

10-10-2000

Eliminating Definitions and Skolem Functions in First-Order Logic

Jeremy Avigad
Carnegie Mellon University

Follow this and additional works at: <http://repository.cmu.edu/philosophy>



Part of the [Philosophy Commons](#)

Recommended Citation

Avigad, Jeremy, "Eliminating Definitions and Skolem Functions in First-Order Logic" (2000). *Department of Philosophy*. Paper 167.
<http://repository.cmu.edu/philosophy/167>

This Technical Report is brought to you for free and open access by the Dietrich College of Humanities and Social Sciences at Research Showcase @ CMU. It has been accepted for inclusion in Department of Philosophy by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

*Eliminating definitions and Skolem
functions in first-order logic*

Jeremy Avigad

October 10, 2000

Technical Report No. CMU-PHIL-113

Philosophy

Methodology

Logic

Carnegie Mellon

Pittsburgh, Pennsylvania 15213

Eliminating definitions and Skolem functions in first-order logic*

Jeremy Avigad

October 10, 2000

Abstract

In any classical first-order theory that proves the existence of at least two elements, one can eliminate definitions with a polynomial bound on the increase in proof length. In any classical first-order theory strong enough to code finite functions, including sequential theories, one can also eliminate Skolem functions with a polynomial bound on the increase in proof length.

1 Introduction

When working with a first-order theory, it is often convenient to use definitions. That is, if $\varphi(\vec{x})$ is a first-order formula with the free variables shown, one can introduce a new relation symbol R to abbreviate φ , with defining axiom $\forall \vec{x} (R(\vec{x}) \leftrightarrow \varphi(\vec{x}))$. Of course, this definition can later be eliminated from a proof, simply by replacing every instance of R by φ .

But suppose the proof involves nested definitions, with a sequence of relation symbols R_0, \dots, R_k abbreviating formulae $\varphi_0, \dots, \varphi_k$, where each φ_i may have multiple occurrences of R_0, \dots, R_{i-1} . In that case, the naive elimination procedure described above can yield an exponential increase in the length of the proof.

In Section 2, I show that if the underlying theory proves that there are at least two elements in the universe, a more careful translation allows one to eliminate the new definitions with at most a polynomial increase in length. A similar trick has been used by Solovay in simulating iterated definitions efficiently, as discussed in [10, Section 3.2]. The essence of Section 2 is therefore the observation that the construction works more generally. As a corollary, the question as to whether there is an efficient means of eliminating definitions from proofs in pure first-order logic is equivalent to the corresponding question for propositional logic.

*Partially supported by NSF grant DMS 0070600. This is a not a final version; it has been submitted for publication.

It is also sometimes convenient, in a first-order setting, to introduce Skolem functions. If $\varphi(\vec{x}, y)$ is any formula with the free variables shown and f is a new function symbol, one can add an axiom, $\forall \vec{x}, y (\varphi(\vec{x}, y) \rightarrow \varphi(\vec{x}, f(\vec{x})))$, asserting, in words, “if any y satisfies $\varphi(\vec{x}, y)$, $f(\vec{x})$ does.” There is an easy model-theoretic proof of the fact that this does not alter the set of consequences in the original language: any first-order model of the original theory can be expanded to a model where f denotes such a choice function.

Explicit syntactic proofs of this fact are, however, somewhat more difficult. The first such proof appears in Hilbert and Bernays’ *Grundlagen der Mathematik* [8], using the epsilon substitution method; a proof by Maehara using cut-elimination is discussed in [13]; and another proof due to Shoenfield is found in [12] (see also the discussion in [11]). All these procedures are, unfortunately, worse than exponential.

In Section 3, I show that if the underlying theory allows for a modicum of coding, one can also eliminate Skolem functions with at most a polynomial increase in proof length. The idea is to use an internal, iterated forcing argument to add the new functions. Similar forcing arguments can be found in [1, 2, 3, 4, 5]. In the case at hand, the forcing conditions are finite approximations to the Skolem functions being added, so the constraint on the underlying theory is that it provides an adequate representation of finite functions. The specific requirements are spelled out below; any sequential theory of arithmetic meets these criteria.

In Section 4, I discuss some questions that remain.

I would like to thank Samuel Buss for advice and suggestions.

2 Eliminating definitions

If d is a proof of a sentence ψ from a set of axioms Γ in first-order logic, then $|d|$ denotes the length of d , according to the number of symbols. Krajíček [9] and Pudlák [10] provide good general references on the lengths of proofs.

In this section and the next I will show that in certain circumstances one can eliminate definitions and/or Skolem functions from a proof d in such a way that the length of the resulting proof is bounded by a polynomial in $|d|$. In doing so, I will not make an effort to compute the exact polynomial; rather, I will repeatedly appeal to the fact that the set of polynomials in $|d|$ is closed under addition, multiplication, and composition. It will be clear from the proofs that in fact all the translations considered can be carried out in polynomial time.

By “first-order logic,” I mean first-order logic with equality, in any of the standard natural deduction calculi, Hilbert-style calculi, or sequent calculi with cut described in [14]. By a theorem due to Krajíček, up to polynomial-time equivalence it does not matter whether we take proofs to be given by trees or sequences of lines (see [10, Section 4], or [9, Section 4.5] for the propositional case). In fact, the proof of Theorem 2.2 only assumes that there is a representation of $\varphi \rightarrow \psi$ which uses φ only once. If \leftrightarrow is assumed to be one of the basic connectives, one can simplify the central argument somewhat; but the proof

below works in either case.

I will use the following conventions: \vec{x} and \vec{t} denote sequences of variables and terms, respectively, and typically their lengths can be inferred from the context. Introducing a formula as $\varphi(\vec{x})$ only serves to distinguish the sequence of variables \vec{x} , after which $\varphi(\vec{t})$ denotes the result of simultaneously substituting \vec{t} for \vec{x} , renaming bound variables in φ if necessary.

Definition 2.1 *Let Γ be a set of first-order sentences in a language L . Say that Γ has an efficient elimination of definitions if there is a polynomial $p(x)$ such that the following holds: whenever $R_0(\vec{x}_0), \dots, R_k(\vec{x}_k)$ are new relation symbols of various arities, $\varphi_0(\vec{x}_0), \dots, \varphi_k(\vec{x}_k)$ are formulae such that each φ_i is in the language $L \cup \{R_0, \dots, R_{i-1}\}$, and d is a proof of a formula ψ in L from*

$$\Gamma \cup \{\forall \vec{x}_0 (R_0(\vec{x}_0) \leftrightarrow \varphi_0(\vec{x}_0)), \dots, \forall \vec{x}_k (R_k(\vec{x}_k) \leftrightarrow \varphi_k(\vec{x}_k))\},$$

then there is a proof d' of ψ from Γ using only formulae in L , with $|d'| \leq p(|d|)$.

This definition is monotone in Γ : if Γ has an efficient elimination of definitions and $\Gamma' \supseteq \Gamma$ then, by the deduction theorem, Γ' has an efficient elimination of definitions as well. The main theorem in this section is the following:

Theorem 2.2 *$\{\exists x, y (x \neq y)\}$ has an efficient elimination of definitions.*

Proof. The proof will occupy most of this section. Let $R_0, \dots, R_k, \varphi_0, \dots, \varphi_k, \psi$, and d be as in the definition. We can assume that each of the defining axioms occurs at least once in the proof, since if the axiom for R_i does not occur in the proof we can replace each occurrence of R_i by an arbitrary sentence, say $\forall x (x = x)$. As a result, we can assume that k and $|\varphi_0|, \dots, |\varphi_k|$ are all less than $|d|$, and so it suffices to bound the length of the final proof by a polynomial in these values.

Let a and b be new constant symbols. It suffices to find a short (i.e. polynomially bounded) proof of ψ from $\{a \neq b\}$. For, if we can find a short proof of $a \neq b \rightarrow \psi$, we can replace a and b by variables and obtain a short proof of ψ from $\exists x, y (x \neq y)$.

First, note that without loss of generality we can assume that all the definitions are given by prenex formulae. If the propositional connectives are among $\{\wedge, \vee, \rightarrow, \neg\}$ this is so because any formula involving these connectives can be proved equivalent to a prenex formula, with a proof whose length is bounded by a polynomial in the length of the original formula. On the other hand, if, say, \leftrightarrow is a propositional connective, one can introduce additional definitions to abbreviate subformulae and ensure that all the definitions are prenex. Alternatively, one can first use definitions to eliminate \leftrightarrow as in the proof of Corollary 2.5, and then proceed as before.

In the following argument, if θ is a formula with a relation symbol $R(\vec{y})$ and $\eta(\vec{y})$ is a formula with the free variables shown, it will be convenient to write $\theta[\eta/R]$ for the result of replacing each atomic formula $R(\vec{t})$ by $\eta(\vec{t})$. At other times, I will write $\theta[R(t_1, \dots, t_m)]$ to indicate that an atomic formula

$R(t_1, \dots, t_m)$ occurs in the quantifier-free formula θ ; thereafter, $\theta[\eta]$ denotes the result of replacing $R(t_1, \dots, t_m)$ by η . While this notation is potentially problematic, the intention should always be clear from the context.

For notational convenience, we may assume that all of the relations R_i have the same arity. We will need a way of representing the numbers $0, \dots, k$. Let z_0, \dots, z_k be a sequence of variables, write $\bar{0}$ for the sequence a, b, b, b, \dots , $\bar{1}$ for the sequence b, a, b, b, \dots , and, more generally, \bar{j} for the sequence of length $k+1$ that has an a in the j th position and b 's elsewhere.

Our strategy will be to define a sequence of formulae $\hat{\varphi}_0(\vec{z}, u, \vec{x}), \dots, \hat{\varphi}_k(\vec{z}, u, \vec{x})$, with length bounded by a polynomial in $|d|$, such that for each $i \leq k$ the following equivalences are all provable from $a \neq b$:

- $\hat{\varphi}_i(\bar{j}, a, \vec{x}) \leftrightarrow \hat{\varphi}_{i-1}(\bar{j}, a, \vec{x})$, for each $j < i$
- $\hat{\varphi}_i(\bar{j}, b, \vec{x}) \leftrightarrow \neg \hat{\varphi}_{i-1}(\bar{j}, a, \vec{x})$, for each $j < i$
- $\forall \vec{x} (\hat{\varphi}_i(\bar{i}, a, \vec{x}) \leftrightarrow \varphi_i(\vec{x})[\hat{\varphi}_{i-1}(\bar{0}, a, \vec{x})/R_0, \dots, \hat{\varphi}_{i-1}(\bar{i-1}, a, \vec{x})/R_{i-1}])$
- $\forall \vec{x} (\hat{\varphi}_i(\bar{i}, b, \vec{x}) \leftrightarrow \neg \varphi_i(\vec{x})[\hat{\varphi}_{i-1}(\bar{0}, a, \vec{x})/R_0, \dots, \hat{\varphi}_{i-1}(\bar{i-1}, a, \vec{x})/R_{i-1}])$.

In other words, for each i and $j \leq i$, $\hat{\varphi}_i(\bar{j}, a, \vec{x})$ is an efficient representation of R_j , and $\hat{\varphi}_i(\bar{j}, b, \vec{x})$ is an efficient representation of $\neg R_j$. The idea is to use quantifiers wisely so that only a single instance of $\hat{\varphi}_i$ is used in the definition of $\hat{\varphi}_{i+1}$. Note that the clauses above imply that for each i and $j \leq i$, we have $\hat{\varphi}_i(\bar{j}, a, \vec{x}) \leftrightarrow \neg \hat{\varphi}_i(\bar{j}, b, \vec{x})$.

The sequence $\hat{\varphi}_0, \dots, \hat{\varphi}_k$ is defined recursively. Start by taking $\hat{\varphi}_0(\vec{z}, u, \vec{x})$ to be the formula

$$(u = a \rightarrow \varphi_0(\vec{x})) \wedge (u = b \rightarrow \neg \varphi_0(\vec{x})).$$

For $i > 0$, assuming $\hat{\varphi}_0, \dots, \hat{\varphi}_{i-1}$ have been defined, the following shows how to determine $\hat{\varphi}_i$. Since we are assuming that all the definitions are prenex, $\varphi_i(\vec{x})$ is of the form

$$Q_1 y_1 \dots Q_m y_m \tilde{\varphi}[R_0(\vec{t}_{0,0}), \dots, R_0(\vec{t}_{0,l_0}), \dots, R_{i-1}(\vec{t}_{i-1,0}), \dots, R_{i-1}(\vec{t}_{i-1,l_{i-1}})],$$

where $\tilde{\varphi}$ is quantifier-free and the sequence in square brackets shows all instances of atomic formulae in $\tilde{\varphi}$ involving R_0, \dots, R_{i-1} . In general, the sequences of terms $\vec{t}_{j,p}$ depend on the quantified variables y_1, \dots, y_m as well as the free variables \vec{x} of φ_i , but I will not display these variables explicitly. Our task is to write down a formula $\hat{\varphi}_i(\vec{z}, u, \vec{x})$ such that

1. for each $j < i$, $\hat{\varphi}_i(\bar{j}, a, \vec{x})$ is equivalent to $\varphi_{i-1}(\bar{j}, a, \vec{x})$;
2. for each $j < i$, $\hat{\varphi}_i(\bar{j}, b, \vec{x})$ is equivalent to $\neg \varphi_{i-1}(\bar{j}, a, \vec{x})$;
3. $\hat{\varphi}_i(\bar{i}, a, \vec{x})$ is equivalent to the displayed formula above, with each $R_j(\vec{t}_{j,p})$ replaced by $\hat{\varphi}_{i-1}(\bar{j}, a, \vec{t}_{j,p})$;
4. $\hat{\varphi}_i(\bar{i}, b, \vec{x})$ is equivalent to the negation of the formula just described; and

5. in the definition of $\hat{\varphi}_i$, $\hat{\varphi}_{i-1}$ is used only once.

In order to do 3 and 4 simultaneously, we need duplicate copies of some of the variables and terms. Let Q'_1, \dots, Q'_m denote the quantifiers dual to Q_1, \dots, Q_m . Pick a new sequence of variables y'_1, \dots, y'_m , and let

$$\vec{t}_{0,0}, \dots, \vec{t}_{0,l_0}, \dots, \vec{t}_{i-1,0}, \dots, \vec{t}_{i-1,l_{i-1}}$$

denote the sequences of terms obtained by replacing the y_1, \dots, y_m by y'_1, \dots, y'_m in each $\vec{t}_{j,p}$. Finally, let

$$\begin{aligned} &v_{0,0}, \dots, v_{0,l_0}, \dots, v_{i-1,0}, \dots, v_{i-1,l_{i-1}} \\ &v'_{0,0}, \dots, v'_{0,l_0}, \dots, v'_{i-1,0}, \dots, v'_{i-1,l_{i-1}} \\ &v''_0, \dots, v''_{i-1} \end{aligned}$$

be sequences of new variables. We will use the variables $v_{j,p}$ to represent the truth values of $\hat{\varphi}_{i-1}(\vec{j}, a, \vec{t}_{j,p})$, the variables $v'_{j,p}$ to represent the truth values of $\hat{\varphi}_{i-1}(\vec{j}, a, \vec{t}'_{j,p})$, and the variables v''_j to represent the truth values of $\hat{\varphi}_{i-i}(\vec{j}, a, \vec{x})$, where the “truth value” is a if the corresponding formula is true, and b if it is false.

The formula $\hat{\varphi}_i(\vec{z}, u, \vec{x})$ is defined to be

$$\begin{aligned} &Q_1 y_1 \dots Q_m y_m Q'_1 y'_1 \dots Q'_m y'_m \forall \vec{v}, \vec{v}', \vec{v}'' \left(Eval(\vec{v}, \vec{v}', \vec{v}'') \rightarrow \right. \\ &\quad \bigwedge_{j < i} (\vec{z} = \vec{j} \wedge u = a \rightarrow v''_j = a) \wedge \\ &\quad \bigwedge_{j < i} (\vec{z} = \vec{j} \wedge u = b \rightarrow v''_j \neq a) \wedge \\ &\quad (\vec{z} = \vec{i} \wedge u = a \rightarrow \tilde{\varphi}[v_{0,0} = a, \dots, v_{0,l_0} = a, \dots, v_{i-1,0} = a, \dots, v_{i-1,l_{i-1}} = a]) \wedge \\ &\quad \left. (\vec{z} = \vec{i} \wedge u = b \rightarrow \neg \tilde{\varphi}[v'_{0,0} = a, \dots, v'_{0,l_0} = a, \dots, v'_{i-1,0} = a, \dots, v'_{i-1,l_{i-1}} = a]) \right) \end{aligned}$$

where $Eval(\vec{v}, \vec{v}', \vec{v}'')$ is the formula

$$\begin{aligned} &\forall \vec{r} \forall s \in \{a, b\} \forall \vec{w} \left(\hat{\varphi}_{i-1}(\vec{r}, s, \vec{w}) \rightarrow \right. \\ &\quad \bigwedge_{j < i} (\vec{r} = \vec{j} \wedge \vec{w} = \vec{x} \rightarrow v''_j = s) \wedge \\ &\quad \bigwedge_{j < i} \bigwedge_{p < l_j} (\vec{r} = \vec{j} \wedge \vec{w} = \vec{t}_{j,p} \rightarrow v_{j,p} = s) \wedge \\ &\quad \left. \bigwedge_{j < i} \bigwedge_{p < l_j} (\vec{r} = \vec{j} \wedge \vec{w} = \vec{t}'_{j,p} \rightarrow v'_{j,p} = s) \right). \end{aligned}$$

Here $\forall s \in \{a, b\} \theta$ abbreviates $\forall s (s = a \vee s = b \rightarrow \theta)$. Note that $Eval(\vec{v}, \vec{v}', \vec{v}'')$ also depends on the free variables $\vec{x}, \vec{y}, \vec{y}'$ (because the terms $t_{j,p}$ and $t'_{j,p}$ do), but I will continue to leave these variables implicit.

First, let us check that each $\hat{\varphi}_i(\vec{x}_i, u)$ satisfies the right equivalences, and then let us worry about the length. Inductively we know, for each $j \leq i-1$, that

$$\forall \vec{x} (\hat{\varphi}_{i-1}(\vec{j}, a, \vec{x}) \leftrightarrow \neg \hat{\varphi}_{i-1}(\vec{j}, b, \vec{x}))$$

is provable from $a \neq b$. We can use this to show

$$\forall \vec{x}, \vec{y}, \vec{y}' \exists \vec{v}, \vec{v}', \vec{v}'' \text{ Eval}(\vec{v}, \vec{v}', \vec{v}'')$$

as well as

$$\begin{aligned} \forall \vec{x}, \vec{y}, \vec{y}', \vec{v}, \vec{v}', \vec{v}'' & \left(\text{Eval}(\vec{v}, \vec{v}', \vec{v}'') \rightarrow \right. \\ & \bigwedge_{j < i} (v_j'' = a \leftrightarrow \hat{\varphi}_{i-1}(\vec{j}, a, \vec{x})) \wedge \\ & \bigwedge_{j < i} \bigwedge_{p < l_j} (v_{j,p} = a \leftrightarrow \hat{\varphi}_{i-1}(\vec{j}, a, \vec{t}_{j,p})) \wedge \\ & \left. \bigwedge_{j < i} \bigwedge_{p < l_j} (v'_{j,p} = a \leftrightarrow \hat{\varphi}_{i-1}(\vec{j}, a, \vec{t}'_{j,p})) \right). \end{aligned}$$

But then, going back to the definition of $\hat{\varphi}_i$, we see that for $j < i$, $\hat{\varphi}_i(\vec{j}, a, \vec{x})$ is equivalent to $\hat{\varphi}_{i-1}(\vec{j}, a, \vec{x})$, and $\hat{\varphi}_i(\vec{j}, b, \vec{x})$ is equivalent to $\neg \hat{\varphi}_{i-1}(\vec{j}, a, \vec{x})$. Also, $\hat{\varphi}_i(\vec{i}, a, \vec{x})$ is equivalent to

$$\begin{aligned} Q_1 y_1 \dots Q_m y_m \tilde{\varphi}[\hat{\varphi}_{i-1}(\vec{0}, a, \vec{t}_{0,0}), \dots, \hat{\varphi}_{i-1}(\vec{0}, a, \vec{t}_{0,l_0}), \dots, \\ \hat{\varphi}_{i-1}(\vec{i-1}, a, \vec{t}_{i-1,0}), \dots, \hat{\varphi}_{i-1}(\vec{i-1}, a, \vec{t}_{i-1,l_{i-1}})] \end{aligned}$$

and so we have

$$\hat{\varphi}_i(\vec{i}, a, \vec{x}) \leftrightarrow \varphi_i(\vec{x})[\hat{\varphi}_{i-1}(\vec{0}, a, \vec{x})/R_0, \dots, \hat{\varphi}_{i-1}(\vec{i-1}, a, \vec{x})/R_{i-1}];$$

and $\hat{\varphi}_i(\vec{i}, b, \vec{x})$ is equivalent to

$$\begin{aligned} Q'_1 y'_1 \dots Q'_m y'_m \neg \tilde{\varphi}[\hat{\varphi}_{i-1}(\vec{0}, a, \vec{t}_{0,0}), \dots, \hat{\varphi}_{i-1}(\vec{0}, a, \vec{t}_{0,l_0}), \dots, \\ \hat{\varphi}_{i-1}(\vec{i-1}, a, \vec{t}_{i-1,0}), \dots, \hat{\varphi}_{i-1}(\vec{i-1}, a, \vec{t}_{i-1,l_{i-1}})] \end{aligned}$$

and so we have

$$\hat{\varphi}_i(\vec{i}, b, \vec{x}) \leftrightarrow \neg \varphi_i(\vec{x})[\hat{\varphi}_{i-1}(\vec{0}, a, \vec{x})/R_0, \dots, \hat{\varphi}_{i-1}(\vec{i-1}, a, \vec{x})/R_{i-1}],$$

as required.

As far as length is concerned, it is not hard to check that the number of symbols occurring in $\hat{\varphi}_i$ apart from the instance of $\hat{\varphi}_{i-1}$ can be bounded by a polynomial in $|d|$ (in fact, even a linear one). In other words, there is a polynomial p such that for each i we have $|\hat{\varphi}_i| \leq p(|d|) + |\hat{\varphi}_{i-1}|$, and hence $|\hat{\varphi}_i| \leq (i+1)p(|d|) \leq |d|p(|d|)$. Similarly, it is not hard to find polynomial bounds on the lengths of the proofs of the needed equivalences, and there are only polynomially many of them.

This completes the proof of Theorem 2.2. \square

We have handled the case where there are at least two elements in the universe. On the other hand, on the assumption that there is only one element of the universe, we are reduced to propositional logic.

Proposition 2.3 *$\{\forall x, y (x = y)\}$ has efficient elimination of definitions if and only if the corresponding assertion holds for propositional logic.*

Proof. Assuming $\forall x, y (x = y)$, every atomic formula $R(t_1, \dots, t_k)$ is equivalent to $R(c, \dots, c)$, where c is the only element of the universe; $t_1 = t_2$ is always true; and quantifiers have no effect. To be more precise, let “the propositional simplification of ψ ” denote the result of deleting all the quantifiers in ψ , replacing all atomic formulae $R(t_1, \dots, t_k)$ by a propositional variable R , and replacing $t_1 = t_2$ by a fixed tautology. Then any first-order proof of $\forall x, y (x = y) \rightarrow \psi$ can be translated efficiently to a propositional proof of the propositional simplification of ψ , and vice-versa. \square

This implies that the general problem of eliminating definitions from proofs in pure first-order logic is as hard (and as easy) as the propositional case.

Theorem 2.4 *\emptyset has an efficient elimination of definitions if and only if the corresponding assertion holds for propositional logic.*

Proof. It is a straightforward exercise to check that $\{\varphi \vee \psi\}$ has an efficient elimination of definitions if and only if $\{\varphi\}$ and $\{\psi\}$ both do. In particular, \emptyset has an efficient elimination of definitions if and only if $\{\forall x, y (x = y)\}$ and $\{\exists x, y (x \neq y)\}$ do. \square

As a corollary of Theorem 2.2, we have that one can eliminate \leftrightarrow from standard proof systems with at most a polynomial increase in proof length. For propositional proof systems the proof (due to Reckhow, using a method by Spira; see [9]) is considerably more difficult.

Corollary 2.5 *With any of the standard proof systems for first-order logic with equality given in [14], one can eliminate the propositional connective \leftrightarrow with at most a polynomial increase in proof length.*

Proof. By Theorem 2.2, it suffices to show that one can eliminate \leftrightarrow efficiently in the corresponding proof systems with definitions. Use definitions to translate formulae in the language with \leftrightarrow to the language without: translate $\varphi(\vec{w}) \leftrightarrow \psi(\vec{z})$ to $(R_\varphi(\vec{w}) \rightarrow R_\psi(\vec{z})) \wedge (R_\psi(\vec{z}) \rightarrow R_\varphi(\vec{w}))$, where R_0 and R_1 are defined to be equivalent to the translations of φ and ψ , respectively. By induction one can show that each axiom and rule of inference can then be simulated, with polynomial bounds on the lengths. \square

3 Eliminating Skolem functions

The following is the analogue of Definition 2.1 for Skolem functions.

Definition 3.1 *Let Γ be a set of first-order sentences in a language L . Say that Γ has an efficient elimination of Skolem functions if there is a polynomial $p(x)$ such that the following holds: whenever $f_0(\vec{x}_0), \dots, f_k(\vec{x}_k)$ are new function symbols of various arities, $\varphi_0(\vec{x}_0, y), \dots, \varphi_k(\vec{x}_k, y)$ are formulae such that each φ_i is in the language $L \cup \{f_0, \dots, f_{i-1}\}$, and d is a proof of a formula ψ in L from*

$$\Gamma \cup \{\forall \vec{x}_0, y (\varphi_0(\vec{x}_0, y) \rightarrow \varphi(\vec{x}_0, f_0(\vec{x}_0))), \dots, \forall \vec{x}_k, y (\varphi_k(\vec{x}_k, y) \rightarrow \varphi(\vec{x}_k, f_k(\vec{x}_k)))\},$$

then there is a proof d' of ψ from Γ using only formulae in L , with $|d'| \leq p(|d|)$.

Right off the bat, we have the following.

Proposition 3.2 $\{\forall x, y (x = y)\}$ *has an efficient elimination of Skolem functions.*

Proof. Roughly speaking, if c is the only element of the universe, every term can be replaced by c . \square

By way of motivation, note that is not hard to show that, say, Zermelo-Fraenkel set theory has an efficient elimination of Skolem functions. Argue as follows. Suppose d is a proof of a formula ψ from the axioms of ZF and some Skolem functions. Let k be a bound on the complexity of the formulae occurring in this proof. In ZF , one can prove that the set of true sentences of complexity at most $k + 1$ is consistent, and hence has a countable model. This countable model has Skolem functions, which can then be used to interpret the proof d .

This example suggests that one way to proceed is to try to determine how little one can get away with in carrying out an internal semantic argument of this kind. The answer turns out to be: very little.

Definition 3.3 *Say a set of sentences Γ codes finite functions (efficiently) if for each n there are*

- *a definable element, " \emptyset_n ";*
- *a definable relation, " $x_0, \dots, x_{n-1} \in \text{dom}_n(p)$ ";*
- *a definable function, " $\text{eval}_n(p, x_0, \dots, x_{n-1})$ "; and*
- *a definable function, " $p \oplus_n (x_0, \dots, x_{n-1} \mapsto y)$ "*

such that, for each n , Γ proves

- $\vec{x} \notin \text{dom}(\emptyset_n)$
- $\vec{w} \in \text{dom}(p \oplus_n (\vec{x} \mapsto y)) \leftrightarrow (\vec{w} \in \text{dom}(p) \vee \vec{w} = \vec{x})$

- $eval_n(p \oplus_n (\vec{x} \mapsto y), \vec{x}) = y$
- $\vec{w} \neq \vec{x} \rightarrow eval_n(p \oplus_n (\vec{x} \mapsto y), \vec{w}) = eval_n(p, \vec{w})$,

and such that the lengths of all the definitions and proofs are bounded by a polynomial in n .

Of course, the intuition is that elements of the universe are assumed to code finite partial functions p ; \emptyset_n is the function that is nowhere defined; $eval_n(p, \vec{x})$ returns the value of p at \vec{x} ; $p \oplus_n (\vec{x} \mapsto y)$ is the modification of p which maps \vec{x} to y ; and so on. One could, more generally, assume that the codes are elements of a definable set; but then nothing is lost by taking the other elements of the universe to code the empty function. If one wants polynomial-time translations (and not just bounds on the lengths of proofs) one needs to add the constraint that the definitions and proofs above are polynomial-time computable in n .

These requirements are not strong ones. For example, any sequential theory of arithmetic (in the terminology of [7]) codes finite functions, since one can take such functions to be sequences of tuples $\langle \vec{x}, y \rangle$. Below I will drop the subscripts n in \emptyset_n , dom_n , etc. and I will write $p(\vec{x})$ instead of $eval(p, \vec{x})$. Clearly it does not hurt to assume that all these are actually given by symbols in the language.

Theorem 3.4 *Suppose Γ codes finite functions. Then Γ has an efficient elimination of Skolem functions.*

Proof. The proof will occupy most of the remainder of this section. By Proposition 3.2 we can assume that there are at least two elements in the universe, and so, by Theorem 2.2, we can use definitions freely. By way of exposition, I will first focus on the case where $k = 0$, i.e. there is only one Skolem function to eliminate. (This part does not require definitions.) Then I will discuss the steps necessary to eliminate multiple, possibly nested instances Skolem functions. (This is the part that requires the definitions.)

Suppose we want to eliminate the use of a single Skolem function, with defining axiom $\forall \vec{x}, y (\varphi(\vec{x}, y) \rightarrow \varphi(\vec{x}, f(x)))$. Let L_f denote the language $L \cup \{f\}$. I will define a forcing relation in L , for formulae in L_f . I will then show that Γ proves that the Skolem axiom is forced; and that anything in the original language is forced if and only if it is true. Given a proof d of ψ from Γ together with the Skolem axiom, then, Γ proves that ψ is forced, and hence true.

Now for the details. Let the formula $Cond(p)$ in the language L assert that p is a finite approximation to a Skolem function for φ , that is,

$$\forall \vec{x} \in dom(p) \forall y (\varphi(\vec{x}, y) \rightarrow \varphi(\vec{x}, f(x))).$$

Let t be a term in L_f , and let p be a variable not occurring in t . Inductively we will define a term t^p in the language of L , whose free variables are those of t together with p . Intuitively, t^p is the value of t , when f is interpreted by p . At the same time, we will define a relation “ t^p is defined,” asserting that the value of t^p makes sense. Let

- $x^p \equiv x$, for each variable x (other than p),
- $(g(t_0, \dots, t_m))^p \equiv g(t_0^p, \dots, t_m^p)$, for each function symbol g of L , and
- $(f(t_0, \dots, t_n))^p \equiv p(t_0^p, \dots, t_n^p)$.

Define “ t^p is defined” inductively as follows:

- “ x^p is defined” is always true.
- “ $(g(t_0, \dots, t_m))^p$ is defined,” where g is a function symbol of L , is true if and only if t_0^p, \dots, t_m^p are all defined.
- “ $(f(t_0, \dots, t_n))^p$ is defined” is true if and only if t_0^p, \dots, t_n^p are all defined and $t_0^p, \dots, t_n^p \in \text{dom}(p)$.

If p and q are conditions, say $p \preceq q$, “ p is stronger than or equal to q ”, if p extends q as a function:

$$\forall \vec{x} (\vec{x} \in \text{dom}(q) \rightarrow \vec{x} \in \text{dom}(p) \wedge p(\vec{x}) = q(\vec{x})).$$

Now we can define the relation $p \Vdash \theta$ inductively. We can assume that the language has connectives \wedge , \rightarrow , \forall , and \neg , with \exists and \vee defined from these in the usual way.

1. $p \Vdash R(t_0, \dots, t_m)$ if and only if $\forall q \preceq p \exists r \preceq q$ (t_0^r, \dots, t_m^r are all defined and $R(t_0^r, \dots, t_m^r)$).
2. $p \Vdash \theta \wedge \eta$ if and only if $p \Vdash \theta$ and $p \Vdash \eta$.
3. $p \Vdash \theta \rightarrow \eta$ if and only if $\forall q \preceq p (q \Vdash \theta \rightarrow q \Vdash \eta)$.
4. $p \Vdash \neg \theta$ if and only if $\forall q \preceq p q \nVdash \theta$.
5. $p \Vdash \forall x \theta$ if and only if $\forall x p \Vdash \theta$.

The quantifiers involving q and r above are intended to range over conditions; so, for example, $\forall q \preceq p \dots$ abbreviates $\forall q (\text{Cond}(q) \wedge q \preceq p \rightarrow \dots)$. For each θ , the relation $p \Vdash \theta$ is a formula in the language of L whose free variables are those of θ together with p . Note that the length of $p \Vdash \theta$ can be bounded by a polynomial in $|\theta|$ (as well as in $|\varphi|$, which is being held fixed for the moment).

The phrase “ θ is forced” and the notation $\Vdash \theta$ abbreviate $\forall p (\text{Cond}(p) \rightarrow p \Vdash \theta)$. In the lemmata that follow, $p, q, r \dots$ are assumed to range over conditions. Most of the proofs are routine and standard, modulo the additional notes provided below. It is important to recognize that the lengths of all the proofs alluded to in the statement of the lemmata can be bounded by a polynomial in the length of the assertion being proved, but having stated this up front, I will not bother to repeat it each time.

Lemma 3.5 (monotonicity) *For each formula θ of L_f , Γ proves*

$$p \Vdash \theta \wedge q \preceq p \rightarrow q \Vdash \theta.$$

Lemma 3.6 For each formula θ of L_f , Γ proves

$$p \Vdash \theta \leftrightarrow \forall q \preceq p \exists r \preceq q \ r \Vdash \theta.$$

Corollary 3.7 For each formula θ of L_f , Γ proves

$$p \Vdash (\theta \leftrightarrow \neg\neg\theta).$$

Lemma 3.8 For any term t of L_f , Γ proves

$$\forall q \exists r \preceq q \ (t^r \text{ is defined}).$$

Proof. Use induction on the term t . The only interesting case is where t is of the form $f(s_0, \dots, s_k)$. By the induction hypothesis, we can find an $r' \preceq q$ such that $s_0^{r'}, \dots, s_k^{r'}$ are all defined. If $s_0^{r'}, \dots, s_k^{r'} \in \text{dom}(r')$, take $r = r'$. Otherwise, if $\exists y \ \varphi(s_0^{r'}, \dots, s_k^{r'}, y)$, let $r = r' \oplus (s_0^{r'}, \dots, s_k^{r'} \mapsto y)$, for any such y ; and if $\forall y \ \neg\varphi(s_0^{r'}, \dots, s_k^{r'}, y)$, let $r = r' \oplus (s_0^{r'}, \dots, s_k^{r'} \mapsto y)$, for any y at all. \square

The next two lemmata are proved by induction on s and θ , respectively.

Lemma 3.9 If t and $s(x)$ are any terms of L_f , Γ proves

$$t^p = z \rightarrow (s(t)^p = s(z)^p)$$

Lemma 3.10 If $\theta(x)$ is any formula of L_f and t is any term of L_f then Γ proves

$$(t^p \text{ is defined} \wedge t^p = z) \rightarrow (p \Vdash \theta(t) \leftrightarrow p \Vdash \theta(z)).$$

Lemma 3.11 For each formula θ of L_f , if θ is provable in classical first-order logic, then Γ proves $\Vdash \theta$.

Proof. The proof is for the most part standard and routine, though one has to be a little bit careful with the quantifier axioms and rules since terms might not always be “defined.” To show $\forall x \ \theta(x) \rightarrow \theta(t)$ is forced, let us argue in first-order logic from assumptions in Γ . Suppose $p \Vdash \forall x \ \theta(x)$. By Lemma 3.6 it suffices to show $\forall q \preceq p \exists r \preceq q \ \theta(t)$. So suppose $q \preceq p$, and by Lemma 3.8 let $r \preceq q$ be such that t^r is defined. Let $z = t^r$. By monotonicity, $r \Vdash \forall x \ \theta(x)$, so $r \Vdash \theta(z)$. By Lemma 3.10, $r \Vdash \theta(t)$. \square

A formula in the original language is forced if and only if it is true.

Lemma 3.12 For each formula θ of L , Γ proves $(p \Vdash \theta) \leftrightarrow \theta$.

Proof. Induction on θ . \square

The next lemma is the important one: it asserts that the Skolem axiom is forced.

Lemma 3.13 Γ proves $\Vdash \forall \vec{x}, y \ (\varphi(\vec{x}, y) \rightarrow \varphi(\vec{x}, f(\vec{x})))$.

Proof. Once again, argue in first-order logic, from Γ . Suppose for some \vec{x}, y we have $p \Vdash \varphi(\vec{x}, y)$. By Lemma 3.12, $\varphi(\vec{x}, y)$. By Lemma 3.6, it suffices to show $\forall q \preceq p \exists r \preceq q q \Vdash \varphi(\vec{x}, f(\vec{x}))$, so suppose $q \preceq p$. If $\vec{x} \in \text{dom}(q)$, the fact that q is a condition guarantees $\varphi(\vec{x}, f(\vec{x}))$, and we can take $r = q$; otherwise, take $r = q \oplus (\vec{x} \mapsto y)$. Either way, as above, we have $r \Vdash \varphi(\vec{x}, f(\vec{x}))$, as required. \square

Proof of Theorem 3.4, for a single Skolem function f . Suppose there is a proof d of a formula ψ in the language L from finitely many sentences in $\Gamma \cup \{\forall \vec{x}, y (\varphi(\vec{x}, y) \rightarrow \varphi(\vec{x}, f(x)))\}$. By Lemma 3.11, Γ proves that this implication is forced. By Lemmata 3.12 and 3.13, Γ proves that all the hypotheses are forced, so Γ proves that ψ is forced as well. By Lemma 3.12, Γ proves ψ .

Since each the length of each component of the derivation just described can be bounded by a polynomial in $|d|$, so can the entire proof. \square

To extend the proof to arbitrary nested definitions of Skolem functions, we need to iterate the forcing definition. A similar iteration was used in [1]; the situation here is easier, since we only have to deal with finite iterations.

Let $d, f_0, \dots, f_k, \varphi_0, \dots, \varphi_k$ be as in Definition 3.1. For each $i \leq k$, we will define the notion of an i -condition, an ordering \preceq_i on i -conditions, and a forcing relation \Vdash_i between i -conditions and formulae θ in the language $L \cup \{f_0, \dots, f_i\}$. An i -condition consists of a sequence p_0, \dots, p_i of finite functions, with arities corresponding to those of f_0, \dots, f_i . As expected, $p_0, \dots, p_i \preceq_i q_0, \dots, q_i$ means that each p_j extends q_j , as above.

The notions Cond_i and \Vdash_i are defined simultaneously, by recursion on i . $\text{Cond}_0(p)$ and $p \Vdash_0 \theta$ are defined as above, in the case where there is only one Skolem function. Assuming Cond_i and \Vdash_i have been defined, the relation $\text{Cond}_{i+1}(p_0, \dots, p_{i+1})$ is defined by

$$\text{Cond}_i(p_0, \dots, p_i) \wedge p_0, \dots, p_i \Vdash_i \forall \vec{x}_{i+1}, y (\vec{x}_{i+1} \in \text{dom}(p_{i+1}) \wedge \varphi(\vec{x}_{i+1}, y) \rightarrow \varphi(\vec{x}_{i+1}, p(\vec{x}))).$$

In the atomic case, assuming t_0, \dots, t_m are terms in the language of $L \cup \{f_0, \dots, f_{i+1}\}$, the relation $p_0, \dots, p_{i+1} \Vdash_{i+1} A(t_0, \dots, t_m)$ is defined by

$$\forall \vec{q} \preceq \vec{p} \exists \vec{r} \preceq \vec{q} (t_0^{\vec{r}}, \dots, t_m^{\vec{r}} \text{ are defined and } A(t_0^{\vec{r}}, \dots, t_m^{\vec{r}})).$$

The forcing relation is then extended to arbitrary formulae in the language as above. Notice that the relation \Vdash_i is used in the definition of Cond_{i+1} , which is in turn used to define \Vdash_{i+1} . By introducing new relation symbols to represent the definitions of $\text{Cond}_0, \dots, \text{Cond}_k$, we can bound the lengths of all the formulae involved by a polynomial.

Lemma 3.14 *For each $i \leq k$, Lemmata 3.5–3.11 hold for i -conditions, \preceq_i , and \Vdash_i .*

Lemma 3.15 *For each $i \leq k$, if θ is in the language $L \cup \{f_0, \dots, f_i\}$, then Γ proves the following:*

$$p_0, \dots, p_k \Vdash_k \theta \leftrightarrow p_0, \dots, p_i \Vdash_i \theta.$$

Lemma 3.16 *For each $i \leq k$, Γ proves that the i th Skolem axiom is k -forced.*

Once again, the lengths of the relevant proofs can be bounded by a polynomial in $|d|$. The proof of Theorem 3.4 now follows exactly as in the case of a single Skolem function. \square

If a and b are distinct and f is a Skolem function for $(\varphi(\vec{x}) \wedge y = a) \vee (\neg\varphi(\vec{x}) \wedge y = b)$, then $f(\vec{x}) = a$ serves as a definition for $\varphi(\vec{x})$. As a corollary to Theorem 3.4 we have the following:

Corollary 3.17 *Suppose Γ codes finite functions and proves $\exists x, y (x \neq y)$. Then one can eliminate arbitrary nested instances of definitions and Skolem functions from proofs in Γ , with a polynomial bound on the increase in the lengths of proofs.*

4 Questions

In standard terminology (e.g. [9, 10]), Section 2 shows that one can eliminate definitions from proofs in first-order logic in polynomial time if and only if extended Frege systems for propositional logic can be p-simulated by Frege systems. Of course, whether or not this is the case is still a major open question. Section 2 shows that Theorem 2.2 and Corollary 2.5 hold for first-order logic with equality. I do not know whether or not they hold when equality is absent.

It is also still open as to whether one can efficiently eliminate even a single Skolem function from proofs in pure logic, or from theories which do not code finite functions. The question as to the increase in proof length when eliminating a single Skolem function from a proof in pure first-order logic is listed as open problem 22 in [6].

The elimination of definitions in Section 2 used the law of the excluded middle. As a result, it is open as to whether one has an efficient elimination of definitions in intuitionistic first-order logic. (See also [11] for a discussion of choice functions in the intuitionistic setting.)

References

- [1] Jeremy Avigad. Formalizing forcing arguments in subsystems of second-order arithmetic. *Annals of Pure and Applied Logic*, 82:165–191, 1996.
- [2] Jeremy Avigad. Interpreting classical theories in constructive ones. To appear in the *Journal of Symbolic Logic*.
- [3] Jeremy Avigad. Saturated models of universal theories. Submitted.
- [4] Jeremy Avigad. Algebraic proofs of the cut-elimination theorems. Submitted.

- [5] Jeremy Avigad and Jeffrey Helzner. Transfer principles for intuitionistic nonstandard arithmetic. Submitted.
- [6] Peter Clote and Jan Krajíček. *Arithmetic, Proof Theory, and Computational Complexity*. Oxford University Press, 1993.
- [7] Petr Hájek and Pavel Pudlák. *Metamathematics of first-order arithmetic*. Springer, Berlin, 1993.
- [8] David Hilbert and Paul Bernays. *Grundlagen der Mathematik*, volume 2. Springer, Berlin, second edition, 1970.
- [9] Jan Krajíček. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*. Cambridge University Press, Cambridge, 1995.
- [10] Pavel Pudlák. The lengths of proofs. In Samuel Buss, editor, *The Handbook of Proof Theory*, North-Holland, Amsterdam, 1998, 547–637.
- [11] Helmut Schwichtenberg. Logic and the axiom of choice. In M. Boffa et al, editor, *Logic Colloquium '78*, North-Holland, 1979, 351–356.
- [12] Joseph Shoenfield. *Mathematical Logic*. Addison Wesley, Reading, Massachusetts, 1967.
- [13] Gaisi Takeuti. *Proof Theory*. North-Holland, Amsterdam, second edition, 1987.
- [14] A. S. Troelstra and Helmut Schwichtenberg. *Basic Proof Theory*. Cambridge University Press, Cambridge, 1996.