Alpha-Structural Recursion and Induction

Andrew M. Pitts Cambridge University Computer Laboratory Cambridge CB3 0FD, UK

Version distributed at the International Summer School On Applied Semantics Frauenchiemsee, Germany, 8-12 September 2005

Abstract

The *nominal* approach to abstract syntax deals with the issues of bound names and α -equivalence by considering constructions and properties that are invariant with respect to permuting names. The use of permutations gives rise to an attractively simple formalisation of common, but often technically incorrect uses of structural recursion and induction for abstract syntax modulo α -equivalence. At the heart of this approach is the notion of *finitely supported* mathematical objects. This paper explains the idea in as concrete a way as possible and gives a new derivation within higher-order logic of principles of α -structural recursion and induction for α -equivalence classes from the ordinary versions of these principles for abstract syntax trees.

| | Index of Defintions | 2 |
|---|--|----|
| 1 | Introduction | 3 |
| 2 | Nominal Syntax | 7 |
| 3 | Finite Support | 11 |
| 4 | Recursion and Induction Principles for α -Terms | 20 |
| 5 | Second α -Structural Recursion & Induction Theorems | 25 |
| 6 | Extended Example: Normalisation by Evaluation | 35 |
| 7 | Assessment | 45 |
| A | Proof of Second α -Structural Induction Theorem | 48 |
| В | Proof of Second α -Structural Recursion Theorem | 51 |
| | References | 53 |

Index of Definitions

ground type, 36

initial, 30

identity substitution, 38

initial algebra property, 29 initial environment, 43

 α -equivalence, 10 length of an α -term, 23 α -term. 11 logical relation, 44 $\beta\eta$ -conversion, 37 NBE, 35 $\beta\eta$ -long normal form, 36 neutral, 36 λ -calculus with letrec, 8 nominal set, 12 π -calculus, 9 nominal signature, 7 nominal subset, 17 action, 12 normalisation by evaluation, 35 algebra, 29 normalisation function, 37, 43 arity, 7 atom-abstraction, 29 primitive recursion, 24 atom-permutations, 12 atom-sorts, 7 recursion with varying parameters, 32 atoms, 7 reflection function, 42 reification function, 42 capture-avoiding substitution, 32, 37, 38 result sort, 7 choice function, 15 composition, 38 second α -structural induction theorem, 28 constructors, 7 second α -structural recursion theorem, 26 separated product, 27 data-sorts, 7 simple type symbols, 36 denotation, 39 some/any theorem, 18 discrete nominal set, 13 structural induction for nominal terms, 10 domain, 38 structural recursion for nominal terms, 10 substitution, 38 environment, 39 support of x, 12 equivariant subset, 17 supports, 12 FCB, 21 terms, 8 finite support principle, 16 transpositions, 12 finitely supported functions, 14 trivial action, 13 finitely supported inductive definitions, 16 finitely supported subsets, 14 updated environment, 39 first α -structural induction theorem, 24 updated substitution, 39 first α -structural recursion theorem, 21 free atoms, 13 vector-transpositions, 49 fresh, 18 freshness condition for binders, 21 freshness theorem, 19 fundamental property, 45

1 Introduction

"They [previous approaches to operational semantics] do not in general have any great claim to being *syntax-directed* in the sense of defining the semantics of compound phrases in terms of the semantics of their components."

—GD Plotkin, A Structural Approach to Operational Semantics, p 21 (Aarhus, 1981; reprinted as [23, p 32])

The above quotation and the title of the work from which it comes indicate the important role played by structural recursion and structural induction in programming language semantics. These are the forms of recursion and induction that fit the commonly used "algebraic" treatment of syntax. In this approach one specifies the syntax of a language at the level of abstract syntax trees (ASTs) by giving an algebraic signature. This consists of a collection of sorts s (one for each syntactic category of the language), and a collection of constructors K (also called "operators" or "function symbols"). Each such K comes with an arity consisting of a finite list (s_1, \ldots, s_n) of sorts and with a result-sort s. Then the ASTs over the signature can be described by inductively generated terms t: if K has arity (s_1, \ldots, s_n) and result sort s, and if t_i is a term of sort s_i for i = 1..n, then $K(t_1, ..., t_n)$ is a term of sort s. One gets off the ground in this inductive definition with the n=0instance of the rule for forming terms covering the case of *constants*, C (and one usually writes the term C() just as C). Recursive definitions and inductive proofs about programs following the structure of their ASTs are both clearer and less prone to error than ones using non-structural methods. However, this treatment of syntax does not take into account the fact that most languages that one deals with in programming semantics involve binding constructors. In the presence of binders many syntax-manipulating operations only make sense, or at least only have good properties, when we operate on syntax at a level of abstraction represented not by ASTs themselves, but by α -equivalence classes of ASTs.

It is true that this level of abstraction, which identifies terms differing only in the names of bound entities, can be reconciled with an algebraic treatment of syntax by using de Bruijn indexes [6]. The well-known disadvantage of this device is that it necessitates a calculus of operations on de Bruijn indexes that does not have much to do with our intuitive view of the structure of syntax. As a result there can be a big "coding gap" between statements of results involving binding syntax we would like to make and their de Bruijn versions; and (hence) it is easy to get things wrong. For this reason, de Bruijn-style representations of syntax may be more suitable for language implementations than for work on language semantics.

In any case, most of the work on semantics which is produced by humans rather than by computers sticks with ordinary ASTs involving explicit bound names and uses an informal approach to α -equivalence classes. This approach is signalled by

¹This includes the metatheory of "higher-order abstract syntax" [19], where the questions we are addressing are pushed up one meta-level to a single binding-form, λ -abstraction.

4 Andrew Pitts

a form of words such as "we identify expressions up to α -equivalence" and means that: (a) occurrences of "t" now really mean its α -equivalence class " $[t]_{\alpha}$ "; and (b) if the representative t for the class $[t]_{\alpha}$ is later used in some context where the particular bound names of t clash in some way with those in the context, then t will be changed to an α -variant whose bound names are fresh (i.e. are ones not used in the current context). In other words it is assumed that the "Barendregt variable convention" [2, Appendix C] is maintained dynamically. In the literature, the ability to change bound names "on the fly" is usually justified by the assertion that final results of constructions involving ASTs are independent of choice of bound names. A fully formal treatment has to prove such independence results and in this paper we examine ways, arising from the results of [13, 20], to reduce the burden of such proofs.

However, proving that pre-existing functions respect α -equivalence is only part of the story; in most cases a prior (or simultaneous) problem is to prove the existence of the required functions in the first place. To see why, consider the familiar example of capture-avoiding substitution (x:=t)t' of a λ -term t for all free occurrences of a variable x in a λ -term t'. To bring out the issues with binders more clearly, let us consider this operation not for the pure λ -calculus, but for an applied calculus that also has expressions for local recursive function declarations. Thus the terms are either variables (x, f, y, \ldots) , applications $(t_1 t_2)$, function abstractions $(\lambda x.t)$, or local recursive function declarations of the form letrec $f x = t_1$ in t_2 . The leftmost occurrence of the variable f in letrec f f in f binds all free occurrences of f in both f and f whereas the left-most occurrence of the variable f only binds free occurrences of f in f in f. A systematic way of specifying such patterns of binding and the associated notion of f equivalence is given in Section 2.2; for the moment I assume the reader can supply a suitable definition of f equivalence for f terms involving such letrec-expressions.

How does one define capture-avoiding substitution for such terms up to α -equivalence? In the vernacular of programming semantics, one might specify (x:=t)(-) by saying it has the following properties, where fv(t) indicates the finite set of free variables of t.

$$(x := t)y = \begin{cases} t & \text{if } y = x \\ y & \text{if } y \neq x \end{cases}$$
 (1)

$$(x := t)(t_1 t_2) = (x := t)t_1 (x := t)t_2$$
(2)

$$y \notin fv(t) \cup \{x\} \Rightarrow (x := t)\lambda y. t_1 = \lambda y. (x := t)t_1$$
(3)

$$y \notin fv(t_2) \cup \{f\} \& f, y \notin fv(t) \cup \{x\} \Rightarrow$$

 $(x := t) let rec f y = t_1 in t_2 = let rec f y = (x := t) t_1 in (x := t) t_2$ (4)

The condition on the equation in (3) should be familiar enough: there is no need to say what happens when y occurs free in t or when y = x, since we are working "up to α -equivalence" and can change λy . t_1 to an α -variant satisfying these conditions.

The same goes for the more complicated condition on the equation in (4): given x and t, we can change $letrec\ f\ y = t_1\ in\ t_2$ up to α -equivalence to ensure that f and g are distinct variables not occurring free in g and not equal to g; but we can also assume that g does not occur free in g, since that term lies outside the binding scope of g in the term g in the term g and g in the term g int

To see what this specification of (x := t)(-) really amounts to, let us restore the usually-invisible notation $[t]_{\alpha}$ for the α -equivalence class of a term t. Writing Λ for the set of terms and $\Lambda/=_{\alpha}$ for its quotient by α -equivalence $=_{\alpha}$, then capture-avoiding substitution of an α -equivalence class e for a variable x is a function $\hat{s}_{x,e} \in \Lambda/=_{\alpha} \to \Lambda/=_{\alpha}$. Every such function corresponds to a function $s_{x,e} \in \Lambda \to \Lambda/=_{\alpha}$ respecting $=_{\alpha}$, i.e. satisfying

$$t_1 =_{\alpha} t_2 \implies s_{x,e}(t_1) = s_{x,e}(t_2)$$
 (5)

(enabling us to define $\hat{s}_{x,e}([t]_{\alpha})$ as $[s_{x,e}(t)]_{\alpha}$). The requirements (1)–(4) mean that we want $s_{x,e}$ to satisfy:

$$s_{x,e}(y) = \begin{cases} e & \text{if } y = x \\ [y]_{\alpha} & \text{if } y \neq x \end{cases}$$
 (6)

$$s_{x,e}(t_1 t_2) = [t'_1 t'_2]_{\alpha}$$
 where $s_{x,e}(t_i) = [t'_i]_{\alpha}$ for $i = 1, 2$ (7)

$$y \notin fv(e) \cup \{x\} \Rightarrow$$

$$s_{x,e}(\lambda y.t_1) = [\lambda y.t_1']_{\alpha} \quad \text{where } s_{x,e}(t_1) = [t_1']_{\alpha}$$
(8)

$$y \notin \mathit{fv}(t_2) \cup \{f\} \ \& \ f, y \notin \mathit{fv}(e) \cup \{x\} \ \Rightarrow$$

$$s_{x,e}(letrec\ f\ y = t_1\ in\ t_2) = [letrec\ f\ y = t_1'\ in\ t_2']_{\alpha}$$
where $s_{x,e}(t_i) = [t_i']_{\alpha}$ for $i = 1, 2$. (9)

The problem is not one of proving that a certain well-defined function $s_{x,e}$ respects α -equivalence, but rather of proving that a function exists satisfying (5)—(9). Note that (6)—(9) do not constitute a definition of $s_{x,e}(t)$ by recursion on the structure of the AST t: even if we patch up the "where" conditions in clauses (7)–(9) by using some enumeration of ASTs to make the choices t'_i definite functions of $s_{x,e}(t_i)$, the fact still remains that clauses (8) and (9) and only specify what to do for certain pairs (y, t_1) , rather than for all such pairs. Of course it is possible to complicate the specification by saying what to do for λ - and letrec-terms that do not meet the preconditions in (8) and (9), thereby arriving at a way of constructing $s_{x,e}(t)$ for any t (either by giving up structural properties and using a less natural recursion on the height of trees; or by fixing an enumeration of variables and using structural recursion to define a more general operation of simultaneous substitution [27]). An alternative approach, and one that works with the original simple specification, is to construct functions by giving rule-based inductive definitions of their graphs, with the rules encoding the required properties of the function. One then has to prove (using rule-based induction) that the resulting relations are single-valued, total and respect $=_{\alpha}$. This is in principle a fully formal and widely applicable approach to constructing functions like $s_{x,e}$ using tools that in any case are part and parcel of structural operational semantics; but one that is extremely tedious to carry out. It would be highly preferable to establish a recursion principle that goes straight from definitions like (1)–(4) to the existence of the function $(x:=t)(-) \in \Lambda/=_{\alpha} \to \Lambda/=_{\alpha}$. We provide such a principle here for a general class of signatures in which binding information can be declared. We call it α -structural recursion and it comes with an associated induction principle, α -structural induction.

These recursion and induction principles for α -equivalence classes of ASTs are simplifications and generalisations of the ones introduced by Gabbay and the author in [13] as part of a new mathematical model of fresh names and name binding. That paper expresses its results in terms of an axiomatic set theory, based on the classical Fraenkel-Mostowski permutation model of set theory. In my experience this formalism impedes the take up within computer science of the new ideas contained in [13]. There is an essentially equivalent, but more concrete description of the model as standard sets equipped with some simple extra structure. These so-called *nominal sets* are introduced in [20] and I will use them here to express α -structural recursion and induction within "ordinary mathematics", or more precisely, within Church's higher-order logic [5].

How to read this paper

Having read the Introduction this far, impatient readers may wish to turn to Theorem 25 to see the statement of the α -structural recursion principle for λ -calculus with letrec-terms and how it is used to define $(x := t)(-) \in \Lambda/=_{\alpha} \to \Lambda/=_{\alpha}$ satisfying (1)-(4). To understand the statement of this theorem, they must then look up the definitions of "nominal set", "finite support" and the freshness relation (-) # (-) in Section 3. This recursion principle and the corresponding induction principle are generalised to arbitrary signatures with binding information in Section 5. The particular way of specifying binding information that we use (via nominal signatures [28]) is explained in Section 2. Section 4 gives a first, simple version of the α -structural recursion and induction principles that are derived from ordinary structural recursion/induction for ASTs by, roughly speaking, taking into account an implicit parameterisation by name-permutations. The reduction of the practically more useful principles of Section 5 to the simpler ones of Section 4 is quite involved and is relegated to the Appendices. Section 6 contains an extended example (on normalisation by evaluation for the simply-typed λ -calculus [3]) that not only uses α -structural recursion and induction, but also shows off some of the power of nominal sets and the notion of freshness of names that they support. The final Section 7 assesses this paper's "nominal" approach to abstract syntax in the context of related work, both from a mathematical perspective and from the perspective of automated theorem proving.

2 Nominal Syntax

The usual principles of structural recursion and induction are parameterised by an algebraic signature that specifies the allowed constructors for forming abstract syntax trees (ASTs) of each sort. In order to state principles of recursion and induction for α -equivalence classes of ASTs, we need to fix a notion of signature that also specifies the forms of binding that occur in the ASTs. As explained in the Introduction, we stick with the usual "nominal" approach in which bound entities are explicitly named. Any generalisation of the notion of algebraic signature to encompass constructors that bind names needs to specify how bound occurrences of names in an AST are associated with a binding site further up the syntax tree. There are a number of such mechanisms in the literature of varying degrees of generality [10, 15, 17, 22, 28]. Here we will use the notion of nominal signature [28]. It has the advantage of dealing with binding and α -equivalence independently of any considerations to do with variables, substitution and β -equivalence: bound names in a nominal signature may be of several different sorts and not just variables that can be substituted for. In common with the other cited approaches, nominal signatures only allow for constructors that bind a fixed number of names (and without loss of much generality, we can take that number to be one). There are certainly forms of binding occurring "in the wild" that do not fit comfortably into this framework (for example, in the full version of $F_{<:}$ with records and pattern-matching used in Part 2B of the "POPLMARK challenge" [1]). I believe that the notion of α -structural recursion given here can be extended to cover more general forms of statically scoped binding, such as those used by Pottier in his C α ml library [24]; but for simplicity's sake I will stick with constructors binding a fixed number of names.

2.1 Atoms

From a logical point of view (as opposed to a pragmatic one that also encompasses issues of parsing and pretty-printing), the names we use for making localised bindings in formal languages only need to be atomic, in the sense that the structure of names (of the same kind) is immaterial compared with the distinctions between names. Therefore we will use the term atom for such names. Throughout this paper we fix two sets: the set $\mathbb A$ of all atoms and the set $\mathbb A\mathbb S$ of all atom-sorts. We also fix a function $sort \in \mathbb A \to \mathbb A\mathbb S$ assigning sorts to atoms and assume that the sets $\mathbb A\mathbb S$ and $\mathbb A_a \triangleq \{a \in \mathbb A \mid sort(a) = a\}$ for each $a \in \mathbb A\mathbb S$, are all countably infinite.

2.2 Nominal signatures

A nominal signature Σ consists of a subset of the atom-sorts, $\Sigma_A \subseteq \mathbb{AS}$, a set Σ_D of **data-sorts** and a set Σ_C of **constructors**. Each constructor $K \in \Sigma_C$ comes with an **arity** σ and a **result sort** $s \in \Sigma_D$, and we write $K : \sigma \to s$ to indicate this

information. The arities σ of Σ are given as follows:

Atom-sorts: every atom-sort $a \in \Sigma_A$ is an arity.

Data-sorts: every data-sort $s \in \Sigma_D$ is an arity.

Unit arity: 1 is an arity.

Pair arities: if σ_1 and σ_2 are arities, then $\sigma_1 * \sigma_2$ is an arity.

Atom-binding arities: if $a \in \Sigma_A$ and σ is an arity, then $a \gg \sigma$ is an arity.

The **terms** t over Σ of each arity are defined as follows, where we write t: σ to indicate that t has arity σ .²

Atoms: If $a \in \mathbb{A}_a$ is an atom of sort a, then a : a.

Constructed terms: If $K : \sigma \to s$ is in Σ_C and $t : \sigma$, then K t : s.

Unit: $\langle \rangle$: 1 is the unique term of unit arity.

Pairs: If $t_1 : \sigma_1$ and $t_2 : \sigma_n$, then $\langle t_1, t_2 \rangle : \sigma_1 * \sigma_2$.

Atom-binding: If $a \in \mathbb{A}_a$ and $t : \sigma$, then $\langle a \rangle t : \langle a \rangle \sigma$.

We write $Ar(\Sigma)$ for the set of all arities over a nominal signature Σ , $\mathsf{T}(\Sigma)$ for the set of all terms over Σ , and $ar \in \mathsf{T}(\Sigma) \to Ar(\Sigma)$ for the function assigning to each term t the unique arity σ such that $t : \sigma$ holds. For each $\sigma \in Ar(\Sigma)$, we write $\mathsf{T}(\Sigma)_{\sigma}$ for the subset $\{t \in \mathsf{T}(\Sigma) \mid ar(t) = \sigma\}$ of terms of arity σ .

Example 1 (λ -calculus with letrec). Here is a nominal signature for the untyped λ -calculus [2]. There is a single atom-sort v for variables, and a single data-sort t for λ -terms.

| atom-sorts | data-sorts | constructors |
|------------|------------|--|
| V | t | V:v 	o t |
| | | $A:t*t\tot$ |
| | | $L: 	extsf{	iny v} 	extsf{	iny t} 	o 	extsf{	iny t}$ |

To illustrate the inter-mixing of the arity-formers for pairing and atom-binding that is allowed in a nominal signature, consider augmenting λ -calculus with the local recursive function declarations, letrec $f x = t_1$ in t_2 , that were used in the discussion of capture-avoiding substitution in the Introduction. Recall that free occurrences of x in t_1 are bound in letrec $f x = t_1$ in t_2 ; and free occurrences of f in either of f or f are bound in the term. To get the effect of this we can add to the above nominal signature a constructor

$$Letrec : \ll v \gg ((\ll v \gg t) \gg t) \rightarrow t$$
.

²Compared with [28, Definition 2.3] we only define *ground* terms, since we do not need to consider variables ranging over terms here.

So for example, the expression letrec $f x = f x in f(\lambda y.y)$ corresponds to the nominal term

$$Letrec \ll f \gg \langle \ll x \gg A \langle Vf, Vx \rangle, A \langle Vf, L \ll y \gg Vy \rangle \rangle$$

of arity t over this signature (where $f, x, y \in \mathbb{A}_{\mathsf{v}}$).

Example 2 (π -calculus). Here is a nominal signature for the version of the Milner-Parrow-Walker π -calculus given in [25, Definition 1.1.1]. There is an atom-sort chan for channel names and a data-sort proc for process expressions; but there are also auxiliary data-sorts gsum, for processes that are guarded sums, and pre, for prefixed processes.

| atom-sorts | data-sorts | constructors |
|------------|------------|-------------------------------------|
| chan | proc | Gsum: gsum 	o proc |
| | gsum | Par: proc * proc 	o proc |
| | pre | Res: «chan»proc $ ightarrow$ proc |
| | | Rep:proc	oproc |
| | | Zero: 1 ightarrow gsum |
| | | Pre:pre	ogsum |
| | | Plus: gsum * gsum 	o gsum |
| | | $Out: (chan * chan) * proc \to pre$ |
| | | In: chan * «chan »proc 	o pre |
| | | Tau:proc	opre |
| | | Match: (chan * chan) * pre 	o pre |

For example, the π -calculus process expression $\nu x((\overline{x}u.0 + \overline{y}v.0)|x(z).\overline{z}w.0)$ corresponds to the following nominal term of arity proc over this signature (where $x, u, y, v, z, w \in \mathbb{A}_{\mathsf{chan}}$):

```
Res < x > Par \langle Gsum Plus \langle Pre Out \langle \langle x, u \rangle, Gsum Zero \langle \rangle \rangle, \\ Pre Out \langle \langle y, v \rangle, Gsum Zero \langle \rangle \rangle \rangle, \\ Gsum Pre In \langle x, < z > Gsum Pre Out \langle \langle z, w \rangle, Gsum Zero \langle \rangle \rangle \rangle.
```

2.3 Ordinary structural recursion and induction

The terms over a nominal signature Σ are just the abstract syntax trees determined by an ordinary signature associated with Σ whose sorts are the arities of Σ , whose constructors are those of Σ , plus constructors for unit, pairs and atom-binding, and with atoms regarded as particular constants. Consequently we can use ordinary structural recursion to define functions on the set $\mathsf{T}(\Sigma)$ of terms over Σ ; and we can use ordinary structural induction to prove properties of those terms. The following two theorems give versions of these principals that we use later. We regard their proofs as standard and omit them.³

³Each theorem can be used to prove the other; and either of them can be proved using induction for the natural numbers once one has fixed upon a particular construction of abstract syntax trees.

10 Andrew Pitts

Theorem 3 (structural recursion for nominal terms). Let Σ be a nominal signature. Suppose we are given sets S_{σ} , for each $\sigma \in Ar(\Sigma)$, and elements

$$\begin{split} g_{\mathbf{a}} &\in \mathbb{A}_{\mathbf{a}} \to S_{\mathbf{a}} & (\mathbf{a} \in \Sigma_{\mathbf{A}}) \\ g_{K} &\in S_{\sigma} \to S_{\mathbf{s}} & ((K:\sigma \to \mathbf{s}) \in \Sigma_{\mathbf{C}}) \\ g_{1} &\in S_{1} & \\ g_{\sigma_{1}*\sigma_{2}} &\in S_{\sigma_{1}} \times S_{\sigma_{2}} \to S_{\sigma_{1}*\sigma_{2}} & (\sigma_{1}, \sigma_{2} \in Ar(\Sigma)) \\ g_{\alpha \mathbf{a} \Rightarrow \sigma} &\in \mathbb{A}_{\mathbf{a}} \times S_{\sigma} \to S_{\alpha \mathbf{a} \Rightarrow \sigma} & (\mathbf{a} \in \Sigma_{\mathbf{A}}, \sigma \in Ar(\Sigma)) \end{split}$$

(We write $X \times Y$ for the cartesian product of two sets X and Y; and write $X \to Y$ for the set of functions from X to Y.) Then there is a unique family of functions $(\hat{g}_{\sigma} \in \mathsf{T}(\Sigma)_{\sigma} \to S_{\sigma} \mid \sigma \in Ar(\Sigma))$ satisfying the following properties

$$\hat{g} a = g_{\mathsf{a}}(a) \tag{10}$$

$$\hat{g}(Kt) = g_K(\hat{g}t) \tag{11}$$

$$\hat{g}\langle\rangle = g_1 \tag{12}$$

$$\hat{g}\langle t_1, t_2 \rangle = g_{\sigma_1 * \sigma_2} \langle \hat{g} \, t_1, \hat{g} \, t_2 \rangle \tag{13}$$

$$\hat{g} \ll a \gg t = g_{\ll a \gg \sigma}(a, \hat{g} t) \tag{14}$$

where we have abbreviated $\hat{g}_{\sigma}(t)$ to \hat{g} t (since $\sigma = ar(t)$ is determined by t). \square

Theorem 4 (structural induction for nominal terms). Let Σ be a nominal signature and $S \subseteq \mathsf{T}(\Sigma)$ a set of terms over Σ . To prove that S is the whole of $\mathsf{T}(\Sigma)$ it suffices to show

$$(\forall \mathsf{a} \in \Sigma_{\mathsf{A}}, a \in \mathbb{A}_{\mathsf{a}}) \ a \in S \tag{15}$$

$$(\forall (K: \sigma \to \mathsf{s}) \in \Sigma_{\mathsf{C}}, t: \sigma) \ t \in S \ \Rightarrow \ K \ t \in S \tag{16}$$

$$\langle \rangle \in S \tag{17}$$

$$(\forall (\sigma_i \in Ar(\Sigma), t_i : \sigma_i \mid i = 1, 2)) \ t_1 \in S \ \& \ t_2 \in S \ \Rightarrow \langle t_1, t_2 \rangle \in S$$
 (18)

$$(\forall \mathsf{a} \in \Sigma_\mathsf{A}, a \in \mathbb{A}_\mathsf{a}, \sigma \in Ar(\Sigma), t : \sigma) \ t \in S \ \Rightarrow \ «a»t \in S \ . \tag{19}$$

2.4 α -Equivalence and α -terms

So far we have taken no account of the fact that atom-binding terms $* a * s * t$ should be identified up to renaming the bound atom a. Given a nominal signature Σ , the relation of α -equivalence, $t =_{\alpha} t' : \sigma$ (where $\sigma \in Ar(\Sigma)$ and $t, t' \in \mathsf{T}(\Sigma)_{\sigma}$) makes such identifications. It is inductively defined by the rules in Figure 1. They generalise to terms over a nominal signature a version of the definition of α -equivalence of λ -terms [16, p. 36] that is conveniently syntax-directed compared with the classic version [2, Definition 2.1.11]. It is easy to see that $=_{\alpha}$ is reflexive, symmetric and respects the various term-forming constructions for nominal syntax. Less

$$\begin{split} (=_{\alpha}\text{-}1) & \frac{\mathsf{a} \in \Sigma_{\mathsf{A}} \quad a \in \mathbb{A}_{\mathsf{a}}}{a =_{\alpha} \ a : \, \mathsf{a}} \qquad (=_{\alpha}\text{-}2) \ \frac{(K : \sigma \to \mathsf{s}) \in \Sigma_{\mathsf{C}} \quad t =_{\alpha} t' : \sigma}{K \ t =_{\alpha} K \ t' : \, \mathsf{s}} \\ \\ (=_{\alpha}\text{-}3) & \frac{\mathsf{a} \in \Sigma_{\mathsf{A}} \quad (=_{\alpha}\text{-}4) \ \frac{t_1 =_{\alpha} t'_1 : \sigma_1 \quad t_2 =_{\alpha} t'_2 : \sigma_2}{\langle t_1, t_2 \rangle =_{\alpha} \langle t'_1, t'_2 \rangle : \sigma_1 * \sigma_2} \\ \\ & \frac{\mathsf{a} \in \Sigma_{\mathsf{A}} \quad a, a', a'' \in \mathbb{A}_{\mathsf{a}} \quad a'' \notin atm \ \langle a, t, a', t' \rangle}{t \{a''/a\} =_{\alpha} t' \{a''/a'\} : \sigma} \\ \\ & \frac{t \{a''/a\} =_{\alpha} t' \{a''/a'\} : \sigma}{\langle a \rangle t' : \langle a \rangle \sigma} \\ \end{split}$$

In rule (= $_{\alpha}$ -5), $atm\ t$ indicates the finite set of atoms occurring in t; and $t\{a'/a\}$ indicates the term resulting from replacing all occurrences in t of the atom a by the atom a' (assumed to be of the same sort); both can defined by structural recursion (Theorem 3):

$$\begin{array}{rcl} atm \ a & = & \{a\} \\ atm(K \ t) & = & atm \ t \\ atm\langle\rangle & = & \emptyset \\ atm\langle t_1, t_2\rangle & = & atm \ t_1 \cup atm \ t_2 \\ atm \ and \ and$$

Figure 1: α -Equivalence of nominal terms

straightforward is the fact that $=_{\alpha}$ is transitive. This can be proved in a number of ways. My favourite way makes good use of the techniques we will be using later, based on the action of atom-permutations on terms; see [20, Example 1].

Definition 5. For each $\sigma \in Ar(\Sigma)$, we write $\mathsf{T}_{\alpha}(\Sigma)_{\sigma}$ for the quotient of $\mathsf{T}(\Sigma)_{\sigma}$ by the equivalence relation $(-) =_{\alpha} (-) : \sigma$. Thus the elements of $\mathsf{T}_{\alpha}(\Sigma)_{\sigma}$ are α -equivalence classes of terms of arity σ ; we write $[t]_{\alpha}$ for the class of t and refer to $[t]_{\alpha}$ as an α -term of arity σ over the nominal signature Σ .

3 Finite Support

The crucial ingredient in the formulation of structural recursion and induction for α -terms over a nominal signature is the notion of finite⁴ *support*. It gives a well-behaved way, phrased in terms of atom-permutations, of expressing the fact that

⁴Both Gabbay [12] and Cheney [4] develop more general notions of "small" supports. As Cheney's work shows, such a generalisation is necessary for some techniques of classical model theory to be applied; but finite supports suffice here.

atoms are fresh for mathematical objects. It turns out to agree with the obvious definition when the objects are finite data such as abstract syntax trees, but allows us to deal with freshness for the not so obvious case of infinite sets and functions. For example, the identity function on \mathbb{A} "mentions" every atom in its graph; nevertheless, it has empty support and any atom is fresh for it.

3.1 Nominal sets

Let Perm denote the set of all (finite, sort-respecting) **atom-permutations**; by definition, its elements are bijections $\pi: \mathbb{A} \leftrightarrow \mathbb{A}$ such that $\{a \in \mathbb{A} \mid \pi(a) \neq a\}$ is finite and $sort(\pi(a)) = sort(a)$ for all $a \in \mathbb{A}$. The operation of composing bijections gives a binary operation $\pi, \pi' \in Perm \mapsto \pi \circ \pi' \in Perm$

$$(\pi \circ \pi')(a) \triangleq \pi(\pi'(a)) \qquad (a \in \mathbb{A})$$

that makes Perm into a group; we write ι for the identity atom-permutation and π^{-1} for the inverse of π . Among the elements of Perm we single out **transpositions** $(a\ a')$ given by a pair of atoms of the same sort: $(a\ a')$ is the atom-permutation mapping a to a', mapping a' to a and leaving all other atoms fixed. It is a basic fact of group theory that every $\pi \in Perm$ is equal to a finite composition of such transpositions.

An **action** of Perm on a set X is a function $Perm \times X \to X$, whose effect on $(\pi,x) \in Perm \times X$ we write as $\pi \cdot x$ (with X understood), and which is required to have the properties: $\iota \cdot x = x$ and $\pi \cdot (\pi' \cdot x) = (\pi \circ \pi') \cdot x$, for all $x \in X$ and $\pi, \pi' \in Perm$. Given such an action and an element $x \in X$, we say that a set $A \subseteq \mathbb{A}$ of atoms **supports** x if $(a\ a') \cdot x = x$ holds for all atoms a and a' (of the same sort) that are not in A.

Definition 6. A **nominal set** is by definition a set X equipped with an action of Perm such that every element $x \in X$ is supported by some *finite* set of atoms.

If X is a nominal set and A_1 and A_2 are both finite sets of atoms supporting $x \in X$, then it is the case that $A_1 \cap A_2$ also supports x. To see this, suppose that a and a' are atoms of the same sort not in $A_1 \cap A_2$; we have to show $(a\ a') \cdot x = x$. This is certainly the case if a = a' (because $(a\ a) = \iota$); and if $a \neq a'$, picking any a'' in the infinite set $\mathbb{A} - (A_1 \cup A_2 \cup \{a,a'\})$, then $(a\ a') = (a\ a'') \circ (a'\ a'') \circ (a\ a'')$ is a composition of transpositions each of which fixes x (since for each of the three pairs of atoms, each element of the pair is either not in A_1 , or not in A_2), so itself fixes x, as required. It follows immediately from this intersection property of finite supports that in a nominal set X, each element $x \in X$ possesses a smallest finite support, which we write as $supp_X(x)$, or just supp(x) if X is clear from the context, and call the **support of** x in X.

Example 7. (i) Each set \mathbb{A}_a of atoms of a particular sort a is a nominal set once we endow it with the atom-permutation action given by $\pi \cdot a = \pi(a)$; as one might expect, $supp(a) = \{a\}$. It is not hard to see that the disjoint union

of nominal sets is again a nominal set. So since the set of all atoms is the disjoint union of \mathbb{A}_a as a ranges over atom-sorts, \mathbb{A} is a nominal set with atom-permutation action and support sets as for each individual \mathbb{A}_a .

(ii) Let Σ be a nominal signature. Using Theorem 3 we can define an atompermutation action on the sets $\mathsf{T}(\Sigma)_{\sigma}$ of terms over Σ of each arity $\sigma \in Ar(\Sigma)$:

$$\pi \cdot a \triangleq \pi(a)$$

$$\pi \cdot K t \triangleq K(\pi \cdot t)$$

$$\pi \cdot \langle \rangle \triangleq \langle \rangle$$

$$\pi \cdot \langle t_1, t_2 \rangle \triangleq \langle \pi \cdot t_1, \pi \cdot t_2 \rangle$$

$$\pi \cdot \langle a \rangle t \triangleq \langle \pi \cdot a \rangle \langle \pi \cdot t \rangle$$

Using Theorem 4 one can prove that this has the properties required of an atom-permutation action, that $a, a' \notin atm \, t \Rightarrow (a \ a') \cdot t = t$, and that $a \in atm \, t \, \& \, (a \ a') \cdot t = t \Rightarrow a = a'$. From these facts it follows that each $\mathsf{T}(\Sigma)_\sigma$ is a nominal set, with $supp(t) = atm \, t$, the finite set of atoms occurring in t.

(iii) Turning next to α -terms over Σ (Section 2.4), first note that the action of atom-permutations on terms preserves α -equivalence: this is a consequence of a general property (Theorem 12) of rule-based inductive definitions that we will establish at the end of Section 3.2. Therefore we get a well-defined action on α -terms by defining: $\pi \cdot [t]_{\alpha} = [\pi \cdot t]_{\alpha}$. For this action one finds that $T_{\alpha}(\Sigma)_{\sigma}$ is a nominal set with $supp([t]_{\alpha}) = fa \, t$, the finite set of **free atoms** of any representative t of the class $[t]_{\alpha}$, defined (using Theorem 3) by:

$$fa \ a \triangleq \{a\}$$

$$fa(K \ t) \triangleq fa \ t$$

$$fa\langle\rangle \triangleq \emptyset$$

$$fa\langle t_1, t_2\rangle \triangleq fa \ t_1 \cup fa \ t_n$$

$$fa \ll x \triangleq fa \ t - \{a\}.$$

(iv) Each set S becomes a nominal set, called the **discrete nominal set** on S, when we endow it with the **trivial action** of atom-permutations, given by $\pi \cdot s = s$ for each $\pi \in Perm$ and $s \in S$; in this case the support of each element is empty. In particular, we will regard the one-element set $1 = \{()\}$, the set of booleans $\mathbb{B} = \{true, false\}$ and the set of natural numbers $\mathbb{N} = \{0, 1, 2, \ldots\}$ as nominal sets in this way.

3.2 Products, functions and powersets

If X_1 and X_2 are nominal sets, then we get an action of atom-permutations on their cartesian product $X_1 \times X_2$ by defining $\pi \cdot (x_1, x_2)$ to be $(\pi \cdot x_1, \pi \cdot x_2)$, for each $(x_1, x_2) \in X_1 \times X_2$. If A_i supports $x_i \in X_i$ for i = 1, 2, then it is not hard to see that $A_1 \cup A_2$ supports $(x_1, x_2) \in X_1 \times X_2$. Thus $X_1 \times X_2$ is also a nominal set. Note that

$$supp((x_1, x_2)) = supp(x_1) \cup supp(x_2)$$
(20)

since we have already observed that $supp(x_1) \cup supp(x_2)$ supports (x_1, x_2) , so that $supp((x_1, x_2)) \subseteq supp(x_1) \cup supp(x_2)$; and conversely, if A supports (x_1, x_2) then it also supports each x_i , so that $supp(x_i) \subseteq supp((x_1, x_2))$.

Turning next to functions, if X and Y are nominal sets, then we get an action of atom-permutations on the set $X \to Y$ of all functions from X to Y by defining $\pi \cdot f$ to be the function mapping each $x \in X$ to $\pi \cdot (f(\pi^{-1} \cdot x)) \in Y$. If you have not seen this definition before, it may look more complicated than expected; however, note that it is equivalent to the requirement that function application be respected by atom-permutations:

$$\pi \cdot (f(x)) = (\pi \cdot f)(\pi \cdot x) . \tag{21}$$

More precisely, the definition of the action on functions is forced by the requirement that $X \to Y$ together with the usual application function be the *exponential* of X and Y in the cartesian closed category whose objects are sets equipped with an atom-permutation action and whose morphisms are functions preserving the action. Unlike the situation for cartesian product, not every element $f \in X \to Y$ is necessarily finitely supported with respect to this action (see Example 9 below). However, note that if f is supported by a finite set of atoms A, then $\pi \cdot f$ is supported by $\{\pi(a) \mid a \in A\}$. Therefore the set

$$X \to_{\mathrm{fs}} Y \triangleq \{ f \in X \to Y \mid (\exists \text{ finite } A \subseteq \mathbb{A}) \text{ } A \text{ supports } f \}$$

of **finitely supported functions** from X to Y is closed under the atom-permutation action and is a nominal set.

Given a nominal set X, we can use the usual bijection between subsets of X and functions in $X \to \mathbb{B}$ (where $\mathbb{B} = \{true, false\}$) to transfer the action of atompermutations on $X \to \mathbb{B}$ to one on subsets of X. From the definition of the action of atom-permutations on functions and using the fact that the action on \mathbb{B} is trivial (see Example 7(iv)), one can calculate that this action sends $\pi \in Perm$ and $S \subseteq X$ to the subset

$$\pi \cdot S \triangleq \{\pi \cdot x \mid x \in S\} .$$

Note that if S is supported by a set of atoms A with respect to this action, then $\pi \cdot S$ is supported by $\{\pi(a) \mid a \in A\}$. So the set

$$P_{\mathrm{fs}}(X) \; \triangleq \; \{S \subseteq X \mid (\exists \, \mathrm{finite} \, A \subseteq \mathbb{A}) \, \, A \, \, \mathrm{supports} \, \, S\}$$

of **finitely supported subsets** of the nominal set X is closed under the atom-permutation action on all subsets of X and hence is a nominal set.

Example 8. Recall that the elements of Perm are bijections from $\mathbb A$ to itself that respect sorts and leave fixed all but finitely many atoms. So each $\pi \in Perm$ is in particular a function $\mathbb A \to \mathbb A$. Regarding $\mathbb A$ as a nominal set as in Example 7(i), the action of atom-permutations on π qua function turns out to be the operation of conjugation: $\pi' \cdot \pi = \pi' \circ \pi \circ (\pi')^{-1}$. Hence the action of atom-permutations on $\mathbb A \to \mathbb A$ restricts to an action on Perm. One can prove that the finite set $\{a \in \mathbb A \mid \pi(a) \neq a\}$ supports π with respect to this action (and is in fact the smallest such set); so Perm is a nominal set.

Example 9. Not every function between nominal sets is finitely supported. For example, since the set A of atoms is countable, there are surjective functions in $\mathbb{N} \to \mathbb{A}$; but it is not hard to see that any $f \in \mathbb{N} \to_{fs} A$ must have a finite image (which is in fact the support of f). A more subtle example of a non-finitelysupported function is any **choice function**⁵ for the set \mathbb{A} of atoms, i.e. any function $choose \in (\mathbb{A} \to_{fs} \mathbb{B}) \to \mathbb{A}$ (where $\mathbb{B} = \{true, false\}$) satisfying $f(a) = true \Rightarrow$ f(choose(f)) = true, for all $f \in \mathbb{A} \to_{fs} \mathbb{B}$ and $a \in \mathbb{A}$. To see this, we suppose that *choose* is supported by some finite set $A \subseteq \mathbb{A}$ and derive a contradiction. Let $f \in \mathbb{A} \to \mathbb{B}$ be the function mapping $a \in \mathbb{A}$ to true if $a \notin A$ and to false if $a \in A$. It is not hard to see that A supports f; in particular $f \in \mathbb{A} \to_{\mathrm{fs}} \mathbb{B}$ and we can apply *choose* to obtain an atom $a_0 \triangleq choose(f)$. Let $a = sort(a_0)$. Since \mathbb{A}_{a} is infinite and $A \cup \{a_0\}$ is finite, there is some atom $a_1 \in \mathbb{A}_{\mathsf{a}}$ with $a_1 \neq a_0$ and $a_1 \notin A$. Since $a_1 \notin A$, $f(a_1) = true$ by definition of f; and so since *choose* is a choice function, we also have f(choose(f)) = true. By definition of f and a_0 this means that $a_0 = choose(f) \notin A$. Since $a_0, a_1 \notin A$ and A supports both choose and f, we have $(a_0 \ a_1) \cdot choose = choose$ and $(a_0 \ a_1) \cdot f = f$. Thus by (21), $a_1 = (a_0 \ a_1) \cdot a_0 = (a_0 \ a_1) \cdot choose(f) = ((a_0 \ a_1) \cdot choose)((a_0 \ a_1) \cdot f) =$ $choose(f) = a_0$, contradicting the fact that we picked a_1 to be different from a_0 and completing the proof.

We note a property of support with respect to function application that we use in what follows.

Lemma 10. If X and Y are nominal sets, $f \in X \to_{fs} Y$ and $x \in X$, then $supp(f(x)) \subseteq supp(f) \cup supp(x)$.

Proof. For all atoms a, a' of the same sort, if $a, a' \notin supp(f) \cup supp(x)$ then $(a\ a') \cdot f = f$ and $(a\ a') \cdot x = x$; so by (21) $(a\ a') \cdot f \ x = ((a\ a') \cdot f)((a\ a') \cdot x) = f\ x$. Thus $supp(f) \cup supp(x)$ supports $f\ x$ and hence $supp(f\ x)$ is contained in this finite set.

A corollary of this is that for nominal sets X and Y, the operation of function application $app_{X,Y}(f,x) \triangleq f x$ is an element of $(X \to_{\mathrm{fs}} Y) \times X \to_{\mathrm{fs}} Y$; indeed $app_{X,Y}$ is supported by the empty set of atoms. Similarly, it is not hard to see

⁵It was the lack of finite support for choice functions that motivated the original construction of the permutation model of set theory by Fraenkel and Mostowski.

that currying, $cur_{X,Y,Z}(f) \triangleq \lambda x \in X.\lambda y \in Y.f(x,y)$, determines an element $cur_{X,Y,Z} \in ((X \times Y \to_{\mathrm{fs}} Y) \to_{\mathrm{fs}} (X \to_{\mathrm{fs}} (Y \to_{\mathrm{fs}} Z)))$ with empty support. The constantly true function $true_X(x) \triangleq true$ and the equality function $eq_X(x,x') \triangleq if \ x = x' \ then \ true \ else \ false$ also determine elements $true_X \in (X \to_{\mathrm{fs}} \mathbb{B})$ and $eq_X \in (X \times X \to_{\mathrm{fs}} \mathbb{B})$ with empty support. It is for these reasons that the following general principle holds good.

Theorem 11 (**finite support principle**). Any function or relation that is defined from finitely supported functions and relations using higher-order logic is itself finitely supported.

Because of this, the collection of finitely supported functions and subsets of nominal sets forms a very rich collection that is closed under the usual constructions of informal mathematics.⁶ If we remain within pure higher-order logic over ground types for numbers and booleans, then we only get elements with empty support. However, if we add a ground type for the set A of atoms, a constant for the function $sort \in \mathbb{A} \to \mathbb{AS}$ (taking \mathbb{AS} to be a copy of \mathbb{N}) and constants for each atom, then the terms and formulas of higher-order logic describe functions and subsets which may have non-empty, finite support. Such a "higher-order logic with atoms" has been developed by Gabbay [12]. In this paper I stick with ordinary higher-order logic: by considering all functions and subsets rather than just finitely supported ones, one sometimes gets more information about a construction. A good example of this is provided by a cornerstone of programming language semantics, namely rule-based inductive definitions. Given a nominal set X, let R be a finitely supported set of rules for defining a subset of X; more precisely, let R be an element of the nominal set $P_{fs}(P_{fs}(X) \times X)$. As usual, a subset $S \subseteq X$ is closed under the rules in R if

$$(\forall (H,c) \in R) \ H \subseteq S \implies c \in S$$

and the smallest such subset, ind(R), is given by the intersection of all such closed subsets. If we worked systematically in "FM-HOL" [12], rather than using arbitrary subsets of X, we would only consider *finitely supported* subsets that are closed under the rules, and would replace ind(R) by $\bigcap \{S \in P_{\mathrm{fs}}(X) \mid (\forall (H,c) \in R) \mid H \subseteq S \Rightarrow c \in S\}$. However these two subsets coincide, as the following theorem shows.

Theorem 12 (finitely supported inductive definitions). Let X be a nominal set. For any set of rules $R \in P_{fs}(P_{fs}(X) \times X)$, the subset $ind(R) \subseteq X$ inductively defined by R is a finitely supported subset of X; indeed $supp(ind(R)) \subseteq supp(R)$.

Proof. Suppose a, a' are atoms of the same sort that are not in the support of R. We have to show that $(a \ a') \cdot ind(R) = ind(R)$. It suffices to just show $ind(R) \subseteq$

⁶The only exception being that the finite support property is not conserved by all uses of choice: see Example 9.

 $(a\ a')\cdot ind(R).$ (For applying $(a\ a')$ to both sides then gives the reverse inclusion.) For this it suffices to show that $(a\ a')\cdot ind(R)$ is closed under the rules in R, since ind(R) is the smallest such subset. But if $(H,c)\in R$, then $((a\ a')\cdot H,(a\ a')\cdot c)$ is also a rule in R, because $(a\ a')\cdot (H,c)\in (a\ a')\cdot R=R$ since $a,a'\notin supp(R).$ So if $H\subseteq (a\ a')\cdot ind(R)$, then $(a\ a')\cdot H\subseteq (a\ a')\cdot (a\ a')\cdot ind(R)=ind(R)$, so $(a\ a')\cdot c\in ind(R)$ since ind(R) is closed under the rule $((a\ a')\cdot H,(a\ a')\cdot c)$; and hence $c=(a\ a')\cdot (a\ a')\cdot c\in (a\ a')\cdot ind(R).$

In this paper we will confine ourselves to finitely supported *finitary* rules, i.e. those R only containing (hypothesis, conclusion)-pairs (H,c) for which H is a finite subset of X. Every finite subset of a nominal set is in particular a finitely supported subset: clearly $\{x_1,\ldots,x_n\}$ is supported by $supp(x_1)\cup\cdots\cup supp(x_n)$. Furthermore the action of atom-permutations on subsets of X clearly sends finite subsets to finite subsets. So the finite powerset $P_{\text{fin}}(X)$ is a nominal set when X is. For finitary rules we just have to check that R is a finitely supported subset of the nominal set $P_{\text{fin}}(X)\times X$ to conclude from the above theorem that the subset ind(R) it inductively defines is again finitely supported. For example, it is not hard to see that the rule set given schematically in Figure 1 is a subset of $P_{\text{fin}}(\mathsf{T}(\Sigma)\times\mathsf{T}(\Sigma))\times(\mathsf{T}(\Sigma)\times\mathsf{T}(\Sigma))$ that is supported by the empty set of atoms. Therefore by the theorem, the relation $=_{\alpha}$ of α -equivalence is supported by the empty set. So $=_{\alpha}$ is preserved by all atom-transpositions and hence also by any $\pi\in Perm$ (since each π is a composition of transpositions):

$$t =_{\alpha} t' : \sigma \implies \pi \cdot t =_{\alpha} \pi \cdot t' : \sigma . \tag{22}$$

We used this property in Example 7(iii) when discussing the nominal set structure of $T_{\alpha}(\Sigma)$.

3.3 Nominal subsets and quotients

If X is a nominal set and $S \in P_{\mathrm{fs}}(X)$ is a finitely supported subset of it, then S is not necessarily itself a nominal set, because for any $x \in S$ and $\pi \in Perm$ we have no guarantee that $\pi \cdot x$ again lies in S. But if $supp(S) = \emptyset$, then $(a\ a') \cdot S = S$ for all atoms a, a' of the same sort; and since each $\pi \in Perm$ is the composition of transpositions, in this case it follows that $\pi \cdot S = S$. From this it is not hard to see that the condition $supp(S) = \emptyset$ is equivalent to

$$(\forall \pi \in Perm, x \in X) \ x \in S \ \Rightarrow \ \pi \cdot x \in S \ . \tag{23}$$

So when (23) holds the action of atom-permutations on elements of X restricts to an action on S; and furthermore, the support of each $x \in S$ is the same as its support as an element of X. Therefore S is a nominal set. We call such an S a **nominal subset** of X. (The term **equivariant subset** is also commonly used for this)

Another useful way of forming nominal sets is by taking quotients. If \sim is an equivalence relation on a nominal set X, then so long as \sim is a nominal subset of

 $X \times X$, the usual set X/\sim of equivalence classes inherits a well-defined atompermutation action, given by $\pi \cdot [x] = [\pi \cdot x]$. Furthermore, an equivalence class is supported by any set of atoms that supports a representative of the class. So X/\sim is a nominal set. The construction of the nominal set $\mathsf{T}_\alpha(\Sigma)_\sigma$ of α -terms (of arity σ over a nominal signature Σ) from the nominal set $\mathsf{T}(\Sigma)_\sigma$ is an example of this construction, since we saw in (22) that α -equivalence is a nominal subset of $\mathsf{T}(\Sigma)_\sigma \times \mathsf{T}(\Sigma)_\sigma$.

3.4 Freshness

Given an element of a nominal set, most of the time we are interested not so much in its support as in the (infinite) set of atoms that are *not* in its support. More generally, if $x \in X$ and $y \in Y$ are elements of nominal sets, we write x # y when $supp_X(x) \cap supp_Y(y) = \emptyset$ and say that x is **fresh** for y.

Lemma 13. Let X and Y be nominal sets. For any $x \in X$, $y \in Y$, $\pi \in Perm$, and atoms $a, a' \in \mathbb{A}$ of the same sort,

$$\pi \cdot ((a \ a') \cdot x) = (\pi(a) \ \pi(a')) \cdot (\pi \cdot x) \tag{24}$$

$$x \# y \Rightarrow \pi \cdot x \# \pi \cdot y . \tag{25}$$

Proof. Equation (24) follows immediately from the fact that the atom-permutations $\pi \circ (a \ a')$ and $(\pi(a) \ \pi(a')) \circ \pi$ are always equal.

Given a set of atoms A, recall from the previous section that we write $\pi \cdot A$ for the set $\{\pi(a) \mid a \in A\}$. Note that since each permutation is in particular a bijection, it is the case that $\pi \cdot (A \cap A') = \pi \cdot A \cap \pi \cdot A'$. Therefore to prove (25) it suffices to show that $\pi \cdot supp(x) = supp(\pi \cdot x)$. But if A supports x, then for any atoms $a, a' \notin \pi \cdot A$ (of the same sort), we have $\pi^{-1}(a), \pi^{-1}(a') \notin A$ and hence $(\pi^{-1}(a) \pi^{-1}(a')) \cdot x = x$. Applying $\pi \cdot (-)$ to both sides of this equation and using (24), we get $(a \cdot a') \cdot (\pi \cdot x) = \pi \cdot x$. Thus $\pi \cdot A$ supports $\pi \cdot x$ when A supports x. So

$$(\forall \pi \in Perm)(\forall x \in X) \ supp(\pi \cdot x) \subseteq \pi \cdot supp(x)$$
.

Hence $supp(x) = supp(\pi^{-1} \cdot \pi \cdot x) \subseteq \pi^{-1} \cdot supp(\pi \cdot x)$ and thus we also have $\pi \cdot supp(x) \subseteq supp(\pi \cdot x)$. So we do indeed have $\pi \cdot supp(x) = supp(\pi \cdot x)$ and hence also (25).

Recall that the set of atoms \mathbb{A} is a nominal set as in Example 7(i). The following simple property of finitely supported sets of atoms is extremely useful when dealing with properties of fresh atoms; it subsumes [13, Proposition 4.10] and [20, Proposition 4].

Theorem 14 (some/any theorem). Let $S \in P_{fs}(\mathbb{A})$ be a set of atoms supported by some finite set of atoms A. For each atom-sort $a \in \mathbb{AS}$, the following are equivalent:

$$(\forall a \in \mathbb{A}_{\mathsf{a}}) \ a \notin A \ \Rightarrow \ a \in S \tag{26}$$

$$(\exists a \in \mathbb{A}_{\mathsf{a}}) \ a \notin A \ \& \ a \in S \ . \tag{27}$$

Proof. Since $\mathbb{A}_a - A$ is infinite, it is in particular non-empty; thus (26) implies (27). Conversely, suppose $a \in \mathbb{A}_a - A$ satisfies a # S. Given any other $a' \in \mathbb{A}_a - A$, we have to show a' # S; but $(a \ a') \cdot S = S$ (since $a, a' \notin A \supseteq supp(S)$) and hence $a' = (a \ a') \cdot a \in (a \ a') \cdot S = S$.

Example 15. Recall from Example 7(ii) that in the nominal set $\mathsf{T}(\Sigma)_\sigma$ of terms of arity σ over a nominal signature Σ , the support of a term t is just the finite set $atm\ t$ of atoms that occur in t. Furthermore, if $a' \notin atm\ t$, then the terms $t\{a'/a\}$ (replace a by a' throughout t) and $(a\ a) \cdot t$ (swap a and a' throughout t) are the same. Therefore the crucial rule $(=_\alpha$ -5) in the definition of α -equivalence of nominal terms can be rewritten as:

In particular we have

Applying Theorem 14 to the set of atoms $S = \{a'' \in \mathbb{A}_{\mathsf{a}} \mid (a \ a'') \cdot t =_{\alpha} (a' \ a'') \cdot t' : \sigma \}$, which (by Lemma 13) is supported by A = supp(a,t,a',t'), we get from (28) a useful property of α -equivalence of atom-binding terms:

$$(a*)t =_{\alpha} (a'*)t' : (a*)\sigma \Leftrightarrow$$

$$(\forall a'' \in \mathbb{A}_{\mathbf{a}}) \ a'' \# (a,t,a',t') \Rightarrow (a\ a'') \cdot t =_{\alpha} (a'\ a'') \cdot t' : \sigma .$$
 (29)

The next result provides a very general criterion for when a construction that "picks a fresh atom" is independent of which fresh atom is chosen.

Theorem 16 (freshness theorem). Given an atom-sort $a \in \mathbb{AS}$ and a nominal set X, if a finitely supported function $h \in \mathbb{A}_a \to_{fs} X$ satisfies

$$(\exists a \in \mathbb{A}_{\mathsf{a}}) \ a \# h \ \& \ a \# h(a) \tag{30}$$

then there is a unique element $fresh(h) \in X$ satisfying

$$(\forall a \in \mathbb{A}_{\mathsf{a}}) \ a \# h \ \Rightarrow \ h(a) = fresh(h) \ . \tag{31}$$

Furthermore, $supp(fresh(h)) \subseteq supp(h)$.

Proof. Given (30) we have to prove that h is constant on the non-empty set $\mathbb{A}_{\mathsf{a}} - supp(h)$. First note that by Theorem 14 (with $S \triangleq \{a \in \mathbb{A}_{\mathsf{a}} \mid a \# h(a)\}$ and $A \triangleq supp(h) \supseteq supp(S)$), if (30) holds then

$$(\forall a \in \mathbb{A}_a) \ a \# h \ \Rightarrow \ a \# h(a) \ . \tag{32}$$

Suppose $a, a' \in \mathbb{A}_{a} - supp(h)$. To see that h(a) = h(a'), without loss of generality we may assume $a \neq a'$. By Lemma 10, a # h(a') (since a # h and a # a'); and a' # h(a') holds by (32). Hence

$$h(a') = (a \ a') \cdot h(a') \qquad \text{since } (a, a') \# h(a')$$

$$= ((a \ a') \cdot h)((a \ a') \cdot a') \qquad \text{by (21)}$$

$$= h((a \ a') \cdot a') \qquad \text{since } (a, a') \# h$$

$$= h(a) .$$

So there is a unique element $fresh(h) \in X$ satisfying (31). To see that it is supported by supp(h), if a, a' are atoms (of the same sort) satisfying $(a, a') \notin h$, choosing any a'' in the infinite set $\mathbb{A}_{\mathsf{a}} - supp(h, a, a')$, we have $(a \ a') \cdot fresh(h) = (a \ a') \cdot h(a'') = ((a \ a') \cdot h)((a \ a') \cdot a'') = h(a'') = fresh(h)$. Thus supp(h) supports fresh(h).

4 Recursion and Induction Principles for α -Terms

Recall from Definition 5 that $T_{\alpha}(\Sigma)_{\sigma}$ denotes the set of α -terms of arity σ over a nominal signature Σ ; by definition these are α -equivalence classes $[t]_{\alpha}$ of terms $t:\sigma$. Elementary properties of the relation $=_{\alpha}$ of α -equivalence yield the following structural properties of α -terms; at the same time we introduce some concrete syntax for α -terms mirroring the informal notation for α -equivalence classes mentioned in the Introduction.

Atoms: if $a \in \Sigma_A$ and $e \in T_\alpha(\Sigma)_a$, then there is a unique $a \in \mathbb{A}_a$ such that $e = [a]_\alpha$. In this case we write e just as a.

Constructed α -terms: if $s \in \Sigma_D$ and $e \in T_{\alpha}(\Sigma)_s$, then there are unique $(K : \sigma \to s) \in \Sigma_C$ and $e' \in T_{\alpha}(\Sigma)_{\sigma}$ such that there exists t' with $e' = [t']_{\alpha}$ and $e = [K t']_{\alpha}$. In this case we write e as K e'.

Unit: $T_{\alpha}(\Sigma)_1$ contains a unique equivalence class, $[\langle \rangle]_{\alpha}$, which we write as ().

Pairs: if $\sigma_1, \sigma_2 \in Ar(\Sigma)$ and $e \in \mathsf{T}_\alpha(\Sigma)_{\sigma_1 * \sigma_2}$, then there are unique $e_i \in \mathsf{T}_\alpha(\Sigma)_{\sigma_i}$ for i=1,2 such that there exist t_i with $e_i=[t_i]_\alpha$ (i=1,2) and $e=[\langle t_1,t_2\rangle]_\alpha$. In this case we write e as (e_1,e_2) .

Atom-binding: if $a \in \Sigma_A$, $\sigma \in Ar(\Sigma)$ and $e \in T_\alpha(\Sigma)_{\text{«a»}\sigma}$, then for each $a \in \mathbb{A}_a$ with a # e (i.e. with a not a free atom of e—cf. Example 7(iii)), there is a unique $e' \in T_\alpha(\Sigma)_\sigma$ such that there exists t' with $e' = [t']_\alpha$ and $e = [\text{«a»}t']_\alpha$. In this case we write e as a.e'.

Using this notation we now give a first version of structural recursion for α -terms over a nominal signature. Compared with Theorem 3, the principle uses nominal sets rather than ordinary sets, and requires a common finite support for

the collection of functions in its hypothesis. Furthermore, the function supplied for each atom-binding arity must satisfy a *freshness condition for binders* (FCB) saying, roughly, that for *some* sufficiently fresh choice of the atom being bound, the result of the function can never contain that atom in its support. These conditions ensure that there is a unique (finitely supported) arity-indexed family of functions that is well-defined on α -equivalence classes and satisfies the required recursion equations—for *all* sufficiently fresh bound atoms, in the case of the recursion equation for binders.

Theorem 17 (first α -structural recursion theorem). Let Σ be a nominal signature. Suppose we are given an arity-indexed family of nominal sets $(X_{\sigma} \mid \sigma \in Ar(\Sigma))$ and elements

$$\begin{split} f_{\mathsf{a}} \in \mathbb{A}_{\mathsf{a}} \to_{\mathrm{fs}} X_{\mathsf{a}} & (\mathsf{a} \in \Sigma_{\mathsf{A}}) \\ f_{K} \in X_{\sigma} \to_{\mathrm{fs}} X_{\mathsf{s}} & ((K:\sigma \to_{\mathrm{fs}} \mathsf{s}) \in \Sigma_{\mathsf{C}}) \\ f_{1} \in X_{1} & \\ f_{\sigma_{1}*\sigma_{2}} \in X_{\sigma_{1}} \times X_{\sigma_{2}} \to_{\mathrm{fs}} X_{\sigma_{1}*\sigma_{2}} & (\sigma_{1},\sigma_{2} \in Ar(\Sigma)) \\ f_{\text{``a} \Rightarrow \sigma} \in \mathbb{A}_{\mathsf{a}} \times X_{\sigma} \to_{\mathrm{fs}} X_{\text{``a} \Rightarrow \sigma} & (\mathsf{a} \in \Sigma_{\mathsf{A}}, \sigma \in Ar(\Sigma)) \end{split}$$

$$(\exists a' \in \mathbb{A}_{\mathsf{a}}) \ a' \notin A \ \& \ (\forall x \in X_{\sigma}) \ a' \# f_{\mathsf{s}(\mathsf{a}),\sigma}(a',x) \ .$$
 (FCB)

Then there is a unique family of finitely supported functions $(\hat{f}_{\sigma} \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma} \to_{\mathrm{fs}} X_{\sigma} \mid \sigma \in Ar(\Sigma))$ with $supp(\hat{f}_{\sigma}) \subseteq A$ and satisfying the following properties for all a, e, e_1, \ldots, e_n of appropriate arity:

$$\hat{f}a = f_{\mathsf{a}}(a) \tag{33}$$

$$\hat{f}(Ke) = f_K(\hat{f}e) \tag{34}$$

$$\hat{f}() = f_1 \tag{35}$$

$$\hat{f}(e_1, e_2) = f_{\sigma_1 * \sigma_2}(\hat{f}e_1, \hat{f}e_2)$$
(36)

$$a \notin A \Rightarrow \hat{f}(a.e) = f_{\text{``a} \to \sigma}(a, \hat{f}e)$$
 (37)

where we have abbreviated $\hat{f}_{\sigma}(e)$ to $\hat{f}e$ and used the notation for α -terms introduced above.

Proof. We can reduce the proof of the theorem to an application of Theorem 3, taking advantage of the fact that we are working (informally) in higher-order logic. From the $Ar(\Sigma)$ -indexed family of nominal sets X_{σ} we define another such family: $S_{\sigma} \triangleq Perm \rightarrow_{fs} X_{\sigma}$ (regarding Perm as a nominal set as in Example 8 and using

⁷In other words the theorem is reducible to primitive recursion at higher types.

the \rightarrow_{fs} construct from Section 3.2). Now define elements g_a , g_K , g_1 , $g_{\sigma_1*\sigma_2}$ and $g_{\langle a \rangle \sigma}$ as in the statement of Theorem 3, as follows.

$$g_{\mathsf{a}} \ a \ \triangleq \lambda \pi \in Perm. \ f_{\mathsf{a}}(\pi(a))$$

$$g_{K} \ s \triangleq \lambda \pi \in Perm. \ f_{K}(s(\pi))$$

$$g_{1} \triangleq \lambda \pi \in Perm. \ f_{1}$$

$$g_{\sigma_{1} * \sigma_{2}}(s_{1}, s_{2}) \triangleq \lambda \pi \in Perm. \ f_{\sigma_{1} * \sigma_{2}}(s_{1}(\pi), s_{2}(\pi))$$

$$g_{\mathsf{\alpha} \mathsf{a} \times \sigma}(a, s) \triangleq \lambda \pi \in Perm. \ fresh(\lambda a' \in \mathbb{A}_{\mathsf{a}}. f_{\mathsf{\alpha} \mathsf{a} \times \sigma}(a', s(\pi \circ (a \ a'))))$$

The crucial clause in this definition is the last one, where we are using the fresh functional from Theorem 16 applied to the function $h \triangleq \lambda a' \in \mathbb{A}_a$. $f_{\text{"a} \to \sigma}(a', s(\pi \circ (a \ a')))$. For this to make sense it has to be the case that h is finitely supported and satisfies condition (30) of that lemma; let us see why this is so. Since $supp(f_{\text{"a} \to \sigma}) \subseteq A$ by assumption, it follows that h is supported by the finite set $A \cup supp(s,\pi,a)$. To see that (30) holds of h, let a' be the atom whose existence is asserted by (FCB); thus $a' \notin A$ and $a' \# f_{\text{"a} \to \sigma}(a',x)$ for any $x \in X_{\sigma}$. For any other $a'' \in \mathbb{A}_a - A$, we have $(a' \ a'') \cdot f_{\text{"a} \to \sigma} = f_{\text{"a} \to \sigma}$ (since $a', a'' \notin supp(f_{\text{"a} \to \sigma})$); hence applying $(a' \ a'')$ to $a' \# f_{\text{"a} \to \sigma}(a'',x)$, from Lemma 13 we get $a'' \# f_{\text{"a} \to \sigma}(a'',(a' \ a'') \cdot x)$ for any $x \in X_{\sigma}$. Choosing a'' to be in the infinite set $\mathbb{A}_a - (A \cup supp(s,\pi,a))$ and $x = (a' \ a'') \cdot s(\pi \circ (a \ a''))$, we conclude that a'' # h and $a'' \# f_{\text{"a} \to \sigma}(a'',s(\pi \circ (a \ a''))) = h(a'')$, as required for (30).

Applying Theorem 3 with this data, we get a family of functions

$$\hat{g}_{\sigma} \in \mathsf{T}(\Sigma)_{\sigma} \to (Perm \to_{\mathrm{fs}} X_{\sigma})$$

satisfying the recursion equations (10)–(14) of that theorem. Next one proves that these functions respect α -equivalence:

$$t_1 =_{\alpha} t_2 : \sigma \Rightarrow \hat{g}_{\sigma} t_1 = \hat{g}_{\sigma} t_2. \tag{38}$$

This is done by induction over the derivation of $t_1 =_{\alpha} t_2 : \sigma$ from the rules in Figure 1; the induction step for rule $(=_{\alpha}$ -5) uses the following property of \hat{g} , which follows by induction on the structure of t, i.e. using Theorem 4:

$$(\forall \sigma \in Ar(\Sigma), t : \sigma)(\forall \pi, \pi' \in Perm) \ \hat{q}_{\sigma} \ t \ (\pi \circ \pi') = \hat{q}_{\sigma}(\pi' \cdot t) \ \pi. \tag{39}$$

In view of (38), the functions \hat{g}_{σ} induce functions $\hat{f}_{\sigma} \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma} \to X_{\sigma}$ given by $\hat{f}_{\sigma}[t]_{\alpha} \triangleq \hat{g}_{\sigma} t \iota$ for any $t : \sigma$ (recalling that ι stands for the identity permutation). One proves that these functions \hat{f}_{σ} are all supported by A by first proving that the functions \hat{g}_{σ} are so supported; the latter follows from the uniqueness part of Theorem 3: if a, a' are atoms of the same sort not in A, then one can show that $(a \ a') \cdot \hat{g}_{\sigma}$ satisfies the same recursion equations as \hat{g}_{σ} and hence is equal to that function. The fact that the \hat{f}_{σ} satisfy the required recursion equations (33)–(37) follows from the recursion equations (10)–(14) satisfied by the \hat{g}_{σ} . That concludes the existence part of the proof of Theorem 17.

For the uniqueness part, suppose functions $f'_{\sigma} \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma} \to_{\mathrm{fs}} X_{\sigma}$ are all supported by A and satisfy the recursion equations (33)–(37) for \hat{f}_{σ} . Define $g'_{\sigma} \in \mathsf{T}(\Sigma)_{\sigma} \to S_{\sigma}$ by $g'_{\sigma} t \pi \triangleq f'_{\sigma}[\pi \cdot t]_{\alpha}$ ($\sigma \in Ar(\Sigma), t : \sigma, \pi \in Perm$). One can show that the g'_{σ} satisfy the same recursion equations (10)–(14) from Theorem 3 as the functions \hat{g}_{σ} ; so by the uniqueness part of that theorem, $g'_{\sigma} = \hat{g}_{\sigma}$. Therefore for all $t : \sigma, f'_{\sigma}[t]_{\alpha} = f'_{\sigma}[\iota \cdot t]_{\alpha} \triangleq g'_{\sigma} t \iota = \hat{g}_{\sigma} t \iota \triangleq \hat{f}_{\sigma}[t]_{\alpha}$; hence $f'_{\sigma} = \hat{f}_{\sigma}$.

Example 18 (length of an α -term). In [14, Section 3.3] Gordon and Melham give the usual recursion scheme for defining the length of a λ -term, remark that it is not a direct instance of the scheme developed in that paper (their Axiom 4) and embark on a detour via simultaneous substitutions to define the length function. This difficulty is analysed by Norrish [18, Section 3] on the way to his improved version of Gordon and Melham's recursion scheme (discussed further in Example 27 and Section 7). Pleasingly, the usual recursive definition of the length of a λ -term, or more generally of an α -term over any nominal signature, is a very simple application of the First α -Structural Recursion Theorem.⁸ Thus in Theorem 17 we take X_{σ} to be the discrete nominal set \mathbb{N} of natural numbers and

$$\begin{split} f_{\mathbf{a}} &\triangleq \lambda a \in \mathbb{A}_{\mathbf{a}}. \ 1 \\ f_{K} &\triangleq \lambda k \in \mathbb{N}. \ k+1 \\ f_{1} &\triangleq 0 \\ f_{\sigma_{1}*\sigma_{2}} &\triangleq \lambda(k_{1},k_{2}) \in \mathbb{N} \times \mathbb{N}. \ k_{1}+k_{2} \\ f_{\text{"a} > \sigma} &\triangleq \lambda(a,k) \in \mathbb{A}_{\mathbf{a}} \times \mathbb{N}. \ k+1 \ . \end{split}$$

These functions are all supported by $A=\emptyset$ and (FCB) holds trivially, because a # k holds for any $a \in \mathbb{A}$ and $k \in \mathbb{N}$. So the theorem gives us functions $\hat{f}_{\sigma} \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma} \to_{\mathsf{fs}} \mathbb{N}$. Writing $\operatorname{length} e$ for $\hat{f}_{\sigma} e$, we have the expected properties of a length function on α -terms:

$$length \ a=1$$
 $length(K \ e) = length \ e+1$ $length() = 0$ $length(e_1, e_2) = length \ e_1 + length \ e_2$ $length(a. \ e) = length \ e+1$.

Note that the last clause holds for all a, because in (37) the condition " $a \notin A$ " is vacuously true (since $A = \emptyset$).

Remark 19. In Theorem 17 we gave (FCB) as an existential statement. It is in fact equivalent to the universal statement

$$(\forall a' \in \mathbb{A}_{\mathsf{a}}) \ a \notin A \Rightarrow (\forall x \in X_{\sigma}) \ a' \# f_{\mathsf{a} \mathsf{a} \mathsf{b} \sigma}(a', x) \ .$$

⁸The same goes for Norrish's stripc function, used to illustrate the limitations of Gordon and Melham's workaround for the *length* function [18, p. 247].

This follows from the "some/any" Theorem 14 by taking S to be $\{a' \in \mathbb{A}_{\mathsf{a}} \mid (\forall x \in X_{\sigma}) \ a' \# f_{\mathsf{(a)},\sigma}(a',x)\}$ and checking that A supports S.

Remark 20 (**primitive recursion**). Theorem 17 gives a simple "iterative" form of structural recursion for α -terms, rather than a more complicated "primitive recursive" form with recursion equations

$$\begin{split} \hat{f}a &= f_{\mathsf{a}}(a) \\ \hat{f}(K\,e) &= f_K(e,\hat{f}e) \\ \hat{f}() &= f_1 \\ \hat{f}(e_1,e_2) &= f_{\sigma_1*\sigma_2}(e_1,e_2,\hat{f}e_1,\hat{f}e_2) \\ a \not\in A \Rightarrow \hat{f}(a.\,e) &= f_{\ll \mathsf{a} \rtimes \sigma}(a,e,\hat{f}e) \;. \end{split}$$

In fact this more general form can be deduced from the simple one given in the theorem by adapting to our nominal setting a similar result for ordinary structural recursion: defining $X'_{\sigma} \triangleq \mathsf{T}_{\alpha}(\Sigma)_{\sigma} \times X_{\sigma}$ and functions

$$\begin{split} f_{\mathsf{a}}'(a) &\triangleq (a, f_{\mathsf{a}}(a)) \\ f_{K}'(e, x) &\triangleq (K \, e, f_{K}(e, x)) \\ f_{1}' &\triangleq ((), f_{1}) \\ f_{\sigma_{1}*\sigma_{2}}'((e_{1}, x_{1}), (e_{2}, x_{2})) &\triangleq ((e_{1}, e_{2}), f_{\sigma_{1}*\sigma_{2}}(e_{1}, e_{2}, x_{1}, , x_{2})) \\ f_{\mathsf{``a} > \sigma}'(a, e, x) &\triangleq (a. \, e, f_{\mathsf{``a} > \sigma}(a, e, x)) \end{split}$$

we first apply Theorem 17 to get functions $\hat{f}'_{\sigma} \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma} \to_{\mathrm{fs}} \mathsf{T}_{\alpha}(\Sigma)_{\sigma} \times X_{\sigma}$. The uniqueness part of the theorem allows us to deduce that the first components of these functions are all identity functions; it follows from this that the second component of \hat{f}'_{σ} is a function $\hat{f}_{\sigma} \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma} \to_{\mathrm{fs}} X_{\sigma}$ satisfying the above scheme of primitive recursion (and is the unique such).

The next theorem gives a version of structural induction for α -terms. Just as Theorem 17 was derived from ordinary structural recursion (Theorem 3), we prove this theorem as a corollary of ordinary structural induction (Theorem 4).

Theorem 21 (first α -structural induction theorem). Let Σ be a nominal signature. Suppose we are given a finitely supported set $S \in P_{fs}(\mathsf{T}_{\alpha}(\Sigma))$ of α -terms over Σ . To prove that S is the whole of $\mathsf{T}_{\alpha}(\Sigma)$, it suffices to show

$$(\forall \mathsf{a} \in \Sigma_{\mathsf{A}}, a \in \mathbb{A}_{\mathsf{a}}) \ a \in S \tag{40}$$

$$(\forall (K : \sigma \to \mathsf{s}) \in \Sigma_C, e \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma}) \ e \in S \ \Rightarrow \ K \ e \in S \tag{41}$$

$$() \in S \tag{42}$$

 $(\forall (\sigma_i \in Ar(\Sigma), e_i \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma_i} \mid i = 1, 2))$

$$e_1 \in S \& e_2 \in S \implies (e_1, e_2) \in S$$
 (43)

 $(\forall a \in \Sigma_A, \sigma \in Ar(\Sigma))(\exists a \in \mathbb{A}_a) \ a \# S \ \&$

$$(\forall e \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma}) \ e \in S \ \Rightarrow \ a. \, e \in S \ . \tag{44}$$

Proof. Let \overline{S} be the set of nominal terms over Σ whose α -equivalence classes lie in S no matter how we permute the atoms occurring in the term:

$$\overline{S} \triangleq \{t \in \mathsf{T}(\Sigma) \mid (\forall \pi \in Perm) \ [\pi \cdot t]_{\alpha} \in S\} \ .$$

Clearly $S = \mathsf{T}_{\alpha}(\Sigma)$ if $\overline{S} = \mathsf{T}(\Sigma)$; and to prove the latter it suffices to check that \overline{S} satisfies conditions (15)–(19) of Theorem 4. The first four of these follow immediately from (40)–(43) respectively. So it just remains to show that (44) implies that \overline{S} satisfies condition (19). First note that by Theorem 14 applied to the set of atoms $\{a \in \mathbb{A}_{\mathsf{a}} \mid (\forall e \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma}) \ e \in S \ \Rightarrow \ a.e \in S\}$, which is supported by supp(S), (44) is equivalent to

$$(\forall \mathsf{a} \in \Sigma_{\mathsf{A}}, \sigma \in Ar(\Sigma))(\forall a \in \mathbb{A}_{\mathsf{a}}) \ a \# S \Rightarrow (\forall e \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma}) \ e \in S \Rightarrow a. \ e \in S. \tag{45}$$

Given $a \in \mathbb{A}_a$ and $t \in \overline{S}$, we have to prove that $(a) \in \overline{S}$, i.e. that $[\pi \cdot (a) t]_\alpha \in S$ for any $\pi \in Perm$. Choosing any atom a' in the infinite set $\mathbb{A}_a - supp(S, \pi, a, t)$, we have

$$\pi \cdot \ll a \gg t = \ll \pi(a) \gg (\pi \cdot t)$$

$$=_{\alpha} \ll a' \gg ((\pi \cdot t) \{a'/\pi(a)\}) \qquad \text{by definition of } =_{\alpha} \text{(Section 2.4)}$$

$$= \ll a' \gg ((\pi(a) \ a') \cdot (\pi \cdot t)) \qquad \text{since } a' \notin atm(\pi \cdot t)$$

$$= \ll a' \gg (\pi' \cdot t) \qquad \text{where } \pi' \triangleq (\pi(a) \ a') \circ \pi.$$

So $[\pi \cdot a \cdot a]_{\alpha} = a' \cdot [\pi' \cdot t]_{\alpha} \in S$ by (45), since a' # S (by choice of a') and $[\pi' \cdot t]_{\alpha} \in S$, because $t \in \overline{S}$. So it is indeed the case that $a \cdot b \in \overline{S}$ when $a \in \mathbb{A}_a$ and $a \cdot b \in \overline{S}$.

5 Second α -Structural Recursion & Induction Theorems

Theorem 17 is an "arity-directed" recursion principle for α -terms: one has to specify nominal sets X_{σ} for each arity σ , and give functions $f_{(_)}$ for atom-sorts, unit, pair and atom-binding arities in addition to ones for constructors. Although this gives flexibility over how to treat atom, unit, pair and atom-binding α -terms when giving an α -structurally recursive definition of some functions, this flexibility is often more of a hindrance than a help. In most cases one is primarily interested in defining functions only on α -terms whose arities are data-sorts $s \in \Sigma_D$, with α -terms of other kinds of arity (atom-sorts, unit, pair and atom-binding arities) playing an auxiliary role. For example, when Σ is the nominal signature for λ -terms with local recursive function declarations (Example 1), to define the capture-avoiding substitution function $\hat{s}_{x,e} \in T_{\alpha}(\Sigma)_t \to_{fs} T_{\alpha}(\Sigma)_t$ discussed in the Introduction, we should only have to specify finitely supported functions corresponding to the right-hand sides of the defining equations (6)–(9), i.e. one function for each of the signature's four constructors V, A, L and Letrec. But as it stands, to define

 $\hat{s}_{x,e}$ using Theorem 17 we have to work out suitable choices for X_{σ} and for the functions $f_{\mathsf{v}}, f_1, f_{\sigma_1 * \sigma_2}, f_{\mathsf{v} \mathsf{v} \mathsf{v} \sigma}$ for any $\sigma, \sigma_1, \sigma_2 \in Ar(\Sigma)$.

So we will develop a second, "sort-directed" version of α -structural recursion in which one only has to give X_{σ} when $\sigma=s$ is a data-sort, and only has to give the functions $f_{(_)}$ for constructors. Here is the statement of the new form of the recursion principle; the notations used in it are defined in Figure 2 and discussed below.

Theorem 22 (second α -structural recursion theorem). Let Σ be a nominal signature. Suppose we are given a family of nominal sets $X = (X_s \mid s \in \Sigma_D)$ indexed by the data-sorts of Σ , a finite set A of atoms, and functions

$$f_K \in X^{(\sigma)} \to_{\mathrm{fs}} X^{(\mathsf{s})} \quad ((K : \sigma \to \mathsf{s}) \in \Sigma_{\mathrm{C}})$$

all of which are supported by A and satisfy

$$(\exists \bar{a} \in \mathbb{A}^{\sigma}) \ \bar{a} \# A \& (\forall \bar{x} \in X^{|\sigma|}) \ \bar{a} \sharp_{\sigma} \bar{x} \Rightarrow \bar{a} \# f_{K}(\bar{a}, \bar{x})_{\sigma}.$$
 (FCB_K)

Then there is a unique family of finitely supported functions $(\hat{f}_s \in T_\alpha(\Sigma)_s \to_{fs} X_s \mid s \in \Sigma_D)$ with $supp(\hat{f}_s) \subseteq A$ and satisfying

$$(\forall \bar{a} \in \mathbb{A}^{\sigma}) \ \bar{a} \# A \Rightarrow (\forall \bar{e} \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}) \ \bar{a} \sharp_{\sigma} \bar{e} \Rightarrow \hat{f}_{\mathsf{s}}(K \ \bar{a}. \ \bar{e}) = f_{K}(\bar{a}, \hat{f}^{|\sigma|} \ \bar{e})_{\sigma}$$
(46)

for each $(K : \sigma \to s) \in \Sigma_{\mathbf{C}}$.

In this theorem we start with a family of nominal sets $X = (X_s \mid s \in \Sigma_D)$ indexed by the data-sorts of the signature Σ and with a family of finitely supported functions $(f_K \mid K \in \Sigma_C)$ indexed by the constructors of Σ . The domain $X^{(\sigma)}$ of f_K is obtained from the arity σ of K by interpreting each atom-sort as the corresponding nominal set of atoms (Example 7(i)), each data-sort as given by X, the unit arity as the one-element discrete nominal set (Example 7(iv)), pair arities using products of nominal sets (Section 3.2), and atom-binding arities just using product with nominal sets of atoms. The aim is to use this data to specify some functions \hat{f}_s mapping α -terms e: s to elements \hat{f}_s $e \in X_s$ by giving recursion equations as in (46), with one (conditional) equation for each way of forming α terms of data-sort, i.e. for each constructor $K: \sigma \to s$ in Σ . The conditional equation for K specifies the effect of \hat{f}_s not for arbitrary α -terms constructed with K, but rather just for those of the form $K \bar{a}.\bar{e}$ where $\bar{a} \# A$ and $\bar{a} \sharp_{\sigma} \bar{e}$ hold. Here $\bar{a} \in \mathbb{A}^{\sigma}$ is a nested tuple of distinct atoms matching the binding occurrences of atom-sorts in the arity σ ; $\bar{e} \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}$ is a nested tuple of α -terms matching the non-binding occurrences of atom-sorts and the occurrences of data-sorts in σ ; and the operation $\bar{a}, \bar{e} \mapsto \bar{a}, \bar{e}$ assembles these two nested tuples into an α -term of arity σ , to which K can be applied to get a constructed α -term $K \bar{a}.\bar{e}:$ s. The equation in (46) is restricted by two conditions:

| σ | $X^{(\sigma)}$ | \mathbb{A}^{σ} | $X^{ \sigma }$ |
|---------------------------|--|---|---------------------------------------|
| $a \in \Sigma_A$ | \mathbb{A}_{a} | 1 | \mathbb{A}_{a} |
| $\textbf{s} \in \Sigma_D$ | X_{s} | 1 | X_{s} |
| 1 | 1 | 1 | 1 |
| $\sigma_1 * \sigma_2$ | $X^{(\sigma_1)} \times X^{(\sigma_2)}$ | $\mathbb{A}^{\sigma_1}\otimes\mathbb{A}^{\sigma_2}$ | $X^{ \sigma_1 } 	imes X^{ \sigma_2 }$ |
| «a» σ_1 | $\mathbb{A}_{a} 	imes X^{(\sigma_1)}$ | $\mathbb{A}_{a}\otimes\mathbb{A}^{\sigma_1}$ | $X^{ \sigma_1 }$ |

The third column of the above table uses the **separated product** of nominal sets: $X \otimes Y \triangleq \{(x,y) \in X \times Y \mid x \# y\}$ (with atom-permutation action the same as for the product $X \times Y$).

| σ | $(\bar{a},\bar{x})\in\mathbb{A}^{\sigma}\times X^{ \sigma }$ | \bar{a} | $\sharp_{\sigma} \bar{x} \in \mathbb{B}$ | | $(\bar{a}, \bar{x})_{\sigma} \in X^{(\sigma)}$ |
|-----------------------|---|---|---|--------------|--|
| $a \in \Sigma_A$ | ((),a) | true | | | a |
| $s\in\Sigma_D$ | ((),x) | true | | | x |
| 1 | ((),()) | true | | | () |
| $\sigma_1 * \sigma_2$ | $((\bar{a}_1, \bar{a}_2), (\bar{x}_1, \bar{x}_2))$ | $ \begin{array}{cccccccccccccccccccccccccccccccccccc$ | | ((ā | $(\bar{a}_1,\bar{x}_1)_{\sigma_1},(\bar{a}_2,\bar{x}_2)_{\sigma_2})$ |
| «a» σ_1 | $((a,\bar{a}_1),\bar{x}_1)$ | $\bar{a}_1 \sharp_{\sigma} \bar{x}_1$ | | | $(a,(\bar{a}_1,\bar{x}_1)_{\sigma})$ |
| σ | $(\bar{a},\bar{e}) \in \mathbb{A}^{\sigma} \times T_{\alpha}(\Sigma)$ | $\Sigma)^{ \sigma }$ | $\bar{a}.\bar{e} \in T_{\!\alpha}(\Sigma$ | $)_{\sigma}$ | $\hat{f}^{ \sigma }\bar{e}\in X^{ \sigma }$ |
| $a \in \Sigma_A$ | ((),a) | | | | |
| | • (()) | | a | | a |
| $s\in\Sigma_{\Gamma}$ | | | $egin{array}{c} a \ e \end{array}$ | | $\hat{f_{s}}e$ |
| $s\in\Sigma_{D}$ | | | | | |
| | ((),e) $((),())$ |)) | e | $ar{z}_2)$ | $\hat{f}_{s} e$ |

In the second column of the above table, $\mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}$ indicates the construction $X \mapsto X^{|\sigma|}$ from the first table applied to the data-sort indexed family $X = (\mathsf{T}_{\alpha}(\Sigma)_{\mathsf{s}} \mid \mathsf{s} \in \Sigma_{\mathsf{D}})$.

Figure 2: Definitions used in Theorem 22

- the first, $\bar{a} \# A$, just requires that the atoms in \bar{a} do not occur in the common finite support of the functions f_K ;
- the second, $\bar{a} \not\parallel_{\sigma} \bar{e}$, ensures that in addition to the atoms in \bar{a} being mutually distinct (by virtue of the definition of \mathbb{A}^{σ}), they avoid the support of the constituents of \bar{e} in an appropriate way (the precise definition of "appropriate" being given in Figure 2).

The theorem guarantees the unique existence of such functions on α -terms provided the functions f_K satisfy the *freshness condition on binders* given by (FCB_K). This asserts the existence of a nested tuple \bar{a} of distinct atoms that can appear in binding positions in elements of $X^{(\sigma)}$ (i.e. $\bar{a} \in \mathbb{A}^{\sigma}$) such that:

- the atoms are distinct from A, i.e. $\bar{a} \# A$ (they are also mutually distinct by definition of \mathbb{A}^{σ});
- whenever \bar{x} is a nested tuple of atoms and X-elements that can appear in non-binding positions in elements of $X^{(\sigma)}$ (i.e. $\bar{x} \in X^{|\sigma|}$) for which \bar{a} is suitably fresh, i.e. satisfying $\bar{a} \sharp_{\sigma} \bar{x}$, then assembling \bar{a} and \bar{x} into an element $(\bar{a},\bar{x})_{\sigma} \in X^{(\sigma)}$, f_K maps this element to one in $X^{(s)} = X_s$ whose support does not contain any of the atoms in \bar{a} (i.e. $\bar{a} \# f_K(\bar{a},\bar{x})_{\sigma}$).

The easiest way I know of proving Theorem 22 is to derive it from the following "sort-directed" version of α -structural induction, which uses notations that are defined in Figures 2 and 3, and which are discussed below.

Theorem 23 (second α -structural induction theorem). Let Σ be a nominal signature. Suppose we are given a family of finitely supported subsets $(S_s \in P_{fs}(T_{\alpha}(\Sigma)_s) \mid s \in \Sigma_D)$ indexed by the data-sorts of Σ and all supported by a finite set of atoms A. Then to prove that S_s is the whole of $T_{\alpha}(\Sigma)_s$ for all $s \in \Sigma_D$, it suffices to show for each constructor $(K : \sigma \to s) \in \Sigma_C$ that

$$(\exists \bar{a} \in \mathbb{A}^{\sigma}) \ \bar{a} \ \# \ A \ \&$$

$$(\forall \bar{e} \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}) \ \bar{a} \ \sharp_{\sigma} \ \bar{e} \ \& \ \bar{e} \in S^{|\sigma|} \ \Rightarrow \ K \ \bar{a}. \ \bar{e} \in S_{\mathsf{s}} \ . \ (\mathsf{IH}_K)$$

In this theorem we start with a family of subsets $S_s \subseteq T_\alpha(\Sigma)_s$ of α -terms whose arities are data-sorts and that are all supported by some finite set of atoms A. We wish to prove that every t:s is in S_s (for all $s\in \Sigma_D$). The theorem guarantees this provided each constructor $(K:\sigma\to s)\in \Sigma_C$ satisfies the induction hypothesis given by (IH_K) . This asserts the existence of a nested tuple \bar{a} of distinct atoms that can appear in binding positions in α -terms of arity σ (i.e. $\bar{a}\in \mathbb{A}^\sigma$) such that:

– the atoms are distinct from A, i.e. $\bar{a} \# A$ (they are also mutually distinct by definition of \mathbb{A}^{σ});

⁹This relation \sharp_{σ} , defined in Figure 2, is a subtlety of the freshness condition on binders that is not apparent in the simpler First α -Structural Recursion Theorem.

| σ | $S^{ \sigma } \in P_{\mathrm{fs}}(T_{\!lpha}(\Sigma)^{ \sigma })$ | $S_{\sigma} \in P_{\mathrm{fs}}(T_{\!lpha}(\Sigma)_{\sigma})$ |
|---------------------------|---|---|
| $a \in \Sigma_\mathrm{A}$ | \mathbb{A}_{a} | $\{a \mid a \in \mathbb{A}_{a}\}$ |
| $s \in \Sigma_\mathrm{D}$ | S_{s} | S_{s} |
| 1 | 1 | $\{()\}$ |
| $\sigma_1 * \sigma_2$ | $S^{ \sigma_1 } \times S^{ \sigma_2 }$ | $\{(e_1, e_2) \mid e_1 \in S_{\sigma_1} \& e_2 \in S_{\sigma_2}\}$ |
| «a» σ_1 | $S^{ \sigma_1 }$ | $\{a. e_1 \mid a \in \mathbb{A}_{a} - supp(S_{\sigma_1}) \& e_1 \in S_{\sigma_1}\}$ |

Figure 3: Definitions used in Theorem 23

- whenever \bar{e} is a nested tuple of α -terms that can appear in non-binding positions in an α -term of arity σ (i.e. $\bar{e} \in T_{\alpha}(\Sigma)^{|\sigma|}$) for which \bar{a} is suitably fresh, i.e. satisfying $\bar{a} \sharp_{\sigma} \bar{e}$, then assembling \bar{a} and \bar{e} into the α -term $\bar{a}.\bar{e}:\sigma$, the constructed α -term $K \bar{a}.\bar{e}$ must lie in the subset S_s .

The proof of Theorem 23 is given in Appendix A and the proof of Theorem 22 in Appendix B. In Sections 5.1 and 5.2 we explore what these principles look like for particular nominal signatures, using the examples from Section 2.2.

Remark 24 (atom-abstraction and the initial algebra property). For each atomsort $a \in \mathbb{AS}$ and each nominal set X, let $[\mathbb{A}_a]X$ denote the set of equivalence classes of pairs $(a,x) \in \mathbb{A}_a \times X$ for the equivalence relation $(a,x) \sim (a',x')$ given by

$$(\exists a'' \in \mathbb{A}_{\mathsf{a}}) \ a'' \# (a, x, a', x') \ \& \ (a \ a'') \cdot x = (a' \ a'') \cdot x' \ . \tag{47}$$

The relation \sim is evidently reflexive and symmetric; to see that it is also transitive, one first applies Theorem 14 with $S = \{a'' \in \mathbb{A}_{\mathsf{a}} \mid (a\ a'') \cdot x = (a'\ a'') \cdot x'\}$ and A = supp(a, x, a', x') to show that $(a, x) \sim (a', x')$ holds if and only if

$$(\forall a'' \in \mathbb{A}_{a}) \ a'' \# (a, x, a, x') \ \Rightarrow \ (a \ a'') \cdot x = (a' \ a'') \cdot x' \ . \tag{48}$$

We write [a]x for the \sim -equivalence class of the pair (a,x) and call it an **atomabstraction**. It follows from Lemma 13 that \sim is a nominal subset of $\mathbb{A}_a \times X$. So the quotient $[\mathbb{A}_a]X$ is a nominal set as in Section 3.3. One can calculate that the support of each element [a]x of $[\mathbb{A}_a]X$ is $supp(x) - \{a\}$.

Using these atom-abstraction nominal sets, it is possible to give an *initial algebra* characterisation of $(T_{\alpha}(\Sigma)_s \mid s \in \Sigma_D)$ that is equivalent to Theorem 22. Consider the category whose objects are families of nominal sets $X = (X_s \mid s \in \Sigma_D)$ indexed by the data-sorts of Σ , and whose morphisms $f: X \to X'$ are indexed families $f = (f_s \in X_s \to_{fs} X'_s \mid s \in \Sigma_D)$ of functions with empty support (i.e. functions that respect the action of all atom-permutations). The constructors of Σ determine a functor F_{Σ} from this category to itself, defined in Figure 4. An F_{Σ} -algebra is simply an object I equipped with a morphism $i: F_{\Sigma}I \to I$. Such

Action of F_{Σ} on objects:

$$(F_{\Sigma}X)_{\mathsf{s}} \triangleq \sum_{(K:\sigma \to \mathsf{s}) \in \Sigma_{\mathbf{C}}} X^{[\sigma]} = \{(K,x) \mid (K:\sigma \to \mathsf{s}) \in \Sigma_{\mathbf{C}} \& x \in X^{[\sigma]}\} \ .$$

Action of F_{Σ} on morphisms:

Figure 4: Functor F_{Σ} associated with a nominal signature Σ

an algebra is **initial** if for any other such algebra $f: F_{\Sigma}X \to X$, there is a unique morphism $\hat{f}: I \to X$ so that

$$F_{\Sigma} I \xrightarrow{i} I$$

$$F_{\Sigma} \hat{f} \downarrow \qquad \qquad \downarrow \hat{f}$$

$$F_{\Sigma} X \xrightarrow{f} X$$

$$(49)$$

commutes, i.e. satisfying

$$(\forall \mathsf{s} \in \Sigma_{\mathrm{D}})(\forall (K: \sigma \to \mathsf{s}) \in \Sigma_{\mathrm{C}})(\forall x \in I^{[\sigma]}) \ \hat{f}_{\mathsf{s}}(i_{\mathsf{s}}(K, x)) = f_{\mathsf{s}}(K, \hat{f}^{[\sigma]} x) \ .$$

Standard category-theoretic results give that i is an isomorphism and that the initial algebra (I,i) is unique up to isomorphism. Theorem 22 can be used to prove that $(\mathsf{T}_{\alpha}(\Sigma)_{\mathsf{s}} \mid \mathsf{s} \in \Sigma_{\mathsf{D}})$ is the (object part of) an initial F_{Σ} -algebra. Conversely, one can give a direct inductive construction of an initial F_{Σ} -algebra and use the initial algebra property (49) to deduce Theorem 22: see [13, Section 6].

5.1 Example: λ -calculus (with letrec)

Let Σ be the nominal signature from Example 1. Thus $\mathsf{T}(\Sigma)_\mathsf{t}$ is the set Λ of abstract syntax trees for λ -calculus with *letrec*; and $\mathsf{T}_{\alpha}(\Sigma)_\mathsf{t}$ is the quotient $\Lambda/=_{\alpha}$ of

that set by the usual notion of α -equivalence—in other words $T_{\alpha}(\Sigma)_t$ is what is normally meant by the set of all (open or closed) untyped λ -terms with local recursive function declarations. Suppose we are given a nominal set X and functions

$$f_V \in \mathbb{A}_{\mathsf{v}} \to_{\mathsf{fs}} X$$
 (50)

$$f_A \in X \times X \to_{\mathsf{fs}} X \tag{51}$$

$$f_L \in \mathbb{A}_{\mathsf{v}} \times X \to_{\mathrm{fs}} X \tag{52}$$

$$f_{Letrec} \in \mathbb{A}_{\mathsf{v}} \times ((\mathbb{A}_{\mathsf{v}} \times X) \times X) \to_{\mathrm{fs}} X$$
 (53)

all supported by a finite set of atoms A. Applying the definitions in Figure 2, one finds that the conditions (FCB_K) for K = V, A are equivalent to true, that (FCB_{Lam}) is equivalent to

$$(\exists a \in \mathbb{A}_{\mathsf{V}}) \ a \notin A \ \& \ (\forall x \in X) \ a \# f_L(a, x) \tag{54}$$

and that (FCB_{Letrec}) is equivalent to

$$(\exists a, a' \in \mathbb{A}_{\mathsf{v}}) \ a \neq a' \ \& \ a, a' \notin A \ \& \ (\forall x, x' \in X) \ a' \# x \Rightarrow (a, a') \# f_{Letrec}(a, ((a', x'), x)) \ .$$
 (55)

So for this nominal signature, Theorem 22 gives us the following recursion principle. We state it using the usual concrete syntax for λ -calculus and using the fact (noted in Example 7(iii)) that the support of a term $e \in \Lambda/=_{\alpha}$ is its finite set fv(e) of free variables.

Theorem 25. Let $\Lambda/=_{\alpha}$ be the nominal set of α -equivalence classes of λ -terms with local recursive function declarations:

$$e ::= x \mid e \mid \lambda x.e \mid let rec \mid x \mid x = e \mid in \mid e$$

where the variables x are drawn from the nominal set \mathbb{A}_{v} of atoms of some fixed sort v . Given any nominal set X and functions as in (50)–(53) all supported by some finite set of atoms A and with f_L and f_{Letrec} satisfying (54) and (55), then there is a unique function $\hat{f} \in (\Lambda/=_{\alpha} \to_{\mathsf{fs}} X)$ supported by A and satisfying

$$\hat{f}(x) = f_V(x) \tag{56}$$

$$\hat{f}(e_1 e_2) = f_A(\hat{f} e_1, \hat{f} e_2)$$
(57)

$$x \notin A \Rightarrow \hat{f}(\lambda x.e) = f_L(x, \hat{f}e)$$
 (58)

$$x, y \notin A \& x \notin fv(e_2) \cup \{y\} \Rightarrow$$

$$\hat{f}(letrec\ y\ x = e_1\ in\ e_2) = f_{Letrec}(y, ((x, \hat{f}\ e_1), \hat{f}\ e_2))\ .$$
 (59)

Turning to the induction principle for this signature, if $S \subseteq \Lambda/=_{\alpha}$ is a set of terms supported by a finite set of atoms A, then (IH_V), (IH_A), (IH_L) and (IH_{Letrec}) are equivalent to

$$(\forall x \in \mathbb{A}_{\mathsf{v}}) \ x \in S \tag{60}$$

$$(\forall e_1, e_2 \in S) \ e_1 \ e_2 \in S \tag{61}$$

$$(\exists x \in \mathbb{A}_{\mathsf{v}}) \ x \notin A \ \& \ (\forall e \in S) \ \lambda x. \ e \in S \tag{62}$$

$$(\exists x, y \in \mathbb{A}_{\mathsf{v}}) \ x \neq y \ \& \ x, y \notin A \ \&$$
$$(\forall e_1, e_2 \in S) \ x \notin fv(e_2) \ \Rightarrow \ letter \ y \ x = e_1 \ in \ e_2 \in S$$
 (63)

respectively. So for this signature Theorem 23 says that S contains all terms if it satisfies (60)–(63).

Example 26 (capture-avoiding substitution). The example mentioned in the Introduction of capture-avoiding substitution of λ -terms, $\hat{s}_{x,e} \in \Lambda/=_{\alpha} \to \Lambda/=_{\alpha}$, is obtained from the above theorem by taking X to be the nominal set $\Lambda/=_{\alpha}$. Given $x \in \mathbb{A}_{\mathsf{v}}$ and $e \in X$, then $\hat{s}_{x,e}$ is given by \hat{f} where

$$f_V(y) \triangleq \begin{cases} e & \text{if } y = x \\ y & \text{if } y \neq x \end{cases}$$

$$f_A(e_1, e_2) \triangleq e_1 e_2$$

$$f_L(y, e_1) \triangleq \lambda y. e_1$$

$$f_{Letrec}(z, ((y, e_1), e_2)) \triangleq letrec \ z \ y = e_1 \ in \ e_2$$

$$A \triangleq fv(e) \cup \{x\} \ .$$

Condition (54) is satisfied because, as noted in Example 7(iii), for each $e_1 \in X = \Lambda/=_{\alpha}$, $supp(e_1)$ is the finite set $fv(e_1)$ of free variables of e_1 ; in particular we have $y \# f_{Lam}(y,e_1) = \lambda y$. e_1 simply because $y \notin fv(\lambda y.e_1) = fv(e_1) - \{y\}$. Similarly, condition (55) is satisfied because

$$fv(letrec\ z\ y = e_1\ in\ e_2) = (fv(e_1) - \{y, z\}) \cup (fv(e_2) - \{z\})$$

so that $(y,z) \# f_{Letrec}(z,((y,e_1),e_2)) = letrec zy = e_1 in e_2$ provided $y \notin fv(e_2)$. Note that f_A , f_L and f_{Letrec} are all supported by the empty set of atoms, whereas f_V is supported by $fv(e) \cup \{x\}$; so this is what we take for the common support A. Thus the conditions on the recursion equations in (58) and (59) correspond precisely to the conditions in (8) and (9).

Example 27 (recursion with varying parameters). Norrish [18, p 245] considers a variant of capture-avoiding substitution whose definition involves recursion with varying parameters; it motivates the parametrised recursion principle he presents in that paper. The α -structural recursion principles we have given here do not involve extra parameters, let alone varying ones; nevertheless it is possible to derive

parameterised versions from them. In the case of ordinary structural recursion, one can derive a parameterised version from an unparameterised one by currying parameters and defining maps into function sets using Theorem 3. In the presence of binders, one has to do something slightly more complicated, involving the Freshness Theorem 16, to derive a parameterised (FCB) from the unparameterised version of the condition.

Let us see how this works for Norrish's example, using the nominal signature for the pure λ -calculus obtained from Example 1 by deleting the Letrec constructor. Fixing on a pair of atoms $x_1, x_2 \in \mathbb{A}_{\mathsf{v}}$, we seek a function $s \in (\Lambda/=_{\alpha}) \to_{\mathrm{fs}} (\Lambda/=_{\alpha}) \to_{\mathrm{fs}} (\Lambda/=_{\alpha})$ satisfying for all y, e, e_1, e_2 :

$$s(y)(e) = \begin{cases} e & \text{if } y = x_1 \\ y & \text{if } y \neq x_1 \end{cases}$$
 (64)

$$s(e_1 e_2)(e) = (s(e_1)(e))(s(e_2)(e))$$
 (65)

$$y \notin fv(e) \cup \{x_1, x_2\} \Rightarrow s(\lambda y. e_1)(e) = \lambda y. s(e_1)(x_2 e)$$
 (66)

Thus the same parameter e appears in each clause defining $s(e_1)(e)$ by recursion on the structure of e_1 except for clause (66), where the application term $x_2 e$ appears instead. Such a function s can be obtained from Theorem 25 (restricted to pure λ -terms) as $s=\hat{f}$ if we take X to be the nominal set $(\Lambda/=_{\alpha}) \to_{\mathrm{fs}} (\Lambda/=_{\alpha})$ and use the functions

$$f_{V} \triangleq \lambda y \in \mathbb{A}_{\mathsf{v}}.\lambda e \in (\Lambda/=_{\alpha}). \ if \ y = x_{1} \ then \ e \ else \ y$$

$$f_{A} \triangleq \lambda(\xi_{1}, \xi_{2}) \in X \times X. \ \lambda e \in (\Lambda/=_{\alpha}). \ (\xi_{1} \ e) \ (\xi_{2} \ e)$$

$$f_{L} \triangleq \lambda(y, \xi_{1}) \in \mathbb{A}_{\mathsf{v}} \times X.\lambda e \in (\Lambda/=_{\alpha}). \ fresh(h(y, \xi_{1}, e))$$

where the last clause uses Theorem 16 applied to the finitely supported function $h(y, \xi_1, e) \in \mathbb{A}_{\mathsf{V}} \to_{\mathrm{fs}} (\Lambda/=_{\alpha})$ that maps each $y' \in \mathbb{A}_{\mathsf{V}}$ to

$$h(y, \xi_1, e)(y') \triangleq \lambda y' \cdot ((y \ y') \cdot \xi_1)(x_2 \ e)$$

This function is easily seen to satisfy the property (30) needed to apply the theorem. All the above functions are supported by $A \triangleq \{x_1, x_2\}$. Properties (56) and (57) of \hat{f} give (64) and (65) respectively. When $y \neq x_1, x_2$, property (58) gives us $\hat{f}(\lambda y. e_1) = f_L(y, \hat{f} e_1) = fresh(h(y, \hat{f} e_1, e))$. So if $y \# (x_1, x_2, e)$, picking any $y' \# (x_1, x_2, e, e_1, h)$, then by Theorem 16 we have $fresh(h(y, \hat{f} e_1, e)) = h(y, \hat{f} e_1, e)(y') \triangleq \lambda y'. ((y y') \cdot (\hat{f} e_1))(x_2 e) = \lambda y'. (y y') \cdot (\hat{f} e_1 (x_2 e))$. Hence by definition of $=_{\alpha}$, $\hat{f}(\lambda y. e_1) = \lambda y. \hat{f} e_1 (x_2 e)$, as required for (66).

5.2 Example: π -calculus

Let Σ be the nominal signature from Example 2. Suppose we are given nominal sets X_{proc} , X_{gsum} , X_{pre} and functions

all supported by a finite set of atoms A. The conditions (FCB_{Res}) and (FCB_{In}) are equivalent to

$$(\exists a \in \mathbb{A}_{\mathsf{chan}}) \ a \notin A \ \& \ (\forall x \in X_{\mathsf{proc}}) \ a \# f_{Res}(a, x)$$

$$(\exists a \in \mathbb{A}_{\mathsf{chan}}) \ a \notin A \ \&$$

$$(\forall a' \in \mathbb{A}_{\mathsf{chan}}) (\forall x \in X_{\mathsf{proc}}) \ a \neq a' \ \Rightarrow \ a \# f_{In}(a', (a, x))$$

$$(68)$$

respectively; and conditions (FCB_K) for $K \neq Res$, In are all equivalent to true. So if f_{Res} and f_{In} satisfy (67) and (68), then by Theorem 22, there are unique finitely supported functions $\hat{f}_s \in T_\alpha(\Sigma)_s \to_{fs} X_s$ (for s = proc, gsum, pre) all supported by A and satisfying for all $e, e_1, e_2, a_1, a_2, a, a'$ of suitable arity

$$\hat{f}(Ke) = f_K(\hat{f}e)$$
 $(K = Gsum, Rep, Pre)$ (69)

$$\hat{f}(K(e_1, e_2)) = f_K(\hat{f} e_1, \hat{f} e_2)$$
 (K = Par, Plus) (70)

$$\hat{f}(K()) = f_K() \qquad (K = Zero, Tau) \qquad (71)$$

$$\hat{f}(K((a_1, a_2), e)) = f_K((a_1, a_2), \hat{f}e)$$
 $(K = Out, Match)$ (72)

$$a \notin A \Rightarrow \hat{f}(Res \ a. \ e) = f_{Res}(a, \hat{f} \ e)$$
 (73)

$$a \notin A \& a \neq a' \Rightarrow \hat{f}(In(a', a. e)) = f_{In}(a', (a, \hat{f} e))$$
 (74)

where we have abbreviated $\hat{f}_{s}(e)$ to \hat{f} e.

We leave the reader to work out what induction principle Theorem 23 gives for this signature.

6 Extended Example: Normalisation by Evaluation

One of the most important aspects of nominal sets is that they provide, via the notion of finite support, a notion of *freshness* of names with respect to mathematical structures. This notion generalises the usual "not a free variable of" relation from finite syntactical structures to infinite objects (sets, functions, ...) where there is no obvious notion of free name. The theory comes into its own in situations where syntax and semantics have to be considered together and yet one still needs a workable notion of fresh name. We give an example in this section by treating **normalisation by evaluation** (**NBE**) for simply typed λ -calculus [3]. This produces $\beta\eta$ -long normal forms for typed λ -terms by first taking their denotational semantics in the standard, extensional functions model of the calculus over a ground type of syntax trees and then composing with a *reification* function that turns elements of the denotational model back into syntax (in normal form). When reifying an extensional function into a λ -abstraction one wants to choose a fresh name v for the λ -bound variable; but as the survey [8, p 157] eloquently puts it when discussing an informal version of the reification function

"The problem is the "v fresh" condition; what exactly does it mean? Unlike such conditions as "x does not occur free in E", it is not even locally checkable whether a variable is fresh; freshness is a global property, defined with respect to a term that may not even be fully constructed yet."

As a result, treatments of NBE in the literature adopt some device for making the current finite context of used names explicit and threading it through all the mathematical definitions involved in the denotational and reification functions used for NBE: see [8, section 3.3] and [9], for example. This tends to obscure the simple, but informal idea behind reification. The authors of [8] go on to mention after the above quote that "freshness" can be characterised rigorously in the framework of [13]. The details are presented here for the first time, using nominal sets rather than the FM-set theory of *loc. cit.* The point is not just that this setting provides a rigorous explanation of "freshness" (since the formal approaches mentioned above also do that), but that it allows us to retain the essential simplicity of an informal account such as in [8, section 3.2].

6.1 Typed λ -terms and their $\beta\eta$ -long normal forms

Rather than give a signature for raw λ -terms and then cut down to the well-typed ones using typing contexts, we make do with a simpler, but less extensible ¹⁰ treatment using explicitly typed variables. Let

$$Ty \triangleq \{\tau ::= \iota \mid \tau \to \tau\} \tag{75}$$

¹⁰Such an approach is fine for simply-typed terms, but becomes unworkable for calculi with type variables or dependent types.

be the set of **simple type symbols** over a single **ground type** ι . We assume given an injective function $\tau \in Ty \mapsto \mathsf{v}_\tau \in \mathbb{AS}$ that codes the simple type symbols as atom-sorts. We use atoms of sort v_τ to stand for variables of type τ in the simply typed λ -calculus. Note that when τ and τ' are different simple type symbols, v_τ and $\mathsf{v}_{\tau'}$ are different atom-sorts; so recalling the assumptions we made in section 2.1, the sets of atoms $\mathbb{A}_{\mathsf{v}_\tau}$ and $\mathbb{A}_{\mathsf{v}_{\mathsf{v}'}}$ are disjoint.

Consider the nominal signature Σ^{STL} with

| atom-sorts | data-sorts | constructors |
|------------|------------|--|
| $v_{	au}$ | $t_{	au}$ | $Vr_{	au}:v_{	au}	ot_{	au}$ |
| | • | $Ap_{	au,	au'}: t_{	au ightarrow 	au'} * t_{	au} ightarrow t_{	au'}$ |
| | | $ \begin{array}{c} Ap_{\tau,\tau'}: t_{\tau \dot{\to} \tau'} * t_{\tau} \to t_{\tau'} \\ Lm_{\tau,\tau'}: «v_{\tau} » t_{\tau'} \to t_{\tau \dot{\to} \tau'} \end{array} $ |

as τ and τ' range over Ty. Thus the (nominal) set

$$\Lambda(\tau) \triangleq \mathsf{T}_{\alpha}(\Sigma^{\mathsf{STL}})_{\mathsf{t}_{\tau}} \tag{76}$$

of α -terms of arity t_{τ} over Σ^{STL} is precisely the usual set of α -equivalence classes of abstract syntax trees for λ -terms of simple type $\tau \in Ty$, using variables that are explicitly tagged with types.

Next we give a nominal signature Σ^{LNF} for $\beta\eta$ -long normal forms:

| atom-sorts | data-sorts | constructors |
|-------------------|------------|---|
| $V_{\mathcal{T}}$ | $n_{	au}$ | $V_	au:v_	au	ou_	au$ |
| | $u_{	au}$ | $ \begin{vmatrix} A_{\tau,\tau'} : u_{\tau \to \tau'} * n_{\tau} \to u_{\tau'} \\ L_{\tau,\tau'} : «v_{\tau} » n_{\tau'} \to n_{\tau \to \tau'} \end{vmatrix} $ |
| | | $\mid L_{	au,	au'}: «v_{	au} » n_{	au'} ightarrow n_{	au ightarrow 	au'}$ |
| | | $I:u_\iota	on_\iota$ |

where τ and τ' range over Ty. The (nominal) set

$$N(\tau) \triangleq \mathsf{T}_{\alpha}(\Sigma^{\mathsf{LNF}})_{\mathsf{n}_{\tau}} \tag{77}$$

corresponds to the set of α -equivalence classes of abstract syntax trees for simply-typed λ -terms of type τ in $\beta\eta$ -long normal form, whereas

$$U(\tau) \triangleq \mathsf{T}_{\alpha}(\Sigma^{\mathsf{LNF}})_{\mathsf{u}_{\tau}} \tag{78}$$

corresponds to the set of α -equivalence classes of **neutral** (or *atomic*) terms of type τ ; see [8, p 155], for example.

Notation 28. From now on we will use the following concrete, overloaded, but hopefully more familiar notations for α -terms over the signatures Σ^{STL} and Σ^{LNF} .

- Typical elements of $\mathbb{A}_{v_{\tau}}$, $\Lambda(\tau)$, $N(\tau)$ and $U(\tau)$ will be written x, e, n and u respectively.
- $Vr_{\tau} x$ will be written just as x, $Ap_{\tau,\tau'}(e_1,e_2)$ as $e_1 e_2$, and $Lm_{\tau,\tau'} x.e$ as $\lambda x: \tau.e$.

• $V_{\tau} x$ will be written just as x, $A_{\tau,\tau'}(u,n)$ as u n, $L_{\tau,\tau'} x$. n as $\lambda x : \tau$. n and I u just as u.

Every $\beta\eta$ -long normal form and every neutral term can be regarded as a λ -term of the corresponding type; indeed a very simple application of the second α -structural recursion theorem for the nominal signature $\Sigma^{\rm LNF}$ tells us that there are functions

$$i_{\tau} \in N(\tau) \to_{\mathrm{fs}} \Lambda(\tau)$$
 $j_{\tau} \in U(\tau) \to_{\mathrm{fs}} \Lambda(\tau)$ (79)

supported by the empty set of atoms and satisfying for all $\tau, \tau' \in Ty$ and x, u, n of suitable arity

$$j_{\tau} x = x$$

$$j_{\tau'}(u n) = (j_{\tau \to \tau'} u)(i_{\tau} n)$$

$$i_{\tau \to \tau'}(\lambda x : \tau \cdot n) = \lambda x : \tau \cdot i_{\tau'} n$$

$$i_{t} u = j_{t} u .$$

We aim to use the technique of NBE [3] to define the normalisation function

$$norm_{\tau} \in \Lambda(\tau) \to_{\mathsf{fS}} N(\tau)$$
 (80)

with the following properties:

$$(\forall \tau \in Ty, e_1, e_2 \in \Lambda(\tau)) \ e_1 =_{\beta\eta} e_2 \Rightarrow norm_{\tau} e_1 = norm_{\tau} e_2$$
 (81)

$$(\forall \tau \in Ty, n \in N(\tau)) \ norm_{\tau}(i_{\tau} n) = n \tag{82}$$

$$(\forall \tau \in Ty, e \in \Lambda(\tau)) \ i_{\tau}(norm_{\tau} \ e) =_{\beta n} e \ . \tag{83}$$

Here $=_{\beta\eta} \subseteq \Lambda(\tau) \times \Lambda(\tau)$ is the usual relation of $\beta\eta$ -conversion between (α -equivalence classes of) λ -terms of the same simple type τ . It is by definition the smallest congruence relation satisfying

$$(\forall \tau, \tau' \in Ty, x \in \mathbb{A}_{\mathsf{V}_{\tau}}, e_1 \in \Lambda(\tau'), e_2 \in \Lambda(\tau))$$
$$(\lambda x : \tau. e_1)e_2 =_{\beta\eta} e_1[x := e_2]$$
(84)

$$(\forall \tau, \tau' \in Ty, e \in \Lambda(\tau \to \tau'), x \in \mathbb{A}_{\mathsf{v}_{\tau}}) \ x \# e \Rightarrow e =_{\beta\eta} \lambda x : \tau . e x . \tag{85}$$

Here $e_1[x := e_2]$ indicates the **capture-avoiding substitution** of e_2 for all free occurrences of x in e_1 . It can be defined using the second α -structural recursion theorem for the nominal signature Σ^{STL} much as in Example 26; instead we will regard it as a special case of the simultaneous substitution functions defined in the next section.

Once (81)–(83) are proved, then it follows that for every $e \in \Lambda(\tau)$ there is a unique $n \in N(\tau)$ with $e =_{\beta\eta} i_{\tau} n$; and, modulo some considerations about computability, deciding $\beta\eta$ -conversion is reduced to the decidable relation of α -equivalence on $\mathsf{T}(\Sigma^{\mathrm{LNF}})_{\mathsf{n}_{\tau}}$ by applying the $norm_{\tau}$ function. We aim to show how to use α -structural recursion to define $norm_{\tau}$; and how to prove (81)–(83) using, among other things, α -structural induction.

 $\tau \in Ty, e \in \Lambda(\tau), \sigma \in Sub \mapsto [e]\sigma \in \Lambda(\tau)$:

$$[x]\sigma = \sigma x \tag{86}$$

$$[e_1 e_2]\sigma = ([e_1]\sigma)([e_2]\sigma)$$
 (87)

$$x \# \sigma \Rightarrow [\lambda x : \tau . e]\sigma = \lambda x : \tau . [e]\sigma$$
 (88)

 $\tau \in \mathit{Ty}, e \in \Lambda(\tau), \rho \in \mathit{Env} \mapsto [\![e]\!] \rho \in D(\tau) \text{:}$

$$[\![x]\!]\rho = \rho x \tag{89}$$

$$[e_1 e_2] \rho = [e_1] \rho ([e_2] \rho)$$
 (90)

$$[\![\lambda x : \tau.e]\!] \rho = \lambda d \in D(\tau). [\![e]\!] (\rho \{x \mapsto d\})$$

$$\tag{91}$$

Figure 5: Substitution and denotation

6.2 Substitution

A (simultaneous) **substitution** σ is a function that maps atoms in $\mathbb{A}_{v_{\tau}}$ to α -terms in $\Lambda(\tau)$ (for any $\tau \in Ty$) and that has the property that its **domain**

$$dom(\sigma) \triangleq \{x \in \bigcup_{\tau \in Ty} \mathbb{A}_{\mathsf{V}_{\tau}} \mid \sigma \, x \neq x\}$$

is a finite set. We let atom-permutations $\pi \in Perm$ act on such functions in the usual way (Section 3.2): the substitution $\pi \cdot \sigma$ maps $x \in \mathbb{A}_{\mathsf{v}_\tau}$ to $\pi \cdot (\rho(\pi^{-1}(x)))$ and (therefore) has domain $dom(\pi \cdot \sigma) = \pi \cdot dom(\sigma) = \{\pi(x) \mid x \in dom(\sigma)\}$. It is not hard to see that with respect to this action, each substitution σ is supported by the finite set of atoms $dom(\sigma) \cup \bigcup_{x \in dom(\sigma)} supp(\sigma x)$. Therefore the collection of substitutions forms a nominal set that we write as Sub.

Given $\tau \in Ty$, $e \in \Lambda(\tau)$ and $\sigma \in Sub$, we let $[e]\sigma \in \Lambda(\tau)$ denote the result of carrying out on e the simultaneous, **capture-avoiding substitution** given by σ . This is specified by the recursion equations (86)–(88) in Figure 5. For each $\sigma \in Sub$, we can use the second α -structural recursion theorem for $\Sigma^{\rm STL}$ to define the functions $([-]\sigma \in \Lambda(\tau) \to_{\rm fs} \Lambda(\tau) \mid \tau \in Ty)$ satisfying these equations, much as in Example 26, but using $supp(\sigma)$ as the common finite support A. (In particular, in Theorem 22 the only non-trivial freshness condition on binders, (FCB $_{Lm_{\tau,\tau'}}$), is easily verified.)

The **identity substitution** $\sigma_0 \in Sub$ maps each $x \in \mathbb{A}_{\mathsf{V}_{\tau}}$ to $x \in \Lambda(\tau)$. The **composition** $\sigma_1; \sigma_2$ of $\sigma_1, \sigma_2 \in Sub$ is the element of Sub that maps each $x \in \mathbb{A}_{\mathsf{V}_{\tau}}$ to $[\sigma_1 \, x] \sigma_2$. These operations satisfy:

$$[e]\sigma_0 = e \tag{92}$$

$$[e](\sigma_1; \sigma_2) = [[e]\sigma_1]\sigma_2. \tag{93}$$

Both properties can be proved easily by applying the second α -structural induction principle (Theorem 23) for the nominal signature $\Sigma^{\rm STL}$. For example, to prove (93), for each $\sigma_1,\sigma_2\in Sub$ we take S_{t_τ} to be $\{e\in\Lambda(\tau)\mid [e](\sigma_1;\sigma_2)=[[e]\sigma_1]\sigma_2\}$, which is supported by $A=supp(\sigma_1,\sigma_2)$; then (IH_{Vr_τ}) and $(\mathrm{IH}_{Ap_{\tau,\tau'}})$ are easy to verify, using (86) and (87) respectively; for $(\mathrm{IH}_{Lm_{\tau,\tau'}})$, which is the statement

$$(\exists x \in \mathbb{A}_{\mathsf{v}_{\tau}}) \ x \notin A \ \& \ (\forall e \in S_{\mathsf{t}_{\tau'}}) \ \lambda x : \tau. \ e \in S_{\mathsf{t}_{\tau \dot{\to} \tau'}} \ ,$$

we can choose any x in the infinite set $\mathbb{A}_{\mathsf{v}_{\tau}}-A$ (so that $x \# (\sigma_1,\sigma_2)$): if $e \in S_{\mathsf{t}_{\tau'}}$, then

$$\begin{split} &[\lambda x:\tau.e](\sigma_1;\sigma_2)\\ &=[[\lambda x:\tau.e]\sigma_1]\sigma_2 & \text{by definition of }\sigma_1;\sigma_2\\ &=[\lambda x:\tau.[e]\sigma_1]\sigma_2 & \text{by (88), since }x \;\#\;\sigma_1\\ &=\lambda x:\tau.[[e]\sigma_1]\sigma_2 & \text{by (88), since }x\;\#\;\sigma_2\\ &=\lambda x:\tau.[e](\sigma_1;\sigma_2) & \text{since }e\in S_{\mathsf{t}_{\tau'}}\\ &=[\lambda x:\tau.e](\sigma_1;\sigma_2) & \text{by (88), since }x\;\#\;(\sigma_1;\sigma_2)\;(\text{because }\sigma_1;\sigma_2\text{ is supported by }supp(\sigma_1)\cup supp(\sigma_2)\subseteq A) \end{split}$$

and hence $\lambda x : \tau. e \in S_{\mathsf{t}_{\tau \to \tau'}}$.

The single-variable substitution used in the definition of $=_{\beta\eta}$ in the previous section can be defined as:

$$e_1[x := e_2] \triangleq [e_1](\sigma_0\{x \mapsto e_2\})$$
 (94)

where in general the **updated substitution** $\sigma\{x \mapsto e\} \in Sub$ maps $x \in \mathbb{A}_{v_{\tau}}$ to $e \in \Lambda(\tau)$ and otherwise acts like $\sigma \in Sub$.

6.3 Denotation

We interpret each simple type $\tau \in Ty$ as a nominal set $D(\tau)$ as follows:

$$D(\iota) \triangleq N(\iota) \tag{95}$$

$$D(\tau \stackrel{.}{\rightarrow} \tau') \triangleq D(\tau) \rightarrow_{fs} D(\tau')$$
 (96)

An **environment** ρ is a function that, for any $\tau \in Ty$, maps the atoms in $\mathbb{A}_{v_{\tau}}$ to elements of $D(\tau)$. We let atom-permutations $\pi \in Perm$ act on such functions in the usual way (Section 3.2): the environment $\pi \cdot \rho$ maps $x \in \mathbb{A}_{v_{\tau}}$ to $\pi \cdot (\rho(\pi^{-1}(x)))$. Let Env be the nominal set of environments that are finitely supported with respect to this action.

Given $\tau \in \mathit{Ty}, \ e \in \Lambda(\tau)$ and $\rho \in \mathit{Env}$, we wish to define the **denotation** $\llbracket e \rrbracket \rho \in D(\tau)$, satisfying equations (89)–(91) in Figure 5. It is clear from the form of these equations that we should try to define $\llbracket e \rrbracket \rho$ by α -structural recursion for Σ^{STL} for all ρ simultaneously, because of the use of an **updated environment** in

(91): $\rho\{x\mapsto d\}$ is by definition the function mapping x to d and otherwise acting like ρ (and is finitely supported by $supp(\rho)\cup\{x\}\cup supp(d)$). So in Theorem 22 it seems that we should take $X_{\mathsf{t}_{\tau}}$ to be $Env\to_{\mathrm{fs}} D(\tau)$ and use the functions

$$f_{Vr_{\tau}} \triangleq \lambda x \in \mathbb{A}_{\mathsf{V}_{\tau}}.\lambda \rho \in Env. \, \rho \, x$$
 (97)

$$f_{Ap_{\tau,\tau'}} \triangleq \lambda(\xi,\xi') \in X_{\mathsf{t}_{\tau \to \tau'}} \times X_{\mathsf{t}_{\tau}}.\lambda\rho \in Env.\,\xi\,\rho\,(\xi'\rho) \tag{98}$$

$$f_{Lm_{\tau,\tau'}} \triangleq \lambda(x,\xi') \in \mathbb{A}_{\mathsf{v}_{\tau}} \times X_{\mathsf{t}_{\tau'}}.\lambda\rho \in Env.\lambda d \in D(\tau).\,\xi'(\rho\{x \mapsto d\}) \ . \tag{99}$$

These functions are all supported by $A=\emptyset$ and the only non-trivial freshness condition on binders is (FCB $_{Lm_{\tau,\tau'}}$), which in view of Theorem 14 (the "some/any" theorem) is the requirement that for all $x\in\mathbb{A}_{\mathsf{V}_{\tau}}$ and $\xi\in Env\to_{\mathrm{fs}}D(\tau')$

$$x \# \lambda \rho \in Env. \lambda d \in D(\tau). \, \xi(\rho \{x \mapsto d\}) . \tag{100}$$

If we could prove that, then the theorem gives us functions $\hat{f}_{t_{\tau}} \in \Lambda(\tau) \to_{\mathrm{fs}} (Env \to_{\mathrm{fs}} D(\tau))$ satisfying (46)—from which it follows that $[\![e]\!] \rho \triangleq \hat{f}_{t_{\tau}} e \rho$ satisfies (89)–(91). The problem is that (100) is not true for all elements ξ of $Env \to_{\mathrm{fs}} D(\tau')$! We have to strengthen the "recursion hypothesis" by suitably restricting the class of functions ξ that we consider.

This is the first time in this paper we have encountered a really non-trivial "freshness condition on binders". Equations (89)–(91) are typical of many such definitions in denotational semantics; but why is it the case that the right-hand side of equation (91) is independent of the choice of bound variable x on the left-hand side? Compared with the similar question for equation (88) a few lines above it, there seems no quick answer to this question. However, the freshness of x for $\lambda \rho \in Env. \lambda d \in D(\tau). [e](\rho\{x \mapsto d\})$ does follow from two expected properties of denotations:

- 1. The denotation of e with respect to an environment ρ only depends on the value of ρ at the free variables of e.
- 2. The denotation of a permuted version $\pi \cdot e$ of e with respect to an environment ρ is the denotation of e with respect to the composition $\rho \circ \pi$.

It is possible to take account of these facts in advance when applying Theorem 22 to construct denotations.¹¹ To do so we cut down to the following nominal subset of $Env \rightarrow_{fs} D(\tau)$:

$$X_{\mathsf{t}_{\tau}} \triangleq \{ \xi \in Env \to_{\mathsf{fs}} D(\tau) \mid \Phi_{1}(\xi) \& \Phi_{2}(\xi) \}$$
 (101)

where

$$\Phi_{1}(\xi) \triangleq (\exists A \in P_{fin}(\mathbb{A}))(\forall \tau \in Ty, x \in \mathbb{A}_{\mathsf{v}_{\tau}}, d \in D(\tau), \rho \in Env) \\ x \notin A \Rightarrow \xi(\rho\{x \mapsto d\}) = \xi \rho$$

$$\tag{102}$$

$$\Phi_2(\xi) \triangleq (\forall \pi \in Perm, \rho \in Env) (\pi \cdot \xi) \rho = \xi(\rho \circ \pi) \}.$$
 (103)

¹¹We are in effect carrying out a simultaneous inductive-recursive definition: see [7].

The slightly elaborate form of Φ_1 compared with property 1 above is needed to show that the functions defined in (97)–(99) satisfy

$$\Phi_{1}(f_{Vr_{\tau}} x)
\Phi_{1}(\xi) & \Phi_{1}(\xi') \Rightarrow \Phi_{1}(f_{Ap_{\tau,\tau'}}(\xi, \xi'))
\Phi_{1}(\xi') \Rightarrow \Phi_{1}(f_{Lm_{\tau,\tau'}}(x, \xi')) .$$

Similar properties hold for Φ_2 . Therefore $f_{Vr_{\tau}} \in \mathbb{A}_{\mathsf{v}_{\tau}} \to_{\mathsf{fs}} X_{\mathsf{t}_{\tau}}$, $f_{Ap_{\tau,\tau'}} \in X_{\mathsf{t}_{\tau \to \tau'}} \times X_{\mathsf{t}_{\tau}} \to_{\mathsf{fs}} X_{\mathsf{t}_{\tau'}}$ and $f_{Lm_{\tau,\tau'}} \in \mathbb{A}_{\mathsf{v}_{\tau}} \times X_{\mathsf{t}_{\tau'}} \to_{\mathsf{fs}} X_{\mathsf{t}_{\tau \to \tau'}}$. Furthermore, it is now the case that if $x \in \mathbb{A}_{\mathsf{v}_{\tau}}$ and $\xi' \in X_{\mathsf{t}_{\tau'}}$, then (100) holds. To see this, first note that since $\Phi_1(\xi')$ holds, there is a finite set of atoms A such that

$$(\forall \tau \in Ty, x' \in \mathbb{A}_{\mathsf{v}_{\tau}}, d' \in D(\tau), \rho' \in Env) \ x' \notin A \Rightarrow \xi'(\rho'\{x' \mapsto d'\}) = \xi'\rho' \ . \tag{104}$$

Choose any x' in the infinite set $\mathbb{A}_{v_{\tau}} - supp(x, \xi', A)$. Hence $x' \# \lambda \rho \in Env. \lambda d \in D(\tau). \xi'(\rho\{x \mapsto d\})$; and so applying the transposition $(x \ x')$ to this we get

$$\begin{split} x &= (x \; x') \cdot x' \\ &\# \; (x \; x') \cdot \lambda \rho \in Env. \\ \lambda d \in D(\tau).\xi'(\rho\{x \mapsto d\}) \\ &= \lambda \rho \in Env. \\ \lambda d \in D(\tau).((x \; x') \cdot \xi')(\rho\{x' \mapsto d\}) \\ &= \lambda \rho \in Env. \\ \lambda d \in D(\tau).\xi'(\rho\{x' \mapsto d\} \circ (x \; x')) \quad \text{because } \Phi_2(\xi') \text{ holds} \\ &= \lambda \rho \in Env. \\ \lambda d \in D(\tau).\xi'(\rho\{x \mapsto d\} \{x' \mapsto \rho \, x\}) \quad \text{because } x' \neq x \\ &= \lambda \rho \in Env. \\ \lambda d \in D(\tau).\xi'(\rho\{x \mapsto d\}) \quad \text{by (104), since } x' \notin A. \end{split}$$

So (100) does indeed hold. Therefore we can apply Theorem 22 to these nominal sets and functions to obtain $\hat{f}_{\mathsf{t}_{\tau}} \in X_{\mathsf{t}_{\tau}}$ satisfying (46). Defining $[\![e]\!] \rho \triangleq \hat{f}_{\mathsf{t}_{\tau}} e \rho$, we get the required properties (89)–(91).

Denotations respect $\beta\eta$ -conversion:

$$e_1 =_{\beta n} e_2 \implies [\![e_1]\!] = [\![e_2]\!].$$
 (105)

To see this it suffices to show that $[\![-]\!] = [\![-]\!]$ is a congruence relation equating β -convertible (84) and η -convertible (85) terms. Congruence is immediate from the defining properties (89)–(91) of $[\![-]\!]$. To see that $[\![-]\!]$ respects β -conversion one has to show

$$[e'[x := e]] \rho = [e'] (\rho \{x \mapsto [e] \rho \})$$
 (106)

and for η -conversion one has to show

$$x \notin fv(e) \Rightarrow \llbracket e \rrbracket (\rho \{x \mapsto d\}) = \llbracket e \rrbracket \rho . \tag{107}$$

These properties can be proved using the second α -structural induction principle (Theorem 23) for Σ^{STL} , with the second one used in the proof of the first.¹² For

¹²It might seem that (107) is a consequence of the property (102) that we built in to the construction of [-]; but unfortunately that property only tells us that $[e](\rho\{x\mapsto d\})$ and $[e]\rho$ are equal when x avoids some finite set of atoms, rather than the particular finite set supp(e).

(107) one can use the subsets

$$S_{\mathsf{t}_{\tau}} \triangleq \{e \in \Lambda(\tau) \mid (\forall x \in \mathbb{A}_{\mathsf{v}_{\tau}}, \rho \in Env, d \in D(\tau)) \ x \ \# \ e \Rightarrow \\ \llbracket e \rrbracket (\rho \{x \mapsto d\}) = \llbracket e \rrbracket \rho \}$$

which are all supported by the empty set. Properties (89)–(91) of [-] allow one to show that these subsets satisfy the induction hypotheses (IH $_{Vr_{\tau}}$), (IH $_{Ap_{\tau,\tau'}}$) and (IH $_{Lm_{\tau,ty'}}$); hence by Theorem 23 each $S_{t_{\tau}}$ is the whole of $\Lambda(\tau)$ and (107) holds. For (106), for each $x \in \mathbb{A}_{v_{\tau}}$ and $e \in \Lambda(\tau)$ we apply Theorem 23 to the subsets

$$S'_{\mathsf{t}_{\tau}} \triangleq \{e' \in \Lambda(\tau) \mid (\forall \rho \in Env) \, \llbracket e' [x := e] \rrbracket \rho = \llbracket e' \rrbracket (\rho \{x \mapsto \llbracket e \rrbracket \rho \}) \}$$

which are all supported by $A \triangleq \{x\} \cup supp(e)$. To conclude that $S'_{\mathsf{t}_{\tau}} = \Lambda(\tau)$ we have to prove the these subsets $S'_{\mathsf{t}_{\tau}}$ satisfy the induction hypotheses $(\mathrm{IH}_{Vr_{\tau}})$, $(\mathrm{IH}_{Ap_{\tau,\tau'}})$ and $(\mathrm{IH}_{Lm_{\tau,ty'}})$. The first two follow easily from (89) and (90). For $(\mathrm{IH}_{Lm_{\tau,ty'}})$ we have to show $(\exists x' \in \mathbb{A}_{\mathsf{v}_{\tau}})$ $x' \neq (x,e)$ & $(\forall e' \in S_{\mathsf{t}_{\tau}})$ $\lambda x' : \tau.e' \in S_{\mathsf{t}_{\tau}}$; but choosing any x' in the infinite set $\mathbb{A}_{\mathsf{v}_{\tau}} - supp(x,e)$, for each $e' \in S'_{\mathsf{t}_{\tau}}$ and $\rho \in Env$ we have:

$$\begin{split} & \llbracket (\lambda x' : \tau. e')[x := e] \rrbracket \rho \\ &= \llbracket \lambda x' : \tau. \left(e'[x := e] \right) \rrbracket \rho & \text{since } x' \notin fv(x, e) \\ &= \lambda d \in D(\tau). \llbracket e'[x := e] \rrbracket \left(\rho \{ x' \mapsto d \} \right) & \text{by (91)} \\ &= \lambda d \in D(\tau). \llbracket e' \rrbracket \left(\rho \{ x' \mapsto d \} \{ x \mapsto \llbracket e \rrbracket \left(\rho \{ x' \mapsto d \} \right) \} \right) & \text{since } e' \in S_{\mathsf{t}_\tau} \\ &= \lambda d \in D(\tau). \llbracket e' \rrbracket \left(\rho \{ x' \mapsto d \} \{ x \mapsto \llbracket e \rrbracket \rho \} \right) & \text{by (107), since } x' \notin fv(e) \\ &= \lambda d \in D(\tau). \llbracket e' \rrbracket \left(\rho \{ x \mapsto \llbracket e \rrbracket \rho \} \{ x' \mapsto d \} \right) & \text{since } x' \neq x \\ &= \llbracket \lambda x' : \tau. e' \rrbracket \left(\rho \{ x \mapsto \llbracket e \rrbracket \rho \} \right) & \text{by (91)} \end{split}$$

so that $\lambda x' : \tau. e' \in S_{\mathsf{t}_{\tau}}$, as required.

6.4 Reification, reflection and normalisation

The **reification function** $\downarrow_{\tau} \in D(\tau) \to_{\mathrm{fs}} N(\tau)$ turns elements of the denotational model into $\beta\eta$ -long normal forms, whereas the **reflection function** $\uparrow_{\tau} \in U(\tau) \to_{\mathrm{fs}} D(\tau)$ turns neutral terms into denotational elements. Both functions are defined simultaneously by ordinary structural recursion for simple type symbols $\tau \in Ty$ in Figure 6, where for clarity we have written α -terms over the signature Σ^{LNF} without the conventions introduced by Notation 28. The interesting part of the definition in this figure is clause (109), where we make use of the fresh construct from Theorem 16. To do so we have to check that the conditions of that theorem are satisfied; but in this case that is easy: given $f \in D(\tau \to \tau')$, choosing any atom x in the infinite set $\mathbb{A}_{\mathsf{V}_{\tau}} - supp(\downarrow_{\tau'}, \uparrow_{\tau}, f)$, \mathbb{A}^{13} then the element

$$h \triangleq \lambda x \in \mathbb{A}_{\mathsf{v}_{\tau}}.L_{\tau,\tau'} x.\downarrow_{\tau'} (f(\uparrow_{\tau}(V_{\tau} x)))$$

¹³In fact $supp(\downarrow_{\tau'}, \uparrow_{\tau}, f) = supp(f)$ because the reification and reflection functions turn out to have empty support.

 $\tau \in Ty, d \in D(\tau) \mapsto \downarrow_{\tau}(d) \in N(\tau)$:

$$\downarrow_{I}(n) \triangleq n \tag{108}$$

$$\downarrow_{\tau \to \tau'}(f) \triangleq fresh(\lambda x \in \mathbb{A}_{\mathsf{V}_{\tau}}. L_{\tau,\tau'} x. \downarrow_{\tau'} (f(\uparrow_{\tau}(V_{\tau} x)))$$
 (109)

 $\tau \in Ty, u \in U(\tau) \mapsto \uparrow_{\tau}(u) \in D(\tau)$:

$$\uparrow_{\iota}(u) \triangleq I u \tag{110}$$

$$\uparrow_{\tau \to \tau'}(u) \triangleq \lambda d \in D(\tau). \uparrow_{\tau'}(A_{\tau,\tau'}(u,\downarrow_{\tau}(d))) \tag{111}$$

Figure 6: Reification (\downarrow_{τ}) and reflection (\uparrow_{τ})

of $\mathbb{A}_{\mathsf{v}_{\tau}} \to_{\mathsf{fs}} N(\tau')$ satisfies x # h (by choice of x) and x # h(x) (because $x \notin fv(L_{\tau,\tau'} x.n)$ for any $n \in N(\tau')$); so we can form fresh(h) as in the theorem.

Definition 29. The **initial environment** ρ_0 maps each $x \in \mathbb{A}_{v_{\tau}}$ to $\uparrow_{\tau}(V_{\tau}x) \in D(\tau)$, for all $\tau \in Ty$. This has empty support (because the \uparrow_{τ} functions do) and hence in particular it is an element of the nominal set Env of finitely supported environments. Using it, we define the **normalisation function** $norm_{\tau} \in \Lambda(\tau) \to_{fs} N(\tau)$ by:

$$norm_{\tau}(e) \triangleq \downarrow_{\tau}(\llbracket e \rrbracket \rho_0) .$$
 (112)

Recall that we wish to show that $norm_{\tau}$ has the properties (81)–(83). Property (81) follows immediately from (105). Property (82) is the first half of the following result.

Lemma 30. For all $\tau \in Ty$, $n \in N(\tau)$ and $u \in U(\tau)$

$$\downarrow_{\tau}(\llbracket i_{\tau} \, n \rrbracket \rho_0) = n \tag{113}$$

$$[\![j_\tau \, u]\!] \rho_0 = \uparrow_\tau(u) . \tag{114}$$

Proof. These properties can be proved by using the second α -structural induction principle (Theorem 23) for $\Sigma^{\rm LNF}$ to show that the subsets

$$S_{\mathsf{n}_{\tau}} \triangleq \{ n \in N(\tau) \mid \downarrow_{\tau} (\llbracket i_{\tau} \, n \rrbracket \rho_0) = n \}$$

$$S_{\mathsf{u}_{\tau}} \triangleq \{ u \in U(\tau) \mid \llbracket j_{\tau} \, u \rrbracket \rho_0 = \uparrow_{\tau} (u) \}$$

are equal to $N(\tau)$ and $U(\tau)$ respectively (for all $\tau \in Ty$). Since these subsets are

44 Andrew Pitts

$$\begin{split} \sim_{\tau} &\subseteq D(\tau) \times \Lambda(\tau) : \\ &n \sim_{\iota} e \; \triangleq \; (i_{\iota} \, n =_{\beta\eta} e) \\ &f \sim_{\tau \stackrel{.}{\rightarrow} \tau'} e \; \triangleq \; (\forall d_1 \in D(\tau), e_1 \in \Lambda(\tau)) \; d_1 \sim_{\tau} e_1 \; \Rightarrow \; f \, d \sim_{\tau'} A p_{\tau,\tau'}(e,e_1) \\ &\sim \subseteq Env \times Sub : \\ &\rho \sim \sigma \; \triangleq \; (\forall \tau \in Ty, x \in \mathbb{A}_{\mathsf{V}_{\tau}}) \; \rho \, x \sim_{\tau} \sigma \, x \end{split}$$

Figure 7: Logical relation

supported by the empty set of atoms one can take $A = \emptyset$ in the theorem and prove

$$(\forall x \in \mathbb{A}_{\mathsf{V}_{\tau}}) \ V_{\tau} \ x \in S_{\mathsf{U}_{\tau}}$$

$$(\mathsf{IH}_{V_{\tau}})$$

$$(\forall u \in S_{\mathsf{U}_{\tau \to \tau'}}, n \in S_{\mathsf{n}_{\tau}}) \ A_{\tau,\tau'}(u,n) \in S_{\mathsf{U}_{\tau'}}$$

$$(\mathsf{IH}_{A_{\tau,\tau'}})$$

$$(\exists x \in \mathbb{A}_{\mathsf{V}_{\tau}}) (\forall n \in S_{\mathsf{n}_{\tau'}}) \ L_{\tau,\tau'} \ x. \ n \in S_{\mathsf{n}_{\tau \to \tau'}}$$

$$(\mathsf{IH}_{L_{\tau,\tau'}})$$

$$(\forall u \in S_{\mathsf{U}_{\iota}}) \ I \ u \in S_{\mathsf{U}_{\iota}} .$$

$$(\mathsf{IH}_{I})$$

These nearly all follow directly from the definitions of $i_{\tau}, j_{\tau}, \downarrow_{\tau}, \uparrow_{\tau}, \llbracket - \rrbracket$ and ρ_0 . The only tricky case is for $(\mathrm{IH}_{L_{\tau,\tau'}})$, because of the use of fresh in $\downarrow_{\tau \to \tau'}$. Picking any atom x in $\mathbb{A}_{\mathsf{V}_{\tau}}$, suppose $n \in S_{\mathsf{n}_{\tau'}}$. We wish to prove that $L_{\tau,\tau'} x. n \in S_{\mathsf{n}_{\tau \to \tau'}}$. Note that $x \# L_{\tau,\tau'} x. n$; so since $\llbracket - \rrbracket, i_{\tau \to \tau'}$ and ρ_0 have empty support, it is the case that $x \# \llbracket i_{\tau \to \tau'} (L_{\tau,\tau'} x. n) \rrbracket \rho_0 = \llbracket \lambda x : \tau. i_{\tau'} n \rrbracket \rho_0 = f$, where $f \triangleq \lambda d \in D(\tau)$. $\llbracket i_{\tau'} n \rrbracket (\rho_0 \{x \mapsto d\})$. Hence

$$x \# (\lambda x' \in \mathbb{A}_{\mathsf{v}_{\tau}}. L_{\tau,\tau'} x'. \downarrow_{\tau'} (f(\uparrow_{\tau}(V_{\tau} x'))))$$

and therefore by definition of fresh in Theorem 16

$$\begin{split} &\downarrow_{\tau \dot{\to} \tau'} (\llbracket i_{\tau \dot{\to} \tau'} (L_{\tau,\tau'} \, x. \, n) \rrbracket \rho_0) \\ &= fresh(\lambda x' \in \mathbb{A}_{\mathsf{V}_\tau}. \, L_{\tau,\tau'} \, x'. \, \downarrow_{\tau'} (f(\uparrow_\tau(V_\tau \, x')))) \\ &= L_{\tau,\tau'} \, x. \, \downarrow_{\tau'} (f(\uparrow_\tau(V_\tau \, x))) \\ &= L_{\tau,\tau'} \, x. \, \downarrow_{\tau'} (\llbracket i_{\tau'} \, n \rrbracket (\rho_0 \{x \mapsto \uparrow_\tau(V_\tau \, x)\})) \\ &= L_{\tau,\tau'} \, x. \, \downarrow_{\tau'} (\llbracket i_{\tau'} \, n \rrbracket \rho_0) & \text{by definition of } \rho_0 \\ &= L_{\tau,\tau'} \, x. \, n & \text{since } n \in S_{\mathsf{n}_{\tau'}} \end{split}$$

so that $L_{\tau,\tau'} x. n$ is indeed an element of $S_{n_{\tau \to \tau'}}$.

It just remains to prove that $norm_{\tau}$ has the property (83). For this we use a **logical relation** $\sim_{\tau} \subseteq D(\tau) \times \Lambda(\tau)$ between elements of the denotational model

and terms. It is defined by ordinary structural recursion for simple type symbols $\tau \in Ty$ in Figure 7, which also extends the relation to one between (finitely supported) environments and substitutions.

Lemma 31 (fundamental property of \sim). For all $\tau \in Ty$, $e \in \Lambda(\tau)$, $\rho \in Env$ and $\sigma \in Sub$

$$\rho \sim \sigma \implies \llbracket e \rrbracket \rho \sim_{\tau} [e] \sigma . \tag{115}$$

Proof. This can be proved by applying the second α -structural induction theorem for the nominal signature Σ^{STL} to show that the subsets

$$S_{\tau} \triangleq \{e \in \Lambda(\tau) \mid (\forall \rho \in Env, \sigma \in Sub) \mid \rho \sim \sigma \Rightarrow \llbracket e \rrbracket \rho \sim_{\tau} [e] \sigma \}$$

are equal to $\Lambda(\tau)$, for all $\tau \in \mathit{Ty}$. Proving the induction hypotheses $(\mathrm{IH}_{\mathit{Vr}_{\tau}})$ and $(\mathrm{IH}_{\mathit{Ap}_{\tau,\tau'}})$ is straightforward; and for $(\mathrm{IH}_{\mathit{Lm}_{\tau,\tau'}})$ one first proves

$$d \sim_{\tau} e \& e =_{\beta \eta} e' \Rightarrow d \sim_{\tau} e' \tag{116}$$

$$([\lambda x. e]\sigma) e' =_{\beta \eta} [e](\sigma\{x \mapsto e'\})$$
(117)

where $\sigma\{x \mapsto e'\}$ indicates an updated substitution mapping x to e' and otherwise acting like σ .

One can prove by ordinary structural induction for types $\tau \in Ty$ that

$$(\forall \tau \in Ty, u \in U(\tau)) \uparrow_{\tau}(u) \sim_{\tau} j_{\tau} u \tag{118}$$

$$(\forall \tau \in Ty, d \in D(\tau), e \in \Lambda(\tau)) \ d \sim_{\tau} e \ \Rightarrow \ i_{\tau}(\downarrow_{\tau}(d)) =_{\beta\eta} e \ . \tag{119}$$

(Both properties are proved simultaneously, and one needs to make use of (116).) Because of (118), the identity substitution $\sigma_0 \in Sub$ satisfies $\rho_0 \sim \sigma_0$. So by Lemma 31 and (92) we have $[\![e]\!] \rho_0 \sim_{\tau} [\![e]\!] \sigma_0 = e$, for all $e \in \Lambda(\tau)$. Thus (119) gives $i_{\tau}(\downarrow_{\tau}([\![e]\!] \rho_0)) =_{\beta\eta} e$, i.e. $i_{\tau}(norm_{\tau}(e)) =_{\beta\eta} e$, as required for (83).

7 Assessment

7.1 Mathematical perspective

The results of this paper are directly inspired by my joint work with Gabbay on "FM-set" theory [13] and by his PhD thesis [11]; in particular those works contain structural recursion and induction principles for an inductively defined FM-set isomorphic to λ -terms modulo α -equivalence. Here I have taken an approach that is both a bit more general and more concrete: more general, because the particular signature for λ -terms has been replaced by an arbitrary nominal signature (a notion which comes from joint work with Urban and Gabbay [28] and is developed further in Cheney's thesis [4]); and more concrete in two respects. First, the key notion of (finite) *support* has been developed using nominal sets within the framework of ordinary higher-order logic, rather than being axiomatised within FM-set

46 Andrew Pitts

theory; see Cheney [4, Chapter 3] for a more leisurely and generalised account of the theory of nominal sets. Secondly, the recursion and induction principles developed here refer directly to α -terms, i.e. standard α -equivalence classes of abstract syntax trees, rather than using an intial algebra that is merely isomorphic to the set of α -terms; see Remark 24. This is also the approach taken by Norrish [18], building on Gordon and Melham's five axioms for α -equivalence [14]; and also by Urban and Tasson [29]. Norrish's recursion principle [18, Fig. 1] has sideconditions requiring that the function being defined be well-behaved with respect to variable-permutations and with respect to fresh name generation. In effect these side-conditions build in just enough of the theory of nominal sets to yield a welldefined and total function, while having to specify how binders with only fresh names are mapped by the function. Along with Urban and Tasson [29], I prefer to develop the theory of nominal sets in its own right and then give a simple-looking ¹⁴ recursion principle within that theory. One advantage of such an approach is that it makes it easier to identify and use properties of name freshness, such as Theorem 16, independently of the recursion principle. We used Theorem 16 in the reduction of Theorem 17 to Theorem 3, in the reduction of "varying parameters" to "no parameters" (Example 27) and in the definition of the reification functions in the extended example in Section 6; another good example of its use occurs (implicitly) in the denotational semantics of FreshML's fresh expression [26, Section 3].

How easy is it to apply these principles of α -structural recursion and induction? Just as for the work of Gordon-Melham, Norrish and Urban-Tasson [14, 18, 29], to use them one does not have to change to an unfamiliar logic (we remain in higher-order logic), or a new way of representing syntax (we use the familiar notion of α -equivalence classes of abstract syntax trees). One does have to get used to thinking in terms of permutations and finite support; and the latter is undoubtedly a subtle concept at higher types. However, the relativisation from arbitrary mathematical objects to finitely supported ones called for by this approach is made easier by the fact (Theorem 11) that the finite support property is conserved by all the usual constructs of higher-order logic except for uses of the axiom of choice. Thus if some language of interest has been specified as the α -terms for a particular nominal signature and one wishes to define a function on those α -terms specified by an instance of the recursion scheme (46) in Theorem 22 (for suitable functions f_K), then there are three tasks involved in applying the theorem to this data:

- (I) Show that the sets X_s that one is mapping into have the structure of nominal sets.
- (II) Show that the functions f_K are all supported by a single finite set of atoms A.
- (III) Show that each function f_K satisfies the "freshness condition on binders" (FCB $_K$).

¹⁴Once one gets used to the distinctive concepts of nominal sets, I believe that principles such as Theorems 17, 22 and 25 *are* quite simple.

Task (I) is usually carried out by showing how the X_s are built up from some standard nominal sets (such as those in Example 7) using the constructions described in Sections 3.2 and 3.3. Task (II) might seem quite difficult, but in fact the Finite Support Principle (Theorem 11) usually reduces it to seeing how the f_K are constructed within higher-order logic. So really the main difficulty is task (III). In some cases, such as the capture-avoiding substitution function in Example 26, (FCB $_K$) is very easily checked. In other cases, such as in the definition of the denotation functions $[\![-]\!]$ in the extended example of Section 6, one has to work hard to verify the freshness condition on binders.

Applying the α -structural induction principles is somewhat easier. For example, for the second α -structural induction principle (Theorem 23) one still has the analogues of the easy tasks (I) and (II); and then one just has to verify the induction hypothesis (IH $_K$) for each constructor K, which is in fact a more restricted property than the corresponding induction hypothesis in an ordinary structural induction for terms rather than α -terms.

7.2 Automated theorem proving perspective

Based on my experience with other formalisms, I claim that the use of permutations and finitely supported objects advocated here is a simple, effective and yet rigorous way of dealing with binders and α -equivalence in "paper-and-pencil" proofs in programming language semantics. But how easy is it to provide computer support for reasoning with α -structural recursion and induction? Of the three types (I–III) of task involved in applying these principles in any particular case that were mentioned above, task (III) will require human-intervention; but in view of Theorem 11, there is the possibility of automating tasks (I) and (II). One way of attempting that is to develop a new higher-order logic in which types only denote nominal sets and that axiomatises properties of permutations and finite support; this is the route taken by Gabbay with his FM-HOL [12]. The disadvantage of such a "new logic" approach is that one does not have easy access to the legacy of alreadyproved results in systems such as HOL4 and Isabelle/HOL. To what extent tasks (I) and (II) can be automated within these "legacy" mechanised logics remains to be seen. The work of Norrish [18] provides a starting point within the HOL4 system; and Urban and Tasson [29] have already developed a theory equivalent to nominal sets within Isabelle/HOL up to and including an induction principle (but not yet a recursion principle) for the particular nominal signature for λ -terms.¹⁵ It is to be hoped that these works will lead to an augmentation of the treatment of data types withing HOL4 or Isabelle/HOL, allowing the user to declare a nominal signature and then have the principles of α -structural recursion and induction for that signature proved and ready to be applied. 16

¹⁵Their proof of validity of the induction principle follows a different route from the one used here to prove the α -structural recursion and induction principles.

¹⁶How best in such a package to deal with the relativisation to nominal sets is certainly an issue; Urban and Tasson report that Isabelle's *axiomatic type classes* are helpful in this respect.

Acknowledgements. I am grateful to James Cheney, Murdoch Gabbay, Michael Norrish, Mark Shinwell and Christian Urban for their many contributions to the subject of this paper which is a revised and much extended version of [21].

A Proof of Second α-Structural Induction Theorem

Given a nominal signature Σ , we prove the second α -structural induction principle (Theorem 23) as a corollary of the first one, Theorem 21. To do so we must first develop some properties of the operation $\bar{a} \in \mathbb{A}^{\sigma}, \bar{e} \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|} \mapsto \bar{a}.\bar{e} \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma}$ and the relation $\sharp_{\sigma} \subseteq \mathbb{A}^{\sigma} \times \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}$ defined in Figure 2.

Lemma 32. If $\bar{a} \in \mathbb{A}^{\sigma}$, $\bar{e} \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}$ and $\bar{a} \sharp_{\sigma} \bar{e}$, then $supp(\bar{a}.\bar{e}) = supp(\bar{e}) - supp(\bar{a})$.

Proof. This can be proved by ordinary structural recursion for arities. The induction step for pair arities uses the fact (20) that the support of a pair is the union of the supports of its two components. The induction step for atom-binding arities uses the fact, noted in Example 7(iii), that the support of an α -term is its finite set of free atoms (so that in particular $supp(a.e) = supp(e) - \{a\}$).

Lemma 33. Let $(S_s \in P_{fs}(T_{\alpha}(\Sigma)_s) \mid s \in \Sigma_D)$ be a family of finitely supported subsets of α -terms indexed by the data-sorts of Σ . Extend this to arity-indexed families

$$(S_{\sigma} \in P_{\mathrm{fs}}(\mathsf{T}_{\alpha}(\Sigma)_{\sigma}) \mid \sigma \in Ar(\Sigma))$$
$$(S^{|\sigma|} \in P_{\mathrm{fs}}(\mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}) \mid \sigma \in Ar(\Sigma))$$

as in Figure 3. Then for any finite set A of atoms, each element of S_{σ} is of the form $\bar{a}.\bar{e}$ for some $\bar{a} \in \mathbb{A}^{\sigma}$ and $\bar{e} \in S^{|\sigma|}$ satisfying $\bar{a} \# A$ and $\bar{a} \sharp_{\sigma} \bar{e}$.

Proof. This can be proved by ordinary structural induction for arities, for all A simultaneously. The induction steps when $\sigma=a\in\Sigma_A$, $\sigma=s\in\Sigma_D$ and $\sigma=1$ are trivial.

Case $\sigma=\sigma_1*\sigma_2$. By definition of S_σ , each $e\in S_\sigma$ is of the form $e=(e_1,e_2)$ with $e_i\in S_{\sigma_i}$ for i=1,2. Given A, by induction hypothesis for σ_1 applied to the finite set of atoms $A\cup supp(e_2)$, we can find $\bar{a}_1\in \mathbb{A}^{\sigma_1}$ and $\bar{e}_1\in S^{|\sigma_1|}$ with $e_1=\bar{a}_1.\bar{e}_1,\ \bar{a}_1\ \#\ (A,e_2)$ and $\bar{a}_1\ \sharp_{\sigma_1}\ \bar{e}_1$. Then by induction hypothesis for σ_2 , we can find $\bar{a}_2\in \mathbb{A}^{\sigma_2}$ and $\bar{e}_2\in S^{|\sigma_2|}$ with $e_2=\bar{a}_2.\bar{e}_2,\ \bar{a}_2\ \#\ (A,\bar{a}_1,\bar{e}_1)$ and $\bar{a}_2\ \sharp_{\sigma_2}\ \bar{e}_2$. Since $\bar{a}_1\ \#\ \bar{a}_2$, we have $(\bar{a}_1,\bar{a}_2)\in \mathbb{A}^{\sigma_1}\otimes \mathbb{A}^{\sigma_2}=\mathbb{A}^\sigma$. So $e=(e_1,e_2)=(\bar{a}_1,\bar{a}_2).(\bar{e}_1,\bar{e}_2)$ with $(\bar{a}_1,\bar{a}_2)\ \#\ A$ and $(\bar{e}_1,\bar{e}_2)\in S^{|\sigma_1|}\times S^{|\sigma_2|}=S^{|\sigma|}$. So it just remains to show $(\bar{a}_1,\bar{a}_2)\ \#\ A$ and $(\bar{e}_1,\bar{e}_2)\in S^{|\sigma_1|}\times S^{|\sigma_2|}=S^{|\sigma|}$. So it just remains to show $(\bar{a}_1,\bar{a}_2)\ \#\ A$ in this, since $\bar{a}_i\ \#\ a_i$ if $\bar{a}_i\ i$ in the have its support disjoint from $supp(e_2)$, which by Lemma 32 is $supp(\bar{e}_2)-supp(\bar{a}_2)$; since $supp(\bar{a}_1)$ is also disjoint from $supp(\bar{a}_2)$, it follows that $supp(\bar{a}_1)\cap supp(\bar{e}_2)=\emptyset$, i.e. $\bar{a}_1\ \#\ \bar{e}_2$, as required.

 $\begin{array}{l} \textit{Case}\ \sigma = \text{$<$a$>$$\sigma_1$. By definition of S_σ, each $e \in S_\sigma$ is of the form $e = a$. e_1 with $a \in \mathbb{A}_{\mathsf{a}} - supp(S_{\sigma_1})$ and $e_1 \in S_{\sigma_1}$. Given A, choosing any atom a' in the infinite set $\mathbb{A}_{\mathsf{a}} - supp(A, a, e_1)$, it follows from the characterisation of α-equivalence in (28) that $e = a'$. e'_1 where $e'_1 \triangleq (a\ a') \cdot e_1 \in (a\ a') \cdot S_{\sigma_1} = S_{\sigma_1}$, since $a, a' \# S_{\sigma_1}$. Then by induction hypothesis for σ_1 applied to the finite set of atoms $A \cup \{a'\}$, we can find $\bar{a}_1 \in \mathbb{A}^{\sigma_1}$ and $\bar{e}_1 \in S^{|\sigma_1|}$ with $e'_1 = \bar{a}_1.\bar{e}_1$, $\bar{a}_1 \# (A,a')$ and $\bar{a}_1 \sharp_{\sigma_1} \bar{e}_1$. since $a' \# \bar{a}_1$, we have $(a',\bar{a}_1) \in \mathbb{A}_{\mathsf{a}} \otimes \mathbb{A}^{\sigma_1} = \mathbb{A}^{\sigma}$. So $e = a'$. $e'_1 = (a',\bar{a}_1)$. \bar{e}_1 with $(a',\bar{a}_1) \# A$, $\bar{e}_1 \in S^{|\sigma_1|} = S^{|\sigma|}$ and $(a',\bar{a}_1) \sharp_{\sigma} \bar{e}_1$.} \label{eq:case_sigma}$

To prove Theorem 23 we also need to see that its induction hypotheses (IH $_K$) are equivalent to ones that universally rather than existentially quantify over suitably fresh nested tuples \bar{a} of atoms. To do so we make use of the following generalisation of permutations that transpose a pair of atoms.

Definition 34 (vector-transpositions). Let Σ be a nominal signature and $\sigma \in Ar(\Sigma)$ an arity over the signature. Let the nominal set \mathbb{A}^{σ} of nested tuples of distinct atoms be defined as in Figure 2. For each pair of elements $\bar{a}, \bar{a}' \in \mathbb{A}^{\sigma}$ with $\bar{a} \# \bar{a}'$, let $\tau_{\bar{a},\bar{a}'} \in Perm$ be the atom-permutation that transposes the atoms at corresponding positions in the nested tuples \bar{a} and \bar{a}' . More formally, $\tau_{\bar{a},\bar{a}'}$ is defined by ordinary structural recursion on the arity σ as follows:

$$\begin{array}{c|c} \sigma & (\bar{a},\bar{a}') \in \mathbb{A}^{\sigma} \otimes \mathbb{A}^{\sigma} \mapsto \tau_{\bar{a},\bar{a}'} \in Perm \\ \hline \mathbf{a} \in \Sigma_{\mathbf{A}} & ((),()) & \mapsto \quad \iota \\ \mathbf{s} \in \Sigma_{\mathbf{D}} & ((),()) & \mapsto \quad \iota \\ \mathbf{1} & ((),()) & \mapsto \quad \iota \\ \sigma_1 * \sigma_2 & ((\bar{a}_1,\bar{a}_2),(\bar{a}'_1,\bar{a}'_2)) \mapsto \tau_{\bar{a}_1,\bar{a}'_1} \circ \tau_{\bar{a}_2,\bar{a}'_2} \\ \mathbf{``a``} \sigma_1 & ((a,\bar{a}_1),(a',\bar{a}'_1)) & \mapsto (a\ a') \circ \tau_{\bar{a}_1,\bar{a}'_1} \end{array}$$

Although the clauses defining $\tau_{\bar{a},\bar{a}'}$ make sense for any pair \bar{a},\bar{a}' , we restrict to pairs satisfying $\bar{a} \# \bar{a}'$ to ensure that $\tau_{\bar{a},\bar{a}'}$ enjoys many properties of single-atom transpositions, $(a\ a')$. In particular one can easily prove by ordinary structural induction for arities that:

$$\pi \circ \tau_{\bar{a},\bar{a}'} \circ \pi^{-1} = \tau_{(\pi \cdot \bar{a}),(\pi \cdot \bar{a}')} \qquad (\text{any } \pi \in Perm)$$
 (120)

$$supp(\tau_{\bar{a},\bar{a}'}) \subseteq supp(\bar{a}) \cup supp(\bar{a}')$$
 (121)

$$\tau_{\bar{a},\bar{a}'} = \tau_{\bar{a}',\bar{a}} \tag{122}$$

$$\tau_{\bar{a},\bar{a}'} \circ \tau_{\bar{a},\bar{a}'} = \iota \tag{123}$$

$$\tau_{\bar{a}.\bar{a}'} \cdot \bar{a} = \bar{a}' \ . \tag{124}$$

Part (ii) of the following result generalises the "some/any" Theorem 14 from single atoms to elements of \mathbb{A}^{σ} .

Lemma 35. (i) For all finite sets of atoms A, there is some $\bar{a} \in \mathbb{A}^{\sigma}$ with $\bar{a} \# A$.

(ii) For all finitely supported subsets $S \in P_{fs}(\mathbb{A}^{\sigma})$, if S is supported by the finite set of atoms A, then the following statements are equivalent.

$$(\forall \bar{a} \in \mathbb{A}^{\sigma}) \ \bar{a} \# A \Rightarrow \bar{a} \in S \tag{125}$$

$$(\exists \bar{a} \in \mathbb{A}^{\sigma}) \ \bar{a} \# A \& \ \bar{a} \in S \ . \tag{126}$$

Proof. Part (i) follows easily by ordinary structural induction for arities $\sigma \in Ar(\Sigma)$. For part (ii), first note that in view of part (i), one just has to show that (126) implies (125). Suppose $\bar{a} \in \mathbb{A}^{\sigma}$ satisfies $\bar{a} \# A \& \bar{a} \in S$. Given any other $\bar{a}' \in \mathbb{A}^{\sigma}$ with $\bar{a}' \# A$, we have to show $\bar{a}' \in S$. We can use part (i) to find $\bar{a}'' \in \mathbb{A}^{\sigma}$ with $\bar{a}'' \# (A, \bar{a}, \bar{a}')$ and then use the vector-transpositions of Definition 34 to define $\pi \triangleq \tau_{\bar{a}',\bar{a}''} \circ \tau_{\bar{a},\bar{a}''} \in Perm.^{17}$ By (124), we have $\pi \cdot \bar{a} = \bar{a}'$; and because by (121) π is supported by $supp(\bar{a}, \bar{a}', \bar{a}'')$, which is disjoint from A and hence from supp(S), we have $\pi \cdot S = S$. So applying π to $\bar{a} \in S$ we get $\bar{a}' \in S$, as required.

The following result generalises the characterisation of α -equivalence mentioned in Example 15. Note that in view of the previous lemma, the right-hand side of the bi-implication in (127) could be replaced by $(\forall \bar{a}'' \in \mathbb{A}^{\sigma}) \ \bar{a}'' \# A \Rightarrow \tau_{\bar{a},\bar{a}''} \cdot \bar{e} = \tau_{\bar{a}',\bar{a}''} \cdot \bar{e}' \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}$.

Lemma 36. Suppose $\bar{a}, \bar{a}' \in \mathbb{A}^{\sigma}$ and $\bar{e}, \bar{e}' \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}$ satisfy $\bar{a} \not\parallel_{\sigma} \bar{e}$ and $\bar{a}' \not\parallel_{\sigma} \bar{e}'$. If A is a finite set of atoms supporting $(\bar{a}, \bar{a}', \bar{e}, \bar{e}')$, then

$$\bar{a}.\,\bar{e} = \bar{a}'.\,\bar{e}' \in \mathsf{T}_{\alpha}(\Sigma)_{\sigma} \Leftrightarrow (\exists \bar{a}'' \in \mathbb{A}^{\sigma})\,\bar{a}'' \# A \& \tau_{\bar{a},\bar{a}''} \cdot \bar{e} = \tau_{\bar{a}',\bar{a}''} \cdot \bar{e}' \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|} . \tag{127}$$

Proof. This can be proved by ordinary structural induction for arities σ . The induction step for atom-binding arities uses the corresponding property (28) of $=_{\alpha}$.

We can now complete the proof of Theorem 23. We are given subsets $(S_s \in P_{fs}(\mathsf{T}_\alpha(\Sigma)_s) \mid s \in \Sigma_D)$ supported by the finite set of atoms A and satisfying (IH_K) for each $K \in \Sigma_C$. Define $(S_\sigma \in P_{fs}(\mathsf{T}_\alpha(\Sigma)_\sigma) \mid \sigma \in Ar(\Sigma))$ as in the right-hand column in Figure 3. It is not hard to see that all of these subsets are also supported by A and hence so is their union $S \triangleq \bigcup_{\sigma \in Ar(\Sigma)} S_\sigma$. If $S = \mathsf{T}_\alpha(\Sigma)$, then $S_s = \mathsf{T}_\alpha(\Sigma)_s$ for each $s \in \Sigma_D$; and to prove $S = \mathsf{T}_\alpha(\Sigma)$ we just have to prove that this S satisfies the conditions (40)–(44) of Theorem 21. Conditions (40) and (42)–(44) are immediate from the definition of $(S_\sigma \mid \sigma \in Ar(\Sigma))$ in Figure 3. For condition (41), given $(K : \sigma \to s) \in \Sigma_C$ and $e \in S_\sigma$, we have to show that $K e \in S_s$. First note that by applying Lemma 35(ii) to the subset $\{\bar{a} \in \mathbb{A}^\sigma \mid (\forall \bar{e} \in \mathsf{T}_\alpha(\Sigma)^{|\sigma|}) \ \bar{a} \not \models_\sigma \bar{e} \& \bar{e} \in S^{|\sigma|} \Rightarrow K \bar{a}. \bar{e} \in S_s\}$, which is supported by A, to see that (IH_K) is equivalent to

$$(\forall \bar{a} \in \mathbb{A}^{\sigma}) \ \bar{a} \# A \Rightarrow (\forall \bar{e} \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}) \ \bar{a} \sharp_{\sigma} \bar{e} \& \bar{e} \in S^{|\sigma|} \Rightarrow K \ \bar{a}. \bar{e} \in S_{\mathsf{c}}. \tag{128}$$

¹⁷Since it may not be the case that $\bar{a} \# \bar{a}'$, we cannot use $\tau_{\bar{a},\bar{a}'}$; so (unlike in the proof of Theorem 14) we resort to swapping via an intermediate, fresh \bar{a}'' .

By Lemma 33, $e = \bar{a}.\bar{e}$ with $\bar{a} \in \mathbb{A}^{\sigma}$, $\bar{e} \in S^{|\sigma|}$, $\bar{a} \# A$ and $\bar{a} \sharp_{\sigma} \bar{e}$; so (128) gives $Ke \in S_s$. Thus we can indeed apply Theorem 21 to conclude that S is the whole of $T_{\alpha}(\Sigma)$ and hence in particular that each S_s is equal to the whole of $T_{\alpha}(\Sigma)_s$, completing the proof of Theorem 23.

B Proof of Second α -Structural Recursion Theorem

We will establish the second α -structural recursion principle (Theorem 22) via a common strategy for reducing recursion to induction: we first construct relations (using a rule-based inductive definition) that would be the graphs of the required functions \hat{f}_s were they single-valued and total relations—and then prove they are so by applying Theorem 23. A further application of Theorem 23 is needed to show that the functions \hat{f}_s are unique with the stated properties.

So suppose we are given a family of nominal sets $X=(X_{\mathsf{s}}\mid \mathsf{s}\in\Sigma_{\mathsf{D}})$ indexed by the data-sorts of a nominal signature Σ , and functions $(f_K\in X^{(\sigma)}\to_{\mathsf{fs}} X^{(\mathsf{s})}\mid (K:\sigma\to\mathsf{s})\in\Sigma_{\mathsf{C}})$ all of which are supported by a finite set A of atoms and satisfy (FCB $_K$).

Existence of the functions $\hat{f}_s \in T_\alpha(\Sigma)_s \to_{fs} X_s$

Consider the subsets $F_s \subseteq T_\alpha(\Sigma)_s \times X_s$ (for $s \in \Sigma_D$) that are inductively defined by the following rule

$$\frac{K: \sigma \to \mathsf{s} \quad \bar{a} \in \mathbb{A}^{\sigma} \quad \bar{e} \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}}{\bar{a} \# A \quad \bar{a} \sharp_{\sigma} \bar{e} \quad (\bar{e}, \bar{x}) \in F^{|\sigma|} \quad \bar{a} \sharp_{\sigma} \bar{x}} \\
(K \bar{a}, \bar{e}, f_{K}(\bar{a}, \bar{x})_{\sigma}) \in F_{\mathsf{s}}} \tag{129}$$

where $F^{|\sigma|} \subseteq \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|} \times X^{|\sigma|}$ is defined from any family $(F_{\mathsf{s}} \subseteq \mathsf{T}_{\alpha}(\Sigma)_{\mathsf{s}} \times X_{\mathsf{s}} \mid \mathsf{s} \in \Sigma_{\mathrm{D}})$ by ordinary recursion on the structure of arities as follows.

$$\begin{array}{c|c|c} \sigma & F^{|\sigma|} \subseteq \mathsf{T}_{\!\alpha}(\Sigma)^{|\sigma|} \times X^{|\sigma|} \\ \hline \mathbf{a} \in \Sigma_{\mathbf{A}} & \{(a,a) \mid a \in \mathbb{A}_{\mathbf{a}}\} \\ \mathbf{s} \in \Sigma_{\mathbf{D}} & F_{\mathbf{s}} \\ \mathbf{1} & \{((),())\} \\ \sigma_1 * \sigma_2 & \{((\bar{e}_1,\bar{e}_2),(\bar{x}_1,\bar{x}_2)) \mid (\bar{e}_1,\bar{x}_1) \in F^{|\sigma_1|} \ \& \\ & (\bar{e}_2,\bar{x}_2) \in F^{|\sigma_2|}\} \\ \mathbf{a} \Rightarrow \sigma_1 & F^{|\sigma_1|} & F^{|\sigma_1|} \end{array}$$

It is not hard to see that the set of rules determined by (129) is supported by A. Hence by Theorem 12, so are all the subsets F_s (and the subsets $F^{|\sigma|}$ defined from them). For the existence part of Theorem 22 it suffices to show that each F_s is the graph of a function \hat{f}_s from $T_{\alpha}(\Sigma)_s$ to X_s . For then:

– each \hat{f}_s and the functions $\hat{f}^{|\sigma|}$ defined from them as in Figure 2 are supported by A, because F_s is;

- the graphs of the functions $\hat{f}^{|\sigma|}$ are the relations $F^{|\sigma|}$, because of the above definition of $F^{|\sigma|}$;
- if $\bar{a} \# A$ and $\bar{a} \sharp_{\sigma} \bar{e}$, then $\bar{a} \sharp_{\sigma} \hat{f}^{|\sigma|} \bar{e}$, because of the definition of $\hat{f}^{|\sigma|}$;
- hence $(\hat{f}_s \mid s \in \Sigma_D)$ satisfies the recursion property (46): for if $\bar{a} \# A$ and $\bar{a} \sharp_{\sigma} \bar{e}$, then $\bar{x} \triangleq \hat{f}^{|\sigma|}\bar{e}$ satisfies $(\bar{e},\bar{x}) \in F^{|\sigma|}$ and $\bar{a} \sharp_{\sigma} \bar{x}$; so $(K\bar{a}.\bar{e},f_K(\bar{a},\bar{x})_{\sigma}) \in F_s$ by rule (129) and hence $\hat{f}_s(K\bar{a}.\bar{e}) = f_K(\bar{a},\bar{x})_{\sigma} = f_K(\bar{a},\hat{f}^{|\sigma|}\bar{e})_{\sigma}$.

To prove that each F_s is the graph of a function, i.e. that

$$S_{\mathsf{s}} \triangleq \{e \in \mathsf{T}_{\alpha}(\Sigma)_{\mathsf{s}} \mid (\exists ! \, x \in X_{\mathsf{s}}) \, (e, x) \in F_{\mathsf{s}} \}$$

is the whole of $\mathsf{T}_\alpha(\Sigma)_{\mathsf{s}}$, we apply Theorem 23. Note that from the way it is defined, each S_{s} is supported by A, because F_{s} is. So we just have to prove that $(S_{\mathsf{s}} \mid \mathsf{s} \in \Sigma_{\mathsf{D}})$ satisfies (IH_K) for each $(K:\sigma\to\mathsf{s})\in\Sigma_{\mathsf{C}}$. By Lemma 35(i), there is some $\bar{a}\in\mathbb{A}^\sigma$ with $\bar{a}\#A$. We prove (IH_K) for this \bar{a} by showing that for each $\bar{e}\in S^{|\sigma|}$ with $\bar{a}\sharp_\sigma\bar{e}$ it is the case that $K\bar{a}.\bar{e}\in S_{\mathsf{s}}$.

First note that by the definitions of $S^{|\sigma|}$ from $(S_s \mid s \in \Sigma_D)$ (in Figure 3) and $F^{|\sigma|}$ from $(F_s \mid s \in \Sigma_D)$, we have

$$(\forall \bar{e} \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}) \; \bar{e} \in S^{|\sigma|} \; \Rightarrow \; (\exists ! \, \bar{x} \in X^{|\sigma|}) \; (\bar{e}, \bar{x}) \in F^{|\sigma|} \; . \tag{130}$$

In particular it follows that if $\bar{e} \in S^{|\sigma|}$ and $(\bar{e}, \bar{x}) \in F^{|\sigma|}$, then $supp(\bar{x}) \subseteq A \cup supp(\bar{e})$ (since A supports $F^{|\sigma|}$); and this in turn implies (by induction on the structure of arities σ) that

$$(\forall \bar{a} \in \mathbb{A}^{\sigma}, \bar{e} \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|}, \bar{x} \in X^{|\sigma|}) \ \bar{a} \ \sharp_{\sigma} \ \bar{e} \ \& \ \bar{e} \in S^{|\sigma|} \ \& \ (\bar{e}, \bar{x}) \in F^{|\sigma|} \Rightarrow \bar{a} \ \sharp_{\sigma} \ \bar{x} \ . \tag{131}$$

So by rule (129), if $\bar{e} \in S^{|\sigma|}$ with $\bar{a} \sharp_{\sigma} \bar{e}$, then there is some $\bar{x} \in X^{|\sigma|}$ with $(K \bar{a}. \bar{e}, f_K(\bar{e}, \bar{x})_{\sigma}) \in F_s$. Thus to see that $K \bar{a}. \bar{e} \in S_s$, it just remains to show that if $(K \bar{a}. \bar{e}, x) \in F_s$, then $x = f_K(\bar{e}, \bar{x})_{\sigma}$. But if $(K \bar{a}. \bar{e}, x) \in F_s$ holds, it must have been deduced by an application of rule (129) to

$$\bar{a}' \# A \& \bar{a}' \sharp_{\sigma} \bar{e}' \& (\bar{e}', \bar{x}') \in F^{|\sigma|} \& \bar{a}' \sharp_{\sigma} \bar{x}'$$

with $K \bar{a}. \bar{e} = K \bar{a}'. \bar{e}'$ and $x = f_K(\bar{a}', \bar{x}')_{\sigma}$. So $\bar{a}. \bar{e} = \bar{a}'. \bar{e}'$ and by Lemma 36, $\tau_{\bar{a},\bar{a}''}.\bar{e} = \tau_{\bar{a}',\bar{a}''}.\bar{e}'$, for some $\bar{a}'' \# (A,\bar{a},\bar{e},\bar{a}',\bar{e}')$. By (121), the atom-permutations $\tau_{\bar{a},\bar{a}''}$ and $\tau_{\bar{a}',\bar{a}''}$ have their support disjoint from A and hence from the support of $F^{[\sigma]}$ and $S^{[\sigma]}$. Therefore

$$\begin{split} (\tau_{\bar{a},\bar{a}''} \cdot \bar{e},\tau_{\bar{a},\bar{a}''} \cdot \bar{x}) &= \tau_{\bar{a},\bar{a}''} \cdot (\bar{e},\bar{x}) \in \tau_{\bar{a},\bar{a}''} \cdot F^{|\sigma|} = F^{|\sigma|} \;, \\ (\tau_{\bar{a}',\bar{a}''} \cdot \bar{e}',\tau_{\bar{a}',\bar{a}''} \cdot \bar{x}') &= \tau_{\bar{a}',\bar{a}''} \cdot (\bar{e}',\bar{x}') \in \tau_{\bar{a}',\bar{a}''} \cdot F^{|\sigma|} = F^{|\sigma|} \;, \\ \tau_{\bar{a}',\bar{a}''} \cdot \bar{e}' &= \tau_{\bar{a},\bar{a}''} \cdot \bar{e} \in \tau_{\bar{a},\bar{a}''} \cdot S^{|\sigma|} = S^{|\sigma|} \end{split}$$

and hence by (130), $\tau_{\bar{a},\bar{a}''}\cdot\bar{x}=\tau_{\bar{a}',\bar{a}''}\cdot\bar{x}'$. Note that from (FCB_K), by Lemma 35(ii) we have

$$(\forall \bar{a} \in \mathbb{A}^{\sigma}) \ \bar{a} \# A \Rightarrow (\forall \bar{x} \in X^{|\sigma|}) \ \bar{a} \sharp_{\sigma} \bar{x} \Rightarrow \bar{a} \# f_{K}(\bar{a}, \bar{x})_{\sigma} . \tag{132}$$

So we have $\bar{a} \# f_K(\bar{a},\bar{x})_{\sigma}$ and $\bar{a}' \# f_K(\bar{a}',\bar{x}')_{\sigma}$; and by choice of \bar{a}'' we also have $\bar{a}'' \# (f_K(\bar{a},\bar{x})_{\sigma},f_K(\bar{a}',\bar{x}')_{\sigma})$. Since the atom-permutations $\tau_{\bar{a},\bar{a}''}$ and $\tau_{\bar{a}',\bar{a}''}$ have their supports in A and hence disjoint from f_K , $f_K(\bar{a},\bar{x})_{\sigma}$ and $f_K(\bar{a}',\bar{x}')_{\sigma}$, we have

$$f_K(\bar{a}, \bar{x})_{\sigma} = \tau_{\bar{a}, \bar{a}''} \cdot f_K(\bar{a}, \bar{x})_{\sigma} = f_K(\tau_{\bar{a}, \bar{a}''} \cdot \bar{a}, \tau_{\bar{a}, \bar{a}''} \cdot \bar{x})_{\sigma}$$

$$f_K(\bar{a}', \bar{x}')_{\sigma} = \tau_{\bar{a}', \bar{a}''} \cdot f_K(\bar{a}', \bar{x}')_{\sigma} = f_K(\tau_{\bar{a}', \bar{a}''} \cdot \bar{a}', \tau_{\bar{a}', \bar{a}''} \cdot \bar{x}')_{\sigma}.$$

From above $\tau_{\bar{a},\bar{a}''} \cdot \bar{e} = \tau_{\bar{a}',\bar{a}''} \cdot \bar{e}'$; and by (124), $\tau_{\bar{a},\bar{a}''} \cdot \bar{a} = \bar{a}'' = \tau_{\bar{a}',\bar{a}''} \cdot \bar{a}'$. So $f_K(\bar{a},\bar{x})_{\sigma} = f_K(\bar{a}',\bar{x}')_{\sigma} = x$, as required.

Uniqueness of the functions $\hat{f}_s \in T_\alpha(\Sigma)_s \to_{fs} X_s$

Suppose $(\hat{f}'_{s} \in \mathsf{T}_{\alpha}(\Sigma)_{s} \to_{\mathrm{fs}} X_{s} \mid s \in \Sigma_{\mathrm{D}})$ is also supported by A and satisfies property (46). To see that $\hat{f}'_{s} = \hat{f}_{s}$ if suffices to show that $S_{s} \triangleq \{e \in \mathsf{T}_{\alpha}(\Sigma)_{s} \mid \hat{f}_{s} e = \hat{f}'_{s} e\}$ is the whole of $\mathsf{T}_{\alpha}(\Sigma)$ for each $s \in \Sigma_{\mathrm{D}}$. This follows almost immediately by Theorem 23; one just has to check that the subsets $S^{|\sigma|}$ defined from this $(S_{s} \mid s \in \Sigma_{\mathrm{D}})$ satisfy $S^{|\sigma|} \subseteq \{\bar{e} \in \mathsf{T}_{\alpha}(\Sigma)^{|\sigma|} \mid \hat{f}^{|\sigma|} \bar{e} = \hat{f}'^{|\sigma|} \bar{e}\}$.

References

- [1] B. E. Aydemir, A Bohannon, M. Fairbairn, J. N. Foster, B. C. Pierce, P. E. Sewell, D. Vytiniotis, G. Washburn, S. Weirich, and S. Zdancewic. Mechanised metatheory for the masses: The POPLMARK challenge. In J. Hurd and T. Melham, editors, 18th International Conference on Theorem Proving in Higher Order Logics: TPHOLs 2005, volume 3603 of Lecture Notes in Computer Science, pages ?—? Springer-Verlag, 2005. www.cis.upenn.edu/group/proj/plclub/mmm/.
- [2] H. P. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. North-Holland, revised edition, 1984.
- [3] U. Berger and H. Schwichtenberg. An inverse of the evaluation functional for typed λ -calculus. In *6th Annual Symposium on Logic in Computer Science*, pages 203–211. IEEE Computer Society Press, Washington, 1991.
- [4] J. Cheney. *Nominal Logic Programming*. PhD thesis, Cornell University, August 2004.
- [5] A. Church. A formulation of the simple theory of types. *Journal of Symbolic Logic*, 5:56–68, 1940.

[6] N. G. de Bruijn. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indag. Math.*, 34:381–392, 1972.

- [7] P. Dybjer. A general formulation of simultaneous inductive-recursive definitions in type theory. *Journal of Symbolic Logic*, 65(2), 2000.
- [8] P. Dybjer and A. Filinski. Normalization and partial evaluation. In G. Barthe, P. Dybjer, and J. Saraiva, editors, *Applied Semantics*, *Advanced Lectures*, volume 2395 of *Lecture Notes in Computer Science, Tutorial*, pages 137–192. Springer-Verlag, 2002. International Summer School, APPSEM 2000, Caminha, Portugal, September 9–15, 2000.
- [9] M. P. Fiore. Semantic analysis of normalisation by evaluation for typed lambda calculus. In *Proceedings of the 4th iNternational ACM SIGPLAN* Conference on Principles and Practice of Declarative Programming, October 6-8, (PPDP 2002), Pittsburgh, PA, USA, pages 26–37. ACM Press, October 2002.
- [10] M. P. Fiore, G. D. Plotkin, and D. Turi. Abstract syntax and variable binding. In 14th Annual Symposium on Logic in Computer Science, pages 193–202. IEEE Computer Society Press, Washington, 1999.
- [11] M. J. Gabbay. A Theory of Inductive Definitions with α-Equivalence: Semantics, Implementation, Programming Language. PhD thesis, University of Cambridge, 2000.
- [12] M. J. Gabbay. FM-HOL, a higher-order theory of names. In F. Kamareddine, editor, *Workshop on Thirty Five years of Automath, Informal Proceedings*. Heriot-Watt University, Edinburgh, Scotland, April 2002.
- [13] M. J. Gabbay and A. M. Pitts. A new approach to abstract syntax with variable binding. *Formal Aspects of Computing*, 13:341–363, 2002.
- [14] A. D. Gordon and T. Melham. Five axioms of alpha-conversion. In *Theorem Proving in Higher Order Logics*, 9th International Conference, volume 1125 of Lecture Notes in Computer Science, pages 173–191. Springer-Verlag, 1996.
- [15] T. G. Griffin. Notational definition a formal account. In *3rd Annual Symposium on Logic in Computer Science*, pages 372–383. IEEE Computer Society Press, 1988.
- [16] C. A. Gunter. *Semantics of Programming Languages: Structures and Techniques*. Foundations of Computing. MIT Press, 1992.
- [17] F. Honsell, M. Miculan, and I. Scagnetto. An axiomatic approach to metareasoning on nominal algebras in HOAS. In F. Orejas, P. G. Spirakis, and

- J. van Leeuwen, editors, 28th International Colloquium on Automata, Languages and Programming, ICALP 2001, Crete, Greece, July 2001. Proceedings, volume 2076 of Lecture Notes in Computer Science, pages 963–978. Springer-Verlag, Heidelberg, 2001.
- [18] M. Norrish. Recursive function definition for types with binders. In *Theorem Proving in Higher Order Logics*, 17th International Conference, volume 3223 of Lecture Notes in Computer Science, pages 241–256. Springer-Verlag, 2004.
- [19] F. Pfenning and C. Elliott. Higher-order abstract syntax. In *Proc. ACM-SIGPLAN Conference on Programming Language Design and Implementation*, pages 199–208. ACM Press, 1988.
- [20] A. M. Pitts. Nominal logic, a first order theory of names and binding. *Information and Computation*, 186:165–193, 2003.
- [21] A. M. Pitts. Alpha-structural recursion and induction (extended abstract). In J. Hurd and T. Melham, editors, *Theorem Proving in Higher Order Logics*, 18th International Conference, TPHOLs 2005, Oxford UK, August 2005, Proceedings, volume 3603 of Lecture Notes in Computer Science, pages 17–34. Springer-Verlag, 2005.
- [22] G. D. Plotkin. An illative theory of relations. In R. Cooper, Mukai, and J. Perry, editors, *Situation Theory and its Applications, Volume 1*, volume 22 of *CSLI Lecture Notes*, pages 133–146. Stanford University, 1990.
- [23] G. D. Plotkin. A structural approach to operational semantics. *Journal of Logic and Algebraic Programming*, 60–61:17–139, 2004.
- [24] F. Pottier. An overview of C α ml. Submitted, June 2005.
- [25] D. Sangiorgi and D. Walker. *The* π -calculus: a Theory of Mobile Processes. Cambridge University Press, 2001.
- [26] M. R. Shinwell and A. M. Pitts. On a monadic semantics for freshness. *Theoretical Computer Science*, 2005. To appear.
- [27] A. Stoughton. Substitution revisited. *Theoretical Computer Science*, 59:317–325, 1988.
- [28] C. Urban, A. M. Pitts, and M. J. Gabbay. Nominal unification. *Theoretical Computer Science*, 323:473–497, 2004.
- [29] C. Urban and C. Tasson. Nominal techniques in Isabelle/HOL. In 20th International Conference on Automated Deduction, CADE-20, Tallinn, Estonia, July 2005, volume ???? of Lecture Notes in Computer Science, pages ?—? Springer-Verlag, 2005.