

Alpha-Structural Recursion and Induction

ANDREW M. PITTS

Cambridge University, Cambridge, UK

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 formalization 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 article explains the idea in as concrete a way as possible and gives a new derivation within higher-order classical logic of principles of α -structural recursion and induction for α -equivalence classes from the ordinary versions of these principles for abstract syntax trees.

Categories and Subject Descriptors: D.3.1 [Programming Languages]: Formal Definitions and Theory—Syntax; F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages—Algebraic approaches to semantics; operational semantics; denotational semantics; F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—Mechanical theorem proving

General Terms: Languages, Theory, Verification

Additional Key Words and Phrases: Abstract syntax, binders, induction, names, recursion

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 Plotkin [2004, 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

This research supported by UK EPSRC grants GR/R07615/1 and EP/D000459/1.

Author’s address: University of Cambridge Computer Laboratory, William Gates Building, 15 JJ Thomson Avenue, Cambridge CB3 0FD UK.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2006 ACM 0004-5411/06/0500-0459 \$5.00

algebraic signature. This consists of a collection of *sorts* \mathbf{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 $(\mathbf{s}_1, \dots, \mathbf{s}_n)$ of sorts and with a *result-sort* \mathbf{s} . Then, the ASTs over the signature can be described by inductively generated terms t : if K has arity $(\mathbf{s}_1, \dots, \mathbf{s}_n)$ and result sort \mathbf{s} , and if t_i is a term of sort \mathbf{s}_i for $i = 1..n$, then $K(t_1, \dots, t_n)$ is a term of sort \mathbf{s} . One gets off the ground in this inductive definition with the $n = 0$ instance of the rule for forming terms, which covers 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 nonstructural 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 indexes, as in de Bruijn [1972]. 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.¹ This approach is signalled by, 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” [Barendregt 1984, 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 article we examine ways, arising from the results of Gabbay and Pitts [2002] and Pitts [2003], 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

¹This includes the metatheory of “higher-order abstract syntax” [Pfenning and Elliott 1988], where the questions we are addressing are pushed up one meta-level to a single binding-form, λ -abstraction.

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, \dots), 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* $f x = t_1$ in t_2 binds all free occurrences of f in both t_1 and t_2 ; whereas the leftmost occurrence of the variable x only binds free occurrences of x in t_1 . A systematic way of specifying such patterns of binding and the associated notion of α -equivalence is given in Section 2.2; for the moment I assume the reader can supply a suitable definition of α -equivalence for λ -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 .

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

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

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

$$y \notin fv(t_2) \cup \{f\} \ \& \ f, y \notin fv(t) \cup \{x\} \Rightarrow \\ (x := t)\textit{letrec } f y = t_1 \textit{ in } t_2 = \textit{letrec } f y = (x := t)t_1 \textit{ in } (x := t)t_2. \quad (4)$$

The condition on Eq. (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 $\lambda y. t_1$ to an α -variant satisfying these conditions. The same goes for the more complicated condition on Eq. (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 y are distinct variables not occurring free in t and not equal to x ; less crucially, we can also assume that y does not occur free in t_2 , because that term lies outside the binding scope of y in the term *letrec* $f y = t_1$ in t_2 .

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 Λ/\equiv_α for its quotient by α -equivalence \equiv_α , then capture-avoiding substitution of an α -equivalence class e for a variable x is a function $\hat{s}_{x,e} \in \Lambda/\equiv_\alpha \rightarrow \Lambda/\equiv_\alpha$. Every such function corresponds to a function $s_{x,e} \in \Lambda \rightarrow \Lambda/\equiv_\alpha$ respecting \equiv_α , that is, satisfying

$$t_1 \equiv_\alpha t_2 \Rightarrow s_{x,e}(t_1) = s_{x,e}(t_2) \quad (5)$$

(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:

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

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

$$y \notin \text{fv}(e) \cup \{x\} \Rightarrow s_{x,e}(\lambda y.t_1) = [\lambda y.t'_1]_\alpha \quad \text{where } s_{x,e}(t_1) = [t'_1]_\alpha \quad (8)$$

$$y \notin \text{fv}(t_2) \cup \{f\} \ \& \ f, y \notin \text{fv}(e) \cup \{x\} \Rightarrow s_{x,e}(\text{letrec } f \ y = t_1 \ \text{in } t_2) = [\text{letrec } f \ y = t'_1 \ \text{in } t'_2]_\alpha \\ \text{where } s_{x,e}(t_i) = [t'_i]_\alpha \ \text{for } i = 1, 2. \quad (9)$$

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 [Stoughton 1988]). 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 $(x := t)(-) \in \Lambda / =_\alpha \rightarrow \Lambda / =_\alpha$. 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 generalizations of the ones introduced by Gabbay and Pitts [2002] as part of a new mathematical model of fresh names and name binding. That article 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 that article. 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 by Pitts [2003], and I will use them here to express α -structural recursion and induction within “ordinary mathematics”, or more precisely, within Church’s higher-order classical logic [Church 1940].

1.1. HOW TO READ THIS ARTICLE. Having read the Introduction this far, impatient readers may wish to turn to Theorem 5.4 to see the statement of the α -structural recursion principle for λ -calculus with *letrec*-terms and how it is used to define $(x := t)(-) \in \Lambda / =_\alpha \rightarrow \Lambda / =_\alpha$ 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* [Urban et al. 2004]) is explained in Section 2. Section 4 gives a first, simple version of the α -structural recursion and induction principles that is 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 *normalization by evaluation* for the simply-typed λ -calculus [Berger and Schwichtenberg 1991]) 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 article'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 [Fiore et al. 1999; Griffin 1988; Honsell et al. 2001; Plotkin 1990; Urban et al. 2004]. Here we will use the notion of *nominal signature* [Urban et al. 2004]. 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" [Aydemir et al. 2005]). 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 [2005] in his *C α ml* library; 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 article, we fix two sets: the set \mathbb{A} of all *atoms* and the set \mathbb{AS} of all *atom-sorts*. We also fix a function $sort \in \mathbb{A} \rightarrow \mathbb{AS}$ assigning sorts to atoms and assume that the sets \mathbb{AS} and $\mathbb{A}_a \triangleq \{a \in \mathbb{A} \mid sort(a) = \mathbf{a}\}$ for each $\mathbf{a} \in \mathbb{AS}$, 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* $\mathbf{s} \in \Sigma_D$, and we write $K : \sigma \rightarrow \mathbf{s}$ to indicate this information. The arities σ of Σ are given as follows:

Atom-sorts. Every atom-sort $\mathbf{a} \in \Sigma_A$ is an arity.

Data-sorts. Every data-sort $\mathbf{s} \in \Sigma_D$ is an arity.

Unit arity. $\mathbf{1}$ 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 $\mathbf{a} \in \Sigma_A$ and σ is an arity, then $\llbracket \mathbf{a} \rrbracket \sigma$ is an arity.

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

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

Constructed terms. If $K : \sigma \rightarrow \mathbf{s}$ is in Σ_C and $t : \sigma$, then $K t : \mathbf{s}$.

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

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

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

We write $Ar(\Sigma)$ for the set of all arities over a nominal signature Σ , $\mathbb{T}(\Sigma)$ for the set of all terms over Σ , and $ar \in \mathbb{T}(\Sigma) \rightarrow Ar(\Sigma)$ 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 $\mathbb{T}(\Sigma)_\sigma$ for the subset $\{t \in \mathbb{T}(\Sigma) \mid ar(t) = \sigma\}$ of terms of arity σ .

Example 2.1 (λ -calculus with letrec). Here is a nominal signature for the untyped λ -calculus [Barendregt 1984]. There is a single atom-sort \mathbf{v} for variables, and a single data-sort \mathbf{t} for λ -terms.

atom-sorts	data-sorts	constructors
\mathbf{v}	\mathbf{t}	$V : \mathbf{v} \rightarrow \mathbf{t}$
		$A : \mathbf{t} * \mathbf{t} \rightarrow \mathbf{t}$
		$L : \llbracket \mathbf{v} \rrbracket \mathbf{t} \rightarrow \mathbf{t}$

To illustrate the intermixing 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 t_1 or t_2 are bound in the term. To get the effect of this, we can add to the above nominal signature a constructor

$$Letrec : \llbracket \mathbf{v} \rrbracket (\llbracket \mathbf{v} \rrbracket \mathbf{t}) * \mathbf{t} \rightarrow \mathbf{t}.$$

² Compared with Urban et al. [2004, 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 $\text{letrec } f\ x = f\ x\ \text{in } f(\lambda y.y)$ corresponds to the nominal term

$$\text{Letrec}\langle\langle f \rangle\rangle\langle\langle x \rangle\rangle A(Vf, Vx), A(Vf, L\langle\langle y \rangle\rangle Vy))$$

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

Example 2.2 (π -calculus). Here is a nominal signature for the version of the Milner–Parrow–Walker π -calculus given by Sangiorgi and Walker [2001, 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 pre	$Gsum : \text{gsum} \rightarrow \text{proc}$ $Par : \text{proc} * \text{proc} \rightarrow \text{proc}$ $Res : \langle\langle \text{chan} \rangle\rangle \text{proc} \rightarrow \text{proc}$ $Rep : \text{proc} \rightarrow \text{proc}$ $Zero : 1 \rightarrow \text{gsum}$ $Pre : \text{pre} \rightarrow \text{gsum}$ $Plus : \text{gsum} * \text{gsum} \rightarrow \text{gsum}$ $Out : (\text{chan} * \text{chan}) * \text{proc} \rightarrow \text{pre}$ $In : \text{chan} * \langle\langle \text{chan} \rangle\rangle \text{proc} \rightarrow \text{pre}$ $Tau : \text{proc} \rightarrow \text{pre}$ $Match : (\text{chan} * \text{chan}) * \text{pre} \rightarrow \text{pre}$

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

$$\begin{aligned} &Res\langle\langle x \rangle\rangle Par\langle GsumPlus\langle PreOut\langle\langle x, u \rangle\rangle, GsumZero\langle\rangle\rangle, \\ &\quad PreOut\langle\langle y, v \rangle\rangle, GsumZero\langle\rangle\rangle\rangle, \\ &GsumPreIn\langle x, \langle\langle z \rangle\rangle GsumPreOut\langle\langle z, w \rangle\rangle, GsumZero\langle\rangle\rangle\rangle. \end{aligned}$$

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 $\mathbb{T}(\Sigma)$ 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 principles that we use later. We regard their proofs as standard and omit them.³

THEOREM 2.3 (STRUCTURAL RECURSION FOR NOMINAL TERMS). *Let Σ be a nominal signature. Suppose we are given sets S_σ , for each $\sigma \in \text{Ar}(\Sigma)$, and*

³ Each 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.

elements

$$\begin{array}{ll}
g_a \in \mathbb{A}_a \rightarrow S_a & (a \in \Sigma_A) \\
g_K \in S_\sigma \rightarrow S_s & ((K : \sigma \rightarrow \mathbf{s}) \in \Sigma_C) \\
g_1 \in S_1 & \\
g_{\sigma_1 * \sigma_2} \in S_{\sigma_1} \times S_{\sigma_2} \rightarrow S_{\sigma_1 * \sigma_2} & (\sigma_1, \sigma_2 \in \text{Ar}(\Sigma)) \\
g_{\ll a \gg \sigma} \in \mathbb{A}_a \times S_\sigma \rightarrow S_{\ll a \gg \sigma} & (a \in \Sigma_A, \sigma \in \text{Ar}(\Sigma))
\end{array}$$

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

$$\hat{g} a = g_a(a) \quad (10)$$

$$\hat{g}(K t) = g_K(\hat{g} t) \quad (11)$$

$$\hat{g} \langle \rangle = g_1 \quad (12)$$

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

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

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

THEOREM 2.4 (STRUCTURAL INDUCTION FOR NOMINAL TERMS). *Let Σ be a nominal signature and $S \subseteq \mathbb{T}(\Sigma)$ a set of terms over Σ . To prove that S is the whole of $\mathbb{T}(\Sigma)$, it suffices to show*

$$(\forall a \in \Sigma_A, a \in \mathbb{A}_a) a \in S \quad (15)$$

$$(\forall (K : \sigma \rightarrow \mathbf{s}) \in \Sigma_C, t : \sigma) t \in S \Rightarrow K t \in S \quad (16)$$

$$\langle \rangle \in S \quad (17)$$

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

$$(\forall a \in \Sigma_A, a \in \mathbb{A}_a, \sigma \in \text{Ar}(\Sigma), t : \sigma) t \in S \Rightarrow \ll a \gg t \in S. \quad (19)$$

2.4. α -EQUIVALENCE AND α -TERMS. So far, we have taken no account of the fact that atom-binding terms $\ll a \gg t$ 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 $\sigma \in \text{Ar}(\Sigma)$ and $t, t' \in \mathbb{T}(\Sigma)_\sigma$) 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 [Gunter 1992, p. 36] that is conveniently syntax-directed compared with the classic version [Barendregt 1984, 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 straightforward is the fact that $=_\alpha$ is transitive. This can be proved in a number of ways. My favorite way makes good use of the techniques we will be using later, based on the action of atom-permutations on terms; see Pitts [2003, Example 1].

Definition 2.5. For each $\sigma \in \text{Ar}(\Sigma)$, we write $\mathbb{T}_\alpha(\Sigma)_\sigma$ for the quotient of $\mathbb{T}(\Sigma)_\sigma$ by the equivalence relation $(-) =_\alpha (-) : \sigma$. Thus, the elements of $\mathbb{T}_\alpha(\Sigma)_\sigma$ 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 .

$$\begin{array}{l}
(=\alpha-1) \frac{\mathbf{a} \in \Sigma_A \quad a \in \mathbb{A}_a}{a =_\alpha a : \mathbf{a}} \qquad (=\alpha-2) \frac{(K : \sigma \rightarrow \mathbf{s}) \in \Sigma_C \quad t =_\alpha t' : \sigma}{K t =_\alpha K t' : \mathbf{s}} \\
(=\alpha-3) \frac{}{\langle \rangle =_\alpha \langle \rangle : \mathbf{1}} \qquad (=\alpha-4) \frac{t_1 =_\alpha t'_1 : \sigma_1 \quad t_2 =_\alpha t'_2 : \sigma_2}{\langle t_1, t_2 \rangle =_\alpha \langle t'_1, t'_2 \rangle : \sigma_1 * \sigma_2} \\
(=\alpha-5) \frac{\mathbf{a} \in \Sigma_A \quad a, a', a'' \in \mathbb{A}_a \quad a'' \notin atm(\langle a, t, a', t' \rangle) \quad t\{a''/a\} =_\alpha t'\{a''/a'\} : \sigma}{\langle a \rangle t =_\alpha \langle a' \rangle t' : \langle \mathbf{a} \rangle \sigma}
\end{array}$$

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 be defined by structural recursion (Theorem 2.3):

$$\begin{array}{ll}
atm(a) = \{a\} & a''\{a''/a\} = \begin{cases} a' & \text{if } a'' = a \\ a'' & \text{if } a'' \neq a \end{cases} \\
atm(K t) = atm(t) & (K t)\{a''/a\} = K(t\{a''/a\}) \\
atm(\langle \rangle) = \emptyset & \langle \rangle\{a''/a\} = \langle \rangle \\
atm(\langle t_1, t_2 \rangle) = atm(t_1) \cup atm(t_2) & \langle t_1, t_2 \rangle\{a''/a\} = \langle t_1\{a''/a\}, t_2\{a''/a\} \rangle \\
atm(\langle a \rangle t) = \{a\} \cup atm(t) & \langle a \rangle t\{a''/a\} = \langle a'' \rangle t\{a''/a\}
\end{array}$$

FIG. 1. α -Equivalence of nominal terms.

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 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 finite⁵ and $sort(\pi(a)) = sort(a)$ for all $a \in \mathbb{A}$. The operation of composing bijections gives a binary operation $\pi, \pi' \in Perm \mapsto \pi \circ \pi' \in Perm$

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

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.

⁴ Both Gabbay [2006] and Cheney [2004] 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.

⁵ A very similar theory of nominal sets can be developed without this restriction to finite permutations; but the restriction does ensure that $Perm$ is itself a nominal set (see Example 3.3)—a fact that we exploit in the proof of the α -structural recursion theorem.

An *action* of $Perm$ on a set X is a function $Perm \times X \rightarrow X$, 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 3.1. 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 atom a'' of the same sort as a and a' not in the finite set $A_1 \cup A_2 \cup \{a, a'\}$, 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 3.2

- (1) 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 .
- (2) Let Σ be a nominal signature. Using Theorem 2.3, we can define an atom-permutation action on the sets $\mathbb{T}(\Sigma)_\sigma$ of terms over Σ of each arity $\sigma \in Ar(\Sigma)$:

$$\begin{aligned} \pi \cdot a &\triangleq \pi(a) \\ \pi \cdot Kt &\triangleq K(\pi \cdot t) \\ \pi \cdot \langle \rangle &\triangleq \langle \rangle \\ \pi \cdot \langle t_1, t_2 \rangle &\triangleq \langle \pi \cdot t_1, \pi \cdot t_2 \rangle \\ \pi \cdot \langle\langle a \rangle\rangle t &\triangleq \langle\langle \pi \cdot a \rangle\rangle (\pi \cdot t). \end{aligned}$$

Using Theorem 2.4, 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 $\mathbb{T}(\Sigma)_\sigma$ is a nominal set, with $supp(t)$ equal to the finite set $atm(t)$ of atoms occurring in t .

- (3) 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 3.6) 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 $\mathbb{T}_\alpha(\Sigma)_\sigma$

is a nominal set with $\text{supp}([t]_\alpha)$ equal to the finite set $\text{fa}(t)$ of *free atoms* of any representative t of the class $[t]_\alpha$, defined (using Theorem 2.3) by:

$$\begin{aligned} \text{fa}(a) &\triangleq \{a\} \\ \text{fa}(Kt) &\triangleq \text{fa}(t) \\ \text{fa}(\langle \rangle) &\triangleq \emptyset \\ \text{fa}(\langle t_1, t_2 \rangle) &\triangleq \text{fa}(t_1) \cup \text{fa}(t_2) \\ \text{fa}(\langle\langle a \rangle\rangle t) &\triangleq \text{fa}(t - \{a\}). \end{aligned}$$

- (4) 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 $\pi \in \text{Perm}$ and $s \in S$; in this case the support of each element is empty. In particular, we will regard the one-element set $1 = \{\langle \rangle\}$, the set of Booleans $\mathbb{B} = \{\text{true}, \text{false}\}$ and the set of natural numbers $\mathbb{N} = \{0, 1, 2, \dots\}$ 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

$$\text{supp}((x_1, x_2)) = \text{supp}(x_1) \cup \text{supp}(x_2). \quad (20)$$

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

Turning next to functions, if X and Y are nominal sets, then we get an action of atom-permutations on the set $X \rightarrow Y$ 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:

$$\pi \cdot (f(x)) = (\pi \cdot f)(\pi \cdot x). \quad (21)$$

More precisely, the definition of the action on functions is forced by the requirement that $X \rightarrow Y$ 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 $f \in X \rightarrow Y$ is necessarily finitely supported with respect to this action (see Example 3.4 below). However, if f is supported by a finite set of atoms A , then $\pi \cdot f$ is supported by $\{\pi(a) \mid a \in A\}$. (This follows as in the proof of property (25) below.) Therefore, the set

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

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 $X \rightarrow \mathbb{B}$ (where $\mathbb{B} = \{\text{true}, \text{false}\}$) to transfer the action of

atom-permutations on $X \rightarrow \mathbb{B}$ 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 3.2(4)), 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

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

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 3.3 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 $\mathbb{A} \rightarrow \mathbb{A}$. Regarding \mathbb{A} as a nominal set, as in Example 3.2(1), 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 $\mathbb{A} \rightarrow \mathbb{A}$ 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 3.4 Not every function between nominal sets is finitely supported. For example, since the set \mathbb{A} of atoms is countable, there are surjective functions from \mathbb{N} to \mathbb{A} ; but it is not hard to see that any $f \in \mathbb{N} \rightarrow_{fs} \mathbb{A}$ must have a finite image (which is in fact the support of f). A more subtle example of a non-finitely-supported function is any *choice function*⁶ for the set \mathbb{A} of atoms, that is, any function $choose \in (\mathbb{A} \rightarrow_{fs} \mathbb{B}) \rightarrow \mathbb{A}$ (where $\mathbb{B} = \{true, false\}$) satisfying $f(a) = true \Rightarrow f(choose(f)) = true$, for all $f \in \mathbb{A} \rightarrow_{fs} \mathbb{B}$ 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 $f \in \mathbb{A} \rightarrow \mathbb{B}$ 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, $f \in \mathbb{A} \rightarrow_{fs} \mathbb{B}$ and we can apply $choose$ to obtain an atom $a_0 \triangleq choose(f)$. Let $\mathfrak{a} = sort(a_0)$. Since $\mathbb{A}_{\mathfrak{a}}$ is infinite and $A \cup \{a_0\}$ is finite, there is some atom $a_1 \in \mathbb{A}_{\mathfrak{a}}$ 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.

For nominal sets X and Y , the operation of function application

$$app_{X,Y}(f, x) \triangleq f x$$

⁶It 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 (see Jech [1977]).

is an element of $(X \rightarrow_{\text{fs}} Y) \times X \rightarrow_{\text{fs}} Y$; indeed, it follows from (21) that $\text{app}_{X,Y}$ is supported by the empty set of atoms. Similarly, it is not hard to see that currying,

$$\text{cur}_{X,Y,Z}(f) \triangleq \lambda x \in X. \lambda y \in Y. f(x, y)$$

determines an element $\text{cur}_{X,Y,Z} \in ((X \times Y \rightarrow_{\text{fs}} Y) \rightarrow_{\text{fs}} (X \rightarrow_{\text{fs}} (Y \rightarrow_{\text{fs}} Z)))$ with empty support. The constantly true function and the equality function

$$\begin{aligned} \text{true}_X(x) &\triangleq \text{true} \\ \text{eq}_X(x, x') &\triangleq \text{if } x = x', \text{ then true else false} \end{aligned}$$

also determine elements $\text{true}_X \in (X \rightarrow_{\text{fs}} \mathbb{B})$ and $\text{eq}_X \in (X \times X \rightarrow_{\text{fs}} \mathbb{B})$ with empty support. It is for these reasons that the following general principle holds good.

THEOREM 3.5 (FINITE SUPPORT PRINCIPLE). *Any function or relation that is defined from finitely supported functions and relations using higher-order, classical logic without choice principles, 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, classical mathematics.⁷ If we remain within pure higher-order, classical 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 \mathbb{A} of atoms, a constant for the function $\text{sort} \in \mathbb{A} \rightarrow \mathbb{A}\mathbb{S}$ (taking $\mathbb{A}\mathbb{S}$ 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 that may have nonempty, finite support. Such a “higher-order logic with atoms” has been developed by Gabbay [2002]. In this article, we stick with ordinary higher-order, classical 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 $P_{\text{fs}}(P_{\text{fs}}(X) \times X)$. As usual, a subset $S \subseteq X$ is closed under the rules in R if

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

and the smallest such subset, $\text{ind}(R)$, is given by the intersection of all such closed subsets. If we worked systematically in “FM-HOL” [Gabbay 2002], rather than using arbitrary subsets of X , we would only consider *finitely supported* subsets that are closed under the rules, and would replace $\text{ind}(R)$ by $\bigcap \{S \in P_{\text{fs}}(X) \mid (\forall (H, c) \in R) H \subseteq S \Rightarrow c \in S\}$. However, these two subsets coincide, as the following theorem shows.

THEOREM 3.6 (FINITELY SUPPORTED INDUCTIVE DEFINITIONS). *Let X be a nominal set. For any set of rules $R \in P_{\text{fs}}(P_{\text{fs}}(X) \times X)$, the subset $\text{ind}(R) \subseteq X$ inductively defined by R is a finitely supported subset of X ; indeed, $\text{supp}(\text{ind}(R)) \subseteq \text{supp}(R)$.*

⁷The only exception being that the finite support property is not conserved by all uses of choice: see Example 3.4. The extent to which a useful theory of finite support can be developed in *constructive* rather than classical mathematics is a very interesting question that is not pursued here.

PROOF. Suppose a, a' are atoms of the same sort that are not in the support of R . We have to show that $(a a') \cdot \text{ind}(R) = \text{ind}(R)$. It suffices to just show $\text{ind}(R) \subseteq (a a') \cdot \text{ind}(R)$. (For applying $(a a')$ to both sides, then gives the reverse inclusion.) For this, it suffices to show that $(a a') \cdot \text{ind}(R)$ is closed under the rules in R , since $\text{ind}(R)$ 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 $a, a' \notin \text{supp}(R)$. So if $H \subseteq (a a') \cdot \text{ind}(R)$, then $(a a') \cdot H \subseteq (a a') \cdot (a a') \cdot \text{ind}(R) = \text{ind}(R)$, so $(a a') \cdot c \in \text{ind}(R)$ since $\text{ind}(R)$ is closed under the rule $((a a') \cdot H, (a a') \cdot c)$; and hence $c = (a a') \cdot (a a') \cdot c \in (a a') \cdot \text{ind}(R)$. \square

In this article we will confine ourselves to finitely supported *finitary* rules, that is, 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, $\{x_1, \dots, x_n\}$ is supported by $\text{supp}(x_1) \cup \dots \cup \text{supp}(x_n)$. 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 $\text{ind}(R)$ 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 $P_{\text{fin}}(\mathbb{T}(\Sigma) \times \mathbb{T}(\Sigma)) \times (\mathbb{T}(\Sigma) \times \mathbb{T}(\Sigma))$ 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 $\pi \in \text{Perm}$ (since each π is a composition of transpositions):

$$t =_\alpha t' : \sigma \Rightarrow \pi \cdot t =_\alpha \pi \cdot t' : \sigma. \quad (22)$$

We used this property in Example 3.2(3) when discussing the nominal set structure of $\mathbb{T}_\alpha(\Sigma)$.

3.3. NOMINAL SUBSETS AND QUOTIENTS. If X is a nominal set and $S \in P_{\text{fs}}(X)$ is a finitely supported subset of it, then S is not necessarily itself a nominal set, because for any $x \in S$ and $\pi \in \text{Perm}$ we have no guarantee that $\pi \cdot x$ again lies in S . But if $\text{supp}(S) = \emptyset$, then $(a a') \cdot S = S$ for all atoms a, a' of the same sort; and since each $\pi \in \text{Perm}$ 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 $\text{supp}(S) = \emptyset$ is equivalent to

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

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 atom-permutation 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 $\mathbb{T}_\alpha(\Sigma)_\sigma$ of α -terms (of arity σ over a nominal signature Σ) from the nominal set $\mathbb{T}(\Sigma)_\sigma$ is an example of

this construction, since we saw in (22) that α -equivalence is a nominal subset of $\mathbb{T}(\Sigma)_\sigma \times \mathbb{T}(\Sigma)_\sigma$.

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 $\text{supp}_X(x) \cap \text{supp}_Y(y) = \emptyset$ and say that x is *fresh* for y .

LEMMA 3.7. *Let X, Y and Z be nominal sets. For any $x \in X, y \in Y, f \in Y \rightarrow_{\text{fs}} Z, \pi \in \text{Perm}$ and atoms $a, a' \in \mathbb{A}$ of the same sort,*

$$\pi \cdot ((a \ a') \cdot x) = (\pi(a) \ \pi(a')) \cdot (\pi \cdot x) \quad (24)$$

$$x \# y \Rightarrow \pi \cdot x \# \pi \cdot y. \quad (25)$$

$$x \# f \ \& \ x \# y \Rightarrow x \# (f \ y). \quad (26)$$

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 $\pi \cdot \text{supp}(x) = \text{supp}(\pi \cdot x)$. 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 $(a \ a') \cdot (\pi \cdot x) = \pi \cdot x$. Thus, $\pi \cdot A$ supports $\pi \cdot x$ when A supports x . So

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

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

Finally, for property (26), note that for all atoms a, a' of the same sort, if $a, a' \notin \text{supp}(f) \cup \text{supp}(y)$ then $(a \ a') \cdot f = f$ and $(a \ a') \cdot y = y$; so by (21) $(a \ a') \cdot f \ y = ((a \ a') \cdot f)((a \ a') \cdot y) = f \ y$. Thus, $\text{supp}(f) \cup \text{supp}(y)$ supports $f \ y$ and hence $\text{supp}(f \ y)$ is contained in this finite set. Therefore, if $\text{supp}(x)$ is disjoint from both $\text{supp}(f)$ and $\text{supp}(y)$, then it is also disjoint from $\text{supp}(f \ y)$. \square

Recall that the set of atoms \mathbb{A} is a nominal set as in Example 3.2(1). The following simple property of finitely supported sets of atoms is extremely useful when dealing with properties of fresh atoms; it subsumes Gabbay and Pitts [2002, Proposition 4.10] and Pitts [2003, Proposition 4].

THEOREM 3.8 (SOME/ANY THEOREM). *Let $S \in P_{\text{fs}}(\mathbb{A})$ be a set of atoms supported by some finite set of atoms A . For each atom-sort $\mathbf{a} \in \mathbb{A}\mathbb{S}$, the following are equivalent:*

$$(\forall a \in \mathbb{A}_{\mathbf{a}}) a \notin A \Rightarrow a \in S \quad (27)$$

$$(\exists a \in \mathbb{A}_{\mathbf{a}}) a \notin A \ \& \ a \in S. \quad (28)$$

PROOF. Since $\mathbb{A}_{\mathbf{a}} - A$ is infinite, it is in particular nonempty; thus, (27) implies (28). Conversely, suppose $a \in \mathbb{A}_{\mathbf{a}} - A$ satisfies $a \in S$. Given any other $a' \in \mathbb{A}_{\mathbf{a}} - A$, we have to show $a' \in S$; but $(a \ a') \cdot S = S$ (since $a, a' \notin A \supseteq \text{supp}(S)$) and hence $a' = (a \ a') \cdot a \in (a \ a') \cdot S = S$. \square

Example 3.9 Recall from Example 3.2(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 $\mathit{atm}(t)$ of atoms that occur in t . Furthermore, if $a' \notin \mathit{atm}(t)$, 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:

$$\frac{\mathbf{a} \in \Sigma_A \quad a, a', a'' \in \mathbb{A}_a \quad a'' \# (a, t, a', t') \quad (a''\ a) \cdot t =_\alpha (a''\ a') \cdot t' : \sigma}{\llbracket a \rrbracket t =_\alpha \llbracket a' \rrbracket t' : \llbracket \mathbf{a} \rrbracket \sigma}$$

In particular, we have

$$\begin{aligned} \llbracket a \rrbracket t =_\alpha \llbracket a' \rrbracket t' : \llbracket \mathbf{a} \rrbracket \sigma &\Leftrightarrow \\ (\exists a'' \in \mathbb{A}_a) a'' \# (a, t, a', t') &\ \& \ (a\ a'') \cdot t =_\alpha (a'\ a'') \cdot t' : \sigma. \end{aligned} \quad (29)$$

Applying Theorem 3.8 to the set of atoms $S = \{a'' \in \mathbb{A}_a \mid (a\ a'') \cdot t =_\alpha (a'\ a'') \cdot t' : \sigma\}$, which (by Lemma 3.7) is supported by $A = \mathit{supp}(a, t, a', t')$, we get from (29) a useful property of α -equivalence of atom-binding terms:

$$\begin{aligned} \llbracket a \rrbracket t =_\alpha \llbracket a' \rrbracket t' : \llbracket \mathbf{a} \rrbracket \sigma &\Leftrightarrow \\ (\forall a'' \in \mathbb{A}_a) a'' \# (a, t, a', t') &\Rightarrow (a\ a'') \cdot t =_\alpha (a'\ a'') \cdot t' : \sigma. \end{aligned} \quad (30)$$

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 3.10 (FRESHNESS THEOREM). *Given an atom-sort $\mathbf{a} \in \mathbb{AS}$ and a nominal set X , if a finitely supported function $h \in \mathbb{A}_a \rightarrow_{\text{fs}} X$ satisfies*

$$(\exists a \in \mathbb{A}_a) a \# h \ \& \ a \# h(a), \quad (31)$$

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

$$(\forall a \in \mathbb{A}_a) a \# h \Rightarrow h(a) = \mathit{fresh}(h). \quad (32)$$

Furthermore, $\mathit{supp}(\mathit{fresh}(h)) \subseteq \mathit{supp}(h)$.

PROOF. Given (31) we have to prove that h is constant on the nonempty set $\mathbb{A}_a - \mathit{supp}(h)$. First, note that, by Theorem 3.8 (with $S \triangleq \{a \in \mathbb{A}_a \mid a \# h(a)\}$ and $A \triangleq \mathit{supp}(h) \supseteq \mathit{supp}(S)$), if (31) holds, then

$$(\forall a \in \mathbb{A}_a) a \# h \Rightarrow a \# h(a). \quad (33)$$

Suppose $a, a' \in \mathbb{A}_a - \mathit{supp}(h)$. To see that $h(a) = h(a')$, without loss of generality we may assume $a \neq a'$. By (26), $a \# h(a')$ (since $a \# h$ and $a \# a'$); and $a' \# h(a')$ holds by (33). Hence,

$$\begin{aligned} h(a') &= (a\ a') \cdot h(a') && \text{since } (a, a') \# h(a') \\ &= ((a\ a') \cdot h)((a\ a') \cdot a') && \text{by (21)} \\ &= h((a\ a') \cdot a') && \text{since } (a, a') \# h \\ &= h(a). \end{aligned}$$

So there is a unique element $\mathit{fresh}(h) \in X$ satisfying (32). To see that it is supported by $\mathit{supp}(h)$, if a, a' are atoms (of the same sort) satisfying $(a, a') \# h$, choosing any

a'' in the infinite set $\mathbb{A}_a - \text{supp}(h, a, a')$, we have $(a a') \cdot \text{fresh}(h) = (a a') \cdot h(a'') = ((a a') \cdot h)((a a') \cdot a'') = h(a'') = \text{fresh}(h)$. Thus, $\text{supp}(h)$ supports $\text{fresh}(h)$. \square

4. Recursion and Induction Principles for α -Terms

Recall from Definition 2.5 that $\mathbb{T}_\alpha(\Sigma)_\sigma$ 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 $\mathbf{a} \in \Sigma_A$ and $e \in \mathbb{T}_\alpha(\Sigma)_{\mathbf{a}}$, 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 $\mathbf{s} \in \Sigma_D$ and $e \in \mathbb{T}_\alpha(\Sigma)_{\mathbf{s}}$, then there are unique $(K : \sigma \rightarrow \mathbf{s}) \in \Sigma_C$ and $e' \in \mathbb{T}_\alpha(\Sigma)_\sigma$ such that there exists t' with $e' = [t']_\alpha$ and $e = [K t']_\alpha$. In this case, we write e as $K e'$.

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

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

Atom-binding. If $\mathbf{a} \in \Sigma_A$, $\sigma \in \text{Ar}(\Sigma)$ and $e \in \mathbb{T}_\alpha(\Sigma)_{\ll \mathbf{a} \gg \sigma}$, then for each $a \in \mathbb{A}_a$ with $a \# e$ (i.e., with a not a free atom of e —cf. Example 3.2(iii)), there is a unique $e' \in \mathbb{T}_\alpha(\Sigma)_\sigma$ such that there exists t' with $e' = [t']_\alpha$ and $e = [\ll a \gg t']_\alpha$. In this case, we write e as $a \cdot e'$.

Using this notation, we now give a first version of structural recursion for α -terms over a nominal signature. Compared with Theorem 2.3, 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 4.1 (FIRST α -STRUCTURAL RECURSION THEOREM). *Let Σ be a nominal signature. Suppose we are given an arity-indexed family of nominal sets $(X_\sigma \mid \sigma \in \text{Ar}(\Sigma))$ and elements*

$$\begin{array}{ll} f_{\mathbf{a}} \in \mathbb{A}_a \rightarrow_{\text{fs}} X_{\mathbf{a}} & (\mathbf{a} \in \Sigma_A) \\ f_K \in X_\sigma \rightarrow_{\text{fs}} X_{\mathbf{s}} & ((K : \sigma \rightarrow_{\text{fs}} \mathbf{s}) \in \Sigma_C) \\ f_1 \in X_1 & \\ f_{\sigma_1 * \sigma_2} \in X_{\sigma_1} \times X_{\sigma_2} \rightarrow_{\text{fs}} X_{\sigma_1 * \sigma_2} & (\sigma_1, \sigma_2 \in \text{Ar}(\Sigma)) \\ f_{\ll \mathbf{a} \gg \sigma} \in \mathbb{A}_a \times X_\sigma \rightarrow_{\text{fs}} X_{\ll \mathbf{a} \gg \sigma} & (\mathbf{a} \in \Sigma_A, \sigma \in \text{Ar}(\Sigma)) \end{array}$$

all of which are supported by a finite set of atoms A and satisfy the freshness condition for binders (FCB): for each atom-binding arity $\ll \mathbf{a} \gg \sigma \in \text{Ar}(\Sigma)$, the

function $f_{\llbracket a \rrbracket \sigma}$ satisfies

$$(\exists a' \in \mathbb{A}_a) a' \notin A \ \& \ (\forall x \in X_\sigma) a' \# f_{\llbracket a \rrbracket \sigma}(a', x). \quad (\text{FCB})$$

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

$$\hat{f}a = f_a(a) \quad (34)$$

$$\hat{f}(K e) = f_K(\hat{f}e) \quad (35)$$

$$\hat{f}() = f_1 \quad (36)$$

$$\hat{f}(e_1, e_2) = f_{\sigma_1 * \sigma_2}(\hat{f}e_1, \hat{f}e_2) \quad (37)$$

$$a \notin A \Rightarrow \hat{f}(a.e) = f_{\llbracket a \rrbracket \sigma}(a, \hat{f}e) \quad (38)$$

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 2.3, taking advantage of the fact that we are working (informally) in higher-order logic.⁸ From the $\text{Ar}(\Sigma)$ -indexed family of nominal sets X_σ we define another such family: $S_\sigma \triangleq \text{Perm} \rightarrow_{\text{fs}} X_\sigma$ (regarding Perm as a nominal set as in Example 3.3 and using the \rightarrow_{fs} construct from Section 3.2). Now define elements $g_a, g_K, g_1, g_{\sigma_1 * \sigma_2}$ and $g_{\llbracket a \rrbracket \sigma}$ as in the statement of Theorem 2.3, as follows.

$$g_a a \triangleq \lambda \pi \in \text{Perm}. f_a(\pi(a))$$

$$g_K s \triangleq \lambda \pi \in \text{Perm}. f_K(s(\pi))$$

$$g_1 \triangleq \lambda \pi \in \text{Perm}. f_1$$

$$g_{\sigma_1 * \sigma_2}(s_1, s_2) \triangleq \lambda \pi \in \text{Perm}. f_{\sigma_1 * \sigma_2}(s_1(\pi), s_2(\pi))$$

$$g_{\llbracket a \rrbracket \sigma}(a, s) \triangleq \lambda \pi \in \text{Perm}. \text{fresh}(\lambda a' \in \mathbb{A}_a. f_{\llbracket a \rrbracket \sigma}(a', s(\pi \circ (a a')))).$$

The crucial clause in this definition is the last one, where we are using the *fresh* functional from Theorem 3.10 applied to the function $h \triangleq \lambda a' \in \mathbb{A}_a. f_{\llbracket a \rrbracket \sigma}(a', s(\pi \circ (a a')))$. For this to make sense, it has to be the case that h is finitely supported and satisfies condition (31) of that lemma, let us see why this is so. Since $\text{supp}(f_{\llbracket a \rrbracket \sigma}) \subseteq A$ by assumption, it follows that h is supported by the finite set $A \cup \text{supp}(s, \pi, a)$. To see that (31) holds of h , let a' be the atom whose existence is asserted by (FCB); thus, $a' \notin A$ and $a' \# f_{\llbracket a \rrbracket \sigma}(a', x)$ for any $x \in X_\sigma$. For any other $a'' \in \mathbb{A}_a - A$, we have $(a' a'') \cdot f_{\llbracket a \rrbracket \sigma} = f_{\llbracket a \rrbracket \sigma}$ (since $a', a'' \notin \text{supp}(f_{\llbracket a \rrbracket \sigma})$); hence, applying $(a' a'')$ to $a' \# f_{\llbracket a \rrbracket \sigma}(a', x)$, from Lemma 3.7, we get $a'' \# f_{\llbracket a \rrbracket \sigma}(a'', (a' a'') \cdot x)$ for any $x \in X_\sigma$. Choosing a'' to be in the infinite set $\mathbb{A}_a - (A \cup \text{supp}(s, \pi, a))$ and $x = (a' a'') \cdot s(\pi \circ (a a'))$, we conclude that $a'' \# h$ and $a'' \# f_{\llbracket a \rrbracket \sigma}(a'', s(\pi \circ (a a'))) = h(a'')$, as required for (31).

⁸In other words, the theorem is reducible to primitive recursion at higher types.

Applying Theorem 2.3 with this data, we get a family of functions

$$\hat{g}_\sigma \in \mathbb{T}(\Sigma)_\sigma \rightarrow (\text{Perm} \rightarrow_{\text{fs}} X_\sigma)$$

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

$$t_1 =_\alpha t_2 : \sigma \Rightarrow \hat{g}_\sigma t_1 = \hat{g}_\sigma t_2. \quad (39)$$

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 , that is, using Theorem 2.4:

$$(\forall \sigma \in \text{Ar}(\Sigma), t : \sigma)(\forall \pi, \pi' \in \text{Perm}) \hat{g}_\sigma t (\pi \circ \pi') = \hat{g}_\sigma (\pi' \cdot t) \pi. \quad (40)$$

In view of (39), the functions \hat{g}_σ induce functions $\hat{f}_\sigma \in \mathbb{T}_\alpha(\Sigma)_\sigma \rightarrow X_\sigma$ 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 2.3: 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 (34)–(38) follows from the recursion equations (10)–(14) satisfied by the \hat{g}_σ . That concludes the existence part of the proof of Theorem 4.1.

For the uniqueness part, suppose functions $f'_\sigma \in \mathbb{T}_\alpha(\Sigma)_\sigma \rightarrow_{\text{fs}} X_\sigma$ are all supported by A and satisfy the recursion equations (34)–(38) for \hat{f}_σ . Define $g'_\sigma \in \mathbb{T}(\Sigma)_\sigma \rightarrow S_\sigma$ by $g'_\sigma t \pi \triangleq f'_\sigma[\pi \cdot t]_\alpha$ ($\sigma \in \text{Ar}(\Sigma), t : \sigma, \pi \in \text{Perm}$). One can show that the g'_σ satisfy the same recursion equations (10)–(14) from Theorem 2.3 as the functions \hat{g}_σ ; so by the uniqueness part of that theorem, $g'_\sigma = \hat{g}_\sigma$. Therefore, for all $t : \sigma$, $f'_\sigma[t]_\alpha = f'_\sigma[t \cdot \iota]_\alpha \triangleq g'_\sigma t \iota = \hat{g}_\sigma t \iota \triangleq \hat{f}_\sigma[t]_\alpha$; hence, $f'_\sigma = \hat{f}_\sigma$. \square

Example 4.2 (Length of an α -Term). Gordon and Melham [1996, Section 3.3] 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 article (their Axiom 4) and embark on a detour via simultaneous substitutions to define the length function. This difficulty is analyzed by Norrish [2004, Sec. 3] on the way to his improved version of Gordon and Melham's recursion scheme (discussed further in Example 5.6 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.⁹ Thus, in Theorem 4.1, we take X_σ

⁹ The same goes for Norrish's `stripc` function, used to illustrate the limitations of Gordon and Melham's [1996] workaround for the `length` function [Norrish 2004, p. 247].

to be the discrete nominal set \mathbb{N} of natural numbers and

$$\begin{aligned} f_a &\triangleq \lambda a \in \mathbb{A}_a. 1 \\ f_K &\triangleq \lambda k \in \mathbb{N}. k + 1 \\ f_1 &\triangleq 0 \\ f_{\sigma_1 * \sigma_2} &\triangleq \lambda (k_1, k_2) \in \mathbb{N} \times \mathbb{N}. k_1 + k_2 \\ f_{\ll a \gg \sigma} &\triangleq \lambda (a, k) \in \mathbb{A}_a \times \mathbb{N}. k + 1. \end{aligned}$$

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 $\hat{f}_\sigma \in \mathbb{T}_\alpha(\Sigma)_\sigma \rightarrow_{\text{fs}} \mathbb{N}$. Writing $\text{length } e$ for $\hat{f}_\sigma e$, we have the expected properties of a length function on α -terms:

$$\begin{aligned} \text{length } a &= 1 \\ \text{length}(K e) &= \text{length } e + 1 \\ \text{length}() &= 0 \\ \text{length}(e_1, e_2) &= \text{length } e_1 + \text{length } e_2 \\ \text{length}(a. e) &= \text{length } e + 1. \end{aligned}$$

Note that the last clause holds for all a , because in (38) the condition “ $a \notin A$ ” is vacuously true (since $A = \emptyset$).

Remark 4.3 In Theorem 4.1, we gave (FCB) as an existential statement. It is in fact equivalent to the universal statement

$$(\forall a' \in \mathbb{A}_a) a' \notin A \Rightarrow (\forall x \in X_\sigma) a' \# f_{\ll a \gg \sigma}(a', x).$$

This follows from the “some/any” Theorem 3.8 by taking S to be $\{a' \in \mathbb{A}_a \mid (\forall x \in X_\sigma) a' \# f_{\ll a \gg \sigma}(a', x)\}$ and checking that A supports S .

Remark 4.4 (Primitive Recursion). Theorem 4.1 gives a simple “iterative” form of structural recursion for α -terms, rather than a more complicated “primitive recursive” form with recursion equations

$$\begin{aligned} \hat{f} a &= f_a(a) \\ \hat{f}(K e) &= f_K(e, \hat{f} e) \\ \hat{f}() &= f_1 \\ \hat{f}(e_1, e_2) &= f_{\sigma_1 * \sigma_2}(e_1, e_2, \hat{f} e_1, \hat{f} e_2) \\ a \notin A &\Rightarrow \hat{f}(a. e) = f_{\ll a \gg \sigma}(a, e, \hat{f} e). \end{aligned}$$

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 $X'_\sigma \triangleq \mathbb{T}_\alpha(\Sigma)_\sigma \times X_\sigma$ and functions

$$\begin{aligned} f'_a(a) &\triangleq (a, f_a(a)) \\ f'_K(e, x) &\triangleq (K e, f_K(e, x)) \\ f'_1 &\triangleq ((), f_1) \\ f'_{\sigma_1 * \sigma_2}((e_1, x_1), (e_2, x_2)) &\triangleq ((e_1, e_2), f_{\sigma_1 * \sigma_2}(e_1, e_2, x_1, x_2)) \\ f'_{\ll a \gg \sigma}(a, e, x) &\triangleq (a, e, f_{\ll a \gg \sigma}(a, e, x)), \end{aligned}$$

we first apply Theorem 4.1 to get functions $\hat{f}'_\sigma \in \mathbb{T}_\alpha(\Sigma)_\sigma \rightarrow_{\text{fs}} \mathbb{T}_\alpha(\Sigma)_\sigma \times X_\sigma$. 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 $\hat{f}_\sigma \in \mathbb{T}_\alpha(\Sigma)_\sigma \rightarrow_{\text{fs}} X_\sigma$ 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 4.1 was derived from ordinary structural recursion (Theorem 2.3), we prove this theorem as a corollary of ordinary structural induction (Theorem 2.4).

THEOREM 4.5 (FIRST α -STRUCTURAL INDUCTION THEOREM). *Let Σ be a nominal signature. Suppose we are given a finitely supported set $S \in P_{\text{fs}}(\mathbb{T}_\alpha(\Sigma))$ of α -terms over Σ . To prove that S is the whole of $\mathbb{T}_\alpha(\Sigma)$, it suffices to show*

$$(\forall a \in \Sigma_A, a \in \mathbb{A}_a) a \in S \quad (41)$$

$$(\forall (K : \sigma \rightarrow \mathbf{s}) \in \Sigma_C, e \in \mathbb{T}_\alpha(\Sigma)_\sigma) e \in S \Rightarrow K e \in S \quad (42)$$

$$() \in S \quad (43)$$

$$(\forall (\sigma_i \in \text{Ar}(\Sigma), e_i \in \mathbb{T}_\alpha(\Sigma)_{\sigma_i} \mid i = 1, 2)) e_1 \in S \ \& \ e_2 \in S \Rightarrow (e_1, e_2) \in S \quad (44)$$

$$(\forall a \in \Sigma_A, \sigma \in \text{Ar}(\Sigma)) (\exists a \in \mathbb{A}_a) a \# S \ \& \ (\forall e \in \mathbb{T}_\alpha(\Sigma)_\sigma) e \in S \Rightarrow a.e \in S. \quad (45)$$

PROOF. Let \bar{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:

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

Clearly, $S = \mathbb{T}_\alpha(\Sigma)$ if $\bar{S} = \mathbb{T}(\Sigma)$; and to prove the latter, it suffices to check that \bar{S} satisfies conditions (15)–(19) of Theorem 2.4. The first four of these follow immediately from (41)–(44) respectively. So it just remains to show that (45) implies that \bar{S} satisfies condition (19). First note that by Theorem 3.8 applied to the set of atoms $\{a \in \mathbb{A}_a \mid (\forall e \in \mathbb{T}_\alpha(\Sigma)_\sigma) e \in S \Rightarrow a.e \in S\}$, which is supported by $\text{supp}(S)$, (45) is equivalent to

$$\begin{aligned} (\forall a \in \Sigma_A, \sigma \in \text{Ar}(\Sigma)) (\forall a \in \mathbb{A}_a) a \# S \Rightarrow \\ (\forall e \in \mathbb{T}_\alpha(\Sigma)_\sigma) e \in S \Rightarrow a.e \in S. \end{aligned} \quad (46)$$

Given $a \in \mathbb{A}_a$ and $t \in \bar{S}$, we have to prove that $\ll a \gg t \in \bar{S}$, that is, that $[\pi \cdot \ll a \gg t]_\alpha \in S$ for any $\pi \in \text{Perm}$. Choosing any atom a' in the infinite set

$\mathbb{A}_a - \text{supp}(S, \pi, a, t)$, we have

$$\begin{aligned} \pi \cdot \ll a \gg t &= \ll \pi(a) \gg (\pi \cdot t) \\ &=_{\alpha} \ll a' \gg ((\pi \cdot t)\{a'/\pi(a)\}) && \text{by definition of } =_{\alpha} \text{ (Section 2.4)} \\ &= \ll a' \gg ((\pi(a) a') \cdot (\pi \cdot t)) && \text{since } a' \notin \text{atm}(\pi \cdot t) \\ &= \ll a' \gg (\pi' \cdot t) && \text{where } \pi' \triangleq (\pi(a) a') \circ \pi. \end{aligned}$$

So $[\pi \cdot \ll a \gg t]_{\alpha} = a' \cdot [\pi' \cdot t]_{\alpha} \in S$ by (46), since $a' \# S$ (by choice of a') and $[\pi' \cdot t]_{\alpha} \in S$, because $t \in \bar{S}$. So it is indeed the case that $\ll a \gg t \in \bar{S}$ when $a \in \mathbb{A}_a$ and $t \in \bar{S}$. \square

5. Second α -Structural Recursion & Induction Theorems

Theorem 4.1 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 $\mathbf{s} \in \Sigma_{\mathbf{D}}$, 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 2.1), to define the capture-avoiding substitution function $\hat{s}_{x,e} \in \mathbb{T}_{\alpha}(\Sigma)_t \rightarrow_{\text{fs}} \mathbb{T}_{\alpha}(\Sigma)_t$ discussed in the Introduction, we should only have to specify finitely supported functions corresponding to the right-hand sides of the defining Eqs. (6)–(9), that is, 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 4.1 we have to work out suitable choices for X_{σ} and for the functions $f_{\mathbf{v}}, f_1, f_{\sigma_1 * \sigma_2}, f_{\ll \mathbf{v} \gg \sigma}$ for any $\sigma, \sigma_1, \sigma_2 \in \text{Ar}(\Sigma)$.

So we will develop a second, ‘‘sort-directed’’ version of α -structural recursion in which one only has to give X_{σ} when $\sigma = \mathbf{s}$ 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 5.1 (SECOND α -STRUCTURAL RECURSION THEOREM). *Let Σ be a nominal signature. Suppose we are given a family of nominal sets $X = (X_{\mathbf{s}} \mid \mathbf{s} \in \Sigma_{\mathbf{D}})$ indexed by the data-sorts of Σ , a finite set A of atoms, and functions*

$$f_K \in X^{(\sigma)} \rightarrow_{\text{fs}} X^{(\mathbf{s})} \quad ((K : \sigma \rightarrow \mathbf{s}) \in \Sigma_{\mathbf{C}})$$

all of which are supported by A and satisfy

$$(\exists \bar{a} \in \mathbb{A}^{\sigma}) \bar{a} \# A \quad \& \quad (\forall \bar{x} \in X^{|\sigma|}) \bar{a} \#_{\sigma} \bar{x} \Rightarrow \bar{a} \# f_K(\bar{a}, \bar{x})_{\sigma}. \quad (\text{FCB}_K)$$

Then there is a unique family of finitely supported functions $(\hat{f}_{\mathbf{s}} \in \mathbb{T}_{\alpha}(\Sigma)_{\mathbf{s}} \rightarrow_{\text{fs}} X_{\mathbf{s}} \mid \mathbf{s} \in \Sigma_{\mathbf{D}})$ with $\text{supp}(\hat{f}_{\mathbf{s}}) \subseteq A$ and satisfying

$$(\forall \bar{a} \in \mathbb{A}^{\sigma}) \bar{a} \# A \Rightarrow (\forall \bar{e} \in \mathbb{T}_{\alpha}(\Sigma)^{|\sigma|}) \bar{a} \#_{\sigma} \bar{e} \Rightarrow \hat{f}_{\mathbf{s}}(K \bar{a} \cdot \bar{e}) = f_K(\bar{a}, \hat{f}^{|\sigma|} \bar{e})_{\sigma} \quad (47)$$

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

σ	$X^{(\sigma)}$	\mathbb{A}^σ	$X^{ \sigma }$
$\mathbf{a} \in \Sigma_A$	\mathbb{A}_a	1	\mathbb{A}_a
$\mathbf{s} \in \Sigma_D$	X_s	1	X_s
1	1	1	1
$\sigma_1 * \sigma_2$	$X^{(\sigma_1)} \times X^{(\sigma_2)}$	$\mathbb{A}^{\sigma_1} \otimes \mathbb{A}^{\sigma_2}$	$X^{ \sigma_1 } \times X^{ \sigma_2 }$
$\llbracket \mathbf{a} \rrbracket \sigma_1$	$\mathbb{A}_a \times X^{(\sigma_1)}$	$\mathbb{A}_a \otimes \mathbb{A}^{\sigma_1}$	$X^{ \sigma_1 }$

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	$(\bar{a}, \bar{x}) \in \mathbb{A}^\sigma \times X^{ \sigma }$	$\bar{a} \#_\sigma \bar{x} \in \mathbb{B}$	$(\bar{a}, \bar{x})_\sigma \in X^{(\sigma)}$
$\mathbf{a} \in \Sigma_A$	$((), a)$	<i>true</i>	a
$\mathbf{s} \in \Sigma_D$	$((), x)$	<i>true</i>	x
1	$((), ())$	<i>true</i>	$()$
$\sigma_1 * \sigma_2$	$((\bar{a}_1, \bar{a}_2), (\bar{x}_1, \bar{x}_2))$	$\bar{a}_1 \#_{\sigma_1} \bar{x}_1 \ \& \ \bar{a}_2 \#_{\sigma_2} \bar{x}_2 \ \& \ \bar{a}_1 \# \bar{x}_2 \ \& \ \bar{a}_2 \# \bar{x}_1$	$((\bar{a}_1, \bar{x}_1)_{\sigma_1}, (\bar{a}_2, \bar{x}_2)_{\sigma_2})$
$\llbracket \mathbf{a} \rrbracket \sigma_1$	$((a, \bar{a}_1), \bar{x}_1)$	$\bar{a}_1 \#_{\sigma_1} \bar{x}_1$	$(a, (\bar{a}_1, \bar{x}_1)_{\sigma_1})$

σ	$(\bar{a}, \bar{e}) \in \mathbb{A}^\sigma \times \mathbb{T}_\alpha(\Sigma)^{ \sigma }$	$\bar{a}. \bar{e} \in \mathbb{T}_\alpha(\Sigma)_\sigma$	$\hat{f}^{ \sigma } \bar{e} \in X^{ \sigma }$
$\mathbf{a} \in \Sigma_A$	$((), a)$	a	a
$\mathbf{s} \in \Sigma_D$	$((), e)$	e	$\hat{f}_s e$
1	$((), ())$	$()$	$()$
$\sigma_1 * \sigma_2$	$((\bar{a}_1, \bar{a}_2), (\bar{e}_1, \bar{e}_2))$	$(\bar{a}_1. \bar{e}_1, \bar{a}_2. \bar{e}_2)$	$(\hat{f}^{ \sigma_1 } \bar{e}_1, \hat{f}^{ \sigma_2 } \bar{e}_2)$
$\llbracket \mathbf{a} \rrbracket \sigma_1$	$((a, \bar{a}_1), \bar{e}_1)$	$a. (\bar{a}_1. \bar{e}_1)$	$\hat{f}^{ \sigma_1 } \bar{e}_1$

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

FIG. 2. Definitions used in Theorem 5.1.

In this theorem, we start with a family of nominal sets $X = (X_s \mid \mathbf{s} \in \Sigma_D)$ 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 3.2(i)), each data-sort as given by X , the unit arity as the one-element discrete nominal set (Example 3.2(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 $e : \mathbf{s}$ to elements $\hat{f}_s e \in X_s$ by giving recursion equations as in (47), with one (conditional) equation for each way of forming α -terms of data-sort, that is, for each constructor $K : \sigma \rightarrow \mathbf{s}$ 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 $\bar{a} \#_\sigma \bar{e}$ 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 ; $\bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}$ is a nested tuple of α -terms matching the *nonbinding* 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

σ	$S^{ \sigma } \in P_{\text{fs}}(\mathbb{T}_\alpha(\Sigma)^{ \sigma })$	$S_\sigma \in P_{\text{fs}}(\mathbb{T}_\alpha(\Sigma)_\sigma)$
$\mathbf{a} \in \Sigma_A$	\mathbb{A}_a	$\{a \mid a \in \mathbb{A}_a\}$
$\mathbf{s} \in \Sigma_D$	S_s	S_s
1	1	$\{()\}$
$\sigma_1 * \sigma_2$	$S^{ \sigma_1 } \times S^{ \sigma_2 }$	$\{(e_1, e_2) \mid e_1 \in S_{\sigma_1} \ \& \ e_2 \in S_{\sigma_2}\}$
$\langle \mathbf{a} \rangle \sigma_1$	$S^{ \sigma_1 }$	$\{a.e_1 \mid a \in \mathbb{A}_a - \text{supp}(S_{\sigma_1}) \ \& \ e_1 \in S_{\sigma_1}\}$

FIG. 3. Definitions used in Theorem 5.2.

K can be applied to get a constructed α -term $K \bar{a}. \bar{e} : \mathbf{s}$. The equation in (47) is restricted by two conditions:

- 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, $\bar{a} \#_\sigma \bar{e}$, 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 , that is, $\bar{a} \# A$ (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, that is, satisfying $\bar{a} \#_\sigma \bar{x}$,¹⁰ 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 $X^{(\mathbf{s})} = X_s$ 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 5.1 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 5.2 (SECOND α -STRUCTURAL INDUCTION THEOREM). *Let Σ be a nominal signature. Suppose we are given a family of finitely supported subsets $(S_s \in P_{\text{fs}}(\mathbb{T}_\alpha(\Sigma)_s) \mid \mathbf{s} \in \Sigma_D)$ 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 $\mathbb{T}_\alpha(\Sigma)_s$ for all $\mathbf{s} \in \Sigma_D$, it suffices to show for each constructor $(K : \sigma \rightarrow \mathbf{s}) \in \Sigma_C$ that*

$$(\exists \bar{a} \in \mathbb{A}^\sigma) \bar{a} \# A \ \& \ (\forall \bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}) \bar{a} \#_\sigma \bar{e} \ \& \ \bar{e} \in S^{|\sigma|} \Rightarrow K \bar{a}. \bar{e} \in S_s. \quad (\text{IH}_K)$$

In this theorem, we start with a family of subsets $S_s \subseteq \mathbb{T}_\alpha(\Sigma)_s$ 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 $t : \mathbf{s}$ is in S_s (for all $\mathbf{s} \in \Sigma_D$). The theorem guarantees this provided each constructor $(K : \sigma \rightarrow \mathbf{s}) \in \Sigma_C$ satisfies the induction hypothesis given by (IH_K) . This asserts the existence of a nested tuple \bar{a} of distinct atoms that

¹⁰This relation $\#_\sigma$, 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.

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 , that is, $\bar{a} \# A$ (they are also mutually distinct by definition of \mathbb{A}^σ);
- whenever \bar{e} is a nested tuple of α -terms that can appear in nonbinding positions in an α -term of arity σ (i.e., $\bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}$) for which \bar{a} is suitably fresh, that is, satisfying $\bar{a} \#_\sigma \bar{e}$, 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 5.2 is given in Appendix A and the proof of Theorem 5.1 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 5.3 (Atom-Abstraction and the Initial Algebra Property). For each atom-sort $\mathbf{a} \in \mathbb{A}\mathbb{S}$ and each nominal set X , let $[\mathbb{A}_\mathbf{a}]X$ denote the set of equivalence classes of pairs $(a, x) \in \mathbb{A}_\mathbf{a} \times X$ for the equivalence relation $(a, x) \sim (a', x')$ given by

$$(\exists a'' \in \mathbb{A}_\mathbf{a}) a'' \# (a, x, a', x') \ \& \ (a a'') \cdot x = (a' a'') \cdot x'. \quad (48)$$

The relation \sim is evidently reflexive and symmetric; to see that it is also transitive, one first applies Theorem 3.8 with $S = \{a'' \in \mathbb{A}_\mathbf{a} \mid (a a'') \cdot x = (a' a'') \cdot x'\}$ and $A = \text{supp}(a, x, a', x')$ to show that $(a, x) \sim (a', x')$ holds if and only if

$$(\forall a'' \in \mathbb{A}_\mathbf{a}) a'' \# (a, x, a, x') \Rightarrow (a a'') \cdot x = (a' a'') \cdot x'. \quad (49)$$

We write $[a]x$ for the \sim -equivalence class of the pair (a, x) and call it an *atom-abstraction*. It follows from Lemma 3.7 that \sim is a nominal subset of $\mathbb{A}_\mathbf{a} \times X$. So the quotient $[\mathbb{A}_\mathbf{a}]X$ is a nominal set as in Section 3.3. One can calculate that the support of each element $[a]x$ of $[\mathbb{A}_\mathbf{a}]X$ is $\text{supp}(x) - \{a\}$.

Using these atom-abstraction nominal sets, it is possible to give an *initial algebra* characterisation of $(\mathbb{T}_\alpha(\Sigma)_\mathbf{s} \mid \mathbf{s} \in \Sigma_D)$ that is equivalent to Theorem 5.1. Consider the category whose objects are families of nominal sets $X = (X_\mathbf{s} \mid \mathbf{s} \in \Sigma_D)$ indexed by the data-sorts of Σ , and whose morphisms $f : X \rightarrow X'$ are indexed families $f = (f_\mathbf{s} : X_\mathbf{s} \rightarrow X'_\mathbf{s} \mid \mathbf{s} \in \Sigma_D)$ 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 $i : F_\Sigma I \rightarrow I$. Such an algebra is *initial* if for any other such algebra $f : F_\Sigma X \rightarrow X$, there is a unique morphism $\hat{f} : I \rightarrow X$ so that

$$\begin{array}{ccc} F_\Sigma I & \xrightarrow{i} & I \\ F_\Sigma \hat{f} \downarrow & & \downarrow \hat{f} \\ F_\Sigma X & \xrightarrow{f} & X \end{array} \quad (50)$$

commutes, that is, satisfying

$$(\forall \mathbf{s} \in \Sigma_D)(\forall (K : \sigma \rightarrow \mathbf{s}) \in \Sigma_C)(\forall x \in I^{|\sigma|}) \hat{f}_\mathbf{s}(i_\mathbf{s}(K, x)) = f_\mathbf{s}(K, \hat{f}^{[\sigma]} x).$$

Standard category-theoretic results give that i is an isomorphism and that the initial algebra (I, i) is unique up to isomorphism. Theorem 5.1 can be used to prove that $(\mathbb{T}_\alpha(\Sigma)_\mathbf{s} \mid \mathbf{s} \in \Sigma_D)$ is the (object part of) an initial F_Σ -algebra. Conversely, one

Action of F_Σ on objects:

$$(F_\Sigma X)_s \triangleq \sum_{(K:\sigma \rightarrow s) \in \Sigma_C} X^{[\sigma]} = \{(K, x) \mid (K : \sigma \rightarrow s) \in \Sigma_C \ \& \ x \in X^{[\sigma]}\}.$$

Action of F_Σ on morphisms:

$$(F_\Sigma f)_s(K, x) \triangleq (K, f^{[\sigma]} x).$$

σ	$X^{[\sigma]}$	$x \in X^{[\sigma]} \mapsto f^{[\sigma]} x \in X'^{[\sigma]}$
$\mathbf{a} \in \Sigma_A$	\mathbb{A}_a	$a \mapsto a$
\mathbf{s}	S_s	$x \mapsto f_s x$
$\mathbf{1}$	1	$() \mapsto ()$
$\sigma_1 * \sigma_2$	$X^{[\sigma_1]} \times X^{[\sigma_2]}$	$(x_1, x_2) \mapsto (f^{[\sigma_1]} x_1, f^{[\sigma_2]} x_2)$
$\langle \mathbf{a} \rangle \sigma_1$	$[\mathbb{A}_a](X^{[\sigma_1]})$	$[a]x_1 \mapsto [a](f^{[\sigma_1]} x_1)$

FIG. 4. Functor F_Σ associated with a nominal signature Σ .

can give a direct inductive construction of an initial F_Σ -algebra and use the initial algebra property (50) to deduce Theorem 5.1: see Gabbay and Pitts [2002, Sect. 6].

5.1. EXAMPLE: λ -CALCULUS (WITH *letrec*). Let Σ be the nominal signature from Example 2.1. Thus, $\mathbb{T}(\Sigma)_t$ is the set Λ of abstract syntax trees for λ -calculus with *letrec*; and $\mathbb{T}_\alpha(\Sigma)_t$ is the quotient Λ / \equiv_α of that set by the usual notion of α -equivalence—in other words $\mathbb{T}_\alpha(\Sigma)_t$ 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

$$f_V \in \mathbb{A}_V \rightarrow_{fs} X \quad (51)$$

$$f_A \in X \times X \rightarrow_{fs} X \quad (52)$$

$$f_L \in \mathbb{A}_V \times X \rightarrow_{fs} X \quad (53)$$

$$f_{Letrec} \in \mathbb{A}_V \times ((\mathbb{A}_V \times X) \times X) \rightarrow_{fs} X \quad (54)$$

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_L) is equivalent to

$$(\exists a \in \mathbb{A}_V) a \notin A \ \& \ (\forall x \in X) a \# f_L(a, x) \quad (55)$$

and that (FCB_{Letrec}) is equivalent to

$$(\exists a, a' \in \mathbb{A}_V) a \neq a' \ \& \ a, a' \notin A \ \& \ (\forall x, x' \in X) a' \# x \Rightarrow (a, a') \# f_{Letrec}(a, ((a', x'), x)). \quad (56)$$

So for this nominal signature, Theorem 5.1 gives us the following recursion principle. We state it using the usual concrete syntax for λ -calculus and using the fact (noted in Example 3.2(iii)) that the support of a term $e \in \Lambda / \equiv_\alpha$ is its finite set $fv(e)$ of free variable

THEOREM 5.4. *Let Λ / \equiv_α be the nominal set of α -equivalence classes of λ -terms with local recursive function declarations:*

$$e ::= x \mid e e \mid \lambda x. e \mid \text{letrec } x x = e \text{ in } e$$

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 (51)–(54) all supported by some finite set of atoms A and with f_L and f_{Letrec} satisfying (55) and (56), then there is a unique function $\hat{f} \in (\Lambda / =_\alpha \rightarrow_{fs} X)$ supported by A and satisfying

$$\hat{f}(x) = f_v(x) \quad (57)$$

$$\hat{f}(e_1 e_2) = f_A(\hat{f} e_1, \hat{f} e_2) \quad (58)$$

$$x \notin A \Rightarrow \hat{f}(\lambda x. e) = f_L(x, \hat{f} e) \quad (59)$$

$$x, y \notin A \ \& \ x \notin fv(e_2) \cup \{y\} \Rightarrow \\ \hat{f}(letrec\ y\ x = e_1\ in\ e_2) = f_{Letrec}(y, ((x, \hat{f} e_1), \hat{f} e_2)). \quad (60)$$

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

$$(\forall x \in \mathbb{A}_v) x \in S \quad (61)$$

$$(\forall e_1, e_2 \in S) e_1 e_2 \in S \quad (62)$$

$$(\exists x \in \mathbb{A}_v) x \notin A \ \& \ (\forall e \in S) \lambda x. e \in S \quad (63)$$

$$(\exists x, y \in \mathbb{A}_v) x \neq y \ \& \ x, y \notin A \ \& \\ (\forall e_1, e_2 \in S) x \notin fv(e_2) \Rightarrow letrec\ y\ x = e_1\ in\ e_2 \in S \quad (64)$$

respectively. So for this signature Theorem 5.2 says that S contains all terms if it satisfies (61)–(64).

Example 5.5 (Capture-Avoiding Substitution). The example mentioned in the Introduction of capture-avoiding substitution of λ -terms, $\hat{s}_{x,e} \in \Lambda / =_\alpha \rightarrow \Lambda / =_\alpha$, is obtained from the above theorem by taking X to be the nominal set $\Lambda / =_\alpha$. Given $x \in \mathbb{A}_v$ 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 (55) is satisfied because, as noted in Example 3.2(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 $y \# f_L(y, e_1) = \lambda y. e_1$ simply because $y \notin fv(\lambda y. e_1) = fv(e_1) - \{y\}$. Similarly, condition (56) is satisfied because

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

so that $(y, z) \# f_{Letrec}(z, ((y, e_1), e_2)) = letrec\ z\ y = e_1\ in\ e_2$ 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 (59) and (60) correspond precisely to the conditions in (8) and (9).

Example 5.6 (Recursion with Varying Parameters). Norrish [2004, p. 245] considers a variant of capture-avoiding substitution whose definition involves recursion with varying parameters; it motivates the parametrized recursion principle that he presents in that article. The α -structural recursion principles we have given here do not involve extra parameters, let alone varying ones; nevertheless it is possible to derive parameterized versions from them. In the case of ordinary structural recursion, one can derive a parameterized version from an unparameterized one by currying parameters and defining maps into function sets using Theorem 2.3. In the presence of binders, one has to do something slightly more complicated, involving the Freshness Theorem 3.10, to derive a parameterized (FCB) from the unparameterized 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 2.1 by deleting the *Letrec* constructor. Fixing on a pair of atoms $x_1, x_2 \in \mathbb{A}_v$, we seek a function $s \in (\Lambda / =_\alpha) \rightarrow_{\text{fs}} (\Lambda / =_\alpha) \rightarrow_{\text{fs}} (\Lambda / =_\alpha)$ satisfying for all y, e, e_1, e_2 :

$$s(y)(e) = \begin{cases} e & \text{if } y = x_1 \\ y & \text{if } y \neq x_1 \end{cases} \quad (65)$$

$$s(e_1 e_2)(e) = (s(e_1)(e))(s(e_2)(e)) \quad (66)$$

$$y \notin \text{fv}(e) \cup \{x_1, x_2\} \Rightarrow s(\lambda y. e_1)(e) = \lambda y. s(e_1)(x_2 e) \quad (67)$$

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 (67), where the application term $x_2 e$ appears instead. Such a function s can be obtained from Theorem 5.4 (restricted to pure λ -terms) as $s = \hat{f}$ if we take X to be the nominal set $(\Lambda / =_\alpha) \rightarrow_{\text{fs}} (\Lambda / =_\alpha)$ and use the functions

$$f_V \triangleq \lambda y \in \mathbb{A}_v. \lambda e \in (\Lambda / =_\alpha). \text{if } y = x_1, \text{ then } e \text{ else } y$$

$$f_A \triangleq \lambda (\xi_1, \xi_2) \in X \times X. \lambda e \in (\Lambda / =_\alpha). (\xi_1 e) (\xi_2 e)$$

$$f_L \triangleq \lambda (y, \xi_1) \in \mathbb{A}_v \times X. \lambda e \in (\Lambda / =_\alpha). \text{fresh}(h(y, \xi_1, e))$$

where the last clause uses Theorem 3.10 applied to the finitely supported function $h(y, \xi_1, e) \in \mathbb{A}_v \rightarrow_{\text{fs}} (\Lambda / =_\alpha)$ that maps each $y' \in \mathbb{A}_v$ to

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

This function is easily seen to satisfy the property (31) needed to apply the theorem. All the above functions are supported by $A \triangleq \{x_1, x_2\}$. Properties (57) and (58) of \hat{f} give (65) and (66) respectively. When $y \neq x_1, x_2$, property (59) gives us $\hat{f}(\lambda y. e_1) = f_L(y, \hat{f} e_1) = \text{fresh}(h(y, \hat{f} e_1, e))$. So if $y \# (x_1, x_2, e)$, picking any $y' \# (x_1, x_2, e, e_1, h)$, then by Theorem 3.10 we have $\text{fresh}(h(y, \hat{f} e_1, e)) = h(y, \hat{f} e_1, e)(y') \triangleq \lambda y'. ((y y') \cdot (\hat{f} e_1))(x_2 e) = \lambda y'. (y y') \cdot (\hat{f} e_1)(x_2 e)$. Hence by definition of $=_\alpha$, $\hat{f}(\lambda y. e_1) = \lambda y. \hat{f} e_1(x_2 e)$, as required for (67).

5.2. EXAMPLE: π -CALCULUS. Let Σ be the nominal signature from Example 2.2. Suppose we are given nominal sets $X_{\text{proc}}, X_{\text{gsum}}, X_{\text{pre}}$ and functions

$$\begin{aligned}
f_{Gsum} &\in X_{gsum} \rightarrow_{fs} X_{proc} \\
f_{Par} &\in X_{proc} \times X_{proc} \rightarrow_{fs} X_{proc} \\
f_{Res} &\in \mathbb{A}_{chan} \times X_{proc} \rightarrow_{fs} X_{proc} \\
f_{Rep} &\in X_{proc} \rightarrow_{fs} X_{proc} \\
\\
f_{Zero} &\in \mathbf{1} \rightarrow_{fs} X_{gsum} \\
f_{Pre} &\in X_{pre} \rightarrow_{fs} X_{gsum} \\
f_{Plus} &\in X_{gsum} \times X_{gsum} \rightarrow_{fs} X_{gsum} \\
\\
f_{Out} &\in (\mathbb{A}_{chan} \times \mathbb{A}_{chan}) \times X_{proc} \rightarrow_{fs} X_{pre} \\
f_{In} &\in \mathbb{A}_{chan} \times (\mathbb{A}_{chan} \times X_{proc}) \rightarrow_{fs} X_{pre} \\
f_{Tau} &\in \mathbf{1} \rightarrow_{fs} X_{pre} \\
f_{Match} &\in (\mathbb{A}_{chan} \times \mathbb{A}_{chan}) \times X_{pre} \rightarrow_{fs} X_{pre}
\end{aligned}$$

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

$$(\exists a \in \mathbb{A}_{chan}) a \notin A \ \& \ (\forall x \in X_{proc}) a \# f_{Res}(a, x) \quad (68)$$

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

$$(\forall a' \in \mathbb{A}_{chan})(\forall x \in X_{proc}) a \neq a' \Rightarrow a \# f_{In}(a', (a, x)) \quad (69)$$

respectively; and conditions (FCB_K) for $K \neq Res, In$ are all equivalent to *true*. So if f_{Res} and f_{In} satisfy (68) and (69), then by Theorem 5.1, there are unique finitely supported functions $\hat{f}_s \in \mathbb{T}_\alpha(\Sigma)_s \rightarrow_{fs} X_s$ (for $s = \text{proc}, \text{gsum}, \text{pre}$) all supported by A and satisfying for all $e, e_1, e_2, a_1, a_2, a, a'$ of suitable arity

$$\hat{f}(Ke) = f_K(\hat{f}e) \quad (K = Gsum, Rep, Pre) \quad (70)$$

$$\hat{f}(K(e_1, e_2)) = f_K(\hat{f}e_1, \hat{f}e_2) \quad (K = Par, Plus) \quad (71)$$

$$\hat{f}(K()) = f_K() \quad (K = Zero, Tau) \quad (72)$$

$$\hat{f}(K((a_1, a_2), e)) = f_K((a_1, a_2), \hat{f}e) \quad (K = Out, Match) \quad (73)$$

$$a \notin A \Rightarrow \hat{f}(Res a. e) = f_{Res}(a, \hat{f}e) \quad (74)$$

$$a \notin A \ \& \ a \neq a' \Rightarrow \hat{f}(In(a', a. e)) = f_{In}(a', (a, \hat{f}e)) \quad (75)$$

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

We leave the reader to work out what induction principle Theorem 5.2 gives for this signature.

6. Extended Example: Normalization 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 generalizes 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 *normalization by evaluation (NBE)* for simply typed λ -calculus [Berger and Schwichtenberg 1991; Berger et al. 2003]. This produces

so-called “ $\beta\eta$ -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 by Dybjer and Filinski [2002, 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 Berger et al. [2003, sect. 2.5], Dybjer and Filinski [2002, sect. 3.3] and Fiore [2002], for example. This tends to obscure the simple, but informal idea behind reification. Dybjer and Filinski [2002] go on to mention after the above quote that “freshness” can be characterised rigorously in the framework of Gabbay and Pitts [2002]. 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 Dybjer and Filinski [2002, sect. 3.2].

6.1. TYPED λ -TERMS AND THEIR $\beta\eta$ -LONG NORMAL FORMS. We assume the reader is familiar with simply typed λ -calculus; if not, see Barendregt [1992] for example. 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 treatment using explicitly typed variables.¹¹ Let

$$Ty \triangleq \{\tau ::= \iota \mid \tau \rightarrow \tau\} \quad (76)$$

be the set of *simple type symbols* over a single *ground type* ι . We assume given an injective function $\tau \in Ty \mapsto v_\tau \in \mathbb{A}\mathbb{S}$ 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 $v_{\tau'}$ are different atom-sorts; so recalling the assumptions we made in Section 2.1, the sets of atoms \mathbb{A}_{v_τ} and $\mathbb{A}_{v_{\tau'}}$ are disjoint.

¹¹ Such an approach is fine for simply typed terms, but becomes unworkable for calculi with type variables or dependent types.

Consider the nominal signature Σ^{STL} with

atom-sorts	data-sorts	constructors
\mathbf{v}_τ	\mathbf{t}_τ	$V_{r_\tau} : \mathbf{v}_\tau \rightarrow \mathbf{t}_\tau$ $Ap_{\tau, \tau'} : \mathbf{t}_{\tau \rightarrow \tau'} * \mathbf{t}_\tau \rightarrow \mathbf{t}_{\tau'}$ $Lm_{\tau, \tau'} : \llbracket \mathbf{v}_\tau \rrbracket \mathbf{t}_{\tau'} \rightarrow \mathbf{t}_{\tau \rightarrow \tau'}$

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

$$\Lambda(\tau) \triangleq \mathbb{T}_\alpha(\Sigma^{\text{STL}})_{\mathbf{t}_\tau} \quad (77)$$

of α -terms of arity \mathbf{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 for $\beta\eta$ -long normal forms. In the system we are considering, the general form of a type is $\tau_1 \rightarrow (\tau_2 \rightarrow \dots (\tau_k \rightarrow \iota) \dots)$ for some $k \geq 0$; using conventional notation for typed λ -terms, a term of this type is in $\beta\eta$ -long normal form if it takes the form

$$\lambda x_1 : \tau_1. \lambda x_2 : \tau_2 \dots \lambda x_k : \tau_k. x n_1 \dots n_k$$

where x is a variable and n_1, \dots, n_k are $\beta\eta$ -long normal forms. Accordingly, we can use the nominal signature Σ^{LNF} with

atom-sorts	data-sorts	constructors
\mathbf{v}_τ	\mathbf{n}_τ \mathbf{u}_τ	$V_\tau : \mathbf{v}_\tau \rightarrow \mathbf{u}_\tau$ $A_{\tau, \tau'} : \mathbf{u}_{\tau \rightarrow \tau'} * \mathbf{n}_\tau \rightarrow \mathbf{u}_{\tau'}$ $L_{\tau, \tau'} : \llbracket \mathbf{v}_\tau \rrbracket \mathbf{n}_{\tau'} \rightarrow \mathbf{n}_{\tau \rightarrow \tau'}$ $I : \mathbf{u}_\iota \rightarrow \mathbf{n}_\iota$

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

$$N(\tau) \triangleq \mathbb{T}_\alpha(\Sigma^{\text{LNF}})_{\mathbf{n}_\tau} \quad (78)$$

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

$$U(\tau) \triangleq \mathbb{T}_\alpha(\Sigma^{\text{LNF}})_{\mathbf{u}_\tau} \quad (79)$$

corresponds to the set of α -equivalence classes of *neutral* (or *atomic*) terms of type τ ; see Dybjer and Filinski [2002, p. 155], for example.

Notation 6.1. 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 $\mathbb{A}_{\mathbf{v}_\tau}$, $\Lambda(\tau)$, $N(\tau)$ and $U(\tau)$ will be written x , e , n and u respectively.

— $V_{r_\tau} x$ 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 Σ^{LNF} tells us that there are functions

$$i_\tau \in N(\tau) \rightarrow_{\text{fs}} \Lambda(\tau) \quad j_\tau \in U(\tau) \rightarrow_{\text{fs}} \Lambda(\tau) \quad (80)$$

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

$$\begin{aligned} j_\tau x &= x \\ j_{\tau'}(u n) &= (j_{\tau \rightarrow \tau'} u)(i_\tau n) \\ i_{\tau \rightarrow \tau'}(\lambda x : \tau. n) &= \lambda x : \tau. i_\tau n \\ i_i u &= j_i u. \end{aligned}$$

We aim to use the technique of NBE [Berger and Schwichtenberg 1991] to define the *normalization function*

$$norm_\tau \in \Lambda(\tau) \rightarrow_{fs} N(\tau) \quad (81)$$

with the following properties:

$$(\forall \tau \in Ty, e_1, e_2 \in \Lambda(\tau)) e_1 =_{\beta\eta} e_2 \Rightarrow norm_\tau e_1 = norm_\tau e_2 \quad (82)$$

$$(\forall \tau \in Ty, n \in N(\tau)) norm_\tau(i_\tau n) = n \quad (83)$$

$$(\forall \tau \in Ty, e \in \Lambda(\tau)) i_\tau(norm_\tau e) =_{\beta\eta} e. \quad (84)$$

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

$$(\forall \tau, \tau' \in Ty, x \in \mathbb{A}_{v_\tau}, e_1 \in \Lambda(\tau'), e_2 \in \Lambda(\tau)) (\lambda x : \tau. e_1)e_2 =_{\beta\eta} e_1[x := e_2] \quad (85)$$

$$(\forall \tau, \tau' \in Ty, e \in \Lambda(\tau \rightarrow \tau'), x \in \mathbb{A}_{v_\tau}) x \# e \Rightarrow e =_{\beta\eta} \lambda x : \tau. e x. \quad (86)$$

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 5.5; instead, we will regard it as a special case of the simultaneous substitution functions defined in the next section.

Once (82)–(84) 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 $\mathbb{T}(\Sigma^{LNF})_{n_\tau}$ by applying the $norm_\tau$ function. We aim to show how to use α -structural recursion to define $norm_\tau$; and how to prove (82)–(84) using, among other things, α -structural induction.

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

$$dom(\sigma) \triangleq \left\{ x \in \bigcup_{\tau \in Ty} \mathbb{A}_{v_\tau} \mid \sigma x \neq x \right\}$$

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 $x \in \mathbb{A}_{v_\tau}$ 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

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

$$[x]\sigma = \sigma x \quad (87)$$

$$[e_1 e_2]\sigma = ([e_1]\sigma)([e_2]\sigma) \quad (88)$$

$$x \# \sigma \Rightarrow [\lambda x : \tau. e]\sigma = \lambda x : \tau. [e]\sigma \quad (89)$$

$\tau \in Ty, e \in \Lambda(\tau), \rho \in Env \mapsto \llbracket e \rrbracket \rho \in D(\tau)$:

$$\llbracket x \rrbracket \rho = \rho x \quad (90)$$

$$\llbracket e_1 e_2 \rrbracket \rho = \llbracket e_1 \rrbracket \rho (\llbracket e_2 \rrbracket \rho) \quad (91)$$

$$\llbracket \lambda x : \tau. e \rrbracket \rho = \lambda d \in D(\tau). \llbracket e \rrbracket (\rho\{x \mapsto d\}) \quad (92)$$

FIG. 5. Substitution and denotation.

is specified by the recursion Eq. (87)–(88) in Figure 5. For each $\sigma \in Sub$, we can use the second α -structural recursion theorem for Σ^{STL} to define the functions ($[-]\sigma \in \Lambda(\tau) \rightarrow_{\text{fs}} \Lambda(\tau) \mid \tau \in Ty$) satisfying these equations, much as in Example 5.5, but using $\text{supp}(\sigma)$ as the common finite support A . (In particular, in Theorem 5.1 the only nontrivial freshness condition on binders, $(\text{FCB}_{Lm_{\tau, \tau'}})$, is easily verified.)

The *identity substitution* $\sigma_0 \in Sub$ maps each $x \in \mathbb{A}_{V_\tau}$ 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 $x \in \mathbb{A}_{V_\tau}$ to $[\sigma_1 x]\sigma_2$. These operations satisfy:

$$[e]\sigma_0 = e \quad (93)$$

$$[e](\sigma_1; \sigma_2) = \llbracket [e]\sigma_1 \rrbracket \sigma_2. \quad (94)$$

Both properties can be proved easily by applying the second α -structural induction principle (Theorem 5.2) for the nominal signature Σ^{STL} . For example, to prove (94), for each $\sigma_1, \sigma_2 \in Sub$ we take $S_{\tau'}$ to be $\{e \in \Lambda(\tau) \mid [e](\sigma_1; \sigma_2) = \llbracket [e]\sigma_1 \rrbracket \sigma_2\}$, which is supported by $A = \text{supp}(\sigma_1, \sigma_2)$; then $(\text{IH}_{V_{\tau'}})$ and $(\text{IH}_{Ap_{\tau, \tau'}})$ are easy to verify, using (87) and (88) respectively; for $(\text{IH}_{Lm_{\tau, \tau'}})$, which is the statement

$$(\exists x \in \mathbb{A}_{V_\tau}) x \notin A \ \& \ (\forall e \in S_{\tau'}) \lambda x : \tau. e \in S_{\tau \rightarrow \tau'},$$

we can choose any x in the infinite set $\mathbb{A}_{V_\tau} - A$ (so that $x \# (\sigma_1, \sigma_2)$): if $e \in S_{\tau'}$, then

$$\begin{aligned} & [\lambda x : \tau. e](\sigma_1; \sigma_2) \\ &= \llbracket [\lambda x : \tau. e]\sigma_1 \rrbracket \sigma_2 && \text{by definition of } \sigma_1; \sigma_2 \\ &= [\lambda x : \tau. [e]\sigma_1]\sigma_2 && \text{by (89), since } x \# \sigma_1 \\ &= \lambda x : \tau. \llbracket [e]\sigma_1 \rrbracket \sigma_2 && \text{by (89), since } x \# \sigma_2 \\ &= \lambda x : \tau. [e](\sigma_1; \sigma_2) && \text{since } e \in S_{\tau'} \\ &= [\lambda x : \tau. e](\sigma_1; \sigma_2) && \text{by (89), since } x \# (\sigma_1; \sigma_2) \text{ (because } \sigma_1; \sigma_2 \text{ is} \\ & && \text{supported by } \text{supp}(\sigma_1) \cup \text{supp}(\sigma_2) \subseteq A) \end{aligned}$$

and hence $\lambda x : \tau. e \in S_{\tau \rightarrow \tau'}$.

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

$$e_1[x := e_2] \triangleq [e_1](\sigma_0\{x \mapsto e_2\}), \quad (95)$$

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:

$$D(\iota) \triangleq N(\iota) \quad (96)$$

$$D(\tau \rightarrow \tau') \triangleq D(\tau) \rightarrow_{fs} D(\tau'). \quad (97)$$

An *environment* ρ is a function that, for any $\tau \in Ty$, maps the atoms in \mathbb{A}_{v_τ} 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 $\tau \in Ty$, $e \in \Lambda(\tau)$ and $\rho \in Env$, we wish to define the *denotation* $\llbracket e \rrbracket \rho \in D(\tau)$, satisfying Eq. (90)–(92) 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 (92): $\rho\{x \mapsto d\}$ 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 5.1 it seems that we should take X_{t_τ} to be $Env \rightarrow_{fs} D(\tau)$ and use the functions

$$f_{v_{t_\tau}} \triangleq \lambda x \in \mathbb{A}_{v_\tau}. \lambda \rho \in Env. \rho x \quad (98)$$

$$f_{Ap_{t_\tau, t'}} \triangleq \lambda (\xi, \xi') \in X_{t_\tau \rightarrow t'} \times X_{t_\tau}. \lambda \rho \in Env. \xi \rho (\xi' \rho) \quad (99)$$

$$f_{Lm_{t_\tau, t'}} \triangleq \lambda (x, \xi') \in \mathbb{A}_{v_\tau} \times X_{t_\tau}. \lambda \rho \in Env. \lambda d \in D(\tau). \xi'(\rho\{x \mapsto d\}). \quad (100)$$

These functions are all supported by $A = \emptyset$ and the only nontrivial freshness condition on binders is $(FCB_{Lm_{t_\tau, t'}})$, which in view of Theorem 3.8 (the “some/any” theorem) is the requirement that for all $x \in \mathbb{A}_{v_\tau}$ and $\xi \in Env \rightarrow_{fs} D(\tau')$

$$x \# \lambda \rho \in Env. \lambda d \in D(\tau). \xi(\rho\{x \mapsto d\}). \quad (101)$$

If we could prove that, then the theorem gives us functions $\hat{f}_{t_\tau} \in \Lambda(\tau) \rightarrow_{fs} (Env \rightarrow_{fs} D(\tau))$ satisfying (47)—from which it follows that $\llbracket e \rrbracket \rho \triangleq \hat{f}_{t_\tau} e \rho$ satisfies (90)–(92). The problem is that (101) is not true for all elements ξ of $Env \rightarrow_{fs} D(\tau')$! 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 article we have encountered a really nontrivial “freshness condition on binders”. Eq. (90)–(92) are typical of many such definitions in denotational semantics; but why is it the case that the right-hand side of Eq. (92) is independent of the choice of bound variable x on the left-hand side? Compared with the similar question for Eq. (89) a few lines above it, there seems no quick answer to this question. However, the freshness of x for $\lambda \rho \in Env. \lambda d \in D(\tau). \llbracket e \rrbracket (\rho\{x \mapsto d\})$ 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 5.1 to construct denotations.¹² To do so, we cut down to the following nominal subset of $Env \rightarrow_{fs} D(\tau)$:

$$X_{t_\tau} \triangleq \{\xi \in Env \rightarrow_{fs} D(\tau) \mid \Phi_1(\xi) \ \& \ \Phi_2(\xi)\}, \quad (102)$$

where

$$\Phi_1(\xi) \triangleq (\exists A \in P_{fin}(\mathbb{A}))(\forall \tau \in Ty, x \in \mathbb{A}_{v_\tau}, d \in D(\tau), \rho \in Env) \quad (103)$$

$$x \notin A \Rightarrow \xi(\rho\{x \mapsto d\}) = \xi \rho$$

$$\Phi_2(\xi) \triangleq (\forall \pi \in Perm, \rho \in Env) (\pi \cdot \xi) \rho = \xi(\rho \circ \pi). \quad (104)$$

The slightly elaborate form of Φ_1 compared with property 1 above is needed to show that the functions defined in (98)–(100) satisfy

$$\begin{aligned} & \Phi_1(f_{v_{r_\tau}} x) \\ & \Phi_1(\xi) \ \& \ \Phi_1(\xi') \Rightarrow \Phi_1(f_{Ap_{\tau,\tau'}}(\xi, \xi')) \\ & \Phi_1(\xi') \Rightarrow \Phi_1(f_{Lm_{\tau,\tau'}}(x, \xi')). \end{aligned}$$

Similar properties hold for Φ_2 . Therefore, $f_{v_{r_\tau}} \in \mathbb{A}_{v_\tau} \rightarrow_{fs} X_{t_\tau}$, $f_{Ap_{\tau,\tau'}} \in X_{t_\tau \rightarrow \tau'} \times X_{t_\tau} \rightarrow_{fs} X_{t_{\tau'}}$ and $f_{Lm_{\tau,\tau'}} \in \mathbb{A}_{v_\tau} \times X_{t_{\tau'}} \rightarrow_{fs} X_{t_\tau \rightarrow \tau'}$. Furthermore, it is now the case that if $x \in \mathbb{A}_{v_\tau}$ and $\xi' \in X_{t_{\tau'}}$, then (101) holds. To see this, first note that since $\Phi_1(\xi')$ holds, there is a finite set of atoms A such that

$$(\forall \tau \in Ty, x' \in \mathbb{A}_{v_\tau}, d' \in D(\tau), \rho' \in Env) x' \notin A \Rightarrow \xi'(\rho'\{x' \mapsto d'\}) = \xi' \rho'. \quad (105)$$

Choose any x' in the infinite set $\mathbb{A}_{v_\tau} - \text{supp}(x, \xi', A)$. 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

$$\begin{aligned} x &= (x \ x') \cdot x' \\ & \# (x \ x') \cdot \lambda \rho \in Env.\lambda d \in D(\tau).\xi'(\rho\{x \mapsto d\}) \\ &= \lambda \rho \in Env.\lambda d \in D(\tau).((x \ x') \cdot \xi')(\rho\{x' \mapsto d\}) \\ &= \lambda \rho \in Env.\lambda d \in D(\tau).\xi'(\rho\{x' \mapsto d\}) \circ (x \ x') \quad \text{because } \Phi_2(\xi') \text{ holds} \\ &= \lambda \rho \in Env.\lambda d \in D(\tau).\xi'(\rho\{x \mapsto d\})\{x' \mapsto \rho x\} \quad \text{because } x' \neq x \\ &= \lambda \rho \in Env.\lambda d \in D(\tau).\xi'(\rho\{x \mapsto d\}) \quad \text{by (105), since } x' \notin A. \end{aligned}$$

So (101) does indeed hold. Therefore, we can apply Theorem 5.1 to these nominal sets and functions to obtain $\hat{f}_{t_\tau} \in X_{t_\tau}$ satisfying (47). Defining $\llbracket e \rrbracket \rho \triangleq \hat{f}_{t_\tau} e \rho$, we get the required properties (90)–(92).

Denotations respect $\beta\eta$ -conversion:

$$e_1 =_{\beta\eta} e_2 \Rightarrow \llbracket e_1 \rrbracket = \llbracket e_2 \rrbracket. \quad (106)$$

To see this, it suffices to show that $\llbracket - \rrbracket = \llbracket - \rrbracket$ is a congruence relation equating β -convertible (85) and η -convertible (86) terms. Congruence is immediate from

¹²We are in effect carrying out a simultaneous inductive-recursive definition: see Dybjer [2000].

the defining properties (90)–(92) of $\llbracket - \rrbracket$. To see that $\llbracket - \rrbracket$ respects β -conversion, one has to show

$$\llbracket e'[x := e] \rrbracket \rho = \llbracket e' \rrbracket (\rho\{x \mapsto \llbracket e \rrbracket \rho\}) \quad (107)$$

and for η -conversion one has to show

$$x \notin \text{fv}(e) \Rightarrow \llbracket e \rrbracket (\rho\{x \mapsto d\}) = \llbracket e \rrbracket \rho. \quad (108)$$

These properties can be proved using the second α -structural induction principle (Theorem 5.2) for Σ^{STL} , with the second one used in the proof of the first.¹³ For (108), one can use the subsets

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

which are all supported by the empty set. Properties (90)–(92) of $\llbracket - \rrbracket$ allow one to show that these subsets satisfy the induction hypotheses $(\text{IH}_{V_{r_{\tau}}})$, $(\text{IH}_{A_{p_{\tau}, \tau'}})$ and $(\text{IH}_{L_{m_{\tau}, \tau'}})$; hence, by Theorem 5.2, each S_{τ} is the whole of $\Lambda(\tau)$ and (108) holds.

For (107), for each $x \in \mathbb{A}_{v_{\tau}}$ and $e \in \Lambda(\tau)$, we apply Theorem 5.2 to the subsets

$$S'_{\tau} \triangleq \{e' \in \Lambda(\tau) \mid (\forall \rho \in \text{Env}) \llbracket e'[x := e] \rrbracket \rho = \llbracket e' \rrbracket (\rho\{x \mapsto \llbracket e \rrbracket \rho\})\},$$

which are all supported by $A \triangleq \{x\} \cup \text{supp}(e)$. To conclude that $S'_{\tau} = \Lambda(\tau)$, we have to prove that these subsets S'_{τ} satisfy the induction hypotheses $(\text{IH}_{V_{r_{\tau}}})$, $(\text{IH}_{A_{p_{\tau}, \tau'}})$ and $(\text{IH}_{L_{m_{\tau}, \tau'}})$. The first two follow easily from (90) and (91). For $(\text{IH}_{L_{m_{\tau}, \tau'}})$, we have to show $(\exists x' \in \mathbb{A}_{v_{\tau}}) x' \# (x, e) \ \& \ (\forall e' \in S'_{\tau}) \lambda x' : \tau. e' \in S'_{\tau}$; but choosing any x' in the infinite set $\mathbb{A}_{v_{\tau}} - \text{supp}(x, e)$, for each $e' \in S'_{\tau}$ and $\rho \in \text{Env}$ we have:

$$\begin{aligned} & \llbracket (\lambda x' : \tau. e')[x := e] \rrbracket \rho \\ &= \llbracket \lambda x' : \tau. (e'[x := e]) \rrbracket \rho && \text{since } x' \notin \text{fv}(x, e) \\ &= \lambda d \in D(\tau). \llbracket e'[x := e] \rrbracket (\rho\{x' \mapsto d\}) && \text{by (92)} \\ &= \lambda d \in D(\tau). \llbracket e' \rrbracket (\rho\{x' \mapsto d\}\{x \mapsto \llbracket e \rrbracket (\rho\{x' \mapsto d\})\}) && \text{since } e' \in S'_{\tau} \\ &= \lambda d \in D(\tau). \llbracket e' \rrbracket (\rho\{x' \mapsto d\}\{x \mapsto \llbracket e \rrbracket \rho\}) && \text{by (108), since } x' \notin \text{fv}(e) \\ &= \lambda d \in D(\tau). \llbracket e' \rrbracket (\rho\{x \mapsto \llbracket e \rrbracket \rho\}\{x' \mapsto d\}) && \text{since } x' \neq x \\ &= \llbracket \lambda x' : \tau. e' \rrbracket (\rho\{x \mapsto \llbracket e \rrbracket \rho\}) && \text{by (92)} \end{aligned}$$

so that $\lambda x' : \tau. e' \in S'_{\tau}$, as required.

6.4. REIFICATION, REFLECTION AND NORMALIZATION. The *reification function* $\downarrow_{\tau} \in D(\tau) \rightarrow_{\text{fs}} N(\tau)$ turns elements of the denotational model into $\beta\eta$ -long normal forms, whereas the *reflection function* $\uparrow_{\tau} \in U(\tau) \rightarrow_{\text{fs}} D(\tau)$ 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

¹³ It might seem that (108) is a consequence of the property (103) that we built in to the construction of $\llbracket - \rrbracket$; but unfortunately that property only tells us that $\llbracket e \rrbracket (\rho\{x \mapsto d\})$ and $\llbracket e \rrbracket \rho$ are equal when x avoids some finite set of atoms, rather than the particular finite set $\text{supp}(e)$.

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

$$\downarrow_\tau(n) \triangleq n \quad (109)$$

$$\downarrow_{\tau \rightarrow \tau'}(f) \triangleq \text{fresh}(\lambda x \in \mathbb{A}_{V_\tau}. L_{\tau, \tau'} x. \downarrow_{\tau'}(f(\uparrow_\tau(V_\tau x)))) \quad (110)$$

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

$$\uparrow_\tau(u) \triangleq I u \quad (111)$$

$$\uparrow_{\tau \rightarrow \tau'}(u) \triangleq \lambda d \in D(\tau). \uparrow_{\tau'}(A_{\tau, \tau'}(u, \downarrow_\tau(d))) \quad (112)$$

FIG. 6. Reification (\downarrow_τ) and reflection (\uparrow_τ).

introduced by Notation 6.1. The interesting part of the definition in this figure is clause (110), where we make use of the *fresh* construct from Theorem 3.10. To do so, we have to check that the conditions of that theorem are satisfied; but in this case that is easy: given $f \in D(\tau \rightarrow \tau')$, choosing any atom x in the infinite set $\mathbb{A}_{V_\tau} - \text{supp}(\downarrow_{\tau'}, \uparrow_\tau, f)$,¹⁴ then the element

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

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

Definition 6.2 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 *normalization function* $\text{norm}_\tau \in \Lambda(\tau) \rightarrow_{\text{fs}} N(\tau)$ by:

$$\text{norm}_\tau(e) \triangleq \downarrow_\tau(\llbracket e \rrbracket \rho_0). \quad (113)$$

Recall that we wish to show that norm_τ has the properties (82)–(84). Property (82) follows immediately from (106). Property (83) is the first half of the following result.

LEMMA 6.3. *For all $\tau \in Ty, n \in N(\tau)$ and $u \in U(\tau)$*

$$\downarrow_\tau(\llbracket i_\tau n \rrbracket \rho_0) = n \quad (114)$$

$$\llbracket j_\tau u \rrbracket \rho_0 = \uparrow_\tau(u). \quad (115)$$

PROOF. These properties can be proved by using the second α -structural induction principle (Theorem 5.2) for Σ^{LNF} to show that the subsets

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

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

¹⁴In fact, $\text{supp}(\downarrow_{\tau'}, \uparrow_\tau, f) = \text{supp}(f)$ because the reification and reflection functions turn out to have empty support.

$$\begin{aligned}
\sim_\tau &\subseteq D(\tau) \times \Lambda(\tau): \\
n \sim_i e &\triangleq (i_i n =_{\beta\eta} e) \\
f \sim_{\tau \rightarrow \tau'} e &\triangleq (\forall d_1 \in D(\tau), e_1 \in \Lambda(\tau)) d_1 \sim_\tau e_1 \Rightarrow f d \sim_{\tau'} Ap_{\tau, \tau'}(e, e_1) \\
\sim &\subseteq Env \times Sub: \\
\rho \sim \sigma &\triangleq (\forall \tau \in Ty, x \in \mathbb{A}_{V_\tau}) \rho x \sim_\tau \sigma x
\end{aligned}$$

FIG. 7. Logical relation.

are equal to $N(\tau)$ and $U(\tau)$ respectively (for all $\tau \in Ty$). Since these subsets are supported by the empty set of atoms one can take $A = \emptyset$ in the theorem and prove

$$\begin{aligned}
(\forall x \in \mathbb{A}_{V_\tau}) V_\tau x \in S_{u_\tau} & \quad (\text{IH}_{V_\tau}) \\
(\forall u \in S_{u_{\tau \rightarrow \tau'}}, n \in S_{n_\tau}) A_{\tau, \tau'}(u, n) \in S_{u_{\tau'}} & \quad (\text{IH}_{A_{\tau, \tau'}}) \\
(\exists x \in \mathbb{A}_{V_\tau})(\forall n \in S_{n_{\tau'}}) L_{\tau, \tau'} x. n \in S_{n_{\tau \rightarrow \tau'}} & \quad (\text{IH}_{L_{\tau, \tau'}}) \\
(\forall u \in S_{u_i}) Iu \in S_{u_i} & \quad (\text{IH}_I)
\end{aligned}$$

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 $(\text{IH}_{L_{\tau, \tau'}})$, because of the use of *fresh* in $\downarrow_{\tau \rightarrow \tau'}$. Picking any atom x in \mathbb{A}_{V_τ} , suppose $n \in S_{n_{\tau'}}$. We wish to prove that $L_{\tau, \tau'} x. n \in S_{n_{\tau \rightarrow \tau'}}$. Note that $x \# L_{\tau, \tau'} x. n$; so since $\llbracket - \rrbracket, i_{\tau \rightarrow \tau'}$ and ρ_0 have empty support, it is the case that $x \# \llbracket i_{\tau \rightarrow \tau'}(L_{\tau, \tau'} x. n) \rrbracket \rho_0 = \llbracket \lambda x : \tau. i_{\tau'} n \rrbracket \rho_0 = f$, where $f \triangleq \lambda d \in D(\tau). \llbracket i_{\tau'} n \rrbracket (\rho_0 \{x \mapsto d\})$. Hence

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

and therefore by definition of *fresh* in Theorem 3.10

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

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

It just remains to prove that $norm_\tau$ has the property (84). 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 6.4 (FUNDAMENTAL PROPERTY OF \sim). *For all $\tau \in Ty, e \in \Lambda(\tau), \rho \in Env$ and $\sigma \in Sub$*

$$\rho \sim \sigma \Rightarrow \llbracket e \rrbracket \rho \sim_\tau \llbracket e \rrbracket \sigma. \quad (116)$$

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

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

are equal to $\Lambda(\tau)$, for all $\tau \in Ty$. Proving the induction hypotheses (IH_{V_τ}) and $(\text{IH}_{Ap_{\tau,\tau'}})$ is straightforward; and for $(\text{IH}_{Lm_{\tau,\tau'}})$, one first proves

$$d \sim_\tau e \ \& \ e =_{\beta_\eta} e' \Rightarrow d \sim_\tau e' \quad (117)$$

$$([\lambda x. e] \sigma) e' =_{\beta_\eta} [e](\sigma\{x \mapsto e'\}) \quad (118)$$

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

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

$$(\forall \tau \in Ty, u \in U(\tau)) \uparrow_\tau(u) \sim_\tau j_\tau u \quad (119)$$

$$(\forall \tau \in Ty, d \in D(\tau), e \in \Lambda(\tau)) d \sim_\tau e \Rightarrow i_\tau(\downarrow_\tau(d)) =_{\beta_\eta} e. \quad (120)$$

(Both properties are proved simultaneously, and one needs to make use of (117).) Because of (119), the identity substitution $\sigma_0 \in Sub$ satisfies $\rho_0 \sim \sigma_0$. So by Lemma 6.4 and (93) we have $\llbracket e \rrbracket \rho_0 \sim_\tau [e] \sigma_0 = e$, for all $e \in \Lambda(\tau)$. Thus, (120) gives $i_\tau(\downarrow_\tau(\llbracket e \rrbracket \rho_0)) =_{\beta_\eta} e$, that is, $i_\tau(\text{norm}_\tau(e)) =_{\beta_\eta} e$, as required for (84).

7. Assessment

7.1. MATHEMATICAL PERSPECTIVE. The results of this article are directly inspired by my joint work with Gabbay on “FM-set” theory [Gabbay and Pitts 2002] and by his PhD thesis [Gabbay 2000]. 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 et al. [2004] and is developed further in Cheney’s thesis [Cheney 2004]); 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 classical logic, rather than being axiomatised within FM-set theory; see Cheney [2004, Chap. 3] for a more leisurely and generalized account of the theory of nominal sets. Second, the recursion and induction principles developed here refer directly to α -terms, that is, standard α -equivalence classes of abstract syntax trees, rather than using an initial algebra that is merely isomorphic to the set of α -terms; see Remark 5.3. This is also the approach taken by Norrish [2004], building on Gordon and Melham’s five axioms for α -equivalence [Gordon and Melham 1996]; and also by Urban and Tasson [2005]. Norrish’s recursion principle [Norrish 2004, Figure 1] has side-conditions requiring that the function being defined be well-behaved with respect to variable-permutations and with respect to generation of fresh variables. In effect, these side-conditions build in just enough of the theory of nominal sets to yield a well-defined and total function, while only having to specify how binders with fresh variables are mapped by the function. Along with Urban and Tasson [2005], I prefer to develop the theory of nominal sets in its

own right and then give a simple-looking¹⁵ 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 3.10, independently of the recursion principle. We used Theorem 3.10 in the reduction of Theorem 4.1 to Theorem 2.3, in the reduction of “varying parameters” to “no parameters” (Example 5.6) 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 [Shinwell and Pitts 2005, Sect. 3].

How easy is it to apply these principles of α -structural recursion and induction? Just as for the work of Gordon and Melham [1996], Norrish [2004] and Urban and Tasson [2005], to use them one does not have to change to an unfamiliar logic (we remain in higher-order classical 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 3.5) 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 (47) in Theorem 5.1 (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).

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 3.2) 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 3.5) 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 5.5, (FCB_K) is very easily checked. In other cases, such as in the definition of the denotation functions $\llbracket - \rrbracket$ 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 5.2) 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.

¹⁵ Once one gets used to the distinctive concepts of nominal sets, I believe that principles such as Theorems 4.1, 5.1 and 5.4 *are* quite simple.

7.2. AUTOMATED THEOREM PROVING PERSPECTIVE. Based on 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 “article-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 3.5, 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 [2002]. The disadvantage of such a “new logic” approach is that one does not have easy access to the legacy of already-proved 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 [2004] provides a starting point within the HOL4 system. For the Isabelle system, Urban and Tasson [2005] developed a theory equivalent to nominal sets within Isabelle/HOL up to and including an induction principle (but not a recursion principle) for the particular nominal signature for λ -terms.¹⁶ Building upon that foundation, Urban and Berghofer are developing a *Nominal Datatype Package* for Isabelle/HOL that allows its users to declare nominal signatures and then have principles of α -structural recursion and induction for those signatures proved and ready to be applied: see <http://isabelle.in.tum.de/nominal/>.

HOL4 and Isabelle/HOL are theorem-proving systems for higher-order *classical* logic. Can support for α -structural recursion and induction be provided in proof-assistants based on constructive metatheory, such as the Coq [pauillac.inria.fr/coq/] and Twelf [www.cs.cmu.edu/twelf/] systems? Answering this question involves investigating a topic we have not addressed here: the constructive content of the theory of nominal sets and (hence) the integration of atoms, atom-permutations and finite support with the metalogical frameworks underlying Coq and Twelf. First steps towards this have been taken by Schöpp [2006], Schöpp and Stark [2004], and by Cheney (personal communication, 2006).

Appendix

A. Proof of Second α -Structural Induction Theorem

Given a nominal signature Σ , we prove the second α -structural induction principle (Theorem 5.2) as a corollary of the first one, Theorem 4.5. To do so, we must first develop some properties of the operation $\bar{a} \in \mathbb{A}^\sigma$, $\bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|} \mapsto \bar{a} \cdot \bar{e} \in \mathbb{T}_\alpha(\Sigma)_\sigma$ and the relation $\#_\sigma \subseteq \mathbb{A}^\sigma \times \mathbb{T}_\alpha(\Sigma)^{|\sigma|}$ defined in Figure 2.

LEMMA A.1. *If $\bar{a} \in \mathbb{A}^\sigma$, $\bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}$ and $\bar{a} \#_\sigma \bar{e}$, then $\text{supp}(\bar{a} \cdot \bar{e}) = \text{supp}(\bar{e}) - \text{supp}(\bar{a})$.*

¹⁶Their 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.

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 3.2(3), that the support of an α -term is its finite set of free atoms (so that in particular $\text{supp}(a. e) = \text{supp}(e) - \{a\}$). \square

LEMMA A.2. *Let $(S_{\mathbf{s}} \in P_{\text{fs}}(\mathbb{T}_{\alpha}(\Sigma)_{\mathbf{s}}) \mid \mathbf{s} \in \Sigma_{\text{D}})$ be a family of finitely supported subsets of α -terms indexed by the data-sorts of Σ . Extend this to arity-indexed families*

$$\begin{aligned} (S_{\sigma} \in P_{\text{fs}}(\mathbb{T}_{\alpha}(\Sigma)_{\sigma}) \mid \sigma \in \text{Ar}(\Sigma)) \\ (S^{|\sigma|} \in P_{\text{fs}}(\mathbb{T}_{\alpha}(\Sigma)^{|\sigma|}) \mid \sigma \in \text{Ar}(\Sigma)) \end{aligned}$$

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 $\bar{a} \#_{\sigma} \bar{e}$.

PROOF. This can be proved by ordinary structural induction for arities, for all A simultaneously. The induction steps when $\sigma = \mathbf{a} \in \Sigma_{\text{A}}$, $\sigma = \mathbf{s} \in \Sigma_{\text{D}}$ 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 $A \cup \text{supp}(e_2)$, 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 $\bar{a}_1 \#_{\sigma_1} \bar{e}_1$. 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 $\bar{a}_2 \#_{\sigma_2} \bar{e}_2$. 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 $(\bar{a}_1, \bar{a}_2) \#_{\sigma} (\bar{e}_1, \bar{e}_2)$. For this, since $\bar{a}_i \#_{\sigma_i} \bar{e}_i$ ($i = 1, 2$) and $\bar{a}_2 \# \bar{e}_1$ hold by construction, we just need to prove that $\bar{a}_1 \# \bar{e}_2$. We chose \bar{a}_1 to have its support disjoint from $\text{supp}(e_2)$, which by Lemma A.1 is $\text{supp}(\bar{e}_2) - \text{supp}(\bar{a}_2)$; since $\text{supp}(\bar{a}_1)$ is also disjoint from $\text{supp}(\bar{a}_2)$, it follows that $\text{supp}(\bar{a}_1) \cap \text{supp}(\bar{e}_2) = \emptyset$, that is, $\bar{a}_1 \# \bar{e}_2$, as required.

Case $\sigma = \ll \mathbf{a} \gg \sigma_1$. By definition of S_{σ} , each $e \in S_{\sigma}$ is of the form $e = a.e_1$ with $a \in \mathbb{A}_{\mathbf{a}} - \text{supp}(S_{\sigma_1})$ and $e_1 \in S_{\sigma_1}$. Given A , choosing any atom a' in the infinite set $\mathbb{A}_{\mathbf{a}} - \text{supp}(A, a, e_1)$, it follows from the characterization of α -equivalence in (29) 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 $\bar{a}_1 \#_{\sigma_1} \bar{e}_1$. since $a' \# \bar{a}_1$, we have $(a', \bar{a}_1) \in \mathbb{A}_{\mathbf{a}} \otimes \mathbb{A}^{\sigma_1} = \mathbb{A}^{\sigma}$. 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 $(a', \bar{a}_1) \#_{\sigma} \bar{e}_1$. \square

To prove Theorem 5.2, 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 generalization of permutations that transpose a pair of atoms.

Definition A.3 (Vector-Transpositions). Let Σ be a nominal signature and $\sigma \in \text{Ar}(\Sigma)$ 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 $\tau_{\bar{a}, \bar{a}'} \in \text{Perm}$ 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:

σ	$(\bar{a}, \bar{a}') \in \mathbb{A}^\sigma \otimes \mathbb{A}^\sigma \mapsto \tau_{\bar{a},\bar{a}'} \in Perm$
$\mathbf{a} \in \Sigma_A$	$((), ()) \mapsto \iota$
$\mathbf{s} \in \Sigma_D$	$((), ()) \mapsto \iota$
$\mathbf{1}$	$((), ()) \mapsto \iota$
$\sigma_1 * \sigma_2$	$((\bar{a}_1, \bar{a}_2), (\bar{a}'_1, \bar{a}'_2)) \mapsto \tau_{\bar{a}_1, \bar{a}'_1} \circ \tau_{\bar{a}_2, \bar{a}'_2}$
$\ll \mathbf{a} \gg \sigma_1$	$((a, \bar{a}_1), (a', \bar{a}'_1)) \mapsto (a \ a') \circ \tau_{\bar{a}_1, \bar{a}'_1}$

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:

$$\pi \circ \tau_{\bar{a},\bar{a}'} \circ \pi^{-1} = \tau_{(\pi \cdot \bar{a}), (\pi \cdot \bar{a}')} \quad (\text{any } \pi \in Perm) \quad (121)$$

$$supp(\tau_{\bar{a},\bar{a}'}) \subseteq supp(\bar{a}) \cup supp(\bar{a}') \quad (122)$$

$$\tau_{\bar{a},\bar{a}'} = \tau_{\bar{a}',\bar{a}} \quad (123)$$

$$\tau_{\bar{a},\bar{a}'} \circ \tau_{\bar{a},\bar{a}'} = \iota \quad (124)$$

$$\tau_{\bar{a},\bar{a}'} \cdot \bar{a} = \bar{a}' \quad (125)$$

Part (2) of the following result generalizes the “some/any” Theorem 3.8 from single atoms to elements of \mathbb{A}^σ .

LEMMA A.4.

- (1) For all finite sets of atoms A , there is some $\bar{a} \in \mathbb{A}^\sigma$ with $\bar{a} \# A$.
- (2) 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.

$$(\forall \bar{a} \in \mathbb{A}^\sigma) \bar{a} \# A \Rightarrow \bar{a} \in S \quad (126)$$

$$(\exists \bar{a} \in \mathbb{A}^\sigma) \bar{a} \# A \ \& \ \bar{a} \in S. \quad (127)$$

PROOF. Part (1) follows easily by ordinary structural induction for arities $\sigma \in Ar(\Sigma)$. For part (2), first note that in view of part (1), one just has to show that (127) implies (126). 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 (1) to find $\bar{a}'' \in \mathbb{A}^\sigma$ with $\bar{a}'' \# (A, \bar{a}, \bar{a}')$ and then use the vector-transpositions of Definition A.3 to define $\pi \triangleq \tau_{\bar{a}',\bar{a}''} \circ \tau_{\bar{a},\bar{a}''} \in Perm$.¹⁷ By (125), we have $\pi \cdot \bar{a} = \bar{a}'$; and because by (122) 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. \square

The following result generalizes the characterization of α -equivalence mentioned in Example 3.9. Note that in view of the previous lemma, the right-hand side of the

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 3.8) we resort to swapping via an intermediate, fresh \bar{a}'' .

bi-implication in (128) could be replaced by $(\forall \bar{a}'' \in \mathbb{A}^\sigma) \bar{a}'' \# A \Rightarrow \tau_{\bar{a}, \bar{a}''} \cdot \bar{e} = \tau_{\bar{a}', \bar{a}''} \cdot \bar{e}' \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}$.

LEMMA A.5. *Suppose $\bar{a}, \bar{a}' \in \mathbb{A}^\sigma$ and $\bar{e}, \bar{e}' \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}$ satisfy $\bar{a} \#_\sigma \bar{e}$ and $\bar{a}' \#_\sigma \bar{e}'$. If A is a finite set of atoms supporting $(\bar{a}, \bar{a}', \bar{e}, \bar{e}')$, then*

$$\begin{aligned} \bar{a} \cdot \bar{e} = \bar{a}' \cdot \bar{e}' \in \mathbb{T}_\alpha(\Sigma)_\sigma &\Leftrightarrow \\ (\exists \bar{a}'' \in \mathbb{A}^\sigma) \bar{a}'' \# A \ \&\ \tau_{\bar{a}, \bar{a}''} \cdot \bar{e} = \tau_{\bar{a}', \bar{a}''} \cdot \bar{e}' \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}. \end{aligned} \quad (128)$$

PROOF. This can be proved by ordinary structural induction for arities σ . The induction step for atom-binding arities uses the corresponding property (29) of $=_\alpha$. \square

We can now complete the proof of Theorem 5.2. We are given subsets $(S_{\mathbf{s}} \in P_{\text{fs}}(\mathbb{T}_\alpha(\Sigma)_{\mathbf{s}}) \mid \mathbf{s} \in \Sigma_{\text{D}})$ supported by the finite set of atoms A and satisfying (IH_K) for each $K \in \Sigma_{\text{C}}$. Define $(S_\sigma \in P_{\text{fs}}(\mathbb{T}_\alpha(\Sigma)_\sigma) \mid \sigma \in \text{Ar}(\Sigma))$ 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 $S \triangleq \bigcup_{\sigma \in \text{Ar}(\Sigma)} S_\sigma$. If $S = \mathbb{T}_\alpha(\Sigma)$, then $S_{\mathbf{s}} = \mathbb{T}_\alpha(\Sigma)_{\mathbf{s}}$ for each $\mathbf{s} \in \Sigma_{\text{D}}$; and to prove $S = \mathbb{T}_\alpha(\Sigma)$ we just have to prove that this S satisfies the conditions (41)–(45) of Theorem 4.5. Conditions (41) and (43)–(45) are immediate from the definition of $(S_\sigma \mid \sigma \in \text{Ar}(\Sigma))$ in Figure 3. For condition (42), given $(K : \sigma \rightarrow \mathbf{s}) \in \Sigma_{\text{C}}$ and $e \in S_\sigma$, we have to show that $Ke \in S_{\mathbf{s}}$. First note that by applying Lemma A.4(2) to the subset $\{\bar{a} \in \mathbb{A}^\sigma \mid (\forall \bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}) \bar{a} \#_\sigma \bar{e} \ \& \ \bar{e} \in S^{|\sigma|} \Rightarrow K\bar{a} \cdot \bar{e} \in S_{\mathbf{s}}\}$, which is supported by A , to see that (IH_K) is equivalent to

$$(\forall \bar{a} \in \mathbb{A}^\sigma) \bar{a} \# A \Rightarrow (\forall \bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}) \bar{a} \#_\sigma \bar{e} \ \& \ \bar{e} \in S^{|\sigma|} \Rightarrow K\bar{a} \cdot \bar{e} \in S_{\mathbf{s}}. \quad (129)$$

By Lemma A.2, $e = \bar{a} \cdot \bar{e}$ with $\bar{a} \in \mathbb{A}^\sigma$, $\bar{e} \in S^{|\sigma|}$, $\bar{a} \# A$ and $\bar{a} \#_\sigma \bar{e}$; so (129) gives $Ke \in S_{\mathbf{s}}$. Thus, we can indeed apply Theorem 4.5 to conclude that S is the whole of $\mathbb{T}_\alpha(\Sigma)$ and hence in particular that each $S_{\mathbf{s}}$ is equal to the whole of $\mathbb{T}_\alpha(\Sigma)_{\mathbf{s}}$, completing the proof of Theorem 5.2.

B. Proof of Second α -Structural Recursion Theorem

We will establish the second α -structural recursion principle (Theorem 5.1) 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 $\hat{f}_{\mathbf{s}}$ were they single-valued and total relations—and then prove they are so by applying Theorem 5.2. A further application of Theorem 5.2 is needed to show that the functions $\hat{f}_{\mathbf{s}}$ are unique with the stated properties.

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

B.1. *Existence of the Functions* $\hat{f}_s \in \mathbb{T}_\alpha(\Sigma)_s \rightarrow_{fs} X_s$

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

$$\frac{K : \sigma \rightarrow s \quad \bar{a} \in \mathbb{A}^\sigma \quad \bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|} \quad \bar{a} \# A \quad \bar{a} \#_\sigma \bar{e} \quad (\bar{e}, \bar{x}) \in F^{|\sigma|} \quad \bar{a} \#_\sigma \bar{x}}{(K \bar{a}. \bar{e}, f_K(\bar{a}, \bar{x})_\sigma) \in F_s} \quad (130)$$

where $F^{|\sigma|} \subseteq \mathbb{T}_\alpha(\Sigma)^{|\sigma|} \times X^{|\sigma|}$ is defined from any family ($F_s \subseteq \mathbb{T}_\alpha(\Sigma)_s \times X_s \mid s \in \Sigma_D$) by ordinary recursion on the structure of arities as follows.

σ	$F^{ \sigma } \subseteq \mathbb{T}_\alpha(\Sigma)^{ \sigma } \times X^{ \sigma }$
$\mathbf{a} \in \Sigma_A$	$\{(a, a) \mid a \in \mathbb{A}_a\}$
$\mathbf{s} \in \Sigma_D$	F_s
$\mathbf{1}$	$\{(\emptyset, \emptyset)\}$
$\sigma_1 * \sigma_2$	$\{((\bar{e}_1, \bar{e}_2), (\bar{x}_1, \bar{x}_2)) \mid (\bar{e}_1, \bar{x}_1) \in F^{ \sigma_1 } \ \& \ (\bar{e}_2, \bar{x}_2) \in F^{ \sigma_2 }\}$
$\ll \mathbf{a} \gg \sigma_1$	$F^{ \sigma_1 }$

It is not hard to see that the set of rules determined by (130) is supported by A . Hence, by Theorem 3.6, so are all the subsets F_s (and the subsets $F^{|\sigma|}$ defined from them). For the existence part of Theorem 5.1, it suffices to show that each F_s is the graph of a function \hat{f}_s from $\mathbb{T}_\alpha(\Sigma)_s$ 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 $\bar{a} \#_\sigma \bar{e}$, then $\bar{a} \#_\sigma \hat{f}^{|\sigma|} \bar{e}$, because of the definition of $\hat{f}^{|\sigma|}$;
- hence $(\hat{f}_s \mid s \in \Sigma_D)$ satisfies the recursion property (47): for if $\bar{a} \# A$ and $\bar{a} \#_\sigma \bar{e}$, then $\bar{x} \triangleq \hat{f}^{|\sigma|} \bar{e}$ satisfies $(\bar{e}, \bar{x}) \in F^{|\sigma|}$ and $\bar{a} \#_\sigma \bar{x}$; so $(K \bar{a}. \bar{e}, f_K(\bar{a}, \bar{x})_\sigma) \in F_s$ by rule (130) 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, that is, that

$$S_s \triangleq \{e \in \mathbb{T}_\alpha(\Sigma)_s \mid (\exists! x \in X_s) (e, x) \in F_s\}$$

is the whole of $\mathbb{T}_\alpha(\Sigma)_s$, we apply Theorem 5.2. 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 $(S_s \mid s \in \Sigma_D)$ satisfies (IH_K) for each $(K : \sigma \rightarrow s) \in \Sigma_C$. By Lemma A.4(1), 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 $\bar{a} \#_\sigma \bar{e}$ it is the case that $K \bar{a}. \bar{e} \in S_s$.

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

$$(\forall \bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}) \bar{e} \in S^{|\sigma|} \Rightarrow (\exists! \bar{x} \in X^{|\sigma|}) (\bar{e}, \bar{x}) \in F^{|\sigma|}. \quad (131)$$

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

$$(\forall \bar{a} \in \mathbb{A}^\sigma, \bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|}, \bar{x} \in X^{|\sigma|}) \bar{a} \#_\sigma \bar{e} \ \& \ \bar{e} \in S^{|\sigma|} \ \& \ (\bar{e}, \bar{x}) \in F^{|\sigma|} \Rightarrow \bar{a} \#_\sigma \bar{x}. \quad (132)$$

So by rule (130), if $\bar{e} \in S^{|\sigma|}$ with $\bar{a} \#_\sigma \bar{e}$, 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 (130) to

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

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 A.5, $\tau_{\bar{a}, \bar{a}''} \cdot \bar{e} = \tau_{\bar{a}', \bar{a}''} \cdot \bar{e}'$, for some $\bar{a}'' \# (A, \bar{a}, \bar{e}, \bar{a}', \bar{e}')$. By (122), 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

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

and hence by (131), $\tau_{\bar{a}, \bar{a}''} \cdot \bar{x} = \tau_{\bar{a}', \bar{a}''} \cdot \bar{x}'$. Note that from (FCB_K), by Lemma A.4(2) we have

$$(\forall \bar{a} \in \mathbb{A}^\sigma) \bar{a} \# A \Rightarrow (\forall \bar{x} \in X^{|\sigma|}) \bar{a} \#_\sigma \bar{x} \Rightarrow \bar{a} \# f_K(\bar{a}, \bar{x})_\sigma. \quad (133)$$

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

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

From above $\tau_{\bar{a}, \bar{a}''} \cdot \bar{e} = \tau_{\bar{a}', \bar{a}''} \cdot \bar{e}'$; and by (125), $\tau_{\bar{a}, \bar{a}''} \cdot \bar{a} = \bar{a}'' = \tau_{\bar{a}', \bar{a}''} \cdot \bar{a}'$. So $f_K(\bar{a}, \bar{x})_\sigma = f_K(\bar{a}', \bar{x}')_\sigma = x$, as required.

B.2. Uniqueness of the Functions $\hat{f}_s \in \mathbb{T}_\alpha(\Sigma)_s \rightarrow_{fs} X_s$

Suppose $(\hat{f}'_s \in \mathbb{T}_\alpha(\Sigma)_s \rightarrow_{fs} X_s \mid s \in \Sigma_D)$ is also supported by A and satisfies property (47). To see that $\hat{f}'_s = \hat{f}_s$, it suffices to show that $S_s \triangleq \{e \in \mathbb{T}_\alpha(\Sigma)_s \mid \hat{f}'_s e = \hat{f}_s e\}$ is the whole of $\mathbb{T}_\alpha(\Sigma)$ for each $s \in \Sigma_D$. This follows almost immediately by Theorem 5.2; one just has to check that the subsets $S^{|\sigma|}$ defined from this $(S_s \mid s \in \Sigma_D)$ satisfy $S^{|\sigma|} \subseteq \{\bar{e} \in \mathbb{T}_\alpha(\Sigma)^{|\sigma|} \mid \hat{f}'^{|\sigma|} \bar{e} = \hat{f}^{|\sigma|} \bar{e}\}$.

ACKNOWLEDGMENT. I am grateful to James Cheney, Murdoch Gabbay, Michael Norrish, Mark Shinwell and Christian Urban for their many contributions to the subject of this article which is a revised and much extended version of Pitts [2005].

REFERENCES

- AYDEMIR, B. E., BOHANNON, A., FAIRBAIRN, M., FOSTER, J. N., PIERCE, B. C., SEWELL, P., VYTINIOTIS, D., WASHBURN, G., WEIRICH, S., AND ZDANCEWIC, S. 2005. Mechanised metatheory for the masses: The POPLmark challenge. In *Proceedings of the 18th International Conference on Theorem Proving in*

- Higher Order Logics (TPHOLs 2005)*, J. Hurd and T. Melham, Eds. Lecture Notes in Computer Science, vol. 3603. Springer-Verlag, New York, 50–65. www.cis.upenn.edu/group/proj/plclub/mmm/.
- BARENDREGT, H. P. 1984. *The Lambda Calculus: Its Syntax and Semantics*, revised ed. North-Holland.
- BARENDREGT, H. P. 1992. Lambda calculi with types. In *Handbook of Logic in Computer Science*, S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, Eds. Vol. 2. Oxford University Press, 117–309.
- BERGER, U., EBERL, M., AND SCHWICHTENBERG, H. 2003. Term rewriting for normalization by evaluation. *Inf. Comput.* 183, 19–42.
- BERGER, U., AND SCHWICHTENBERG, H. 1991. An inverse of the evaluation functional for typed λ -calculus. In *Proceedings of the 6th Annual Symposium on Logic in Computer Science*. IEEE Computer Society Press, Los Alamitos, CA, 203–211.
- CHENEY, J. 2004. Nominal logic programming. Ph.D. dissertation, Cornell University.
- CHURCH, A. 1940. A formulation of the simple theory of types. *J. Symb. Logic* 5, 56–68.
- DE BRUIJN, N. G. 1972. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indag. Math.* 34, 381–392.
- DYBJER, P. 2000. A general formulation of simultaneous inductive-recursive definitions in type theory. *J. Symb. Logic* 65, 2.
- DYBJER, P., AND FILINSKI, A. 2002. Normalization and partial evaluation. In *Applied Semantics, Advanced Lectures*, G. Barthe, P. Dybjer, and J. Saraiva, Eds. Lecture Notes in Computer Science, Tutorial, vol. 2395. Springer-Verlag, New York, 137–192. (International Summer School, APPSEM 2000, Caminha, Portugal, September 9–15, 2000).
- FIORE, M. P. 2002. 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 (PPDP 2002)*, (Pittsburgh, PA), ACM, New York, 26–37.
- FIORE, M. P., PLOTKIN, G. D., AND TURI, D. 1999. Abstract syntax and variable binding. In *Proceedings of the 14th Annual Symposium on Logic in Computer Science*. IEEE Computer Society Press, Los Alamitos, CA, 193–202.
- GABBAY, M. J. 2000. A theory of inductive definitions with α -equivalence: Semantics, implementation, programming language. Ph.D. dissertation, University of Cambridge.
- GABBAY, M. J. 2002. FM-HOL, a higher-order theory of names. In *Workshop on Thirty Five years of Automath, Informal Proceedings*, F. Kamareddine, ed. Heriot-Watt University, Edinburgh, Scotland.
- GABBAY, M. J. 2006. A general mathematics of names in syntax. *Math. Struct. Comput. Sci.* to appear.
- GABBAY, M. J., AND PITTS, A. M. 2002. A new approach to abstract syntax with variable binding. *Form. Asp. Comput.* 13, 341–363.
- GORDON, A. D., AND MELHAM, T. 1996. Five axioms of alpha-conversion. In *Theorem Proving in Higher Order Logics, 9th International Conference*. Lecture Notes in Computer Science, vol. 1125. Springer-Verlag, New York, 173–191.
- GRIFFIN, T. G. 1988. Notational definition—A formal account. In *Proceedings of the 3rd Annual Symposium on Logic in Computer Science*. IEEE Computer Society Press, 372–383.
- GUNTER, C. A. 1992. Semantics of programming languages: Structures and techniques, Los Alamitos, CA. *Foundations of Computing*. MIT Press, Cambridge, MA.
- HONSELL, F., MICULAN, M., AND SCAGNETTO, I. 2001. An axiomatic approach to metareasoning on nominal algebras in HOAS. In *Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP 2001)*, Crete, Greece, July. F. Orejas, P. G. Spirakis, and J. Leeuwen, Eds. Lecture Notes in Computer Science, vol. 2076. Springer-Verlag, New York, 963–978.
- JECH, T. J. 1977. About the axiom of choice. In *Hand book of Mathematical Logic*. J. Barwise, ed. North-Holland, Amsterdam, The Netherlands, 345–370.
- NORRISH, M. 2004. Recursive function definition for types with binders. In *Theorem Proving in Higher Order Logics, 17th International Conference*. Lecture Notes in Computer Science, vol. 3223. Springer-Verlag, New York, 241–256.
- PFENNING, F., AND ELLIOTT, C. 1988. Higher-order abstract syntax. In *Proceedings of the ACM-SIGPLAN Conference on Programming Language Design and Implementation*. ACM, New York, 199–208.
- PITTS, A. M. 2003. Nominal logic, a first order theory of names and binding. *Inf. Comput.* 186, 165–193.
- PITTS, A. M. 2005. Alpha-structural recursion and induction (extended abstract). In *Theorem Proving in Higher Order Logics, 18th International Conference (TPHOLs 2005)*, Oxford UK, Aug. J. Hurd and T. Melham, Eds. Lecture Notes in Computer Science, vol. 3603. Springer-Verlag, New York, 17–34.
- PLOTKIN, G. D. 1990. An illative theory of relations. In *Situation Theory and its Applications, Volume 1*, R. Cooper, Mukai, and J. Perry, Eds. CSLI Lecture Notes, vol. 22. Stanford Univ., Stanford, CA, 133–146.

- PLOTKIN, G. D. 2004. A structural approach to operational semantics. *J. Logic Algeb. Prog.* 60–61, 17–139.
- POTTIER, F. 2005. An overview of Caml. In *Proceedings of the ACM SIGPLAN Workshop on ML* (Tallinn, Estonia). ACM, New York.
- SANGIORGI, D., AND WALKER, D. 2001. *The π -calculus: A Theory of Mobile Processes*. Cambridge University Press, Cambridge, UK.
- SCHÖPP, U. 2006. Names and binding in type theory. Ph.D. dissertation, University of Edinburgh.
- SCHÖPP, U., AND STARK, I. D. B. 2004. A dependent type theory with names and binding. In *Proceedings of the Symposium on Computer Science Logic (CSL04)*, (Karpacz, Poland). Lecture Notes in Computer Science, vol. 3210. Springer-Verlag, New York, 235–249.
- SHINWELL, M. R., AND PITTS, A. M. 2005. On a monadic semantics for freshness. *Theoret. Comput. Sci.* 342, 28–55.
- STOUGHTON, A. 1988. Substitution revisited. *Theoret. Comput. Sci.* 59, 317–325.
- URBAN, C., PITTS, A. M., AND GABBAY, M. J. 2004. Nominal unification. *Theoret. Comput. Sci.* 323, 473–497.
- URBAN, C., AND TASSON, C. 2005. Nominal techniques in Isabelle/HOL. In *Proceedings of the 20th International Conference on Automated Deduction (CADE-20)*, (Tallinn, Estonia, July). Lecture Notes in Computer Science, vol. 3632. Springer-Verlag, New York, 38–53.

RECEIVED SEPTEMBER 2005; REVISED MARCH 2006; ACCEPTED APRIL 2006