

Observational Equality, Now!

Thorsten Altenkirch Conor McBride Wouter Swierstra

University of Nottingham, UK

{txa,ctm,wss}@cs.nott.ac.uk

Abstract

This paper has something new and positive to say about propositional equality in programming and proof systems based on the Curry-Howard correspondence between propositions and types. We have found a way to present a propositional equality type

- which is *substitutive*, allowing us to reason by replacing equal for equal in propositions;
- which reflects the *observable* behaviour of values rather than their construction: in particular, we have *extensionality*—functions are equal if they take equal inputs to equal outputs;
- which retains *strong* normalisation, *decidable* typechecking and *canonicity*—the property that closed normal forms inhabiting datatypes have canonical constructors;
- which allows inductive data structures to be expressed in terms of a standard characterisation of *well-founded trees*;
- which is presented *syntactically*—you can implement it directly, and we are doing so—this approach stands at the core of Epigram 2;
- which you can play with now: we have simulated our system by a shallow embedding in Agda 2, shipping as part of the standard examples package for that system [21].

Until now, it has always been necessary to sacrifice some of these aspects. The closest attempt in the literature is Altenkirch’s construction of a setoid-model for a system with canonicity and extensionality on top of an intensional type theory with proof-irrelevant propositions [4]. Our new proposal simplifies Altenkirch’s construction by adopting McBride’s *heterogeneous* approach to equality [19].

Categories and Subject Descriptors F.4.1 [Mathematical Logic]: Lambda calculus and related systems; D.3.1 [Formal Definitions and Theory]: Semantics

General Terms Languages, Theory

Keywords Type Theory, Equality

1. Introduction

Equations are ubiquitous in mathematical reasoning, and reasoning about programs is no exception. Moreover, notions of equality

play a crucial internal rôle in the programs of any language where datatypes are indexed in some way. Case analysis relies on solving the equation between the type of the scrutinee and the possible types returned by each constructor, specialising the return type of each branch accordingly. For a standard example, consider length-indexed lists. We need to know that a length 4 list cannot be made by the ‘nil’ constructor because 4 does not equal 0, but can be made by ‘cons’, given a head element and a tail of length 3.

Further, computation within types is now commonplace: appending two length 2 lists yields a length $2 + 2$ list, which should somehow be regarded also as a length 4 list. Further still, we sometimes need to exploit algebraic properties of type-level expressions beyond their symbolic evaluation: we may wish to use a length $x + y$ list where length $y + x$ is expected.

Of course, dependent type theories have always had this issue, but it is becoming increasingly prevalent in programming languages, whether they have ‘full-blown’ value-dependency, like Cayenne [5], Epigram [20, 11] and Agda [21], or a separate static language of indices, as in DML [29], Ω mega [25] and Haskell with Generalized Algebraic Datatypes [24].

This paper concerns the strength of equational reasoning available for both behind-the-scenes constraint solving and the explicit manipulation of type-level expressions in Martin-Löf type theories, hence in proof systems and programming languages based on the “propositions-as-types” principle. These tend to be divided into two camps:

extensional type theories (ETTs, as implemented in NuPRL [7]) identify the *definitional* equality (as used in typechecking) with the *propositional* equality (expressed as a type and used for reasoning), resulting in powerful systems with undecidable typechecking;

intensional type theories (ITTs, as implemented in Coq, Agda and Epigram) separate these notions, keeping the definitional equality (and hence typechecking) *decidable*, but at some cost to the power of the propositional equality—until now.

In this paper, we examine an approach which delivers a ‘no compromise’ candidate: Observational Type Theory (OTT), with all the good properties of both. OTT is an intensional type theory with all the key computational properties we expect of such systems, but a notion of propositional equality that identifies values *up to observation*, rather than by construction. It is just as powerful for reasoning as the extensional theories.

This paper also records an attempt to code the phenomena of dependently typed programming in a simple and closed core type theory. In particular, recursive data are introduced uniformly via the so-called ‘W-types’, inhabited by higher-order encodings of *well-founded trees*. This encoding has always been feasible in ETTs, but it has never worked in an ITT—until now. The OTT approach has allowed us to make the encoding, and to discover its price.

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

PLPV’07, October 5, 2007, Freiburg, Germany.

Copyright © 2007 ACM 978-1-59593-677-6/07/0010...\$5.00

Our adventure has a strong practical component. We show how to *construct* an OTT by a shallow embedding in an existing intensional type theory, defined within the Agda 2 framework. By doing so, we expose the mechanics of the approach and gain evidence for key metatheoretical properties. In particular, we obtain at least a sketch proof for strong normalisation for expressions in general, and *canonicity* for closed expressions—the latter always compute to canonical values, as programmers might naturally hope. However, you can also download the OTT construction (now a standard Agda 2 example) and try out observational equality, now!

1.1 The Equality Dilemma

In traditional type theories, equational propositions are typically represented as instances of a family of types:

$$\frac{\vdash A : \text{Set} \quad a, b : A}{\vdash a =_A b : \text{Prop}}$$

where `Set` is the sort of datatypes and `Prop` is the sort of propositional types—sometimes these coincide. Most people agree that equality should be reflexive, and provide the introduction rule:

$$\frac{\vdash a : A}{\vdash \text{refl}_A a : a =_A a}$$

As these theories have nontrivial computations within types, they come with a notion of *definitional* equality, presented here as a typed equality judgment (\equiv). Types are identified up to definitional equality, as expressed by the *conversion* rule:

$$\frac{\vdash s : S \quad \vdash S \equiv T}{\vdash s : T}$$

The definitional equality is a congruence, so the full effect of the `refl` constructor is to embed \equiv into $=$:

$$\text{if } \vdash a \equiv b : A \text{ then } \vdash a =_A a \equiv a =_A b : \text{Prop} \\ \text{and therefore } \vdash \text{refl}_A a : a =_A a$$

The two camps diverge when it comes to the way that propositional equations are *used*. Extensional type theories adopt the *equality reflection* rule, allowing provable equations to impact directly and silently on typechecking:

$$\frac{\vdash q : a =_A b}{\vdash a \equiv b : A}$$

By contrast, intensional type theories require an explicit operation, corresponding to the ‘Leibniz rule’ for equality, transporting values between types which are merely provably—but perhaps not definitionally—equal

$$\frac{\vdash q : a =_A b \quad \vdash T : A \rightarrow \text{Set} \quad \vdash t : T a}{\text{subst}_{A;a;b} q T t : T b}$$

with the associated computational rule

$$\vdash \text{subst}_{A;a;a} (\text{refl}_A a) T t \equiv t : T a$$

making `subst` disappear when the proof `reflA a` guarantees that initial and final types coincide definitionally. This approach preserves the desirable *canonicity* property of ITTs—*closed* expressions have canonical values. In the empty context, the only possible proof of an equation is by `refl`, so all top-level `subst`s are guaranteed to vanish. Compilers for programming languages based on strongly-normalising ITTs can thus erase equation proofs and appeals to `subst` from run-time code.

Clearly, ITT results in terms which are more bureaucratic: there is no need for `subst` in ETT because the proof of $a =_A b$ gives us $t : T b$ by the conversion rule and equality reflection. However, a high price must be paid for this benefit.

Firstly, the equality reflection rule is clearly not *syntax-directed*: if we want a machine to recover the typing derivations for well-typed terms, it must be able to invent the proofs of equations demanded by appeals to equality reflection. In an ITT, the proofs are always explicit.

Secondly, ETTs identify all types in the presence of false hypotheses, making it unsafe to compute under binders. In an ITT, the machine can check *terms* and execute them safely if they are well-typed, regardless of the context. In an ETT, the checkable objects are the *derivations*, recording not only the appeals to typing rules but also every equational step. ETT terms may be smaller, but only because the evidence which justifies them is somewhere else. Gonthier and Werner’s proof of the Four-Colour Theorem is a relatively small ITT term, but its ETT derivation would be unfeasibly large because it relies very heavily on computation.

A big attraction of ETTs is that more desirable equations have proofs. In particular, the *extensionality* law holds for functions, just by congruence, equality reflection and the η -law:

$$\text{if } \forall x. f x = g x \text{ then } (\lambda x. f x) = (\lambda x. g x) \text{ so } f = g$$

This just cannot happen in ITTs following the `refl` and `subst` pattern. To see this, consider the functions $\lambda x. 0+x$ and $\lambda x. x+0$. One may certainly give a closed inductive proof

$$\text{plus0} : \forall x. 0+x = x+0$$

but the two functions are not definitionally equal, hence `refl` cannot show that $\lambda x. 0+x = \lambda x. x+0$ and correspondingly, there can be no closed proof of this mathematically intuitive fact. By the same token, it is difficult to reason equationally about higher-order functional programs in ITT-based systems, where laws like

$$\text{map id} = \text{id} \quad \text{map } (f \circ g) = \text{map } f \circ \text{map } g$$

cannot be proven.

The challenge is to combine the computational power (and decidable typechecking) of intensional theories with the reasoning power of extensional theories. One cannot just add an axiom

$$\frac{q : \forall x. f x = g x}{\text{ext } q : f = g}$$

without losing canonicity:

$$\text{subst } (\text{ext plus0}) (\lambda _ . \text{Nat}) 0 : \text{Nat}$$

is a closed expression which cannot compute a number, because the `subst` has got stuck! The irony of this example is that $\lambda _ . \text{Nat}$ makes no *use* of its argument. In fact, `subst (ext plus0) (λ_. Nat)` is a function from `Nat` to `Nat` which is provably equal (by `ext`) to `id`. If we only looked at the source and target, we could quite easily see how to perform the transportation! It is in this irony that we find some hope of resolution.

1.2 The Observational Approach

Our approach is designed to ensure the effectiveness of the maps induced by equations, transporting values between provably equal sets. We start by explicitly introducing a notion of set equality. This will not be defined inductively, but will rather *compute* what you need to know to transport canonical values from one set to another. Equations between sets induce *coercions*.

$$\frac{S, T : \text{Set}}{S = T : \text{Prop}} \quad \frac{Q : S = T \quad s : S}{s [Q:S=T] : T}$$

Of course, we still need to talk about equality of *values*. We allow the formation of *heterogeneous* equations, relating values in arbitrary sets. This is again a computed notion, telling you how to

ensure that equal values support equal *observations*.

$$\frac{s : S \quad t : T}{(s : S) = (t : T) : \text{Prop}}$$

Heterogeneous equality may seem a little disturbing: when would it make sense to say that $(0 : \text{Nat}) = (\text{id} : \text{Nat} \rightarrow \text{Nat})$? On the other hand, how would you state the fact that some function f preserves equality of arguments if its type is dependent—some $\Pi a : A. B[a]$?¹ With heterogeneous equality, we can write

$$\forall x, y. (x : A) = (y : A) \rightarrow (f x : B[x]) = (f y : B[y])$$

but with the traditional homogeneous equality, it is rather clumsy to cope with the way the results have definitionally distinct types. Using homogeneous equality, with extra dependency and heavy use of `subst`, we would have written:

$$\forall x, y. \forall q : x =_A y. \text{subst } q (\lambda a. B[a]) (f x) =_{B[y]} f y$$

Our approach scales up much more conveniently. The crucial point is that heterogeneous equations are only *useful* if we know why their types are equal.

Heterogeneity also allows us to formulate the *coherence* operator, another OTT primitive, whose definition requires us to ensure that all of our equality-induced coercions are in some sense just the identity:

$$\frac{Q : S = T \quad s : S}{\{s \parallel Q : S = T\} : (s : S) = (s [Q : S = T] : T)}$$

Note that, by construction, the two sides of this equation have equal types. We shall need coherence in order to explain the action of coercion on types which involve binding.

The need to implement coercion between sets will drive the definition of equality on sets. We shall then be in a position to consider what equality on values must be.

1.3 A Brief History of Equality

There has been quite a lot of work addressing the problem that intensional type theories have weak reasoning properties, despite their good computational behaviour. The most significant contribution is that of Martin Hofmann, who showed that equality reflection, in the style of ETT, is conservative over ITT extended by axioms for ‘extensional concepts’ [14, 13]. The idea of extending ITT consistently with extensionality axioms goes back at least to Turner [27], but the problem of losing canonicity was also clear at that time.

Hofmann also pioneered the construction of *setoid* models, encoding a type theory where extensionality holds in an extensional theory. A setoid is a set equipped with its own notion of equivalence; functions between setoids form a setoid with functions equivalent when they take equivalent domain values to equivalent range value. In order to reason by substitution, we need to be sure that the predicates we can form over a setoid respect its equivalence. Hofmann’s machinery ensures that this is always the case.

However, Hofmann’s attempt to build a setoid model of a system with *large elimination*—computation of types from values—ran into technical difficulties. Altenkirch’s setoid model [4] gets around these problems: his meta-logic has a more liberal definitional equality, identifying all proofs of the same proposition, allowing the extension of Hofmann’s work.

Meanwhile, McBride’s contribution of *heterogeneous* equality [19] reduced the bureaucracy of working with equations between values in dependent types, motivated by need to express the unification problems inherent to Coquand’s dependent pattern

¹ We write Π when we range over a `Set` and \forall when we range over a `Prop`.

matching [8]. The idea of using heterogeneous equality goes back at least to Huet and Saïbi’s formalisation of category theory in type theory [16], itself an epic struggle. However, McBride’s contribution was to provide a *homogeneous* substitution principle, equivalent to extending the usual intensional propositional equality with the so-called ‘K-rule’, expressing the uniqueness of identity proofs. The latter is a trivial consequence of pattern matching, but was shown by Hofmann and Streicher to be unprovable from the usual induction principle for equality [15].

Heterogeneous equality was employed more recently by Oury to simplify Hofmann’s proof, showing that an extensional variant of the Calculus of Constructions is conservative over the intensional version with extensional axioms [22]. Our present work is a heterogeneous variant of the theory for which Altenkirch constructed a setoid model. The consequent simplification allows us to give a direct syntactic presentation of OTT, calculating all the necessary machinery from the structure of the types involved.

In an exciting parallel development, the need to manage equational constraints for functional programming with Generalized Algebraic Datatypes in Haskell has resulted in an extension of System F with a notion of equality, inducing coercions which compute through the structure of values [26]. Our system has something of the same flavour, but type dependency involves more interaction between its layers: we have fewer, more general components, but we need a more delicate touch.

2. A Simple Core Type Theory

We shall start from a very basic type theory, then show how to equip it with a notion of observational equality. We should emphasize that we intend this theory to act as a very simple *core* language, with few syntactic conveniences and the minimal feature set required to investigate and illustrate the observational approach. In particular, our basic theory has no hierarchy of universes, although it still has type-level computation. Let us first establish the core syntax of sets and terms.

```

set      S ::= G | B X : S . S | If T Then S Else S
ground  G ::= 0 | 1 | 2
binder  B ::= Π | Σ | W

term    T ::= ⟨ ⟩ | tt | ff | λ X : S . T | ⟨ T , T ⟩Σ X : S . S | T <W X : S . S T
          | T ! S | if T / X . S then T else T
          | T T | fst T | snd T | rec T / X . S with T

```

We have a fixed repertoire of set-formers: ground sets with 0 elements, 1 element $\langle \rangle$ and 2 elements $\{\text{tt}, \text{ff}\}$, and set binders Π for dependent functions (λ -terms), Σ for dependent pairs, and W for well-founded trees whose nodes $s \triangleleft_{W, X, S, T} f$ comprise a choice s of node shape from S and a ‘child-function’ f from recursive positions $T[s]$ to subtrees. Informally, we shall drop type annotations from value constructors where they may readily be inferred.² We shall adopt the standard abbreviations \rightarrow and \times for non-dependent Π - and Σ -types, respectively. We also write $\lambda x. y. t$ for $\lambda x. \lambda y. t$ and $\langle r, s, t \rangle$ for $\langle r, \langle s, t \rangle \rangle$.

Computation happens via the corresponding eliminators, with $-\!|T$ making a ‘magic’ element of T from a mythical element of 0, dependent if-then-else for 2, application for Π , projections for Σ , and dependent recursion for W . We have been careful to consider only functions between sets: dependent eliminators like `if` and `rec` abstract their target sets over a bound variable, leaving no need to internalise a type of set-valued functions. Informally, we shall drop this range annotation when eliminating a *variable*,

² More formally, we can introduce a bidirectional approach to typechecking which ensures that we can always drop these annotations, but that is a story for another time.

in which case the abstraction is readily determined from the type required. Despite the absence of set-valued functions, we shall still see computation within types via ‘large’ conditionals on elements of 2.

The operational semantics of computation is given by structural closure of the following schemes:

$$\begin{aligned}
& (\lambda x. t) v \rightsquigarrow t[v] \\
& \text{fst } \langle s, t \rangle \rightsquigarrow s \\
& \text{snd } \langle s, t \rangle \rightsquigarrow t \\
& \text{if } \mathbf{tt}/x.P \text{ then } t \text{ else } f \rightsquigarrow t \\
& \text{if } \mathbf{ff}/x.P \text{ then } t \text{ else } f \rightsquigarrow f \\
& \text{rec } s \triangleleft f/w.P \text{ with } p \rightsquigarrow p s f (\lambda t. \text{rec } f t/w.P \text{ with } p) \\
& \text{If } \mathbf{tt} \text{ Then } T \text{ Else } F \rightsquigarrow T \\
& \text{If } \mathbf{ff} \text{ Then } T \text{ Else } F \rightsquigarrow F
\end{aligned}$$

This system is just a fragment of many standard-issue type theories (Martin-Löf’s Intensional Type Theory, the Calculus of Inductive Constructions, Luo’s UTT, etc) all of which enjoy metatheoretical properties such as confluence and strong normalisation. We are therefore entitled to presume the existence of an operator $-^\downarrow$, defined on well-typed terms, taking them to *values* (in this case, just the normal forms):

$$\begin{aligned}
\text{set } \quad \bar{\mathbf{S}} &::= \mathbf{G} \mid \mathbf{B}\mathbf{X}:\bar{\mathbf{S}}.\bar{\mathbf{S}} \mid \text{If } \bar{\mathbf{N}} \text{ Then } \bar{\mathbf{S}} \text{ Else } \bar{\mathbf{S}} \\
\text{term } \quad \bar{\mathbf{T}} &::= \bar{\mathbf{N}} \mid \langle \rangle \mid \mathbf{tt} \mid \mathbf{ff} \mid \lambda \mathbf{X}.\bar{\mathbf{T}} \mid \langle \bar{\mathbf{T}}, \bar{\mathbf{T}} \rangle \mid \bar{\mathbf{T}} \triangleleft \bar{\mathbf{T}} \\
\text{neutral } \bar{\mathbf{N}} &::= \mathbf{X} \mid \bar{\mathbf{N}}!\bar{\mathbf{S}} \mid \text{if } \bar{\mathbf{N}}/\mathbf{X}.\bar{\mathbf{S}} \text{ then } \bar{\mathbf{T}} \text{ else } \bar{\mathbf{T}} \\
& \mid \bar{\mathbf{N}}\bar{\mathbf{T}} \mid \text{fst } \bar{\mathbf{N}} \mid \text{snd } \bar{\mathbf{N}} \mid \text{rec } \bar{\mathbf{N}}/\mathbf{X}.\bar{\mathbf{S}} \text{ with } \bar{\mathbf{T}}
\end{aligned}$$

We say that normal terms are *canonical* if they are built by canonical constructors, i.e., $\langle \rangle$, \mathbf{tt} , \mathbf{ff} , $\lambda x. -$, $\langle -, - \rangle$, $- \triangleleft -$; other normal terms are *neutral*. Similarly, the *canonical sets* are the ground sets and those introduced by the three set-binders: the neutral sets are those given by large elimination on neutral elements of 2. As the base case for neutral terms is the *variable* case, we can readily see that there are no neutral terms in the empty context. Hence, in the empty context, all normal forms are canonical, for both sets and values: we call this property *canonicity*.

We specify *definitional* equality (\equiv) on terms by giving an equivalence on their values. Until section 6, when we introduce an important modification, α -equivalence on values will suffice.

As usual, we take contexts to assign sets to variables:

$$\mathbf{Ctxt} ::= \mathcal{E} \mid \mathbf{Ctxt}; \mathbf{X}:\mathbf{S}$$

We can now give the rules of *context validity*, *set formation* and *type synthesis*.

$$\boxed{\mathbf{Ctxt} \vdash}$$

$$\overline{\mathcal{E} \vdash} \quad \frac{\Gamma \vdash S \mathbf{set}}{\Gamma; x:\mathbf{S} \vdash} x \notin \Gamma$$

$$\boxed{\mathbf{Ctxt} \vdash S \mathbf{set}}$$

$$\begin{aligned}
& \frac{\Gamma \vdash}{\Gamma \vdash G \mathbf{set}} G \in \mathbf{G} \quad \frac{\Gamma \vdash S \mathbf{set} \quad \Gamma; x:\mathbf{S} \vdash T \mathbf{set}}{\Gamma \vdash Bx:\mathbf{S}.T \mathbf{set}} B \in \mathbf{B} \\
& \frac{\Gamma \vdash b:2 \quad \Gamma \vdash T \mathbf{set} \quad \Gamma \vdash F \mathbf{set}}{\Gamma \vdash \text{If } b \text{ Then } T \text{ Else } F \mathbf{set}}
\end{aligned}$$

$$\boxed{\mathbf{Ctxt} \vdash \mathbf{T}:\mathbf{S}}$$

$$\frac{\Gamma \vdash s:S \quad S \equiv T}{\Gamma \vdash s:T}$$

$$\frac{\Gamma \vdash}{\Gamma \vdash \langle \rangle:1} \quad \frac{\Gamma \vdash}{\Gamma \vdash \mathbf{tt}:2} \quad \frac{\Gamma \vdash}{\Gamma \vdash \mathbf{ff}:2}$$

$$\frac{\Gamma; x:\mathbf{S} \vdash t:T \quad \Gamma \vdash s:S \quad \Gamma \vdash t:T[s]}{\Gamma \vdash \lambda x:\mathbf{S}.t:\Pi x:\mathbf{S}.T} \quad \frac{\Gamma \vdash s:S \quad \Gamma \vdash t:T[s]}{\Gamma \vdash \langle s, t \rangle_{\Sigma x:\mathbf{S}.T}:\Sigma x:\mathbf{S}.T}$$

$$\frac{\Gamma \vdash s:S \quad \Gamma \vdash f:T[s] \rightarrow Wx:\mathbf{S}.T}{\Gamma \vdash s \triangleleft_{Wx:\mathbf{S}.T} f:Wx:\mathbf{S}.T}$$

$$\frac{\Gamma \vdash}{\Gamma \vdash x:T} x:T \in \Gamma \quad \frac{\Gamma \vdash f:\Pi x:\mathbf{S}.T \quad \Gamma \vdash s:S}{\Gamma \vdash f s:T[s]}$$

$$\frac{\Gamma \vdash p:\Sigma x:\mathbf{S}.T}{\Gamma \vdash \text{fst } p:S} \quad \frac{\Gamma \vdash p:\Sigma x:\mathbf{S}.T}{\Gamma \vdash \text{snd } p:T[\text{fst } p]}$$

$$\begin{aligned}
& \frac{\Gamma \vdash z:0 \quad \Gamma \vdash b:2 \quad \Gamma \vdash t:P[\mathbf{tt}]}{\Gamma \vdash P \mathbf{set}} \quad \frac{\Gamma; x:2 \vdash P \mathbf{set} \quad \Gamma \vdash f:P[\mathbf{ff}]}{\Gamma \vdash z!P:P} \quad \frac{\Gamma \vdash t:P[\mathbf{tt}]}{\Gamma \vdash \text{if } b/x.P \text{ then } t \text{ else } f:P[b]}
\end{aligned}$$

$$\begin{aligned}
& \frac{\Gamma \vdash u:Wx:\mathbf{S}.T \quad \Gamma; w:Wx:\mathbf{S}.T \vdash P \mathbf{set}}{\Gamma \vdash p:\Pi s:\mathbf{S}.P f:T[s] \rightarrow Wx:\mathbf{S}.T} \\
& \frac{(\Pi t:T[s].P[f t]) \rightarrow P[s \triangleleft f]}{\Gamma \vdash \text{rec } u/w.P \text{ with } p:P[u]}
\end{aligned}$$

2.1 Examples, and a Problem

Although this is rather a small theory, lots of familiar constructions can be expressed in terms of it. For example, the binary sum type, $S+T$ can be coded as the pair of a tag choice and a suitable element, whose type is computed from the tag by large elimination.

$$\begin{aligned}
S+T &\mapsto \Sigma b:2. \text{If } b \text{ Then } S \text{ Else } T \\
\text{inl } s &\mapsto \langle \mathbf{tt}, s \rangle \\
\text{inr } t &\mapsto \langle \mathbf{ff}, t \rangle
\end{aligned}$$

In much the same way, the natural numbers can be defined as the *well-founded* structure with two shapes: the ‘successor’ shape has one recursive position and the ‘zero’ shape has none.

$$\begin{aligned}
\text{Tr } b &\mapsto \text{If } b \text{ Then } 1 \text{ Else } 0 \\
\text{Nat} &\mapsto Wb:2. \text{Tr } b \\
\text{zero} &\mapsto \mathbf{ff} \triangleleft \lambda z. z! \text{Nat} \\
\text{suc } n &\mapsto \mathbf{tt} \triangleleft \lambda_. n
\end{aligned}$$

To construct elements, we either indicate the ‘zero’ shape and supply the trivial child-function, or we choose the ‘successor’ shape and make the child-function return the predecessor.

Operators such as addition can now be defined via the recursor:

$$\begin{aligned}
\text{plus} &\mapsto \lambda x y. \text{rec } x \text{ with} \\
& \quad \lambda b. \text{if } b \text{ then } \lambda f h. \text{suc } (h \langle \rangle) \text{ else } \lambda f h. y
\end{aligned}$$

Although it is reassuring that we can implement specific familiar operations, we run into trouble when we try to give a general induction scheme for this coding of Nat. Suppose we have some $n:\text{Nat} \vdash P \mathbf{set}$. We should like to find

$$\text{ind}_P : P[\text{zero}] \rightarrow (\Pi n:\text{Nat}. P[n] \rightarrow P[\text{suc } n]) \rightarrow \Pi n:\text{Nat}. P[n]$$

but the ‘obvious’ program, following the same pattern as plus,

$$\text{ind}_P \mapsto \lambda p z p s n. \text{rec } n \text{ with} \\ \lambda b. \text{if } b \text{ then } \lambda f h. \text{ps } (f \langle \rangle) (h \langle \rangle) \text{ else } \lambda f h. \text{pz}$$

does not typecheck, and for the most infuriating of reasons.

In the ‘zero’ case, we must prove $P[\mathbf{ff} \triangleleft f]$, but the proof we offer is $pz : P[\mathbf{ff} \triangleleft \lambda z. z! \text{Nat}]$. That is, we are obliged to show

that P holds for any implementation of ‘zero’, but we only know that it holds for a specific implementation of ‘zero’. Similarly, in the ‘successor’ case, we need a proof of $P[\mathbf{tt} \triangleleft f]$, but we supply a proof of $P[\mathbf{tt} \triangleleft \lambda \dots f \langle \rangle]$, which is again a more specific thing.³

This is a real problem. There are infinitely many implementations of ‘zero’. For example,

$$\text{zero}' \mapsto \mathbf{ff} \triangleleft \lambda z. \text{suc} (\text{suc } \text{zero})$$

defines ‘zero’ to be ‘the number with no predecessors, all of which are two’! The frustrating thing is that these distinctions cannot be observed from within the theory. They are all *extensionally* equal. In each case of ind_P , we can prove that the child-functions given and required coincide on all inputs, so our program is typable in ETTs. With observational equality, we shall be able to bridge the typing gaps explicitly, with coercions.

2.2 An Inductive-Recursive Universe

We can define our core theory as an inductive-recursive universe in the intensional setting of Agda 2. Firstly, we must declare the type constructors we shall need: 0 is the datatype with no constructors,

data Empty : Set **where**

1 is the record type with no fields

record Unit : Set **where**

and 2 is the usual Boolean datatype

data Bool : Set **where**

$\mathbf{tt} : 2$
 $\mathbf{ff} : 2$

Meanwhile, we can re-use the system’s own Π -sets, written

$$(x : S) \rightarrow T x$$

we can define Σ -sets as records,

record $\Sigma (S : \text{Set})(T : S \rightarrow \text{Set}) : \text{Set} **where**$

$\text{fst} : S$
 $\text{snd} : T \text{ fst}$

and give W -sets as a datatype.

data $W (S : \text{Set})(T : S \rightarrow \text{Set}) : \text{Set} **where**$

$_ \triangleleft _ : (x : S) \rightarrow (T x \rightarrow W S T) \rightarrow W S T$

Now we can build the collection of these sets, simultaneously defining a datatype to ‘name’ them and the function which decodes names as sets

mutual

data ‘set’ : Set **where**

‘0’, ‘1’, ‘2’ : ‘set’
‘ Π ’, ‘ Σ ’, ‘ W ’ : $(S : \text{‘set’}) \rightarrow ([S] \rightarrow \text{‘set’}) \rightarrow \text{‘set’}$

$[_]$: ‘set’ \rightarrow Set
 $[‘0’]$ = Empty
 $[‘1’]$ = Unit
 $[‘2’]$ = Bool
 $[‘\Pi’ S T]$ = $(x : [S]) \rightarrow [T x]$
 $[‘\Sigma’ S T]$ = $\Sigma [S] (\lambda x \mapsto [T x])$
 $[‘W’ S T]$ = $W [S] (\lambda x \mapsto [T x])$

With this encoding, we can readily define all our constructors and eliminators of our core theory as functions in Agda 2, either invoking the corresponding Agda constructors, or implementing the computational behaviour exactly as we have specified it. This gives us a translation $\hat{_}$ taking core sets to elements of ‘set’ and core

³Adding η -laws for Π and 1 solves this particular problem, but does not help the general case.

elements of a core set S to elements of $[S]$, preserving both type and computational behaviour. This is the first step towards the inheritance of strong normalisation and canonicity from an existing intensional theory.

2.3 A Propositional Fragment

We shall shortly construct our propositional equality, but before we do so, let us consider what constitutes a *proposition* in this setting. Of course, we could just identify propositions with sets, but we prefer to be more precise: we want to know that *proofs* have no interesting computational content. Correspondingly, we shall identify a sublanguage of propositions, then explain how to interpret them as sets of proofs.

$$\mathbf{P} ::= \perp \mid \top \mid \mathbf{P} \wedge \mathbf{P} \mid \forall \mathbf{X} : \mathbf{S}. \mathbf{P}$$

Let us be clear that this is by no means the largest sublanguage of content-free sets. Rather, it is the smallest such language fit for our current purpose.

We introduce a judgment form to distinguish the well-formed propositions, and we show how the latter may be interpreted admissibly as sets of proofs,

$\text{Ctx} \vdash \mathbf{P} \text{ prop}$	$\frac{\Gamma \vdash P \text{ prop}}{\Gamma \vdash [P] \text{ set}}$
$\frac{\Gamma \vdash}{\Gamma \vdash \perp \text{ prop}}$	$[\perp] \mapsto 0$
$\frac{\Gamma \vdash}{\Gamma \vdash \top \text{ prop}}$	$[\top] \mapsto 1$
$\frac{\Gamma \vdash P \text{ prop} \quad \Gamma \vdash Q \text{ prop}}{\Gamma \vdash P \wedge Q \text{ prop}}$	$[P \wedge Q] \mapsto \Sigma : [P]. [Q]$
$\frac{\Gamma \vdash S \text{ set} \quad \Gamma; x : S \vdash P \text{ prop}}{\Gamma \vdash \forall x : S. P \text{ prop}}$	$[\forall x : S. P] \mapsto \Pi x : S. [P]$

Propositional implication, $P \Rightarrow Q$, is coded as $\forall _ : [P]. Q$. We say that a term is a *proof* if it inhabits some $[P]$.

To some extent, this is just window-dressing. We could just have worked with sets which happen to be propositional. However, we have an ulterior motive. Consider which observations on proofs can interfere with the world of *non-propositional* sets: \top has no observations; both projections from a proof of $P \wedge Q$ are proofs; applying a proof of $\forall x : S. P$ yields a proof; only from a contradiction can we make *data*!

This has at least two useful consequences. Firstly, it allows us to erase all proof objects from run-time code: if we never compute under λ , we shall never find a proof of \perp , so proofs are just dead code. Secondly, we may extend the language of proofs with whatever propositional laws we like, as long as they are *consistent*, and yet retain canonicity.

We can also express this propositional fragment as an inductive-recursive universe in Agda, this time decoding into our ‘set’ universe rather than the native Set, exactly following the rules above.

mutual

data ‘prop’ : Set **where**

‘ \perp ’, ‘ \top ’ : ‘prop’
‘ \wedge ’ : ‘prop’ \rightarrow ‘prop’ \rightarrow ‘prop’
‘ \forall ’ : $(S : \text{‘set’}) \rightarrow ([S] \rightarrow \text{‘prop’}) \rightarrow \text{‘prop’}$
 $[_]$: ‘prop’ \rightarrow ‘set’
...

3. Equality, Coercion and Laziness

In this section, we shall try to implement the operation which transports values between equal sets. As we do so, we shall find

out what the necessary consequences of set equality should be, and that will tell us how set equality should be defined! Let us introduce new operators for set equality and coercion, by recursion on sets,

$$\frac{\Gamma \vdash S \text{ set} \quad \Gamma \vdash T \text{ set}}{\Gamma \vdash S = T \text{ prop}} \quad \frac{\Gamma \vdash Q : [S = T] \quad \Gamma \vdash s : S}{\Gamma \vdash s [Q : S = T] : T}$$

together with their value-level counterparts, which we shall define in the next section.

$$\frac{\Gamma \vdash s : S \quad \Gamma \vdash t : T}{\Gamma \vdash (s : S) = (t : T) \text{ prop}} \quad \frac{\Gamma \vdash Q : [S = T] \quad \Gamma \vdash s : S}{\Gamma \vdash \{s \parallel Q : S = T\} : [(s : S) = (s [Q : S = T] : T)]}$$

As you can see, coercion apparently introduces a new way for proofs to interfere with data. However, we shall define coercion recursively in terms of the existing operations on proofs, thus retaining both the erasability of propositions and, crucially, our freedom to add consistent laws.

As the sets inhabited by canonical values are themselves canonical, it is sufficient to explain how to compute a coercion between any two canonical sets, given a proof that they are equal. Let us deal with the easy cases first: at identical ground types, there is nothing to do!

$$\begin{aligned} z [Q : 0=0] &\mapsto z \\ u [Q : 1=1] &\mapsto u \\ b [Q : 2=2] &\mapsto b \end{aligned}$$

Correspondingly, we may take

$$\begin{aligned} 0 = 0 &\mapsto \top \\ 1 = 1 &\mapsto \top \\ 2 = 2 &\mapsto \top \end{aligned}$$

The hard work will come when explaining what to do with types made with the same binder. Even harder are the off-diagonal cases, so we shall have to rule them out.

$$\begin{aligned} f_0 [Q : \Pi x_0 : S_0. T_0 = \Pi x_1 : S_1. T_1] &\mapsto ? \\ p_0 [Q : \Sigma x_0 : S_0. T_0 = \Sigma x_1 : S_1. T_1] &\mapsto ? \\ (s_0 \triangleleft f_0) [Q : W x_0 : S_0. T_0 = W x_1 : S_1. T_1] &\mapsto ? \\ x [Q : S = T] &\mapsto Q!T \text{ otherwise} \end{aligned}$$

$$\begin{aligned} (\Pi x_0 : S_0. T_0) = (\Pi x_1 : S_1. T_1) &\mapsto ? \\ (\Sigma x_0 : S_0. T_0) = (\Sigma x_1 : S_1. T_1) &\mapsto ? \\ (W x_0 : S_0. T_0) = (W x_1 : S_1. T_1) &\mapsto ? \\ S = T &\mapsto \perp \text{ for other canonical sets} \end{aligned}$$

We must fill in those ?s. Let us start with Σ -types, as tuples are relatively mundane, compared to functions or well-founded trees. We must solve

$$\begin{aligned} \dots; Q : (\Sigma x_0 : S_0. T_0) = (\Sigma x_1 : S_1. T_1); \\ p_0 : \Sigma x_0 : S_0. T_0 \\ \vdash ? : \Sigma x_1 : S_1. T_1 \end{aligned}$$

Let us work by refinement, following the trail left by the types. Some moves are clear. Let us name the input's projections for convenient access, then construct the output componentwise. We use an informal **let** notation for 'definitions' to be substituted out in the final term. This notation is not part of our theory, but makes our terms a bit more legible.

$$\begin{aligned} \text{let} \\ s_0 &\mapsto \text{fst } p_0 : S_0 \\ t_0 &\mapsto \text{snd } p_0 : T_0[s_0] \\ s_1 &\mapsto ? : S_1 \\ t_1 &\mapsto ? : T_1[s_1] \\ \text{in } \langle s_1, t_1 \rangle \end{aligned}$$

Now, the only way we can make an output equal to the input is if the output components are equal to the input components. We must make s_1 by coercing s_0 and t_1 by coercing t_0 . We get

$$\begin{aligned} \text{let} \\ F s_0 &\mapsto \text{fst } p_0 : S_0 \\ F t_0 &\mapsto \text{snd } p_0 : T_0[s_0] \\ Q_S &\mapsto ? : [S_0 = S_1] \\ s_1 &\mapsto s_0 [Q_S : S_0 = S_1] : S_1 \\ Q_T &\mapsto ? : [T_0[s_0] = T_1[s_1]] \\ t_1 &\mapsto t_0 [Q_T : T_0[s_0] = T_1[s_1]] : T_1[s_1] \\ \text{in } \langle s_1, t_1 \rangle \end{aligned}$$

It remains to find proofs of the equations which justify the coercions. These concern the components of the Σ -types, and must surely be extracted from Q . We shall clearly need Q to tell us that $S_0 = S_1$. However, in the case of Q_T , we are obliged to show that the s_0 instance of T_0 equals the s_1 instance of T_1 , where s_0 and s_1 are unknown when Q 's type is being determined. It is too much to require that *arbitrary* instances of T_0 and T_1 are equal—that would force us to remove any meaningful dependency. We can, however require T_0 and T_1 to coincide whenever they are instantiated *equally*, for s_1 is a coercion of $F s_0$, hence equal by *coherence*.

Let us therefore take

$$\begin{aligned} (\Sigma x_0 : S_0. T_0) = (\Sigma x_1 : S_1. T_1) &\mapsto \\ S_0 = S_1 \wedge \\ \forall x_0 : S_0. \forall x_1 : S_1. (x_0 : S_0) = (x_1 : S_1) &\Rightarrow T_0[x_0] = T_1[x_1] \end{aligned}$$

so that we can have

$$\begin{aligned} Q_S &\mapsto \text{fst } Q \\ Q_T &\mapsto (\text{snd } Q) s_0 s_1 \{s_0 \parallel Q_S : S_0 = S_1\} \end{aligned}$$

completing the case for Σ -types.

For Π -types, we can expect a little contravariant twist. We need to coerce the argument s_1 right-to-left, so that we can apply the function f_0 , then coerce the result left-to-right. Working in the same style, we have

$$\begin{aligned} \dots; Q : (\Pi x_0 : S_0. T_0) = (\Pi x_1 : S_1. T_1); \\ f_0 : \Pi x_0 : S_0. T_0 \\ \vdash \lambda s_1. \text{let} \\ F Q_S &\mapsto ? : [S_1 = S_0] \\ s_0 &\mapsto s_1 [Q_S : S_1 = S_0] : S_0 \\ t_0 &\mapsto f_0 s_0 : T_0[s_0] \\ Q_T &\mapsto ? : [T_0[s_0] = T_1[s_1]] \\ t_1 &\mapsto t_0 [Q_T : T_0[s_0] = T_1[s_1]] : T_1[s_1] \\ \text{in } t_1 \\ : \Pi x_1 : S_1. T_1 \end{aligned}$$

Correspondingly, we should take

$$\begin{aligned} (\Pi x_0 : S_0. T_0) = (\Pi x_1 : S_1. T_1) &\mapsto \\ S_1 = S_0 \wedge \\ \forall x_1 : S_1. \forall x_0 : S_0. (x_1 : S_1) = (x_0 : S_0) &\Rightarrow T_0[x_0] = T_1[x_1] \end{aligned}$$

$$\begin{aligned} Q_S &\mapsto \text{fst } Q \\ Q_T &\mapsto (\text{snd } Q) s_1 s_0 \{s_1 \parallel Q_S : S_1 = S_0\} \end{aligned}$$

In the case of W types, we shall need a recursive coercion, shifting shapes and child-functions left-to-right: for the latter, we shall need to shift child-positions right-to-left, and the resulting child-trees left-to-right recursively. Perhaps you can already guess that we shall take

$$\begin{aligned} (W x_0 : S_0. T_0) = (W x_1 : S_1. T_1) &\mapsto \\ S_0 = S_1 \wedge \\ \forall x_0 : S_0. \forall x_1 : S_1. (x_0 : S_0) = (x_1 : S_1) &\Rightarrow T_1[x_1] = T_0[x_0] \end{aligned}$$

The coercion can then be defined recursively as follows:

```

( $s_0 \triangleleft f_0$ ) [Q:Wx0:S0. T0=Wx1:S1. T1] ↦
let
  QS ↦ fst Q : [S0 = S1]
  s1 ↦ s0 [QS:S0=S1] : S1
  QT ↦ (snd Q) s0 s1 {s0 || QS:S0=S1} : [T1[s1] = T0[s0]]
in s1 <λt1. let t0 ↦ t1 [QT:T1[s1]=T0[s0]] : T0[s0]
      in f0 t0 [Q:Wx0:S0. T0=Wx1:S1. T1]

```

We have completed the first stage of our journey. What we have produced may look complex and detailed, but there was relatively little choice at any point. The structure of the types determined the structure of the coercions between them. This delivered in turn the requirements which determined what set equality should be: impossible off the diagonal; trivial at ground types; for binders, equal domains and equal range instances for equal domain values, twisted according to variance.

3.1 Value Equality

We must now explain when values are equal. The idea is that values should be equal when they support equal observations. For those sets equipped with inductive eliminators (0, 2, W-sets), this amounts to equality of construction. For those equipped only with projective eliminators (1, Σ, Π), we are free to require only that projections coincide.

```

(z0 : 0) = (z1 : 0) ↦ ⊤
(u0 : 1) = (u1 : 1) ↦ ⊤
(Ⓢ : 2) = (Ⓢ : 2) ↦ ⊤
(Ⓢ : 2) = (Ⓢ : 2) ↦ ⊥
(Ⓢ : 2) = (Ⓢ : 2) ↦ ⊥
(Ⓢ : 2) = (Ⓢ : 2) ↦ ⊤

```

```

(f0 : Πx0:S0. T0) = (f1 : Πx1:S1. T1) ↦
  ∀x0:S0. ∀x1:S1. (x0 : S0) = (x1 : S1) ⇒
  (f0 x0 : T0[x0]) = (f1 x1 : T1[x1])

```

```

(p0 : Σx0:S0. T0) = (p1 : Σx1:S1. T1) ↦
  (fst p0 : S0) = (fst p1 : S1) ∧
  (snd p0 : T0[fst p0]) = (snd p1 : T1[fst p1])

```

```

(s0 <f0 : Wx0:S0. T0) = (s1 <f1 : Wx1:S1. T1) ↦
  (s0 : S0) = (s1 : S1) ∧
  ∀y0:T0[s0]. ∀y1:T1[s1]. (y0 : T0[s0]) = (y1 : T1[s0]) ⇒
  (f0 y0 : Wx0:S0. T0) = (f1 y1 : Wx1:S1. T1)
  (- : T0) = (- : T1) ↦ ⊥

```

T_0, T_1 other canonical sets

As you can see, two functions are equal if they take equal inputs to equal outputs; two pairs are equal if they have equal projections. Meanwhile, as the eliminators for 2 and W allow us to observe the construction of their contents, equality for elements of these types is equality of construction. Note that this definition of value equality preserves set equality: if we know that the values on the left inhabit equal sets, then we know that every equation on the right relates values in equal sets.

What can impede the computation of = for values? Only the presence of a neutral set, or of a neutral value in 2 or a W-set.

We have shown how to construct equations between sets and how these equations, if provable, yield coercions. We have also shown how to construct observational justifications for the equality of terms. This does *not* complete our presentation of OTT. We shall need to add more introduction rules for equations which are not directly observational. However, we have now given all of OTT's *computation* schemes. Whatever we add will not change the computational behaviour of the system, only enlarge the collection of

provable equations. Correspondingly, it is a good time to consider the metatheoretical properties of these computation schemes.

3.2 Modelling Equality and Coercion yields Normalisation

The definition of equality and coercion, given above, is entirely representable in our Agda model. We show how our universe can readily be equipped with equality and coercion by *definitional extension*. Moreover, we ensure that all the computation schemes are faithfully simulated. By this means, we shall be able to demonstrate that OTT is strongly normalising. We define

```

mutual
  Eq : 'set' → 'set' → 'prop'
  ...
  eq : (S : 'set') → [S] → (T : 'set') → [T] → 'prop'
  ...

```

with operational behaviour delivering exactly the above computation schemes. These functions cover all the canonical sets and values from our core theory: they introduce no 'new' propositions.

Of course, we must be sure that these programs terminate. The contravariant twists we used to keep coercion simple mean that our recursion is not directly structural: although we appeal to recursion on smaller elements of 'set', we swap the argument positions in which they appear. We shall address this issue shortly.

Once we know what equality is, we may introduce coercion and coherence as a mutual definition:

```

mutual
  coe : (S : 'set')(T : 'set') → [[Eq S T]] → [S] → [T]
  coe '0' '0' Q z = z
  coe '0' '1' Q z = Q!Unit
  :
  coe ('W' S0 T0) ('W' S1 T1) Q (s0 <f0) = ...
  coh : (S : 'set')(T : 'set')(Q : [[Eq S T]])(s : [S]) →
        eq S s T (coe S T Q s)
  coh S T Q s = ?

```

The definition of coercion follows the same twisted recursion scheme on the names of types, combined lexicographically with structural recursion on trees in W-types, when required. Correspondingly, the computation schemes for coercion are faithfully simulated by the embedding. We shall keep you in suspense about the definition of coherence a little while longer.

LEMMA 1 (Strong Normalisation). *OTT is strongly normalising.*

PROOF SKETCH Our approach to strong normalisation is to *simulate* the reduction behaviour of OTT within a strongly normalising intensional type theory. We proceed in two steps.

Firstly, we must account for the non-structural twisting in contravariant positions. There are advanced ways to do this [2], but the basic, clumsy way is to define the equality in the model without twists and close the simulating set modelling each OTT equation under symmetry. If we do this, we must extend the model of coercion to give us both directions at once, and switch direction in contravariant positions. By doing so, we acquire a model in the system of inductive-recursive definitions proposed by Dybjer and Setzer [9]. Although this system has a set-theoretic model and is widely suspected to be normalising, we cannot consider the job done.

Correspondingly, the second step is to code the induction-recursion within a standard type theory such as Luo's UTT [17] or the Calculus of Inductive Constructions [28]. The usual coding trick, discovered independently by Capretta, McBride and doubtless others, is to turn the decoding function into an index. Where

Agda allows us

```
data 'set' : Set ... [-] : 'set' → Set
```

we may instead resort to an indexed-family, given a larger universe

```
HasCode : Set → Type
```

and rearrange our constructions accordingly,

```
'set' = ΣS : Set. HasCode S [-] = fst
```

We may define the operations, all as a package, using just the standard eliminator for structural recursion on HasCode. \square

4. Canonicity from Consistency

We are now in a position to investigate the issue of canonicity for OTT. By adding coercion, we introduced a new way for proofs of propositions to interact with calculations on data, so we must check that coercion never ‘gums up the works’. Our careful separation of propositions from sets allows us to proceed in two stages: first, we check that *logical contradiction* is the only possible source of non-canonical *data*; then, we check that our *proof*-language is free of such contradictions.

LEMMA 2 (Canonicity from Consistency). *Suppose OTT is consistent, i.e. that there is no s such that $\mathcal{E} \vdash s : 0$. Then, for all normal S and s ,*

- if $\mathcal{E} \vdash S \text{ set}$ then S is canonical;
- if $\mathcal{E} \vdash s : S$ then either s is canonical, or s is a proof.

PROOF We proceed by mutual induction on the normal forms S and s . Clearly we need only consider how to rule out neutral sets and values. We have three forms of neutral set:

- $\mathcal{E} \vdash \text{If } b \text{ Then } T \text{ Else } F \text{ set}$ This must follow from some $\mathcal{E} \vdash b : 2$; however, 2 is not a proposition, so inductively, b must be \mathfrak{t} or \mathfrak{f} , hence our set computes further, contradicting the hypothesis that it is in normal form.
- $\mathcal{E} \vdash [S_0 = S_1] \text{ set}$ This must follow from $\mathcal{E} \vdash S_0 \text{ set}$ and $\mathcal{E} \vdash S_1 \text{ set}$. Inductively, S_0 and S_1 are canonical, so the equation computes further.
- $\mathcal{E} \vdash [(s_0 : S_0) = (s_1 : S_1)] \text{ set}$ Inductively, the two sets must be canonical. Moreover, if S_i is 2 or a W -set, then s_i is also canonical. Hence the equation computes further.

Meanwhile, for the values, the result holds for

- canonical term-formers, directly;
- variables, vacuously, because the context is empty;
- $z!S$, by our assumption of consistency;
- if $b/x.P$ then t else f , because $b : 2$ must be canonical, inductively, hence the conditional computes;
- $\text{rec } w/x.P$ with p , because $w : Wx : S. T$ must be canonical, inductively, hence the recursor computes;
- $f a$, because, inductively, f must be either
 - $\lambda x. t$, contradicting normality of $f a$, or
 - a proof, in which case so is $f a$;
- $\text{fst } p$ and $\text{snd } p$, by the same argument—these must be proofs because, inductively, p must be a proof;
- $\{s \parallel Q : S = T\}$, because this is a proof;
- $s [Q : S = T]$, because inductively, S and T are canonical which is enough to make the coercion compute, except in the case where they are both W -sets, in which case we also need the inductive hypothesis that s is canonical. \square

The impact of this lemma is significant: it tells us that we can add introduction rules for any *propositions* we like, provided that they are consistent, without threatening either normalisation (because they do not compute) or canonicity (because this is ensured by consistency). Our hard work in designing this coercion mechanism to ensure the transportation of canonical values between equal types has bought us the freedom to design the propositional fragment of OTT for our convenience, within reason.

4.1 Laws for Equality

Which equational laws shall we have? We should certainly like to recover refl and subst , so let us have

$$\frac{\Gamma \vdash s : S}{\Gamma \vdash s : \overline{S} : [(s : S) = (s : S)]}$$

$$\frac{\Gamma \vdash S \text{ set} \quad \Gamma; x : S \vdash T \text{ set}}{\Gamma \vdash R x : S. T : [\forall y : S. \forall z : S. (y : S) = (z : S) \Rightarrow T[y] = T[z]]}$$

The former is just reflexivity for values. The latter asserts that any set abstracting a value must preserve the value equality—any such abstracted set thus induces a substitution operator for values.

$$\text{subst}[x : S. T] \mapsto \lambda y z q t. t [(R x : S. T) y z q : T[y] = T[z]]$$

$$: \Pi y : S. \Pi z : S. [(y : S) = (z : S)] \rightarrow T[y] \rightarrow T[z]$$

Note that we can now form

$$\text{subst}[x : S. T] s s (\overline{s : S}) : T[s] \rightarrow T[s]$$

but that the computational behaviour of this is not uniformly the identity: it depends on the structure of T , and the value being transported. In particular,

$$\text{subst}[x : 2. \text{If } x \text{ Then } T \text{ Else } F] b b (\overline{b : 2}) t$$

is bound to get stuck if b is neutral, because we can only coerce between canonical types. Altenkirch’s extensional model construction [4] also had this issue, losing some computational properties of ITTs to gain some reasoning properties from ETTs. This is a big issue: the vanishing of trivial subst s is what gives elaborated Epigram programs their expected computational behaviour [20, 12]. Fortunately, as you shall see later, we can recover this intensional behaviour without further adjusting our notion of evaluation.

Of course, *coherence* is enough to tell us that trivial subst s are *provably* equal to the identity. We have not yet defined coherence, but as it is a propositional law, we do not need to. Let us leave our coherence operator without computational behaviour!

We should like our equalities to be equivalences. We have reflexivity for values; for sets, we can take

$$\overline{S} \mapsto (R x : 1. S) \langle \rangle \langle \rangle : S = S$$

We could derive symmetry and transitivity directly if our theory was rich enough to support abstraction over sets, but even with the restricted theory of this paper, it is enough to add that equality respects equality, i.e.

$$\frac{\Gamma \vdash P : [A = C] \quad \Gamma \vdash Q : [B = D]}{\Gamma \vdash P \parallel Q : [(A = B) = (C = D)]}$$

and similarly for values, although we must be careful to consider only equations on values between equal sets

$$\frac{\Gamma \vdash P : [A = C] \quad \Gamma \vdash Q : [B = D] \quad \Gamma \vdash p : [(a : A) = (c : C)] \quad \Gamma \vdash q : [(b : B) = (d : D)]}{\Gamma \vdash P, p \parallel Q, q : [(a : A) = (b : B)] = [(c : C) = (d : D)]}$$

Now, if $Q : X = Y$, then we can derive symmetry

$$Q \rightsquigarrow \overline{X} [Q \parallel \overline{X} : [X = X] = [Y = X]] : [Y = X]$$

and if, further, $Q' : Y = Z$, then we can derive transitivity

$$Q \circ Q' \mapsto Q [\overline{X} \parallel Q' : [X = Y] = [X = Z]] : [X = Z]$$

with similar constructions serving for equal values in equal sets.

Further, we should like every syntactic construct to respect the observational equality, making equal objects from equal subobjects. These laws already hold within OTT for each syntactic construct of our base theory, and we have just added that equality respects equality. We can easily use coherence to prove that coercion respects equality. This leaves us only with the need to prove that all of our propositional laws respect equality. We can get this in one go by adding a *proof-irrelevance* law: “equal propositions have equal proofs”.

$$\frac{\Gamma \vdash P_0 \text{ prop} \quad \Gamma \vdash P_1 \text{ prop} \quad \Gamma \vdash Q : [P_0] = [P_1]}{\text{Irr } Q : [\forall p_0 : [P_0]. \forall p_1 : [P_1]. (p_0 : P_0) = (p_1 : P_1)]}$$

Now, it is the case that many of our laws can be implemented by ‘generic programming’ in our Agda embedding, where we have access to recursion over ‘set’. Certainly, `Irr` holds by induction on the structure of propositions. However, it is also clear that not all such laws can be provable in Agda: we represent OTT functions by Agda functions, but extensionality holds in OTT. To see this, suppose $FG : [\forall x : S. (f x : T[x]) = (g x : T[x])]$, and observe

$$\lambda x y xy. FG x \circ \text{snd } (g : \Pi x : S. T[x]) x y xy \\ : [(f : \Pi x : S. T[x]) = (g : \Pi x : S. T[x])]$$

Extensionality does not hold in Agda, although it is consistent to assume it. Fortunately for us, our Agda embedding delivers all the computational behaviour of OTT without any need to derive all these laws. Coercion never looks at the proof, so why should we trouble to compute it?

4.2 Consistency from the Extensional Theory

There is nothing to stop us embedding OTT into an extensional type theory. The obvious way to achieve this is to embed our existing intensional translation into the corresponding extensional type theory. The latter may have poor computational properties, but it is at least *consistent*, so we can prove consistency of OTT just by finishing our homework: we must derive our propositional laws in the extensional setting.

Let us write \equiv_A for the built-in propositional equality of the extensional theory, which coincides with the definitional equality by the conversion rule in one direction, and the equality reflection rule in the other. The key fact is this:

LEMMA 3 (Equality Interchange). *For all $S_0, S_1 : \text{‘set’}$,*

$$[[[S_0 = S_1]]] \Leftrightarrow S_0 \equiv_{\text{‘set’}} S_1$$

Moreover, if $S_0 \equiv_{\text{‘set’}} S_1$, then for all $s_0 : S_0$ and $s_1 : S_1$,

$$[[[s_0 : S_0] = (s_1 : S_1)]] \Leftrightarrow s_0 \equiv_{[[S_0]]} s_1$$

PROOF Note that we are using equality reflection even to *state* this fact—the value level equation is only homogeneous because we have a proof that the sets coincide, hence their interpretations are definitionally equal.

We proceed by the usual variance-twisted induction on S_0 and S_1 . The off-diagonal cases are trivial, as are the cases for equal ground sets.

For ‘ Π ’, we must show

$$\begin{aligned} & [[[S_1 = S_0 \wedge \\ & \quad \forall x_1 : S_1. \forall x_0 : S_0. (x_1 : S_1) = (x_0 : S_0) \Rightarrow T_0 x_0 = T_1 x_1]]] \\ & \Leftrightarrow \\ & \text{‘}\Pi\text{’ } S_0 T_0 \equiv_{\text{‘set’}} \text{‘}\Pi\text{’ } S_1 T_1 \end{aligned}$$

Going right-to-left, injectivity of constructors allows us to identify $S_0 \equiv_{\text{‘set’}} S_1$ and hence $T_0 \equiv_{[[S_0] \rightarrow \text{‘set’}} T_1$. Using our inductive hypothesis for the domain sets, we may deduce $[[[S_1 = S_0]]]$; for the range, we can take $x_1 \equiv_{[[S_0]]} x_0$ by the inductive hypothesis for domain values (noting that the domains have been identified), hence we have $T_0 x_0 \equiv_{\text{‘set’}} T_1 x_1$, so, by the range inductive hypothesis, we have $[[[T_0 x_0 = T_1 x_1]]]$.

Left-to-right, our inductive hypothesis for domain sets immediately identifies them. To deduce that $T_0 \equiv_{[[S_0] \rightarrow \text{‘set’}} T_1$, we may appeal to extensionality and show

$$(x : S_0) \rightarrow T_0 x \equiv_{\text{‘set’}} T_1 x$$

our domain value hypothesis tells us that $[[[(x : S_1) = (x : S_0)]]]$, hence we may deduce $[[[T_0 x = T_1 x]]]$, ready for the range hypothesis.

On the value level, with domain and range identified, we must show

$$\begin{aligned} & [[[\forall x_0 : S. \forall x_1 : S. (x_0 : S) = (x_1 : S) \Rightarrow \\ & \quad (f_0 x_0 : T x_0) = (f_1 x_1 : T x_1)]]] \\ & \Leftrightarrow \\ & f_0 \equiv_{(x:S) \rightarrow T} f_1 \end{aligned}$$

Right-to-left, introducing the hypotheses allows us to identify $x_0 \equiv_S x_1$, yielding $f_0 x_0 \equiv_{T x_0} f_1 x_1$, ready for the range value hypothesis to complete the derivation. Left to right, we again appeal to extensionality, just as we did with the range sets.

For ‘ Σ ’ and ‘ W ’ the proofs on the set-level go just as they do for ‘ Π ’. On the value-level, pairs are treated componentwise and trees require a further induction. The crux remains that whenever we have a hypothetical equation on two values, we always know their sets are equal, so our inductive hypotheses allow us to identify them. \square

THEOREM 4 (Consistency). *There is no s such that $\mathcal{E} \vdash s : 0$.*

PROOF Inspecting our propositional laws, we have maintained the property that values are equated only when they inhabit provably equal sets. We may therefore appeal to the equality interchange lemma to derive these laws in our extensional model from their counterparts with $=$ replaced by \equiv . All these laws hold trivially in extensional type theory. Correspondingly, every closed term in OTT induces a closed term in ETT inhabiting the corresponding set. As ETT is consistent, no closed OTT term can inhabit 0. \square

COROLLARY 5 (Canonicity). *If $\mathcal{E} \vdash S \text{ set}$ then S is canonical. If $\mathcal{E} \vdash s : S$ then s is either canonical or a proof.*

A more arduous but arguably more satisfying proof would be to complete the missing derivations of the propositional laws in an intensional theory, consistently extended with axioms for “extensional concepts”, as proposed by Hofmann [13]. Such extensions give the strength of the extensional theory, whilst retaining the checkability of constructions. W. Swierstra has done most of the work required, for a minor variation of the theory presented here: his construction leaves only reflexivity unfinished, and that is where the appeal to extensionality is required.

5. Induction for Natural Numbers

Let us now put observational equality to work, implementing the induction principle for natural numbers defined via W -sets:

$$\text{ind}_P : P[\text{zero}] \rightarrow (\prod n : \text{Nat}. P[n] \rightarrow P[\text{suc } n]) \rightarrow \prod n : \text{Nat}. P[n]$$

Recall that the problem was to show that the proof of P given for the standard implementation of a number is good for any implementation of that number. We are now ready to solve this problem.

$$\begin{aligned} \text{ind}_P \mapsto & \lambda p z \text{ ps } n. \text{rec } n \text{ with} \\ & \lambda b. \text{if } b \text{ then } \lambda f h. \text{ps } (f \langle \rangle) (h \langle \rangle) \\ & \quad [?:P[\text{suc } (f \langle \rangle)]] = P[\mathbf{tt} \langle f \rangle] \\ & \text{else } \lambda f h. p z [?:P[\text{zero}]] = P[\mathbf{ff} \langle f \rangle] \end{aligned}$$

To fill the holes, we can appeal to $Rx : P[x]$ and then provide the requisite proofs that the numbers coincide: it is trivial to show that they have the same shape, but the child-functions are equal only up to observation:

$$\begin{aligned} & \langle \rangle, \lambda x y q. \overline{(f : 1 \rightarrow \text{Nat})} \langle \rangle y \langle \rangle \\ & : (\text{suc } (f \langle \rangle) : \text{Nat}) = (\mathbf{tt} \langle f : \text{Nat} \rangle) \\ & \langle \rangle, \lambda x y q. x![(x! \text{Nat} : \text{Nat}) = (f y : \text{Nat})] \\ & : (\text{zero} : \text{Nat}) = (\mathbf{ff} \langle f : \text{Nat} \rangle) \end{aligned}$$

So, coercion via proofs based on extensional reasoning repairs the derivation of the induction principle, without loss of canonicity: $2+2$ is certainly a successor. We can play the same game with other inductive sets: related work on *containers* [1] shows how to encode a wide variety of structures as W -sets in an extensional setting. We may now import all those constructions, wholesale, acquiring at least the inductive families accepted by Luo's schema, as used in Epigram [10, 17].

This is progress, but, of course we want more! We want the usual *computational* behaviour as well, on *open* terms. Unfortunately

$$\begin{aligned} \text{ind}_P \text{ ps } p z \text{ zero} & \mapsto p z [\dots : P[\text{zero}]] = P[\text{zero}] \\ \text{ind}_P \text{ ps } p z (\text{suc } n) & \mapsto \\ & p z n (\text{ind}_P \text{ ps } p z n) [\dots : P[\text{suc } n]] = P[\text{suc } n] \end{aligned}$$

If we are lucky, P will have a form which allows these coercions to compute, but in general, they get stuck, even though the source and target are *definitionally* equal sets.

Notice, of course, that this is not an inherent problem with the observational approach, just with the computational behaviour of higher-order encodings of data. If we add first-order presentations of datatypes to our theory, they behave as they always did.

Even with the higher-order encodings, the fact that we lose the computation rules because of coercions which a computer can tell are unnecessary is somewhat ironic. Let us see how to use computers to solve this problem!

6. Type-Directed Quotation

So far, our definitional equality has been the most basic available— α -equivalence of β -normal forms. We can do better, by post-processing our β -normal values in a type-directed way. In the literature of *normalisation by evaluation*, this process is usually called ‘quotation’, as it is usually the means by which semantically-represented values are reified syntactically.

We define mutually recursive operations, $S \uparrow^\Gamma v$ quoting values in a known set as values, $\Downarrow^\Gamma n$ quoting neutral terms as neutral terms and reconstructing the set which they inhabit, $\uparrow^\Gamma S$ quoting normal sets, $\Downarrow^\Gamma N$ quoting neutral sets. These operations perform η -expansion, much in style of Abel et al. [3]. Moreover, they detect the presence of propositional types and mark the corresponding

proofs with a box. We shall also make \Downarrow^Γ eliminate stuck coercions between equal sets.

Definitional equality, \equiv , for sets and for values in a set now becomes α -equivalence of quoted normal forms, written $\equiv_{\alpha\Box}$. That is, we identify quotations up to renaming of variables and the equivalence of arbitrary boxed proofs. We now have *proof irrelevant* propositions.

Quotation for normal values performs η -expansion where possible, pushes through constructors, marks inhabitants of 0 and proofs of neutral propositions (i.e., unexpanded equations). The remaining terms are necessarily neutral inhabitants of datatypes 2, W -sets or some $\text{If } b \text{ Then } T \text{ Else } F$, so we change direction.

$$\begin{aligned} 0 \uparrow^\Gamma z & \mapsto \boxed{z} \\ 1 \uparrow^\Gamma z & \mapsto \langle \rangle \\ 2 \uparrow^\Gamma \mathbf{tt} & \mapsto \mathbf{tt} \\ 2 \uparrow^\Gamma \mathbf{ff} & \mapsto \mathbf{ff} \\ \prod x : S. T \uparrow^\Gamma f & \mapsto \lambda x. (T \uparrow^{\Gamma; \alpha S} f x) \\ \Sigma x : S. T \uparrow^\Gamma p & \mapsto \langle S \uparrow^\Gamma \text{fst } p, T[\text{fst } p] \uparrow^\Gamma \text{snd } p \rangle \\ Wx : S. T \uparrow^\Gamma s \langle f \rangle & \mapsto (S \uparrow^\Gamma s) \langle T[s] \rightarrow Wx : S. T \uparrow^\Gamma f \rangle \\ [N] \uparrow^\Gamma p & \mapsto \boxed{p} \\ T \uparrow^\Gamma n & \mapsto n' \text{ if } \Downarrow^\Gamma n \mapsto n' : _, \text{ otherwise} \end{aligned}$$

For neutral terms, we follow the typing rules

$$\begin{aligned} & \Downarrow^{\Gamma; \alpha S; \Gamma'} x \mapsto x : S \\ & \Downarrow^\Gamma z!S \mapsto \boxed{z} \uparrow^\Gamma S : S \\ \Downarrow^\Gamma \text{if } b/x.P \text{ then } t \text{ else } f \mapsto & \\ & \text{if } 2 \uparrow^\Gamma b/x. \uparrow^{\Gamma; \alpha 2} P \text{ then } P[\mathbf{tt}] \uparrow^\Gamma t \text{ else } P[\mathbf{ff}] \uparrow^\Gamma f : P[b] \\ & \Downarrow^\Gamma f s \mapsto f' (S \uparrow^\Gamma s) : T[s] \\ & \text{if } \Downarrow^\Gamma f \mapsto f' : \prod x : S. T \\ \Downarrow^\Gamma \text{fst } p \mapsto \text{fst } p' : S & \\ & \text{if } \Downarrow^\Gamma p \mapsto p' : \Sigma x : S. T \\ \Downarrow^\Gamma \text{snd } p \mapsto \text{snd } p' : T[\text{fst } p] & \\ & \text{if } \Downarrow^\Gamma p \mapsto p' : \Sigma x : S. T \\ \Downarrow^\Gamma \text{rec } u/w.P \text{ with } p \mapsto \text{rec } u'/w. \uparrow^{\Gamma; u \Delta \text{Nat}; \Gamma} T P & \\ & \text{with } \prod \dots \uparrow^\Gamma p \\ & \text{if } \Downarrow^\Gamma u \mapsto u' : Wx : S. T \end{aligned}$$

where you can recover the missing type of the recursive method from the typing rule for `rec with`. However, the real excitement is in the quotation of neutral coercions.

$$\begin{aligned} \Downarrow^\Gamma s [Q : S = T] & \mapsto T \uparrow^\Gamma s : T \\ & \text{if } \uparrow^\Gamma S \equiv_{\alpha\Box} \uparrow^\Gamma T \\ & \mapsto (S \uparrow^\Gamma s) [\boxed{Q} : (\uparrow^\Gamma S) = (\uparrow^\Gamma T)] : T \\ & \text{otherwise} \end{aligned}$$

This procedure eliminates stuck coercions between equal sets, solving our problem, but does it make sense? The crucial point is that if our coercion is stuck and the sets are equal, s must be neutral: if s were canonical, then S would be canonical and T would be equal to it, hence the coercion would *reduce*.

Quotation for sets is straightforward:

$$\begin{aligned} \uparrow^\Gamma G & \mapsto G & \text{ground sets} \\ \uparrow^\Gamma Bx : S. T & \mapsto Bx : (\uparrow^\Gamma S). (\uparrow^{\Gamma; \alpha S} T) & \text{binder sets} \\ \uparrow^\Gamma N & \mapsto \Downarrow^\Gamma N \end{aligned}$$

$$\begin{aligned} \Downarrow^\Gamma \text{If } b \text{ Then } T \text{ Else } F & \mapsto \text{If } (2 \uparrow^\Gamma b) \text{ Then } (\uparrow^\Gamma T) \text{ Else } (\uparrow^\Gamma F) \\ \Downarrow^\Gamma [S = T] & \mapsto [(\uparrow^\Gamma S) = (\uparrow^\Gamma T)] \\ \Downarrow^\Gamma [(s : S) = (t : T)] & \mapsto [(S \uparrow^\Gamma s : \uparrow^\Gamma S) = (T \uparrow^\Gamma t : \uparrow^\Gamma T)] \end{aligned}$$

What is remarkable about this is that we did not have to change *evaluation*, just our equivalence on values. In particular, our deletion of coercions happens only within the equational theory of neutral terms. Of course, we had better check that

LEMMA 6 (Coercion Elimination).

$$S \equiv T \text{ implies } s [Q:S=T] \equiv s : T$$

even when the coercion is not stuck.

PROOF By induction on the quoted normal forms of S and T . For ground sets, this is exactly the computational behaviour of coercion. For neutral sets, quotation erases the coercion. The interest lies in binder sets.

If $Bx : S_0. T_0 \equiv_{\alpha\Box} Bx : S_1. T_1$ then $S_0 \equiv_{\alpha\Box} S_1$ and $T_0 \equiv_{\alpha\Box} T_1$. Recall that in each case, the coercion operates by coercing a domain component $s [\text{fst } Q:S_i = S_j]$ with i, j being 0, 1 or 1, 0, according to variance. As $S_0 \equiv_{\alpha\Box} S_1$, we have inductively that $s [\text{fst } Q:S_i = S_j] \equiv s$, and hence that $T_i[s [\text{fst } Q:S_i = S_j]] \equiv T_j[s]$ for either twisting. The latter guarantees inductively that the coercion of the range component disappears. Hence the action of coercion between equal sets is at worst η -expansion, which certainly falls within our new equivalence. \square

The effect of this extension of definitional equality is thus to recover the lost computational behaviour of intensional type theories. Substitution by reflexivity is equal to the identity, and the computation rules for Nat’s induction principle hold definitionally.

7. Conclusions and Further Work

Where do we stand now, with Observational Type Theory? The basic system given in this paper has been coded in Agda 2, which gives some evidence for its computational properties and consistency. We have yet to execute our plan to code the construction in a less convenient but more standard system.

We hope to follow Oury’s proof method to show that ETT is a conservative extension of OTT, and we can certainly validate the key axioms underpinning his translation of extensional derivations to intensional terms [22]. His work gives us good reason to conjecture that OTT has the full reasoning power of the corresponding extensional type theory, and hence that we really have no need to suffer the negative computational consequences of the equality reflection rule in order to obtain its logical benefits.

Moreover, as we have shown, we do not need to compromise any of the computational equalities of intensional type theory to achieve this result. Substituting by a reflexive equation is still the identity and induction principles compute as they should. Indeed, we have shown that we can add proof irrelevance to an intensional type theory without serious modification to its evaluation process. In particular, we can now elaborate Epigram’s pattern matching to OTT, retaining all the same computational behaviour, even on open terms.

For a realistic implementation, it is crucial to erase equality proofs and coercions from *run-time* objects. It is perfectly safe to do so, provided we never compute under a binder. This is nothing new: ETTs, where coercions are invisible, have always supported *weak* normalisation, and *program extraction* from ITTs has always supported the erasure of proofs [23].

Three key pieces of the jigsaw are still missing:

Hierarchical universes. We have illustrated the observational approach with a minimal type theory, excluding abstraction over sets. Clearly, we need to introduce a type hierarchy [17] which allows us to scale up. This certainly does not preclude the extension of the observational approach.

Coinductive data. As with functions, the useful notion of equality for coinductive data is observational in character: equal codata have equal one-step unfoldings. Of course, it is consistent to add the propositional law that a *bisimulation* induces equality on codata: as we have seen, this has no impact on the computational properties of the system.

Quotients. We should very much like to internalise setoids—sets of values up to a programmer-supplied equivalence—as quotient sets. We plan to represent quotients as abstract datatypes, allowing you access to the element of the underlying set only if you can prove that you respect the equivalence. Correspondingly, observational equality on quotient sets should just reduce to the given equivalence.

The potential applications of observational equality are considerable. There are many constructions and developments in the literature which have struggled to cope with the rigidity of intensional equality for functions. McBride’s correctness proof for unification [18] is burdened throughout by the need to show explicitly that predicates respect the observational equality for substitutions represented functionally.

Moreover, we expect formalisations of categorical structure to be greatly simplified by the ability to use sets and observational equality rather than setoids with a hand-cranked equivalence. We hope that Buisse and Dybjer will be able to take advantage of OTT to streamline their recent study of categories-with-families, allowing us to model the mathematical structures underlying dependent type systems [6].

If we are to integrate real programming with proof, we need to be able to reason effectively about effects, higher-order objects, processes, abstract data types and the like. That means we need the monad laws to hold; that means we need to reason by observational congruence for processes; that means we need to exploit the equivalences preserved by encapsulation. This paper shows that Observational Type Theory can provide a computational foundation for dependently typed programming, integrated with a logic which steps up to that challenge.

Acknowledgments

The authors wish to express their sincere gratitude to Peter Hancock, for Agda driving lessons and wise insights. Nicolas Oury has made many helpful contributions to our work and is now the prime force behind the implementation of OTT in Epigram 2. We should also like to thank Thierry Coquand and Peter Dybjer for useful conversations, the anonymous referees for their helpful remarks, and Ulf Norell for his spectacular new version of Agda which made this story so much easier to tell.

References

- [1] Michael Abbott, Thorsten Altenkirch, and Neil Ghani. Containers - constructing strictly positive types. *Theoretical Computer Science*, 342:3–27, September 2005.
- [2] Andreas Abel. *A Polymorphic Lambda-Calculus with Sized Higher-Order Types*. PhD thesis, Ludwig-Maximilians-Universität München, 2006.
- [3] Andreas Abel, Thierry Coquand, and Peter Dybjer. Normalization by evaluation for Martin-Löf type theory with typed equality judgements. In *Logic in Computer Science*, 2007.
- [4] Thorsten Altenkirch. Extensional equality in intensional type theory. In *LICS 99*, 1999.
- [5] Lennart Augustsson. Cayenne – a language with dependent types. In *ICFP ’98: Proceedings of the Third ACM SIGPLAN International Conference on Functional Programming*, pages 239–250, 1998.
- [6] Alexandre Buisse and Peter Dybjer. Formalizing categorical models of type theory in type theory. In *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice*, 2007.
- [7] Robert L. Constable, Stuart F. Allen, H. M. Bromley, W. R. Cleaveland, J. F. Cremer, R. W. Harper, Douglas J. Howe, T. B. Knoblock, N. P. Mendler, P. Panangaden, James T. Sasaki, and Scott F.

- Smith. *Implementing Mathematics with the Nuprl Development System*. Prentice-Hall, 1986.
- [8] Thierry Coquand. Pattern matching with dependent types. In B. Nordström, K. Pettersson, and G. Plotkin, editors, *Informal Proceedings Workshop on Types for Proofs and Programs, Båstad, Sweden, 8–12 June 1992*, pages 71–84. Dept. of Computing Science, Chalmers Univ. of Technology and Göteborg Univ., 1992.
- [9] P. Dybjer and A. Setzer. A finite axiomatization of inductive-recursive definitions. *Typed Lambda Calculi and Applications*, 1581:129–146, 1999.
- [10] Peter Dybjer. Inductive Sets and Families in Martin-Löf’s Type Theory. In Gérard Huet and Gordon Plotkin, editors, *Logical Frameworks*. Cambridge University Press, 1991.
- [11] Conor McBride et al. Epigram, 2004. <http://www.e-pig.org>.
- [12] Healdene Goguen, Conor McBride, and James McKinna. Eliminating dependent pattern matching. In Kokichi Futatsugi, Jean-Pierre Jouannaud, and José Meseguer, editors, *Essays Dedicated to Joseph A. Goguen*, volume 4060 of *Lecture Notes in Computer Science*, pages 521–540. Springer, 2006.
- [13] Martin Hofmann. Conservativity of equality reflection over intensional type theory. In *TYPES 95*, pages 153–164, 1995.
- [14] Martin Hofmann. *Extensional concepts in intensional type theory*. PhD thesis, Laboratory for Foundations of Computer Science, University of Edinburgh, 1995. Available from <http://www.lfcs.informatics.ed.ac.uk/reports/95/ECS-LFCS-95-327/>.
- [15] Martin Hofmann and Thomas Streicher. A groupoid model refutes uniqueness of identity proofs. In *LICS 94*, pages 208–212, 1994.
- [16] G. Huet and A. Saïbi. *Constructive Category Theory*. MIT Press, 1998.
- [17] Zhaohui Luo. *Computation and Reasoning: A Type Theory for Computer Science*. Oxford University Press, 1994.
- [18] Conor McBride. *Dependently Typed Functional Programs and their Proofs*. PhD thesis, University of Edinburgh, 1999. Available from <http://www.lfcs.informatics.ed.ac.uk/reports/00/ECS-LFCS-00-419/>.
- [19] Conor McBride. Elimination with a Motive. In Paul Callaghan, Zhaohui Luo, James McKinna, and Robert Pollack, editors, *Types for Proofs and Programs (Proceedings of the International Workshop, TYPES’00)*, volume 2277. Springer-Verlag, 2002.
- [20] Conor McBride and James McKinna. The view from the left. *Journal of Functional Programming*, 14(1):69–111, 2004.
- [21] Ulf Norell. Agda 2. <http://appserv.cs.chalmers.se/users/ulfn/wiki/agda.php>.
- [22] Nicolas Oury. Extensionality in the calculus of constructions. In *TPHOL 05*, pages 278–293, 2005.
- [23] C. Paulin-Mohring. Extracting F_ω ’s programs from proofs in the Calculus of Constructions. In *Sixteenth Annual ACM Symposium on Principles of Programming Languages*, Austin, January 1989. ACM.
- [24] Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Geoffrey Washburn. Simple unification-based type inference for GADTs. In *Proceedings of the Eleventh ACM SIGPLAN International Conference on Functional Programming*, 2006.
- [25] Tim Sheard. Putting Curry-Howard to work. In *Haskell ’05: Proceedings of the 2005 ACM SIGPLAN workshop on Haskell*, pages 74–85, 2005.
- [26] Martin Sulzmann, Manuel M. T. Chakravarty, Simon Peyton Jones, and Kevin Donnelly. System f with type equality coercions. In *TLDI ’07: Proceedings of the 2007 ACM SIGPLAN international workshop on Types in languages design and implementation*, pages 53–66, New York, NY, USA, 2007. ACM Press.
- [27] David A. Turner. A new formulation of constructive type theory. In Peter Dybjer et al., editor, *Proceedings of the Workshop on Programming Logic*. Programming Methodology Group, University of Gothenburg and Chalmers University of Technology, 1989.
- [28] B. Werner. *Une Théorie des Constructions Inductives*. PhD thesis, Université Paris 7, 1994.
- [29] Hongwei Xi. *Dependent Types in Practical Programming*. PhD thesis, Carnegie Mellon University, 1998.