

Proof Pearl: De Bruijn Terms Really Do Work

Michael Norrish¹ and René Vestergaard²

¹ Canberra Research Lab, Nicta

`Michael.Norrish@nicta.com.au`

² School of Information Science, JAIST, Ishikawa, Japan

`vester@jaist.ac.jp`

Abstract. Placing our result in a web of related mechanised results, we give a direct proof that the de Bruijn λ -calculus (à la Huet, Nipkow and Shankar) is isomorphic to an α -quotiented λ -calculus. In order to establish the link, we introduce an “index-carrying” abstraction mechanism over de Bruijn terms, and consider it alongside a simplified substitution mechanism. Relating the new notions to those of the α -quotiented and the proper de Bruijn formalisms draws on techniques from the theory of nominal sets.

1 Introduction

Implementors and theorists alike have long been in de Bruijn’s debt [4]. Representing λ -terms without names is a relatively straightforward way of implementing binders internally in, e.g., theorem-proving systems such as Isabelle and HOL4. Complementing this, the same idea has made it possible to formalise a great many results about syntax that features binders.

However, we believe that the state-of-the-art for formalisation has moved on. Implementors may well keep de Bruijn terms in their systems’ internals, but it does now seem as if it is possible to mechanise important results about syntax without needing to prove “painful” and “non-obvious” lemmas about indices [11]. The emerging formalisation technologies are appealing for a range of reasons, but here we focus on just two: i) it is now possible to check that multiple formalisations do actually correspond and ii) “non-obvious” lemmas about indices turn out to be related to general nominal properties of de Bruijn terms.

In particular, we can confirm that de Bruijn terms are a model for the λ -calculus, and can do so in an elegant way. By way of contrast, see for example Appendix C of Barendregt [2], where the proofs are given as “implicit in de Bruijn [4]”; or the pioneering work by Shankar [15], where mechanised proofs of complex theorems establish a connection between *unquotiented* λ -terms and de Bruijn terms.

In fact, there are a number of isomorphism results already extant in the literature. In this paper, we summarise these, and provide a new link between de Bruijn terms and what we hope will be an ever-expanding web of results, and body of understanding.

2 Formalising Binding

In this section, we position our result in the space of relationships between ten different formalisations of the λ -calculus, following Figure 1. λ^C is Curry’s unquotiented λ -terms using variable names and fresh-naming substitution [3] (plus explicit α [7]). λ_α^{Hi} , due to Hindley [7], is a quotienting construction saying that an equivalence class reduces to another if there are two syntactic terms that witness the reduction, in this case in λ^C . λ^V is similar to λ^C but uses renaming-free substitution (and explicit α), as studied in Vestergaard [17]. λ_α^V is the “Hindley-quotient” of λ^V . λ_α^{Ho} is a quotient of Curry’s set-up but with substitution rather than reduction lifted from the syntactic to the quotient type, as studied in Homeier [8]. dB_2 is the set of well-formed, two-sorted de Bruijn terms in Gordon [5]. λ_α^G is the type derived from this, in Gordon and Melham [6]. λ_α^N is a quotient of the first order syntax, with substitution defined directly on the quotiented level, as in Norrish [12]. dB is the type of pure de Bruijn terms (with indices representing free and bound variables), as mechanised in Shankar [15], Huet [9] and Nipkow [11]. Finally, dB' is an intermediate formalism that we define here. It uses constructors that are not injective, hence the oval node: conceptually, dB' is syntactic sugar for dB that looks like λ_α^N .




Fig. 1. λ -calculus formalisations and their directly-proved relationships; dashed, dark lines are the results in this paper. Double lines are isomorphisms; arrows α -quotienting.

2.1 Named Abstraction

We use A^{var} to refer to the “raw” syntax of the λ -calculus, that is the free algebra generated by the recursion equation

$$A^{\text{var}} \cong V + A^{\text{var}} \times A^{\text{var}} + V \times A^{\text{var}} \tag{1}$$

with V some countably-infinite set of variable names. We use A_α to refer to the quotient of A^{var} with respect to α -equivalence. The various λ_α^X are reduction systems defined over A_α , while λ^V and λ^C are defined directly over A^{var} .

Table 1. Signatures for formalisms with named abstraction, see Figure 1. The α -collapsed types below a dashed line are bootstrapped from the preceding “raw” type (λ^V/λ^C). For λ_α^G , substitution is bootstrapped from dB_2 . See the text for the [proofs: ...]-comments.

λ^V [proofs: $=_{\Lambda^{\text{var}}}$]	substitution: $\Lambda^{\text{var}} \times V \times \Lambda^{\text{var}} \rightarrow \Lambda^{\text{var}}$ reduction: $\Lambda^{\text{var}} \times \Lambda^{\text{var}}$
λ_α^V	reduction: $\Lambda_\alpha \times \Lambda_\alpha$
λ^C [proofs: $=_\alpha$]	substitution: $\Lambda^{\text{var}} \times V \times \Lambda^{\text{var}} \rightarrow \Lambda^{\text{var}}$ reduction: $\Lambda^{\text{var}} \times \Lambda^{\text{var}}$
λ_α^{Hi}	reduction: $\Lambda_\alpha \times \Lambda_\alpha$
λ_α^{Ho}	substitution: $\Lambda_\alpha \times V \times \Lambda_\alpha \rightarrow \Lambda_\alpha$ reduction: $\Lambda_\alpha \times \Lambda_\alpha$
λ_α^G	substitution: $\left\{ \begin{array}{l} (dB_2 \times V \times dB_2 \rightarrow dB_2) \\ \Lambda_\alpha \times V \times \Lambda_\alpha \rightarrow \Lambda_\alpha \end{array} \right.$ reduction: $\Lambda_\alpha \times \Lambda_\alpha$
λ_α^N	substitution: $\Lambda_\alpha \times V \times \Lambda_\alpha \rightarrow \Lambda_\alpha$ reduction: $\Lambda_\alpha \times \Lambda_\alpha$

There are seven different types identified in Figure 1 that use names under abstractions. Though intuition (and, as we shall see, mechanised results) assure us that they are all the “same”, their constructions are all different to varying degrees. The differences are most clearly seen in the types’ definitions of substitution, see Table 1. At the “raw” level, λ^V defines substitution to be “renaming free”, as used to be the tradition in logic books: where a substitution would cause a free variable to become captured, the substitution is instead defined to be the identity function, with the implicit (and satisfiable) requirement that actual uses avoid such situations. This approach finesses many of the problems of λ^C , where an everywhere-correct capture-avoiding substitution mandates fresh-naming (via either simultaneous substitution, as in Homeier [8], or by well-founded recursion on term sizes) and, thus, only allows us to reason “up to α ” for most properties [7]. Alternatively, and with substitution defined for Λ^{var} , the type itself can be α -quotiented. In the case of λ_α^{Ho} , Homeier lifts the substitution function and proceeds to define β -reduction at the quotiented type. In the case of λ_α^{Hi} and λ_α^V , Hindley and Vestergaard lift β -reductions from Λ^{var} [7,17].

In λ_α^G , substitution is defined with reference to the substitution function in dB_2 . In λ_α^N , substitution is defined directly on Λ_α , using its “native” notion of primitive recursion, as in Norrish [12] and Pitts [14].

2.2 De Bruijn's Anonymous Abstraction, Using Indices

The type of de Bruijn terms is the free algebra generated by the recursion equation

$$\mathbf{dB} \cong \mathbb{N} + \mathbf{dB} \times \mathbf{dB} + \mathbf{dB} \quad (2)$$

We will use constructors \mathbf{dV} , \mathbf{dAPP} , and \mathbf{dABS} to correspond to the three cases in the equation above, with the idea that, e.g., $\mathbf{dABS} (\mathbf{dV} 0)$ is the *identity* function: “ $\lambda x.x$ ”, and $\mathbf{dABS} (\mathbf{dABS} (\mathbf{dAPP} (\mathbf{dV} 1) (\mathbf{dV} 0)))$ is *apply*: “ $\lambda f.\lambda x.f x$ ”. We will use t, u, z to range over values in \mathbf{dB} , while M, N will range over values in Λ_α .

Nipkow [11] (following Huet [9]) defines two functions by primitive recursion over this type (see Figure 2). The first is \mathbf{lift} , which adds one to the value of all free indices that are at least equal to the provided bound. (Shankar [15] calls his version of this function \mathbf{BUMP} ; it is slightly different because he has the first bound index of an abstraction be 1, not 0.)

```

lift (dV i) k      = if i < k then dV i
                   else dV (i + 1)
lift (dAPP t u) k = dAPP (lift t k) (lift u k)
lift (dABS t) k   = dABS (lift t (k + 1))

nsub z k (dV i)   = if k < i then dV (i - 1)
                   else if i = k then z
                   else dV i
nsub z k (dAPP t u) = dAPP (nsub z k t) (nsub z k u)
nsub z k (dABS t)  = dABS (nsub (lift z 0) (k + 1) t)

```

Fig. 2. Functions from Nipkow's formalisation of de Bruijn terms [11]

Second is substitution, which we call \mathbf{nsub} here.¹ In order to make their proofs simpler, Nipkow, Huet and Shankar all define a notion of substitution that simultaneously performs a substitution and decrements the indices of the other free variables that appear in the term in which the substitution is occurring. This reflects the fact that the substitution is intended to be used for defining β -contraction, where a \mathbf{dABS} constructor disappears in the process. As usual, the considered β -reduction relation is given as β -contraction plus congruence rules.

Definition 1 (Nipkow [11]). \rightarrow_a is the least relation satisfying the following rules:

$$\frac{}{\mathbf{dAPP} (\mathbf{dABS} t) u \rightarrow_a \mathbf{nsub} u 0 t}$$

$$\frac{t \rightarrow_a u}{\mathbf{dAPP} t z \rightarrow_a \mathbf{dAPP} u z} \quad \frac{t \rightarrow_a u}{\mathbf{dAPP} z t \rightarrow_a \mathbf{dAPP} z u}$$

$$\frac{t \rightarrow_a u}{\mathbf{dABS} t \rightarrow_a \mathbf{dABS} u}$$

¹ We have altered the order of arguments to be compatible with the other mechanisations: the value being substituted comes first, followed by the index being substituted out, followed by the term in which the substitution is occurring.

2.3 Existing Results

We say that two types of the form we consider, e.g., $\langle X, \rightarrow_x \rangle$, $\langle Y, \rightarrow_y \rangle$, are isomorphic if there exists a bijection, $f : X \rightarrow Y$, that is homomorphic

$$\forall x_1, x_2. x_1 \rightarrow_x x_2 \Leftrightarrow f(x_1) \rightarrow_y f(x_2) \quad (3)$$

In the case of α -quotients, the mappings between the types are surjective and have the same kernels. By the appropriate adaptation of the First Isomorphism Theorem of Algebra,² this means that the various range types of each “raw” type are pairwise isomorphic. Once it has been established that the α -equivalences of λ^V and λ^C coincide, this means that the isomorphism of λ^V and λ^C implies that all the range types, i.e., dB and the λ_α^X , are pairwise isomorphic. (In [13], we further show that a broad class of properties is also shared between the elements of the “raw” syntax and the range types.)

As noted, the λ^V -to- λ_α^V and λ^C -to- λ_α^{Hi} α -quotients of the figure are underpinned by an epimorphism (i.e., (3) plus surjectivity) by construction [7,17]. In the remaining “raw-to-quotient” cases, property (3) plus surjectivity have been verified by Norrish [12] (λ^V -to- λ_α^G , λ^V -to- λ_α^N) and Homeier [8] (λ^C -to- λ_α^{Ho} , implicitly).

Vestergaard has mechanised the isomorphism at the base of the figure, between λ^V and λ^C (including their respective α -relations), in Coq. This follows the proof in Vestergaard [17].

The isomorphism at the named α -equivalent abstractions level between λ_α^G and λ_α^N is between the Gordon-Melham terms and the terms formed by quotienting the raw syntax (Λ^{var}). This isomorphism is uninteresting (and implicitly demonstrated in [12]) because the constructors and definitions in each type are immediate images of each other.

The isomorphism between dB_2 and λ_α^G is between the well-formed terms in Gordon’s two-sorted de Bruijn type [5], and the terms of Gordon and Melham’s subsequent paper [6]. This connection is a bijection by construction: the latter type is established using the HOL principle of type definition, which permits new types to be defined isomorphic to non-empty subsets of old types. Further, the bijection is behaviour-respecting because substitution in the new type is defined by following the bijection back to the old type, performing the substitution there and then returning.

The result most similar to that in this paper is Shankar’s link between dB and λ^V [15]. (Strictly speaking Shankar’s λ^V is not quite the same as the one used by Vestergaard [17]: where Vestergaard’s substitution will go wrong, it returns the term being substituted into unchanged; in the same situation, Shankar returns the result of a recursive descent with free variables incorrectly captured. In neither case does this detail matter: references to both substitution functions are set about with side conditions ensuring they are applied only in situations that make sense.)

Shankar represents the bijection between the two types with a function from λ^V to dB (called `TRANSLATE`) and a function `UNTRANS` for the opposite direction.

² The result has been verified in Coq by Vestergaard.

Neither of these can work directly on the values from the types, but need to be accompanied by contextual information. For example, the statement that TRANSLATE undoes an UNTRANS is given as

117. **Theorem.** UNTRANS-TRANS (rewrite):
 (IMPLIES (AND (TERMP X) (LEQ (FIND-M X N) M))
 (EQUAL (TRANSLATE (UNTRANS X M N)
 (BINDINGS M N))
 X))

The other identity is even more complicated, partly because the equality is actually α -equivalence. The subsequent statements that β -reductions in each type match up are similarly complex. (They are used to good effect: the Church-Rosser result for dB is transferred back to λ^V .)

As we shall see, our own directly-established connection can be expressed more naturally because we build an isomorphism to λ_α^N instead of λ^V . We also provide an isomorphism result for η -reduction (in Section 4.4).

3 Nominal Reasoning

In this section we provide a brief introduction to the nominal technology that has an important rôle to play in the isomorphism proofs to come. (For more on this, see for example, Pitts [14].)

Permutations. Nominal reasoning proceeds over sets that support a permutation action satisfying certain conditions. We call such sets, *nominal sets*.³ At the “base level” permutations are defined in terms of their effect on strings.

Definition 2 (Permutations Applied to Strings). *A permutation (written π , π_1 etc) is a bijection on strings that is the identity on all but finitely many strings. We represent a permutation as a list of pairs of strings, and define the action of a permutation π on a string s (written $\pi(s)$) by primitive recursion on the structure of the list:*

$$\begin{aligned} [](s) &= s \\ ((x, y) :: \pi)(s) &= \text{if } x = \pi(s) \text{ then } y \text{ else if } y = \pi(s) \text{ then } x \\ &\quad \text{else } \pi(s) \end{aligned}$$

If permutation π is the singleton list $[(x, y)]$, then π is also written (xy) .

For a type other than strings to be a nominal set, the type must support a permutation action (written $\pi \cdot x$, for x in the putative nominal set) satisfying the following conditions:

$$\begin{aligned} [] \cdot t &= t \\ \pi_1 \cdot (\pi_2 \cdot t) &= (\pi_1 \circ \pi_2) \cdot t \\ (\forall s. \pi_1(s) = \pi_2(s)) &\Rightarrow \pi_1 \cdot t = \pi_2 \cdot t \end{aligned} \tag{4}$$

³ Pitts only calls sets nominal if each element also has *finite support*, for which see below.

As our permutations are actually lists, permutation composition ($\pi_1 \circ \pi_2$ in the second condition) is actually list concatenation. This is also the reason why the third condition tests for extensional equality (over strings) on the permutations π_1 and π_2 .

Examples

- The natural numbers and the booleans are both nominal sets, with a permutation action that is everywhere the identity.
- If A is a nominal set, then so too are lists of values drawn from A . The permutation action over a list value is the permutation of each element.
- If A and B are nominal sets, then so too are functions from A to B . The permutation action is defined to be

$$\pi \cdot (f : A \rightarrow B) = \lambda(x : A). \pi \cdot (f (\pi^{-1} \cdot x))$$

It follows that $\pi \cdot (f(x)) = (\pi \cdot f)(\pi \cdot x)$

Support. With a permutation action defined, a nominal set immediately gives rise to a notion of “support”.

Definition 3. *We say that the set A supports a value t if*

$$\forall x y \notin A. (xy) \cdot t = t$$

If there is a finite set that supports a value t , then there is a smallest such supporting set, which we refer to as the support of t (written $\text{supp}(t)$). For types such as λ -calculus terms, support corresponds to the free variables of the term.

When a value has empty support, it follows that $(xy) \cdot t = t$ for all possible x and y . When the value is a function, this gives us a nice result about the distribution of permutations over applications:

$$(xy) \cdot (f(v)) = ((xy) \cdot f)((xy) \cdot v) = f((xy) \cdot v)$$

This result extends to arbitrary permutations, so that $\pi \cdot f(x) = f(\pi \cdot x)$. Better yet, if the function is actually a relation or predicate (a function whose range is the boolean type), then the outer π disappears and we have that $R (\pi \cdot x_1) \dots (\pi \cdot x_n) \Leftrightarrow R x_1 \dots x_n$. Such a predicate or relation is said to be *equivariant*.

4 Isomorphisms and De Bruijn Terms

Our goal is to show that **dB** is isomorphic to λ_α^N , *i.e.*, a formalism where reductions occur between α -equivalence classes of the raw syntax. The mechanisation was performed by Norrish using the HOL4 proof assistant.

At the base level, a bijection between the terms of the quotiented λ -calculus (Λ_α) and **dB** requires a bijection between strings and natural numbers because

the existing mechanisation of Λ_α uses strings for its variables. Though occasionally awkward, the fact that the two sorts of variables are of different types is a useful sanity check when proofs are being performed. In what follows, s , s' , v and v' will range over strings, and i , j and k over natural numbers. The bijections between the types will be written with a hat, so that \hat{i} is the string corresponding to number i , while \hat{s} is the number corresponding to string s .

4.1 An Intermediate Formalism: dB'

We define dB' to match λ_α^N , both in its construction and in its reduction relations. In particular, abstraction in dB' will involve an identifier in the style of “ $\lambda 2.2$ ” (for the *identity* function). This new constructor is necessarily not injective (as, otherwise, we would distinguish, e.g., “ $\lambda 2.2$ ” and “ $\lambda 1.1$ ”). Abstraction in λ_α^N is also not injective but, whereas the λ_α^N -construction is based on a quotient of Λ^{var} , the non-injectivity of the new λ -constructor in dB' follows directly from its definition as syntactic sugar for the existing dABS constructor in dB. Though a separate oval in Figure 1, dB' is based on the same underlying HOL type as dB. The difference arises from its use of different constructors to “cover” the type, and a different reduction relation.

We begin by defining dB-substitution without the index-decrementing used in nsub. We shall see that this function is isomorphic to substitution over Λ_α , while nsub has no natural reading in Λ_α .

Definition 4 (A clean notion of substitution for dB).

$$\begin{aligned} \text{sub } z k (\text{dV } i) &= \text{if } i = k \text{ then } z \text{ else dV } i \\ \text{sub } z k (\text{dAPP } t u) &= \text{dAPP } (\text{sub } z k t) (\text{sub } z k u) \\ \text{sub } z k (\text{dABS } t) &= \text{dABS } (\text{sub } (\text{lift } z 0) (k + 1) t) \end{aligned}$$

The (non-injective) dLAM-function that performs abstraction over a particular index can now be defined as follows.

Definition 5 (Abstraction over an index).

$$\text{dLAM } i t = \text{dABS } (\text{sub } (\text{dV } 0) (i + 1) (\text{lift } t 0))$$

All dABS terms can be seen as abstractions over an arbitrary fresh index.

Lemma 6 (dABS terms as dLAM terms). *If i is not among the free indices of dABS t , then there exists a t_0 such that $\text{dABS } t = \text{dLAM } i t_0$. By the definition of dLAM and the injectivity of dABS we further know that*

$$t = \text{sub } (\text{dV } 0) (i + 1) (\text{lift } t_0 0)$$

Proof. Prove by structural induction on t that there exists a t_0 such that $t = \text{sub } (\text{dV } n) i (\text{lift } t_0 n)$, where i is again suitably fresh. \square

Now we define \rightarrow_d , the β -reduction relation for **dB**. It mimics the traditional statement in the λ -calculus for both contraction and congruence rules.

Definition 7 (dB' β -reduction).

$$\frac{}{\text{dAPP } (\text{dLAM } i t) u \rightarrow_d \text{ sub } u i t}$$

$$\frac{t \rightarrow_d u}{\text{dAPP } t z \rightarrow_d \text{ dAPP } u z} \quad \frac{t \rightarrow_d u}{\text{dAPP } z t \rightarrow_d \text{ dAPP } z u}$$

$$\frac{t \rightarrow_d u}{\text{dLAM } i t \rightarrow_d \text{ dLAM } i u}$$

4.2 dB' vs λ_α^N : Nominal Reasoning with Indices

Thus far we have been working entirely with the **dB** type, where functions are trivial to define by primitive recursion. Now we are faced with a choice: to define a function from **dB** into Λ_α , or *vice versa*. The former choice is less appealing because a choice of arbitrary fresh name would need to be made in the recursive **dABS**-clause. One might imagine writing something like

$$f (\text{dABS } t) = \text{LAM } (\text{fresh}(t)) (f (\text{sub } (\text{fresh}(t)) 0 t))$$

where $\text{fresh}(t)$ generates an index fresh for t . But this clause is not primitive recursive, so one would, e.g., have to write an f that took an additional environment parameter specifying how the bound variables encountered so far should be replaced. Instead, we take as our objective a bijective $f : \Lambda_\alpha \rightarrow \text{dB}$, satisfying the following equations

$$\begin{aligned} f (\text{VAR } v) &= \text{dV } \hat{v} \\ f (\text{APP } M N) &= \text{dAPP } (f M) (f N) \\ f (\text{LAM } v M) &= \text{dLAM } \hat{v} (f M) \end{aligned} \tag{5}$$

The existence of such an f is not immediate because Λ_α is not a free algebra (**LAM** is not injective). However, suitable recursion principles for this type do exist, as shown in Pitts [14].

de Bruijn Terms are a Nominal Set. In order to use recursion principles for Λ_α , it remains to be shown that the range of the desired function is a *nominal set*, as in Section 3. A subtlety derives from the fact that we (necessarily must) send free variables to suitably free indices. Permutation on free variables therefore must be reflected on free indices. Conversely, bound indices are what they are and must be left untouched.

Thus we are required to define a permutation action for de Bruijn terms, and it can not be a trivial one (everywhere the identity, say). Inevitably, the complexity arises in the **dABS** clause of the definition. When a permutation passes through an abstraction, the free variables it is acting on are now at different index positions, which means that the permutation has to be similarly adjusted.

Definition 8 (Permutation for dB). Let *inc* (“increment”) over permutations be

$$\begin{aligned} inc [] &= [] \\ inc ((s_1, s_2) :: tl) &= (\widehat{s_1 + 1}, \widehat{s_2 + 1}) :: inc tl \end{aligned}$$

The action of a permutation π on a term in dB (written $\pi \cdot t$) is

$$\begin{aligned} \pi \cdot (\mathbf{dV} \ i) &= \mathbf{dV} \ (\widehat{\pi(i)}) \\ \pi \cdot (\mathbf{dAPP} \ t \ u) &= \mathbf{dAPP} \ (\pi \cdot t) \ (\pi \cdot u) \\ \pi \cdot (\mathbf{dABS} \ t) &= \mathbf{dABS} \ (inc(\pi) \cdot t) \end{aligned}$$

It is straightforward to show that this definition satisfies the requirements to make dB a nominal set. The support for a value in type dB is as one might expect: its set of free indices converted into strings *via* the hat-isomorphism.

Relating dLAM and Permutation. In order to admit (5), we must also characterise the relationship between dLAM (which appears on the RHS of (5)), and dB-permutation. As dLAM is defined in terms of **lift** and **sub**, analogous results are required of these constants.

We first characterise the relationship between **lift** and permutation. This requires the definition of an *inc* that is parameterised by the index below which variables should not be touched, with the already-defined *inc* being *inc*₀ and with *inc*_{*i*} being the obvious generalisation.

Lemma 9 (Permutations and lift).

$$\mathbf{lift} \ (\pi \cdot t) \ n = (inc_n(\pi)) \cdot (\mathbf{lift} \ t \ n)$$

Proof. By induction on the structure of *t*. The result first needs

$$inc_0 \ (inc_n \ \pi) = inc_{n+1} \ (inc_0 \ \pi)$$

to be shown. This latter is just the sort of tedious result that the POPLmark authors inveigh against so convincingly in [1]. (Note that it is also an equality on the representing lists of pairs, not just extensionally in terms of the permutations’ effect on strings.) □

This lemma leads to important results about **sub**, dLAM and \rightarrow_d :

Theorem 10 (sub and dLAM have empty support; \rightarrow_d is equivariant).

$$\begin{aligned} \pi \cdot (\mathbf{sub} \ t \ j \ u) &= \mathbf{sub} \ (\pi \cdot t) \ (\widehat{\pi(j)}) \ (\pi \cdot u) \\ \pi \cdot (\mathbf{dLAM} \ i \ t) &= \mathbf{dLAM} \ (\widehat{\pi(i)}) \ (\pi \cdot t) \\ \pi \cdot t \rightarrow_d, \ \pi \cdot u &\Leftrightarrow t \rightarrow_d, \ u \end{aligned}$$

See Lemma 17 below for the contrast with **nsub**.

Proof. The substitution result follows from a structural induction on t . The result for **dLAM** is then immediate. The result for $\rightarrow_{\mathbf{d}}$, follows from a rule induction showing that $t \rightarrow_{\mathbf{d}'} u \Rightarrow \pi \cdot t \rightarrow_{\mathbf{d}} \pi \cdot u$. The fact that all permutations have inverses gives the other direction of the if-and-only-if. \square

Theorem 11 (dB is in bijection with Λ_α). *There exists a bijective $f : \Lambda_\alpha \rightarrow \mathbf{dB}$ with defining equations as in (5).*

Proof. The existence of f now follows once **dB** is established as a nominal set, and it is shown that

$$\text{supp}(\mathbf{dLAM } i t) = \text{supp}(t) \setminus \{i\}$$

This in turn relies on results specifying the free variables (or support) of **sub** and **lift** (the latter is from Nipkow).

Injectivity of f follows from a structural induction, and standard properties of substitution. Surjectivity of f follows from a complete induction on the size of the **dB** term onto which a value in Λ_α maps.⁴ \square

The Homomorphism between \mathbf{dB}' and λ_α^N . Given its construction, it is not surprising that **dB** with relation $\rightarrow_{\mathbf{d}'}$ is isomorphic to Λ_α and relation \rightarrow_β

Theorem 12.

$$M \rightarrow_\beta N \Leftrightarrow f(M) \rightarrow_{\mathbf{d}'} f(N)$$

Proof. Two straightforward rule inductions, one for each direction. Both results rely on the fact that

$$f(M[v := N]) = \text{sub}(f N) \hat{v}(f M)$$

which is proved by “BVC-compatible” structural induction (see [12]) on M . \square

4.3 The Homomorphism Between **dB** and **dB'**

Now we must grasp the nettle, and get to grips with Nipkow’s definition of $\rightarrow_{\mathbf{d}}$, and the accompanying **nsub** function. We will establish that $\rightarrow_{\mathbf{d}}$ and $\rightarrow_{\mathbf{d}'}$ are in fact the same relation, giving us the identity as a homomorphism between **dB** and **dB'**.

We begin by proving the following:

Lemma 13.

$$n \leq i \Rightarrow \text{sub } u \ i \ t = \text{nsub } u \ n \ (\text{sub } (\mathbf{dV } n) \ (i + 1) \ (\text{lift } t \ n))$$

Proof. By structural induction on t . \square

⁴ Complete induction allows us to assume the induction hypothesis for all smaller terms, not just those of size one smaller.

Using this, we can show that

$$\mathbf{dAPP} (\mathbf{dLAM} \ i \ t) \ u \ \rightarrow_{\mathbf{d}} \ \mathbf{sub} \ u \ i \ t$$

which is to say that Nipkow’s relation $\rightarrow_{\mathbf{d}}$ does indeed reduce $\rightarrow_{\mathbf{d}}$ -redexes in the correct way. If we expand the definition of \mathbf{dLAM} on the LHS, we see that the redex is

$$\mathbf{dAPP} (\mathbf{dABS} (\mathbf{sub} (\mathbf{dV} \ 0) (i + 1) (\mathbf{lift} \ t \ 0)) \ u)$$

Using Lemma 13 with $n = 0$ we get the desired identity. The same lemma is used to show the “converse” relation, that

$$\mathbf{dAPP} (\mathbf{dABS} \ t) \ u \ \rightarrow_{\mathbf{d}'} \ \mathbf{nsub} \ u \ 0 \ t$$

The congruence rules for \mathbf{dAPP} are the same in both relations, so our remaining task is to show that the different congruence rules for \mathbf{dABS} and \mathbf{dLAM} match up.

Abstraction Congruences Correspond. One of the two correspondences is to show that

$$t \rightarrow_{\mathbf{d}'} u \Rightarrow \mathbf{dABS} \ t \rightarrow_{\mathbf{d}'} \mathbf{dABS} \ u$$

In order to apply the abstraction rule for $\rightarrow_{\mathbf{d}'}$, we must show that the \mathbf{dABS} terms correspond to \mathbf{dLAM} terms. Using Lemma 6, we can pick a suitably fresh i , and show the existence of t_0 and u_0 such that

$$\begin{aligned} \mathbf{dABS} \ t &= \mathbf{dLAM} \ i \ t_0 \\ \mathbf{dABS} \ u &= \mathbf{dLAM} \ i \ u_0 \\ t &= \mathbf{sub} (\mathbf{dV} \ 0) (i + 1) (\mathbf{lift} \ t_0 \ 0) \\ u &= \mathbf{sub} (\mathbf{dV} \ 0) (i + 1) (\mathbf{lift} \ u_0 \ 0) \end{aligned}$$

We have $t \rightarrow_{\mathbf{d}'} u$, and want $t_0 \rightarrow_{\mathbf{d}'} u_0$, from which the abstraction rule for $\rightarrow_{\mathbf{d}'}$ would give us $\mathbf{dABS} \ t \rightarrow_{\mathbf{d}'} \mathbf{dABS} \ u$. The problem here is that our assumption is about a large, complicated term, and we want to conclude that the reduction occurs between sub-terms (t_0 and u_0). Standard proofs by rule induction would handle the reverse situation well (results of the form $t \rightarrow u \Rightarrow \mathbf{sub} \ w \ i \ t \rightarrow \mathbf{sub} \ w \ i \ u$), but our goal looks like the converse.

The Nominal “Aha!”. The solution to our problem is to realise that both the substitution and the lifting that occur in terms t and u are no more than variable permutations, and that they are in fact the same permutation applied to t_0 and u_0 . Then, we will express our assumption $t \rightarrow_{\mathbf{d}'} u$ as $\pi \cdot t_0 \rightarrow_{\mathbf{d}'} \pi \cdot u_0$ for some permutation π . Because $\rightarrow_{\mathbf{d}'}$ is equivariant (Theorem 10), this gives $t_0 \rightarrow_{\mathbf{d}'} u_0$ as desired.

That a substitution of a fresh variable into a term is a permutation is an easy induction. The interest comes in showing that a lifting is also a permutation. First we define a function *lift-pm* that constructs a “lifting permutation”. Given a lower bound ℓ and an upper bound m , *lift-pm*(ℓ, m) constructs the permutation that maps \hat{i} to $\widehat{i + 1}$ if $\ell \leq i \leq m$, maps $\widehat{m + 1}$ back to $\hat{\ell}$, and preserves the values of all other arguments.

Lemma 14 (Lifting is a permutation). *If n is at least as large as the largest free index in t , then*

$$\mathbf{lift} \ t \ m = (\mathbf{lift}\text{-}pm(m, n)) \cdot t$$

Proof. By structural induction on t . The **dABS** case requires another tedious result about index-maneuvering, namely that

$$\mathbf{inc}_0(\mathbf{lift}\text{-}pm(m, n)) = \mathbf{lift}\text{-}pm(m + 1, n + 1)$$

Thankfully, the result does not need to be generalised to cover the behaviour of \mathbf{inc}_i . \square

We can now prove that abstraction congruences for any equivariant reduction relation must correspond:

Theorem 15 (Abstraction Congruences Correspond). *If R is equivariant, that is $R(\pi \cdot t)(\pi \cdot u) \Leftrightarrow R t u$, then*

$$\frac{R t u}{R(\mathbf{dABS} \ t)(\mathbf{dABS} \ u)} \quad \Leftrightarrow \quad \frac{R t u}{R(\mathbf{dLAM} \ i \ t)(\mathbf{dLAM} \ i \ u)}$$

where one can read each inference rule as an implication universally quantified (over t , u and i) at the top level. The generality of the result is useful when we consider η -reduction.

Proof. Immediate, given Lemma 14, and Lemma 6. \square

Lemma 16 ($\rightarrow_{\mathbf{d}} \subseteq \rightarrow_{\mathbf{d}'}$).

$$t \rightarrow_{\mathbf{d}} u \Rightarrow t \rightarrow_{\mathbf{d}'} u$$

Proof. A rule induction over the definition of $\rightarrow_{\mathbf{d}}$. The redex case is handled by Lemma 13 (as already discussed). The **dAPP** congruence cases are immediate; the **dABS** congruence case is as just discussed (by Lemma 14, and Theorem 10). \square

The Other Inclusion. Showing $\rightarrow_{\mathbf{d}'} \subseteq \rightarrow_{\mathbf{d}}$, we know that $\rightarrow_{\mathbf{d}}$ can perform both $\rightarrow_{\mathbf{d}'}$'s redex-reductions (from Lemma 13), and the application congruences. To show the abstraction congruence case, we need to show that $\rightarrow_{\mathbf{d}}$ is equivariant in order to apply Theorem 15.

But the definition of $\rightarrow_{\mathbf{d}}$ involves **nsub**, so we have to characterise how **nsub** behaves when it interacts with a permutation. This prompts a rather incomprehensible lemma (note the odd way in which the permutation π does not affect the index i):

Lemma 17.

$$\pi \cdot (\mathbf{nsub} \ t \ i \ u) = \mathbf{nsub} \ (\pi \cdot t) \ i \ (\mathbf{inc}_i(\pi) \cdot u)$$

Proof. By structural induction on u . The \mathbf{dV} case requires a number of splits on the ordering possibilities ($<$, $=$, $>$) between the index of the variable and i . The \mathbf{dABS} case uses Lemma 9. \square

Now we can conclude that $\rightarrow_{\mathbf{d}}$ is equivariant (by rule induction), and follow immediately with

Lemma 18 ($\rightarrow_{\mathbf{d}'} \subseteq \rightarrow_{\mathbf{d}}$).

$$t \rightarrow_{\mathbf{d}'} u \Rightarrow t \rightarrow_{\mathbf{d}} u$$

Theorem 19 (Relations $\rightarrow_{\mathbf{d}}$ and $\rightarrow_{\mathbf{d}'}$ Coincide).

$$t \rightarrow_{\mathbf{d}'} u \Leftrightarrow t \rightarrow_{\mathbf{d}} u$$

Proof. Combining Lemmas 16 and 18. \square

Theorem 20 (Isomorphism). *The types \mathbf{dB} and λ_{α}^N , with their respective reduction relations, $\rightarrow_{\mathbf{d}}$ and \rightarrow_{β} , are isomorphic, as witnessed by the bijective f :*

$$M \rightarrow_{\beta} N \Leftrightarrow f(M) \rightarrow_{\mathbf{d}} f(N)$$

Proof. From Theorems 12 and 19. \square

4.4 Bonus: η -Reduction

Nipkow [11] also defines a notion of η -reduction for his terms. The congruence rules are standard. Writing $\rightarrow_{\mathbf{ne}}$ for “Nipkow’s η ”, the redex rule is

$$\frac{0 \notin \mathbf{dFV}(t)}{\mathbf{dABS}(\mathbf{dAPP} t (\mathbf{dV} 0)) \rightarrow_{\mathbf{ne}} \mathbf{nsub} u 0 t}$$

where the choice of u is immaterial. (Because 0 is not a free index of t , \mathbf{nsub} is used only to decrement the free indices of t , which is required as t moves out from underneath an abstraction.)

Using exactly the same techniques as before, it is easy to define a notion of reduction that better emulates the definition in the λ -calculus. This relation (written $\rightarrow_{\mathbf{e}}$), has contraction rule

$$\frac{i \notin \mathbf{dFV}(t)}{\mathbf{dLAM} i (\mathbf{dAPP} t (\mathbf{dV} i)) \rightarrow_{\mathbf{e}} t}$$

It is straightforward to show that $\rightarrow_{\mathbf{e}}$ corresponds to the η -reduction rule in Λ_{α} . Thanks to Theorem 15, showing that $\rightarrow_{\mathbf{ne}}$ and $\rightarrow_{\mathbf{e}}$ correspond only requires us to show the equivalence of the redex rules, and that $\rightarrow_{\mathbf{ne}}$ and $\rightarrow_{\mathbf{e}}$ are equivariant. These proofs are mainly straightforward. We use some of the lemmas stated in Nipkow, among them

$$\mathbf{nsub} u n (\mathbf{lift} t n) = t$$

We also state and prove the similar

$$n \notin \mathbf{dFV}(t) \Rightarrow \mathbf{lift} (\mathbf{nsub} u n t) n = t$$

5 Conclusion

We have proved what everyone knew (or hoped) to be true from the outset. However, the proofs we have described are not entirely trivial, suggesting that they were indeed worth doing. In fact, the proofs are only made in the least palatable by using insights from the theory of nominal sets.

This theory helps in two important ways: it makes it possible to define the bijection f from Λ_α to \mathbf{dB} in a natural style; and it helps us perform many of the resulting proofs. In particular, the proofs are helped greatly by the realisation that `lift` is a permutation. To best exploit this, we needed to prove that our reduction relations \rightarrow_a and $\rightarrow_{a'}$ were equivariant, which in turn prompted the curious Lemma 17.

We hope that this work encourages others to extend the web of equivalence results in Figure 1, which we believe to be a valuable summary of existing work in its own right. In the near future, we hope to consider, for example, the two-sorted calculus of McKinna and Pollack [10], and the weak HOAS style construction of the λ -calculus in Urban and Tasson [16].

Finally, we hope this work will go some way towards bringing about the day when reasoning about de Bruijn terms is never again necessary.

Acknowledgements. NICTA is funded by the Australian Government's *Backing Australia's Ability* initiative, in part through the Australian Research Council. We would like to thank James Cheney for valuable discussions in this area, and comments on this paper in particular.

Availability. The sources mechanising this paper's proofs are distributed as part of the HOL4 system, available from <http://hol.sourceforge.net>. (The proofs for β are part of the latest public release (January 2007); the η proofs are in the publicly accessible CVS repository.)

References

1. Aydemir, B.E., Bohannon, A., Fairbairn, M., Nathan Foster, J., Pierce, B.C., Sewell, P., Vytiniotis, D., Washburn, G., Weirich, S., Zdancewic, S.: Mechanized metatheory for the masses: the POPLMARK. In: Hurd, J., Melham, T. (eds.) TPHOLs 2005. LNCS, vol. 3603, Springer, Heidelberg (2005)
2. Barendregt, H.P.: The Lambda Calculus: its Syntax and Semantics. Studies in Logic and the Foundations of Mathematics, vol. 103. revised edn. Elsevier, Amsterdam (1984)
3. Curry, H.B., Feys, R.: Combinatory Logic. North-Holland, Amsterdam (1958)
4. de Bruijn, N.G.: Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser Theorem. *Indag. Math.* 34, 381–392 (1972)
5. Gordon, A.D.: A mechanisation of name-carrying syntax up to alpha-conversion. In: Joyce, J.J., Seger, C.-J.H. (eds.) HUG 1993. LNCS, vol. 780, pp. 413–425. Springer, Heidelberg (1994)

6. Gordon, A.D., Melham, T.: Five axioms of alpha conversion. In: von Wright, J., Harrison, J., Grundy, J. (eds.) TPHOLs 1996. LNCS, vol. 1125, pp. 173–190. Springer, Heidelberg (1996)
7. Roger Hindley, J.: The Church-Rosser Property and a Result in Combinatory Logic. PhD thesis, University of Newcastle upon Tyne (1964)
8. Homeier, P.: A proof of the Church-Rosser theorem for the lambda calculus in higher order logic. In: Boulton, R.J., Jackson, P.B. (eds.) TPHOLs 2001. LNCS, vol. 2152, pp. 207–222. Springer, Heidelberg (2001)
9. Huet, G.P.: Residual theory in lambda-calculus: a formal development. *Journal of Functional Programming* 4(3), 371–394 (1994)
10. McKinna, J., Pollack, R.: Some lambda calculus and type theory formalized. *Journal of Automated Reasoning* 23(3–4), 373–409 (1999)
11. Nipkow, T.: More Church-Rosser proofs (in Isabelle/HOL). *Journal of Automated Reasoning* 26, 51–66 (2001)
12. Norrish, M.: Mechanising λ -calculus using a classical first order theory of terms with permutations. *Higher-Order and Symbolic Computation* 19, 169–195 (2006)
13. Norrish, M., Vestergaard, R.: Structural preservation and reflection of diagrams. Technical Report IS-RR-2006-011, JAIST (2006)
14. Pitts, A.M.: Alpha-structural recursion and induction. *Journal of the ACM* 53(3), 459–506 (2006)
15. Shankar, N.: A mechanical proof of the Church-Rosser theorem. *Journal of the ACM* 35(3), 475–522 (1988)
16. Urban, C., Tasson, C.: Nominal techniques in Isabelle/HOL. In: Nieuwenhuis, R. (ed.) *Automated Deduction – CADE-20*. LNCS (LNAI), vol. 3632, pp. 38–53. Springer, Heidelberg (2005)
17. Vestergaard, R.: The Primitive Proof Theory of the λ -Calculus. PhD thesis, School of Mathematical and Computer Sciences, Heriot-Watt University (2003)