Math. Struct. in Comp. Science (2019), *vol.* 29, *pp.* 1250–1274. © Cambridge University Press 2018
doi:10.1017/S0960129518000063 First published online 17 April 2018

A semantics for nabla[†]

JEAN GOUBAULT-LARRECQ

LSV, ENS Paris-Saclay, CNRS, Université
Paris-Saclay,
94230 Cachan, France
Email: goubault@lsv.fr

Received 26 June 2017; revised 15 February 2018

We give a semantics for a classical variant of Dale Miller and Alwen Tiu's logic $FO\lambda^{\nabla}$. Our semantics validates the rule that nabla x implies exists x, but is otherwise faithful to the authors' original intentions. The semantics is based on a category of so-called nabla sets, which are simply strictly increasing sequences of non-empty sets. We show that the logic is sound for that semantics. Assuming there is a unique base type ι, we show that it is complete for Henkin structures, incomplete for standard structures in general, but complete for standard structures in the case of Π_1 formulae, and that includes all first-order formulae.

1. Prolog(ue)

I started my research career in automated deduction, and came to learn about Dale when I touched the subject of proofs in higher order logic. His work on expansion proofs was impressive, and daunting. I kept on hearing of Dale, as he developed λ-Prolog, as he discovered higher order patterns, as he realized the value of uniform proofs, of intuitionism, of hereditary Harrop formulae, as he studied extensions of logic with definitions, as he delved into focusing and linear logic, and so on and so forth.

We finally got in touch on February 14, 2002. I had sent him a rather vague question on his paper (Miller 1992) by email that day. My interest was to encode fresh names (nonces) in cryptographic protocols, and I had seen that Dale had pursued the idea of using the quantifiers of linear logic to this very end. The paper's title ended with the enigmatic phrase 'preliminary results,' and I wanted to know whether he had done any more recent research in this vein. He answered me the same day, despite the fact that he was busy at a Logic and Interaction meeting in Marseilles–Luminy, and that we had never met before. Dale has to be commended for giving me a lucid and candid answer. Who do you know would tell the following to a perfect stranger?

If you map processes to logical formulas directly, you have a lot of exciting things that can happen. My original efforts (an experiment, really) failed, however, for at least two reasons (referring to the paper 'The pi-calculus as a theory in linear logic').

[†] A variant of this paper was presented at Dale Miller's 60th birthday. The current presentation is simpler, although essentially equivalent. We have also taken the opportunity to correct a mistake in the proof of Proposition 8.4.

I am not including any more of his email to me. One of the two reasons he mentions is that, if try to encode $\nu x.P(x)$ (create a fresh name x, then do $P(x)$) as $\forall x.P(x)$ in linear logic, then you cannot make much of a difference between $\nu x.\nu y.P(x,y)$ and $\nu z.P(z,z)$, because $\forall x.\forall y.P(x,y)$ linearly implies $\forall z.P(z,z)$ – so much for y being fresh.

For cryptographic protocols, one can get around that problem (Cervesato et al. 1999), as we learned a few years later. Meanwhile, Dale worked on finding a general way of talking about freshness. Alwen Tiu and he found a simple logical way of answering the question (Miller and Tiu 2005): the *nabla* quantifier ∇. That certainly goes way beyond cryptographic protocols, and has the distinctive quality of good mathematics: simple, elegant and general. It was only natural for me to pay homage to Dale by contributing to the theory of nabla.

2. Introduction

With Alwen Tiu, Dale Miller introduced a logic $\mathrm{FO}\lambda^{\nabla}$ for so-called generic judgements (Miller and Tiu 2005). The main new feature of that logic is the *nabla* quantifier: $\nabla x : \tau.F(x)$ means that $F(x)$ holds for x *generic* of type τ.

Generic stands for 'with no remarkable property,' and is close to the notion of being fresh, but different. Pitts and Gabbay gave nice, deep definitions of the notion of freshness (Gabbay and Pitts 1999), based on the category of nominal sets. Dale Miller's solution came later, and is an elegant proof-theoretic construction. One can define what it means to be fresh, using the nabla quantifier, but there are some differences (Miller and Tiu 2005, Section 8). First, in $\nabla x : \tau.F(x)$, one may request a generic object x of *any* arbitrary type τ. The only fresh thing one can create in Pitts and Gabbay's approach is a name. Second, $\forall x.P(x)$ implies $\reflectbox{N}x.P(x)$, which implies $\exists x.P(x)$, while no such implication holds with ∇ instead of $\reflectbox{N}$. Also, while $\reflectbox{N}x.\reflectbox{N}y.P(x,y)$ and $\reflectbox{N}y.\reflectbox{N}x.P(x,y)$ are equivalent, $\nabla x.\nabla y.P(x,y)$ and $\nabla y.\nabla x.P(x,y)$ are not; but that equivalence was added later (Gacek 2008).

One may hope to understand ∇ better by giving it a semantics, and it is precisely one of the purposes of this paper. Historically, the first semantics of ∇, and of the logic $\mathrm{FO}\lambda^{\nabla}$, was given by Miculan and Yemane's (2005, Section 7), based on the category $\widehat{\mathbb{D}}$ of presheaves over the category $\mathbb{D}$ of so-called distinctions. $\mathrm{FO}\lambda^{\nabla}$ is sound for their semantics, but completeness is not addressed. For simplicity, nabla-quantified variables can only be of one type α (Remark 5, loc. cit.) Schöpp later generalized a similar construction, based on the category $\widehat{\mathbf{L}}$ of presheaves over the category $\mathbf{L}$ of so-called λ-tree contexts and substitutions, in the form of categories with binding structure (Schöpp 2007, Part II). That Part II also generalizes Part I of the same paper, where Schöpp defines a Henkin semantics for a classical variant of $\mathrm{FO}\lambda^{\nabla}$, which he proves sound and complete for that semantics. The ∇ quantifier is again restricted there to apply to variables of certain base types called the lambda-tree types. One may also cite Bucalo et al. (2006), who offer a semantics of higher order abstract syntax – a very closely related problem – based on a glueing construction due to Hofmann (1999), and with a unique type υ of names. We shall make a small guided tour of the semantical differences between the categories involved right at the end of Section 3, when we have enough material.

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

In all those proposals, the types of which one can create fresh objects are designated base types. While that makes a comparison with the И quantifier easier, this ignores one distinctive feature of Miller and Tiu's proposal: the possibility of considering fresh objects of *any* type, even higher order, not just *base* types.

Our semantics will address this. All the semantics mentioned earlier, except for Bucalo et al. (2006) and Hofmann (1999), are presheaf semantics, over various categories, and the properties they enjoy are due to general categorical reasons, notably the fact that presheaf categories are toposes, or the crucial use of the Yoneda lemma for completeness at λ-tree types in Schöpp (2007). Our semantics is close to a presheaf semantics, but is not a presheaf semantics, precisely because of the need to interpret ∇ over non-base types. The properties it has do not seem to be due to any general categorical reason, and our development will therefore be elementary.

To be more precise, our semantics will be given in a category $\boldsymbol{\nabla}$ that is close to the presheaf category $\mathbf{Set}^{\mathbb{N}}$, but is sufficiently different that, for example, it has no terminal object. (We shall explain why in Section 3.) The objects of $\mathbf{Set}^{\mathbb{N}}$ are families of sets $(D_n)_{n\in\mathbb{N}}$ together with maps $\mathrm{old}_n^D : D_n \to D_{n+1}$, $n \in \mathbb{N}$. In $\boldsymbol{\nabla}$, we additionally require each D_n to be non-empty, and each old_n^D to be injective and non-surjective. Non-surjectivity allows us to find fresh elements in D_{n+1}, namely elements that are not of the form $\mathrm{old}_n^D(d)$ for any $d \in D_n$. That is required to ensure that our models have *enough maps*, which in turn is necessary for soundness. Injectivity implies that D_n can be considered as a subset of D_{n+1}, which will lead to the simplified Definition 3.1.

For completeness purposes (but not for soundness), we will assume that there is only one base type ι. (Schöpp makes a similar assumption in his Part II.) As we have just said, this will not prevent us from consider fresh objects of any type, as in ∇x^{ι}, $\nabla y^{\iota\to\iota}$, or $\nabla z^{(\iota\to\iota\to\iota)\to\iota\to\iota}$.

We shall be almost, but not completely faithful to Miller and Tiu (2005). First and foremost, our semantics – and our proof rules – will validate the implication of $\exists x.P(x)$ by $\nabla x.P(x)$. That rule is also valid in Abella (Gacek 2008), and also in Miculan and Yemane's (2005) original semantics. It will also validate the rule that $\nabla x.F$ and $\nabla y.F$ are equivalent when x, y are not free in F, even when x and y have different types. However, and conforming to Miller and Tiu (2005), it will not validate the Abella equivalences between $\nabla x.F$ and F when x is not free in F, or between $\nabla x.\nabla y.P(x,y)$ and $\nabla y.\nabla x.P(x,y)$. Second, our logic will be classical, not intuitionistic, as in Schöpp (2007): semantics is easier in a classical setting.

Outline. We introduce the category $\boldsymbol{\nabla}$ of nabla sets in Section 3. This is the basis of our semantics for nabla, of which the most general form is a kind of Henkin semantics (Section 4), including both standard semantics and a term-based semantics. We show that classical $\mathrm{FO}\lambda^{\nabla}$ is sound for all Henkin structures with generic families that admit enough maps in Section 5, and that it is complete in Section 6, provided there is exactly one base type ι. Completeness is obtained for term structures, by using a construction of Hintikka sets. This has many similarities with what Schöpp does (Schöpp 2007, Section 1.6), and indeed rests on principles that have been well established for decades. Our use of Hintikka

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

sets, instead of Henkin sets, also allows us to show that the cut rule can be eliminated. We examine the question of completeness for standard structures in the rest of the paper, and this is more involved. We notice that the logic is in fact incomplete for standard structures in Section 7, by showing that the axiom of choice is true but unprovable; then, we show that the logic is complete, even without the cut rule again, if we restrict ourselves to so-called Π_1 formulae – and that includes the first-order fragment as a special case. We do this by building a specific retraction of the standard universe onto the term universe, which interacts nicely with a natural Kripke logical relation. We list a few open questions in Section 9, and conclude in Section 10.

3. Nabla sets

Our main object of study is the following.

Definition 3.1. A *nabla set* D is a strictly increasing chain of non-empty sets $D_0 \subsetneq D_1 \subsetneq D_2 \subsetneq \cdots \subsetneq D_n \subsetneq D_{n+1} \subsetneq \cdots$. For convenience, we write D_∞ for $\bigcup_{n\in\mathbb{N}} D_n$.

D_n is meant to be the set of values of some type D, in a context where at most n generic values have been created. The objects in $D_{n+1} \smallsetminus D_n$ are meant to be fresh relative to D_n.

Definition 3.2. A *nabla map* f from a nabla set D to a nabla set E is a map $f : D_\infty \to E_\infty$ such that f maps every element of D_n to an element of E_n, for each $n \in \mathbb{N}$.

Nabla sets and nabla maps form a category ∇. An isomorphism between D and E in ∇ is a bijection that restricts to bijections between D_n and E_n for each $n \in \mathbb{N}$.

∇ has products of all non-empty families $(D_i)_{i\in I}$. The canonical product $D = \prod_{i\in I} D_i$ is defined pointwise: $D_n = \prod_{i\in I} D_{\text{in}}$.

Beware that there is no product of the empty family, namely, no terminal object, as the reader will realize by him/herself. That means that one cannot model unit types such as `unit` in ML, or `Unit` in Haskell, or `void` in Java or in C. This is probably unavoidable: recall that our semantics is meant to allow for the creation of generic, or fresh, objects, of any type; then, how should one interpret the creation of a generic or fresh object of a unit type? We do not.

We shall use the following notations:

— For a product $D = \prod_{i\in I} D_i$, $\pi_i : D \to D_i$ is ith projection, defined by $\pi_i(d_j)_{j\in I} = d_i$.
— $D_1 \times D_2$ stands for $\prod_{i=1,2} D_i$.
— For $f_1 : D_1 \to E_1$, $f_2 : D_2 \to E_2$, $f_1 \times f_2 : D_1 \times D_2 \to E_1 \times E_2$ maps (d_1, d_2) to $(f_1(d_1), f_2(d_2))$.
— For $f_1 : D \to E_1$, $f_2 : D \to E_2$, $\langle f_1, f_2\rangle : D \to E_1 \times E_2$ maps d to $(f_1(d), f_2(d))$.

An object D in a category with binary products is *exponentiable* if and only if $\_ \times D$ has a right adjoint $[D \to \_]$ (usually written $\_^D$). Explicitly, D is exponentiable if and only if, for every object E, there is an object $[D \to E]$, a so-called *application* (a.k.a., evaluation) map $\mathsf{App} : [D \to E] \times D \to E$, and for every morphism $f : C \times D \to E$, a *currified* map $\Lambda(f) : C \to [D \to E]$ satisfying the following equations (Curien 1993):

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

— (β-rule) $\mathsf{App} \circ (\Lambda f \times \mathrm{id}_D) = f$ for every $f : C \times D \to E$.
— (η-rule) $\Lambda(\mathsf{App}) = \mathrm{id}_{[D \to E]}$.
— (Substitution rule) $\Lambda f \circ g = \Lambda(f \circ (g \times \mathrm{id}_D))$, for all $f : C \times D \to E$ and $g : B \to C$.

Proposition 3.3. In $\mathbf{\nabla}$, every object is exponentiable.

Proof. Given two nabla sets D and E, we define $[D \to E]_n$ as the set of maps $f : D_\infty \to E_\infty$ such that f maps every element of D_m to E_m, for every $m \geqslant n$. Clearly, $[D \to E]_n \subseteq [D \to E]_{n+1}$.

Pick e_0 from E_0, and e_{n+1} from $E_{n+1} \smallsetminus E_n$ for each $n \in \mathbb{N}$. For each given $n \in \mathbb{N}$, we define a map $f[n]$ from D_∞ to E_∞ by: $f[n](d) = e_n$ for every $d \in D_n$, $f[n](d) = e_{m+1}$ for every $d \in D_{m+1} \smallsetminus D_m$, $m \geqslant n$. Clearly, $f[n]$ is in $[D \to E]_n$. When $n \geqslant 1$, $f[n]$ is not in $[D \to E]_{n-1}$ because $f[n]$ maps the elements of D_{n-1} (and there are some) to e_n, which is not in E_{n-1}. Therefore, the inclusion $[D \to E]_n \subseteq [D \to E]_{n+1}$ is strict. This shows that the chain $[D \to E]$ of sets $[D \to E]_n$ forms a nabla set.

Define App by $\mathsf{App}(f, d) = f(d)$. This is a morphism from $[D \to E] \times D$ to E. For every morphism $f : C \times D \to E$, define $\Lambda f(c)$ as $f(c, \_)$, the map that sends d to $f(c, d)$, for every $c \in C_\infty$. We check that Λf is a morphism from C to $[D \to E]$. For every $n \in \mathbb{N}$, for every $n \in \mathbb{N}$, for every $c \in C_n$, this amounts to say that for every $m \geqslant n$, for every $d \in D_m$, $\Lambda f(c)(d) = f(c, d)$ is in E_m. Indeed, since c is in $C_n \subseteq C_m$.

The β-rule, the η-rule and the substitution rule are now immediate. □

Nabla sets satisfy several forms of the axiom of choice. Let us call *weak nabla subset* of a nabla set D any monotonic sequence $A_0 \subseteq A_1 \subseteq \cdots \subseteq A_n \subseteq \cdots$ of subsets A_n of D_n, $n \in \mathbb{N}$. We do not require the inclusions $A_n \subseteq A_{n+1}$ to be proper. A *weak nabla relation* R between D and E is a weak nabla subset of $D \times E$.

Proposition 3.4 (Choice). Let D and E be two nabla sets, and R be a weak nabla relation between D and E. If, for all $n \in \mathbb{N}$ and $d \in D_n$, there is an $e \in E_n$ such that $(d, e) \in R_n$, then there is a nabla map $f : D \to E$ such that, for all $n \in \mathbb{N}$ and $d \in D_n$, $(d, f(d)) \in R_n$.

Proof. Define f as follows. We use the set-theoretic axiom of choice. For each $d \in D_0$, pick some $e \in E_0$ such that $(d, e) \in R_0$ and define $f(d)$ as e. For every $n \in \mathbb{N}$, define $f(d)$ for $d \in D_{n+1} \smallsetminus D_n$ by picking some $e \in E_{n+1}$ such that $(d, e) \in R_{n+1}$, and letting $f(d)$ be that e.

That defines a nabla map from D to E. For every $d \in D_\infty$, if n is the smallest natural number such that $d \in D_n$, f is built so that $(d, f(d)) \in R_n$. For every $m \geqslant n$, $(d, f(d))$ is also in R_m because $R_n \subseteq R_m$. This shows that $(d, f(d)) \in R_m$ for every $d \in D_m$, for every $m \in \mathbb{N}$. □

That implies the following, which will be our bane in Section 7.

Corollary 3.5 (Weak choice). Let D, E be two nabla sets. Fix $n \in \mathbb{N}$, and let $R \subseteq D_n \times E_n$. If for every $d \in D_n$, there is an $e \in E_n$ such that $(d, e) \in R$, then there is an element f of $[D \to E]_n$ such that for every $d \in D_n$, $(d, f_n(d))$ is in R.

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

Proof. Let $D' = (D_m)_{m \geqslant n}$, $E' = (E_m)_{m \geqslant n}$. Those are again nabla sets, where $D'_k = D_{n+k}$ and $E'_k = E_{n+k}$. For every $k \in \mathbb{N}$, let R'_k be $R_n \cup ((D'_k \smallsetminus D'_0) \times E'_k)$. Since $R_n \subseteq D_n \times E_n = D'_0 \times E'_0 \subseteq D'_k \times E'_k$, and since clearly $R'_k \subseteq R'_{k+1}$, $(R'_k)_{k \in \mathbb{N}}$ is a weak nabla relation.

For each $d \in D'_k$, there is an e such that $(d, e) \in R'_k$: either $d \in D'_0 = D_n$, and we can find e so that $(d, e) \in R_n$ by assumption, or $d \in D'_k \smallsetminus D'_0$, in which case we can pick any e from E'_k, which is non-empty.

We can therefore use Proposition 3.4, and obtain a nabla map $f : D' \to E'$ such that, for every $k \in \mathbb{N}$, for every $d' \in D'_k$, $(d', f(d')) \in R'_k$. We finally observe that f, being a morphism from D' to E', is in $[D \to E]_n$. □

Seemingly related is the following result, which will however be a boon to us: it will be used to show that our semantics of ∇ is sound. This is exactly the place where we require the inclusion $D_n \subseteq D_{n+1}$ to be proper.

Lemma 3.6. *Let D, E be two nabla sets, $n \in \mathbb{N}$, $d \in D_{n+1}$ and $e \in E_{n+1}$. There is a nabla map $f : D \to E$ such that $f(d) = e$. In particular, there is an $f \in [D \to E]_n$ such that $f(d) = e$.*

Proof. Pick g from the set $[D \to E]_0$. That is non-empty, since $[D \to E]$ is a nabla set. By definition, g is a nabla map from D to E. Define f by $f(d) = e$, and $f(x) = g(x)$ for every $x \neq d$. This is again a nabla map, and one which satisfies the desired constraint. □

At this point, it is useful to make a summary of the main properties that distinguish our category of nabla sets from other categories aimed at giving meaning to names. Recall that a topos is a Cartesian-closed category in particular, and satisfies AC!, the axiom of unique choice. Let us write AC for the full axiom of choice.

- Our category $\boldsymbol{\nabla}$ of nabla sets is not Cartesian-closed, because it does not have a terminal object, but every object is exponentiable. It satisfies some forms of AC, as we have already seen, and as we shall formally demonstrate in Lemma 7.1. It also satisfies the formula $\nabla x.\varphi \supset \exists x.\varphi$.
- The category of nominal sets Gabbay and Pitts (1999), also called the Schanuel topos, is a topos, hence satisfies AC!. It does not satisfy AC, but satisfies $\nabla x.\varphi \supset \exists x.\varphi$ (confusing ∇ with И).
- $\widehat{\mathbb{D}}$ (Miculan and Yemane's 2005), $\widehat{\mathbf{L}}$ (Schöpp 2007, Part I) are presheaf toposes.
- The category of contexts (Bucalo et al. 2006) is a tripos, not a topos, and satisfies neither AC nor AC!.

At the risk of repeating ourselves, $\boldsymbol{\nabla}$ will allow us to make sense of ∇ at all types, contrarily to the other proposals.

4. Standard and Henkin semantics for λ-terms

Let us consider simply typed λ-terms M in Church style, that is, all variables x^τ have a preassigned type τ. There are countably infinitely many variables of each type τ. We shall sometimes omit the subscript τ when it is clear. There are base types β (at least one), and other types are formed using the arrow type former $\to$. Explicitly, the (simply typed)

λ-terms are inductively defined by: x^φ is a λ-term of type φ; if M is of type $\varphi \to \tau$ and N is of type φ, then MN is of type τ; if M is of type τ, then $\lambda x^\varphi.M$ is of type $\varphi \to \tau$.

We shall consider λ-terms modulo $\beta\eta$-equivalence (including α-renaming), and we shall often confuse terms for their equivalence classes. In particular, we shall often write $M = N$ to say that M and N are $\beta\eta$-equivalent, although we shall make that explicit when there is a risk of confusion. The $=$ relation respects types.

If M is of type τ, then $\lambda x^\varphi.M$ is of type $\varphi \to \tau$.

Proposition 3.3 allows us to define a *standard semantics* for λ-terms: we fix nabla sets $S[\![\beta]\!]$ for every base type β, define $S[\![\varphi \to \tau]\!]$ as the exponential object $[S[\![\varphi]\!] \to S[\![\tau]\!]]$, inductively; finally, we define the value of applications through App and the value of λ-abstractions through Λ.

There is a more general construction, which we shall need to obtain completeness results in the style of Henkin's completeness theorems for higher order logic. The following is imitated from the notion of typed combinatory algebra, replacing sets by nabla sets. An apt name would be 'typed extensional combinatory nabla algebra,' but that would be lengthy. Similarly, we should write $\mathsf{App}_{\varphi,\tau}$, $\mathsf{k}_{\varphi,\psi}$, etc., below, but we prefer to drop the subscripts.

Definition 4.1 (Henkin universe). A *Henkin universe* S is the following data:

- For each type τ, a nabla-set $S[\![\tau]\!]$.
- A nabla map $\mathsf{App} : S[\![\varphi \to \tau]\!] \times S[\![\varphi]\!] \to S[\![\tau]\!]$, one for each pair of types φ, τ; we shall write $f \cdot x$ for $\mathsf{App}(f, x)$; $\cdot$ associates to the left.
- An element $\mathsf{k} \in S[\![\varphi \to \psi \to \varphi]\!]_0$, one for each pair of types φ, ψ.
- An element $\mathsf{s} \in S[\![(\varphi \to \psi \to \tau) \to (\varphi \to \psi) \to \varphi \to \tau]\!]_0$, one for each triple of types φ, ψ, τ

satisfying

$$\mathsf{k} \cdot a \cdot b = a, \tag{1}$$

$$\mathsf{s} \cdot a \cdot b \cdot c = a \cdot c \cdot (b \cdot c), \tag{2}$$

$$(\forall n \geqslant m, \forall u \in S[\![\varphi]\!]_n, f \cdot a = g \cdot a) \Rightarrow f = g, \tag{3}$$

where $a \in S[\![\varphi]\!]_m$ and $b \in S[\![\psi]\!]_m$, $m \in \mathbb{N}$ in Equation (1), $a \in S[\![\varphi \to \psi \to \tau]\!]_m$, $b \in S[\![\varphi \to \psi]\!]_m$, $c \in S[\![\varphi]\!]_m$, $m \in \mathbb{N}$ in Equation (2) and $f, g \in S[\![\varphi \to \tau]\!]_m$, $m \in \mathbb{N}$ in Equation (3).

A *generic family* new on S is a family of elements $\mathsf{new}^\varphi_{n+1} \in S[\![\varphi]\!]_{n+1} \smallsetminus S[\![\varphi]\!]_n$, one for each type φ and each $n \in \mathbb{N}$.

Formula (3) means that all our Henkin universes are *extensional.*

Remark 4.2. Condition (3), written in a Kripke style, is equivalent to the following:

$$(\forall a \in S[\![\varphi]\!]_\infty, f \cdot a = g \cdot a) \Rightarrow f = g, \tag{4}$$

for all $f, g \in S[\![\varphi \to \tau]\!]_m$, $m \in \mathbb{N}$. The implication (3)$\Rightarrow$(4) is clear. In the converse direction, assume that for every $n \geqslant m$, for every $a \in S[\![\varphi]\!]_n$, $f \cdot a = g \cdot a$. Then, $f \cdot a = g \cdot a$

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

also holds for every $a \in S[\![\varphi]\!]_n$ with $n < m$, since $S[\![\varphi]\!]_n \subseteq S[\![\varphi]\!]_m$ in that case. We can then apply Equation (4) and conclude $f = g$.

Every Henkin universe S gives rise to an interpretation of simply typed λ-terms in the expected way. Because of the way we introduced them, it is practically to make a detour through typed *combinatory terms*: x^φ is a combinatory term of type φ; if M is of type $\varphi \to \tau$ and N is of type φ, then MN is of type τ; and there are constants K of each type of the form $\varphi \supset \psi \supset \varphi$ and S of each type of the form $(\varphi \supset \psi \supset \tau) \supset (\varphi \supset \psi) \supset (\varphi \supset \tau)$.

Write Env for the product $\prod_{x^\tau} S[\![\tau]\!]$, where x^τ ranges over all variables: Env_n is the set of *environments* ρ *at level* n, namely functions mapping each variable x^τ to an element $\rho(x^\tau) \in S[\![\tau]\!]_n$. For each $d \in S[\![\tau]\!]_n$, we write $\rho[x^\tau \mapsto d]$ for the environment that is like ρ except that it maps x^τ to d. We obtain a semantics of (simply typed) combinatory terms $u : \tau$, given as $S[\![u]\!]\rho$, where

$$\begin{aligned} S[\![x]\!]\rho &= \rho(x), \\ S[\![uv]\!]\rho &= S[\![u]\!]\rho \cdot S[\![v]\!]\rho, \\ S[\![\mathsf{K}]\!]\rho = \mathsf{k} & \qquad S[\![\mathsf{S}]\!]\rho = \mathsf{s}. \end{aligned} \tag{5}$$

We obtain a semantics $S[\![M]\!]\rho$ for (simply typed) λ-terms M by letting $S[\![M]\!]\rho = S[\![M^\circ]\!]\rho$, where $M \mapsto M^\circ$ is the familiar translation from λ-terms M to combinatory terms M°: $x^\circ = x$, $(MN)^\circ = M^\circ N^\circ$, $(\lambda x.M)^\circ = [x]M^\circ$, where for each combinatory term u, $[x]u$ is defined by $[x]x = \mathsf{SKK}$, $[x]u = \mathsf{K}u$ if x is not free in u, $[x](vw) = \mathsf{S}([x]v)([x]w)$ if x is free in vw.

The following three lemmas are standard.

Lemma 4.3. Let S be a Henkin universe. For every $n \in \mathbb{N}$, for every environment ρ at level n, for every simply typed combinatory term u,

1. for every $\rho' \in \mathsf{Env}_n$ that coincides with ρ on the free variables of u, $S[\![u]\!]\rho = S[\![u]\!]\rho'$;
2. for every $d \in S[\![\varphi]\!]_n$, $S[\![[x^\varphi]u]\!]\rho \cdot d = S[\![u]\!](\rho[x^\varphi \mapsto d])$.

Proof. (1) is obvious. (2) is proved by induction on u. When $u = x^\varphi$, $S[\![[x^\varphi]x]\!]\rho \cdot d = S[\![\mathsf{SKK}]\!]\rho \cdot d = \mathsf{s} \cdot \mathsf{k} \cdot \mathsf{k} \cdot d = \mathsf{k} \cdot d \cdot (\mathsf{k} \cdot d) = d = S[\![x^\varphi]\!](\rho[x^\varphi \mapsto d])$. When x^φ is not free in u, $S[\![[x^\varphi]u]\!]\rho \cdot d = S[\![\mathsf{K}u]\!]\rho \cdot d = \mathsf{k} \cdot S[\![u]\!]\rho \cdot d = S[\![u]\!]\rho = S[\![u]\!](\rho[x^\varphi \mapsto d])$, using (1). When $u = vw$ and x^φ is free in u, $S[\![[x^\varphi]u]\!]\rho \cdot d = S[\![\mathsf{S}([x^\varphi]v)([x^\varphi]w)]\!]\rho \cdot d = \mathsf{s} \cdot S[\![[x^\varphi]v]\!]\rho \cdot S[\![[x^\varphi]w]\!]\rho \cdot d = S[\![[x^\varphi]v]\!]\rho \cdot d \cdot (S[\![[x^\varphi]w]\!]\rho \cdot d) = S[\![v]\!](\rho[x^\varphi \mapsto d]) \cdot S[\![w]\!](\rho[x^\varphi \mapsto d]) = S[\![u]\!](\rho[x^\varphi \mapsto d])$, using the induction hypothesis in the next-to-last equality. □

Lemma 4.4. Let S be a Henkin universe. For every $n \in \mathbb{N}$, for every environment ρ at level n, the following holds, and characterizes the semantics $S[\![M]\!]\rho$ of simply typed λ-terms:

1. $S[\![x]\!]\rho = \rho(x)$.
2. $S[\![MN]\!]\rho = S[\![M]\!]\rho \cdot S[\![N]\!]\rho$.
3. $S[\![\lambda x^\varphi.M]\!]\rho$, where $M : \tau$, is the unique $f \in S[\![\varphi \to \tau]\!]_n$ such that for every $m \geqslant n$, for every $d \in S[\![\varphi]\!]_m$, $f \cdot d = S[\![M]\!](\rho[x^\varphi \mapsto d])$.

Proof. (1) and (2) are obvious. For (3), we first check that $S[\![\lambda x^\varphi.M]\!]\rho \cdot d = S[\![M]\!](\rho[x^\varphi \mapsto d])$. Indeed $S[\![\lambda x^\varphi.M]\!]\rho \cdot d = S[\![[x^\varphi]M^\circ]\!]\rho \cdot d = S[\![M^\circ]\!](\rho[x^\varphi \mapsto d])$ by Lemma 4.3 (2),

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

and that is equal to $S[\![M]\!](\rho[x^\varphi \mapsto d])$ by definition. For uniqueness, imagine there is another $f \in S[\![\varphi \to \tau]\!]_n$ such that for every $m \geqslant n$, for every $d \in S[\![\varphi]\!]_m$, $f \cdot d = S[\![M]\!](\rho[x^\varphi \mapsto d])$, equivalently $f \cdot d = S[\![\lambda x^\varphi.M]\!]\rho \cdot d$. Equation (3) then implies $f = S[\![\lambda x^\varphi.M]\!]\rho$. □

Lemma 4.5. Let S be a Henkin universe. The following hold:

1. For all λ-terms $N : \tau$ and $M : \varphi$, for every $n \in \mathbb{N}$, for every environment ρ at level n, $S[\![N[M/x^\varphi]]\!]\rho = S[\![N]\!](\rho[x^\varphi \mapsto S[\![M]\!]\rho])$.
2. For every λ-term $M : \tau$, $S[\![M]\!]\rho$ does not depend on $\rho(y)$ if y is not free in M, namely, if $\rho(z) = \rho'(z)$ for every $z \neq y$, then $S[\![M]\!]\rho = S[\![M]\!]\rho'$.
3. For all $\beta\eta$-convertible λ-terms $M, N : \tau$, $S[\![M]\!]\rho = S[\![N]\!]\rho$.

Proof. 1. It is well known that $(N[M/x^\varphi])^\circ = N^\circ[M^\circ/x^\varphi]$, and it is easy to check that $S[\![u[v/x^\varphi]]\!]\rho = S[\![u]\!](\rho[x^\varphi \mapsto S[\![v]\!]\rho])$ for all combinatory terms u, v of the right type, from which (1) follows.

2. is by Lemma 4.3 (1), realizing that M and M° have the same free variables.

3. We do this in several steps:

— (β) $S[\![(\lambda x^\varphi.M)N]\!]\rho = S[\![M]\!](\rho[x^\varphi \mapsto S[\![N]\!]\rho])$ by Lemma 4.4 (3), and this is equal to $S[\![M[N/x^\varphi]]\!]\rho$ by (1) above.

— (η) We check that $S[\![\lambda x^\varphi.Mx^\varphi]\!]\rho = S[\![M]\!]\rho$, where x^φ is not free in M. For every $m \geqslant n$, for every $d \in S[\![\varphi]\!]_m$, $S[\![\lambda x^\varphi.Mx^\varphi]\!]\rho \cdot d = S[\![Mx^\varphi]\!](\rho[x^\varphi \mapsto d])$ (by Lemma 4.4 (3)) $= S[\![M]\!](\rho[x^\varphi \mapsto d]) \cdot d$ (by Lemma 4.4 (2)) $= S[\![M]\!]\rho \cdot d$ (by (2) above). By the uniqueness part of Lemma 4.4 (3), $S[\![\lambda x^\varphi.Mx^\varphi]\!]\rho = S[\![M]\!]\rho$.

— (ξ) If $S[\![M]\!](\rho[x^\varphi \mapsto d]) = S[\![N]\!](\rho[x^\varphi \mapsto d])$ for every $m \geqslant n$ and $d \in S[\![\varphi]\!]_m$, then $S[\![\lambda x^\varphi.M]\!]\rho = S[\![\lambda x^\varphi.N]\!]\rho$. This is a direct consequence of the uniqueness part of Lemma 4.4 (3).

— If $S[\![M_1]\!]\rho = S[\![M_2]\!]\rho$ and $S[\![N_1]\!]\rho = S[\![N_2]\!]\rho$, then $S[\![M_1N_1]\!]\rho = S[\![M_2N_2]\!]\rho$. This is by Lemma 4.4 (2).

If M reduces to N in one $\beta\eta$-reduction step, then $S[\![M]\!]\rho = S[\![N]\!]\rho$, by induction on the depth of the contracted redex, using the remarks above. (3) follows. □

The next definition is particular to our setting, and adapts Lemma 3.6 to Henkin universes. This will be needed for soundness. We purposefully require f to be in $S[\![\varphi \to \tau]\!]_n$, not in the larger set $S[\![\varphi \to \tau]\!]_{n+1}$.

Definition 4.6 (Enough maps). A Henkin universe S for nabla *has enough maps* with respect to a generic family new if and only if, for all types φ and τ, for every $n \in \mathbb{N}$, for every $d \in S[\![\tau]\!]_{n+1}$, there is an $f \in S[\![\varphi \to \tau]\!]_n$ such that $f \cdot \mathsf{new}^\varphi_{n+1} = d$.

4.1. *Standard universes*

Lemma 4.7 (Standard universe). Given nabla sets D_β, one for each base type β, there is a Henkin universe S such that

— $\mathsf{S}[\![\beta]\!] = D_\beta$ for each base type β, and $\mathsf{S}[\![\varphi \to \tau]\!] = [\mathsf{S}[\![\varphi]\!] \to \mathsf{S}[\![\tau]\!]]$ for all types φ, τ;

— App is the application morphism $(f, x) \mapsto f(x)$ in $\boldsymbol{\nabla}$;

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

— $\mathsf{k} : u \mapsto (v \mapsto u)$;
— $\mathsf{s} : u \mapsto (v \mapsto (w \mapsto u(w)(v(w))))$.

This Henkin universe S has enough maps with respect to every generic family.

We call S the *standard universe* on the nabla sets D_τ.

Proof. That it is a Henkin universe is obvious, except perhaps for Equation (3). We use the equivalent Equation (4): if $f, g \in S[\![\varphi \to \tau]\!]_m$ are such that $f(u) = g(u)$ for every $u \in S[\![\varphi]\!]_\infty$, then $f = g$ as maps from $S[\![\varphi]\!]_\infty$ to $S[\![\tau]\!]_\infty$, hence as nabla maps from $S[\![\varphi]\!]$ to $S[\![\tau]\!]$.

We now claim that S has enough maps with respect to any generic family new. Fix $d \in \mathsf{S}[\![\tau]\!]_{n+1}$. By Lemma 3.6, there is an $f \in [\mathsf{S}[\![\varphi]\!] \to \mathsf{S}[\![\tau]\!]]_n$ such that $f(\mathsf{new}^\varphi_{n+1}) = d$. □

4.2. *The term universe*

We now exhibit another Henkin universe T, built from syntax. This will be useful to show completeness. For that, we assume there is exactly *one* base type ι. T is built from a variant of the λ-calculus we have considered until now, and which we call the λ-calculus *with names*. The variant has infinitely many new constants a_i, $i \geqslant 1$, of type ι, called *names*. Those names are pairwise distinct; a_i is the name *at level i*.

Explicitly, the *(simply typed) λ-terms with names* are defined inductively by: every variable x^τ is a λ-term with names, of type τ; every name a_i is a λ-term with names, of type ι; if M is a λ-term with names of type $\varphi \to \tau$ and N is a λ-term with names of type φ, then MN is a λ-term with names of type τ; if M is a λ-term with names of type τ, and x^φ is a variable, then $\lambda x^\varphi.M$ is a λ-term with names of type $\varphi \to \tau$. Note that names cannot occur bound, since they are not variables.

We consider λ-terms with names modulo $\beta\eta$-conversion (including α-conversion), and by this we mean that a λ-term with names is shorthand for its β-normal η-long form. This convention allows us to make sense of the notions of free variables, and of free names, of a λ-term with names.

Definition 4.8. For each type τ, for every $n \in \mathbb{N}$, $T[\![\tau]\!]_n$ is the set of all λ-terms with names of type τ (up to $\beta\eta$-conversion) in which the only free names are of the form a_i with $1 \leqslant i \leqslant n$.

We let App be syntactic application, namely $M \cdot N = MN$, and define k as $\lambda x^\varphi.\lambda y^\psi.x^\varphi$, s as $\lambda x^{\varphi\to\psi\to\tau}.\lambda y^{\varphi\to\psi}.\lambda z^\varphi.xz(yz)$.

Assume there is exactly one base type ι. The generic family a is defined by $\mathsf{a}^\tau_{n+1} = \lambda x_1^{\tau_1}.\lambda x_2^{\tau_2}.\cdots.\lambda x_m^{\tau_m}.a_{n+1}$, where $\tau = \tau_1 \to \tau_2 \to \cdots \to \tau_m \to \iota$, and $x_1^{\tau_1}$, $x_2^{\tau_2}$, $\ldots x_m^{\tau_m}$ are distinct fresh variables.

In the definition of a^τ_n, we use the fact that every type τ is of the form $\tau_1 \to \tau_2 \to \cdots \to \tau_m \to \beta$, where β is a base type. The assumption that there is a unique base type ι allows us to conclude that $\beta = \iota$.

Remark 4.9. For every type τ, $T[\![\tau]\!]_0$ is just the set of ordinary, not with names, λ-terms of type τ, modulo $\beta\eta$-conversion.

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

Lemma 4.10. T, as defined in Definition 4.8, is a Henkin universe. If there is exactly one base type ι, then T has enough maps with respect to the generic family a.

Proof. Again the fact that this is a Henkin universe is obvious, except perhaps for Equation (3), for which the argument is nonetheless standard. Assume that M, N are in $T[\![\varphi \to \tau]\!]_m$, and that $MP = NP$ (modulo $\beta\eta$) for every $P \in T[\![\varphi]\!]_n$, $n \geqslant m$. Let P be some fresh variable X^φ, and note that this is in $[\![\varphi]\!]_0 \subseteq [\![\varphi]\!]_n$. Then, $MX^\varphi = NX^\varphi$, from which we obtain $M = \lambda X^\varphi.MX^\varphi = \lambda X^\varphi.NX^\varphi = N$, using the η rule.

The enough maps property has to be given some care. Let $N \in T[\![\tau]\!]_{n+1}$. We wish to find an $M \in T[\![\varphi \to \tau]\!]_n$ such that $M\mathsf{a}^\varphi_{n+1} = N$ (up to $\beta\eta$-conversion).

Write φ is a unique way as $\varphi_1 \to \varphi_2 \to \cdots \to \varphi_m \to \iota$, and pick some arbitrary λ-terms $M_1 : \varphi_1$, $M_2 : \varphi_2$, $\ldots$, $M_m : \varphi_m$ – variables, for example. Build a new term $\widetilde{N}$ by replacing all occurrences of a_{n+1} in N by the term $x^\varphi M_1 M_2 \cdots M_m$, where x^φ is a fresh variable of type φ. Finally, define M as $\lambda x^\varphi.\widetilde{N}$. The only names a_i that occur free in M are such that $1 \leqslant i \leqslant n$, by construction, so M is in $T[\![\varphi \to \tau]\!]_n$, and $M\mathsf{a}^\varphi_{n+1} = M(\lambda x_1^{\varphi_1}.\lambda x_2^{\varphi_2}.\cdots.\lambda x_m^{\varphi_m}.a_{n+1}) = \widetilde{N}[\lambda x_1^{\varphi_1}.\lambda x_2^{\varphi_2}.\cdots.\lambda x_m^{\varphi_m}.a_{n+1}/x^\varphi] = N$. □

For any set of variables A, a *substitution* θ at level n of domain A is any function that maps every variable z^ψ to an element of $T[\![\psi]\!]_n$. When A is finite, we define the capture-avoiding application $M\theta$ of θ to the λ-term M in the usual way.

If θ and θ' agree on the set of free variables of M, then $M\theta = M\theta'$. We can therefore extend the notation $M\theta$ to substitutions θ of arbitrary (not necessarily finite) domains, by defining $M\theta$ as $M\theta_{|A}$, where A is any finite subset containing the free variables of M.

Such substitutions at level n are none other than the environments at level n in the Henkin universe T.

Lemma 4.11. For every λ-term $M : \tau$, for every substitution θ at level n, $T[\![M]\!]\theta = M\theta$, up to $\beta\eta$-equivalence.

Proof. By structural induction on M, using Lemma 4.4. Only the case of abstractions $M = \lambda x^\varphi.N$, $N : \tau$, is interesting. Let A be a finite subset containing the free variables of M. By α-renaming, we make sure that x^φ is not in A, and not free in any term $\theta(y)$, $y \in A$. By Lemma 4.4 (3), $T[\![M]\!]\theta$ is the unique $f \in T[\![\varphi \to \tau]\!]_n$ such that for every $m \geqslant n$, for every $P \in T[\![\varphi]\!]_m$, $f \cdot P = T[\![N]\!](\theta[x^\varphi \mapsto P])$, namely $f \cdot P = N(\theta[x^\varphi \mapsto P])$, using the induction hypothesis. The term $M\theta$ (modulo $\beta\eta$) is another such f, since $M\theta \cdot P = (\lambda x^\varphi.N\theta)P = N\theta[P/x^\varphi] = N(\theta[x^\varphi \mapsto P])$. By uniqueness, $T[\![M]\!]\theta = M\theta$. □

5. A semantics for FOλ^∇ and soundness

The logic FOλ^∇ was introduced by Miller and Tiu (2005), as an intuitionistic first-order logic with predicates on higher order terms, together with the ∇ operator. Schöpp (2007) used a classical variant of that logic. We use a close cousin of the latter: the only differences are that $\nabla x^\tau.F$ will imply $\exists x^\tau.F$ in our logic, and that $\nabla x^\tau.F$ and $\nabla y^\varphi.F$ will be equivalent if x^τ and y^φ are not free in F, even when $\tau \neq \varphi$.

Instead of considering all the connectives, we shall restrict ourselves to $\perp$ (false), $\supset$ (implication) and $\forall$ (universal quantification). The other connectives could be dealt

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

$$\frac{}{\Gamma, (\sigma \rhd \bot) \longrightarrow \Delta}\,(\bot L) \qquad \frac{}{\Gamma, J \longrightarrow J, \Delta}\,(Ax) \qquad \frac{\Gamma \longrightarrow J, \Delta \quad \Gamma', J \longrightarrow \Delta'}{\Gamma, \Gamma' \longrightarrow \Delta, \Delta'}\,(Cut)$$

$$\frac{\Gamma, J, J \to \Delta}{\Gamma, J \to \Delta}\,(cL) \qquad \frac{\Gamma \to \Delta}{\Gamma, J \to \Delta}\,(wL) \qquad \frac{\Gamma \to \Delta, J, J}{\Gamma \to \Delta, J}\,(cR) \qquad \frac{\Gamma \to \Delta}{\Gamma \to \Delta, J}\,(wR)$$

$$\frac{\Gamma, J \to \Delta}{\Gamma, J' \to \Delta}\,(J \approx J') \quad (\approx L) \qquad \frac{\Gamma \to \Delta, J}{\Gamma \to \Delta, J'}\,(J \approx J') \quad (\approx R)$$

$$\frac{\Gamma \longrightarrow \Delta, (\sigma \rhd F) \quad \Gamma, (\sigma \rhd G) \longrightarrow \Delta}{\Gamma, (\sigma \rhd F \supset G) \longrightarrow \Delta}\,(\supset L) \qquad \frac{\Gamma, (\sigma \rhd F) \longrightarrow \Delta, (\sigma \rhd G)}{\Gamma \longrightarrow \Delta, (\sigma \rhd F \supset G)}\,(\supset R)$$

$$\frac{M : \tau \quad \Gamma, (\sigma \rhd F[M/x^\tau]) \longrightarrow \Delta}{\Gamma, (\sigma \rhd \forall x^\tau.F) \longrightarrow \Delta}\,(\forall L) \qquad \frac{\Gamma \longrightarrow \Delta, (\sigma \rhd F[h\sigma/x^\tau])}{\Gamma \longrightarrow \Delta, (\sigma \rhd \forall x^\tau.F)}\,(h^{\sigma \to \tau} \text{ fresh}) \quad (\forall R)$$

$$\frac{\Gamma, (\sigma, x : \tau \rhd F) \longrightarrow \Delta}{\Gamma, (\sigma \rhd \nabla x^\tau.F) \longrightarrow \Delta}\,(\nabla L) \qquad \frac{\Gamma \longrightarrow \Delta, (\sigma, x : \tau \rhd F)}{\Gamma \longrightarrow \Delta, (\sigma \rhd \nabla x^\tau.F)}\,(\nabla R)$$

Fig. 1. A sequent calculus formulation of FOλ^∇.

with similarly. Alternatively, in a classical logic, those other connectives are definable: $\neg F = F \supset \bot$, $F \vee G = (\neg F) \supset G$, $F \wedge G = \neg(F \supset \neg G)$, $\exists x^\tau.F = \neg(\forall x^\tau.\neg F)$.

We are given a countable set of so-called *relation symbols* P, each coming with an *arity*, which is a finite list of types $\tau_1, \tau_2, \cdots, \tau_k$. Atomic formulae are of the form $P(M_1, M_2, \cdots, M_k)$ where $M_1 : \tau_1$, $M_2 : \tau_2$, ..., $M_k : \tau_k$ are λ-terms and P is a relation symbol of arity $\tau_1, \tau_2, \cdots, \tau_k$. The formulae are built from atomic formulae and $\bot$ using $\supset$, $\forall$ and the *nabla quantifier* ∇: if F is a formula, then $\nabla x^\tau.F$ is a formula.

Call a *signature* any finite list σ of pairwise distinct variables $x_1^{\tau_1}, x_2^{\tau_2}, \cdots, x_m^{\tau_m}$. To stick with conventional writing, we shall write that signature $x_1 : \tau_1, x_2 : \tau_2, \cdots, x_m : \tau_m$.

A *generic judgement* (or, more simply, a *judgement*) J is an expression of the form $\sigma \rhd F$ where σ is a signature (the *local signature* of the judgement) and F is a formula; $\rhd$ is a separator. The meaning of $x_1 : \tau_1, x_2 : \tau_2, \cdots, x_m : \tau_m \rhd F$ is intended to be the same as $\nabla x_1^{\tau_1}.\nabla x_2^{\tau_2}.\cdots.\nabla x_m^{\tau_m}.F$. We write $\lambda\sigma.F$ for $\lambda x_1 : \tau_1, x_2 : \tau_2, \cdots, x_m : \tau_m.F$. We also write σ, σ' for the concatenation of signatures when that makes sense.

Definition 5.1. Let $\approx$ be the smallest equivalence relation on judgements such that

— if $\lambda\sigma.F$ and $\lambda\sigma'.F'$ are $\beta\eta$-convertible, then $(\sigma \rhd F) \approx (\sigma' \rhd F')$;
— if x^τ and y^φ are not free in F, then $(\sigma, x : \tau, \sigma' \rhd F) \approx (\sigma, y : \varphi, \sigma' \rhd F)$.

The second requirement does not follow from the first one: when $\tau \neq \varphi$, $\lambda\sigma, x : \tau, \sigma'.F$ and $\lambda\sigma, y : \varphi, \sigma'.F$ have different types and are neither α-convertible nor $\beta\eta$-convertible.

A *sequent* of FOλ^∇ is an expression $\Gamma \longrightarrow \Delta$, where Γ and Δ are finite multisets of judgements.

Remark 5.2. Those are slightly different from the sequents of Miller and Tiu (2005), which are of the form $\Sigma; \Gamma \longrightarrow \Delta$, where Σ is a (global) signature. This makes a difference in our way of formulating the $(\forall L)$ rule, which allows us to instantiate x^τ by any term of type τ whatsoever, including non-ground terms; hence, to prove the implication $\forall x^\tau.F \supset \nabla x^\tau.F$, and therefore also (since ∇ commutes with negation), $\nabla x^\tau.F \supset \exists x^\tau.F$.

We write Γ, J for the addition of the judgement J to Γ, and Γ, Θ for the union of the multisets Γ and Θ. We write $M : \tau$ to state that M is a term of type τ, as in the first premise of $(\forall L)$.

The rules of $\mathrm{FO}\lambda^\nabla$ are shown in Figure 5. In the rightmost premise of $(\forall L)$, one can find a judgement $\sigma \rhd F[M/x^\tau]$. $F[M/x^\tau]$ denotes capture-avoiding substitution of M for x^τ in F, but M *is* allowed to capture variables from σ, on purpose. In $(\forall R)$, $h : \sigma \to \tau$ abbreviates $h : \tau_1 \to \tau_2 \to \cdots \to \tau_n \to \tau$, and $h\sigma$ abbreviates $hx_1x_2\cdots x_n$.

We define a semantics of all the objects considered above, as follows.

Definition 5.3. Given a nabla set D, let a *nabla predicate* P on D be a family $(P_n)_{n\in\mathbb{N}}$ of subsets P_n of D_n.

Nabla predicates are not weak nabla subsets: we do not require that P_n be included in P_{n+1}.

Definition 5.4. A *Henkin structure* is a Henkin universe S, together with nabla predicates $S[\![P]\!]$ on $S[\![\tau_1]\!] \times S[\![\tau_2]\!] \times \cdots \times S[\![\tau_k]\!]$ for each relation symbol P of arity $\tau_1, \tau_2, \cdots, \tau_k$.

A *standard structure* is a Henkin structure whose underlying Henkin universe is a standard universe S (see Lemma 4.7).

We now define satisfaction of a formula F at level n as follows, in a Henkin structure S, modulo a generic family new and where ρ is an environment at level n.

$$\begin{aligned}
S, \mathsf{new}; \rho \models_n P(M_1, \cdots, M_k) &\text{ iff } (S[\![M_1]\!]_n(\rho), \cdots, S[\![M_k]\!]_n(\rho)) \in S[\![P]\!]_n \\
S, \mathsf{new}; \rho \models_n \bot &\quad \text{never} \\
S, \mathsf{new}; \rho \models_n F \supset G &\text{ iff } (S, \mathsf{new}; \rho \not\models_n F \text{ or } S, \mathsf{new}; \rho \models_n G) \\
S, \mathsf{new}; \rho \models_n \forall x^\tau.F &\text{ iff } (\text{for every } d \in S[\![\tau]\!]_n,\ S, \mathsf{new}; \rho[x \mapsto d] \models_n F) \\
S, \mathsf{new}; \rho \models_n \nabla x^\tau.F &\text{ iff } S, \mathsf{new}; \rho[x \mapsto \mathsf{new}^\tau_{n+1}] \models_{n+1} F.
\end{aligned}$$

This extends to judgements by letting $S, \mathsf{new}; \rho \models_n x_1 : \tau_1, x_2 : \tau_2, \cdots, x_m : \tau_m \rhd F$ if and only if $S, \mathsf{new}; \rho \models_n \nabla x_1^{\tau_1}.\nabla x_2^{\tau_2}.\cdots.\nabla x_m^{\tau_m}.F$; then, to sequents by letting $S, \mathsf{new}; \rho \models_n \Gamma \longrightarrow \Delta$ if and only if $S, \mathsf{new}; \rho \not\models_n J$ for some J in Γ or $S, \mathsf{new}; \rho \models_n J$ for some J in Δ.

Remark 5.5. It may be worth comparing our semantics to Schöpp (2007). Instead of an index n, Schöpp uses a so-called λ-tree context σ, i.e., a signature where each variable is mapped to a λ-tree type ι (a certain class of base types). The connection is that n is the length of σ. Schöpp does not require a generic family new as we do to give semantics to $\nabla x^\iota.F$ (in some environment ρ), and instead decides to map the generic variable x^ι to itself: explicitly, Schöpp's semantics of $\nabla x^\iota.F$ in environment ρ and λ-tree context σ is the semantics of F in environment $\rho[x^\iota \mapsto x^\iota]$ (and context $\sigma, x^\iota : \iota$). This is only possible because Schöpp's interpretation of terms at a base type ι is restricted to be a term of the

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

same type (see the definition of $||\iota||\sigma$ in Schöpp (2007), Section 1.3): so x^ι is not only a variable, but also a value of the right type, and $\rho[x^\iota \mapsto x^\iota]$ makes sense. We do not make such a restriction, but soundness does not come for free: we require our models to have enough maps for that (see the proof of Lemma 5.7 (3) below).

Lemma 5.6. For every λ-term M of type τ, for every $n \in \mathbb{N}$,

1. $S, \mathsf{new}; \rho \models_n J[M/x^\tau]$ iff $S, \mathsf{new}; \rho[x^\tau \mapsto S[\![M]\!]\rho] \models_n J$;
2. $S, \mathsf{new}; \rho \models_n \Gamma[M/x^\tau] \longrightarrow \Delta[M/x^\tau]$ iff $S, \mathsf{new}; \rho[x^\tau \mapsto S[\![M]\!]\rho] \models_n \Gamma \longrightarrow \Delta$.

Proof. (1) It is enough to prove the claim when J is a formula, by structural induction on it, paying attention to α-renaming in the case of universal quantification and ∇ quantification. We describe the latter case, when $J = \nabla y^\varphi.F$. By α-renaming, y^φ is different from x^τ and not free in M. Write $\widetilde{\rho}$ for $\rho[y^\varphi \mapsto \mathsf{new}^\varphi_{n+1}]$. Then, writing $\rho \models_n J$ instead of the more formal but more cumbersome $S, \mathsf{new}; \rho \models_n J$:

$$\begin{aligned}
\rho \models_n J \ \text{ iff } & \ \widetilde{\rho} \models_{n+1} F[M/x^\tau] \\
\text{iff } & \ \rho[y^\varphi \mapsto \mathsf{new}^\varphi_{n+1}, x^\tau \mapsto S[\![M]\!]\widetilde{\rho}] \models_{n+1} F \qquad \text{(induction hypothesis)} \\
\text{iff } & \ \rho[y^\varphi \mapsto \mathsf{new}^\varphi_{n+1}, x^\tau \mapsto S[\![M]\!]\rho] \models_{n+1} F \qquad \text{(Lemma 4.5 (2))} \\
\text{iff } & \ \rho[x^\tau \mapsto S[\![M]\!]\rho][y^\varphi \mapsto \mathsf{new}^\varphi_{n+1}] \models_{n+1} F \\
\text{iff } & \ \rho[x^\tau \mapsto S[\![M]\!]\rho] \models_n \nabla y^\varphi.F.
\end{aligned}$$

(2) Immediate consequence of (1). □

We say that two formulae F and G are *equivalent* if and only if, for every Henkin structure S and for every generic family new on S, with enough maps, for every $n \in \mathbb{N}$, for every environment ρ at level n, $S, \mathsf{new}; \rho \models_n F$ if and only if $S, \mathsf{new}; \rho \models_n G$.

Lemma 5.7. The following are pairs of equivalent formulae:

1. $\nabla x^\tau.(F \supset G)$ and $(\nabla x^\tau.F) \supset (\nabla x^\tau.G)$.
2. $\nabla x^\tau.F$ and $\nabla y^\varphi.F$, if neither x^τ nor y^φ is free in F.
3. $\nabla x^\tau.\forall y^\varphi.F$ and $\forall h^{\tau\to\varphi}.\nabla x^\tau.F[hx/y]$.

Proof. The first equivalence is a simple verification. The second one follows from the fact that the semantics of a formula F in an environment ρ does not depend on the values $\rho(z^\psi)$ such that z^ψ is not free in F. This an easy induction on F, which uses Lemma 4.5 (2).

Finally, for the third equivalence (writing again $\rho \models_n F$ instead of $S, \mathsf{new}; \rho \models_n F$),

$$\begin{aligned}
\rho \models_n \nabla x^\tau.\forall y^\varphi.F \quad \text{iff} \quad & \rho[x \mapsto \mathsf{new}^\tau_{n+1}] \models_{n+1} \forall y^\varphi.F \\
\text{iff} \quad & (\text{for every } d \in S[\![\varphi]\!]_{n+1},\ \rho[x \mapsto \mathsf{new}^\tau_{n+1}, y \mapsto d] \models_{n+1} F) \qquad (6)
\end{aligned}$$

while $\rho \models_n \forall h^{\tau\to\varphi}.\nabla x^\tau.F[hx/y]$ if and only if

$$\begin{aligned}
& (\text{for every } f \in S[\![\tau \to \varphi]\!]_n,\ \rho[h \mapsto f] \models_n \nabla x^\tau.F[hx/y]) \\
\text{iff} \quad & (\text{for every } f \in S[\![\tau \to \varphi]\!]_n,\ \rho[h \mapsto f][x \mapsto \mathsf{new}^\tau_{n+1}] \models_{n+1} F[hx/y]) \\
\text{iff} \quad & (\text{for every } f \in S[\![\tau \to \varphi]\!]_n, \qquad (7) \\
& \quad \rho[x \mapsto \mathsf{new}^\tau_{n+1}, y \mapsto \mathsf{App}(f, \mathsf{new}^\tau_{n+1})] \models_{n+1} F)
\end{aligned}$$

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

where we have used Lemma 5.6 (2), and the fact that h is not free in F in the last line. The two are equivalent: in one direction, for every $f \in S[\![\tau \to \varphi]\!]_n$, $\mathsf{App}(f, \mathsf{new}^{\tau}_{n+1})$ is a value d in $S[\![\varphi]\!]_{n+1}$, so Equation (6) implies Equation (7). In the converse direction, for every $d \in S[\![\varphi]\!]_{n+1}$, we can find an $f \in S[\![\tau \to \varphi]\!]_n$ such that $\mathsf{App}(f, \mathsf{new}^{\tau}_{n+1}) = d$, because S has enough maps. Hence, Equation (7) implies Equation (6). □

We write $S, \mathsf{new} \models_n \Gamma \longrightarrow \Delta$ if and only if $S, \mathsf{new}; \rho \models_n \Gamma \longrightarrow \Delta$ for every environment ρ at level n, and we say that $\Gamma \longrightarrow \Delta$ is *valid* if and only if this holds for every $n \in \mathbb{N}$, for every Henkin structure S, and for every generic family new, with enough maps.

Proposition 5.8 (Soundness). Every derivable sequent $\Gamma \longrightarrow \Delta$ is valid.

Proof. It suffices to show that $S, \mathsf{new}; \rho \models_n \Gamma \longrightarrow \Delta$ by induction on the given derivation. Again, we drop the S, new prefix, in the name of readability.

In the case of the $(\supset L)/(\supset R)$ rules, we must show that $\rho \models_n \sigma \rhd (F \supset G)$ if and only if $\rho \not\models_n \sigma \rhd F$ or $\rho \models_n \sigma \rhd G$: this is an easy induction on the number of variables in σ, using Lemma 5.7 (1).

In the case of $(\approx L)/(\approx R)$, we must show that $\rho \models_n J$ if and only if $\rho \models_n J'$, assuming $J \approx J'$. It suffices to show that this is the case when J and J' are $\beta\eta$-convertible (that follows from Lemma 4.5 (3)), and when $J = \sigma, x : \tau, \sigma' \rhd F$, $J' = \sigma, y : \varphi, \sigma' \rhd F$, with x^{τ}, y^{φ} not free in F; the latter follows from Lemma 5.7 (2).

In the case of $(\forall R)$, assume that $\rho \models_n \Gamma \longrightarrow \Delta, (\sigma \rhd F[h\sigma/x^{\tau}])$, with h fresh of type $\sigma \to \tau$. Equivalently, $\rho \models_n \Gamma \longrightarrow \Delta, (\rhd \nabla\sigma.F[h\sigma/x^{\tau}])$, where we write $\nabla\sigma$ for $\nabla x_1^{\tau_1}.\nabla x_2^{\tau_2}.\cdots.\nabla x_m^{\tau_m}$, assuming $\sigma = x_1 : \tau_1, x_2 : \tau_2, \cdots, x_m : \tau_m$. Trivially, this implies $\rho \models_n \Gamma \longrightarrow \Delta, (\rhd \forall h^{\sigma\to\tau}.\nabla\sigma.F[h\sigma/x^{\tau}])$, since h is fresh. By iterating Lemma 5.7 (3), we obtain $\rho \models_n \Gamma \longrightarrow \Delta, (\rhd \nabla\sigma.\forall x^{\tau}.F)$, that is, $\rho \models_n \Gamma \longrightarrow \Delta, (\sigma \rhd \forall x^{\tau}.F)$.

In the case of $(\forall L)$, let M be a λ-term of type τ, and assume $\rho_n \models_n \Gamma, (\sigma \rhd F[M/x^{\tau}]) \longrightarrow \Delta$. Assume also that $\rho \models_n J$ for every J in Γ, and $\rho \models_n (\sigma \rhd \forall x^{\tau}.F)$. We aim to show that $\rho \models_n J'$ for some J' in Δ. By Lemma 5.7 (3) again, the latter implies $\rho \models_n \forall h^{\sigma\to\tau}.\nabla\sigma.F[h\sigma/x^{\tau}]$. Instantiate $h^{\sigma\to\tau}$ by $\lambda\sigma.M$. It follows that $\rho \models_n \nabla\sigma.F[M/x^{\tau}]$, hence $\rho \models_n \sigma \rhd F[M/x^{\tau}]$. Since $\rho \models_n J$ for every J in Γ and $\rho_n \models_n \Gamma, (\sigma \rhd F[M/x^{\tau}]) \longrightarrow \Delta$, we conclude.

The other cases are immediate. □

6. Henkin completeness

We shall show that the deduction system of Figure 5 is complete for Henkin structures (under the assumption of a unique base type ι) using a variant of the technique of Hintikka sets, a technique used to show that tableaux calculi are complete for first-order logic (Fitting 1996, Section 3.5). This will also show that the (Cut) rule is not needed for completeness.

Our purpose now is, given an unprovable sequent, to find a model of it.

A *signed judgement* is an expression of the form $+J$ or $-J$, where J is a judgement. Semantically, we understand $+J$ as meaning 'J is true,' and $-J$ as 'J is false.' Syntactically, we see a sequent $J_1, \cdots, J_m \to J'_1, \cdots, J'_n$ as a collection of signed judgements $+J_1, \cdots, +J_m, -J'_1, \cdots, -J'_n$. We extend $\approx$ to signed judgements in the obvious way.

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

Definition 6.1. A *theory* $\mathcal{T}$ is a set of signed judgements.

$\mathcal{T}$ is *inconsistent* if and only if there are finitely many signed judgements $+J_1, \ldots, +J_m$, $-J'_1, \ldots, -J'_n$ in $\mathcal{T}$ such that the sequent $J_1, \cdots, J_m \rightarrow J'_1, \cdots, J'_n$ is derivable in the system of Figure 5, using all rules except the cut rule (Cut). $\mathcal{T}$ is *consistent* otherwise.

$\mathcal{T}$ is a *Hintikka theory* if and only if

1. $\mathcal{T}$ is consistent;
2. if $J \in \mathcal{T}$ and $J \approx J'$, then $J' \in \mathcal{T}$;
3. if $+\sigma \triangleright F \supset G$ is in $\mathcal{T}$, then $-\sigma \triangleright F$ or $+\sigma \triangleright G$ is in $\mathcal{T}$;
4. if $-\sigma \triangleright F \supset G$ is in $\mathcal{T}$, then both $+\sigma \triangleright F$ and $-\sigma \triangleright G$ are in $\mathcal{T}$;
5. if $+\sigma \triangleright \forall x^\tau.F$ is in $\mathcal{T}$, then $+\sigma \triangleright F[M/x^\tau]$ is in $\mathcal{T}$ for every λ-term $M : \tau$;
6. if $-\sigma \triangleright \forall x^\tau.F$ is in $\mathcal{T}$, then $-\sigma \triangleright F[h\sigma/x^\tau]$ is in $\mathcal{T}$ for some variable $h^{\sigma\rightarrow\tau}$ that does not occur in σ;
7. if $+\sigma \triangleright \nabla x^\tau.F$ is in $\mathcal{T}$, then $+\sigma, x : \tau \triangleright F$ is in $\mathcal{T}$;
8. if $-\sigma \triangleright \nabla x^\tau.F$ is in $\mathcal{T}$, then $-\sigma, x : \tau \triangleright F$ is in $\mathcal{T}$.

Fact 1. A consistent theory cannot contain both $+J$ and $-J$ for the same judgement J; otherwise, it would be inconsistent, using rule (Ax). It cannot contain a judgement of the form $+\sigma \triangleright \bot$ either (rule $(\bot L)$).

Lemma 6.2. Let $\mathcal{T}$ be a consistent theory.

1. For every signed judgement $+\sigma \triangleright F \supset G$ in $\mathcal{T}$, $\mathcal{T} \cup \{-\sigma \triangleright F\}$ or $\mathcal{T} \cup \{+\sigma \triangleright G\}$ is consistent.
2. For every signed judgement $-\sigma \triangleright F \supset G$ in $\mathcal{T}$, $\mathcal{T} \cup \{+\sigma \triangleright F, -\sigma \triangleright G\}$ is consistent.
3. For every signed judgement $+\sigma \triangleright \forall x^\tau.F$ in $\mathcal{T}$, for every $M : \tau$, $\mathcal{T} \cup \{+\sigma \triangleright F[M/x^\tau]\}$ is consistent.
4. For every signed judgement $-\sigma \triangleright \forall x^\tau.F$ in $\mathcal{T}$, for every variable $h : \sigma \rightarrow \tau$ that is not free in $\mathcal{T}$ and does not occur in σ, $\mathcal{T} \cup \{-\sigma \triangleright F[h\sigma/x^\tau]\}$ is consistent.
5. For every signed judgement $+\sigma \triangleright \nabla x^\tau.F$ in $\mathcal{T}$, $\mathcal{T} \cup \{+\sigma, x : \tau \triangleright F\}$ is consistent.
6. For every signed judgement $-\sigma \triangleright \nabla x^\tau.F$ in $\mathcal{T}$, $\mathcal{T} \cup \{-\sigma, x : \tau \triangleright F\}$ is consistent.
7. For every signed judgement $+J$ in $\mathcal{T}$, for every $J' \approx J$, $\mathcal{T} \cup \{+J'\}$ is consistent.
8. For every signed judgement $-J$ in $\mathcal{T}$, for every $J' \approx J$, $\mathcal{T} \cup \{-J'\}$ is consistent.

Proof. (1) Assume that both $\mathcal{T} \cup \{-\sigma \triangleright F\}$ and $\mathcal{T} \cup \{+\sigma \triangleright G\}$ are inconsistent. There are cut-free derivations of sequent of the form $\Gamma \rightarrow \underbrace{(\sigma \triangleright F)}_{m \text{ times}}, \Delta$ and $\Gamma', \underbrace{(\sigma \triangleright G)}_{n \text{ times}} \rightarrow \Delta'$, where Γ and Γ' consist of judgements that appear with the $+$ sign in $\mathcal{T}$, Δ and Δ' consist of judgements that appear with the $-$ sign in $\mathcal{T}$, and $m, n \in \mathbb{N}$. Necessarily, $m \neq 0$ since otherwise $\mathcal{T}$ would be inconsistent. Using the contraction rule (cR), we may assume that $m = 1$. Similarly, and using (cL), we may assume that $n = 1$. Using the weakening rules (wL) and (wR), we may assume that $\Gamma = \Gamma'$ and $\Delta = \Delta'$. It now suffices to apply $(\forall L)$ to obtain a cut-free derivation of $\Gamma, (\sigma \triangleright F \supset G) \rightarrow \Delta$. However, $+\sigma \triangleright F \supset G$ is in $\mathcal{T}$, so that contradicts the consistency of $\mathcal{T}$.

(2)–(8). Similar analysis, using rule $(\supset R)$, $(\forall L)$, $(\forall R)$, (∇L), (∇R), $(\approx L)$ or $(\approx R)$ instead. □

Lemma 6.3. Every finite consistent theory is contained in some Hintikka theory.

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

Proof. Since there are only countably many variables and countably many relation symbols, there are only countably many λ-terms (up to $\beta\eta$-conversion), and countably many signed judgements. Call a *task* either: a signed judgement $\pm J$, where J is not of the form $+\sigma \rhd \forall x^\tau.F$; or a pair $(+\sigma \rhd \forall x^\tau.F, M)$ where $M : \tau$; or a pair $(+J, +J')$ or $(-J, -J')$ with $J \approx J'$. Fix an enumeration of all tasks, in such a way that every task occurs infinitely often on the list. The latter is a standard trick, and we shall explain its purpose at the end of the proof.

Let $\mathcal{T}_0$ be a finite consistent theory. We define an increasing sequence of finite consistent theories $\mathcal{T}_n$, $n \in \mathbb{N}$, starting with $\mathcal{T}_0$. Given that $\mathcal{T}_n$ has been built, we build $\mathcal{T}_{n+1}$ by considering the nth task Θ_n on the enumeration.

If Θ_n is of the form $+\sigma \rhd F \supset G$, and is in $\mathcal{T}_n$, then by Lemma 6.2 (1), $\mathcal{T}_n \cup \{-\sigma \rhd F\}$ or $\mathcal{T}_n \cup \{+\sigma \rhd G\}$ is consistent: in the first case, let $\mathcal{T}_{n+1} = \mathcal{T}_n \cup \{-\sigma \rhd F\}$, otherwise let $\mathcal{T}_{n+1} = \mathcal{T}_n \cup \{+\sigma \rhd G\}$. If $\Theta_n = +\sigma \rhd F \supset G$ is not in $\mathcal{T}_n$, then $\mathcal{T}_{n+1} = \mathcal{T}_n$.

If Θ_n is of the form $-\sigma \rhd F \supset G$ and is in $\mathcal{T}_n$, then we let $\mathcal{T}_{n+1} = \mathcal{T}_n \cup \{+\sigma \rhd F, -\sigma \rhd G\}$, using Lemma 6.2 (2). And if $\Theta_n = -\sigma \rhd F \supset G$ is not in $\mathcal{T}_n$, then $\mathcal{T}_{n+1} = \mathcal{T}_n$.

We proceed similarly if Θ_n is of the form $\pm\sigma \rhd \nabla x^\tau.F$, using Lemma 6.2 (5) or (6).

If Θ_n is of the form $(+\sigma \rhd \forall x^\tau.F, M)$ where $+\sigma \rhd \forall x^\tau.F$ is in $\mathcal{T}$, and M is of type τ, then we let $\mathcal{T}_{n+1} = \mathcal{T}_n \cup \{+\sigma \rhd F[M_n/x^\tau]\}$, using Lemma 6.2 (3). If $+\sigma \rhd \forall x^\tau.F$ is not in $\mathcal{T}$, then we let $\mathcal{T}_{n+1} = \mathcal{T}_n$.

If Θ_n is of the form $-\sigma \rhd \forall x^\tau.F$ and is in $\mathcal{T}_n$, then there is a variable h of type $\sigma \to \tau$ that is free neither in $\mathcal{T}_n$ nor in σ since $\mathcal{T}_n$ is finite. Relying on Lemma 6.2 (4), we define $\mathcal{T}_{n+1}$ as $\mathcal{T}_n \cup \{-\sigma \rhd F[h\sigma/x^\tau]\}$. If $\Theta_n = -\sigma \rhd \forall x^\tau.F$ is not in $\mathcal{T}_n$, then $\mathcal{T}_{n+1} = \mathcal{T}_n$.

Finally, if Θ_n is of the form $(+J, +J')$ with $J \approx J'$ (and similarly if it is of the form $(-J, -J')$), either $+J \in \mathcal{T}_n$ and we let $\mathcal{T}_{n+1} = \mathcal{T}_n \cup \{+J'\}$, relying on Lemma 6.2 (7) and (8), or $+J \notin \mathcal{T}_n$ and we let $\mathcal{T}_{n+1} = \mathcal{T}_n$.

Define $\mathcal{T}_\infty$ as $\bigcup_{n \in \mathbb{N}} \mathcal{T}_n$. $\mathcal{T}_\infty$ is a Hintikka theory, as one checks easily. For example, if $+\sigma \rhd \forall x^\tau.F$ is in $\mathcal{T}_\infty$, it must occur in $\mathcal{T}_n$ for some $n \in \mathbb{N}$. Since each task appears infinitely often in the enumeration, for every λ-term $M : \tau$, the task $(+\sigma \rhd \forall x^\tau.F, M)$ occurs at some rank m after n. Our construction then ensures that $+\sigma \rhd F[M/x^\tau]$ is in $\mathcal{T}_{m+1}$, hence in $\mathcal{T}_\infty$. □

Now consider the term universe T of Section 4.2. Recall that it only makes sense provided there is a unique base type ι. For every local signature $\sigma = x_1 : \tau_1, x_2 : \tau_2, \ldots, x_n : \tau_n$ (of *length* n), let θ_σ be the substitution $[\mathsf{a}_1^{\tau_1}/x_1, \mathsf{a}_2^{\tau_2}/x_2, \ldots, \mathsf{a}_n^{\tau_n}/x_n]$. This is a substitution at level n.

Lemma 6.4. Let $\mathcal{T}$ be a Hintikka theory, and assume there is a unique base type ι. Define $T[\![P]\!]_n$, for each relation symbol P, of arity $\tau_1, \tau_2, \ldots, \tau_k$, as the set of k-tuples $(M_1\theta_\sigma, M_2\theta_\sigma, \ldots, M_k\theta_\sigma)$ such that $+\sigma \rhd P(M_1, M_2, \ldots, M_k) \in \mathcal{T}$ for some local signature σ of length n. This defines a Henkin structure such that

1. for every signed judgement $+J \in \mathcal{T}$, $T, \mathsf{a}; \epsilon \models_0 J$,
2. for every signed judgement $-J \in \mathcal{T}$, $T, \mathsf{a}; \epsilon \not\models_0 J$,

where ϵ is the identity substitution (at level 0).

Proof. First look at the case where $J = \sigma \rhd P(M_1, M_2, \ldots, M_k)$, where σ is of length n. If $+J \in \mathcal{T}$, then by definition $(M_1\theta_\sigma, M_2\theta_\sigma, \ldots, M_k\theta_\sigma)$ is in $T\llbracket P \rrbracket_n$. By Lemma 4.11, $(T\llbracket M_1 \rrbracket \theta_\sigma, T\llbracket M_2 \rrbracket \theta_\sigma, \ldots, T\llbracket M_k \rrbracket \theta_\sigma)$ is in $T\llbracket P \rrbracket_n$, so $T, \mathsf{a}; \theta_\sigma \models_n P(M_1, M_2, \ldots, M_k)$, namely $T, \mathsf{a}; \epsilon \models_0 \sigma \rhd P(M_1, M_2, \ldots, M_k)$. If $-J \in \mathcal{T}$, then $+J \notin \mathcal{T}$ (Fact 1), so $(M_1\theta_\sigma, M_2\theta_\sigma, \ldots, M_k\theta_\sigma)$ is not in $T\llbracket P \rrbracket_n$. By a similar argument, $T, \mathsf{a}; \epsilon \not\models_0 \sigma \rhd P(M_1, M_2, \ldots, M_k)$.

Now assume $J = \sigma \rhd \bot$. Since every Hintikka theory is consistent, and using Fact 1, $+J$ is not in $\mathcal{T}$. If $-J$ is in $\mathcal{T}$, we have $T, \mathsf{a}; \epsilon \not\models_0 \sigma \rhd \bot$ anyway.

The case where $J = \sigma \rhd F \supset G$ presents no difficulty. If $+J \in \mathcal{T}$, then $-\sigma \rhd F$ or $+\sigma \rhd G$ is in $\mathcal{T}$, hence by induction hypothesis $T, \mathsf{a}; \epsilon \not\models_0 \sigma \rhd F$ or $T, \mathsf{a}; \epsilon \models_0 \sigma \rhd G$, meaning that $T, \mathsf{a}; \epsilon \models_0 \sigma \rhd F \supset G$. If $-J \in \mathcal{T}$, then $+\sigma \rhd F$ and $-\sigma \rhd G$ are in $\mathcal{T}$, so by induction hypothesis $T, \mathsf{a}; \epsilon \models_0 \sigma \rhd F$ and $T, \mathsf{a}; \epsilon \not\models_0 \sigma \rhd G$, meaning that $T, \mathsf{a}; \epsilon \not\models_0 \sigma \rhd F \supset G$.

Now assume $J = \forall x^\tau.F$. If $+J \in \mathcal{T}$, then $+\sigma \rhd F[M/x^\tau]$ is in $\mathcal{T}$ for every λ-term $M : \tau$. By induction hypothesis, this implies that $T, \mathsf{a}; \epsilon \models_0 \sigma \rhd F[M/x^\tau]$ for every λ-term $M : \tau$. We wish to show that $T, \mathsf{a}; \epsilon \models_0 \sigma \rhd \forall x^\tau.F$. Using Lemma 5.7 (3), we know that the latter is equivalent to $T, \mathsf{a}; \epsilon \models_0 (\rhd \forall h^{\sigma \to \tau}.\nabla \sigma.F[h\sigma/x^\tau])$. Hence, we must show that for every $N \in T\llbracket \sigma \to \tau \rrbracket_0$ (i.e., for every ordinary λ-term $N : \sigma \to \tau$, by Remark 4.9), $T, \mathsf{a}; \epsilon[h \mapsto N] \models_0 \nabla \sigma.F[h\sigma/x^\tau]$. Using Lemma 5.6, and since $T\llbracket N \rrbracket \epsilon = N$ (Lemma 4.11), this boils down to showing that $T, \mathsf{a}; \epsilon \models_n (\nabla \sigma.F[h\sigma/x^\tau])[h \mapsto N]$, that is, $T, \mathsf{a}; \epsilon \models_n \sigma \rhd F[N\sigma/x^\tau]$ for every $N : \sigma \to \tau$ that has no free variable in the list σ. Since $T, \mathsf{a}; \epsilon \models_0 \sigma \rhd F[M/x^\tau]$ for every λ-term $M : \tau$, this is clear.

If $-J \in \mathcal{T}$ for $J = \forall x^\tau.F$, then $-\sigma \rhd F[h\sigma/x^\tau]$ is in $\mathcal{T}$ for some variable $h : \sigma \to \tau$ that does not occur in σ. By induction hypothesis, $T, \mathsf{a}; \epsilon \not\models_0 \sigma \rhd F[h\sigma/x^\tau]$. We wish to show that $T, \mathsf{a}; \epsilon \not\models_0 \sigma \rhd \forall x^\tau.F$, and using the same machinery as above, this is equivalent to showing that $T, \mathsf{a}; \epsilon \not\models_n \sigma \rhd F[N\sigma/x^\tau]$ for some $N : \sigma \to \tau$ that has no free variable in the list σ: we simply take $N = h$.

The cases when $J = \nabla x^\tau.F$ are easy. □

Call any Henkin structure H whose underlying Henkin universe is the term universe T a *Herbrand structure*.

Proposition 6.5. Assume there is a unique base type ι. Let $\Gamma \to \Delta$ be a sequent such that $H, \mathsf{a}; \epsilon \models_0 \Gamma \to \Delta$ for every Herbrand structure H. Then, $\Gamma \to \Delta$ is derivable using the rules of $\mathrm{FO}\lambda^\nabla$, without (Cut).

Proof. Assume $\Gamma \to \Delta$ is not derivable. Let $\mathcal{T}_0$ be the theory containing the signed judgements $+J$, $J \in \Gamma$ and $-J$, $J \in \Delta$. If $\mathcal{T}_0$ were inconsistent, then using the contraction and weakening rules, we would obtain a derivation of $\Gamma \to \Delta$. Therefore, $\mathcal{T}_0$ is consistent. By Lemma 6.3, $\mathcal{T}_0$ is contained in some Hintikka theory $\mathcal{T}$. Using the Henkin structure H defined in Lemma 6.4 – this is a Herbrand structure – we obtain that $H, \mathsf{a}; \epsilon \not\models_0 \Gamma \to \Delta$, a contradiction. □

As a corollary, we obtain:

Theorem 1 (Henkin completeness). Assume there is a unique base type ι. $\mathrm{FO}\lambda^\nabla$ is complete for the Henkin semantics: every valid sequent is derivable in $\mathrm{FO}\lambda^\nabla$, and even by a cut-free proof. □

7. Incompleteness for standard structures

We shall see that $\mathrm{FO}\lambda^{\nabla}$ is *incomplete* for standard structures. This is due to the higher order nature of the terms that $\mathrm{FO}\lambda^{\nabla}$ is based on, and to the fact that the category $\boldsymbol{\nabla}$ validates the weak axiom of choice (Corollary 3.5).

Consider the formula

$$(\forall x^{\varphi}.\exists y^{\tau}.F) \supset (\exists h^{\varphi\to\tau}.\forall x^{\varphi}.F[hx/y]), \tag{AC}$$

where $\exists z^{\psi}.G$ abbreviates $\neg\forall z^{\psi}.\neg G$ and $\neg G$ abbreviates $G \supset \bot$. Explicitly,

$$S, \mathsf{new}; \rho \models_n \exists z^{\psi}.F \text{ iff (for some } e \in S[\![\psi]\!]_n,\ S, \mathsf{new}; \rho[z \mapsto e] \models_n F).$$

Lemma 7.1. (AC) holds in every standard structure S, for every generic family new.

Proof. Assume that $S, \mathsf{new}; \rho \models_n \forall x^{\varphi}.\exists y^{\tau}.F$, in other words, for every $d \in S[\![\varphi]\!]_n$, there is an $e \in S[\![\tau]\!]_n$ such that $S, \mathsf{new}; \rho[x \mapsto d, y \mapsto e] \models_n F$. Let $R \subseteq S[\![\varphi]\!]_n \times S[\![\tau]\!]_n$ be the set of all pairs (d, e) such that $S, \mathsf{new}; \rho[x \mapsto d, y \mapsto e] \models_n F$. Corollary 3.5 applies, so there is an element f of $S[\![\varphi \to \tau]\!]_n$ such that for every $d \in S[\![\varphi]\!]_n$, $S, \mathsf{new}; \rho[x \mapsto d, y \mapsto f(d)] \models_n F$. In other words, $S, \mathsf{new}; \rho \models_n \exists h^{\varphi\to\tau}.\forall x^{\tau}.F[hx/y]$. □

However, (AC) is not provable in $\mathrm{FO}\lambda^{\nabla}$. The following states it for the instance of (AC) where $F = P(x, y)$ and $\varphi = \tau = \iota$.

Lemma 7.2. The sequent $\to \triangleright(\forall x^{\iota}.\exists y^{\iota}.P(x, y)) \supset (\exists h^{\iota\to\iota}.\forall x^{\iota}.P(x, hx))$ is not derivable using the rules of Figure 5.

Proof. We build a Herbrand structure by a diagonal argument. For each $n \in \mathbb{N}$, since $T[\![\iota \to \iota]\!]_n$ is countably infinite, we can enumerate its elements as M_j, $j \in \mathbb{N}$. Enumerate the elements of $T[\![\iota]\!]_n$ as N_j, $j \in \mathbb{N}$, as well. Define $T[\![P]\!]_n \subseteq T[\![\iota]\!]_n \times T[\![\iota]\!]_n$ to be a set of pairs (N_j, N'_j), $j \in \mathbb{N}$, where for each $j \in \mathbb{N}$, N'_j is chosen so as to be different from $M_j N_j$ (remembering that all the terms involved are considered up to $\beta\eta$-conversion). By construction, $T; \epsilon \models_0 \forall x^{\iota}.\exists y^{\iota}.P(x, y)$, but $T; \epsilon \not\models_0 \exists h^{\iota\to\iota}.\forall x^{\iota}.P(x, hx)$, since the latter would mean that there is an element M_j of $T[\![\iota \to \iota]\!]_0$ such that $(N_k, M_j N_k)$ would be in $T[\![P]\!]_0$ for every $k \in \mathbb{N}$, and that fails for $k = j$. We conclude by using Proposition 5.8. □

As a consequence, $\mathrm{FO}\lambda^{\nabla}$ is incomplete for standard structures.

8. Π_1-completeness

We claim that we regain completeness for the fragment consisting of Π_1 formulae (which we define in Definition 8.7 below). This requires some λ-calculus machinery to relate the interpretation $\mathsf{S}[\![M]\!]$ of terms M in a standard universe S and the interpretation $T[\![M]\!]$ in the term universe T.

The standard universe S we choose is the unique standard universe such that $\mathsf{S}[\![\beta]\!] = T[\![\beta]\!]$ for every base type β (see Lemma 4.7). Beware that $\mathsf{S}[\![\tau]\!] = T[\![\tau]\!]$ will fail for non-base types τ: for arrow types, $T[\![\varphi \to \tau]\!]$ is a nabla set of terms, in particular, $T[\![\beta \to \beta]\!]_n$

is countable for every n; on the contrary, $\mathsf{S}[\![\varphi \to \tau]\!] = [\mathsf{S}[\![\varphi]\!] \to \mathsf{S}[\![\tau]\!]]$, and in particular, $\mathsf{S}[\![\beta \to \beta]\!]_0$ is uncountable.

We define the following Kripke logical relation.

Definition 8.1. Define the relations $R[\tau]_n$, for each type τ and each $n \in \mathbb{N}$, between $T[\![\tau]\!]_n$ and $\mathsf{S}[\![\tau]\!]_n$, by

1. $R[\beta]_n$ is equality, for each base type β and every $n \in \mathbb{N}$;
2. for every $n \in \mathbb{N}$, for every $M \in T[\![\varphi \to \tau]\!]_n$, for every $f \in \mathsf{S}[\![\varphi \to \tau]\!]_n$, $M\ R[\varphi \to \tau]_n\ f$ if and only if, for every $m \geqslant n$, for all $N \in T[\![\varphi]\!]_m$ and $d \in \mathsf{S}[\![\varphi]\!]_m$ such that $N\ R[\varphi]_m\ d$, $MN\ R[\tau]_m\ f(d)$.

Call a weak nabla subset $A_0 \subseteq A_1 \subseteq \cdots \subseteq A_n \subseteq \cdots$ of a nabla set D a *nabla subset* of D if and only if $A_{n+1} \cap D_n = A_n$ for every $n \in \mathbb{N}$. A nabla subset is entirely determined by $A_\infty = \bigcup_{n\in\mathbb{N}} A_n$, since we can recover A_n as $A_\infty \cap D_n$. We define *nabla relations* between two nabla sets D and E as the nabla subsets of $D \times E$.

Lemma 8.2. For every type τ, $R[\tau]$ is a nabla relation.

Proof. By induction on τ. This is obvious when τ is a base type. Assume that $R[\varphi]$ and $R[\tau]$ are nabla relations, and let us show that $R[\varphi \to \tau]$ is a nabla relation. Fix $n \in \mathbb{N}$, $M \in T[\![\varphi \to \tau]\!]_n$ and $f \in \mathsf{S}[\![\varphi \to \tau]\!]_n$.

If $M\ R[\varphi \to \tau]_n\ f$, then for every $m \geqslant n$, for all $N \in T[\![\varphi]\!]_m$ and $d \in \mathsf{S}[\![\varphi]\!]_m$ such that $N\ R[\varphi]_m\ d$, $MN\ R[\tau]_m\ f(d)$. This holds in particular for $m \geqslant n+1$, so $M\ R[\varphi \to \tau]_{n+1}\ f$.

Conversely, assume $M\ R[\varphi \to \tau]_{n+1}\ f$. In order to show that $M\ R[\varphi \to \tau]_n\ f$, let $m \geqslant n$, and $N \in T[\![\varphi]\!]_m$ and $d \in \mathsf{S}[\![\varphi]\!]_m$ be such that $N\ R[\varphi]_m\ d$. We wish to show that $MN\ R[\tau]_m\ f(d)$. If $m \geqslant n+1$, this is the assumption. If $m = n$, we use the fact that $N\ R[\varphi]_n\ d$ and the fact that $R[\varphi]$ is a weak nabla relation to obtain $N\ R[\varphi]_{n+1}\ d$. By assumption, $MN\ R[\tau]_{n+1}\ f(d)$. Since MN is in $T[\![\tau]\!]_n$ and $f(d)$ is in $\mathsf{S}[\![\tau]\!]_n$, we can appeal to the fact that $R[\tau]$ is a nabla relation and infer that $MN\ R[\tau]_n\ f(d)$. □

The main result on logical relations is the so-called *basic lemma*, which we now state and prove, in a nabla set-theoretic variant. The argument is standard.

Lemma 8.3 (Basic lemma). For every $n \in \mathbb{N}$, for every substitution θ at level n whose domain $\operatorname{dom}\theta$ is finite, for every environment ρ at level n, we say that $\theta\ R_n\ \rho$ if and only if for every variable $z^\psi \in \operatorname{dom}\theta$, $\theta(z^\psi)\ R[\psi]_n\ \rho(z^\psi)$.

For every λ-term $M : \tau$ with free variables in $\operatorname{dom}\theta$, $\theta\ R_n\ \rho$ implies $M\theta\ R[\tau]_n\ \mathsf{S}[\![M]\!]\rho$.

Beware that M is an ordinary λ-term here, not a λ-term with names.

Proof. By induction on a typing derivation for M. This is clear for variables. If M is an application M_1M_2 with $M_1 : \varphi \to \tau$ and $M_2 : \varphi$, then the induction hypothesis tells us that $M_1\theta\ R[\varphi \to \tau]_n\ f$, where $f = \mathsf{S}[\![M_1]\!]\rho$. It also tells us that $M_2\theta\ R[\varphi]_n\ \mathsf{S}[\![M_2]\!]\rho$. Using the definition of $R[\varphi \to \tau]_n$ with $m = n$, we obtain that $M_1M_2\ R[\tau]_n\ f(\mathsf{S}[\![M_2]\!]\rho) = \mathsf{S}[\![M_1M_2]\!]\rho$.

If M is a λ-abstraction $\lambda x^\varphi.P$ of type $\varphi \to \tau$, then let $f = \mathsf{S}[\![M]\!]\rho$. We must show that, for every $m \geqslant n$, for all $N \in T[\![\varphi]\!]_m$ and $d \in \mathsf{S}[\![\varphi]\!]_m$ such that $N\ R[\varphi]_m\ d$, $(M\theta)N\ R[\tau]_m\ f(d)$.

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

By α-renaming, we may assume that x^φ is not in $\operatorname{dom}\theta$, and not free in any term $\theta(x^\psi)$, $x^\psi \in \operatorname{dom}\theta$. Let $\theta' = \theta[x^\varphi \mapsto N]$ and $\rho' = \rho[x^\varphi \mapsto d]$. Those are at level m, not n. We see that for every variable $z^\psi \in \operatorname{dom}\theta'$, $\theta'(z^\psi)\ R[\psi]_m\ \rho'(z^\psi)$: this follows from $N\ R[\varphi]_m\ d$ when $z^\psi = x^\varphi$, and from $\theta\ R_n\ \rho$ (hence $\theta\ R_m\ \rho$) in the other cases.

By induction hypothesis, $P\theta'\ R[\tau]_m\ \mathsf{S}[\![P]\!]\rho'$. We conclude by noting that $P\theta'$ is equal (up to $\beta\eta$-conversion) to $(M\theta)N$, and that $\mathsf{S}[\![P]\!]\rho' = f(d)$. □

The following is the crux of our argument, and will be used in nearly every forthcoming result.

Proposition 8.4. There are families of nabla maps $s_\tau : T[\![\tau]\!] \to \mathsf{S}[\![\tau]\!]$ and $r_\tau : \mathsf{S}[\![\tau]\!] \to T[\![\tau]\!]$, indexed by types τ, such that the following implications hold for all $M \in T[\![\tau]\!]_n$ and $d \in \mathsf{S}[\![\tau]\!]_n$, $n \in \mathbb{N}$:

$$s_\tau(M) = d \;\Rightarrow\; M\ R[\tau]_n\ d, \tag{8}$$

$$M\ R[\tau]_n\ d \;\Rightarrow\; r_\tau(d) = M. \tag{9}$$

Proof. Those are built by structural induction on τ. For a base type β, we define both s_β and r_β as identities. We define $s_{\varphi\to\tau}$ as $\Lambda(\widetilde{s}_{\varphi\to\tau})$, where $\widetilde{s}_{\varphi\to\tau}$ is the following composition:

$$T[\![\varphi \to \tau]\!] \times \mathsf{S}[\![\varphi]\!] \xrightarrow{\mathrm{id}_{T[\![\varphi\to\tau]\!]} \times r_\varphi} T[\![\varphi \to \tau]\!] \times T[\![\varphi]\!] \xrightarrow{\mathsf{App}} T[\![\tau]\!] \xrightarrow{s_\tau} \mathsf{S}[\![\tau]\!].$$

Here, $\mathsf{App} : T[\![\varphi \to \tau]\!] \times T[\![\varphi]\!] \to T[\![\tau]\!]$ is the nabla map defined by letting $\mathsf{App}(M, N)$ be the term MN (modulo $\beta\eta$); this is application in the term structure. Using the fact that s_τ and r_φ are nabla maps by induction hypothesis, $s_{\varphi\to\tau}$ is a nabla map.

We must show that (8) holds at type $\varphi \to \tau$, that is, that for every $M \in T[\![\varphi \to \tau]\!]_n$ and for $f = s_{\varphi\to\tau}(M) \in \mathsf{S}[\![\varphi \to \tau]\!]_n$, $M\ R[\varphi \to \tau]_n\ f$. To show this, let $m \geqslant n$, and N and d be such that $N\ R[\varphi]_m\ d$. We must show that $MN\ R[\tau]_m\ f(d)$. Since $f = s_{\varphi\to\tau}(M)$, f maps d to $\widetilde{s}_{\varphi\to\tau}(M, d)$, namely, to $s_\tau(\mathsf{App}(M, r_\varphi(d))) = s_\tau(M(r_\varphi(d)))$, where the application of M to $r_\varphi(d)$ is syntactic application. Since $N\ R[\varphi]_m\ d$, by induction hypothesis on φ, $r_\varphi(d) = N$, so $f_m(d) = s_\tau(MN)$. By induction hypothesis on τ, $MN\ R[\tau]_m\ f(d)$.

In order to build $r_{\varphi\to\tau}$, we recall that $R[\varphi \to \tau]$ is a nabla relation (Lemma 8.2), hence is entirely determined by $R[\varphi \to \tau]_\infty = \bigcup_{n\in\mathbb{N}} R[\varphi \to \tau]_n$, in the sense that $R[\varphi \to \tau]_n = R[\varphi \to \tau]_\infty \cap (T[\![\varphi \to \tau]\!]_n \times \mathsf{S}[\![\varphi \to \tau]\!]_n)$.

We claim that for every $f \in \mathsf{S}[\![\varphi \to \tau]\!]_\infty$, there is at most one element $M \in T[\![\varphi \to \tau]\!]_\infty$ such that $M\ R[\varphi \to \tau]_\infty\ f$. Imagine there are two, M_1 and M_2. By abuse of language, consider M_1 and M_2 as terms. Find natural numbers n, m_1, m_2 such that $f \in \mathsf{S}[\![\varphi \to \tau]\!]_n$, $M_1 \in T[\![\varphi \to \tau]\!]_{m_1}$ and $M_2 \in T[\![\varphi \to \tau]\!]_{m_2}$. Let $m = \max(n, m_1, m_2)$, so that f is in $\mathsf{S}[\![\varphi \to \tau]\!]_m$, M_1 and M_2 are in $T[\![\varphi \to \tau]\!]_m$. Since $R[\varphi \to \tau]$ is a nabla relation, we have $M_1\ R[\varphi \to \tau]_m\ f$ and $M_2\ R[\varphi \to \tau]_m\ f$. Pick a variable X^φ that is not free in M_1, and not free in M_2. Let $d = s_\varphi(X^\varphi)$. By induction hypothesis, $X^\varphi\ R[\varphi]_m\ d$, so $M_1X^\varphi\ R[\tau]_m\ f(d)$ and $M_2X^\varphi\ R[\tau]_m\ f(d)$. By induction hypothesis again, $r_\tau(f(d))$ is then equal to both M_1X^φ and to M_2X^φ (up to $\beta\eta$-conversion). Therefore, $\lambda X^\varphi.M_1X^\varphi = \lambda X^\varphi.M_2X^\varphi$, and by η-conversion, $M_1 = M_2$.

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

Using the above claim, we may now define $r_{\varphi\to\tau}(f)$, for each $f \in \mathsf{S}[\![\varphi \to \tau]\!]_\infty$, as the unique $M \in T[\![\varphi \to \tau]\!]_\infty$ such that $M\ R[\varphi \to \tau]_\infty\ f$, if it exists, and as $z^{\varphi\to\tau}$ otherwise, where $z^{\varphi\to\tau}$ is a fixed variable.

It is clear that Equation (9) holds, but the fact that $r_{\varphi\to\tau}(f)$ is a nabla map needs some verification. Let $f \in \mathsf{S}[\![\varphi \to \tau]\!]_n$ for some $n \in \mathbb{N}$. If $r_{\varphi\to\tau}(f) = z^{\varphi\to\tau}$, then that is in $T[\![\varphi \to \tau]\!]_0 \subseteq T[\![\varphi \to \tau]\!]_n$. Otherwise, there is a (unique) $M \in T[\![\varphi \to \tau]\!]_\infty$ such that $M\ R[\varphi \to \tau]_\infty\ f$, and we need to show that M is in $T[\![\varphi \to \tau]\!]_n$. We only know that M is in $T[\![\varphi \to \tau]\!]_m$ for some $m \in \mathbb{N}$. If $m \leqslant n$, we are done since $T[\![\varphi \to \tau]\!]_m \subseteq T[\![\varphi \to \tau]\!]_n$, so assume $m > n$. Fix a variable X^φ, and notice that X^φ is in $T[\![\varphi]\!]_0 \subseteq T[\![\varphi]\!]_m$. By induction hypothesis (8), $X^\varphi\ R[\varphi]_0\ d$ where $d = s(X^\varphi)$, hence also $X^\varphi\ R[\varphi]_m\ d$. By definition of $R[\varphi \to \tau]$, $MX^\varphi\ R[\tau]_m\ f(d)$. Note that $f(d)$ is in $\mathsf{S}[\![\tau]\!]_n$. By induction hypothesis (9), $MX^\varphi = r_\tau(f(d))$, so MX^φ is in $T[\![\tau]\!]_n$, using the fact that r_τ is a nabla map. Recall how we have defined $T[\![\tau]\!]_n$ (Definition 4.8): the set of λ-terms with names of type τ whose β-normal η-long form contains at most $a_1, \ldots, a_n$ as free names. The β-normal η-long form of M, which is of type $\varphi \to \tau$, must be a lambda-abstraction, and by α-renaming we may assume that it is of the form $\lambda X^\varphi.P$ for some β-normal η-long term P of type τ, with the same variable X^φ we have chosen earlier. The β-normal η-long form of MX^φ is then P. Since MX^φ is in $T[\![\tau]\!]_n$, the only free names in P are among $a_1, \ldots, a_n$. Therefore, the same can be said of M, showing that M is in $T[\![\varphi \to \tau]\!]_n$. Therefore, $r_{\varphi\to\tau}$ is a nabla map. □

The previous proof is a modification of a classical proof (Mitchell 1985, Theorem 8.4.2) of an equational completeness result of Friedman (1975). Showing that $r_{\varphi\to\tau}$ is a nabla map is a new difficulty in our case. We immediately obtain the following, similar completeness result, in the category of nabla sets $\boldsymbol{\nabla}$ rather than in **Set**. (For a closed term M, and a given Henkin universe S, $S[\![M]\!]\rho$ does not depend on the environment ρ at level n, and we write $S[\![M]\!]_n$ for $S[\![M]\!]\rho$ in that case.)

Corollary 8.5. The semantics of λ-terms is *equationally complete*: there is a standard universe S such that the following are equivalent, for any two closed λ-terms M, N of the same type τ:

1. M and N are $\beta\eta$-convertible.
2. $\mathsf{S}[\![M]\!]_0 = \mathsf{S}[\![N]\!]_0$.
3. $\mathsf{S}[\![M]\!] = \mathsf{S}[\![N]\!]$.

Proof. (1) ⇒ (3) ⇒ (2) is obvious, considering Lemma 4.5 (3). Assume (2). Let $d = \mathsf{S}[\![M]\!]_0 = \mathsf{S}[\![N]\!]_0$. By the Basic Lemma 8.3, used with $n = 0$ and the empty substitution θ, $M\ R[\tau]_0\ d$ and $N\ R[\tau]_0\ d$. Apply Proposition 8.4 to obtain that $M = r_\tau(d)$ and $N = r_\tau(d)$ (up to $\beta\eta$-conversion), so $M = N$. □

This ends our parenthesis. We now define a specific generic family on S. The map s_τ was introduced in Proposition 8.4.

Lemma 8.6. Assume there is exactly one base type ι. The family new defined by $\mathsf{new}^\tau_{n+1} = s_\tau(\mathsf{a}^\tau_{n+1})$, $n \in \mathbb{N}$, is a generic family on S.

Proof. We must check that $\mathsf{new}^{\tau}_{n+1}$ is not in $\mathsf{S}[\![\tau]\!]_n$. As a consequence of Proposition 8.4, $r_\tau \circ s_\tau = \mathrm{id}_{T[\![\tau]\!]}$. Therefore, $r_\tau(\mathsf{new}^{\tau}_{n+1}) = \mathsf{a}^{\tau}_{n+1}$. Since that is not in $T[\![\tau]\!]_n$, and r_τ is a nabla map, $\mathsf{new}^{\tau}_{n+1}$ cannot be in $\mathsf{S}[\![\tau]\!]_n$. □

Definition 8.7 (Δ_0 formula, Π_1 formula). A Δ_0 *formula* of $\mathrm{FO}\lambda^{\nabla}$ is a formula whose universal and existential quantifiers are first-order, i.e., of the form $\forall x^\iota$ or $\exists x^\iota$, where ι is a base type. (There is no restriction on the nabla quantifier.)

A Π_1 *formula* is a formula of the form $\forall x_1^{\tau_1}, \dots, x_p^{\tau_p}.G$, where G is a Δ_0 formula.

Proposition 8.8. Assume there is a unique base type ι, and let H be a Herbrand structure. Define a standard structure S^H on the standard universe S by letting

$$\mathsf{S}^H[\![P]\!]_n = \{(d_1, d_2, \dots, d_k) \in \prod_{i=1}^{k} \mathsf{S}[\![\tau_i]\!]_n \mid (r_{\tau_1}(d_1), r_{\tau_2}(d_2), \dots, r_{\tau_k}(d_k)) \in H[\![P]\!]_n\},$$

for every relation symbol P of arity $\tau_1, \tau_2, \dots, \tau_k$ and every $n \in \mathbb{N}$.

For every $n \in \mathbb{N}$, for every substitution θ at level n, for every environment ρ such that $\theta \; R_n \; \rho$:

1. for every Δ_0 formula G whose free variables are included in $\mathrm{dom}\,\theta$, $\mathsf{S}^H, \mathsf{new}; \rho \models_n G$ if and only if $H, \mathsf{a}; \theta \models_n G$;
2. for every Π_1-formula F whose free variables are included in $\mathrm{dom}\,\theta$, if $\mathsf{S}^H, \mathsf{new}; \rho \models_n F$ then $H, \mathsf{a}; \theta \models_n F$;

where new is the generic family of Lemma 8.6.

Proof. (1) By structural induction on G.

If G is an atomic formula $P(M_1, M_2, \dots, M_k)$, where each M_i has type τ_i, then $\mathsf{S}^H, \mathsf{new}; \rho \models_n G$ if and only if $(\mathsf{S}[\![M_1]\!]_n\rho, \mathsf{S}[\![M_2]\!]_n\rho, \dots, \mathsf{S}[\![M_k]\!]_n\rho)$ is in $\mathsf{S}^H[\![P]\!]_n$. By the basic Lemma (Lemma 8.3), $M_i\theta \; R[\tau_i]_n \; \mathsf{S}[\![M_i]\!]\rho$, so, using Proposition 8.4 and specifically (9), $r_{\tau_i}(\mathsf{S}[\![M_i]\!]\rho) = M_i\theta$. Using the definition of $\mathsf{S}^H[\![P]\!]_n$, we obtain that $\mathsf{S}^H, \mathsf{new}; \rho \models_n G$ if and only if $(M_1\theta, M_2\theta, \dots, M_k\theta) \in H[\![P]\!]_n$. The latter is equivalent to $(T[\![M_1]\!]_n\theta, T[\![M_2]\!]_n\theta, \dots, T[\![M_k]\!]_n\theta) \in H[\![P]\!]_n$ (Lemma 4.11), hence to $H, \mathsf{a}; \theta \models_n G$.

If G is a first-order quantified formula $\forall x^\iota.G'$, then $\mathsf{S}^H, \mathsf{new}; \rho \models_n G$ if and only if $\mathsf{S}^H, \mathsf{new}; \rho[x \mapsto d] \models_n G'$ for every $d \in \mathsf{S}[\![\iota]\!]_n$. Since $\mathsf{S}[\![\iota]\!] = H[\![\iota]\!]$, and $R[\iota]_n$ is the identity relation, $\theta[x \mapsto d] \; R_n \; \rho[x \mapsto d]$ for every $d \in \mathsf{S}[\![\iota]\!]_n$. Hence, $\mathsf{S}^H, \mathsf{new}; \rho \models_n G$ if and only $H, \mathsf{a}; \theta[x \mapsto d] \models_n G'$ for every $d \in \mathsf{S}[\![\iota]\!]_n = H[\![\iota]\!]_n$, if and only if $H, \mathsf{a}; \theta \models_n G$.

The other cases follow by an easy induction, except perhaps when G is of the form $\nabla x^\tau.G'$. Then, $\mathsf{S}^H, \mathsf{new}; \rho \models_n G$ if and only if $\mathsf{S}^H, \mathsf{new}; \rho[x \mapsto \mathsf{new}^{\tau}_{n+1}] \models_{n+1} G'$. Let $\rho' = \rho[x \mapsto \mathsf{new}^{\tau}_{n+1}]$, $\theta' = \theta[x \mapsto \mathsf{a}^{\tau}_{n+1}]$. Since $\mathsf{new}^{\tau}_{n+1} = s_\tau(\mathsf{a}^{\tau}_{n+1})$, Proposition 8.4 (and specifically (8)) implies that $\mathsf{a}^{\tau}_{n+1} \; R[\tau]_{n+1} \; \mathsf{new}^{\tau}_{n+1}$. Hence, $\theta' \; R_{n+1} \; \rho'$, and we can apply the induction hypothesis: $\mathsf{S}^H, \mathsf{new}; \rho' \models_{n+1} G'$ if and only if $H, \mathsf{a}; \theta' \models_{n+1} G'$, and therefore $\mathsf{S}^H, \mathsf{new}; \rho \models_n G$ if and only if $H, \mathsf{a}; \theta \models G$.

(2) Let now F be a Π_1 formula $\forall x_1^{\tau_1}, \dots, x_p^{\tau_p}.G$, where G is a Δ_0 formula. If $\mathsf{S}^H, \mathsf{new}; \rho \models_n F$, then $\mathsf{S}^H, \mathsf{new}; \rho[x_1 \mapsto d_1, \dots, x_p \mapsto d_p] \models_n G$ for all values $d_1 \in \mathsf{S}[\![\tau_1]\!]_n, \dots, d_p \in \mathsf{S}[\![\tau_p]\!]_n$. This is in particular true if we pick $d_1 = s_{\tau_1}(N_1), \dots, d_p = s_{\tau_p}(N_p)$ for arbitrary elements $N_1 \in T[\![\tau_1]\!]_n, \dots, N_p \in T[\![\tau_p]\!]_n$. Let $\rho' = \rho[x_1 \mapsto d_1, \dots, x_p \mapsto d_p]$, and $\theta' = \theta[x_1 \mapsto$

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

$N_1, \ldots, x_p \mapsto N_p]$. By Proposition 8.4, and specifically (8), $\theta' \; R_n \; \rho'$. By part 1 of the Proposition, we conclude that $H, \mathsf{a}; \theta[x_1 \mapsto N_1, \ldots, x_p \mapsto N_p] \models_n G$ for all $N_1 \in T[\![\tau_1]\!]_n$, $\ldots$, $N_p \in T[\![\tau_p]\!]_n$, that is, that $H, \mathsf{a}; \theta \models_n F$. □

Write $S, \mathsf{new} \models_0 F$ if $S, \mathsf{new}; \rho \models_0 F$, where F is a closed formula; in that case, the environment ρ is irrelevant.

Proposition 8.9 (Π_1-completeness). Assume there is a unique base type ι. Let F be a closed Π_1 formula. If $S, \mathsf{new} \models_0 F$ for every standard structure S and every generic family new on S, then $\rightarrow \triangleright F$ is derivable in $\mathrm{FO}\lambda^\nabla$, by a cut-free proof.

Proof. Let ρ be any environment at level 0: then $\epsilon \; R_0 \; \rho$. Hence, we can use Proposition 8.8 (2) and conclude that $H, \mathsf{a}; \epsilon \models_n F$. By Proposition 6.5, $\rightarrow \triangleright F$ has a cut-free proof in $\mathrm{FO}\lambda^\nabla$. □

Remark 8.10. Proposition 8.9 in particular implies that $\mathrm{FO}\lambda^\nabla$ is complete for all first-order formulae F in standard structures. This is because every first-order formula is a Δ_0 formula, hence a Π_1-formula.

9. Open questions

(1) Is $\mathrm{FO}\lambda^\nabla$ plus (AC) complete for standard models? (2) What would happen if there were more than one base type ι? (3) Can we extend the present results to the logic of Abella (Gacek 2008), which includes such proof principles as the equivalence of $\nabla x.F$ and F when x is not free in F, and of $\nabla x.\nabla y.F(x, y)$ and $\nabla y.\nabla x.F(x, y)$? (4) Does all this extend to intuitionistic versions of $\mathrm{FO}\lambda^\nabla$? I would say (4) is easy, (3) should be doable, (2) is irritatingly difficult and I have no idea about (1) – although I thought I had one, once.

10. Conclusion

Happy 60th, Dale!

I would like to thank the anonymous referees. In particular, one suggested that I include a list of differences between the various categories used to give semantics to fresh name creation in the literature (end of Section 3).

References

Bucalo, A., Honsell, F., Miculan, M., Scagnetto, I. and Hofmann, M. (2006). Consistency of the theory of contexts. *Journal of Functional Programming* **16** (3) 327–372.

Cervesato, I., Durgin, N. A., Lincoln, P. D., Mitchell, J. C. and Scedrov, A. (1999). A meta-notation for protocol analysis. In: *Proceedings of the 12th IEEE Computer Security Foundations Workshop*, IEEE Conference Publications, 55–69.

Curien, P.-L. (1993). *Categorical Combinators, Sequential Algorithms, and Functional Programming*. Birkhäuser, Boston, MA.

Fitting, M. (1996). *First-Order Logic and Automated Theorem Proving*. Graduate Texts in Computer Science. Spring Verlag, 2nd edition.

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press

Friedman, H. (1975). Equality between functionals. In: Parikh, P. (ed.) *Logic Colloquium 1972–73*, Lecture Notes in Mathematics, vol. 453, Springer-Verlag, 22–37.

Gabbay, M. J. and Pitts, A. M. (1999). A new approach to abstract syntax involving binders. In: *14th Annual Symposium on Logic in Computer Science*, IEEE Computer Society Press, Washington, 214–224.

Gacek, A. (2008). The Abella interactive theorem prover (system description). In: Armando,A., Baumgartner, A. and Dowek, G. (eds.) *Proceedings of IJCAR*, Lecture Notes in Artificial Intelligence, vol. 5195, Springer, 154–161.

Hofmann, M. (1999). Semantical analysis of higher-order abstract syntax. In: *Proceedings of the 14th Annual IEEE Symposium on Logics in Computer Science (LICS'99)*, IEEE, 204–213.

Miculan, M. and Yemane, K. (2005). A unifying model of variables and names. In: *Proceedings of 8th Intl. Conf. Foundations of Software Science and Computational Structures (FOSSACS'05), held as part of the Joint European Conferences on Theory and Practice of Software (ETAPS'05)*, LNCS, Springer Verlag, Edinburgh, UK, 3441.

Miller, D. (1992). The pi-calculus as a theory in linear logic: Preliminary results. Technical Report MS-CIS-92-48, University of Pennsylvania (CIS), October.

Miller, D. and Tiu, A. (2005). A proof theory for generic judgments. *Transactions on Computational Logic* **6** (4) 749–783.

Mitchell, J. C. (1985). *Foundations for Programming Languages*, MIT Press.

Schöpp, U. (2007). Modelling generic judgments. *Electronic Notes in Theoretical Computer Science* **174** (5) 19–35.

https://doi.org/10.1017/S0960129518000063 Published online by Cambridge University Press