

Synthetic Undecidability and Incompleteness of First-Order Axiom Systems in Coq

Dominik Kirst^{1*} and Marc Hermes²

¹Saarland Informatics Campus, Saarland University, Germany.

²Department of Mathematics, Saarland University, Germany.

*Corresponding author(s). E-mail(s): kirst@cs.uni-saarland.de;

Abstract

We mechanise the undecidability of various first-order axiom systems in Coq, employing the synthetic approach to computability underlying the growing Coq Library of Undecidability Proofs. Concretely, we cover both semantic and deductive entailment in fragments of Peano arithmetic (PA) as well as ZF and related finitary set theories, with their undecidability established by many-one reductions from solvability of Diophantine equations, i.e. Hilbert's tenth problem (H10), and the Post correspondence problem (PCP), respectively. In the synthetic setting based on the computability of all functions definable in a constructive foundation, such as Coq's type theory, it suffices to define these reductions as meta-level functions with no need for further encoding in a formalised model of computation.

The concrete cases of PA and the considered set theories are supplemented by a general synthetic theory of undecidable axiomatisations, focusing on well-known connections to consistency and incompleteness. Specifically, our reductions rely on the existence of standard models, necessitating additional assumptions in the case of full ZF, and all axiomatic extensions still justified by such standard models are shown incomplete. As a by-product of the undecidability of set theories formulated using only membership and no equality symbol, we obtain the undecidability of first-order logic with a single binary relation.

Keywords: undecidability, synthetic computability, first-order logic, incompleteness, Peano arithmetic, ZF set theory, constructive type theory, Coq

1 Introduction

Being among the mainstream formalisms to underpin mathematics, first-order logic (FOL) has been subject to investigation from many different perspectives since its concretisation in the late 19th century. One of them is concerned with algorithmic properties, prominently pushed by Hilbert and Ackermann with the formulation of the *Entscheidungsproblem* [1], namely the search for a decision procedure determining the formulas φ that are valid in all interpretations, usually written $\models \varphi$. With their groundbreaking work in the 1930s, Turing [2] and Church [3] established that such a general decision procedure cannot exist. However, this outcome can change if one considers validity of φ restricted to interpretations satisfying a given collection \mathcal{A} of axioms, written $\mathcal{A} \models \varphi$. Already in 1929, Presburger presented a decision procedure for an axiomatisation of linear arithmetic [4] and Tarski contributed further instances with his work on Boolean algebras, real-closed ordered fields, and Euclidean geometry in the 1940s [5].

On the other hand, as soon as an axiomatisation \mathcal{A} is strong enough to express computation, the undecidability proof for the Entscheidungsproblem can be replayed within \mathcal{A} , turning its entailed theory undecidable. Used as standard foundations for large branches of mathematics exactly due to their expressiveness, Peano arithmetic (PA) and Zermelo-Fraenkel set theory (ZF) are prime examples of such axiomatisations. In this paper, we use the Coq proof assistant [6] to mechanise the undecidability of PA and ZF and related finitary set theories, based on the synthetic approach to computability results available in Coq's constructive type theory.

As is common in constructive foundations, all functions definable in Coq's type theory are effectively computable. So for instance any Boolean function on natural numbers $f : \mathbb{N} \rightarrow \mathbb{B}$ coinciding with a predicate $P \subseteq \mathbb{N}$ may be understood as a *decider* for P , even without explicitly relating f to some encoding as a Turing machine, μ -recursive function, or untyped λ -term. In this fashion, many positive notions of computability theory can be rendered *synthetically*, disposing of the need for an intermediate formal model of computation [7, 8]. Moreover, negative notions like *undecidability* are mostly established by transport along *reductions*, i.e. computable functions encoding instances of one problem in terms of another problem. Synthetically, the requirement that reductions are computable is again satisfied by construction. In fact, all problems included in the Coq Library of Undecidability Proofs [9] are shown undecidable in the sense that their decidability would entail the decidability of Turing machine halting by synthetic reduction from the latter.

Therefore, revisiting the undecidability of first-order axiom systems using a proof assistant like Coq is worthwhile for several reasons. First, using the synthetic approach to undecidability makes a mechanisation of these fundamental results of metamathematics pleasantly feasible [8, 10]. Our mechanisations follow the informal (and instructive) practice to just define and verify reduction functions while leaving their computability implicit, with the key difference that in our constructive setting this relaxation is formally justified.

Secondly, it is well-known that undecidable axiomatisations \mathcal{A} are negation-incomplete, i.e. admit φ with neither $\mathcal{A} \models \varphi$ nor $\mathcal{A} \models \neg\varphi$. By characterising $\mathcal{A} \models \varphi$ with an enumerable deduction system $\mathcal{A} \vdash \varphi$, this is a consequence of Post's theorem [11] stating that bi-enumerable predicates are decidable. Indeed, assuming negation-completeness, also the complement $\mathcal{A} \not\models \varphi$ would be enumerable via $\mathcal{A} \vdash \neg\varphi$. Based on a synthetic proof of Post's theorem [7, 8], all axiomatisations shown synthetically undecidable in the present paper are incomplete in the sense that their completeness would imply the decidability of the halting problem (for Turing machines). These algorithmic observations complement the otherwise notoriously hard to mechanise incompleteness proofs based on Gödel sentences [12, 13].

Lastly, undecidability of a first-order axiomatisation \mathcal{A} like PA or ZF can only be established in a stronger system, since a reduction from a non-trivial problem yields the consistency of \mathcal{A} . Coq exhibits standard models for PA and ZF (the latter relying on classical assumptions [14]), enabling proofs of their undecidability. In fact, we sharpen the results for fragments Q' and Z' even strictly below Robinson arithmetic Q and Zermelo set theory Z , respectively, with the latter now also admitting a fully constructive standard model.

In summary, the contributions of this paper can be listed as follows:

- We extend the Coq Library of Undecidability Proofs with verified reductions to Q' , Q , PA, Z' , Z , and ZF(-regularity), regarding both Tarski semantics and natural deduction.¹
- We verify a translation of set theory over a convenient signature with function symbols for set operations to smaller signatures just containing one or two binary relation symbols.
- By composition, we obtain the undecidability of the Entscheidungsproblem for a single binary relation, improving on a previous mechanisation with additional symbols [8].
- By isolating a generic theorem, we obtain synthetic undecidability and incompleteness for all axiomatisations extending the fragments Q' and Z' with respect to standard models.

This extended version of [15] adds the following contributions:

- We eliminate the assumption of excluded middle in the treatment of PA by means of a general Gödel-Gentzen-Friedman translation (Section 5).
- We mechanise direct and indirect reductions to various finitary set theories not requiring or actively refuting infinite sets (Section 8).
- We extend on the signature transformation employed for set theory without function symbols to obtain conservativity results (Lemma 54 - Fact 57).
- We analyse the abstract preconditions necessary for the synthetic approach to undecidability and incompleteness of arbitrary formalisms (Section 9).

¹The Coq development is available at www.ps.uni-saarland.de/extras/axiomatisations-ext and systematically hyperlinked with every definition and fact in the PDF version of this document.

After a preliminary discussion of constructive type theory, synthetic undecidability, and first-order logic in Section 2, we proceed with the general results relating undecidability, incompleteness, and consistency of first-order axiom systems in Section 3. This is followed by the case studies concerning arithmetical axiomatisations (Sections 4 and 5), set theory with (Section 6) and without (Section 7) Skolem functions, as well as finitary set theories (Section 8). We conclude with the abstract analysis of undecidability and incompleteness of arbitrary formalisms (Section 9) and with a discussion of the Coq mechanisation as well as related and future work Section 10.

2 Preliminaries

In order to make this paper self-contained and accessible, we briefly outline the synthetic approach to undecidability proofs and the representation of first-order logic in constructive type theory used in previous papers.

2.1 Constructive Type Theory

We work in the framework of a constructive type theory such as the one implemented in Coq, providing a predicative hierarchy of *type universes* above a single impredicative universe \mathbb{P} of *propositions*. On type level, we have the unit type $\mathbb{1}$ with a single element $*$: $\mathbb{1}$, the void type $\mathbb{0}$, function spaces $X \rightarrow Y$, products $X \times Y$, sums $X + Y$, dependent products $\forall(x : X). F x$, and dependent sums $\Sigma(x : X). F x$. On propositional level, these types are denoted by the usual logical notation (\top , \perp , \rightarrow , \wedge , \vee , \forall , and \exists). So-called *large elimination* from \mathbb{P} into computational types is restricted, in particular case distinction on proofs of \vee and \exists to form computational values is disallowed. On the other hand, this restriction is permeable enough to allow large elimination of the equality predicate $= : \forall X. X \rightarrow X \rightarrow \mathbb{P}$ specified by the constructor $\forall(x : X). x = x$, as well as function definitions by well-founded recursion.

We employ the basic inductive types of *Booleans* ($\mathbb{B} := \text{tt} \mid \text{ff}$), *Peano natural numbers* ($n : \mathbb{N} := 0 \mid n + 1$), the *option type* ($\mathbb{O}(X) := \ulcorner x \urcorner \mid \emptyset$), and *lists* ($l : \mathbb{L}(X) := [] \mid x :: l$). We write $|l|$ for the length of a list, $l ++ l'$ for the concatenation of l and l' , $x \in l$ for membership, and just fl for application of the pointwise map function. We denote by X^n the type of *vectors* \vec{v} of length $n : \mathbb{N}$ over X and reuse the definitions and notations introduced for lists.

2.2 Synthetic Undecidability

The base of the synthetic approach to computability theory [7, 16] is the fact that all functions definable in a constructive foundation are computable. This fact applies to many variants of constructive type theory and we let the assumed variant sketched in the previous section be one of those. Of course, we are confident that in particular the polymorphic calculus of cumulative inductive constructions (pCuIC) [17] currently implemented in Coq satisfies this condition although there is no formal proof yet.

Now beginning with positive notions, we can introduce decidability and enumerability of decision problems synthetically, i.e. without reference to a formal model of computation:

Definition 1. Let $P : X \rightarrow \mathbb{P}$ be a predicate over a type X .

- P is decidable if there exists $f : X \rightarrow \mathbb{B}$ with Px iff $fx = \mathbf{tt}$,
- P is enumerable if there exists $f : \mathbb{N} \rightarrow \mathbb{O}(X)$ with Px iff $\exists n. fn = \lceil x \rceil$.

Note that it is commonly accepted practice to mechanise decidability results in this synthetic sense (e.g. [18–20]). In the present paper, however, we mostly consider negative results in the form of undecidability of decision problems regarding first-order axiomatisations. Such negative results cannot be established in form of the actual negation of positive results, since constructive type theory is consistent with strong classical axioms turning every problem (synthetically) decidable (as witnessed by classical models, cf. [21]).

The approximation chosen in the Coq Library of Undecidability Proofs [9] is to call P (synthetically) undecidable if the decidability of P would imply the decidability of a seed problem known to be undecidable, specifically the halting problem for Turing machines. Therefore the negative notion can be turned into a positive notion, namely the existence of a computable reduction function, that again admits a synthetic rendering:

Definition 2. Given predicates $P : X \rightarrow \mathbb{P}$ and $Q : Y \rightarrow \mathbb{P}$, we call a function $f : X \rightarrow Y$ a (many-one) reduction if Px iff $Q(fx)$ for all x . We write $P \preceq Q$ if such a function exists.

Then interpreting reductions from the halting problem for Turing machines as undecidability results is backed by the following fact:

Fact 3. If $P \preceq Q$ and Q is decidable, then so is P .

Such reductions have already been verified for Hilbert’s tenth problem (H_{10}) [22] and the Post correspondence problem (PCP) [23] that we employ in the present paper, so by transitivity it is enough to verify continuing reductions to the axiom systems considered.

2.3 Syntax, Semantics, and Deduction Systems of FOL

We now review the representation of first-order syntax, semantics, and natural deduction systems developed in previous papers [8, 10, 24]. Beginning with the syntax, we describe terms $t : \mathbb{T}$ and formulas $\varphi : \mathbb{F}$ as inductive types over a fixed signature $\Sigma = (\mathcal{F}_\Sigma; \mathcal{P}_\Sigma)$ of function symbols $f : \mathcal{F}_\Sigma$ and relation symbols $P : \mathcal{P}_\Sigma$ with arities f and P :

$$t ::= x_n \mid f \vec{t} \quad (n : \mathbb{N}, \vec{t} : \mathbb{T}^f) \quad \varphi ::= P \vec{t} \mid \perp \mid \varphi \rightarrow \psi \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \forall \varphi \mid \exists \varphi \quad (\vec{t} : \mathbb{T}^P)$$

Negation $\neg \varphi$ and equivalence $\varphi \leftrightarrow \psi$ are obtained by the usual abbreviations.

In the chosen de Bruijn representation [25], a bound variable is encoded as the number of quantifiers shadowing its binder, e.g. $\forall x. \exists y. P x u \rightarrow P y v$ may be represented by $\forall \exists P x_1 x_4 \rightarrow P x_0 x_5$. For the sake of legibility, we write concrete formulas with named binders where instructive and defer de Bruijn encodings to the Coq development. A formula with all occurring variables bound by a quantifier is called *closed*.

Next, we define Tarski semantics providing an interpretation of formulas:

Definition 4. A model \mathcal{M} consists of a type D with functions $f^{\mathcal{M}} : D^f \rightarrow D$ and $P^{\mathcal{M}} : D^P \rightarrow \mathbb{P}$ interpreting the symbols in Σ . Given an assignment $\rho : \mathbb{N} \rightarrow D$ we define term evaluation $\hat{\rho} : \mathbb{T} \rightarrow D$ and formula satisfaction $\rho \models \varphi$ by

$$\hat{\rho} x_n := \rho n \quad \hat{\rho} (f \vec{t}) := f^{\mathcal{M}} (\hat{\rho} \vec{t}) \quad \rho \models P \vec{t} := P^{\mathcal{M}} (\hat{\rho} \vec{t}),$$

the remaining cases of $\rho \models \varphi$ map each connective to its meta-level counterpart.

If a model \mathcal{M} satisfies a formula φ for all variable assignments ρ , we write $\mathcal{M} \models \varphi$. Moreover, given a *theory* $\mathcal{T} : \mathbb{F} \rightarrow \mathbb{P}$, we write $\mathcal{M} \models \mathcal{T}$ if $\mathcal{M} \models \psi$ for all ψ with $\mathcal{T} \psi$ and $\mathcal{T} \models \varphi$ if $\mathcal{M} \models \mathcal{T}$ implies $\mathcal{M} \models \varphi$ for all \mathcal{M} . The same notations apply to (*finite*) *contexts* $\Gamma : \mathbb{L}(\mathbb{F})$.

Finally, we represent deduction systems as inductive predicates of type $\mathbb{L}(\mathbb{F}) \rightarrow \mathbb{F} \rightarrow \mathbb{P}$. We consider intuitionistic and classical natural deduction $\Gamma \vdash_i \varphi$ and $\Gamma \vdash_c \varphi$, respectively, and write $\Gamma \vdash \varphi$ if a statement applies to both variants. The rules of the two systems are standard and listed in Appendix A, here we only highlight the quantifier rules depending on the de Bruijn encoding

$$\frac{\Gamma[\uparrow] \vdash \varphi}{\Gamma \vdash \forall \varphi} \text{ AI} \quad \frac{\Gamma \vdash \forall \varphi}{\Gamma \vdash \varphi[t]} \text{ AE} \quad \frac{\Gamma \vdash \varphi[t]}{\Gamma \vdash \exists \varphi} \text{ EI} \quad \frac{\Gamma \vdash \exists \varphi \quad \Gamma[\uparrow], \varphi \vdash \psi[\uparrow]}{\Gamma \vdash \psi} \text{ EE}$$

where $\varphi[\sigma]$ denotes the *capture-avoiding instantiation* of a formula φ with a *parallel substitution* $\sigma : \mathbb{N} \rightarrow \mathbb{T}$, where the substitution \uparrow maps n to x_{n+1} , where the substitution $(t; \sigma)$ maps 0 to t and $n + 1$ to σn , and where $\varphi[t]$ is short for $\varphi[t; (\lambda n. x_n)]$. Extending the deduction systems to theories $\mathcal{T} : \mathbb{F} \rightarrow \mathbb{P}$, we write $\mathcal{T} \vdash \varphi$ if there is $\Gamma \subseteq \mathcal{T}$ with $\Gamma \vdash \varphi$.

Constructively, only soundness of the intuitionistic system ($\mathcal{T} \vdash_i \varphi$ implies $\mathcal{T} \models \varphi$) is provable without imposing a restriction on the admitted models (as done in [24]). However, it is easy to verify the usual weakening ($\Gamma \vdash \varphi$ implies $\Delta \vdash \varphi$ for $\Gamma \subseteq \Delta$) and substitution ($\Gamma \vdash \varphi$ implies $\Gamma[\sigma] \vdash \varphi[\sigma]$) properties of both variants by induction on the given derivations. The latter gives rise to named reformulations of (AI) and (EE) helpful in concrete derivations

$$\frac{\Gamma \vdash \varphi[x_n]}{\Gamma \vdash \forall \varphi} \quad x_n \notin \Gamma, \varphi \quad \frac{\Gamma \vdash \exists \varphi \quad \Gamma, \varphi[x_n] \vdash \psi}{\Gamma \vdash \psi} \quad x_n \notin \Gamma, \varphi, \psi$$

where $x_n \notin \Gamma$ denotes that x_n is *fresh*, i.e. does not occur in any formula of Γ .

The concrete signatures used in this paper all contain a reserved binary relation symbol \equiv for equality. Instead of making equality primitive in the syntax, semantics, and deduction systems, we implicitly restrict $\mathcal{M} \models \varphi$ to

extensional models \mathcal{M} interpreting \equiv as actual equality $=$ and define $\mathcal{T} \vdash \varphi$ as derivability from \mathcal{T} augmented with the standard axioms characterising \equiv as an equivalence relation congruent for the symbols in Σ .

3 Undecidable and Incomplete Axiom Systems

In this section, we record some general algorithmic facts concerning first-order axiomatisations and outline the common scheme underlying the undecidability proofs presented in the subsequent two sections. We fix an enumerable and discrete signature Σ for the remainder of this section and begin by introducing the central notion of axiom systems formally.

Definition 5. We call $\mathcal{A} : \mathbb{F} \rightarrow \mathbb{P}$ an axiomatisation if \mathcal{A} is enumerable.

Any given axiomatisation induces two related decision problems, namely semantic entailment $\mathcal{A}^{\varepsilon} := \lambda\varphi. \mathcal{A} \varepsilon \varphi$ and deductive entailment $\mathcal{A}^{\vdash} := \lambda\varphi. \mathcal{A} \vdash \varphi$. Since in our constructive setting we can show the classical deduction system \vdash_c neither sound nor complete (cf. [24]), we mostly consider a combined notion of classical semantics and intuitionistic deduction:

Definition 6. We say that a predicate $P : X \rightarrow \mathbb{P}$ reduces to \mathcal{A} , written $P \preceq \mathcal{A}$, if there is a function $f : X \rightarrow \mathbb{F}$ witnessing both $P \preceq \mathcal{A}^{\varepsilon}$ and $P \preceq \mathcal{A}^{\vdash}$.

Assuming the law of excluded middle $\text{LEM} := \forall p : \mathbb{P}. p \vee \neg p$ would be sufficient to obtain $P \preceq \mathcal{A}^{\vdash_c}$ from $P \preceq \mathcal{A}^{\varepsilon}$, since then $\mathcal{A} \vdash_c \varphi$ and $\mathcal{A} \varepsilon \varphi$ coincide. In fact, already the soundness direction is enough for our case studies on PA and ZF, since for them it is still feasible to verify $\mathcal{A} \vdash f x$ given $P x$ by hand without appealing to completeness.

We now formulate two facts stating the well-known connections of undecidability with consistency and incompleteness for our synthetic setting. The first observation is that verifying a reduction from a non-trivial problem is at least as hard as a consistency proof.

Fact 7. If $P \preceq \mathcal{A}^{\vdash}$ and there is x with $\neg P x$, then $\mathcal{A} \not\vdash \perp$.

Proof If $f : X \rightarrow \mathbb{F}$ witnesses $P \preceq \mathcal{A}^{\vdash}$, then by $\neg P x$ we obtain $\mathcal{A} \not\vdash f x$. This prohibits a derivation $\mathcal{A} \vdash \perp$ by the explosion rule (E). \square

The second observation is a synthetic version of (negation-)incompleteness for all axiomatisations strong enough to express an undecidable problem. We follow the common practice to focus on incompleteness of the classical deduction system, see Section 10.1 for a discussion.

Definition 8. We call \mathcal{A} complete if for all closed φ , $\mathcal{A} \vdash_c \varphi$ or $\mathcal{A} \vdash_c \neg\varphi$.

Fact 9. If \mathcal{A} is complete and $\mathcal{A} \not\vdash_c \perp$, then \mathcal{A}^{\vdash_c} is decidable on closed formulas. Hence, if f witnesses $P \preceq \mathcal{A}^{\vdash_c}$ such that all $f x$ are closed, then P is decidable.

8 *Synthetic Undecidability and Incompleteness of First-Order Axiom Systems in Coq*

Proof By a synthetic version of Post's theorem ([8, Lemma 2.15]) it suffices to show that \mathcal{A}^{\vdash_c} is bi-enumerable, i.e. both $\lambda\varphi. \mathcal{A} \vdash_c \varphi$ and $\lambda\varphi. \mathcal{A} \not\vdash_c \varphi$ are enumerable, and logically decidable, i.e. $\mathcal{A} \vdash_c \varphi$ or $\mathcal{A} \not\vdash_c \varphi$ for all φ . This follows by enumerability of \vdash_c and since by completeness $\mathcal{A} \not\vdash_c \varphi$ iff $\mathcal{A} \vdash_c \neg\varphi$. The consequence is by Fact 3. \square

Note that this fact is an approximation of the usual incompleteness theorem in two ways. First, similar to the synthetic rendering of undecidability, axiomatisations \mathcal{A} subject to a reduction $P \preceq \mathcal{A}^{\vdash_c}$ for P known to be undecidable are only shown incomplete in the sense that their completeness would imply decidability of P . Deriving an actual contradiction would rely on computability axioms (e.g. Church's thesis [26, 27] or an undecidability assumption [8]) or extraction to a concrete model (e.g. a weak call-by-value λ -calculus [28]). Secondly, the fact does not produce a witness of an independent formula the way a more informative proof based on Gödel sentences does. Also note that inconsistent axiomatisations are trivially decidable, so the requirement $\mathcal{A} \not\vdash_c \perp$ is inessential (especially given Fact 7).

Next, we outline the general pattern of the reductions verified in this paper:

1. We choose an undecidable seed problem $P : X \rightarrow \mathbb{P}$ easy to encode in the target axiomatisation. This will be H_{10} for PA and PCP for ZF.
2. We define the translation function $X \rightarrow \mathbb{F}$ mapping instances $x : X$ to formulas φ_x in a way compact enough to be stated without developing much of the internal theory of \mathcal{A} .
3. We isolate a finite fragment $A \subseteq \mathcal{A}$ of axioms that suffices to implement the main argument. This yields a reusable factorisation and is easier to mechanise.
4. We verify the semantic part locally by showing for every \mathcal{M} with $\mathcal{M} \models A$ that Px iff $\mathcal{M} \models \varphi_x$. For the backwards direction, we in fact need to restrict \mathcal{M} to satisfy a suitable property of standardness allowing us to reconstruct an actual solution of P .
5. We construct standard models for A and \mathcal{A} , possibly relying on assumptions.
6. We verify the deductive part by establishing that Px implies $A \vdash \varphi_x$, closely following the semantic proof from before. The backwards direction follows from soundness.
7. We conclude that A , \mathcal{A} , and any $\mathcal{B} \supseteq A$ are undecidable:

Theorem 10. *Let a problem $P : X \rightarrow \mathbb{P}$, an axiomatisation \mathcal{A} , a notion of standardness on models $\mathcal{M} \models \mathcal{A}$, and a function $\varphi_- : X \rightarrow \mathbb{F}$ be given with:*

1. Px implies $\mathcal{A} \models \varphi_x$.
2. Every standard model $\mathcal{M} \models \mathcal{A}$ with $\mathcal{M} \models \varphi_x$ yields Px .
3. Px implies $\mathcal{A} \vdash \varphi_x$.

Then $P \preceq \mathcal{B}$ for all $\mathcal{B} \supseteq \mathcal{A}$ admitting a standard model. If we additionally assume LEM, then also $P \preceq \mathcal{B}^{\vdash_c}$.

Proof We begin with $P \preceq \mathcal{B}^{\models}$. That Px implies $\mathcal{B} \models \varphi_x$ is direct by (i) since every model of \mathcal{B} is a model of \mathcal{A} . Conversely, if $\mathcal{B} \models \varphi_x$ then in particular the assumed standard model $\mathcal{M} \models \mathcal{B}$ satisfies φ_x . Thus we obtain Px by (ii).

Turning to $P \preceq \mathcal{B}^{+i}$, the first direction is again trivial, this time by (iii) and weakening. For the converse, we assume that $\mathcal{B} \vdash_i \varphi_x$ and hence $\mathcal{B} \vDash \varphi_x$ by soundness. Thus we conclude $P x$ with the previous argument relying on (ii).

Finally, with LEM we obtain $P \preceq \mathcal{B}^{+c}$ since then $\mathcal{B} \vdash_c \varphi_x$ implies $\mathcal{B} \vDash \varphi_x$. \square

Of course (i) follows from (iii) via soundness, so the initial semantic verification could be eliminated from Theorem 10 and the informal strategy outlined before. However, we deem it more instructive to first present a self-contained semantic verification without the overhead introduced by working in a syntactic deduction system, mostly apparent in the Coq mechanisation. Also note that the necessity of a standard model will be no burden in the treatment of PA but in the case of ZF this will require a careful analysis of preconditions.

We end this section with the unsurprising but still instructive fact that the decision problem for finite axiomatisations A reduces to the general Entscheidungsproblem of first-order logic concerning validity and provability in the empty context [1].

Fact 11. For $A : \mathbb{L}(\mathbb{F})$ we have $A^{\vDash} \preceq (\lambda\varphi. \vDash \varphi)$ and $A^{+} \preceq (\lambda\varphi. \vdash \varphi)$.

Proof It is straightforward to verify that the function $\lambda\varphi. \bigwedge A \rightarrow \varphi$ prefixing φ with the conjunction of all formulas in A establishes both reductions. \square

So the reductions to finite fragments of PA and ZF presented in the next sections in particular complement the direct reductions to the Entscheidungsproblem given in [8]. More general variants of this insight can be formulated as follows:

Fact 12. Let A be finite and \mathcal{B} be an arbitrary axiomatisation.

1. If $A \vdash \mathcal{B}$, then $A \preceq \mathcal{B}$.
2. If $\mathcal{B} \subseteq A$, then $A \preceq \mathcal{B}$.
3. $\mathcal{B} \cup A \preceq \mathcal{B}$.

Proof All witnessed by the reduction $\lambda\varphi. \bigwedge A \rightarrow \varphi$, (2) is a special case of (1). \square

4 Peano Arithmetic

We begin with a rather simple case study to illustrate our general approach to undecidability and incompleteness. For the theory of Peano arithmetic (PA) we use a signature containing symbols for the constant zero, the successor function, addition, multiplication and equality:

$$(O, S_-, - \oplus -, - \otimes -; - \equiv -)$$

The core of PA consists of axioms characterising addition and multiplication:

$$\begin{array}{ll} \oplus\text{-base: } \forall x. O \oplus x \equiv x & \oplus\text{-recursion: } \forall xy. (Sx) \oplus y \equiv S(x \oplus y) \\ \otimes\text{-base: } \forall x. O \otimes x \equiv O & \otimes\text{-recursion: } \forall xy. (Sx) \otimes y \equiv y \oplus x \otimes y \end{array}$$

The finite list Q' consisting of these four axioms is strong enough to be undecidable. Undecidability (and incompleteness) then transport in particular to the (infinite) axiomatisation PA adding

$$\text{Disjointness: } \forall x. Sx \equiv O \rightarrow \perp \quad \text{Injectivity: } \forall xy. Sx \equiv Sy \rightarrow x \equiv y$$

and the [axiom scheme of induction](#), which we define as a function on formulas:

$$\lambda\varphi. \varphi[O] \rightarrow (\forall x. \varphi[x] \rightarrow \varphi[Sx]) \rightarrow \forall x. \varphi[x]$$

Another typical reference point for incompleteness is Robinson arithmetic Q , obtained by replacing the induction scheme by $\forall x. x \equiv O \vee \exists y. x \equiv Sy$.

Turning to undecidability, Hilbert's 10th problem (H_{10}) is concerned with the solvability of Diophantine equations and comes as a natural seed problem for showing the undecidability of PA, since the equations are a syntactic fragment of PA formulas. To be more precise, H_{10} consists of deciding whether a Diophantine equation $p = q$ has a solution in the natural numbers \mathbb{N} , where p, q are polynomials constructed by parameters, variables, addition, and multiplication:

$$p, q ::= a_n \mid \text{var } k \mid \text{add } p \ q \mid \text{mult } p \ q \quad (n, k : \mathbb{N})$$

Evaluation $\llbracket p \rrbracket_\alpha$ of a polynomial p for an assignment $\alpha : \mathbb{N} \rightarrow \mathbb{N}$ is defined by

$$\llbracket a_n \rrbracket_\alpha := n \quad \llbracket \text{var } k \rrbracket_\alpha := \alpha k \quad \llbracket \text{add } p \ q \rrbracket_\alpha := \llbracket p \rrbracket_\alpha + \llbracket q \rrbracket_\alpha \quad \llbracket \text{mult } p \ q \rrbracket_\alpha := \llbracket p \rrbracket_\alpha \times \llbracket q \rrbracket_\alpha$$

and a Diophantine equation $p = q$ then has a solution, if there is α with $\llbracket p \rrbracket_\alpha = \llbracket q \rrbracket_\alpha$. Given their similarity, it is easy to encode H_{10} into PA, beginning with numerals:

Definition 13. We define $\nu(n) : \mathbb{T}$ by $\nu(0) := O$ and $\nu(n+1) := S(\nu(n))$.

We now translate polynomials into PA terms by defining $p^* : \mathbb{T}$ recursively:

$$a_n^* := \nu(n) \quad (\text{var } k)^* := x_k \quad (\text{add } p \ q)^* := p^* \oplus q^* \quad (\text{mult } p \ q)^* := p^* \otimes q^*$$

A Diophantine equation with greatest free variable N can now be encoded as the formula $\varphi_{p,q} := \exists^N p^* \equiv q^*$ where we use N leading existential quantifiers to internalise the solvability condition. The formula $\varphi_{p,q}$ thus asserts the existence of a solution for $p = q$ which gives us a natural encoding from Diophantine equations into PA.

We prepare the verification of the three requirements (Facts 20, 22 and 25) necessary for Theorem 10 with the following lemma about existential formulas:

Lemma 14. If $\exists^N \varphi$ is closed, then

1. $\mathcal{M} \models \exists^N \varphi$ iff there is $\rho : \mathbb{N} \rightarrow \mathcal{M}$ such that $\rho \models \varphi$,
2. $\Gamma \vdash \exists^N \varphi$ if there is $\sigma : \mathbb{N} \rightarrow \mathbb{T}$ such that $\Gamma \vdash \varphi[\sigma]$.

Proof We only provide some intuition for (i). For the implication from left to right, the assumption $\mathcal{M} \models \exists^N \varphi$ gives us $x_1, \dots, x_N : \mathcal{M}$ such that $\forall \rho. x_1; \dots; x_N; \rho \models \varphi$, so in particular we have $\rho' \models \varphi$ for $\rho' := x_1; \dots; x_N; (\lambda x. O^{\mathcal{M}})$, showing the claim. For the other implication, we get ρ with $\rho \models \varphi$. By setting $\rho' := \lambda x. \rho(x + N)$ we have $\rho = \rho(0); \dots; \rho(N); \rho'$ and hence there are $x_1, \dots, x_N : \mathcal{M}$ such that $x_1; \dots; x_N; \rho' \models \varphi$. Since φ has at most N free variables, ρ' can be exchanged with any other $\tau : \mathbb{N} \rightarrow \mathcal{M}$. \square

By Lemma 14, showing $\varphi_{p,q}$ is equivalent to finding a satisfying environment $\rho : \mathbb{N} \rightarrow \mathcal{M}$ for $p^* \equiv q^*$ in a model \mathcal{M} or deductively showing that a substitution $\sigma : \mathbb{N} \rightarrow \mathbb{T}$ solves it. This enables us to transport a solution for $p = q$ to both the model and the deduction system.

We now verify the semantic part of the reduction for the axiomatic fragment \mathbf{Q}' . To this end, we fix a model $\mathcal{M} \models \mathbf{Q}'$ for the next definitions and lemmas.

Definition 15. We define $\mu(n) : D$ by $\mu(0) := O^{\mathcal{M}}$ and $\mu(n + 1) := S^{\mathcal{M}}(\mu(n))$.

The axioms in \mathbf{Q}' are sufficient to prove that μ is a homomorphism.

Lemma 16. $\mu(n + m) = \mu(n) \oplus^{\mathcal{M}} \mu(m)$ and $\mu(n + m) = \mu(n) \otimes^{\mathcal{M}} \mu(m)$.

Proof The proof for addition is done by induction on $n : \mathbb{N}$ and using the axioms for addition in \mathbf{Q}' . The proof for multiplication is done in the same fashion, using the axioms for multiplication and the previous result for addition. \square

Lemma 17. For any $\rho : \mathbb{N} \rightarrow \mathcal{M}$ and $n : \mathbb{N}$ we have $\hat{\rho}(\nu(n)) = \mu(n)$.

Given an assignment $\alpha : \mathbb{N} \rightarrow \mathbb{N}$, we can transport the evaluation of a polynomial $\llbracket p \rrbracket_{\alpha}$ to any \mathbf{Q}' model by applying μ . The homomorphism property of μ now makes it easy to verify that we get the same result by evaluating the encoded version p^* with the composition $\mu \circ \alpha$.

Lemma 18. For any p and $\alpha : \mathbb{N} \rightarrow \mathbb{N}$ we have $\widehat{(\mu \circ \alpha)}(p^*) = \mu(\llbracket p \rrbracket_{\alpha})$.

Proof By induction on p , using Lemmas 17 and 18. \square

Corollary 19. If $p = q$ has solution α , then in any \mathbf{Q}' model $(\mu \circ \alpha) \models p^* \equiv q^*$.

Proof $\mu(\llbracket p \rrbracket_{\alpha}) = \mu(\llbracket q \rrbracket_{\alpha}) \xrightarrow{L.18} \widehat{(\mu \circ \alpha)}(p^*) = \widehat{(\mu \circ \alpha)}(q^*) \implies (\mu \circ \alpha) \models p^* \equiv q^*$ \square

Fact 20. If $p = q$ has a solution, then $\mathbf{Q}' \models \varphi_{p,q}$.

Proof Let α be the solution of $p = q$, then $(\mu \circ \alpha) \models p^* \equiv q^*$ holds by Corollary 19 and since $\exists^N p^* \equiv q^*$ is closed by construction, the goal follows by Lemma 14. \square

For the converse direction, we employ the type \mathbb{N} as standard model.

Lemma 21. \mathbb{N} is a model of \mathbf{Q}' , \mathbf{Q} , and PA.

It is easy to extract a solution of $p = q$ if $\mathbb{N} \models \varphi_{p,q}$ by the previous lemmas.

Fact 22. *If $\mathbb{N} \models \varphi_{p,q}$ then $p = q$ has a solution.*

Proof By assumption we have $\mathbb{N} \models \varphi_{p,q}$ which by Lemma 14 gives us $\alpha : \mathbb{N} \rightarrow \mathbb{N}$ with

$$\alpha \models p^* \equiv q^* \implies (\widehat{\mu \circ \alpha})(p^*) = (\widehat{\mu \circ \alpha})(q^*) \xrightarrow{L.18} \mu(\llbracket p \rrbracket_\alpha) = \mu(\llbracket q \rrbracket_\alpha).$$

Since over \mathbb{N} the function μ is simply the identity, we conclude $\llbracket p \rrbracket_\alpha = \llbracket q \rrbracket_\alpha$. \square

The deductive part of the reduction can be shown analogously to Fact 20, encoding the proofs of all intermediate results as ND derivations. We just list the relevant statements and refer to the Coq code for more detail.

Lemma 23. $\mathbb{Q}' \vdash \nu(n + m) \equiv \nu(n) \oplus \nu(m)$ and $\mathbb{Q}' \vdash \nu(n \times m) \equiv \nu(n) \otimes \nu(m)$.

Lemma 24. *If $p = q$ has a solution α , then we can deduce $\mathbb{Q}' \vdash (p^* \equiv q^*)[\nu \circ \alpha]$.*

Fact 25. *If $p = q$ has a solution then $\mathbb{Q}' \vdash \varphi_{p,q}$.*

Now we have all facts in place to verify the reductions with Theorem 10.

Theorem 26. $H_{10} \preceq \mathbb{Q}'$, $H_{10} \preceq \mathbb{Q}$, and $H_{10} \preceq \text{PA}$.

Proof Since \mathbb{N} is a standard model for \mathbb{Q}' , \mathbb{Q} , and PA, the claims follow by Theorem 10 since we have shown the three necessary conditions in Facts 20, 22 and 25. \square

As a consequence of the reductions, we conclude incompleteness as follows:

Theorem 27. *Assuming LEM, completeness of any extension $\mathcal{A} \supseteq \mathbb{Q}'$ satisfied by the standard model \mathbb{N} would imply the decidability of the halting problem.*

Proof By Theorems 10 and 26, Fact 9, and the reductions verified in [22]. \square

In fact, all axiomatisations satisfied by \mathbb{N} are undecidable and incomplete:

Fact 28. $H_{10} \preceq \mathcal{A}$ for any axiomatisation \mathcal{A} satisfied by the standard model \mathbb{N} .

Proof Theorems 10 and 26 imply $H_{10} \preceq \mathcal{A} \cup \mathbb{Q}'$ and by Fact 12 also $\mathcal{A} \cup \mathbb{Q}' \preceq \mathcal{A}$. \square

We close this section with a few remarks about the theories \mathbb{Q}' , \mathbb{Q} , and PA. The theory \mathbb{Q}' is trivially incomplete under LEM: using soundness of classical deduction, we have $\mathbb{Q}' \not\vdash_c \forall xy. x = y$ because of the standard model \mathbb{N} and $\mathbb{Q}' \not\vdash_c \neg \forall xy. x = y$ because of the trivial model. Similarly, the formula $\forall x. Sx \neq x$ is independent in \mathbb{Q} , for instance violated by the model \mathbb{N}^∞ extending \mathbb{N} with a maximal number ∞ . Note that these models in particular show that the theories \mathbb{Q}' , \mathbb{Q} , and PA are all distinct.

5 Eliminating the Law of Excluded Middle

We can strengthen the result of Theorem 27 and remove its reliance on LEM by utilising a combination of the double negation and Friedman translations [29]. Given any signature $\Sigma = (\mathcal{F}_\Sigma; \mathcal{P}_\Sigma)$ we add a new 0-ary predicate F to \mathcal{P}_Σ , giving us the new signature $\Sigma^F := (\mathcal{F}_\Sigma, \mathcal{P}_\Sigma \cup \{F\})$. This way of setting up the Friedman transform is easier to mechanise compared to the syntactic version where \perp is replaced by a formula, and sufficient for our purpose here.

Definition 29. We recursively define the F -translation $(\cdot)^F : \mathbb{F}_\Sigma \rightarrow \mathbb{F}_{\Sigma^F}$ by:

$$\begin{aligned} \perp^F &:= F & (\varphi \rightarrow \psi)^F &:= \varphi^F \rightarrow \psi^F & (\forall \varphi)^F &:= \forall \varphi^F \\ (P\vec{t})^F &:= \dot{\neg} \dot{\neg} (P\vec{t}) & (\varphi \wedge \psi)^F &:= \varphi^F \wedge \psi^F & (\exists \varphi)^F &:= \dot{\neg} \dot{\neg} \exists \varphi^F \\ & & (\varphi \vee \psi)^F &:= \dot{\neg} \dot{\neg} (\varphi^F \vee \psi^F) & & \end{aligned}$$

where $\dot{\neg} \varphi$ is short for $\varphi \rightarrow F$. We extend $(\cdot)^F$ to contexts Γ as expectable.

We will state the crucial results concerning the F -translation with respect to minimal natural deduction $\Gamma \vdash_m \varphi$, which is natural deduction \vdash_i without the explosion rule and restricted to formulas without the \perp symbol.

Lemma 30. For any formula φ we have $\vdash_m \dot{\neg} \dot{\neg} \varphi^F \rightarrow \varphi^F$.

Proof By induction on the size of φ . □

Lemma 31. For any formula φ and context Γ , if $\Gamma \vdash_c \varphi$ then $\Gamma^F \vdash_m \varphi^F$.

Proof By induction on the deduction $\Gamma \vdash_c \varphi$, some cases need Lemma 30. □

Definition 32. Given a proposition $P : \mathbb{P}$ and model \mathcal{M} of the signature Σ , we can extend \mathcal{M} to a model \mathcal{M}^F of the extended signature Σ^F by setting $F^{\mathcal{M}} := P$ and following the interpretation of \mathcal{M} in all other cases.

We now apply the F -translation to the particular case of the PA signature to derive an improved version of Theorem 27, eliminating the usage of LEM.

Lemma 33. If $\mathcal{M}^P \models (\varphi_{p,q})^F$ then $\mathcal{M}^P \models \dot{\neg} \dot{\neg} \varphi_{p,q}$.

Proof By $\mathcal{M}^P \models \exists^N \dot{\neg} \dot{\neg} (p^* \equiv q^*) \rightarrow \dot{\neg} \dot{\neg} \exists^N p^* \equiv q^*$, proved inductively on N . □

Definition 34. We call a theory \mathcal{T} preserved under F -translation if for every model $\mathcal{M} \models \mathcal{T}$ and $P : \mathbb{P}$ we have $\mathcal{M}^P \models \mathcal{T}^F$, denoted by $\mathcal{T} \models \mathcal{T}^F$.

It is easy to show that Q', Q and PA are preserved under the F -translation, since for every axiom $\alpha \in \text{PA}$ one can verify that $\mathcal{M} \models \alpha \rightarrow \alpha^F$.

Theorem 35. *Any axiomatisation $\mathcal{A} \supseteq \mathcal{Q}'$ satisfied by the standard model \mathbb{N} and preserved under the F -translation admits a reduction $\mathsf{H}_{10} \preceq \mathcal{A}^{\perp c}$. Hence, its completeness would imply the decidability of the halting problem.*

Proof First we will show $\mathsf{H}_{10} \preceq \mathcal{A}^{\perp c}$, by verifying that $\varphi_{p,q}$ is a reduction, where the first part of the verification follows from Fact 25. In the converse we are given $\Gamma \subseteq \mathcal{A}$ with $\Gamma \vdash_c \varphi_{p,q}$ and need to find a solution for $p = q$ or equivalently (Fact 22) need to show $\mathbb{N} \models \varphi_{p,q}$. Utilising Lemma 31 we get $\Gamma^F \vdash_m (\varphi_{p,q})^F$ which by soundness gives $\mathcal{M}^P \models \bigwedge \Gamma^F \rightarrow (\varphi_{p,q})^F$ in every model \mathcal{M}^P . Since $\mathcal{A} \models \mathcal{A}^F$ we have $\mathbb{N}^P \models \Gamma^F$ and therefore $\mathbb{N}^P \models (\varphi_{p,q})^F$. By Lemma 33 this gives us $\mathbb{N}^P \models \neg\neg\varphi_{p,q}$, which reduces to $(\mathbb{N} \models \varphi_{p,q} \rightarrow P) \rightarrow P$. The model with $P := \mathbb{N} \models \varphi_{p,q}$ then proves that $\mathbb{N} \models \varphi_{p,q}$.

Secondly, we can show that \mathcal{A} is consistent (with respect to \vdash_c) by the fact that $\mathcal{A} \vdash_c \perp$ together with Lemma 31 and soundness implies $\mathbb{N}^P \models \perp^F$, which reduces to \perp in the model with $P := \perp$. Therefore by Fact 9, completeness of \mathcal{A} would imply the decidability of H_{10} and thus also of the halting problem. \square

6 ZF Set Theory with Skolem Functions

Turning to set theory, we first work in a signature providing function symbols for the operations of ZF. So for the rest of this section we fix the signature

$$\Sigma := (\emptyset, \{-, \}, \cup -, \mathcal{P}(-), \omega ; - \equiv -, - \in -)$$

with function symbols denoting the empty set, pairing, union, power set, the set of natural numbers, next to the usual relation symbols for equality and membership. Using such Skolem functions for axiomatic and other definable operations is common practice in set-theoretic literature and eases the definition and verification of the undecidability reduction in our case. That the undecidability result can be transported to minimal signatures just containing equality and membership, or even just the latter, is subject of the next section.

We do not list all axioms in detail but refer the reader to Appendix B, the Coq code, and standard literature (eg. [30]). The only point worth mentioning again is the representation of axiom schemes as functions $\mathbb{F} \rightarrow \mathbb{F}$, for instance by the [separation scheme](#) expressed as

$$\lambda\varphi. \forall x. \exists y. \forall z. z \in y \leftrightarrow z \in x \wedge \varphi[x].$$

We then distinguish the following axiomatisations:

- Z' contains extensionality and the specifications of the function symbols.
- Z is obtained by adding all instances of the separation scheme.
- ZF is obtained by further adding all instances of the replacement scheme.

Note that in ZF we do not include the axiom of regularity since this would force the theory classical and would require to extend Coq's type theory even further to obtain a model [31]. Alternatively, one could add the more constructive axiom for ϵ -induction, but instead we opt for staying more general and just leave the well-foundedness of sets unspecified.

Following the general outline for the undecidability proofs in this paper, we first focus on verifying a reduction to the base theory Z' and then extend to the stronger axiomatisations by use of Theorem 10. As a seed problem for this reduction, we could naturally pick just any decision problem since set theory is a general purpose foundation expressive enough for most standard mathematics. However, the concrete choice has an impact on the mechanisation overhead, where formalising Turing machine halting directly is tricky enough in Coq's type theory itself, and even a simple problem like H_{10} used in the previous section would presuppose a modest development of number theory and recursion in the axiomatic framework. We therefore base our reduction to Z' on the Post correspondence problem (PCP) which has a simple inductive characterisation expressing a matching problem given a finite stack S of pairs (s, t) of Boolean strings:

$$\frac{(s, t) \in S}{S \triangleright (s, t)} \qquad \frac{S \triangleright (u, v) \quad (s, t) \in S}{S \triangleright (su, tv)} \qquad \frac{S \triangleright (s, s)}{\text{PCP } S}$$

Informally, S is used to derive pairs (s, t) , written $S \triangleright (s, t)$, by repeatedly appending the pairs from the stack componentwise in any order or multitude. S admits a solution, written $\text{PCP } S$, if a matching pair (s, s) can be derived.

Encoding data like numbers and Booleans in set theory is standard, using usual notations for binary union $x \cup y$, singletons $\{x\}$, and ordered pairs (x, y) :

- Numbers: $\bar{0} := \emptyset$ and $\overline{n+1} := \bar{n} \cup \{\bar{n}\}$
- Booleans: $\bar{\text{tt}} := \{\emptyset\}$ and $\bar{\text{ff}} := \emptyset$
- Strings: $\overline{b_1, \dots, b_n} := (\overline{b_1}, (\dots (\overline{b_n}, \emptyset) \dots))$
- Stacks: $\overline{S} := \{(\overline{s_1}, \overline{t_1}), \dots, (\overline{s_m}, \overline{t_m})\}$

Starting informally, the solvability condition of PCP can be directly expressed in set theory by just asserting the existence of a set encoding a match for S :

$$\exists x. (x, x) \in \bigcup_{k \in \omega} \overline{S}^k \quad \text{where } \overline{S}^0 = \overline{S} \quad \text{and } \overline{S}^{k+1} = S \boxtimes \overline{S}^k = \bigcup_{s/t \in S} \{(\overline{sx}, \overline{ty}) \mid (x, y) \in \overline{S}^k\}$$

Unfortunately, formalizing this idea is not straightforward, since the iteration operation \overline{S}^k is described by recursion on set-theoretic numbers $k \in \omega$ missing a native recursion principle akin to the one for type-theoretic numbers $n : \mathbb{N}$. Such a recursion principle can of course be derived but in our case it is simpler to inline the main construction.

The main construction used in the recursion theorem for ω is a sequence of finite approximations f accumulating the first k steps of the recursive equations. Since in our case we do not need to form the limit of this sequence requiring the approximations to agree, it suffices to ensure that at least the first k steps are contained without cutting off, namely

$$f \gg k := (\emptyset, \overline{S}) \in f \wedge \forall (l, B) \in f. l \in k \rightarrow (l \cup \{l\}, S \boxtimes B) \in f$$

where we reuse the operation $S \boxtimes B$ appending the encoded elements of the list S component-wise to the elements of the set B as specified above. Note

that this operation is not really definable as a function $\mathbb{L}(\mathbb{B}) \rightarrow \mathbb{T} \rightarrow \mathbb{T}$ and needs to be circumvented by quantifying over candidate sets satisfying the specification. However, for the sake of a more accessible explanation, we leave this subtlety to the Coq code and continue using $S \boxtimes B$ as a function.

Now solvability of S can be expressed formally as the existence of a functional approximation f of length k containing a match (x, x) :

$$\varphi_S := \exists k, f, B, x. k \in \omega \wedge (\forall (l, B), (l, B') \in f. B = B') \wedge f \gg k \wedge (k, B) \in f \wedge (x, x) \in B$$

We proceed with the formal verification of the reduction function $\lambda S. \varphi_S$ by proving the three facts necessary to apply Theorem 10. Again beginning with the semantic part for clarity, we fix a model $\mathcal{M} \models Z'$ for the next lemmas in preparation of the facts connecting PCP S with $\mathcal{M} \models \varphi_S$. We skip the development of basic set theory in \mathcal{M} reviewable in the Coq code and only state lemmas concerned with encodings and the reduction function:

Lemma 36. *Let $n, m : \mathbb{N}$ and $s, t : \mathbb{L}(\mathbb{B})$ be given, then the following hold:*

1. $\mathcal{M} \models \bar{n} \in \omega$
2. $\mathcal{M} \models \bar{n} \notin \bar{n}$
3. $\mathcal{M} \models \bar{n} \equiv \bar{m}$ implies $n = m$
4. $\mathcal{M} \models \bar{s} \equiv \bar{t}$ implies $s = t$

Proof 1. By induction on n , employing the infinity axiom characterising ω .

2. Again by induction on n , using the fact that numerals \bar{n} are transitive sets.

3. By trichotomy we have $n < m$, $m < n$, or $n = m$ as desired. If w.l.o.g. it were $n < m$, then $\mathcal{M} \models \bar{n} \in \bar{m}$ would follow by structural induction on the derivation of $n < m$. But then the assumption $\mathcal{M} \models \bar{n} \equiv \bar{m}$ would yield $\mathcal{M} \models \bar{n} \in \bar{n}$ in conflict with (ii).

4. By induction on the given strings, employing injectivity of $\bar{\cdot}$. □

In order to match the structure of iterated derivations encoded in φ_S , we reformulate $S \triangleright (s, t)$ by referring to the composed derivations S^n of length n , now definable by recursion on $n : \mathbb{N}$ via $S^0 := S$ and $S^{n+1} := S \boxtimes S^n$ reusing the operation \boxtimes for lists as expected.

Lemma 37. *$S \triangleright (s, t)$ iff there is $n : \mathbb{N}$ with $(s, t) \in S^n$.*

Then S^n can be encoded as set-level functions $f_S^n := \{(\emptyset, \bar{S}), \dots, (\bar{n}, \bar{S}^n)\}$ that are indeed recognised by the model \mathcal{M} as correct approximations:

Lemma 38. *For every $n : \mathbb{N}$ we have $\mathcal{M} \models f_S^n \gg \bar{n}$.*

Proof In this proof we work inside of \mathcal{M} to simplify intermediate statements. For the first conjunct, we need to show that $(\emptyset, \bar{S}) \in f_S^n$ which is straightforward since $(\emptyset, \bar{S}) \in f_S^0$ and $f_S^m \subseteq f_S^n$ whenever $m \leq n$. Regarding the second conjunct, we assume $(k, B) \in f_S^n$ with $k \in \bar{n}$ and need to show $(k \cup \{k\}, S \boxtimes B) \in f_S^n$. From $(k, B) \in f_S^n$ we obtain that there is m with $k \in \bar{m}$ and $B = \bar{S}^m$. Then from $\bar{m} \in \bar{n}$ and hence $m < n$ we deduce that also $(\overline{m+1}, \overline{S^{m+1}}) \in f_S^n$. The claim follows since $\overline{m+1} = k \cup \{k\}$ and

$$\overline{S^{m+1}} = \overline{S \boxtimes S^m} = S \boxtimes \overline{S^m} = S \boxtimes B$$

using that \boxtimes on lists respectively sets interacts well with string encodings. □

With these lemmas in place, we can now conclude the first part of the semantic verification.

Fact 39. *If $\text{PCP } \overline{S}$ then $Z' \models \varphi_S$.*

Proof Assuming $\text{PCP } \overline{S}$, there are $s : \mathbb{L}(\mathbb{B})$ and $n : \mathbb{N}$ with $(s, s) \in S^n$ using Lemma 37. Now to prove $Z' \models \varphi_S$ we assume $\mathcal{M} \models Z'$ and need to show $Z' \models \varphi_S$. Instantiating the leading existential quantifiers of φ_S with \overline{n} , f_S^n , \overline{S}^n , and \overline{s} leaves the following facts to verify:

- $\mathcal{M} \models \overline{n} \in \omega$, immediate by (i) of Lemma 36.
- Functionality of f_S^n , straightforward by construction of f_S^n .
- $\mathcal{M} \models f_S^n \gg \overline{n}$, immediate by Lemma 38.
- $\mathcal{M} \models (\overline{n}, \overline{S}^n) \in f_S^n$, again by construction of f_S^n .
- $\mathcal{M} \models (\overline{s}, \overline{s}) \in \overline{S}^n$, by the assumption $(s, s) \in S^n$. □

For the converse direction, we again need to restrict to models \mathcal{M} only containing standard natural numbers, i.e. satisfying that any $k \in \omega$ is the numeral $k = \overline{n}$ for some $n : \mathbb{N}$. Then the internally recognised solutions correspond to actual external solutions of PCP.

Lemma 40. *If in a standard model \mathcal{M} there is a functional approximation $f \gg k$ for $k \in \omega$ with $(k, B) \in f$, then for all $p \in B$ there are $s, t : \mathbb{L}(\mathbb{B})$ with $p = (\overline{s}, \overline{t})$ and $S \triangleright (s, t)$.*

Proof Since \mathcal{M} is standard, there is $n : \mathbb{N}$ with $k = \overline{n}$, so we have $f \gg \overline{n}$ and $(\overline{n}, B) \in f$. In any model with $f \gg \overline{n}$ we can show that $(\overline{k}, \overline{S}^k) \in f$ by induction on k , so in particular $(\overline{n}, \overline{S}^n) \in f$ in \mathcal{M} . But then by functionality of f it must be $B = \overline{S}^n$, so for any $p \in B$ we actually have $p \in \overline{S}^n$ for which it is easy to extract $s, t : \mathbb{L}(\mathbb{B})$ with $p = (\overline{s}, \overline{t})$ and $(s, t) \in S^n$. We then conclude $S \triangleright (s, t)$ with Lemma 37. □

Fact 41. *Every standard model $\mathcal{M} \models Z'$ with $\mathcal{M} \models \varphi_S$ yields $\text{PCP } \overline{S}$.*

Proof A standard model of Z' with $\mathcal{M} \models \varphi_S$ yields a functional approximation $f \gg k$ for $k \in \omega$ with some $(k, B) \in f$ and $(x, x) \in B$. Then by Lemma 40 there are $s, t : \mathbb{L}(\mathbb{B})$ with $(x, x) = (\overline{s}, \overline{t})$ and $S \triangleright (s, t)$. By the injectivity of ordered pairs and string encodings ((iv) of Lemma 36) we obtain $s = t$ and thus $S \triangleright (s, s)$. □

Finally, we just record the fact that the semantic argument in Fact 41 can be repeated deductively with an analogous intermediate structure.

Fact 42. *If $\text{PCP } \overline{S}$ then $Z' \vdash \varphi_S$.*

With the three facts verifying φ_S , we conclude reductions as follows:

Theorem 43. *We have the following reductions.*

- $\text{PCP} \preceq Z'$, provided a standard model of Z' exists.
- $\text{PCP} \preceq Z$, provided a standard model of Z exists.
- $\text{PCP} \preceq \text{ZF}$, provided a standard model of ZF exists.

18 *Synthetic Undecidability and Incompleteness of First-Order Axiom Systems in Coq*
Proof By Facts 39, 41 and 42 as well as Theorem 10. \square

In a previous paper [14] based on Aczel's sets-as-trees interpretation [21, 32, 33], we analyse assumptions necessary to obtain models of higher-order set theories in Coq's type theory. The two relevant axioms concerning the type \mathcal{T} of well-founded trees can be formulated as the extensionality of classes, i.e. unary predicates, on trees (CE), and the existence of a description operator for isomorphism classes $[t]_{\approx}$ of trees (TD):

$$\begin{aligned} \text{CE} &:= \forall(P, P' : \mathcal{T} \rightarrow \mathbb{P}). (\forall t. P t \leftrightarrow P' t) \rightarrow P = P' \\ \text{TD} &:= \exists(\delta : (\mathcal{T} \rightarrow \mathbb{P}) \rightarrow \mathcal{T}). \forall P. (\exists t. P = [t]_{\approx}) \rightarrow P(\delta P) \end{aligned}$$

Then Theorem 43 can be reformulated as follows.

Corollary 44. *Assuming CE implies both $\text{PCP} \preceq Z'$ and $\text{PCP} \preceq Z$, and assuming both CE and TD implies $\text{PCP} \preceq \text{ZF}$.*

Proof By Fact 5.4 and Theorem 5.9 of [14] CE and $\text{CE} \wedge \text{TD}$ yield models of higher-order Z and ZF set theory, respectively. It is easy to show that they are standard models and satisfy the first-order axiomatisations Z and ZF. \square

Note that assuming CE to obtain a model of higher-order Z is unnecessary if we allow the interpretation of equality by any equivalence relation congruent for membership, backed by the fully constructive model given in Theorem 4.6 of [14]. This variant is included in the [Coq development](#) but we focus on the simpler case of extensional models in this text.

By these reductions, we can conclude the incompleteness of ZF.

Theorem 45. *Assuming LEM, completeness of any extension $\mathcal{A} \supseteq Z'$ satisfied by a standard model would imply the decidability of the halting problem.*

Proof By Corollary 44, Theorem 10, Fact 9, and the reductions verified in [23]. \square

7 ZF Set Theory without Skolem Functions

We now work in the signature $\tilde{\Sigma} := (- \equiv -, - \in -)$ only containing equality and membership. To express set theory in this syntax, we [reformulate the axioms](#) specifying the Skolem symbols used in the previous signature Σ to just assert the existence of respective sets, for instance:

$$\begin{aligned} \emptyset & : & \forall x. x \notin \emptyset & \rightsquigarrow & \exists u. \forall x. x \notin u \\ \mathcal{P}(x) & : & \forall xy. y \in \mathcal{P}(x) \leftrightarrow y \subseteq x & \rightsquigarrow & \forall x. \exists u. \forall y. y \in u \leftrightarrow y \subseteq x \end{aligned}$$

In this way we obtain axiomatisations \tilde{Z}' , \tilde{Z} , and $\tilde{\text{ZF}}$ as the respective counterparts of Z' , Z, and ZF. In this section, we show that these symbol-free axiomatisations admit the same reduction from PCP.

Instead of reformulating the reduction given in the previous section to the smaller signature, which would require us to replace the natural encoding of

numbers and strings as terms by a more obscure construction, we define a general translation $\tilde{\varphi} : \mathbb{F}_{\tilde{\Sigma}}$ of formulas $\varphi : \mathbb{F}_{\Sigma}$. We then show that $\tilde{Z}' \models \tilde{\varphi}$ implies $Z' \models \varphi$ (Fact 49) and that $Z' \vdash \varphi$ implies $\tilde{Z}' \vdash \tilde{\varphi}$ (Fact 52), which is enough to deduce the undecidability of \tilde{Z}' , \tilde{Z} , and $\tilde{Z}\tilde{F}$ (Theorem 53).

The informal idea of the translation function is to replace terms $t : \mathbb{T}_{\Sigma}$ by formulas $\varphi_t : \mathbb{F}_{\tilde{\Sigma}}$ characterising the index x_0 to behave like t , for instance:

$$x_n \rightsquigarrow x_0 \equiv x_{n+1} \quad \emptyset \rightsquigarrow \forall x_0 \notin x_1 \quad \mathcal{P}(t) \rightsquigarrow \exists \varphi_t[x_0; \uparrow^2] \wedge \forall x_0 \in x_2 \leftrightarrow x_0 \subseteq x_1$$

The formula expressing $\mathcal{P}(t)$ first asserts that there is a set satisfying φ_t (where the substitution \uparrow^n shifts all indices by n) and then characterises x_0 (appearing as x_2 given the two quantifiers) as its power set. Similarly, formulas are translated by descending recursively to the atoms, which are replaced by formulas asserting the existence of characterised sets being in the expected relation, for instance:

$$t \in t' \rightsquigarrow \exists \varphi_t[x_0; \uparrow^2] \wedge \exists \varphi_{t'}[x_0; \uparrow^3] \wedge x_1 \in x_0$$

We now verify that the translation $\tilde{\varphi}$ satisfies the two desired facts, starting with the easier semantic implication. To this end, we denote by $\tilde{\mathcal{M}}$ the $\tilde{\Sigma}$ -model obtained from a Σ -model \mathcal{M} by forgetting the interpretation of the function symbols not present in $\tilde{\Sigma}$. Then for a model $\mathcal{M} \models Z'$, satisfiability is preserved for translated formulas, given that the term characterisations are uniquely satisfied over the axioms of Z' :

Lemma 46. $x = \hat{\rho}t$ iff $(x; \rho) \models_{\tilde{\mathcal{M}}} \varphi_t$ in all models $\mathcal{M} \models Z'$.

Proof By induction on t with x generalised. We only consider the cases x_n and \emptyset :

- We need to show $x = \hat{\rho}x_n$ iff $(x; \rho) \models_{\tilde{\mathcal{M}}} x_0 \equiv x_{n+1}$ which is immediate by definition.
- First assuming $x = \emptyset$, we need to show that $\forall y. y \notin x$, which is immediate since \mathcal{M} satisfies the empty set axiom. Conversely assuming $\forall y. y \notin x$ yields $x = \emptyset$ by using the extensionality axiom also satisfied by \mathcal{M} . \square

Lemma 47. $\rho \models_{\mathcal{M}} \varphi$ iff $\rho \models_{\tilde{\mathcal{M}}} \tilde{\varphi}$ in all models $\mathcal{M} \models Z'$.

Proof By induction on φ with ρ generalised, all cases but atoms are directly inductive. Considering the case $t \in t'$, we first need to show that if $\hat{\rho}t \in \hat{\rho}t'$, then there are x and x' with $x \in x'$ satisfying φ_t and $\varphi_{t'}$, respectively. By Lemma 46 the choice $x := \hat{\rho}t$ and $x' := \hat{\rho}t'$ is enough. Now conversely, if there are such x and x' , by Lemma 46 we know that $x = \hat{\rho}t$ and $x' = \hat{\rho}t'$ and thus conclude $\hat{\rho}t \in \hat{\rho}t'$. The case of $t \equiv t'$ is analogous. \square

Then the semantic implication follows since pruned models $\tilde{\mathcal{M}}$ satisfy \tilde{Z}' :

Lemma 48. If $\mathcal{M} \models Z'$ then $\tilde{\mathcal{M}} \models \tilde{Z}'$.

Proof We only need to consider the axioms concerned with set operations, where we instantiate the existential quantifiers introduced in \tilde{Z}' with the respective operations

available in \mathcal{M} . For instance, to show $\tilde{\mathcal{M}} \models \exists u. \forall x. x \notin u$ it suffices to show that $\forall x. x \notin \emptyset$ in $\tilde{\mathcal{M}}$, which is exactly the empty set axiom satisfied by \mathcal{M} . \square

Fact 49. $\tilde{Z}' \models \tilde{\varphi}$ implies $Z' \models \varphi$.

Proof Straightforward by Lemmas 47 and 48. \square

We now turn to the more involved deductive verification of the translation, beginning with the fact that \tilde{Z}' proves the unique existence of sets satisfying the term characterisations:

Lemma 50. *For all $t : \mathbb{T}$ we have $\tilde{Z}' \vdash \exists \varphi_t$ and $\tilde{Z}' \vdash \varphi_t[x] \rightarrow \varphi_t[x'] \rightarrow x \equiv x'$.*

Proof Both claims are by induction on t , the latter with x and x' generalised. The former is immediate for variables and \emptyset , we discuss the case of $\mathcal{P}(t)$. By induction we know $\tilde{Z}' \vdash \exists \varphi_t$ yielding a set x simulating t and need to show $\tilde{Z}' \vdash \exists \exists \varphi_t[x_0; \uparrow^2]$ $\wedge \forall x_0 \in x_2 \leftrightarrow x_0 \subseteq x_1$. After instantiating the first quantifier with the set u guaranteed by the existential power set axiom for the set x and the second quantifier with x itself, it remains to show $\varphi_t[x]$ and $\forall x_0 \in u \leftrightarrow x_0 \subseteq x$ which are both straightforward by the choice of x and u .

The second claim follows from extensionality given that the characterisation φ_t specifies its satisfying sets exactly by their elements. So in fact the axioms concerning the set operations are not even used in the proof of uniqueness. \square

During translation, term can be simulated by variables:

Lemma 51. *For all $\varphi : \mathbb{F}$ and $t : \mathbb{T}$ we have $\tilde{Z}' \vdash \varphi_t[x] \rightarrow (\tilde{\varphi}[x] \leftrightarrow \widetilde{\varphi}[t])$.*

Proof By induction on φ , all cases but the atoms are straightforward, relying on the fact that the syntax translation interacts well with variable renamings in the quantifier cases. The proof for atoms relies on a similar lemma for terms stating that $\varphi_s[y; x]$ and $\varphi_{s[t]}[y]$ are interchangeable whenever $\varphi_t[x]$, the rest is routine. \square

This is the main ingredient to verify the desired proof transformation:

Fact 52. $Z' \vdash \varphi$ implies $\tilde{Z}' \vdash \tilde{\varphi}$.

Proof We prove the more general claim that $\Gamma \Vdash Z' \vdash \varphi$ implies $\tilde{\Gamma} \Vdash \tilde{Z}' \vdash \tilde{\varphi}$ by induction on the first derivation. All rules but the assumption rule (A), \forall -elimination (AE), and \exists -elimination (EE) are straightforward, we explain the former two.

- If $\varphi \in \Gamma \Vdash Z'$, then either $\varphi \in \Gamma$ or $\varphi \in Z'$. In the former case we have $\tilde{\varphi} \in \tilde{\Gamma}$, so $\tilde{\Gamma} \Vdash \tilde{Z}' \vdash \tilde{\varphi}$ by (A). Regarding the latter case, we can verify $\tilde{Z}' \vdash \tilde{\varphi}$ for all $\varphi \in Z'$ by rather tedious derivations given the sheer size of some axiom translations.
- If $\Gamma \Vdash Z' \vdash \varphi[t]$ was derived from $\Gamma \Vdash Z' \vdash \forall \varphi$, then by the inductive hypothesis we know $\tilde{\Gamma} \Vdash \tilde{Z}' \vdash \forall \tilde{\varphi}$. Given Lemma 50 we may assume $\varphi_t[x]$ for a fresh variable x . Then by instantiating the inductive hypothesis to x via (AE) we obtain $\tilde{\Gamma} \Vdash \tilde{Z}' \vdash \tilde{\varphi}[x]$ and conclude the claim $\tilde{\Gamma} \Vdash \tilde{Z}' \vdash \widetilde{\varphi}[t]$ with Lemma 51. \square

Now we obtain the undecidability of the symbol-free axiomatisations.

Theorem 53. *Assuming CE implies both $\text{PCP} \preceq \tilde{Z}'$ and $\text{PCP} \preceq \tilde{Z}$, and assuming both CE and TD implies $\text{PCP} \preceq \tilde{ZF}$.*

Proof As Theorem 10, using Facts 49 and 52 and the reduction from Section 6. \square

Note that Fact 52 almost yields *deductive conservativity*, i.e. the fact that if Z' proves a symbol-free formula over $\tilde{\Sigma}$ then so does \tilde{Z}' . The missing lemma is that from \tilde{Z}' such a formula φ is provably equivalent to its translation $\tilde{\varphi}$ (after tacitly embedding φ into the full signature Σ):

Lemma 54. $\tilde{Z}' \vdash \varphi \leftrightarrow \tilde{\varphi}$ for all φ over $\tilde{\Sigma}$.

Proof By induction on φ , all composite cases are trivial. For the atom $x \in y$, we have to show its equivalence to $\exists x'. x \equiv x' \wedge \exists y'. y \equiv y' \wedge x \in y$, similarly for $x \equiv y$. \square

We can then record conservativity results as follows:

Fact 55. *If $Z'/Z/ZF$ proves a formula φ over $\tilde{\Sigma}$, then so does $\tilde{Z}'/\tilde{Z}/\tilde{ZF}$.*

Proof First let $Z' \vdash \varphi$. Then by Fact 52 we have $\tilde{Z}' \vdash \tilde{\varphi}$ and thus $\tilde{Z}' \vdash \varphi$ by Lemma 54.

If we instead suppose $Z \vdash \varphi$, we have in particular $Z' \vdash \Gamma \vdash \varphi$, where Γ contains finitely many instances of the separation scheme. Then by the generalised goal used in the proof of Lemma 54 also $\tilde{Z}' \vdash \tilde{\Gamma} \vdash \tilde{\varphi}$ and therefore $\tilde{Z}' \vdash \tilde{\Gamma} \vdash \varphi$ again using Lemma 54. We hence conclude $\tilde{Z} \vdash \varphi$ since every translated instance of separation for a formula ψ can be proved from the respective instance for $\tilde{\psi}$ available in \tilde{Z} .

The case for ZF is analogous by further decomposing into the finitely many used instances of the replacement scheme. \square

For the sake of completeness, we also establish the converse directions. To this end, we first verify a deductive counterpart of Lemma 48:

Lemma 56. $Z' \vdash \tilde{Z}'$, i.e. Z' proves every axiom from \tilde{Z}' (embedded into Σ).

Proof By instantiating every existentially formulated axiom from \tilde{Z}' with the respective symbol available in Z' . \square

Fact 57. *If $\tilde{Z}'/\tilde{Z}/\tilde{ZF}$ proves a formula φ over $\tilde{\Sigma}$, then so does $Z'/Z/ZF$.*

Proof If $\tilde{Z}' \vdash \varphi$, we obtain the same deduction if we consider both \tilde{Z}' and φ embedded into the full signature. Then by Lemma 56 we can conclude that $Z' \vdash \varphi$.

The respective results for \tilde{Z} and \tilde{ZF} follow by similar decompositions regarding the axiom schemes as used in the proof of Fact 55. \square

Note that in the absence of unique choice there is no direct proof for *semantic conservativity*, i.e. the fact that if Z' validates a symbol-free formula over $\tilde{\Sigma}$ then so does \tilde{Z}' , since this would involve constructing a Σ -model from a $\tilde{\Sigma}$ -model only existentially exhibiting the set operations.

We conclude this section with a brief observation concerning the further reduced signature $\tilde{\Sigma} := (- \in -)$, full detail can be found in the [Coq development](#). Since equality is expressible in terms of membership by $x \equiv y := \forall z. x \in z \leftrightarrow y \in z$, we can rephrase the above translation to yield formulas $\tilde{\varphi} : \mathbb{F}_{\tilde{\Sigma}}$

satisfying the same properties as stated in Facts 49 and 52 for a corresponding axiomatisation \check{Z}' . Moreover, since \check{Z}' does not refer to primitive equality, we can freely interpret it with the fully constructive model given in Theorem 4.6 of [14] and therefore obtain $\text{PCP} \preceq \check{Z}'$ without assumptions. This allows us to deduce the undecidability of the Entscheidungsproblem in its sharpest possible form:

Theorem 58. *FOL with a single binary relation symbol is undecidable.*

Proof By Fact 11 and the reduction $\text{PCP} \preceq \check{Z}'$. □

8 Finitary Set Theories

In this section, we consider various finitary set theories, i.e. axiomatisations of set theory that do not guarantee infinite sets or do even refute their existence. Given our setting, the undecidability and incompleteness of such systems can be established either by indirectly reducing from set theories such as Z' or by modifying the direct reduction function $\text{PCP} \preceq Z'$. We discuss both of these strategies where applicable.

A first way to axiomatise finite set theory is to work in the full signature used in Section 6 and simply leave the set ω unspecified. Then on top, one can add an axiom ruling out any inductive sets like ω , i.e. sets containing \emptyset and being closed under successors $x \cup \{x\}$.

- FZ' denotes Z' without the axioms specifying ω as the least inductive set.
- $\text{FZ}' + \neg\text{Inf}$ denotes FZ' plus the axiom that no set is inductive.

That FZ' as a mere subset of Z' is undecidable follows immediately by Fact 12:

Fact 59. $Z' \preceq \text{FZ}'$ and therefore, provided CE, also $\text{PCP} \preceq \text{FZ}'$.

Proof By (2) of Fact 12 and Corollary 44. □

However, this direct result is unsatisfactory by the reliance on the extensional standard model \mathcal{T} of Z' requiring CE and containing infinite sets. So in order to show $\text{FZ}' + \neg\text{Inf}$ undecidable and dispense with CE, we have to rework the reduction $\text{PCP} \preceq Z'$ from Section 6 to avoid mention of ω such that the constructive model of hereditarily finite sets [34] can be employed.

In this model, the numerals are exactly the hereditarily transitive sets (i.e. sets x that are transitive, meaning $y \subseteq x$ for all $y \in x$, and every element of x is transitive, written $\text{HT}(x)$), allowing us to modify the reduction formula φ_S given a PCP-instance as follows:

$$\varphi_S := \exists k, f, B, x. k \in \omega \wedge f \gg k \wedge \dots \quad \Rightarrow \quad \psi_S := \exists k, f, B, x. \text{HT}(k) \wedge f \gg k \wedge \dots$$

Note that the bound $k \in \omega$ was only used to express that k is a natural number such that (at least in standard models) the approximation $f \gg k$ corresponds to a faithful accumulation of PCP-solutions. This bound can be replaced by any defining property of numerals in the intended model and in the present case, $\text{HT}(x)$ is particularly easy to express.

By according modification of the proofs for φ_S we can verify the new reduction ψ_S with respect to all standard models, i.e. models where every hereditarily transitive set is a numeral:

Lemma 60. *The following facts about ψ_S hold:*

1. If PCP S then $\text{FZ}' \vDash \psi_S$.
2. Every standard model $\mathcal{M} \vDash \text{FZ}'$ with $\mathcal{M} \vDash \psi_S$ yields PCP S .
3. If PCP S then $\text{FZ}' \vdash \psi_S$.

Proof Analogous to Facts 39, 41 and 42, using the fact that $\text{HT}(\bar{n})$ for all $n : \mathbb{N}$. \square

Following the construction from [34], adopted more recently for [35], a model \mathcal{T}_2 of FZ' can be obtained by taking the inductive type of binary trees quotiented by tree equivalence and implementing the set operations by suitable tree manipulations. In particular, this model is standard in the above sense and does not contain inductive sets:

Lemma 61. *\mathcal{T}_2 is a standard model of $\text{FZ}' + \neg\text{Inf}$.*

Proof To establish that \mathcal{T}_2 is standard, we show that for every $x : \mathcal{T}_2$ we can compute a number $n_x : \mathbb{N}$ such that $x = \bar{n}_x$. By induction on the well-foundedness of x we may assume that every element $y \in x$ is a numeral \bar{n}_y . Since x is finite, we can compute a bound n such that $n_y < n$ for all $y \in x$. Then we can obtain that x is a numeral (and in fact compute n_x) since x is a transitive subset of the numeral \bar{n} by induction on n .

Regarding the second claim, suppose x were inductive. By finiteness of x we obtain the cardinality N of distinct elements in x . But since x is inductive, it must contain the set of the first $N + 1$ numerals that are distinct by construction, yielding a contradiction. \square

So we can conclude the undecidability of FZ' and $\text{FZ}' + \neg\text{Inf}$ as usual:

Theorem 62. $\text{PCP} \preceq \text{FZ}'$ and $\text{PCP} \preceq \text{FZ}' + \neg\text{Inf}$.

Proof By applying Theorem 10 to Lemmas 60 and 61. \square

An alternative, more incisive formulation of finitary set theory just axiomatises the empty set in addition to the adjunction operation $\{x\} \cup y$ (usually definable from union and pairing) [36], i.e. we work in the signature

$$\Sigma_{\text{PS}} := (\emptyset, \dots; _ \equiv _, _ \in _)$$

where the term $x.y$ is enforced to behave like $\{x\} \cup y$ by the axiom

$$\forall z. z \in x.y \leftrightarrow z \equiv x \vee z \in y.$$

Moreover, to rule out infinite sets, one can require an induction scheme on top:

$$\lambda\varphi. \varphi[\emptyset] \rightarrow (\forall xy. \varphi[x] \rightarrow \varphi[y] \rightarrow \varphi[x.y]) \rightarrow \forall x. \varphi[x]$$

- **PS** denotes the axioms characterising \emptyset and $x.y$ as well as extensionality.
- **PS + Ind** denotes PS plus all instances of the induction scheme.

We again begin with the indirect argument to establish undecidability of the core axiomatisation PS still compatible with Z' . First note that, while the usual ZF-operations can define adjunction, the converse does not hold since the ZF-operations are strictly stronger on infinite models. We can therefore not directly translate formulas in the ZF-signature to the new signature Σ_{PS} . Instead, the translation has to go through the function-free signature $\tilde{\Sigma} := (- \equiv -, - \in -)$ used in Section 7, reusing the verified translation $\tilde{\varphi}$.

Fact 63. $\text{PCP} \preceq \text{PS}$

Proof We use the reduction formula $\varphi_S^{\text{PS}} := \bigwedge \tilde{Z}' \rightarrow \tilde{\varphi}_S$ tacitly embedding the translated formulas from \tilde{Z}' and $\tilde{\varphi}_S$ in $\tilde{\Sigma}$ into the signature Σ_{PS} . Then the sufficient facts are that $\text{PCP } S$ implies $\text{PS} \vdash \varphi_S^{\text{PS}}$ and, conversely, that $\text{PS} \vDash \varphi_S^{\text{PS}}$ implies $\text{PCP } S$.

Regarding the former, from $\text{PCP } S$ we obtain $\tilde{Z}' \vdash \tilde{\varphi}_S$ from Facts 42 and 52. So in particular $\vdash \bigwedge \tilde{Z}' \rightarrow \tilde{\varphi}_S$ and by weakening (and correctness of the tacit embedding) $\text{PS} \vdash \varphi_S^{\text{PS}}$.

Regarding the latter, suppose $\text{PS} \vDash \varphi_S^{\text{PS}}$. The (intensional) standard model \mathcal{T} from Fact 39 interprets the full ZF-signature, so in particular Σ_{PS} and the axioms of PS. We therefore obtain that $\mathcal{T} \vDash \varphi_S^{\text{PS}}$. Then by Lemmas 47 and 48 we can deduce that \mathcal{T} (now equipped with the full ZF-structure again) satisfies φ_S and conclude $\text{PCP } S$ with Fact 41. \square

As with Fact 59 before, this indirect method does not extend to the axiomatisation $\text{PS} + \text{Ind}$, which is not satisfied by the standard model \mathcal{T} . We therefore sketch the direct reduction from PCP obtained by further modifying the formula ψ_S , full detail is given in the [Coq formalisation](#).

First, the encodings of numbers and strings is mostly unaffected since the adjunction operation is exactly the natural successor function and can define unordered pairs $\{x, y\}$ by $x.y.\emptyset$, from which we obtained the ordered pairs used for strings. Secondly, the only other usage of a ZF-function in ψ_S is the (binary) union used to implement the operation $S \boxtimes B$ recursively, which can be replaced by any set enforced to behave accordingly. Thus we obtain a formula ψ_S^{PS} in the signature Σ_{PS} that we can verify to capture PCP as usual:

Lemma 64. *The following facts about ψ_S^{PS} hold:*

1. If $\text{PCP } S$ then $\text{PS} \vDash \psi_S^{\text{PS}}$.
2. Every standard model $\mathcal{M} \vDash \text{PS}$ with $\mathcal{M} \vDash \psi_S^{\text{PS}}$ yields $\text{PCP } S$.
3. If $\text{PCP } S$ then $\text{PS} \vdash \psi_S^{\text{PS}}$.

Proof Analogous to Lemma 60 with the expectable differences regarding the altered data encodings and the elimination of binary unions. \square

Lemma 65. \mathcal{T}_2 is a standard model of $\text{PS} + \text{Ind}$.

Proof That \mathcal{T}_2 is standard was already part of Lemma 61 and that it models PS was shown in [34]. They also established the higher-order induction principle

$$\forall P : \mathcal{T}_2 \rightarrow \mathbb{P}. P \emptyset \rightarrow (\forall xy. P x \rightarrow P y \rightarrow P(x.y)) \rightarrow \forall x. P x$$

which is easily seen to entail the first-order induction scheme. □

Theorem 66. $\text{PCP} \preceq \text{PS}$ and $\text{PCP} \preceq \text{PS} + \text{Ind}$.

Proof By applying Theorem 10 to Lemmas 64 and 65. □

We conclude with a formulation of PS in the binary signature $\check{\Sigma} := (- \in -)$ introduced in Section 7. As done with Z' to obtain \check{Z}' , we can replace the two axioms from PS specifying \emptyset and $x.y$ by existentially quantified versions, express equality via membership, and hence obtain the axiomatisation $\check{\text{PS}}$ over $\check{\Sigma}$. This is a particularly compact system showing a single binary relation symbol undecidable, by virtue of the following reduction:

Fact 67. $\check{Z}' \preceq \check{\text{PS}}$ and thus also $\text{PCP} \preceq \check{\text{PS}}$.

Proof To obtain $\check{Z}' \preceq \check{\text{PS}}$ we use (1) of Fact 12, so we have to show $\check{Z}' \vdash \check{\text{PS}}$. The only axiom of $\check{\text{PS}}$ not already present in \check{Z}' is the existential specification of adjunction, which can be established by the existential specification of union and pairing available in \check{Z}' . The full reduction $\text{PCP} \preceq \check{\text{PS}}$ is obtained by composition with the reduction $\text{PCP} \preceq \check{Z}'$ underlying Theorem 58. □

9 Abstract Undecidability and Incompleteness

We conclude the technical part of this paper by isolating the synthetic arguments underlying Fact 9 and Theorem 10, abstracting from the concrete formalism of FOL. This abstraction is in the spirit of Popescu and Traytel's [37] analysis of the abstract preconditions for Gödel's two incompleteness theorems. Given our computational approach, much less internal structure like substitution or numerals needs to be assumed, at the cost of essential incompleteness and Gödel's second incompleteness theorem remaining out of reach.

Overwriting all notation from before, our base setup is to assume an arbitrary discrete type \mathbb{F} representing formulas as well as an enumerable predicate $\lambda\varphi : \mathbb{F}. \vdash \varphi$ considered the provable formulas. We do not have to commit to \mathbb{F} only containing a specific sort of formulas (e.g. the closed formulas) or to \vdash being defined over a particular context (e.g. an axiomatisation of arithmetic) or coming in a specific flavour (e.g. intuitionistic or classical).

If we add a reasonably well-behaved negation operation, we obtain an abstract version of the fact that negation-completeness implies decidability:

Fact 68. We assume a negation operation $\neg : \mathbb{F} \rightarrow \mathbb{F}$ as follows:

- *Discriminability:* given φ it is decidable if φ is a negation $\neg\psi$ for some ψ .
- *Injectivity:* we have $\varphi = \psi$ whenever $\neg\varphi = \neg\psi$.
- *Consistency:* there is no φ with both $\vdash \varphi$ and $\vdash \neg\varphi$.

Then if \vdash is complete (i.e. either $\vdash \varphi$ or $\vdash \neg\varphi$ for all φ), then it is decidable.

Proof As in the proof of Fact 9 we use Post's theorem, leaving us to show logical decidability and co-enumerability of provability (given enumerability by assumption):

- Given φ , to (logically) decide whether $\vdash \varphi$ or $\not\vdash \varphi$ is the case, we analyse completeness for φ . In the non-trivial case where $\vdash \neg\varphi$ we obtain $\not\vdash \varphi$ by consistency.
- For co-enumerability, by completeness and consistency it suffices to enumerate $\lambda\varphi. \vdash \neg\varphi$ instead of $\lambda\varphi. \not\vdash \varphi$. This is obtained by the enumerator of \vdash , using discriminability to check for each φ if it is a negation, and injectivity for the correctness proof. \square

If instead of a negation operation we add an abstract notion of (standard) models, we obtain an abstract undecidability result analogous to Theorem 10:

Fact 69. *We assume a type \mathbb{M} of models together with the following data:*

- *Satisfaction: a relation $\mathcal{M} \models \varphi$ inducing validity $\models \varphi$ as $\mathcal{M} \models \varphi$ for all \mathcal{M} .*
- *Soundness: all provable formulas are valid, i.e. $\vdash \varphi$ implies $\models \varphi$.*
- *Standardness: a predicate $S : \mathbb{M} \rightarrow \mathbb{P}$ with at least one standard model $S \mathcal{M}$.*

If we further assume $P : X \rightarrow \mathbb{P}$ and $F : X \rightarrow \mathbb{F}$ satisfying

- *Whenever $P x$ holds, we have a derivation $\vdash F x$, and*
- *Whenever $\mathcal{M} \models F x$ in a standard model $S \mathcal{M}$, we obtain $P x$,*

then the function F induces reductions $P \preceq (\lambda\varphi. \vdash \varphi)$ and $P \preceq (\lambda\varphi. \models \varphi)$.

Proof The assumed standard model justifies that $P x$ whenever $\models F x$. We hence obtain the two reductions, with soundness used for the missing directions. \square

Note that if we extend the setting of Fact 69 with the negation operation from Fact 68, we arrive at the conclusion that completeness of \vdash would entail the decidability of P .

It is easy to [instantiate](#) Fact 69 to obtain Theorem 10 concerning first-order axiomatisations \mathcal{B} . We simply let \mathbb{F} be the first-order formulas, \vdash the formulas (intuitionistically) provable from \mathcal{B} , and \mathbb{M} be the type of first-order models \mathcal{M} with environments ρ such that $\rho \models \mathcal{B}$. Then the remaining assumptions of Theorem 10 imply the assumptions of Fact 69.

Slightly more involved (at least on mechanisation level) is the [instantiation](#) of Fact 69 to Fact 9, since this time we pick \mathbb{F} as the type of closed first-order formulas, to which we have to adopt the negation operation and the (classical) deduction system as well as the discreteness and enumerability proofs for arbitrary formulas.

Although these comments only show the applicability of our abstract analysis to the case of first-order logic as examined in this paper, we remark that Facts 68 and 69 could as well be instantiated to extended formalisms such as second- or higher-order logic, or systems based on completely different primitives such as dependent type theories.

10 Discussion

10.1 General Remarks

In this paper, we have described a synthetic approach to the formalisation and mechanisation of undecidability and incompleteness results in first-order logic. The general approach was then instantiated to case-studies concerned with arithmetical theories in the family of PA as the typical systems considered in the investigation of incompleteness, and with various formulations of set theory as one of the standard foundations of mathematics. The chosen strategy complements the considerably harder to mechanise proofs relying on Gödel sentences, and for ZF the choice of PCP as seed problem instead of H_{10} or PA itself is a slight simplification since only a single recursion needs to be simulated. We use this section for some additional remarks based on the helpful feedback by the anonymous reviewers.

As formally stated in Definition 8, we only consider incompleteness as a property of the *classical* deduction system. This is simply owing to the fact that much of the literature on incompleteness seems focused on classical logic, with a notable exception of the more agnostic treatment in [37]. Although perhaps weaker in general, incompleteness of the *intuitionistic* deduction system can also be considered a meaningful property and follows in an analogous way. Concretely, a corresponding version of Fact 9 holds for the intuitionistic notion, yielding variants of Theorems 27 and 45 provable without LEM.

In alignment with [8] but in contrast to [24], we define semantic entailment $\mathcal{T} \models \varphi$ without restricting to *classical models*, i.e. models that satisfy all first-order instances of LEM. In our constructive meta-theory this relaxation is necessary to be able to use the standard models of PA and ZF, which would only be classical in a classical meta-theory. Leaving $\mathcal{T} \models \varphi$ in this sense constructively underspecified seems like a reasonable trade for a more economical usage of LEM.

Similarly, we leave it underspecified whether PA and ZF are seen as classical theories or their intuitionistic counterparts, namely Heyting arithmetic and a variant of intuitionistic set theory, respectively. By the choice not to distinguish these explicitly by LEM as a first-order axiom scheme, we leave it to the deduction system to discriminate between both views while the Tarski-style semantics leans towards the classical interpretation (especially in the presence of LEM). For simplicity, we decided to only speak of PA and ZF in the main body of the text, especially since a discussion of intuitionistic set theories would involve choosing a particular system. While IZF is an extension of Z' close to ZF with collection instead of replacement, the more predicative CZF does not have power sets as included in Z' .

10.2 Coq Mechanisation

Our axiom-free [mechanisation](#) contributes about 10k lines of code (loc) to the Coq Library of Undecidability Proofs [9], on top of about 1500loc that could

be reused from previous developments [14, 24]. Remarkably, the axiomatisation, undecidability, and incompleteness of PA add up to only 800loc, while already the initial reduction from PCP to ZF in the skolemised signature is above 1800loc. The remaining development is mostly concerned with the signature reduction for ZF (2500loc) and the material on finitary set theories (3000loc). Both contain files with very similar proofs, especially the reduction files for Z' and FZ' are nearly identical and therefore it should be possible to reduce the development size by reorganisation (at the cost of a less transparent presentation). The abstract development outlined in Section 9 is below 300loc, including the instantiation to FOL.

Our mechanisation of first-order logic unifies ideas from previous versions [8, 10, 24] and is general enough to be reused in other use cases. Notably, we refrained from including equality as a syntactic primitive to treat both intensional and extensional interpretations without changing the underlying signature. On the other hand, with primitive equality, the extensionality of models would hold definitionally and the deduction system could be extended with the Leibniz rule, making the additional axiomatisation of equality obsolete.

Furthermore, manipulating deductive goals of the form $\Gamma \vdash \varphi$ benefitted a lot from custom tactics, mostly to handle substitution and the quantifier rules. The former tactics approximate the automation provided by the Auto-subst 2 framework unfortunately relying on functional extensionality [38] and the latter are based on the named reformulations of (AI) and (EE) given in Section 2.3. We are currently working on a more scalable proof mode for deductive goals including a HOAS input language hiding de Bruijn encodings [39], implementing a two-level approach in comparison to the one-level compromise proposed by Laurent [40].

10.3 Related Work

We report on other mechanisations concerned with incompleteness and undecidability results in first-order logic. Regarding the former, a fully mechanised proof of Gödel's first incompleteness theorem was first given by Shankar [41] using the Nqthm prover. O'Connor [12] implements the same result fully constructively in Coq, and Paulson [13] provides an Isabelle/HOL mechanisation of both incompleteness theorems using the theory of hereditarily finite sets instead of a fragment of PA. Moreover, there are several partial mechanisations [42–44], and Popescu and Traytel [37] investigate the abstract preconditions of the incompleteness theorems using Isabelle/HOL. With the independence of the continuum hypothesis, Han and van Doorn [45] mechanise a specific instance of incompleteness for ZF in Lean. None of these mechanisations approach incompleteness via undecidability.

Turning to undecidability results, Forster, Kirst, and Smolka [8] mechanise the undecidability of the Entscheidungsproblem in Coq, using a convenient signature to encode PCP, and Kirst and Larchey-Wendling [10] give a Coq mechanisation of Trakhtenbrot's theorem [46], stating the undecidability of

finite satisfiability. They also begin with a custom signature for the encoding of PCP but provide the transformations necessary to obtain the undecidability result for the small signature containing a single binary relation symbol. We are not aware of any previous mechanisations of the undecidability of PA or ZF.

10.4 Future Work

There are two ways how our incompleteness results (Theorems 27 and 45) could be strengthened. First, while we were able to eliminate the use of LEM in the case of PA (Section 5), it is unclear whether the same technique applies to the concrete reduction formulas used for ZF and the related systems. It might be necessary to reformulate (and streamline) the reduction to make the argument feasible for mechanisation. Secondly, that supposed negation-completeness only implies synthetic decidability of a halting problem instead of a provable contradiction could be sharpened by extracting all reduction functions to a concrete model of computation like the weak call-by-value λ -calculus L [28]. Then the actual contradiction of an L-decider for L-halting could be derived.

We plan to continue the work on PA with a constructive analysis of Tennenbaum's theorem [47], stating that no computable non-standard model of PA exists. Translated to the synthetic setting where all functions are computable by construction, this would mean that no non-standard model of PA can be defined in Coq's type theory as long as function symbols are interpreted with type-theoretic functions. It would be interesting to investigate which assumptions of synthetic computability [7] are necessary to derive this observation as an actual theorem inside of Coq.

Complementing Theorem 58 and Fact 67, it would be interesting to find a more elementary characterisation of an undecidable binary relation usable for the sharp formulations of the Entscheidungsproblem and Trakhtenbrot's theorem. This might well work without an intermediate axiomatisation of set theory and express an undecidable decision problem more directly.

Regarding the signature translations and conservativity results for ZF discussed in Section 7, it should be possible to mechanise similar results for arbitrary axiom systems with definable extensions. Results like these would path the way for an abstract mechanisation of undecidable theories as outlined by Tarski [48].

Finally, we plan to mechanise similar undecidability and incompleteness results for second-order logic. Since second-order PA is categorical, in particular the incompleteness of any sound and enumerable deduction system for second-order logic would then follow easily.

Acknowledgements. The authors want to thank Andrej Dudenhefner, Yannick Forster, Lennard Gäher, Julian Rosemann, Gert Smolka, and the anonymous reviewers for helpful comments and suggestions.

References

- [1] Hilbert, D., Ackermann, W.: *Grundzüge der Theoretischen Logik*. Springer, Berlin, Heidelberg (1928)
- [2] Turing, A.M.: On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London mathematical society* **2**(1), 230–265 (1937)
- [3] Church, A., *et al.*: A note on the Entscheidungsproblem. *J. Symb. Log.* **1**(1), 40–41 (1936)
- [4] Presburger, M.z., Jabcquette, D.: On the completeness of a certain system of arithmetic of whole numbers in which addition occurs as the only operation. *History and Philosophy of Logic* **12**(2), 225–233 (1991)
- [5] Doner, J., Hodges, W.: Alfred Tarski and decidable theories. *The Journal of symbolic logic* **53**(1), 20–35 (1988)
- [6] Team, T.C.D.: The Coq Proof Assistant, version 8.12.0. Zenodo (2020). <https://doi.org/10.5281/zenodo.4021912>. <https://doi.org/10.5281/zenodo.4021912>
- [7] Bauer, A.: First steps in synthetic computability theory. *Electronic Notes in Theoretical Computer Science* **155**, 5–31 (2006)
- [8] Forster, Y., Kirst, D., Smolka, G.: On synthetic undecidability in coq, with an application to the entscheidungsproblem. In: *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs* (2019)
- [9] Forster, Y., Larchey-Wendling, D., Dudenhefner, A., Heiter, E., Kirst, D., Kunze, F., Smolka, G., Spies, S., Wehr, D., Wuttke, M.: A Coq library of undecidable problems. In: *CoqPL 2020, New Orleans, LA, United States* (2020). <https://github.com/uds-psl/coq-library-undecidability>
- [10] Kirst, D., Larchey-Wendling, D.: Trakhtenbrot’s theorem in Coq: a constructive approach to finite model theory. In: *International Joint Conference on Automated Reasoning (IJCAR 2020), Paris, France*. Springer, Paris, France (2020)
- [11] Post, E.L.: Recursively enumerable sets of positive integers and their decision problems. *bulletin of the American Mathematical Society* **50**(5), 284–316 (1944)
- [12] O’Connor, R.: Essential incompleteness of arithmetic verified by Coq. In: Hurd, J., Melham, T. (eds.) *Theorem Proving in Higher Order Logics*,

- pp. 245–260. Springer, Berlin, Heidelberg (2005)
- [13] Paulson, L.C.: A mechanised proof of Gödel’s incompleteness theorems using Nominal Isabelle. *Journal of Automated Reasoning* **55**(1), 1–37 (2015)
 - [14] Kirst, D., Smolka, G.: Large model constructions for second-order ZF in dependent type theory. *Certified Programs and Proofs - 7th International Conference, CPP 2018, Los Angeles, USA, 2018* (2018)
 - [15] Kirst, D., Hermes, M.: Synthetic undecidability and incompleteness of first-order axiom systems in coq. In: *12th International Conference on Interactive Theorem Proving (ITP 2021)* (2021). Schloss Dagstuhl-Leibniz-Zentrum für Informatik
 - [16] Richman, F.: Church’s thesis without tears. *The Journal of symbolic logic* **48**(3), 797–803 (1983)
 - [17] Sozeau, M., Anand, A., Boulier, S., Cohen, C., Forster, Y., Kunze, F., Malecha, G., Tabareau, N., Winterhalter, T.: The MetaCoq Project. *Journal of Automated Reasoning* (2020)
 - [18] Braibant, T., Pous, D.: An efficient Coq tactic for deciding Kleene algebras. In: *International Conference on Interactive Theorem Proving*, pp. 163–178. Springer, Berlin, Heidelberg (2010)
 - [19] Maksimović, P., Schmitt, A.: HOCore in Coq. In: *International Conference on Interactive Theorem Proving*, pp. 278–293. Springer, International (2015)
 - [20] Schäfer, S., Smolka, G., Tebbi, T.: Completeness and decidability of de Bruijn substitution algebra in Coq. In: *Proceedings of the 2015 Conference on Certified Programs and Proofs*, pp. 67–73. ACM, New York, NY, USA (2015)
 - [21] Werner, B.: Sets in types, types in sets. In: *Theoretical Aspects of Computer Software*, pp. 530–546. Springer, Berlin, Heidelberg (1997)
 - [22] Larchey-Wendling, D., Forster, Y.: Hilbert’s tenth problem in Coq. In: *4th International Conference on Formal Structures for Computation and Deduction. LIPIcs*, vol. 131, pp. 27–12720 (2019)
 - [23] Forster, Y., Heiter, E., Smolka, G.: Verification of PCP-related computational reductions in Coq. In: *International Conference on Interactive Theorem Proving*, pp. 253–269 (2018). Springer
 - [24] Forster, Y., Kirst, D., Wehr, D.: Completeness theorems for first-order

- logic analysed in constructive type theory: Extended version. *Journal of Logic and Computation* **31**(1), 112–151 (2021)
- [25] de Bruijn, N.G.: Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae (Proceedings)* **75**(5), 381–392 (1972)
- [26] Kreisel, G.: Church’s thesis: a kind of reducibility axiom for constructive mathematics. In: *Studies in Logic and the Foundations of Mathematics* vol. 60, pp. 121–150 (1970)
- [27] Forster, Y.: Church’s Thesis and related axioms in Coq’s type theory. In: Baier, C., Goubault-Larrecq, J. (eds.) *29th EACSL Annual Conference on Computer Science Logic (CSL 2021)*. *LIPIcs*, vol. 183, pp. 21–12119. Dagstuhl, Germany (2021)
- [28] Forster, Y., Kunze, F.: A Certifying Extraction with Time Bounds from Coq to Call-By-Value Lambda Calculus. In: Harrison, J., O’Leary, J., Tolmach, A. (eds.) *10th International Conference on Interactive Theorem Proving*. *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 141, pp. 17–11719. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2019). <https://doi.org/10.4230/LIPIcs.ITP.2019.17>. <http://drops.dagstuhl.de/opus/volltexte/2019/11072>
- [29] Friedman, H.: Classically and intuitionistically provably recursive functions. In: *Higher Set Theory*, pp. 21–27. Springer, Berlin, Heidelberg (1978)
- [30] Smullyan, R.M., Fitting, M.: *Set Theory and the Continuum Problem*. Dover Publications, Mineola, N.Y. (2010)
- [31] Myhill, J.: Some properties of intuitionistic Zermelo-Frankel set theory. In: *Cambridge Summer School in Mathematical Logic*, pp. 206–231. Springer, Berlin, Heidelberg (1973)
- [32] Aczel, P.: The type theoretic interpretation of constructive set theory. In: *Studies in Logic and the Foundations of Mathematics* vol. 96, pp. 55–66 (1978)
- [33] Barras, B.: Sets in Coq, Coq in sets. *Journal of Formalized Reasoning* **3**(1), 29–48 (2010)
- [34] Smolka, G., Stark, K.: Hereditarily finite sets in constructive type theory. In: *Interactive Theorem Proving - 7th International Conference, ITP 2016, Nancy, France, August 22-27, 2016*. *LNCS*, vol. 9807, pp. 374–390. Springer, Cham (2016)

- [35] Kirst, D., Larchey-Wendling, D.: Trakhtenbrot’s theorem in coq. In: International Joint Conference on Automated Reasoning, pp. 79–96 (2020). Springer
- [36] Kirby, L.: Finitary set theory. *Notre Dame Journal of Formal Logic* **50**(3), 227–244 (2009)
- [37] Popescu, A., Traytel, D.: A formally verified abstract account of Gödel’s incompleteness theorems. In: International Conference on Automated Deduction, pp. 442–461 (2019). Springer
- [38] Stark, K., Schäfer, S., Kaiser, J.: Autosubst 2: reasoning with multi-sorted de Bruijn terms and vector substitutions. In: International Conference on Certified Programs and Proofs, pp. 166–180 (2019). ACM
- [39] Hostert, J., Koch, M., Kirst, D.: A toolbox for mechanised first-order logic. In: *Coq Workshop*, vol. 2021 (2021)
- [40] Laurent, O.: An anti-locally-nameless approach to formalizing quantifiers. In: *Proceedings of the 10th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pp. 300–312 (2021)
- [41] Shankar, N.: *Proof-checking Metamathematics*, The University of Texas at Austin (1986). PhD Thesis
- [42] Quaipe, A.: Automated proofs of Löb’s theorem and Gödel’s two incompleteness theorems. *Journal of Automated Reasoning* **4**(2), 219–231 (1988)
- [43] Bundy, A., Giunchiglia, F., Villafiorita, A., Walsh, T.: *An incompleteness theorem via abstraction*. Technical report (1996)
- [44] Sieg, W., Field, C.: Automated search for Gödel’s proofs. In: *Deduction, Computation, Experiment*, pp. 117–140. Springer, ??? (2008)
- [45] Han, J., van Doorn, F.: A formal proof of the independence of the continuum hypothesis. In: *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pp. 353–366 (2020)
- [46] Trakhtenbrot, B.A.: The impossibility of an algorithm for the decidability problem on finite classes. *Dokl. Akad. Nok. SSSR* **70**(4), 569–572 (1950)
- [47] Tennenbaum, S.: Non-Archimedean models for arithmetic. *Notices of the American Mathematical Society* **6**(270), 44 (1959)
- [48] Tarski, A.: I: A general method in proofs of undecidability. In: Tarski, A. (ed.) *Undecidable Theories*. *Studies in Logic and the Foundations of Mathematics*, vol. 13, pp. 1–34 (1953)

A Deduction Systems

Intuitionistic natural deduction $\Gamma \vdash_i \varphi$ is defined by the following rules:

$$\begin{array}{c}
 \frac{\varphi \in \Gamma}{\Gamma \vdash \varphi} \text{ C} \quad \frac{\Gamma \vdash \perp}{\Gamma \vdash \varphi} \text{ E} \quad \frac{\Gamma, \varphi \vdash \psi}{\Gamma \vdash \varphi \rightarrow \psi} \text{ II} \quad \frac{\Gamma \vdash \varphi \rightarrow \psi \quad \Gamma \vdash \varphi}{\Gamma \vdash \psi} \text{ IE} \\
 \\
 \frac{\Gamma \vdash \varphi \quad \Gamma \vdash \psi}{\Gamma \vdash \varphi \wedge \psi} \text{ CI} \quad \frac{\Gamma \vdash \varphi \wedge \psi}{\Gamma \vdash \varphi} \text{ CE}_1 \quad \frac{\Gamma \vdash \varphi \wedge \psi}{\Gamma \vdash \psi} \text{ CE}_2 \\
 \\
 \frac{\Gamma \vdash \varphi}{\Gamma \vdash \varphi \vee \psi} \text{ DI}_1 \quad \frac{\Gamma \vdash \psi}{\Gamma \vdash \varphi \vee \psi} \text{ DI}_2 \quad \frac{\Gamma \vdash \varphi \vee \psi \quad \Gamma, \varphi \vdash \theta \quad \Gamma, \psi \vdash \theta}{\Gamma \vdash \theta} \text{ DE} \\
 \\
 \frac{\Gamma[\uparrow] \vdash \varphi}{\Gamma \vdash \forall \varphi} \text{ AI} \quad \frac{\Gamma \vdash \forall \varphi}{\Gamma \vdash \varphi[t]} \text{ AE} \quad \frac{\Gamma \vdash \varphi[t]}{\Gamma \vdash \exists \varphi} \text{ EI} \quad \frac{\Gamma \vdash \exists \varphi \quad \Gamma[\uparrow], \varphi \vdash \psi[\uparrow]}{\Gamma \vdash \psi} \text{ EE}
 \end{array}$$

The classical variant $\Gamma \vdash_c \varphi$ adds the Peirce rule $((\varphi \rightarrow \psi) \rightarrow \varphi) \rightarrow \varphi$.

B Axioms of Set Theory

We list the ZF axioms over $\Sigma := (\emptyset, \{-, -\}, \bigcup, \mathcal{P}(-), \omega; _ \equiv _, - \in _)$:

Structural axioms

Extensionality: $\forall xy. x \subseteq y \rightarrow y \subseteq x \rightarrow x \equiv y$

Set operations

Empty set: $\forall x. x \notin \emptyset$

Unordered pair: $\forall xyz. z \in \{x, y\} \leftrightarrow x \equiv y \vee x \equiv z$

Union: $\forall xy. y \in \bigcup x \leftrightarrow \exists z \in x. y \in z$

Power set: $\forall xy. y \in \mathcal{P}(x) \leftrightarrow y \subseteq x$

Infinity: $(\emptyset \in \omega \wedge \forall x. x \in \omega \rightarrow x \cup \{x\} \in \omega) \wedge (\forall y. (\emptyset \in y \wedge \forall x. x \in y \rightarrow x \cup \{x\} \in y) \rightarrow \omega \subseteq y)$

Axiom schemes

Separation: $\lambda \varphi. \forall x. \exists y. \forall z. z \in y \leftrightarrow z \in x \wedge \varphi[x]$

Replacement: $\lambda \varphi. (\forall xy y'. \varphi[x, y] \rightarrow \varphi[x, y'] \rightarrow y \equiv y') \rightarrow \forall x. \exists y. \forall z. z \in y \leftrightarrow \exists u \in x. \varphi[u, z]$

Equality axioms

Reflexivity: $\forall x. x \equiv x$

Symmetry: $\forall xy. x \equiv y \rightarrow y \equiv x$

Transitivity: $\forall xyz. x \equiv y \rightarrow y \equiv z \rightarrow x \equiv z$

Congruence: $\forall xx' yy'. x \equiv x' \rightarrow y \equiv y' \rightarrow x \in y \rightarrow x' \in y'$

The core axiomatisation Z' contains extensionality and the set operation axioms, Z adds the separation scheme, and ZF also adds the replacement scheme. The equality axioms are added when working with the deduction system or in an intensional model.