Provability logic: models within models in Peano Arithmetic

Alessandro Berarducci and Marcello Mamino

CUNY Graduate Center, NY
MOPA Seminars, 24 October 2023
(online)

The problem

We were inspired by Jech's proof of Gödel's second incompleteness theorem for ZF.

Jech (1994) showed, by model theoretic means, that ZF does not prove that ZF has a model.

Can we do something similar for PA?
We want to avoid the syntactic notion of proof and the completeness theorem.

Difficulty: within PA we cannot speak about models of PA.

The idea

Work with models of bounded complexity.

Theorem: the assertion that PA does not have a Σ^0_2-model is independent of PA.

What is a Σ^0_2-model?

First attempt

A Σ_2^0-model is a model whose underlying set is $\mathbb{N}$ and such that addition and multiplication are defined by Σ_2^0 predicates.

It does not work. The standard model has these properties, but its truth predicate is non-arithmetic. We need additional constrains to be able to quantify over models of PA within PA.

Second attempt

A Σ^0_2-model is a maximal consistent theory T extending PA such that the set $\{\ulcorner\phi\urcorner : \phi \in T\} \subseteq \mathbb{N}$ is Σ^0_2-definable.

Kotlarski (2019) had already done this.

We need something different if we want to avoid the syntactic notion of consistency.

Final choice: use formulas with parameters

Let M be a model of PA with domain $\mathbb{N}$.

M is a Σ^0_2-model if the set of pairs $(\ulcorner\varphi\urcorner, s) \in \mathbb{N} \times \mathbb{N}$ such that

$$M \models \varphi[s]$$

can be defined by a formula $\Theta(x, y)$ of complexity Σ^0_2.
($s \in \mathbb{N}$ codes the list of parameters).

Thus $$M \models \varphi[s] \iff \mathbb{N} \models \Theta(\ulcorner\varphi\urcorner, s)$$

If $U(x, y, z) \in \Sigma^0_2$ is a universal Σ^0_2-predicate and $m = \ulcorner\Theta\urcorner$,

$$M \models \varphi[s] \iff \mathbb{N} \models U(m, \ulcorner\varphi\urcorner, s)$$

We call $m \in \mathbb{N}$ the code of the model M.

Next level: non-standard non-standard [sic] models

So far we have

$$M \models \varphi[s] \iff \mathbb{N} \models U(m, \ulcorner\varphi\urcorner, s)$$

Now we want to allow m, φ, s to be non-standard.

There is a Π^0_3-formula MODEL(m) saying

"m codes a Σ^0_2-model of PA"

namely MODEL(m) says that $\{(\varphi, s) \mid U(m, \varphi, s)\}$ satisfies Tarski's truth conditions (and contains the axioms of PA).

If MODEL(m) holds, we write "$m \models \varphi[s]$" for $U(m, \varphi, s)$.

The modal operator

In provability logic $\Box\phi$ is usually interpreted as "ϕ is provable in PA".

Now we redefine $\Box\phi$ as the Π^0_4-formula expressing "ϕ holds in all Σ^0_2-models of PA":

$$\Box\phi :\iff \forall m(\mathsf{MODEL}(m) \to \forall s \text{``} m \models \phi[s]\text{''}).$$

Digression: Kripke models

In the Kripke semantic for provability logic, a modal formula is valid if it is true in every well founded transitive Kripke model.

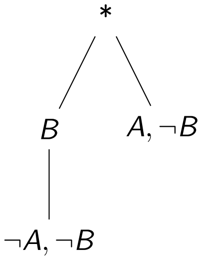

$$* \Vdash \Diamond A \wedge \Box(B \to \Diamond\neg A)$$

It is tempting to interpret the nodes of the Kripke frame as models of PA. Our interpretation of $\Box\phi$ as truth in all Σ^0_2-models is in tune with this intuition.

The derivability conditions

For closed formulas ϕ, ψ we have:

1. $\mathsf{PA} \vdash \phi \implies \mathsf{PA} \vdash \Box\phi$,
2. $\mathsf{PA} \vdash \Box\phi \to \Box\Box\phi$,
3. $\mathsf{PA} \vdash \Box(\phi \to \psi) \to (\Box\phi \to \Box\psi)$.
4. $\mathbb{N} \models \Box\phi \implies \mathsf{PA} \vdash \phi$.

If we had interpreted $\Box\phi$ as "ϕ is provable", then 1 and 2 would follow from the provable Σ^0_1 completeness of PA.

In the new interpretation of $\Box\phi$ as "ϕ is true in all Σ^0_2-models" we must follow a different approach which we discuss below.

A model inside a model

Let $\mathcal{X}$ be a model of PA with domain $X = \mathbb{N}$ (not necessarily a Σ^0_2-model). Given $y \in X$ such that $\mathcal{X} \models \mathrm{MODEL}(y)$, there is a model ${}^{\mathcal{X}}y \models \mathrm{PA}$ such that, for all pairs (ϕ, s) we have

$$ {}^{\mathcal{X}}y \models \phi[s] \iff \mathcal{X} \models \text{“}y \models \phi[s]\text{”}. $$

If $\mathcal{X}$ is a Σ^0_2-model, then so is ${}^{\mathcal{X}}y$ and there is a function $x, y \mapsto {}^{x}y$ of complexity Π^0_3 which computes a code of ${}^{\mathcal{X}}y$ given y and a code x for $\mathcal{X}$. So we have, provably in PA,

$$ {}^{x}y \models \phi[s] \iff x \models \text{“}y \models \phi[s]\text{”} $$

for all x, y satisfying $\mathrm{MODEL}(x)$ and “$x \models \mathrm{MODEL}(y)$”.

Missing details

All the above "can obviously be done", but to actually do it requires some care because we are dealing with three kinds of objects:

- A formula ϕ is a description: it has a different meaning in different models.
- A parameter is not a description: 2 means 2, not $1+1$.
- A formula with parameters consist of a formula and a *description* of a sequence of parameters.

If we are dealing with formulas with a non-standard number of parameters, passing this information from a model to a model M inside the model requires (roughly) three different functions:

- $x \mapsto x$ (parameters, names, rigid designators),
- $x \mapsto \dot{x}$ (numerals, formulas, descriptions),
- $x \mapsto [x, M]$ (sequences of parameters).

In quantified provability logic, one has formulas of the form $\forall x \Box A(\dot{x})$, but here it is more complicated.

Proof of the derivability conditions

1. $\mathsf{PA} \vdash \phi \implies \mathsf{PA} \vdash \Box\phi$,
2. $\mathsf{PA} \vdash \Box\phi \to \Box\Box\phi$,
3. $\mathsf{PA} \vdash \Box(\phi \to \psi) \to (\Box\phi \to \Box\psi)$.
4. $\mathbb{N} \models \Box\phi \implies \mathsf{PA} \vdash \phi$.

For 1 we use ${}^{\mathcal{X}}y \models \psi[s] \iff \mathcal{X} \models$ "$y \models \psi[s]$".

For 2 we use ${}^{x}y \models \psi[s] \iff x \models$ "$y \models \psi[s]$".

3 is obvious.

Proof of 4:
if $\mathsf{PA} \nvdash \phi$, there is a Σ^0_2-model $M \models \mathsf{PA}$ where ϕ fails.
A code m of M witnesses $\mathbb{N} \not\models \Box\phi$.

The existence of a Σ^0_2-model is independent of PA

Let $\mathrm{PA} \vdash G \leftrightarrow \neg\Box G$.

By the derivability conditions and the usual arguments, G is independent of PA and $\mathrm{PA} \vdash G \leftrightarrow \neg\Box\perp$.

Dialog with the editor

Editor: - the referee complained about the proof of condition
4. $\mathbb{N} \models \Box\phi \implies \mathsf{PA} \vdash \phi$.

Authors: - Did they find a mistake?

Editor: - No, but they say that you used the fact that if a recursive theory has a model, then it has a Σ^0_2-model.

Authors: - Isn't this true? Just use König's lemma to find a Δ^0_2-path in the tree of finite consistent extensions with Henkin constants.

Editor: - They say that this proof uses syntactic consistency, so it is not model theoretic.

Authors: - But we only use it in the metatheory! Moreover, 4. is only used for $\mathsf{PA} \not\vdash \Box\bot$, not for $\mathsf{PA} \not\vdash \neg\Box\bot$.

Editor: - Cool down. The referee was kind enough to provide some bibliographical pointers, so maybe you can fix it.

Authors: - But there is nothing to be fixed! Damn it!

Existence of Σ^0_2-models

We want to show that if a recursive theory T has a model, then it has a Σ^0_2-model.

This can be derived proof-theoretically from the usual proof of the completeness theorem based on König's lemma (the left-most branch is Δ^0_2), but we want a model-theoretic proof.

The idea is to construct models as limits of finite structures as in Skolem (1922).

A nice way to formalize this is through Kripke's notion of fulfilment. A variant of this is presented below.

Some complications arise because the theory has infinitely many axioms and we want the elementary diagram of the model, not just the atomic diagram, to be Σ^0_2.

Skolem 1922, Shelah 1982, Putnam 2000, Kripke...

An n-structure is a sequence of relational structures $\bar{M} = (M_0, \ldots, M_n)$ with M_i a substructure of M_{i+1}.

Given $\bar{M}$, $i \leq n$ and ϕ with parameters from M_i, we define:

$$
\begin{array}{lcl}
\bar{M}^i \Vdash \exists x\phi & \iff & i = n \vee \exists a \in M_{i+1}\ ,\ \bar{M}^{i+1} \Vdash \phi(a/x) \\
\bar{M}^i \Vdash \forall x\phi & \iff & \forall j \geq i\ \forall a \in M_j\ ,\ \bar{M}^j \Vdash \phi(a/x) \\
\bar{M}^i \Vdash \phi \wedge \psi & \iff & \bar{M}^i \Vdash \phi \text{ and } \bar{M}^i \Vdash \psi \\
\bar{M}^i \Vdash \phi \vee \psi & \iff & \bar{M}^i \Vdash \phi \text{ or } \bar{M}^i \Vdash \psi \\
\bar{M}^i \Vdash \phi & \iff & M_i \models \phi \quad \text{for } \phi \text{ atomic or negated atomic}
\end{array}
$$

Negations are only applied to atomic formulas.

Similarly we define an ω-structure $\bar{M} = (M_n : n < \omega)$ and put

$$\bar{M}^i \Vdash \exists x\phi \iff \exists a \in M_{i+1}\ ,\ \bar{M}^{i+1} \Vdash \phi(a/x)$$

$\bar{M}$ fulfils ϕ if $\bar{M}^i \Vdash \phi$ where M_i contains the parameters of ϕ.

If $i = 0$, write $\bar{M} \Vdash \phi$.

Properties of fulfillment

φ is satisfiable if and only if it is n-fulfillable for each n. In fact:

- n-fulfillable for each $n \implies \omega$-fulfillable (compactness).
- ω-fullfillable $\implies$ satisfiable (in the union).
- satisfiable $\implies \omega$-fulfillable $\implies n$-fulfillable.

More formally:

$$
\begin{array}{lcl}
(M_i : i < \omega) \Vdash \phi & \implies & (M_0, \ldots, M_n) \Vdash \phi \\
(M_0, \ldots, M_n, M_{n+1}) \Vdash \phi & \implies & (M_0, \ldots, M_n) \Vdash \phi \\
(M_0, M_1, \ldots, M_n) \Vdash \phi & \implies & (M_1, \ldots, M_n) \Vdash \phi \\
(M_i : i < \omega) \Vdash \phi & \implies & \bigcup_n M_n \models \phi \\
\textit{etc.} & &
\end{array}
$$

Note that

$$\bigcup_n M_n \models \phi \;\not\Longrightarrow\; (M_i : i < \omega) \Vdash \phi$$

Bounded n-structures

Let Δ be a finite set of formulas closed under subformulas.

Let $\bar{M} = (M_0, \ldots, M_n)$ be an n-structure.

We say that $\bar{M}$ is Δ-bounded if:

- $|M_0|$ at most equal to the number of closed formulas in Δ beginning with $\exists$.
- $|M_{i+1}|$ is at most $|M_i| + c|M_i|^k$ where c is the number of formulas in Δ beginning with $\exists$ and k is the largest number of free variables in any such formulas.

We say that $\bar{M} = (M_0, \ldots, M_n)$ is a substructure of $\bar{N} = (N_0, \ldots, N_n)$ if M_i is a substructure of N_i for each i.

If $\phi \in \Delta$ is fulfilled in an n-structure, then it is fulfilled in a Δ-bounded n-substructure.

Extending an n-structure

We want to prolong an n-structure, enlarge its domains and expand the language.

Consider an n-structure $\bar{M} = (M_0, \ldots, M_n)$ in a language L and an $n+k$-structure $(N_0, \ldots, N_{n+k})$ in a language $L' \supseteq L$.

$(N_0, \ldots, N_{n+k})$ extends $(M_0, \ldots, M_n)$ if for each $i \leq n$, M_i is a substructure of $N_{i|L}$.

$(M_0, \ldots, M_n)$ is initial if the domain of its last element M_n has the form $\{n \in \mathbb{N} \mid n \leq k\}$ for some $k \in \mathbb{N}$.

Remark: every Δ-bounded n-structure is isomorphic to an initial one (and there are finitely many such).

n-fulfillable for each $n \implies \Sigma^0_2$-Satisfiable

Let T be a recursively axiomatized L-theory.

Let T_n be the conjunction of the first n-axioms of T and let $L_n \subseteq L$ be the language of T_n. Suppose that, for each n, T_n is n-fulfillable.

Let F be the finitely branching tree whose nodes at level n are the T_n-bounded initial n-models $(M_0, \dots, M_n)$ of T_n which extend a node at level $n-1$.

The children of $(M_0, \dots, M_n)$ are the T_{n+1}-bounded initial models $(N_0, \dots, N_n, N_{n+1})$ extending $(M_0, \dots, M_n)$.

Then F is an infinite finitely branching recursive tree, so it has a Δ^0_2 definable infinite branch.

The union of the models in that path is a model of T whose atomic diagram is Δ^0_2.

If T has effective elimination of quantifiers, the elementary diagram is also Δ^0_2. We can reduce to this case by expanding the language.