One-and-a-halfth-order Logic

Murdoch J. Gabbay *
Dept. of Computer Science
Heriot-Watt University
Riccarton, Edinburgh EH14 4AS
UK
murdoch.gabbay@gmail.com

Aad Mathijssen
Dept. of Mathematics and Computer Science
Eindhoven University of Technology (TU/e)
P.O. Box 513, 5600 MB Eindhoven
The Netherlands
aad.mathijssen@gmail.com

Abstract

The practice of first-order logic is replete with meta-level concepts. Most notably there are the meta-variables themselves (ranging over predicates, variables, and terms), assumptions about freshness of variables with respect to these meta-variables, alpha-equivalence and capture-avoiding substitution. We present one-and-a-halfth-order logic, in which these concepts are made explicit. We exhibit both algebraic and sequent specifications of one-and-a-halfth-order logic derivability, show them equivalent, show that the derivations satisfy cut-elimination, and prove correctness of an interpretation of first-order logic within it.

We discuss the technicalities in a wider context as a case-study for nominal algebra, as a logic in its own right, as an algebraisation of logic, as an example of how other systems might be treated, and also as a theoretical foundation for future implementation.

Categories and Subject Descriptors F.4.1 [*Mathematical Logic and Formal Languages*]: Mathematical Logic—Proof theory

General Terms Theory

Keywords First-order logic, α-conversion, meta-variables, nominal terms, Fraenkel-Mostowski techniques, higher-order logic.

1. Introduction

Consider the following valid assertions about first-order predicate logic with equality (FOL) [2, 4], written in standard notation also explained later in this document:

- $\phi \supset (\psi \supset \phi)$,
- if $a \notin \mathit{fn}(\phi)$ then $\phi \supset (\phi[\![a \mapsto t]\!])$,
- if $a \notin \mathit{fn}(\phi)$ then $\phi \supset \forall a.\phi$,
- if $a, b \notin \mathit{fn}(\phi)$ then $(\forall a.\phi) \supset \forall b.\phi$,
- if $a \notin \mathit{fn}(\phi)$ then $\forall a.(\phi \supset \psi) \supset (\phi \supset \forall a.\psi)$,
- if $a \notin \mathit{fn}(\phi)$ then $\psi \supset (\phi \supset \forall a.\phi)$,
- $\forall b.\forall a.\phi \supset \forall a.\phi[\![b \mapsto a]\!]$, are derivable.

* Partially supported by EPSRC grant number EP/C013573/1.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PPDP'06 July 10–12, 2006, Venice, Italy.
Copyright © 2006 ACM 1-59593-388-3/06/0007…$5.00.

These assertions cannot be proved in FOL, since FOL derivations involve FOL syntax, while the syntax of the assertions just given contains meta-variables ϕ, ψ, a, b and t. These are *not* FOL syntax, they *vary over* FOL syntax. Also we refer to properties of syntax when we write '$a \notin \mathit{fn}(\phi)$' and '$\phi[\![a \mapsto t]\!]$', but FOL syntax cannot represent these explicitly.

Of course to us humans this is all obvious. One reason is that the derivations fall into a limited number of schema. For example the 'derivation' below on the left:

$$\dfrac{\dfrac{\dfrac{}{\psi, \phi \vdash \phi}(\mathbf{Ax})}{\phi \vdash \psi \supset \phi}(\supset\mathbf{R})}{\vdash \phi \supset (\psi \supset \phi)}(\supset\mathbf{R}) \qquad \dfrac{\dfrac{\dfrac{}{\bot, \bot \vdash \bot}(\mathbf{Ax})}{\bot \vdash \bot \supset \bot}(\supset\mathbf{R})}{\vdash \bot \supset (\bot \supset \bot)}(\supset\mathbf{R})$$

is not a derivation, but it obviously represents a schema of derivations of which the (real) derivation on the right is an instance setting ϕ and ψ to $\bot$. But is there a logic in which the beast on the left is a derivation too?

Note that it has been pointed out before that meta-variables varying over syntax are not themselves syntax, and that schematic derivations are not real derivations [20, page 7] (Hodges calls them 'argument schema'). Many authors do leave meta-variables at the meta-level. Some suggest that this is where they belong!

We feel that as mathematical computer scientists it is reasonable, nay our *duty*, to pursue formalisation whenever possible. Logic teaches us that reasoning can and should be formalised, not only its conclusions. So if we use meta-variables in reasoning, we can and should ask 'what is the mathematics of this reasoning'?

This paper presents **one-and-a-halfth-order logic**, a generalisation of first-order logic in which meta-variables and properties of syntax are made explicit.

We briefly mention the main technical barriers involved:

- $\forall a.\phi$ and $\forall b.\phi$ need not be α-convertible if ϕ mentions a and b free, so any syntactic representation which represents the meta-variable ϕ must sacrifice α-equivalence or some part of it. What is a suitable representation of meta-variables and how does it interact with binding?
- In the presence of meta-variables substitution becomes nontrivial, since the $\forall$-left intro rule (see Subsection 8.2) demands we reason about $\phi[\![a \mapsto t]\!]$ where $[\![a \mapsto t]\!]$ means 'replace a by t'. What is a suitable and correct representation of substitution, and what are its properties?
- Once these problems are solved, what derivation rules manage the extra complexity involved so that derivations remain faithful to 'first-order logic style', and cut-elimination is (fairly) easily provable?

Without further ado we give derivation rules of one-and-a-halfth-order logic in Figure 1. Also, Figure 2 includes one-and-a-halfth-order logic derivations of the assertions above (in the last three examples, we left out the use of $(\supset\mathbf{R})$ on the top-level implications).

Explanations and technical machinery follow in the rest of this paper. In Sections 2 to 4 we introduce the syntax and an equational axiomatisation of one-and-a-halfth-order logic in terms of nominal algebra. In Sections 5 to 7 we develop the sequent calculus of Figure 1 and establish properties including cut-elimination and equivalence with the axiomatisation. In Section 8 we show that a subset of one-and-a-halfth-order logic is equivalent to first-order logic. We discuss related and future work in the Conclusions.

2. Nominal algebra

We need a syntax in which $\forall a.\phi$, $a \notin fn(\phi)$, and $\phi[\![a \mapsto t]\!]$ may be explicitly represented. For this we use Nominal Terms [29], which offer built-in support for meta-variables, abstraction, and freshness. In this section, we describe nominal terms and the framework of Nominal Algebra [15], which is a theory of equational equality for nominal terms.

2.1 Sorts and terms

Fix two **base sorts** $\mathbb{P}$ of **predicates** and $\mathbb{T}$ of **terms**; we may indicate these with δ. Fix a sort of **atoms** $\mathbb{A}$. Define **sorts** τ:

$$\tau ::= \delta \mid \mathbb{A} \mid [\mathbb{A}]\tau$$

The intuition of $[\mathbb{A}]\tau$ is 'elements of τ with an atom abstracted'. This has no intuitive functional denotation, e.g. $[\tau']\tau$ is not a valid sort. ($[\mathbb{A}]\tau$ behaves more like the set of α-equivalence classes of elements of τ with a distinguished bound atom.)

Fix countably infinite disjoint collections of **atoms** a, b, c, and **unknowns** X, Y, Z.

Atoms represent object-level variable symbols, for examples see a, b in the Introduction. They will have sort $\mathbb{A}$. Unknowns represent meta-level variables, for examples see ϕ, ψ, t in the Introduction. Unknowns may have any sort and we assume that they are inherently sorted (and that there are infinitely many of each sort). We may write $X : \tau$ as shorthand for 'X, which has sort τ'. We tend to give unknowns of sort $\mathbb{P}$ names P, Q, R and unknowns of sort $\mathbb{T}$ names T, U.

A **term-former** is a syntactic token f with an associated **arity** $\rho = (\tau_1, \ldots, \tau_n)\tau$, where $n \geq 0$. We may write $\mathsf{f} : \rho$ as shorthand for 'f, which has arity ρ'. Fix the following term-formers:

$$\bot : ()\mathbb{P} \qquad \supset : (\mathbb{P}, \mathbb{P})\mathbb{P} \qquad \forall : ([\mathbb{A}]\mathbb{P})\mathbb{P} \qquad \approx : (\mathbb{T}, \mathbb{T})\mathbb{P}$$
$$\mathsf{var} : (\mathbb{A})\mathbb{T} \qquad \mathsf{sub} : ([\mathbb{A}]\tau, \mathbb{T})\tau \quad (\tau \in \{\mathbb{T}, [\mathbb{A}]\mathbb{T}, \mathbb{P}, [\mathbb{A}]\mathbb{P}\})$$

We discuss the intuitive meanings of these term-formers after we have defined terms.

We call term-formers with arities $(\mathbb{T}, \ldots, \mathbb{T})\mathbb{T}$ **object-level term-formers**; one example could be $+ : (\mathbb{T}, \mathbb{T})\mathbb{T}$. We call term-formers with arities $(\mathbb{T}, \ldots, \mathbb{T})\mathbb{P}$ **atomic predicate-formers**; one example is $\approx$, others are $\mathsf{socrates} : (\mathbb{T})\mathbb{P}$ and $\mathsf{greek} : (\mathbb{T})\mathbb{P}$. These too can be added and cause no difficulties for the results which follow.

We generally let f vary over *all* term-formers and (later in the paper) we let of vary over object-level term-formers, and op vary over atomic predicate-formers ('o' for 'object-level').

Terms t, u, v are inductively defined by:

$$t ::= a_{\mathbb{A}} \mid (\pi \cdot X_\tau)_\tau \mid ([a_{\mathbb{A}}]t_\tau)_{[\mathbb{A}]\tau} \mid (\mathsf{f}_{(\tau_1,\ldots,\tau_n)\tau}(t^1_{\tau_1}, \ldots, t^n_{\tau_n}))_\tau$$

Here subscripts indicate sorting rules. We repeat the definition without them, just for clarity:

$$t ::= a \mid \pi \cdot X \mid [a]t \mid \mathsf{f}(t_1, \ldots, t_n)$$

In case $n = 0$, we write f instead of $\mathsf{f}()$.

π is a permutation on atoms, we discuss π in the next subsection. We call $\pi \cdot X$ a **moderated unknown**; syntactically this is just a pair of a permutation and an unknown, but intuitively this represents the permutation π acting on an 'unknown term'. This intuition is made concrete in the definition of substitution on moderated unknowns, below.

We may write $t : \tau$ as shorthand for 't of sort τ'. We may call terms $\phi : \mathbb{P}$ **predicates** (not to be confused with *atomic* predicates like $\mathsf{greek}(\mathsf{var}(a))$).

We sugar the term-formers fixed earlier, and give their intuitive meaning. For any atoms a, b, terms $t, u : \mathbb{T}$ and predicates ϕ, ψ:

- $\bot$ represents falsity.
- $\phi \supset \psi$ is $\supset(\phi, \psi)$. Intuitively, this is an implication.
- $\forall[a]\phi$ is $\forall([a]\phi)$. Intuitively, this is universal quantification (which takes an *abstraction* of a formula and yields a formula). Accordingly we call the syntax-fragment $[a]_$ an **abstractor**. It has no functional semantics; $[a]t$ is intuitively merely t with a bound. The fact that $\forall$ takes an abstraction of a term and gives a term is imposed by its sort $([\mathbb{A}]\mathbb{P})\mathbb{P}$.
- $t \approx u$ is $\approx(t, u)$. Intuitively, this is equality in the object-language.
- a is $\mathsf{var}(a)$. Intuitively, this term-former connects an atom a, which has sort $\mathbb{A}$, and an object-level variable symbol $\mathsf{var}(a)$, which has sort $\mathbb{T}$.
- $v[a \mapsto t]$ is $\mathsf{sub}([a]v, t)$ for any term v of sort $\mathbb{T}$, $[\mathbb{A}]\mathbb{T}$, $\mathbb{P}$ or $[\mathbb{A}]\mathbb{P}$. Intuitively, this is substitution.

We use standard classical logic sugar:

$$\neg\phi \text{ is } \phi \supset \bot \qquad \top \text{ is } \neg\bot$$
$$\phi \wedge \psi \text{ is } \neg(\phi \supset \neg\psi) \qquad \phi \vee \psi \text{ is } \neg\phi \supset \psi$$
$$\phi \Leftrightarrow \psi \text{ is } (\phi \supset \psi) \wedge (\psi \supset \phi) \qquad \exists[a]\phi \text{ is } \neg\forall[a]\neg\phi$$

To save on (unnecessary) parentheses, take $[a]_$, $_[_ \mapsto _]$, $\approx$, $\{\neg, \forall, \exists\}$, $\{\wedge, \vee\}$, $\supset$, $\Leftrightarrow$ as the descending order of precedence. Also let $\wedge$, $\vee$, $\supset$ and $\Leftrightarrow$ associate to the right.

We write $a \in t$ (or $X \in t$) for 'a (or X) **occurs in (the syntax of)** t'. Occurrence is literal, e.g. $a \in [a]a$ and $a \in \pi \cdot X$ if $\pi(a) \neq a$. We omit inductive definitions. Similarly we may write $a \notin t$ and $X \notin t$ for 'does not occur in the syntax of t'.

Call t **closed** when t mentions no unknowns — t may still mention atoms, e.g. the term a is closed.

Write **syntactic identity** of terms t, u as $t \equiv u$. This emphasises the difference from provable equality $t = u$, which is a logical assertion, and object-level equality $t \approx u$, which is a term.

2.2 Permutations and substitutions

A **permutation** π of atoms is a total bijection $\mathbb{A} \to \mathbb{A}$ with **finite support**, meaning that for some finite set of atoms (which may be empty) $\pi(a) \neq a$, *but* for all atoms not in that set, $\pi(a) = a$. This is a mathematical notion of 'most': $\pi(a) = a$ for most a.

As usual, we write $\mathbf{Id}$ for the **identity** permutation, π^{-1} for the **inverse** of π, and $\pi \circ \pi'$ for the **composition** of π and π', i.e. $(\pi \circ \pi')(a) = \pi(\pi'(a))$. $\mathbf{Id}$ is also the identity of composition, i.e. $\mathbf{Id} \circ \pi = \pi$ and $\pi \circ \mathbf{Id} = \pi$. We may abbreviate $\mathbf{Id} \cdot X$ to X.

We write $\pi \cdot t$ for the **action** of a permutation on a term, defined inductively on syntax by:

$$\pi \cdot a \equiv \pi(a) \quad \pi \cdot (\pi' \cdot X) \equiv (\pi \circ \pi') \cdot X$$
$$\pi \cdot [a]t \equiv [\pi(a)](\pi \cdot t) \quad \pi \cdot \mathsf{f}(t_1, \ldots, t_n) \equiv \mathsf{f}(\pi \cdot t_1, \ldots, \pi \cdot t_n)$$

LEMMA 2.1. *$(\pi \circ \pi') \cdot t \equiv \pi \cdot (\pi' \cdot t)$ and $\mathbf{Id} \cdot t \equiv t$.*

$$\frac{}{\phi,\ \Phi \vdash_\Delta \Psi,\ \phi}\ (\mathbf{Ax}) \qquad \frac{}{\bot,\ \Phi \vdash_\Delta \Psi}\ (\bot\mathbf{L}) \qquad \frac{\Phi \vdash_\Delta \Psi,\ \phi \quad \psi,\ \Phi \vdash_\Delta \Psi}{\phi \supset \psi,\ \Phi \vdash_\Delta \Psi}\ (\supset\mathbf{L}) \qquad \frac{\phi,\ \Phi \vdash_\Delta \Psi,\ \psi}{\Phi \vdash_\Delta \Psi,\ \phi \supset \psi}\ (\supset\mathbf{R})$$

$$\frac{\phi[a \mapsto t],\ \Phi \vdash_\Delta \Psi}{\forall[a]\phi,\ \Phi \vdash_\Delta \Psi}\ (\forall\mathbf{L}) \qquad \frac{\Phi \vdash_\Delta \Psi,\ \psi}{\Phi \vdash_\Delta \Psi,\ \forall[a]\psi}\ (\forall\mathbf{R}) \quad (\Delta \vdash a\#\Phi, \Psi) \qquad \frac{\phi[a \mapsto t'],\ \Phi \vdash_\Delta \Psi}{t' \approx t,\ \phi[a \mapsto t],\ \Phi \vdash_\Delta \Psi}\ (\approx\mathbf{L}) \qquad \frac{}{\Phi \vdash_\Delta \Psi,\ t \approx t}\ (\approx\mathbf{R})$$

$$\frac{\phi',\ \Phi \vdash_\Delta \Psi}{\phi,\ \Phi \vdash_\Delta \Psi}\ (\mathbf{StructL}) \quad (\Delta \vdash_{\mathsf{SUB}} \phi' = \phi) \qquad \frac{\Phi \vdash_\Delta \Psi,\ \psi'}{\Phi \vdash_\Delta \Psi,\ \psi}\ (\mathbf{StructR}) \quad (\Delta \vdash_{\mathsf{SUB}} \psi' = \psi)$$

$$\frac{\Phi \vdash_{\Delta, a\#X_1, \ldots, a\#X_n} \Psi}{\Phi \vdash_\Delta \Psi}\ (\mathbf{Fr}) \quad (a \notin \Phi, \Psi, \Delta) \qquad \frac{\Phi \vdash_\Delta \Psi,\ \phi \quad \phi',\ \Phi \vdash_\Delta \Psi}{\Phi \vdash_\Delta \Psi}\ (\mathbf{Cut}) \quad (\Delta \vdash_{\mathsf{SUB}} \phi = \phi')$$

Figure 1. Sequent calculus for one-and-a-halfth-order logic

$$\dfrac{\dfrac{\dfrac{}{Q, P \vdash_\emptyset P}\ (\mathbf{Ax})}{P \vdash_\emptyset Q \supset P}\ (\supset\mathbf{R})}{\vdash_\emptyset P \supset (Q \supset P)}\ (\supset\mathbf{R}) \qquad \dfrac{\dfrac{\dfrac{}{P \vdash_{a\#P} P}\ (\mathbf{Ax})}{P \vdash_{a\#P} P[a \mapsto T]}\ (\mathbf{StructR}) \quad (a\#P \vdash_{\mathsf{SUB}} P = P[a \mapsto T])}{\vdash_{a\#P} P \supset (P[a \mapsto T])}\ (\supset\mathbf{R})$$

$$\dfrac{\dfrac{\dfrac{}{P \vdash_{a\#P} P}\ (\mathbf{Ax})}{P \vdash_{a\#P} \forall[a]P}\ (\forall\mathbf{R}) \quad (a\#P \vdash a\#P)}{\vdash_{a\#P} P \supset \forall[a]P}\ (\supset\mathbf{R}) \qquad \dfrac{\dfrac{\dfrac{}{\forall[a]P \vdash_{a\#P, b\#P} \forall[a]P}\ (\mathbf{Ax})}{\forall[a]P \vdash_{a\#P, b\#P} \forall[b]P}\ (\mathbf{StructR}) \quad (a\#P, b\#P \vdash_{\mathsf{SUB}} \forall[a]P = \forall[b]P)}{\vdash_{a\#P, b\#P} (\forall[a]P) \supset \forall[b]P}\ (\supset\mathbf{R})$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{P \vdash_{a\#P} Q, P}\ (\mathbf{Ax}) \qquad \dfrac{}{Q, P \vdash_{a\#P} Q}\ (\mathbf{Ax})}{P, P \supset Q \vdash_{a\#P} Q}\ (\supset\mathbf{L})}{P, (P \supset Q)[a \mapsto a] \vdash_{a\#P} Q}\ (\mathbf{StructL}) \quad (a\#P \vdash_{\mathsf{SUB}} P \supset Q = (P \supset Q)[a \mapsto a])}{P, \forall[a](P \supset Q) \vdash_{a\#P} Q}\ (\forall\mathbf{L})}{P, \forall[a](P \supset Q) \vdash_{a\#P} \forall[a]Q}\ (\forall\mathbf{R}) \quad (a\#P \vdash a\#P,\ \forall[a](P \supset Q))$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{P, Q \vdash_{a\#P, b\#P, Q} P}\ (\mathbf{Ax})}{P, Q \vdash_{a\#P, b\#P, Q} (a\ b) \cdot P}\ (\mathbf{StructR}) \quad (a\#P, b\#P, Q \vdash_{\mathsf{SUB}} P = (a\ b) \cdot P)}{P, Q \vdash_{a\#P, b\#P, Q} \forall[b](a\ b) \cdot P}\ (\forall\mathbf{R}) \quad (a\#P, b\#P, Q \vdash b\#P, Q)}{P, Q \vdash_{a\#P, b\#P, Q} \forall[a]P}\ (\mathbf{StructR}) \quad (a\#P, b\#P, Q \vdash_{\mathsf{SUB}} \forall[b](a\ b) \cdot P = \forall[a]P)}{P, Q \vdash_{a\#P} \forall[a]P}\ (\mathbf{Fr}) \quad (b \notin P,\ Q,\ \forall[a]P,\ a\#P)$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{P[b \mapsto c][a \mapsto c] \vdash_{c\#P} P[b \mapsto c][a \mapsto c]}\ (\mathbf{Ax})}{\forall[a](P[b \mapsto c]) \vdash_{c\#P} P[b \mapsto c][a \mapsto c]}\ (\forall\mathbf{L})}{(\forall[a]P)[b \mapsto c] \vdash_{c\#P} P[b \mapsto a][a \mapsto c]}\ (\mathbf{StructL}) \quad (c\#P \vdash_{\mathsf{SUB}} \forall[a](P[b \mapsto c]) = (\forall[a]P)[b \mapsto c])}{\forall[b]\forall[a]P \vdash_{c\#P} P[b \mapsto c][a \mapsto c]}\ (\forall\mathbf{L})}{\forall[b]\forall[a]P \vdash_{c\#P} \forall[c](P[b \mapsto c][a \mapsto c])}\ (\forall\mathbf{R}) \quad (c\#P \vdash c\#\forall[b]\forall[a]P)}{\forall[b]\forall[a]P \vdash_{c\#P} \forall[a](P[b \mapsto a])}\ (\mathbf{StructR}) \quad (c\#P \vdash_{\mathsf{SUB}} \forall[c](P[b \mapsto c][a \mapsto c]) = \forall[a](P[b \mapsto a]))}{\forall[b]\forall[a]P \vdash_\emptyset \forall[a](P[b \mapsto a])}\ (\mathbf{Fr}) \quad (c \notin \forall[b]\forall[a]P,\ \forall[a](P[b \mapsto a]))$$

Figure 2. Example derivations in one-and-a-halfth-order logic

In this section we omit proofs of lemmas; they are all quite routine. Full proofs are available elsewhere [7, 29].

A **substitution** σ is a finitely supported sort-respecting function from unknowns to terms. Here, **finitely supported** means that for some finite set of unknowns $\sigma(X) \not\equiv \mathbf{Id} \cdot X$, but for all other unknowns $\sigma(X) \equiv \mathbf{Id} \cdot X$. **Sort-respecting** means that for each X the term $\sigma(X)$ should have the same sort as X.

Write $[t_1/X_1, \ldots, t_n/X_n]$ for the substitution σ such that $\sigma(X_i) \equiv t_i$ and $\sigma(Y) \equiv \mathbf{Id} \cdot Y$, for all $Y \not\equiv X_i$, $1 \leq i \leq n$. Write $[]$ for the **empty** substitution, which maps each X to $\mathbf{Id} \cdot X$.

Write $a \in \sigma$ if there exists an X such that $a \in \sigma(X)$, and similarly write $a \notin X$ if there is no such X.

A substitution σ has a natural **action** on terms t, inductively defined by:

$$a\sigma \equiv a \quad (\pi \cdot X)\sigma \equiv \pi \cdot \sigma(X)$$
$$([a]t)\sigma \equiv [a](t\sigma) \quad \mathsf{f}(t_1, \ldots, t_n)\sigma \equiv \mathsf{f}(t_1\sigma, \ldots, t_n\sigma)$$

Give substitution and permutation actions higher precedence than abstraction and any of the sugared term-formers, and put substitution before permutation.

Note how substitution interacts with permutation in the case of an unknown, for example $((a\ b) \cdot X)[b/X] \equiv (a\ b) \cdot b \equiv a$. So π in X is 'waiting for a substitution to arrive', as also made formal in the following property:

LEMMA 2.2. $\pi \cdot t\sigma \equiv (\pi \cdot t)\sigma$.

Another permutation action will be useful. Write t^π for the **meta-level action** of π on t, which is defined by:

$$a^\pi \equiv \pi(a) \quad (\pi' \cdot X)^\pi \equiv (\pi'^\pi) \cdot X$$
$$([a]t)^\pi \equiv [\pi(a)](t^\pi) \quad \mathsf{f}(t_1, \ldots, t_n)^\pi \equiv \mathsf{f}({t_1}^\pi, \ldots, {t_n}^\pi)$$

where $\pi'^\pi = \pi \circ \pi' \circ \pi^{-1}$ (this is the **conjugation action** [9]).

LEMMA 2.3. *Fix t and π, and let σ map $X \in t$ to $\pi \cdot X$, and σ' map $X \in t$ to $\pi^{-1} \cdot X$. Then $\pi \cdot t \equiv t^\pi\sigma$ and $t^\pi \equiv (\pi \cdot t)\sigma'$.*

So the two permutation actions are interdefinable in the presence of substitution σ; however, sometimes one is more natural than the other, we shall point out how, later.

2.3 Assertions, axioms and derivations

Nominal algebra has two forms of **assertions**:

1. A **freshness (assertion)** is a pair $a\#t$ of an atom and a term. If $t \equiv X$ we call the assertion **primitive**.
2. An **equality (assertion)** is a pair $t = u$ where t and u are terms of the same sort, we discuss them below.

Write Δ for a (possibly infinite) set of *primitive* freshnesses and call it a **freshness context**. We may drop set brackets in freshnesses, e.g. writing $a\#t, b\#u$ for $\{a\#t, b\#u\}$. Also, we may write $a\#t, u$ for $a\#t, a\#u$.

Extend the notions of occurrence, closedness, permutation actions and substitution action pointwise to assertions and freshness contexts. Note that for a freshness context Δ, $\Delta\sigma$ need not be a freshness context since each unknown X in Δ is replaced by $\sigma(X)$, which is a term and need not be an unknown.

Call a pair $\Delta \to t = u$ of a finite freshness context Δ and an equality assertion $t = u$ an **axiom**. If $\Delta = \emptyset$, we may write the axiom just $t = u$.

Nominal algebra (**NA**) [15] is the logic of equality between nominal terms. We define the derivation rules of NA by the rules in Figure 3.

In this figure f ranges over term-formers,[1] t, u, v and $t_1, \ldots, t_n$ range over terms, X ranges over unknowns, π over permutations, σ over substitutions, Δ over freshness contexts, and a and b *permutatively* range over atoms, i.e. a and b represent any two *distinct* atoms. We use similar conventions henceforth. $C[_]$ is a context, it is introduced later.

Write $\Delta \vdash a\#t$ when a derivation of a freshness assertion $a\#t$ exists using the elements of Δ as assumptions, according to the rules above. Say that Δ **entails** $a\#t$ or $a\#t$ **is derivable from** Δ.

Call a (possibly infinite) set of axioms T a **theory**. We write $\Delta \vdash_{\mathsf{T}} t = u$ when a derivation of $t = u$ exists using the rules above, such that every assumption used is a freshnesses from Δ, and for every use of $(\mathbf{ax_A})$ A is an axiom of T. Say that Δ **entails** $t = u$ or $t = u$ **is derivable from** Δ.

For example, taking $A \equiv a = b$ and $B \equiv [a]X = [b]Y$ as axioms, the derivations

$$\frac{}{b = c}\,(\mathbf{ax_A}) \qquad \frac{}{[b]b = [a]a}\,(\mathbf{ax_B})$$

are valid taking $\pi = (a\ b\ c)$ and any σ, and $\pi = (a\ b)$ and $\sigma = [b/X,\ a/Y]$, respectively. Note that it is not possible to derive $a = a$ using $(\mathbf{ax_A})$.

Taking $C \equiv a\#X \to [a]X = [b]X$, of the derivations

$$\frac{\dfrac{}{a\#b}\,(\#\mathbf{ab})}{[a]b = [b]b}\,(\mathbf{ax_C}) \qquad \frac{a\#a}{[a]a = [b]a}\,(\mathbf{ax_C})$$

the left one is valid, but the right one is *not*, because $a\#a$ is not derivable.

So now we appear to have *two* derivation systems; the sequent calculus for one-and-a-halfth-order logic from Figure 1 and nominal algebra from Figure 3, both using nominal terms! We shall soon show that a particular nominal algebra theory gives rise to one-and-a-halfth-order logic.

For the rest of this subsection we discuss the derivation rules of Figure 3.

Contexts $C[_]$ is the usual notion of context as being a 'term with a hole'. Note that this notion may be represented *directly in our syntax* taking a term C with a distinguished unknown X — call it '**the hole**' — which occurs in C only as $\mathbf{Id} \cdot X$ (so not under a moderation, e.g. not as $(a\ b) \cdot X$ or somesuch). We can then define $C[t] \equiv C[t/X]$. The restriction on the moderation is to ensure that t really *does* occur in $C[t]$ and not some renamed version of it.

Note that C may contain abstractors, e.g. $C \equiv [a][b]X$, and that the substitution $[t/X]$ can capture under these abstractors, e.g. $C[a] = [a][b]a$. See [7] for a fuller treatment of this observation.

For example, $(a\ b) \cdot X$ and $\mathsf{g}(X, X)$ (X is a shorthand for $\mathbf{Id} \cdot X$) *are not* contexts. X and $\mathsf{h}(X)$ *are* contexts. Here g and h are term-formers with arities of the form $(\tau, \tau)\tau'$ and $(\tau)\tau'$, respectively.

The rule $(\mathbf{fr})$ In $(\mathbf{fr})$ square brackets denote *discharge* in the sense of natural deduction (as in implication introduction [20]); Δ denotes the other assumptions of the derivation of $t = u$.

The rule generates 'fresh atoms'. Clearly $a \notin t, u, \Delta$ manifests this intuition, but also we must account for unknowns which (intuitively) represent unknown terms. Thus to generate an atom that is *really* fresh, also for the unknown terms we use, we insist on this explicitly with freshness conditions.

The axiom rule $(\mathbf{ax_A})$ Where the axiom A is understood or irrelevant we may write just $(\mathbf{ax})$ (and $(\mathbf{ax}')$, see below).

In $(\mathbf{ax})$ atoms are *permuted* and unknowns are *instantiated*. Thus atoms stand for *any* atoms but in a way which preserves their distinctness, whereas unknowns stand for *any* term at all.

[1] More precisely, f is a *meta-variable* ranging over term-formers.

$$\frac{}{a\#b}\,(\#\mathbf{ab}) \qquad \frac{\pi^{-1}(a)\#X}{a\#\pi\cdot X}\,(\#\mathbf{X}) \qquad \frac{}{a\#[a]t}\,(\#[]\mathbf{a}) \qquad \frac{a\#t}{a\#[b]t}\,(\#[]\mathbf{b}) \qquad \frac{a\#t_1\ \cdots\ a\#t_n}{a\#\mathsf{f}(t_1,\ldots,t_n)}\,(\#\mathbf{f})$$

$$\frac{}{t=t}\,(\mathbf{refl}) \qquad \frac{t=u}{u=t}\,(\mathbf{symm}) \qquad \frac{t=u \quad u=v}{t=v}\,(\mathbf{tran}) \qquad \frac{t=u}{C[t]=C[u]}\,(\mathbf{cong}) \qquad \frac{a\#t \quad b\#t}{(a\ b)\cdot t=t}\,(\mathbf{perm})$$

$$\frac{\Delta^{\pi}\sigma}{t^{\pi}\sigma=u^{\pi}\sigma}\,(\mathbf{ax_A})\ \ A\ \equiv\ \Delta\to t=u \qquad \frac{\begin{array}{c}[a\#X_1,\ldots,a\#X_n] \qquad \Delta\\ \vdots\\ t=u\end{array}}{t=u}\,(\mathbf{fr}) \quad (a\notin t,u,\Delta)$$

Figure 3. Derivation rules of nominal algebra

Recall the discussion of $\pi\cdot t$ versus t^{π} above. Another axiom rule is possible:

$$\frac{\pi\cdot\Delta\sigma}{\pi\cdot t\sigma=\pi\cdot u\sigma}\,(\mathbf{ax}'_{\mathbf{A}}) \qquad A\ \equiv\ \Delta\to t=u$$

however in this case, atoms in the substitution σ are renamed according to permutation π, which turns out to be rather mind-bending. For example, from the axiom $[a]X=[b]X$ it is immediate that $\vdash [b]a=[a]a$ is derivable using $(\mathbf{ax})$ where we choose $\pi=(b\ a)$ and $\sigma=[a/X]$. If we use $(\mathbf{ax}')$ we must choose $\pi=(b\ a)$ and $\sigma=[b/X]$.

3. Theories

Recall that a theory is a (possibly infinite) set of axioms. Write CORE for the theory which is the empty set. Other theories of interest are listed in Figures 4 and 5.

We use a shorthand in each of those figures that the theory includes the axioms listed in the figure *and* the axioms of previous theories:

$$\mathsf{CORE}\subset\mathsf{SUB}\subset\mathsf{FOL}.$$

In the figures, f ranges over all term-formers of appropriate sort, a,b are particular, but arbitrary, distinct atoms, P,Q,R are unknowns of sort $\mathbb{P}$, and T,U are unknowns of sort $\mathbb{T}$, and $X,X_1,\ldots,X_n$ are unknowns of appropriate sorts.

The theories have the following intuitive meaning:

1. CORE is a theory of α-equivalence.
2. SUB is a theory of capture-avoiding substitution.
3. FOL **is first-order logic — with unknowns!**

We will show how and why in the rest of this paper.

Theory SUB is discussed in detail in [14]. The rest of this section discusses theory FOL.

In Figure 5 we should think of a term of the form '$\phi=\top$' as meaning intuitively 'ϕ is true'. Thus, it expresses a Hilbert-style axiom. The first block of rules are then standard axioms [21] of classical propositional logic, we call them **propositional axioms**. The second block adds quantifiers, we call them **quantifier axioms**; they exploit NA freshness conditions. The third block adds object-level equality, we call them **equational axioms**. Together, we call them the **logical axioms**.

Note that the quantifier axioms are not new! They appear in the literature [10, page 5 (2)], just like the propositional ones. What *is* new is that our axioms are not axiom-schemes but *individual axioms*; this, because we have meta-variables. Using nominal terms we can also represent such axiom schemes faithful to their 'usual' syntactic form. The next step is to show that they *also* admit a proof-theory.

There are interesting theories besides FOL (see the Conclusions).

4. Sequent-like admissible rules

For our first technical results we show that the rules like those of Figure 1, are valid in theory FOL.

LEMMA 4.1. *For all freshness contexts Δ and predicates ϕ,ψ:*

1. $\Delta\vdash_{\mathsf{FOL}}\phi\Leftrightarrow\psi\ =\ \top$ *if and only if* $\Delta\vdash_{\mathsf{FOL}}\phi=\psi$;
2. $\Delta\vdash_{\mathsf{FOL}}\phi\supset\psi\ =\ \top$ *if and only if* $\Delta\vdash_{\mathsf{FOL}}\phi=\phi\wedge\psi$.

PROOF. We only prove the first part; the second is similar. For the right-to-left implication, we reduce the consequent to $\phi\Leftrightarrow\phi\ =\top$ by congruence using $\phi=\psi$; then the result follows easily.

For the left-to-right implication, suppose $\phi\Leftrightarrow\psi\ =\top$ is derivable using Δ. Then

$$\phi\ =\ \phi\wedge\top\ =\ \phi\wedge(\phi\Leftrightarrow\psi)\ =\ \phi\wedge\psi.$$

Similarly we can derive $\psi=\phi\wedge\psi$ whence $\phi=\psi$ as required. □

LEMMA 4.2. $\vdash_{\mathsf{FOL}}\forall[a]\bot=\bot$ *is derivable.*

PROOF. It suffices to derive $\forall[a]\bot\Leftrightarrow\bot\ =\top$, by part 1 of Lemma 4.1. Or without sugar for $\Leftrightarrow$:

$$(\forall[a]\bot\supset\bot)\wedge(\bot\supset\forall[a]\bot)\ =\top.$$

By standard calculations using the propositional axioms, this follows from $\forall[a]\bot\supset\bot\ =\top$ and $\bot\supset\forall[a]\bot\ =\top$. The latter follows directly from the last propositional axiom. The former is derived as follows:

$$\dfrac{\dfrac{\dfrac{}{\bot[a\mapsto T]=\bot}\,(\mathbf{ax})}{\bot=\bot[a\mapsto T]}\,(\mathbf{symm})}{\forall[a]\bot\supset\bot=\forall[a]\bot\supset\bot[a\mapsto T]}\,(\mathbf{cong}) \qquad \dfrac{}{\forall[a]\bot\supset\bot[a\mapsto T]=\top}\,(\mathbf{ax})$$
$$\dfrac{\qquad\qquad}{\forall[a]\bot\supset\bot=\top}\,(\mathbf{tran})$$

□

Note that by this result our intuition of FOL is that the denotation of $\mathbb{T}$ is a *non-empty* set; if $\mathbb{T}$ has (intuitively) an empty set as denotation then $\forall[a]\bot\ =\top$ would be possible for some models and $\forall[a]\bot\ =\bot$ should not be derivable. This is due to an interac-

$$
\begin{array}{rr}
(\mathbf{var}\mapsto) & a[a\mapsto T] = T \\
(\#\mapsto) & a\#X \to X[a\mapsto T] = X \\
(\mathbf{f}\mapsto) & \mathsf{f}(X_1,\ldots,X_n)[a\mapsto T] = \mathsf{f}(X_1[a\mapsto T],\ldots,X_n[a\mapsto T]) \\
(\mathbf{abs}\mapsto) & b\#T \to ([b]X)[a\mapsto T] = [b](X[a\mapsto T]) \\
(\mathbf{ren}\mapsto) & b\#X \to X[a\mapsto b] = (b\ a)\cdot X
\end{array}
$$

Figure 4. Axioms of SUB

$$P \supset Q \supset P = \top \qquad \neg\neg P \supset P = \top \qquad (P \supset Q) \supset (Q \supset R) \supset (P \supset R) = \top \qquad \bot \supset P = \top \qquad \textbf{(Props)}$$

$$\forall[a]P \supset P[a\mapsto T] = \top \qquad \forall[a](P \wedge Q) \Leftrightarrow \forall[a]P \wedge \forall[a]Q = \top \qquad a\#P \to \forall[a](P \supset Q) \Leftrightarrow P \supset \forall[a]Q = \top \qquad \textbf{(Quants)}$$

$$T \approx T = \top \qquad U \approx T \wedge P[a\mapsto T] \supset P[a\mapsto U] = \top \qquad \textbf{(Eq)}$$

Figure 5. Axioms of FOL

tion of the first quantifier axiom with the rest of the theory, and is also present in most treatments of first-order logic [20].[2]

LEMMA 4.3. *For all predicates $\phi, \phi', \psi, \psi', \theta, \varepsilon$, atoms a, terms $t, t' : \mathbb{T}$, and unknowns $X_1, \ldots, X_n$:*

1. $\Delta \vdash_{\mathsf{FOL}} \phi \wedge \theta \supset \varepsilon \vee \phi = \top$
2. $\Delta \vdash_{\mathsf{FOL}} \bot \wedge \theta \supset \varepsilon = \top$
3. *if* $\Delta \vdash_{\mathsf{FOL}} \theta \supset \varepsilon \vee \phi = \top$ *and* $\Delta \vdash_{\mathsf{FOL}} \psi \wedge \theta \supset \varepsilon = \top$ *then* $\Delta \vdash_{\mathsf{FOL}} (\phi \supset \psi) \wedge \theta \supset \varepsilon = \top$
4. *if* $\Delta \vdash_{\mathsf{FOL}} \phi \wedge \theta \supset \varepsilon \vee \psi = \top$ *then* $\Delta \vdash_{\mathsf{FOL}} \theta \supset \varepsilon \vee (\phi \supset \psi) = \top$
5. *if* $\Delta \vdash_{\mathsf{FOL}} \phi[a\mapsto t] \wedge \theta \supset \varepsilon = \top$ *then* $\Delta \vdash_{\mathsf{FOL}} \forall[a]\phi \wedge \theta \supset \varepsilon = \top$
6. *if* $\Delta \vdash_{\mathsf{FOL}} \theta \supset \varepsilon \vee \psi = \top$ *and* $\Delta \vdash a\#\theta, \varepsilon$ *then* $\Delta \vdash_{\mathsf{FOL}} \theta \supset \varepsilon \vee \forall[a]\psi = \top$
7. *if* $\Delta \vdash_{\mathsf{FOL}} \phi[a\mapsto t'] \wedge \theta \supset \varepsilon = \top$ *then* $\Delta \vdash_{\mathsf{FOL}} (t' \approx t) \wedge \phi[a\mapsto t] \wedge \theta \supset \varepsilon = \top$
8. $\Delta \vdash_{\mathsf{FOL}} \theta \supset \varepsilon \vee (t \approx t) = \top$
9. *if* $\Delta \vdash_{\mathsf{FOL}} \phi' \wedge \theta \supset \varepsilon = \top$ *and* $\Delta \vdash_{\mathsf{SUB}} \phi' = \phi$ *then* $\Delta \vdash_{\mathsf{FOL}} \phi \wedge \theta \supset \varepsilon = \top$
10. *if* $\Delta \vdash_{\mathsf{FOL}} \theta \supset \varepsilon \vee \psi' = \top$ *and* $\Delta \vdash_{\mathsf{SUB}} \psi' = \psi$ *then* $\Delta \vdash_{\mathsf{FOL}} \theta \supset \varepsilon \vee \psi = \top$
11. *if* $\Delta, a\#X_1, \ldots, a\#X_n \vdash_{\mathsf{FOL}} \theta \supset \varepsilon = \top$ *and* $a \notin \theta, \varepsilon, \Delta$ *then* $\Delta \vdash_{\mathsf{FOL}} \theta \supset \varepsilon = \top$
12. *if* $\Delta \vdash_{\mathsf{FOL}} \theta \supset \varepsilon \vee \phi = \top$, $\Delta \vdash_{\mathsf{FOL}} \phi' \wedge \theta \supset \varepsilon = \top$ *and* $\Delta \vdash_{\mathsf{SUB}} \phi = \phi'$ *then* $\Delta \vdash_{\mathsf{FOL}} \theta \supset \varepsilon = \top$

PROOF. The proofs of the first four items are standard and establish the equivalence of Hilbert- and sequent-style presentations of propositional logic.

For item 5 we use item 2 of Lemma 4.1. Then we need to show that $\Delta \vdash_{\mathsf{FOL}} \forall[a]\phi \wedge \theta = \forall[a]\phi \wedge \theta \wedge \varepsilon$ follows from the assumption $\Delta \vdash_{\mathsf{FOL}} \phi[a\mapsto t] \wedge \theta = \phi[a\mapsto t] \wedge \theta \wedge \varepsilon$. Also, by this item and the first quantifier axiom, $\Delta \vdash_{\mathsf{FOL}} \forall[a]\phi = \forall[a]\phi \wedge \phi[a\mapsto t]$ holds.

We reason algebraically, implicitly using associativity of $\wedge$:

$$
\begin{array}{ll}
 & \forall[a]\phi \wedge \theta \\
= & \quad \{\ \Delta \vdash_{\mathsf{FOL}} \forall[a]\phi = \forall[a]\phi \wedge \phi[a\mapsto t]\ \} \\
 & \forall[a]\phi \wedge \phi[a\mapsto t] \wedge \theta \\
= & \quad \{\ \Delta \vdash_{\mathsf{FOL}} \phi[a\mapsto t] \wedge \theta = \phi[a\mapsto t] \wedge \theta \wedge \varepsilon\ \} \\
 & \forall[a]\phi \wedge \phi[a\mapsto t] \wedge \theta \wedge \varepsilon \\
= & \quad \{\ \Delta \vdash_{\mathsf{FOL}} \forall[a]\phi = \forall[a]\phi \wedge \phi[a\mapsto t]\ \} \\
 & \forall[a]\phi \wedge \theta \wedge \varepsilon
\end{array}
$$

The calculations for items 6 to 12 are in the same spirit, except for item 11 which is trivially deduced using (**fr**). □

5. Sequent presentation

We are now ready to directly confront Figure 1.

Let **(predicate) contexts** Φ, Ψ be finite (possibly empty) sets of predicates. A **sequent** is a triple $\Phi \vdash_{\Delta} \Psi$ where Δ is a freshness context and Φ and Ψ are predicate contexts; when a context appears to the right of $\vdash$ we may call it a **co-context**.

We may write ϕ for $\{\phi\}$, ϕ, Φ for $\{\phi\} \cup \Phi$, and Φ, Φ' for $\Phi \cup \Phi'$, and we may omit empty predicate contexts, e.g. writing $\vdash_{\Delta}$ for $\emptyset \vdash_{\Delta} \emptyset$.

Extend the notions of occurrence, closedness, permutation actions and substitution action elementwise to predicate contexts.

Define the **sequent calculus** for one-and-a-halfth-order logic to be the set of sequents inductively specified by the derivation rules in Figure 1. We may also call this set an **entailment relation** (that of one-and-a-halfth-order logic, to be precise).

Call $(\mathbf{StructL})$ and $(\mathbf{StructR})$ **structural rules**. $(\mathbf{Cut})$ can emulate them, but we would lose cut-elimination. To see why $(\mathbf{Fr})$ is useful, consider the last two examples in Figure 2.

Our rules resemble those of Gentzen's sequent calculus for classical first-order logic with equality [6, 17, 26], but with the following distinctive features:

- There is an explicit notion of *unknown predicates* given by the unknowns of sort $\mathbb{P}$, which models meta-variables in the sense that, for example, $\forall[a]P$ with $[(a \approx a)/P]$ is $\forall[a](a \approx a)$ (and not $\forall[a'](a \approx a)$).
- Freshness, α-equivalence and capture-avoiding substitution are now explicit; they are represented by *derivability* of freshnesses and equality in SUB as side-conditions.

As we used standard classical logic sugar like $\neg P$ for $P \supset \bot$, so we also 'use' the standard sequent rules for them (as sugar for 'macros' of sequent rules) without comment.

[2] In item (4) on page 48, Hodges states that "Most authors require it to have at least one member.", where 'it' denotes the domain in terms of a pure set. Also see Remark 6 on page 110 for a discussion on the implications for Hilbert-style proof calculi.

The sequent calculus is able to mimick the logical axioms from Figure 5:

LEMMA 5.1. *The following are derivable:*

1. $\vdash_\Delta \phi \supset \psi \supset \phi$. *2.* $\vdash_\Delta \neg\neg\phi \supset \phi$.

3. $\vdash_\Delta (\phi \supset \psi) \supset (\psi \supset \rho) \supset (\phi \supset \rho)$. *4.* $\vdash_\Delta \bot \supset \phi$.

5. $\vdash_\Delta \forall[a]\phi \supset \phi[a \mapsto t]$.

6. $\vdash_\Delta \forall[a](\phi \wedge \psi) \Leftrightarrow \forall[a]\phi \wedge \forall[a]\psi$.

7. If $\Delta \vdash a\#\phi$ *then* $\vdash_\Delta \forall[a](\phi \supset \psi) \Leftrightarrow \phi \supset \forall[a]\psi$.

8. $\vdash_\Delta t \approx t$. *9.* $\vdash_\Delta u \approx t \wedge \phi[a \mapsto t] \supset \phi[a \mapsto u]$.

PROOF. We consider just item 7. By $(\Leftrightarrow\mathbf{R})$ and $(\supset\mathbf{R})$ it suffices to derive

(a) $\phi, \forall[a](\phi \supset \psi) \vdash_\Delta \forall[a]\psi$. (b) $\phi \supset \forall[a]\psi \vdash_\Delta \forall[a](\phi \supset \psi)$.

We consider (a); showing (b) derivable follows similar lines:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\phi \vdash_\Delta \psi, \phi}(\mathbf{Ax}) \quad \dfrac{}{\phi, \psi \vdash_\Delta \psi}(\mathbf{Ax})}{\phi, \phi \supset \psi \vdash_\Delta \psi}(\supset\mathbf{L})}{\phi, (\phi \supset \psi)[a \mapsto a] \vdash_\Delta \psi}(\mathbf{StructL})}{\phi, \forall[a](\phi \supset \psi) \vdash_\Delta \psi}(\forall\mathbf{L})}{\phi, \forall[a](\phi \supset \psi) \vdash_\Delta \forall[a]\psi}(\forall\mathbf{R})$$

The uses of $(\mathbf{StructL})$ and $(\forall\mathbf{R})$ are valid because the following hold:

$$\Delta \vdash_{\mathsf{SUB}} (\phi \supset \psi)[a \mapsto a] = (\phi \supset \psi)$$
$$\Delta \vdash a\#\phi, \forall[a](\phi \supset \psi)$$

The corresponding calculations are interesting but have to do with NA and SUB, not FOL. See elsewhere for details [14]. □

We conclude this section with two theorems describing how derivations and their structure (for example the ones in Figure 2) interact with atoms and unknowns. In brief: atoms can be permuted, unknowns can be instantiated.

Extending notation for permutation action, we write Π^π for the result of applying π to the terms in the syntax of Π.

THEOREM 5.2. *If* Π *is a valid derivation of* $\Phi \vdash_\Delta \Psi$ *then* Π^π *is a valid derivation of* $\Phi^\pi \vdash_{\Delta^\pi} \Psi^\pi$.

Call this property **meta-level equivariance**.

PROOF. The statement

'Π is a valid derivation of $\Phi \vdash_\Delta \Psi$'

has four free variables and so by FM equivariance [16] is invariant under permuting atoms. The result follows. □

Write $\Pi(\sigma, \Delta')$ for the **substitution action** on derivations. It is inductively defined on the structure of Π:

- If Π concludes with a rule $(\mathbf{R})$ different from $(\mathbf{Fr})$, it is of the form
$$\dfrac{\Pi_1 \quad \cdots \quad \Pi_k}{\Phi \vdash_\Delta \Psi}(\mathbf{R}) \quad (cond)$$
where $k \in \{0, 1, 2\}$ and $cond$ is $\Delta \vdash_{\mathsf{SUB}} \phi = \psi$, $\Delta \vdash a\#\Phi'$ or empty.
Then $\Pi(\sigma, \Delta')$ is
$$\dfrac{\Pi_1(\sigma, \Delta') \quad \cdots \quad \Pi_k(\sigma, \Delta')}{\Phi\sigma \vdash_{\Delta'} \Psi\sigma}(\mathbf{R}) \quad (cond')$$
where $cond'$ is $\Delta' \vdash_{\mathsf{SUB}} \phi\sigma = \psi\sigma$, $\Delta' \vdash a\#\Phi'\sigma$ or empty, respectively.
- Otherwise, the derivation concludes in
$$\dfrac{\Pi'}{\Phi \vdash_\Delta \Psi}(\mathbf{Fr}) \quad (a \notin \Phi, \Psi, \Delta)$$
where Π' is a derivation of $\Phi \vdash_{\Delta, a\#X_1, \ldots, a\#X_n} \Psi$.
Then $\Pi(\sigma, \Delta')$ is
$$\dfrac{\Pi'^{(a'\ a)}(\sigma, \Delta'')}{\Phi\sigma \vdash_{\Delta'} \Psi\sigma}(\mathbf{Fr}) \quad (a' \notin \Phi\sigma, \Psi\sigma, \Delta')$$
where a' is chosen fresh (i.e. $a' \notin a$, Φ, Ψ, Δ, Δ', σ) and $\Delta'' = \Delta', a'\#Y_1, \ldots, a'\#Y_m$, in which $Y_1, \ldots, Y_m$ are all unknowns mentioned in $\sigma(X_i)$, for $1 \leq i \leq n$.

So σ is consistently applied throughout the predicate contexts occurring in Π, Δ' replaces Δ, and $(\mathbf{Fr})$ may generate slightly different freshness assumptions.

LEMMA 5.3. *For any* Δ', Δ, σ, *if* $\Delta' \vdash \Delta\sigma$ *then:*

1. if $\Delta \vdash a\#t$ *then* $\Delta' \vdash a\#t\sigma$;

2. if $\Delta \vdash_{\mathsf{T}} t = u$ *then* $\Delta' \vdash_{\mathsf{T}} t\sigma = u\sigma$.

THEOREM 5.4. *If* $\Delta' \vdash \Delta\sigma$ *and* Π *is a valid derivation of* $\Phi \vdash_\Delta \Psi$ *then* $\Pi(\sigma, \Delta')$ *is a valid derivation of* $\Phi\sigma \vdash_{\Delta'} \Psi\sigma$.

Call this property **meta-level substitution**.

PROOF. By induction on Π. We only treat the cases $(\forall\mathbf{R})$ and $(\mathbf{Fr})$. The other cases are similar to the $(\forall\mathbf{R})$ case or simpler.

1. The case of $(\forall\mathbf{R})$: Suppose $\Phi \vdash_\Delta \Psi, \forall[a]\psi$ is derived using $(\forall\mathbf{R})$. Then $\Delta \vdash a\#\Phi, \Psi$ holds and Π' is a derivation of $\Phi \vdash_\Delta \Psi, \psi$. By Lemma 5.3 $\Delta' \vdash a\#\Phi\sigma, \Psi\sigma$, and by inductive hypothesis $\Pi'(\sigma, \Delta')$ is a derivation of $\Phi\sigma \vdash_{\Delta'} \Psi\sigma, \psi\sigma$. Then we conclude that $\Phi\sigma \vdash_{\Delta'} \Psi\sigma, \forall[a]\psi\sigma$ is derivable, by extending $\Pi'(\sigma, \Delta')$ with $(\forall\mathbf{R})$, as required.
2. The case of $(\mathbf{Fr})$: Suppose $\Phi \vdash_\Delta \Psi$ is derived using $(\mathbf{Fr})$. Then Π' is a derivation of $\Phi \vdash_{\Delta, a\#X_1, \ldots, a\#X_n} \Psi$ where we assume $a \notin \Phi, \Psi, \Delta$. By meta-level equivariance (Theorem 5.2) also $\Pi'^{(a'\ a)}$ is a derivation of $\Phi \vdash_{\Delta, a'\#X_1, \ldots, a'\#X_n} \Psi$, where a' is chosen fresh (i.e. $a' \notin a, \Phi, \Psi, \Delta, \Delta', \sigma$).
By FM equivariance validity of the property

'Π' has the inductive hypothesis'

is itself invariant under permuting atoms. So

'$\Pi'^{(a'\ a)}$ has the inductive hypothesis'

is also valid.
Take $\Delta'' = \Delta', a'\#Y_1, \ldots, a'\#Y_m$, where $Y_1, \ldots, Y_m$ are all the unknowns mentioned in $\sigma(X_i)$, $1 \leq i \leq n$. It is easy to deduce $\Delta'' \vdash (\Delta, a'\#X_1, \ldots, a'\#X_n)\sigma$. By inductive hypothesis $\Pi'^{(a'\ a)}(\sigma, \Delta'')$ is a derivation of $\Phi\sigma \vdash_{\Delta''} \Psi\sigma$. Since $a' \notin \Phi\sigma, \Psi\sigma, \Delta'$ we may deduce $\Phi\sigma \vdash_{\Delta'} \Psi\sigma$ using $(\mathbf{Fr})$, as required.

□

This use of FM equivariance is a powerful and general technique which we do not expand on here. To our knowledge its first use to obtain theorems in an actual paper (as opposed to being the object of investigation itself [16, 11]) is in a paper on 'Fresh Logic' [12]. Here we shall use it repeatedly to rename atoms in the presence of unknowns.

6. Cut-elimination

This technical section establishes a cut-elimination result (Theorem 6.8) and a consistency corollary (Corollary 6.9).

We need some notation and technical lemmas. Call the **depth** of a derivation the greatest number of derivation steps not counting rules $(\mathbf{Fr})$, $(\mathbf{StructL})$ and $(\mathbf{StructR})$ between its conclusion and its leaves, over all paths. We do not count NA derivations of freshnesses and equalities that occur as side-conditions. For example, the last two derivations of Figure 2 have depth 2 and 4, respectively.

LEMMA 6.1. *If $\Delta \vdash a\#\Phi, \Psi$ and $\Delta \vdash b\#\Phi, \Psi$ then*

$\Phi, \Phi' \vdash_\Delta \Psi, \Psi'$ *if and only if* $\Phi, (a\ b)\cdot\Phi' \vdash_\Delta \Psi, (a\ b)\cdot\Psi'$.

The derivation has the same depth as the original one, and no more instances of cut.[3]

Call this property **object-level equivariance**.

PROOF. By repeated use of $(\mathbf{StructL})$ and $(\mathbf{StructR})$. □

The following results are not normally problematic but we have internalised both α-equivalence *and* being fresh — so renaming and freshening must be represented in the derivation. First, a technical lemma:

LEMMA 6.2. *If $\Delta \subseteq \Delta'$ then:*

1. *if $\Delta \vdash a\#t$ then $\Delta' \vdash a\#t$;*
2. *if $\Delta \vdash_{\mathsf{T}} t = u$ then $\Delta' \vdash_{\mathsf{T}} t = u$.*

LEMMA 6.3. *If $\Phi \vdash_\Delta \Psi$ and $\Delta \subseteq \Delta'$ then $\Phi \vdash_{\Delta'} \Psi$. The derivation has the same depth as the original one, and no more instances of cut.*

Call this property **meta-level weakening**.

PROOF. We work by induction on the structure of the derivation. The conditions on preserving depth and number of cuts can easily be verified from the structure of the reasoning which follows, and we do not mention them further.

We only treat the cases $(\forall\mathbf{R})$ and $(\mathbf{Fr})$. The other cases are trivial or similar to the $(\forall\mathbf{R})$ case.

1. The case of $(\forall\mathbf{R})$: Suppose $\Phi \vdash_\Delta \Psi, \forall[a]\psi$ is derived using $(\forall\mathbf{R})$. Then $\Delta \vdash a\#\Phi, \Psi$ and $\Phi \vdash_\Delta \Psi, \psi$ are derivable. Then $\Delta' \vdash a\#\Phi, \Psi$ by Lemma 6.2, and $\Phi \vdash_{\Delta'} \Psi, \psi$ is derivable by inductive hypothesis. Extending derivations of the latter with $(\forall\mathbf{R})$ we conclude that $\Phi \vdash_{\Delta'} \Psi, \forall[a]\psi$ is derivable, as required.
2. The case of $(\mathbf{Fr})$: Suppose $\Phi \vdash_{\Delta, a\#X_1, \ldots, a\#X_n} \Psi$ where $a \notin \Phi, \Psi, \Delta$. Choose a' fresh (i.e. $a' \notin \Phi, \Psi, \Delta'$ — note the prime on the Δ'), then $\Phi \vdash_{\Delta, a'\#X_1, \ldots, a'\#X_n} \Psi$ is derivable by meta-level equivariance .
   By FM equivariance we retain the inductive hypothesis, so it follows that $\Phi \vdash_{\Delta', a'\#X_1, \ldots, a'\#X_n} \Psi$. Then we may deduce $\Phi \vdash_{\Delta'} \Psi$ using $(\mathbf{Fr})$, as required.

□

Write $\mathcal{U}(\mathit{Stuff})$ for the unknowns mentioned in the Stuff.

LEMMA 6.4. *If $a \notin u$ and $a\#\mathcal{U}(u) \subseteq \Delta$ then $\Delta \vdash a\#u$.*

LEMMA 6.5. *If $\Phi \vdash_\Delta \Psi$ and $\Phi \subseteq \Phi'$ and $\Psi \subseteq \Psi'$ then $\Phi' \vdash_\Delta \Psi'$. The new derivation has the same depth as the original one, and no more instances of cut.*

[3] It may mention more structural rules but by an astounding coincidence we have excluded them from our notion of depth.

Call this property **object-level weakening**.

PROOF. We work by strong induction on the *pair* of the depth of the derivation and its structure, lexicographically ordered. We consider only nontrivial cases.

1. The case of $(\forall\mathbf{R})$: Suppose $\Phi \vdash_\Delta \Psi, \forall[a]\psi$ is derived using $(\forall\mathbf{R})$ and suppose the inductive hypothesis of all strictly lesser derivations.
   By assumption $\Phi \vdash_\Delta \Psi$, ψ has a derivation of strictly lesser depth, and also $\Delta \vdash a\#\Phi, \Psi$ holds.
   Choose some a'' fresh (i.e. $a'' \notin a, \psi, \Phi', \Psi', \Delta$), and take $\Delta'' = \Delta, a''\#\mathcal{U}(\Phi', \Psi', \Delta, \psi)$.
   By NA weakening (Lemma 6.2) $\Delta'' \vdash a\#\Phi, \Psi$ and by meta-level weakening (Lemma 6.3)[4] $\Phi \vdash_{\Delta''} \Psi$, ψ. Then by object-level equivariance (Lemma 6.1) also $\Phi \vdash_{\Delta''} \Psi$, $(a''\ a)\cdot\psi$, and by inductive hypothesis (the derivation still has strictly lesser depth) there exists a derivation Π of
   $$\Phi' \vdash_{\Delta''} \Psi',\ (a''\ a)\cdot\psi.$$
   By Lemma 6.4 also $\Delta'' \vdash a''\#\Phi', \Psi'$, and by simple calculations we observe $\Delta'' \vdash_{\mathsf{SUB}} \forall[a'']((a''\ a)\cdot\psi) = \forall[a]\psi$ (we use $(\mathbf{perm})$, and the freshness information we have assumed of a'').
   Now we can conclude $\Phi' \vdash_\Delta \Psi', \forall[a]\psi$ as follows:
   $$\dfrac{\dfrac{\dfrac{\Pi}{\Phi' \vdash_{\Delta''} \Psi',\ \forall[a'']((a''\ a)\cdot\psi)}\,(\forall\mathbf{R})}{\Phi' \vdash_{\Delta''} \Psi',\ \forall[a]\psi}\,(\mathbf{StructR})}{\Phi' \vdash_\Delta \Psi',\ \forall[a]\psi}\,(\mathbf{Fr})$$
2. The case of $(\mathbf{Fr})$: Suppose $\Phi \vdash_{\Delta, a\#X_1, \ldots, a\#X_n} \Psi$ where $a \notin \Phi, \Psi, \Delta$. In case $a \notin \Phi', \Psi'$ then things are easy and we use $(\mathbf{Fr})$. If however $a \in \Phi', \Psi'$ then we use FM equivariance to rename a to some $a' \notin \Phi', \Psi', \Delta$ in the whole derivation to obtain one of $\Phi \vdash_{\Delta, a'\#X_1, \ldots, a\#X_n} \Psi$. We can now apply the inductive hypothesis (which, as discussed above, is also preserved by the permutative renaming) to weaken to Φ' and Ψ', and finish off with $(\mathbf{Fr})$.

□

Write $\Phi[b \mapsto u]$ for the elementwise application of the explicit substitution to the elements of predicate context Φ.

LEMMA 6.6. *If $\Phi \vdash_\Delta \Psi$ then $\Phi[b \mapsto u] \vdash_\Delta \Psi[b \mapsto u]$. The depth of the derivation does not increase, and neither does the number of cuts it contains.*

Call this property **object-level substitution**.

PROOF. By induction on the depth and structure of the derivation of $\Phi \vdash_\Delta \Psi$, lexicographically ordered. Most cases are easy, we consider only the case of $(\forall\mathbf{R})$.

Suppose $\Phi \vdash_\Delta \Psi, \forall[a]\psi$ is derived using $(\forall\mathbf{R})$, and suppose the inductive hypothesis of all strictly lesser derivations.

By assumption $\Phi \vdash_\Delta \Psi$, ψ is derivable and $\Delta \vdash a\#\Phi, \Psi$ holds.

Choose some a' fresh (i.e. $a' \notin a, b, u, \psi, \Phi, \Psi, \Delta$) and let $\Delta' = \Delta, a'\#\mathcal{U}(u, \psi, \Phi, \Psi, \Delta)$. We use meta-level weakening, object-level equivariance and the inductive hypothesis to conclude that there exists a derivation Π of

$$\Phi[b \mapsto u] \vdash_{\Delta'} \Psi[b \mapsto u],\ ((a'\ a)\cdot\psi)[b \mapsto u],$$

[4] It appears convenient to prove meta-level weakening first separately; we do not want to weaken Φ and Ψ to Φ' and Ψ' until we have renamed a to a'', in a moment.

which has the same depth and number of cuts.

By Lemma 6.4 also $\Delta' \vdash a'\#((a'\ a)\cdot\psi)[b\mapsto u]$. Furthermore, by simple calculations we observe

$$\Delta' \vdash_{\mathsf{SUB}} \forall[a']((a'\ a)\cdot\psi)[b\mapsto u] = (\forall[a]\psi)[b\mapsto u].$$

We finish the derivation:

$$\dfrac{\dfrac{\dfrac{\Pi}{\Phi[b\mapsto u]\vdash_{\Delta'} \Psi[b\mapsto u],\ \forall[a']((a'\ a)\cdot\psi)[b\mapsto u]}(\forall\mathbf{R})}{\Phi[b\mapsto u]\vdash_{\Delta'} \Psi[b\mapsto u],\ (\forall[a]\psi)[b\mapsto u]}(\mathbf{StructR})}{\Phi[b\mapsto u]\vdash_{\Delta} \Psi[b\mapsto u],\ (\forall[a]\psi)[b\mapsto u]}(\mathbf{Fr})$$

□

LEMMA 6.7. ($\mathbf{Fr}$) *may be commuted down through all other rules. The transformations involved do not increase the depth of a derivation or its number of cuts.*

PROOF. We consider only one case. Suppose ($\mathbf{Fr}$) is followed by ($\approx\mathbf{L}$). It is not immediate that we may swap the derivation rules round, since perhaps t in ($\approx\mathbf{L}$) mentions a and t' does not (we use notation from the rules in Figure 1). However, we may rename atoms in the derivation up to the use of ($\mathbf{Fr}$) changing a to some a' which does not occur also in t, and then proceed. That the inductive hypothesis is preserved follows by FM equivariance. □

THEOREM 6.8 (Cut-elimination). *If $\Phi\vdash_{\Delta}\Psi$ has a derivation in the system above, then it has one which does not use* ($\mathbf{Cut}$).

PROOF. The commutation cases and essential cases for the propositional part are standard [17, 26], we use Lemma 6.5 for the essential case for $\supset$. The essential case for $\forall$ is handled by Lemma 6.6. Commutation cases are standard (except for the extra case of ($\mathbf{Fr}$), which is handled by Lemma 6.7). □

COROLLARY 6.9. *The sequent calculus of one-and-a-halfth-order logic is **consistent**, i.e. $\vdash_{\Delta}$ can never be derived.*

PROOF. By contradiction. Suppose $\vdash_{\Delta}$ is derivable, then by Theorem 6.8 a cut-free derivation exists. Let Π be the shortest derivation of $\vdash_{\Delta}$ for all possible Δ. We check through all possible derivation rules and see by their syntax-directed nature that the derivation must conclude in ($\mathbf{Fr}$). But then we have a shorter derivation of some $\vdash_{\Delta'}$, which is a contradiction. □

7. Equivalence of $\vdash_{\Delta}$ and $\vdash_{\mathsf{FOL}}$

This section shows how derivability in the sequent calculus of one-and-a-halfth-order logic relates to derivability in theory FOL (Theorem 7.5).

We need some notation.

For a context $\Phi=\{\phi_1,\ldots,\phi_n\}$, define its **conjunctive form** $\Phi^{\wedge}$ to be $\top$ when $n=0$, and $\phi_1\wedge\cdots\wedge\phi_n$ when $n>0$. Analogously, define the **disjunctive form** $\Phi^{\vee}$ to be $\bot$ when $n=0$, and $\phi_1\vee\cdots\vee\phi_n$ when $n>0$. The order of the ϕ_i is irrelevant; we (promise) never (to) do anything where it matters.

Sequent derivability translates to FOL derivability in the following way:

LEMMA 7.1. *If $\Phi\vdash_{\Delta}\Psi$ then $\Delta\vdash_{\mathsf{FOL}}\Phi^{\wedge}\supset\Psi^{\vee}=\top$.*

PROOF. By induction on derivations of $\Phi\vdash_{\Delta}\Psi$. For every rule ($\mathbf{R}$), the derivation has the following format:

$$\dfrac{\Pi_1\quad\cdots\quad\Pi_k}{\Phi\vdash_{\Delta}\Psi}(\mathbf{R})\quad(cond)$$

Here $k\in\{0,1,2\}$, Π_i are derivations of $\Phi_i\vdash_{\Delta_i}\Psi_i$, $1\le i\le k$, and $cond$ is a (possibly empty) side-condition.

So $\Phi_i\vdash_{\Delta_i}\Psi_i$ are derivable, then by inductive hypothesis $\Delta_i\vdash_{\mathsf{FOL}}{\Phi_i}^{\wedge}\supset{\Psi_i}^{\vee}=\top$ holds. We use this together with $cond$ to prove $\Delta\vdash_{\mathsf{FOL}}\Phi^{\wedge}\supset\Psi^{\vee}=\top$. For each inference rule ($\mathbf{R}$), this is an instance of an item of Lemma 4.3.

For example, in case ($\mathbf{R}$) is ($\mathbf{Cut}$) $\Delta\vdash_{\mathsf{FOL}}\Phi^{\wedge}\supset\Psi^{\vee}=\top$ should follow from the assumptions $\Delta\vdash_{\mathsf{FOL}}\Phi^{\wedge}\supset\Psi^{\vee}\vee\phi=\top$, $\Delta\vdash_{\mathsf{FOL}}\phi'\wedge\Phi^{\wedge}\supset\Psi^{\vee}=\top$ and $\Delta\vdash_{\mathsf{SUB}}\phi=\phi'$. This is an instance of item 12 of Lemma 4.3, using $\theta\equiv\Phi^{\wedge}$ and $\varepsilon\equiv\Psi^{\vee}$.

In case ($\mathbf{R}$) is ($\forall\mathbf{R}$) $\Delta\vdash_{\mathsf{FOL}}\Phi^{\wedge}\supset\Psi^{\vee}\vee\forall[a]\psi=\top$ should follow from $\Delta\vdash_{\mathsf{FOL}}\Phi^{\wedge}\supset\Psi^{\vee}\vee\psi=\top$ and $\Delta\vdash a\#\Phi^{\wedge},\Psi^{\vee}$. This is an instance of item 6, again using $\theta\equiv\Phi^{\wedge}$ and $\varepsilon\equiv\Psi^{\vee}$. □

For the reverse of the above, we need a number of technical lemmas.

LEMMA 7.2. *Bi-implication $\Leftrightarrow$ is an equivalence relation.*[5] *Also, $\vdash_{\Delta}\top\Leftrightarrow\phi$ if and only if $\vdash_{\Delta}\phi$ if and only if $\vdash_{\Delta}\phi\Leftrightarrow\top$.*

LEMMA 7.3. *For all sorts τ, terms $t,u:\tau$, freshness contexts Δ and contexts $C[_]:\mathbb{P}$:*

$$\textit{if}\quad\Delta\vdash_{\mathsf{FOL}}t=u\quad\textit{then}\quad\vdash_{\Delta}C[t]\Leftrightarrow C[u]$$

PROOF. By induction on the structure of FOL derivations of $t=u$ from Δ.
($\mathbf{refl}$): $\vdash_{\Delta}C[t]\Leftrightarrow C[t]$ follows by reflexivity of $\Leftrightarrow$.
($\mathbf{symm}$): $\vdash_{\Delta}C[u]\Leftrightarrow C[t]$ follows from $\vdash_{\Delta}C[t]\Leftrightarrow C[u]$ by symmetry of $\Leftrightarrow$. By inductive hypothesis this follows from the assumption. The case of ($\mathbf{tran}$) is similar.
($\mathbf{cong}$): $\vdash_{\Delta}C[D[t]]\Leftrightarrow C[D[v]]$ follows by inductive hypothesis from the assumption taking $C[_]:=C[D[_]]$.
($\mathbf{perm}$): we show $\vdash_{\Delta}C[(a\ b)\cdot t]\Leftrightarrow C[t]$ as follows:

$$\dfrac{\dfrac{}{\vdash_{\Delta}C[(a\ b)\cdot t]\Leftrightarrow C[(a\ b)\cdot t]}(\mathbf{Ax})}{\vdash_{\Delta}C[(a\ b)\cdot t]\Leftrightarrow C[t]}(\mathbf{StructR})$$

where $\Delta\vdash_{\mathsf{SUB}}C[(a\ b)\cdot t]\Leftrightarrow C[(a\ b)\cdot t]=C[(a\ b)\cdot t]\Leftrightarrow C[t]$ is the side-condition of ($\mathbf{StructR}$). By ($\mathbf{cong}$), this follows from the assumption $\Delta\vdash_{\mathsf{SUB}}(a\ b)\cdot t=t$.
($\mathbf{fr}$): By ($\mathbf{Fr}$) derivability of $\vdash_{\Delta}C[t]\Leftrightarrow C[u]$ follows from that of $\vdash_{\Delta,a\#X_1,\ldots,a\#X_n}C[t]\Leftrightarrow C[u]$ and $a\notin C[t]\Leftrightarrow C[u],\Delta$. The former follows by inductive hypothesis, and the latter from the assumption $a\notin C[t],C[u],\Delta$.
($\mathbf{ax_A}$): If A is an axiom of SUB (one from Figure 4), the proof is analogous to the ($\mathbf{perm}$) case. If A is a logical axiom (one from Figure 5) then, looking at the structure of the axioms, the derivation is of the form

$$\dfrac{\Pi}{\phi^{\pi}\sigma=\top}(\mathbf{ax_A})$$

where A is $\Delta'\to\phi=\top$ and Π is a derivation of ${\Delta'}^{\pi}\sigma$. We need to show $\vdash_{\Delta}C[\phi^{\pi}\sigma]\Leftrightarrow C[\top]$. By congruence and right identity of $\Leftrightarrow$, this follows from $\vdash_{\Delta}\phi^{\pi}\sigma$. For each logical axiom A, this is an instance of an item of Lemma 5.1, using the assumption $\Delta\vdash{\Delta'}^{\pi}\sigma$. □

LEMMA 7.4. *If $\vdash_{\Delta}\Phi^{\wedge}\supset\Psi^{\vee}$ then $\Phi\vdash_{\Delta}\Psi$.*

PROOF. By ($\mathbf{Cut}$) $\Phi\vdash_{\Delta}\Psi$ follows from $\Phi\vdash_{\Delta}\Psi$, $\Phi^{\wedge}\supset\Psi^{\vee}$ and $\Phi^{\wedge}\supset\Psi^{\vee}$, $\Phi\vdash_{\Delta}\Psi$. The former follows from the assumption using Lemma 6.5. The latter follows from $\Phi\vdash_{\Delta}\Psi$, $\Phi^{\wedge}$ and $\Psi^{\vee}$, $\Phi\vdash_{\Delta}\Psi$ using ($\supset\mathbf{L}$). We prove the former, the latter is analogous. We know $\Phi=\{\phi_1,\ldots,\phi_n\}$, where $n\ge0$. If $n=0$

[5] $\Leftrightarrow$ is a reflexive symmetric transitive congruence.

then $\Phi^\wedge = \top$ and we derive $\vdash_\Delta \Psi, \top$ using $(\supset\mathbf{R})$ and $(\mathbf{Ax})$. If $n > 0$ then $\Phi^\wedge = \phi_1 \wedge \ldots \wedge \phi_n$, and it suffices to derive $\Phi \vdash_\Delta \Psi, \phi_i$ for all $i < n$ using $(\wedge\mathbf{R})$ ($n-1$ times). For each i, this follows by $(\mathbf{Ax})$, since $\phi_i \in \Phi$. □

Sequent derivability is equivalent to FOL derivability:

THEOREM 7.5. *$\Phi \vdash_\Delta \Psi$ if and only if $\Delta \vdash_{\mathsf{FOL}} \Phi^\wedge \supset \Psi^\vee = \top$.*

PROOF. The left-to-right part is handled by Lemma 7.1.

For the right-to-left part, we assume $\Delta \vdash_{\mathsf{FOL}} \Phi^\wedge \supset \Psi^\vee = \top$. Then by Lemma 7.3 $\vdash_\Delta \Phi^\wedge \supset \Psi^\vee \Leftrightarrow \top$ is derivable. By the right identity of $\Leftrightarrow$, also $\vdash_\Delta \Phi^\wedge \supset \Psi^\vee$. By Lemma 7.4 we obtain the consequent $\Phi \vdash_\Delta \Psi$. □

This theorem has some nice corollaries.

COROLLARY 7.6. *For any Δ, ϕ, ψ:*

$$\Delta \vdash_{\mathsf{FOL}} \phi = \psi \text{ if and only if } \phi \vdash_\Delta \psi \text{ and } \psi \vdash_\Delta \phi.$$

PROOF. By Theorem 7.5 $\phi \vdash_\Delta \psi$ and $\psi \vdash_\Delta \phi$ are equivalent to $\Delta \vdash_{\mathsf{FOL}} \phi \supset \psi = \top$ and $\Delta \vdash_{\mathsf{FOL}} \psi \supset \phi = \top$. These can easily be shown equivalent to $\Delta \vdash_{\mathsf{FOL}} \phi = \psi$ using item 2 of Lemma 4.1. □

COROLLARY 7.7. FOL *is consistent, i.e. $\Delta \vdash_{\mathsf{FOL}} \top = \bot$ does not hold for any Δ.*

PROOF. By contradiction. Suppose $\Delta \vdash_{\mathsf{FOL}} \top = \bot$. Using the propositional axioms and some simple equational reasoning, also $\Delta \vdash_{\mathsf{FOL}} \top \supset \bot = \top$. Note that $\top \equiv \emptyset^\wedge$ and $\bot \equiv \emptyset^\vee$, so by Theorem 7.5 $\vdash_\Delta$ is derivable, which contradicts Corollary 6.9. □

8. First-order logic

Call a term **ground** if it does not mention unknowns (it is closed) and it does not mention explicit substitutions.

In this section we show formally how a syntax for a first-order logic 'lives inside' one-and-a-halfth-order logic, given by the ground terms of sort $\mathbb{P}$ (the **predicates**) taken up to α-equivalence of $\forall$-abstracted atoms, and the ground terms of sort $\mathbb{T}$ (the **term-language**). The precise term-language depends on the set of object-level term-formers and atomic predicate-formers with which we built our one-and-a-halfth-order logic in Section 2. As mentioned before, we let of and op vary over object-level term-formers and atomic predicate-formers.

8.1 Properties of ground terms

We may write the term $\forall[a]\phi$ where ϕ is ground just as $\forall a.\phi$ (consistent with standard notation). Recall that a in $\forall[a]P$ may *not* be renamed in general, e.g. to $\forall[b]P$. Intuitively P represents an unknown formula which *might* mention a (if we know $b\#P$ we can at least rename to $\forall[b](b\ a).P$). To emphasise this we retained the notation $[a]_$ until now. In $\forall a.\phi$ where ϕ is ground, we know all atoms in ϕ and this issue does not arise.

Write $fn(t)$ and $fn(\phi)$ for the **free names** of ground terms $t : \mathbb{T}$ and $\phi : \mathbb{P}$ respectively, inductively defined by:

$$fn(a) = \{a\} \qquad fn(\mathsf{of}(t_1, \ldots, t_k)) = \bigcup_{1 \le i \le k} fn(t_i)$$

$$fn(\bot) = \emptyset \qquad fn(\phi \supset \psi) = fn(\phi) \cup fn(\psi)$$

$$fn(\forall a.\phi) = fn(\phi) \setminus \{a\} \qquad fn(\mathsf{op}(t_1, \ldots, t_k)) = \bigcup_{1 \le i \le k} fn(t_i)$$

LEMMA 8.1. *$\vdash a\#\phi$ if and only if $a \notin fn(\phi)$, for all ground predicates ϕ.*

PROOF. By simple induction on derivations of $a\#\phi$ on the one hand, and by induction on the definition of fn on the other. □

Define **α-equivalence** $=_\alpha$ as syntactic identity plus

$$\frac{b \notin fn(\phi) \quad (a\ b)\cdot\phi =_\alpha \psi}{\forall a.\phi =_\alpha \forall b.\psi}.$$

The reader might have expected the clause for $\forall$ to read something like $\dfrac{\phi_c =_\alpha \psi_c}{\forall a.\phi_a =_\alpha \forall b.\psi_b}$ where here ϕ_c is informal notation for ϕ with every a replaced throughout by a freshly chosen c, and similarly for ψ_c. The two notions of α-equivalence are identical [11]. The definition we adopt gives a closer match to how equality is defined in NA (specifically to $(\mathbf{perm})$).

LEMMA 8.2. *$\vdash_{\mathsf{CORE}} \phi = \psi$ if and only if $\phi =_\alpha \psi$, for all ground predicates ϕ, ψ.*

PROOF. By known arguments of nominal results [15, 16]. □

For each finite set of atoms make an arbitrary but canonical choice of a fresh (that is, not in the set) atom. In a given context of some finite collection of terms and predicates, which being finite mentions finitely many atoms, write 'a fresh' for *the* canonical but arbitrarily chosen a which is fresh for the atoms in that collection.

In the following definition we elide the context, which is the terms and formulae mentioned to the left of $\equiv$.

For ground terms $t, u : \mathbb{T}$ and $\phi : \mathbb{P}$, write $u[\![a \mapsto t]\!]$ and $\phi[\![a \mapsto t]\!]$ for u and ϕ **with** a **replaced by** t, inductively defined by:

$$a[\![a \mapsto t]\!] \equiv t \qquad b[\![a \mapsto t]\!] \equiv b$$

$$\mathsf{of}(t_1, \ldots, t_k)[\![a \mapsto t]\!] \equiv \mathsf{of}(t_1[\![a \mapsto t]\!], \ldots, t_k[\![a \mapsto t]\!])$$

$$\bot[\![a \mapsto t]\!] \equiv \bot \qquad (\phi \supset \psi)[\![a \mapsto t]\!] \equiv \phi[\![a \mapsto t]\!] \supset \psi[\![a \mapsto t]\!]$$

$$(\forall a.\phi)[\![a \mapsto t]\!] \equiv \forall a.\phi$$

$$(\forall b.\phi)[\![a \mapsto t]\!] \equiv \forall b'.\phi[\![b \mapsto b']\!][\![a \mapsto t]\!] \quad (b' \text{ fresh})$$

$$\mathsf{op}(t_1, \ldots, t_k)[\![a \mapsto t]\!] \equiv \mathsf{op}(t_1[\![a \mapsto t]\!], \ldots, t_k[\![a \mapsto t]\!])$$

LEMMA 8.3. *For all ground terms $t, u, v : \mathbb{T}$ and $\phi, \psi : \mathbb{P}$:*

$$\vdash_{\mathsf{SUB}} u[a \mapsto t] = u[\![a \mapsto t]\!] \quad \text{and} \quad \vdash_{\mathsf{SUB}} \phi[a \mapsto t] = \phi[\![a \mapsto t]\!]$$

PROOF. By induction on the depths of u and ϕ. For the case $u \equiv \mathsf{f}(t_1, \ldots, t_n)$, we must prove

$$\vdash_{\mathsf{SUB}} \mathsf{f}(t_1, \ldots, t_n)[a \mapsto t] = \mathsf{f}(t_1[\![a \mapsto t]\!], \ldots, t_n[\![a \mapsto t]\!]),$$

for which we use axiom $(\mathbf{f}\mapsto)$ of SUB and the inductive hypothesis.

The only difficult case is $\phi \equiv \forall b.\phi'$ because there is no directly corresponding axiom of SUB. By calculation of $(\forall b.\phi')[\![a \mapsto t]\!]$, we obtain $\forall b'.\phi'[\![b \mapsto b']\!][\![a \mapsto t]\!]$, where b' is fresh. By the inductive hypothesis we have

$$\vdash_{\mathsf{SUB}} \phi'[b \mapsto b'] = \phi'[\![b \mapsto b']\!]$$
$$\vdash_{\mathsf{SUB}} \phi'[\![b \mapsto b']\!][a \mapsto t] = \phi'[\![b \mapsto b']\!][\![a \mapsto t]\!].$$

We need to prove

$$\vdash_{\mathsf{SUB}} (\forall b.\phi')[a \mapsto t] = \forall b'.\phi'[\![b \mapsto b']\!][\![a \mapsto t]\!],$$

which follows by easy calculations using the above assumptions. □

LEMMA 8.4. *For all ground terms $t, u, v : \mathbb{T}$ and $\phi, \psi : \mathbb{P}$:*

1. *$\vdash_{\mathsf{SUB}} u[a \mapsto t] = v$ if and only if $u[\![a \mapsto t]\!] \equiv v$;*[6]
2. *$\vdash_{\mathsf{SUB}} \phi[a \mapsto t] = \psi$ if and only if $\phi[\![a \mapsto t]\!] =_\alpha \psi$.*

[6] If we had binders in terms of sort $\mathbb{T}$ then the $\equiv$ would become an α-equivalence.

PROOF. By lemmas 8.2 and 8.3, using the fact that SUB is conservative over CORE (see [14]). □

8.2 Derivability in First-Order Logic

A **first-order context** is a finite and possibly empty set of ground predicates Φ or Ψ. A **first-order sequent** is a pair $\Phi \vdash \Psi$. The **valid judgements** of Gentzen's sequent calculus for first-order logic are inductively derived by:

$$\frac{}{\phi,\ \Phi \vdash \Psi,\ \phi}\ (\mathbf{Ax}) \qquad \frac{}{\bot,\ \Phi \vdash \Psi}\ (\bot\mathbf{L})$$

$$\frac{\Phi \vdash \Psi,\ \phi \quad \psi,\ \Phi \vdash \Psi}{\phi \supset \psi,\ \Phi \vdash \Psi}\ (\supset\mathbf{L}) \qquad \frac{\phi,\ \Phi \vdash \Psi,\ \psi}{\Phi \vdash \Psi,\ \phi \supset \psi}\ (\supset\mathbf{R})$$

$$\frac{\phi[\![a \mapsto t]\!],\ \Phi \vdash \Psi}{\forall a.\phi,\ \Phi \vdash \Psi}\ (\forall\mathbf{L}) \qquad \frac{\Phi \vdash \Psi,\ \phi}{\Phi \vdash \Psi,\ \forall a.\phi}\ (\forall\mathbf{R})\ (a \not\in \mathit{fn}(\Phi,\Psi))$$

$$\frac{\phi[\![a \mapsto t']\!],\ \Phi \vdash \Psi}{t' \approx t,\ \phi[\![a \mapsto t]\!],\ \Phi \vdash \Psi}\ (\approx\mathbf{L}) \qquad \frac{}{\Phi \vdash \Psi,\ t \approx t}\ (\approx\mathbf{R})$$

Here we take predicates up to α-equivalence, e.g. if $\mathsf{p} : (\mathbb{T})\mathbb{P}$ is an atomic predicate term-former then $\forall a.\mathsf{p}(a) \vdash \forall b.\mathsf{p}(b)$ follows directly by $(\mathbf{Ax})$ since $\forall a.\mathsf{p}(a) =_\alpha \forall b.\mathsf{p}(b)$.

THEOREM 8.5. *$\Phi \vdash \Psi$ is derivable in the system above, if and only if $\Phi \vdash_\emptyset \Psi$ is derivable in the sequent calculus for one-and-a-halfth-order logic.*

PROOF. By induction on the structure of derivations, using cut-elimination (Theorem 6.8) and the results from Subsection 8.1 (Lemmas 8.1 and 8.4). □

COROLLARY 8.6. *$\vdash_{\mathsf{FOL}} \phi = \psi$ if and only if $\phi \vdash \psi$ and $\psi \vdash \phi$ are derivable in system above.*

PROOF. By Corollary 7.6 and Theorem 8.5. □

9. Conclusions

Explicitly representing meta-variables has a long pedigree.

Monadic second-order logic [5] enriches first-order logic explicitly with n-ary relation variables, representing 'unknown n-ary predicates'. The stronger second-order and higher-order logics [30, 27] represent unknowns as function variables.

These approaches share two characteristic features inherited from their intended functional semantics. First, you have to choose the arity of your unknown in advance, e.g. $f : \overbrace{\mathbb{T} \to \cdots \to \mathbb{T}}^{n} \to \mathbb{P}$ can be interpreted as an unknown n-ary predicate — but *which* n? — thus these logics distribute 'unknown predicates' across many types. Second, and perhaps more importantly, instantiation of these variables avoids capture. Instantation of our unknowns does not, which accurately reflects our intention when we write $\forall a.\phi$, where ϕ may be instantiated in a capturing manner. This lends a distinctive style to our nominal algebra theory FOL (see Figure 5), and to the sequent rules of one-and-a-halfth-order logic (see Figure 1), which accurately reflects informal practice (see [10, 20] and the examples of the Introduction) and as we have seen has allowed us to import elements of first-order proof theory quite directly into an augmented setting.

Note that second-order logic is *far* more expressive. One-and-a-halfth-order logic by design expresses universal quantification at top level only. For example, the following second-order theorem cannot be expressed: $\forall P.(\forall P.P) \supset P$. On the other hand one-and-a-halfth-order logic can express some things which second-order logic cannot, e.g. $\forall[a]P \vdash_{\mathsf{a\#P}} P$ is derivable. It is not presently clear what the intersection of these two logics is.

Our ambient nominal algebra framework *can* express *certain* strong 'second-order' principles, for instance an inductive principle on natural numbers may be expressed as a single axiom (assuming suitable term-formers 0 and succ)

$$P[a \mapsto 0] \wedge \forall[a](P \supset P[a \mapsto \mathit{succ}(a)]) \supset \forall[a]P \ = \ \top$$

rendered in second-order logic as

$$\forall P.(P(0) \wedge \forall a.(P(a) \supset P(\mathit{succ}(a))) \supset \forall a.P).$$

Our work is one (more) element in a very long line of investigations into algebraic logic [1]; for example cylindric [19], polyadic [18], and quantifier [25] algebra. There too, unknowns are syntax representing unknown elements quantified universally at top level, and abstraction $[a]$- (our notation) is clearly visible, e.g. as the c_i of cylindric algebra. Our treatment of substitution is perhaps cleaner; cylindric and quantifier algebras axiomatise logic whose terms are in a suitable sense restricted to being atoms (our terminology), whereas we could easily extend our logic to talk about, say, λ-calculus terms just by postulating term-formers $\lambda : ([\mathbb{A}]\mathbb{T})\mathbb{T}$ and $\mathit{app} : (\mathbb{T},\mathbb{T})\mathbb{T}$ plus a few nominal equalities [15, 7] (here, interesting work by Beeson is also relevant [3]). Arguably our treatment of substitution is also more systematic than polyadic algebras, at least in the sense that we define it in terms of more primitive constructs and can so study substitution in its own right [14].

Probably more important is our use of *freshness contexts* and *permutations*, by means of which we can express properties such as $b\#P \vdash_{\mathsf{FOL}} \forall[a]P = \forall[b](b\ a) \cdot P$, and of course the examples of the Introduction, directly. These assertions are valid in, say, cylindric algebra, but only by recourse to quantification over closed terms, thus, they cannot be stated *within* that framework. The extra expressivity we enjoy thanks to a slightly richer term-language and judgement-form is significant, because we can exploit it to extract a sequent presentation of derivability, namely the sequent system of Figure 1. In other words, our *nominal* management of binding permits us to construct an account of first-order logic which in a suitable sense is *simultaneously* algebraic *and* sequent-style.

It is possible to represent the syntax of a logic in a 'framework' logical system, at 'object-level', i.e. as an inductive datatype. Then meta-variables are easily representable as meta-variables of the framework. This path is taken by Higher-Order Abstract Syntax [24], Fraenkel-Mostowski syntax [16], and other systems ([23] is just one of very many). That is a separate enterprise from that undertaken in this paper; one-and-a-halfth-order logic is about extending the syntax of the logic itself so it contains something which behaves very much like a meta-variable ranging over unknown formulae, without losing logical properties such as cut-elimination. Perhaps one day one-and-a-halfth-order logic too will be formalised in a framework!

The technical tools used in this paper were developed based on work on Nominal Unification by the first author with Urban and Pitts [29], which introduced the theory of nominal terms up to CORE (our terminology). This was extended with Fernández [7] to Nominal Rewriting, a theory of rewriting on nominal terms, again up to CORE, and recently investigated with the second author, as a general framework of nominal algebra [15, 14].

For future work we are particularly interested in the following topics:

We can return to theory and be inspired by higher-order logic to ask whether we could permit abstraction over meta-variables, introducing an infinite hierarchy of stronger meta-variables such that at each level a meta-variable of higher level behaves to the lower level as X behaves to a (see the NEW calculus of contexts [13]). This might recover some or all of the power which one-and-

a-halfth-order logic lacks compared to higher-order logic, but in a different way. In short, we envisage two- three- four- and ω-and-a-halfth-order logic. This would involve interesting extensions to the 'nominal theme'. Another direction is to allow unknowns ranging over derivations of sequents, which may have interesting interactions with $(\forall\mathbf{R})$, which would abstract in such an unknown.

The semantics of one-and-a-halfth-order logic are interesting and raise the questions 'what is an appropriate semantics for X', and 'what is an appropriate semantics for a'? Note that it is not possible to directly evaluate X to an element of a set underlying domain, because intuitively X 'can mention a'. Thus we can use domains in which atoms *can* appear (à la Fraenkel-Mostowski sets [15, 16] or other approaches [3, 8]). The simplest solution, and perhaps the best one, is to evaluate X to terms (a 'substitutional semantics' [22, Section 2] faithful to its intuition as an 'unknown term') and then a to elements of a set underlying domain.

References

[1] H. Andréka, I. Németi, and I. Sain. Algebraic logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, 2nd Edition*, volume 2, pages 133–249. Kluwer, 2001.

[2] J. Barwise. An introduction to first-order logic. In J. Barwise, editor, *Handbook of Mathematical Logic*, pages 5–46. North Holland, 1977.

[3] M. Beeson. Lambda logic. In *Second International Joint Conference on Automated Reasoning (IJCAR 2004)*, volume 3097 of *LNCS*, pages 460–474. Springer, 2004.

[4] J. Bell and M. Machover. *A course in mathematical logic*. North-Holland, 1977.

[5] C. Bruno. The expression of graph properties in some fragments of monadic second-order logic. *DIMACS Series in Discrete Mathematics*, 31, 1997.

[6] A. Degtyarev and A. Voronkov. Equality reasoning in sequent-based calculi. In J. A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, pages 611–706. Elsevier and MIT Press, 2001.

[7] M. Fernández and M. J. Gabbay. Nominal rewriting. Journal version, submitted Information and Computation, 2005.

[8] K. Fine. *Reasoning with Arbitrary Objects*. Blackwell, 1985.

[9] J. B. Fraleigh. *A First Course in Abstract Algebra*. Addison-Wesley, 7th edition, 2002.

[10] D. M. Gabbay and G. Malod. Naming worlds in modal and temporal logic. *Journal of Logic, Language and Information*, 11(1):29–65, 2002.

[11] M. J. Gabbay. *A Theory of Inductive Definitions with alpha-Equivalence*. PhD thesis, Cambridge, UK, 2000.

[12] M. J. Gabbay. Fresh logic. *Journal of Logic and Computation*, July 2003. Accepted for publication.

[13] M. J. Gabbay. A new calculus of contexts. In *Proc. 7th Int. ACM SIGPLAN Conf. on Principles and Practice of Declarative Programming (PPDP'2005)*. ACM, 2005.

[14] M. J. Gabbay and A. Mathijssen. Capture-avoiding substitution as a nominal algebra. Submitted ICTAC'06.

[15] M. J. Gabbay and A. Mathijssen. Nominal algebra. Submitted CSL'06.

[16] M. J. Gabbay and A. M. Pitts. A new approach to abstract syntax with variable binding. *Formal Aspects of Computing*, 13(3–5):341–363, 2001.

[17] G. Gentzen. Untersuchungen über das logische schließen [Investigations into logical deduction]. *Mathematische Zeitschrift 39*, pages 176–210,405–431, 1935. Translated in [28], pages 68–131.

[18] P. Halmos. Algebraic logic, ii. homogeneous locally finite polyadic boolean algebras of infinite degree. *Fundamenta Mathematicae*, 43:255–325, 1956.

[19] L. Henkin, J. D. Monk, and A. Tarski. *Cylindric Algebras*. North Holland. Part I (1971), Part II (1985).

[20] W. Hodges. Elementary predicate logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, 2nd Edition*, volume 1, pages 1–131. Kluwer, 2001.

[21] P. T. Johnstone. *Notes on logic and set theory*. Cambridge University Press, 1987.

[22] H. Leblanc. Alternatives to standard first-order semantics. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, 2nd Edition*, volume 2, pages 53–132. Kluwer, 2001.

[23] M. Miculan. Developing (meta)theory of lambda-calculus in the theory of contexts. *ENTCS*, 1(58), 2001.

[24] F. Pfenning and C. Elliot. Higher-order abstract syntax. In *SIGPLAN Conference on Programming Language Design and Implementation*, pages 199–208, 1988.

[25] C. Pinter. A simple algebra of first-order logic. *Notre Dame Journal of Formal Logic*, 14(3):361–366, 1973.

[26] D. Prawitz. *Natural Deduction: A Proof Theoretical Study*. Almqvist and Wiksell, Stockholm, 1965.

[27] S. Shapiro. Systems between first-order and second-order logics. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, 2nd Edition*, volume 1, pages 131–188. Kluwer, 2001.

[28] M. Szabo, editor. *Collected Papers of Gerhard Gentzen*. North Holland, 1969.

[29] C. Urban, A. M. Pitts, and M. J. Gabbay. Nominal unification. *Theoretical Computer Science*, 323(1–3):473–497, 2004.

[30] J. van Benthem. Higher-order logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, 2nd Edition*, volume 1, pages 189–244. Kluwer, 2001.