

MURDOCH J. GABBAY

NOMINAL TERMS AND NOMINAL LOGICS: FROM FOUNDATIONS TO META-MATHEMATICS

Nominal techniques concern the study of names using mathematical semantics. Whereas in much previous work names in abstract syntax were studied, here we will study them in meta-mathematics. More specifically, we survey the application of nominal techniques to languages for unification, rewriting, algebra, and first-order logic.

What characterises the languages of this chapter is that they are first-order in character, and yet they can specify and reason on names. In the languages we develop, it will be fairly straightforward to give first-order ‘nominal’ axiomatisations of name-related things like alpha-equivalence, capture-avoiding substitution, beta- and eta-equivalence, first-order logic with its quantifiers—and as we shall see, also arithmetic. The formal axiomatisations we arrive at will closely resemble ‘natural behaviour’; the specifications we see typically written out in normal mathematical usage.

This is possible because of a novel name-carrying semantics in nominal sets, through which our languages will have name-permutations and term-formers that can bind as primitive built-in features.

This chapter draws together material from several papers to deliver a coherent account of a journey from the foundations of a mathematics with names, via logical systems based on those foundations, to concrete applications in axiomatising systems with binding. Definitions and proofs have been improved, generalised, and shortened, and placed into an overall narrative.

On the way we touch on a variety of definitions and results. These include: the nominal unification algorithm; nominal rewriting and its confluence proofs; nominal algebra, its soundness, completeness, and an HSP theorem; permissive-nominal logic and its soundness and completeness; various axiomatisations with pointers to proofs of their correctness; and we conclude with a case study stating and proving correct a finite first-order axiomatisation of arithmetic in permissive-nominal logic.

1 INTRODUCTION

Nominal sets for meta-mathematics Suppose we want to axiomatise the λ -calculus or first-order logic. Then we need to express properties like this:

- If $y \notin fv(t)$ then $\forall x.t =_{\alpha} \forall y.(t[y/x])$.

- If $x \notin fv(u)$ then $(\lambda x.t)[u/y] = \lambda x.(t[u/y])$.
- If $x \notin fv(t)$ then $\lambda x.(tx) =_{\eta} t$.

x , y , t , and u here are what we would call *names*. A linguist might call them *referents*, a mathematician might call them *variables*. But the words ‘referent’ and ‘variable’ carry connotations (a referent should refer to something, a variable should vary), so we prefer the more neutral term ‘name’. So for us, a name is just an atomic symbol, to which we may then associate further properties, at our discretion, using additional axioms.

The axioms above are typical of a certain kind of specification. Mathematical specification is nothing new. First-order logic can specify, to choose a classic trio of examples, groups, rings, and fields. But the λ -calculus, first-order logic itself, the π -calculus, and a very great many other examples, are different. They have *names*.

By adding names to first-order logic in the correct way, we can axiomatise the specifications above, cleanly and in a manner very close to the informal specification. How should we do this? Using a recent application of mathematical foundations originating in computer science: *nominal sets* [Gabbay and Pitts, 2001; Gabbay, 2011b], to which we will use *nominal terms* [Urban *et al.*, 2003; Urban *et al.*, 2004] as a corresponding formal syntax. To survey and update the state of the art of logics based on nominal terms and taking semantics in nominal sets, is our goal here.

In nominal terms, term-formers can bind names and freshening renamings like the $[y/x]$ or $[u/y]$ above are taken as primitive.

Here are the informal statements above, rewritten in permissive-nominal algebra—an algebraic logic based on nominal terms with a sound and complete semantics in nominal sets:

- If $b \notin \text{supp}(X)$ then $\forall([a]X) = \forall([b](b a) \cdot X)$.
- If $a \notin \text{supp}(Y)$ then $\lambda([a]X)[b \rightarrow Y] = \lambda([a](X[b \rightarrow Y]))$.
- If $a \notin \text{supp}(X)$ then $\lambda([a](Xa)) = X$.

In this chapter we will briefly consider nominal sets, then survey nominal terms, unification, rewriting, algebra, and permissive-nominal logic. We cover the nominal unification algorithm, confluence proofs for nominal rewriting, soundness and completeness results for nominal algebra and permissive-nominal logic, an HSP theorem, and a finite axiomatisation of first-order logic.

By doing this we aim to give an overview of the applications of nominal sets to meta-mathematical syntax. We cannot be exhaustive, but we can try to be representative of what can be achieved.

As we shall see, nominal syntax is more expressive than first-order syntax (for instance we can give a finite first-order axiomatisation of arithmetic), because term-formers that can explicitly manipulate names. Yet, it remains first-order in flavour, preserving theoretical and computational properties like completeness and most general unifiers.

A few words on atoms What nominal sets add to ‘ordinary’ structures is an assumption of a distinguished class of symmetric atomic elements called *atoms*: these are also called *urelemente* or *names*. We will use these terms more-or-less synonymously.

Indeed, nominal sets are a special case Zermelo-Fraenkel sets with atoms, and are instances of the structures considered by Fraenkel and Mostowski in their celebrated independence proof of the the Axiom of Choice from the other axioms of set theory with atoms. For detailed references see [Gabbay, 2011b, Remark 2.22]. So this chapter really does describe a journey from mathematical foundations to meta-mathematics, and that is representative of how the maths we describe here was arrived at.

We can view the underlying philosophy of nominal techniques is as the following informal inequality, where ‘smaller’ means ‘greater generality’:

$$\text{atoms} = \text{urelemente} = \text{names} \leq \text{referents} \leq \text{variables}$$

Discovering to what extent these intuitions can be made precise, concrete, and useful, is the topic of much ongoing research, some of which is reported on here.

Names induce automorphisms generated by permuting them. We shall see that if we model variables as a special case of atoms, then α -renaming becomes a special case of a much more general fact that nominal sets are symmetric under permuting atoms. This generalisation turns out to have powerful consequences, including the atoms-abstraction and \mathcal{N} -quantifier introduced by the author with Pitts in [Gabbay and Pitts, 2001]. So the point of view described above has led to and continues to lead to new reasoning principles.

If we identify a thing with the properties of that thing, then the ‘nominal’ model suggests that names are equal to the following set of three properties:

$$\text{names} = \{\text{atomic, symmetric, generative}\}$$

The reader familiar with nominal techniques can identify these three properties with the use of: atomic symbols a (an atom, name, or urelement, with a distinct existence in the denotation), permutations π (symmetries under permutation of names), and the \mathcal{N} -quantifier (‘choose a fresh name’). These three properties will appear directly in this chapter as atoms, permutations, and permission sets.¹ Full definitions appear below.

This material in the literature This paper surveys existing literature on logics based on nominal terms, and adds a few new results. Very broadly, Section 2 is based on [Gabbay and Pitts, 1999; Gabbay and Pitts, 2001]

¹In other papers, such as [Urban *et al.*, 2004], permission sets are presented instead as syntactic freshness assumptions.

(nominal sets; they were called *equivariant FM sets* there); Sections 3 and 4 are based on [Urban *et al.*, 2003; Urban *et al.*, 2004; Dowek *et al.*, 2009; Dowek *et al.*, 2010; Gabbay, 2011c] (nominal terms and unification); Sections 5 and 6 are based on [Fernández *et al.*, 2004; Fernández and Gabbay, 2007; Gabbay, 2011c] (rewriting and closed terms); Section 7 is based on [Gabbay, 2005; Gabbay and Mathijssen, 2006a; Gabbay and Mathijssen, 2007; Gabbay and Mathijssen, 2009] (nominal algebra); Sections 9 to 11 are based on [Dowek and Gabbay, 2010; Dowek and Gabbay, 2011] (permissive-nominal logic).

Definitions and proofs may have changed from the original presentations. In particular:

- The semantics is *permissive*-nominal, meaning that it is based on possibly infinitely supported nominal sets with co-infinite support. In [Gabbay and Pitts, 2001] a nominal semantics based on finite and co-infinite support was used.
- Unlike [Urban *et al.*, 2004] and [Dowek *et al.*, 2010] we use nominal abstract syntax to build our nominal terms. That is, in this paper nominal terms atoms-abstraction is directly equal to Gabbay-Pitts atoms-abstraction. Thus, nominal terms here are an instance of nominal abstract syntax and come quotiented by α -equivalence by construction.
- Permutation may be stronger than usual, and we parameterise over the group of permutations.
We consider (as usual) finite permutations (generated by *swappings*, also called *transpositions*) as standard, but in particular we also find *shift*-permutations δ useful, which shift infinitely many atoms. The *shift*-permutation δ corresponds to a de Bruijn shift function \uparrow and presheaf reindexing map up , though δ is not equal to them since it is a permutation and so invertible.
- Syntax includes non-equivariant constant symbols. In [Urban *et al.*, 2004] all term-formers/function-symbols (including 0-ary ones, i.e. constants) were equivariant. This does not matter for finite support but it does make a difference with infinite support.
- Nominal unknowns are modelled as arbitrary elements of a strongly-supported nominal set. This means that the X and Y in this paper correspond to *moderated unknowns* from [Urban *et al.*, 2004]: see Example 3.1.7.
- Because unknowns have support, there are no freshness contexts and substitutions are characterised as equivariant functions (the freshness conditions normally attached to substitutions follow from equivariance: see Proposition 3.4.3). The theories of nominal unification, rewriting, and algebra are reformulated to reflect this.
- The simplification rules for unification problems (Figure 2) are new

and the treatments of closed terms and closed nominal rewriting (Section 6) are entirely revised with respect to [Fernández and Gabbay, 2007].

DRAFT

Part I

Nominal sets and nominal terms

2 NOMINAL SETS

We open with a brief presentation of nominal sets, which are the semantic basis for this work: this is the universe that the logics we define will describe, and be sound (and complete) for.

Nominal sets were developed with Pitts and introduced in the author's thesis [Gabbay, 2001], a conference paper [Gabbay and Pitts, 1999], and journal paper [Gabbay and Pitts, 2001]. The nominal sets here are more general than in [Gabbay and Pitts, 2001]: following [Dowek *et al.*, 2010] we are *permissive*, meaning that we split the set of atoms into two infinite halves and consider infinite support. This specific idea was developed jointly with Dowek,² but shades of it appear also in Cheney's paper [Cheney, 2006] and in the author's study of infinite atoms-abstraction [Gabbay, 2007b].

In addition we parameterise over a group of permutations which need not just be finitely-supported permutations. This is new.

2.1 Atoms, permutations, permission sets

In Definition 2.1.2 we need several sets of atoms. This is to model the several sorts of names that will appear in our syntax later on.

Following [Dowek *et al.*, 2010] our development will be *permissive*-nominal. A permission set S splits a set of atoms into two halves $\mathbb{A}^<$ and $\mathbb{A}^>$. One intuition for $\mathbb{A}^<$ is 'the atoms that have been generated so far', and for $\mathbb{A}^>$ is 'the atoms that might be generated later'.

DEFINITION 2.1.1. Write $\mathbb{N} = \{0, 1, 2, 3, \dots\}$ for the natural numbers and $\mathbb{Z} = \{0, -1, 1, -2, 2, \dots\}$ for the integers.

DEFINITION 2.1.2. For each $i \in \mathbb{N}$ fix a pair of disjoint countably infinite sets of **atoms** $\mathbb{A}_i^<$ and $\mathbb{A}_i^>$.

Write

$$\mathbb{A}_i = \mathbb{A}^< \uplus \mathbb{A}^> \quad \mathbb{A}^< = \bigcup \mathbb{A}_i^< \quad \mathbb{A}^> = \bigcup \mathbb{A}_i^> \quad \mathbb{A} = \bigcup \mathbb{A}_i$$

²The development here is a little different from that in [Dowek *et al.*, 2010] because we take permission sets to be sets of the form $\pi \cdot \mathbb{A}^<$ instead of sets of the form $(\mathbb{A}^< \setminus A) \cup B$.

a, b, c, \dots will range over *distinct* atoms: we call this the **permutative** convention.

REMARK 2.1.3 (*Comments on splitting the set of atoms*). The different sets of atoms \mathbb{A}_i are different ‘types’ of atoms. Thus, later on in Definitions 3.1.1 and 3.2.1 we can give each name sort its own distinct population of atoms.

The reasons for splitting the set of atoms into $\mathbb{A}^<$ and $\mathbb{A}^>$ will become clear as the maths develops. It might help to think of $\mathbb{A}^<$ as ‘atoms that can be captured’ and of $\mathbb{A}^>$ as ‘atoms that cannot be captured’, or as ‘atoms that might have been generated in the past’ and ‘atoms that may be generated in the future’—but with reservations. In Definition 2.1.10 we see that this is only true up to permuting atoms.

The real purpose of Definition 2.1.2 is to ensure that we have plenty—countably infinitely many—of ‘capturable’ and ‘non-capturable’ atoms. Permutations (below) can and will move atoms between these worlds, but no permutation can move them *all at once*. So the interest of $\mathbb{A}^<$ is not just for the set itself but for its orbit under permutations; this is a property of the set as a whole, and not of its individual elements.

REMARK 2.1.4 (*Comments on the permutative convention*). While visiting Tel-Aviv University in 2006 I gave talks on nominal techniques and Arnon Avron asked: “Do a and b refer to specific atoms (e.g. in the axioms in the Introduction), or to any two atoms?”. In other words, are a and b constants or variables?

In response I started using a *permutative convention* that a and b are variables, but they range over distinct atoms so that variables with distinct names refer to distinct objects (the first uses were in [Gabbay and Mathijssen, 2006c; Gabbay and Mathijssen, 2006a]; the convention was explicitly named in [Gabbay and Mathijssen, 2008c]).

For a while this was resisted by some anonymous referees. Yet, we typically apply the permutative convention informally; e.g. we silently assume that $\lambda x. \lambda y. xy$ is never the same term as $\lambda x. \lambda x. xx$. I would claim that the permutative convention expresses something about the foundational origins of the nominal view of names as *urelemente*—constants that are distinguishable yet symmetric—in an underlying set theory.

Perhaps this is why the referees did not like it: the permutative convention may seem unnatural if we are committed to standard (nameless) Zermelo-Fraenkel foundations, since names are then just some set, and like any set should be varied over non-permutatively by variables. Thus the fact that we accept that $\lambda x. \lambda y. xy$ and $\lambda x. \lambda x. xx$ always signify distinct λ -terms to us, can be taken as a sign that we inhabit a nameful foundation, so that the permutative convention is a signpost on the way to something more extensive.

A formal reflection of the permutative convention appears explicitly in

the formal logics of this paper: it lives in the π of the $\pi \cdot X$ in Definition 5.2.1.

DEFINITION 2.1.5. Given $a, b \in \mathbb{A}_i$ for some $i \in \mathbb{N}$ write $(a \ b)$ for the **swapping** bijection on atoms mapping a to b , b to a , and any other $c \in \mathbb{A} \setminus \{a, b\}$ to c .

Another standard name for a swapping is a **transposition**.

By convention $(a \ a)$ will denote the **identity** function on atoms id .

If π is a bijection on atoms define

$$nontriv(\pi) = \{a \mid \pi(a) \neq a\}.$$

DEFINITION 2.1.6. A **nominal permutation group** is any set of bijections \mathbb{P} of \mathbb{A} such that:

1. If $a \in \mathbb{A}_i$ and $b \in \mathbb{A}_i$ then $(a \ b) \in \mathbb{P}$.
2. If $\pi \in \mathbb{P}$ then $a \in \mathbb{A}_i$ if and only if $\pi(a) \in \mathbb{A}_i$.
3. There exists some infinite $S \subseteq \mathbb{A}$ such that $nontriv(\pi) \cap S$ is finite for every $\pi \in \mathbb{P}$.

Call a bijection on atoms π a **finite permutation** when it is in the subgroup generated by swappings. (π is finite when $\pi(a) \in \mathbb{A}_i$ if and only if $a \in \mathbb{A}_i$ and $nontriv(\pi)$ is finite.)

Write $\pi \circ \pi'$ for the **composition** of π and π' (so $(\pi \circ \pi')(a) = \pi(\pi'(a))$).

Write id for the **identity** permutation (so $id(a) = a$ always).

The purpose of conditions 1 to 3 of Definition 2.1.6 are as follows:

1. Swappings make sure we can always rename a to b (and b to a).
2. Condition 2 is a standard typing condition, that we do not try to turn an atom of one sort, into an atom of another sort.
3. This condition guarantees that we can still always choose a fresh atom for any finite set of permutations (see for instance Lemma 3.2.9).

EXAMPLE 2.1.7.


1. The set of all finite permutations is a nominal permutation group.
2. For each i fix a bijection f_i between \mathbb{A}_i and the integers \mathbb{Z} , such that $\{f(i) \mid i \leq 0\} = \mathbb{A}_i^<$ and (consequently) $\{f(i) \mid i > 0\} = \mathbb{A}_i^>$. We can do this because we assumed atoms are countable.

Write δ_i for the permutation mapping

- $f_i(j)$ to $f_i(j-1)$ for $j \leq 0$,
- $f_i(2j)$ to $f_i(2(j-1))$ and $f_i(2j-1)$ to $f_i(2j-1)$ for $j \geq 1$, and

- any other $c \in \mathbb{A} \setminus \mathbb{A}_i$ to c .

This is an example of a *shift*-permutation, considered in more generality in Definition 3.6.1 and throughout Subsection 3.6. We illustrate fragments of the actions of a swapping $(f(0) f(1))$ and a δ_i :



The atoms corresponding to positive odd integers are taken to be fixed points of δ_i in order to satisfy condition 3 of Definition 2.1.6, so that these atoms can be taken fresh for δ_i if we need to.

The set of permutations generated by swappings and δ_i , is a nominal permutation group.

REMARK 2.1.8. The nominal permutation group \mathbb{P} determines the symmetries of our nominal syntax and semantics. We consider permutations designed to guarantee (in Definition 2.2.3) symmetry up to equality/inequality of atoms. We will get sets with atoms that are atomic, symmetric (up to equality and inequality of names), and generative—the main further design choice we care about is whether or not to include a *shift* (Example 2.1.7), which goes strictly beyond what can be achieved with finite permutations as considered e.g. in [Gabbay and Pitts, 2001].

Other notions of permutation may lead to other symmetries, so an interesting topic of future research is to weaken the conditions in Definition 2.1.6.

For instance, if we only allow permutations generated by $f(i) \mapsto f(i + 1)$ and $f(i) \mapsto f(-i)$ then we preserve a notion of ‘distance’ between atoms.³ In a similar vein, we can identify atoms with points in a plane and consider Euclidian transformations. It is not known how much of ‘nominal techniques’ would hold of such examples.

More generally of course, presheaves are a forum within which sets with symmetry structure can be expressed. Indeed, nominal sets can be viewed as a category of presheaves [Gabbay and Pitts, 1999] and a similar presheaf category was considered at the same time [Fiore *et al.*, 1999] (see also the later related *nominal renaming sets* [Gabbay and Hofmann, 2008], which are in some sense half-way in between those two systems).

There is no shortage of research into this kind of structure [Mac Lane and Moerdijk, 1992]. It remains, however, to understand what are the abstract properties that make a set with a group action, or a presheaf, into something ‘nominal’.

³This example modified from an example by Bartek Klin; private communication from Alexander Kurz.

DEFINITION 2.1.9. If $A \subseteq \mathbb{A}$ define the **pointwise** action by

$$\pi \cdot A = \{\pi(a) \mid a \in A\}.$$

DEFINITION 2.1.10. A **permission set** S is a set of the form $\pi \cdot \mathbb{A}^<$. S, T will range over permission sets.

REMARK 2.1.11. Some preliminary comments on permission sets:

- The notion of permission set used in some previous work, for instance in [Dowek *et al.*, 2010, Definition 2.2], was slightly different: a permission set was taken to be a set of the form $(S \setminus A) \cup B$ for finite $A \subseteq \mathbb{A}^<$ and $B \subseteq \mathbb{A}^>$. In the presence of *shift*-permutations we can do this using a permutation, and any $(S \setminus A) \cup B$ can be written as $\pi \cdot S$ for suitable π (cf. Remark 3.6.2 and $\delta_{X-a} \cdot X$ in **(IF)** of Figure 2). Given that the designs are equivalent for the cases we will care about, we chose Definition 2.1.10 because it is somewhat simpler to do mathematics with.
- In the semantics, permission sets are used in the definition of support Definition 2.2.3; if permutations specify *symmetry*, permission sets specify *capturability* and *generativity*.
- In the syntax, permission sets are used to control capture (see Remark 3.4.9); atoms in S are intuitively ‘capturable’ and atoms not in S are intuitively ‘not capturable’.

This is reminiscent of some treatments of syntax where a formal distinction is made between ‘names that exist to be bound’ and ‘names that exist to be free’. See for instance the *freie* and *gebundene Gegenstandsvariable* of Gentzen [Gentzen, 1935, Section 1], and the *individual variables* and *parameters* of Prawitz [Prawitz, 1965, Section 1], or Smullyan [Smullyan, 1968, Chapter IV, Section 1].

However, note that here, for any $a \in S$ and $b \notin S$, also $a \notin (b a) \cdot S$ and $b \in (b a) \cdot S$. That is, for any given atom there is no fixed sense in which it is capturable or not capturable. Each individual permission sets defines its own world of capturable/non-capturable atoms, which differs by a permutation π from what is really a fixed but entirely arbitrary representative $\mathbb{A}^<$.

2.2 Permissive-nominal sets

DEFINITION 2.2.1. A set with a (\mathbb{P} -)permutation action X is a pair $(|X|, \cdot)$ of

- a **carrier set** $|X|$ and
- a group action $(\mathbb{P} \times |X|) \rightarrow |X|$, written infix as $\pi \cdot x$.
So, $id \cdot x = x$ and $\pi \cdot (\pi' \cdot x) = (\pi \circ \pi') \cdot x$ for every π, π' , and $x \in |X|$.

DEFINITION 2.2.2. Given a set with a \mathbb{P} -permutation action X say that $A \subseteq \mathbb{A}$ **supports** $x \in |X|$ when for all permutations $\pi \in \mathbb{P}$, if $\pi(a) = a$ for all $a \in A$ then $\pi \cdot x = x$.

Also, call $A \subseteq \mathbb{A}$ **small** when $A \subseteq S$ for some permission set S .

DEFINITION 2.2.3. A **permissive-nominal set** is a set with a permutation action such that every element has a unique least small supporting set $supp(x)$. We call this the **support** of x .
 X, Y will range over permissive-nominal sets.

Note in Definition 2.2.3 that $supp(x)$ must be *small*, that is, included in some permission set. For instance, $a \in \mathbb{A}$ —with \mathbb{A} having the natural permutation action given by $\pi \cdot x = \pi(x)$ for $x \in \mathbb{A}$ —is supported by $\{a\}$ and $\mathbb{A} \setminus \{a\}$, but the former is small while the latter is not.

REMARK 2.2.4. The difference between a set with a permutation action and a ‘nominal’ set is that nominal sets guarantee for any element, infinitely many atoms fresh for that element.

A mild generalisation of Definition 2.2.3 is possible, where we insist there is a supporting set but do not insist on the existence of a unique *least* such set. It is possible to do a surprising amount just with that; see for instance Fiore, Plotkin and Turi’s paper [Fiore *et al.*, 1999] based on presheaves, and the ‘nominal’ study of infinite permutations and infinite atoms-abstraction in [Gabbay, 2007b].⁴

⁴ If all permutations in \mathbb{P} are finite then we have as a *Technical Lemma* that the existence of *some* supporting set implies the existence of a unique least small supporting set.

In the more general case where infinite permutations are allowed, it is possible to construct a set with a permutation action X and $x \in |X|$ such that x has a supporting set but does not have a unique least small supporting set. See [Gabbay, 2007b, Lemma 21] for an example.

An intermediate state is to admit infinite permutations but restrict the notion of support to consider only the finite ones. We do this in Definitions 3.1 and 3.2 and Remark 3.3 of [Dowek and Gabbay, 2011].

For this paper, none of this will matter directly.

EXAMPLE 2.2.5.

- First-order syntax with variable symbols (modelled as atoms) is a permissive-nominal set, where the permutation action permutes variable symbols directly in syntax so that e.g. $\pi \cdot \lambda a.t = \lambda \pi(a).\pi \cdot t$. A term t is supported by the variable symbols it contains. In this and the following examples the precise nature of the permutation group is not important.
- First-order syntax up to α -equivalence is a permissive-nominal set. The α -equivalence class of t is supported by the free variable symbols of t . A full proof is in [Gabbay, 2011b, Theorem 5.18].
- Traces of π -calculus processes with channel names (atoms) taken from some permission set S , form a permissive-nominal set. A trace is supported by the set of channel names it mentions (which may be infinite in number).
- Given a permissive-nominal nominal set X the set of subsets $U \subseteq |\mathsf{X}|$ with the pointwise action $\pi \cdot U = \{\pi \cdot u \mid u \in U\}$ is a set with a permutation action (this generalises Definition 2.1.9). The subset of this consisting of those subsets $U \subseteq |\mathsf{X}|$ that have a supporting permission set under this action, forms a permissive-nominal set $\text{pow}(\mathsf{X})$.⁵

LEMMA 2.2.6. Suppose X is a permissive-nominal set and $x \in |\mathsf{X}|$. Then $\text{supp}(\pi \cdot x) = \pi \cdot \text{supp}(x)$.

Proof. By a routine calculation using the group action. ■

We conclude with a useful condition for checking whether $a \in \text{supp}(x)$:

COROLLARY 2.2.7. Suppose X is a permissive-nominal set and $x \in |\mathsf{X}|$. Suppose $b \notin \text{supp}(x)$. Then $(b \ a) \cdot x = x$ if and only if $a \notin \text{supp}(x)$.

Proof. Suppose $b \notin \text{supp}(x)$. The right-to-left implication is by the definition of support. For the left-to-right implication, we prove the contrapositive. Suppose $a \in \text{supp}(x)$. By Lemma 2.2.6 $\text{supp}((b \ a) \cdot x) = (b \ a) \cdot \text{supp}(x)$. By our suppositions, $(b \ a) \cdot \text{supp}(x) \neq \text{supp}(x)$. It follows that $(b \ a) \cdot x \neq x$. ■

2.3 Equivariance

DEFINITION 2.3.1. Suppose X and Y are permissive-nominal sets.

⁵Using possibly repeated powersets, arbitrarily complex structures may be constructed. Thus this example guarantees an inexhaustible supply of arbitrarily large and complex structures with which to model . . . almost anything we can imagine. The survey [Gabbay, 2011b] explores this in detail.

Call $x \in |X|$ **equivariant** when $\text{supp}(x) = \emptyset$. (So x is equivariant when $\pi \cdot x = x$ for all π .)

Call $F \in |X| \rightarrow |Y|$ **equivariant** when

$$\forall \pi \in \mathbb{P}. \forall x \in |X|. \pi \cdot (F(x)) = F(\pi \cdot x).$$

F will range over equivariant functions between pairs of permissive-nominal sets.

REMARK 2.3.2. The second notion of equivariance in Definition 2.3.1 is a special case of the first. For details, see e.g. Definition 9.3 and Lemma 9.4 of [Gabbay, 2011b].

LEMMA 2.3.3. If F from $|X|$ to $|Y|$ is equivariant then $\text{supp}(F(x)) \subseteq \text{supp}(x)$ for all $x \in |X|$.

Proof. Suppose $\pi \in \text{fix}(\text{supp}(x))$. By assumption $\pi \cdot F(x) = F(\pi \cdot x)$, and $\pi \cdot x = x$. ■

DEFINITION 2.3.4. Write PmsPrm for the category with objects permissive-nominal sets and arrows equivariant functions between them.

So X, Y range over objects in PmsPrm (Definition 2.2.3).

2.4 Examples of permissive-nominal sets

Throughout the rest of this document we will need the following examples of permissive-nominal sets: atoms, booleans, lists, product, equivariant elements, permutation orbits, and atoms-abstraction. We consider each in turn now.

Atoms, Booleans, infinite lists

DEFINITION 2.4.1 (Atoms). \mathbb{A} the set of all atoms can be considered a permissive-nominal set with a natural permutation action $\pi \cdot a = \pi(a)$. So can each \mathbb{A}_ν .

DEFINITION 2.4.2. If X is a permissive-nominal set say the permutation action is **trivial** when $\pi \cdot x = x$ for all $x \in |X|$ and all $\pi \in \mathbb{P}$.

So X is trivial if and only if all its elements are equivariant.

DEFINITION 2.4.3. Any ‘ordinary’ set can be made into a permissive-nominal set by giving it the trivial permutation action such that $\pi \cdot x = x$ always.

In particular, the set $\mathbb{B} = \{0, 1\}$ can be considered a permissive-nominal set with the trivial permutation action; so can \mathbb{N} and \mathbb{Z} from Definition 2.1.1.

In the cases of \mathbb{A} and $\{0, 1\}$ only, we will be lax about the distinction between the set, and the permissive-nominal set with its natural permutation action.

DEFINITION 2.4.4 (Infinite lists). Define a permissive-nominal set \mathbb{L} by:

- $|\mathbb{L}|$ is the set of infinite sequences of distinct atoms $L = [a_1, a_2, a_3, \dots]$ such that $\text{atms}(L) = \{a_1, a_2, a_3, \dots\}$ is a permission set.
- $\pi \cdot L = [\pi(a_1), \pi(a_2), \pi(a_3), \dots]$.

Product

DEFINITION 2.4.5. Suppose I is an indexing set.⁶ If X_i are permissive-nominal sets for $i \in I$ then define $\prod_i X_i$ by:

- $|\prod_i X_i|$ is the set of I -tuples $(x_i)_i$ such that $\forall i. x_i \in |X_i|$ and there exists a permission set S such that $\forall i. \text{supp}(x_i) \subseteq S$.
- $\pi \cdot (x_i)_i = (\pi \cdot x_i)_i$ (the **elementwise** or **pointwise** action).

Permutation orbits

Permutation orbits will serve us later in Definition 3.3.2 (free unknowns of a term).

If X is a nominal set then $\text{orb}(X)$ is ‘ X quotiented by the permutation action’.

DEFINITION 2.4.6. If X is a permissive-nominal set define $\text{orb}(X)$ by:

- If $x \in X$ then define its **permutation orbit** by $\text{orb}(x) = \{\pi \cdot x \mid \pi \in \mathbb{P}\}$.
- $|\text{orb}(X)| = \{\text{orb}(x) \mid x \in X\}$.
- $\pi \cdot \text{orb}(x) = \text{orb}(x)$.

LEMMA 2.4.7.

- $\text{supp}(\text{orb}(x)) = \emptyset$. That is, $\text{orb}(x)$ is *equivariant* (Definition 2.3.1).
- $\text{orb}(x) = \text{orb}(y)$ if and only if $y = \pi \cdot x$ for some π .

⁶For clarity, note that we intend this set to *not* have a permutation action. Or, we can take this to be a nominal set with the trivial action (Definition 2.4.2). We have in mind \mathbb{N} .

Atoms-abstraction

DEFINITION 2.4.8. Suppose \mathbf{X} is a permissive-nominal set and \mathbb{A}_i is a set of atoms. Define **atoms-abstraction** $[\mathbb{A}_i]\mathbf{X}$ by:

$$\begin{aligned} [a]x &= \{(a, x)\} \cup \{(b, (b a) \cdot x \mid b \in \mathbb{A}_i \setminus \text{supp}(x))\} \\ |[\mathbb{A}_i]\mathbf{X}| &= |[\mathbb{A}_i]\mathbf{X}| = \{[a]x \mid a \in \mathbb{A}_i, x \in |\mathbf{X}|\} \\ \pi \cdot [a]x &= [\pi(a)]\pi \cdot x \end{aligned}$$

LEMMA 2.4.9.

1. $[\mathbb{A}_i]\mathbf{X}$ is a permissive-nominal set.
2. $[a]x = [a]x'$ if and only if $x = x'$, for $a \in \mathbb{A}_i$ and $x \in |\mathbf{X}|$.
3. $[a]x = [a']x'$ if and only if $a' \notin \text{supp}(x)$ and $(a' a) \cdot x = x'$, for $a, a' \in \mathbb{A}_i$ and $x, x' \in |\mathbf{X}|$.

LEMMA 2.4.10. Suppose a function F from $|\mathbb{A} \times \mathbf{X}|$ to $|\mathbf{Y}|$ is equivariant and suppose $\forall a, x. a \notin \text{supp}(F(a, x))$. Then there is a unique equivariant function \hat{F} from $|[\mathbb{A}]\mathbf{X}|$ to $|\mathbf{Y}|$ such that $\forall a, x. \hat{F}([a]x) = F(a, x)$.

Proof. It suffices to show that if $b \notin \text{supp}(x) \cup \text{supp}(F(a, x))$ then $F(b, (b a) \cdot x) = F(a, x)$. By assumption $a \notin \text{supp}(F(a, x))$, so $(b a) \cdot F(a, x) = F(a, x)$. The result follows by equivariance. ■

Here are some basic properties of support:

LEMMA 2.4.11.

- $\text{supp}(a) = \{a\}$.
- $\text{supp}([a]x) = \text{supp}(x) \setminus \{a\}$.
- $\text{supp}((x_1, \dots, x_n)) = \bigcup \{\text{supp}(x_i) \mid 1 \leq i \leq n\}$.

Proof. Proofs are as in [Gabbay and Pitts, 2001] or [Gabbay, 2011b]. ■

The fine design of PmsPrm

Studying PmsPrm (Definition 2.3.4) is not the point of this paper, but for the benefit of the interested reader we will discuss a few aspects of its behaviour.

- If \mathbb{P} consists of finite permutations then PmsPrm is a Boolean topos, directly generalising the category of nominal sets (equivariant FM sets) from [Gabbay and Pitts, 2001; Gabbay, 2011b]. The proof proceeds much as in [Gabbay, 2011b, Corollary 9.11].

- If \mathbb{P} contains infinite permutations then \mathbf{PmsPrm} is cartesian (has products) but is not necessarily cartesian closed (may not have exponentials). This is the *fuzzy support* observed in [Gabbay, 2007b]; see [Gabbay, 2007b, Lemma 21] for the concrete construction. This is reasonable, and it happens because it is possible to construct a function f on $\omega + \omega$ which satisfies $f(0) = 0$ and $f(i+1) = f(i)$ yet which is not a constant function (it returns 0 on finite cardinals and 1 on infinite ones).
- If \mathbb{P} contains infinite permutations but we follow [Dowek *et al.*, 2010] and take the notions of support in Definition 2.2.2 and equivariance to consider only *finite* permutations, then the category we obtain is a Boolean topos but we only have $\text{supp}(x) \cap \text{nontriv}(\pi) = \emptyset$ implies $\pi \cdot x = x$ for finite π . In other words, an element can be fixed by all finite permutations and have empty support, but be shifted by some infinite permutation. Again, this is reasonable; it is no surprise that infinite permutations can ‘observe’ more than finite ones.
- If \mathbb{P} contains infinite permutations and we work with presheaves (in essence, we lose the ‘unique least supporting set’ assumption in Definition 2.2.3), then we get a topos, though it is not Boolean.

In this paper we do not attempt to reason inside \mathbf{PmsPrm} so we do not care whether it is a topos; and we do want the possibility of infinite permutations because these let us write nice algorithms and they give our logics some useful extra expressive power (see e.g. rule **(IF)** of Figure 2, Subsection 3.6, and Remark 9.2.4).

So we admit the possibility of infinite permutations in Definition 2.1.6, we let Definition 2.2.2 consider all $\pi \in \mathbb{P}$ (even infinite ones), and we insist in Definition 2.2.3 that every x have a unique least small supporting set.

In another paper, another set of design decisions might be appropriate.

The reader who does not care about these considerations need not worry; they are all swept under the carpet henceforth.

2.5 Strong support

Strong support exists in nominal terms, though this is implicit. Consider in [Urban *et al.*, 2004] the \approx -suspension rule in Figure 2, and Lemma 2.8. We call this *strong support*, following [Tzevelekos, 2007, Definition 1].

A possibly useful intuition is that an element $x \in \mathbf{X}$ has strong support when the atoms in its support occur *in order*. Formally, the notion of strong support enters into the mathematics in this paper via Proposition 2.5.5, Lemma 3.4.6, and Lemma 7.4.4.

DEFINITION 2.5.1. Suppose \mathbf{X} is a permissive-nominal set. Say $A \subseteq \mathbb{A}$ **strongly supports** $x \in |\mathbf{X}|$ when $\pi \cdot x = x$ if and only if $\forall a \in A. \pi(a) = a$.

If x has some strongly supporting set, call x **strongly supported**.

If every $x \in |\mathbb{X}|$ is strongly supported then call \mathbb{X} **strongly supported**.

LEMMA 2.5.2. $x \in \mathbb{X}$ is strongly supported if and only if

$$\forall \pi, \pi'. (\pi \cdot x = \pi' \cdot x \Leftrightarrow (\forall a \in \text{supp}(x). \pi(a) = \pi'(a))).$$

Proof. From Definition 2.5.1 by considering $\pi^{-1} \circ \pi'$. ■

EXAMPLE 2.5.3.

- The pair $(a, b) \in \mathbb{A} \times \mathbb{A}$ is strongly supported by $\{a, b\}$.
- The unordered pair $\{a, b\} \subseteq \mathbb{A}$ with the pointwise permutation action (Definition 2.1.9) is not strongly supported, because $(a \ b) \cdot \{a, b\} = \{a, b\}$.
- The infinite sequences $[a_1, a_2, a_3, \dots]$ in \mathbb{L} from Definition 2.4.4 are strongly supported.

DEFINITION 2.5.4. Suppose \mathbb{X} and \mathbb{Y} are permissive-nominal sets and \mathbb{X} is strongly-supported. Suppose we are given the following data:

- For each $x \in |\text{orb}(\mathbb{X})|$ a fixed but arbitrary choice of representative $X_x \in x$.
- For each $x \in |\text{orb}(\mathbb{X})|$ a choice of $y_x \in |\mathbb{Y}|$ such that $\text{supp}(y_x) \subseteq \text{supp}(X_x)$.

Define the **equivariant extension** F of this data, which is a function from $|\mathbb{X}|$ to $|\mathbb{Y}|$, by:

$$F(\pi \cdot X_x) = \pi \cdot y_x$$

PROPOSITION 2.5.5.

1. The equivariant extension is well-defined and is an equivariant function from $|\mathbb{X}|$ to $|\mathbb{Y}|$.
2. Every equivariant f is an equivariant extension.

Proof. For the first part, by properties of orbits every $x \in |\mathbb{X}|$ has the form $\pi \cdot X_x$ for some π and for precisely one X_x . This is equivariant by construction, if it is well-defined. So suppose $\pi \cdot X_x = \pi' \cdot X_x$. By assumption X_x is strongly supported so $\pi(a) = \pi'(a)$ for every $a \in \text{supp}(X_x)$. By assumption $\text{supp}(y_x) \subseteq \text{supp}(X_x)$. The result follows by the definition of support.

The second part is easy, noting that $\text{supp}(F(x)) \subseteq \text{supp}(x)$ by Lemma 2.3.3. ■

3 THE SYNTAX OF NOMINAL TERMS

Nominal terms were introduced in [Urban *et al.*, 2004]. The development here is permissive, following [Dowek *et al.*, 2010], but with some additional ingredients: We allow non-equivariant constant symbols and we parameterise over a set of unknowns which is a *strongly-supported* [Tzevelekos, 2007].

Some example permissive-nominal terms are given in Example 3.2.4. See also how nominal terms are used in rewrite theories (Example 5.1.3), algebra (Example 7.1.3), and first-order logic (Subsection 10.1).

3.1 Signatures

DEFINITION 3.1.1. A **sort-signature** is a tuple $(\mathcal{A}, \mathcal{B})$ of **name** and **base** sorts $\mathcal{A} \subseteq \mathbb{N}$ and \mathcal{B} .

ν will range over name sorts; τ will range over base sorts.

A **sort language** is defined by

$$\alpha ::= \nu \mid \tau \mid (\alpha, \dots, \alpha) \mid [\nu]\alpha.$$

EXAMPLE 3.1.2. Example base sorts are: ‘ λ -terms’, ‘formulae’, ‘ π -calculus processes’, and ‘program environments’, ‘functions’, ‘truth-values’, ‘behaviours’, and ‘valuations’.

Base sorts τ are arbitrary; later on when we build denotations they will be populated by elements of arbitrary permissive-nominal sets, see Definition 7.3.1.

Examples of name sorts are ‘variable symbols’, ‘channel names’, ‘thread identifiers’, or ‘memory locations’. Name sorts ν are populated by the atoms we fixed in Definition 2.1.2 and which we used to build permutations and permissive-nominal sets.

REMARK 3.1.3. (α, \dots, α) is a product sort and behaves as expected.

$[\nu]\alpha$ is an *atoms-abstraction sort*; this is different. The behaviour of a term of sort $[\nu]\alpha$ corresponds to ‘ α -abstract a name of sort ν in a term of sort α ’. This is *binding without functions*: we will use atoms-abstractions (Definition 2.4.8) to populate atoms-abstraction sorts.

REMARK 3.1.4. In Definition 3.1.1 we insist that a name sort ν is a natural number; this is not necessary but it makes it easier for us to identify name sorts with sets of atoms from Definition 2.1.2, which are also indexed by numbers.

DEFINITION 3.1.5. A **(nominal) term-signature** over a sort-signature $(\mathcal{A}, \mathcal{B})$ is a tuple $(\mathcal{C}, \mathcal{X}, \mathcal{F}, ar)$ where:

- \mathcal{C} is a permissive-nominal set of **constants**.
- \mathcal{X} is a strongly supported (Definition 2.5.1) permissive-nominal set of **unknowns**.
- \mathcal{F} is a set of equivariant **term-formers**.
- ar assigns
 - to each constant $C \in \mathcal{C}$ a base sort τ which we may write $sort(C)$,
 - to each unknown $X \in \mathcal{X}$ a sort α which we may write $sort(X)$, and
 - to each $f \in \mathcal{F}$ a **term-former arity** $(\alpha)\tau$, where α and τ are in the sort-language determined by $(\mathcal{A}, \mathcal{B})$.

A **(nominal terms) signature** Σ is then a tuple $(\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{X}, \mathcal{F}, ar)$.

The support $supp(X)$ of an unknown $X \in \mathcal{X}$ is intuitively the atoms that may occur free in a term we substitute for that unknown, and $\mathbb{A} \setminus supp(X)$ is the atoms which may not occur free. See Proposition 3.4.3.

NOTATION 3.1.6. We may write $((\alpha_1, \dots, \alpha_n))\tau$ just as $(\alpha_1, \dots, \alpha_n)\tau$.

We write $f : (\alpha)\tau$ for $ar(f) = (\alpha)\tau$ and similarly we write $P : \alpha$ for $ar(P) = \alpha$.

EXAMPLE 3.1.7. Here are some examples of suitable \mathcal{X} .

1. For each sort α and permission set S choose a disjoint countably infinite set of **unknown symbols** $X_\alpha^S, Y_\alpha^S, \dots$. Define $\pi \cdot X_\alpha^S = \{(\pi', X_\alpha^S) \mid \forall \mathbf{a} \in S. \pi(\mathbf{a}) = \pi'(\mathbf{a})\}$. Let $\mathcal{X} = \{\pi \cdot X_\alpha^S \mid \text{all } X_\alpha^S, \pi\}$ with permutation action $\pi \cdot (\pi' \cdot X_\alpha^S) = (\pi \circ \pi') \cdot X_\alpha^S$. Define $ar(\pi \cdot X_\alpha^S) = \alpha$. Essentially this \mathcal{X} was used in [Dowek *et al.*, 2010].
2. For each sort α choose a disjoint countably infinite set of **unknown symbols** $X_\alpha, Y_\alpha, \dots$. Define $\pi \cdot X_\alpha = \{(\pi', X_\alpha) \mid \forall \mathbf{a} \in \mathbb{A}^<. \pi(\mathbf{a}) = \pi'(\mathbf{a})\}$. Let $\mathcal{X} = \{\pi \cdot X_\alpha \mid \text{all } X_\alpha, \pi\}$ with permutation action $\pi \cdot (\pi' \cdot X_\alpha) = (\pi \circ \pi') \cdot X_\alpha$. Define $ar(\pi \cdot X_\alpha) = \alpha$.
3. Take $X = (\alpha, (a_0, a_1, a_2, \dots))$ where $\{a_i \mid i \in \mathbb{N}\}$ is a permission set and let \mathcal{X} be the set of all possible X . Give this the pointwise permutation action $\pi \cdot X = (\alpha, (\pi(a_0), \pi(a_1), \dots))$ and define $ar(X) = \alpha$. This \mathcal{X} is mathematically simple, eliminating the need to take quotients over π .
4. Take $\mathcal{X} = \{0, 1, 2, \dots\}$ with the trivial action $\pi \cdot x = x$, so every $x \in \mathcal{X}$ has $supp(x) = \emptyset$. This example illustrates that our framework is

Vanilla	Permissive
X	Unknown with permission set $\mathbb{A}^<$
$a\#X$	$a \notin \text{supp}(X)$
$a\#r$	$a \notin \text{fa}(r)$
$\nabla \vdash r \rightarrow s$ or $\Delta \vdash r = s$	$r \rightarrow s$ or $r = s$
Extend freshness context	<i>shift</i> -permutation (approx)
Finite support	Small support

Figure 1: Cheat sheet relating ‘vanilla’ nominal terms concepts with ‘permissive’ ones

general enough to include the possibility of unknowns ranging over closed elements (a possibility also mooted in [Fernández and Gabbay, 2007, Subsection 9.2]). By adding further structure to \mathcal{X} , further possibilities can be explored. See also [Gabbay, 2011d] and [Gabbay, 2011c].

In all cases it can be verified that \mathcal{X} is strongly supported.

REMARK 3.1.8. In the case that \mathcal{X} the set of unknowns is as described in parts 1 or 2 of Example 3.1.7, $\text{orb}(X)$ (Definition 2.4.6) may be identified with \mathbb{X}_α^S or \mathbb{X}_α respectively.

The \mathcal{X} of part 1 above may be equivalent to that of \mathcal{X} of part 2, if there exists $\pi \in \mathbb{P}$ bijecting S with $S \setminus \{a\}$ for $a \in S$. This is a *shift*-permutation; see Definition 3.6.1 and subsequent discussion.

For the benefit of the reader familiar with ‘vanilla’ nominal terms as used e.g. in [Urban *et al.*, 2004; Fernández and Gabbay, 2007; Gabbay and Mathijssen, 2009], Figure 1 gives a cheat sheet suggesting how concepts in those papers map to the ‘permissive’ context.

EXAMPLE 3.1.9. A nominal terms signature for the λ -calculus would have one name sort ν , one base sort τ , and term-formers $\text{lam} : ([\nu]\tau)\tau$, $\text{app} : (\tau, \tau)\tau$, and $\text{var} : (\nu)\tau$. The set of constants is empty, and for unknowns we can consider Example 3.1.7.

Usually we assume ‘plenty’ of variable symbols. Definition 3.1.10 makes that formal:

DEFINITION 3.1.10. Say that a signature $\Sigma = (\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{X}, \mathcal{F}, ar)$ has **enough unknowns** when for every sort α in $(\mathcal{A}, \mathcal{B})$ and every permission set S , the set $\{\text{orb}(X) \mid X \in \mathcal{X}, \text{sort}(X) = \alpha, \text{supp}(X) = S\}$ is infinite.

All the examples in Example 3.1.7 have enough unknowns.

3.2 Terms

DEFINITION 3.2.1. For each signature $\Sigma = (\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{X}, \mathcal{F}, ar)$ (Definition 3.1.5) define (**permissive-nominal**) **terms** over Σ by:

$\frac{(a \in \mathbb{A}_\nu, \nu \in \mathcal{A})}{a : \nu}$	$\frac{(sort(C) = \tau)}{C : \tau}$	$\frac{(sort(X) = \alpha)}{X : \alpha}$
$\frac{r : \alpha \quad (ar(f) = (\alpha)\tau)}{f(r) : \tau}$	$\frac{r_1 : \alpha_1 \ \dots \ r_n : \alpha_n}{(r_1, \dots, r_n) : (\alpha_1, \dots, \alpha_n)}$	$\frac{r : \alpha \quad (a \in \mathbb{A}_\nu, \nu \in \mathcal{A})}{[a]r : [\nu]\alpha}$

NOTATION 3.2.2. We may write $f((r_1, \dots, r_n))$ as $f(r_1, \dots, r_n)$.

REMARK 3.2.3. Definition 3.2.1 is *nominal abstract syntax*: terms come pre-quotiented by α -equivalence by construction by virtue of our use of atoms-abstraction $[a]r$. That is, if $a \in \mathbb{A}_\nu$ and $r : \alpha$ then $[a]r$ is not a pair (a, r) , it is a set $\{(a, r)\} \cup \{(b, (b a) \cdot r) \mid b \in \mathbb{A}_\nu \setminus supp(r)\}$ (Definition 2.4.8).

EXAMPLE 3.2.4. Recall the signature for the λ -calculus from Example 3.1.9. In that signature we can form terms as illustrated in the following table, where $a : \nu$ and $X : \tau$:

$a : \nu$	This is not a λ -term.
$var(a) : \tau$	If we want an atom to behave like a λ -term variable, we use <code>var</code> to ‘inject’ it into τ .
$lam([a]var(a)) : \tau$	A typical λ -term as one might meet in the street.
$[a]a : [\nu]\nu$	An atoms-abstraction. Do not be fooled! This is not a λ -term!
$[a]var(a) : [\nu]\tau$	An atoms-abstraction of a λ -term—if we apply <code>lam</code> to this then we get a λ -term.
$lam([a]app(X, var(a))) : \tau$	An open nominal term, representing what might informally be written ‘ $\lambda x.tx$, where t ranges over λ -terms’. Depending on whether $a \notin supp(X)$, we may model a side-condition ‘where x is not free in t ’.

LEMMA 3.2.5. Support and the permutation action are characterised on

terms r as follows:

$$\begin{array}{ll}
\text{supp}(a) = \{a\} & \text{supp}(f(r)) = \text{supp}(r) \\
\text{supp}(C) = \text{supp}(C) & \text{supp}((r_1, \dots, r_n)) = \bigcup_{1 \leq i \leq n} \text{supp}(r_i) \\
\text{supp}(X) = \text{supp}(X) & \text{supp}([a]r) = \text{supp}(r) \setminus \{a\} \\
\pi \cdot a = \pi(a) & \pi \cdot f(r) = f(\pi \cdot r) \\
\pi \cdot C = \pi \cdot C & \pi \cdot (r_1, \dots, r_n) = (\pi \cdot r_1, \dots, \pi \cdot r_n) \\
\pi \cdot X = \pi \cdot X & \pi \cdot [a]r = [\pi(a)]\pi \cdot r
\end{array}$$

Proof. By facts of the permutation action and Lemma 2.4.11. ■

REMARK 3.2.6. Lemma 3.2.5 is important because it verifies that ‘support of r ’ coincides with the usual definition of ‘free variables (atoms) of r ’. This is false of nominal terms; for instance the support of the structure $[a]\mathbf{X}$ as constructed in [Urban *et al.*, 2004] is $\{a\}$, and that of $(a\ b)\mathbf{X}$ is $\{a, b\}$.

What makes Lemma 3.2.5 work is the very specific way in which we constructed our permissive-nominal terms syntax, so that it coincides with the nominal abstract syntax of [Gabbay and Pitts, 2001]. In this sense, what Lemma 3.2.5 expresses is a unification (no pun intended) of the mathematics of [Gabbay and Pitts, 2001] and [Urban *et al.*, 2004].

In Lemma 3.2.5 the clauses for C and X are uninformative, of course. This is because support and the permutation action are determined by the choice of \mathcal{C} and \mathcal{X} . If we assume further internal structure of $C \in \mathcal{C}$ or $X \in \mathcal{X}$ then we can be more specific: for instance in the case of part 1 of Example 3.1.7, $fa(\pi \cdot X^S) = \{\pi(a) \mid a \in S\}$.

Because of Lemma 3.2.5, we are entitled to use the following notation:

NOTATION 3.2.7. In the case of syntax r , we may write $fa(r)$ for $\text{supp}(r)$ and call this the **free atoms** of r .

LEMMA 3.2.8. $fa(\pi \cdot r) = \pi \cdot fa(r)$.

Proof. By a routine induction on r . ■

LEMMA 3.2.9. If $\pi(a) = \pi'(a)$ for all $a \in fa(r)$ then $\pi \cdot r = \pi' \cdot r$. The reverse implication also holds, provided that all constant symbols in r are strongly supported.

Proof. The first part is immediate from Notation 3.2.7 and the definition of support in Definition 2.2.2.

The reverse implication is by a nominal abstract syntax induction on r . For the case of $r = [a]r'$ we α -convert a to be fresh so that $a \notin \text{nontriv}(\pi) \cup \text{nontriv}(\pi')$; by assumption 3 in Definition 2.1.6 we can do this. We then

use part 2 of Lemma 2.4.9. The case of $r = X \in \mathcal{X}$ uses the assumption of strong support in Definition 3.1.5.⁷ ■

3.3 Free unknowns of a term

REMARK 3.3.1. Defining a notion of ‘the free unknowns of r ’ is not entirely evident.

Consider for example $[a]X$ where $a \in \text{supp}(X)$. If ‘ X appears in $[a]X$ ’ is true then so is ‘ $(b\ a) \cdot X$ appears in $[a]X$ ’ for any $b \notin \text{supp}(X)$, since $[a]X = [b](b\ a) \cdot X$. We deal with this in Definition 3.3.2 using *permutation orbits* from Definition 2.4.6; we simply quotient out all permutations. We take a more refined look at this later in Remark 3.7.1.

DEFINITION 3.3.2. Define (**free**) **unknowns** $fv(r)$ by:

$fv(a) = \emptyset$	$fv(f(r)) = fv(r)$
$fv(C) = \emptyset$	$fv((r_1, \dots, r_n)) = \bigcup_i fv(r_i)$
$fv(X) = \{orb(X)\}$	$fv([a]r) = fv(r)$

By abuse of notation we write $X \in fv(r)$ for $orb(X) \in fv(r)$ and $X \notin fv(r)$ for $orb(X) \notin fv(r)$, and so forth.

LEMMA 3.3.3. $fv(r)$ is well-defined.

Proof. Using Lemmas 2.4.7 and 2.4.10. ■

NOTATION 3.3.4. Call a term r **ground** when $fv(r) = \emptyset$. Otherwise, call r **open**.

3.4 Substitutions

REMARK 3.4.1. Substitutions are of course how unknowns ‘stand for’ terms. Somewhat later we will develop a denotational theory for nominal terms, and so valuations for unknowns will appear, in Definition 7.3.3. Between now and then, substitutions are king.

The permissive-nominal framework we work with allows us an elegant definition:

⁷Details of how induction on nominal abstract syntax allows us to α -convert and make freshness assumptions, are the topic of [Gabbay, 2011b]. A less fancy proof of both implications by a standard induction—so not this new-fangled nominal nonsense—on terms not quotiented by α -equivalence, is in Appendix A of [Dowek *et al.*, 2010], proof of Lemma 4.15 on page 50. We leave it to the reader to judge which is the nicer proof.

DEFINITION 3.4.2. Suppose Σ is a signature. A **substitution** θ in Σ is an equivariant function from \mathcal{X} to terms in Σ such that $\text{sort}(\theta(X)) = \text{sort}(X)$ always.
 θ will range over substitutions.
 Write id for the **identity** substitution mapping X to X always. It will always be clear whether id means the identity substitution or permutation.

The reader familiar with nominal terms will expect a ‘freshness’ condition on substitutions corresponding to ‘ $\nabla' \vdash \nabla\theta$ ’, as in for example Equation (11) or Lemma 2.14 of [Urban *et al.*, 2004], or ‘ $\text{fa}(\theta(X)) \subseteq \text{supp}(X)$ ’ as in Definition 3.1 of [Dowek *et al.*, 2010]. This follows immediately from equivariance:

PROPOSITION 3.4.3. If θ is a substitution then $\forall X \in \mathcal{X}. \text{fa}(\theta(X)) \subseteq \text{supp}(X)$.

Proof. Direct from Lemma 2.3.3. ■

Putting Propositions 3.4.3 and 2.5.5 together with a concrete \mathcal{X} recovers the notion of substitution used in [Dowek *et al.*, 2010]:

LEMMA 3.4.4. If \mathcal{X} is equal to example 1 of Example 3.1.7 then the construction in Definition 2.5.4 describes a 1-1 correspondence between substitutions and maps from unknowns X_α^S to terms $t : \alpha$ such that $\text{fa}(t) \subseteq S$.

DEFINITION 3.4.5. Suppose $\text{fa}(t) \subseteq \text{supp}(X)$ and $\text{sort}(t) = \text{sort}(X)$. Write $[X:=t]$ for the **atomic substitution** equivariantly extending the assignment $X \mapsto t$, so that

$$\begin{aligned} [X:=t](\pi \cdot X) &= \pi \cdot t & \text{and} \\ [X:=t](Y) &= Y & \text{for all other } Y. \end{aligned}$$

By Proposition 2.5.5 we have:

LEMMA 3.4.6. Definition 3.4.5 is well-defined. That is, if $\pi \cdot X = \pi' \cdot X$ then $\pi \cdot t = \pi' \cdot t$.

REMARK 3.4.7. The ‘moderated unknown’ $\pi \cdot X$ in Definition 3.4.5 is an artefact of our writing $[X:=t]$ instead of a mathematically equal $[\pi \cdot X := \pi \cdot t]$ for some other π .

Since θ is equivariant its behaviour on $\pi \cdot X$ is already determined by its behaviour on X and so we could unambiguously specify $[X:=t]$ succinctly as $[X:=t](X) = t$ and $[X:=t](Y) = Y$.

DEFINITION 3.4.8. Define a **substitution action** on terms by:

$ \begin{array}{ll} a\theta = a & f(r)\theta = f(r\theta) \\ C\theta = C & (r_1, \dots, r_n)\theta = (r_1\theta, \dots, r_n\theta) \\ X\theta = \theta(X) & ([a]r)\theta = [a](r\theta) \end{array} $

Note that $X\theta$ refers to θ acting on X as a term whereas $\theta(X)$ refers the value of the function θ at X . The substitution action is well-defined by Lemmas 2.4.10 and 2.4.11.

REMARK 3.4.9. Famously, the nominal terms substitution is capturing [Urban *et al.*, 2004, Definition 2.13]. We spell out how this works in our permissive-nominal context: Suppose $\text{supp}(X)$ is equal to a permission set S and $a \in S$ and $b \notin S$ (where we assume appropriate sorts). Then:

- $([a]X)[X:=a] = [a]a$. The a in the substitution $[X:=a]$ has been captured by the $[a]X$.
- $([b]X)[X:=a] = [b]a$.
- It is impossible to even ask what $([b]X)[X:=b]$ is equal to because $[X:=b]$ is not even a substitution, since $b \notin S$. So $b \notin S$ cannot be captured by a substitution $[X:=b]$, because that substitution does not exist. This is no *ad hoc* restriction: by Proposition 3.4.3 it *cannot* exist.
- Also, $[b](b a) \cdot X = [a]X$. By construction in Definition 3.4.5

$$([b](b a) \cdot X)[X:=a] = [b](b a) \cdot a = [b]b = [a]a.$$

Also $[X:=a] = [(b a) \cdot X := b]$ and $([b](b a) \cdot X)[(b a) \cdot X := b] = [b]b$.

That is, the choice of representative of $[a]X$ and $[X:=a]$ does not matter for capture to occur.

It is interesting to note that in our setting, $[X:=a]$ is *equivariant* and that $a \notin \text{supp}([a]X)$. If a is fresh for both $[X:=a]$ and $[a]X$, how can it be captured?

What allows a to get captured is the *strong support property* of X . Because X is strongly supported, we can think of it as ‘containing’ a list of its supporting atoms in some order, so that the a in $[X:=a]$ is bound by $\text{supp}(X)$ but in being bound it points to a ‘position’ in X .

Viewed from this interesting perspective, the nominal substitution action is not capturing at all: it is simply a compact way to present an ‘infinite raising’ (terminology from higher-order logic), or a de Bruijn index.

LEMMA 3.4.10. $\pi \cdot (r\theta) = (\pi \cdot r)\theta$.

Proof. By a routine induction on r using equivariance. ■

LEMMA 3.4.11. $fa(r\theta) \subseteq fa(r)$.

Proof. From Lemmas 2.3.3 and 3.4.10. ■

LEMMA 3.4.12. $r\theta = r\theta'$ if and only if $\forall X \in fv(r). \theta(X) = \theta'(X)$.

Proof. By a routine induction on r . We consider two cases:

- *The case $[a]r$.* Suppose $\theta(X) = \theta'(X)$ for every $X \in fv([a]r)$. $fv([a]r) = fv(r)$ so by inductive hypothesis $r\theta = r\theta'$. The result follows from the definitions.

The reverse implication is similar.

- *The case X .* Suppose $\theta(\pi \cdot X) = \theta'(\pi \cdot X)$ for all π . Then taking $\pi = id$ we have $X\theta = \theta(X) = \theta'(X) = X\theta'$.

Conversely if $X\theta = X\theta'$ then using equivariance (Definition 3.4.2) $\theta(\pi \cdot X) = \theta'(\pi \cdot X)$ for all π . ■

REMARK 3.4.13. Recall from Definition 3.3.2 that we write $X \in fv(r)$ for $orb(X) \in fv(r)$. It might seem that the condition $\forall X \in fv(r). \theta(X) = \theta'(X)$ in Lemma 3.4.12 would require checking $\theta(X) = \theta'(X)$ for infinitely many X provided that $fv(r) \neq \emptyset$. In fact, this is not the case: by equivariance of θ , we only need to check equality for one representative X of each permutation orbit: $X \in orb(X) \in fv(r)$.

3.5 Composition and invertibility of substitutions

DEFINITION 3.5.1. Define **composition** of substitutions $\theta_1 \circ \theta_2$ by

$$(\theta_1 \circ \theta_2)(X) = (\theta_1(X))\theta_2.$$

LEMMA 3.5.2. $(r\theta)\theta' = r(\theta \circ \theta')$.

Proof. By induction on r . ■

DEFINITION 3.5.3. Call θ **invertible** when there exists θ^{-1} such that $\theta \circ \theta^{-1} = \theta^{-1} \circ \theta = id$.

LEMMA 3.5.4. θ is invertible if and only if θ is a bijection on \mathcal{X} the set of all unknowns. Furthermore, if θ is invertible then $supp(\theta(X)) = supp(X)$ always.

Proof. Substitution cannot make syntax smaller, or (by Lemma 3.4.11) make free atoms larger. ■

So an invertible θ must biject unknowns of a particular sort and permission set with other unknowns of that same sort and permission set. So, like atoms, we can rename unknowns to ‘be fresh’ (provided we have given ourselves enough of them). Invertible substitutions will be useful later, and they are also one manifestation of a more general framework of *two-level nominal sets* [Gabbay, 2011d].

3.6 shift-permutations

The reader may be familiar with nominal freshness conditions $a\#X$ from [Urban *et al.*, 2004]. In that paper, $a\#X$ indicated that X should be substituted only for terms for which a is fresh.

In [Urban *et al.*, 2004; Fernández and Gabbay, 2007], we might have to extend a freshness context in order to give ourselves more fresh atoms. This is what rules like **(Fr)** from [Gabbay and Mathijssen, 2008c, Figure 2] or **(fr)** from [Gabbay and Mathijssen, 2009, Figure 2] do; see also [Fernández and Gabbay, 2010] where the issue of extending nominal freshness contexts is made very explicit.

In principle, permission sets guarantee an infinite supply of fresh atoms, so the problem of extending a freshness context should not arise. But this may rely on oracular knowledge of what the permission set should be, which we might prefer not to assume. The choice of nominal permutation group \mathbb{P} gives us the power to implicitly parameterise over this decision.

Suppose we have some X such that $a \in \text{supp}(X)$ and we perhaps we are solving a unification problem and the information that a should be fresh for X has just been revealed by an algorithm; so we want to remove a from the permission set of X . This arises in the unification algorithm of Section 4.

Suppose alternatively we would like to make the permission set *larger*, e.g. if we know $\forall X.\phi$ and want to deduce $\phi[X:=t]$ where $fa(t) \not\subseteq \text{supp}(X)$, or we have a rewrite rule $X \rightarrow X$ and want to deduce $t \rightarrow t$ where again $fa(t) \not\subseteq \text{supp}(X)$. This arises in the nominal rewriting, algebra and permissive-nominal logic which we construct later.

This is where *shift*-permutations can help.

DEFINITION 3.6.1. Call a permutation $\delta \in \mathbb{P}$ a **shift-permutation** when there exists a permission set S and atom $a \in S$ such that $S \setminus \{a\} = \delta \cdot S$.
 Say that a nominal permutation group \mathbb{P} has **shift-permutations** when for every permission set S and atom $a \in \mathbb{A}$ there exists a permutation $\pi \in \mathbb{P}$ such that $\pi \cdot S = S \setminus \{a\}$.

REMARK 3.6.2. Another way to read Definition 3.6.1 is that \mathbb{P} has *shift*-permutations when, if S is a permission set and A is finite, then $S \setminus A$ and

$S \cup A$ are permission sets. Stronger versions allowing infinite A are certainly imaginable.

EXAMPLE 3.6.3. The nominal permutation group in part 2 of Example 2.1.7 has *shift*-permutations.

δ_i bijects $\mathbb{A}_i^<$ with $\mathbb{A}_i^< \setminus \{f(0)\}$. Using swappings we can now generate a π to biject any permission set S with $S \setminus \{a\}$ for $a \in S$. We give the concrete constructions below, culminating with Lemma 3.6.9.

For the rest of this subsection we work concretely with the nominal permutation group from part 2 of Example 2.1.7; the reader only interested in the high-level picture can skip this. Recall the bijections f_i from integers to atoms from part 2 of Example 2.1.7. For simplicity drop the subscript i and consider just one set of atoms.


NOTATION 3.6.4. By abuse of notation write 0 for the atom $f(0)$.

DEFINITION 3.6.5. Suppose that A is a finite set of atoms $a \in \mathbb{A} \setminus A$. Define δ_A^a inductively on the size of A as follows:

$$\begin{aligned} \delta_{\emptyset}^a &= (a \ 0) \circ \delta \circ (a \ 0) \\ \delta_{b,A}^a &= (b \ \delta_A^a(b)) \circ \delta_A^a \quad (b \notin A) \end{aligned}$$

We may write δ_{\emptyset}^a as δ^a .

EXAMPLE 3.6.6. We illustrate δ^a and $\delta_{\{b,c\}}^a$ where $a = f(3)$, $b = f(-1)$, and $c = f(-2)$.



LEMMA 3.6.7. Suppose $A \subseteq \mathbb{A}$ is finite and $a \in \mathbb{A} \setminus A$. Then δ_A^a is well-defined (does not choose in which order we take the atoms in A in Definition 3.6.5) and:

1. δ^a bijects $(a \ 0) \cdot \mathbb{A}^<$ with $((a \ 0) \circ \delta) \cdot \mathbb{A}^<$.
2. δ_A^a bijects $((a \ 0) \cdot \mathbb{A}^<) \setminus A$ with $((a \ 0) \circ \delta) \cdot \mathbb{A}^< \setminus A$, and fixes every atom in A .

DEFINITION 3.6.8. Given an unknown X and an atom $a \in \text{supp}(X)$ define

$$\begin{aligned} \delta_{X-a} &= \delta_{((\mathbb{A}^< \setminus \text{supp}(X)) \cup (\text{supp}(X) \setminus \mathbb{A}^<)) \setminus \{a\}}^a \\ X-a &= \delta_{X-a} \cdot X. \end{aligned}$$

Here $(\mathbb{A}^< \setminus \text{supp}(X)) \cup (\text{supp}(X) \setminus \mathbb{A}^<)$ is the *exclusive or* of $\mathbb{A}^<$ and $\text{supp}(X)$; the atoms in precisely one of these sets (this is a measure of their difference).

If D is a finite list of distinct atoms d_1, \dots, d_n then we may write

$$X-D = (\dots(X-d_1)\dots-d_n).$$

The details of δ_{X-a} and δ_{X-D} are only interesting insofar as they give us Lemma 3.6.9. Many permutations have this property but that does not matter; we only need that one exists:

LEMMA 3.6.9. δ_{X-a} bijects $\text{supp}(X)$ with $\text{supp}(X)\setminus\{a\}$ and so if D is a list d_1, \dots, d_n then $\text{supp}(X-D) = \text{supp}(X)\setminus\{d_1, \dots, d_n\}$.

REMARK 3.6.10. The reader may be familiar with the de Bruijn *shift* function \uparrow [Abadi *et al.*, 1991, Subsection 2.2]. This maps \mathbb{N} to $\mathbb{N}\setminus\{0\}$ by mapping $j \in \mathbb{N}$ to $j+1 \in \mathbb{N}$, and in doing so it ‘creates a fresh number’ 0. The reader familiar with presheaf techniques may know of a functor δ and arrow up , which work the same way, as exemplified in [Fiore *et al.*, 1999, Section 1].

δ_i from part 2 of Example 2.1.7 is in the same spirit. It shifts ‘down’ instead of ‘up’, but δ_i^{-1} shifts ‘up’.

Note that δ is *invertible* (\uparrow and up are not). This is consistent with the general preference of nominal techniques for using permutations where possible.

3.7 Occurrences

REMARK 3.7.1. As discussed in Remark 3.3.1 we have to be careful if we wish to say ‘ X appears in r ’; this might not quite mean what we think it does.

For example if ‘ X appears in $[a]X$ ’ where $a \in \text{supp}(X)$ then also ‘ $(b a) \cdot X$ appears in $[a]X$ ’ for any $b \notin \text{supp}(X)$. We dealt with this in Definition 3.3.2 by quotienting out all permutations.

But this is a little drastic. For instance, ‘ $(b a) \cdot X$ appears in $[a]X$ ’ is not true for $b \in \text{supp}(X)$; it is not the case that if ‘ X appears in r ’ then ‘ $\pi \cdot X$ appears in r ’ for any π .

We did not need to quotient out *all* permutations—only some of them—and so returning $\text{orb}(X)$ in Definition 3.3.2 throws out more information than necessary.

Definitions 3.7.2 and 3.7.3 develop a more refined notion of occurrence, based on an intuition of ‘ X appears in r under a list of abstractions D ’. This will be useful later.

DEFINITION 3.7.2. D will range over finite lists of distinct atoms. A **(level 2) occurrence** is a term of the form $[D]X$ where $\square X$ is X and $[a, D]X$ is $[a][D]X$.

DEFINITION 3.7.3. Define the **occurrences in** r inductively by:

$\begin{aligned} occ(a) &= \emptyset \\ occ(C) &= \emptyset \\ occ(X) &= X \end{aligned}$	$\begin{aligned} occ(\mathbf{f}(r)) &= occ(r) \\ occ((r_1, \dots, r_n)) &= \bigcup occ(r_i) \\ occ([a]r) &= \{[a]x \mid x \in occ(r)\} \end{aligned}$
---	---

EXAMPLE 3.7.4.

- X occurs in X .
- $[a]X$ occurs in $[a]X$ and also in $[a](X, Y)$; so does $[a]Y$. X does not occur in $[a]X$ or $[a](X, Y)$.
- $[a][b]X$ and $[a][a]X$ occur in $[a]([b]X, [a]X)$.

We write occurrences as $[D]X$ for D a finite list of distinct atoms. Note that $[a][a]X$ is an occurrence since it is equal to $[a][b](b a) \cdot X$ where $b \notin \text{supp}(X)$. This is an equality, not an equivalence imposed on terms after they are constructed, because of our use of atoms-abstraction (Definition 2.4.8) in syntax (Definition 3.2.1).

Part II

Rewrites, equations, and algebras

4 UNIFICATION

We want to write rewrite rules and equality axioms using nominal terms. In order to do this, we have to unify nominal terms (answer the question: “given r and s what substitutions θ make them equal?”). Unification makes variables ‘come alive’ and represent unknown terms.

Therefore, we now create a nominal unification algorithm. One notable property of nominal unification is that it has most general (principal) unifiers Theorem 4.4.6. Contrast this with higher-order unification, which does not [Dowek, 2001, Section 4]. This is one reason we say that the nominal approach to names and binding has a ‘first-order’ flavour.

The algorithm we use follows the spirit of [Urban *et al.*, 2004] but the design is different. In [Urban *et al.*, 2004] a solution to $[a]X \stackrel{?}{=} [b]Y$ would be $(b\#X, [Y := (b\ a)\cdot X])$; that is, the unification algorithm returns a pair of some freshness side-conditions and some equalities.⁸

Here, solutions are equalities only, without freshness conditions. The extra power resides in the notion of an *shift*-permutation (Definition 3.6.1).

A solution to $[a]X = [b]Y$ where $b \in \text{supp}(X) = \text{supp}(Y)$ would be $[X := \delta'\cdot X, Y := ((b\ a) \circ \delta')\cdot X]$ where δ' bijects $\text{supp}(X)$ with $\text{supp}(X) \setminus \{b\}$ (and by this bijection ‘internally freshens’ X with respect to b).

In another design [Dowek *et al.*, 2010, Section 5] we use permission sets and fresh unknowns; a solution to $[a]X = [b]Y$ where $b \in \text{supp}(X) = \text{supp}(Y)$ is $[X := Z, Y := (b\ a)\cdot Z]$ where $\text{supp}(Z) = \text{supp}(X) \setminus \{b\}$. Generating Z fresh requires us to solve problems in a context of ‘known unknowns’ \mathcal{V} . This introduces a notion of state and sequentiality into the algorithm of [Dowek *et al.*, 2010] which we avoid here.

Nothing forces us to feed the unification algorithm syntax with *shift*-permutations, even if the solutions it returns might mention them; similarly in [Urban *et al.*, 2004] we may obtain a solution with freshness side-conditions to a unification problem with only equalities. So use of *shift*-

⁸We write typewriter font to avoid confusion between the symbols used in [Urban *et al.*, 2004] (which have no support) and the elements $X \in \mathcal{X}$ used in this paper (which do have support). To see how to travel between these two worlds see part 2 of Example 3.1.7, or [Dowek *et al.*, 2010].

permutation in Definition 4.0.5 should not be read as a commitment to using them everywhere (though we do note empirically that *shift* seems to be useful elsewhere too).

The main definition of this section is Definition 4.1.7. The main result is Theorem 4.4.6.

DEFINITION 4.0.5. Throughout this Section we fix some signature Σ and we work with syntax over Σ . We assume a nominal permutation group \mathbb{P} with *shift*-permutations and a set of unknowns \mathcal{X} such that every unknown is supported by a permission set (see e.g. part 2 of Example 3.1.7).

4.1 The unification algorithm

DEFINITION 4.1.1. A **(unification) equality** is a unordered pair $r \stackrel{?}{=} s$ (so $r \stackrel{?}{=} s$ is identical to $s \stackrel{?}{=} r$) such that:

1. $sort(r) = sort(s)$.
2. If $[D]X$ and $[D']\pi \cdot X$ are both in $occ(r) \cup occ(s)$ then π is finite.
So we exclude an equality like $X \stackrel{?}{=} \delta \cdot X$, where δ is a shift permutation and $nontriv(\delta) \cap supp(X)$ is not finite.

A **(unification) freshness** is an ordered pair $a\#?r$.

Let ef range over equalities or freshnesses and define $ef\theta$ by:

- $(r \stackrel{?}{=} s)\theta = (r\theta \stackrel{?}{=} s\theta)$.
- $(a\#?r)\theta = (a\#?(r\theta))$.

A **nominal unification problem** Pr is a finite list ef_1, \dots, ef_n .

We (ab)use standard sets notation and write $ef \in Pr$ as shorthand for ‘ ef appears in the list Pr ’.

REMARK 4.1.2. Condition 2 in Definition 4.1.1 protects $(\stackrel{?}{=} \mathbf{X})$ in Figure 2 from an ‘infinite freshness explosion’, if $nontriv(\pi) \cap supp(X)$ is not finite. This condition exists implicitly in [Urban *et al.*, 2004], in the sense that all permutations there are finite. However, condition 2 is not only computation-ally motivated. Given constants C and D with $supp(C) = \emptyset = supp(D)$, $X \stackrel{?}{=} \delta \cdot X$ may have solutions C and D but have no principal solution. We discuss the implications of this condition to nominal rewriting, at the end of Section 6.

DEFINITION 4.1.3. If $Pr = ef_1, \dots, ef_n$ is a problem then define $Pr\theta$ by:

$$Pr\theta = ef_1\theta, \dots, ef_n\theta$$

Say θ **solves** Pr and call θ a **solution** to Pr when

$ \begin{aligned} r\theta = s\theta & \quad \text{for every } r \stackrel{?}{=} s \in Pr, \quad \text{and} \\ a \notin fa(r\theta) & \quad \text{for every } a\#?r \in Pr. \end{aligned} $
--

$(\stackrel{?}{=}a)$	$a \stackrel{?}{=} a, Pr$	\implies	Pr
$(\stackrel{?}{=}C)$	$C \stackrel{?}{=} C, Pr$	\implies	Pr
$(\stackrel{?}{=}f)$	$f(r) \stackrel{?}{=} f(s), Pr$	\implies	$r \stackrel{?}{=} s, Pr$
$(\stackrel{?}{=}())$	$(r_1, \dots, r_n) \stackrel{?}{=} (s_1, \dots, s_n), Pr$	\implies	$r_1 \stackrel{?}{=} s_1, \dots, r_n \stackrel{?}{=} s_n, Pr$
$(\stackrel{?}{=}[])$	$[a]r \stackrel{?}{=} [a]s, Pr$	\implies	$r \stackrel{?}{=} s, Pr$
$(\stackrel{?}{=}X)$	$X \stackrel{?}{=} \pi \cdot X, Pr$	\implies	$a_1 \#_{?} X, \dots, a_n \#_{?} X, Pr$ $(\{a_1, \dots, a_n\} = \text{nontriv}(\pi) \cap \text{supp}(X))$
(F)	$r \stackrel{?}{=} X, Pr$	\implies	$a \#_{?} r, r \stackrel{?}{=} X, Pr$ $(a \in \text{fa}(r) \setminus \text{supp}(X))$
(F#)	$a \#_{?} r, Pr$	\implies	Pr $(a \notin \text{fa}(r))$
(Ff)	$a \#_{?} f(r), Pr$	\implies	$a \#_{?} r, Pr$
(F())	$a \#_{?} (r_1, \dots, r_n), Pr$	\implies	$a \#_{?} r_1, \dots, a \#_{?} r_n, Pr$
(F[])	$a \#_{?} [b]r, Pr$	\implies	$a \#_{?} r, Pr$
(IE)	$r \stackrel{?}{=} X, Pr$	$\xRightarrow{[X:=r]}$	$Pr[X:=r]$ $(X \notin \text{fv}(r), \text{fa}(r) \subseteq \text{supp}(X))$
(IF)	$a \#_{?} X, Pr$	$\xRightarrow{[X:=\delta_{X-a} \cdot X]}$	$Pr[X:=\delta_{X-a} \cdot X]$

Figure 2: Simplification rules for problems

Write $Sol(Pr)$ for the set of solutions to Pr and call Pr **solvable** when $Sol(Pr)$ is non-empty.

Recall the definition of $\theta \circ \theta'$ from Definition 3.5.1.

LEMMA 4.1.4. $\theta \circ \theta' \in Sol(Pr)$ if and only if $\theta' \in Sol(Pr\theta)$.

Proof. By unpacking Definition 4.1.3 and using Lemma 3.5.2. ■

DEFINITION 4.1.5. Define a **simplification** rewrite relation $Pr \implies Pr'$ on unification problems by the rules in Figure 2.

We call rules **(IF)** and **(IE)** **instantiating rules**. We call all the other rules **non-instantiating rules**.

In **(IF)** δ_{X-a} is some permutation bijecting $\text{supp}(X)$ with $\text{supp}(X) \setminus \{a\}$. We can do this because we assumed *shift*-permutations in Definition 4.0.5.⁹

Write \implies^* for the transitive and reflexive closure of \implies .

REMARK 4.1.6. Compare Figure 2 with Figure 3 of [Urban *et al.*, 2004]. Note of $(\stackrel{?}{=}[])$ that we do not consider the case $[a]r \stackrel{?}{=} [b]s$. This is because

⁹The specific choice does not matter. Intuitively this is because permutations are invertible so any one choice and be undone and redone at will. A more formal statement of this is Theorem 4.3.6. For an example of a *shift*-permutation concretely constructed, see Definition 3.6.8.

This algorithm generates *shifts* just like in [Urban *et al.*, 2004] we generated freshness conditions, and for the same reason.

α -equivalence is handled automatically by nominal abstract syntax, specifically by Definition 2.4.8. So α -renaming is pushed into the background (just as is usually the case for first-order syntax) and these rules are somewhat higher-level than those of [Urban *et al.*, 2004].

We also do not require a rule $a\#?[a]r, Pr \Longrightarrow Pr$ because the abstracted atom in $[a]r$ is α -convertible; more formally, $[a]r = [b](b a)\cdot r$ for some/any fresh b (so $b \notin fa(r)$).

Finally, in $(\stackrel{?}{=}X)$ we do not need to write $\pi\cdot X \stackrel{?}{=} \pi'\cdot X$ (though we could) because unknowns are just a strongly-supported nominal set. We know that $\text{nontriv}(\pi) \cap \text{supp}(X)$ is finite by a routine argument based on condition 2 of Definition 4.1.1. It is not hard to check that the instantiating rules **(IF)** and **(IE)** do indeed preserve these conditions—**(IF)** involves a *shift* permutation, but in a manner that is applied uniformly to the whole problem.

DEFINITION 4.1.7. If Pr is a problem, define a **unification algorithm** by:

1. Rewrite Pr using the rules of Definition 4.1.5 where possible, with top-down precedence (so apply $(\stackrel{?}{=}a)$ before $(\stackrel{?}{=}f)$, and so on).
2. If we reduce to \emptyset then we succeed and return θ where θ is the composition of all the substitutions labelling rewrites (we take $\theta = id$ if there are none). Otherwise, we fail.

REMARK 4.1.8. Note in Definition 4.1.7 that we apply each rule to the head of the list Pr . This is to prevent ‘unfair’ looping, e.g. repeatedly applying **(F)** to some equality $r \stackrel{?}{=} X$ wherever it appears in Pr .

Note also that the rule **(F#)** is equivalent—in the presence of the other rules—to three rules as follows:

$$\begin{array}{lll}
 \mathbf{(Fa)} & a\#?b, Pr \Longrightarrow Pr & \\
 \mathbf{(FC)} & a\#?C, Pr \Longrightarrow Pr & (a \notin \text{supp}(C)) \\
 \mathbf{(FX)} & a\#?X, Pr \Longrightarrow Pr & (a \notin \text{supp}(X))
 \end{array}$$

PROPOSITION 4.1.9. The algorithm of Definition 4.1.7 always terminates.

Proof. It is not hard to generate an inductive quantity which is reduced by the reductions in Figure 2. ■

4.2 Examples of the algorithm

We assume the permutation group from part 2 of Example 2.1.7 and we recall the definition of X - D from Definition 3.6.8.

Example one (succeeds).

Suppose $a, c \in \mathbb{A}^<$ and $d \notin \mathbb{A}^<$. Take $\text{supp}(X) = \mathbb{A}^<$ and suppose a term-former g . We apply the algorithm to $\{g([a]X, [a]a) \stackrel{?}{=} g([d]c, [d]d)\}$:

$$\begin{array}{ll}
 g([a]X, [a]a) \stackrel{?}{=} g([d]c, [d]d) & \Longrightarrow \quad (\stackrel{?}{=}g), (\stackrel{?}{=}()) \\
 [a]X \stackrel{?}{=} [d]c, [a]a \stackrel{?}{=} [d]d & \Longrightarrow \quad (\stackrel{?}{=}[]), [a]X = [d](d \ a) \cdot X \\
 (d \ a) \cdot X \stackrel{?}{=} c, [a]a \stackrel{?}{=} [d]d & \xRightarrow{[X:=c]} \quad \text{(IE)} \\
 [a]a \stackrel{?}{=} [d]d & \Longrightarrow \quad (\stackrel{?}{=}[]), [a]a = [d]d \\
 d \stackrel{?}{=} d & \Longrightarrow \quad (\stackrel{?}{=}a) \\
 \emptyset & \text{Success, with } [X:=c]
 \end{array}$$

Example two (succeeds).

Suppose $a, c \in \mathbb{A}^<$ and $b, d \notin \mathbb{A}^<$. Take $\text{supp}(X) = \mathbb{A}^< \cup \{b, d\}$, $\text{supp}(Y) = \mathbb{A}^< \cup \{f\}$, and $\text{supp}(Z) = \mathbb{A}^<$. Suppose a term-former f .

We apply the algorithm to $\{f([a]b, Z, X) \stackrel{?}{=} f([d]b, [a]a, Y)\}$:

$$\begin{array}{ll}
 f([a]b, Z, X) \stackrel{?}{=} f([d]b, [a]a, Y) & \Longrightarrow \quad (\stackrel{?}{=}f), (\stackrel{?}{=}()) \\
 [a]b \stackrel{?}{=} [d]b, Z \stackrel{?}{=} [a]a, X \stackrel{?}{=} Y & \Longrightarrow \quad (\stackrel{?}{=}[]), [a]b = [d]b \\
 b \stackrel{?}{=} b, Z \stackrel{?}{=} [a]a, X \stackrel{?}{=} Y & \Longrightarrow \quad (\stackrel{?}{=}a) \\
 Z \stackrel{?}{=} [a]a, X \stackrel{?}{=} Y & \xRightarrow{[Z:=a]a} \quad \text{(IE)} \\
 X \stackrel{?}{=} Y & \Longrightarrow \quad \text{(F)} \\
 b \#_? X, X \stackrel{?}{=} Y & \xRightarrow{[X:=X-b]} \quad \text{(IF)} \\
 X-b \stackrel{?}{=} Y & \Longrightarrow \quad \text{(F)} \\
 d \#_? X-b, X-b \stackrel{?}{=} Y & \xRightarrow{[X-b:=X-b,d]} \quad \text{(IF)} \\
 X-b, d \stackrel{?}{=} Y & \Longrightarrow \quad \text{(F)} \\
 f \#_? Y, X-b, d \stackrel{?}{=} Y & \xRightarrow{[Y:=Y-f]} \quad \text{(IF)} \\
 X-b, d \stackrel{?}{=} Y-f & \xRightarrow{[Y-f:=X-b,d]} \quad \text{(IE)} \\
 \emptyset & \text{Success, with } [X:=X-b, d], [Y:=X-b, d], [Z:=a]a
 \end{array}$$

Example three (fails).

Take $\text{supp}(X) = \mathbb{A}^<$. We run the algorithm on $\{[a][b]X \stackrel{?}{=} [a]X\}$:

$$\begin{array}{ll}
 [a][b]X \stackrel{?}{=} [a]X & \Longrightarrow \quad (\stackrel{?}{=}[])) \\
 [b]X \stackrel{?}{=} X & \text{Failure}
 \end{array}$$

The algorithm fails because the precondition of rule **(IE)**, $X \notin fv([b]X)$ is not satisfied.

Example four (succeeds).

Take $supp(X) = \mathbb{A}^<$ and take $a, b \in \mathbb{A}^<$. We run the algorithm on $\{X \stackrel{?}{=} (a\ b) \cdot X\}$:

$$\begin{array}{lcl}
 X \stackrel{?}{=} (a\ b) \cdot X & \Longrightarrow & (\stackrel{?}{=} \mathbf{X}) \\
 a \#? X, b \#? X & \xRightarrow{[X := X - a]} & \mathbf{(IF)} \\
 b \#? X - a & \xRightarrow{[X - a := (X - a) - b]} & \\
 \emptyset & \text{Success, with } [X := (X - a) - b] &
 \end{array}$$

Later we will prove Theorem 4.4.6, which tells us that failure here implies that no solution to the unification problem exists.

4.3 Preservation of solutions

... under non-instantiating rules

LEMMA 4.3.1. If $Pr \Longrightarrow Pr'$ by a non-instantiating rule (Definition 4.1.5) then $Sol(Pr) = Sol(Pr')$.

Proof. The empty set cannot be simplified, so suppose $Pr = r \stackrel{?}{=} s, Pr'$ where the simplification rule acts on $r \stackrel{?}{=} s$. We consider two cases:

- *The case $(\stackrel{?}{=}[])$.* Suppose $Pr = [a]r \stackrel{?}{=} [a]s, Pr'$ and $[a]r \stackrel{?}{=} [a]s, Pr' \Longrightarrow r \stackrel{?}{=} s, Pr'$ by $(\stackrel{?}{=}[])$. By Definition 3.4.8 and properties of equality, $[a](r\theta) = [a](s\theta)$ if and only if $r\theta = s\theta$.
- *The case $(\mathbf{F}())$.* Suppose $Pr = a \#?(r_1, \dots, r_n), Pr'$ and suppose that $a \#?(r_1, \dots, r_n), Pr' \Longrightarrow a \#?r_1, \dots, a \#?r_n, Pr'$ by $(\mathbf{F}())$. By Definition 3.4.8 and Lemma 3.2.5, $a \notin fa((r_1, \dots, r_n)\theta)$ if and only if $a \notin fa(r_1\theta), \dots, a \notin fa(r_n\theta)$.

■

LEMMA 4.3.2. Suppose $\theta(X) = \theta'(X)$ for all $X \in fv(Pr)$. Then $\theta \in Sol(Pr)$ if and only if $\theta' \in Sol(Pr)$.

Proof. From Definition 4.1.3 it suffices to show that $r\theta = s\theta$ if and only if $r\theta' = s\theta'$, for every $(r \stackrel{?}{=} s) \in Pr$, and $a \notin fa(r\theta)$ if and only if $a \notin fa(r\theta')$, for every $(a \#?r) \in Pr$. This is immediate using Lemma 3.4.12. ■

... under **(IE)**

Recall from Remark 3.4.7 the discussion of why we write $\pi \cdot X$ when we have chosen a representative element X of an equivalence class of unknowns under permutations.

DEFINITION 4.3.3. Write $\theta - X$ for the substitution such that

$$\begin{aligned} (\theta - X)(\pi \cdot X) &= \pi \cdot X \\ (\theta - X)(Y) &= \theta(Y) \quad \text{for all other } Y. \end{aligned}$$

In the right circumstances, a substitution θ can be factored as ‘a part of θ that does not touch X ’ and ‘a single substitution for X ’:

THEOREM 4.3.4. If $X\theta = s\theta$ and $X \notin fv(s)$ then

$$\theta = [X := s] \circ (\theta - X).$$

That is:

$$\begin{aligned} \theta(X) &= X([X := s] \circ (\theta - X)) \quad \text{and} \\ \theta(Y) &= Y([X := s] \circ (\theta - X)). \end{aligned}$$

Proof. We reason as follows:

$$\begin{aligned} (\pi \cdot X)([X := s] \circ (\theta - X)) &= (\pi \cdot s)(\theta - X) && \text{Definition 3.4.8, Lemma 3.5.2} \\ &= (\pi \cdot s)\theta && X \notin fv(s), \text{ Lemma 3.4.12} \\ &= (\pi \cdot X)\theta && \text{Assumption} \\ \\ Y([X := s] \circ (\theta - X)) &= Y(\theta - X) && \text{Definition 3.4.8, Lemma 3.5.2} \\ &= Y\theta && \text{Definition 4.3.3} \end{aligned}$$

■

... under **(IF)**

DEFINITION 4.3.5. Suppose θ is a substitution. Suppose $a \in \text{supp}(X)$ and $a \notin fa(\theta(X))$. Let δ_{X-a} be a *shift* permutation bijecting $\text{supp}(X)$ with $\text{supp}(X) \setminus \{a\}$.

Then define a substitution $\theta_{[X-a:=X]}(X)$ by:

$$\begin{aligned} (\theta_{[X-a:=X]})(\pi \cdot X) &= (\delta_{X-a}^{-1} \circ \pi) \cdot \theta(X) \\ (\theta_{[X-a:=X]})(Y) &= \theta(Y) \quad \text{for all other } Y. \end{aligned}$$

It is routine to verify that Definition 4.3.5 is well-defined and a substitution.

THEOREM 4.3.6. If $a \in \text{supp}(X)$ and $a \notin \text{fa}(\theta(X))$ then

$$\theta = [X:=X-a] \circ (\theta_{[X-a:=X]}).$$

That is:

$$\begin{aligned} \theta(\pi \cdot X) &= ([X:=X-a] \circ \theta_{[X-a:=X]})(\pi \cdot X) \quad \text{and} \\ \theta(Y) &= ([X:=X-a] \circ \theta_{[X-a:=X]})(Y). \end{aligned}$$

Proof. A fact of the group action. ■

4.4 Simplification rewrites calculate principal solutions

DEFINITION 4.4.1. Write $\theta_1 \leq \theta_2$ when there exists some θ' such that $X\theta_2 = X(\theta_1 \circ \theta')$ always. Call \leq the **instantiation ordering**.

DEFINITION 4.4.2. A **principal** (or **most general**) solution to a problem Pr is a solution $\theta \in \text{Sol}(Pr)$ such that $\theta \leq \theta'$ for all other $\theta' \in \text{Sol}(Pr)$.

Our main result is Theorem 4.4.5: the unification algorithm from Definition 4.1.7 calculates a principal solution.

LEMMA 4.4.3. If $\theta_1 \leq \theta_2$ then $\theta \circ \theta_1 \leq \theta \circ \theta_2$.

Proof. By Definition 4.4.1, θ' exists such that $X\theta_2 = X(\theta_1 \circ \theta')$ always. Then:

$$\begin{aligned} X(\theta \circ \theta_2) &= (X\theta)\theta_2 && \text{Lemma 3.5.2} \\ &= (X\theta)(\theta_1 \circ \theta') && \text{Lemma 3.4.12} \\ &= X((\theta \circ \theta_1) \circ \theta') && \text{Lemma 3.5.2} \end{aligned}$$

LEMMA 4.4.4.

1. Suppose $\text{fa}(s) \subseteq \text{supp}(X)$ and $X \notin \text{fv}(s)$. Write $\chi = [X:=s]$. If $Pr \xrightarrow{X} Pr'$ with **(IE)** then $\theta \in \text{Sol}(Pr)$ implies $\theta - X \in \text{Sol}(Pr')$.
2. Suppose $a \in \text{supp}(X)$. Write $\rho = [X:=X-a]$. If $Pr \xrightarrow{\rho} Pr'$ with **(IF)** then $\theta \in \text{Sol}(Pr)$ implies $\theta_{[X-a:=X]} \in \text{Sol}(Pr')$.

Proof.

1. Suppose $Pr = X \stackrel{?}{=} s$, Pr'' so that $X \stackrel{?}{=} s$, $Pr'' \xrightarrow{X} Pr''\chi$. Now suppose $\theta \in \text{Sol}(Pr)$. By Theorem 4.3.4 $\chi \circ (\theta - X) \in \text{Sol}(Pr)$. By Lemma 4.1.4, $\theta - X \in \text{Sol}(Pr\chi)$. It follows that $\theta - X \in \text{Sol}(Pr''\chi)$ as required.
2. Suppose $Pr = a \# ? X$, Pr'' and $a \in \text{supp}(X)$ so that $Pr \xrightarrow{\rho} Pr\rho$. Now suppose $\theta \in \text{Sol}(Pr)$. By Theorem 4.3.6 $\rho \circ \theta_{[X-a:=X]} \in \text{Sol}(Pr)$. By Lemma 4.1.4, $\theta_{[X-a:=X]} \in \text{Sol}(Pr\rho)$ as required.

■

THEOREM 4.4.5. If $Pr \xRightarrow{\theta} \emptyset$ then θ is a principal solution to Pr (Definition 4.4.2).

Proof. By induction on the path of $Pr \xRightarrow{\theta} \emptyset$.

- *The empty path.* So $Pr = \emptyset$ and $\theta = id$. By Definition 4.4.1, $id \leq \theta'$.
- *The non-instantiating case.* Suppose

$$Pr \Longrightarrow Pr' \xRightarrow{\theta} \emptyset$$

where $Pr \Longrightarrow Pr'$ by a non-instantiating rule. By inductive hypothesis θ is a principal solution of Pr' . It follows from Lemma 4.3.1 that θ is also a principal solution of Pr .

- *The case (IE).* Suppose $fa(r) \subseteq \text{supp}(X)$ and $X \notin \text{fv}(r)$. Write $\chi = [X:=r]$. Suppose $Pr = r \stackrel{?}{=} X, Pr''$ so that

$$r \stackrel{?}{=} X, Pr'' \xrightarrow{X} Pr''\chi \xRightarrow{\theta'} \emptyset.$$

Further, consider any other $\theta' \in \text{Sol}(Pr)$.

By Lemma 4.4.4 $(\theta' - X) \in \text{Sol}(Pr''\chi)$ and by inductive hypothesis $\theta'' \in \text{Sol}(Pr''\chi)$ and $\theta'' \leq \theta' - X$. By Lemma 4.4.3, $\chi \circ \theta'' \leq \chi \circ (\theta' - X)$. By Theorem 4.3.4 $\chi \circ (\theta' - X) = \theta'$.

- *The case (IF).* Suppose $a \in \text{supp}(X)$. Write $\rho = [X:=X-a]$, so that

$$Pr \xRightarrow{\rho} Pr\rho \xRightarrow{\theta''} \emptyset,$$

Further, consider any other $\theta' \in \text{Sol}(Pr)$.

By Lemma 4.4.4, $\theta'_{[X-a:=X]} \in \text{Sol}(Pr\rho)$ and by inductive hypothesis $\theta'' \in \text{Sol}(Pr\rho)$ and $\theta'' \leq \theta'_{[X-a:=X]}$. By Lemma 4.4.3, $\rho \circ \theta'' \leq \rho \circ \theta'_{[X-a:=X]}$. By Theorem 4.3.6 $\rho \circ \theta'_{[X-a:=X]} = \theta'$.

■

THEOREM 4.4.6 (Correctness of algorithm). Given a problem Pr , if the algorithm of Definition 4.1.7 succeeds then it returns a principal solution; if it fails then there is no solution.

Proof. If the algorithm succeeds we use Theorem 4.4.5. Otherwise, the algorithm generates an element of the form $f(r) \stackrel{?}{=} g(s)$, $a \stackrel{?}{=} b$, $a \#_? a$, $a \#_? C$ where $a \in \text{supp}(C)$, or $X \stackrel{?}{=} s$ where $X \in \text{fv}(s)$ and s is not of the form $\pi \cdot X$. By arguments on syntax and size of syntax, no solution to the reduced problem exists. It follows by Lemma 4.4.4 that no solution to Pr exists. ■

DEFINITION 4.4.7. Fix terms r and s .

- Call **nominal unification** the problem of finding a θ to make $r\theta = s\theta$.
- Call **nominal matching** the problem of finding a θ to make $r\theta = s$.

COROLLARY 4.4.8. Providing that equality of \mathcal{C} (constants), \mathcal{X} (unknowns), and \mathbb{P} (permutations) are decidable, nominal unification and nominal matching over signatures using them are also decidable.

Proof. An algorithm for unification is sketched in Definition 4.1.7; furthermore by Theorem 4.4.6 it calculates a most general θ which represents all other solutions.

For matching, we substitute unknowns in s with fresh (non-equivariant) constants of the same sorts and permission sets—we extend the signature if we need to—and run the unification algorithm. We then replace the constants by the original unknowns.¹⁰ It is not hard to see that this calculates a most general matching solution. ■

REMARK 4.4.9. The matching and unification algorithms might generate solutions with *shift*-permutations. If we prefer to eliminate them then—provided that \mathcal{X} has enough unknowns (Definition 3.1.10)—we may do so by appending an invertible substitution (Definition 3.5.3) mapping each shifted $\delta \cdot X$ in the solution to a fresh unknown Y such that $\text{supp}(Y) = \delta \cdot \text{supp}(X)$.

5 REWRITING

Nominal rewriting was the first logical system designed to study theories (sets of axioms, i.e. *rewrite rules*) over nominal terms. It was introduced by Fernández and the author in [Fernández *et al.*, 2004; Fernández and Gabbay, 2007]. Nominal terms allow us to express rewrite rules involving binding, like substitution and the λ -calculus (see Example 5.1.3).

The presentation of nominal rewriting here differs from that in [Fernández and Gabbay, 2007], and is more concise. Partly this is optimisation, but this is also due to the permissive-nominal approach. We compare and contrast nominal rewriting from [Fernández and Gabbay, 2007] with nominal rewriting here, in Subsection 5.6.

¹⁰We do not make this formal, but since constants are structurally just like unknowns the definitions can easily be constructed by proceeding exactly as we did when we defined substitution for unknowns.

5.1 Rewrite rules

DEFINITION 5.1.1. A **rewrite rule** in a signature $\Sigma = (\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{F}, ar)$ is a pair of terms $l \rightarrow m$ in Σ such that $sort(l) = sort(m) \in \mathcal{B}$ and $fv(m) \subseteq fv(l)$.

R will range over rewrite rules.

A **rewrite theory** $\mathbf{R} = (\Sigma, Rew)$ is a pair of a signature Σ (Definition 3.1.5) and a (possibly infinite) set of rewrite rules Rew in Σ .

NOTATION 5.1.2. Write $(l \rightarrow m) \in \mathbf{R}$ to mean ‘ l and m are terms in Σ and $(l \rightarrow m) \in Rew$ ’.

The notion of rewrite rule and rewrite theory in Definition 5.1.1 is much like the first-order case, but because of the ‘nominal’ aspects of our syntax we can handle names and binding.

EXAMPLE 5.1.3. Here are some example rewrite theories:

- **nrSUB** expresses the usual capture-avoiding substitution action on λ -calculus terms.

Let Σ have a base sort τ and the following term-formers:

$$\text{sub} : ([\nu]\tau, \tau)\tau \quad \text{lam} : ([\nu]\tau)\tau \quad \text{app} : (\tau, \tau)\tau \quad \text{var} : (\nu)\tau$$

Rewrite rules are as follows:

$$\begin{array}{lll} (\text{var} \rightarrow) & \text{var}(a)[a \mapsto X] & \rightarrow X \\ (\text{var} \rightarrow') & \text{var}(b)[a \mapsto X] & \rightarrow \text{var}(b) \\ (\text{lam} \rightarrow) & \text{lam}([\![a]\!]X)[b \mapsto Y] & \rightarrow \text{lam}([\![a]\!](X[b \mapsto Y])) \quad (a \notin \text{supp}(Y)) \\ (\text{app} \rightarrow) & \text{app}(X, X')[b \mapsto Y] & \rightarrow \text{app}(X[b \mapsto Y], X'[b \mapsto Y]) \end{array}$$

Here and in the next example we sugar $\text{sub}([\![a]\!]r, t)$ to $r[a \mapsto t]$. Every permission set contains b and every permission set contains a except for $\text{supp}(Y)$, as indicated above.

- **nrLAM** extends the previous theory with two more rewrites:

$$\begin{array}{lll} (\beta \rightarrow) & (\lambda[\![a]\!]Z)X & \rightarrow Z[a \mapsto X] \\ (\eta \rightarrow) & \lambda[\![a]\!](Ya) & \rightarrow Y \quad (a \notin \text{supp}(Y)) \end{array}$$

Sugar $\text{lam}(r)$ to λr , $\text{app}(r, s)$ to rs , and $\text{var}(a)$ to a . We anticipate Subsection 5.2 and sketch how one might rewrite $(\lambda[b](\lambda[a]ab))a$ to $\lambda[a']a'a$:

$$\begin{aligned} (\lambda[b](\lambda[a]ab))a &\rightarrow (\lambda[a]ab)[b \mapsto a] \\ &= (\lambda[a']a'b)[b \mapsto a] \\ &\rightarrow \lambda[a']((a'b)[b \mapsto a]) \\ &\rightarrow^* \lambda[a']a'a \end{aligned}$$

5.2 Rewrite steps

DEFINITION 5.2.1. Define the terms s in which X occurs **only once** by:

$$s ::= \pi \cdot X \mid [a]s \mid f(r_1, \dots, r_{i-1}, s, r_{i+1}, \dots, r_n) \\ (X \notin fv(r_1), \dots, fv(r_{i-1}), fv(r_{i+1}), \dots, fv(r_n))$$

A **position** P is a pair (s, X) of a nominal term and an unknown X which occurs only once in s .

Our notion of *position* is also sometimes called a **context**; the idea goes back to at least [?].

In Definition 5.2.1, $\pi \cdot X$ denotes an unknown in the same permutation orbit as X .

NOTATION 5.2.2. If $P = (s, X)$ is a position write $supp(P)$ for $supp(X)$ and $sort(P)$ for $sort(X)$.

If $sort(r) = sort(P)$ and $fa(r) \subseteq supp(P)$ (so that $[X:=r]$ is a substitution) write $P[r]$ for $s[X:=r]$.

DEFINITION 5.2.3. The **one-step rewrite relation** $r \xrightarrow{R} s$ is the least relation such that for every $(l \rightarrow m) \in R$, position P , and substitution θ , if $sort(r) = sort(P)$ and $fa(l\theta) \cup fa(m\theta) \subseteq supp(P)$ (so that $P[l\theta]$ and $P[m\theta]$ are well-defined) then

$$P[l\theta] \xrightarrow{R} P[m\theta].$$

The **multi-step rewrite relation** $r \xrightarrow{R^*} s$ is the reflexive transitive closure of the one-step rewrite relation.

We consider decidability and complexity of the rewrite relation in Section 6.

EXAMPLE 5.2.4. Let T have one name sort ν , one base sort τ , one term-former $triv$ and one axiom $triv(a) \rightarrow triv(b)$.

Then $triv(a) \rightarrow triv(b)$ but also (using positions $(\pi \cdot X, X)$ for any π) $triv(b) \rightarrow triv(a)$ and $triv(a') \rightarrow triv(b')$ for any pair of distinct atoms a' and b' .

Thus atoms in rewrite rules range over ‘any atom’ analogously to how unknowns in rewrite rules range over ‘any term’.

EXAMPLE 5.2.5. Recall the rule $(\eta \rightarrow) = (\lambda[a](Ya) \rightarrow Y)$ where $a \notin supp(Y)$ from Example 5.1.3. Suppose also $b \notin supp(Y)$.


1. To deduce $\lambda[a](ba) \rightarrow b$ we take $P = ((b \ c) \cdot Y, Y)$ for some $c \in supp(Y)$ and we take $\theta = [Y:=c]$.

2. To deduce $\lambda[a'](ba') \rightarrow b$ for any other a' we *also* take $P = ((b\ c)\cdot Y, Y)$ and $\theta = [Y:=c]$. This is because $\lambda[a'](ba')$ and $\lambda[a](ba)$ are the same term (Lemma 2.4.9).
3. To deduce $\lambda[a](Ya) \rightarrow Y$ we take $P = (Y, Y)$ and $\theta = id$.
4. Suppose $supp(Y') = supp(Y) \cup \{a\}$.
 Suppose we have *shift*-permutations so there exists a permutation, write it $\delta_{Y'-a}$, bijecting $supp(Y')$ with $supp(Y)$. To deduce $\lambda[a](Y'a) \rightarrow Y'$ we take $P = ((\delta_{Y'-a})^{-1}\cdot Y, Y)$ and $\theta = [Y:=\delta_{Y'-a}\cdot Y']$.
 Without *shift* we cannot deduce $\lambda[a](Y'a) \rightarrow Y'$; we can still deduce $\lambda[a](Ya) \rightarrow Y$.
5. We cannot deduce $\lambda[a](aa) \rightarrow a$, because $[Y:=a]$ is not a substitution: no function mapping Y to a can be equivariant, since $(b\ a)\cdot Y = Y$ but $(b\ a)\cdot a = b \neq a$ (also $a \notin supp(Y)$: see Proposition 3.4.3).
6. A rewrite $X \rightarrow X$ only entails rewrites for t with $fa(t) \subseteq \pi \cdot supp(X)$ for some π . With *shift*, the effect of this may be that we can deduce $t \rightarrow t$ from $X \rightarrow X$ for any t . We make no claim to there being a ‘right’ or ‘wrong’ answer here: the issue is purely a design question of how much expressivity we want permutations to have. Our results are parameterised over this choice.

DEFINITION 5.2.6.

- Call **R locally confluent** when $r \xrightarrow{R} s_1$ and $r \xrightarrow{R} s_2$ implies there exists some s' such that $s_1 \xrightarrow{R^*} s'$ and $s_2 \xrightarrow{R^*} s'$.
- Call **R confluent** when $r \xrightarrow{R^*} s_1$ and $r \xrightarrow{R^*} s_2$ implies there exists some s' such that $s_1 \xrightarrow{R^*} s'$ and $s_2 \xrightarrow{R^*} s'$.

We illustrate this below.



5.3 Peaks, critical pairs, joinability

We now begin to investigate criteria for deducing confluence of nominal rewrite systems. Our first observation is that things are not quite as simple as in first-order rewriting [Baader and Nipkow, 1998, Section 6.2]: by Lemma 5.3.5, trivial critical pairs are not always joinable.

DEFINITION 5.3.1. Write $r \rightarrow s_1, s_2$ when $r \rightarrow s_1$ and $r \rightarrow s_2$ and call this a **peak**. Call this peak **joinable** when there exists a t such that $s_1 \rightarrow^* t$ and $s_2 \rightarrow^* t$.

So \mathbf{R} is locally confluent when every peak is joinable.

DEFINITION 5.3.2. Consider two rewrite rules $R_1 = (l_1 \rightarrow m_1)$ and $R_2 = (l_2 \rightarrow m_2)$. Call R_1 a **copy** of R_2 when there exists an invertible substitution θ such that $(l_2\theta \rightarrow m_2\theta) = R_1$.

Clearly, if R_1 is a copy of R_2 then R_2 is also a copy of R_1 . Furthermore:

LEMMA 5.3.3. If R_1 and R_2 are copies of the same rule then $l \xrightarrow{R_1} m$ if and only if $l \xrightarrow{R_2} m$.

Proof. Unpacking Definition 5.2.3 and exploiting the existence of an inverse θ^{-1} . ■

DEFINITION 5.3.4. Suppose that $R_i = (l_i \rightarrow m_i)$ for $i = 1, 2$ and $fv(R_1) \cap fv(R_2) = \emptyset$. Suppose $l_1 = P[l'_1]$ for some l'_1 , and suppose $l'_1 \stackrel{?}{=} l_2$ has a principal solution θ . Call the pair $(m_1\theta, P[m_2]\theta)$ a **critical pair**.

Call $(m_1\theta, P[m_2]\theta)$ **trivial** when at least one of the following hold:

1. $P = (\pi \cdot X, X)$ and R_1 and R_2 are copies of the same rule.
2. $l'_1 = X$ for some unknown X .

LEMMA 5.3.5. Peaks that are instances of trivial critical pairs, are not always joinable.

Proof. It suffices to provide a counterexample. Fix term-formers 0 and f and take $R_1 = (0 \rightarrow a)$ and $R_2 = (X \rightarrow f(a))$ where $a \notin \text{supp}(X)$.

There is a critical pair $(a, f(a))$ between R_1 and R_2 .

Also, $0 \xrightarrow{R_1} a$ and $0 \xrightarrow{R_2} f(a)$ and it is a fact that this peak cannot be joined—we ‘want’ to close this peak by rewriting a to $f(a)$ using R_2 , but the fact that $a \notin \text{supp}(X)$ blocks this. ■

5.4 Uniform rewriting

The proof of Lemma 5.3.5 suggests a simple cure:

DEFINITION 5.4.1. Call a rule $R = (l \rightarrow m)$ **uniform** when

$$fa(m) \subseteq fa(l).$$

Call a rewrite theory \mathbf{R} **uniform** when every $R \in \mathbf{R}$ is uniform.

Definition 5.4.1 mirrors the condition in Definition 5.1.1 that $fv(m) \subseteq fv(l)$, but for atoms instead of unknowns. This condition is sufficient to

obtain Theorem 5.4.7, which is a nominal rewriting version of the well-known critical pair lemma from first-order rewriting [Baader and Nipkow, 1998, Theorem 6.2.4].

EXAMPLE 5.4.2. Let R have one name sort ν , one base sort τ , two term-formers $\text{triv} : (\nu)\tau$ and $\text{abs} : ([\nu]\tau)\tau$, and rewrite rules

$$\text{triv}(a) \rightarrow \text{triv}(a) \quad \text{triv}(a) \rightarrow \text{triv}(b) \quad \text{abs}([a]X) \rightarrow X.$$

$fa(\text{triv}(a)) \subseteq fa(\text{triv}(a))$ and $fa(\text{triv}(b)) \not\subseteq fa(\text{triv}(a))$. Also $fa(X) \not\subseteq fa(\text{abs}([a]X))$ if and only if $a \notin \text{supp}(X)$.

So the first rule is uniform, the second is not, and the third is uniform if and only if $a \notin \text{supp}(X)$.

The rewrite rules of nrSUB and nrLAM in Example 5.1.3 are uniform.¹¹

LEMMA 5.4.3. If $fa(m) \subseteq fa(l)$ then $fa(P[m]) \subseteq fa(P[l])$.

Proof. Routine induction using Lemmas 3.2.8 and 3.2.5. ■

COROLLARY 5.4.4. $R = (l \rightarrow m)$ is uniform if and only if $\forall r, s. (r \xrightarrow{R} s \Rightarrow fa(s) \subseteq fa(r))$.

Proof. From Lemmas 3.2.8 and 3.4.11. ■

LEMMA 5.4.5. Suppose $R = (l \rightarrow m)$ is uniform and $X \notin fv(R)$. Suppose $\theta(X) = l\theta$. Specify θ' by $\theta'(\pi \cdot X) = \pi \cdot (m\theta)$ and $\theta'(Y) = \theta(Y)$. Then $r\theta \xrightarrow{*} r\theta'$ for any r .

Proof. θ' is a substitution by Lemmas 3.4.11 and 3.2.8. The result follows by a routine induction on r . ■

Because of Lemma 5.3.3, we can be relaxed about the particular (orbits of) unknowns that are used in a rewrite rule, if we only care about the rewrites that they generate. We do this in Theorems 5.4.6 and 5.4.7. This can always be made formal by inserting invertible ‘freshening’ substitutions as appropriate.

THEOREM 5.4.6. If a rewrite theory R (Definition 5.1.1) is uniform then peaks that are instances of trivial critical pairs, are joinable.

Proof. Consider two rules $R_i = (l_i \rightarrow m_i) \in R$ for $i = 1, 2$. Taking copies if necessary, suppose $fv(R_1) \cap fv(R_2)$. Suppose they have a critical pair $(m_1\theta, P[m_2]\theta)$. That is, there exists l'_1 such that $l_1 = P[l'_1]$ and θ is a principal solution to $l'_1 \stackrel{?}{=} l_2$.

There are two cases:

¹¹There is a deeper reason for this: they are also closed. See Example 6.2.2 and Theorem 6.2.3.

- The case $P = (\pi \cdot X, X)$ and R_1 and R_2 are copies of the same rule $l \rightarrow m$. The peak we want to join is $l_1\theta = \pi \cdot l_2\theta \rightarrow m_1\theta, \pi \cdot m_2\theta$, where the rules $l_1 \rightarrow m_1$ and $l_2 \rightarrow m_2$ are identical aside from their free variables which are renamed disjoint. We use Lemma 3.4.12 and the assumption in Definition 5.1.1 that $fv(m) \subseteq fv(l)$.
- The case of $(m_1\theta, P[m_2]\theta)$ where $l_1 = P[X]$ and $\theta(X) = l_2$. Specify θ' by $\theta'(\pi \cdot X) = \pi \cdot m_2$ and $\theta'(Y) = \theta(Y)$ for all other Y ; note that θ' is a substitution since $fa(m_2) \subseteq fa(l_2)$ by uniformity and $fa(l_2) \subseteq supp(X)$ by our assumption that θ is a substitution.
By Lemma 5.4.5 $m_1\theta \rightarrow^* m_1\theta'$. By definition $P[m_2]\theta = l_1\theta' \xrightarrow{R_1} m_1\theta'$, so we have joined the peak. ■

THEOREM 5.4.7. Suppose all non-trivial critical pairs of R are joinable and suppose R is uniform. Then R is locally confluent.

Proof. Suppose $r \xrightarrow{R_1} s_1$ and $r \xrightarrow{R_2} s_2$. Write P_1 and P_2 for the positions at which the two rewrites occur. Taking copies if necessary, suppose $fv(R_1) \cap fv(R_2) = \emptyset$.

If P_1 and P_2 identify distinct subterms of r then local confluence holds by a standard diagrammatic argument (see for instance [Baader and Nipkow, 1998]).

Otherwise it must be that $P_2 = (P_1[P], X)$ for some position P ; that is, P_2 identifies a point in r beneath the point identified by P_1 (or the symmetric case that $P_1 = (P_2[P], X)$, which is similar and we elide). There are now three possibilities:

1. X in P_2 replaces an unknown in r . This is an instance of a trivial critical pair; we use Theorem 5.4.6.
2. $P = (\pi \cdot X, X)$ and R_1 and R_2 are copies of the same rule. Then again this is an instance of a trivial critical pair and we use Theorem 5.4.6.
3. Otherwise, this is an instance of a non-trivial critical pair at it may be joined using our assumption that non-trivial critical pairs are joinable. ■

DEFINITION 5.4.8. Call a rewrite system R **terminating** when all rewrite sequences are finite. Call a term r a **normal form** (with respect to a rewrite system R) when $\forall s. \neg(r \xrightarrow{R} s)$, that is, when r does not R -rewrite to anything.

EXAMPLE 5.4.9. It can be proved that nrSUB in Example 5.1.3 is terminating. nrLAM (famously) is not terminating, because of $(\beta \mapsto)$.

COROLLARY 5.4.10. Suppose R is terminating, uniform, and suppose non-trivial critical pairs in R are joinable. Then:

1. R is confluent.
2. If $r \twoheadrightarrow^* s$ and $r \twoheadrightarrow^* s'$ and s and s' are normal forms, then $s = s'$.

5.5 Orthogonal rewrite systems

We now treat another standard criterion in rewriting: orthogonality [Der-showitz and Jouannaud, 1989; Baader and Nipkow, 1998]. By Theorem 5.5.7 orthogonality implies not only local confluence, but the stronger property of confluence (Definition 5.2.6). The proof is not direct: it turns out that it is easier to consider an auxilliary *parallel reduction* relation \Rightarrow (Definition 5.5.4). The reflexive transitive closure of \Rightarrow is equal to that of \rightarrow (Lemma 5.5.5), but \Rightarrow allows (intuitively) multiple reductions provided that they do not occur ‘one after the other, in the same position’. This is the kind of multiple reduction generated in the second case of the proof of Theorem 5.4.6, when we rewrite $m_1\theta$ to $m_1\theta'$.

DEFINITION 5.5.1. Call $R = (l \rightarrow m)$ **left-linear** when each unknown occurring in l occurs only once (Definition 5.2.1).

For example $f(X) \rightarrow g(X, X)$ is left-linear but $g(X, X) \rightarrow f(X)$ and $g(\pi \cdot X, x) \rightarrow f(X)$ are not. Note that $(a, a) \rightarrow a$ is left-linear.

DEFINITION 5.5.2. Call R **orthogonal** when every $R \in R$ is uniform and left-linear, and all critical pairs are trivial.

(Note that we insist that R is uniform, as well as the standard condition that it be left-linear.)

DEFINITION 5.5.3. Suppose $R = (l \rightarrow m)$. Write $r \xrightarrow{R}_\epsilon s$ when $r \xrightarrow{R} s$ and the rewrite occurs at a position $P = (\pi \cdot X, X)$. We say that the rewrite with R occurs at **root position**.

Expanding Definition 5.5.3, $r \xrightarrow{R}_\epsilon s$ when there exists θ and π such that $r = \pi \cdot (l\theta)$ and $s = \pi \cdot (m\theta)$. For example: if $R = (a \rightarrow a)$ then $a \xrightarrow{R}_\epsilon a$ but not $[a]a \xrightarrow{R}_\epsilon [a]a$.

DEFINITION 5.5.4. We define a **parallel reduction** relation \Rightarrow by the rules in Figure 3.

LEMMA 5.5.5. $r \twoheadrightarrow^* s$ if and only if $r \Rightarrow^* s$.

Proof. By routine inductions. ■

LEMMA 5.5.6. If R is orthogonal then \Rightarrow is confluent.

$$\begin{array}{c}
\frac{r_1 \Rightarrow s_1 \cdots r_n \Rightarrow s_n}{f(r_1, \dots, r_n) \Rightarrow f(s_1, \dots, s_n)} (\Rightarrow f) \\
\\
\frac{r_1 \Rightarrow s_1 \cdots r_n \Rightarrow s_n \quad f(s_1, \dots, s_n) \xrightarrow{R}_\epsilon s'}{f(r_1, \dots, r_n) \Rightarrow s'} (\Rightarrow f') \\
\\
\frac{s \Rightarrow t}{[a]s \Rightarrow [a]t} (\Rightarrow \mathbf{abs}) \quad \frac{r \Rightarrow s \quad [a]s \xrightarrow{R}_\epsilon s'}{[a]r \Rightarrow s'} (\Rightarrow \mathbf{abs}') \\
\\
\frac{}{r \Rightarrow r} (\mathbf{refl}) \quad \frac{a \xrightarrow{R}_\epsilon s'}{a \Rightarrow s'} (\Rightarrow \mathbf{a}') \quad \frac{X \xrightarrow{R}_\epsilon s'}{X \Rightarrow s'} (\Rightarrow \mathbf{X}')
\end{array}$$

Figure 3: Parallel reduction relation

Proof. We prove by induction on the derivation of $r \Rightarrow s$ that a stronger property holds, often called the **diamond property**: for all s' if $r \Rightarrow s'$ then there exists some s'' such that $s \Rightarrow s''$ and $s' \Rightarrow s''$. From this, confluence easily follows by a standard diagrammatic argument.

We consider a selection of cases:

- *The derivations of $r \Rightarrow s$ and $r \Rightarrow s'$ both end in $(\Rightarrow f)$.* We use the inductive hypotheses and $(\Rightarrow f)$.
- *The derivation of $r \Rightarrow s$ ends in $(\Rightarrow f)$ and that of $r \Rightarrow s'$ ends in $(\Rightarrow f')$.* So $r_i \Rightarrow s_i$ and $r_i \Rightarrow s'_i$ for $1 \leq i \leq n$, and $f(s'_1, \dots, s'_n) = \pi \cdot (l\theta) \xrightarrow{R}_\epsilon \pi \cdot (m\theta)$ for some π and $R = (l \rightarrow m) \in \mathbf{R}$. By inductive hypothesis there exist s''_i such that $s_i \Rightarrow s''_i$ and $s'_i \Rightarrow s''_i$. We now proceed as illustrated and explained below:

$$\begin{array}{ccc}
f(r_1, \dots, r_n) \Longrightarrow f(s'_1, \dots, s'_n) & = & \pi \cdot (l\theta) \xrightarrow{R}_\epsilon \pi \cdot (m\theta) \\
\Downarrow & & \Downarrow \\
f(s_1, \dots, s_n) \Longrightarrow f(s''_1, \dots, s''_n) & = & \pi \cdot (l\theta') \xrightarrow{R}_\epsilon \pi \cdot (m\theta')
\end{array}$$

Either l is an unknown X or the rewrite $f(s'_1, \dots, s'_n) \Rightarrow f(s''_1, \dots, s''_n)$ takes place in the substitution θ .

If l is an unknown then by uniformity we may rewrite $f(s''_1, \dots, s''_n)$ using R and close the diagram by rewriting corresponding instances of $\theta(X)$ in $\pi \cdot (m\theta)$.

Otherwise, by uniformity there is a substitution θ' such that $\theta(X) \Rightarrow \theta'(X)$ for every X and $f(s''_1, \dots, s''_n) = \pi \cdot (l\theta')$. Rules are also left-linear

so R still applies to $\pi \cdot (l\theta)$: $f(s''_1, \dots, s''_n) \xrightarrow{R}_\epsilon \pi \cdot (m\theta')$ and therefore $f(s_1, \dots, s_n) \Rightarrow s\theta'$ by $(\Rightarrow f')$ for R .

The other cases are no harder. ■

THEOREM 5.5.7. If a theory R is orthogonal (Definition 5.5.2) then R is confluent (Definition 5.2.6).

Proof. If the uniform rewrite system has only left-linear rules and only trivial critical pairs, then \Rightarrow is confluent by Lemma 5.5.6. It follows that \Rightarrow^* is confluent. By Lemma 5.5.5 the result follows. ■

5.6 Nominal rewriting with freshness contexts versus permissive-nominal rewriting

As mentioned in the introduction to this Section, the presentation of this paper differs from that of [Fernández and Gabbay, 2007] in being permissive-nominal.

For clarity, let us call the nominal rewrite framework from [Fernández and Gabbay, 2007] ‘System ∇ ’ and the nominal rewrite framework here ‘System S ’.

In system ∇ a rewrite rule takes the form $\nabla \vdash t \rightarrow u$ where ∇ is a set of assumptions $a\#X$ called a *freshness context*. X is an *unknown*. This is not typed by a permission set; freshness information is given by ∇ .

Here are $(\lambda \rightarrow)$ from Example 5.1.3, and how it would look in System ∇ :

$$\begin{array}{l} \text{System } S \quad \text{lam}([a]X)[b \rightarrow Y] \rightarrow \text{lam}([a](X[b \rightarrow Y])) \quad (a \notin \text{supp}(Y)) \\ \text{System } \nabla \quad a\#Y \vdash \text{lam}([a]X)[b \rightarrow Y] \rightarrow \text{lam}([a](X[b \rightarrow Y])) \end{array}$$

$a \notin \text{supp}(Y)$ is a fact (we must choose Y so that this is true). It does not matter *which* permission set we give Y because using δ and swappings we can build a π to map $\text{supp}(Y)$ to every other permission set $\pi \cdot \text{supp}(Y)$ —which will contain $\pi(a)$.

Conversely $a\#Y$ is a freshness condition. It directly controls the terms to which we may instantiate Y ; they must not contain a free. Here we attain this effect using Proposition 3.4.3.

Freshness conditions are elementary: they mean what they say and what they mean be quickly understood. Permission sets are still finitely representable, but somewhat harder to understand. So from the point of view of keeping a gentle learning curve, System ∇ may be preferable to System S .

However, System S rewards us with some advantages: we can use nominal abstract syntax and the freshness conditions which must be explicitly stated (repeatedly) in [Fernández and Gabbay, 2007] are handled in the background by equivariance of substitutions (as Proposition 3.4.3 makes formal).

This also has some effects on mathematical properties. In System ∇ from [Fernández and Gabbay, 2007] it was not in general the case that

if $\nabla \vdash r \approx_\alpha r'$ and $\nabla \vdash r \xrightarrow{R} s$ then $\nabla \vdash r' \xrightarrow{R} s$ (see the end of Subsection 5.2 in [Fernández and Gabbay, 2007]). It was also not in general the case that nominal rewriting coincides with nominal algebra (Section 7), essentially because any fixed freshness contexts might not be ‘big enough’. Fernández and the author wrote a paper on how to adjust for this [Fernández and Gabbay, 2010]. In a permissive-nominal context, these issues do not arise in the first place.

This author’s feeling is that nominal-terms-with-freshness-contexts and permissive-nominal terms can be considered as essentially the same thing. However, if our goal is to prove theorems then we get closer to what is ‘really going on’ via the permissive-nominal presentation.

6 CLOSED TERMS

Equivariant unification—the problem of finding θ and π such that $\pi \cdot (r\theta) = s\theta$ —is NP complete [Cheney, 2004; Cheney, 2010]. The same applies to corresponding matching problems. This matters to us because the rewrite relation in Definition 5.2.3 is equivariant; to determine whether r rewrites with a rule ($l \rightarrow r$), we must solve an equivariant matching problem.

Fernández and the author introduced a notion of *closed term* such that for closed terms, equivariant matching/unification coincides with ‘ordinary’ matching/unification [Fernández and Gabbay, 2007]. That is, for closed terms we can throw away the π .

We now develop corresponding definitions and results. The definitions and proofs in this paper are significantly different from those in [Fernández and Gabbay, 2007].¹²

6.1 The definition

DEFINITION 6.1.1. Define **explicit atoms** $ea(r)$ inductively by:

$ea(a) = \{a\}$	$ea(C) = \text{supp}(C)$	$ea(X) = \emptyset$
$ea(\mathbf{f}(r)) = ea(r)$	$ea((r_1, \dots, r_n)) = \bigcup ea(r_i)$	$ea([a]r) = ea(r) \setminus \{a\}$

REMARK 6.1.2. Intuitions for $ea(r)$ versus $fa(r)$ are as follows:

- The explicit atoms of r are the atoms that actually appear in r (unbound). That is, we can read ‘ $a \in ea(r)$ ’ as ‘ a appears in r ’.

¹²The interested reader can begin by comparing our notion of closed terms in Definition 6.1.7, based on two simpler inductive definitions, with that used in [Fernández and Gabbay, 2007, Definition 68], based on a renamed variant of a term and an equality derivable in an extended freshness context. See also an inductive characterisation of closed terms in unpublished notes [Clouston, 2007].

- The free atoms of r are the atoms that can appear in $r\theta$ for some θ .

For instance, $ea(X) = \emptyset \neq supp(X) = fa(X)$.

This is an intuition, not a fact. $fa(r) = \bigcup_{\theta} ea(r\theta)$ is not true in general (but see Lemma 6.1.5). For instance in a signature with one base sort τ and no term formers, terms containing atoms simply do not populate the sort τ .

Recall the notion of *occurrences* $occ(r)$ from Definition 3.7.3.

NOTATION 6.1.3. Write $\pi \cdot occ(r) = \{\pi \cdot x \mid x \in occ(r)\}$. Also if $D = [d_1, \dots, d_n]$ and S is a permission set define $S \setminus D = S \setminus \{d_1, \dots, d_n\}$.

LEMMA 6.1.4. $ea(\pi \cdot r) = \pi \cdot ea(r)$ and $occ(\pi \cdot r) = \pi \cdot occ(r)$. In addition, $ea(r) \subseteq ea(r\theta)$.

Proof. By routine inductions on r . ■

LEMMA 6.1.5. $fa(r) = ea(r) \cup \bigcup \{supp(x) \mid x \in occ(r)\}$.

As an easy corollary using Lemma 3.2.5, $fa(r) = ea(r) \cup \bigcup \{supp(X) \setminus D \mid [D]X \in occ(r)\}$.

Proof. By a routine induction on r . We consider one case:

- *The case $[a]r$.* Suppose $fa(r) = ea(r) \cup \bigcup \{supp(x) \mid x \in occ(r)\}$. By definition $fa([a]r) = fa(r) \setminus \{a\}$, and $ea([a]r) = ea(r) \setminus \{a\}$ and $occ([a]r) = \{[a]x \mid x \in occ(r)\}$. The result follows by an easy sets calculation. ■

DEFINITION 6.1.6. Call r ***fa-functional*** when if $[D_1]X \in occ(r)$ and $[D_2]X \in occ(r)$ then $fa([D_1]X) = fa([D_2]X)$ (equivalently, when D_1 and D_2 contain the same atoms but not necessarily in the same order).

DEFINITION 6.1.7. Call r **closed** when r is *fa-functional* and $ea(r) = \emptyset$.

EXAMPLE 6.1.8.

- a is not closed (ea is non-empty).
- X is closed, so note that ‘closed’ does not mean ‘ $fv(r) = \emptyset$ ’. Our terminology is consistent with [Fernández and Gabbay, 2007] and the subsequent literature.
- $([a]X, X)$ is not closed (occ is not *fa-functional*).
- $[a](X, a)$ is closed.

LEMMA 6.1.9. Suppose $ea(r) = \emptyset$. Then $\pi \cdot (r\theta) = r\theta'$ if and only if $\pi \cdot ([D]X)\theta = ([D]X)\theta'$ for every $[D]X \in occ(r)$.

Proof. By a routine induction on r . ■

THEOREM 6.1.10. r is closed if and only if

$$\exists S. fa(r) \subseteq S \wedge \forall \pi, \theta. \pi \cdot fa(r\theta) \subseteq S \Rightarrow \exists \theta'. \pi \cdot (r\theta) = r\theta'.$$

Proof. Suppose there is a permission set $S \supseteq fa(r)$ such that if $\pi \cdot fa(r\theta) \subseteq S$ then there exists θ' such that $\pi \cdot (r\theta) = r\theta'$. There are two things to prove:

- *ea(r) is empty.* Suppose there exists $a \in ea(r)$. Pick $b \in S \setminus ea(r)$. By assumption taking $\theta = id$ there exists θ' such that $(b a) \cdot (r\theta) = r\theta'$. By Lemma 6.1.4 $ea((b a) \cdot r) = (b a) \cdot ea(r) \not\ni a$ and $a \in ea(r) \subseteq ea(r\theta')$, a contradiction.
- *occ(r) is fa-functional.* Consider $[D_1]X$ and $[D_2]X$ in $occ(r)$; choose D_i such that $D_i \cap fa(r) = \emptyset$ for $i = 1, 2$. Suppose there exists $a \in fa([D_2]X) \setminus fa([D_1]X)$, and choose any $b \in fa([D_1]X)$ (since $supp(X)$ is infinite and D_1 is finite, such a b exists). By Lemma 6.1.5 $a, b \in fa(r)$ so by assumption taking $\theta = id$ there exists θ' such that $(b a) \cdot r = r\theta'$. We proved above that $ea(r) = \emptyset$, so by Lemma 6.1.9 $(b a) \cdot [D_1]X = ([D_1]X)\theta$. By Lemma 3.2.8 a is free in the left-hand side, and by Lemma 3.4.11 a is not free in the right-hand side; a contradiction.

Suppose $occ(r)$ is *fa-functional* and $ea(r) = \emptyset$ and choose some permutation π and substitution θ .

If $occ(r) = \emptyset$ then by Lemma 6.1.5 $fa(r) = \emptyset$ so by Lemmas 3.2.9 and 3.4.12 $\pi \cdot (r\theta) = r$ and $r\theta' = r$, so there is nothing to prove.

Otherwise take $S = fa(r)$. For every element of in $occ(r)$ make a fixed but arbitrary choice of representation as $[D]X$ where the atoms in D are disjoint from the atoms in $nontriv(\pi)$. We take θ' to equivariantly extend this choice (Definition 2.5.4), so we map $\pi' \cdot X$ to $(\pi' \circ \pi) \cdot \theta(X)$ for the choice of representing X above, and otherwise to map Y to Y . Using Proposition 2.5.5 this is a substitution and $\pi \cdot ([D]X)\theta = ([D]X)\theta'$ for every $[D]X \in occ(r)$. We use Lemma 6.1.9. ■

6.2 Closed rewrite rules

DEFINITION 6.2.1. Call a rewrite rule $l \rightarrow m$ **closed** when (l, m) is closed.

EXAMPLE 6.2.2. Let R have one name sort ν , one base sort τ , two term-formers $\text{triv} : (\nu)\tau$ and $\text{abs} : ([\nu]\tau)\tau$, and rewrite rules

$$\text{triv}(a) \rightarrow \text{triv}(a) \quad \text{triv}(a) \rightarrow \text{triv}(b) \quad \text{abs}([a]X) \rightarrow X.$$

The terms $\text{triv}(a)$ and $\text{triv}(b)$ are not closed; the terms $\text{abs}([a]X)$ and X are closed. The terms $(\text{triv}(a), \text{triv}(a))$ and $(\text{triv}(a), \text{triv}(b))$ are not closed. The term $(\text{abs}([a]X), X)$ is closed if and only if $a \notin \text{supp}(X)$. So the first two rules are not closed and the third is closed if and only if $a \notin \text{supp}(X)$.

The rewrite rules of nrSUB and nrLAM in Example 5.1.3 are closed.

Recall that uniform rules have good properties like Theorems 5.4.7 and 5.5.7. Closed rules inherit these good properties, because:

THEOREM 6.2.3. If $R = (l \rightarrow m)$ is closed then it is uniform.

Proof. By assumption $fv(m) \subseteq fv(l)$. Also (l, m) is fa -functional; it follows that $occ(m) \subseteq occ(l)$. The result follows from Lemma 6.1.5. ■

LEMMA 6.2.4. Suppose r and l are terms and l is closed. Then

1. $\exists \pi, \theta. r = \pi \cdot (l\theta)$ implies
2. $\forall \pi. fa(r) \subseteq \pi \cdot fa(l) \Rightarrow \exists \theta. r = \pi \cdot (l\theta)$

Proof. Suppose $fa(r) \subseteq \pi \cdot fa(l)$ and $fa(r) \subseteq \pi' \cdot fa(l)$ and $r = \pi \cdot (l\theta)$.

We need a θ' such that $r = \pi' \cdot (l\theta')$. It follows from the above that $(\pi'^{-1} \circ \pi) \cdot fa(l\theta) \subseteq fa(l)$. We use Theorem 6.1.10. ■

THEOREM 6.2.5. If R is closed then \xrightarrow{R} can be checked as follows, where for simplicity we suppose $R = \{(l \rightarrow m)\}$:

1. We try to match r against $\pi \cdot l$ for some π such that $fa(r) \subseteq \pi \cdot fa(l)$, if such a π exists.
2. If we fail then by Lemma 6.2.4 we must fail for instantiating for any $\pi \cdot l$. We descend into subterms of r and repeat the previous step.

Whether step 1 of the algorithm above is decidable depends on the decidability of \mathbb{P} , \mathcal{X} , and \mathcal{C} ; obviously, if equality of the syntax is undecidable then matching will also be undecidable. So assuming that we have not been *silly*, closed rules are useful because we only need to compute one π and consider matching, rather than consider an equivariant matching problem.

To use the matching algorithm of Section 4, we need terms to satisfy condition 2 of Definition 4.1.1. So, we could forbid *shift* permutations altogether. The algorithm might reintroduce them but as noted in Remark 4.4.9, *shift* can be eliminated once a solution is found. Thus, if we care about decidability and not so much about infinite permutations—which was the case

e.g. in [Fernández *et al.*, 2004; Fernández and Gabbay, 2007]—then *shift* can be viewed as an internal mechanism of our unification/matching algorithm. However we have designed the mathematics to allow the possibility of exploring other, more liberal (and perhaps still decidable) choices, if we wish. More on this in [Gabbay, 2011c].

7 EQUALITY: (PERMISSIVE-)NOMINAL ALGEBRA

Permissive-nominal algebra has one judgement form: an equality $r = s$. This is just an unoriented nominal rewriting rule, so what makes algebra different from rewriting is not so much the judgement form as the properties we care about: instead of confluence and decidability, we primarily care about soundness and completeness. These are Theorems 7.4.6 and Corollary 7.5.12.

This different emphasis affects the axioms we write. The rewrites in Example 5.1.3 are designed to work on λ -terms without unknowns (since we expect to ‘evaluate’ closed terms using rewrites). The analogous axioms in Example 7.1.3 are designed to work also on open terms (since we expect to reason about arbitrary denotations).

Permissive-nominal algebra simplifies and streamlines the nominal algebra logic of [Gabbay and Mathijssen, 2009] (which was based on nominal terms). Essentially, these two logics do the same thing, but there are significant differences which we discuss in Subsection 7.7. Nominal Algebra (NA) was developed with Mathijssen and presented in [Gabbay, 2005; Gabbay and Mathijssen, 2006b; Gabbay and Mathijssen, 2007; Gabbay and Mathijssen, 2009]. It was used to axiomatise substitution, first-order logic, and the λ -calculus [Gabbay and Mathijssen, 2006a; Gabbay and Mathijssen, 2008a; Gabbay and Mathijssen, 2006c; Gabbay and Mathijssen, 2008c; Gabbay and Mathijssen, 2008b; Gabbay and Mathijssen, 2010]. The interest of these papers was not merely to write down the axioms—which all take advantage of atoms-abstraction to axiomatise various binding operators—but also to prove these axioms sound and complete. These proofs are not included here; see the presentations in [Gabbay and Mathijssen, 2008a; Gabbay and Mathijssen, 2008c; Gabbay and Mathijssen, 2010]. Or, to see a much more sophisticated instance of the same general idea, the reader can examine the permissive-nominal logic axiomatisation of arithmetic which is proved correct in the case study of Section 10.

7.1 Judgement form, axioms, theories

DEFINITION 7.1.1. A (nominal algebra) **equality judgement** is a pair $r = s$.

DEFINITION 7.1.2. A **theory** $\mathbb{T} = (\Sigma, Ax)$ is a pair of a signature Σ and

a possibly infinite set of *equality* judgements Ax in that signature; we call them the **axioms**.

EXAMPLE 7.1.3. Here are some example nominal algebra theories:

- **naSUB** axiomatises capture-avoiding substitution on, say, the λ -calculus.

Let Σ have a base sort τ and the following term-formers:

$$\text{sub} : ([\nu]\tau, \tau)\tau \quad \text{lam} : ([\nu]\tau)\tau \quad \text{app} : (\tau, \tau)\tau \quad \text{var} : (\nu)\tau$$

Axioms are as follows:

$$\begin{array}{lll} (\text{var} \mapsto) & \text{var}(a)[a \mapsto X] & = X \\ (\# \mapsto) & Y[a \mapsto X] & = Y \quad (a \notin \text{supp}(Y)) \\ (\text{f} \mapsto) & \text{f}(Y)[a \mapsto X] & = \text{f}(Y[a \mapsto X]) \quad (\text{f} \in \{\text{lam}, \text{app}, \text{var}, \text{sub}\}) \\ (\text{tup} \mapsto) & (X_1, \dots, X_n)[a \mapsto X] & = (X_1[a \mapsto X], \dots, X_n[a \mapsto X]) \\ (\text{abs} \mapsto) & ([b]Y)[a \mapsto X] & = [b](Y[a \mapsto X]) \quad (a \notin \text{supp}(Y)) \\ (\text{id} \mapsto) & Y[b \mapsto \text{var}(b)] & = Y \\ (\eta \mapsto) & [a]\text{sub}(Y, \text{var}(a)) & = Y \quad (a \notin \text{supp}(Y)) \end{array}$$

Here and in the next example we sugar $\text{sub}([a]r, t)$ to $r[a \mapsto t]$. Every permission set contains b and every permission set contains a except for $\text{supp}(Y)$, as indicated above. Sorts are filled in as appropriate.

This theory is studied in [Gabbay and Mathijssen, 2008a] (actually, a family of theories parameterised over the signature Σ). We prove the axioms above sound and complete for a syntactic model in which $Z[a := X]$ is interpreted as capture-avoiding substitution.

Note that the completeness result from Corollary 7.5.12 is valid but is also weaker: the extra work in [Gabbay and Mathijssen, 2008a] is to prove completeness with respect to a model in which $Z[a := X]$ is interpreted specifically by capture-avoiding substitution.

- **naLAM** extends the previous theory with two more axioms:

$$\begin{array}{lll} (\beta) & (\lambda[a]Y)X & = Y[a \mapsto X] \\ (\eta) & \lambda[a](Xa) & = X \quad (a \notin \text{supp}(X)) \end{array}$$

This theory is studied in [Gabbay and Mathijssen, 2010]. Analogously to **naSUB**, we prove the axioms sound and complete for a syntactic model where substitution *is* substitution and β - and η -conversion *are* β - and η -conversion.

REMARK 7.1.4. Compare and contrast Example 7.1.3 with Example 5.1.3. Clearly, one is an equality theory and another a rewrite theory, but we obtain a nominal algebra theory from Example 5.1.3 by replacing \rightarrow by $=$, and conversely we can replace $=$ with \rightarrow in Example 7.1.3.

So why are they different? They demonstrate different design priorities.

The rewrites in Example 5.1.3 are designed to operate on ground terms ($fv(r) = \emptyset$), following an intuition that rewriting is about ‘executing programs’. The equalities in Example 7.1.3 are designed to operate on possibly open terms, following an intuition that algebra is about models, not all of whose elements need be referenced by ground terms.

What we gain in deductive power we lose in computational properties. For instance, **nrSUB** is terminating whereas (an oriented version of) **naSUB** is not terminating, because explicit substitutions can ‘churn’ by distributing repeatedly over one another (this is essentially the idea behind Melliès’s counterexample in [Melliès, 1995]). On the other hand while the effect of $(\# \mapsto)$ from **naSUB** can be obtained on ground terms using the rules in **nrSUB**, by pushing the substitution down to the atoms, the rules of **nrSUB** are not deductively powerful enough to do this for open terms (or arbitrary models). More on this in [Gabbay and Mathijssen, 2008a; Gabbay and Mathijssen, 2010].

7.2 Derivable equality

DEFINITION 7.2.1. Suppose \mathbb{T} is a theory. **Derivable equality** $\mathbb{T} \vdash r = s$ is the least transitive reflexive symmetric relation such that for every $(r = s) \in \mathbb{T}$, position P , and substitution θ , if $sort(r) = sort(P)$ and $fa(r\theta) \cup fa(s\theta) \subseteq supp(P)$ (so that $P[l\theta]$ and $P[s\theta]$ are well-defined) then

$$\mathbb{T} \vdash P[r\theta] = P[s\theta].$$

REMARK 7.2.2. Definition 7.2.1 is rather compact; it might be useful to expand it a little. This is Figure 4, given in natural deduction style.

The reader familiar with nominal terms (see for instance Figure 2 of [Urban *et al.*, 2004]) should note of **(Cong3)** that we do not need to consider the case $[a]r = [b]s$, because α -equivalence is handled automatically for us by nominal abstract syntax. It is built in by Definition 2.4.8. In other words, thanks to how we set up our permissive-nominal terms syntax, we can always rename abstracted atoms so that they are equal. We noted an analogous point earlier on, in Remark 4.1.6.

LEMMA 7.2.3. Suppose $\mathbb{T} = (\Sigma, Ax)$ is a theory. Then:

- $\mathbb{T} \vdash a = b$ is impossible.
- $\mathbb{T} \vdash [a]r = [b]s$ if and only if $b \notin fa(r)$ and $(b\ a) \cdot r = s$.
- $\mathbb{T} \vdash (r_1, \dots, r_n) = (s_1, \dots, s_n)$ if and only if $\mathbb{T} \vdash r_i = s_i$ for $1 \leq i \leq n$.

Proof. In axiom $(r = s) \in Ax$, r and s must have base sort τ ; thus it is not possible to assert equalities between atoms, abstractions, or tuples (unless

$\frac{}{r = r} \text{ (Refl)}$	$\frac{r = s \quad s = t}{r = t} \text{ (Trans)}$
$\frac{r = s}{s = r} \text{ (Symm)}$	$\frac{r_1 = r'_1 \quad \dots \quad r_n = r'_n}{(r_1, \dots, r_n) = (r'_1, \dots, r'_n)} \text{ (Cong1)}$
$\frac{r = r'}{f(r) = f(r')} \text{ (Cong2)}$	$\frac{r = r'}{[a]r = [a]r'} \text{ (Cong3)}$
$\frac{((r=s) \in \mathbb{T})}{\pi \cdot (r\theta) = \pi \cdot (s\theta)} \text{ (Ax}_{r=s}\text{)}$	

Figure 4: Derivable entailment in Permissive-Nominal Algebra (PNA)

wrapped in a term-former and so injected into a base sort). The second part additionally uses Lemma 2.4.9. ■

LEMMA 7.2.4. Suppose $\mathbb{T} \vdash r = s$. Then:

1. $\mathbb{T} \vdash \pi \cdot r = \pi \cdot s$.
2. $\mathbb{T} \vdash r\theta = s\theta$.

Proof. Both parts are by a routine argument on derivations. We consider one case:

- *The case $(r' = s') \in \mathbb{T}$ and $r = P[r'\theta']$ and $s = P[s'\theta']$ and $P = (t, X)$.*
 For the first part we use a position $(\pi \cdot t, X)$.
 For the second part we consider a position $P' = (t(\theta - X), X)$ and consider $P'[r'\theta'\theta]$ and $S'[r'\theta'\theta]$ ($\theta - X$ defined in Definition 4.3.3). It is not hard to check that $P'[r'\theta'\theta] = P[r'\theta']\theta$ and $P'[s'\theta'\theta] = P[s'\theta']\theta$, and the result follows. ■

7.3 Interpretation of signatures and terms

DEFINITION 7.3.1. Suppose $(\mathcal{A}, \mathcal{B})$ is a sort-signature (Definition 3.1.1).

An **interpretation** \mathcal{I} for $(\mathcal{A}, \mathcal{B})$ consists of an assignment of a nonempty permissive-nominal set $\llbracket \alpha \rrbracket^{\mathcal{I}}$ to each sort α in $(\mathcal{A}, \mathcal{B})$, along with equivariant maps

- for each $\nu \in \mathcal{A}$ an equivariant and injective map $\mathbb{A}_\nu \rightarrow \llbracket \nu \rrbracket^{\mathcal{I}}$ which we write a_ν ,

- for each $\nu \in \mathcal{A}$ and α an equivariant and injective map $[\mathbb{A}_\nu][[\alpha]]^\mathcal{S} \rightarrow [[[\nu]\alpha]]^\mathcal{S}$ which we write $[a]^\mathcal{S}x$, and
- for each α_i for $1 \leq i \leq n$ an equivariant and injective map $\Pi_i[[\alpha_i]]^\mathcal{S} \rightarrow [[(\alpha_1, \dots, \alpha_n)]]^\mathcal{S}$ which we write $(x_1, \dots, x_n)^\mathcal{S}$.

DEFINITION 7.3.2. Suppose $\Sigma = (\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{F}, ar)$ is a signature (Definition 3.1.5).

An **interpretation** \mathcal{S} for Σ , or Σ -**algebra**, consists of the following data:

- An interpretation for the sort-signature $(\mathcal{A}, \mathcal{B})$ (Definition 7.3.1).
- For every $f \in \mathcal{F}$ with $ar(f) = (\alpha)\tau$ an equivariant function $f^\mathcal{S}$ from $[[\alpha]]^\mathcal{S}$ to $[[\tau]]^\mathcal{S}$.
- An equivariant assignment from $C \in \mathcal{C}$ to $C^\mathcal{S} \in [[sort(C)]]^\mathcal{S}$. (That is, $(\pi \cdot C)^\mathcal{S} = \pi \cdot (C^\mathcal{S})$.)

DEFINITION 7.3.3. Suppose \mathcal{S} is a Σ -algebra. A **valuation** ς to \mathcal{S} is an equivariant function on unknowns \mathcal{X} such that for each unknown X , $\varsigma(X) \in [[sort(X)]]^\mathcal{S}$.

ς will range over valuations.

DEFINITION 7.3.4. Suppose \mathcal{S} is a Σ -algebra. Suppose ς is a valuation to \mathcal{S} .

Extend \mathcal{S} to an **interpretation** on terms $[[r]]_\varsigma^\mathcal{S}$ (where of course r is a term in the signature Σ) by:

$\begin{aligned} [[a]]_\varsigma^\mathcal{S} &= a^\mathcal{S} & [[f(r)]]_\varsigma^\mathcal{S} &= f^\mathcal{S}([[r]]_\varsigma^\mathcal{S}) \\ [[C]]_\varsigma^\mathcal{S} &= C^\mathcal{S} & [[(r_1, \dots, r_n)]]_\varsigma^\mathcal{S} &= ([[r_1]]_\varsigma^\mathcal{S}, \dots, [[r_n]]_\varsigma^\mathcal{S})^\mathcal{S} \\ [[X]]_\varsigma^\mathcal{S} &= \varsigma(X) & [[a]r]_\varsigma^\mathcal{S} &= [a]^\mathcal{S}[[r]]_\varsigma^\mathcal{S} \end{aligned}$
--

Lemma 7.3.5 is a basic sanity check and an important soundness result:

LEMMA 7.3.5. If $r : \alpha$ then $[[r]]_\varsigma^\mathcal{S} \in [[\alpha]]^\mathcal{S}$.

Proof. By a routine induction on r . ■

LEMMA 7.3.6. $\pi \cdot [[r]]_\varsigma^\mathcal{S} = [[\pi \cdot r]]_\varsigma^\mathcal{S}$.

Proof. By a routine induction on r . We consider one case:

- *The case X .* By Definition 7.3.4 $[[X]]_\varsigma^\mathcal{S} = \varsigma(X)$. Therefore $\pi \cdot [[X]]_\varsigma^\mathcal{S} = \pi \cdot \varsigma(X)$. By assumption $\pi \cdot \varsigma(X) = \varsigma(\pi \cdot X) = [[\pi \cdot X]]_\varsigma^\mathcal{S}$. ■

LEMMA 7.3.7. $supp([[r]]_\varsigma^\mathcal{S}) \subseteq fa(r)$.

Proof. From Lemmas 2.3.3 and 7.3.6. ■

7.4 Models and soundness

DEFINITION 7.4.1. For a theory $\mathbb{T} = (\Sigma, Ax)$ and interpretation \mathcal{I} of \mathbb{T} call $(r = s)$ **valid** in \mathcal{I} when $\llbracket r \rrbracket_{\zeta}^{\mathcal{I}} = \llbracket s \rrbracket_{\zeta}^{\mathcal{I}}$ for every valuation ζ to \mathcal{I} . Call \mathcal{I} a **model** of \mathbb{T} when every axiom $(r = s) \in Ax$ is valid in \mathcal{I} . Write $\mathbb{T} \models r = s$ when $(r = s)$ is valid in every model of \mathbb{T} .

LEMMA 7.4.2. If $\zeta(X) = \zeta'(X)$ for all $X \in fv(r)$ then $\llbracket r \rrbracket_{\zeta}^{\mathcal{I}} = \llbracket r \rrbracket_{\zeta'}^{\mathcal{I}}$.

Proof. By a routine induction on r . ■

DEFINITION 7.4.3. Suppose ζ is a valuation to \mathcal{I} . Suppose X is an unknown and $x \in \llbracket sort(X) \rrbracket_{\zeta}^{\mathcal{I}}$ is such that $supp(x) \subseteq supp(X)$. Define a function $\zeta[X:=x]$ by

$$(\zeta[X:=x])(\pi \cdot X) = \pi \cdot x \quad \text{and} \quad (\zeta[X:=x])(Y) = \zeta(Y) \quad \text{all other } Y$$

LEMMA 7.4.4. $\zeta[X:=x]$ in Definition 7.4.3 is well-defined and a valuation to \mathcal{I} .

Proof. As that of Proposition 2.5.5. ■

LEMMA 7.4.5. $\llbracket r \rrbracket_{\zeta[X:=t]}^{\mathcal{I}} = \llbracket r[X:=t] \rrbracket_{\zeta}^{\mathcal{I}}$. As corollaries we have:

1. If $\llbracket r \rrbracket_{\zeta}^{\mathcal{I}} = \llbracket s \rrbracket_{\zeta}^{\mathcal{I}}$ then $\llbracket P[r] \rrbracket_{\zeta}^{\mathcal{I}} = \llbracket P[s] \rrbracket_{\zeta}^{\mathcal{I}}$.
2. If $\llbracket r \rrbracket_{\zeta}^{\mathcal{I}} = \llbracket s \rrbracket_{\zeta}^{\mathcal{I}}$ then $\llbracket r\theta \rrbracket_{\zeta}^{\mathcal{I}} = \llbracket s\theta \rrbracket_{\zeta}^{\mathcal{I}}$.

Proof. By a routine induction on the definition of $\llbracket r \rrbracket_{\zeta}^{\mathcal{I}}$. We consider one case:

- *The case of $\llbracket \pi \cdot X \rrbracket_{\zeta[X:=t]}^{\mathcal{I}}$.* We reason as follows:

$$\begin{aligned} \llbracket \pi \cdot X \rrbracket_{\zeta[X:=t]}^{\mathcal{I}} &= \pi \cdot \llbracket t \rrbracket_{\zeta}^{\mathcal{I}} && \text{Definition 7.3.4} \\ &= \llbracket \pi \cdot t \rrbracket_{\zeta}^{\mathcal{I}} && \text{Lemma 7.3.6} \\ &= \llbracket (\pi \cdot X)[X:=t] \rrbracket_{\zeta}^{\mathcal{I}} && \text{Definition 3.4.8.} \end{aligned}$$

For the two corollaries we reason as follows:

1. By definition where $P = (t, X)$, $P[r] = t[X:=r]$ and $P[s] = t[X:=s]$. Using the assumptions,

$$\llbracket t[X:=r] \rrbracket_{\zeta}^{\mathcal{I}} = \llbracket t \rrbracket_{\zeta[X:=r]}^{\mathcal{I}} = \llbracket t \rrbracket_{\zeta[X:=s]}^{\mathcal{I}} = \llbracket t[X:=s] \rrbracket_{\zeta}^{\mathcal{I}}.$$

2. It is a fact of syntax that $fv(r)$ and $fv(s)$ are finite. Using Lemma 3.4.12 we may represent the effect of θ on r and s as a sequence of atomic substitutions (Definition 3.4.5). The result follows. ■

THEOREM 7.4.6 (Soundness). For any $\mathbb{T} = (\Sigma, Ax)$ if $\mathbb{T} \vdash r = s$ then $\mathbb{T} \models r = s$.

Proof. Let \mathcal{I} be a model of \mathbb{T} and ζ be a valuation to \mathcal{I} .

Identity in the denotation is reflexive, transitive, and symmetric so it suffices to check the theorem for axioms. That is, suppose $(r = s) \in Ax$ and assume a position P and substitution θ such that $sort(r) = sort(P)$ and $fa(r\theta) \cup fa(s\theta) \subseteq supp(P)$. We must show that $\llbracket P[r\theta] \rrbracket_\zeta^\mathcal{I} = \llbracket P[s\theta] \rrbracket_\zeta^\mathcal{I}$.

\mathcal{I} is a model so $\llbracket r \rrbracket_\zeta^\mathcal{I} = \llbracket s \rrbracket_\zeta^\mathcal{I}$. We use parts 1 and 2 of Lemma 7.4.5. ■

7.5 Free term models and completeness

In this subsection fix a signature Σ and a theory $\mathbb{T} = (\Sigma, Ax)$.

The proof of completeness follows a standard method: we construct a model out of syntax in which by construction two terms denote equal elements if and only if they are derivably equal.

The subtlety occurs in Lemma 7.5.9. We want to eliminate ζ in $\llbracket r \rrbracket_\zeta^{\mathcal{I}(\mathbb{T})}$ by converting it into a substitution θ . This ‘should’ be easy, since for each X , $\zeta(X)$ is a provably equivalent class of terms. We need only choose some representative term in $\zeta(X)$ for each X and set $\theta(X)$ to be that representative.

If we are naive in our construction then this could be impossible, as outlined in Example 7.5.10: there might be ‘too many atoms’ in the available representatives. We enrich our syntax with ‘enough’ extra constant symbols, to guarantee ‘enough’ representatives of every element of the model. Nominal algebra without the constant symbols is complete for the same semantics, but the proof would be more complex.

DEFINITION 7.5.1. For each sort α in Σ define $[r]_\mathbb{T}$ and $\mathcal{F}(\mathbb{T})_\alpha$ by

$$\begin{aligned} [r]_\mathbb{T} &= \{r' : \alpha \mid \mathbb{T} \vdash r = r'\} & (r : \alpha) \\ \mathcal{F}(\mathbb{T})_\alpha &= \{[r]_\mathbb{T} \mid r : \alpha\}. \end{aligned}$$

Make each $\mathcal{F}(\mathbb{T})_\alpha$ into a permissive-nominal set by giving it a permutation action

$$\pi \cdot [r]_\mathbb{T} = [\pi \cdot r]_\mathbb{T}.$$

$\mathcal{F}(\mathbb{T})$ stands for ‘Free terms in the signature of \mathbb{T} , up to derivable equality in \mathbb{T} ’. Lemmas 7.5.2 and 7.5.3 relate permutation and support to the natural notions from nominal sets:

LEMMA 7.5.2. The permutation action on $[r]_{\mathbb{T}}$ is pointwise on $[r]_{\mathbb{T}}$ as a set: that is, $\pi \cdot [r]_{\mathbb{T}} = \{\pi \cdot r' \mid r' \in [r]_{\mathbb{T}}\}$.

Proof. From Definition 7.5.1 and Lemma 7.2.4. ■

LEMMA 7.5.3. $\text{supp}([r]_{\mathbb{T}}) \subseteq \text{fa}(r)$.

Proof. From Definition 7.5.1 and Lemma 2.3.3. ■

DEFINITION 7.5.4. We construct the **free term interpretation** $\mathcal{F}(\mathbb{T})$ of \mathbb{T} as follows:

- Take $\mathcal{F}(\mathbb{T})_{\alpha}$ as in Definition 7.5.1.
- $a^{\mathcal{F}(\mathbb{T})} = [a]_{\mathbb{T}}$, $[a]^{\mathcal{F}(\mathbb{T})}[r]_{\mathbb{T}} = [[a]r]_{\mathbb{T}}$, and $([r_1]_{\mathbb{T}}, \dots, [r_n]_{\mathbb{T}})^{\mathcal{F}(\mathbb{T})} = [(r_1, \dots, r_n)]_{\mathbb{T}}$.
- $f^{\mathcal{F}(\mathbb{T})}([r]_{\mathbb{T}}) = [f(r)]_{\mathbb{T}}$ for each term-former $f : (\alpha)\tau$ in Σ and each $r : \alpha$.
- $C^{\mathcal{F}(\mathbb{T})} = [C]_{\mathbb{T}}$ for each constant in Σ .

LEMMA 7.5.5. Definition 7.5.4 is well-defined and is an interpretation. That is:

- The choice of representative of $[r]_{\mathbb{T}}$ does not matter in any of the clauses.
- The choice of abstracted atom in the clause for $[a]^{\mathcal{F}(\mathbb{T})}[r]_{\mathbb{T}}$ does not matter.
- The maps $a^{\mathcal{F}(\mathbb{T})}$, $[a]^{\mathcal{F}(\mathbb{T})}[r]_{\mathbb{T}}$, and $([r_1]_{\mathbb{T}}, \dots, [r_n]_{\mathbb{T}})^{\mathcal{F}(\mathbb{T})}$ are injective.

Proof. The first part follows by congruence properties of derivable equality. The second part additionally uses Lemmas 2.4.10 and 2.4.11. The third part uses Lemma 7.2.3. ■

DEFINITION 7.5.6. Define a theory $\mathbb{T}^+ = (\Sigma^+, Ax^+)$ to be equal to \mathbb{T} except that we adjoin $\bigcup_{\tau} \mathcal{F}(\mathbb{T})_{\tau}$ to the set of constants in Σ , and we add axioms equating r with $[r]_{\mathbb{T}}$ in Ax .

That is, for every $r : \tau$ there is a constant $C_r = [r]_{\mathbb{T}} \in \Sigma^+$, and an axiom $(C_r = r) \in \mathcal{F}(\mathbb{T})^+$.

LEMMA 7.5.7. $\mathcal{F}(\mathbb{T})$ extends to an interpretation $\mathcal{F}(\mathbb{T})^+$ of \mathbb{T}^+ , where for each $r : \tau$ we take $C_r^{\mathcal{F}(\mathbb{T})^+} = [r]_{\mathbb{T}}$. Furthermore, $\mathcal{F}(\mathbb{T})^+$ is a model of \mathbb{T}^+ .

DEFINITION 7.5.8. Write ς_{id} for the valuation to $\mathcal{F}(\mathbb{T})$ mapping each X to $C_X = [X]_{\mathbb{T}}$.

LEMMA 7.5.9. For every valuation ς to $\mathcal{F}(\mathbb{T})$ there exists a substitution θ in \mathbb{T}^+ such that $\llbracket r \rrbracket_{\varsigma}^{\mathcal{F}(\mathbb{T})} = \llbracket r\theta \rrbracket_{\varsigma_{id}}^{\mathcal{F}(\mathbb{T})^+}$.

Proof. For each orbit $x \in |\text{orb}(\mathcal{X})|$ choose a representative $X \in x$. Define θ by $\theta(\pi \cdot X) = \pi \cdot C_X$. Recall that $C_X = [X]_{\top}$ and by Lemma 7.5.3 $\text{supp}([X]_{\top}) \subseteq \text{supp}(X)$. By Proposition 2.5.5 θ is well-defined and is a substitution

It is not hard to check by induction on r that $\llbracket r \rrbracket_{\zeta}^{\mathcal{T}} = \llbracket r\theta \rrbracket_{\zeta_{\text{id}}}^{\mathcal{T}^+}$. ■

EXAMPLE 7.5.10. To see why Lemma 7.5.9 is non-trivial and how \mathbb{T}^+ helps, suppose \mathbb{T} has one name sort ν , two base sorts τ and τ' , one term-former $\text{abs} : (\nu, \tau)\tau'$, and one axiom $\text{abs}(b, (b \ a) \cdot X) = \text{abs}(a, X)$ where $a \in \text{supp}(X)$ and $b \notin \text{supp}(X)$.

Then it is a fact that there is no $r \in [\text{abs}(a, X)]_{\top}$ such that $\text{fa}(r) \subseteq \text{supp}([\text{abs}(a, X)]_{\top})$ and it follows that there is no θ such that $\llbracket X' \rrbracket_{[X':=[\text{abs}(a, X)]_{\top}]}^{\mathcal{T}} = \llbracket X'\theta \rrbracket_{\zeta_{\text{id}}}^{\mathcal{T}^+}$ (recall that substitutions must be equivariant).

THEOREM 7.5.11. $\mathcal{F}(\mathbb{T})$ is a model of \mathbb{T} .

Proof. We must show that $\mathcal{F}(\mathbb{T})$ validates the axioms.

Suppose $(r = s) \in Ax$. Suppose ζ is a valuation to $\mathcal{F}(\mathbb{T})$. We must show that $\llbracket r \rrbracket_{\zeta}^{\mathcal{T}} = \llbracket s \rrbracket_{\zeta}^{\mathcal{T}}$.

By Lemma 7.5.9 there exists θ to \mathbb{T}^+ such that $\llbracket r \rrbracket_{\zeta}^{\mathcal{T}} = \llbracket r\theta \rrbracket_{\zeta_{\text{id}}}^{\mathcal{T}^+}$ and $\llbracket s \rrbracket_{\zeta}^{\mathcal{T}} = \llbracket s\theta \rrbracket_{\zeta_{\text{id}}}^{\mathcal{T}^+}$.

By assumption $\mathbb{T}^+ \vdash r\theta = s\theta$. By Lemma 7.5.7, $\llbracket r\theta \rrbracket_{\zeta_{\text{id}}}^{\mathcal{T}^+} = \llbracket s\theta \rrbracket_{\zeta_{\text{id}}}^{\mathcal{T}^+}$. The result follows. ■

COROLLARY 7.5.12 (Completeness). If $\mathbb{T} \models r = s$ then $\mathbb{T} \vdash r = s$.

Proof. Suppose $\mathbb{T} \models r = s$. By Theorem 7.5.11 $\llbracket r \rrbracket_{\zeta_{\text{id}}}^{\mathcal{T}} = \llbracket s \rrbracket_{\zeta_{\text{id}}}^{\mathcal{T}}$ (ζ_{id} is defined in Definition 7.5.8).

It is not hard to prove by induction that $\llbracket r \rrbracket_{\zeta_{\text{id}}}^{\mathcal{T}} = [r]_{\top}$ and $\llbracket s \rrbracket_{\zeta_{\text{id}}}^{\mathcal{T}} = [s]_{\top}$. It follows that $\mathbb{T} \vdash r = s$ as required. ■

7.6 Freshness

Nominal terms freshness conditions $a \# X$ and $a \# r$ from [Urban *et al.*, 2004] correspond in this paper to ‘free atoms of’ $a \notin \text{supp}(X)$ and $a \notin \text{fa}(r)$. See Notation 3.2.7 and Lemma 3.2.5. Call this **syntactic** freshness.

Nominal sets freshness $a \notin \text{supp}(\llbracket r \rrbracket)$ is a distinct notion which can be expressed using equality; call this **semantic** freshness. The two are not identical, but they are connected in various ways which we briefly explore.

Proposition 7.6.1 corresponds to Theorem 5.5 from [Gabbay and Mathijssen, 2007] and Lemma 4.51 from [Gabbay and Mathijssen, 2009]:

PROPOSITION 7.6.1. Suppose $b \notin \text{fa}(r)$.

Then $\mathbb{T} \vdash (b \ a) \cdot r = r$ if and only if for every model \mathcal{S} of \mathbb{T} and valuation ζ to \mathcal{S} , $a \notin \text{supp}(\llbracket r \rrbracket_{\zeta}^{\mathcal{T}})$.

Proof. By Theorem 7.4.6 and Corollary 7.5.12 $\top \vdash (b\ a)\cdot r = r$ if and only if $\top \models (b\ a)\cdot r = r$, which unpacking definitions means that for every \mathcal{I} and ζ , $\llbracket (b\ a)\cdot r \rrbracket_{\zeta}^{\mathcal{I}} = \llbracket r \rrbracket_{\zeta}^{\mathcal{I}}$. By Lemma 7.3.6 $\llbracket (b\ a)\cdot r \rrbracket_{\zeta}^{\mathcal{I}} = (b\ a)\cdot \llbracket r \rrbracket_{\zeta}^{\mathcal{I}}$, and by Lemma 7.3.7 $b \notin \text{supp}(\llbracket r \rrbracket_{\zeta}^{\mathcal{I}})$. The result follows by Corollary 2.2.7. ■

Lemmas 7.5.3 (and also Lemma 7.3.7) express that syntactic freshness implies semantic freshness. A partial converse is Proposition 7.6.3, which is based on a technical property of nominal sets:

LEMMA 7.6.2. Suppose X is a nominal set and $U \subseteq |X|$ is finitely-supported (so $U \in |\text{pow}(X)|$ from Example 2.2.5) and nonempty.

Then if $a\#U$ then there exists some $x \in U$ with $a\#x$.

Proof. U is nonempty so choose any $x' \in U$. Choose fresh b (so $b \notin \text{supp}(U) \cup \text{supp}(x')$) and set $x = (b\ a)\cdot x'$. By the definition of support $(b\ a)\cdot U = U$. By the pointwise action (Example 2.2.5) $x \in U$. By Lemma 2.2.6 $a \notin \text{supp}(x)$. ■

PROPOSITION 7.6.3. $a\#[r]_{\top}$ implies there exists some r' such that $\top \vdash r = r'$ and $a \notin \text{fa}(r')$.

Proof. By Lemmas 7.5.2 and Lemma 7.6.2. ■

7.7 Design of nominal algebra

When we designed nominal algebra, we encountered two design decisions: whether to include freshness axioms, and whether to include atoms-abstraction as primitive.

We were aware that these decisions do not matter for expressivity, because of the following two equalities from [Gabbay and Pitts, 2001]:

$$\mathbb{I}b.a\#x \Leftrightarrow (b\ a)\cdot x = x \quad \mathbb{I}b.[b](b\ a)\cdot x = [a]x$$

\mathbb{I} is the *new*-quantifier meaning ‘for some/any fresh atom’ [Gabbay and Pitts, 2001; Gabbay, 2011b].¹³

\mathbb{I} does not care which fresh atom we choose (the some/any property [Gabbay, 2011b, Theorem 6.5]). So, we do not have to be exact about $\text{supp}(x)$ when we choose some fresh b ; *any* will do. Thus e.g. Proposition 7.6.1 is an ‘if and only if’ even though we chose $b \notin \text{fa}(r)$ (syntactic freshness) instead of $b \notin \text{supp}(\llbracket r \rrbracket)$ (semantic freshness), and it may be that $\text{supp}(\llbracket r \rrbracket) \subsetneq \text{fa}(r)$.

So nominal algebra gets the power of \mathbb{I} -quantifier style freshness, from its syntactic freshnesses side-conditions. In this paper we do the same thing using permission-sets; see Subsection 11.1.

¹³In words: ‘ a is fresh for x if for some/any fresh b , $(b\ a)\cdot x = x$ ’ and ‘for some/any fresh b , $[b](b\ a)\cdot x = [a]x$ ’.

We left out freshness axioms but kept atoms-abstraction because we knew readers would expect to see it; a design choice inherited by the nominal algebra of this paper. Note that atoms-abstraction is isolated by the sort system, so that if there are no term-formers injecting atoms-abstraction into base sorts then it cannot interact with the rest of the logic.

Clouston and Pitts develop a notion of *nominal Lawvere theory* [Clouston, 2009; Clouston, 2011]. In this chapter we have used a nominal sets semantics; it would be interesting to try to translate this to nominal Lawvere theories.

The nominal algebra of this paper differs from the nominal algebra of [Gabbay and Mathijssen, 2009] in the following respects:

- The system here is sorted, the system in [Gabbay and Mathijssen, 2009] is not.
- We use permissive-nominal terms and semantics here, and ‘vanilla’ nominal terms and nominal sets in [Gabbay and Mathijssen, 2009]. That is, the logic here is *permissive-nominal algebra*. Freshness conditions $a\#X$ and $a\#r$ translate to $a \notin \text{supp}(X)$ and $a \notin \text{fa}(r)$ here.
- Axioms are exactly equalities, with no freshness contexts: permission sets play this role instead.
- The syntax here admits non-equivariant constant symbols, that of [Gabbay and Mathijssen, 2009] does not. That does not matter if we are using finitely-supported models (as is the case in [Gabbay and Mathijssen, 2009]) because finite non-equivariance can be emulated using term-formers applied to finitely many atoms. Here, elements can have infinite support, which cannot be emulated using (finite) equivariant term-formers.
- The syntax here admits the possibility of unknowns with empty support ranging over closed elements (so it includes the $\bullet t$ freshness constraint of [Fernández and Gabbay, 2007, Subsection 9.2]), unknowns with finite support ranging over finitely-supported elements, unknowns with support equal to a permission set, and whatever else we can imagine in-between.
- The development is parameterised over the set of unknowns \mathcal{X} and *also* the group of permutations \mathbb{P} . In particular we admit (but do not insist on) the possibility of infinite permutations, including the *shift*-permutations considered in Subsection 3.6.
- Substitutions and valuations are—rather elegantly—treated as equivariant functions on \mathcal{X} the set of unknowns.

In spite of these many differences, the spirit of the proofs remains the same. The details become simpler, and in particular the non-equivariant constants make construction of the free term model easier.¹⁴

¹⁴In [Gabbay and Mathijssen, 2009] to build the free term model we enriched syntax

8 THE NOMINAL HSP THEOREM

The HSP theorem states that a class of Σ -algebras is equational if and only if it is closed under Homomorphism, Subobject, and Product. Definitions follow below, and the main result is Theorem 8.7.3.

The result was first proved for the case of ‘ordinary’ algebra (using first-order terms and not over nominal sets) by Birkhoff [Birkhoff, 1935]. It is also called *Birkhoff’s theorem* [Burris and Sankappanavar, 1981, Theorem 11.12]. We prefer ‘HSP’ since this is more descriptive and Birkhoff’s name is attached to several other results.

The result was first proved for nominal algebra by the author [Gabbay, 2009], and an elegant alternative proof was provided by Kurz and Petrişan [Kurz and Petrişan, 2010]. The new proof presented here is also rather short.

HSP was interesting for two reasons: first, it is not obvious that nominal algebra is a true logic of equality, because of the freshness side-conditions which give the nominal algebra as presented e.g. in [Gabbay and Mathijssen, 2009] or in Mathijssen’s thesis [Mathijssen, 2007] a *prima facie* flavour of conditional equalities. The HSP result holding for nominal algebra was a way of making formal that this *is* a logic of equality.

The use of permission sets to phrase the logic entirely in terms of equality (freshness migrates to the types, as permission sets) is a step forward from this point of view: the nominal algebra of this paper is more *visibly* an equational logic. Still, HSP along with soundness and completeness (Theorem 7.4.6 and Corollary 7.5.12) form a triumvirate of results of interest for an algebraic reasoning framework.

The proofs here are much shorter and clearer than those of [Gabbay, 2009]—and the final result is strictly stronger than [Gabbay, 2009; Kurz and Petrişan, 2010], which actually proved an HSPA theorem that a class of Σ -algebras is equational if and only if it is closed under Homomorphism, Subobject, Product, and Atoms-abstraction.

That is, we have dropped the ‘atoms-abstraction’ from the closure conditions. How can this be? The use of permission sets gives us finer control over the support of valuations; we needed atoms-abstraction in the proof of [Gabbay, 2009, Theorem 9.8] to eliminate ‘extra’ atoms introduced by a valuation ς —‘extra’ relative to the freshness information in a freshness context Δ . Here, because freshness contexts/permission sets are fixed, this cannot happen.

with n -ary term-formers applied to atoms. This idea goes back to a completeness proof in [Gabbay, 2007a].

8.1 Algebra homomorphisms

DEFINITION 8.1.1. Suppose $\Sigma = (\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{X}, \mathcal{F}, ar)$ is a signature and suppose \mathcal{X} and \mathcal{Y} are interpretations of Σ . A Σ -**homomorphism** Θ from \mathcal{X} to \mathcal{Y} is a family of equivariant functions Θ_α from $[[\alpha]]^{\mathcal{X}}$ to $[[\alpha]]^{\mathcal{Y}}$ for each sort α in the sort-signature $(\mathcal{A}, \mathcal{B})$ such that:

- $\Theta_\nu(a^{\mathcal{X}}) = a^{\mathcal{Y}}$.
- $\Theta_{(\alpha_1, \dots, \alpha_n)}(x_1, \dots, x_n)^{\mathcal{X}} = (\Theta_{\alpha_1}(x_1), \dots, \Theta_{\alpha_n}(x_n))^{\mathcal{Y}}$.
- $\Theta_{[\nu]\alpha}([a]^{\mathcal{X}} x) = [a]^{\mathcal{Y}} \Theta_\alpha(x)$.
- $\Theta_\tau(f^{\mathcal{X}}(x)) = f^{\mathcal{Y}}(\Theta_\alpha(x))$ where $f : (\alpha)\tau$ is in \mathcal{F} .

DEFINITION 8.1.2. Call \mathcal{Y} a **homomorphic image** of \mathcal{X} when there is a Σ -homomorphism Θ from \mathcal{X} to \mathcal{Y} such that Θ_α is surjective for every sort α in $(\mathcal{A}, \mathcal{B})$.

Call Θ **injective** when Θ_α is injective for every sort α in $(\mathcal{A}, \mathcal{B})$.

LEMMA 8.1.3. Suppose Σ is a signature and \mathcal{X} and \mathcal{Y} are Σ -algebras. Suppose Θ is a Σ -algebra homomorphism from \mathcal{X} to \mathcal{Y} .

Suppose that ς is a valuation to \mathcal{X} . Define $\Theta(\varsigma)$ a valuation to \mathcal{Y} by $\Theta(\varsigma)(X) = \Theta_{\text{sort}(X)}(\varsigma(X))$ for every $X \in \mathcal{X}$.

Then for every $r : \alpha$, $\Theta_\alpha([[r]]_\varsigma^{\mathcal{X}}) = [[r]]_{\Theta(\varsigma)}^{\mathcal{Y}}$.

Proof. By an easy induction on r . ■

LEMMA 8.1.4. Suppose Σ is a signature and $\mathbb{T} = (\Sigma, Ax)$ is a theory. Suppose \mathcal{X} and \mathcal{Y} are Σ -algebras and \mathcal{Y} is a homomorphic image of \mathcal{X} under Θ .

Then if \mathcal{X} is a model of \mathbb{T} , then so is \mathcal{Y} .

Proof. Choose $(r = s) \in Ax$ and a valuation ς to \mathcal{Y} . It suffices to show that $[[r]]_\varsigma^{\mathcal{Y}} = [[s]]_\varsigma^{\mathcal{Y}}$.

We construct a valuation ς' to \mathcal{X} as an equivariant extension (Definition 2.5.4) of the following data. For each unknown $X : \alpha$ let $\mathcal{X}_X = \{x \in |\mathcal{X}_\alpha| \mid \Theta(x) = \varsigma(X)\}$. We construct a valuation ς' to \mathcal{X} by for each orbit and representative $X \in \text{orb}(X) \in \mathcal{X}$ setting $\varsigma'(X) = x$ for some choice of $x \in \mathcal{X}_X$.

By construction $\Theta\varsigma' = \varsigma$. By assumption $[[r]]_{\varsigma'}^{\mathcal{X}} = [[s]]_{\varsigma'}^{\mathcal{X}}$. We apply Θ to both sides and use Lemma 8.1.3. ■

8.2 Subalgebras

DEFINITION 8.2.1. For Σ -algebras \mathcal{X} and \mathcal{Y} , call \mathcal{X} a **subalgebra** of \mathcal{Y} when:

- $|\tau^{\mathcal{X}}| \subseteq |\tau^{\mathcal{Y}}|$ for every $\tau \in \mathcal{B}$.

- The subset inclusion maps form a Σ -algebra homomorphism (Definition 8.1.1).¹⁵

LEMMA 8.2.2. For Σ -algebras \mathcal{X} , \mathcal{Y} and a theory $\mathsf{T} = (\Sigma, Ax)$, if \mathcal{Y} is a model of T and \mathcal{X} is a subalgebra of \mathcal{Y} then \mathcal{X} is a model of T .

8.3 Products

DEFINITION 8.3.1. Let I be a (possibly countably infinite) indexing set and $(\mathcal{X}_i)_{i \in I}$ be an I -indexed collection of Σ -algebras. The **product algebra** $\prod_{i \in I} \mathcal{X}_i$ is the Σ -algebra such that:

- For each α in Σ , $\alpha^{\prod_{i \in I} \mathcal{X}_i} = \prod_{i \in I} \alpha^{\mathcal{X}_i}$ as defined in Definition 2.4.5.
- The i th projection map to \mathcal{X}_i is a Σ -algebra homomorphism for every $i \in I$.

LEMMA 8.3.2. For any I -indexed collection of Σ -algebras $(\mathcal{X}_i)_{i \in I}$, if \mathcal{X}_i is a model of $\mathsf{T} = (\Sigma, Ax)$ for every $i \in I$ then so is $\prod_{i \in I} \mathcal{X}_i$.

Proof. Suppose $(r = s) \in Ax$. Suppose ς is a valuation to $\prod_{i \in I} \mathcal{X}_i$. For each $i \in I$ we obtain a valuation ς_i to \mathcal{X}_i by projecting to the i th component. It follows that $\llbracket r \rrbracket_{\varsigma_i}^{\mathcal{X}_i} = \llbracket s \rrbracket_{\varsigma_i}^{\mathcal{X}_i}$, and thus $\llbracket r \rrbracket_{\varsigma}^{\prod_{i \in I} \mathcal{X}_i} = \llbracket s \rrbracket_{\varsigma}^{\prod_{i \in I} \mathcal{X}_i}$. ■

8.4 Ground term models and extending a signature

DEFINITION 8.4.1. Call r **ground** when $fv(r) = \emptyset$.

Definition 8.4.2 exactly follows Definition 7.5.1 (cf. Remark 8.4.6):

DEFINITION 8.4.2. Suppose $\mathsf{T} = (\Sigma, Ax)$ is a theory. For each sort α in Σ define $[r]_{\mathsf{T}}^{gnd}$ and $\mathcal{G}(\mathsf{T})_{\alpha}$ by

$$\begin{aligned} [r]_{\mathsf{T}}^{gnd} &= \{r' : \alpha \mid \mathsf{T} \vdash r = r'\} && (r : \alpha, r \text{ ground}) \\ \mathcal{G}(\mathsf{T})_{\alpha} &= \{[r]_{\mathsf{T}}^{gnd} \mid r : \alpha, r \text{ ground}\}. \end{aligned}$$

Make each $\mathcal{G}(\mathsf{T})_{\alpha}$ into a permissive-nominal set by giving it a permutation action

$$\pi \cdot [r]_{\mathsf{T}}^{gnd} = [\pi \cdot r]_{\mathsf{T}}^{gnd}.$$

LEMMA 8.4.3. $\text{supp}([r]_{\mathsf{T}}^{gnd}) \subseteq \text{fa}(r)$.

¹⁵That is:

- $a^{\mathcal{X}} = a^{\mathcal{Y}}$ for every atom a .
- $(x_1, \dots, x_n)^{\mathcal{X}} = (x_1, \dots, x_n)^{\mathcal{Y}}$ for every $x_1 \in \llbracket [\alpha_1] \rrbracket^{\mathcal{X}}$, \dots , $x_n \in \llbracket [\alpha_n] \rrbracket^{\mathcal{X}}$.
- $[a]^{\mathcal{X}} x = [a]^{\mathcal{Y}} x$ for every $x \in \llbracket [\alpha] \rrbracket^{\mathcal{X}}$.
- For every term-former f in \mathcal{F} , $f^{\mathcal{X}}(x) = f^{\mathcal{Y}}(x)$ for every $x \in \llbracket [\alpha] \rrbracket^{\mathcal{X}}$ where $ar(f) = (\alpha)\tau$.

Proof. From Definition 7.5.1 and Lemma 2.3.3. ■

DEFINITION 8.4.4. We construct the **ground free term interpretation** $\mathcal{G}(\mathbb{T})$ of \mathbb{T} as follows:

- Take $\mathcal{G}(\mathbb{T})_\alpha$ as in Definition 8.4.2.
- $a^{\mathcal{G}(\mathbb{T})} = [a]_{\mathbb{T}}^{gnd}$, $[a]^{\mathcal{G}(\mathbb{T})}[r]_{\mathbb{T}}^{gnd} = [[a]r]_{\mathbb{T}}^{gnd}$, and $([r_1]_{\mathbb{T}}^{gnd}, \dots, [r_n]_{\mathbb{T}}^{gnd})^{\mathcal{G}(\mathbb{T})} = [(r_1, \dots, r_n)]_{\mathbb{T}}^{gnd}$.
- $f^{\mathcal{G}(\mathbb{T})}([r]_{\mathbb{T}}^{gnd}) = [f(r)]_{\mathbb{T}}^{gnd}$ for each term-former $f : (\alpha)\tau$ in Σ and each $r : \alpha$.
- $C^{\mathcal{G}(\mathbb{T})} = [C]_{\mathbb{T}}^{gnd}$ for each constant in Σ .

LEMMA 8.4.5. Definition 8.4.4 is well-defined and is an interpretation.

Proof. As the proof of Lemma 7.5.5. ■

REMARK 8.4.6. Definition 7.5.1 is a special case of Definition 7.5.1. We obtain $\mathcal{F}(\mathbb{T})$ as $\mathcal{G}(\mathbb{T}')$ where \mathbb{T}' is obtained from \mathbb{T} by extending its signature with a copy of \mathcal{X} as constants (the construction is made formal in Definition 8.5.2 below).

Doing this in Definition 7.5.1 would have complicated the presentation for no immediate gain, so it seemed kinder on the reader to build the special case first by hand.

Note that we *need* to use ground terms now, for the proof of Theorem 8.5.3 to work. The reason is that $\mathcal{F}(\mathbb{T})$ has elements in each sort given by the elements $[X]_{\mathbb{T}}$, whereas $\mathcal{G}(\mathbb{T})$ lacks these elements.

8.5 Surjective maps onto algebras

Fix a signature Σ and any collection of Σ -algebras \mathcal{V} .

DEFINITION 8.5.1. Suppose $\mathbb{T} = (\Sigma, Ax)$ and suppose \mathcal{X} and \mathcal{Y}_i for $i \in I$ are models of \mathbb{T} . Suppose $\theta_i \in \mathcal{X} \rightarrow \mathcal{Y}_i$ is a family of homomorphisms.

Write $\Pi_i \theta_i$ for the natural map from \mathcal{X} to $\Pi_i \mathcal{Y}_i$, mapping $x \in |\mathcal{X}_\alpha|$ to $(\theta_i(x))_i \in |\Pi_i \mathcal{Y}_i|$.

It is easy to verify that $\Pi_i \theta_i$ is a Σ -algebra homomorphism.

DEFINITION 8.5.2. Suppose Σ and Σ' are signatures. Say Σ' **extends** Σ **with fresh constants** when $\Sigma = (\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{X}, \mathcal{F}, ar')$ and $\Sigma' = (\mathcal{A}, \mathcal{B}, \mathcal{C} \cup \mathcal{D}, \mathcal{X}, \mathcal{F}, ar')$ where $\mathcal{D} \cap \mathcal{C} = \emptyset$ and $ar'(C) = ar(C)$ for every $C \in \mathcal{C}$.

THEOREM 8.5.3. Suppose $\mathbb{T} = (\Sigma, Ax)$ is a theory and \mathcal{V} is a model of \mathbb{T} . Then there exists a theory $\mathbb{T}' = (\Sigma', Ax)$ where Σ' extends Σ with some fresh constants \mathcal{D} such that \mathcal{V} is a homomorphic image of $\mathcal{G}(\mathbb{T}')$.

Proof. We take $\mathcal{D} = \bigcup_\alpha |\mathcal{V}_\alpha|$ and construct a homomorphism based on mapping $x \in \mathcal{V}_\alpha$ (as a constant in \mathcal{D}) to itself (as an element of $|\mathcal{V}_\alpha|$). ■

8.6 Injections out of free algebras

DEFINITION 8.6.1. Suppose Σ is a signature and \mathcal{V} is a set of Σ -algebras. Let $\mathbb{T} = (\Sigma, Ax)$ where Ax is the collection of judgements valid in all $\mathcal{V} \in \mathcal{V}$ for all valuations. Call \mathbb{T} the (Σ) -theory **generated by \mathcal{V}** .

REMARK 8.6.2. So $(r = s) \in Ax$ in Definition 8.6.1 when for every $\mathcal{V} \in \mathcal{V}$ and every valuation ς to \mathcal{V} , it is the case that $\llbracket r \rrbracket_{\varsigma}^{\mathcal{V}} = \llbracket s \rrbracket_{\varsigma}^{\mathcal{V}}$.

DEFINITION 8.6.3. Define the **constants** of a term $consts(r)$ just as Definition 3.3.2 except that we take $consts(C) = \{orb(C)\}$ and $consts(X) = \emptyset$.

LEMMA 8.6.4. Suppose Σ is a signature and Σ' extends Σ with some fresh constants \mathcal{D} . Suppose Σ has enough unknowns (Definition 3.1.10).

If g is a ground term in Σ' then there exists a term g^{-1} in Σ and substitution θ such that $g^{-1}\theta = g$.

Proof. For each orbit in $consts(r)$ choose a representative $C \in orb(C) \in consts(r)$, and some distinct unknown X_C with $sort(X_C) = sort(C)$ and $supp(C) \subseteq supp(X_C)$ —we can do this because we have assumed enough unknowns and it is a fact that $consts(r)$ is finite. Define θ to be the equivariant extension of this choice, so $\theta(\pi \cdot X_C) = \pi \cdot C$ and (for all the other unknowns) $\theta(Y) = Y$. This is well-defined by Proposition 2.5.5.

It is now easy to generate g^{-1} by replacing each C in g with X_C (modulo some permutations). ■

THEOREM 8.6.5. Suppose \mathcal{V} is a collection of Σ -algebras and Σ has enough unknowns. Let $\mathbb{T} = (\Sigma, Ax)$ be the Σ -theory generated by \mathcal{V} . Suppose Σ' extends Σ with some fresh constants \mathcal{D} and write $\mathbb{T}' = (\Sigma', Ax)$.

Then there exists some indexing set I , set of algebras $\{\mathcal{V}_i \in \mathcal{V} \mid i \in I\}$, and an injective Σ -algebra homomorphism Θ from $\mathcal{G}(\Sigma')$ to $\prod_{i \in I} \mathcal{V}_i$.

Proof. Take I to be the set of all pairs (g, h) of ground terms in Σ' such that $\mathbb{T}' \not\vdash g = h$.

Consider some $i = (g, h) \in I$. By Lemma 8.6.4 there exist g^{-1} , h^{-1} , and θ_i such that $g^{-1}\theta_i = g$ and $h^{-1}\theta_i = h$. We assumed that $\mathbb{T}' \not\vdash g = h$ and it follows using Lemma 7.2.4 that $\mathbb{T} \not\vdash g^{-1} = h^{-1}$. Since \mathbb{T} is the theory generated by \mathcal{V} there exists a model $\mathcal{V}_i \in \mathcal{V}$ and a valuation ς such that $\llbracket g^{-1} \rrbracket_{\varsigma}^{\mathcal{V}_i} \neq \llbracket h^{-1} \rrbracket_{\varsigma}^{\mathcal{V}_i}$. We define a Σ -homomorphism Θ_i from $\mathcal{G}(\mathbb{T}')$ to \mathcal{V}_i as an equivariant extension of mapping $C \in \mathcal{D}$ to $\varsigma(X_C)$, where C and X_C are as chosen in the proof of Lemma 8.6.4.

It follows by the choice of \mathcal{V}_i that $\prod_{i \in I} \Theta_i$ from $\mathcal{G}(\mathbb{T}')$ to $\prod_{i \in I} \mathcal{V}_i$ is injective as a map on underlying sets. ■

8.7 Proof of the HSP theorem

DEFINITION 8.7.1. Suppose Σ is a signature. Suppose \mathcal{V} is a collection of Σ -algebras. Then:

- Call \mathcal{V} a **(Σ -)variety** when it is closed under Homomorphic image (Definition 8.1.1), Subalgebra (Definition 8.2.1), and countable Product (Definition 8.3.1).
- Call \mathcal{V} **(Σ -)equational** when it is the collection of Σ -algebras that are models of $\mathbb{T} = (\Sigma, Ax)$ for some set of axioms Ax .

LEMMA 8.7.2. Suppose Σ is a signature with enough unknowns. Suppose \mathcal{V} is a Σ -variety and let $\mathbb{T} = (\Sigma, Ax)$ be the Σ -theory generated by \mathcal{V} . Suppose Σ' extends Σ with some fresh constants \mathcal{D} and write $\mathbb{T}' = (\Sigma', Ax)$. Then $\mathcal{G}(\mathbb{T}') \in \mathcal{V}$.

Proof. By Theorem 8.6.5 there is some indexing set I , set of Σ -algebras $\{\mathcal{V}_i \in \mathcal{V} \mid i \in I\}$, and injective Σ -algebra homomorphism Θ from $\mathcal{G}(\mathbb{T}')$ to $\prod_{i \in I} \mathcal{V}_i$. \mathcal{V} is closed under products so $\prod_{i \in I} \mathcal{V}_i \in \mathcal{V}$. The image of $|\mathcal{G}(\mathbb{T}')|$ is a subalgebra of $\prod_{i \in I} \mathcal{V}_i$, and is a homomorphic image of that subalgebra (by inverting Θ). \mathcal{V} is closed under subalgebras and homomorphic images, so the result follows. ■

THEOREM 8.7.3. Suppose Σ is a signature with enough unknowns. A collection of Σ -algebras \mathcal{V} is equational if and only if it is a variety.

Proof. Suppose \mathcal{V} is equational. By Lemma 8.3.2 \mathcal{V} is closed under products. By Lemma 8.1.4 \mathcal{V} is closed under homomorphic images. By Lemma 8.2.2 \mathcal{V} is closed under subalgebras. Therefore \mathcal{V} is a variety.

Conversely, suppose \mathcal{V} is a variety. Let \mathbb{T} be the theory on Σ generated by \mathcal{V} as described in Definition 8.6.1. Let \mathcal{V} be any model of \mathbb{T} . By Theorem 8.5.3 there exists a signature Σ' extending Σ with some fresh constants \mathcal{D} such that \mathcal{V} is a homomorphic image of $\mathcal{G}(\mathbb{T}')$. By Lemma 8.7.2 $\mathcal{G}(\mathbb{T}') \in \mathcal{V}$. Since \mathcal{V} is closed under homomorphisms, $\mathcal{V} \in \mathcal{V}$ as required. Therefore \mathcal{V} is equational. ■

Part III

Permissive-nominal logic: $\forall X$

9 PERMISSIVE-NOMINAL LOGIC

We now add quantification over unknowns—that is, $\forall X$ —to permissive-nominal terms. Permissive nominal techniques makes the theory of α -equivalence easy here: if we used ‘vanilla’ nominal terms, then the development below might not be impossible, but we believe that it would be harder. We obtain a first-order logic which we call *permissive-nominal logic*.

Permissive-nominal logic (PNL) is just first-order logic, enriched with nominal-style names. Thus, the derivation rules in Figure 5 are (virtually) identical to those of first-order logic. Only the term language is really changed.

In this section we set up PNL as a sequent derivation system (Figure 5), interpret it in permissive-nominal sets (Definition 9.3.2), and prove soundness and completeness (Theorems 9.3.6 and 9.4.16).

In Section 10 we undertake as an extended case study a sound and complete finite axiomatisation of arithmetic inside PNL.

9.1 Syntax

The notions of sort-signature $(\mathcal{A}, \mathcal{B})$ and sort language are as in Definition 3.1.1. An interpretation \mathcal{I} for $(\mathcal{A}, \mathcal{B})$ consists of an assignment of a permissive-nominal set $\tau^{\mathcal{I}}$ to each $\tau \in \mathcal{B}$, and we extend \mathcal{I} to sorts as in Definition 7.3.1.

DEFINITION 9.1.1. For this section it is convenient to take \mathcal{X} to be specifically example 2 of Example 3.1.7.

REMARK 9.1.2. So an unknown X takes the form

$$\pi \cdot \mathbf{X}_\alpha = \{(\pi', \mathbf{X}_\alpha) \mid \forall a \in \mathbb{A}^<. \pi(a) = \pi'(a)\}.$$

In this case, in the light of Remark 3.1.8, we may take $fv(r)$ to be equal to the set of \mathbf{X}_α occurring in r .

It is now easy to define binding of level 2 variables in terms. An abstract account of level 2 binding is available [Gabbay, 2011d], but it is not worth introducing it for now.

DEFINITION 9.1.3. A **PNL signature** over a sort-signature $(\mathcal{A}, \mathcal{B})$ is a tuple $(\mathcal{C}, \mathcal{F}, \mathcal{P}, ar)$ where:

- \mathcal{C} is a permissive-nominal set of **constants**.
 - \mathcal{F} is a set of equivariant **term-formers**.
 - \mathcal{P} is a set of equivariant **predicate-formers**.
 - ar assigns
 - to each constant $C \in \mathcal{C}$ an arity τ ,
 - to each $f \in \mathcal{F}$ a **term-former arity** $(\alpha)\tau$, and
 - to each $P \in \mathcal{P}$ a **proposition-former arity** α , where
- α and τ are in the sort-language determined by $(\mathcal{A}, \mathcal{B})$.

A **(PNL) signature** \mathcal{S} is then a tuple $(\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{F}, \mathcal{P}, ar)$.

DEFINITION 9.1.4. Suppose $\mathcal{S} = (\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{F}, \mathcal{P}, ar)$ is a PNL signature.

Terms are defined as in Definition 3.2.1 for the signature $(\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{X}, \mathcal{F}, ar)$.¹⁶

Propositions are defined as follows:

\perp proposition	ϕ proposition ψ proposition $\phi \Rightarrow \psi$ proposition
$r : \alpha \quad (ar(P) = \alpha)$ $P(r)$ proposition	ϕ proposition $\forall X_\alpha. \phi$ proposition

Here $\forall X_\alpha$ binds X_α in ϕ . We can use nominal abstract syntax to do this.

NOTATION 9.1.5. Write $\forall X. \phi$ as shorthand for $\forall X_\alpha. \phi$ where $X = \{(\pi', X_\alpha) \mid \forall a \in \mathbb{A}^\lt . \pi'(a) = \pi(a)\}$ for some π .

LEMMA 9.1.6. Support and the permutation action as characterised for terms on Lemma 3.2.5 extend to propositions as follows:

$$\begin{array}{ll}
 \text{supp}(\perp) = \emptyset & \text{supp}(P(r)) = \text{supp}(r) \\
 \text{supp}(\phi \Rightarrow \psi) = \text{supp}(\phi) \cup \text{supp}(\psi) & \text{supp}(\forall X. \phi) = \text{supp}(\phi) \\
 \pi \cdot \perp = \perp & \pi \cdot P(r) = P(\pi \cdot r) \\
 \pi \cdot (\phi \Rightarrow \psi) = (\pi \cdot \phi) \Rightarrow \pi \cdot \psi & \pi \cdot \forall X. \phi = \forall X. \pi \cdot \phi
 \end{array}$$

NOTATION 9.1.7. We may write $fa(\phi)$ for $\text{supp}(\phi)$.

¹⁶Yes. We have to adjust ar to remove \mathcal{P} and add \mathcal{X} mapping X to α where $(\pi, X_\alpha) \in X$. 10 points.

$\frac{}{\Phi, \phi \vdash \pi \cdot \phi, \Psi} \text{ (Ax)}$	$\frac{}{\Phi, \perp \vdash \Psi} \text{ (\perp L)}$
$\frac{\Phi \vdash \phi, \Psi \quad \Phi, \psi \vdash \Psi}{\Phi, \phi \Rightarrow \psi \vdash \Psi} \text{ (\Rightarrow L)}$	$\frac{\Phi, \phi \vdash \psi, \Psi}{\Phi \vdash \phi \Rightarrow \psi, \Psi} \text{ (\Rightarrow R)}$
$\frac{\Phi, \phi[X:=r] \vdash \Psi \quad (fa(r) \subseteq \text{supp}(X), r:\text{sort}(X))}{\Phi, \forall X. \phi \vdash \Psi} \text{ (\forall L)}$	$\frac{\Phi \vdash \phi, \Psi \quad (X \notin \text{fv}(\Phi, \Psi))}{\Phi \vdash \forall X. \phi, \Psi} \text{ (\forall R)}$
$\frac{\Phi \vdash \phi, \Psi \quad \Phi, \phi \vdash \Psi}{\Phi \vdash \Psi} \text{ (Cut)}$	

Figure 5: Sequent derivation rules of Permissive-Nominal Logic

9.2 Derivability

DEFINITION 9.2.1. Φ and Ψ will range over sets of propositions. We may write ϕ, Φ and Φ, ϕ as shorthand for $\{\phi\} \cup \Phi$. We may write Φ, Ψ as shorthand for $\Phi \cup \Psi$.

Write $\text{fv}(\Phi) = \bigcup \{\text{fv}(\phi) \mid \phi \in \Phi\}$.

A **sequent** is a pair $\Phi \vdash \Psi$.

DEFINITION 9.2.2 (Derivable sequents). The **derivable sequents** are defined in Figure 5.

REMARK 9.2.3. As standard, the intuition of $\Phi \vdash \Psi$ being derivable is “the conjunction (logical and) of the propositions in Φ entails the disjunction (logical or) of the propositions in Ψ ”. So for instance, intuitively the axiom rule (Ax) expresses that ϕ if and only if $\pi \cdot \phi$.

The permutation π in (Ax) is deliberate and represents equivariance (preservation of truth under permuting atoms) within permissive-nominal logic. Because of this permutation π , free atoms can behave like variables ranging over distinct atoms.

Thus in PNL we can express a theory of atoms-inequality as follows: Suppose a name sort Atm and a proposition-former $\text{neq} : (\text{Atm}, \text{Atm})$, along with a single proposition $\text{neq}(a, b)$ for two distinct atoms in Atm —and, if we wish, another proposition $\text{neq}(a, a) \Rightarrow \perp$. The permutation π in (Ax) ensures that a and b represent *any* two distinct atoms.

REMARK 9.2.4. The condition $fa(r) \subseteq \text{supp}(X)$ in (\forall L) might suggest

that $\forall X.\phi$ means “ $\phi[X:=r]$ for every r such that $fa(r) \subseteq \text{supp}(X)$ ”. This is so, but the π in (\mathbf{Ax}) means that what $\text{supp}(X)$ in $\forall X.\phi$ really restricts is *capture*, as we now discuss.

- Suppose a name sort Atm and suppose $X : \text{Atm}$ and $P : \text{Atm}$. Suppose $b \in \text{pmss}(X)$. By considering the swapping $(b\ a)$ and (\mathbf{Ax}) , and $(\forall\mathbf{L})$, $\forall X.P(X) \vdash P(b)$ for *all* a , even if $a \notin \text{supp}(X)$, as follows:

$$\frac{\frac{}{P(b) \vdash P(a)} (\mathbf{Ax}) \quad \pi = (b\ a)}{\forall X.P(X) \vdash P(a)} (\forall\mathbf{L}) \quad [X:=b]$$

So we can derive $P(a)$ from $\forall X.P(X)$, even if a is not permitted in X .

- This may not work if we have to ‘shift’ infinitely many atoms; e.g. there is no finite π such that $fa(\pi \cdot (X, a)) \subseteq \text{supp}(X)$ where $a \notin \text{supp}(X)$. If \mathbb{P} has *shift*-permutations (Definition 3.6.1), then we can use them.

Consider any sort α and suppose $X : \alpha$ and $\text{supp}(X) = S$. Suppose $Q : \alpha$. Consider any other $Y : \alpha$ with $\text{supp}(Y) = S \cup \{a\}$ where $a \notin S$. We will show that given *shift*-permutations, $\forall X.Q(X) \vdash Q(Y)$ is derivable.

Suppose $S \cup \{a\} = \pi \cdot S$. We derive as follows:

$$\frac{\frac{}{Q(\pi \cdot Y) \vdash Q(Y)} (\mathbf{Ax})}{\forall X.Q(X) \vdash Q(Y)} (\forall\mathbf{L}) \quad [X:=\pi \cdot Y]$$

- Nevertheless, $\forall X.\phi$ does not mean “ $\phi[X:=r]$ for every r ”. This is because permutations are bijective. For example, suppose $X : \text{Atm}$, $a \notin \text{supp}(X)$, and $P : ([\text{Atm}]\text{Atm})$. Then $\forall X.P([a]X) \vdash P([a]r)$ for all r such that $a \notin fa(r)$, and also $\forall X.P([b]X) \vdash P([b]r)$ for all r and all b such that $b \notin fa(r)$. However,

$$\forall X.P([a]X) \not\vdash P([a]a), \quad \text{and for all } b, \quad \forall X.P([a]X) \not\vdash P([b]b).$$

The fact that $a \notin \text{supp}(X)$ forbids capture by an instantiation, in a suitable sense.

9.3 Interpretation and soundness

DEFINITION 9.3.1. Suppose $\mathcal{S} = (\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{F}, \mathcal{P}, ar)$ is a signature.

A **(PNL) interpretation** \mathcal{I} for \mathcal{S} consists of the following data:

- An interpretation for the sort-signature $(\mathcal{A}, \mathcal{B})$ (Definition 7.3.1).
- For every $f \in \mathcal{F}$ with $ar(f) = (\alpha')\alpha$ an equivariant function $f^\mathcal{I}$ from $\llbracket \alpha' \rrbracket_\zeta^\mathcal{I}$ to $\llbracket \alpha \rrbracket_\zeta^\mathcal{I}$ (Definition 2.3.1).
- For every $P \in \mathcal{P}$ with $ar(P) = \alpha$ an equivariant function $P^\mathcal{I}$ from $\llbracket \alpha \rrbracket_\zeta^\mathcal{I}$ to $\{0, 1\}$ (Definition 2.4.1).

DEFINITION 9.3.2. Suppose that \mathcal{I} is an interpretation. Define an **interpretation of propositions** by:

$$\begin{aligned} \llbracket P(r) \rrbracket_\zeta^\mathcal{I} &= P^\mathcal{I}(\llbracket r \rrbracket_\zeta^\mathcal{I}) \\ \llbracket \perp \rrbracket_\zeta^\mathcal{I} &= 0 \\ \llbracket \phi \Rightarrow \psi \rrbracket_\zeta^\mathcal{I} &= \max\{1 - \llbracket \phi \rrbracket_\zeta^\mathcal{I}, \llbracket \psi \rrbracket_\zeta^\mathcal{I}\} \\ \llbracket \forall X.\phi \rrbracket_\zeta^\mathcal{I} &= \min\{\llbracket \phi \rrbracket_{\zeta[X:=x]}^\mathcal{I} \mid x \in \llbracket \text{sort}(X) \rrbracket_\zeta^\mathcal{I}, \text{supp}(x) \subseteq \text{supp}(X)\} \end{aligned}$$

LEMMA 9.3.3. $\llbracket \phi \rrbracket_\zeta^\mathcal{I} = \llbracket \pi \cdot \phi \rrbracket_\zeta^\mathcal{I}$ always.

Proof. By induction on ϕ . We consider two cases:

- The case $\forall X.\phi$. Suppose $\llbracket \forall X.\phi \rrbracket_\zeta^\mathcal{I} = 1$. This means that $\llbracket \phi \rrbracket_{\zeta[X:=x]}^\mathcal{I} = 1$ for all $x \in \llbracket \alpha \rrbracket_\zeta^\mathcal{I}$ such that $\text{supp}(x) \subseteq \text{supp}(X)$. By inductive hypothesis $\llbracket \pi \cdot \phi \rrbracket_{\zeta[X:=x]}^\mathcal{I} = 1$ for all $x \in \llbracket \alpha \rrbracket_\zeta^\mathcal{I}$ such that $\text{supp}(x) \subseteq \text{supp}(X)$. Therefore $\llbracket \forall X.\pi \cdot \phi \rrbracket_\zeta^\mathcal{I} = 1$. The result follows, since $\pi \cdot (\forall X.\phi) = \forall X.\pi \cdot \phi$.
- The case $P(r)$. We have $\llbracket P(r) \rrbracket_\zeta^\mathcal{I} = P^\mathcal{I}(\llbracket r \rrbracket_\zeta^\mathcal{I})$. As $P^\mathcal{I}$ is equivariant, we get $\llbracket P(r) \rrbracket_\zeta^\mathcal{I} = P^\mathcal{I}(\pi \cdot \llbracket r \rrbracket_\zeta^\mathcal{I})$. By Lemma 7.3.6 $\pi \cdot \llbracket r \rrbracket_\zeta^\mathcal{I} = \llbracket \pi \cdot r \rrbracket_\zeta^\mathcal{I}$. Thus $\llbracket P(r) \rrbracket_\zeta^\mathcal{I} = P^\mathcal{I}(\llbracket \pi \cdot r \rrbracket_\zeta^\mathcal{I}) = \llbracket \pi \cdot P(r) \rrbracket_\zeta^\mathcal{I}$.

■

LEMMA 9.3.4. $\llbracket \phi \rrbracket_{\zeta[X:=\llbracket t \rrbracket_\zeta^\mathcal{I}]}^\mathcal{I} = \llbracket \phi[X:=t] \rrbracket_\zeta^\mathcal{I}$.

Proof. By a routine induction on the definition of $\llbracket \phi \rrbracket_\zeta^\mathcal{I}$ in Definition 9.3.2. We consider one case:

- The case of $\llbracket P(r) \rrbracket_{\zeta[X:=t]}^\mathcal{I}$. We reason as follows:

$$\begin{aligned} \llbracket P(r) \rrbracket_{\zeta[X:=\llbracket t \rrbracket_\zeta^\mathcal{I}]}^\mathcal{I} &= P^\mathcal{I}(\llbracket r \rrbracket_{\zeta[X:=\llbracket t \rrbracket_\zeta^\mathcal{I}]}^\mathcal{I}) && \text{Definition 9.3.2} \\ &= P^\mathcal{I}(\llbracket r[X:=t] \rrbracket_\zeta^\mathcal{I}) && \text{Lemma 7.4.5} \\ &= \llbracket P(r)[X:=t] \rrbracket_\zeta^\mathcal{I} && \text{Definition 9.3.2.} \end{aligned}$$

■

Validity and soundness

DEFINITION 9.3.5 (Validity). Call the proposition ϕ **valid** in \mathcal{I} when $\llbracket \phi \rrbracket_{\varsigma}^{\mathcal{I}} = 1$ for all ς .

Call the sequent $\phi_1, \dots, \phi_n \vdash \psi_1, \dots, \psi_p$ **valid** in \mathcal{I} when $(\phi_1 \wedge \dots \wedge \phi_n) \Rightarrow (\psi_1 \vee \dots \vee \psi_p)$ is valid.

THEOREM 9.3.6 (Soundness). If $\Phi \vdash \Psi$ is derivable, then it is valid in all interpretations.

Proof. By induction on derivations (Figure 5). The case of **(Ax)** uses Lemma 9.3.3. The case of **($\forall\mathbf{L}$)** uses Lemma 9.3.4. The case of **($\forall\mathbf{R}$)** uses Lemma 7.4.2. Other rules are routine by unpacking definitions. ■

9.4 Completeness

In this subsection we prove Theorem 9.4.16: a converse to Theorem 9.3.6, that if ϕ is valid in all interpretations, then ϕ it is derivable.

For this subsection fix the following data:

- A signature $\mathcal{S} = (\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{F}, \mathcal{P}, ar)$.
- A formula ϕ such that $\not\vdash \phi$.

We will construct an interpretation \mathcal{I} and a valuation ς (Definition 7.3.1) such that $\llbracket \phi \rrbracket_{\varsigma}^{\mathcal{I}} = 0$. This suffices to prove the result.

Maximally consistent set of propositions

DEFINITION 9.4.1. Make a fixed but arbitrary order on propositions $\xi_1, \xi_2, \xi_3, \dots$

Define $\Phi_1 = \{\neg\phi\}$ (where ϕ was fixed above). For each $i \geq 1$ we define Φ_{i+1} as follows:

- If $\Phi_i \vdash \xi_i$ then write $\xi = \xi_i$.
- If $\Phi_i \vdash \neg\xi_i$ then write $\xi = \neg\xi_i$.
- If $\Phi_i \not\vdash \xi_i$ and $\Phi_i \not\vdash \neg\xi_i$ then write $\xi = \xi_i$.

There are now two cases:

- If ξ has the form $\neg\forall X.\xi'$ then we define $\Phi_{i+1} = \Phi_i \cup \{\xi, \neg\xi'[X:=Z]\}$ where Z is some fixed but arbitrary choice of unknown that is not free in any proposition in Φ_i and is such that $supp(Z) = supp(X)$ and $sort(Z) = sort(X)$.
- Otherwise, we define $\Phi_{i+1} = \Phi_i \cup \{\xi\}$.

Finally, we define Φ_{ω} by $\Phi_{\omega} = \bigcup_i \Phi_i$.

LEMMA 9.4.2. For every i , $\Phi_i \not\vdash \perp$.

Proof. By induction on i :

- By definition $\Phi_1 = \{\neg\phi\}$. As $\not\vdash \phi$, we have $\neg\phi \not\vdash \perp$
- Suppose $\Phi_i \not\vdash \perp$.
 Either $\Phi_{i+1} = \Phi_i \cup \{\neg\xi\}$ or $\Phi_{i+1} = \Phi_i \cup \{\neg\xi, \neg\xi'[X:=Z]\}$ —we consider the first, simpler case; the second case is similar. Suppose $\Phi_i, \xi \vdash \perp$. It follows that $\Phi_i \vdash \neg\xi$. So we are not in the third case of Definition 9.4.1 and we are either in the first or the second. So $\Phi_i \vdash \xi$ and thus $\Phi_i \vdash \perp$; a contradiction. ■

LEMMA 9.4.3. $\Phi_\omega \not\vdash \perp$.

Proof. Assume $\Phi_\omega \vdash \perp$. So there exists a finite subset Γ of Φ_ω such that $\Gamma \vdash \perp$. As Γ is finite it is included in some Φ_i , and $\Phi_i \vdash \perp$, contradicting Proposition 9.4.2. ■

REMARK 9.4.4. It is well-known that in nominal sets, least upper bounds can sometimes not exist if there are ‘too many’ atoms; so sometimes we have to be careful to not make too many arbitrary choices.¹⁷

The reader familiar with nominal techniques will be alert to the possibility that Φ_ω might fail to have a supporting permission set, and that this could cause problems. In fact, in this particular case this is a non-issue: **(Ax)** from Figure 5 ensures that Φ_ω is not only supported, but in fact equivariant.

LEMMA 9.4.5. For every ξ , at least one of $\xi \in \Phi_\omega$ and $\neg\xi \in \Phi_\omega$ holds.

Proof. We check of Definition 9.4.1 that for every i , either $\xi_i \in \Phi_{i+1}$ or $\neg\xi_i \in \Phi_{i+1}$. The result follows. ■

LEMMA 9.4.6. For every ξ , if $\neg\forall X.\xi \in \Phi_\omega$ then there exists a Z such that $\neg\xi[X:=Z] \in \Phi_\omega$.

Proof. There exists an i such that $\xi_i = \neg\forall X.\xi$. Since $\Phi_\omega \vdash \xi_i$ and $\Phi_\omega \not\vdash \perp$, we have that $\Phi_\omega \not\vdash \neg\xi_i$, and so $\Phi_i \not\vdash \neg\xi_i$. Thus $\Phi_{i+1} = \Phi_i \cup \{\neg\forall X.\xi, \neg\xi[X:=Z]\}$. The result follows. ■

LEMMA 9.4.7. If $\Phi_\omega \vdash \phi$ then $\phi \in \Phi_\omega$.

Proof. As, by Lemma 9.4.3, $\Phi_\omega \not\vdash \perp$, if $\Phi_\omega \vdash \phi$ then $\neg\phi \notin \Phi_\omega$. Thus by Lemma 9.4.5, $\phi \in \Phi_\omega$. ■

¹⁷For instance, in permissive-nominal sets it is possible represent a well-order of each permission set, but not to represent a well-ordering of the set of all atoms (which is a limit of permission sets). This is also the feature which Fraenkel and Mostowski used to prove the independence of the axiom of choice from the other axioms of set theory.

COROLLARY 9.4.8.

1. $(\phi \Rightarrow \psi) \in \Phi_\omega$ if and only if $(\phi \notin \Phi_\omega \text{ or } \psi \in \Phi_\omega)$.
2. $\forall X.\phi \in \Phi_\omega$ if and only if
(for every r such that $r : \text{sort}(X)$ and $\text{fa}(r) \subseteq \text{supp}(X)$, $\phi[X:=r] \in \Phi_\omega$).

Proof.

1. Suppose $(\phi \Rightarrow \psi) \in \Phi_\omega$ and $\phi \in \Phi_\omega$. Then $\Phi_\omega \vdash \psi$ and so by Lemma 9.4.7 $\psi \in \Phi_\omega$.
Suppose $\phi \notin \Phi_\omega$. By Lemma 9.4.5 $\neg\phi \in \Phi_\omega$. So $\Phi_\omega \vdash \neg\phi$ and therefore $\Phi_\omega \vdash \phi \Rightarrow \psi$. By Lemma 9.4.7 $(\phi \Rightarrow \psi) \in \Phi_\omega$.
Suppose $\psi \in \Phi_\omega$. Then $\Phi_\omega \vdash \psi$ and so $\Phi_\omega \vdash \phi \Rightarrow \psi$. By Lemma 9.4.7 $(\phi \Rightarrow \psi) \in \Phi_\omega$.
2. Suppose $\forall X.\phi \in \Phi_\omega$. By Lemma 9.4.7, if $r : \text{sort}(X)$ and $\text{fa}(r) \subseteq \text{supp}(X)$ then $\phi[X:=r] \in \Phi_\omega$.
Conversely, suppose $\phi[X:=r] \in \Phi_\omega$ for every r such that $r : \text{sort}(X)$ and $\text{fa}(r) \subseteq \text{supp}(X)$. We proceed by contradiction: suppose $\forall X.\phi \notin \Phi_\omega$. By Lemma 9.4.5 $\neg\forall X.\phi \in \Phi_\omega$ and by Lemma 9.4.6, there exists a Z such that $\neg\phi[X:=Z] \in \Phi_\omega$. So $\Phi_\omega \vdash \neg\phi[X:=Z]$, and so $\Phi_\omega \vdash \phi[X:=Z]$, and so $\Phi_\omega \vdash \perp$, contradicting Lemma 9.4.3. ■

The term interpretation

DEFINITION 9.4.9. Define the **term interpretation** \mathcal{I} by:

- $\llbracket \alpha \rrbracket^{\mathcal{I}} = \{r \mid r : \alpha\}$.
- $\mathbf{f}^{\mathcal{I}}$ maps r to $\mathbf{f}(r)$.
- $\mathbf{P}^{\mathcal{I}}$ maps r_1, \dots, r_n to 1 if $\mathbf{P}(r_1, \dots, r_n) \in \Phi_\omega$ and to 0 otherwise.

Define ς by $\varsigma(X) = X$ for all $X \in \mathcal{X}$ and endow $\llbracket \alpha \rrbracket^{\mathcal{I}}$ with a permutation action given by the action on terms.

REMARK 9.4.10. In Definition 7.5.4 we built an interpretation to prove completeness of nominal algebra (Corollary 7.5.12). There, we built our interpretation out of terms quotiented by derivable equality; here we just use terms. Why the difference?

In nominal algebra the judgement-form of the logic *is* equality—so it makes sense to build an interpretation such that equality maps to denotational identity.

LEMMA 9.4.11.

1. If $ar(f) = (\alpha)\tau$ then $f^\mathcal{I}$ is well-defined, equivariant, and maps $\llbracket \alpha \rrbracket^\mathcal{I}$ to $\llbracket \tau \rrbracket^\mathcal{I}$.
2. If $ar(P) = \alpha$ then $P^\mathcal{I}$ is well-defined, equivariant, and maps $\llbracket \alpha \rrbracket^\mathcal{I}$ to $\{0, 1\}$.

PROPOSITION 9.4.12. \mathcal{I} is an interpretation of the signature $\mathcal{S} = (\mathcal{A}, \mathcal{B}, \mathcal{F}, \mathcal{P}, ar)$ which we fixed at the beginning of this subsection. In addition, ς is a valuation to \mathcal{I} .

Proof. By Lemma 9.4.11 for each $f : (\alpha')\alpha \in \mathcal{F}$, $f^\mathcal{I}$ is an equivariant map from $\llbracket \alpha' \rrbracket^\mathcal{I}$ to $\llbracket \alpha \rrbracket^\mathcal{I}$ and for each $P : \alpha \in \mathcal{P}$, $P^\mathcal{I}$ is an equivariant function from $\llbracket \alpha \rrbracket^\mathcal{I}$ to $\{0, 1\}$.

By construction $\varsigma(X) \in \llbracket sort(X) \rrbracket^\mathcal{I}$ always. Equivariance is easy. \blacksquare

LEMMA 9.4.13. $\llbracket r \rrbracket_\varsigma^\mathcal{I} = r$.

LEMMA 9.4.14. $\llbracket \xi \rrbracket_\varsigma^\mathcal{I} = 1$ if and only if $\xi \in \Phi_\omega$.

Proof. By induction on the definition of $\llbracket \xi \rrbracket_\varsigma^\mathcal{I}$ (Definition 9.3.2):

- The case of $\llbracket P(r) \rrbracket_\varsigma^\mathcal{I}$. We reason as follows:

$\llbracket P(r) \rrbracket_\varsigma^\mathcal{I} = 1$	\Leftrightarrow	$P^\mathcal{I}(\llbracket r \rrbracket_\varsigma^\mathcal{I}) = 1$	Definition 9.3.2
	\Leftrightarrow	$P^\mathcal{I}(r) = 1$	Lemma 9.4.13
	\Leftrightarrow	$P(r) \in \Phi_\omega$	Definition 9.4.9
- The case of $\llbracket \perp \rrbracket_\varsigma^\mathcal{I}$. By definition $\llbracket \perp \rrbracket_\varsigma^\mathcal{I} = 0$. By part 1 of Corollary 9.4.8, $\perp \notin \Phi_\omega$.
- The case of $\llbracket \phi \Rightarrow \psi \rrbracket_\varsigma^\mathcal{I}$. We reason as follows:

$\llbracket \phi \Rightarrow \psi \rrbracket_\varsigma^\mathcal{I} = 1$	\Leftrightarrow	$\llbracket \phi \rrbracket_\varsigma^\mathcal{I} = 0$ or $\llbracket \psi \rrbracket_\varsigma^\mathcal{I} = 1$	Definition 9.3.2
	\Leftrightarrow	$\phi \notin \Phi_\omega$ or $\psi \in \Phi_\omega$	ind. hyp.
	\Leftrightarrow	$\phi \Rightarrow \psi \in \Phi_\omega$	Cor. 9.4.8, part 2
- The case of $\llbracket \forall X. \phi \rrbracket_\varsigma^\mathcal{I}$, where $\alpha = sort(X)$ and $S = supp(X)$.

$\llbracket \forall X. \phi \rrbracket_\varsigma^\mathcal{I} = 1$	\Leftrightarrow	$\forall t \in \llbracket \alpha \rrbracket^\mathcal{I}. supp(t) \subseteq S \Rightarrow \llbracket \phi \rrbracket_{\varsigma[X:=t]}^\mathcal{I} = 1$	Definition 9.3.2
	\Leftrightarrow	$\forall t \in \llbracket \alpha \rrbracket^\mathcal{I}. supp(t) \subseteq S \Rightarrow \llbracket \phi[X:=t] \rrbracket_\varsigma^\mathcal{I} = 1$	Lems. 7.4.2, 9.4.13
	\Leftrightarrow	$\llbracket \phi[X:=t] \rrbracket_\varsigma^\mathcal{I} = 1$ every $t:\alpha$ s.t. $fa(t) \subseteq S$	$supp(t) = fa(t)$
	\Leftrightarrow	$\phi[X:=t] \in \Phi_\omega$ every $t:\alpha$ s.t. $fa(t) \subseteq S$	ind. hyp.
	\Leftrightarrow	$\forall X. \phi \in \Phi_\omega$	Cor. 9.4.8, part 4

LEMMA 9.4.15. If $\not\vdash \phi$, then there exists an interpretation \mathcal{I} and a valuation ς such that $\llbracket \phi \rrbracket_\varsigma^\mathcal{I} = 0$.

Proof. As $\neg\phi \in \Phi_0 \subseteq \Phi_\omega$ and $\Phi_\omega \not\vdash \perp$, we have $\phi \notin \Phi_\omega$. By Lemma 9.4.14, we get $\llbracket \phi \rrbracket_\zeta^\varepsilon = 0$. ■

As a corollary we get Theorem 9.4.16:

THEOREM 9.4.16 (Completeness). If ϕ is valid in all interpretations, then ϕ is derivable.

10 CASE STUDY: ARITHMETIC IN PERMISSIVE-NOMINAL LOGIC

Because term-formers in PNL can bind, we can axiomatise first-order logic. Thus assume a sort o whose terms reflect formulas of first-order logic. Then PNL quantification $\forall Z$ where $Z : o$ has the quality of an *axiom schema*, and we can use those terms to axiomatise arithmetic (a theory which in first-order logic famously involves an axiom schema).

So, we should be able to use PNL to give a finite, first-order axiomatisation of arithmetic. Writing down some plausible-looking axioms is one thing—proving they do what we expect them to do, is another. In this section, as a case study of an application of PNL, we do just that.

We proceed as follows, starting with the following PNL definitions:

- Figure 6 gives $\dot{\mathcal{L}}$ a signature for a shallow embedding of terms and formulas of first-order logic as PNL terms of sort ι and o respectively.
- Figure 7 gives equality axioms, as a transitive reflexive symmetric congruence for the term-formers in $\dot{\mathcal{L}}$.
- Figure 8 axiomatises substitution. We can have some confidence in this axiomatisation because it was already considered for nominal algebra in [Gabbay and Mathijssen, 2008a] and proven correct.
- Figure 9 gives axioms for first-order logic.
- Finally, Figure 10 gives axioms for arithmetic. As discussed above, the induction axiom schema is captured using a universal quantification (the $\forall Z$ in (**PInd**)).

Subsection 10.4 briefly recalls the syntax and derivability relation of ‘real’ first-order logic. Then Subsection 10.5 maps this into the PNL theory we just constructed. Subsection 10.6 briefly recalls Peano arithmetic in the ‘real’ first-order logic.

Finally, in Subsection 10.7 by arguments on models we show our main result of this section: Theorem 10.7.7. A formula is derivable in ‘real’ Peano arithmetic if and only if its translation in PNL is derivable in the PNL theory for arithmetic.

The permissive-nominal terms, PNL, permission-sets, and permissive-nominal sets semantics, all work together, and at the end of it all we really *can* embed a non-trivial theory with binding in PNL, and know it is correct.

We assume one atomic sort ν and two base sorts ι and o .
 We assume term-formers and proposition-formers as follows:

$$\begin{array}{llll}
 \dot{0} : \iota & \text{succ} : (\iota)\iota & \dot{+} : (\iota, \iota)\iota & \dot{*} : (\iota, \iota)\iota \\
 \dot{\perp} : o & \Rightarrow : (o, o)o & \dot{\forall} : ([\nu]o)o & \approx : (\iota, \iota)o \\
 \text{var} : (\nu)\iota & \text{sub}_\iota : ([\nu]\iota, \iota)\iota & \text{sub}_o : ([\nu]o, \iota)o & \\
 \approx_\iota : (\iota, \iota) & \approx_o : (o, o) & \epsilon : (o) &
 \end{array}$$

Figure 6: Signature $\dot{\mathcal{L}}$ suitable for a PNL specification of arithmetic

10.1 The signature $\dot{\mathcal{L}}$ and the axioms

DEFINITION 10.1.1. A signature $\dot{\mathcal{L}}$ suitable for writing out a PNL theory of first-order logic is given in Figure 6.

NOTATION 10.1.2. We introduce the following syntactic sugar:

- We may write $\text{sub}_o([a]r, t)$ as $r[a \mapsto t]$.
- We may write $\text{sub}_\iota([a]r, t)$ as $r[a \mapsto t]$.
- We may write both \approx_ι and \approx_o just as \approx .

Examples of this in use, follow immediately below.

10.2 The axioms: equality, substitution, first-order logic, and arithmetic

Equality

Axioms for equality $\approx : (\iota, \iota)$ and equality $\approx : (o, o)$ are given in Figure 7.

Substitution

Axioms for substitution sub_ι and sub_o are given in Figure 8.

We arguably abuse notation in Figure 8 when we use variables of sort ι and o as appropriate not necessarily giving them distinct names (e.g. in $(\text{sub}^*) X$ has sort ι , whereas in $(\text{sub} \Rightarrow)$ we use another variable also written X with sort o).

First-order logic

Axioms for (a shallow reflection of) first-order formulas as terms in PNL (the $\dot{\perp}$, \Rightarrow , and $\dot{\forall}$) are given in Figure 9.

(≈ 2)	$\forall X', X, Y', Y. (X' \approx X \wedge Y' \approx Y) \Rightarrow X' \text{ op } Y' \approx X \text{ op } Y$ ($op \in \{\dot{+}, \dot{*}, \dot{\Rightarrow}, \dot{\approx}\}$)
(≈ 1)	$\forall X', X. X' \approx X \Rightarrow op(X') \approx op(X)$ ($op \in \{\text{succ}\}$)
(≈ 0)	$\forall X. X \approx X$
($\approx \checkmark$)	$\forall Z', Z. Z' \approx Z \Rightarrow \checkmark([a]Z') \approx \checkmark([a]Z)$
($\approx \text{sub}$)	$\forall X', X, Y', Y. (X' \approx X \wedge Y' \approx Y) \Rightarrow op([a]X', Y') \approx op([a]X, Y)$ ($op \in \{\text{sub}_\iota, \text{sub}_o\}$)
($\approx o$)	$\forall Z', Z. Z' \approx Z \Rightarrow (\epsilon(Z') \Leftrightarrow \epsilon(Z))$
($\approx \iota$)	$\forall X', X. X' \approx X \Rightarrow \epsilon(X') \approx X$

We fill in sorts as appropriate. Thus, $\dot{\perp} \approx_o \dot{\perp}$ whereas $0 \approx_\iota 0$, and so on. The permission sets of all variables are equal to $\mathbb{A}^<$, and $a \in \mathbb{A}^<$.

Figure 7: EQU: axioms for equality as a PNL theory

(subvar)	$\forall X. \text{var}(a)[a \mapsto X] \approx X$
(sub#)	$\forall X, Z. Z[a \mapsto X] \approx Z$ ($\text{supp}(Z) = (b \ a) \cdot \mathbb{A}^<$)
(subsucc)	$\forall X', X. \text{succ}(X')[a \mapsto X] \approx \text{succ}(X'[a \mapsto X])$
(sub+)	$\forall X'', X', X. (X'' \dot{+} X')[a \mapsto X] \approx (X''[a \mapsto X] \dot{+} X'[a \mapsto X])$
(sub*)	$\forall X'', X', X. (X'' \dot{*} X')[a \mapsto X] \approx (X''[a \mapsto X] \dot{*} X'[a \mapsto X])$
(sub\Rightarrow)	$\forall X'', X', X. (X'' \dot{\Rightarrow} X')[a \mapsto X] \approx (X''[a \mapsto X] \dot{\Rightarrow} X'[a \mapsto X])$
(sub\approx)	$\forall X'', X', X. (X'' \dot{\approx} X')[a \mapsto X] \approx (X''[a \mapsto X] \dot{\approx} X'[a \mapsto X])$
(sub\checkmark)	$\forall X, Z. (\checkmark([b]Z))[a \mapsto X] \approx \checkmark([b](Z[a \mapsto X]))$ ($\text{supp}(Z) = (b \ a) \cdot \mathbb{A}^<$)
(subid)	$\forall X. X[a \mapsto \text{var}(a)] \approx X$

$a \in \mathbb{A}^<$ and $b \notin \mathbb{A}^<$. The permission set of X'' , X' , and X is equal to $\mathbb{A}^<$. The permission set of Z is equal to $(b \ a) \cdot \mathbb{A}^<$.

Figure 8: SUB: axioms for substitution as a PNL theory

Arithmetic

Given EQU, SUB, and FOL, it is not hard to write axioms for arithmetic in PNL. This is in Figure 10. Later on in Theorem 10.7.7 we prove that this *is* an axiomatisation of arithmetic in PNL.

10.3 Comments on the axioms

REMARK 10.3.1. SUB is a PNL rendering of the nominal algebra theory **naSUB** from Example 7.1.3; the universal quantifiers which are implicit in the nominal algebraic judgement-form are made explicit here. This is

$$\begin{array}{ll}
 (\dot{\Rightarrow}) & \forall Z', Z. \epsilon(Z' \dot{\Rightarrow} Z) \Leftrightarrow (\epsilon(Z') \Rightarrow \epsilon(Z)) \\
 (\dot{\forall}) & \forall Z. (\epsilon(\dot{\forall}([a]Z)) \Leftrightarrow \forall X. \epsilon(Z[a \mapsto X])) \\
 (\dot{\perp}) & \epsilon(\dot{\perp}) \Rightarrow \perp
 \end{array}$$

Here Z' and Z have sort o , permission set $\mathbb{A}^<$, and $a \in \mathbb{A}^<$.

Figure 9: FOL: axioms for first-order formulas as a PNL theory

$$\begin{array}{ll}
 \text{(PS0)} & \forall X. \text{succ}(X) \approx \dot{0} \Rightarrow \perp \\
 \text{(PSS)} & \forall X', X. \text{succ}(X') \approx \text{succ}(X) \Rightarrow X' \approx X \\
 \text{(P+0)} & \forall X. X \dot{+} \dot{0} \approx X \\
 \text{(P+succ)} & \forall X', X. X' \dot{+} \text{succ}(X) \approx \text{succ}(X') \dot{+} X \\
 \text{(P*0)} & \forall X. X \dot{*} \dot{0} \approx \dot{0} \\
 \text{(P*succ)} & \forall X', X. X' \dot{*} \text{succ}(X) \approx (X' \dot{*} X) \dot{+} X \\
 \text{(PInd)} & \forall Z. (\epsilon(Z[a \mapsto \dot{0}]) \Rightarrow \\
 & \quad (\forall X. (\epsilon(Z[a \mapsto X]) \Rightarrow \epsilon(Z[a \mapsto \text{succ}(X)]))) \Rightarrow \\
 & \quad \forall X. \epsilon(Z[a \mapsto X]))
 \end{array}$$

All variables have permission set $\mathbb{A}^<$, and $a \in \mathbb{A}^<$.

Figure 10: ARITH: axioms for arithmetic as a PNL theory

essentially the same axiomatisation as studied in [Gabbay and Mathijssen, 2006a; Gabbay and Mathijssen, 2008a]. Soundness and completeness are proved, so providing some formal sense in which the axioms of SUB are ‘right’.

In [Gabbay and Mathijssen, 2008c] first-order logic is equationally axiomatised using nominal algebra (so the axioms involve only equality). Because PNL is already a first-order logic, we can use \perp , \Rightarrow , and \forall directly to capture the behaviour of $\dot{\perp}$, $\dot{\Rightarrow}$, and $\dot{\forall}$. So note that FOL here is *not* the axiomatisation of [Gabbay and Mathijssen, 2008c]; there we had to work a little harder because the ambient logic, nominal algebra, was purely equational.

REMARK 10.3.2. Instead of the axioms for equality EQU, we could directly extend PNL by adding derivation rules Figure 5 as follows:

$$\frac{\Phi, r \approx s, \phi[X:=r], \phi[X:=s] \vdash \Psi}{(fa(r) \cup fa(s) \subseteq \text{supp}(X)) \quad (\approx\mathbf{S})} \quad \frac{\Phi, r \approx r \vdash \Psi}{\Phi \vdash \Psi} (\approx\mathbf{R})$$

REMARK 10.3.3. Every unknown has a sort, and a permission set.

Different choices of permission set may yield logically equivalent results. For example, in (**sublam**) it is not vital that $\text{supp}(Z)$ is *exactly* $(b \ a) \cdot \mathbb{A}^<$. The important point is that $a \notin \text{supp}(Z)$.

Similarly, in (**subapp**) it is not vital that $\text{supp}(X'') = \text{supp}(X')$; when we use the axiom we can instantiate X'' and X' to r'' and r' such that $fa(r'') \neq fa(r')$, and conversely if we take $\text{supp}(X'') \neq \text{supp}(X')$ then we can still instantiate X'' and X' to r'' and r' such that $fa(r'') = fa(r') \subseteq \text{supp}(X'') \cap \text{supp}(X')$.

10.4 First-order logic \mathcal{L}

We will use the atoms \mathbb{A}_ν from $\dot{\mathcal{L}}$ in Section 10 as variables of our first-order logic (this is not necessary, but it is convenient). So for this section, a, b, c, \dots will range over distinct atoms in \mathbb{A}_ν .

DEFINITION 10.4.1. Define **terms** and **formulas** of \mathcal{L} by:

$$\begin{aligned} t &::= a \mid 0 \mid \text{succ}(t) \mid t + t \mid t * t \\ \xi &::= t \approx t \mid \perp \mid \xi \Rightarrow \xi \mid \forall a. \xi \end{aligned}$$

Substitution $t'[a:=t]$ and $\xi[a:=t]$ is as usual for first-order logic. We write sequents $\Xi \vdash \chi$ where Ξ and χ are sets of formulas. Derivability is as usual for first-order logic.

DEFINITION 10.4.2. Define a mapping $(-)\cdot$ from terms and formulas of \mathcal{L} to terms of $\dot{\mathcal{L}}$ by:

$$\begin{aligned} (a)\cdot &= a & (0)\cdot &= \dot{0} \\ (\text{succ}(t))\cdot &= \text{succ}((t)\cdot) & (t' + t)\cdot &= (t')\cdot \dot{+} (t)\cdot \\ (t' * t)\cdot &= (t')\cdot \dot{*} (t)\cdot \\ (t' \approx t)\cdot &= (t')\cdot \dot{\approx} (t)\cdot & (\perp)\cdot &= \dot{\perp} \\ (\xi' \Rightarrow \xi)\cdot &= (\xi')\cdot \dot{\Rightarrow} (\xi)\cdot & (\forall a. \xi)\cdot &= \dot{\forall}[a](\xi)\cdot \end{aligned}$$

DEFINITION 10.4.3. Extend $(-)\cdot$ to first-order logic sequents $\Xi \vdash \chi$ as follows:

$$(\Xi \vdash \chi)\cdot = \epsilon(\dot{\forall}[a_1] \dots \dot{\forall}[a_n]((\xi_1 \wedge \dots \wedge \xi_k) \Rightarrow (\chi_1 \vee \dots \vee \chi_l)))\cdot$$

Here, $\Xi = \{\xi_1, \dots, \xi_k\}$, $\chi = \{\chi_1, \dots, \chi_l\}$, and the free variables of Ξ and χ are $\{a_1, \dots, a_n\}$ (in some order).

NOTATION 10.4.4. Write **S** for $\text{EQU} \cup \text{SUB} \cup \text{FOL}$.

LEMMA 10.4.5. $\text{S} \vdash (t'[a:=t])\cdot \approx (t')\cdot [a \mapsto (t)\cdot]$ and
 $\text{S} \vdash (\xi[a:=t])\cdot \approx (\xi)\cdot [a \mapsto (t)\cdot]$.

Proof. By routine inductions on t and ξ . ■

THEOREM 10.4.6 (Correctness). If $\Xi \vdash \chi$ is derivable in first-order logic then $S \vdash (\Xi \vdash \chi)$ is derivable in PNL.

Proof. By a long but routine inspection we can check that EQU, SUB, and FOL allow us to model the behaviour of ‘real’ first-order logic. We use Lemma 10.4.5. \blacksquare

10.5 Interpretation of first-order logic

We recall the usual definition of interpretations in first-order logic:

DEFINITION 10.5.1. A nominal **(first-order logic) interpretation** \mathcal{M} is a **carrier set** M , and elements:

$$0'' \in M, \quad succ'' \in M \rightarrow M,$$

$$+'' \in (M \times M) \rightarrow M, \quad \text{and} \quad *'' \in (M \times M) \rightarrow M.$$

It is convenient to fix some \mathcal{M} from here until Theorem 10.7.7.

DEFINITION 10.5.2. Define $Valu_{\mathbb{A}_\nu}(M)$ by:

$$Valu_{\mathbb{A}_\nu}(M) = \{\varepsilon \in \mathbb{A}_\nu \rightarrow M \mid \exists A \subseteq \mathbb{A}_\nu. A \text{ finite} \wedge \forall a, b \notin A. \varepsilon(a) = \varepsilon(b)\}$$

Call elements of $Valu_{\mathbb{A}_\nu}(M)$ \mathbb{A}_ν -**valuations** (to M). ε will range over \mathbb{A}_ν -valuations.

If $x \in M$ write $\varepsilon[a:=x]$ for the valuation mapping b to $\varepsilon(b)$ and mapping a to x :

$$\begin{aligned} \varepsilon[a:=x](a) &= x \\ \varepsilon[a:=x](b) &= \varepsilon(b) \end{aligned}$$

Give $\varepsilon \in Valu_{\mathbb{A}_\nu}(M)$ and $X \subseteq Valu_{\mathbb{A}_\nu}(M)$ a **pointwise** permutation action:

$$\begin{aligned} (\pi \cdot \varepsilon)(a) &= \varepsilon(\pi^{-1}(a)). \\ \pi \cdot X &= \{\pi \cdot \varepsilon \mid \varepsilon \in X\}. \end{aligned}$$

U, V will range over *finitely-supported* subsets of $Valu_{\mathbb{A}_\nu}(M)$ —so there exists some finite $A \subseteq \mathbb{A}_\nu$ such that for all π , if $\pi(a) = a$ for all $a \in A$ then $\pi \cdot U = U$.

REMARK 10.5.3. $Valu_{\mathbb{A}_\nu}(M)$ would normally just be called ‘the set of valuations’. We are more specific since we separately also have valuations on unknowns X (Definition 7.3.3).

PNL atoms are serving as variable symbols of \mathcal{L} . To conveniently apply nominal techniques, it is useful to restrict to valuations that are finite in the sense given in Definition 10.5.2. In any case, any term or formula will only contain finitely many atoms.

$$\begin{array}{ll}
(\mathbf{ps0}) & \forall a. \text{succ}(a) \approx 0 \Rightarrow \perp \\
(\mathbf{pss}) & \forall a', a. \text{succ}(a) \approx \text{succ}(a') \Rightarrow a \approx a' \\
(\mathbf{p+0}) & \forall a. a + 0 \approx a \\
(\mathbf{p+succ}) & \forall a', a. a' + \text{succ}(a) \approx \text{succ}(a') + a \\
(\mathbf{p*0}) & \forall a. a * 0 \approx 0 \\
(\mathbf{p*succ}) & \forall a', a. a' * \text{succ}(a) \approx (a' * a) + a \\
(\mathbf{pind}) & \xi[a:=0] \Rightarrow \\
& \quad \forall a. (\xi \Rightarrow \xi[a:=\text{succ}(a)]) \Rightarrow \\
& \quad \forall a. \xi \qquad \qquad \qquad \text{(every } \xi, \text{ every } a)
\end{array}$$

Figure 11: arithmetic: axioms for arithmetic in first-order logic

DEFINITION 10.5.4. We extend the interpretation to first-order logic syntax as follows:

$$\begin{array}{l}
\llbracket a \rrbracket_{\varepsilon}^{\#} = \varepsilon(a) \\
\llbracket 0 \rrbracket_{\varepsilon}^{\#} = 0^{\#} \\
\llbracket \text{succ}(t) \rrbracket_{\varepsilon}^{\#} = \text{succ}^{\#}(\llbracket t \rrbracket_{\varepsilon}^{\#}) \\
\llbracket t' + t \rrbracket_{\varepsilon}^{\#} = +^{\#}(\llbracket t' \rrbracket_{\varepsilon}^{\#}, \llbracket t \rrbracket_{\varepsilon}^{\#}) \\
\llbracket t' * t \rrbracket_{\varepsilon}^{\#} = *^{\#}(\llbracket t' \rrbracket_{\varepsilon}^{\#}, \llbracket t \rrbracket_{\varepsilon}^{\#}) \\
\llbracket \perp \rrbracket_{\varepsilon}^{\#} = 0 \\
\llbracket \xi' \Rightarrow \xi \rrbracket_{\varepsilon}^{\#} = \max\{1 - \llbracket \xi' \rrbracket_{\varepsilon}^{\#}, \llbracket \xi \rrbracket_{\varepsilon}^{\#}\} \\
\llbracket \forall a. \xi \rrbracket_{\varepsilon}^{\#} = \min\{\llbracket \xi \rrbracket_{\varepsilon[a:=x]}^{\#} \mid x \in M\} \\
\llbracket t' \approx t \rrbracket_{\varepsilon}^{\#} = 1 \text{ if } \llbracket t' \rrbracket_{\varepsilon}^{\#} = \llbracket t \rrbracket_{\varepsilon}^{\#} \text{ and } 0 \text{ otherwise}
\end{array}$$

DEFINITION 10.5.5. Call the formula ξ **valid** in \mathcal{M} when $\llbracket \xi \rrbracket_{\varepsilon}^{\#} = 1$ for all ε .

Call $\xi_1, \dots, \xi_k \vdash \chi_1, \dots, \chi_l$ **valid** in \mathcal{M} when $(\xi_1 \wedge \dots \wedge \xi_k) \Rightarrow (\chi_1 \vee \dots \vee \chi_l)$ is valid.

10.6 A theory of arithmetic in \mathcal{L}

DEFINITION 10.6.1. Define a first-order theory of **arithmetic** by the axioms in Figure 11.

An interpretation \mathcal{M} is a **model** of arithmetic when $\llbracket \xi \rrbracket_{\varepsilon}^{\#} = 1$ for ξ each of **(ps0)**, **(pss)**, **(p+0)**, **(p+succ)**, **(p*0)**, **(p*succ)**, and every instance of **(pind)**.

REMARK 10.6.2. (**pin**d) the induction axiom-scheme is of course of particular interest. We therefore unpack what its validity

$$\llbracket \xi[a:=0] \Rightarrow \forall a.(\xi \Rightarrow \xi[a:=succ(a)]) \Rightarrow \forall a.\xi \rrbracket^\# = 1 \quad (\text{every } \xi, \text{ every } a)$$

means, in a little more detail. For every a and ξ :

- If $\llbracket \xi[a:=0] \rrbracket_\varepsilon^\# = 1$, and
- if for every $x \in M$, $\llbracket \xi \rrbracket_{\varepsilon[a:=x]}^\# = 1$ implies that $\llbracket \xi[a:=succ(a)] \rrbracket_{\varepsilon[a:=x]}^\# = 1$,
- then for every $x \in M$, $\llbracket \xi \rrbracket_{\varepsilon[a:=x]}^\# = 1$.

In (**pin**d) we take ‘every a ’, and in (**PI**nd) we do not. This is because in (**PI**nd), a is α -convertible,

10.7 Building an interpretation for $\dot{\mathcal{L}}$ out of one for \mathcal{L}

Recall the PNL signature $\dot{\mathcal{L}}$ from Section 10. Suppose \mathcal{M} is a model of arithmetic. We use it to build an interpretation \mathcal{N} of $\dot{\mathcal{L}}$.

DEFINITION 10.7.1. Extend \mathcal{L} to $\mathcal{L}+M$ where we add all elements of M as constants, and extend the interpretation to interpret these constants as themselves in M . (So if $x \in M$ then x is a constant symbol in $\mathcal{L}+M$ and $\llbracket x \rrbracket_\varepsilon^\# = x$.)

Define an \mathbb{A}_ν -valuation $\varepsilon_0 \in \text{Valu}_{\mathbb{A}_\nu}(M)$ by

$$\varepsilon_0(a) = 0^\# \quad \text{always.}$$

If t is a term, we write $\llbracket t \rrbracket_\varepsilon^\#$ for the function $\lambda\varepsilon.\llbracket t \rrbracket_\varepsilon^\#$. If ξ is a formula, we write $\llbracket \xi \rrbracket_\varepsilon^\#$ for the function $\lambda\varepsilon.\llbracket \xi \rrbracket_\varepsilon^\#$.

We now define an interpretation \mathcal{N} for $\dot{\mathcal{L}}$. We give a denotation to the base sorts ι and o of $\dot{\mathcal{L}}$, as follows:

$$\begin{aligned} \iota^\# &= \{\llbracket t \rrbracket_\varepsilon^\# \mid t \text{ a term of } \mathcal{L}+M\} \\ o^\# &= \{\llbracket \xi \rrbracket_\varepsilon^\# \mid \xi \text{ a formula of } \mathcal{L}+M\} \end{aligned}$$

We give a denotation to the term-formers and proposition-formers of $\dot{\mathcal{L}}$, as follows:

$$\begin{array}{ll} \text{var}^\# a \varepsilon = \varepsilon(a) & \text{sub}_o^\# ([a]u, v) \varepsilon = u(\varepsilon[a:=v\varepsilon]) \\ \dot{0}^\# \varepsilon = 0^\# & \dot{\Rightarrow}^\# (U, V) \varepsilon = \max\{1-U(\varepsilon), V(\varepsilon)\} \\ \text{succ}^\# u \varepsilon = \text{succ}^\#(u\varepsilon) & \dot{\forall}^\# [a]U \varepsilon = \min\{U(\varepsilon[a:=x]) \mid x \in M\} \\ \dot{+}^\# (u, v) \varepsilon = +^\#(u\varepsilon, v\varepsilon) & \dot{\approx}^\# (u, v) \varepsilon = \approx^\#(u\varepsilon, v\varepsilon) \\ \dot{*}^\# (u, v) \varepsilon = *^\#(u\varepsilon, v\varepsilon) & \approx_\iota^\# (u, v) = 1 \text{ if } u=v \text{ and } 0 \text{ otherwise} \\ \text{sub}_\iota^\# ([a]u, v) \varepsilon = u(\varepsilon[a:=v\varepsilon]) & \approx_o^\# (U, V) = 1 \text{ if } U=V \text{ and } 0 \text{ otherwise} \\ \dot{\perp}^\# \varepsilon = 0 & \varepsilon^\# U = U(\varepsilon_0) \end{array}$$

Here, u and v range over ι^ν and U and V range over σ^ν .

LEMMA 10.7.2.

1. $\llbracket t'[a:=t] \rrbracket_\varepsilon^\nu = \llbracket t' \rrbracket_{\varepsilon[a:=\llbracket t \rrbracket_\varepsilon^\nu]}^\nu$.
2. $\llbracket \xi[a:=t] \rrbracket_\varepsilon^\nu = 1$ if and only if $\llbracket \xi \rrbracket_{\varepsilon[a:=\llbracket t \rrbracket_\varepsilon^\nu]}^\nu = 1$.

LEMMA 10.7.3. The following equalities all hold:

$$\begin{array}{ll}
\text{var}^\nu(a) = \llbracket a \rrbracket^\nu & \text{sub}_\iota^\nu(\llbracket a \rrbracket^\nu, \llbracket t \rrbracket^\nu) = \llbracket t'[a:=t] \rrbracket^\nu \\
\dot{0}^\nu = \llbracket 0 \rrbracket^\nu & \text{sub}_\sigma^\nu(\llbracket a \rrbracket^\nu, \llbracket s \rrbracket^\nu) = \llbracket \xi[a:=s] \rrbracket^\nu \\
\text{succ}^\nu(\llbracket t \rrbracket^\nu) = \llbracket \text{succ}(t) \rrbracket^\nu & \dot{\perp}^\nu = \llbracket \perp \rrbracket^\nu \\
\dot{+}^\nu(\llbracket t' \rrbracket^\nu, \llbracket t \rrbracket^\nu) = \llbracket t' + t \rrbracket^\nu & \dot{\Rightarrow}^\nu(\llbracket \xi' \rrbracket^\nu, \llbracket \xi \rrbracket^\nu) = \llbracket \xi' \Rightarrow \xi \rrbracket^\nu \\
\dot{*}^\nu(\llbracket t' \rrbracket^\nu, \llbracket t \rrbracket^\nu) = \llbracket t' * t \rrbracket^\nu & \dot{\forall}^\nu(\llbracket a \rrbracket^\nu, \llbracket \xi \rrbracket^\nu) = \llbracket \forall a. \xi \rrbracket^\nu \\
& \dot{\approx}^\nu(\llbracket r \rrbracket^\nu, \llbracket s \rrbracket^\nu) = \llbracket r \approx s \rrbracket^\nu
\end{array}$$

Proof. We compare Definitions 10.7.1 and 10.5.4. Most cases are immediate; we consider only the slightly less trivial ones:

$$\begin{array}{ll}
\text{var}^\nu(a) = (\lambda a. \lambda \varepsilon. \varepsilon(a))a & \text{Definition 10.7.1} \\
= (\lambda a. \llbracket a \rrbracket^\nu)a & \text{Definition 10.5.4} \\
= \llbracket a \rrbracket^\nu & \text{fact} \\
\text{sub}_\iota^\nu(\llbracket a \rrbracket^\nu, \llbracket t \rrbracket^\nu) = \lambda \varepsilon. \llbracket t' \rrbracket^\nu(\varepsilon[a:=\llbracket t \rrbracket^\nu \varepsilon]) & \text{Definition 10.7.1} \\
= \lambda \varepsilon. \llbracket t'[a:=t] \rrbracket^\nu & \text{Lemma 10.7.2}
\end{array}$$

Other cases are no harder. ■

LEMMA 10.7.4. \mathcal{N} (Definition 10.7.1) is a PNL interpretation.

Proof. We must check that:

- ι^ν and σ^ν are *permissive-nominal sets*.
By routine calculations. (In fact, ι^ν and σ^ν are *nominal sets*; that is, their elements all have finite support.)
 - *The functions defined in Definition 10.7.1 map elements of ι^ν , σ^ν , $[\mathbb{A}]\iota^\nu$, and $[\mathbb{A}]\sigma^\nu$ correctly to the appropriate sets.*
By Lemma 10.7.3.
 - ε^ν is *equivariant from σ^ν to $\{0, 1\}$* .
By routine calculations using the fact that $(a \ b) \cdot \varepsilon_0 = \varepsilon_0$.
-

LEMMA 10.7.5. If $(\Xi \vdash \chi)$ is valid in \mathcal{N} , then $\Xi \vdash \chi$ is valid in \mathcal{M} .

Proof. We calculate that if $(\Xi \vdash \chi) \cdot$ is valid in \mathcal{N} , then

$$\llbracket (\xi_1 \wedge \dots \wedge \xi_k) \Rightarrow (\chi_1 \vee \dots \vee \chi_l) \rrbracket_{\varepsilon_0}^M = 1$$

But the proposition written out above is closed, so for all valuations ε , $\llbracket (\xi_1 \wedge \dots \wedge \xi_k) \Rightarrow (\chi_1 \vee \dots \vee \chi_l) \rrbracket_{\varepsilon}^{\mathcal{M}} = 1$. ■

Recall from Notation 10.4.4 that we write **S** for **EQU** \cup **SUB** \cup **FOL**.

PROPOSITION 10.7.6. The axioms of **S** \cup **ARITH** are valid in \mathcal{N} .

Proof. By a routine verification. We consider the axiom (\forall) from Figure 9. We unpack definitions and see that we must prove that for every ξ in $\mathcal{L}+M$,

- $\forall x \in M. \varepsilon_0[a:=x] \in (\xi) \cdot$ if and only if
- $\varepsilon_0[a:=(t)] \in (\xi) \cdot$ for every t a term of $\mathcal{L}+M$.

This follows, because $\mathcal{L}+M$ has a constant symbol for every $x \in M$. Validity of the other axioms is no harder. ■

THEOREM 10.7.7. arithmetic, $\Xi \vdash \chi$ in first-order logic if and only if **S** \cup **ARITH** $\vdash (\Xi \vdash \chi) \cdot$ in PNL.

Proof. We prove two implications. The top-to-bottom implication follows using Theorem 10.4.6.

For the bottom-to-top implication, we reason as follows: Suppose **S** \cup **ARITH** $\vdash (\Xi \vdash \chi) \cdot$ in PNL. Choose an arbitrary interpretation \mathcal{M} of first-order logic that is a model of arithmetic, with carrier set M . By Soundness (Theorem 9.3.6) and Proposition 10.7.6, $(\Xi \vdash \chi) \cdot$ is valid in \mathcal{N} . By Lemma 10.7.5 $\Xi \vdash \chi$ is valid in \mathcal{M} . \mathcal{M} was arbitrary, so by completeness of first-order logic [Shoenfield, 1967, §4.2] it follows that $\Xi \vdash \chi$ is derivable. ■

11 FURTHER PROPERTIES OF PNL

11.1 More PNL theories

We briefly mention on how to express some familiar ‘nominal’ constructs in PNL.

Inductive types

Permissive-nominal logic can express the principles of nominal abstract syntax developed in [Gabbay and Pitts, 2001].

Suppose a base sort ι , a name sort ν , and term-formers

$$\text{var} : (\nu)\iota, \quad \text{app} : (\iota, \iota)\iota, \quad \text{and} \quad \text{lam} : ([\nu]\iota)\iota.$$

Fix an unknown $U : \iota$ and for brevity write $\phi[U:=r]$ as $\phi(r)$ for every ϕ . Suppose an axiom-scheme, for every ϕ :

$$\begin{aligned} \phi(\text{var}(a)) &\Rightarrow \\ \forall X.(\phi(X) &\Rightarrow \phi(\text{lam}([a]X))) \Rightarrow \\ \forall X, Y.(\phi(X) &\Rightarrow \phi(Y) \Rightarrow \text{app}(X, Y)) \Rightarrow \\ &\forall X.(\phi(X)) \end{aligned}$$

Here X and Y have sort ι and we make a fixed but arbitrary choice of atom $a \in \text{supp}(X)$.

We can also express this finitely, if we axiomatise a sort for predicates (as we did for arithmetic). Here is the axiom-scheme above made finite by using the theories EQU, SUB, and FOL from Section 10:

$$\begin{aligned} \forall Z.\epsilon(Z[a \mapsto \text{var}(a)]) &\Rightarrow \\ \forall X.(\epsilon(Z[a \mapsto X]) &\Rightarrow \epsilon(Z[a \mapsto \text{lam}([a]X)])) \Rightarrow \\ \forall X, Y.(\epsilon(Z[a \mapsto X]) &\Rightarrow \epsilon(Z[a \mapsto Y]) \Rightarrow \epsilon(Z[a \mapsto \text{app}(X, Y)])) \Rightarrow \\ &\forall X.\epsilon(Z[a \mapsto X]) \end{aligned}$$

The \mathbb{N} quantifier

Nominal sets support the \mathbb{N} -quantifier [Gabbay and Pitts, 2001]. PNL also includes the \mathbb{N} -quantifier; the way in which it does this is quite interesting, as we shall see in a moment.

\mathbb{N} has some distinctive properties which are reflected in nominal logic (NL) and the logic of FM sets (FM):

$$\frac{\forall x.(P(x) \Rightarrow \mathbb{N}a.Q(a, x))}{\forall x.\mathbb{N}a.(P(x) \Rightarrow Q(a, x))} \quad \frac{\forall x.\mathbb{N}a.\mathbb{N}b.(b a) \cdot x \approx x}{\mathbb{N}a.\mathbb{N}b.\forall x.(a \# x \Rightarrow b \# x \Rightarrow (b a) \cdot x \approx x)}$$

Here and below we write a double horizontal line for ‘is provably equivalent to’. \mathbb{N} appears absent from Permissive-Nominal Logic (PNL). It is ‘hiding’ in the permission sets. Corresponding propositions are, where $a, b \notin \text{supp}(X)$

$$\frac{\forall X.(P(X) \Rightarrow Q(a, X))}{\forall X.(P(X) \Rightarrow Q(a, X))} \quad \frac{\forall X.(b a) \cdot X \approx X}{\forall X.(b a) \cdot X \approx X}$$

We see from these examples that two things are happening: first, freshness conditions are hard-coded into the syntax by permission sets—and second, so is the \mathcal{U} -quantifier.

It is interesting to consider another example. In NL/FM:

$$\frac{\mathcal{U}a.P(a) \wedge \mathcal{U}a.Q(b)}{\mathcal{U}a.\mathcal{U}b.(P(a) \wedge Q(b))} \qquad \frac{\mathcal{U}a.P(a) \wedge \mathcal{U}a.Q(b)}{\mathcal{U}a.(P(a) \wedge Q(a))}$$

Correspondingly in PNL we have:

$$\frac{P(a) \wedge Q(b)}{P(a) \wedge Q(b)} \qquad \frac{P(a) \wedge Q(b)}{P(a) \wedge Q(a)}$$

It is easy to use the rule **(Ax)** from Figure 5 to construct a derivation proving that $P(a) \wedge Q(b)$ and $P(a) \wedge Q(a)$ are indeed logically equivalent in Permissive-Nominal Logic.

The π in **(Ax)** expresses that truth is preserved by permutative renaming, or in symbols: $\vdash \phi \Leftrightarrow \pi \cdot \phi$ always.

A permission set S can be viewed in two ways: as giving permission to instantiate using free atoms in S —but also as a form of \mathcal{U} for the atoms not in S .

11.2 Admissibility of Cut

We indicate how **(Cut)** is admissible in the presence of the other rules in Figure 5.

DEFINITION 11.2.1. Suppose $fa(r) \subseteq \text{supp}(X)$ and $r : \text{sort}(X)$. Define $\Phi[X:=r]$ by

$$\Phi[X:=r] = \{\phi[X:=r] \mid \phi \in \Phi\}.$$

Lemma 11.2.2 is proved by routine arguments like those in [Dowek *et al.*, 2010; Urban *et al.*, 2004]:

LEMMA 11.2.2. Suppose $Y \notin fV(t)$. Then

$$r[Y:=u][X:=t] = r[X:=t][Y:=u[X:=t]].$$

LEMMA 11.2.3. Suppose $fa(r) \subseteq \text{supp}(X)$ and $r : \text{sort}(X)$. Then

$$\Phi \vdash \Psi \quad \text{implies} \quad \Phi[X:=r] \vdash \Psi[X:=r].$$

Proof. By a routine induction on derivations. The case of **(Ax)** uses Lemmas 3.4.10 and 11.2.2. The case of **($\forall\mathbf{L}$)** uses Lemma 11.2.2. \blacksquare

LEMMA 11.2.4.

1. If there exists a derivation Δ of $\Phi \vdash \psi, \Psi$ then there exists a derivation of $\Phi \vdash \pi \cdot \psi, \Psi$.
2. If there exists a derivation Δ of $\Phi, \phi \vdash \Psi$ then there exists a derivation of $\Phi, \pi \cdot \phi \vdash \Psi$.

Proof. By a simultaneous induction on Δ . The case of $(\forall\mathbf{L})$ uses Lemma 3.4.10. (We need the *simultaneous* induction for $(\Rightarrow\mathbf{L})$ and $(\Rightarrow\mathbf{R})$, since parts of the proposition move between left and right.) ■

NOTATION 11.2.5. An instance of **(Cut)** rests on two sub-derivations. It is convenient to call them the **left branch** and **right branch** as illustrated:

$$\frac{\begin{array}{c} \vdots \textit{Left branch} \\ \Phi, \phi \vdash \Psi \end{array} \quad \begin{array}{c} \vdots \textit{Right branch} \\ \Phi \vdash \phi, \Psi \end{array}}{\Phi \vdash \Psi} \textit{(Cut)}$$

THEOREM 11.2.6 (Cut-elimination). If $\Phi \vdash \Psi$ is derivable with a derivation that uses **(Cut)**, then it is derivable with a derivation that does not use **(Cut)**.

Proof. The proof is as for first-order logic. The only differences are a π in **(Ax)** and a side-condition $fa(r) \subseteq \textit{supp}(X)$ in $(\forall\mathbf{L})$. These affect terms and have no effect on the structure of derivations; for the purposes of this proof they are irrelevant.

We commute instances of **(Cut)** upwards, as usual, following the method of [Dummett, 1977, pages 139-145] or [Gabbay, 2011a]. At each step, the following measure based on the depth of subderivations and the size of the cut formula, decreases:

- The size of the cut formula, and
- the longest path up the derivation the cut, that the formula persists, lexicographically ordered.
- The commutation cases between rules for \Rightarrow and \forall are as standard for first-order logic.
- The essential case for \Rightarrow is as standard.
- For the essential case for \forall , suppose the subderivation has the following form:

$$\frac{\frac{\Phi, \phi[X:=r] \vdash \Psi}{\Phi, \forall X.\phi \vdash \Psi} (\forall\mathbf{L}) \quad \frac{\begin{array}{c} \vdots \Delta \\ \Phi \vdash \phi, \Psi \end{array}}{\Phi \vdash \forall X.\phi, \Psi} (\forall\mathbf{R})}{\Phi \vdash \Psi} \textit{(Cut)}$$

By Lemma 11.2.3 there is a derivation $\Delta[X:=r]$ of $\Phi \vdash \phi[X:=r]$, Ψ . We eliminate the essential case as follows:

$$\frac{\Phi, \phi[X:=r] \vdash \Psi \quad \begin{array}{c} \vdots \\ \Delta[X:=r] \\ \Phi \vdash \phi[X:=r], \Psi \end{array}}{\Phi \vdash \Psi} \text{ (Cut)}$$

- Suppose the subderivation has the following form:

$$\frac{\frac{}{\Phi, \phi \vdash \pi \cdot \phi, \Psi} \text{ (Ax)} \quad \begin{array}{c} \vdots \\ \Delta \\ \Phi, \pi \cdot \phi \vdash \Psi \end{array}}{\Phi, \phi \vdash \Psi} \text{ (Cut)}$$

We use Lemma 11.2.4 to obtain a derivation Δ' of $\Phi, \phi \vdash \Psi$ (the transformations involved in the proof of Lemma 11.2.4 do not increase the inductive measure).

■

11.3 Exhausting the available atoms

We conclude with a brief discussion on a subtle point in the PNL design. Suppose a name sort ν , a base sort τ , and a proposition former $\# : (\nu, \tau)$. Suppose an atom a and an unknown $X : \tau$ with $\text{supp}(X) = \mathbb{A}^<$. Suppose an unknown $Y : \nu$ with $\text{supp}(Y) = \mathbb{A}^<$. Consider an interpretation in which $\#(a, X)$ is interpreted as $a \notin \text{supp}(\zeta(X))$ and τ is interpreted as \mathbb{L} (Definition 2.4.4).

That is, $\#$ is interpreted as freshness and τ is interpreted as well-orderings of permission-sets.

In the PNL of this paper, the interpretation of the proposition $\phi = \forall X. \exists Y. \#(Y, X)$ is false: we take $\zeta(X)$ to well-order $\mathbb{A}^<$ and there is no $a \in \text{supp}(Y)$ such that $a \notin \text{supp}(\zeta(X))$.

Suppose we decide that we want a version of PNL in which ϕ is true. In this case, we can consider denotations such that every element has support of the form $\pi \cdot \mathbb{A}^{\ll}$ where \mathbb{A}^{\ll} is infinite and $\mathbb{A}^{\ll} \subseteq \mathbb{A}^<$ and $\mathbb{A}^< \setminus \mathbb{A}^{\ll}$ is also infinite. In this way, an unknown X cannot ‘exhaust’ $\mathbb{A}^<$.

The lesson we draw from this small example is that nominal semantics offer a host of interesting and inspiring design options. In this paper, we have cut one of many possible paths through this design space; we make no claim to have arrived at a definitive truth. We hope that others will follow us in exploring this fruitful area.

12 CONCLUSIONS

This paper reflects a research arc by the author in collaboration with others, roughly from 2005 to the present day. Thanks to improvements in presentation and the use of permissive-nominal techniques, definitions and proofs are simpler than in previous literature, and new properties emerge.

In a sense this paper is a sequel to the survey of [Gabbay, 2011b] (written in 2008 and submitted in early 2009). But whereas [Gabbay, 2011b] concentrated on applications of nominal sets to syntax with binding, this paper considers nominal sets as a basis for meta-mathematics.

Hints of this appeared in nominal rewriting [Fernández *et al.*, 2004; Fernández and Gabbay, 2007], which allowed arbitrary (oriented) equality theories over nominal terms. Since then we have seen syntax, axioms, soundness and completeness results, and notions of algebra, all with nominal-style names and binding, and based on nominal sets. Perhaps unwisely, we shall succumb to following wordplay: [Gabbay, 2011b; Gabbay and Pitts, 2001] give denotations to syntax-with-binding whereas here, we give syntax to denotations with binding.

We have briefly constructed nominal sets (a permissive variant thereof). We saw nominal terms and explored their computational properties in nominal unification and rewriting. We considered algebra and proved soundness, completeness, and HSP over permissive-nominal sets. We gave nominal terms a \forall -quantifier for unknowns and used this to build a first-order logic. Finally, in an extended case study we gave finite axiomatisations of first-order logic and arithmetic and proved correctness.

Mathematical foundations influence language, and (famously) language influences thought. Nominal sets are a foundation with a model of names which is different from what has been considered before, so the question is: what new languages, and new thoughts, can emerge? This chapter attempts to address that question by illustrating the broad sweeps of what a ‘nominal’ meta-mathematics might look like.

We are not and cannot be encyclopaedic or exhaustive. For other work we should mention α Prolog, which allows Horn clauses [Cheney and Urban, 2008] (this preceded PNL, and could be viewed as a subset of it). The author in collaboration has proved correctness for several non-trivial theories in nominal syntaxes, including equational treatments of substitution [Gabbay and Mathijssen, 2006a; Gabbay and Mathijssen, 2008a], λ -calculus [Gabbay and Mathijssen, 2008b; Gabbay and Mathijssen, 2010], and first-order logic [Gabbay and Mathijssen, 2006c; Gabbay and Mathijssen, 2008c], as well as the finite first-order nominal axiomatisation of arithmetic [Dowek and Gabbay, 2010; Dowek and Gabbay, 2011] which we considered in Section 10. There are translations from nominal terms to λ -terms by Levy and Villaret and by Dowek, the author, and Mulligan [Levy and Villaret, 2008; Dowek *et al.*, 2010], including a translation of algebraic reasoning (so, not just

unification) [Gabbay and Mulligan, 2009]. There is also a translation to many-sorted first-order syntax by Kurz and Petrişan [Kurz and Petrişan, 2010], and a categorical treatment of nominal Lawvere theories in [Clouston, 2009]. It may also prove useful to consider nominal languages over nominal structures other than sets, for instance over nominal domains [Turner, 2009]. See also the ‘atlas of nominal languages’ in Appendix A.

This is all developing a topic which this author believes could become an immense field, the informal meta-level having been relatively unformalised until now for want of a denotation with names, which is what nominal sets provide.

The logic of FM sets, nominal logic, and the Nominal Isabelle package [Gabbay and Pitts, 2001; Pitts, 2003; Urban, 2008] are all first-order axiomatisations of nominal sets.¹⁸ What matters here is that in all cases, syntax is that of ‘ordinary’ first- or higher-logic.¹⁹ These are denotations for syntax-with-binding.

The logical systems in this paper take nominal sets as a given (literally; we use nominal abstract syntax) and build from there. This chapter has surveyed the author’s attempts, via results both syntactic and semantic, to outline what meta-mathematics could look like if it were based on a foundation of nominal sets.

So in this document we have explored the applications of names to meta-mathematics. But even that does not exhaust the potential applications of names. Mathematics and computer science are evolving in ways which will make understanding names more and more useful and relevant; nominal techniques have arisen as a consequence of this evolution.

Rather beautifully, in doing this we also revisit the roots of design decisions of mathematical foundations; whether to admit atoms—to sound more mathematical, we say *urelemente* and to sound less mathematical, we say *names*—into our mathematical universe, and if we do, what properties these elements should have. On the other hand, linguists would call these things *referents*, and they have been studying them for a long time.

Whatever these things are called, they exist and we use them all the time. So we will conclude with two slogans:

- *Names are data.*

¹⁸Essentially, [Gabbay and Pitts, 2001] is the first third of the author’s thesis; [Pitts, 2003] is the same but minus the cumulative sets hierarchy; and [Urban, 2008] is an extensive implementation in higher-order logic, with a large library of powerful macros. One reason this is non-trivial has to do with automatically deriving the *equivariance* properties described e.g. in [Gabbay, 2011b, Subsection 4.2].

¹⁹Sometimes, authors write ‘nominal logic’ for that logic obtained by adding for each atom a constant symbol to the syntax of first-order logic, and adding infinitely many axioms reflecting nominal sets (equalities of swapping atoms, fresh atoms, and so on). This is nominal sets wearing a ‘syntactic disguise’: consider by analogy a theory of arithmetic with a constant symbol for each number and an axiom for every arithmetic equality.

- *Names with additional properties are ubiquitous.*

This chapter has studied formal languages with which to specify some of the possible additional properties of names, such as ‘having a substitution action’ or ‘being universally quantifiable’. But more generally, by this combination of a new point of view and a rigorous mathematics, nominal techniques have the potential to simplify, factor out common properties, and help control some of a modern mathematics of logic and computation. The problem of names is not just a technical issue. It is a philosophical, foundational, linguistic, and computational issue; this is mathematics in the best sense of the word.

Dov Gabbay wrote in his preface to the second edition that

the researcher ... is having more and more in common with the traditional philosopher who has been analysing such questions for centuries (unrestricted by the capabilities of any hardware).
... I believe the day is not far away in the future when the computer scientist will wake up one morning with the realisation that he is actually a kind of formal philosopher!

We would add “and philosophers, linguists—and some artists too—may wake up one morning with the realisation that they are actually a kind of abstract computer scientist”. Amen.

Acknowledgements

We are grateful to Arnon Avron, Ranald Clouston, Nachum Dershowitz, Gilles Dowek, Dov Gabbay, Joanne Gabbay, Alexander Kurz, Dominic Mulligan, and Daniela Petrişan for research collaborations, support, advice, corrections, and suggestions. We thank the many anonymous referees who took the time to give detailed constructive criticism and to suggest improvements for the papers on which this survey is based. We are grateful to the Leverhulme Trust, to DIGITEO, and to the École Polytechnique in Paris. We acknowledge the support of grant RYC-2006-002131 at the Polytechnic University of Madrid. Thank you.

BIBLIOGRAPHY

- [Abadi *et al.*, 1991] Martín Abadi, Luca Cardelli, Pierre-Louis Curien, and Jean-Jacques Lévy. Explicit substitutions. *Journal of Functional Programming*, 1(4):375–416, 1991.
- [Baader and Nipkow, 1998] Franz Baader and Tobias Nipkow. *Term rewriting and all that*. Cambridge University Press, Great Britain, 1998.
- [Birkhoff, 1935] Garrett Birkhoff. On the structure of abstract algebras. *Proceedings of the Cambridge Philosophical Society*, 31:433–454, 1935.
- [Burris and Sankappanavar, 1981] S. Burris and H. Sankappanavar. *A Course in Universal Algebra*. Graduate texts in mathematics. Springer, 1981.

- [Calvès, 2010] Christophe Calvès. *Complexity and implementation of nominal algorithms*. PhD thesis, King's College London, 2010.
- [Cheney and Urban, 2003] James Cheney and Christian Urban. System description: Alpha-Prolog, a fresh approach to logic programming modulo alpha-equivalence. In *UNIF'03*, pages 15–19. Universidad Politécnica de Valencia, 2003.
- [Cheney and Urban, 2008] James Cheney and Christian Urban. Nominal logic programming. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 30(5):1–47, 2008.
- [Cheney, 2004] James Cheney. The complexity of equivariant unification. In *Proceedings of the 31st International Colloquium on Automata, Languages and Programming (ICALP 2004)*, volume 3142 of *Lecture Notes in Computer Science*, pages 332–344. Springer, 2004.
- [Cheney, 2005a] James Cheney. Relating nominal and higher-order pattern unification. In *Proceedings of the 19th International Workshop on Unification (UNIF 2005)*, pages 104–119, 2005.
- [Cheney, 2005b] James Cheney. A simpler proof theory for nominal logic. In *FoSSaCS*, volume 3441 of *Lecture Notes in Computer Science*, pages 379–394. Springer, 2005.
- [Cheney, 2006] James Cheney. Completeness and Herbrand theorems for nominal logic. *Journal of Symbolic Logic*, 71:299–320, 2006.
- [Cheney, 2010] James Cheney. Equivariant unification. *Journal of Automated Reasoning*, 45(3):267–300, October 2010.
- [Clouston and Pitts, 2007] Ranald A. Clouston and Andrew M. Pitts. Nominal equational logic. In *Computation, Meaning and Logic: Articles dedicated to Gordon Plotkin*, volume 172 of *Electronic Notes in Theoretical Computer Science*, pages 223–257. Elsevier Science, 2007.
- [Clouston, 2007] Ranald Clouston. Closed terms (unpublished notes). <http://users.cecs.anu.edu.au/~rclouston/closedterms.pdf>, 2007.
- [Clouston, 2009] Ranald Clouston. *Equational logic for names and binding*. PhD thesis, University of Cambridge, UK, 2009.
- [Clouston, 2011] Ranald Clouston. Nominal lawvere theories. In *Proceedings of the 18th International Workshop on Logic, Language, and Information (WoLLIC)*, volume 6642 of *Lecture Notes in Computer Science*. Springer, 2011.
- [Dershowitz and Jouannaud, 1989] Nachum Dershowitz and Jean-Pierre Jouannaud. Rewrite Systems. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science: Formal Methods and Semantics*, volume B. North-Holland, 1989.
- [Dowek and Gabbay, 2010] Gilles Dowek and Murdoch J. Gabbay. Permissive Nominal Logic. In *Proceedings of the 12th International ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming (PPDP 2010)*, pages 165–176, 2010.
- [Dowek and Gabbay, 2011] Gilles Dowek and Murdoch J. Gabbay. Permissive Nominal Logic (journal version). *Transactions on Computational Logic*, 2011. In press.
- [Dowek et al., 2009] Gilles Dowek, Murdoch J. Gabbay, and Dominic P. Mulligan. Permissive Nominal Terms and their Unification. In *Proceedings of the 24th Italian Conference on Computational Logic (CILC'09)*, 2009.
- [Dowek et al., 2010] Gilles Dowek, Murdoch J. Gabbay, and Dominic P. Mulligan. Permissive Nominal Terms and their Unification: an infinite, co-infinite approach to nominal techniques (journal version). *Logic Journal of the IGPL*, 18(6):769–822, 2010.
- [Dowek, 2001] Gilles Dowek. Higher-order unification and matching. In *Handbook of automated reasoning*, pages 1009–1062. Elsevier, 2001.
- [Dummett, 1977] Michael Dummett. *Elements of intuitionism*. Clarendon Press, 1 edition, 1977.
- [Fernández and Gabbay, 2007] Maribel Fernández and Murdoch J. Gabbay. Nominal rewriting (journal version). *Information and Computation*, 205(6):917–965, June 2007.
- [Fernández and Gabbay, 2010] Maribel Fernández and Murdoch J. Gabbay. Closed nominal rewriting and efficiently computable nominal algebra equality. In *Proceedings of the 5th International Workshop on Logical Frameworks and Meta-Languages (LFMTP 2010)*, 2010.

- [Fernández *et al.*, 2004] Maribel Fernández, Murdoch J. Gabbay, and Ian Mackie. Nominal Rewriting Systems. In *Proceedings of the 6th ACM SIGPLAN symposium on Principles and Practice of Declarative Programming (PPDP 2004)*, pages 108–119. ACM Press, August 2004.
- [Fiore and Hur, 2010] Marcelo Fiore and Chung-Kil Hur. Second-order equational logic. In *Proceedings of the 19th EACSL Annual Conference on Computer Science Logic (CSL 2010)*, Lecture Notes in Computer Science, 2010.
- [Fiore *et al.*, 1999] Marcelo P. Fiore, Gordon D. Plotkin, and Daniele Turi. Abstract syntax and variable binding. In *Proceedings of the 14th IEEE Symposium on Logic in Computer Science (LICS 1999)*, pages 193–202. IEEE Computer Society Press, 1999.
- [Gabbay and Cheney, 2004] Murdoch J. Gabbay and James Cheney. A Sequent Calculus for Nominal Logic. In *Proceedings of the 19th IEEE Symposium on Logic in Computer Science (LICS 2004)*, pages 139–148. IEEE Computer Society, July 2004.
- [Gabbay and Hofmann, 2008] Murdoch J. Gabbay and Martin Hofmann. Nominal renaming sets. In *Proceedings of the 15th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2008)*, pages 158–173. Springer, November 2008.
- [Gabbay and Mathijssen, 2006a] Murdoch J. Gabbay and Aad Mathijssen. Capture-avoiding Substitution as a Nominal Algebra. In *ICTAC 2006: Theoretical Aspects of Computing*, volume 4281 of *Lecture Notes in Computer Science*, pages 198–212, November 2006.
- [Gabbay and Mathijssen, 2006b] Murdoch J. Gabbay and Aad Mathijssen. Nominal Algebra. In *18th Nordic Workshop on Programming Theory*, October 2006.
- [Gabbay and Mathijssen, 2006c] Murdoch J. Gabbay and Aad Mathijssen. One-and-a-halfth-order logic. In *Proceedings of the 8th ACM-SIGPLAN International Symposium on Principles and Practice of Declarative Programming (PPDP 2006)*, pages 189–200. ACM, July 2006.
- [Gabbay and Mathijssen, 2007] Murdoch J. Gabbay and Aad Mathijssen. A Formal Calculus for Informal Equality with Binding. In *WoLLIC'07: 14th Workshop on Logic, Language, Information and Computation*, volume 4576 of *Lecture Notes in Computer Science*, pages 162–176. Springer, July 2007.
- [Gabbay and Mathijssen, 2008a] Murdoch J. Gabbay and Aad Mathijssen. Capture-Avoiding Substitution as a Nominal Algebra. *Formal Aspects of Computing*, 20(4-5):451–479, June 2008.
- [Gabbay and Mathijssen, 2008b] Murdoch J. Gabbay and Aad Mathijssen. The lambda-calculus is nominal algebraic. In Christoph Benzmüller, Chad Brown, Jörg Siekmann, and Rick Statman, editors, *Reasoning in simple type theory: Festschrift in Honour of Peter B. Andrews on his 70th Birthday*, Studies in Logic and the Foundations of Mathematics. IFCoLog, December 2008.
- [Gabbay and Mathijssen, 2008c] Murdoch J. Gabbay and Aad Mathijssen. One-and-a-halfth-order Logic. *Journal of Logic and Computation*, 18(4):521–562, August 2008.
- [Gabbay and Mathijssen, 2009] Murdoch J. Gabbay and Aad Mathijssen. Nominal universal algebra: equational logic with names and binding. *Journal of Logic and Computation*, 19(6):1455–1508, December 2009.
- [Gabbay and Mathijssen, 2010] Murdoch J. Gabbay and Aad Mathijssen. A nominal axiomatisation of the lambda-calculus. *Journal of Logic and Computation*, 20(2):501–531, April 2010.
- [Gabbay and Mulligan, 2009] Murdoch J. Gabbay and Dominic P. Mulligan. Universal algebra over lambda-terms and nominal terms: the connection in logic between nominal techniques and higher-order variables. In *Proceedings of the 4th International Workshop on Logical Frameworks and Meta-Languages (LFMTP 2009)*, pages 64–73. ACM, August 2009.
- [Gabbay and Pitts, 1999] Murdoch J. Gabbay and Andrew M. Pitts. A New Approach to Abstract Syntax Involving Binders. In *Proceedings of the 14th Annual Symposium on Logic in Computer Science (LICS 1999)*, pages 214–224. IEEE Computer Society Press, July 1999.

- [Gabbay and Pitts, 2001] Murdoch J. Gabbay and Andrew M. Pitts. A New Approach to Abstract Syntax with Variable Binding. *Formal Aspects of Computing*, 13(3–5):341–363, July 2001.
- [Gabbay, 2001] Murdoch J. Gabbay. *A Theory of Inductive Definitions with α -Equivalence*. PhD thesis, University of Cambridge, UK, March 2001.
- [Gabbay, 2005] Murdoch J. Gabbay. Axiomatisation of first-order logic (talk). In *Second workshop on Computational Aspects of Nominal sets (CANS'05)*. King's College, London, December 2005.
- [Gabbay, 2007a] Murdoch J. Gabbay. Fresh Logic. *Journal of Applied Logic*, 5(2):356–387, June 2007.
- [Gabbay, 2007b] Murdoch J. Gabbay. A General Mathematics of Names. *Information and Computation*, 205(7):982–1011, July 2007.
- [Gabbay, 2009] Murdoch J. Gabbay. Nominal Algebra and the HSP Theorem. *Journal of Logic and Computation*, 19(2):341–367, April 2009.
- [Gabbay, 2011a] Michael Gabbay. A proof-theoretic treatment of lambda-reduction with cut-elimination: lambda calculus as a logic programming language. *Journal of Symbolic Logic*, 76(2):673–699, June 2011.
- [Gabbay, 2011b] Murdoch J. Gabbay. Foundations of nominal techniques: logic and semantics of variables in abstract syntax. *Bulletin of Symbolic Logic*, 17(2):161–229, 2011.
- [Gabbay, 2011c] Murdoch J. Gabbay. Meta-variables as infinite lists in nominal terms unification and rewriting. *Logic Journal of the IGPL*, 2011. Accepted for publication.
- [Gabbay, 2011d] Murdoch J. Gabbay. Two-level nominal sets and semantic nominal terms: an extension of nominal set theory for handling meta-variables. *Mathematical Structures in Computer Science*, 2011. Published online.
- [Gentzen, 1935] Gerhard Gentzen. Untersuchungen über das logische Schließen [Investigations into logical deduction]. *Mathematische Zeitschrift* 39, pages 176–210, 405–431, 1935. Translated in [Szabo, 1969], pages 68–131.
- [Kurz and Petrişan, 2010] Alexander Kurz and Daniela Petrişan. On universal algebra over nominal sets. *Mathematical Structures in Computer Science*, 20:285–318, 2010.
- [Levy and Villaret, 2008] Jordi Levy and Mateu Villaret. Nominal unification from a higher-order perspective. In *Rewriting Techniques and Applications, Proceedings of RTA 2008*, volume 5117 of *Lecture Notes in Computer Science*. Springer, 2008.
- [Levy and Villaret, 2010] Jordi Levy and Mateu Villaret. An efficient nominal unification algorithm. In *Proceedings of the 21st International Conference on Rewriting Techniques and Applications (RTA 2010)*, volume 6 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 209–226. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2010.
- [Levy and Villaret, 2011] Jordi Levy and Mateu Villaret. Nominal unification from a higher-order perspective. *Transactions on Computational Logic (TOCL)*, 2011.
- [Mac Lane and Moerdijk, 1992] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Universitext. Springer, 1992.
- [Mathijssen, 2007] Aad Mathijssen. *Logical Calculi for Reasoning with Binding*. PhD thesis, Technische Universiteit Eindhoven, 2007.
- [Melliès, 1995] Paul-André Melliès. Typed lambda-calculi with explicit substitutions may not terminate. In Mariangiola Dezani-Ciancaglini and Gordon D. Plotkin, editors, *Proceedings of the 2nd International Conference on Typed Lambda Calculi and Applications, (TLCA 1995)*, volume 902 of *Lecture Notes in Computer Science*, pages 328–334. Springer, April 1995.
- [Miller *et al.*, 1989] Dale Miller, Gopalan Nadathur, Frank Pfenning, and Andre Scedrov. Uniform proofs as a foundation for logic programming. Technical report, Durham, NC, USA, 1989.
- [Pitts, 2001] Andrew M. Pitts. Nominal logic: A first order theory of names and binding. In N. Kobayashi and B. C. Pierce, editors, *Proc. 4th Int'l Symposium on Theoretical Aspects of Computer Software (TACS 2001)*, volume 2215 of *Lecture Notes in Computer Science*, pages 219–242. Springer, 2001.

- [Pitts, 2003] Andrew M. Pitts. Nominal logic, a first order theory of names and binding. *Information and Computation*, 186(2):165–193, 2003.
- [Pitts, 2011] Andrew Pitts. Nominal sets and their applications. In *Midlands Graduate School (MGS 2011)*, 11-15 April 2011. available online at cl.cam.ac.uk/~amp12/talks/MGS2011_nominal_sets_slides.pdf.
- [Prawitz, 1965] Dag Prawitz. *Natural deduction: a proof-theoretical study*. Almqvist and Wiksell, 1965. Reprinted by Dover, 2006.
- [Shoenfield, 1967] Joseph Shoenfield. *Mathematical Logic*. Addison-Wesley, 1967.
- [Smullyan, 1968] Raymond Smullyan. *First-order logic*. Springer, 1968. Reprinted by Dover, 1995.
- [Szabo, 1969] M. E. Szabo, editor. *Collected Papers of Gerhard Gentzen*. North Holland, 1969.
- [Turner, 2009] David C. Turner. *Nominal Domain Theory for Concurrency*. PhD thesis, University of Cambridge, 2009.
- [Tzevelekos, 2007] Nikos Tzevelekos. Full abstraction for nominal general references. In *Proceedings of the 22nd IEEE Symposium on Logic in Computer Science (LICS 2007)*, pages 399–410. IEEE Computer Society Press, 2007.
- [Urban *et al.*, 2003] Christian Urban, Andrew M. Pitts, and Murdoch J. Gabbay. Nominal Unification. In *CSL*, volume 2803 of *Lecture Notes in Computer Science*, pages 513–527. Springer, December 2003.
- [Urban *et al.*, 2004] Christian Urban, Andrew M. Pitts, and Murdoch J. Gabbay. Nominal Unification. *Theoretical Computer Science*, 323(1–3):473–497, September 2004.
- [Urban, 2008] Christian Urban. Nominal reasoning techniques in Isabelle/HOL. *Journal of Automatic Reasoning*, 40(4):327–356, 2008.

Part IV

Appendix

A AN ATLAS OF NOMINAL LANGUAGES

The reader coming to the nominal literature could be forgiven for finding it perplexing. What are ‘Fraenkel-mostowski sets’, ‘nominal sets’, ‘nominal terms’, ‘nominal logic’, ‘nominal rewriting and algebra’, ‘ α Prolog’, ‘nominal equational logic’, ‘permissive-nominal algebra’, ‘permissive-nominal logic’ (with/without *shift*-permutations)? In this Appendix we will give a brief annotated bibliography covering, loosely, the relevant publications. This list is not meant to be exhaustive.

Traditionally, nominal sets are understood as a tool for the mathematical analysis of syntax, as described for instance in the author’s previous survey/research paper [Gabbay, 2011b], or in slides of an excellent course of lectures by Pitts [Pitts, 2011]. This author takes a view of nominal sets not just as a foundation for syntax with binding, but as a foundation for mathematics itself—names and binding, after all, appear everywhere. The atlas below surveys relevant publications.

For each item in the list below, we reference where the idea was introduced to the ‘nominal’ literature, and any other relevant conference and journal papers.

FM set theory [Gabbay and Pitts, 1999; Gabbay and Pitts, 2001]. Fraenkel-Mostowski set theory (FM) and nominal sets (called ‘equivariant FM sets’ in that paper) are the foundational semantics for nominal techniques.

Fraenkel-Mostowski sets were already known and had been used for other purposes; see [Gabbay, 2011b, Remark 2.22] for more detailed historical comments. Nominal sets were familiar as e.g. the Schanuel topos. So both semantics were known.

What was new to [Gabbay and Pitts, 2001] was the observation by the author and Pitts of the notions of support, atoms-abstraction, the self-dual behaviour of the \mathbb{N} quantifier, and the application to what is now called *nominal abstract syntax*.²⁰

²⁰At the same time, Fiore Plotkin and Turi developed an approach to abstract syntax

Nominal logic [Pitts, 2001; Pitts, 2003]. The constructions of [Gabbay and Pitts, 2001] are repeated, but in a first-order axiomatisation of nominal sets rather than one of the FM cumulative hierarchy. Pitts also coined the catchy label ‘nominal’.

Sometimes authors identify the nominal logic of [Pitts, 2003] with nominal techniques in general. This is limiting, and it gets the mathematical development the wrong way round. Nominal logic is a Hilbert-style axiomatisation in first-order logic. These axioms have meaning because of the underlying nominal sets models, and not the other way around; nor does the axiomatisation *per se* contribute to new syntax or proof-theory with which to study names.

In order to make progress, we needed new syntax that more explicitly represents atoms and their properties.

Thus for instance the *nominal logic programming* developed by Cheney and Urban [Cheney and Urban, 2008] (also referenced below) is called logic-programming in nominal logic, but we also see from Figures 6, 7, and 8 of [Cheney and Urban, 2008] that the syntax and axioms used are a variant of nominal terms.

Proof-theories for the \mathbb{N} -quantifier [Gabbay, 2007a; Gabbay and Cheney, 2004; Cheney, 2005b]. Some attempts have been made to give the distinctive \mathbb{N} -quantifier of nominal techniques, a proof-theory. In arguably increasing order of elegance these are [Gabbay, 2007a] (this was received by the journal in 2003 but took four years to get printed), [Gabbay and Cheney, 2004] (written with Cheney to develop on [Gabbay, 2007a]), and [Cheney, 2005b].

The permissive-nominal logic (PNL) of this survey is another item on that list, and perhaps it is one of the nicest; certainly the PNL treatment of \mathbb{N} is very different from what has come before, see Subsection 11.1.

Complete semantics for this family of logics are in [Gabbay, 2007a], [Cheney, 2006], and in [Dowek and Gabbay, 2011]. See also Subsection 9.4 of this survey.

Nominal terms [Urban *et al.*, 2003; Urban *et al.*, 2004]. This new syntax introduced the distinctive freshness side-conditions and the nominal

which was really exactly the same thing [Fiore *et al.*, 1999]. The key difference turned out to be that nominal sets admit a relatively elementary sets-based interpretation of the presheaves. As argued in [Gabbay and Hofmann, 2008] there are ‘fewer presheaves’ in the nominal semantics, we feel that an elementary presentation of the mathematics—where this is possible—is a powerful advantage not just for the reader but also for the practicing theorist.

Fiore has continued this line of research in collaboration and produced logics which in some sense which has never been made formal, parallel the development here. For an example of this see [Fiore and Hur, 2010].

terms syntax, with its separation of atoms a and unknowns X into two syntactic classes. [Urban *et al.*, 2004] is where the syntactic ideas of this survey were born, if not the specific ‘permissive’ implementation, which came later (*permissive-nominal terms* below).

There is now quite a substantial body of work devoted to computing efficiently on nominal terms; notably [Calvès, 2010; Levy and Villaret, 2010]. There is also a body of work devoted to translating between nominal terms and *higher-order patterns* [Miller *et al.*, 1989]. We are far from exhaustive, but good places to start reading are [Cheney, 2005a], [Levy and Villaret, 2008; Levy and Villaret, 2011], and [Gabbay and Mulligan, 2009; Dowek *et al.*, 2010].

Nominal rewriting [Fernández *et al.*, 2004; Fernández and Gabbay, 2007] and α Prolog [Cheney and Urban, 2003; Cheney and Urban, 2008]. These were the first logical languages using nominal terms as a general-purpose assertion language; nominal rewriting was designed explicitly to allow us to assert (directed) equalities between terms such as β or η -equivalence. α Prolog was intended by its designers for reasoning on nominal abstract syntax, and explicitly presented as such—but in retrospect it can also be viewed as a general-purpose ‘nominal’ reasoning system in the same family as nominal rewriting and later work.²¹

Nominal algebra [Gabbay, 2005; Gabbay and Mathijssen, 2006a; Gabbay and Mathijssen, 2007; Gabbay and Mathijssen, 2009]. Nominal algebra is simply the undirected version of nominal rewriting.²² What makes nominal algebra interesting above and beyond nominal rewriting is the different theorems we prove about equality instead of rewriting; for instance the HSPA theorem of [Gabbay, 2009] (much simplified here in Section 8), and various correctness results for nominal algebra axiomatisations of e.g. substitution, λ -calculus, and first-order logic [Gabbay and Mathijssen, 2008a; Gabbay and Mathijssen, 2010; Gabbay and Mathijssen, 2008c].

The paper [Gabbay and Mathijssen, 2006a] is where the *permutative convention* of Definition 2.1.2 was introduced, used by the author consistently since then. This comes from the author’s work formalising nominal reasoning in Isabelle in [Gabbay, 2001] and spares us from having to explicitly enumerate all inequalities between atoms. Thus, if pressed to be entirely formal, ‘ a and b ’ refers to two meta-variables ranging over *distinct* atoms.

²¹James Cheney, private communication.

²²Actually, this is a simplification. There is a significant difference, which is described in [Fernández and Gabbay, 2010]: nominal rewriting does not have an explicit rule to generate fresh atoms, whereas nominal algebra does. To the level of detail we wish to go into here, this does not matter. The permissive-nominal syntax of this survey makes the issue obsolete because fresh atoms are a structural fact of the permission sets.

Kurz and Petrişan proved an HSP theorem for nominal algebra by treating nominal algebra as a kind of many-sorted first-order logic [Kurz and Petrişan, 2010]—the sorts are finite sets of atoms and come from the categorical view of nominal sets as presheaves. The effect of nominal theories can thus be attained in many-sorted first-order syntax. That syntax is just standard first-order syntax is potentially a big advantage, for instance if one wants to transfer results directly from universal algebra. This offers alternative and effective methods of semantic proof; e.g. [Kurz and Petrişan, 2010] significantly simplifies the proofs of [Gabbay, 2009]. We pay for this convenience with infinities; e.g. even the simplest theory is infinite since equalities are replicated at every sort. Of course, the theory may still be finitely presentable. Section 8 of the current paper contains another, further simplified, HSP proof.

Nominal equational logic [Clouston and Pitts, 2007; Clouston, 2009]. Clouston and Pitts developed a similar system to nominal algebra, which they called *Nominal Equational Logic (NEL)*. In [Clouston and Pitts, 2007] Clouston and Pitts claimed that NEL was significantly more complete than NA because it had something called ‘semantic freshness judgements’. However, they had misunderstood the Nominal Algebra judgement $\Delta \vdash a \# r$ as a logical judgement when it is a syntactic judgement (corresponding to $a \notin fa(r)$ in this permissive-nominal paper), and they had not noticed that semantic freshness is already expressible using equality ([Gabbay and Mathijssen, 2007, Theorem 5.5], [Gabbay and Mathijssen, 2009, Lemma 4.51]; in this paper see Proposition 7.6.1).

In [Clouston, 2009], Clouston and Pitts correctly retracted their previous claim.

The NEL family of logics has its own syntactic freshness judgement, written ‘ $supp(\vec{a}) \# (\nabla, \vec{a}, t)$ ’ in (most recently) [Clouston, 2011, Remark 3.2]. In the permissive-nominal context of this paper, this and the NA freshness judgement $\Delta \vdash a \# r$ become equal to each other and to $a \notin fa(r)$.

Permissive-nominal terms [Gabbay and Mulligan, 2009; Dowek et al., 2009; Dowek et al., 2010]. These simplify and improve classical nominal terms in two ways: we give explicitly the (countably infinite) atoms that may be free / are guaranteed to be fresh in every unknown, and since freshness information is stored directly we eliminate the need for freshness contexts. Thus good properties emerge: permissive-nominal terms can be constructed as nominal abstract syntax, we can directly choose a name fresh for a term (which is not possible in nominal terms without expanding the freshness context), and properties and proofs can then be expressed for terms alone, rather than for terms-in-freshness-context.

For instance, in classical nominal terms a solution of a nominal unifica-

tion problem is a pair of a substitution and a freshness context; a nominal rewrite rule is a left and a right-hand side term and a freshness context; the proof-theory of nominal algebra requires an explicit freshness rule to generate fresh atoms, and so on. In fact, manipulating nominal terms almost always requires us to manipulate an external structure representing freshness constraints.

In contrast permissive terms are ‘self-sufficient’, like ordinary syntax. Proofs and algorithms have more of the look and feel of ordinary syntax. We have seen how, in the body of this survey. A detailed treatment of permissive-nominal syntax, including a simple translation from the nominal terms of [Urban *et al.*, 2004] into permissive-nominal terms, is also in [Dowek *et al.*, 2010].

Permissive-nominal algebra ([Gabbay and Mulligan, 2009], and **Section 7**). The permissive-nominal algebra of Section 7 is another variant of nominal algebra. It uses permissive-nominal terms and has a significantly different and simpler proof-theory. The notable differences are, aside from being permissive-nominal, the inclusion (if we want them) of infinitely-supported constant symbols and of infinitely-supported permutations. So previous work is a special case of the general framework of this survey, but what we do here goes strictly beyond what was possible in previous work, also in some significant mathematical properties such as satisfying an HSP instead of an HSPA result; see the discussion opening Section 8.

Permissive-nominal logic ([Dowek and Gabbay, 2010; Dowek and Gabbay, 2011] and **Section 9**). As we discuss in this survey, permissive-nominal logic (PNL) adds universal quantification over unknowns X . This is non-evident for nominal terms because of their freshness contexts; in nominal terms X behaves like an element with cofinite support so we lose α -equivalence whereas in permissive-nominal terms X has coinfinite support and we can always α -rename bound atoms. We get a proof-theory which is pleasingly close to that of first-order logic, a sound and complete semantics, and we can axiomatise and prove correct a non-trivial and mathematically relevant theory, such as arithmetic.

Name	Intended model	Refs	Notes
FM set theory	Cumulative hierarchy / Cat. of FM sets	[Gabbay and Pitts, 1999; Gabbay and Pitts, 2001]	Previously used to prove independence of AC
Nominal / FM sets	Themselves	[Gabbay and Pitts, 1999; Gabbay and Pitts, 2001]	Nominal sets called ‘equivariant FM sets’ in these papers
Nom. logic	Schanuel topos / Cat. of nom. sets	[Pitts, 2001; Pitts, 2003]	States axioms of nominal sets, used word ‘nominal’
Nom. terms	Nom. sets	[Urban <i>et al.</i> , 2003; Urban <i>et al.</i> , 2004]	Introduced $a, X, a\#X, [a]X, \pi \cdot X$
Nom. rewriting	Nom. terms	[Fernández <i>et al.</i> , 2004; Fernández and Gabbay, 2007]	First framework for asserting general theories on nominal terms
α Prolog	Nom. terms	[Cheney and Urban, 2003; Cheney and Urban, 2008]	Intended as logic programming language for abstract syntax, but can be viewed more generally
Nom. algebra	Nom. sets	[Gabbay, 2005; Gabbay and Mathijssen, 2006a; Gabbay and Mathijssen, 2007; Gabbay and Mathijssen, 2009]	Axiomatisation & models for binders like $[a \mapsto t], \lambda a, \forall a$
Nom. equational logic	Nom. sets / Nom. Lawvere Theories	[Clouston and Pitts, 2007; Clouston, 2009]	Also provides semantics for nominal algebra
Permissive-nom. terms	Nom. sets or permissive-nom. sets	[Dowek <i>et al.</i> , 2009; Gabbay and Mulligan, 2009; Dowek <i>et al.</i> , 2010]	Eliminate freshness contexts; add <i>shift</i> -permutation; standardise α -equivalence
Permissive-nom. algebra	Permissive-nom. sets or nom. sets	[Gabbay and Mulligan, 2009], this survey	More expressive, esp. in presence of <i>shift</i> -permutation
Permissive-nom. logic	Permissive-nom. sets or nom. sets	[Dowek and Gabbay, 2010; Dowek and Gabbay, 2011], this survey	A first-order logic for nom. terms

Figure 12: Cheat-sheet of nominal languages