Adding Equations to System F Types

Neelakantan R. Krishnaswami¹ and Nick Benton²

¹ Max Planck Institute for Software Systems

`neelk@mpi-sws.org`

² Microsoft Research

`nick@microsoft.com`

Abstract. We present an extension of System F with types for term-level equations. This internalization of the rich equational theory of the polymorphic lambda calculus yields an expressive core language, suitable for formalizing features such as Haskell’s rewriting rules mechanism or Extended ML signatures.

1 Introduction

Abstraction, modularity and information hiding are fundamental principles of software engineering and language design. Yet programming against an interface is often difficult, as conventional type systems can express only the most basic of the many assumptions and guarantees made by a module. The problem is that *too much* information is hidden, or only present in informal, ambiguous documentation. Dependent types allow much richer interfaces but come with their own costs, including a significant increase in the complexity of the language. Hoare-style program logics also allow much more expressive interfaces, but as well as being highly complex are arguably too decoupled from the underlying programming language; logics have their own syntax and typing cannot exploit logical specifications.

Compilers face problems similar to those of developers. If an optimizing compiler respects the modular structure of a program, it loses vital information that it could use to generate better code. But looking through abstractions to implementations is expensive and non-modular. Some compilers produce enriched interfaces for compiled modules that summarize the results of static analyses but the connection between this metadata and the program is often somewhat ad hoc, making it hard to soundly exploit the extra information in non-trivial transformations of client code.

In this paper, we address the question of how to incorporate more precise module specifications into a language via an extension of System F, the paradigmatic calculus for studying data abstraction. We restrict attention to specifications of a particular form, viz. equations between terms, understood as contextual equivalences. Equations are made part of the type system, making precise how they may be scoped, passed around and exploited. In this regard, our system resembles a restricted form of dependent types. However, equations are not actually proved within the language. A denotational model makes precise what semantic

conditions must be verified in order to establish equalities; various techniques, from automatic analyses to interactive manual proofs, could be soundly used to check the conditions. This aspect resembles the treatment of entailment checking as a delegated semantic side-condition in Hoare-style program logics. The contributions of the paper may be summarized as follows:

- We extend the type system of F with a type of term-level equations.
- We illustrate how expressive this language is by encoding some of patterns of programming with dependent types and GADTs.
- We give a very simple parametric [20,22] relational semantics to our extended language.
- We prove that our extended language is type-safe and terminating.
- We illustrate how the addition of equations enables useful new reasoning patterns in parametricity proofs, such as proofs of the equivalence of existential ML-style module interfaces with Church-style datatype encodings.

2 Informal Overview

This section describes our language, $F_{=}$, semi-formally, and gives examples of its use. The grammar is shown in Figure 1. The types of $F_{=}$ are the usual ones of System F , extended with a new type former $e \equiv_A e'$, which asserts the equality of e and e' at type A . Since types now contain terms, and in particular term-level variables, we change the syntax of the function type from $A \rightarrow B$ to $(x : A) \rightarrow B$, so that function arguments of type A can be referenced within B . This resembles the dependent product of dependent type theory, and ensures that the argument of a function can be mentioned in any returned equality types.

The terms of the language include those of System F , including variables x ; type abstractions $\lambda\alpha. e$ and type applications $e [A]$; and term abstractions $\lambda x. e$ and term application $e e'$. We defer giving the precise typing rules until Section 3, but they very closely resemble their counterparts in System F .

There are also two new term-level constants: $\bullet$ and **abort**. The term $\bullet$ is the sole inhabitant of the equality type $e \equiv_A e'$ when the equation holds (the type is uninhabited otherwise). We remark that $\bullet$ does not provide any intrinsic evidence of the equality — the right to introduce a $\bullet$ arises as a semantic side-condition. The absence of evidence keeps the term language very simple, since the $\bullet$ is merely a placeholder for the reflection of a fact in the semantic model back into the language's type system. The **abort** constant allows equational information to influence typing: it has arbitrary type, but only if the context is (semantically) inconsistent. One can thus do deep semantic proofs to justify complicated equations, inject those into the types, and then use simple syntactic means to handle the plumbing which pushes facts around to other parts of the program.

Figure 2 presents the notational abbreviations that we use in the remainder of the paper. These are mostly well-known Church encodings, but note particularly the weak sum type $\exists x : A. B$. We'll also assume standard syntactic sugar (e.g. projections, case analysis) associated with these abbreviations in examples. (The model we present in Section 3 can be used to show that these suggestive abbreviations actually have the intended semantics.)

Types	$A ::= \alpha \mid \forall \alpha. A \mid (x : A) \rightarrow B \mid e \equiv_A e'$
Terms	$e ::= \Lambda \alpha. e \mid e[A] \mid \lambda x. e \mid e e' \mid \bullet \mid \mathbf{abort} \mid x$
Values	$v ::= \Lambda \alpha. e \mid \lambda x. e \mid \bullet$
Contexts	$\Gamma ::= \cdot \mid \Gamma, \alpha \mid \Gamma, x : A$

Fig. 1. Syntax

$A \rightarrow B$	$= (x : A) \rightarrow B$, when $x \notin \text{FV}(B)$
$\forall x : A. B$	$= (x : A) \rightarrow B$
$\forall x_1 : A_1, \dots, x_n : A_n. B$	$= \forall x_1 : A_1. \dots \forall x_n : A_n. B$
$A \times B$	$= \forall \alpha. (A \rightarrow B \rightarrow \alpha) \rightarrow \alpha$
1	$= \forall \alpha. \alpha \rightarrow \alpha$
$\exists x : A. B$	$= \forall \alpha. ((x : A) \rightarrow B \rightarrow \alpha) \rightarrow \alpha$
$\exists x_1 : A_1, \dots, x_n : A_n. B$	$= \exists x_1 : A_1. \dots \exists x_n : A_n. B$
$\exists \alpha. A$	$= \forall \beta. (\forall \alpha. A \rightarrow \beta) \rightarrow \beta$
$A + B$	$= \forall \alpha. (A \rightarrow \alpha) \rightarrow (B \rightarrow \alpha) \rightarrow \alpha$
$\perp$	$= \forall \alpha. \alpha$
$\neg A$	$= A \rightarrow \perp$

Fig. 2. Type Abbreviations

2.1 Examples

Refined Typings. Equations allow extra constraints to be imposed on arguments and extra guarantees to be given for results. For example, a function that should only be called with commutative binary operations on a type A might be given a type like

$$(f : A \rightarrow A \rightarrow A) \rightarrow (\forall a : A, a' : A. f a a' \equiv_A f a' a) \rightarrow A$$

For producing values together with assertions of their equational properties, one makes use of existential packages. For example, a function yielding idempotent unary operations on A from arguments of type B could be typed as

$$B \rightarrow (\exists f : A \rightarrow A. \forall a : A. f(f a) \equiv_A f a)$$

Clients can project both the underlying value and the equational information from such packages and use them to justify their own equations.

Datatype Encodings. More useful examples of $F_=$ involve enriching the signatures of modules, encoded using second-order existential types in the standard way [17,21]. We now give a few examples of how one can type abstract datatypes in $F_=$, encapsulating types, operations on those types and algebraic properties satisfied by those operations.

Booleans We begin by giving an $F_=$ encoding of an interface to the Booleans.

```

1  ∃bool,
2    true : bool,
3    false : bool,
4    if : ∀α. bool → α → α → α.
5    ∀α, a : α, a' : α. if [α] true a a' ≡α a ×
6    ∀α, a : α, a' : α. if [α] false a a' ≡α a'

```

The signature exposes constructors and eliminators for the boolean type, together with their β -theory as equational properties. A natural question is how this module type relates to the traditional Church encoding of booleans. One can certainly give an implementation of the abstract datatype in terms of that encoding, in which **bool** is instantiated with $\forall\alpha. \alpha \rightarrow \alpha \rightarrow \alpha$. More surprisingly perhaps, as we will show more formally in Section 4, the extended module type uniquely characterizes the booleans, and $F=$ allows clients to exploit the consequences of this semantic fact. Note that the interface does not explicitly state any of the properties which ordinarily characterize datatypes, such as the disjointness of *true* and *false*, or that they are the only constructors for the boolean type. However, the presence of the eliminator and its equational theory, plus parametricity, are sufficient to derive these properties: parametricity ensures that *true* and *false* are the only way to construct the booleans, and furthermore, they must be disjoint, or else we could use the equational theory of *if* to derive a contradiction.

Natural Numbers As an example of a non-finite type, we can give an interface for the type of natural numbers:

```

1  ∃nat,
2    z : nat,
3    s : nat → nat,
4    iter : ∀α. α → (α → α) → nat → α.
5    ∀α, i, f. iter [α] i f z ≡α i ×
6    ∀α, i, f, x. iter [α] i f (s x) ≡α f(iter [α] i f x)

```

This interface says that we have an abstract type **nat**, with constructors *z* and *s*. We have an eliminator *iter*, and two equations explaining how to eliminate zero and successor.

This signature does *not* expose the representation of **nat**, nor does it specify the implementation of *iter*. We are free to implement the natural numbers in any way we like – for example, with a binary (rather than unary) representation. Furthermore, in Section 4 we will prove that this signature is isomorphic to the Church encoding of the natural numbers, which means that any implementation of this type actually is an implementation of the natural numbers.

Lists Here is a possible interface for the type of lists of booleans.

```

1  ∃Listbool,
2    nil : Listbool,

```

```

3  cons : bool → Listbool → Listbool,
4  map : (bool → bool) → Listbool → Listbool,
5  fold : ∀α. α → (bool → α → α) → Listbool → α,
6  map idbool ≡Listbool idListbool ×
7  ∀f. map f nil ≡Listbool nil ×
8  ∀f, b, bs. map f (cons b bs) ≡Listbool cons (f b) (map f bs) ×
9  ∀f, g : bool → bool. (map f) ∘ (map g) ≡Listbool map (f ∘ g) ×
10 ∀α, a, f. fold [α] a f nil ≡α a ×
11 ∀α, a, f, b, bs. fold [α] a f (cons b bs) ≡α f b (fold [α] a f bs)

```

This example illustrates that the interface does not have to precisely match the constructors of the Church encoding for lists — in this interface, we include the *map* operation. However, since the interface tells us what the behavior of *map* is, we can still prove that all values of list type can be built up just from *nil* and *cons*.

This lets us greatly extend the interface to a module, without having to give up natural reasoning principles for it. For example, suppose we extend lists with a left fold operator, characterized by the following signature:

```

12 foldl : ∀α. α → (bool → α → α) → Listbool → α
13   ∀α, i, f. foldl [α] i f nil ≡α i
14   ∀α, i, f, b, bs. foldl [α] i f (cons b bs) ≡α foldl [α] (f b i) f bs

```

With this definition in hand, clients can establish things like the fact that if a function is associative and commutative, then the left and right folds coincide — a fact which we can encode in types:

```

15 assocA(f) ≜ ∀a1, a2, a3. f (f a1 a2) a3 ≡A f a1 (f a2 a3)
16 commA(f) ≜ ∀a1, a2. f a1 a2 ≡A f a2 a1
17   ∀α, f, i, bs. assocα(f) → commα(i) → fold f i bs ≡α foldl f i bs

```

Here we exploit the usual Curry-Howard pun, using the function type to do the duty of a logical implication. The proof of the equation above goes by induction on the *nil*/*cons* structure of lists, which is only possible because we can prove that this structure exists via parametricity.

2.2 Applications

Deforestation of Higher-Order Programs. Consider the following sequence of equational rewrites.

$$\begin{aligned}
 (\text{map not}) \circ (\text{map not}) &= \text{map } (\text{not} \circ \text{not}) \\
 &= \text{map } id_{\text{bool}} \\
 &= id_{\text{List}_{\text{bool}}}
 \end{aligned}$$

This is a standard example of deforestation [15], in which two intermediate data structures (a negated list and a double-negated list) are not generated, in

order to save memory usage. Deforestation offers many interesting examples, since it appeals to equational properties which go well beyond simple inlining and other forms of compile-time β -reduction.

However, now consider the following expression h :

$$h \triangleq \lambda \text{map}. (\text{map } \text{not}) \circ (\text{map } \text{not})$$

The question is, can the body of h be simplified? In general, the answer will be no — unless we can prove that only *map* functionals satisfying the right equational properties flow into the lambda, we are not justified in rewriting this expression. Hence we are in the somewhat ironic position that deforestation — an optimization invented to make higher-order programming more efficient — is often inapplicable in client programs which make use of higher-order functions.

One way around this difficulty would be if we could prove the soundness of these transformations in *open* contexts, where we don't know exactly which lambda-terms might flow into a higher-order program. By adding the necessary information as type data to the programming language, we can rely on the necessary equational properties to hold without having to make the concrete bindings of terms like *map* and *not* visible.

So by rewriting h to pass in the desired properties, we can ensure that a rewriting is performable *in an open context*:

$$\begin{aligned} h : & ((\text{map} : \text{bool} \rightarrow \text{bool}) \rightarrow \text{List}_{\text{bool}} \rightarrow \text{List}_{\text{bool}}) \\ & \rightarrow \text{map } \text{id}_{\text{bool}} \equiv_{\text{List}_{\text{bool}} \rightarrow \text{List}_{\text{bool}}} \text{id}_{\text{List}_{\text{bool}}} \\ & \rightarrow \forall f, g : \text{bool} \rightarrow \text{bool}. (\text{map } f) \circ (\text{map } g) \equiv_{\text{List}_{\text{bool}} \rightarrow \text{List}_{\text{bool}}} \text{map } (f \circ g) \\ & \rightarrow \text{List}_{\text{bool}} \rightarrow \text{List}_{\text{bool}} \\ h \triangleq & \lambda \text{map}. \lambda p f. \lambda p f'. (\text{map } \text{not}) \circ (\text{map } \text{not}) \end{aligned}$$

Here, we do not need to know what the actual implementation of *map* is, since we have stipulated that the function h must be called only with functions which have the equational properties we need them to satisfy, and hence we can conclude that h is equivalent to *id*.

GADT-style Encodings and Making Unsafe Operations Safe. Generalized algebraic data types [11] extend ordinary algebraic datatypes with index information permitting static types to contain information about the specific data constructors used to build them. This lets programmers use dynamic run-time tests to gain additional information about the static type of terms. Since our type system lets us directly place information about terms into types, many of these encodings can be fit into our framework.

Concretely, consider the following specification of an option type.

$$\begin{aligned} 1 \quad & \exists \text{option}_A, \\ 2 \quad & \text{none} : \text{option}_A, \\ 3 \quad & \text{some} : A \rightarrow \text{option}_A, \\ 4 \quad & \text{case} : \forall \alpha. \text{option}_A \rightarrow \alpha \rightarrow (A \rightarrow \alpha) \rightarrow \alpha. \end{aligned}$$

$$\begin{array}{ll}
5 & \forall \alpha. (v : \alpha) \rightarrow (f : A \rightarrow \alpha) \rightarrow \text{case } [\alpha] \text{ none } v f \equiv_{\alpha} v \times \\
6 & \forall \alpha. (a : A) \rightarrow (v : \alpha) \rightarrow (f : A \rightarrow \alpha) \rightarrow \text{case } [\alpha] (\text{some } a) v f \equiv_{\alpha} f a
\end{array}$$

This specification follows the pattern of the boolean and list types earlier, containing the *none* and *some* constructors as well as the *case* eliminator for them, plus equations describing the β -theory of *case*.

We can use this specification to write refined case programs which return type-level evidence of equalities. First, we define a variant case function that returns not only a value, but also a proof that the returned value is equal to the input.

$$\begin{aligned}
\text{case}' & : (x : \text{option}_A) \rightarrow (x \equiv_{\text{option}_A} \text{none}) + (\exists a : A. x \equiv_{\text{option}_A} \text{some } a) \\
\text{case}' & \triangleq \lambda x. \text{case } [\dots] (\text{inl } \bullet) (\lambda a. \text{inr } (a, \bullet))
\end{aligned}$$

As can be seen from the type, *case'* takes an option and returns a sum type. If the argument is *none*, it returns the left branch containing the static fact that the argument is *none*, and if the argument is a *some*, then it returns a value *a* such that *some a* is equal to the argument of *case*. All of the equality type terms are witnessed by $\bullet$ terms.

$$\begin{aligned}
\text{valOf} & : (x : \text{option}_A) \rightarrow \neg(x \equiv_{\text{option}_A} \text{none}) \rightarrow A \\
\text{valOf} & \triangleq \lambda x. \lambda pf. \text{case } [A] \text{ abort } (\lambda a. a)
\end{aligned}$$

This operator takes an option and a proof that it is not equal to *none*. This lets us pass **abort** as an argument in the untaken branch, since we know the *case* can only be reduced when it has a proof that its argument is not *none*.

As before, these operations are provably correct only when injectivity and disjointness hold, and again, these properties are provable from the interface specification. As a result, we can define these apparently-unsafe operators *outside* the body of the module, since our program *valOf* only relies on the equational properties specified in the interface, and not on the specifics of the implementation.

3 Syntax and Semantics

Before proceeding to the metatheory, we give a high-level overview of the structure of this section.

1. We give the syntax of terms and types, and an *untyped* operational semantics for our programming language. This language contains an **abort** construct which can get stuck.
2. We define a “pre-typing” relation, which judges whether terms and types are syntactically well-formed. Unlike a true type system, our pretyping system is (by design) unsound: there are no restrictions on the use of the **abort** connective.
3. However, the pretyping relation offers enough structure that we can define a binary logical relation giving semantics to each of the type constructors, by structural induction on the pretyping relation.

$$\frac{}{v \Downarrow v} \quad \frac{e_0 \Downarrow \lambda x. e'_0 \quad [e_1/x]e'_0 \Downarrow v}{e_0 \ e_1 \Downarrow v} \quad \frac{e_0 \Downarrow \Lambda \alpha. e'_0 \quad [A/\alpha]e'_0 \Downarrow v}{e_0 \ [A] \Downarrow v}$$

Fig. 3. Operational Semantics

$$\begin{array}{c} \boxed{\Gamma \triangleright \text{ok}} \quad \boxed{\Gamma \triangleright A} \quad \boxed{\Gamma \triangleright e : A} \\[10pt] \frac{}{\cdot \triangleright \text{ok}} \quad \frac{\Gamma \triangleright \text{ok}}{\Gamma, \alpha \triangleright \text{ok}} \quad \frac{\Gamma \triangleright \text{ok} \quad \Gamma \triangleright A}{\Gamma, x : A \triangleright \text{ok}} \\[10pt] \frac{\Gamma \triangleright \text{ok} \quad \alpha \in \Gamma}{\Gamma \triangleright \alpha} \quad \frac{\Gamma, \alpha \triangleright A}{\Gamma \triangleright \forall \alpha. A} \quad \frac{\Gamma \triangleright A \quad \Gamma, x : A \triangleright B}{\Gamma \triangleright (x : A) \rightarrow B} \\[10pt] \frac{\Gamma \triangleright A \quad \Gamma \triangleright e : A \quad \Gamma \triangleright e' : A}{\Gamma \triangleright e \equiv_A e'} \\[10pt] \frac{\Gamma, x : A \triangleright e : B}{\Gamma \triangleright \lambda x. e : (x : A) \rightarrow B} \quad \frac{\Gamma \triangleright e : (x : A) \rightarrow B \quad \Gamma \triangleright e' : A}{\Gamma \triangleright e e' : [e/x]B} \\[10pt] \frac{\Gamma, \alpha \triangleright e : A}{\Gamma \triangleright \Lambda \alpha. e : \forall \alpha. A} \quad \frac{\Gamma \triangleright e : \forall \alpha. B \quad \Gamma \triangleright A}{\Gamma \triangleright e [A] : [A/\alpha]B} \quad \frac{\Gamma \triangleright \text{ok} \quad x : A \in \Gamma}{\Gamma \triangleright x : A} \\[10pt] \frac{\Gamma \triangleright e \equiv_A e'}{\Gamma \triangleright \bullet : e \equiv_A e'} \text{ (DANGER1)} \quad \frac{\Gamma \triangleright A}{\Gamma \triangleright \text{abort} : A} \text{ (DANGER2)} \end{array}$$

Fig. 4. Pretyping Relation

4. Then, we give the true typing relation, which refines the pretyping relation to include semantic side-conditions on equality formation and the use of **abort**.
5. Finally, we prove the identity extension lemma for the true typing relation.

Readers familiar with PER models for System F (e.g., [6]) will find this proof structure quite familiar. We begin with an untyped model of computation as a universe, and then define a semantics of types as relations on the universe by induction on the derivation of the pretyping relation. The main technical novelty in our approach is that our types may contain terms, and we thus need to interpret types in a context containing interpretations of the terms.

The operational semantics for our programming language is given in Figure 3, and is a standard call-by-name semantics. There is no evaluation rule for the constant $\bullet$, since it has no explicit elimination form. There is no reduction rule for **abort** — this term creates a stuck computation, since it indicates unreachable code.

In Figure 4, we give the pretyping rules. We have three judgements:

- The $\Gamma \triangleright \text{ok}$ judgement judges whether a context is well-formed, and
- the $\Gamma \triangleright A$ judgement judges whether a type A is well-formed in context Γ , and
- the $\Gamma \triangleright e : A$ judgement judges whether a term e is well-formed with respect to pretype A in context Γ .

All three of these judgements are mutually-recursive, since the equality type $e \equiv_A e'$ contains terms, and its well-formedness rule asserts that e and e' must have pretype A . The rules mostly resemble F's rules, with the variation that both term and type applications need to perform a substitution (rather than solely type application, as in ordinary System F).

The two surprising rules of the system are DANGER1 and DANGER2, which are the rules for introducing the equality type and the **abort** keyword. As a result, the pretype system is obviously unsound, since we can freely introduce the stuck term **abort** wherever we like.

Of course, we will eventually refine these two rules so that equalities can only be used to introduce true equalities, and **abort** can only be used in contexts under which we can prove that evaluation can never reach that point. Do note, however, that in this setting, it is the existence of **abort** which gives the equality type its force. There are no elimination rules for equality types, and so the only way that programs can make use of equalities is to exploit the equations in context to write **abort** at certain places.

Now, we state the basic syntactic substitution properties of the calculus.

Theorem 1. (*Syntactic Substitution*) *Suppose $\Gamma \triangleright A$ and $\Gamma \triangleright e : B$. Then*

- *If $\Gamma, \alpha, \Gamma' \triangleright \text{ok}$ then $\Gamma, [A/\alpha]\Gamma' \triangleright \text{ok}$.*
- *If $\Gamma, \alpha, \Gamma' \triangleright B$ then $\Gamma, [A/\alpha]\Gamma' \triangleright [A/\alpha]B$.*
- *If $\Gamma, \alpha, \Gamma' \triangleright e' : C$ then $\Gamma, [A/\alpha]\Gamma' \triangleright [A/\alpha]e' : [A/\alpha]C$.*
- *If $\Gamma, x : A, \Gamma' \triangleright \text{ok}$ then $\Gamma, [e/x]\Gamma' \triangleright \text{ok}$.*
- *If $\Gamma, x : A, \Gamma' \triangleright B$ then $\Gamma, [e/x]\Gamma' \triangleright [e/x]B$.*
- *If $\Gamma, x : A, \Gamma' \triangleright e' : C$ then $\Gamma, [e/x]\Gamma' \triangleright [e/x]e' : [e/x]C$.*

The proofs of these theorems are a routine structural induction.

To add semantic side-conditions to the DANGER1 and DANGER2 rules, we need to give a relational semantics of types, since we need to be able to compare terms for equality. In Figure 6, we give the logical relation defining the relational interpretation of types, as a structural recursion over the pretyping derivations $\Gamma \triangleright A$. For each type constructor, we give the relation defining equality at that type. Furthermore, since we are defining our relations by induction on the structure of the pretyping derivation $\Gamma \triangleright A$, we also parameterize this relation by a grounding substitution γ .

The two key judgements in this relation begin with $\text{Env}(\Gamma \triangleright \text{ok})$, which defines the set of well-formed grounding substitutions for the environment Γ . As a context consists of a sequence of type variables α and term variables $x : A$, the grounding substitutions consist of sequences of triples (A, A', R) of closed types and the relations between the terms of those types, which ground the type

$$\begin{aligned}
(\cdot)_0 &= \cdot \\
(\gamma, (e, e')/x)_0 &= \gamma_0, (e/x) \\
(\gamma, (A, A', R)/\alpha)_0 &= \gamma_0, (A/\alpha) \\
\\
(\cdot)_1 &= \cdot \\
(\gamma, (e, e')/x)_1 &= \gamma_1, (e'/x) \\
(\gamma, (A, A', R)/\alpha)_1 &= \gamma_1, (A'/\alpha) \\
\\
\gamma(e) &= (\gamma_0(e), \gamma_1(e)) \\
\gamma(A) &= (\gamma_0(A), \gamma_1(A))
\end{aligned}$$

Fig. 5. Operations on Relational Substitutions

variables α , and pairs of expressions (e, e') which lie in the relation for A to ground each term variable.

As γ is a relational substitution, we also need operations to extract ordinary substitutions from it. These operations are defined in Figure 5. Given γ , the substitution γ_0 takes the left components of the relational substitution, and γ_1 takes the right components. We write $\gamma(e)$ as shorthand for the pair $(\gamma_0(e), \gamma_1(e))$, and similarly we write $\gamma(A)$ for $(\gamma_0(A), \gamma_1(A))$.

The environment relation is used mutually-recursively to define the relation $Val(\Gamma \triangleright A)(\gamma)$, which relates pairs of closed values of type A in the context Γ closed by the substitution γ . This definition follows the usual pattern of logical relations: type variables α look up the appropriate relation in the context γ , and the value relation for function space $(x : A) \rightarrow B$ relates two functions f and g if they take related arguments to related results.

The interpretation of the universal quantifier $\forall \alpha. B$ says that two terms are related if for all value relations between pretypes A and A' the type application preserves the relation. By quantifying over relations between arbitrary values, we avoid recursively mentioning the definition of the logical relation, and thereby avoid circularity. This is a syntactic version of the techniques used in PER models of polymorphism: fixing a universe ahead of time lets us consider the intersection of all relations on that universe, without running afoul of the apparent circularity of impredicative quantification.

Finally, we define the value relation for equality types $Val(\Gamma \triangleright e \equiv_A e')(\gamma)$ as the pair $(\bullet, \bullet)$, but *only if* the pair $(\gamma_0(e), \gamma_1(e'))$ is in the relation for A . Otherwise the relation is empty. This gives the semantic sense in which the equality type is an equality type: it is a type containing a single unit value when the equality is true, and is the empty type when it is not.

We also include the definition $Exp(\Gamma \triangleright A)(\gamma)$, which are pairs of expressions reducing to values related by $Val(\Gamma \triangleright A)(\gamma)$. This is an auxiliary definition simplifying the definitions of values and environments.

Having fixed the semantics of types, we give the true typing rules in Figure 7. As before, we have three mutually-recursive judgements, $\Gamma \vdash \text{ok}$, for well-formed contexts, $\Gamma \vdash A$, for well-formed types, and $\Gamma \vdash e : A$ for well-typing. All of

the typing rules precisely mirror the pretyping rules, with the exception of the equality and **abort** rules.

$$\begin{aligned}
 Val(\Gamma \triangleright \alpha)(\gamma) &= \text{let } (A, B, R) = \gamma(\alpha) \text{ in } R \\
 Val(\Gamma \triangleright (x : A) \rightarrow B)(\gamma) &= \\
 &\left\{ \langle \lambda x. e, \lambda x. e' \rangle \left| \begin{array}{l} \cdot \triangleright \lambda x. e : \gamma_0((x : A) \rightarrow B) \text{ and} \\ \cdot \triangleright \lambda x. e' : \gamma_1((x : A) \rightarrow B) \text{ and} \\ \forall e_0, e'_0 \in Exp(\Gamma \triangleright A)(\gamma). \\ \langle [e_0/x]e, [e'_0/x]e' \rangle \in Exp(\Gamma, x : A \triangleright B)(\gamma, \langle e_0, e'_0 \rangle / x) \end{array} \right. \right\} \\
 Val(\Gamma \triangleright \forall \alpha. B)(\gamma) &= \\
 &\left\{ \langle \Lambda \alpha. e, \Lambda \alpha'. e' \rangle \left| \begin{array}{l} \cdot \triangleright \Lambda \alpha. e : \gamma_0(\forall \alpha. B) \text{ and} \\ \cdot \triangleright \Lambda \alpha. e' : \gamma_1(\forall \alpha. B) \text{ and} \\ \forall A, A', R. \\ \cdot \triangleright A \text{ and } \cdot \triangleright A' \text{ and} \\ R \subseteq \{ \langle v, v' \rangle \mid \cdot \triangleright v : A \text{ and } \cdot \triangleright v' : A' \} \text{ and} \\ \langle [A/\alpha]e, [A'/\alpha']e' \rangle \in Exp(\Gamma, \alpha \triangleright B)(\gamma, (A, A', R)/\alpha) \end{array} \right. \right\} \\
 Val(\Gamma \triangleright e_0 \equiv_A e_1)(\gamma) &= \{ \langle \bullet, \bullet \rangle \mid \langle \gamma_0(e_0), \gamma_1(e_1) \rangle \in Exp(\Gamma \triangleright A)(\gamma) \} \\
 Exp(\Gamma \triangleright A)(\gamma) &= \\
 &\left\{ \langle e_0, e_1 \rangle \left| \begin{array}{l} \cdot \triangleright e_0 : \gamma_0(A) \text{ and } \cdot \triangleright e_1 : \gamma_1(A) \text{ and} \\ \exists v_0, v_1. \quad \gamma_0(e_0) \Downarrow v_0 \text{ and } \gamma_1(e_1) \Downarrow v_1 \text{ and} \\ \langle v_0, v_1 \rangle \in Val(\Gamma \triangleright A)(\gamma) \end{array} \right. \right\} \\
 Env(\cdot \triangleright \text{ok}) &= \{ \langle \rangle \} \\
 Env(\Gamma, x : A \triangleright \text{ok}) &= \{ (\gamma, \langle e, e' \rangle / x) \mid \gamma \in Env(\Gamma \triangleright \text{ok}) \text{ and } (e, e') \in Exp(\Gamma \triangleright A)(\gamma) \} \\
 Env(\Gamma, \alpha \triangleright \text{ok}) &= \\
 &\left\{ (\gamma, (A, A', R)/\alpha) \left| \begin{array}{l} \gamma \in Env(\Gamma \triangleright \text{ok}) \text{ and } \cdot \triangleright A \text{ and } \cdot \triangleright A' \text{ and} \\ R \subseteq \{ \langle v, v' \rangle \mid \cdot \triangleright v : A \text{ and } \cdot \triangleright v' : A' \} \end{array} \right. \right\}
 \end{aligned}$$

Fig. 6. Relational Semantics

Each of these has a semantic side-condition controlling when they can be used. These side-conditions mean that the type-checking problem is not decidable, since potentially arbitrary mathematical reasoning may be needed to show that the rule applies. However, the soundness theorem for the language ensures that once the side-conditions are discharged, then evaluation cannot alter the typability of the program under reduction.

The premise of the equality rule contains the non-syntactic premise that $\Gamma \models e = e' : A$. This means that in all semantic environments $\gamma \in Env(\Gamma \triangleright \text{ok})$, the pair $(\gamma_0(e), \gamma_1(e'))$ must lie in the expression relation for the type A . This means that the two expressions must be equivalent to introduce an equality type.

Similarly, the premise of the **abort** rule is that $\Gamma \models \perp$ must hold, which means that there are *no* environments in $Env(\Gamma \triangleright \text{ok})$. This means that the context Γ is a contradictory one, with no environments that can satisfy it.

Now, we can prove a semantic version of the substitution theorem.

Theorem 2. (*Semantic Substitution*)

- Suppose $\Gamma \triangleright A$ and $(\gamma, \gamma(A)/\alpha, \gamma') \in \text{Env}(\Gamma, \alpha, \Gamma' \triangleright \text{ok})$. Then
 - $(\gamma, \gamma') \in \text{Env}(\Gamma, [A/\alpha]\Gamma' \triangleright \text{ok})$
 - If $(v, v') \in \text{Val}(\Gamma, \alpha, \Gamma' \triangleright e)(B)(\gamma, \gamma(A)/\alpha, \gamma')$, then $(v, v') \in \text{Val}(\Gamma, [A/\alpha]\Gamma' \triangleright [A/\alpha]B)(\gamma, \gamma')$.
 - If $(e, e') \in \text{Exp}(\Gamma, \alpha, \Gamma' \triangleright e)(B)(\gamma, \gamma(A)/\alpha, \gamma')$, then $(e, e') \in \text{Exp}(\Gamma, [A/\alpha]\Gamma' \triangleright [A/\alpha]B)(\gamma, \gamma')$.
- Suppose $\Gamma \triangleright e : A$ and $(\gamma, \gamma(e)/x, \gamma') \in \text{Env}(\Gamma, x : A, \Gamma' \triangleright \text{ok})$. Then
 - $(\gamma, \gamma) \in \text{Env}(\Gamma, [e/x]\Gamma' \triangleright \text{ok})$.
 - If $(v, v') \in \text{Val}(\Gamma, x : A, \Gamma' \triangleright e)(A)(\gamma, \gamma(e)/x, \gamma')$, then $(v, v') \in \text{Val}(\Gamma, [e/x]\Gamma' \triangleright [e/x]A)(\gamma, \gamma')$.
 - If $(e, e') \in \text{Exp}(\Gamma, x : A, \Gamma' \triangleright e)(A)(\gamma, \gamma(e)/x, \gamma')$, then $(e, e') \in \text{Exp}(\Gamma, [e/x]\Gamma' \triangleright [e/x]A)(\gamma, \gamma')$.

These theorems follow from induction on the context and type pre-well-formedness judgements. We can use these theorems to prove Reynolds' abstraction theorem for our language.

Theorem 3. (*Abstraction Theorem*) If $\Gamma \vdash e : A$, then $\Gamma \models e = e : A$.

This theorem follows from a structural induction on the typing derivation, making use of the semantic substitution principles. Normalization and type-preservation follow immediately.

Corollary 1. (*Normalization*) If $\cdot \vdash e : A$, then $\exists v$ such that $e \Downarrow v$.

Corollary 2. (*Type Preservation*) If $\cdot \vdash e : A$ and $e \Downarrow v$, then $\cdot \vdash v : A$.

It is worth noting that the type preservation lemma is *exact* — the type of the result is exactly the same as the type of the original. We do not need any notion of type equality beyond the same syntactic equality (modulo α) that System F needed.

4 Existential Representations of Inductive Datatypes

A surprising feature of the examples in Section 2 is that we gave an apparently *existential* encoding of inductive datatypes such as the booleans. This is a little surprising, since the Church encodings of these types in System F are *universal*.

As a concrete example, recall the Church encoding of the boolean type.

- The type of Church booleans $\text{cbool} = \forall \alpha. \alpha \rightarrow \alpha \rightarrow \alpha$.
- Truth is defined as $\Lambda \alpha. \lambda a. \lambda a'. a$.
- Falsity is defined as $\Lambda \alpha. \lambda a. \lambda a'. a'$.
- The conditional is $\text{if} : \text{cbool} \rightarrow \forall \alpha. \alpha \rightarrow \alpha \rightarrow \alpha \triangleq \lambda b. b$.

Contrast this with the interface we gave for the boolean type:

$$\begin{array}{c}
\boxed{\Gamma \models e = e' : A} \qquad \boxed{\Gamma \models \perp} \\
\\
\Gamma \models e_0 = e_1 : A \iff \forall \gamma \in \text{Env}(\Gamma \triangleright \text{ok}). (\gamma_0(e_0), \gamma_1(e_1)) \in \text{Exp}(\Gamma \triangleright A)(\gamma) \\
\Gamma \models \perp \iff \text{Env}(\Gamma \triangleright \text{ok}) = \emptyset \\
\\
\boxed{\Gamma \vdash \text{ok}} \qquad \boxed{\Gamma \vdash A} \qquad \boxed{\Gamma \vdash e : A} \\
\\
\frac{}{\cdot \vdash \text{ok}} \qquad \frac{\Gamma \vdash \text{ok}}{\Gamma, \alpha \vdash \text{ok}} \qquad \frac{\Gamma \vdash \text{ok} \quad \Gamma \vdash A}{\Gamma, x : A \vdash \text{ok}} \\
\\
\frac{\Gamma \vdash \text{ok} \quad \alpha \in \Gamma}{\Gamma \vdash \alpha} \qquad \frac{\Gamma, \alpha \vdash A}{\Gamma \vdash \forall \alpha. A} \qquad \frac{\Gamma \vdash A \quad \Gamma, x : A \vdash B}{\Gamma \vdash (x : A) \rightarrow B} \\
\\
\frac{\Gamma \vdash A \quad \Gamma \vdash e : A \quad \Gamma \vdash e' : A}{\Gamma \vdash e \equiv_A e'} \\
\\
\frac{\Gamma, x : A \vdash e : B}{\Gamma \vdash \lambda x. e : (x : A) \rightarrow B} \qquad \frac{\Gamma \vdash e : (x : A) \rightarrow B \quad \Gamma \vdash e' : A}{\Gamma \vdash e e' : [e/x]B} \qquad \frac{\Gamma, \alpha \vdash e : A}{\Gamma \vdash \Lambda \alpha. e : \forall \alpha. A} \\
\\
\frac{\Gamma \vdash e : \forall \alpha. B \quad \Gamma \vdash A}{\Gamma \vdash e [A] : [A/\alpha]B} \qquad \frac{\Gamma \vdash e \equiv_A e' \quad \Gamma \models e = e' : A}{\Gamma \vdash \bullet : e \equiv_A e'} \qquad \frac{\Gamma \vdash A \quad \Gamma \models \perp}{\Gamma \vdash \text{abort} : A} \\
\\
\frac{\Gamma \vdash \text{ok} \quad x : A \in \Gamma}{\Gamma \vdash x : A}
\end{array}$$

Fig. 7. Typing

```

1  B ≡
2    ∃ bool
3    true : bool,
4    false : bool,
5    if : ∀α. bool → α → α → α.
6    ∀α, a : α, a' : α. if [α] true a a' ≡α a ×
7    ∀α, a : α, a' : α. if [α] false a a' ≡α a'

```

Unlike the Church encoding, the interface completely conceals the representation type of the booleans, as well as the implementations of truth, falsity and if-then-else. The only constraint we place in the interface is to require the β -theory of the booleans to hold.

Now we will show that these two implementations of the booleans are actually the same. To do this, first note that we somehow need to compare an *arbitrary* element of the existential type to a particular set of elements of the Church type. Luckily, we have precisely the tools we need with the equality types of our calculus. The Church booleans can be represented as elements of the type

```

1   B' ≡
2   ∃ true : cbool,
3     false : cbool,
4     true ≡cbool λα. λa. λa'. a ×
5     false ≡cbool λα. λa. λa'. a'

```

By using equality types, we ensure that we have a tuple whose first element is Church truth and whose second element is Church falsity.

This gives us the material we need to prove the following theorem:

Theorem 4. (*Equivalence of Boolean Types*) *We have an isomorphism between the types B and B' .*

Proof. To show this holds, we will give explicit maps $i : B \rightarrow B'$ and $j : B' \rightarrow B$. Then we will show that $\cdot \models i \circ j = id : B'$ and that $\cdot \models j \circ i = id : B$. We give the definitions below, using the syntax for tuples and existentials for clarity.

$$\begin{aligned}
 i : B \rightarrow B' &\triangleq \lambda b. \langle \lambda\alpha. \lambda a. \lambda a'. a, \lambda\alpha. \lambda a. \lambda a'. a', \bullet, \bullet \rangle \\
 &\quad \text{let } t = \lambda\alpha. \lambda a. \lambda a'. a \text{ in} \\
 j : B' \rightarrow B &\triangleq \lambda b'. \text{let } f = \lambda\alpha. \lambda a. \lambda a'. a' \text{ in} \\
 &\quad \text{let } if = \lambda\alpha. \lambda x. \lambda y. \lambda b. b [\alpha] x y \text{ in} \\
 &\quad \text{pack } \langle \text{cbool}, t, f, if, \bullet, \bullet \rangle
 \end{aligned}$$

The $B \rightarrow B'$ direction ignores its argument, and simply returns the obvious tuple inhabiting B' . The $B' \rightarrow B$ direction also ignores its argument, and returns an instance of the existential representation which uses the Church booleans as the representation type.

Therefore, each composition is a constant function, and so showing that it is equivalent to the identity function means showing that all elements of B are equivalent, and similarly for B' . The case for B' is easy, and the interesting case for B reduces to the problem of showing that any element of the existential boolean type is equivalent to element using the Church booleans as its representation type.

This follows from unwinding the definitions. To do this, we introduce a relation that (unsurprisingly) relates Church truth to the true value of the hidden existential type, and Church falsity to the false value of the hidden existential type. Then, the equations for the hidden existential implementation of B can be used to show that the hidden implementation of if is equivalent to the Church implementation.

Similarly, we can relate (an extended version of) the existential natural number interface given in Section 2 with the Church encoding `churchnat` = $\forall\alpha. \alpha \rightarrow (\alpha \rightarrow \alpha) \rightarrow \alpha$.

```

1   N ≡
2   ∃ nat
3   z : nat,

```

```

4      s : nat → nat,
5      pred : nat → optionnat,
6      iter : ∀α. nat → α → (α → α) → α.
7      pred z ≡optionnat none ×
8      ∀n. pred (s n) ≡optionnat some n ×
9      ∀α, i, f. iter [α] z i f ≡α i ×
10     ∀n, α, i, f. iter [α] (s n) i f ≡α f (iter [α] n i f)

1  N' ≡
2  ∃z : churchnat,
3  s : churchnat → churchnat,
4  - : z ≡, Λα. λi. λf. i
5  - : s ≡churchnat→churchnat λn. Λα. λi. λf. f (n α i f)

```

We can then prove the equivalence of these two types.

Theorem 5. (*Equivalence of Natural Number Types*) *There exists an isomorphism between N and N' .*

Proof. The proof of this theorem follows exactly the same pattern as for the booleans. Ultimately we will end up showing that arbitrary elements of N are equivalent to the representation using the Church natural numbers. To do this, we will also need to define the predecessor function $pred$ on the Church naturals, which is a linear time operation.

The most interesting thing about this theorem is not the proof, which is standard, but rather the fact that we extended the natural number interface with the predecessor $pred$. The fact that the representation of natural numbers is completely hidden in the existential style means that we can (for example) use a representation of the natural numbers in which the predecessor is cheap to compute. This contrasts with the explicit unary representation of the Church encoding, in which the predecessor is necessarily linear time. As a result, we can relate this slow implementation to fast ones without any difficulties.

This all relies critically on the equations. In the absence of equations specifying the behavior of the predecessor, there is no way to have this constructor while ensuring that the type really does represent the natural numbers object, since there could be many implementations which are type-correct (in F) but lack the necessary equational properties. However, with equations we can add operations for efficiency without ruining the reasoning properties of the datatype, by cutting down the set of reasonable implementations until only ones equivalent to the intended datatype are possible. (We made extensive use of this in our list example in Section 2.)

This is why we have not proven a general representation theorem for all polynomial datatypes. While a representation theorem does not seem hard to come by in the case where the constructors and fold-style eliminators constitute the interface, it seems that we should consider a representation theorems in the more interesting case in which the interfaces are augmented with extra operations that improve the computational efficiency of implementation. However, it remains unclear to us what such interfaces should be, in general.

5 Related Work

5.1 System R and Plotkin-Abadi Logic

The two most prominent systems for reasoning about parametricity are System R [1,4] and Plotkin-Abadi logic [19]. These logics can be viewed as program logics à la Hoare logic, in that they fix a programming logic (System F), and then give a logical system for reasoning about terms in that language.

Our language can be understood as an attempt to take a small fragment of these logics, and then reflect them back into the types of F. This naturally suggests two directions. First, might it be worthwhile to add more of these logics to the type system of $F_{=}$? In this first paper, we wished to illustrate just how much is achievable with a very modest addition to the type theory of System F, but the extension is a very natural question.

In particular, all of the semantic side-conditions have been discharged by working directly with the relational semantics. By building a logic for parametricity, we could potentially use it to give a proof system for equalities and aborts. However, the presence of **abort** in our language means that such a logic could not be a simple replay of the developments of [1] or [19], though.

5.2 Dependent Types

The appearance of terms in types in our calculus is rather reminiscent of systems of dependent types, such as Martin-Löf type theory [16] or the calculus of constructions [10]. Indeed, the realizability semantics we use for $F_{=}$ is quite similar to the semantics of extensional type theories such as Nuprl [9]. Furthermore, we share with extensional type theory the property that typechecking is not syntax-directed: our proof term for equality, $\bullet$, does not contain the evidence of equality. This is similar to the equality reflection property of extensional type theory, in which proof terms for introducing equalities may depend on propositional equality proofs not evident in the proof term.

However, the semantics of our equality type is a bit different from the equality of dependent type theory. In type theory the elimination form for the equality type $e \equiv_A e'$ is used to cast terms of type $B[e]$ into ones of type $B[e']$. As a result, actually deriving a contradiction (i.e., a terms of type $\perp$) from an impossible equality (e.g., a proof of $0 \equiv_{\mathbb{N}} 1$) requires using a large elimination to turn contradictions into proofs of falsity.

In our setting, we instead admit the use of the **abort** keyword in any inconsistent context, which allows us to make use of contradictions without having to explicitly support large eliminations.

5.3 The Haskell Rules Mechanism

The Glasgow Haskell compiler contains a mechanism called *rules* [12], which allow programmers to specify equational rewrite rules (such as $(\text{map } f) \circ (\text{map } g) \mapsto \text{map } (f \circ g)$) as part of library interfaces. However, these rewrite rules are restricted

to referring to top-level module identifiers, and rewriting cannot be applied to an expression unless the term in question refers to *exactly* the same variables as the rules definition referred to. This restriction means that rules — which were a feature whose purpose is to lower the cost of good higher-order style — are much less effective when applied to higher-order code (where operators such as *map* may flow in as arguments to functions).

Our type theory illustrates that it is possible to integrate Haskell-style rules into a simple type theory treating rewrite rules as first-class types. One particularly interesting direction to investigate is adding rules to type classes, which would permit stating the equational assumptions about polymorphic terms. E.g., Haskell’s `Functor` typeclass has a method with type

$$fmap : \text{Functor } F \Rightarrow (\alpha \rightarrow \beta) \rightarrow F \alpha \rightarrow F \beta$$

It is intended that *fmap* is functorial — that is, that *fmap id* = *id*, and that *(fmap f) ∘ (fmap g)* = *fmap (f ∘ g)*. By placing these equations into the `Functor` interface and verifying the typeclass instances, they could even be used to drive optimizations of client code.

5.4 Extended ML

One of the earliest serious attempts to extend a functional language with equational specifications was the Extended ML [13] project. In this work, SML module signatures were extended with algebraic signatures stating the intended equational properties of the abstract datatypes.

This work was quite ambitious, and it involved a rather large fragment of ML including features such as exceptions and non-termination. Furthermore, the concept of algebraic signature was generalized well beyond equational properties to include full logical predicates. However, the technical ambition of this approach meant that its semantics were never fully settled (the question of polymorphism was especially vexing, as was the specification of imperative ML code).

In this paper, we have avoided effects to maximize the force of parametricity. This lets us specify quite sophisticated properties (e.g., initiality) with a bare minimum of additional syntactic and semantic machinery. One especially nice feature of our work is that the presence of equation makes it very natural to connect Church-style datatype encodings with the existential style of data abstraction more common in ML (and exploited by EML).

These days, there are quite well-developed semantic frameworks in place to model polymorphic languages with features like nontermination, recursive types, and higher order state [18,7]. However, in spite of this machinery, it is simply an unavoidable fact that fewer equations hold when effects are present. To what extent the reduced of validity equational reasoning limits the use of equality types is unclear. One approach to this problem may be to encapsulate effects in a monadic type, and then use other techniques (such as Hoare logic [14]) to reason about the monadic code.

6 Future Work

There are two strands of future work. First, there is the theoretical strand. The first question is whether our termination result can be strengthened into a strong normalization result, which would require a more sophisticated logical relation [2].

Second, it may be possible to give a logic for this calculus along the lines of Plotkin-Abadi logic, and then use the rules of that calculus to give proof terms for the equality type. This would make typechecking decidable, and might make an interesting basis for a dependent type theory with parametricity, along the lines of [5]. While this is a challenging problem, the extreme simplicity of our semantics offers reasonable grounds for hope.

ML-style modules support the “strong” dot-notation elimination form [8], whereas our existential encoding uses F-style existentials with a “weak” let-binding eliminator. Recently, Rossberg, Russo and Dreyer have shown [21] how to translate ML-style modules into System F, and it would be interesting to study if a similar translation could take ML signatures extended with equations and translate into $F_{=}$.

On a practical note, how can equation types be profitably employed in optimizations? Connecting equations to optimizations is an intriguing problem.

Finally, our type system emits proof obligations at each introduction of an equality or use of an `abort`. It would be useful to ship these proof obligations off to a theorem prover such as Coq. Doing so will require a certain amount of care, since parametricity is essential to the arguments we make, and we will need to make use of recent work [3] on representing the semantics of polymorphism in type theory.

References

1. Abadi, M., Cardelli, L., Curien, P.-L.: Formal parametric polymorphism. In: Principles of Programming Languages, pp. 157–170 (1993)
2. Abel, A.: Weak $\beta\eta$ -Normalization and Normalization by Evaluation for System F. In: Cervesato, I., Veith, H., Voronkov, A. (eds.) LPAR 2008. LNCS (LNAI), vol. 5330, pp. 497–511. Springer, Heidelberg (2008)
3. Atkey, R.: Syntax for Free: Representing Syntax with Binding Using Parametricity. In: Curien, P.-L. (ed.) TLCA 2009. LNCS, vol. 5608, pp. 35–49. Springer, Heidelberg (2009)
4. Bellucci, R., Abadi, M., Curien, P.-L.: A Model for Formal Parametric Polymorphism: A PER Interpretation for System R. In: Dezani-Ciancaglini, M., Plotkin, G. (eds.) TLCA 1995. LNCS, vol. 902, pp. 32–46. Springer, Heidelberg (1995)
5. Bernardy, J.-P., Jansson, P., Paterson, R.: Parametricity and dependent types. In: Hudak, P., Weirich, S. (eds.) Proceeding of the 15th ACM SIGPLAN International Conference on Functional Programming, ICFP 2010, pp. 345–356. ACM (2010)
6. Birkedal, L., Møgelberg, R.E., Petersen, R.L.: Domain-theoretical models of parametric polymorphism. Theoretical Computer Science 388(1-3), 152–172 (2007)
7. Birkedal, L., Støvring, K., Thamsborg, J.: Realisability semantics of parametric polymorphism, general references and recursive types. Mathematical Structures in Computer Science 20(4), 655–703 (2010)

8. Cardelli, L., Leroy, X.: Abstract types and the dot notation. In: Broy, M., Jones, C.B. (eds.) *Proceedings IFIP TC2 Working Conference on Programming Concepts and Methods*, pp. 479–504. North-Holland (1990)
9. Constable, R.L., Allen, S.F., Bromley, M., Cleaveland, R., Cremer, J.F., Harper, R.W., Howe, D.J., Knoblock, T.B., Mendler, N.P., Panangaden, P., Sasaki, J.T., Smith, S.F.: *Implementing mathematics with the Nuprl proof development system*. Prentice Hall (1986)
10. Coquand, T., Huet, G.P.: The calculus of constructions. *Inf. Comput.* 76(2/3), 95–120 (1988)
11. Jones, S.P., Vytiniotis, D., Weirich, S., Washburn, G.: Simple unification-based type inference for GADTs. In: Reppy, J.H., Lawall, J.L. (eds.) *Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming, ICFP 2006*, Portland, Oregon, USA, September 16–21, pp. 50–61. ACM (2006)
12. Jones, S.P., Tolmach, A., Hoare, T.: Playing by the rules: rewriting as a practical optimisation technique in GHC. In: *Haskell Workshop* (2001)
13. Kahrs, S., Sannella, D., Tarlecki, A.: The definition of Extended ML: A gentle introduction. *Theoretical Computer Science* 173(2), 445–484 (1997)
14. Krishnaswami, N.R.: *Verifying Higher-Order Imperative Programs with Higher-Order Separation Logic*. PhD thesis, Carnegie Mellon University (2011)
15. Marlow, S., Wadler, P.: Deforestation for higher-order functions. In: Launchbury, J., Sansom, P.M. (eds.) *Functional Programming, Workshops in Computing*, pp. 154–165. Springer, Heidelberg (1992)
16. Martin-Löf, P.: *Intuitionistic Type Theory*. Bibliopolis (1984)
17. Mitchell, J.C., Plotkin, G.D.: Abstract types have existential type. *ACM Trans. Program. Lang. Syst.* 10, 470–502 (1988)
18. Neis, G., Dreyer, D., Rossberg, A.: Non-parametric parametricity. In: Hutton, G., Tolmach, A.P. (eds.) *Proceeding of the 14th ACM SIGPLAN International Conference on Functional Programming*, pp. 135–148. ACM (2009)
19. Plotkin, G.D., Abadi, M.: A Logic for Parametric Polymorphism. In: Bezem, M., Groote, J.F. (eds.) *TLCA 1993*. LNCS, vol. 664, pp. 361–375. Springer, Heidelberg (1993)
20. Reynolds, J.C.: Types, abstraction and parametric polymorphism. In: *IFIP Congress*, pp. 513–523 (1983)
21. Rossberg, A., Russo, C.V., Dreyer, D.: F-ing modules. In: *Proceedings of the 5th ACM SIGPLAN Workshop on Types in Language Design and Implementation, TLDI 2010*, ACM (2010)
22. Wadler, P.: *The Girard-Reynolds isomorphism (second edition)*. *Theoretical Computer Science* 375(1–3), 201–226 (2007)