA proves $\mathrm{Con}(T)$;

PA has some nice properties

As befitting an important foundational theory, PA enjoys some nice properties.

- PA isn't finitely axiomatizable;
- For each formula $\varphi(x)$, PA proves $\varphi(x)$ admits a definable Skolem function;
- (Mostowski) For each finite $T \subseteq \mathsf{PA}$, PA proves $\mathrm{Con}(T)$;
- (Visser) If T_0, T_1 are extensions of PA, then T_0 and T_1 are bi-interpretable iff they have the same deductive closure.

Tightness

Definition

A theory T is tight if any two deductively complete extensions of T in the same language are bi-interpretable iff they are identical.

T is semantically tight if any two bi-interpretable models of T are isomorphic.

Tightness

Definition

A theory T is tight if any two deductively complete extensions of T in the same language are bi-interpretable iff they are identical.

T is semantically tight if any two bi-interpretable models of T are isomorphic.

(Without the same language restriction this is trivial. Consider e.g. PA + "the new unary predicate is the evens" versus PA + "the new unary predicate is the odds".)

Tightness

Definition

A theory T is tight if any two deductively complete extensions of T in the same language are bi-interpretable iff they are identical.
T is semantically tight if any two bi-interpretable models of T are isomorphic.

(Without the same language restriction this is trivial. Consider e.g. PA + "the new unary predicate is the evens" versus PA + "the new unary predicate is the odds".)

The following theories are both tight and semantically tight:.

- PA (Visser)
- ZF (Enayat)
- Z_2 (Enayat)
- KM (Enayat)

Tightness

Definition

A theory T is tight if any two deductively complete extensions of T in the same language are bi-interpretable iff they are identical.
T is semantically tight if any two bi-interpretable models of T are isomorphic.

The following theories are both tight and semantically tight:.

- PA (Visser)
- ZF (Enayat)
- Z_2 (Enayat)
- KM (Enayat)

(Without the same language restriction this is trivial. Consider e.g. PA + "the new unary predicate is the evens" versus PA + "the new unary predicate is the odds".)

For example, ZF + CH and ZF + ¬CH are mutually interpretable. (ZF + CH can be interpreted as L, and ZF + ¬CH can be interpreted through the boolean ultrapower approach to forcing.) But these interpretations lose information, and there is no way to produce a bi-interpretation.

The main question

Each of these tight theories have a natural hierarchy of increasingly stronger fragments.

$$\mathsf{I}\Sigma_0 \subseteq \mathsf{I}\Sigma_1 \subseteq \cdots \subseteq \mathsf{I}\Sigma_k \subseteq \cdots \subseteq \mathsf{PA}$$

$$\mathsf{ACA}_0 \subseteq \Pi^1_1\text{-}\mathsf{CA}_0 \subseteq \cdots \subseteq \Pi^1_k\text{-}\mathsf{CA}_0 \subseteq \cdots \subseteq \mathsf{Z}_2$$

Do we need the full strength of the theory to get tightness? Or are these fragments also tight?

The main question

Each of these tight theories have a natural hierarchy of increasingly stronger fragments.

$$\mathsf{I}\Sigma_0 \subseteq \mathsf{I}\Sigma_1 \subseteq \cdots \subseteq \mathsf{I}\Sigma_k \subseteq \cdots \subseteq \mathsf{PA}$$

$$\mathsf{ACA}_0 \subseteq \Pi^1_1\text{-}\mathsf{CA}_0 \subseteq \cdots \subseteq \Pi^1_k\text{-}\mathsf{CA}_0 \subseteq \cdots \subseteq \mathsf{Z}_2$$

Do we need the full strength of the theory to get tightness? Or are these fragments also tight?

We addressed this question for Z_2 and KM.

The main theorem

Theorem (Freire–W.)

The following theories are not tight.

- ACA;
- Π^1_k-CA, *for* $k \geq 1$;
- GB;
- $\mathrm{GB} + \Sigma^1_k$-*Comprehension, for* $k \geq 1$.

The main theorem

Theorem (Freire–W.)

The following theories are not tight.

- ACA;
- Π^1_k-CA, *for* $k \geq 1$;
- GB;
- $\mathsf{GB} + \Sigma^1_k$-*Comprehension, for* $k \geq 1$.

The constructions for arithmetic versus set theory are very similar.
I will talk about the arithmetic case, since this is MOPA.

A warm-up: ACA is not semantically tight

To prove this, it suffices to demonstrate two models of ACA which satisfy different theories but are bi-interpretable.

We will show that the minimum ω-model of ACA is bi-interpretable with a carefully chosen extension by Cohen forcing.

Since these two models satisfy a different theory, we will get the desired failure.

Identifying the minimum ω-model of ACA

- A model of second-order arithmetic is of the form $(M, \mathcal{X})$ where M are the numbers of the model and $\mathcal{X} \subseteq \mathcal{P}(M)$ are the sets.
- If $M \cong \omega$ then we call it an ω-model.

- ACA is axiomatized by:
  - the axioms of discretely ordered semirings;
  - induction in the full language, i.e. not just for arithmetical formulae; and
  - arithmetical comprehension.

Identifying the minimum ω-model of ACA

- A model of second-order arithmetic is of the form $(M, \mathcal{X})$ where M are the numbers of the model and $\mathcal{X} \subseteq \mathcal{P}(M)$ are the sets.
- If $M \cong \omega$ then we call it an ω-model.
- Any ω-model automatically satisfies full induction.

- ACA is axiomatized by:
  - the axioms of discretely ordered semirings;
  - induction in the full language, i.e. not just for arithmetical formulae; and
  - arithmetical comprehension.

Identifying the minimum ω-model of ACA

- A model of second-order arithmetic is of the form $(M, \mathcal{X})$ where M are the numbers of the model and $\mathcal{X} \subseteq \mathcal{P}(M)$ are the sets.
- If $M \cong \omega$ then we call it an ω-model.
- Any ω-model automatically satisfies full induction.
- It's easy to see that the minimum ω-model of ACA is $(\omega, \mathrm{Def}(\omega))$, the finite ordinals equipped with their arithmetically definable subsets.

I will write $\mathcal{D}$ for $\mathrm{Def}(\omega)$.

- ACA is axiomatized by:
  - the axioms of discretely ordered semirings;
  - induction in the full language, i.e. not just for arithmetical formulae; and
  - arithmetical comprehension.

Identifying the minimum ω-model of ACA

Let T denote the Tarskian satisfaction class for ω. By the undefinability of truth, $\mathrm{T} \notin \mathcal{D}$. Nevertheless, T is definable over $(\omega, \mathcal{D})$.

Identifying the minimum ω-model of ACA

Let T denote the Tarskian satisfaction class for ω. By the undefinability of truth, $\mathrm{T} \notin \mathcal{D}$. Nevertheless, T is definable over $(\omega, \mathcal{D})$.

- For each $k \in \omega$, the restriction T_k of T to Σ_k formulae is in $\mathcal{D}$.
- So we can define that $\varphi[a]$ is in T iff there exists k so that there exists a set satisfying the definition of T_k which judges $\varphi[a]$ to be true.
- (The T_k are not *uniformly* arithmetically definable, but the property of being a T_k is uniformly recognizable.)

Identifying the minimum ω-model of ACA

Let T denote the Tarskian satisfaction class for ω. By the undefinability of truth, $\mathrm{T} \notin \mathcal{D}$. Nevertheless, T is definable over $(\omega, \mathcal{D})$.

- For each $k \in \omega$, the restriction T_k of T to Σ_k formulae is in $\mathcal{D}$.
- So we can define that $\varphi[a]$ is in T iff there exists k so that there exists a set satisfying the definition of T_k which judges $\varphi[a]$ to be true.
- (The T_k are not *uniformly* arithmetically definable, but the property of being a T_k is uniformly recognizable.)

- This gives a Σ^1_1 definition of T.
- There's also Π^1_1 definition—any set that looks like a T_k which has $\varphi[a]$ in its domain judges $\varphi[a]$ to be true.
- So this is absolute between ω-models of ACA. They all define T the same.

Identifying the minimum ω-model of ACA, and codes for higher order sets

Because T is definable, so is the property "$X \in \mathcal{D}$":

- $X \in \mathcal{D}$ iff there is $\varphi[a, x]$ so that $X = \{x : \varphi[a, x] \in \mathrm{T}\}$.

So "every set is arithmetically definable" is a single second-order assertion, and the only ω-model of ACA which satisfies it is the minimum ω-model.

Identifying the minimum ω-model of ACA, and codes for higher order sets

Because T is definable, so is the property "$X \in \mathcal{D}$":

- $X \in \mathcal{D}$ iff there is $\varphi[a, x]$ so that $X = \{x : \varphi[a, x] \in \mathrm{T}\}$.

So "every set is arithmetically definable" is a single second-order assertion, and the only ω-model of ACA which satisfies it is the minimum ω-model.

$\mathcal{D}$ is a set of sets of integers, but it can be coded by a single set of integers. The elements of $\mathcal{D}$ are the slices of T.

Because ω has a canonical well-order, we have a canonical enumeration of the element of $\mathcal{D}$: order them by the order of their smallest index in T.

Relativizing truth and definability

Consider $C \subseteq \omega$.

- $\mathrm{T}(C)$ is the truth predicate with C as a predicate;
- $\mathcal{D}(C)$ is the sets arithmetically definable from C.

The facts about T and $\mathcal{D}$ generalize to give:

- If $\mathcal{X}$ is an ω-model of ACA with $C \in \mathcal{X}$ then $\mathrm{T}(C)$ is definable over $\mathcal{X}$ and so is the predicate "$X \in \mathcal{D}(C)$".

Relativizing truth and definability

Consider $C \subseteq \omega$.

- $\mathrm{T}(C)$ is the truth predicate with C as a predicate;
- $\mathcal{D}(C)$ is the sets arithmetically definable from C.

The facts about T and $\mathcal{D}$ generalize to give:

- If $\mathcal{X}$ is an ω-model of ACA with $C \in \mathcal{X}$ then $\mathrm{T}(C)$ is definable over $\mathcal{X}$ and so is the predicate "$X \in \mathcal{D}(C)$".

If $C \notin \mathcal{D}$, then $\mathrm{T}(C)$ in general needn't be definable over $\mathcal{D}$. (Quick proof: there are continuum many different C but only countably many definitions.)

Relativizing truth and definability

Consider $C \subseteq \omega$.

- $\mathrm{T}(C)$ is the truth predicate with C as a predicate;
- $\mathcal{D}(C)$ is the sets arithmetically definable from C.

The facts about T and $\mathcal{D}$ generalize to give:

- If $\mathcal{X}$ is an ω-model of ACA with $C \in \mathcal{X}$ then $\mathrm{T}(C)$ is definable over $\mathcal{X}$ and so is the predicate "$X \in \mathcal{D}(C)$".

If $C \notin \mathcal{D}$, then $\mathrm{T}(C)$ in general needn't be definable over $\mathcal{D}$. (Quick proof: there are continuum many different C but only countably many definitions.)

But if C is definable over $\mathcal{D}$ and generic over $\mathcal{D}$ for Cohen forcing then the truth lemma implies $\mathrm{T}(C)$ is definable over $\mathcal{D}$.

- An arithmetical formula $\varphi(C)$ is true iff there is $p \in C$ such that $p \Vdash \varphi(\dot{C})$.
- So we can define $\mathrm{T}(C)$ over $\mathcal{D}$ as: $\varphi[x, C] \in \mathrm{T}(C)$ iff there is $p \in C$ which forces $\varphi(x, \dot{C})$.

Defining a Cohen generic

Recall:

- Cohen forcing $\mathbb{P} = \mathrm{Add}(\omega, 1)$ is the infinite binary tree.
- A filter $C \subseteq \mathbb{P}$ is generic over $\mathcal{D}$ if it meets every dense subset of $\mathbb{P}$ from $\mathcal{D}$.

Defining a Cohen generic

Recall:

- Cohen forcing $\mathbb{P} = \mathrm{Add}(\omega, 1)$ is the infinite binary tree.
- A filter $C \subseteq \mathbb{P}$ is generic over $\mathcal{D}$ if it meets every dense subset of $\mathbb{P}$ from $\mathcal{D}$.

From T we have a canonical enumeration of the ω many dense subsets. Now follow the usual proof of the Rasiowa–Sikorski lemma:

- Start with $p_0 = \emptyset$;
- At stage $n+1$, extend p_n to the least condition in the n-th dense set which is below p_n, get $p_n + 1$
- Then define $C = \{q : q \geq p_n \text{ for some } n\}$.

Defining a Cohen generic

Recall:

- Cohen forcing $\mathbb{P} = \mathrm{Add}(\omega, 1)$ is the infinite binary tree.
- A filter $C \subseteq \mathbb{P}$ is generic over $\mathcal{D}$ if it meets every dense subset of $\mathbb{P}$ from $\mathcal{D}$.

From T we have a canonical enumeration of the ω many dense subsets. Now follow the usual proof of the Rasiowa–Sikorski lemma:

- Start with $p_0 = \emptyset$;
- At stage $n + 1$, extend p_n to the least condition in the n-th dense set which is below p_n, get $p_n + 1$
- Then define $C = \{q : q \geq p_n \text{ for some } n\}$.

Because we have a definable enumeration of the dense sets and we always pick the least condition, there is a uniform definition of the p_n. So C is definable. Note the definition quantifies over sets in $\mathcal{D}$.

Defining a Cohen generic

Recall:

- Cohen forcing $\mathbb{P} = \mathrm{Add}(\omega, 1)$ is the infinite binary tree.
- A filter $C \subseteq \mathbb{P}$ is generic over $\mathcal{D}$ if it meets every dense subset of $\mathbb{P}$ from $\mathcal{D}$.

From T we have a canonical enumeration of the ω many dense subsets. Now follow the usual proof of the Rasiowa–Sikorski lemma:

- Start with $p_0 = \emptyset$;
- At stage $n+1$, extend p_n to the least condition in the n-th dense set which is below p_n, get $p_n + 1$
- Then define $C = \{q : q \geq p_n \text{ for some } n\}$.

Because we have a definable enumeration of the dense sets and we always pick the least condition, there is a uniform definition of the p_n. So C is definable. Note the definition quantifies over sets in $\mathcal{D}$.

Because $\mathcal{D}$ is uniformly definable over any ω-model of ACA, any ω-model of ACA defines C the same.

Putting it all together

Let $\mathcal{U} = \mathcal{D}(C)$.

Theorem (Freire-W., independently Enayat)

$(\omega, \mathcal{D})$ and $(\omega, \mathcal{U})$ are bi-interpretable but satisfy different extensions of ACA.

Putting it all together

Let $\mathcal{U} = \mathcal{D}(C)$.

Theorem (Freire-W., independently Enayat)

$(\omega, \mathcal{D})$ and $(\omega, \mathcal{U})$ are bi-interpretable but satisfy different extensions of ACA.

That $(\omega, \mathcal{U}) \models$ ACA is because forcing preserves arithmetical comprehension. And it satisfies "there is a set which is not arithmetically definable" whereas $(\omega, \mathcal{D})$ satisfies "every set is arithmetically definable".

Finally, since you know that $\mathrm{T}(C)$ is definable over $\mathcal{D}$ it's easy to build the interpretations. Interpreting $\mathcal{D}$ in $\mathcal{U}$ is just restricting the domain of the sets, and for the other direction you can represent sets by their least index in $\mathrm{T}(C)$. And they form a bi-interpretation because the two models agree on T, C, and $\mathrm{T}(C)$.

From ω-models to non-tightness

To get a failure of tightness, we need a construction that works uniformly across any model (of an appropriate theory).

From ω-models to non-tightness

To get a failure of tightness, we need a construction that works uniformly across any model (of an appropriate theory).

It turns out essentially the same construction works.

From ω-models to non-tightness

To get a failure of tightness, we need a construction that works uniformly across any model (of an appropriate theory).
It turns out essentially the same construction works.

- If $(M, \mathcal{X}) \models$ ACA, then $\mathcal{X}$ has a Σ_k-satisfaction class for every $k \in M$.

(Because the set of such k is inductive.)

From ω-models to non-tightness

To get a failure of tightness, we need a construction that works uniformly across any model (of an appropriate theory).
It turns out essentially the same construction works.

- If $(M, \mathcal{X}) \models \mathrm{ACA}$, then $\mathcal{X}$ has a Σ_k-satisfaction class for every $k \in M$.

(Because the set of such k is inductive.)

- Corollary: If $(M, \mathcal{X}) \models \mathrm{ACA}$ then there is an inductive full satisfaction class over M. In particular, $M \models \mathrm{Con(PA)}$ and if M is ω-nonstandard then it is recursively saturated.

From ω-models to non-tightness

To get a failure of tightness, we need a construction that works uniformly across any model (of an appropriate theory).
It turns out essentially the same construction works.

- If $(M, \mathcal{X}) \models \mathrm{ACA}$, then $\mathcal{X}$ has a Σ_k-satisfaction class for every $k \in M$.

(Because the set of such k is inductive.)

- Corollary: If $(M, \mathcal{X}) \models \mathrm{ACA}$ then there is an inductive full satisfaction class over M. In particular, $M \models \mathrm{Con(PA)}$ and if M is ω-nonstandard then it is recursively saturated.

If M is countable and recursively saturated it admits continuum many different full satisfaction classes, so we cannot expect that all M-models of ACA will define T the same.

But if two M-models have the same Σ_k-satisfaction classes, then they define T the same. For example, this happens if one is a forcing extension of the other.

From ω-models to non-tightness

To get a failure of tightness, we need a construction that works uniformly across any model (of an appropriate theory).
It turns out essentially the same construction works.

- If $(M, \mathcal{X}) \models \mathrm{ACA}$, then $\mathcal{X}$ has a Σ_k-satisfaction class for every $k \in M$.

(Because the set of such k is inductive.)

- Corollary: If $(M, \mathcal{X}) \models \mathrm{ACA}$ then there is an inductive full satisfaction class over M. In particular, $M \models \mathrm{Con(PA)}$ and if M is ω-nonstandard then it is recursively saturated.

If M is countable and recursively saturated it admits continuum many different full satisfaction classes, so we cannot expect that all M-models of ACA will define T the same.

But if two M-models have the same Σ_k-satisfaction classes, then they define T the same. For example, this happens if one is a forcing extension of the other.

Observation: Any model of ACA has a minimum ω-submodel (= submodel that agrees on ω) of ACA.

Definability and truth in arbitrary models of ACA

- T is the union of the Σ_k-satisfaction classes.
- $\mathcal{D}$ consists of the slices of T.

So there is a second-order axiom expressing "every set is in $\mathcal{D}$".

Definability and truth in arbitrary models of ACA

- T is the union of the Σ_k-satisfaction classes.
- $\mathcal{D}$ consists of the slices of T.

So there is a second-order axiom expressing "every set is in $\mathcal{D}$".

- C is the Cohen generic over $\mathcal{D}$ constructed from the canonical enumeration of the dense subsets in $\mathcal{D}$ arising from T.
- (Full induction is what guarantees the construction of the p_n continues even for nonstandard n.)

So there is a second-order axiom expressing "C exists and every set is in $\mathcal{D}(C)$".

Definability and truth in arbitrary models of ACA

- T is the union of the Σ_k-satisfaction classes.
- $\mathcal{D}$ consists of the slices of T.

So there is a second-order axiom expressing "every set is in $\mathcal{D}$".

- C is the Cohen generic over $\mathcal{D}$ constructed from the canonical enumeration of the dense subsets in $\mathcal{D}$ arising from T.
- (Full induction is what guarantees the construction of the p_n continues even for nonstandard n.)

So there is a second-order axiom expressing "C exists and every set is in $\mathcal{D}(C)$".

Let $D = \mathrm{ACA} +$ "every set is in $\mathcal{D}$" and $U = \mathrm{ACA} +$ "C exists and every set is in $\mathcal{D}(C)$".

Theorem (Freire-W., independently Enayat)

The theories D and U are bi-interpretable. Consequently, ACA *is not tight.*

From ACA to Π^1_k-CA

Abstractly, the strategy to prove the non-tightness of ACA was this:

- There is a minimum model of ACA (the arithmetically definable sets).
- There is a second-order axiom to characterize this minimum model.
- We can define a canonical Cohen generic over this minimum model, and thereby get a definable choice for an extension of the minimum model.
- The minimum model and its canonical extension are bi-interpretable.
- The construction machinery for the bi-interpretation works even over ω-nonstandard models.

From ACA to Π^1_k-CA

Abstractly, the strategy to prove the non-tightness of ACA was this:

- There is a minimum model of ACA (the arithmetically definable sets).
- There is a second-order axiom to characterize this minimum model.
- We can define a canonical Cohen generic over this minimum model, and thereby get a definable choice for an extension of the minimum model.
- The minimum model and its canonical extension are bi-interpretable.
- The construction machinery for the bi-interpretation works even over ω-nonstandard models.

To prove the non-tightness of Π^1_k-CA we will adopt the same strategy.

Second-order arithmetic is set theory in disguise

Strong subsystems of Z_2 are bi-interpretable with fragments of ZFC^-+ "every set is countable". (The minus in ZFC^- means *minus Powerset*).

- $Z_2 +$ the AC schema is bi-interpretable with ZFC^-+ "every set is countable".
- For Z_2 alone, drop Collection from the set theory side.
- For Π^1_k-CA_0, $k \geq 2$, restrict Separation to Π^1_{k-1} formulae.

The set theory $\to$ arithmetic direction is simple—restrict to subsets of ω. The arithmetic $\to$ set theory direction is based on the idea, going back to Aczel and Scott, of coding sets as trees and constructing an appropriate membership relation between trees. A key observation, due to Simpson, is that ATR_0 suffices to carry out this interpretation.

Minimum β-models of arithmetic

An ω-model of arithmetic is a β-model if it is correct about which of its relations are well-founded.

- (Harrison 1968) The hyperarithmetic sets do not form a β-model.

Any β-model of arithmetic is bi-interpretable with a transitive model of set theory. (Any β-model automatically satisfies ATR and so is strong enough to carry out the sets as trees construction.)

Minimum β-models of arithmetic

An ω-model of arithmetic is a β-model if it is correct about which of its relations are well-founded.

- (Harrison 1968) The hyperarithmetic sets do not form a β-model.

Any β-model of arithmetic is bi-interpretable with a transitive model of set theory. (Any β-model automatically satisfies ATR and so is strong enough to carry out the sets as trees construction.)

- (Set theoretical fact) Levels of Gödel's constructible universe L give minimum transitive models of set theories.
- Important point! L has a definable global well-order, so we can use it to make canonical choices.

This translates over to arithmetic to give minimum β-models of subsystems of second-order arithmetic.

Minimum β-models of arithmetic

The following theories have minimum β-models arising as the set of reals in a level L_α of the constructible universe.

- For Π^1_1-CA: the supremum of the first ω many admissible ordinals.
- For Π^1_k-CA, $k \geq 2$: the least ordinal α so that $L_\alpha \models \Pi_{k-1}$-Comprehension.
- For Z_2: the ordinal of ramified analysis β_0—the least ordinal so that $L_{\beta_0} \models \mathsf{ZFC}^-$

Moreover, these minimum β-models are bi-interpretable with their level of L.

Minimum β-models of arithmetic

The following theories have minimum β-models arising as the set of reals in a level L_α of the constructible universe.

- For Π^1_1-CA: the supremum of the first ω many admissible ordinals.
- For Π^1_k-CA, $k \geq 2$: the least ordinal α so that $\mathrm{L}_\alpha \models \Pi_{k-1}$-Comprehension.
- For Z_2: the ordinal of ramified analysis β_0—the least ordinal so that $\mathrm{L}_{\beta_0} \models \mathsf{ZFC}^-$

Moreover, these minimum β-models are bi-interpretable with their level of L.

These ordinals increase as the strength of the theory increases.

Minimum β-models of arithmetic

The following theories have minimum β-models arising as the set of reals in a level L_α of the constructible universe.

- For Π^1_1-CA: the supremum of the first ω many admissible ordinals.
- For Π^1_k-CA, $k \geq 2$: the least ordinal α so that $\mathrm{L}_\alpha \models \Pi_{k-1}$-Comprehension.
- For Z_2: the ordinal of ramified analysis β_0—the least ordinal so that $\mathrm{L}_{\beta_0} \models \mathsf{ZFC}^-$

Moreover, these minimum β-models are bi-interpretable with their level of L.

These ordinals increase as the strength of the theory increases.

In particular, if L_α gives the minimum β-model of Π^1_k-CA then L_α will not satisfy the full Replacement schema.

So there is an increasing cofinal map $f : \omega \to \alpha$ definable over L_α. (Because L_α thinks every set is countable, any failure of Replacement can be ported to have domain ω.)

Getting a code for the minimum β-model

For the next few slides, fix $k \geq 1$ and let $\mathcal{B} = \mathcal{P}(\omega) \cap \mathrm{L}_\alpha$ be the minimum β-model of Π^1_k-CA. Fix a definable increasing cofinal map $f : \omega \to \alpha$.

Getting a code for the minimum β-model

For the next few slides, fix $k \geq 1$ and let $\mathcal{B} = \mathcal{P}(\omega) \cap \mathrm{L}_\alpha$ be the minimum β-model of Π^1_k-CA. Fix a definable increasing cofinal map $f : \omega \to \alpha$.

- For each n, L_α sees a bijection $\omega \to \mathrm{L}_{f(n)}$. Pick the L-least, call it b_n.
- Define $\mathrm{T}_\mathcal{B} \subseteq \omega^3$ to consist of the triples (n, i, x) so that $x \in b_n(i)$. We can think of $\mathrm{T}_\mathcal{B}$ as a subset of ω.

Getting a code for the minimum β-model

For the next few slides, fix $k \geq 1$ and let $\mathcal{B} = \mathcal{P}(\omega) \cap \mathrm{L}_\alpha$ be the minimum β-model of Π^1_k-CA. Fix a definable increasing cofinal map $f : \omega \to \alpha$.

- For each n, L_α sees a bijection $\omega \to \mathrm{L}_{f(n)}$. Pick the L-least, call it b_n.
- Define $\mathrm{T}_\mathcal{B} \subseteq \omega^3$ to consist of the triples (n, i, x) so that $x \in b_n(i)$. We can think of $\mathrm{T}_\mathcal{B}$ as a subset of ω.
- Every element of $\mathcal{B}$ is some slice of $\mathrm{T}_\mathcal{B}$.

$\mathrm{T}_\mathcal{B}$ is definable over L_α, since I just defined it. Note that I used the global well-order of L to make choices for the definition.

Getting a code for the minimum β-model

For the next few slides, fix $k \geq 1$ and let $\mathcal{B} = \mathcal{P}(\omega) \cap \mathrm{L}_\alpha$ be the minimum β-model of Π^1_k-CA. Fix a definable increasing cofinal map $f : \omega \to \alpha$.

- For each n, L_α sees a bijection $\omega \to \mathrm{L}_{f(n)}$. Pick the L-least, call it b_n.
- Define $\mathrm{T}_\mathcal{B} \subseteq \omega^3$ to consist of the triples (n, i, x) so that $x \in b_n(i)$. We can think of $\mathrm{T}_\mathcal{B}$ as a subset of ω.
- Every element of $\mathcal{B}$ is some slice of $\mathrm{T}_\mathcal{B}$.

$\mathrm{T}_\mathcal{B}$ is definable over L_α, since I just defined it. Note that I used the global well-order of L to make choices for the definition.

Claim: $\mathrm{T}_\mathcal{B}$ is second-order definable over $\mathcal{B}$.

Getting a code for the minimum β-model

For the next few slides, fix $k \geq 1$ and let $\mathcal{B} = \mathcal{P}(\omega) \cap \mathrm{L}_\alpha$ be the minimum β-model of Π^1_k-CA. Fix a definable increasing cofinal map $f : \omega \to \alpha$.

- For each n, L_α sees a bijection $\omega \to \mathrm{L}_{f(n)}$. Pick the L-least, call it b_n.
- Define $\mathrm{T}_\mathcal{B} \subseteq \omega^3$ to consist of the triples (n, i, x) so that $x \in b_n(i)$. We can think of $\mathrm{T}_\mathcal{B}$ as a subset of ω.
- Every element of $\mathcal{B}$ is some slice of $\mathrm{T}_\mathcal{B}$.

$\mathrm{T}_\mathcal{B}$ is definable over L_α, since I just defined it. Note that I used the global well-order of L to make choices for the definition.

Claim: $\mathrm{T}_\mathcal{B}$ is second-order definable over $\mathcal{B}$.

At root, this is because $\mathrm{T}_\mathcal{B}$ is definable over L_α. While $\mathcal{B}$ doesn't have direct access to every set in L_α it is bi-interpretable with L_α. It has trees coding each set in L_α, so it can mimic definitions over L_α by quantifying over these trees.

Absoluteness of the code $\mathrm{T}_\mathcal{B}$

Say that a β-model $\mathcal{Y}$ is an outer model of $\mathcal{B}$ if $\mathcal{B} \subseteq \mathcal{Y}$ and $\mathcal{Y}$ doesn't have any new ordertypes for a well-order. More precisely, if $\Gamma \in \mathcal{Y}$ is a well-order then $\mathcal{Y}$ sees an isomorphism of Γ to some $\Gamma' \in \mathcal{B}$.

- Outer models of $\mathcal{B}$ are bi-interpretable with outer models of L_α—transitive models of set theory with the same ordinals.

Absoluteness of the code $\mathrm{T}_{\mathcal{B}}$

Say that a β-model $\mathcal{Y}$ is an outer model of $\mathcal{B}$ if $\mathcal{B} \subseteq \mathcal{Y}$ and $\mathcal{Y}$ doesn't have any new ordertypes for a well-order. More precisely, if $\Gamma \in \mathcal{Y}$ is a well-order then $\mathcal{Y}$ sees an isomorphism of Γ to some $\Gamma' \in \mathcal{B}$.

- Outer models of $\mathcal{B}$ are bi-interpretable with outer models of L_α—transitive models of set theory with the same ordinals.

Claim: $\mathrm{T}_{\mathcal{B}}$ is uniformly definable across all outer models of $\mathcal{B}$.

Absoluteness of the code $\mathrm{T}_{\mathcal{B}}$

Say that a β-model $\mathcal{Y}$ is an outer model of $\mathcal{B}$ if $\mathcal{B} \subseteq \mathcal{Y}$ and $\mathcal{Y}$ doesn't have any new ordertypes for a well-order. More precisely, if $\Gamma \in \mathcal{Y}$ is a well-order then $\mathcal{Y}$ sees an isomorphism of Γ to some $\Gamma' \in \mathcal{B}$.

- Outer models of $\mathcal{B}$ are bi-interpretable with outer models of L_α—transitive models of set theory with the same ordinals.

Claim: $\mathrm{T}_{\mathcal{B}}$ is uniformly definable across all outer models of $\mathcal{B}$.

By the absoluteness of L. Relativize the definition of $\mathrm{T}_{\mathcal{B}}$ to L and then all outer models of L_α will define it the same.

Absoluteness of the code $\mathrm{T}_{\mathcal{B}}$

Say that a β-model $\mathcal{Y}$ is an outer model of $\mathcal{B}$ if $\mathcal{B} \subseteq \mathcal{Y}$ and $\mathcal{Y}$ doesn't have any new ordertypes for a well-order. More precisely, if $\Gamma \in \mathcal{Y}$ is a well-order then $\mathcal{Y}$ sees an isomorphism of Γ to some $\Gamma' \in \mathcal{B}$.

- Outer models of $\mathcal{B}$ are bi-interpretable with outer models of L_α—transitive models of set theory with the same ordinals.

Claim: $\mathrm{T}_{\mathcal{B}}$ is uniformly definable across all outer models of $\mathcal{B}$.

By the absoluteness of L. Relativize the definition of $\mathrm{T}_{\mathcal{B}}$ to L and then all outer models of L_α will define it the same.

In particular, there's a definition of $\mathrm{T}_{\mathcal{B}}$ absolute between $\mathcal{B}$ and its Cohen extensions.

(Because forcing extensions of a model of KP + Mostowski's collapse lemma cannot add new ordinals, Cohen extensions are outer models.)

Defining a Cohen extension of $\mathcal{B}$

Claim: There is a definition for a Cohen generic C over $\mathcal{B}$ which is absolute between outer models of $\mathcal{B}$.

Defining a Cohen extension of $\mathcal{B}$

Claim: There is a definition for a Cohen generic C over $\mathcal{B}$ which is absolute between outer models of $\mathcal{B}$.

- It's the same construction as before.
- From $\mathrm{T}_{\mathcal{B}}$ we can define an enumeration of the dense sets in $\mathcal{B}$.
- We define a sequence $\langle p_n \rangle$ of stronger and stronger conditions, at each stage choosing the least condition which gets in the next dense set.
- Finally set $C = \{q : q \geq p_n \text{ for some } n\}$.

Defining a Cohen extension of $\mathcal{B}$

Claim: There is a definition for a Cohen generic C over $\mathcal{B}$ which is absolute between outer models of $\mathcal{B}$.

- It's the same construction as before.
- From $\mathrm{T}_{\mathcal{B}}$ we can define an enumeration of the dense sets in $\mathcal{B}$.
- We define a sequence $\langle p_n \rangle$ of stronger and stronger conditions, at each stage choosing the least condition which gets in the next dense set.
- Finally set $C = \{q : q \geq p_n \text{ for some } n\}$.

This definition is absolute between outer models of $\mathcal{B}$ because we have an absolute definition for $\mathrm{T}_{\mathcal{B}}$.

Defining a Cohen extension of $\mathcal{B}$

Claim: There is a definition for a Cohen generic C over $\mathcal{B}$ which is absolute between outer models of $\mathcal{B}$.

- It's the same construction as before.
- From $\mathrm{T}_\mathcal{B}$ we can define an enumeration of the dense sets in $\mathcal{B}$.
- We define a sequence $\langle p_n \rangle$ of stronger and stronger conditions, at each stage choosing the least condition which gets in the next dense set.
- Finally set $C = \{q : q \geq p_n \text{ for some } n\}$.

This definition is absolute between outer models of $\mathcal{B}$ because we have an absolute definition for $\mathrm{T}_\mathcal{B}$.

- Using C we can define a code $\mathrm{T}_\mathcal{B}$ for $\mathcal{B}[C]$: this works similar to the definition of $\mathrm{T}_\mathcal{B}$, except instead of directly at a level $\mathrm{L}_{f(n)}$ we look for conditions $p \in C$ which force behavior about $\mathrm{L}_{f(n)}[C]$.

Semantic non-tightness of Π^1_k-CA

Let $\mathcal{B}[C]$ denote the Cohen extension by C defined as on the previous slide.

Theorem (Freire-W.)

The ω-models $(\omega, \mathcal{B})$ and $(\omega, \mathcal{B}[C])$ of Π^1_k-CA satisfy different theories but are bi-interpretable.

Semantic non-tightness of Π^1_k-CA

Let $\mathcal{B}[C]$ denote the Cohen extension by C defined as on the previous slide.

Theorem (Freire-W.)

The ω-models $(\omega, \mathcal{B})$ and $(\omega, \mathcal{B}[C])$ of Π^1_k-CA satisfy different theories but are bi-interpretable.

- $\mathcal{B}$ satisfies "every set is in $\mathcal{B}$", which is expressible using $\mathrm{T}_\mathcal{B}$, whereas $\mathcal{B}[C]$ does not satisfy this axiom.
- $\mathcal{L}[C] \models \Pi^1_k$-CA because this is preserved by forcing.
- Interpreting $\mathcal{B}$ in $\mathcal{B}[C]$ is just restricting the domain of the sets. In the other direction, you can represent sets by their index in $\mathrm{T}_\mathcal{B}(C)$.

From semantic non-tightness to non-tightness for Π^1_k-CA

We follow the ACA strategy, doing the same construction, but in a formal setting rather than working over a specific model.

- Again, we can write down an axiom expressing "I am the minimum model".
- This comes from a (possibly ill-founded!) level of the constructible universe.
- A canonical Cohen generic can be defined, and our two theories will include the assertions "I am the minimum model" and "I am the canonical Cohen extension of the minimum model".
- Full induction is essential to ensure constructions go all the way through.

From semantic non-tightness to non-tightness for Π^1_k-CA

We follow the ACA strategy, doing the same construction, but in a formal setting rather than working over a specific model.

- Again, we can write down an axiom expressing "I am the minimum model".
- This comes from a (possibly ill-founded!) level of the constructible universe.
- A canonical Cohen generic can be defined, and our two theories will include the assertions "I am the minimum model" and "I am the canonical Cohen extension of the minimum model".
- Full induction is essential to ensure constructions go all the way through.

Most of this is straightforward, and is just like the ACA case, but there's one sticking point.

From semantic non-tightness to non-tightness for Π^1_k-CA

We follow the ACA strategy, doing the same construction, but in a formal setting rather than working over a specific model.

- Again, we can write down an axiom expressing "I am the minimum model".
- This comes from a (possibly ill-founded!) level of the constructible universe.
- A canonical Cohen generic can be defined, and our two theories will include the assertions "I am the minimum model" and "I am the canonical Cohen extension of the minimum model".
- Full induction is essential to ensure constructions go all the way through.

Most of this is straightforward, and is just like the ACA case, but there's one sticking point.

From semantic non-tightness to non-tightness for Π^1_k-CA

We follow the ACA strategy, doing the same construction, but in a formal setting rather than working over a specific model.

- Again, we can write down an axiom expressing "I am the minimum model".
- This comes from a (possibly ill-founded!) level of the constructible universe.
- A canonical Cohen generic can be defined, and our two theories will include the assertions "I am the minimum model" and "I am the canonical Cohen extension of the minimum model".
- Full induction is essential to ensure constructions go all the way through.

Most of this is straightforward, and is just like the ACA case, but there's one sticking point.

To define $\mathrm{T}_\mathcal{B}$, I used that L_α didn't satisfy Replacement, and so there was some definable cofinal map $f : \omega \to \alpha$. That's not good enough now. We need an explicit construction, one which works uniformly.

A warm-up: Π^1_1-CA

The minimum β-model of Π^1_1-CA consists of the reals which appear in $\mathrm{L}_{\omega_\omega^{\mathrm{CK}}}$.

A warm-up: Π^1_1-CA

The minimum β-model of Π^1_1-CA consists of the reals which appear in $\mathrm{L}_{\omega_\omega^{\mathrm{CK}}}$.

Axiomatize B_1 as:

- Π^1_1-CA;
- Every set is constructible;
- There are ω many admissible ordinals.

A warm-up: Π^1_1-CA

The minimum β-model of Π^1_1-CA consists of the reals which appear in $\mathrm{L}_{\omega_\omega^{\mathrm{CK}}}$.

Axiomatize B_1 as:

- Π^1_1-CA;
- Every set is constructible;
- There are ω many admissible ordinals.

The only β-model which satisfies B_1 is the minimum β-model of Π^1_1-CA.

A warm-up: Π^1_1-CA

The minimum β-model of Π^1_1-CA consists of the reals which appear in $L_{\omega_\omega^{CK}}$.

Axiomatize B_1 as:

- Π^1_1-CA;
- Every set is constructible;
- There are ω many admissible ordinals.

The only β-model which satisfies B_1 is the minimum β-model of Π^1_1-CA.

Among non-β-models there isn't a minimum model of B_1. But every model of Π^1_1-CA has a minimum β-submodel (= submodel which agrees about which relations are well-founded), which is a model of B_1.

In particular this happens if $(M, \mathcal{B}[C])$ is an extension of $(M, \mathcal{B}) \models B_1$ by Cohen forcing.

A warm-up: Π^1_1-CA

What remains is to define a code $\mathrm{T}_{\mathcal{B}}$ which over Π^1_1-CA gives the minimum β-submodel of Π^1_1-CA.

- The set of n for which there are at least n many admissible ordinals is inductive, so by full induction must contain all n.
- Define $\mathrm{T}_{\mathcal{B}}$ to consist of triples (n, i, x) so that x is in the i-th set (according to the L-least enumeration) of $\mathrm{L}_{\omega_n^{\mathrm{CK}}}$.

A warm-up: Π^1_1-CA

What remains is to define a code $\mathrm{T}_{\mathcal{B}}$ which over Π^1_1-CA gives the minimum β-submodel of Π^1_1-CA.

- The set of n for which there are at least n many admissible ordinals is inductive, so by full induction must contain all n.
- Define $\mathrm{T}_{\mathcal{B}}$ to consist of triples (n, i, x) so that x is in the i-th set (according to the L-least enumeration) of $\mathrm{L}_{\omega_n^{\mathrm{CK}}}$.

Fact: If $(M, \mathcal{X})$ is a β-submodel of $(M, \mathcal{Y})$, then they define $\mathrm{T}_{\mathcal{B}}$ the same.

A warm-up: Π^1_1-CA

What remains is to define a code $\mathrm{T}_{\mathcal{B}}$ which over Π^1_1-CA gives the minimum β-submodel of Π^1_1-CA.

- The set of n for which there are at least n many admissible ordinals is inductive, so by full induction must contain all n.
- Define $\mathrm{T}_{\mathcal{B}}$ to consist of triples (n, i, x) so that x is in the i-th set (according to the L-least enumeration) of $\mathrm{L}_{\omega_n^{\mathrm{CK}}}$.

Fact: If $(M, \mathcal{X})$ is a β-submodel of $(M, \mathcal{Y})$, then they define $\mathrm{T}_{\mathcal{B}}$ the same.

Using the code $\mathrm{T}_{\mathcal{B}}$ we can canonically define a Cohen generic C over the minimum β-submodel of Π^1_1-CA, and we can define a code $\mathrm{T}_{\mathcal{B}}(C)$ for the extension by C. (Again, full induction is used to define C.)

A warm-up: Π^1_1-CA

What remains is to define a code $\mathrm{T}_\mathcal{B}$ which over Π^1_1-CA gives the minimum β-submodel of Π^1_1-CA.

- The set of n for which there are at least n many admissible ordinals is inductive, so by full induction must contain all n.
- Define $\mathrm{T}_\mathcal{B}$ to consist of triples (n, i, x) so that x is in the i-th set (according to the L-least enumeration) of $\mathrm{L}_{\omega_n^{\mathrm{CK}}}$.

Fact: If $(M, \mathcal{X})$ is a β-submodel of $(M, \mathcal{Y})$, then they define $\mathrm{T}_\mathcal{B}$ the same.

Using the code $\mathrm{T}_\mathcal{B}$ we can canonically define a Cohen generic C over the minimum β-submodel of Π^1_1-CA, and we can define a code $\mathrm{T}_\mathcal{B}(C)$ for the extension by C. (Again, full induction is used to define C.)

Let U_1 be Π^1_1-CA + "every set is a slice of $\mathrm{T}_\mathcal{B}(C)$".

Theorem (Freire-W.)

B_1 and U_1 are bi-interpretable. Hence, Π^1_1-CA is not tight.

Going from Π^1_1-CA to Π^1_k-CA

—But first I have to talk about fine structure theory!

- In studying combinatorics in L, Jensen needed a fine understanding of how Σ_ℓ elementarity behaves in L. For this he invented fine structure theory.

—But first I have to talk about fine structure theory!

- In studying combinatorics in L, Jensen needed a fine understanding of how Σ_ℓ elementarity behaves in L. For this he invented fine structure theory.
- Σ_ℓ uniformization theorem (Jensen): If α is a limit ordinal then there is a Σ_ℓ-definable over L_α Skolem function for Σ_ℓ properties.

—But first I have to talk about fine structure theory!

- In studying combinatorics in L, Jensen needed a fine understanding of how Σ_ℓ elementarity behaves in L. For this he invented fine structure theory.
- Σ_ℓ uniformization theorem (Jensen): If α is a limit ordinal then there is a Σ_ℓ-definable over L_α Skolem function for Σ_ℓ properties.
- Using this you can give a fine structural characterization of the minimum β-models: The minimum β-model of Π^1_k-CA, $k \geq 2$, is the reals in L_α where α is the smallest ordinal whose Σ_{k-1} projectum is bigger than ω.

—But first I have to talk about fine structure theory!

- In studying combinatorics in L, Jensen needed a fine understanding of how Σ_ℓ elementarity behaves in L. For this he invented fine structure theory.
- Σ_ℓ uniformization theorem (Jensen): If α is a limit ordinal then there is a Σ_ℓ-definable over L_α Skolem function for Σ_ℓ properties.
- Using this you can give a fine structural characterization of the minimum β-models: The minimum β-model of Π^1_k-CA, $k \geq 2$, is the reals in L_α where α is the smallest ordinal whose Σ_{k-1} projectum is bigger than ω.

Lemma: Let L_α be the first level of L which satisfies Σ_ℓ-Replacement. Then over L_α can be defined a set $\mathrm{T}_\mathcal{B} \subseteq \omega^3$ which codes $\mathcal{P}(\omega) \cap \mathrm{L}_\alpha$.

—But first I have to talk about fine structure theory!

- In studying combinatorics in L, Jensen needed a fine understanding of how Σ_ℓ elementarity behaves in L. For this he invented fine structure theory.
- Σ_ℓ uniformization theorem (Jensen): If α is a limit ordinal then there is a Σ_ℓ-definable over L_α Skolem function for Σ_ℓ properties.
- Using this you can give a fine structural characterization of the minimum β-models: The minimum β-model of Π^1_k-CA, $k \geq 2$, is the reals in L_α where α is the smallest ordinal whose Σ_{k-1} projectum is bigger than ω.

Lemma: Let L_α be the first level of L which satisfies Σ_ℓ-Replacement. Then over L_α can be defined a set $\mathrm{T}_\mathcal{B} \subseteq \omega^3$ which codes $\mathcal{P}(\omega) \cap \mathrm{L}_\alpha$.

- It is enough to define a sequence $\langle \alpha_n : n < \omega \rangle$ cofinal in α. From such a sequence we can put $(n, i, x) \in \mathrm{T}_\mathcal{B}$ if x is in the i-th real of L_{α_n}, according to the L-least enumeration.

—But first I have to talk about fine structure theory!

- In studying combinatorics in L, Jensen needed a fine understanding of how Σ_ℓ elementarity behaves in L. For this he invented fine structure theory.
- Σ_ℓ uniformization theorem (Jensen): If α is a limit ordinal then there is a Σ_ℓ-definable over L_α Skolem function for Σ_ℓ properties.
- Using this you can give a fine structural characterization of the minimum β-models: The minimum β-model of Π^1_k-CA, $k \geq 2$, is the reals in L_α where α is the smallest ordinal whose Σ_{k-1} projectum is bigger than ω.

Lemma: Let L_α be the first level of L which satisfies Σ_ℓ-Replacement. Then over L_α can be defined a set $\mathrm{T}_\mathcal{B} \subseteq \omega^3$ which codes $\mathcal{P}(\omega) \cap \mathrm{L}_\alpha$.

- It is enough to define a sequence $\langle \alpha_n : n < \omega \rangle$ cofinal in α. From such a sequence we can put $(n, i, x) \in \mathrm{T}_\mathcal{B}$ if x is in the i-th real of L_{α_n}, according to the L-least enumeration.
- Given α_n set α_{n+1} to be the least ordinal so that $\mathrm{L}_{\alpha_{n+1}}$ is closed under the Σ_ℓ Skolem function with inputs from L_{α_n}.
- This sequence must be cofinal by leastness of α, as $\mathrm{L}_{\sup_n \alpha_n}$ satisfies Σ_ℓ-Replacement.

Arithmetization

All this can be formalized.

Let ZFC^-_ℓ be the theory axiomatized by $\mathsf{KP} + \Sigma_\ell$-Replacement $+ \Pi_\ell$-Separation $+ \Pi_\ell$-Foundation.

Arithmetization

All this can be formalized.

Let ZFC^-_ℓ be the theory axiomatized by KP $+$ Σ_ℓ-Replacement $+$ Π_ℓ-Separation $+$ Π_ℓ-Foundation.

- Formal Σ_ℓ uniformization: The theory $\mathsf{ZFC}^-_\ell + \mathrm{V} = \mathrm{L}$ proves the existence of a Σ_ℓ Skolem function.

Arithmetization

All this can be formalized.

Let ZFC^-_ℓ be the theory axiomatized by $\mathsf{KP} + \Sigma_\ell$-Replacement $+ \Pi_\ell$-Separation $+ \Pi_\ell$-Foundation.

- Formal Σ_ℓ uniformization: The theory $\mathsf{ZFC}^-_\ell + \mathrm{V} = \mathrm{L}$ proves the existence of a Σ_ℓ Skolem function.

Axiomatize the theory B_k, $k \geq 2$ by

- Π^1_k-CA;
- Every set is constructible;
- There is no ordinal ξ so that $\mathrm{L}_\xi \models \Pi^1_k$-CA.

Arithmetization

All this can be formalized.

Let ZFC^-_ℓ be the theory axiomatized by $\mathsf{KP} + \Sigma_\ell$-Replacement + Π_ℓ-Separation + Π_ℓ-Foundation.

- Formal Σ_ℓ uniformization: The theory $\mathsf{ZFC}^-_\ell + \mathrm{V} = \mathrm{L}$ proves the existence of a Σ_ℓ Skolem function.

Axiomatize the theory B_k, $k \geq 2$ by

- Π^1_k-CA;
- Every set is constructible;
- There is no ordinal ξ so that $\mathrm{L}_\xi \models \Pi^1_k$-CA.

Fact: ("probably well-known" –Simpson): Because we include the statement that every set is constructible, we get for free the Σ^1_k-AC schema.

Fact: The set theory bi-interpretable with B_k contains $\mathsf{ZFC}^-_{k-1} + \mathrm{V} = \mathrm{L}$ + "every set is countable".

Arithmetization

All this can be formalized.

Let ZFC^-_ℓ be the theory axiomatized by KP + Σ_ℓ-Replacement + Π_ℓ-Separation + Π_ℓ-Foundation.

- Formal Σ_ℓ uniformization: The theory $\mathsf{ZFC}^-_\ell + \mathrm{V} = \mathrm{L}$ proves the existence of a Σ_ℓ Skolem function.
- Formal definition of $\mathrm{T}_{\mathcal{B}}$: Same as before, but work in B_k. We use full induction to know the sequence goes all the way through the model's ω.

Axiomatize the theory B_k, $k \geq 2$ by

- Π^1_k-CA;
- Every set is constructible;
- There is no ordinal ξ so that $\mathrm{L}_\xi \models \Pi^1_k$-CA.

Fact: ("probably well-known" –Simpson): Because we include the statement that every set is constructible, we get for free the Σ^1_k-AC schema.

Fact: The set theory bi-interpretable with B_k contains $\mathsf{ZFC}^-_{k-1} + \mathrm{V} = \mathrm{L}$ + "every set is countable".

Arithmetization

All this can be formalized.

Let ZFC^-_ℓ be the theory axiomatized by $\mathsf{KP} + \Sigma_\ell$-Replacement $+$ Π_ℓ-Separation $+$ Π_ℓ-Foundation.

- Formal Σ_ℓ uniformization: The theory $\mathsf{ZFC}^-_\ell + \mathrm{V} = \mathrm{L}$ proves the existence of a Σ_ℓ Skolem function.
- Formal definition of $\mathrm{T}_{\mathcal{B}}$: Same as before, but work in B_k. We use full induction to know the sequence goes all the way through the model's ω.
- Relativizing the definition to L we get a definition that's uniform between a model of B_k and its forcing extensions.

Axiomatize the theory B_k, $k \geq 2$ by

- Π^1_k-CA;
- Every set is constructible;
- There is no ordinal ξ so that $\mathrm{L}_\xi \models \Pi^1_k$-CA.

Fact: ("probably well-known" –Simpson): Because we include the statement that every set is constructible, we get for free the Σ^1_k-AC schema.

Fact: The set theory bi-interpretable with B_k contains $\mathsf{ZFC}^-_{k-1} + \mathrm{V} = \mathrm{L} +$ "every set is countable".

Non-tightness of Π^1_k-CA

Now that we have a definition for a code $\mathrm{T}_{\mathcal{B}}$ for the "minimum model" of B_k, we can do the same construction.

Non-tightness of Π^1_k-CA

Now that we have a definition for a code $\mathrm{T}_{\mathcal{B}}$ for the "minimum model" of B_k, we can do the same construction.

- Define a canonical Cohen generic C over the minimum model.
- Then $\mathrm{T}_{\mathcal{B}}(C)$, a code for the extension by C, is definable over the minimum model.
- All this is absolute between any model of B_k and its forcing extensions.

Non-tightness of Π^1_k-CA

Now that we have a definition for a code $\mathrm{T}_\mathcal{B}$ for the "minimum model" of B_k, we can do the same construction.

- Define a canonical Cohen generic C over the minimum model.
- Then $\mathrm{T}_\mathcal{B}(C)$, a code for the extension by C, is definable over the minimum model.
- All this is absolute between any model of B_k and its forcing extensions.

Let U_k be Π^1_k-CA + "every set is a slice of $\mathrm{T}_\mathcal{B}(C)$".

Theorem (Freire-W.)

B_k and U_k are bi-interpretable. Hence, Π^1_k-CA is not tight.

Related results

Alfredo and I were originally interested in the case of class theory, and only realized our constructions could be ported to arithmetic after the fact.

Theorem (Freire-W.)

The theories GB *and* $\mathsf{GB} + \Pi^1_k\text{-CA}$ *are not tight.*

Related results

Alfredo and I were originally interested in the case of class theory, and only realized our constructions could be ported to arithmetic after the fact.

Theorem (Freire-W.)

The theories GB *and* $\mathrm{GB} + \Pi^1_k$-CA *are not tight.*

Independently to us, Ali Enayat has been working on closely related questions.

Theorem (Enayat)

No finitely axiomatized subtheory of PA, ZF, Z_2, *or* KM *is tight.*

Thank you!