

Polymorphic Contracts

João Filipe Belo*, Michael Greenberg*,
Atsushi Igarashi†, and Benjamin C. Pierce*

*University of Pennsylvania

†Kyoto University

Abstract. *Manifest contracts* track precise properties by refining types with predicates—e.g., $\{x:\text{Int} \mid x > 0\}$ denotes the positive integers. Contracts and *polymorphism* make a natural combination: programmers can give abstract types strong contracts, precisely stating pre- and postconditions while hiding implementation details—for example, an abstract type of stacks might specify that the `pop` operation has input type $\{x:\alpha \text{ Stack} \mid \text{not}(\text{empty } x)\}$. We formalize this combination by defining F_H , a polymorphic calculus with manifest contracts, and establishing its fundamental properties, including type soundness and relational parametricity. Our development relies on a significant technical improvement over earlier presentations of contracts: instead of introducing a denotational model to break a problematic circularity between typing, subtyping, and evaluation, we develop the metatheory of contracts in a completely syntactic fashion, omitting subtyping from the core system and recovering it *post facto* as a derived property.

Keywords: contracts, refinement types, preconditions, postconditions, dynamic checking, parametric polymorphism, abstract datatypes, syntactic proof, logical relations, subtyping

1 Introduction

Software contracts allow programmers to state precise properties—e.g., that a function takes a non-empty list to a positive integer—as concrete predicates written in the same language as the rest of the program; these predicates can be checked dynamically as the program executes or, more ambitiously, verified statically with the assistance of a theorem prover. Findler and Felleisen [5] introduced “higher-order contracts” for functional languages, which can take one of two forms: predicate contracts like $\{x:\text{Int} \mid x > 0\}$, denoting the positive numbers, and function contracts like $x:\text{Int} \rightarrow \{y:\text{Int} \mid y \geq x\}$, denoting monotonically increasing functions over the integers.

Greenberg, Pierce, and Weirich [7] contrast two different approaches to contracts: in the *manifest* approach, contracts are types—the type system itself makes contracts ‘manifest’; in the *latent* approach, contracts and types live in different worlds (indeed, there may be no types at all, as in PLT Racket’s contract system [1]). These two presentations lead to different ways of checking contracts. Latent systems run contracts with checks: for example, $\langle \{x:\text{Int} \mid x > 0\} \rangle^l n$

checks that $n > 0$. If the check succeeds, then the application will just return n . If it fails, then the entire program will “blame” the label l , raising an uncatchable exception $\uparrow l$, pronounced “blame l ”. Manifest systems use casts, $\langle \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rangle^l$ to convert values from one type to another (the left-hand side is the *source* type and the right-hand side is the *target* type). For predicate contracts, a cast will behave just like a check on the target type: applied to n , the cast either returns n or raises $\uparrow l$. Checks and casts differ when it comes to function contracts. A function check $\langle \langle T_1 \rightarrow T_2 \rangle^l v \rangle v'$ will reduce to $\langle T_2 \rangle^l (v (\langle T_1 \rangle^l v'))$, giving v the argument checked at the domain contract and checking that the result satisfies the codomain contract. A function cast $\langle \langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle^l v \rangle v'$ will reduce to $\langle T_{12} \Rightarrow T_{22} \rangle^l (v (\langle T_{21} \Rightarrow T_{11} \rangle^l v'))$, wrapping the argument v' in a (contravariant) cast between the domain types and wrapping the result of the application in a (covariant) cast between the codomain types. The differences between checks and casts are discussed at length in [7]. Both presentations have their pros and cons: latent contracts are simpler to design and extend, while manifest contracts make a clearer connection between the static constraints captured by types and the dynamic checks performed by casts. In this work, we adopt the manifest approach and endeavor to tame its principal drawback: the complexity of its metatheory.

Subtyping is the source of complexity in the most expressive manifest calculi—those with arbitrary predicates and dependent functions [7, 10]. These calculi have subtyping for two reasons. First, subtyping helps preserve types when evaluating casts with predicate contracts: if $\langle \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rangle^l n \longrightarrow^* n$, then we need to type n at $\{x:\text{Int} \mid x > 0\}$. Subtyping allows us to type n at any predicate contract it satisfies. Second, subtyping can show the equivalence of types with different but related term substitutions. Consider the standard dependent-function application rule:

$$\frac{\Gamma \vdash e_1 : (x:T_1 \rightarrow T_2) \quad \Gamma \vdash e_2 : T_1}{\Gamma \vdash e_1 e_2 : T_2[e_2/x]}$$

If $e_2 \longrightarrow e'_2$, how do $T_2[e_2/x]$ and $T_2[e'_2/x]$ relate? (This is an important question in a proof of preservation!) Subtyping shows that these types are really the same: the first type parallel reduces to the second, and it can be shown that parallel reduction between types implies mutual subtyping—that is, equivalence.

Subtyping brings its own challenges, though. A naïve treatment of subtyping introduces a circularity that leaves the type system undefined. Existing systems break the circularity by defining their systems in a careful order: first, the evaluation relation and the corresponding parallel reduction relation; then a denotational semantics based on the evaluation relation and subtyping based on the denotational semantics; and finally the syntactic type system. Making this carefully sequenced series of definitions hold together requires a long series of tedious and fiddly lemmas relating evaluation and parallel reduction. In short, existing manifest calculi have taken a great deal of effort to construct.

We have developed a simpler approach to manifest calculi that greatly simplifies their definition and metatheory. Rather than using subtyping, we define

a type conversion relation based on parallel reduction—there is no risk of circularity, obviating the need for denotational semantics. We can use this type conversion to give a completely syntactic account of type soundness—with only a few easy lemmas relating evaluation and parallel reduction. We should add that eliminating subtyping doesn’t weaken our approach: it is possible to define a subtyping relation and prove its soundness *post facto*.

We bring this new technique to bear in F_H , a manifest calculus with parametric polymorphism. We should be clear: researchers have already studied the *dynamic* enforcement of parametric polymorphism in languages that mix (conventional, un-refined) static and dynamic typing; we study the *static* enforcement of parametric polymorphism in languages that go beyond conventional static types by adding refinement types and dependent function contracts. We discuss those systems more in Section 6. Concretely, we offer three main contributions:

1. We devise a simpler approach to manifest contract calculi and apply it to F_H , proving type soundness using straightforward syntactic methods [19].
2. We prove that F_H is relationally parametric—establishing that contract checking does not interfere with this desirable property.
3. We define a *post facto* subtyping relation and prove that “upcasts” from subtypes to supertypes always succeed in F_H , i.e., that subtyping is sound.

We begin our technical work with some examples in Section 2. We then describe F_H and prove type soundness in Section 3. We prove parametricity in Section 4 and the upcast lemma in Section 5. We discuss related work in Section 6 and conclude with ideas for future work in Section 7. A long version of this paper with proofs included is available from: http://www.cis.upenn.edu/~mgree/papers/fh_tr.pdf

2 Examples

Like other manifest calculi, F_H checks contracts with casts: the cast $\langle T_1 \Rightarrow T_2 \rangle^l$ takes a value of type T_1 (the source type) and ensures that it behaves (and is treated) like a T_2 (the target type). The l superscript is a *blame label*, used to differentiate between different casts and identify the source of failures. How we check $\langle T_1 \Rightarrow T_2 \rangle^l v$ depends on the structure of T_1 and T_2 . Checking predicate contracts with casts is easy: if v satisfies the predicate of the target type, the entire application goes to v ; if not, then the program aborts, “raising” blame, written $\uparrow l$. For example, $\langle \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rangle^l 5 \longrightarrow^* 5$, since $5 > 0$. But $\langle \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rangle^l 0 \longrightarrow^* \uparrow l$, since $0 \not> 0$. When checking predicate contracts, only the target type matters—the type system guarantees that whatever value we have is well typed at the source type. Checking function contracts is a little trickier: what should $\langle \text{Int} \rightarrow \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rightarrow \{x:\text{Int} \mid x > 0\} \rangle^l v$ do? We can’t just open up v and check whether it always returns positives. The solution is to decompose the cast into its parts:

$$\langle \text{Int} \rightarrow \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rightarrow \{x:\text{Int} \mid x > 0\} \rangle^l v \longrightarrow \lambda x:\{x:\text{Int} \mid x > 0\}. (\langle \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rangle^l (v (\langle \{x:\text{Int} \mid x > 0\} \Rightarrow \text{Int} \rangle^l x)))$$

Note that the domain cast is contravariant, while the codomain is covariant—the context will be forced by the type system to provide a positive number, so we need to cast the input to an appropriate type for v . (In this example, the contravariant cast $\langle\{x:\text{Int} \mid x > 0\} \Rightarrow \text{Int}\rangle^l$ will always succeed.) After v returns, we run the covariant codomain cast to ensure that v didn’t misbehave. So:

$$\begin{aligned} & \langle\text{Int} \rightarrow \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rightarrow \{x:\text{Int} \mid x > 0\}\rangle^l (\lambda x:\text{Int}. x) 5 \longrightarrow^* 5 \\ & \qquad \qquad \qquad \langle\cdots\rangle^l (\lambda x:\text{Int}. 0) 5 \longrightarrow^* \uparrow l \\ & \qquad \qquad \langle\cdots\rangle^l (\lambda x:\text{Int}. 0) (\langle\text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\}\rangle^{l'} 0) \longrightarrow^* \uparrow l' \end{aligned}$$

Note that we omitted the case where a cast function is applied to 0. It is an important property of our system that 0 isn’t well typed at $\{x:\text{Int} \mid x > 0\}$!

As for polymorphism, the standard encodings of existential and product types transfer over to F_H without a problem. Indeed, our dependent functions allow us to go one step further and encode even dependent products such as $(x : \text{Int}) \times \{y:\alpha \text{ List} \mid \text{length } y = x\}$, which represents lists paired with their lengths.

With these preliminaries out of the way, let’s look at an example combining contracts and polymorphism—an abstract datatype of natural numbers.

$$\begin{aligned} \text{NAT} : \exists \alpha. & (\text{zero} : \alpha) \times (\text{succ} : (\alpha \rightarrow \alpha)) \times (\text{iszero} : (\alpha \rightarrow \text{Bool})) \times \\ & (\text{pred} : \{x:\alpha \mid \text{not}(\text{iszero } x)\} \rightarrow \alpha) \end{aligned}$$

(We omit the implementation, a standard Church encoding.) The NAT interface hides our encoding of the naturals behind an existential type, but it also guarantees that `pred` is only ever applied to terms of type $\{x:\alpha \mid \text{not}(\text{iszero } x)\}$. Assuming that $\text{iszero } v \longrightarrow^* \text{true}$ iff $v = \text{zero}$, we can infer that `pred` is never given `zero` as an argument. Consider the following expression, where I is the interface we specified for NAT and we omit the term binding for brevity:

$$\text{unpack NAT} : \exists \alpha. I \text{ as } \alpha, _ \text{ in } \text{pred} (\langle\langle \alpha \Rightarrow \{x:\alpha \mid \text{not}(\text{iszero } x)\} \rangle\rangle^l \text{zero}) : \alpha$$

The application of `pred` directly to `zero` would not be well typed, since $\text{zero} : \alpha$. On the other hand, this term is well typed, since we cast `zero` to the type we need. This program will ultimately raise $\uparrow l$, because $\text{not}(\text{iszero } \text{zero}) \longrightarrow^* \text{false}$.

The example so far imposes constraints only on the *use* of the abstract datatype, in particular on the use of `pred`. To have constraints imposed also on the *implementation* of the abstract data type, consider the extension of the interface with a subtraction operation, `sub`, and a “less than or equal” predicate, `leq`. We now have the interface:

$$I' = I \times (\text{leq} : \alpha \rightarrow \alpha \rightarrow \text{Bool}) \times (\text{sub} : (x:\alpha \rightarrow \{y:\alpha \mid \text{leq } y \ x\} \rightarrow \{z:\alpha \mid \text{leq } z \ x\}))$$

The `sub` function requires that its second argument isn’t greater than the first, and it promises to return a result that isn’t greater than the first argument.

We get contracts in interfaces by putting casts in the implementations. For example, the contracts on `pred` and `sub` are imposed when we “pack up” NAT; we write `nat` for the implementation type:

$$\text{pack } \langle \text{nat}, (\text{zero}, \text{succ}, \text{iszero}, \text{pred}, \text{leq}, \text{sub}) \rangle \text{ as } \exists \alpha. I'$$

Types and contexts

$$T ::= B \mid \alpha \mid x:T_1 \rightarrow T_2 \mid \forall\alpha. T \mid \{x:T \mid e\} \qquad \Gamma ::= \emptyset \mid \Gamma, x:T \mid \Gamma, \alpha$$

Terms

$$\begin{aligned} e &::= x \mid k \mid \text{op}(e_1, \dots, e_n) \mid \lambda x:T. e \mid \Lambda\alpha. e \mid e_1 e_2 \mid e T \mid \\ &\quad \langle T_1 \Rightarrow T_2 \rangle^l \mid \uparrow^l \mid \langle \{x:T \mid e_1\}, e_2, v \rangle^l \\ v &::= k \mid \lambda x:T. e \mid \Lambda\alpha. e \mid \langle T_1 \Rightarrow T_2 \rangle^l \qquad r ::= v \mid \uparrow^l \\ E &::= [] e_2 \mid v_1 [] \mid [] T \mid \langle \{x:T \mid e\}, [], v \rangle^l \mid \text{op}(v_1, \dots, v_{i-1}, [], e_{i-1}, \dots, e_n) \end{aligned}$$

Fig. 1. Syntax for F_H

where:

$$\begin{aligned} \text{pred} &= \langle \text{nat} \rightarrow \text{nat} \Rightarrow \{x:\text{nat} \mid \text{not}(\text{iszero } x)\} \rightarrow \text{nat} \rangle^l \text{pred}' \\ \text{sub} &= \langle \text{nat} \rightarrow \text{nat} \rightarrow \text{nat} \Rightarrow x:\text{nat} \rightarrow \{y:\text{nat} \mid \text{leq } y x\} \rightarrow \{z:\text{nat} \mid \text{leq } z x\} \rangle^l \text{sub}' \end{aligned}$$

That is, the existential type dictates that we must pack up *cast* versions of our implementations, pred' and sub' . Note, however, that the cast on pred' will never actually check anything at runtime: if we unfold the domain contract contravariantly, we see that $\langle \{x:\text{nat} \mid \text{not}(\text{iszero } x)\} \Rightarrow \text{nat} \rangle^l$ is a no-op. Instead, clients of NAT can only call pred with terms that are typed at $\{x:\text{nat} \mid \text{not}(\text{iszero } x)\}$, i.e., by checking that values are nonzero with a cast into pred' 's input type. The story is the same for the contract on sub' 's second argument—the contravariant cast won't actually check anything. The codomain contract on sub , however, could fail if sub' mis-implemented subtraction.

We can sum up the situation for contracts in interfaces as follows: the positive parts of the interface type are checked and can raise blame—these parts are the responsibility of the implementation; the negative parts of the interface type are not checked by the implementation—clients must check these themselves before calling functions from the ADT. This pattern is reminiscent of the original idea of client and server blame found in Findler and Felleisen's seminal work [5].

3 Defining F_H

The syntax of F_H is given in Figure 1. For unrefined types we have: base types B , which must include **Bool**; type variables α ; dependent function types $x:T_1 \rightarrow T_2$ where x is bound in T_2 ; and universal types $\forall\alpha. T$, where α is bound in T . Aside from dependency in function types, these are just the types of the standard polymorphic lambda calculus. We also have predicate contracts, or *refinement types*, written $\{x:T \mid e\}$. Conceptually, $\{x:T \mid e\}$ denotes values v of type T for which $e[v/x]$ reduces to **true**. For each B , we fix a set \mathcal{K}_B of the constants in that type; we require our typing rules for constants and our typing and evaluation rules for operations to respect this set. We also require that $\mathcal{K}_{\text{Bool}} = \{\text{true}, \text{false}\}$.

In the syntax of terms, the first line is standard for a call-by-value polymorphic language: variables, constants, first-order operations (on one or more arguments), term and type abstractions, and term and type applications. The

second line offers the standard constructs of a manifest contract calculus [6, 7, 10], with a few alterations, discussed below.

Casts are the distinguishing feature of manifest contract calculi. When applied to a value of type T_1 , the cast $\langle T_1 \Rightarrow T_2 \rangle^l$ ensures that its argument behaves—and is treated—like a value of type T_2 . When a cast detects a problem, it raises blame, a label-indexed uncatchable exception written $\uparrow l$. The label l allows us to trace blame back to a specific cast. (While our labels here are drawn from an arbitrary set, in practice l will refer to a source-code location.) Finally, we use active checks $\langle \{x:T \mid e_1\}, e_2, v \rangle^l$ to support a small-step semantics for checking casts into refinement types. In an active check, $\{x:T \mid e_1\}$ is the refinement being checked, e_2 is the current state of checking, and v is the value being checked. If checking succeeds, the check will return v ; if checking fails, the check will blame its label, raising $\uparrow l$. Active checks and blame are not intended to occur in source programs—they are runtime devices. (In a real programming language based on this calculus, casts will probably not appear explicitly either, but will be inserted by an elaboration phase. The details of this process are beyond the scope of the present work.)

The values in F_H are constants, term and type abstractions, and casts. We also define *results*, which are either values or blame. (Type soundness—a consequence of Theorems 2 and 3 below—will show that evaluation produces a result, but not necessarily a value.) In some earlier work [7, 8], casts between function types applied to values were themselves considered values. We make the other choice here: excluding applications from the possible syntactic forms of values simplifies our inversion lemmas.

There are two notable features relative to existing manifest calculi: first, *any* type (even a refinement type) can be refined, not just base types (as in [6–8, 10, 12]); second, the third part of the active check form $\langle \{x:T \mid e_1\}, e_2, v \rangle^l$ can be any value, not just a constant. Both of these changes are motivated by the introduction of polymorphism. In particular, to support refinement of type variables we must allow refinements of *all* types, since any type can be substituted in for a variable.

Operational semantics

The call-by-value operational semantics in Figure 2 are given as a small-step relation, split into two sub-relations: one for reductions (\rightsquigarrow) and one for congruence and blame lifting (\longrightarrow).

The latter relation is standard. The `E_REDUCE` rule lifts \rightsquigarrow reductions into \longrightarrow ; the `E_COMPAT` rule turns \longrightarrow into a congruence over our evaluation contexts; and the `E_BLAKE` rule lifts blame, treating it as an uncatchable exception. The reduction relation \rightsquigarrow is more interesting. There are four different kinds of reductions: the standard lambda calculus reductions, structural cast reductions, cast staging reductions, and checking reductions.

The `E_OP`, `E_BETA`, and `E_TBETA` rules should need no explanation—these are the standard call-by-value polymorphic lambda calculus reductions.

Reduction rules $\boxed{e_1 \rightsquigarrow e_2}$

$$\begin{array}{l}
\text{op}(v_1, \dots, v_n) \rightsquigarrow [\text{op}](v_1, \dots, v_n) \quad \text{E_OP} \\
(\lambda x:T_1. e_{12}) v_2 \rightsquigarrow e_{12}[v_2/x] \quad \text{E_BETA} \\
(\Lambda \alpha. e) T \rightsquigarrow e[T/\alpha] \quad \text{E_TBETA} \\
\langle T \Rightarrow T \rangle^l v \rightsquigarrow v \quad \text{E_REFL} \\
\langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l v \rightsquigarrow \\
\lambda x:T_{21}. (\langle T_{12}[(\langle T_{21} \Rightarrow T_{11} \rangle^l x/x) \Rightarrow T_{22}]^l (v(\langle T_{21} \Rightarrow T_{11} \rangle^l x)) \rangle \\
\text{when } x:T_{11} \rightarrow T_{12} \neq x:T_{21} \rightarrow T_{22} \quad \text{E_FUN} \\
\langle \forall \alpha. T_1 \Rightarrow \forall \alpha. T_2 \rangle^l v \rightsquigarrow \Lambda \alpha. (\langle T_1 \Rightarrow T_2 \rangle^l (v \alpha)) \quad \text{E_FORALL} \\
\text{when } \forall \alpha. T_1 \neq \forall \alpha. T_2 \\
\langle \{x:T_1 \mid e\} \Rightarrow T_2 \rangle^l v \rightsquigarrow \langle T_1 \Rightarrow T_2 \rangle^l v \quad \text{E_FORGET} \\
\text{when } T_2 \neq \{x:T_1 \mid e\} \text{ and } T_2 \neq \{y:\{x:T_1 \mid e\} \mid e_2\} \\
\langle T_1 \Rightarrow \{x:T_2 \mid e\} \rangle^l v \rightsquigarrow \langle T_2 \Rightarrow \{x:T_2 \mid e\} \rangle^l (\langle T_1 \Rightarrow T_2 \rangle^l v) \quad \text{E_PRECHECK} \\
\text{when } T_1 \neq T_2 \text{ and } T_1 \neq \{x:T' \mid e'\} \\
\langle T \Rightarrow \{x:T \mid e\} \rangle^l v \rightsquigarrow \langle \{x:T \mid e\}, e[v/x], v \rangle^l \quad \text{E_CHECK} \\
\langle \{x:T \mid e\}, \text{true}, v \rangle^l \rightsquigarrow v \quad \text{E_OK} \\
\langle \{x:T \mid e\}, \text{false}, v \rangle^l \rightsquigarrow \uparrow l \quad \text{E_FAIL}
\end{array}$$

Evaluation rules $\boxed{e_1 \longrightarrow e_2}$

$$\frac{e_1 \rightsquigarrow e_2}{e_1 \longrightarrow e_2} \quad \text{E_REDUCE} \quad \frac{e_1 \longrightarrow e_2}{E[e_1] \longrightarrow E[e_2]} \quad \text{E_COMPAT} \quad \frac{}{E[\uparrow l] \longrightarrow \uparrow l} \quad \text{E_BLAME}$$

Fig. 2. Operational semantics

The E_REFLECT, E_FUN, and E_FORALL rules are structural cast reductions. E_REFLECT eliminates a cast from a type to itself; intuitively, such a cast should always succeed anyway. (We discuss this rule more in Section 4.) When a cast between function types is applied to a value v , the E_FUN rule produces a new lambda, wrapping v with a contravariant cast on the domain and covariant cast on the codomain. The extra substitution in the left-hand side of the codomain cast may seem suspicious, but in fact the rule must be this way in order for type preservation to hold (see [7] for an explanation). The E_FORALL rule is similar to E_FUN, generating a type abstraction with the necessary covariant cast. Side conditions on E_FORALL and E_FUN ensure that these rules apply only when E_REFLECT doesn't.

The E_FORGET, E_PRECHECK, and E_CHECK rules are cast-staging reductions, breaking a complex cast down to a series of simpler casts and checks. All of these rules require that the left- and right-hand sides of the cast be different—if they are the same, then E_REFLECT applies. The E_FORGET rule strips a layer of refinement off the left-hand side; in addition to requiring that the left- and right-hand sides are different, the preconditions require that the right-hand side isn't a refinement of the left-hand side. The E_PRECHECK rule breaks a cast into two

parts: one that checks exactly one level of refinement and another that checks the remaining parts. This rule applies when the two sides of the cast are different and when the left-hand side isn't a refinement. The `E_CHECK` rule applies when the right-hand side refines the left-hand side; it takes the cast value and checks that it satisfies the right-hand side. (We don't have to check the left-hand side, since that's the type we're casting *from*.)

Before explaining how these rules interact in general, we offer a few examples. First, here's a reduction using `E_CHECK`, `E_COMPAT`, `E_OP`, and `E_OK`:

$$\begin{aligned} \langle \text{Int} \Rightarrow \{x:\text{Int} \mid x \geq 0\} \rangle^l 5 &\longrightarrow \langle \{x:\text{Int} \mid x \geq 0\}, 5 \geq 0, 5 \rangle^l \\ &\longrightarrow \langle \{x:\text{Int} \mid x \geq 0\}, \text{true}, 5 \rangle^l \longrightarrow 5 \end{aligned}$$

A failed check will work the same way until the last reduction, which will use `E_FAIL` rather than `E_OK`:

$$\begin{aligned} \langle \text{Int} \Rightarrow \{x:\text{Int} \mid x \geq 0\} \rangle^l (-1) &\longrightarrow \langle \{x:\text{Int} \mid x \geq 0\}, -1 \geq 0, -1 \rangle^l \\ &\longrightarrow \langle \{x:\text{Int} \mid x \geq 0\}, \text{false}, -1 \rangle^l \longrightarrow \uparrow^l \end{aligned}$$

Notice that the blame label comes from the cast that failed. Here's a similar reduction that needs some staging, using `E_FORGET` followed by the first reduction we gave:

$$\begin{aligned} \langle \{x:\text{Int} \mid x = 5\} \Rightarrow \{x:\text{Int} \mid x \geq 0\} \rangle^l 5 &\longrightarrow \langle \text{Int} \Rightarrow \{x:\text{Int} \mid x \geq 0\} \rangle^l 5 \\ &\longrightarrow \langle \{x:\text{Int} \mid x \geq 0\}, 5 \geq 0, 5 \rangle^l \longrightarrow^* 5 \end{aligned}$$

There are two cases where we need to use `E_PRECHECK`. First, when multiple refinements are involved:

$$\begin{aligned} \langle \{y:\text{Int} \mid y \geq 0\} \Rightarrow \{x:\{y:\text{Int} \mid y \geq 0\} \mid x = 5\} \rangle^l (\langle \text{Int} \Rightarrow \{y:\text{Int} \mid y \geq 0\} \rangle^l 5) &\longrightarrow^* \\ \langle \{y:\text{Int} \mid y \geq 0\} \Rightarrow \{x:\{y:\text{Int} \mid y \geq 0\} \mid x = 5\} \rangle^l (\langle \text{Int} \Rightarrow \{y:\text{Int} \mid y \geq 0\} \rangle^l 5) &\longrightarrow^* \\ \langle \{y:\text{Int} \mid y \geq 0\} \Rightarrow \{x:\{y:\text{Int} \mid y \geq 0\} \mid x = 5\} \rangle^l 5 &\longrightarrow^* \\ \langle \{x:\{y:\text{Int} \mid y \geq 0\} \mid x = 5\}, 5 = 5, 5 \rangle^l &\longrightarrow^* \\ 5 & \end{aligned}$$

Second, when casting a function or universal type into a refinement of a *different* function or universal type.

$$\begin{aligned} \langle \text{Bool} \rightarrow \{x:\text{Bool} \mid x\} \Rightarrow \{f:\text{Bool} \rightarrow \text{Bool} \mid f \text{ true} = f \text{ false}\} \rangle^l v &\longrightarrow \\ \langle \text{Bool} \rightarrow \text{Bool} \Rightarrow \{f:\text{Bool} \rightarrow \text{Bool} \mid f \text{ true} = f \text{ false}\} \rangle^l & \\ \langle \langle \text{Bool} \rightarrow \{x:\text{Bool} \mid x\} \Rightarrow \text{Bool} \rightarrow \text{Bool} \rangle^l v & \end{aligned}$$

`E_REFL` is necessary for simple cases, like $\langle \text{Int} \Rightarrow \text{Int} \rangle^l 5 \longrightarrow 5$. Hopefully, such a silly cast would never be written, but it could arise as a result of `E_FUN` or `E_FORALL`. (We also need `E_REFL` in our proof of parametricity; see Section 4.)

Cast evaluation follows a regular schema:

$$\text{REFL} \mid (\text{FORGET}^* (\text{REFL} \mid (\text{PRECHECK}^* (\text{REFL} \mid \text{FUN} \mid \text{FORALL})? \text{CHECK}^*)))$$

Let's consider the cast $\langle T_1 \Rightarrow T_2 \rangle^l v$. First, if $T_1 = T_2$, we can apply `E_REFL` and be done with it. If that doesn't work, we'll reduce by `E_FORGET` until the

left-hand side doesn't have any refinements. (N.B. we may not have to make any of these reductions.) Either all of the refinements will be stripped away from the source type, or `E_REFL` eventually applies and the entire cast disappears. Assuming `E_REFL` doesn't apply, we now have $\langle \text{unref}(T_1) \Rightarrow T_2 \rangle^l v$. Next, we apply `E_PRECHECK` until the cast is completely decomposed into one-step casts, once for each refinement in T_2 :

$$\langle \text{unref}_1(T_2) \Rightarrow T_2 \rangle^l (\langle \text{unref}_2(T_2) \Rightarrow \text{unref}_1(T_2) \rangle^l (\dots (\langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l v) \dots))$$

(We write unref_n to strip away exactly n levels of refinements.) As our next step, we apply whichever structural cast rule applies to $\langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l v$, one of `E_REFL`, `E_FUN`, or `E_FORALL`. Now all that remains are some number of refinement checks, which can be dispatched by the `E_CHECK` rule (and other rules, of course, during the predicate checks themselves).

Static typing

The type system comprises three mutually recursive judgments: context well formedness, type well formedness, and term well typing. The rules for contexts and types are unsurprising. The rules for terms are mostly standard. First, the `T_APP` rule is dependent, to account for dependent function types. The `T_CAST` rule is standard for manifest calculi, allowing casts between compatibly structured well formed types. Compatibility of type structures is defined in Figure 4; in short, compatible types erase to identical simple type skeletons. Note that we assign casts a non-dependent function type.

Some of the typing rules—`T_CHECK`, `T_BLAKE`, `T_EXACT`, `T_FORGET`, and `T_CONV`—are “runtime only”. We don't expect to use these rules to type check source programs, but we need them to guarantee preservation. Note that the conclusions of these rules use a context Γ , but their premises don't use Γ at all. Even though runtime terms and their typing rules should only ever occur in an empty context, the `T_APP` rule substitutes terms into types—so a runtime term could end up under a binder. We therefore allow the runtime typing rules to apply in any well formed context, so long as the terms they type check are closed. The `T_BLAKE` rule allows us to give any type to blame—this is necessary for preservation. The `T_CHECK` rule types an active check, $\langle \{x:T \mid e_1\}, e_2, v \rangle^l$. Such a term arises when a term like $\langle T \Rightarrow \{x:T \mid e_1\} \rangle^l v$ reduces by `E_CHECK`. The premises of the rule are all intuitive except for $e_1[v/x] \longrightarrow^* e_2$, which is necessary to avoid nonsensical terms like $\langle \{x:T \mid x \geq 0\}, \text{true}, -1 \rangle^l$, where the wrong predicate gets checked. The `T_EXACT` rule allows us retype a closed value of type T at $\{x:T \mid e\}$ if $e[v/x] \longrightarrow^* \text{true}$. This typing rule guarantees type preservation for `E_OK`: $\langle \{x:T \mid e_1\}, \text{true}, v \rangle^l \longrightarrow v$. If the active check was well typed, then we know that $e_1[v/x] \longrightarrow^* \text{true}$, so `T_EXACT` applies. Finally, the `T_CONV` rule allows us to retype expressions at convertible types: if $\emptyset \vdash e : T$ and $T \equiv T'$, then $\emptyset \vdash e : T'$ (or in any well formed context Γ). We define \equiv as the symmetric, transitive closure of call-by-value respecting parallel reduction,

Context well formedness $\boxed{\vdash \Gamma}$

$$\frac{}{\vdash \emptyset} \text{WF_EMPTY} \quad \frac{\vdash \Gamma \quad \Gamma \vdash T}{\vdash \Gamma, x:T} \text{WF_EXTENDVAR} \quad \frac{\vdash \Gamma}{\vdash \Gamma, \alpha} \text{WF_EXTENDTVAR}$$

Type well formedness $\boxed{\Gamma \vdash T}$

$$\frac{\vdash \Gamma}{\Gamma \vdash B} \text{WF_BASE} \quad \frac{\vdash \Gamma \quad \alpha \in \Gamma}{\Gamma \vdash \alpha} \text{WF_TVAR} \quad \frac{\Gamma, \alpha \vdash T}{\Gamma \vdash \forall \alpha. T} \text{WF_FORALL}$$

$$\frac{\Gamma \vdash T_1 \quad \Gamma, x:T_1 \vdash T_2}{\Gamma \vdash x:T_1 \rightarrow T_2} \text{WF_FUN} \quad \frac{\Gamma \vdash T \quad \Gamma, x:T \vdash e : \mathbf{Bool}}{\Gamma \vdash \{x:T \mid e\}} \text{WF_REFINE}$$

Term typing $\boxed{\Gamma \vdash e : T}$

$$\frac{\vdash \Gamma \quad x:T \in \Gamma}{\Gamma \vdash x : T} \text{T_VAR} \quad \frac{\vdash \Gamma}{\Gamma \vdash k : \text{ty}(k)} \text{T_CONST} \quad \frac{\emptyset \vdash T \quad \vdash \Gamma}{\Gamma \vdash \uparrow l : T} \text{T_BLAME}$$

$$\frac{\Gamma, x:T_1 \vdash e_{12} : T_2}{\Gamma \vdash \lambda x:T_1. e_{12} : x:T_1 \rightarrow T_2} \text{T_ABS} \quad \frac{\Gamma \vdash e_1 : (x:T_1 \rightarrow T_2) \quad \Gamma \vdash e_2 : T_1}{\Gamma \vdash e_1 e_2 : T_2[e_2/x]} \text{T_APP}$$

$$\frac{\vdash \Gamma \quad \text{ty}(\text{op}) = x_1 : T_1 \rightarrow \dots \rightarrow x_n : T_n \rightarrow T \quad \Gamma \vdash e_i[e_1/x_1, \dots, e_{i-1}/x_{i-1}] : T_i[e_1/x_1, \dots, e_{i-1}/x_{i-1}]}{\Gamma \vdash \text{op}(e_1, \dots, e_n) : T[e_1/x_1, \dots, e_n/x_n]} \text{T_OP}$$

$$\frac{\Gamma, \alpha \vdash e : T}{\Gamma \vdash \Lambda \alpha. e : \forall \alpha. T} \text{T_TABS} \quad \frac{\Gamma \vdash e_1 : \forall \alpha. T \quad \Gamma \vdash T_2}{\Gamma \vdash e_1 T_2 : T[T_2/\alpha]} \text{T_TAPP}$$

$$\frac{\Gamma \vdash T_1 \quad \Gamma \vdash T_2 \quad T_1 \parallel T_2}{\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle^l : _ : T_1 \rightarrow T_2} \text{T_CAST}$$

$$\frac{\vdash \Gamma \quad \emptyset \vdash \{x:T \mid e_1\} \quad \emptyset \vdash v : T \quad \emptyset \vdash e_2 : \mathbf{Bool} \quad e_1[v/x] \longrightarrow^* e_2}{\Gamma \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle^l : \{x:T \mid e_1\}} \text{T_CHECK}$$

$$\frac{\vdash \Gamma \quad \emptyset \vdash e : T \quad \emptyset \vdash T' \quad T \equiv T'}{\Gamma \vdash e : T'} \text{T_CONV} \quad \frac{\emptyset \vdash v : \{x:T \mid e\} \quad \vdash \Gamma}{\Gamma \vdash v : T} \text{T_FORGET}$$

$$\frac{\vdash \Gamma \quad \emptyset \vdash v : T \quad \emptyset \vdash \{x:T \mid e\} \quad e[v/x] \longrightarrow^* \mathbf{true}}{\Gamma \vdash v : \{x:T \mid e\}} \text{T_EXACT}$$

Fig. 3. Typing rules

$$\begin{array}{c}
\text{Type compatibility } \boxed{T_1 \parallel T_2} \\
\\
\frac{}{T \parallel T} \text{ C_REFL} \qquad \frac{T_1 \parallel T_2}{\{x:T_1 \mid e\} \parallel T_2} \text{ C_REFINEL} \qquad \frac{T_1 \parallel T_2}{T_1 \parallel \{x:T_2 \mid e\}} \text{ C_REFINER} \\
\\
\frac{T_{11} \parallel T_{21} \quad T_{12} \parallel T_{22}}{x:T_{11} \rightarrow T_{12} \parallel x:T_{21} \rightarrow T_{22}} \text{ C_FUN} \qquad \frac{T_1 \parallel T_2}{\forall\alpha. T_1 \parallel \forall\alpha. T_2} \text{ C_FORALL}
\end{array}$$

Operations on refinements

$$\begin{array}{l}
\text{unref}(T) = \begin{cases} \text{unref}(T') & \text{if } T = \{x:T' \mid e\} \\ T & \text{otherwise} \end{cases} \\
\text{casts}(T) = \begin{cases} \langle T' \Rightarrow \{x:T' \mid e\} \rangle^l \circ \text{casts}(T') & \text{if } T = \{x:T' \mid e\} \\ \lambda x:T. x & \text{otherwise} \end{cases}
\end{array}$$

Fig. 4. Operations on types

which we write \Rightarrow . This rule is necessary to prove preservation in the case where $e_1 e_2 \longrightarrow e_1 e_2'$. Why? The first term is typed at $T_2[e_2/x]$ (by T_APP), but reapplying T_APP types the second term at $T_2[e_2'/x]$. Conveniently, $T_2[e_2/x] \Rightarrow T_2[e_2'/x]$, so the two are convertible if we take parallel reduction as our type conversion. Naturally, we have to take the transitive closure so we can string together conversion derivations. We take the symmetric closure, since it is easier for us to work with an equivalence. In previous work, subtyping is used instead of the \equiv relation; one of our contributions is the insight that subtyping—with its accompanying metatheoretical complications—is not an essential component of manifest calculi.

We define type compatibility and a few metatheoretically useful operators in Figure 4. We define $\text{unref}(T)$ as T without any outer refinements (though refinements on, e.g., the domain of a function would be unaffected); we write $\text{unref}_n(T)$ when we remove only the n outermost refinements. We define $\text{casts}(T)$ as the composition of casts necessary to cast from $\text{unref}(T)$ to T .

Lemma 1 (Canonical forms). *If $\emptyset \vdash v : T$, then:*

1. *If $\text{unref}(T) = B$ then $v = k \in \mathcal{K}_B$*
2. *If $\text{unref}(T) = x:T_1 \rightarrow T_2$ then v is*
 - (a) *$\lambda x:T_1'. e_{12}$ and $T_1' \equiv T_1$, or*
 - (b) *$\langle T_1' \Rightarrow T_2' \rangle^l$ and $T_1' \equiv T_1$ and $T_2' \equiv T_2$*
3. *If $\text{unref}(T) = \forall\alpha. T$ then v is $\Lambda\alpha. v'$*

Theorem 2 (Progress). *If $\emptyset \vdash e : T$, then either $e \longrightarrow e'$ or $e = r$.*

Theorem 3 (Preservation). *If $\emptyset \vdash e : T$ and $e \longrightarrow e'$, then $\emptyset \vdash e' : T$.*

Requiring standard weakening, substitution, and inversion lemmas, the syntactic proof of type soundness is straightforward. It is easy to restrict F_H to a simply typed calculus with a similar type soundness proof.

$$\begin{array}{l}
\text{Closed terms } \boxed{r_1 \sim r_2 : T; \theta; \delta \text{ and } e_1 \simeq e_2 : T; \theta; \delta} \\
k \sim k : B; \theta; \delta \iff k \in \mathcal{K}_B \\
v_1 \sim v_2 : \alpha; \theta; \delta \iff \alpha \mapsto R, T_1, T_2 \in \theta \wedge v_1 R v_2 \\
v_1 \sim v_2 : (x:T_1 \rightarrow T_2); \theta; \delta \iff \forall v'_1 \sim v'_2 : T_1; \theta; \delta. v_1 v'_1 \simeq v_2 v'_2 : T_2; \theta; \delta[v'_1, v'_2/x] \\
v_1 \sim v_2 : \forall \alpha. T; \theta; \delta \iff \forall R, T_1, T_2. v_1 T_1 \simeq v_2 T_2 : T; \theta[\alpha \mapsto R, T_1, T_2]; \delta \\
v_1 \sim v_2 : \{x:T \mid e\}; \theta; \delta \iff v_1 \sim v_2 : T; \theta; \delta \wedge \\
\theta_1(\delta_1(e))[v_1/x] \longrightarrow^* \text{true} \wedge \theta_2(\delta_2(e))[v_2/x] \longrightarrow^* \text{true} \\
\uparrow l \sim \uparrow l : T; \theta; \delta \\
e_1 \simeq e_2 : T; \theta; \delta \iff e_1 \longrightarrow^* r_1 \wedge e_2 \longrightarrow^* r_2 \wedge r_1 \sim r_2 : T; \theta; \delta \\
\text{Types } \boxed{T_1 \simeq T_2 : *; \theta; \delta} \\
B \simeq B : *; \theta; \delta \\
\alpha \simeq \alpha : *; \theta; \delta \\
x:T_{11} \rightarrow T_{12} \simeq x:T_{21} \rightarrow T_{22} : *; \theta; \delta \iff T_{11} \simeq T_{21} : *; \theta; \delta \wedge \\
\forall v_1 \sim v_2 : T_{11}; \theta; \delta. \\
T_{12} \simeq T_{22} : *; \theta; \delta[v_1, v_2/x] \\
\forall \alpha. T_1 \simeq \forall \alpha. T_2 : *; \theta; \delta \iff \forall R, T'_1, T'_2. T_1 \simeq T_2 : *; \theta[\alpha \mapsto R, T'_1, T'_2]; \delta \\
\{x:T_1 \mid e_1\} \simeq \{x:T_2 \mid e_2\} : *; \theta; \delta \iff T_1 \simeq T_2 : *; \theta; \delta \wedge \\
\forall v_1 \sim v_2 : T_1; \theta; \delta. \theta_1(\delta_1(e_1))[v_1/x] \simeq \theta_2(\delta_2(e_2))[v_2/x] : \text{Bool}; \theta; \delta \\
\text{Open terms and types } \boxed{\Gamma \vdash \theta; \delta \text{ and } \Gamma \vdash e_1 \simeq e_2 : T \text{ and } \Gamma \vdash T_1 \simeq T_2 : *} \\
\Gamma \vdash \theta; \delta \iff \forall x:T \in \Gamma. \theta_1(\delta_1(x)) \simeq \theta_2(\delta_2(x)) : T; \theta; \delta \wedge \\
\forall \alpha \in \Gamma. \alpha \mapsto R, T_1, T_2 \in \theta \\
\Gamma \vdash e_1 \simeq e_2 : T \iff \forall \Gamma \vdash \theta; \delta. \theta_1(\delta_1(e_1)) \simeq \theta_2(\delta_2(e_2)) : T; \theta; \delta \\
\Gamma \vdash T_1 \simeq T_2 : * \iff \forall \Gamma \vdash \theta; \delta. T_1 \simeq T_2 : *; \theta; \delta
\end{array}$$

Fig. 5. The logical relation for parametricity

4 Parametricity

We prove relational parametricity for two reasons: (1) it gives us powerful reasoning techniques such as free theorems [17], and (2) it indicates that contracts don't interfere with type abstraction. Our proof is standard: we define a (syntactic) logical relation where each type is interpreted as a relation on terms and the relation at type variables is given as a parameter. In the next section, we will define a subtype relation and show that an upcast—a cast whose source type is a subtype of the target type—is logically related to the identity function. Since logically related terms are contextually equivalent, this property means that upcasts can be eliminated without changing the meaning of a program.

We begin by defining a relation on closed results, $r_1 \sim r_2 : T; \theta; \delta$, as a fixpoint on types; we then extend it to a relation on closed expressions, $e_1 \simeq e_2 : T; \theta; \delta$. The definitions are shown in Figure 5.¹ Both relations have three indices: a type T , a substitution θ for type variables, and a substitution δ for term variables. A type substitution θ , which gives the interpretation of free type variables in

¹ To save space, we write $\uparrow l \sim \uparrow l : T; \theta; \delta$ separately instead of manually adding such a clause for each type.

T , maps a type variable to a triple (R, T_1, T_2) comprising a binary relation R on terms and two closed types T_1 and T_2 . We require that R be closed under parallel reduction (the \Rightarrow relation). A term substitution δ maps from variables to a pair of closed values. We write θ_i ($i = 1, 2$) for a substitution that maps a type variable α to T_i where $\theta(\alpha) = (R, T_1, T_2)$. We denote projections δ_i similarly.

With these definitions out of the way, the term relation is mostly straightforward. First, $\uparrow l$ is related to itself at every type. A base type B gives the identity relation on \mathcal{K}_B , the set of constants of type B . A type variable α simply uses the relation assumed in the substitution θ . Related functions map related arguments to related results. Type abstractions are related when their bodies are parametric in the interpretation of the type variable. Finally, two values are related at a refinement type when they are related at the underlying type and both satisfy the predicate; here, the predicate e gets closed by applying the substitutions. The \sim relation on results is extended to the relation \simeq closed terms in a straightforward manner: terms are related if and only if they both terminate at related results. We extend the relation to open terms, written $\Gamma \vdash e_1 \simeq e_2 : T$, relating open terms that are related when closed by any “ Γ -respecting” pair of substitutions θ and δ (written $\Gamma \vdash \theta; \delta$, also defined in Figure 5).

To show that a (well-typed) cast is logically related to itself, we also need a relation on types $T_1 \simeq T_2 : *; \theta; \delta$; we define this relation in Figure 5. We use the logical relation on terms to handle the arguments of function types and refinement types. Note that T_1 and T_2 are not necessarily closed; terms in refinement types, which should be related at **Bool**, are closed by applying substitutions. In the function/refinement type cases, the relation on a smaller type is universally quantified over logically related values. There are two choices of the type at which they should be related (for example, T_{11} on the second line could be T_{21} in the function type case), but it does not really matter which to choose since they are related types. Here, we have chosen the type from the left-hand side; in our proof, we justify this choice by proving a “type exchange” lemma that lets us replace a type index T_1 in the term relation by T_2 when $T_1 \simeq T_2 : *$. Finally, we lift our type relation to open terms: we write $\Gamma \vdash T_1 \simeq T_2 : *$ when two types are equivalent for any Γ -respecting substitutions.

It is worth discussing a few points peculiar to our formulation. First, we allow any relation on terms closed under parallel reduction to be used in θ ; terms related at T need not be well typed at T . The standard formulation of a logical relation is well typed throughout, requiring that the relation R in every triple be well typed, only relating values of type T_1 to values of type T_2 (e.g., [14]). We have two motivations for leaving our relations untyped. First, functions of type $x:T_1 \rightarrow T_2$ must take related values ($v_1 \sim v_2 : T_1$) to related results...but at what type? While $T_2[v_1/x]$ and $T_2[v_2/x]$ are related in our type logical relation, terms that are well typed at one type won't necessarily be well typed at the other. Second, we prove in Section 5 that upcasts have no effect: if $T_1 <: T_2$, then $\langle T_1 \Rightarrow T_2 \rangle^l \sim \lambda x:T_1. x : T_1 \rightarrow T_2$. That is, we want a cast $\langle T_1 \Rightarrow T_2 \rangle^l$, of type $T_1 \rightarrow T_2$, to be related to the identity $\lambda x:T_1. x$, of type $T_1 \rightarrow T_1$: the cast and the identity won't (in general) have the same type. We therefore don't

demand that two expressions related at T be well typed at T , and we allow *any* relation to be chosen as R , so long as it is closed under parallel reduction. Another peculiarity is in our treatment of substitutions and type indices. Just as the interpretation of free type variables in the logical relation’s type index are kept in a substitution θ , we keep δ as a substitution for the free term variables that can appear in type indices. Keeping this substitution separate avoids a problem in defining the logical relation at function types. Consider a function type $x:T_1 \rightarrow T_2$: our *logical* relation says that values v_1 and v_2 are related at this type when they take related values to related results, i.e. if $v'_1 \sim v'_2 : T_1; \theta; \delta$, then we should be able to find $v_1 v'_1 \simeq v_2 v'_2$. The question here is which type index we should use. If we keep our type indices closed (with respect to term variables), we cannot use T_2 on its own—we have to choose a binding for x ! Knowles and Flanagan [10] deal with this problem by introducing the “wedge product” operator, which merges two types—one with v'_1 substituted for x and the other with v'_2 for x —into one. Instead of substituting eagerly, we put both bindings in δ and apply them when needed—the refinement type case. We think our formulation is more uniform with regard to free term/type variables, since eager substitution is a non-starter for type variables, anyway.

As we developed our proof, we found that the E_REFL rule

$$\langle T \Rightarrow T \rangle^l v \rightsquigarrow v$$

is not just a convenient way to skip decomposition of a trivial cast into smaller trivial casts (when T is a polymorphic or dependent function type); E_REFL is, in fact, crucial to obtaining parametricity in our syntactic setting. For example, by parametricity, we expect every value of type $\forall \alpha. \alpha$ to behave the same as the polymorphic identity function. One of the values of this type is $\lambda \alpha. \langle \alpha \Rightarrow \alpha \rangle^l$. Without E_REFL, however, applying this type abstraction to a compound type, say $\text{Bool} \rightarrow \text{Bool}$, and a function f of type $\text{Bool} \rightarrow \text{Bool}$ would return, by E_FUN, a value that is syntactically different from f , breaking parametricity!² With E_REFL, $\langle T \Rightarrow T \rangle^l$ returns the input immediately, regardless of T , just as the identity function. So, this rule is a technical necessity, ensuring that casts containing type variables behave parametrically. (Naturally, the evaluation of well-typed programs never encounters casts with uninstantiated type variables.)

We have relational parametricity—every well-typed term (under T) is related to itself for any T -respecting substitutions.

Theorem 4 (Parametricity).

1. If $\Gamma \vdash e : T$ then $\Gamma \vdash e \simeq e : T$, and
2. If $\Gamma \vdash T$ then $\Gamma \vdash T \simeq T : *$.

The proof is mostly standard, although—like the proof of semantic type soundness in Greenberg, Pierce, and Weirich [7]—it requires a separate reflexivity

² Intuitively, we expect the returned value should behave the same as the input, though. Moreover, the subtyping we define is reflexive, so the upcast lemma we prove applies, as well—though, of course, we used E_REFL to prove it!

$$\boxed{\Gamma \vdash T_1 <: T_2}$$

$$\frac{}{\Gamma \vdash B <: B} \text{S_BASE} \quad \frac{}{\Gamma \vdash \alpha <: \alpha} \text{S_TVAR} \quad \frac{\Gamma, \alpha \vdash T_1 <: T_2}{\Gamma \vdash \forall \alpha. T_1 <: \forall \alpha. T_2} \text{S_FORALL}$$

$$\frac{\Gamma \vdash T_{21} <: T_{11} \quad \Gamma, x:T_{21} \vdash T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] <: T_{22}}{\Gamma \vdash x:T_{11} \rightarrow T_{12} <: x:T_{21} \rightarrow T_{22}} \text{S_FUN}$$

$$\frac{\Gamma \vdash \text{unref}(T_1) <: \text{unref}(T_2) \quad \Gamma, x:\text{unref}(T_1) \vdash \text{casts}(T_1) x \supset \text{casts}(T_2) (\langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l x)}{\Gamma \vdash T_1 <: T_2} \text{S_REFINE}$$

$$\boxed{\Gamma \vdash e_1 \supset e_2}$$

$$\frac{\forall \Gamma \vdash \theta; \delta. (\exists v. \theta_1(\delta_1(e_1)) \longrightarrow^* v) \text{ implies } (\exists v. \theta_1(\delta_1(e_2)) \longrightarrow^* v)}{\Gamma \vdash e_1 \supset e_2} \text{IMP}$$

Fig. 6. Subtyping, implication, and closing substitutions

lemma for casts, as mentioned above. We make one small disclaimer: we have not completed the standard but tedious proof showing that parallel reduction implies cotermination at similar values, i.e., if $e_1 \Rightarrow e_2$ and $e_1 \longrightarrow^* r_1$, then $e_2 \longrightarrow^* r_2$ such that $r_1 \Rightarrow r_2$, and vice versa. We expect that our existing Coq proof of this fact for a similar operational semantics (from [7]) will adapt readily. Note that our proof of type soundness in Section 3 relies on much simpler properties of parallel reduction, which we *have* proved.

5 Subtyping and Upcast Elimination

Knowles and Flanagan [10] define a subtyping relation for their manifest calculus, λ_{H} , as a primitive notion of the system. Furthermore, they prove that upcast elimination is sound: if $T_1 <: T_2$, then $\langle T_1 \Rightarrow T_2 \rangle^l$ is equivalent to the identity function. This is, at heart, an optimization: since the cast can never fail, there's no point in running it. In this section, we define a subtyping relation for F_{H} and prove that upcast elimination is sound. To be clear, the type system of F_{H} doesn't have subtyping or a subsumption rule at all; we simply show that upcasts are logically related—and therefore contextually equivalent—to the identity.

We define subtyping in Figure 6. Our subtyping rules are similar to those in λ_{H} . The first three rules are standard. The rule for dependent function types is mostly usual: contravariant on argument types and covariant on return types. Here, we need to be careful about the type of x . Return types T_{12} and T_{22} should be compared under the assumption that x has T_{21} , which is a subtype of the other argument type T_{11} [4]. However, x in T_{12} has a different type, i.e.,

T_{11} , so we need to insert a cast to keep the subtyping relation well typed— F_H doesn't have subsumption!

Our rule for subtyping of refinements differs substantially from λ_H 's, mostly because F_H allows refinements of arbitrary types, while λ_H only refines base types. The `S_REFINE` rule essentially says T_1 is a subtype of T_2 if (1) T_1 without the (outermost) refinements is a subtype of T_2 without the (outermost) refinements, and (2) for any v of type $\text{unref}(T_1)$, if $\text{casts}(T_1) v$ reduces to a value, so does $\text{casts}(T_2) \langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l v$, for any l . The intuition behind the second condition is that, for T_1 to be a subtype of T_2 , the predicates in T_1 (combined by conjunction) should be stronger than those in T_2 . Recall that $\text{casts}(T)$ is defined in Figure 4 as the composition of casts necessary to cast from $\text{unref}(T)$ to T . So, if application of $\text{casts}(T)$ to a value of $\text{unref}(T)$ does not raise blame, then the value can be typed at T by repeated use of `T_EXACT`.

If the implication in `S_REFINE` holds for a value v of type $\text{unref}(T_1)$, then either: (1) v did not pass the checks in $\text{casts}(T_1)$, so this value is not in T_1 ; or (2) v passed the checks in $\text{casts}(T_1)$ and $\langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l v$ passed all of the checks in $\text{casts}(T_2)$. So, if (1) or (2) hold for all values of type $\text{unref}(T_1)$, then it means that all values of type T_1 can be safely treated as if they had type T_2 , i.e., T_1 a subtype of T_2 .

Finally, we need a source of closing substitutions to compare the evaluation of the two casts. We use the closing substitutions from the logical relation at T as the source of “values of type T ”. (Arbitrarily, we take the values and types from the left.) There is a similar but more dire situation in the manifest calculi of Knowles and Flanagan [10] and Greenberg, Pierce, and Weirich [7]. They both define a denotational semantics for use in their refinement subtyping rule—but they *need* to do so, in order to avoid a circularity. We have no such issues, and make the decision because it is expedient.

We formulate our implication judgment in terms of cotermination at values rather than cotermination at `true` (as in [7, 10]) because we have to contend with multiple layers of refinement in types—using cotermination at values reduces the amount of predicate bookkeeping we have to do.

Having defined subtyping, we are able to show that upcast elimination is sound.

Lemma 5 (Upcast lemma). *If $\Gamma \vdash T_1 <: T_2$ and $\Gamma \vdash T_1$ and $\Gamma \vdash T_2$, then $\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle^l \simeq \lambda x:T_1. x : T_1 \rightarrow T_2$.*

6 Related Work

We discuss the related work in two parts. We first distinguish our work from the untyped contract systems that enforce parametric polymorphism *dynamically*, rather than statically as F_H does; we then discuss how F_H differs from existing manifest contract calculi in greater detail.

Dynamically checked polymorphism

The F_H type system enforces parametricity with type abstractions and type variables, while refinements are dynamically checked. Another line of work omits refinements, seeking instead to dynamically enforce parametricity—typically with some form of sealing (à la Pierce and Sumii [13]).

Guha et al. [9] define contracts with polymorphic signatures, maintaining abstraction with sealed “coffers”; they do not prove parametricity. Matthews and Ahmed [11] prove parametricity for a polymorphic multi-language system with a similar policy. Ahmed et al. [3] prove parametricity for a gradual typing [15] calculus which enforces polymorphism with a set of global runtime seals. Strickland et al. add support for dynamically checked variable-arity polymorphism to Typed Racket [16]. Ahmed et al. [2] define a polymorphic calculus for gradual typing, using local syntactic “barriers” instead of global seals. We believe that it is possible to combine F_H with the barrier calculus of Ahmed et al., but we leave it to future work.

Manifest systems

Wadler and Findler [18] gave a simple syntactic account of a calculus combining refinement types and gradual types [15]; they, like us, define subtyping *post facto*, proving theorems similar to the upcast lemma. They do not, however, support dependent function types. Gronski and Flanagan [8] compares non-dependent latent and manifest contract calculi.

Four existing manifest calculi have dependent function types (such as [6, 7, 10, 12]) use subtyping and theorem provers as part of the definition of the type system. All four of these calculi have complicated metatheory. Ou et al. [12] restrict refinements and arguments of dependent functions to a conservative approximation of pure terms; they also place strong requirements on their prover. Knowles and Flanagan [10] as well as Greenberg, Pierce, and Weirich [7] use denotational semantics to give a firm foundation to Flanagan’s earlier work [6]. We consider three systems in more detail: Knowles and Flanagan’s λ_H (KF) [10]; Greenberg, Pierce, and Weirich’s λ_H (GPW) [7]; and F_H . The rest of this subsection addresses the differences between KF, GPW, and F_H .

In Section 1, we discussed in general terms some of the complexity that KF and GPW encountered. What made KF and GPW so complicated? Both systems share the same two impediments in the preservation proof: preservation after active checks and after congruence steps in the argument position of applications. KF and GPW resolve both of these with subtyping, using a rule like the following for refinements:³

$$\frac{\forall \Gamma, x:\{x:B \mid \text{true}\} \vdash \sigma. \sigma(e_1) \longrightarrow^* \text{true} \text{ implies } \sigma(e_2) \longrightarrow^* \text{true}}{\Gamma \vdash \{x:B \mid e_1\} <: \{x:B \mid e_2\}}$$

³ Readers familiar with the systems will recognize that we’ve folded the implication judgment into the relevant subtyping rule.

Subtyping and the requirement that constants be assigned most specific types, —i.e., if $e[k/x] \longrightarrow^* \text{true}$ for $k \in \mathcal{K}_B$ then $\emptyset \vdash \text{ty}(k) <: \{x:B \mid e\}$ —are used to show preservation of active checks. The two systems use subtyping to relate substituted types in different ways. KF use full beta reduction, showing that subtyping is closed under reduction. GPW use call-by-value reduction, showing that subtyping is closed under parallel reduction. Once these two difficulties are resolved, both preservation proofs are standard, given appropriate subtyping inversion lemmas.

So much for subtyping. Why do KF and GPW need denotational semantics? Spelled out pedantically, the subtyping rule above has the following premise:

$$\forall \sigma. \Gamma, x:\{x:B \mid \text{true}\} \vdash \sigma \text{ implies } (\sigma(e_1) \longrightarrow^* \text{true} \text{ implies } \sigma(e_2) \longrightarrow^* \text{true})$$

That is, the well formedness of the closing substitution σ is in a negative position. Where do closing substitutions come from? We cannot use the typing judgment itself, as this would be ill-defined: term typing requires subtyping via subsumption; subtyping requires closing substitutions in a negative position via the refinement case; but closing substitutions require typing. We need another source of values: hence, denotational semantics. Both KF and GPW define syntactic term models of types to use as a source of values for closing substitutions, though the specifics differ.

After adding subtyping and denotational semantics, both KF and GPW are well defined and have syntactic proofs of type soundness. But in the process of proving syntactic type soundness, both languages proved semantic soundness theorems:

$$\Gamma \vdash e : T \text{ implies } \forall \Gamma \vdash \sigma. \sigma(e) \in \llbracket \sigma(T) \rrbracket$$

This theorem suffices for soundness of the language... so why bother with a syntactic proof? In light of this, GPW only proves semantic soundness. The situation in KF and GPW is unsatisfying: the syntactic proof of type soundness motivated subtyping, which motivated denotational semantics, which obviated the need for syntactic proof. Beyond this, the proofs are hard to scale: adding in polymorphism or state is a non-trivial task, since we must—before defining the type system!—construct an appropriate denotational semantics, which itself depends on our evaluation relation.

F_H solves the problem by avoiding subtyping—which is what forced the presence of closing substitutions and denotational semantics in the first place. The first issue in preservation—that of preserving refinement types after checks have finished—was resolved in KF and GPW with subtyping. We instead resolve it with a runtime rule that allows us to type values with any refinement they satisfy:

$$\frac{\vdash \Gamma \quad \emptyset \vdash v : T \quad \emptyset \vdash \{x:T \mid e\} \quad e[v/x] \longrightarrow^* \text{true}}{\Gamma \vdash v : \{x:T \mid e\}} \quad \text{T_EXACT}$$

This eliminates one use of subtyping and also the “most-specific type” restriction. If we “bit the bullet” and allowed non-empty contexts in T_EXACT, then we

would need to apply a closing substitution to $e[v/x]$ before checking if it reduces to true. But the circularity in subtyping alluded to in Section 1 was caused by closing substitutions—we must avoid them! The second issue in preservation—that of conversion between $T_2[e_2/x]$ and $T_2[e'_2/x]$ —can be resolved in a similar fashion. We define another runtime rule that allows us to convert types:

$$\frac{\vdash \Gamma \quad \emptyset \vdash e : T \quad \emptyset \vdash T' \quad T \equiv T'}{\Gamma \vdash e : T'} \quad \text{T_CONV}$$

The conversion we use, \equiv , is defined as the symmetric, transitive closure of CBV-respecting parallel reduction. This is only as much equivalence as we need: if $e_2 \longrightarrow e'_2$, then $T_2[e_2/x] \equiv T_2[e'_2/x]$. These two rules suffice to keep subtyping out of F_H , which in turn avoids denotational semantics.

7 Future Work

We presented a simpler approach to manifest contract calculi, which we applied to construct F_H , a parametrically polymorphic manifest contract calculus. Two other standard extensions are natural next steps: general recursion and state.

Acknowledgments

Stephanie Weirich provided many insights throughout. A conversation about parametricity with Amal Ahmed and Stephanie Weirich was particularly illuminating. This work was supported in part by the National Science Foundation under grant 0915671, *Contracts for Precise Types*, in part by the JSPS Grant-in-Aid for Young Scientists (A) No. 21680002, and in part by the Portuguese Foundation for Science and Technology, POPH - QREN, under grant SFRH / BPD / 46065 / 2008.

References

1. PLT Racket Contracts, <http://pre.plt-scheme.org/docs/html/guide/contracts.html>
2. Ahmed, A., Findler, R.B., Siek, J., Wadler, P.: Blame for all. In: Principles of Programming Languages (POPL) (2011)
3. Ahmed, A., Findler, R.B., Siek, J.G., Wadler, P.: Blame for all. In: Workshop on Script-to-Program Evolution (STOP) (2009)
4. Aspinall, D., Compagnoni, A.: Subtyping dependent types. Theor. Comput. Sci. 266(1-2), 273–309 (Sep 2001)
5. Findler, R.B., Felleisen, M.: Contracts for higher-order functions. In: International Conference on Functional Programming (ICFP). pp. 48–59 (2002)
6. Flanagan, C.: Hybrid type checking. In: POPL. pp. 245–256 (2006)
7. Greenberg, M., Pierce, B.C., Weirich, S.: Contracts made manifest. In: Principles of Programming Languages (POPL) 2010 (2010)
8. Gronski, J., Flanagan, C.: Unifying hybrid types and contracts. In: Trends in Functional Programming (TFP) (2007)

9. Guha, A., Matthews, J., Findler, R.B., Krishnamurthi, S.: Relationally-parametric polymorphic contracts. In: DLS. pp. 29–40 (2007)
10. Knowles, K., Flanagan, C.: Hybrid type checking (2010), to appear in TOPLAS.
11. Matthews, J., Ahmed, A.: Parametric polymorphism through run-time sealing or, theorems for low, low prices! In: European Symposium on Programming. pp. 16–31. Springer-Verlag, Berlin, Heidelberg (2008)
12. Ou, X., Tan, G., Mandelbaum, Y., Walker, D.: Dynamic typing with dependent types. In: IFIP TCS. pp. 437–450 (2004)
13. Pierce, B., Sumii, E.: Relating cryptography and polymorphism (Jul 2000)
14. Pitts, A.M.: Typed operational reasoning. In: Pierce, B.C. (ed.) *Advanced Topics in Types and Programming Languages*, chap. 7, pp. 245–289. The MIT Press (2005)
15. Siek, J.G., Taha, W.: Gradual typing for functional languages. In: *Scheme and Functional Programming Workshop* (September 2006)
16. Strickland, T.S., Tobin-Hochstadt, S., Felleisen, M.: Practical variable-arity polymorphism. In: *ESOP '09: Proceedings of the 18th European Symposium on Programming Languages and Systems*. pp. 32–46. Springer-Verlag, Berlin, Heidelberg (2009)
17. Wadler, P.: Theorems for free! In: *Proceedings of ACM Conference on Functional Programming and Computer Architecture (FPCA'89)*. pp. 347–359. London, UK (Sep 1989)
18. Wadler, P., Findler, R.B.: Well-typed programs can't be blamed. In: *European Symposium on Programming (ESOP)*. pp. 1–16 (2009)
19. Wright, A.K., Felleisen, M.: A syntactic approach to type soundness. *Information and Computation* 115, 38–94 (1992)