

A Simpler Proof Theory for Nominal Logic

James Cheney

University of Edinburgh
(jcheney@inf.ed.ac.uk)

Abstract. Nominal logic is a variant of first-order logic equipped with a “fresh-name quantifier” \mathbb{N} and other features useful for reasoning about languages with bound names. Its original presentation was as a Hilbert axiomatic theory, but several attempts have been made to provide more convenient Gentzen-style sequent or natural deduction calculi for nominal logic. Unfortunately, the rules for \mathbb{N} in these calculi involve complicated side-conditions, so using and proving properties of these calculi is difficult. This paper presents an improved sequent calculus NL^{\Rightarrow} for nominal logic. Basic results such as cut-elimination and conservativity with respect to nominal logic are proved. Also, NL^{\Rightarrow} is used to solve an open problem, namely relating nominal logic’s \mathbb{N} -quantifier and the self-dual ∇ -quantifier of Miller and Tiu’s $FO\lambda^{\nabla}$.

1 Introduction

Gabbay and Pitts [8] have introduced a new way of reasoning about names and binding, in which α -equivalence and capture-avoiding substitution can be defined in terms of the basic concepts of *swapping* and *freshness*. This approach provides a cleaner treatment of α -equivalence than the classical first-order approach in which α -equivalence and capture-avoiding substitution are defined by mutual recursion. On the other hand, unlike higher-order techniques for dealing with names and binding, the semantics of this model of name-binding is relatively straightforward, so well-understood mathematical tools like structural induction can be used to reason about syntax with bound names.

These ideas have been incorporated into a logic called *nominal logic* [12]. Nominal logic is typed, first-order equational logic augmented with:

- *name-types* ν, ν', \dots inhabited by countably many *names* a, b, \dots ;
- a *swapping operation* $(- -) \cdot - : \nu \rightarrow \nu \rightarrow \tau \rightarrow \tau$ for each name-type ν and type τ , which acts on values by exchanging occurrences of names;
- a *freshness relation* $- \# - : \nu \rightarrow \tau \rightarrow o^1$ for each name-type ν and type τ , that holds between a name and a value independent of the name;
- an *abstraction type constructor* $\langle - \rangle -$ and *abstraction function symbol* $\langle - \rangle - : \nu \rightarrow \tau \rightarrow \langle \nu \rangle \tau$ which constructs values equal up to consistent renaming, axiomatized as follows:

$$\forall a, b, x, y. \langle a \rangle x = \langle b \rangle y \iff (a = b \wedge x = y) \vee (a \# y \wedge x = (a b) \cdot y);$$

- a *some/any fresh-name quantifier* \mathbb{N} that is self-dual ($\neg \mathbb{N}a. \varphi \iff \mathbb{N}a. \neg \varphi$);
- and *freshness* and *equivariance* principles which state that *fresh names can always be chosen* and *truth is preserved by name-swapping*, respectively.

¹ o is the type of propositions

1.1 The Problem

This paper is concerned with developing simple rules for reasoning with the \mathbb{N} -quantifier. Pitts’ original formalization of nominal logic was a Hilbert-style collection of first-order axioms (which we call NL). There were no new inference rules for \mathbb{N} . Instead, \mathbb{N} was defined using the axiom scheme $\forall \bar{x}. (\mathbb{N}a.\varphi \iff \exists a.a \# \bar{x} \wedge \varphi)$, where $FV(\varphi) \subseteq \{a, \bar{x}\}$. While admirable from a reductionist point of view, Hilbert systems have well-known deficiencies for modeling actual reasoning. Instead, Gentzen-style *natural deduction* and *sequent* systems provide a more intuitive approach to formal reasoning in which logical connectives are explained as *proof-search* operations. Gentzen systems are especially useful for computational applications, such as automated deduction and logic programming. A sequent calculus formalization would also be convenient for relating nominal logic with other logics by proof-theoretic translations.

Gentzen-style rules for \mathbb{N} have been considered in previous work. Pitts [12] proposed sequent and natural deduction rules for \mathbb{N} based on the observation that

$$\forall a.(a \# \bar{x} \supset \varphi(a, \bar{x})) \supset \mathbb{N}a.\varphi(a, \bar{x}) \supset \exists a.(a \# \bar{x} \wedge \varphi(a, \bar{x})) .$$

These rules (see Figure 1(NL)) are symmetric, emphasizing \mathbb{N} ’s self-duality. However, they are not closed under substitution, which greatly complicates the the proof of cut-elimination or proof-normalization properties.

Gabbay [6] introduced Fresh Logic (FL), an intuitionistic natural deduction calculus for nominal logic, and studied semantic issues including soundness and completeness as well as proving proof-normalization. Gabbay and Cheney [7] presented a similar sequent calculus called FL_{Seq} . In FL , Gabbay introduced a technical device called *slices* for obtaining rules that are closed under substitution. Technically, a slice $\varphi[a\#\bar{u}]$ of a formula φ is a decomposition of the formula as $\varphi(a, \bar{x})[\bar{u}/\bar{x}]$ for fresh variables \bar{x} , such that a does not appear in any of the \bar{u} . Slices were also used in the FL_{Seq} rules (see Figure 1(FL_{Seq})). The slice-based rules shown in Figure 1(FL_{Seq}) are closed under substitution, so proving cut-elimination for these rules is relatively straightforward once several technical lemmas involving slices have been proved. Noting that the FL_{Seq} rules are structurally similar to $\forall L$ and $\exists R$, respectively, Gabbay and Cheney observed that alternate rules in which $\mathbb{N}L$ was similar to $\exists L$ and $\mathbb{N}R$ similar to $\forall R$ were possible (see Figure 1(FL'_{Seq})). These rules seem simpler and more deterministic; however, they still involve slices.

Gabbay and Cheney presented a proof-theoretic semantics for nominal logic programming based on FL_{Seq} . However, this analysis suggested an interpretation of \mathbb{N} -quantified formulas that was radically different from the approach used in the α Prolog nominal logic programming language [2]. The proof-search interpretation of $\mathbb{N}a.\varphi$ suggested by FL_{Seq} is “search for a slice $\varphi[a\#\bar{u}]$ of φ and substitution t for a such that $t \# \bar{u}$ and solve $\varphi(t, \bar{u})$ ”, while in α Prolog, the interpretation of $\mathbb{N}a.\varphi$ is “generate a fresh name a' and solve $\varphi(a')$ ”. The approach motivated by the FL_{Seq} proof-theoretic semantics seems much more complicated than experience with α Prolog suggests.

Gabbay and Cheney also gave a translation from $FO\lambda^\nabla$, a logic introduced by Miller and Tiu that also includes a self-dual quantifier, ∇ [9] into FL_{Seq} . This translation was sound (mapped derivable sequents to derivable sequents), but incomplete

$$\begin{array}{c}
\frac{\Gamma, a \# \bar{x} \Rightarrow \varphi, \Delta \quad (\dagger)}{\Gamma \Rightarrow \mathbb{I}a.\varphi, \Delta} \mathbb{I}R \qquad \frac{\Gamma, a \# \bar{x}, \varphi \Rightarrow \Delta \quad (\dagger)}{\Gamma, \mathbb{I}a.\varphi \Rightarrow \Delta} \mathbb{I}L \quad (NL) \\
\\
\frac{\Gamma \vdash u \# \bar{t} \quad \Gamma \vdash \varphi[u/a] \quad (*)}{\Gamma \vdash \mathbb{I}a.\varphi} \mathbb{I}I \qquad \frac{\Gamma \vdash \mathbb{I}a.\varphi \quad \Gamma \vdash u \# \bar{t}}{\Gamma, \varphi[u/a] \vdash \psi \quad (*)} \mathbb{I}E \quad (FL) \\
\\
\frac{\Gamma, u \# \bar{t} \Rightarrow \varphi[u/a] \quad (*)}{\Gamma, u \# \bar{t} \Rightarrow \mathbb{I}a.\varphi} \mathbb{I}R \qquad \frac{\Gamma, u \# \bar{t}, \varphi[u/a] \Rightarrow \psi \quad (*)}{\Gamma, u \# \bar{t}, \mathbb{I}a.\varphi \Rightarrow \psi} \mathbb{I}L \quad (FL_{Seq}) \\
\\
\frac{\Gamma, a \# \bar{t} \Rightarrow \varphi \quad (*), (**)}{\Gamma \Rightarrow \mathbb{I}a.\varphi} \mathbb{I}R \qquad \frac{\Gamma, a \# \bar{t}, \varphi \Rightarrow \psi \quad (*), (**)}{\Gamma, \mathbb{I}a.\varphi \Rightarrow \psi} \mathbb{I}L \quad (FL'_{Seq}) \\
\\
\frac{\Sigma \# a : \Gamma \Rightarrow \varphi \quad (a \notin \Sigma)}{\Sigma : \Gamma \Rightarrow \mathbb{I}a.\varphi} \mathbb{I}R \qquad \frac{\Sigma \# a : \Gamma, \varphi \Rightarrow \psi \quad (a \notin \Sigma)}{\Sigma : \Gamma, \mathbb{I}a.\varphi \Rightarrow \psi} \mathbb{I}L \quad (NL^{\Rightarrow}) \\
\\
(\dagger) \bar{x} = FV(\Gamma, \mathbb{I}a.\varphi, \Delta) \quad (*) \varphi = \varphi[a \# \bar{t}] \quad (**) a \notin FV(\Gamma, \psi)
\end{array}$$

Fig. 1. Evolution of rules for \mathbb{I}

(mapped some non-derivable sequents to derivable ones). Gabbay and Cheney conjectured that their translation would be complete relative to $FO\lambda^{\nabla}$ extended with weakening and exchange for ∇ .

In this paper we present a simplified sequent calculus for nominal logic, called NL^{\Rightarrow} , in which slices are not needed in the rules for \mathbb{I} (or anywhere else), and which seems more compatible with the proof-search reading of \mathbb{I} in α Prolog. Following Urban, Pitts, and Gabbay [14, 6], we employ a new syntactic class of *name-symbols* a, b, \dots . Like variables, such name-symbols may be bound (by \mathbb{I}), but unlike variables, two distinct name-symbols are always regarded as denoting distinct name values. In place of slices, we introduce variable contexts that encode information about freshness. Specifically, contexts $\Sigma \# a : \nu$ may be formed by adjoining a *fresh name-symbol* a which is also assumed to be semantically fresh for any value mentioned in Σ . Our rules for \mathbb{I} (Figure 1(NL^{\Rightarrow})) are in the spirit of the original rules and are very simple.

Besides the sequent calculus itself, we present two applications. First, we verify that NL^{\Rightarrow} and Pitts' axiomatization NL are equivalent. Second, we present and prove the soundness and completeness of a new translation from $FO\lambda^{\nabla}$ to nominal logic, solving a problem left unsolved by Gabbay and Cheney. We have also found that the original translation is complete relative to $FO\lambda^{\nabla}$ extended with ∇ -weakening and contraction.

The structure of this paper is as follows: Section 2 presents the sequent calculus NL^{\Rightarrow} along with proofs of structural properties and conservativity of NL^{\Rightarrow} relative to NL . In Section 3, we present sound and complete translations from $FO\lambda^{\nabla}$ (with and without ∇ -weakening and exchange) to NL^{\Rightarrow} . Section 4 discusses additional related and future work, and Section 5 concludes.

2 Sequent Calculus

The sequent calculus in this section is a generalization of the one presented in Chapter 4 of the author's dissertation [5]. Full proofs can be found there and in a companion technical report [3].

2.1 Syntax and Well-Formedness

The types τ , terms t , and formulas φ of NL^{\Rightarrow} are generated by the following grammar:

$$\begin{aligned} \tau &::= o \mid \delta \mid \nu \mid \tau \rightarrow \tau' \mid \langle \nu \rangle \tau & t, u &::= c \mid a \mid \lambda x:\tau.t \mid t u \mid x \\ \varphi, \psi &::= \top \mid \perp \mid t \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \supset \psi \mid \forall x:\tau.\varphi \mid \exists x:\tau.\varphi \mid \mathbf{I}a:\nu.\varphi \end{aligned}$$

The base types are datatypes δ , name-types ν , and the type o of propositions; additional types are formed using the function and abstraction type constructors. Variables x, y are drawn from a countably infinite set V ; also, name-symbols a, b are drawn from a disjoint countably infinite set A . The letters a, b are typically used for terms of some name-sort ν . Note that λ -terms are included in this language and are handled in a traditional fashion. In particular, terms are considered equal up to $\alpha\beta\eta$ -equivalence. Similarly, \forall, \exists , and \mathbf{I} -quantified formulas are identified up to α -equivalence. We assume given a signature that maps constant symbols c to types τ , and containing at least the following declarations:

$$eq_{\tau} : \tau \rightarrow \tau \rightarrow o \quad fresh_{\nu\tau} : \nu \rightarrow \tau \rightarrow o \quad swap_{\nu\tau} : \nu \rightarrow \nu \rightarrow \tau \rightarrow \tau \quad abs_{\nu\tau} : \nu \rightarrow \tau \rightarrow \langle \nu \rangle \tau$$

for all name-types ν and types τ . The notations $t \approx u$, $t \# u$, $(t u) \cdot v$, and $\langle t \rangle u$ are syntactic sugar for $eq t u$, $fresh t u$, $swap t u v$, and $abs t u$, respectively.

The *contexts* used in NL^{\Rightarrow} are generated by the grammar:

$$\Sigma ::= \cdot \mid \Sigma, x:\tau \mid \Sigma \# a:\nu$$

We often abbreviate $\cdot, x:\tau$ and $\cdot \# a:\nu$ to $x:\tau$ and $a:\nu$ respectively, and may omit type declarations when no ambiguity ensues. We write ω for a term that may be either a name-symbol a or a variable x . The functions $FV(-)$, $FN(-)$, $FVN(-)$ calculate the sets of free variables, name-symbols, or both variables and name-symbols of a term or formula. Note that abstraction $\langle - \rangle$ is just a function symbol and does not bind its first argument (which may be any term of type ν), and so $FN(\langle a \rangle t) = FN(a) \cup FN(t)$, whereas $\mathbf{I}a.\varphi$ does bind a , so $FN(\mathbf{I}a.\varphi) = FN(\varphi) - \{a\}$. We write $\omega:\tau \in \Sigma$ if the binding $\omega:\tau$ is present in Σ . We write $\Sigma; \Sigma'$ for the result of concatenating two contexts such that $FVN(\Sigma) \cap FVN(\Sigma') = \emptyset$.

Remark 1. The inclusion of λ -terms and identification of terms and formulas with bound names up to α -equivalence may be objectionable because it appears that we are circularly attempting to define binding in terms of binding. This is not the case. A key contribution of Gabbay and Pitts' approach is that it shows how one can formally justify a traditional, informal approach to binding syntax by constructing syntax trees modulo α -equivalence as simple mathematical objects in a particularly clever way [8][5, Ch. 3–4]. We assume that this or some other standard technique for dealing with binding in nominal logic's terms and formulas is acting behind the scenes.

We write $\Sigma \vdash t : \tau$ or $\Sigma \vdash \varphi : o$ to indicate that t is a well-formed term of type τ or φ is a well-formed formula. From the point of view of typechecking, the freshness information given by the context is irrelevant. There are only two nonstandard rules for typechecking:

$$\frac{\omega : \tau \in \Sigma}{\Sigma \vdash \omega : \tau} \quad \frac{\Sigma \# a : \nu \vdash \varphi : o}{\Sigma \vdash \forall a : \nu. \varphi : o}$$

Terms viewed as formulas must, as usual, be of type o . Quantification using \forall and \exists is only allowed over types not mentioning o ; \forall -quantification is only allowed over name-types.

Let $Tm_\Sigma = \{t \mid \Sigma \vdash t : \tau\}$ be the set of well-formed terms in context Σ . We associate a set of freshness formulas $|\Sigma|$ to each context Σ as follows:

$$|\cdot| = \emptyset \quad |\Sigma, x : \tau| = |\Sigma| \quad |\Sigma \# a : \nu| = |\Sigma| \cup \{a \# t \mid t \in Tm_\Sigma\}$$

For example, $a \# x, b \# a$, and $b \# f x y \in |x : \tau \# a : \nu, y : \tau' \# b : \nu'|$ (provided $f : \tau \rightarrow \tau' \rightarrow \sigma$ is a function symbol). We say that Σ is stronger than Σ' ($\Sigma' \leq \Sigma$) if $Tm_{\Sigma'} \subseteq Tm_\Sigma$ and $|\Sigma'| \subseteq |\Sigma|$. For example, $a, x \leq x \# a, y$.

Lemma 1 (Term Weakening). *If $\Sigma \vdash t : \tau$ and $\Sigma \leq \Sigma'$ then $\Sigma' \vdash t : \tau$.*

Lemma 2 (Term Substitution). *If $\Sigma \vdash t : \tau$ and $\Sigma, x : \tau; \Sigma' \vdash u : \tau'$ then $\Sigma; \Sigma' \vdash u[t/x] : \tau'$.*

2.2 The Rules

Judgments are of the form $\Sigma : \Gamma \Rightarrow \Delta$, where Σ is a context and Γ, Δ are multisets of formulas. We define classical and intuitionistic versions of NL^\Rightarrow . *Classical* NL^\Rightarrow is based on the classical sequent calculus **G3c** [11] (see Figure 2), whereas *Intuitionistic* NL^\Rightarrow (INL^\Rightarrow) is based on the intuitionistic calculus **G3im** (in which $\supset, \forall R$, and $\exists L$ -rules are restricted to a single-conclusion form). Both versions include two additional *logical rules*, $\forall L$ and $\forall R$, shown in Figure 1(NL^\Rightarrow). In addition, NL^\Rightarrow includes several *nonlogical rules* (Figure 4) defining the properties of swapping, equality, freshness and abstraction. Figure 5 lists some admissible rules.

Many of the nonlogical rules correspond to first-order universal axioms of nominal logic (Figure 3), which may be incorporated into sequent systems in a uniform fashion using the Ax rule without affecting cut-elimination [11]. The remaining nonlogical rules are as follows. Rule A_2 expresses an invertibility property for abstractions: two abstractions are equal only if they are structurally equal or equal by virtue of A_1 . A_3 says that all values of abstraction type are formed using the abstraction function symbol. The F rule expresses the freshness principle: that a name fresh for a given context may always be chosen. Finally, the $\Sigma \#$ rule allows freshness information to be extracted from the context Σ . It states that in context Σ , any constraint in $|\Sigma|$ is valid.

2.3 Structural Properties

We now list some routinely-verified properties of NL^\Rightarrow derivations. We write $\vdash_n J$ to indicate that judgment J has a derivation of height at most n .

$$\begin{array}{c}
\frac{}{\Sigma : \Gamma, p \bar{t} \Rightarrow p \bar{t}, \Delta} \text{hyp} \\
\frac{}{\Sigma : \Gamma \Rightarrow \top, \Delta} \top R \\
\frac{\Sigma : \Gamma \Rightarrow \varphi, \Delta \quad \Sigma : \Gamma \Rightarrow \psi, \Delta}{\Sigma : \Gamma \Rightarrow \varphi \wedge \psi, \Delta} \wedge R \\
\frac{\Sigma : \Gamma \Rightarrow \varphi_1, \varphi_2, \Delta}{\Sigma : \Gamma \Rightarrow \varphi_1 \vee \varphi_2, \Delta} \vee R \\
\frac{\Sigma : \Gamma, \varphi \Rightarrow \psi, \Delta}{\Sigma : \Gamma \Rightarrow \varphi \supset \psi, \Delta} \supset R \\
\frac{\Sigma, x : \Gamma \Rightarrow \varphi, \Delta \quad (x \notin \Sigma)}{\Sigma : \Gamma \Rightarrow \forall x. \varphi, \Delta} \forall R \\
\frac{\Sigma \vdash t : \sigma \quad \Sigma : \Gamma \Rightarrow \exists x : \tau. \varphi, \varphi[t/x], \Delta}{\Sigma : \Gamma \Rightarrow \exists x : \tau. \varphi, \Delta} \exists R \\
\frac{\Sigma : \Gamma, t \approx t \Rightarrow \Delta}{\Sigma : \Gamma \Rightarrow \Delta} \approx R \\
\frac{}{\Sigma : \Gamma, \perp \Rightarrow \Delta} \perp L \\
\frac{\Sigma : \Gamma, \varphi_1, \varphi_2 \Rightarrow \Delta}{\Sigma : \Gamma, \varphi_1 \wedge \varphi_2 \Rightarrow \Delta} \wedge L \\
\frac{\Sigma : \Gamma, \varphi \Rightarrow \Delta \quad \Gamma, \psi \Rightarrow \Delta}{\Sigma : \Gamma, \varphi \vee \psi \Rightarrow \Delta} \vee L \\
\frac{\Sigma : \Gamma \Rightarrow \varphi, \Delta \quad \Sigma : \Gamma, \psi \Rightarrow \Delta}{\Sigma : \Gamma, \varphi \supset \psi \Rightarrow \Delta} \supset L \\
\frac{\Sigma \vdash t : \sigma \quad \Sigma : \Gamma, \forall x : \tau. \varphi, \varphi[t/x] \Rightarrow \Delta}{\Sigma : \Gamma, \forall x : \tau. \varphi \Rightarrow \Delta} \forall L \\
\frac{\Sigma, x : \Gamma, \varphi \Rightarrow \Delta \quad (x \notin \Sigma)}{\Sigma : \Gamma, \exists x. \varphi \Rightarrow \Delta} \exists L \\
\frac{\Sigma : \Gamma, t \approx u, P(t), P(u) \Rightarrow \Delta}{\Sigma : \Gamma, t \approx u, P(t) \Rightarrow \Delta} \approx S
\end{array}$$

Fig. 2. Classical typed first-order equational logic (**G3c**)

Lemma 3 (Weakening). If $\vdash_n \Sigma : \Gamma \Rightarrow \Delta$ is derivable then so is $\vdash_n \Sigma : \Gamma, \varphi \Rightarrow \Delta$.

Lemma 4 (Context Weakening). If $\vdash_n \Sigma : \Gamma \Rightarrow \Delta$ and $\Sigma \leq \Sigma'$ then $\vdash_n \Sigma' : \Gamma \Rightarrow \Delta$

Lemma 5 (Substitution). If $\vdash_n \Sigma \vdash t : \tau$ and $\Sigma, x : \tau; \Sigma' : \Gamma \Rightarrow \Delta$ then $\vdash_n \Sigma; \Sigma' : \Gamma[t/x] \Rightarrow \Delta[t/x]$.

The remaining structural transformations do not preserve the height of derivations. However, they do preserve the logical height of the derivation, which is defined as follows.

Definition 1. The logical height of a derivation is the maximum number of logical rules in any branch of the derivation. We write $\vdash_n^l J$ to indicate that J has a derivation of logical height $\leq n$.

Lemma 6 (Admissibility of EVL, EVR). If $\vdash_n^l \Sigma : \Gamma, (a \ b) \cdot \varphi \Rightarrow \Delta$, then so is $\vdash_n^l \Sigma : \Gamma, \varphi \Rightarrow \Delta$. Similarly, if $\vdash_n^l \Sigma : \Gamma \Rightarrow (a \ b) \cdot \varphi, \Delta$ is derivable, then so is $\vdash_n^l \Sigma : \Gamma \Rightarrow \varphi, \Delta$.

Lemma 7 (Admissibility of hyp^*). The judgment $\Sigma : \Gamma, \varphi \Rightarrow \varphi, \Delta$ is derivable for any φ .

Proof (Sketch). Induction on the construction of φ . The only new case is for $\varphi = \mathcal{V}a. \psi(a, \bar{x})$. By induction we know that $\Sigma \# a \# \mathbf{b} : \Gamma, \psi(\mathbf{b}, \bar{x}) \Rightarrow \psi(\mathbf{b}, \bar{x})$. Using equivariance we have $\Sigma \# a \# \mathbf{b} : \Gamma, (a \ \mathbf{b}) \cdot \psi(a, \bar{x}) \Rightarrow \psi(\mathbf{b}, \bar{x})$. Since $\bar{x} \subset FV(\Sigma)$, using $\Sigma \#$ we know that $a \ \# \ \bar{x}, \mathbf{b} \ \# \ \bar{x}$, hence $(a \ \mathbf{b}) \cdot \bar{x} \approx \bar{x}$, so using equational reasoning we have $(a \ \mathbf{b}) \cdot \psi(a, \bar{x}) \approx \psi(\mathbf{b}, \bar{x})$. Then using $\mathcal{V}L$ and $\mathcal{V}R$ we can conclude $\Sigma : \Gamma, \mathcal{V}a. \psi \Rightarrow \mathcal{V}a. \psi, \Delta$. \square

$(S_1) \quad (a a) \cdot x \approx x$	$(E_4) \quad (a b) \cdot \lambda x. e[x] \approx \lambda x. (a b) \cdot e[(a b) \cdot x]$
$(S_2) \quad (a b) \cdot (a b) \cdot x \approx x$	$(F_1) \quad a \# x \wedge b \# x \supset (a b) \cdot x \approx x$
$(S_3) \quad (a b) \cdot a \approx b$	$(F_2) \quad a \# b \quad (a : \nu, b : \nu', \nu \neq \nu')$
$(E_1) \quad (a b) \cdot c \approx c$	$(F_3) \quad a \# a \supset \perp$
$(E_2) \quad (a b) \cdot (t u) \approx ((a b) \cdot t) ((a b) \cdot u)$	$(F_4) \quad a \# b \vee a \approx b$
$(E_3) \quad p(\bar{x}) \supset p((a b) \cdot \bar{x})$	$(A_1) \quad a \# y \wedge x \approx (a b) \cdot y \supset \langle a \rangle x \approx \langle b \rangle y$

Fig. 3. Equational and freshness axioms

$$\begin{array}{c}
\frac{\Sigma : \Gamma, \bar{P}, Q_1 \Rightarrow \Delta \quad \dots \quad \Sigma : \Gamma, \bar{P}, Q_n \Rightarrow \Delta}{\Sigma : \Gamma, \bar{P} \Rightarrow \Delta} Ax \quad \wedge \bar{P} \supset \bigvee \bar{Q} \text{ an axiom instance} \\
\frac{\Sigma : \Gamma, \langle a \rangle t \approx \langle b \rangle u, a \approx b, t \approx u \Rightarrow \Delta \quad \Sigma : \Gamma, \langle a \rangle t \approx \langle b \rangle u, a \# u, t = (a b) \cdot u \Rightarrow \Delta}{\Sigma : \Gamma, \langle a \rangle t \approx \langle b \rangle u \Rightarrow \Delta} A_2 \\
\frac{\Sigma \vdash t : \langle \nu \rangle \sigma \quad \Sigma, a : \nu, x : \sigma : \Gamma, t \approx \langle a \rangle x \Rightarrow \Delta \quad (a, x \notin \Sigma)}{\Sigma \vdash t : \langle \nu \rangle \sigma} A_3 \\
\frac{\Sigma \# a : \Gamma \Rightarrow \Delta \quad (a \notin \Sigma)}{\Sigma : \Gamma \Rightarrow \Delta} F \quad \frac{\Sigma : \Gamma \Rightarrow \Delta \quad \Sigma : \Gamma, t \# u \Rightarrow \Delta \quad (t \# u \in |\Sigma|)}{\Sigma : \Gamma \Rightarrow \Delta} \Sigma \#
\end{array}$$

Fig. 4. Nonlogical rules

$$\begin{array}{c}
\frac{\Sigma : \Gamma \Rightarrow \Delta}{\Sigma : \Gamma, \varphi \Rightarrow \Delta} W \quad \frac{}{\Sigma : \Gamma, \varphi \Rightarrow \varphi, \Delta} hyp^* \quad \frac{\Sigma : \Gamma \Rightarrow \varphi, \Delta \quad \Sigma : \Gamma', \varphi \Rightarrow \Delta'}{\Sigma : \Gamma, \Gamma' \Rightarrow \Delta, \Delta'} cut \\
\frac{\Sigma : \Gamma, \varphi, \varphi \Rightarrow \Delta}{\Sigma : \Gamma, \varphi \Rightarrow \Delta} C \quad \frac{\Sigma : \Gamma, (a b) \cdot \varphi \Rightarrow \Delta}{\Sigma : \Gamma, \varphi \Rightarrow \Delta} EVL \quad \frac{\Sigma : \Gamma \Rightarrow (a b) \cdot \varphi, \Delta}{\Sigma : \Gamma \Rightarrow \Delta, \varphi} EVR
\end{array}$$

Fig. 5. Some admissible rules of NL^{\Rightarrow}

Lemma 8 (Inversion). *The $\supset L$, $\exists L$, $\wedge L$, and $\vee L$ rules are invertible, in the sense of lemma 2.3.5 and 4.2.8 of Negri and von Plato [11]. In addition, $\forall L$ is invertible: if $\vdash_n^l \Sigma : \Gamma, \forall a. \varphi \Rightarrow \Delta$ is derivable then so is $\vdash_n^l \Sigma \# a : \Gamma, \varphi \Rightarrow \Delta$ for fresh a .*

Lemma 9 (Contraction). *If $\vdash_n^l \Sigma : \Gamma, \varphi, \varphi \Rightarrow \Delta$ is derivable then so is $\vdash_n^l \Sigma : \Gamma, \varphi \Rightarrow \Delta$.*

2.4 Cut-Elimination

Lemma 10 (Admissibility of Cut). *If $\Sigma : \Gamma \Rightarrow \Delta, \varphi$ and $\Sigma : \Gamma', \varphi \Rightarrow \Delta'$ have cut-free derivations then so does $\Sigma : \Gamma, \Gamma' \Rightarrow \Delta, \Delta'$.*

Proof (Sketch). We show the most interesting case, that for principal cuts on \forall -quantified formulas. In this case, the derivations are of the form

$$\frac{\Sigma \# a : \Gamma \Rightarrow \varphi, \Delta}{\Sigma : \Gamma \Rightarrow \forall a. \varphi, \Delta} \forall R \quad \frac{\Pi \quad \Sigma \# a : \Gamma', \varphi \Rightarrow \Delta'}{\Sigma : \Gamma', \forall a. \varphi \Rightarrow \Delta'} \forall L$$

where without loss of generality we assume that the same fresh name $a \notin \Sigma$ was used in both sub-derivations. Since φ is smaller than $\mathcal{V}a.\varphi$, we can obtain a derivation Π'' of $\Sigma\#a : \Gamma, \Gamma' \Rightarrow \Delta, \Delta'$ from Π and Π' by the induction hypothesis. Then

$$\frac{\Pi''}{\Sigma\#a : \Gamma, \Gamma' \Rightarrow \Delta, \Delta'} F$$

follows using rule F . □

Theorem 1 (Cut-elimination). *If $\Sigma : \Gamma \Rightarrow \Delta$ has any derivation then it has a cut-free derivation.*

Corollary 1 (Consistency). *There is no derivation of $\Sigma : \cdot \Rightarrow \perp$.*

Corollary 2 (Orthogonality). *Suppose $\Sigma : \Gamma \Rightarrow \Delta$ and Γ, Δ have no subterms of the form $\langle a \rangle t$ (respectively, $\lambda x.t$). Then there is a derivation of $\Sigma : \Gamma \Rightarrow \Delta$ that does not use any nonlogical rules involving abstraction (respectively, λ).*

2.5 Conservativity

In this section, we show that NL^{\Rightarrow} is conservative relative to Pitts' original axiomatization NL [12]. That is, every theorem of NL is provable in NL^{\Rightarrow} , and no new theorems become provable. For convenience, we assume that the same underlying first-order sequent calculus is used for NL and NL^{\Rightarrow} .

Write $\vdash_{NL} \Sigma : \Gamma \Rightarrow \Delta$ if there is a first-order equational sequent proof of $\Sigma : \Gamma, \Gamma' \Rightarrow \Delta$ for some set of NL axioms Γ' . Write $\vdash_{NL^{\Rightarrow}} \Sigma : \Gamma \Rightarrow \Delta$ if $\Sigma : \Gamma \Rightarrow \Delta$ is derivable in NL^{\Rightarrow} without using any rules involving λ . Write \vdash_{IX} for the intuitionistic version of provability in system X , that is, provability using only single-conclusion sequents.

We translate NL formulas φ to NL^{\Rightarrow} formulas φ^* by replacing all subformulas of the form $\mathcal{V}a.\varphi(a)$ with $\mathcal{V}a.\varphi^*(a)$, for fresh name-symbols a . This translation is uniquely defined up to α -equivalence. For example, $(\mathcal{V}a.\mathcal{V}b.p(a, b))^* = \mathcal{V}a.\mathcal{V}b.p(a, b)$.

To prove the reverse direction of conservativity, it is necessary to show that NL^{\Rightarrow} sequents involving fresh name-symbols and contexts $\Sigma\#a$ are equivalent to sequents involving only variables.

Lemma 11 (Name-Elimination). *Suppose Σ mentions only variables and $\vdash_n^l \Sigma\#a : \Gamma[a] \Rightarrow \Delta[a]$. Then $\vdash_n^l \Sigma, a : \Gamma[a], a \# \Sigma \Rightarrow \Delta[a]$, where $a \# \Sigma$ is an abbreviation for $\{a \# x \mid x \in \Sigma\}$.*

Theorem 2 (Conservativity). $\vdash_{(I)NL} \Sigma : \Gamma \Rightarrow \Delta$ if and only if $\vdash_{(I)NL^{\Rightarrow}} \Sigma : \Gamma^* \Rightarrow \Delta^*$

Remark 2 (Semantics). Conservativity justifies NL^{\Rightarrow} 's description as a sequent calculus for nominal logic. Although this paper focuses exclusively on proof theory at the expense of more traditional model theoretic semantics, conservativity guarantees that NL^{\Rightarrow} inherits Pitts' nominal set semantics for nominal logic (as well as suffering from the same completeness problem). Space constraints preclude further discussion; however, these issues are considered in detail in Cheney's dissertation and a paper in preparation.

3 A Sound and Complete Translation of $FO\lambda^\nabla$

Miller and Tiu introduced a sequent calculus called $FO\lambda^\nabla$, which abbreviates “First-order Logic with λ -terms and the ∇ -quantifier” [9]. Like \mathcal{N} , the ∇ quantifier is self-dual. However, \mathcal{N} and ∇ have distinctly different properties. Nominal logic and $FO\lambda^\nabla$ have similar aims (reasoning about languages in which binding and fresh name-generation play an important role), so it is of interest to determine the relationship between $FO\lambda^\nabla$ and INL^\Rightarrow . Also, $FO\lambda^\nabla$ has only been studied using proof theory, but nominal logic has a well-understood semantics [12], so relating the two systems may also elucidate the semantics of $FO\lambda^\nabla$.

In $FO\lambda^\nabla$, formulas are generalized to *formulas-in-context* $\sigma \triangleright \varphi$, where σ is a list of *local parameters* (variables introduced by ∇) and φ is a formula built out of first-order connectives and quantifiers or $\nabla x.\psi$. We abbreviate “formula-in-context” to “c-formula”. Local parameter contexts are subject to α -renaming, so that $a \triangleright p(a)$ and $b \triangleright p(b)$ are considered equal c-formulas. However, c-formulas are not considered equivalent up to reordering or extension of the contexts. Thus, $a, b \triangleright p(a)$, $a \triangleright p(a)$, and $b, a \triangleright p(a)$ are all considered different c-formulas.

The sequent calculus rules dealing with ∇ are as follows:

$$\frac{\Sigma : \Gamma \Rightarrow (\sigma, x) \triangleright \varphi}{\Sigma : \Gamma \Rightarrow \sigma \triangleright \nabla x.\varphi} \nabla R \quad \frac{\Sigma : \Gamma, (\sigma, x) \triangleright \varphi \Rightarrow \mathcal{A}}{\Sigma : \Gamma, \sigma \triangleright \nabla x.\varphi \Rightarrow \mathcal{A}} \nabla L$$

where in either case x must not already appear in σ or Σ . However, x may appear in some other local context.

Most of the other sequent rules of $FO\lambda^\nabla$ are standard, except for the presence of local contexts. For example,

$$\frac{\Sigma : \Gamma, \sigma \triangleright \varphi, \sigma \triangleright \psi \Rightarrow \mathcal{A}}{\Sigma : \Gamma, \sigma \triangleright \varphi \wedge \psi \Rightarrow \mathcal{A}} \wedge L \quad \frac{\Sigma : \Gamma \Rightarrow \sigma \triangleright \varphi \quad \Sigma : \Gamma \Rightarrow \sigma \triangleright \psi}{\Sigma : \Gamma \Rightarrow \sigma \triangleright \varphi \wedge \psi} \wedge R$$

are the rules dealing with \wedge . The only exceptions are the \forall and \exists rules. In $\forall R$ and $\exists L$, the bound variable is “lifted” to show its dependence on local parameters. Dually, in $\forall L$ and $\exists R$, the term substituted for the bound variable may depend on local parameters. Here are the \forall -rules; the rules for \exists are similar.

$$\frac{\Sigma, h:\bar{\tau}_\sigma \rightarrow \tau : \Gamma \Rightarrow \sigma \triangleright A[h\bar{\sigma}/x]}{\Sigma : \Gamma \Rightarrow \sigma \triangleright \forall_\tau x.A} \forall R \quad \frac{\Sigma, \sigma \vdash t : \tau \quad \Sigma : \Gamma, \sigma \triangleright A[t/x] \Rightarrow \mathcal{C}}{\Sigma : \Gamma, \sigma \triangleright \forall_\tau x.A \Rightarrow \mathcal{C}} \forall L$$

Although ∇ and \mathcal{N} have some properties in common and seem to have similar motivations, the relation between them is not obvious. For example, INL^\Rightarrow includes name-types, and \mathcal{N} may only quantify over them; $FO\lambda^\nabla$ has no name-types, and ∇ may quantify over any simple type. In addition, \mathcal{N} admits weakening ($\varphi \iff \mathcal{N}a.\varphi$ where $a \notin FN(\varphi)$) and exchange ($\mathcal{N}a.\mathcal{N}b.\varphi \iff \mathcal{N}b.\mathcal{N}a.\varphi$), and satisfies $\forall x.\varphi(x) \supset \mathcal{N}a.\varphi(a) \supset \exists x.\varphi(x)$. None of these inferences are derivable with ∇ substituted for \mathcal{N} . On the other hand, ∇ commutes with all propositional connectives, \forall , and \exists , while \mathcal{N} only commutes with propositional connectives.

Gabbay and Cheney studied the problem of embedding $FO\lambda^\nabla$ into nominal logic. They presented a translation (which we call T_{GC}) from $FO\lambda^\nabla$ to FL_{Seq} satisfying

a soundness property: if J is derivable in $FO\lambda^\nabla$ then its translation $\llbracket J \rrbracket$ is derivable in FL_{Seq} . However, their translation did not satisfy the corresponding completeness property: some non-derivable judgments of $FO\lambda^\nabla$ were translated to derivable FL_{Seq} judgments. In particular, the translation failed to reconcile the different behavior of \mathbb{I} and ∇ with respect to weakening and exchange principles.

In the rest of this section, we present a modified translation and prove its soundness and completeness. We also sketch a proof that the original translation is complete with respect to $FO\lambda^\nabla$ with ∇ -weakening and exchange. Full proofs will be given in a companion technical report [4].

Our translation T departs from T_{GC} in two ways. First, T_{GC} translated c-formulas such as $x \triangleright \varphi \wedge \psi$ by first using \mathbb{I} -quantifiers for the local context, then translating $\varphi \wedge \psi$, and finally substituting $n(\mathbf{a})$ for x , resulting in $\mathbb{I}\mathbf{a}.\llbracket \varphi \rrbracket[n(\mathbf{a})/x] \wedge \llbracket \psi \rrbracket[n(\mathbf{a})/x]$. In this approach, the head symbol of a translated c-formula was hidden beneath a sequence of \mathbb{I} -quantifiers, which made T_{GC} difficult to analyze. Instead, our translation delays \mathbb{I} -quantification as long as possible and preserves the head symbol for most formulas: for example, the prior example translates to $\llbracket x \triangleright \varphi \rrbracket \wedge \llbracket x \triangleright \psi \rrbracket$. Any \mathbb{I} -quantification is delayed as long as possible, that is, until the base case for atomic formulas.

The second change is the translation of atomic formulas. As noted earlier, the validity of c-formulas is sensitive to both the *order* and *number* of local parameters in context. To deal with this, we relativize atomic formulas to their local contexts. This is accomplished by adding an argument to each atomic formula symbol for a list of names representing the local context. Let ν^* be a type with constructors $nil : \nu^*$ and $cons : \nu \rightarrow \nu^* \rightarrow \nu^*$, that is, a type of lists of names. We use a conventional comma-separated list notation for lists: $[a, b, c] = cons(a, cons(b, cons(c, nil)))$. The translation of an atomic c-formula $\sigma \triangleright p\bar{t}$ is $\mathbb{I}\bar{\mathbf{a}}.p^* [\bar{\mathbf{a}}] \bar{t}[n_\tau(\mathbf{a})/\sigma]$, where if $p : \bar{\tau} \rightarrow o$ then $p^* : \nu^* \rightarrow \bar{\tau} \rightarrow o$.

Otherwise, T is similar to T_{GC} . Ordinary \forall and \exists -quantified values are lifted to *equivariant* functions applied to lists of names. For example, $\sigma \triangleright \forall x:\tau'.p(x)$ was translated to $\mathbb{I}\bar{\mathbf{a}}.\forall h:\tau_1 \rightarrow \dots \tau_n \rightarrow \tau'.ev(h) \supset p(h \ n_\tau(\bar{\mathbf{a}}))$, where each \mathbf{a}_i is the name representing x_i , and $ev(x) = \forall a : \nu.a \# x$.

The new translation is shown in full in Figure 6. The function $\llbracket \cdot \rrbracket$ translates judgments, contexts, and c-formulas of $FO\lambda^\nabla$ to judgments, formula multisets, and formulas of INL^\Rightarrow respectively. Note that the context Σ is translated to a set of hypotheses $ev(x)$, one for each $x \in \Sigma$. Here are two examples of the new translation. The formula $\nabla x.p \iff p$ is translated to $\mathbb{I}\mathbf{a}.p^* [\mathbf{a}] \iff p^* []$. Likewise, we translate $\nabla x,y.p \ x \ y \iff \nabla y,x.p \ x \ y$ to $\mathbb{I}\mathbf{a}, \mathbf{b}.p^* [\mathbf{a}, \mathbf{b}] (n(\mathbf{a}))(n(\mathbf{b})) \iff \mathbb{I}\mathbf{b}, \mathbf{a}.p^* [\mathbf{b}, \mathbf{a}] (n(\mathbf{a}))(n(\mathbf{b}))$. Neither of these translated formulas is derivable in nominal logic.

Lemma 12. *If $\Sigma \vdash_{FO\lambda^\nabla} t : \tau$ then $\Sigma \vdash_{INL^\Rightarrow} t : \tau$; in addition, $\Sigma : \llbracket \Sigma \rrbracket \Rightarrow ev(t)$. Also, if $\Sigma : \Gamma \Rightarrow \mathcal{A}$ is well-formed then so is $\llbracket \Sigma : \Gamma \Rightarrow \mathcal{A} \rrbracket$.*

Proposition 1 (Soundness). *If $\Sigma : \Gamma \Rightarrow \mathcal{A}$ is derivable in $FO\lambda^\nabla$ then $\llbracket \Sigma : \Gamma \Rightarrow \mathcal{A} \rrbracket$ is derivable in INL^\Rightarrow .*

Proof. Similar to, but simpler than, the proof for T_{GC} . □

$$\begin{array}{l}
\llbracket \sigma \triangleright \top \rrbracket = \top \\
\llbracket \sigma \triangleright \perp \rrbracket = \perp \\
\llbracket \sigma \triangleright p \bar{t} \rrbracket = \mathbf{I}\bar{a}.p^* [\bar{a}] (\bar{t}[\overline{n_\tau(\mathbf{a})}/\sigma]) \\
\llbracket \sigma \triangleright \varphi \wedge \psi \rrbracket = \llbracket \sigma \triangleright \varphi \rrbracket \wedge \llbracket \sigma \triangleright \psi \rrbracket \\
\llbracket \cdot \rrbracket = \cdot
\end{array}
\qquad
\begin{array}{l}
\llbracket \sigma \triangleright \varphi \vee \psi \rrbracket = \llbracket \sigma \triangleright \varphi \rrbracket \vee \llbracket \sigma \triangleright \psi \rrbracket \\
\llbracket \sigma \triangleright \varphi \supset \psi \rrbracket = \llbracket \sigma \triangleright \varphi \rrbracket \supset \llbracket \sigma \triangleright \psi \rrbracket \\
\llbracket \sigma \triangleright \forall x:\tau.\varphi \rrbracket = \forall h:\overline{\tau_\sigma} \rightarrow \tau. ev(h) \supset \llbracket \sigma \triangleright \varphi[h\sigma/x] \rrbracket \\
\llbracket \sigma \triangleright \exists x:\tau.\varphi \rrbracket = \exists h:\overline{\tau_\sigma} \rightarrow \tau. ev(h) \wedge \llbracket \sigma \triangleright \varphi[h\sigma/x] \rrbracket \\
\llbracket \sigma \triangleright \nabla x:\tau.\varphi \rrbracket = \llbracket \sigma, x:\tau \triangleright \varphi \rrbracket \\
\llbracket \Sigma, x:\tau \rrbracket = \llbracket \Sigma \rrbracket, ev(x) \quad (ev(x) = \forall a:\nu.a \# x)
\end{array}$$

$$\llbracket \Sigma : \Gamma \Rightarrow \mathcal{A} \rrbracket = \Sigma : \llbracket \Sigma \rrbracket, \llbracket \Gamma \rrbracket \Rightarrow \llbracket \mathcal{A} \rrbracket$$

Fig. 6. Translation T from $FO\lambda^\nabla$ to INL^\Rightarrow

Theorem 3 (Completeness). *If $\llbracket \Sigma : \Gamma \Rightarrow \mathcal{A} \rrbracket$ is derivable in INL^\Rightarrow then $\Sigma : \Gamma \Rightarrow \mathcal{A}$ is derivable in $FO\lambda^\nabla$.*

Proof (Sketch). We break the proof into the following steps:

1. Identify two normal forms for INL^\Rightarrow proofs, and show that proofs of translated sequents can be normalized.
2. Show that proofs of the first normal form are proofs of initial sequents.
3. Show that proofs of the second normal form correspond to applications of $FO\lambda^\nabla$ rules.

In the analysis to follow, it simplifies matters to eliminate as many nonlogical rules as possible from derivations. By the orthogonality property, we need not consider the rules for abstraction in translated derivations, since abstractions are not used in the translation. In addition, the nonlogical rules F_3 and F_4 can also be eliminated, as we shall now show.

Lemma 13. *Suppose Σ has no name-variables. If $\Sigma \vdash a : \nu$, then for some $\mathbf{a} \in \Sigma$, $\Sigma : \cdot \Rightarrow a \approx \mathbf{a}$.*

Proposition 2. *If $\llbracket \Sigma : \Gamma \Rightarrow \mathcal{A} \rrbracket$ is derivable then it has a derivation that does not use F_3 or F_4 .*

Proof. To show that F_3 cannot be used in a derivation of a translated sequent, note that $\llbracket \Gamma \rrbracket$ and $\llbracket \mathcal{A} \rrbracket$ do not mention equality or freshness, and the formulas $\llbracket \Sigma \rrbracket = \forall a.a \# x_1, \dots, \forall a.a \# x_n$ cannot be instantiated to $x_i \# x_i$ since the variables x_i are not of name-type. We can therefore show that no sequent occurring in the derivation of a translated sequent can contain $a \# a$ using methods similar to those used for consistency and orthogonality.

Consider a subderivation ending with F_4 , of the form

$$\frac{\Sigma : \Gamma, a \# b \Rightarrow \varphi \quad \Sigma : \Gamma, a \approx b \Rightarrow \varphi}{\Sigma : \Gamma \Rightarrow \varphi}$$

Name-variables are never introduced in translated derivations, so by Lemma 13, we have $\Sigma \Rightarrow a \approx \mathbf{a}$, $\Sigma : \cdot \Rightarrow b \approx \mathbf{b}$ for some $\mathbf{a}, \mathbf{b} \in \Sigma$. If $\mathbf{a} = \mathbf{b}$ then clearly $\Sigma : \cdot \Rightarrow a \approx b$, so we can use the second subderivation and cut to derive $\Sigma : \Gamma \Rightarrow \varphi$. On the other hand, if $\mathbf{a} \neq \mathbf{b}$ then clearly $\Sigma : \cdot \Rightarrow a \# \mathbf{b}$ and also $\Sigma : \cdot \Rightarrow a \# b$. Using cut and the subderivation $\Sigma : \Gamma, a \# b$ we can derive $\Sigma : \Gamma \Rightarrow \varphi$. \square

Definition 2. A derivation is in first normal form if it uses only the rules $\forall L$, $\forall R$, hyp , and nonlogical rules.

A derivation beginning with a left- or right-rule is in second normal form provided that if the top-level rule is $\forall L$, $\forall R$, $\exists L$, or $\exists R$, then the next rule used is $\supset L$, $\supset R$, $\wedge L$, or $\wedge R$, respectively.

Before proving that translated derivations always have normal forms, we need some additional technical machinery. We write $\hat{\varphi}(t)$ for the formula $ev(t) \supset \varphi(t)$; translations of universal c-formulas are always of the form $\forall x.\hat{\varphi}(x)$. We write $\hat{\Gamma}(\bar{t})$ for a set of formulas $\hat{\varphi}_1(t_1), \dots, \hat{\varphi}_n(t_n)$ such that $\forall x.\hat{\varphi}_i(x) \in [\Gamma]$ for each i .

Lemma 14. If Σ is a $FO\lambda^\nabla$ context, $\Sigma \# \bar{a} \vdash t : \tau$ and $\Sigma \# \bar{a} : [\Sigma] \Rightarrow ev(t)$ then $\Sigma \vdash t : \tau$.

Lemma 15. If Σ is a $FO\lambda^\nabla$ context, $\Sigma \# \bar{a} : [\Sigma], [\Gamma], \hat{\Gamma}(\bar{t}) \Rightarrow ev(t)$ then either $\Sigma : [\Gamma] \Rightarrow \varphi$ has a normal derivation for any formula φ , or $\Sigma \# \bar{a} : [\Sigma] \Rightarrow ev(t)$.

Lemma 16. If $\Sigma : \Gamma \Rightarrow \varphi$ has a derivation using only nonbranching nonlogical rules, then it has either a first normal form derivation or one that starts with F or a logical rule.

Proposition 3. If $[\Sigma : \Gamma \Rightarrow \mathcal{A}]$ is derivable, then it has a normal derivation.

Proof (Sketch). First, by Corollary 2 and Proposition 2, $[\Sigma : \Gamma \Rightarrow \mathcal{A}]$ must have a derivation that does not use the rules A_1, A_2, A_3, F_3 or F_4 .

Because of subtleties involved in the interaction between the F and $\forall L$ rule, we need a stronger induction hypothesis. We prove that if $\Sigma \# \bar{a} : [\Sigma], [\Gamma], \hat{\Gamma}(\bar{t}) \Rightarrow [\mathcal{A}]$ has a derivation, then $\Sigma : [\Sigma], [\Gamma] \Rightarrow [\mathcal{A}]$ has a normal derivation.

Using Lemma 16, the sequent either has a first normal form derivation (in which case we are done) or begins with F or a logical rule. If it starts with a propositional rule applied to an element of $[\Gamma]$, then we are done. The induction steps for F and $\forall L$ are immediate. For $\exists L, \forall R$, we can use the invertibility of $\wedge L$ and $\supset R$ respectively and then use $\forall L$. This leaves the cases for $\exists R$ and for $\supset L$ applied to an element of $\hat{\Gamma}$. For $\supset L$ we must have subderivations of $\Sigma \# \bar{a} \vdash t : \tau$ and $\Sigma \# \bar{a} : [\Gamma], \hat{\Gamma}(\bar{t}), \varphi(t) \Rightarrow [\mathcal{A}]$. Using the lemmas we can show that the witnessing term t does not mention any names, and so we can construct a derivation starting with $\forall L$ and $\supset L$. In the similar case of $\exists R$, we also need the invertibility of $\wedge R$. \square

We next show that if the derivation is in first normal form, then the $FO\lambda^\nabla$ sequent is derivable. We need two auxiliary facts.

Lemma 17. Suppose $\bar{x} \# \bar{a} \vdash t : \tau$ and $\pi \cdot [\bar{a}] = [\bar{b}]$. Then $\bar{x} \# \bar{a} \# \bar{b} : \cdot \Rightarrow \pi \cdot t \approx t[\mathbf{b}_1/\mathbf{a}_1, \dots, \mathbf{b}_n/\mathbf{a}_n]$

Lemma 18. Suppose that Σ has no name-variables and Γ consists of freshness and equality formulas only. If $\Sigma : \Gamma, p \bar{t} \Rightarrow p \bar{u}$ then for some permutation π of names in Σ , we have $\Sigma : \Gamma \Rightarrow \pi \cdot \bar{t} \approx \bar{u}$.

Proof. The proof is by induction on the structure of the derivation. Only the hypothesis and nonbranching nonlogical rules can be involved, of these cases, only F poses a challenge. In the case for F , the π obtained by induction may mention the fresh name a introduced by F ; however, a cannot appear in t or u , so $b = \pi^{-1}(a)$ must not appear in t , and so $\pi' = \pi \circ (a \ b)$ also works since $\pi' \cdot t = \pi \cdot (a \ b) \cdot t = \pi \cdot t = u$. \square

Proposition 4. *Let $\llbracket \Sigma : \Gamma \Rightarrow \mathcal{A} \rrbracket$ have a first-normal form derivation. Then $\Sigma : \Gamma \Rightarrow \mathcal{A}$ is derivable.*

Proof. If $\llbracket \Sigma : \Gamma \Rightarrow \mathcal{A} \rrbracket$ has a first normal form derivation, then \mathcal{A} and some element \mathcal{B} of Γ must be of the form $\sigma \triangleright \nabla \bar{x}.p \bar{t}$. Without loss of generality, we consider the case where no ∇ -quantifiers appear. After stripping off the initial sequence of \mathcal{NL} and \mathcal{NR} rules, there must be a subderivation of

$$\Sigma \# \bar{a} \# \bar{b} : \llbracket \Sigma \rrbracket, \llbracket \Gamma \rrbracket, p^* [\bar{a}] \theta(\bar{t}) \Rightarrow p^* [\bar{b}] \theta'(\bar{u})$$

for some names \bar{a}, \bar{b} , where $\theta = [\bar{n}(\bar{a})/\sigma]$ and $\theta' = [\bar{n}(\bar{b})/\sigma']$. Note that θ and θ' are one-to-one and so invertible on their ranges, and that $\Sigma \# \bar{a} \vdash \theta(\bar{t}) : \bar{\tau}$ (that is, none of the b appear in $\theta(\bar{t})$).

By Lemma 18, there must be a ground permutation π such that $\Sigma : \cdot \Rightarrow \pi \cdot ([\bar{a}] \theta(\bar{t})) \approx [\bar{b}] \theta'(\bar{u})$. Clearly, $\pi \cdot [\bar{a}] = [\bar{b}]$, so by Lemma 17 we have $\bar{u}[\bar{n}(\bar{b})/\sigma'] = \theta'(\bar{u}) \approx \pi \cdot \theta(\bar{t}) \approx \theta(\bar{t})[b_1/a_1, \dots, b_n/a_n] = \bar{t}[\bar{n}(\bar{b})/\sigma]$. Since $[\bar{n}(\bar{b})/\sigma']$ is invertible, we have $\bar{u} \approx \bar{t}[\bar{n}(\bar{b})/\sigma][\sigma'/\bar{n}(\bar{b})] = \bar{t}[\sigma'/\sigma]$, which implies $\sigma \triangleright p \bar{t} \equiv_{\alpha} \sigma' \triangleright p \bar{u}$. \square

Proof (Completeness Theorem). In $FO\lambda^{\nabla}$, ∇ commutes with all propositional connectives, \forall , and \exists . Therefore, every judgment is equivalent to one in which ∇ -quantifiers only occur around atomic formulas, that is, in subformulas of the form $\nabla \bar{x}.p \bar{t}$. So it suffices to consider only judgments of this form.

The proof is by induction on the complexity of the judgment $\Sigma : \Gamma \Rightarrow \mathcal{A}$. If the normalized derivation is of the first form, then by Proposition 4, the sequent is derivable. If the normalized derivation is of the second form, there are many subcases, one for each possible starting left- or right-rule. The cases for propositional rules are straightforward. The remaining cases are those for \forall and \exists . We will show that translated sequents derived using $\forall L/R$, $\exists L/R$ in INL^{\Rightarrow} can be derived using $\forall L/R$ and $\exists L/R$ in $FO\lambda^{\nabla}$.

If the final step of the derivation is $\forall R$, then the derivation must be of the form

$$\frac{\frac{\Sigma, h : \llbracket \Sigma \rrbracket, ev(h), \llbracket \Gamma \rrbracket \Rightarrow \llbracket \sigma \triangleright \varphi[h\sigma/x] \rrbracket}{\Sigma, h : \llbracket \Sigma \rrbracket, \llbracket \Gamma \rrbracket \Rightarrow ev(h) \supset \llbracket \sigma \triangleright \varphi[h\sigma/x] \rrbracket} \supset R}{\Sigma : \llbracket \Sigma \rrbracket, \llbracket \Gamma \rrbracket \Rightarrow \forall h. ev(h) \supset \llbracket \sigma \triangleright \varphi[h\sigma/x] \rrbracket} \forall R$$

Note that $\llbracket \Sigma \rrbracket, ev(h) = \llbracket \Sigma, h \rrbracket$, so the topmost sequent is of the form $\llbracket \Sigma, h : \Gamma \Rightarrow \sigma \triangleright \varphi[h\sigma/x] \rrbracket$. By induction, $\Sigma, h : \Gamma \Rightarrow \sigma \triangleright \varphi[h\sigma/x]$ is derivable, and using $\forall R$, we conclude $\Sigma : \Gamma \Rightarrow \sigma \triangleright \forall x. \varphi$. The $\exists L$ case is similar.

If the final inference is $\forall L$, then the derivation must be of the form

$$\frac{\frac{\Sigma : \llbracket \Sigma \rrbracket, \llbracket \Gamma \rrbracket \Rightarrow ev(t) \quad \Sigma : \llbracket \Sigma \rrbracket, \llbracket \Gamma \rrbracket, \llbracket \sigma \triangleright \varphi[h\sigma/x] \rrbracket[t/h] \Rightarrow \llbracket \mathcal{A} \rrbracket}{\Sigma \vdash t : \bar{\tau}_{\sigma} \rightarrow \tau \quad \Sigma : \llbracket \Sigma \rrbracket, \llbracket \Gamma \rrbracket, ev(t) \supset \llbracket \sigma \triangleright \varphi[h\sigma/x] \rrbracket \Rightarrow \llbracket \mathcal{A} \rrbracket} \supset L}{\Sigma : \llbracket \Sigma \rrbracket, \llbracket \Gamma \rrbracket, \forall h. ev(h) \supset \llbracket \sigma \triangleright \varphi[h\sigma/x] \rrbracket \Rightarrow \llbracket \mathcal{A} \rrbracket} \forall L$$

Since Σ does not mention name-constants, we have $\Sigma \vdash t : \bar{\tau}_\sigma \rightarrow \tau$ and also $\Sigma, \sigma \vdash t \sigma : \tau$ in $FO\lambda^\nabla$. Note that $\llbracket \sigma \triangleright \varphi[h\sigma/x] \rrbracket [t/h] = \llbracket \sigma \triangleright \varphi[t \sigma/x] \rrbracket$ so we also have $\Sigma : \llbracket \Sigma \rrbracket, \llbracket \Gamma \rrbracket, \llbracket \sigma \triangleright \varphi[t \sigma/x] \rrbracket \Rightarrow \llbracket \mathcal{A} \rrbracket$, which is the same as $\llbracket \Sigma : \Gamma, \sigma \triangleright \varphi[t \sigma/x] \Rightarrow \mathcal{A} \rrbracket$. By induction, $\Sigma : \Gamma, \sigma \triangleright \varphi[t \sigma/x] \Rightarrow \mathcal{A}$ is derivable, and since $\Sigma, \sigma \vdash t \sigma : \tau$, we can use $\forall L$ to conclude that $\Sigma : \Gamma, \sigma \triangleright \forall x. \varphi \Rightarrow \mathcal{A}$. The $\exists R$ case is similar. \square

Remark 3. If we modify the translation step for atomic formulas by defining $\llbracket \sigma \triangleright p \bar{t} \rrbracket = \mathbb{V}\bar{a}. p \bar{t}[\bar{n}(\bar{a})/\sigma]$ then we obtain a translation T_{WX} that is essentially the same as T_{GC} , and is complete with respect to $FO\lambda^\nabla$ with ∇ -weakening and exchange principles.

We write $\theta : \sigma \hookrightarrow \sigma'$ to indicate that θ is a partial injective renaming mapping σ to σ' . We say that c-formulas are WX -equivalent ($\sigma \triangleright A \equiv_{WX} \sigma' \triangleright B$) if there is a $\theta : \sigma \hookrightarrow \sigma'$ such that $\theta(A) = B$. For example, $x, y \triangleright p(x, y) \equiv_{WX} y, x, z \triangleright p(x, y)$. Note that \equiv_{WX} subsumes α -equivalence. Let $FO\lambda_{WX}^\nabla$ be $FO\lambda^\nabla$ except that atomic c-formulas are considered equal modulo \equiv_{WX} .

It is not difficult to show that the formulas $\nabla x. \varphi \iff \varphi$ (where $x \notin FV(\varphi)$) and $\nabla x. \nabla y. \varphi \iff \nabla y. \nabla x. \varphi$ are derivable in $FO\lambda_{WX}^\nabla$ for any formula φ . In addition, using the same techniques as above, we can show that the translation is sound and complete relative to $FO\lambda_{WX}^\nabla$. The proof is the same as that for completeness relative to $FO\lambda^\nabla$, except that we need to show that Proposition 4 holds for atomic c-formulas equal modulo \equiv_{WX} instead of α -equivalence.

4 Related and Future Work

Besides previous formalizations of nominal logic by Pitts, Gabbay, and Cheney (surveyed in Section 1.1), several other logics and type systems have considered rules for \mathbb{V} -quantified formulas or types. Caires and Cardelli [1] investigated a logic incorporating proof rules for \mathbb{V} -quantified formulas based on maintaining a set of side-conditions involving freshness constraints. However, the freshness constraints are not formulas of their logic. These rules are similar in spirit to (and partly inspired) the slice-based rules of FL and FL_{Seq} . Another related system is the type system of Nanevski [10], which includes rules similar to those of FL for \mathbb{V} -quantified types. A third closely related system is Schöpp and Stark's dependent type theory for names and binding [13], in which a bunched context is used to store freshness information. Our freshness contexts and rules for \mathbb{V} are simpler special cases of the contexts and rules in their theory.

There are several directions for future work. NL^{\Rightarrow} may be useful for developing an improved proof-theoretic semantics for nominal logic programming. Natural deduction calculi or type theories for nominal logic based on our approach could be used as the basis of proof checkers and interactive theorem provers for nominal logic. The existence of translations from $FO\lambda^\nabla$ to NL^{\Rightarrow} suggest that $FO\lambda^\nabla$ can be interpreted using the semantics of nominal logic. Moreover, a semantic approach may lead to a simpler proof of the completeness of the translations.

5 Conclusions

This paper makes two contributions. First, we present a new sequent calculus for nominal logic which avoids the *slices* used in the rules for \mathbb{V} in FL and FL_{Seq} . Instead,

our calculus deals with \mathbb{N} using *freshness contexts* that encode freshness information as well as typing information. Although this is partly a matter of taste, we believe that our approach is easier to use and analyze and provides a more transparent reading of \mathbb{N} as a proof search operation than any previous system. In particular, the proofs of cut-elimination and conservativity relative to Pitts' axiomatization seem simpler and require fewer technical lemmas than previous attempts.

The second contribution of this paper is an improved translation from $FO\lambda^\nabla$ to intuitionistic nominal logic (INL^{\Rightarrow}), which explains the behavior of the ∇ -quantifier in terms of \mathbb{N} . We show that $FO\lambda^\nabla$ can be soundly and completely interpreted in INL^{\Rightarrow} , so any argument carried out in $FO\lambda^\nabla$ can also safely be carried out in INL^{\Rightarrow} . In addition, we argued that the translation originally proposed by Gabbay and Cheney is complete relative to $FO\lambda^\nabla$ with weakening and exchange for ∇ .

Acknowledgments Discussions with Ian Stark and Uli Schöpp and the anonymous reviewers' comments were of great value in improving this paper.

References

1. Luís Caires and Luca Cardelli. A spatial logic for concurrency–II. *Theoretical Computer Science*, 322(3):517–565, September 2004.
2. J. Cheney and C. Urban. Alpha-Prolog: A logic programming language with names, binding and alpha-equivalence. In *Proc. 20th Int. Conf. on Logic Programming (ICLP 2004)*, number 3132 in LNCS, pages 269–283, 2004.
3. James Cheney. A simpler proof theory for nominal logic. Technical Report EDI-INF-RR-0237, LFCS, University of Edinburgh, November 2004.
4. James Cheney. A sound and complete translation of generic judgments into nominal logic. Technical report, LFCS, University of Edinburgh, 2005. In preparation.
5. James R. Cheney. *Nominal Logic Programming*. PhD thesis, Cornell University, Ithaca, NY, August 2004.
6. M. J. Gabbay. Fresh logic: A logic of FM, 2003. Submitted.
7. M. J. Gabbay and J. Cheney. A proof theory for nominal logic. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science (LICS 2004)*, pages 139–148, Turku, Finland, 2004.
8. M. J. Gabbay and A. M. Pitts. A new approach to abstract syntax with variable binding. *Formal Aspects of Computing*, 13:341–363, 2002.
9. Dale Miller and Alwen Tiu. A proof theory for generic judgments: extended abstract. In *Proc. 18th Symp. on Logic in Computer Science (LICS 2003)*, pages 118–127. IEEE Press, 2003.
10. Aleksandar Nanevski. Meta-programming with names and necessity. In *Proc. 8th ACM SIGPLAN Int. Conf. on Functional Programming*, pages 206–217. ACM Press, 2002.
11. Sara Negri and Jan von Plato. *Structural Proof Theory*. Cambridge University Press, 2001.
12. A. M. Pitts. Nominal logic, a first order theory of names and binding. *Information and Computation*, 183:165–193, 2003.
13. Ulrich Schöpp and Ian Stark. A dependent type theory with names and binding. In *Proceedings of the 2004 Computer Science Logic Conference*, number 3210 in Lecture notes in Computer Science, pages 235–249, Karpacz, Poland, 2004.
14. C. Urban, A. M. Pitts, and M. J. Gabbay. Nominal unification. *Theoretical Computer Science*, 323(1–3):473–497, 2004.