A Semantics for Nabla

Jean Goubault-Larrecq[1]

LSV, ENS Cachan, CNRS, Université Paris-Saclay, 94235 Cachan, France
goubault@lsv.fr

Abstract

We give a semantics for a classical variant of Dale Miller and Alwen Tiu's logic $FO\lambda^{\nabla}$. No such semantics seems to have existed for the nabla operator, except for one given by U. Schöpp. 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 category of so-called nabla-sets, which we define as presheaves over the poset of natural numbers, with additional generic elements at each level. The semantics is sound, complete for Henkin structures, and complete for standard structures in the case of Π_1 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 the 14th of February, 2002. I had sent him a rather vague question on his paper [5] 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").*

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.

Dale later found a proper logical way of talking about freshness, or genericity, with Alwen Tiu [6]: the *nabla* quantifier ∇. 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 for so-called generic judgments [6]. The distinctive characteristic 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 [3], 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. I have repeatedly nagged Dale, notably during INRIA project committees, asking him what the relation was between his ∇ operator and Pitts and Gabbay's $\not\vert$ ("new") operator.

Dale and Alwen themselves had answered the question [6, 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 $\not\vert x.P(x)$, which implies $\exists x.P(x)$, while no such implication holds with ∇ instead of $\not\vert$. Also, while $\not\vert x.\not\vert y.P(x,y)$ and $\not\vert y.\not\vert x.P(x,y)$ are equivalent, $\nabla x.\nabla y.P(x,y)$ and $\nabla y.\nabla x.P(x,y)$ are not.

One may hope to understand ∇ better by giving it a semantics, and it is precisely the purpose of this paper. The only other attempt I know of a semantics for ∇ is due to Ulrich Schöpp [8]. This is an elaborate construction, based on categories with binding structure, of which nominal sets provide an example. I feel my proposal is more elementary. More importantly, Schöpp's ∇ operator is restricted to create fresh objects of type ι only, where ι ranges over a dedicated family of base types he calls the lambda-tree types. While this makes a comparison with the $\not\vert$ quantifier easier, this, I feel, ignores one distinctive feature of Dale and Alwen's proposal: the possibility of considering fresh objects of *any* type. Our semantics will address this.

We shall not be completely faithful to [6]. 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 [4]. It will also validate the rule that $\nabla x.F$ and $\nabla y.F$ are equivalent when x, y are not free in F. However, and conforming to [6], 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 [8]: 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 $FO\lambda^{\nabla}$ is sound for all Henkin structures with enough maps in Section 5, and complete in Section 6. Completeness is obtained for term structures, assuming exactly one base type ι, by using a construction of Hintikka sets. This also shows that the cut rule can be eliminated. We examine the question of completeness for standard structures in the rest of the paper. 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

We write $\operatorname{Im} f$ for the image of a map f. Our main object of study is the following.

Definition 3.1. *A* nabla-set D *is the following data:*

- *a family of non-empty sets* $(D_n)_{n\in\mathbb{N}}$ *indexed by natural numbers,*
- *a family of injective maps* $\mathrm{old}_n^D : D_n \to D_{n+1}$, $n \in \mathbb{N}$,

- *a family of elements* $\mathrm{new}^D_{n+1} \in D_n$, $n \in \mathbb{N}$, *such that* $\mathrm{new}^D_{n+1} \notin \mathrm{Im}\,\mathrm{old}^D_n$, $n \in \mathbb{N}$.

We shall write old_n for old^D_n when no confusion may arise. We shall also write $\mathrm{old}^D_{n\to m}$, or $\mathrm{old}_{n\to m}$, for the composition $\mathrm{old}_{m-1} \circ \cdots \circ \mathrm{old}_{n+1} \circ \mathrm{old}_n \colon D_n \to D_m$, $n \leq m$.

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. To convert a value $d \in D_n$ to an element of D_{n+1}, we produce $\mathrm{old}^D_n(d)$. The generic object $\mathrm{new}^D_{n+1} \in D_{n+1}$ is required to be fresh, that is, different from all old objects $\mathrm{old}^D_n(d)$, $d \in D_n$.

Definition 3.2. *A* nabla-map f *from a nabla-set* D *to a nabla-set* E *is a family of maps* $f_n \colon D_n \to E_n$, $n \in \mathbb{N}$, *such that* $\mathrm{old}_n \circ f_n = f_{n+1} \circ \mathrm{old}_n$.

We do *not* require that $f_{n+1}(\mathrm{new}^D_{n+1}) = \mathrm{new}^E_{n+1}$. The elements new^D_{n+1} are generic, and should have no specific property, including the latter.

Nabla-sets and nabla-maps form a category $\boldsymbol{\nabla}$. An isomorphism between D and E in $\boldsymbol{\nabla}$ is a collection of bijections $f_n \colon D_n \to E_n$ such that $\mathrm{old}_n \circ f_n = f_{n+1} \circ \mathrm{old}_n$. The bijections will usually fail to map new^D_{n+1} to new^E_{n+1}. Hence the following defines isomorphic nabla-sets, obtained by moving the generic elements new^D_{n+1} around:

Definition 3.3 (Variant)**.** *A* variant *of a nabla-set* D *is a nabla-set* D' *such that* $D'_n = D_n$ *and* $\mathrm{old}^{D'}_n = \mathrm{old}^D_n$ *for every* $n \in \mathbb{N}$.

$\boldsymbol{\nabla}$ is very close to a familiar presheaf category. Let $\mathbb{N}$ be the set of natural numbers, with the usual ordering. Any poset can be considered as a category, whose objects are the elements of the poset, and where there is one morphism from m to n if $m \leq n$, no morphism otherwise. One can then form the category $\mathbf{Set}^{\mathbb{N}}$ of all functors from $\mathbb{N}$ to the category $\mathbf{Set}$ of sets, with natural transformations as morphisms. One can check that $\mathbf{Set}^{\mathbb{N}}$ is equivalently defined much like $\boldsymbol{\nabla}$: objects of $\mathbf{Set}^{\mathbb{N}}$ are given by a family of sets $(D_n)_{n\in\mathbb{N}}$ (possibly empty), together with maps $\mathrm{old}^D_n \colon D_n \to D_{n+1}$ (not necessarily injective), $n \in \mathbb{N}$, and morphisms f from D to E are families of maps $f_n \colon D_n \to E_n$, $n \in \mathbb{N}$, such that $\mathrm{old}_n \circ f_n = f_{n+1} \circ \mathrm{old}_n$.

$\mathbf{Set}^{\mathbb{N}}$ is Cartesian-closed and complete. Almost the same thing can be said for $\boldsymbol{\nabla}$, except for one important point: $\boldsymbol{\nabla}$ does not have a terminal object (i.e., a 0-ary product; I will leave that as an exercise to the reader).

Nonetheless, most of the structure of $\mathbf{Set}^{\mathbb{N}}$ is preserved. Notably, $\boldsymbol{\nabla}$ 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: $D_n = \prod_{i\in I} D[i]_n$, old^D_n maps $(d_i)_{i\in I}$ to $(\mathrm{old}^{D[i]}_n(d_i))_{i\in I}$, and we may define new^D_{n+1} as $(\mathrm{new}^{D[i]}_{n+1})_{i\in I}$. All other products are, as usual, obtained as isomorphic copies.

We shall use the following notations:

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

A Cartesian-closed category is one with all finite products (i.e., terminal objects and binary products), and where every object D is *exponentiable*. The latter means, abstractly, that the functor $\_ \times D$ 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 $\mathsf{App}\colon [D \to E] \times D \to E$, and for every morphism $f\colon C \times D \to E$, a *currified* map $\Lambda(f)\colon C \to [D \to E]$ satisfying the following equations [1]:

- (β-rule) $\mathsf{App} \circ (\Lambda f \times \mathrm{id}_D) = f$ for every $f\colon C \times D \to E$;
- (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 $f\colon C \times D \to E$ and $g\colon B \to C$.

Proposition 3.4. *In $\boldsymbol{\nabla}$, every object is exponentiable.*

Proof. Given two nabla-sets D and E, we define $[D \to E]$ as in $\mathbf{Set}^{\mathbb{N}}$. Let $[D \to E]_n$ be the set of all families of maps $(f_m)_{m\geq n}$ such that $\mathrm{old}^E_m \circ f_m = f_{m+1} \circ \mathrm{old}^D_m$ for every $m \geq n$. We let $\mathrm{old}^{[D\to E]}_n$ map $(f_m)_{m\geq n} \in [D \to E]_n$ to $(f_m)_{m\geq n+1} \in [D \to E]_{n+1}$.

We claim that $[D \to E]_n$ is non-empty for every $n \in \mathbb{N}$. Pick $e \in E_0$, define f_m as the constant map with value $\mathrm{old}^E_{0\to m}(e)$, and check that $(f_m)_{m\geq n}$ is an element of $[D \to E]_n$.

Next, the maps $\mathrm{old}^{[D\to E]}_n$ are injective. Indeed, given $(f_m)_{m\geq n+1} \in [D \to E]_{n+1}$, there is at most one map f_n such that $\mathrm{old}^E_n \circ f_n = f_{n+1} \circ \mathrm{old}^D_n$, because old^E_n is injective.

The generic element $\mathrm{new}^{[D\to E]}_{n+1}$ must be a collection of maps $\left(\mathrm{new}^{[D\to E]}_{(n+1)m}\right)_{m\geq n+1}$, which we define by induction on m. For $m = n+1$, we let $\mathrm{new}^{[D\to E]}_{(n+1)(n+1)}$ be the constant map with value new^E_{n+1}. For $m > n$, we define $\mathrm{new}^{[D\to E]}_{(n+1)m}(d)$, for every $d \in D_m$, as follows. If $d \in \operatorname{Im} \mathrm{old}^D_{m-1}$, then there is a unique $d' \in D_{m-1}$ such that $\mathrm{old}^D_{m-1}(d') = d$, since old^D_{m-1} is injective, and we define $\mathrm{new}^{[D\to E]}_{(n+1)m}(d)$ as $\mathrm{old}^E_{m-1}(\mathrm{new}^{[D\to E]}_{(n+1)(m-1)}(d'))$, using the induction hypothesis. This part of the construction ensures that $\mathrm{new}^{[D\to E]}_{n+1}$ is an element of $[D \to E]_{n+1}$.

It remains to define $\mathrm{new}^{[D\to E]}_{(n+1)m}(d)$ when $d \notin \operatorname{Im} \mathrm{old}^D_{m-1}$, and we set it to new^E_m. That ensures that $\mathrm{new}^{[D\to E]}_{n+1}$ is not in the image of $\mathrm{old}^{[D\to E]}_n$, for $n \in \mathbb{N}$, as we now show. Otherwise, there would be an element $f = (f_m)_{m\geq n}$ of $[D \to E]_n$ such that $f_m = \mathrm{new}^{[D\to E]}_{(n+1)m}$ for every $m \geq n+1$. Recall that $f_{n+1} \circ \mathrm{old}^D_n = \mathrm{old}^E_n \circ f_n$. Since $f_{n+1} = \mathrm{new}^{[D\to E]}_{(n+1)(n+1)}$, this would imply that $\mathrm{new}^E_{n+1} = (\mathrm{old}^E_n \circ f_n)(d)$ for every $d \in D_n$; since D_n is non-empty, we reach a contradiction.

That $[D \to E]$ is an exponential object either follows from a tedious verification, or as a consequence of the corresponding statement that $\mathbf{Set}^{\mathbb{N}}$ is Cartesian-closed.

Explicitly, $\mathsf{App}\colon [D \to E] \times D \to E$ is defined by $\mathsf{App}_n((f_m)_{m\geq n}, d) = f_n(d)$. For every morphism $f = (f_n)_{n\in\mathbb{N}}$ from $C \times D$ to E, the currified morphism Λf from C to $[D \to E]$ is defined by $(\Lambda f)_n(c) = (f_m(\mathrm{old}^D_{n\to m}(c), \_))_{m\geq n}$ for every $c \in C_n$. (We write $f_m(a, \_)$ for the map that sends d to $f_m(a, d)$.) We leave it as an exercise to the reader to check that the β, η and substitution rules are satisfied. □

Unlike most other presheaf categories, $\mathbf{Set}^{\mathbb{N}}$ satisfies the so-called external axiom of choice, and similarly for $\boldsymbol{\nabla}$: every epi splits. We prefer the following formulation, which is closer to what we think the axiom of choice should state, and also slightly more general.

Proposition 3.5 (Choice)**.** *A* nabla-subset $(A_n)_{n\in\mathbb{N}}$ *of a nabla-set* $(D_n)_{n\in\mathbb{N}}$ *is a collection of subsets A_n of D_n, $n \in \mathbb{N}$, such that for every $n \in \mathbb{N}$, for every $a \in A_n$, $\mathrm{old}^D_n(a)$ is in A_{n+1}.*

A binary nabla-relation *R between D and E is a nabla-subset of $D \times E$.*

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 $f = (f_n)_{n\in\mathbb{N}}$ from D to E such that, for all $n \in \mathbb{N}$ and $d \in D_n$, $(d, f_n(d))$ is in R_n.

Proof. We build f_n by induction on n. The function f_0 is simply obtained by applying the set-theoretic axiom of choice to select an $e \in E_0$ such that $(d, e) \in R_0$, for each $d \in D_0$, and defining $f_0(d)$ as e.

At level $n+1$, we make cases depending on whether d is in the image of old_n^D or not.

If $d = \mathrm{old}_n^D(d')$ for some (unique) $d' \in D_n$, then we define $f_{n+1}(d)$ as $\mathrm{old}_n^E(f_n(d'))$. Note that, since $(d', f_n(d')) \in R_n$, $\mathrm{old}_n^{D\times E}(d', f_n(d')) = (d, f_{n+1}(d))$ is in R_{n+1}.

If $d \notin \mathrm{Im}\,\mathrm{old}_n^D$, then we let $f_{n+1}(d)$ be some $e \in E_{n+1}$ such that $(d, e) \in R_{n+1}$. It is clear that $f = (f_n)_{n\in\mathbb{N}}$ is a nabla-map, and that $(d, f_n(d)) \in R_n$ for all $n \in \mathbb{N}$, $d \in D_n$. □

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

Corollary 3.6 (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 $f = (f_m)_{m\geq n}$ of $[D \to E]_n$ such that for every $d \in D_n$, $(d, f_n(d))$ is in R.*

Proof. Let $D' = (D_m)_{m\geq n}$, $E' = (E_m)_{m\geq n}$, with the obvious nabla-set structure. For every $k \in \mathbb{N}$, let R'_k be the set of pairs $(d', e') \in D'_k \times E'_k = D_{n+k} \times E_{n+k}$ such that $d' = \mathrm{old}_{n\to k}^D(d)$ and $e' = \mathrm{old}_{n\to k}^E(e)$ for some $d \in D_n$ and $e \in E_n$ such that $(d, e) \in R$, or such that $d' \notin \mathrm{Im}\,\mathrm{old}_{n\to k}^D$ and e' is arbitrary. Then $(R'_k)_{k\in\mathbb{N}}$ is a nabla-subset of $D' \times E'$, and one satisfying the assumptions of Proposition 3.5. Hence there is a nabla-map $f' = (f'_k)_{k\in\mathbb{N}}$ from D' to E' such that, for every $k \in \mathbb{N}$, for every $d' \in D'_k$, $(d', f'_k(d'_k)) \in R'_k$. The claim follows by looking at the case $k = 0$, and by letting $f_m = f'_{m-n}$ for every $m \geq n$. □

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 need new_{n+1}^D to be fresh, that is, outside $\mathrm{Im}\,\mathrm{old}_n^D$ (Definition 3.1, third item).

Lemma 3.7. *Let D, E be two nabla-sets, $n \in \mathbb{N}$, and $e \in E_{n+1}$. There is a nabla-map $f = (f_m)_{m\in\mathbb{N}} \colon D \to E$ such that $f_{n+1}(\mathrm{new}_{n+1}^D) = e$.*

Proof. Let $g = \mathrm{new}_0^{[D\to E]}$. This is a collection of maps $(g_m)_{m\geq 0}$ such that $\mathrm{old}_m^E \circ g_m = g_{m+1} \circ \mathrm{old}_m^D$ for every $m \geq 0$. In other words, it is a nabla-map from D to E. We build a nabla-map $f = (f_m)_{m\in\mathbb{N}}$ from D to E by patching g.

For every $m \leq n$, we let $f_m = g_m$. For $m = n+1$, we let f_{n+1} map new_{n+1}^D to e, and every element $d \neq \mathrm{new}_{n+1}^D$ to $g_{n+1}(d)$. The relation $\mathrm{old}_n^E \circ f_n = f_{n+1} \circ \mathrm{old}_n^D$ is satisfied vacuously, because $\mathrm{old}_n^D(d) \neq \mathrm{new}_{n+1}^D$ for every $d \in D_n$. We then proceed to define f_m for every $m > n+1$ by induction on m: it maps $\mathrm{old}_{n+1\to m}^D(\mathrm{new}_{n+1}^D)$ to $\mathrm{old}_{n+1\to m}^E(e)$, and every element $d \neq \mathrm{old}_{n+1\to m}^D(\mathrm{new}_{n+1}^D)$ to $g_m(d)$. □

4 Standard and Henkin Semantics for λ-Terms

Let use consider simply-typed λ-terms M in Church style, that is, all variables have a pre-assigned type. We agree that given a variable x_τ, its type is 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 including the type ι of individuals, and other types are formed using the arrow type former $\to$. Later, and for the purposes of completeness, we shall require that there be *exactly* one type ι of individuals.

Proposition 3.4 allows us to define a *standard semantics* for λ-terms: we fix nabla-sets $S[\![\tau]\!]$ for every base type, define $S[\![\varphi \to \tau]\!]$ as the exponential object $[S[\![\varphi]\!] \to S[\![\tau]\!]]$ or one of its

variants, 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 call a *Henkin semantics* for nabla, and which we shall need to establish completeness. This is simply a listing of our basic requirements.

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

- *for each type τ, a nabla-set $S[\![\tau]\!]$;*
  We 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 type τ, a set $S(\tau)$ of nabla-maps from Env to τ, containing all the projections π_{x_τ}—where $(\pi_{x_\tau})_n(\rho) = \rho(x_\tau)$ for every environment ρ at level n;*
- *for each pair of types φ, τ, a nabla-map $\mathsf{App}\colon S[\![\varphi \to \tau]\!] \times S[\![\varphi]\!] \to S[\![\tau]\!]$, with the property that for every $f \in S(\varphi \to \tau)$ and for every $g \in S(\varphi)$, $\mathsf{App} \circ \langle f, g\rangle$ is in $S(\tau)$;*
- *for every variable x_φ and each type τ, a function $\Lambda_{x_\varphi}\colon S(\tau) \to S(\varphi \to \tau)$;*

such that, defining:

$$\begin{aligned} S[\![x_\tau]\!] &= \pi_{x_\tau} \\ S[\![MN]\!] &= \mathsf{App} \circ \langle S[\![M]\!], S[\![N]\!]\rangle \\ S[\![\lambda x_\varphi.M]\!] &= \Lambda_{x_\varphi}(S[\![M]\!]) \end{aligned}$$

then:

1. *for all $\beta\eta$-convertible λ-terms $M, N : \tau$, $S[\![M]\!] = S[\![N]\!]$;*
2. *for every λ-term $M : \tau$, for every $n \in \mathbb{N}$, $S[\![M]\!]_n\rho$ 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 $S[\![M]\!]_n\rho = S[\![M]\!]_n\rho'$;*
3. *for all λ-terms $N : \tau$ and $M : \varphi$, for every $n \in \mathbb{N}$, for every environment ρ at level n, $S[\![N[M/x_\varphi]]\!]_n\rho = S[\![N]\!]_n(\rho[x_\varphi \mapsto S[\![M]\!]_n\rho])$;*

Adapting Lemma 3.7 in view of our upcoming proof of soundness, we also define:

Definition 4.2 (Enough Maps)**.** *A Henkin universe S for nabla* has enough maps *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 $\mathsf{App}_{n+1}(\mathrm{old}_n(f), \mathrm{new}_{n+1}^{S[\![\varphi]\!]}) = d$.*

4.1 The Standard Universe

Lemma 4.3 (The Standard Universe)**.** *Given nabla-sets D_τ, one for each base type τ, there is a Henkin universe S such that:*

- *$\mathsf{S}[\![\varphi \to \tau]\!]$ is a variant of $[\mathsf{S}[\![\varphi]\!] \to \mathsf{S}[\![\tau]\!]]$ for all types φ, τ (see Definition 3.3);*
- *$\mathsf{S}(\tau)$ is the set of all nabla-maps from Env to τ, for each type τ;*
- *App is the application morphism in $\boldsymbol{\nabla}$;*
- *for every $f\colon \mathsf{Env} \to \mathsf{S}[\![\tau]\!]$, $\Lambda_{x_\varphi}(f) = \Lambda(f \circ \mathsf{bind}_{x_\varphi})$, where Λ is currification in $\boldsymbol{\nabla}$ and $\mathsf{bind}_{x_\varphi} : \mathsf{Env} \times \mathsf{S}[\![\varphi]\!] \to \mathsf{Env}$ is defined by $(\mathsf{bind}_{x_\varphi})_n(\rho, d) = \rho[x_\varphi \mapsto d]$, the environment that maps x_φ to d and every variable $y \neq x$ to $\rho(y)$.*

This Henkin universe S has enough maps.

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

Proof. 1. The fact that for all $\beta\eta$-convertible λ-terms $M, N : \tau$, $S[\![M]\!] = S[\![N]\!]$, owes to the properties of exponentiable objects (see Proposition 3.4), and is immediate.

2. If $\rho(z) = \rho'(z)$ for every $z \neq y$, then $S[\![M]\!]_n\rho = S[\![M]\!]_n\rho'$: this is an easy structural induction on M.

3. We show that for all λ-terms $N : \tau$ and $M : \varphi$, for every $n \in \mathbb{N}$, for every environment ρ at level n, $S[\![N[M/x_\varphi]]\!]_n\rho = S[\![N]\!]_n(\rho[x_\varphi \mapsto S[\![M]\!]_n\rho])$.

We first notice that: (a) for every $m \geq n$, $\mathrm{old}^{S[\![\tau]\!]}_{n\to m}(S[\![M]\!]_n\rho) = S[\![M]\!]_m(\mathrm{old}^{\mathsf{Env}}_{n\to m}(\rho))$. This is merely the expression that $S[\![M]\!]$ is a nabla-map, by definition.

We now show that $S[\![N[M/x_\varphi]]\!]_n\rho = S[\![N]\!]_n(\rho[x_\varphi \mapsto S[\![M]\!]_n\rho])$, by structural induction on N. The only interesting case is when N is a λ-abstraction $\lambda y_\varphi.P$. Then, assuming that $y_\varphi \neq x_\tau$ and that y_φ is not free in M, by α-renaming:

$$\begin{aligned}
S[\![N[M/x_\tau]]\!]_n\rho &= (\lambda d \in S[\![\varphi]\!]_m.S[\![P[M/x_\tau]]\!]_m(\mathrm{old}^{\mathsf{Env}}_{n\to m}(\rho)[y_\varphi \mapsto d]))_{m\geq n} \\
&= (\lambda d \in S[\![\varphi]\!]_m. && \text{(by induction hypothesis)} \\
&\quad S[\![P]\!]_m(\mathrm{old}^{\mathsf{Env}}_{n\to m}(\rho)[y_\varphi \mapsto d, x_\tau \mapsto S[\![M]\!]_m(\mathrm{old}^{\mathsf{Env}}_{n\to m}(\rho)[y_\varphi \mapsto d])]))_{m\geq n} \\
&= (\lambda d \in S[\![\varphi]\!]_m. && \text{(by 2.)} \\
&\quad S[\![P]\!]_m(\mathrm{old}^{\mathsf{Env}}_{n\to m}(\rho)[y_\varphi \mapsto d, x_\tau \mapsto S[\![M]\!]_m(\mathrm{old}^{\mathsf{Env}}_{n\to m}(\rho))]))_{m\geq n} \\
&= (\lambda d \in S[\![\varphi]\!]_m. && \text{(by (a))} \\
&\quad S[\![P]\!]_m(\mathrm{old}^{\mathsf{Env}}_{n\to m}(\rho)[y_\varphi \mapsto d, x_\tau \mapsto \mathrm{old}^{S[\![\varphi]\!]}_{n\to m}(S[\![M]\!]_n\rho)])) \\
&= (\lambda d \in S[\![\varphi]\!]_m. && \text{(by definition of old}^{\mathsf{Env}}\text{)} \\
&\quad S[\![P]\!]_m(\mathrm{old}^{\mathsf{Env}}_{n\to m}(\rho[x_\tau \mapsto S[\![M]\!]_n\rho])[y_\varphi \mapsto d])) \\
&= S[\![N]\!]_n(\rho[x_\tau \mapsto S[\![M]\!]_n\rho]).
\end{aligned}$$

We now claim that S has enough maps. Fix $d \in \mathsf{S}[\![\tau]\!]_{n+1}$. By Lemma 3.7, there is a nabla-map $(f_m)_{m\in\mathbb{N}} \colon \mathsf{S}[\![\varphi]\!] \to \mathsf{S}[\![\tau]\!]$ such that $f_{n+1}(\mathrm{new}^{\mathsf{S}[\![\varphi]\!]}_{n+1}) = d$. Let $f = (f_m)_{m\geq n}$. Then $\mathrm{old}_n(f) = (f_m)_{m\geq n+1}$, and $\mathsf{App}_{n+1}(\mathrm{old}_n(f), \mathrm{new}^{S[\![\varphi]\!]}_{n+1}) = f_{n+1}(\mathrm{new}^{\mathsf{S}[\![\varphi]\!]}_{n+1}) = d$. □

Remark 4.4. *A standard universe* S *is uniquely determined by choosing nabla-sets* $\mathsf{S}[\![\tau]\!]$ *for each base type* τ*, and by choosing generic elements* $\mathrm{new}^{\mathsf{S}[\![\varphi\to\tau]\!]}_n$*,* $n \in \mathbb{N}$*, for each arrow type* $\varphi \to \tau$*.*

4.2 The Term Universe

We now exhibit another Henkin universe T, built from syntax. This will be useful to show completeness. Here we require that there be *exactly* one base type ι.

The universe T is built from an extension of the λ-calculus we have considered until now, obtained by adding a countably infinite supply of new constants a_i, $i \geq 1$, all of type ι, and called *names*. We assume that those names are pairwise distinct; a_i is the name *at level* i.

We build *nominal (simply-typed)* λ*-terms* inductively by: every variable x_τ is a nominal λ-term, of type τ; every name a_i is a nominal λ-term, of type ι; if M is a nominal λ-term of type $\varphi \to \tau$ and N is a nominal λ-term of type φ, then MN is a nominal λ-term of type τ; if M is a nominal λ-term of type τ, and x_φ is a variable, then $\lambda x_\varphi.M$ is a nominal λ-term of type $\varphi \to \tau$.

In other words, nominal λ-terms are ordinary λ-terms on an enlarged set of variables, consisting of variables and names, and restricted so that names cannot occur bound. We will not take this view, and we will enforce a strict separation between variables and names.

We consider nominal λ-terms modulo $\beta\eta$-conversion, and by this we mean a nominal λ-term is shorthand for its $\beta\eta$-normal form. This convention allows us to make sense of the notions of free variables, and of free names, of a nominal λ-term.

Definition 4.5. *For each type τ, for every $n \in \mathbb{N}$, $T[\![\tau]\!]_n$ is the set of all nominal λ-terms of type τ (up to $\beta\eta$-conversion) in which the only free names are of the form a_i with $1 \leq i \leq n$.*

The maps $\mathrm{old}_n^{T[\![\tau]\!]}$ *map M to M, and, writing τ in a unique way as $\tau_1 \to \tau_2 \to \cdots \to \tau_m \to \iota$,* $\mathrm{new}_n^{T[\![\tau]\!]} = \lambda x_{1\,\tau_1}.\lambda x_{2\,\tau_2}.\cdots.\lambda x_{m\,\tau_m}.a_n$, *where $x_{1\,\tau_1}$, $x_{2\,\tau_2}$, $\ldots x_{m\,\tau_m}$ are distinct fresh variables.*

Remark 4.6. *For every type τ, $T[\![\tau]\!]_0$ is just the set of ordinary, not nominal, λ-terms of type τ, modulo $\beta\eta$-conversion.*

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 domains, by defining $M\theta$ as $M\theta_{|A}$, where A is any finite subset containing the free variables of M.

Define again Env as $\prod_{x_\tau} T[\![\tau]\!]$. An element θ of Env_n is a substitution at level n, so that every λ-term M defines a map $\widehat{M}_n \colon \mathsf{Env}_n \to T[\![\tau]\!]_n$, which sends θ to $M\theta$. Then $\widehat{M} = (\widehat{M}_n)_{n\in\mathbb{N}}$ is a nabla-map from Env to $T[\![\tau]\!]$.

Definition 4.7. *For every type τ, let $T(\tau)$ be the set of all nabla-maps of the form $\widehat{M}$, where M ranges over the λ-terms of type τ.*

$T(\tau)$ contains all the projections π_{x_τ}, since $\pi_{x_\tau} = \widehat{x_\tau}$.

It is also clear that every element of $T(\tau)$ is of the form $\widehat{M}$ for a *unique* λ-term M (up to $\beta\eta$-conversion): $M = \widehat{M}(\theta)$, where θ is the identity substitution. So the following makes sense.

Definition 4.8. *Let* $\mathsf{App} \colon T(\varphi \to \tau) \times T(\varphi) \to T(\tau)$ *be defined by* $\mathsf{App}(\widehat{M}, \widehat{N}) = \widehat{MN}$, *and* $\Lambda_{x_\varphi} \colon T(\tau) \to T(\varphi \to \tau)$ *map* $\widehat{M}$ *to* $\widehat{\lambda x_\varphi.M}$.

The following fact is immediate.

Fact 4.9. *For every λ-term $M : \tau$, for every substitution θ at level n, $T[\![M]\!]_n\theta = \widehat{M}(\theta) = M\theta$.*

Lemma 4.10. *Assume there is exactly one base type ι. T, defined in Definitions 4.5–4.8, is a Henkin universe with enough maps.*

Proof. Clearly, $\mathrm{old}_n^{T[\![\tau]\!]}$ is injective. Properties 1–3 of Definition 4.1 are clear, given Fact 4.9. Let $N \in T[\![\tau]\!]_{n+1}$. We wish to find an $M \in T[\![\varphi \to \tau]\!]_n$ such that $M\mathrm{new}_{n+1}^{T[\![\varphi]\!]} = N$ (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 $M_1 : \varphi_1$, $M_2 : \varphi_2$, …, $M_n : \varphi_n$—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 $1 \leq i \leq n$, by construction, so M is in $T[\![\varphi \to \tau]\!]_n$, and $M\mathrm{new}_{n+1}^{T[\![\varphi]\!]} = M(\lambda x_{1\,\varphi_1}.\lambda x_{2\,\varphi_2}.\cdots.\lambda x_{m\,\varphi_m}.a_{n+1}) = \widetilde{N}[\lambda x_{1\,\varphi_1}.\lambda x_{2\,\varphi_2}.\cdots.\lambda x_{m\,\varphi_m}.a_{n+1}/x_\varphi]$ is equal to N. □

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

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

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

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

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

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

Figure 1: A Sequent Calculus Formulation of $FO\lambda^\nabla$

5 A Semantics for $FO\lambda^\nabla$, and Soundness

The logic $FO\lambda^\nabla$ was introduced by Miller and Tiu [6], as an intuitionistic first-order logic with predicates on higher-order terms, together with the ∇ operator. Schöpp [8] 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.

Instead of considering all the connectives, we shall restrict ourselves to $\bot$ (false), $\supset$ (implication) and $\forall$ (universal quantification). The other connectives could be dealt with similarly. We profit from the fact that our logic is classical, so that 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 judgment* (or, more simply, a *judgment*) J is an expression of the form $\sigma \triangleright F$ where σ is a signature (the *local signature* of the judgment) and F is a formula. The meaning of $x_1 : \tau_1, x_2 : \tau_2, \cdots, x_m : \tau_m \triangleright F$ 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 $\lambda x_1 : \tau_1, x_2 : \tau_2, \cdots, x_m : \tau_m \triangleright F$. We also write σ, σ' for the concatenation of signatures when this makes sense.

Definition 5.1. *Let $\approx$ be the smallest equivalence relation on judgments such that:*

- *if* $\lambda\sigma.F$ *and* $\lambda\sigma'.F'$ *are* $\beta\eta$*-convertible, then* $(\sigma \triangleright F) \approx (\sigma' \triangleright F')$*;*
- *if* x_τ *and* y_φ *are not free in* F*, then* $(\sigma, x : \tau, \sigma' \triangleright F) \approx (\sigma, y : \varphi, \sigma' \triangleright F)$*.*

A *sequent* of $FO\lambda^\nabla$ is an expression $\Gamma \longrightarrow \Delta$, where Γ, Δ are finite multisets of judgments.

Remark 5.2. *Those are slightly different from the sequents of [6], 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 judgment 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 $FO\lambda^\nabla$ are shown in Figure 5. In the rightmost premise of $(\forall L)$, one can find a judgment $\sigma \triangleright F[M/x_\tau]$. $F[M/x_\tau]$ 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 nabla-relations, as defined in Proposition 3.5: we do not require that for every $n \in \mathbb{N}$ and for every $d \in P_n$, $\mathrm{old}_n^D(d)$ is 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 4.3).*

We now define satisfaction of a formula F at level n as follows, in a Henkin structure S, where ρ is a Σ-environment at level n.

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

This extends to judgments by letting $S;\rho \models_n x_1 : \tau_1, x_2 : \tau_2, \cdots, x_m : \tau_m \triangleright F$ if and only if $S;\rho \models_n \nabla x_{1\,\tau_1}.\nabla x_{2\,\tau_2}.\cdots.\nabla x_{m\,\tau_m}.F$; then, to sequents by letting $S;\rho \models_n \Gamma \longrightarrow \Delta$ if and only if $S;\rho \not\models_n J$ for some J in Γ or $S;\rho \models_n J$ for some J in Δ.

Lemma 5.5. *For every* λ*-term* M *of type* τ*, for every* $n \in \mathbb{N}$*,*

1. $S;\rho \models_n J[M/x_\tau]$ *iff* $S;\rho[x_\tau \mapsto S[\![M]\!]_n\rho] \models_n J$ *for every judgment* J*;*
2. $S;\rho \models_n \Gamma[M/x_\tau] \longrightarrow \Delta[M/x_\tau]$ *iff* $S;\rho[x_\tau \mapsto S[\![M]\!]_n\rho] \models_n \Gamma \longrightarrow \Delta$*.*

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 $\mathrm{old}_n^{\mathsf{Env}}(\rho)[y_\varphi \mapsto \mathrm{new}_{n+1}^{S[\![\varphi]\!]}]$. Then $S; \rho \models_n J$ if and only if $S; \widetilde{\rho} \models_{n+1} F[M/x_\tau]$,

$$
\begin{array}{ll}
\text{iff} & S; \mathrm{old}_n^{\mathsf{Env}}(\rho)[y_\varphi \mapsto \mathrm{new}_{n+1}^{S[\![\varphi]\!]}, x_\tau \mapsto S[\![M]\!]_{n+1}\widetilde{\rho}] \models_{n+1} F \qquad \text{(by induction hypothesis)} \\
\text{iff} & S; \mathrm{old}_n^{\mathsf{Env}}(\rho)[y_\varphi \mapsto \mathrm{new}_{n+1}^{S[\![\varphi]\!]}, x_\tau \mapsto S[\![M]\!]_{n+1}(\mathrm{old}_n^{\mathsf{Env}}(\rho))] \models_{n+1} F \text{ (prop. 2 of Henkin universes)} \\
\text{iff} & S; \mathrm{old}_n^{\mathsf{Env}}(\rho)[y_\varphi \mapsto \mathrm{new}_{n+1}^{S[\![\varphi]\!]}, x_\tau \mapsto \mathrm{old}_n^{S[\![\tau]\!]}(S[\![M]\!]_n\rho)] \models_{n+1} F \quad (S[\![M]\!] \text{ is a nabla-map}) \\
\text{iff} & S; \mathrm{old}_n^{\mathsf{Env}}(\rho[x_\tau \mapsto S[\![M]\!]_n\rho])[y_\varphi \mapsto \mathrm{new}_{n+1}^{S[\![\varphi]\!]}] \models_{n+1} F \\
\text{iff} & S; \rho[x_\tau \mapsto S[\![M]\!]_n\rho] \models_n \nabla y_\varphi.F.
\end{array}
$$

2. Immediate consequence of 1. □

We say that two formulae F and G are *equivalent* if and only if, for every nabla-structure S, for every $n \in \mathbb{N}$, for every environment ρ at level n, $S; \rho \models_n F$ if and only if $S; \rho \models_n G$.

Lemma 5.6. *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 property 2 of Henkin structures in the base case.

Finally, for the third equivalence, we have:

$$
\begin{array}{rl}
S; \rho \models_n \nabla x_\tau.\forall y_\varphi.F \text{ iff } & S; \mathrm{old}_n^{\mathsf{Env}}(\rho)[x \mapsto \mathrm{new}_{n+1}] \models_{n+1} \forall y_\varphi.F \\
& \text{iff (for every } d \in S[\![\varphi]\!]_{n+1},\ S; \mathrm{old}_n^{\mathsf{Env}}(\rho)[x \mapsto \mathrm{new}_{n+1}, y \mapsto d] \models_{n+1} F) \quad (1)
\end{array}
$$

while $S; \rho \models_n \forall h_{\tau\to\varphi}.\nabla x_\tau.F[hx/y]$ if and only if:

$$
\begin{array}{ll}
 & (\text{for every } f \in S[\![\tau \to \varphi]\!]_n,\ S; \rho[h \mapsto f] \models_n \nabla x_\tau.F[hx/y]) \\
\text{iff} & (\text{for every } f \in S[\![\tau \to \varphi]\!]_n,\ S; \mathrm{old}_n^{\mathsf{Env}}(\rho[h \mapsto f])[x \mapsto \mathrm{new}_{n+1}] \models_{n+1} F[hx/y]) \\
\text{iff} & (\text{for every } f \in S[\![\tau \to \varphi]\!]_n, \qquad (2) \\
 & \quad S; \mathrm{old}_n^{\mathsf{Env}}(\rho)[x \mapsto \mathrm{new}_{n+1}, y \mapsto \mathsf{App}_{n+1}(\mathrm{old}_n(f), \mathrm{new}_{n+1})] \models_{n+1} F)
\end{array}
$$

where we have used Lemma 5.5, item 2, and the fact that h is not free in F in the last line. The two are equivalent: in one direction, for every $f \in S[\![\tau \to \varphi]\!]_n$, $\mathsf{App}_{n+1}(\mathrm{old}_n(f), \mathrm{new}_{n+1})$ is a value d in $S[\![\varphi]\!]_{n+1}$, so (1) implies (2). 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 $\mathsf{App}_{n+1}(\mathrm{old}_n(f), \mathrm{new}_{n+1}^{S[\![\varphi]\!]}) = d$, because S has enough maps. Hence (2) implies (1). □

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

Proposition 5.7 (Soundness)**.** *Every derivable sequent* $\Gamma \longrightarrow \Delta$ *is valid.*

Proof. It suffices to show that $S; \rho \models_n \Gamma \longrightarrow \Delta$ by induction on the given derivation.

In the case of the $(\supset L)/(\supset R)$ rules, we must show that $S; \rho \models_n \sigma \triangleright (F \supset G)$ if and only if $S; \rho \not\models_n \sigma \triangleright F$ or $S; \rho \models_n \sigma \triangleright G$: this is an easy induction on the number of variables in σ, using Lemma 5.6, item 1.

In the case of $(\approx L)/(\approx R)$, we must show that $S; \rho \models_n J$ if and only if $S; \rho \models_n J'$, assuming $J \approx J'$. It suffices to show that this is the case when J and J' are $\beta\eta$-convertible (which follows from property 1 of Henkin universes), and when $J = \sigma, x : \tau, \sigma' \triangleright F$, $J' = \sigma, y : \varphi, \sigma' \triangleright F$, with x_τ, y_φ not free in F (that follows from Lemma 5.6, item 2).

In the case of $(\forall R)$, assume that $S; \rho \models_n \Gamma \longrightarrow \Delta, (\sigma \triangleright F[h\sigma/x_\tau])$, with h fresh of type $\sigma \to \tau$. Equivalently, $S; \rho \models_n \Gamma \longrightarrow \Delta, (\triangleright \nabla\sigma.F[h\sigma/x_\tau])$, 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 $S; \rho \models_n \Gamma \longrightarrow \Delta, (\triangleright \forall h_{\sigma\to\tau}.\nabla\sigma.F[h\sigma/x_\tau])$, since h is fresh. By iterating Lemma 5.6, item 3, m times, we obtain $S; \rho \models_n \Gamma \longrightarrow \Delta, (\triangleright \nabla\sigma.\forall x_\tau.F)$, that is, $S; \rho \models_n \Gamma \longrightarrow \Delta, (\sigma \triangleright \forall x_\tau.F)$.

In the case of $(\forall L)$, let M be a λ-term of type τ, and assume $S; \rho_n \models_n \Gamma, (\sigma \triangleright F[M/x_\tau]) \longrightarrow \Delta$. Assume also that $S; \rho \models_n J$ for every J in Γ, and $S; \rho \models_n (\sigma \triangleright \forall x_\tau.F)$. We aim to show that $S; \rho \models_n J'$ for some J' in Δ. By Lemma 5.6, item 3 again, the latter implies $S; \rho \models_n \forall h_{\sigma\to\tau}.\nabla\sigma.F[h\sigma/x_\tau]$. Instantiate $h_{\sigma\to\tau}$ by $\lambda\sigma.M$. It follows that $S; \rho \models_n \nabla\sigma.F[M/x_\tau]$, hence $S; \rho \models_n \sigma \triangleright F[M/x_\tau]$. Since $S; \rho \models_n J$ for every J in Γ and $S; \rho_n \models_n \Gamma, (\sigma \triangleright F[M/x_\tau]) \longrightarrow \Delta$, we conclude.

The other cases are immediate. □

6 Henkin Completeness

We shall show that the deduction system of Figure 5 is complete using a variant of the technique of Hintikka sets, a technique used to show that tableaux calculi are complete for first-order logic. 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 judgment* is an expression of the form $+J$ or $-J$, where J is a judgment. On the semantic side, we understand $+J$ as meaning "J is true", and $-J$ as "J is false". On the syntactic side, we see a sequent $J_1, \cdots, J_m \to J'_1, \cdots, J'_n$ as a collection of signed judgments $+J_1, \cdots, +J_m, -J'_1, \cdots, -J'_n$. We extend $\approx$ to signed judgments in the obvious way.

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

$\mathcal{T}$ *is* inconsistent *if and only if there are finitely many signed judgments* $+J_1, \ldots, +J_m, -J'_1, \ldots, -J'_n$ *in* $\mathcal{T}$ *such that the sequent* $J_1, \cdots, J_m \to J'_1, \cdots, J'_n$ *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* $h_{\sigma\to\tau}$ *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 6.2. *A consistent theory cannot contain both $+J$ and $-J$ for the same judgment J; otherwise it would be inconsistent, using rule (Ax). It cannot contain a judgment of the form $+\sigma \triangleright \bot$ either (rule $(\bot L)$).*

Lemma 6.3. *Let $\mathcal{T}$ be a consistent theory.*

1. *For every signed judgment of the form $+\sigma \triangleright F \supset G$ in $\mathcal{T}$, $\mathcal{T} \cup \{-\sigma \triangleright F\}$ or $\mathcal{T} \cup \{+\sigma \triangleright G\}$ is consistent;*
2. *for every signed judgment of the form $-\sigma \triangleright F \supset G$ in $\mathcal{T}$, $\mathcal{T} \cup \{+\sigma \triangleright F, -\sigma \triangleright G\}$ is consistent;*
3. *for every signed judgment of the form $+\sigma \triangleright \forall x_\tau.F$ in $\mathcal{T}$, for every $M : \tau$, $\mathcal{T} \cup \{+\sigma \triangleright F[M/x_\tau]\}$ is consistent;*
4. *for every signed judgment of the form $-\sigma \triangleright \forall x_\tau.F$ in $\mathcal{T}$, for every variable $h : \sigma \to \tau$ 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 judgment of the form $+\sigma \triangleright \nabla x_\tau.F$ in $\mathcal{T}$, $\mathcal{T} \cup \{+\sigma, x : \tau \triangleright F\}$ is consistent;*
6. *for every signed judgment of the form $-\sigma \triangleright \nabla x_\tau.F$ in $\mathcal{T}$, $\mathcal{T} \cup \{-\sigma, x : \tau \triangleright F\}$ is consistent;*
7. *for every signed judgment $+J$ in $\mathcal{T}$, for every $J' \approx J$, $\mathcal{T} \cup \{+J'\}$ is consistent;*
8. *for every signed judgment $-J$ 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 $\Gamma \to \underbrace{(\sigma \triangleright F), \ldots, (\sigma \triangleright F)}_{m \text{ times}}, \Delta$ and $\Gamma', \underbrace{(\sigma \triangleright G), \ldots, (\sigma \triangleright G)}_{n \text{ times}} \to \Delta'$, where Γ and Γ' consist of judgments that appear with the $+$ sign in $\mathcal{T}$, Δ and Δ' consist of judgments 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 $\Gamma, (\sigma \triangleright F \supset G) \to \Delta$. 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 6.4. *Every finite consistent theory is contained in some Hintikka theory.*

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 judgments. Call a *task* either: a signed judgment $\pm J$, where J is not of the form $+\sigma \triangleright \forall x_\tau.F$; or a pair $(+\sigma \triangleright \forall x_\tau.F, M)$ 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.

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 $+\sigma \triangleright F \supset G$, and is in $\mathcal{T}_n$, then by Lemma 6.3, item 1, $\mathcal{T}_n \cup \{-\sigma \triangleright F\}$ or $\mathcal{T}_n \cup \{+\sigma \triangleright G\}$ is consistent: in the first case, let $\mathcal{T}_{n+1} = \mathcal{T}_n \cup \{-\sigma \triangleright F\}$, otherwise let $\mathcal{T}_{n+1} = \mathcal{T}_n \cup \{+\sigma \triangleright G\}$. If $\Theta_n = +\sigma \triangleright F \supset G$ is not in $\mathcal{T}_n$, then $\mathcal{T}_{n+1} = \mathcal{T}_n$.

If Θ_n is of the form $-\sigma \triangleright F \supset G$ and is in $\mathcal{T}_n$, then we let $\mathcal{T}_{n+1} = \mathcal{T}_n \cup \{+\sigma \triangleright F, -\sigma \triangleright G\}$, using Lemma 6.3, item 2. And if $\Theta_n = -\sigma \triangleright F \supset G$ is not in $\mathcal{T}_n$, then $\mathcal{T}_{n+1} = \mathcal{T}_n$.

We proceed similarly if Θ_n is of the form $\pm\sigma \triangleright \nabla x_\tau.F$, using item 5 or 6 of Lemma 6.3.

If Θ_n is of the form $(+\sigma \triangleright \forall x_\tau.F, M)$ where $+\sigma \triangleright \forall x_\tau.F$ is in $\mathcal{T}$, and M_n is of type τ, then we let $\mathcal{T}_{n+1} = \mathcal{T}_n \cup \{+\sigma \triangleright F[M_n/x_\tau]\}$, using Lemma 6.3, item 3. If $+\sigma \triangleright \forall x_\tau.F$ is not in $\mathcal{T}$, then we let $\mathcal{T}_{n+1} = \mathcal{T}_n$.

If Θ_n is of the form $-\sigma \triangleright \forall x_\tau.F$ and is in $\mathcal{T}_n$, then there is a variable h of type $\sigma \to \tau$ that is not free in $\mathcal{T}_n$ since $\mathcal{T}_n$ is finite. Relying on Lemma 6.3, item 4, we define $\mathcal{T}_{n+1}$ as $\mathcal{T}_n \cup \{-\sigma \triangleright F[h\sigma/x_\tau]\}$. If $\Theta_n = -\sigma \triangleright \forall x_\tau.F$ 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 6.3, items 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. □

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 $\sigma = x_1 : \tau_1, x_2 : \tau_2, \cdots, x_n : \tau_n$ (of *length* n), let θ_σ be the substitution $[\mathrm{new}_1^{T[\![\tau_1]\!]}/x_1, \mathrm{new}_2^{T[\![\tau_2]\!]}/x_2, \cdots, \mathrm{new}_n^{T[\![\tau_n]\!]}/x_n]$. This is a substitution at level n.

Lemma 6.5. *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 $\tau_1, \tau_2, \cdots, \tau_k$, as the set of k-tuples $(M_1\theta_\sigma, M_2\theta_\sigma, \cdots, M_k\theta_\sigma)$ such that $+\sigma \triangleright P(M_1, M_2, \cdots, M_k) \in \mathcal{T}$ for some local signature σ of length n. This defines a Henkin structure such that:*

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

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

Proof. First look at the case where $J = \sigma \triangleright P(M_1, M_2, \cdots, M_k)$, where σ is of length n. If $+J \in \mathcal{T}$ then by definition $(M_1\theta_\sigma, M_2\theta_\sigma, \cdots, M_k\theta_\sigma)$ is in $T[\![P]\!]_n$. By Lemma 4.9, $(T[\![M]\!]_1\theta_\sigma, T[\![M]\!]_2\theta_\sigma, \cdots, T[\![M]\!]_k\theta_\sigma)$ is in $T[\![P]\!]_n$, so $T; \theta_\sigma \models_n P(M_1, M_2, \cdots, M_k)$. In other words, $T; \epsilon \models_0 \sigma \triangleright P(M_1, M_2, \cdots, M_k)$. If $-J \in \mathcal{T}$, then $+J$ in not not $\mathcal{T}$ (Fact 6.2), so $(M_1\theta_\sigma, M_2\theta_\sigma, \cdots, M_k\theta_\sigma)$ is not in $T[\![P]\!]_n$. By a similar argument, $T; \epsilon \not\models_0 \sigma \triangleright P(M_1, M_2, \cdots, M_k)$.

Now assume $J = \sigma \triangleright \bot$. Since every Hintikka theory is consistent, and using Fact 6.2, $+J$ is not in $\mathcal{T}$. If $-J$ is in $\mathcal{T}$, $T; \epsilon \not\models_0 \sigma \triangleright \bot$.

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

Now assume $J = \forall x_\tau.F$. If $+J \in \mathcal{T}$, then $+\sigma \triangleright F[M/x_\tau]$ is in $\mathcal{T}$ for every λ-term $M : \tau$. By induction hypothesis, this implies that $T; \epsilon \models_0 \sigma \triangleright F[M/x_\tau]$ for every λ-term $M : \tau$. We wish to show that $T; \epsilon \models_0 \sigma \triangleright \forall x_\tau.F$. Using Lemma 5.6, item 3, we know that the latter is equivalent to $T; \epsilon \models_0 (\triangleright \forall h_{\sigma\to\tau}.\nabla\sigma.F[h\sigma/x_\tau])$. Hence we must show that for every $N \in T[\![\sigma \to \tau]\!]_0$ (i.e., for every ordinary λ-term $N : \sigma \to \tau$, by Remark 4.6), $T; \epsilon[h \mapsto N] \models_0 \nabla\sigma.F[h\sigma/x_\tau]$. Using Lemma 5.5, and since $T[\![N]\!]\epsilon = \widehat{N}(\epsilon) = N$ (Lemma 4.9), this boils down to showing that $T; \epsilon \models_n (\nabla\sigma.F[h\sigma/x_\tau])[h \mapsto N]$, that is, $T; \epsilon \models_n \sigma \triangleright F[N\sigma/x_\tau]$ for every $N : \sigma \to \tau$ that has no free variable in the list σ. Since $T; \epsilon \models_0 \sigma \triangleright F[M/x_\tau]$ for every λ-term $M : \tau$, this is clear.

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

The cases when $J = \nabla x_\tau.F$ is easy. □

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

Proposition 6.6. *Assume there is a unique base type ι. Let $\Gamma \to \Delta$ be a sequent such that $H; \epsilon \models_0 \Gamma \to \Delta$ for every Herbrand structure H. Then $\Gamma \to \Delta$ is derivable using the rules of $FO\lambda^{\nabla}$, without (Cut).*

Proof. Assume $\Gamma \to \Delta$ is not derivable. Let $\mathcal{T}_0$ be the theory containing the signed judgments $+J$, $J \in \Gamma$ 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 6.4, $\mathcal{T}_0$ is contained in some Hintikka theory $\mathcal{T}$. Using the Henkin structure H defined in Lemma 6.5—this is a Herbrand structure—we obtain that $H; \epsilon \not\models_0 \Gamma \to \Delta$, a contradiction. □

As a corollary, we obtain:

Theorem 6.7 (Henkin Completeness)**.** *Assume there is a unique base type ι. The Henkin semantics is complete for $FO\lambda^{\nabla}$: every valid sequent is derivable in $FO\lambda^{\nabla}$, and even by a cut-free proof.*

7 Incompleteness for Standard Structures

Standard structures are *incomplete* for $FO\lambda^{\nabla}$. This has nothing to do with the nabla quantifier, and is only due to the higher-order nature of the terms that $FO\lambda^{\nabla}$ is based on, and to the fact that $\boldsymbol{\nabla}$ validates the weak axiom of choice (Corollary 3.6).

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:

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

Lemma 7.1. *(AC) holds in every standard structure S.*

Proof. Assume that $S; \rho \models_n \forall x_{\varphi}.\exists y_{\tau}.F$, in other words, for every $d \in S[\![\varphi]\!]_n$, there is an $e \in S[\![\tau]\!]_n$ such that $S; \rho[x \mapsto d, y \mapsto e] \models_n F$. Let $R \subseteq S[\![\varphi]\!]_n \times S[\![\tau]\!]_n$ be the set of all pairs (d, e) such that $S; \rho[x \mapsto d, y \mapsto e] \models_n F$. Corollary 3.6 applies, so there is an element $f = (f_m)_{m \geq n}$ of $S[\![\varphi \to \tau]\!]_n$ such that for every $d \in S[\![\varphi]\!]_n$, $S; \rho[x \mapsto d, y \mapsto f_n(d)] \models_n F$. In other words, $S; \rho \models_n \exists h_{\varphi\to\tau}.\forall x_{\tau}.F[hx/y]$. □

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

Lemma 7.2. *The sequent $\to \triangleright(\forall x_{\iota}.\exists_{\iota}.P(x, y)) \supset (\exists h_{\iota\to\iota}.\forall x_{\iota}.P(x, hx))$ 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 5.7. □

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

8 Π_1-Completeness

However, we claim that we regain completeness for the fragment consisting of Π_1 formulae (which we define later). This requires some λ-calculus machinery to relate the interpretation $\mathsf{S}_1[\![M]\!]$ of terms M in a specific standard universe S_1 and the interpretation $T[\![M]\!]$ in the term universe (which we do now).

We start with another standard universe S_0, which is defined by specifying:

$$\mathsf{S}_0[\![\tau]\!] \quad = \quad T[\![\tau]\!] \tag{3}$$

for every base type τ, and letting $\mathsf{S}_0[\![\varphi \to \tau]\!]$ be $[\mathsf{S}_0[\![\varphi]\!] \to \mathsf{S}_0[\![\tau]\!]]$ for all arrow types. We shall define S_1 later, by specifying $\mathsf{S}_1[\![\tau]\!]$ as well-chosen variants (Definition 3.3) of $\mathsf{S}_0[\![\tau]\!]$.

Beware that (3) will fail for non-base types τ: for arrow types, $T[\![\varphi \to \tau]\!]$ is a nabla-set of terms, in particular $T[\![\iota \to \iota]\!]_n$ is countable for every n; on the contrary, $\mathsf{S}_0[\![\varphi \to \tau]\!] = [\mathsf{S}_0[\![\varphi]\!] \to \mathsf{S}_0[\![\tau]\!]]$, and in particular $\mathsf{S}_0[\![\iota \to \iota]\!]_0$ is uncountable.

Recall the notion of nabla-relation from Proposition 3.5. We define the following Kripke logical relation.

Definition 8.1. *Define the nabla-relations $R[\tau]$, for each type τ, between $T[\![\tau]\!]$ and $\mathsf{S}_0[\![\tau]\!]$, by:*

1. *$R[\tau]_n$ is equality, for each base type τ and every $n \in \mathbb{N}$;*
2. *for every $n \in \mathbb{N}$, for every $M \in T[\![\varphi \to \tau]\!]_n$, for every $f = (f_m)_{m \geq n} \in \mathsf{S}_0[\![\varphi \to \tau]\!]_n$, $M\ R[\varphi \to \tau]_n\ f$ if and only if, for every $m \geq n$, for all $N \in T[\![\varphi]\!]_m$ and $d \in \mathsf{S}_0[\![\varphi]\!]_m$ such that $N\ R[\varphi]_m\ d$, $MN\ R[\tau]_m\ f_m(d)$.*

We check that this indeed defines nabla-relations, by induction on types. In the second case, if $M\ R[\varphi \to \tau]_n\ f$ was obtained by checking that for every $m \geq n$, for all $N \in T[\![\varphi]\!]_m$ and $d \in \mathsf{S}_0[\![\varphi]\!]_m$ such that $N\ R[\varphi]_m\ d$, $MN\ R[\tau]_m\ f_m(d)$, then that is true in particular for every $m \geq n+1$. Recalling that $\mathrm{old}_n^{[\varphi \to \tau]}(M) = M$ and $\mathrm{old}_n^{\mathsf{S}_0[\![\varphi \to \tau]\!]}(f) = (f_m)_{m \geq n+1}$, we obtain $\mathrm{old}_n^{[\varphi \to \tau]}(M)\ R[\varphi \to \tau]_{n+1}\ \mathrm{old}_n^{\mathsf{S}_0[\![\varphi \to \tau]\!]}(f)$.

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 8.2 (Basic Lemma of Logical Relations)**.** *For every $n \in \mathbb{N}$, for every substitution θ at level n whose domain $\mathrm{dom}\,\theta$ is finite, for every environment ρ at level n, we say that $\theta\ R\ \rho$ if and only if for every variable $z_\psi \in \mathrm{dom}\,\theta$, $\theta(z_\psi)\ R[\psi]_n\ \rho(z_\psi)$.*

For every λ-term M of type τ whose free variables are in $\mathrm{dom}\,\theta$, if $\theta\ R\ \rho$ then $M\theta\ R[\tau]_n\ \mathsf{S}_0[\![M]\!]_n\rho$.

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

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 $f = \mathsf{S}_0[\![M_1]\!]_n\rho$. It also tells us that $M_2\theta\ R[\varphi]_n\ \mathsf{S}_0[\![M_2]\!]_n\rho$. Using the definition of $R[\varphi \to \tau]_n$ with $m = n$, we obtain that $M_1M_2\ R[\tau]_n\ f_n(\mathsf{S}_0[\![M_2]\!]_n\rho) = \mathsf{S}_0[\![M_1M_2]\!]_n\rho$.

If M is a λ-abstraction $\lambda x_\varphi.P$ of type $\varphi \to \tau$, then let $f = (f_m)_{m \geq n} = \mathsf{S}_0[\![M]\!]_n\rho$. We must show that, for every $m \geq n$, for all $N \in [\varphi]_m$ and $d \in \mathsf{S}_0[\![\varphi]\!]_m$ such that $N\ R[\varphi]_m\ d$, $(M\theta)N\ R[\tau]_m\ f_m(d)$.

By α-renaming, we may assume that x_φ is not in $\mathrm{dom}\,\theta$, and not free in any term $\theta(x_\psi)$, $x_\psi \in \mathrm{dom}\,\theta$. Let $\theta' = \theta[x_\varphi \mapsto N]$, and $\rho' = \mathrm{old}_{n \to m}^{\mathtt{Env}}(\rho)[x_\varphi \mapsto d]$. 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 the fact that $R[\psi]$ is a nabla-relation in the other cases.

By induction hypothesis, $P\theta'\ R[\tau]_m\ \mathsf{S}_0[\![P]\!]_m\rho'$. We conclude by noting that $P\theta'$ is equal (up to $\beta\eta$-conversion) to $(M\theta)N$, and that $\mathsf{S}_0[\![P]\!]_m\rho' = f_m(d)$. □

Proposition 8.3. *There are families of nabla-maps $s_\tau\colon T[\![\tau]\!] \to \mathsf{S}_0[\![\tau]\!]$ and $r_\tau\colon \mathsf{S}_0[\![\tau]\!] \to T[\![\tau]\!]$, indexed by types τ, such that the following implications hold for all $M \in T[\![\tau]\!]_n$ and $d \in \mathsf{S}_0[\![\tau]\!]_n$:*

$$(s_\tau)_n(M) = d \quad\Rightarrow\quad M\ R[\tau]_n\ d \tag{4}$$

$$M\ R[\tau]_n\ d \quad\Rightarrow\quad (r_\tau)_n(d) = M. \tag{5}$$

Proof. Those are built by structural induction on τ. For a base type τ, we define both s_τ and r_τ 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:

$$T[\![\varphi\to\tau]\!] \times \mathsf{S}_0[\![\varphi]\!] \xrightarrow{\mathrm{id}_{T[\![\varphi\to\tau]\!]}\times r_\varphi} T[\![\varphi\to\tau]\!] \times T[\![\varphi]\!] \xrightarrow{\mathsf{App}} T[\![\tau]\!] \xrightarrow{s_\tau} \mathsf{S}_0[\![\tau]\!].$$

Here $\mathsf{App}\colon T[\![\varphi\to\tau]\!] \times T[\![\varphi]\!] \to T[\![\tau]\!]$ is the nabla-map defined by letting $\mathsf{App}_n(M,N)$ 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 (4) holds at type $\varphi\to\tau$, that is, that for every $M \in T[\![\varphi\to\tau]\!]_n$ and for $f = (s_{\varphi\to\tau})_n(M) \in \mathsf{S}_0[\![\varphi\to\tau]\!]_n$, $M\ R[\varphi\to\tau]_n\ f$. To show this, let $m \geq n$, and N and d be such that $N\ R[\varphi]_m\ d$. We must show that $MN\ R[\tau]_m\ f_m(d)$, where $f = (f_m)_{m\geq n}$. Since $f = (s_{\varphi\to\tau})_n(M)$, f_m maps d to $(\widetilde{s}_{\varphi\to\tau})_m(\mathrm{old}_{n\to m}^{T[\![\varphi\to\tau]\!]}(M), d)$, namely, to $(s_\tau)_m(\mathsf{App}(\mathrm{old}_{n\to m}^{T[\![\varphi\to\tau]\!]}(M), (r_\varphi)_m(d))) = (s_\tau)_m(M((r_\varphi)_m(d)))$, where the application of M to $(r_\varphi)_m(d)$ is syntactic application. Since $N\ R[\varphi]_m\ d$, by induction hypothesis on φ, $(r_\varphi)_m(d) = N$, so $f_m(d) = (s_\tau)_m(MN)$. By induction hypothesis on τ, $MN\ R[\tau]_m\ f_m(d)$.

In order to build $r_{\varphi\to\tau}$, we show that, for every $f \in \mathsf{S}_0[\![\varphi\to\tau]\!]_n$, there is at most one element $M \in T[\![\varphi\to\tau]\!]_n$ such that $M\ R[\varphi\to\tau]_n\ f$. Imagine there are two, M_1 and M_2. By abuse of language, consider M_1 and M_2 as terms, and pick a variable X_φ that is not free in M_1, and not free in M_2. Let $d = (s_\varphi)_n(X_\varphi)$. By induction hypothesis, $X_\varphi\ R[\varphi]_n\ d$, so $M_1X_\varphi\ R[\tau]_n\ f_n(d)$ and $M_2X_\varphi\ R[\tau]_n\ f_n(d)$. By induction hypothesis again, $(r_\tau)_n(f_n(d))$ 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, $M_1 = M_2$.

This would lead us to define $(r_{\varphi\to\tau})_n(f)$ as the unique $M \in T[\![\varphi\to\tau]\!]_n$ such that $M\ R[\varphi\to\tau]_n\ f$ if such an M exists. That would not define a nabla-map. Instead, we define $(r_{\varphi\to\tau})_n(f)$ as follows. If there is an $m \geq n$ and an $M \in T[\![\varphi\to\tau]\!]_n$ such that $M\ R[\varphi\to\tau]_m\ \mathrm{old}_{n\to m}^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f)$, then we let $(r_{\varphi\to\tau})_n(f)$ be this M. (To show that this makes sense, we ought to write $\mathrm{old}_{n\to m}^{T[\![\varphi\to\tau]\!]}(M)$ instead of M on the left of $R[\varphi\to\tau]_m$, but of course $\mathrm{old}_{n\to m}^{T[\![\varphi\to\tau]\!]}(M) = M$.) Otherwise, we let $(r_{\varphi\to\tau})_n(f)$ be some fixed term of type $\varphi\to\tau$, independent of n, say a variable $z_{\varphi\to\tau}$.

This is well-defined: if there are two natural numbers $m_1, m_2 \geq n$ and two elements M_1, M_2 in $T[\![\varphi\to\tau]\!]_n$ such that $M_1\ R[\varphi\to\tau]_{m_1}\ \mathrm{old}_{n\to m_1}^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f)$ and $M_2\ R[\varphi\to\tau]_{m_2}\ \mathrm{old}_{n\to m_2}^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f)$, then we would have $M_1\ R[\varphi\to\tau]_m\ \mathrm{old}_{n\to m}^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f)$ and $M_2\ R[\varphi\to\tau]_m\ \mathrm{old}_{n\to m}^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f)$ where $m = \max(m_1, m_2)$, using the fact that $R[\varphi\to\tau]$ is a nabla-relation. We have seen that this implies $M_1 = M_2$ (up to $\beta\eta$).

By construction, (5) is satisfied at type $\varphi\to\tau$; in other words, $M\ R[\varphi\to\tau]_n\ f$ implies $(r_{\varphi\to\tau})_n(f) = M$. This is by definition, taking $m = n$. It remains to check that $r_{\varphi\to\tau}$ is a nabla-map. Since old_n works as the identity map on syntactic nabla-sets, that amounts to checking

that $(r_{\varphi\to\tau})_n(f) = (r_{\varphi\to\tau})_{n+1}(\mathrm{old}_n^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f))$. If there is an $m \geq n$ and an $M \in T[\![\varphi\to\tau]\!]_n$ such that $M\ R[\varphi\to\tau]_m\ \mathrm{old}_{n\to m}^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f)$, then $(r_{\varphi\to\tau})_n(f) = M$. In that case, using the fact that $R[\varphi\to\tau]$ is a nabla-map, $M\ R[\varphi\to\tau]_{m+1}\ \mathrm{old}_{n\to m+1}^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f) = \mathrm{old}_{n+1\to m+1}^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(\mathrm{old}_n^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f))$, so $(r_{\varphi\to\tau})_{n+1}(\mathrm{old}_n^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f)) = M$. In case there is no such $m \geq n$ and no such M, then $(r_{\varphi\to\tau})_n(f)$ and $(r_{\varphi\to\tau})_{n+1}(\mathrm{old}_n^{\mathsf{S}_0[\![\varphi\to\tau]\!]}(f))$ are both equal to $z_{\varphi\to\tau}$. □

Remark 8.4. *It would be tempting to produce a different proof of Proposition 8.3 by giving an explicit formula for $r_{\varphi\to\tau}$. The following seems to work—but does not:*

$$(r_{\varphi\to\tau})_n(f) \quad = \quad \lambda X_\varphi.(r_\tau)_n(f_n((s_\varphi)_n(X_\varphi))) \tag{6}$$

for every $f = (f_m)_{m\geq n} \in \mathsf{S}_0[\![\varphi\to\tau]\!]_n$. One can check that this defines a nabla-map. If you try to prove (5) at type $\varphi\to\tau$ with that formula, you will obtain that $(r_{\varphi\to\tau})_n(d)$ is equal to $\lambda X_\varphi.MX_\varphi$. That is only η-convertible to M provided X_φ is not free in M, and α-renaming X_φ into a fresh variable should do the trick... but (6) is not invariant under α-renaming! If you pick a different variable X_φ, you will in general get a different term. This is why we defined $r_{\varphi\to\tau}$ in a roundabout way. A similar difficulty occurs in the classical proof [7] of a similar result by Harvey Friedman [2].

The paper we have just cited by H. Friedman shows the following. Define an interpretation of simply-typed λ-terms up to $\beta\eta$-conversion in **Set** by defining a set $[\![\tau]\!]$ for each basic type τ, and letting $[\![\varphi\to\tau]\!]$ be the set of all functions from $[\![\varphi]\!]$ to $[\![\tau]\!]$. Interpret λ-terms in the obvious way. A λ-term is *closed* if and only if it has no free variable. If M and N are $\beta\eta$-equivalent closed λ-terms, then $[\![M]\!] = [\![N]\!]$, and Friedman's result states that there is a way of fixing $[\![\tau]\!]$ for each basic type τ so that the converse implication holds.

As a parenthesis, a similar result holds in $\boldsymbol{\nabla}$, using our notion of interpretation of λ-terms, as we now claim. For a closed term M, and a given nabla-universe S, $S[\![M]\!]_n\rho$ does not depend on the environment ρ, and we write $S[\![M]\!]_n$ for $S[\![M]\!]_n\rho$ in that case. This allows us to state:

Corollary 8.5. *The semantics of λ-terms is* equationally complete*: there is a nabla-universe S_0 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. *$\mathsf{S}_0[\![M]\!]_0 = \mathsf{S}_0[\![N]\!]_0$.*

Proof. $1 \Rightarrow 2$ is obvious. For the converse implication, assume 2. Let $d = \mathsf{S}_0[\![M]\!]_0 = \mathsf{S}_0[\![N]\!]_0$. By the Basic Lemma 8.2, used with $n = 0$ and the empty substitution θ, $M\ R[\tau]_0\ d$ and $N\ R[\tau]_0\ d$. Apply Proposition 8.3 to obtain that $M = (r_\tau)_0(d)$ and $N = (r_\tau)_0(d)$ (up to $\beta\eta$-conversion), so $M = N$. □

This ends our parenthesis.

Definition 8.6. *Let S_1 be the standard universe defined so that $\mathsf{S}_1[\![\tau]\!]$ is the variant of $\mathsf{S}_0[\![\tau]\!]$ where $\mathrm{new}_{n+1}^{\mathsf{S}_1[\![\tau]\!]} = (s_\tau)_{n+1}(\mathrm{new}_{n+1}^{T[\![\tau]\!]})$ for every $n \in \mathbb{N}$ and every type τ.*

It is easy to see that $\mathsf{S}_1[\![\varphi\to\tau]\!]$ is a variant of $[\mathsf{S}_1[\![\varphi]\!] \to \mathsf{S}_1[\![\tau]\!]]$ for all types φ and τ. What is perhaps less obvious is that $\mathrm{new}_{n+1}^{\mathsf{S}_1[\![\tau]\!]}$ is indeed outside $\mathrm{Im}\,\mathrm{old}_n^{\mathsf{S}_1[\![\tau]\!]}$, as required in the definition of a nabla-set. For that, note that $(r_\tau)_{n+1}(\mathrm{new}_{n+1}^{\mathsf{S}_1[\![\tau]\!]}) = (r_\tau \circ s_\tau)_{n+1}(\mathrm{new}_{n+1}^{T[\![\tau]\!]}) = \mathrm{new}_{n+1}^{T[\![\tau]\!]}$, since $r_\tau \circ s_\tau = \mathrm{id}_{T\tau}$ is a consequence of Proposition 8.3. If $\mathrm{new}_{n+1}^{\mathsf{S}_1[\![\tau]\!]}$ was equal to $\mathrm{old}_n^{\mathsf{S}_1[\![\tau]\!]}(d)$ for some d, then $(r_\tau)_{n+1}(\mathrm{new}_{n+1}^{\mathsf{S}_1[\![\tau]\!]})$ would be equal to $\mathrm{old}_n^{T[\![\tau]\!]}((r_\tau)_n(d))$, since r_τ is a nabla-map, and that would entail $\mathrm{new}_{n+1}^{T[\![\tau]\!]} = \mathrm{old}_n^{T[\![\tau]\!]}((r_\tau)_n(d))$, a contradiction.

Definition 8.7. *A* Δ_0 formula *of* $FO\lambda^\nabla$ *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* $\forall x_{1\,\tau_1}, \cdots, x_{p\,\tau_p}.G$*, where* G *is a* Δ_0 *formula.*

Proposition 8.8. *Let* H *be a Herbrand structure. Define a standard structure* S_1^H *on the standard universe* S_1 *by letting:*

$$\mathsf{S}_1^H[\![P]\!]_n \quad = \quad \{(d_1, d_2, \cdots, d_k) \in \prod_{i=1}^{k} \mathsf{S}_1[\![\tau_i]\!] \mid (r_{\tau_1}(d_1), r_{\tau_2}(d_2), \cdots, r_{\tau_k}(d_k)) \in H[\![P]\!]_n\}.$$

for every relation symbol P *of arity* $\tau_1, \tau_2, \cdots, \tau_k$ *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* $\theta\ R\ \rho$*:*

1. *for every* Δ_0 *formula* G *whose free variables are included in* $\operatorname{dom}\theta$*,* $\mathsf{S}_1^H; \rho \models_n G$ *if and only if* $H; \theta \models_n G$*;*
2. *for every* Π_1*-formula* F *whose free variables are included in* $\operatorname{dom}\theta$*, if* $\mathsf{S}_1^H; \rho \models_n F$ *then* $H; \theta \models_n F$*.*

Proof. 1. By structural induction on G.

If G is an atomic formula $P(M_1, M_2, \cdots, M_k)$, where each M_i has type τ_i, then $\mathsf{S}_1^H; \rho \models_n G$ if and only if $(\mathsf{S}_1[\![M_1]\!]_n\rho, \mathsf{S}_1[\![M_2]\!]_n\rho, \cdots, \mathsf{S}_1[\![M_k]\!]_n\rho)$ is in $\mathsf{S}_1^H[\![P]\!]_n$. By the Basic Lemma (Lemma 8.2), $M_i\theta\ R[\tau_i]_n\ \mathsf{S}_1[\![M]\!]_i\rho$, so, using Proposition 8.3 and specifically (5), $r_{\tau_i}(\mathsf{S}_1[\![M]\!]_i\rho) = M_i\theta$. Using the definition of $\mathsf{S}_1^H[\![P]\!]_n$, we obtain that $\mathsf{S}_1^H; \rho \models_n G$ if and only if $(M_1\theta, M_2\theta, \cdots, M_k\theta) \in H[\![P]\!]_n$. The latter is equivalent to $(T[\![M_1]\!]_n\theta, T[\![M_2]\!]_n\theta, \cdots, T[\![M_k]\!]_n\theta) \in H[\![P]\!]_n$ (Fact 4.9), hence to $H; \theta \models_n G$.

If G is a first-order quantified formula $\forall x_\iota.G'$, then $\mathsf{S}_1^H; \rho \models_n G$ if and only if $\mathsf{S}_1^H; \rho[x \mapsto d] \models_n G'$ for every $d \in \mathsf{S}_1[\![\iota]\!]$. Since $\mathsf{S}_1[\![\iota]\!] = H[\![\iota]\!]$, and $R[\iota]_n$ is the identity relation, $\theta[x \mapsto d]\ R\ \rho[x \mapsto d]$ for every $d \in \mathsf{S}_1[\![\iota]\!]$. Hence $\mathsf{S}_1^H; \rho \models_n G$ if and only $H; \theta[x \mapsto d] \models_n G'$ for every $d \in \mathsf{S}_1[\![\iota]\!] = H[\![\iota]\!]$, if and only if $H; \theta \models_n G$.

The other cases follow by an easy induction, except perhaps when G is of the form $\nabla x_\tau.G'$. Then $\mathsf{S}_1^H; \rho \models_n G$ if and only if $\mathsf{S}_1^H; \mathrm{old}^{\mathsf{Env}}(\rho)[x \mapsto \mathrm{new}_{n+1}^{\mathsf{S}_1[\![\tau]\!]}] \models_{n+1} G'$. Let $\rho' = \mathrm{old}^{\mathsf{Env}}(\rho)[x \mapsto \mathrm{new}_{n+1}^{\mathsf{S}_1[\![\tau]\!]}]$, $\theta' = \theta[x \mapsto \mathrm{new}_{n+1}^{T[\![\tau]\!]}]$. Since we chose $\mathrm{new}_{n+1}^{\mathsf{S}_1[\![\tau]\!]} = (s_\tau)_{n+1}(\mathrm{new}_{n+1}^{T[\![\tau]\!]})$, Proposition 8.3 (and specifically (4)) implies that $\mathrm{new}_{n+1}^{T[\![\tau]\!]}\ R[\tau]_{n+1}\ \mathrm{new}_{n+1}^{\mathsf{S}_1[\![\tau]\!]}$. Hence $\theta'\ R\ \rho'$, and we can apply the induction hypothesis: $\mathsf{S}_1^H; \rho' \models_{n+1} G'$ if and only if $H; \theta' \models_{n+1} G'$, and therefore $\mathsf{S}_1^H; \rho \models_n G$ if and only if $H; \theta \models G$.

2. Let now F be a Π_1 formula $\forall x_{1\,\tau_1}, \cdots, x_{p\,\tau_p}.G$, where G is a Δ_0 formula. If $\mathsf{S}_1^H; \rho \models_n F$, then $\mathsf{S}_1^H; \rho[x_1 \mapsto d_1, \cdots, x_p \mapsto d_p] \models_n G$ for all values $d_1 \in \mathsf{S}_1[\![\tau_1]\!]_n$, ..., $d_p \in \mathsf{S}_1[\![\tau_p]\!]_n$. This is in particular true if we pick $d_1 = (s_{\tau_1})_n(N_1)$, ..., $d_p = (s_{\tau_p})_n(N_p)$ for arbitrary elements $N_1 \in T[\![\tau_1]\!]_n$, ..., $N_p \in T[\![\tau_p]\!]_n$. Let $\rho' = \rho[x_1 \mapsto d_1, \cdots, x_p \mapsto d_p]$, and $\theta' = \theta[x_1 \mapsto N_1, \cdots, x_p \mapsto N_p]$. By Proposition 8.3, and specifically (4), $\theta'\ R\ \rho'$. By part 1 of the Proposition, we conclude that $H; \theta[x_1 \mapsto N_1, \cdots, x_p \mapsto N_p] \models_n G$ for all $N_1 \in T[\![\tau_1]\!]_n$, ..., $N_p \in T[\![\tau_p]\!]_n$, that is, that $H; \theta \models_n F$. □

Write $S \models_0 F$ if $S; \rho \models_0 F$, 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* $S \models_0 F$ *for every standard structure* S*, then* $\to \triangleright F$ *is derivable in* $FO\lambda^\nabla$*, and even by a cut-free proof.*

Proof. Let ρ be any environment at level 0: then $\epsilon\ R\ \rho$. Hence we can use Proposition 8.8, item 2, and conclude that $H; \epsilon \models_n F$. By Theorem 6.7, $\rightarrow \triangleright F$ has a cut-free proof in $FO\lambda^{\nabla}$. $\square$

Remark 8.10. *Proposition 8.9 in particular implies that $FO\lambda^{\nabla}$ 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

If $FO\lambda^{\nabla}$ plus (AC) complete for standard models? What would happen if there were more than one base type ι? Can we extend the present results to the logic of Abella, 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 $\nabla y.\nabla x.F(x,y)$? Does all this extend to intuitionistic versions of $FO\lambda^{\nabla}$?

10 Conclusion

Happy 60th, Dale!

References

[1] Pierre-Louis Curien. *Categorical Combinators, Sequential Algorithms, and Functional Programming.* Birkhäuser, Boston, MA, 1993.

[2] Harvey Friedman. Equality between functionals. In Rohit Parikh, editor, *Logic Colloquium 1972-73*, volume 453 of *Lecture Notes in Mathematics*, pages 22–37. Springer-Verlag, 1975.

[3] Murdoch Jamie Gabbay and Andrew M. Pitts. A new approach to abstract syntax involving binders. In *14th Annual Symposium on Logic in Computer Science*, pages 214–224. IEEE Computer Society Press, Washington, 1999.

[4] Andrew Gacek. The Abella interactive theorem prover (system description). In A. Armando, P. Baumgartner, and G. Dowek, editors, *Proceedings of IJCAR 2008*, volume 5195 of *Lecture Notes in Artificial Intelligence*, pages 154–161. Springer, August 2008.

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

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

[7] John C. Mitchell. *Foundations for Programming Languages.* MIT Press, 1985.

[8] Ulrich Schöpp. Modelling generic judgments. *Electronic Notes in Theoretical Computer Science*, 174(5):19–35, 2007. Proceedings of the First International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2006).