

MANIFEST CONTRACTS

Michael Greenberg
A DISSERTATION

in

Computer and Information Science

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy

2013

Supervisor of Dissertation

Benjamin C. Pierce
Henry Salvatori Professor of CIS

Graduate Group Chairperson

Val Tannen
Professor of CIS

Dissertation Committee

Stephanie Weirich (Associate Professor of CIS; Committee Chair)

Rajeev Alur (Zisman Family Professor of CIS)

Greg Morrisett (Allen B. Cutting Professor of CS at Harvard)

Steve Zdancewic (Associate Professor of CIS)

Acknowledgments

I know you have faith in me; let me take it with me on the road
before me and I will be able to go on forever.

“Tang Qiaodi”

The Gourmet and Other Stories of Modern China

Lu Wenfu

Let me begin by thanking Benjamin Pierce for being a formative collaborator and ally—for his teaching, advice, guidance, and forbearance.

Many people—members of the Penn PL Club and otherwise—have provided helpful and insightful comments along the way: Brian Aydemir and João Belo; Amal Ahmed, Cătălin Hrițcu, Taro Sekiyama, Jianzhou Zhao; Ron Garcia, Fritz Henglein, Jeremy Siek, Phil Wadler, and Brent Yorgey. My fellow travelers Chris Casinghino and Adam Aviv have been welcome company. Nate Foster has, on several occasions, helped me remember to stay on the wagon. I hope I have been equally supportive of Arjun Ravi Narayan and Scott Kilpatrick.

Atsushi Igarashi and Stephanie Weirich have been a delight to work with as co-authors, and I have learned a great deal from both. Atsushi and his student, Taro Sekiyama, helped quash some serious bugs while being a pleasure to work with.

I am grateful for the attention, advice, and patience of my committee—Stephanie Weirich (again), Steve Zdancewic, Rajeev Alur, and Greg Morrisett. Their efforts have seriously improved the quality of my thesis.

I am interested in programming languages because of Dan Friedman. Dan insisted that I let loose and have fun in college, and perhaps I’ve taken his advice too much to heart, but he also insisted that I meet with Shriram Krishnamurthi, who formed so much of my early research taste—and helped me find Benjamin in turn.

My friends here in Philadelphia and beyond form a broad family who, along with my actual family, soothed and encouraged me. B2 was effectively my office, and its *baristi* kindly put up with me.

Finally, Hannah de Keijzer literally and figuratively brings music into my life every day, and there is no way for me to do that justice—not that I won’t try.

ABSTRACT

MANIFEST CONTRACTS

Michael Greenberg

Benjamin Pierce

Eiffel [47] popularized *design by contract*, a software design philosophy where programmers specify the requirements and guarantees of functions via executable pre- and post-conditions written in code. Findler and Felleisen [26] brought contracts to higher-order programming, inspiring the PLT Racket implementation of contracts [56]. Existing approaches for runtime checking lack reasoning principles and stop short of their full potential—most Racket contracts check only simple types. Moreover, the standard algorithm for higher-order contract checking can lead to unbounded space consumption and can destroy tail recursion. In this dissertation, I develop so-called *manifest* contract systems which integrate more coherently in the type system, and relate them to Findler-and-Felleisen-style *latent* contracts. I extend a manifest system with type abstraction and relational parametricity, and also show how to integrate dynamic types and contracts in a *space efficient* way, i.e., in a way that doesn't destroy tail recursion. I put manifest contracts on a firm type-theoretic footing, showing that they support extensions necessary for real programming. Developing these principles is the first step in designing and implementing higher-order languages with contracts and refinement types.

My work has been supported by the National Science Foundation under grants 0534592 *Linguistic Foundations for XML View Update*, and 0915671 *Contracts for Precise Types*.

Some of the material in Chapter 4 was supported in part by the DARPA CRASH program through the United States Air Force Research Laboratory (AFRL) under Contract No. FA8650-10-C-7090. The views expressed are my own and do not reflect the official policy or position of the Department of Defense or the U.S. Government—and the more daylight between us, the better.

Contents

1	Introduction	1
1.1	Latent and manifest contracts	5
1.2	Contracts and abstraction	7
1.3	Efficiency and contract checking	8
1.4	Summary	10
1.5	Notation and other conventions	10
2	Contracts made manifest	11
2.1	The nondependent languages	14
2.1.1	The language λ_C	14
2.1.2	The language λ_H	17
2.2	The nondependent translations	21
2.3	The dependent languages	23
2.3.1	Dependent λ_C	23
2.3.2	Dependent λ_H	27
2.4	The translations	37
2.4.1	Translating λ_C to λ_H : ϕ	38
2.4.2	Translating λ_H to λ_C : ψ	39
2.5	Exact translations	40
2.5.1	Translating picky λ_C to λ_H	41
2.5.2	Translating λ_H to lax λ_C	44
2.6	Inexact translations	50
2.6.1	Translating lax λ_C to λ_H	50
2.6.2	Translating λ_H to picky λ_C	56
2.6.3	Alternative calculi	61
2.7	Conclusion	61
3	Polymorphic manifest contracts	63
3.1	Examples	65
3.2	Defining F_H	72
3.3	Parametricity	96
3.4	Subtyping and Upcast Elimination	113
3.5	Type conversion: parallel reduction vs. common subexpression reduction	123

3.6	Conclusion	127
4	Space-efficient manifest contracts	129
4.1	Design philosophy	134
4.2	A cast calculus	136
4.2.1	Syntax and typing	136
4.2.2	Operational semantics	139
4.2.3	Proofs	143
4.3	A naïve coercion calculus	148
4.3.1	Syntax and typing	149
4.3.2	Operational semantics	153
4.3.3	Proofs	157
4.4	Soundness of NAIVE with regard to CAST	167
4.5	A space-efficient coercion calculus	178
4.5.1	Space-efficient coercions	180
4.5.2	Operational semantics	196
4.5.3	Proofs	196
4.6	Soundness of EFFICIENT with regard to NAIVE	200
4.7	Space efficiency	215
4.8	Conclusion	217
5	Related work	218
5.1	Contracts: a survey	218
5.1.1	Refinement types and contracts	220
5.1.2	Situating λ_H	221
5.2	F_H : polymorphism and manifest metatheory	222
5.2.1	Dynamically checked polymorphism	222
5.2.2	F_H and other manifest calculi	223
5.3	Space efficiency and gradual types	225
5.3.1	Space efficiency, gradual typing, and refinement types	225
5.3.2	Coercions	229
6	Conclusion and future work	230
6.1	Future work	232
6.1.1	State and effects	232
6.1.2	Datatypes	232
6.1.3	Coercion insertion	233
6.1.4	Theoretical curiosities	233
6.1.5	Extensions for EFFICIENT	234
	Bibliography	237

List of Tables

4.1	CAST cast reductions, by type	142
4.2	Canonical coercions	188
5.1	Comparison between contract systems	219
5.2	Comparison between gradual typing systems	226

List of Figures

1.1	A queue module in Racket	3
1.2	Contract checking destroys tail recursion	8
2.1	The axis of blame	12
2.2	Base types and constants for λ_C and λ_H	14
2.3	Syntax and semantics for λ_C	15
2.4	Syntax and semantics for λ_H	18
2.5	Typing rules for λ_H	19
2.6	Syntax and semantics for dependent λ_C	24
2.7	Typing rules for dependent λ_C	25
2.8	Syntax and operational semantics for dependent λ_H	27
2.9	Typing rules for dependent λ_H	29
2.10	Type and kind semantics for dependent λ_H	30
2.11	Parallel reduction for dependent λ_H	35
2.12	The translation ϕ from dependent λ_C to dependent λ_H	38
2.13	ψ mapping dependent λ_H to dependent λ_C	40
2.14	A blame-exact result/term correspondence	40
2.15	Blame-exact correspondence for ϕ from picky λ_C	41
2.16	Blame-exact correspondence for ψ into lax λ_C	45
2.17	Blame-inexact correspondence for ϕ from lax λ_C	51
2.18	Blame-inexact correspondence for ψ into picky λ_C	57
3.1	Syntax for F_H	72
3.2	Operational semantics for F_H	74
3.3	Typing rules for F_H	78
3.4	Type compatibility and conversion for F_H	80
3.5	The logical relation for parametricity	97
3.6	Complexity of casts	102
3.7	Subtyping, implication, and closing substitutions	114
3.8	Parallel reduction	124
3.9	Type conversion via common subexpression reduction	125
3.10	Counterexamples to substitutivity of parallel reduction in F_H	126

4.1	The dyn/refine spectrum of cast expressiveness	130
4.2	Space-inefficient reduction	132
4.3	CAST syntax	136
4.4	Typing for CAST, part 1	137
4.5	Typing for CAST, part 2	138
4.6	CAST operational semantics (core rules)	140
4.7	CAST operational semantics (cast rules)	141
4.8	NAIVE syntax	148
4.9	Typing for NAIVE	150
4.10	Coercion typing	151
4.11	Primitive coercions	152
4.12	NAIVE operational semantics (core rules)	154
4.13	NAIVE operational semantics (coercion rules)	155
4.14	NAIVE reduction	158
4.15	Translating from CAST to NAIVE	168
4.16	Relating CAST and NAIVE	169
4.17	Relating casts and NAIVE coercions	172
4.18	Updated syntax for EFFICIENT	179
4.19	Updated typing rules for EFFICIENT	179
4.20	Coercion rewriting rules	181
4.21	Merging coercions	186
4.22	EFFICIENT operational semantics	195
4.23	Space-efficient reduction	196
4.24	Relating NAIVE and EFFICIENT	201
4.25	Canonicalizing NAIVE terms	203
4.26	Relating NAIVE coercions to canonical EFFICIENT coercions	207

The technical development in this dissertation comprises three papers: Greenberg et al. [35] (itself an extended version of Greenberg et al. [34]), an extended and corrected version of Belo et al. [8], and a rejected POPL 2013 submission, Greenberg [33].

This document is typeset in \LaTeX using Computer Modern. There would be even more errors had I not used Sewell and Zappa Nardelli's OTT tool [62]. I also used Aydemir and Weirich's LGen tool [7] for the Coq development of parallel reduction in Chapter 2.

Chapter 1

Introduction

Master, I've filled my contract, wrought in Thy many lands;
Not by my sins wilt Thou judge me, but by the work of my hands.

The Song of the Wage-slave
Robert W. Service

Program specifications allow programmers to check their work. Data structure invariants for, e.g., red-black trees, help programmers write correct implementations; input/output specifications, also called pre- and post-conditions, offer a finer grained checking, at the level of individual functions, e.g., is this function an ϵ -approximation of the square-root function? Temporal logics allow programmers to specify the behavior of long-running and non-terminating programs, e.g., does every request packet sent to a server get a response?

It is folklore that the mere act of writing down a specification is enough to help with program correctness. Computer assisted checking adds a stronger guarantee: sound specification checkers can, in an assisted or even entirely automated way, prove that programs meet their specifications. Naturally, there are trade-offs: sound and complete checkers (of nontrivial properties) can't exist because of the halting problem. As a result, program specification languages exist on a spectrum, where less precise specifications are more robustly checkable—and, therefore, more often used.

Types are perhaps the most common form of program specification, even though their use isn't universal. Types occupy a sweet spot: they are easy for programmers to understand, allow for compositional reasoning, and are machine checkable, if not inferrable. Types also have a beautiful mathematical theory, and they are compelling to study in their own right.

Assertions and pre/post-conditions are also common forms of specification during the development cycle; they are very common when debugging, but perhaps less so in deployed code. Written in code, typically as part of the program itself, assertions and pre- and post-conditions allow for direct feedback at runtime. For example, the `pop` operation of a stack might have a pre-condition checking that the stack

argument is non-empty. When the program calls `pop` with an empty stack, an informative error message can be printed to the screen—or the debugger can kick in. Eiffel [47] makes pre- and post-conditions an essential part of program development in a method called *design by contract*. They use software contracts—pre- and post-conditions—extensively, stating precise properties as concrete predicates written in the same language as the rest of the program.

Eiffel’s contracts are *first order*, designed for an imperative, object-oriented language. Findler and Felleisen [26] introduced “higher-order contracts” for functional languages. These can take one of two forms: predicate contracts (also known as *refinement types*) like $\{x:\text{Int} \mid x > 0\}$, which denotes the positive numbers; and function contracts like $(x:\text{Int}) \rightarrow \{y:\text{Int} \mid y > x\}$, which denotes functions over the integers that return numbers larger than their inputs. Other examples might require that divisors are non-zero:

$$\text{div} : \text{Real} \rightarrow \{m:\text{Real} \mid m \neq 0\} \rightarrow \text{Real}$$

or that the square root function produces answers that are correct (up to some ϵ):

$$\text{sqrt} : (x:\text{Real}) \rightarrow \{y:\text{Real} \mid \text{abs}(x - y^2) < \epsilon\}$$

This last is a *very* strong contract; it entirely specifies the input/output behavior of the `sqrt` function. We could have given a similarly strong contract for division:

$$\text{div} : (n:\text{Real}) \rightarrow \{m:\text{Real} \mid m \neq 0\} \rightarrow \{k:\text{Real} \mid m \cdot k = n\}$$

One feature of contracts is that within the framework of contracts, it’s possible to express a wide range of specifications at varying levels of precision.

In this dissertation, I **design higher-order languages combining conventional type systems, like the simply typed lambda calculus and System F, with contracts**. Designing type systems supporting contracts bolsters the power of abstractions defined by types. Before addressing my own work on the topic, I want to discuss the state of the art—PLT Racket’s contract system [55, 56]—and its shortcomings.

PLT Racket contracts are overwhelmingly just simple types. Racket’s large codebase has many contracts, but 82% of the contracts in their standard library merely recover simple types.¹ Findler and Felleisen developed their contract system to add a form of specification to PLT Racket. Racket is a Scheme [69], written in S-expression syntax and *dynamically typed*. (They have gone on to develop slightly more traditional type systems, Typed Racket [74] in particular.) **In using contracts for simple types, Racket contracts fall short of their full potential.**

As an anecdote to bolster and clarify the data I collected, consider a typical Racket file, shown in Figure 1.1. The file, adapted from the Racket standard library, defines a `queue` data structure, and then exports (in Racket parlance, “provides”) several

¹My own count, as of January 26th, 2011.

```

#lang racket/base

(struct queue (head tail) #:mutable)
(struct link (value [tail #:mutable]))

(define (make-queue) (queue #f #f))
(define (queue-empty? q) (not (queue-head q)))
(define (nonempty-queue? v) (and (queue? v) (queue-head v) #t))

(define (enqueue! q v)
  (unless (queue? q) (raise-type-error enqueue! "queue" 0 q))
  (let ([new (link v #f)])
    (if (queue-head q)
        (set-link-tail! (queue-tail q) new)
        (set-queue-head! q new))
    (set-queue-tail! q new)))

(define (dequeue! q)
  (unless (queue? q) (raise-type-error dequeue! "queue" 0 q))
  (let ([old (queue-head q)])
    (unless old (error 'dequeue! "empty queue"))
    (set-queue-head! q (link-tail old))
    (link-value old)))

(define queue/c (flat-named-contract "queue" queue?))
(define nonempty-queue/c (flat-named-contract "nonempty-queue" nonempty-queue?))

(provide/contract
 [queue/c flat-contract?]
 [nonempty-queue/c flat-contract?]
 [queue? (-> any/c boolean?)]
 [make-queue (-> queue/c)]
 [queue-empty? (-> queue/c boolean?)])
(provide enqueue! dequeue!)

```

Taken from `collects/data/queue.rkt`, as of November 6th 2011. Some comments have been removed.

Figure 1.1: A queue module in Racket

things. This is where contracts are typically placed in PLT Racket definitions, for reasons described below in Section 1.3. First, there are `queue/c` and `nonempty-queue/c`, which are “flat” contracts on values, i.e., predicate contracts `a/k/a` refinement types. Then the predicate `queue?` tests whether a given value is a queue: it has a function contract, written in prefix form, that says that `queue?` accepts any value (`any/c`) and returns a boolean (the predicate `boolean?`, which is automatically treated as a contract). The constructor `make-queue` is a function of zero arguments that returns a value that satisfies the contract `queue/c`. The predicate `queue-empty?` checks whether or not the queue is empty: its sole argument must be a queue (`queue/c`) and it returns a boolean. Notice that this interface is *dependent*: we provide `queue/c` and then use it in the contracts for `make-queue` and `queue-empty?`.

The queue module also provides two non-contracted identifiers, `enqueue!` and `dequeue!`.² Looking at their definitions, `enqueue!` expects a queue and returns whatever the structure mutator `set-queue-tail!` returns—a unit value. So a good contract for `enqueue!` would be `(-> queue/c any/c)`. Similarly, `dequeue!` takes in a queue and returns its head. What’s more, `dequeue!` expects a *non-empty* queue—if given an empty queue, it will raise an error. So a good contract for `dequeue!` would be `(-> nonempty-queue/c any/c)`. The return contract of `any/c` makes sense, since we don’t know anything about values in the queue.

As written, **the contracts on the queue interface don’t enforce anything more than a simple type system would**, though they define a refined type (`nonempty-queue/c`). I would reiterate that this isn’t cherry picking examples: four out of five Racket contracts enforce nothing more than simple types.

PLT Racket’s use of contracts to recover simple types doesn’t teach us much for the design of new languages: **it’s already possible to check simple types completely statically**. While Racket’s solution helps with their pre-existing code, why would a new language would follow suit? Why bother paying the runtime cost of contract checking at runtime when static checking is available?

Dynamically typed languages create a “contracts for simple types bias”. Writing specifications can be hard work—as hard as programming. It’s tempting to stop at “good enough”—which in a dynamically typed language, may mean only simple types. Moreover, dynamic typing and contracts leaves us without empty handed. What are the reasoning principles for contracts? How does contract checking affect execution? Are there limits on how much space and time contract checking can take?

In this dissertation, I explore **what contracts look like when we treat contracts as our type system, taking simple types for granted**. Suppose we start with a simply typed lambda calculus. Can we coherently integrate contracts into a type system? In terms of the PLT Racket example above, I’d like to treat `queue/c` as a type (just as it would be in ML)—and I’d also like to treat `nonempty-queue/c` as a type. Treating contracts as types has many benefits. First, it gets rid of the low bar for contracts in PLT Racket: if the programmer wants simple types, they

²The `!` are a Racket naming convention for functions that mutate their inputs.

have them as a static check. Next, having contracts as part of type signatures allows for operations on built-in abstract types to specify their abstractions totally; `dequeue!` is a partial operation, and the type `queue/c → unit` doesn't adequately capture that partiality; the type `nonempty-queue/c → unit` does. A similar case applies to built-in operations, e.g., the division operator. Finally, types and contracts combine to yield strong reasoning principles, giving better guarantees to clients and implementors of, e.g., the queue abstraction.

Two issues stand in the way of integrating contracts in types. What should the semantics of a type system with “manifest” contracts be? Are such semantics practical from an efficiency perspective?

The first issue, determining the semantics, is challenging because contracts contain code. If contracts are to be part of the type system, we are left with a *dependent* type system, one where treating `nonempty-queue/c` as a *type* forces the meaning of types to depend on evaluation. The metatheory of dependent languages is notoriously difficult, and is typically treated in rarified, “pure” settings, e.g., without nontermination or effects. But we are interested in writing specifications for *programs*, where programs may diverge or contracts may be unsatisfied—and therefore must signal errors. Early attempts at hybrid type/contract systems had serious issues with the well-foundedness of their definitions. **If contracts and types are to coexist, the metatheory must scale to impure core calculi.**

The second issue, efficiency, is in fact a challenge for all contract systems. Contract checking can change the asymptotic complexity of programs by accumulating too many checks on functions or on the stack. **If contracts are to be a core part of language design, they must not affect the asymptotic complexity of programs.**

In this dissertation, I (a) situate type/contract systems in the design space, (b) develop abstraction and reasoning principles (specifically, relational parametricity) for a type/contract system, and (c) resolve issues with space consumption affecting runtime checking of contracts. The rest of this introduction covers more specific technical details. At the end of the introduction, I remark briefly on notations and conventions (Section 1.5).

1.1 Latent and manifest contracts

Findler and Felleisen's work sparked a resurgence of interest in contracts, and in the intervening years a bewildering variety of related systems have been studied. Broadly, these come in two different sorts. In systems with *latent* contracts, types and contracts are orthogonal features. Examples of this style include Findler and Felleisen's original system, Hinze et al. [41], Blume and McAllester [11], Chitil and Huch [16], Guha et al. [37], and Tobin-Hochstadt and Felleisen [74]. By contrast, *manifest* contracts are integrated into the type system, which tracks, for each value, the most recently checked contract. *Hybrid types* [28] are a well-known example in this

style; others include the work of Ou et al. [51], Wadler and Findler [78], and Knowles et al. [45]. (I discuss related work at length in Chapter 5.)

Findler and Felleisen’s latent calculus annotates function definitions with contracts; to actually check the contracts, they include terms like $\langle\{x:\text{Int} \mid \text{pos } x\}\rangle^{l,l'} 1$, in which a boolean predicate, `pos`, is applied to a run-time value, 1. This term evaluates to 1, since `pos 1` returns `true`. On the other hand, the term $\langle\{x:\text{Int} \mid \text{pos } x\}\rangle^{l,l'} 0$ evaluates to *blame*, written $\uparrow l$, signaling that a contract with label l has been violated. The other label on the contract, l' , comes into play with *function contracts*, $c_1 \mapsto c_2$. For example, the term

$$\langle\{x:\text{Int} \mid \text{nonzero } x\}\rangle^{l,l'} \mapsto \langle\{x:\text{Int} \mid \text{pos } x\}\rangle^{l,l'} (\lambda x:\text{Int}. x - 1)$$

“wraps” the function $\lambda x:\text{Int}. x - 1$ in a pair of checks: whenever the wrapped function is called, the argument is checked to see whether it is nonzero; if not, the blame term $\uparrow l'$ is produced, signaling that the *context* of the contracted term violated the expectations of the contract. If the argument check succeeds, then the function is run and its result is checked against the contract `pos x`, raising $\uparrow l$ if this fails (e.g., if the wrapped function is applied to 1).

The key feature of manifest systems is that descriptions like $\{x:\text{Int} \mid \text{nonzero } x\}$ are incorporated into the type system as *refinement types*. Values of refinement type are introduced via *casts* like $\langle\{x:\text{Int} \mid \text{true}\}\rangle \Rightarrow \langle\{x:\text{Int} \mid \text{nonzero } x\}\rangle^l n$, which has static type $\{x:\text{Int} \mid \text{nonzero } x\}$ and checks, dynamically, that n really is nonzero, raising $\uparrow l$ otherwise. Similarly, $\langle\{x:\text{Int} \mid \text{nonzero } x\}\rangle \Rightarrow \langle\{x:\text{Int} \mid \text{pos } x\}\rangle^l n$ casts an integer that is statically known to be nonzero to one that is statically known to be positive.

The manifest analogue of function contracts is casts between function types. For example, consider:

$$f = \langle[\text{Int}] \rightarrow [\text{Int}] \Rightarrow \{x:\text{Int} \mid \text{pos } x\} \rightarrow \{x:\text{Int} \mid \text{pos } x\}\rangle^l (\lambda x:[\text{Int}]. x - 1),$$

where $[\text{Int}] = \{x:\text{Int} \mid \text{true}\}$. The sequence of events when f is applied to some argument n (of type P) is similar to what we saw before:

$$f n \longrightarrow_h \langle[\text{Int}] \Rightarrow \{x:\text{Int} \mid \text{pos } x\}\rangle^l ((\lambda x:[\text{Int}]. x - 1) (\langle\{x:\text{Int} \mid \text{pos } x\}\rangle^l n))$$

First, n is cast from $\{x:\text{Int} \mid \text{pos } x\}$ to $[\text{Int}]$ (it happens that in this case the cast cannot fail, since the target predicate is just `true`, but if it did, it would raise $\uparrow l$); then the function body is evaluated; and finally its result is cast from $[\text{Int}]$ to $\{x:\text{Int} \mid \text{pos } x\}$, raising $\uparrow l$ if this fails. The domain cast is contravariant and the codomain cast is covariant.

One point to note here is that casts in the manifest system have just one label, while contract checks in the latent system have two. This difference is not fundamental to the latent/manifest distinction—both latent and manifest systems can be given more or less rich algebras of blame—but rather a question of the pragmatics of assigning responsibility: contract checks (called *obligations* in Findler and Felleisen [26]) use

two labels, while casts use one. Informally, a function contract check $\langle c_1 \mapsto c_2 \rangle^{l,l'}$ f divides responsibility for f 's behavior between its body and its environment: the programmer is saying “If f is ever applied to an argument that does not pass c_1 , I refuse responsibility ($\uparrow l'$), whereas if f 's result for good arguments does not satisfy c_2 , I accept responsibility ($\uparrow l$).” In a system with casts, the programmer who writes $\langle R_1 \rightarrow R_2 \Rightarrow S_1 \rightarrow S_2 \rangle^l$ f is saying “Although all I know statically about f is that its results satisfy R_2 when it is applied to arguments satisfying R_1 , I assert that it's okay to use it on arguments satisfying S_1 [because I believe that S_1 implies R_1] and that its results will always satisfy S_2 [because R_2 implies S_2].” In the latter case, the programmer is taking responsibility for *both* assertions (so $\uparrow l$ makes sense in both cases), while the additional responsibility for checking that arguments satisfy S_1 will be discharged elsewhere (by another cast, with a different label).

We compare and contrast the latent and manifest semantics in Chapter 2. This lays the foundation: what are the semantics of contracts, and how do existing semantics in the literature relate to each other?

1.2 Contracts and abstraction

In Chapter 3, I study how contracts and type abstraction interact. Type abstraction is an essential part of language design, especially for higher-order languages. If we are going to design programming languages with manifest contracts, it's important that we know that type abstraction still follows our intuition.

It turns out that manifest contracts not only enjoy relational parametricity, the gold standard property of type abstraction, but that manifest contracts allow programmers to design particularly rich abstractions. Contracts and *polymorphism* make a natural combination: programmers can give strong contracts to abstract types, precisely stating pre- and post-conditions while hiding implementation details—for example, an abstract type of stacks might specify that the `pop` operation requires non-empty stacks as input, i.e., have the domain type $\{x:\alpha \text{ Stack} \mid \text{not}(\text{empty? } x)\}$. That is, just as the type abstraction of ML was used to force sound reasoning in proofs [32], we can use type abstraction to force sound uses of partial operations.

As one simple example, consider the following abstract type for natural numbers:

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

Contracts have allowed us to specify a very precise type for `pred`—one which ensures the correctness of the function, since `pred zero` isn't a natural. Since these are *manifest* contracts, the type system will track which values are known to be non-zero. In Section 3.1, I develop a more extended example of an embedded domain-specific language of transducers based on Boomerang [12].

Beyond these rich abstractions, subtyping and relational parametricity offer structured reasoning principles that are invaluable to compiler writers (for optimizations

```

odd 3
→ ((Dyn→Dyn ⇒ Int→Bool) even) 2
→ ⟨Dyn ⇒ Bool⟩ (even ((Int ⇒ Dyn) 2))
→ ⟨Dyn ⇒ Bool⟩ (even 2Int!)
→ ⟨Dyn ⇒ Bool⟩ (((Int→Bool ⇒ Dyn→Dyn) (λx:Int. ...)) 2Int!)
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) ((λx:Int. ...) ((Dyn ⇒ Int) 2Int!)))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) ((λx:Int. ...) 2))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) (odd 1))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) (((Dyn→Dyn ⇒ Int→Bool) even) 0))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) ((Dyn ⇒ Bool) (even ((Int ⇒ Dyn) 0))))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) ((Dyn ⇒ Bool) (even 0Int!)))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) ((Dyn ⇒ Bool)
  ((Int→Bool ⇒ Dyn→Dyn) (λx:Int. ...) 0Int!)))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) ((Dyn ⇒ Bool) ((Bool ⇒ Dyn)
  (λx:Int. ...) ((Dyn ⇒ Int) 0Int!))))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) ((Dyn ⇒ Bool) ((Bool ⇒ Dyn)
  (λx:Int. ...) 0))))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) ((Dyn ⇒ Bool) ((Bool ⇒ Dyn) true)))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) ((Dyn ⇒ Bool) trueBool!))
→ ⟨Dyn ⇒ Bool⟩ ((Bool ⇒ Dyn) true)
→ (Dyn ⇒ Bool) trueBool!
→ true

```

Figure 1.2: Contract checking destroys tail recursion

to representations, deforestation, etc.) and programmers (for formally supported reasoning). These principles are one of the payoffs of the manifest approach.

Finally, our development incorporates a significant technical improvement over earlier presentations of contracts (e.g., Flanagan [28] and Greenberg et al. [34], which is in effect Chapter 2): instead of introducing a denotational model to break a problematic circularity between typing, subtyping, and evaluation, I 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.

1.3 Efficiency and contract checking

Since casts and checks have runtime effects, we must ask: what are the costs? Racket’s contracts are essentially checked along the same lines as Findler and Felleisen’s plan, as in the wrap function:³

```

(define (wrap ctc v)
  (if (flat-contract? ctc)
      (if ((flat-contract-pred ctc) v)

```

³For simplicity’s sake, I’ve omitted blame.

```

      v
      (error (flat-contract-err-info ctc)))
    (let ([dom (fun-contract-dom ctc)]
          [cod (fun-contract-cod ctc)])
      (lambda (x)
        (wrap cod (v (wrap dom x)))))))

```

That is, flat contracts (like `nonempty-queue/c`) just check the predicate; function contracts η -expand their argument, wrapping the input in the domain contract and the return value in the codomain contract. For example, wrapping `queue?` with its `(-> any/c boolean?)` contract yields:

```

(lambda (x)
  (wrap boolean? (queue? (wrap any/c x))))

```

This η -expanded lambda is called a *function proxy*. In PLT Racket (and all other formulations of higher-order contracts), functions can accumulate an unbounded number of function proxies—with serious ramifications for asymptotic space bounds. Similarly, the codomain contract check in the function proxy can destroy tail recursion.

One longstanding problem with casts is space efficiency: contract checks, in their naïve formulation, can consume unbounded amounts of space at runtime both through excessive function proxies and through tail-recursion-breaking stack growth. This unbounded use of space can lead to changes in the asymptotic space efficiency of programs. Prior work [39, 40, 65, 68] offers space-efficient solutions exclusively in the domain of *gradual types* [67]—type systems integrating simple types and the dynamic type, `Dyn`.

As an example, take the mutually recursive definition of `even` and `odd` from Herman et al. [39]. Here `odd` is written in a simply typed style, but `even` is written in a more dynamically typed style.

$$\begin{aligned}
 \text{even} &: \text{Dyn} \rightarrow \text{Dyn} = \langle \text{Int} \rightarrow \text{Bool} \Rightarrow \text{Dyn} \rightarrow \text{Dyn} \rangle \lambda x: \text{Int}. \\
 &\quad \text{if } x = 0 \text{ then true else odd } (x - 1) \\
 \text{odd} &: \text{Int} \rightarrow \text{Bool} = \lambda x: \text{Int}. \\
 &\quad \text{if } x = 0 \text{ then false else } (\langle \text{Dyn} \rightarrow \text{Dyn} \Rightarrow \text{Int} \rightarrow \text{Bool} \rangle \text{even}) (x - 1)
 \end{aligned}$$

We show how `odd 3` reduces in Figure 1.2. Note how casts between `Bool` and `Dyn` accumulate. Imagine writing `even` and `odd` in a surface language—how predictable is cast insertion? What was written as tail recursive—and ought to use only $O(1)$ space—instead takes $O(n)$ space.

This problem of contracts breaking tail recursion has very much informed the design philosophy of PLT Racket’s contracts [56], where contracts are typically placed at module boundaries. Calls between functions in the same module aren’t mediated by contracts, so intra-module tail recursion still works. Calls between modules do have contracts, though, so there’s no guarantee about tail recursion. Naturally, this

scheme is somewhat limiting: what about libraries with higher-order functions and callbacks?

Following Herman et al. [39], I develop a space efficient scheme for a language with manifest contracts and type `Dyn` in Chapter 4. We show that space efficiency avoids some checks, failing and diverging less often than naïve calculi—but the two are otherwise observationally equivalent.

1.4 Summary

In this thesis, I show that **higher-order contracts are on a firm type-theoretic footing and that they support extensions necessary for real programming**. Developing these principles is the first step in designing and implementing higher-order languages with contracts and refinement types. In this dissertation, I:

- Study the relationship between the two families of core calculi for contracts (Chapter 2),
- Extend one family so that it admits type abstraction and enjoys relational parametricity (Chapter 3), and
- Resolve a show-stopping space-inefficiency (Chapter 4).

I discuss related work in Chapter 5 and conclude with a discussion of future work in Chapter 6.

1.5 Notation and other conventions

There are many different calculi in this dissertation: five in Chapter 2, one in Chapter 3, and three in Chapter 4. What’s more, all but two of these calculi use manifest contracts, and all but two of those use casts. So there is a significant overlap in notation and names. It is easiest to interpret names as scoped by chapter: there are three different rules named `E_APP`, but each is unique in its chapter. Well formedness and typing rules are named `WF_...` and `T_...` (though sometimes also `S_...`); evaluation rules are named `E_...`, but also `F_...` and `G_...`.

We explain the notation in each chapter, but all of the calculi should be familiar as extensions of the simply typed lambda calculus.

This document is thoroughly hyperlinked—references to chapters, sections, figures, theorems, and citations are all linked. While this document reads best in color and on a device supporting linking, I have not marked the links in order to reduce visual noise.

Chapter 2

Contracts made manifest

And something is happening here
But you don't know what it is

Ballad of a Thin Man
Bob Dylan

Recall the distinction between latent and manifest systems made in Section 1.1: latent contracts are orthogonal to the type system (if it exists), while manifest contracts conflate contracts and types. While contract checks in latent systems may seem intuitively to be much the same thing as casts in manifest systems, the formal correspondence is not immediate. Manifest contracts used casts $\langle S_1 \Rightarrow S_2 \rangle^l$ with contravariant checking for functions; latent contracts used obligations $\langle c \rangle^{l,l'}$ with an invariant rule—but with polarized blame labels. How do these forms of checking and blame relate? These questions have led to some confusion in the community about the nature of contracts. Indeed, as we will see, matters become yet murkier in richer languages with features such as dependency.

Gronski and Flanagan [36] initiated a formal investigation of the connection between the latent and manifest worlds. They defined a core calculus, λ_C , capturing the essence of latent contracts in a simply typed lambda-calculus, and an analogous manifest calculus λ_H . I will follow the naming scheme of their treatment. Both calculi offer contract-like dynamic checks, but they differ in the exact structure of these checks, in their treatment of “blame,” and in how dynamic checks and types interact.

In λ_C , dynamic checks occur when a contract is applied, as in $\langle c \rangle^{l,l'} t$. The contract checks to make sure that t is treated and behaves like a c . To achieve the same effect, λ_H applies a cast, $\langle S_1 \Rightarrow S_2 \rangle^l s$. This cast ensures that s —which used to behave like an S_1 —is now treated and behaves like an S_2 , where S_1 and S_2 are (very precise) types. For example, $\langle \{x:\text{Int} \mid \text{pos } x\} \rangle^{l,l'} t$ ensures that t evaluates to a positive integer; $\langle \{x:\text{Int} \mid \text{nonzero } x\} \Rightarrow \{x:\text{Int} \mid \text{pos } x\} \rangle^l s$ ensures that s —which we know to be nonzero—is also positive.

To compare these systems, they introduced a type-preserving translation ϕ from λ_C to λ_H . What makes ϕ interesting is that it relates the languages *feature for feature*:

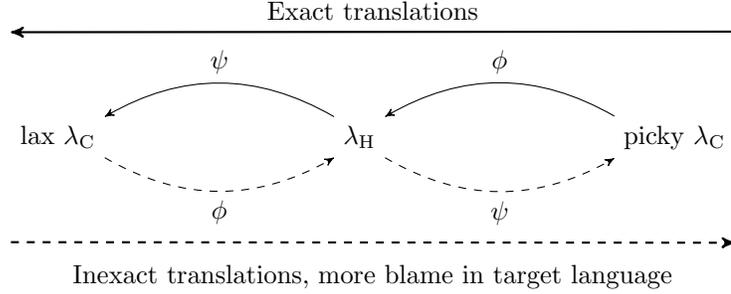


Figure 2.1: The axis of blame

predicate contracts over base types are mapped to casts between refinements of base type, and function contracts are mapped to function casts. The main result is that ϕ preserves behavior, in the sense that if a term t in λ_C evaluates to a constant k or blame $\uparrow l$, then its translation $\phi(t)$ evaluates similarly.

In this chapter, I extend their work in two directions. First, I strengthen their main result by introducing a new feature-for-feature translation ψ from λ_H to λ_C and proving a similar correspondence theorem for ψ . (I also give a new, more detailed, proof of the correspondence theorem for ϕ .) These correspondences show that the manifest and latent approaches are effectively equivalent in the nondependent case.

Second, and more significantly, I extend the whole story to allow dependent function contracts in λ_C and dependent arrow types in λ_H . Dependency is extremely handy in contracts, as it allows for precise specifications of how the results of functions depend on their arguments. For example, here is a contract that we might use with an implementation of vector concatenation:

$$z_1:\text{Vec} \mapsto z_2:\text{Vec} \mapsto \{z_3:\text{Vec} \mid \text{vlen } z_3 = \text{vlen } z_1 + \text{vlen } z_2\}$$

Adding dependent contracts to λ_C is easy: the dependency is all in the contracts and the types stay simple. We have just one significant design choice: should domain contracts be rechecked when the bound variable appears in the codomain contract? This choice leads to two dialects of λ_C , one which does recheck (*picky* λ_C) and one which does not (*lax* λ_C). The choice is not clear, so I consider both. The question of which blame labels belong on this extra check is discussed at length in Dimoulas et al. [22], which introduces *indy* blame. Indy blame is a variant of picky. I do not consider it in depth here, since it does not affect *whether or not* blame is raised, only *which* blame. I discuss this point more in Section 2.6.3. In λ_H , on the other hand, dependency significantly complicates the metatheory, requiring the addition of a denotational semantics for types and kinds (a unary logical relation with a term model) to break a potential circularity in the definitions, plus an intricate sequence of technical lemmas involving parallel reduction to establish type soundness.

Surprisingly, the tight correspondence between λ_C and λ_H breaks down in the dependent case: the natural generalization of the translations does not preserve behavior exactly. Indeed, we can place λ_H between the two variants of λ_C on an “axis

of blame” (Figure 2.1), where evaluation behavior is preserved exactly when moving left on the axis (from picky λ_C to λ_H to lax λ_C), but translated terms can blame more than their pre-images when moving right.¹ It is still the case that when a pre-image raises blame, its translation blames as well—though not necessarily the same label. The discrepancy arises in the case of *abusive* contracts, such as

$$f: (\{x:\text{Int} \mid \text{nonzero } x\} \mapsto \{y:\text{Int} \mid \text{true}\}) \mapsto \{z:\text{Int} \mid f\ 0 = 0\}$$

This rather strange contract has the form $f:c_1 \mapsto c_2$, where c_2 uses f in a way that violates c_1 ! In particular, if we apply it (in lax λ_C) to $\lambda f:\text{Int} \rightarrow \text{Int}. 0$ and then apply the result to $\lambda x:\text{Int}. x$ and 5, the final result will be 5, since $\lambda x:\text{Int}. x$ does satisfy the contract $\{x:\text{Int} \mid \text{nonzero } x\} \mapsto \{y:\text{Int} \mid \text{true}\}$ and 5 satisfies the contract $\{z:\text{Int} \mid (\lambda x:\text{Int}. x)\ 0 = 0\}$. However, running the translation of f in λ_H yields an extra check, wrapping the occurrence of f in the codomain contract with a cast from $\{x:\text{Int} \mid \text{nonzero } x\} \rightarrow \{y:\text{Int} \mid \text{true}\}$ to $\{x:\text{Int} \mid \text{true}\} \rightarrow \{y:\text{Int} \mid \text{true}\}$, which fails when the wrapped function is applied to 0. I discuss this phenomenon in greater detail in Section 2.3.

We should note at the outset that, like Gronski and Flanagan [36], we are interested in translations that relate λ_C and λ_H feature for feature, i.e., mapping predicate contracts to predicate contracts and function contracts to function contracts. Translations which don’t map feature for feature can give an exact treatment of blame. Consider the following dependent version of the wrap operator from Findler and Felleisen [26]. There are two cases: one for refinements of base types B , one for dependent function contracts.

$$\begin{aligned} \phi(\langle\{x:B \mid t\}\rangle^{l,l'}) &= \langle[B] \Rightarrow \{x:B \mid \phi(t)\}\rangle^l \\ \phi(\langle x:c_1 \mapsto c_2 \rangle^{l,l'}) &= \lambda f: [x:c_1 \mapsto c_2]. \\ &\quad \lambda x: [c_1]. \\ &\quad \phi(\langle c_2 \rangle^{l,l'}) (f (\phi(\langle c_1 \rangle^{l,l'}) x)) \end{aligned}$$

We can define a similar mapping function that implements λ_H ’s semantics as predicate contracts in lax or picky λ_C . It is unsurprising that an exact mapping exists: λ_C and λ_H are lambda calculi that feature, among other things, a way to conditionally raise exceptions. That these languages are interencodable is completely unsurprising. But translations like these do not relate function contracts to function casts at all, so they do not do much to tell us about how semantics of contracts and the semantics of casts relate.

In summary, the main contributions of this chapter are (a) the translation ψ and a symmetric version of Gronski and Flanagan’s behavioral correspondence theorem, (b) the basic metatheory of (CBV, blame-sensitive) dependent λ_H , (c) dependent

¹There might, in principle, be some other way of defining ϕ and ψ that (a) preserves types, (b) maps feature for feature, and (c) induces an exact behavioral equivalence. After considering a number of alternatives, we conjecture that no such ϕ and ψ exist.

$$\begin{array}{ll}
B ::= \text{Bool} \mid \dots & \text{base types} \\
k ::= \text{true} \mid \text{false} \mid \dots & \text{first-order constants}
\end{array}$$

Figure 2.2: Base types and constants for λ_C and λ_H

versions of ϕ and ψ and their properties with regard to λ_H and both dialects of λ_C , and (d) a weaker behavioral correspondence in the dependent case. I restrict my attention to strongly normalizing programs, though I believe the results should generalize readily to programs with recursion and nontermination. This chapter is an extended adaptation of Greenberg et al. [34] and Greenberg et al. [35].

2.1 The nondependent languages

We begin in this section by defining the nondependent versions of λ_C and λ_H and continue in Section 2.2 with the translations between them. The dependent languages, dependent translations, and their properties are developed in Sections 2.3, 2.5, and 2.6. Throughout this chapter, rules prefixed with an E or an F are operational rules for λ_C and λ_H , respectively. An initial T is used for λ_C typing rules; typing rules beginning with an S belong to λ_H .

All of the languages will share a set of base types and first-order constants, given in Figure 2.2. Let the set \mathcal{K}_B contain constants of base type B . We assume that `Bool` is among the base types, with $\mathcal{K}_{\text{Bool}} = \{\text{true}, \text{false}\}$.

2.1.1 The language λ_C

The language λ_C is the simply typed lambda calculus straightforwardly augmented with contracts. *Contracts* c come in two forms: predicate contracts $\{x:B \mid t\}$ over a base type B and higher-order contracts $c_1 \mapsto c_2$, which check the arguments and results of functions. We can use contracts in terms with the *contract obligation* $\langle c \rangle^{l,l'}$. Applying a contract obligation $\langle c \rangle^{l,l'}$ to a term t dynamically ensures that t and its surrounding context satisfy c . If t does not satisfy c , then the *positive* label l will be blamed and the whole term will reduce to $\uparrow l$; on the other hand, if the context does not treat $\langle c \rangle^{l,l'}$ t as c demands, then the *negative* label l' will be blamed and the term will reduce to $\uparrow l'$. In contexts where it is unambiguous, I refer to contract obligations simply as contracts.

The syntax and semantics of λ_C appears in Figure 2.3, with some common definitions (shared with λ_H) in Figure 2.2. Besides the contract term $\langle c \rangle^{l,l'}$, λ_C includes first-order constants k , blame, and *active checks* $\langle \{x:B \mid t_1\}, t_2, k \rangle^l$. Active checks do not appear in source programs; they are a technical artifact of the small-step operational semantics, as I explain below. Also, note that my contracts only allow refinements over base types B : we have function contracts, like $\{x:\text{Int} \mid \text{pos } x\} \mapsto \{x:\text{Int} \mid$

Syntax for $\lambda_{\mathbf{C}}$

$T ::= B \mid T_1 \rightarrow T_2$	types
$c ::= \{x:B \mid t\} \mid c_1 \mapsto c_2$	contracts
$t ::= x \mid k \mid \lambda x:T_1. t_2 \mid t_1 t_2 \mid$	terms
$\quad \uparrow l \mid \langle c \rangle^{l,l'} \mid \langle \{x:B \mid t_1\}, t_2, k \rangle^l$	
$v ::= k \mid \lambda x:T_1. t_2 \mid \langle c \rangle^{l,l'} \mid \langle c_1 \mapsto c_2 \rangle^{l,l'} v$	values
$r ::= v \mid \uparrow l$	results
$E ::= [] t \mid v [] \mid \langle \{x:B \mid t\}, [], k \rangle^l$	evaluation contexts

Operational semantics for $\lambda_{\mathbf{C}}$

$(\lambda x:T_1. t_2) v \longrightarrow_c t_2\{x := v\}$	E_BETA
$k v \longrightarrow_c \llbracket k \rrbracket(v)$	E_CONST
$\langle \{x:B \mid t\}^{l,l'} k \rangle^l \longrightarrow_c \langle \{x:B \mid t\}, t\{x := k\}, k \rangle^l$	E_CCHECK
$\langle \{x:B \mid t\}, \text{true}, k \rangle^l \longrightarrow_c k$	E_OK
$\langle \{x:B \mid t\}, \text{false}, k \rangle^l \longrightarrow_c \uparrow l$	E_FAIL
$\langle \langle c_1 \mapsto c_2 \rangle^{l,l'} v \rangle^l v' \longrightarrow_c \langle c_2 \rangle^{l,l'} (v (\langle c_1 \rangle^{l,l'} v'))$	E_CDECOMP
$E[\uparrow l] \longrightarrow_c \uparrow l$	E_BLAZE
$E[t_1] \longrightarrow_c E[t_2] \quad \text{when } t_1 \longrightarrow_c t_2$	E_COMPAT

Typing rules for $\lambda_{\mathbf{C}}$

$\boxed{\Gamma \vdash t : T}$	
$\frac{x:T \in \Gamma}{\Gamma \vdash x : T}$	T_VAR
$\frac{\Gamma, x:T_1 \vdash t_2 : T_2}{\Gamma \vdash \lambda x:T_1. t_2 : T_1 \rightarrow T_2}$	T_LAM
$\frac{}{\Gamma \vdash k : \text{ty}_c(k)}$	T_CONST
$\frac{\Gamma \vdash t_1 : T_1 \rightarrow T_2 \quad \Gamma \vdash t_2 : T_1}{\Gamma \vdash t_1 t_2 : T_2}$	T_APP
$\frac{}{\Gamma \vdash \uparrow l : T}$	T_BLAZE
$\frac{\vdash c : T}{\Gamma \vdash \langle c \rangle^{l,l'} : T \rightarrow T}$	T_CONTRACT
$\frac{\emptyset \vdash k : B \quad \emptyset \vdash t_2 : \text{Bool} \quad \vdash \{x:B \mid t_1\} : B \quad t_2 \longrightarrow_c^* \text{true} \text{ implies } t_1\{x := k\} \longrightarrow_c^* \text{true}}{\emptyset \vdash \langle \{x:B \mid t_1\}, t_2, k \rangle^l : B}$	T_CHECKING
$\boxed{\vdash c : T}$	
$\frac{x:B \vdash t : \text{Bool}}{\vdash \{x:B \mid t\} : B}$	T_BASEC
$\frac{\vdash c_1 : T_1 \quad \vdash c_2 : T_2}{\vdash c_1 \mapsto c_2 : T_1 \rightarrow T_2}$	T_FUNC

Figure 2.3: Syntax and semantics for $\lambda_{\mathbf{C}}$

nonzero x }, but not predicate contracts over functions themselves, like $\{f:\text{Bool} \rightarrow \text{Bool} \mid f \text{ true} = f \text{ false}\}$.

Values v include constants, abstractions, contracts, and function contracts applied to values (more on these later); a *result* r is either a value or $\uparrow l$ for some l . We interpret constants using two constructions: the type-assignment function ty_c , which maps constants to first-order types of the form $B_1 \rightarrow B_2 \rightarrow \dots \rightarrow B_n$ (and which is assumed to agree with \mathcal{K}_B); and the denotation function $\llbracket - \rrbracket$, which maps constants to functions from constants to constants (or blame, to allow for partiality). Denotations must agree with ty_c , i.e., if $\text{ty}_c(k) = B_1 \rightarrow B_2$, then $\llbracket k \rrbracket(k_1) \in \mathcal{K}_{B_2}$ if $k_1 \in \mathcal{K}_{B_1}$.

The operational semantics is given in Figure 2.3. It includes six rules for basic (small-step, call-by-value) reductions, plus two rules that involve evaluation contexts E (Figure 2.3). The evaluation contexts implement left-to-right evaluation for function application. If $\uparrow l$ appears in the active position of an evaluation context, it is propagated to the top level, like an uncatchable exception. As usual, values (and results) do not step.

The first two basic rules are standard, implementing primitive reductions and β -reductions for abstractions. In these rules, arguments must be values v . Since constants are first-order, we know that when `E_CONST` reduces a well-typed application, the argument is not just a value, but a constant.

The rules `E_CCHECK`, `E_OK`, `E_FAIL` and `E_CDECOMP`, describe the semantics of contracts. In `E_CCHECK`, predicate contracts applied to constants step to an active check. Active checks include the original contract, the current state of the check, the constant being checked, and a label to blame if necessary. I hold on to the original contract as a technical device for the translation ϕ from λ_C to λ_H , since λ_H needs to know the target type of an active check. If the check evaluates to `true`, then `E_OK` returns the initial constant. If `false`, the check has failed and a contract has been violated, so `E_FAIL` steps the term to $\uparrow l$. Higher-order contracts on a value v wait to be applied to an additional argument. This is why function contracts applied to values are values. There is no substantial difference between this approach and expanding function contracts into new lambdas. When that argument has also been reduced to a value v' , `E_CDECOMP` decomposes the function cast: the argument value is checked with the argument part of the contract (switching positive and negative blame, since the context is responsible for the argument), and the result of the application is checked with the result contract.

The typing rules for λ_C (Figure 2.3) are mostly standard. We give types to constants using the type-assignment function ty_c . Blame expressions have all types. Contracts are checked for well-formedness using the judgment $\vdash c : T$, comprising the rules `T_BASEC`, which requires that the checking term in a predicate contract return a boolean value when supplied with a term of the right type, and `T_FUNC`. Note that the predicate t in a contract $\{x:B \mid t\}$ can contain at most x free, since we are considering only nondependent contracts for now. Contract application, like function application, is checked using `T_APP`.

The `T_CHECKING` rule only applies in the empty context (active checks are only created at the top level during evaluation). The rule ensures that the contract $\{x:B \mid t_1\}$ has the right base type for the constant k , that the check expression t_2 has a boolean type, and that the check is actually checking the right contract. The latter condition is formalized by the implication: $t_2 \longrightarrow_c^* \text{true}$ implies $t_1\{x := k\} \longrightarrow_c^* \text{true}$ asserts that if t_2 evaluates to `true`, then the original check $t_1\{x := k\}$ must also evaluate to `true`. This requirement is needed for two reasons: first, nonsensical terms like $\langle\{x:\text{Int} \mid \text{pos } x\}, \text{true}, 0\rangle^l$ should not be well typed; and second, I use this property in showing that the translations are type preserving (see Section 2.5). One may think that this rule makes typechecking for the full “internal language” with checks undecidable—I certainly did—but in fact the entire language is strongly normalizing.² I could give a more precise condition—for example, that $t_1\{x := k\} \longrightarrow_c^* t_2$ —but there is no need. We will find that this condition is more useful in the metatheory of Chapters 3 and 4.

The language enjoys standard preservation and progress theorems. Together, these ensure that evaluating a well-typed term to a normal form always yields a result r , which is either blame or a value of the appropriate type.

2.1.2 The language λ_H

The second calculus, nondependent λ_H , extends the simply typed lambda-calculus with *refinement types* and *cast expressions*. The definitions appear in Figure 2.4 (syntax and semantics) and Figure 2.5. Unlike λ_C , which separates contracts from types, λ_H combines them into refined base types $\{x:B \mid s_1\}$ and function types $S_1 \rightarrow S_2$. As for λ_C , I do not allow refinement types over functions nor refinements of refinements. (We add these features to a dependent λ_H in Chapter 3.) Unrefined base types B are *not* valid types; they must be wrapped in a trivial refinement, as the *raw* type $\{x:B \mid \text{true}\}$. The terms of the language are mostly standard, including variables, the same first-order constants as λ_C , blame, abstractions, and applications. The cast expression $\langle S_1 \Rightarrow S_2 \rangle^l$ dynamically checks that a term of type S_1 can be given type S_2 . Like λ_C , active checks are used to give a small-step semantics to cast expressions.

The values of λ_H include constants, abstractions, casts, and function casts applied to values. Results are either values or blame. We give meaning to constants as we did in λ_C , reusing $\llbracket - \rrbracket$. Type assignment is via ty_h , which we assume produces well-formed types (defined in Figure 2.5). To keep the languages in sync, I require that ty_h and ty_c agree on “type skeletons”: if $\text{ty}_c(k) = B_1 \rightarrow B_2$, then $\text{ty}_h(k) = \{x:B_1 \mid s_1\} \rightarrow \{x:B_2 \mid s_2\}$.

The small-step, call-by-value semantics in Figure 2.4 comprises six basic rules and two rules involving evaluation contexts F . Each rule corresponds closely to its λ_C counterpart.

²As observed by Stephan Zdancewic at my defense.

Syntax for $\lambda_{\mathbf{H}}$

$S ::= \{x:B \mid s_1\} \mid S_1 \rightarrow S_2$	types/contracts
$s ::= x \mid k \mid \lambda x:S_1. s_2 \mid s_1 s_2 \mid \uparrow l \mid \langle S_1 \Rightarrow S_2 \rangle^l \mid \langle \{x:B \mid s_1\}, s_2, k \rangle^l$	terms
$w ::= k \mid \lambda x:S_1. s_2 \mid \langle S_1 \Rightarrow S_2 \rangle^l \mid \langle S_{11} \rightarrow S_{12} \Rightarrow S_{21} \rightarrow S_{22} \rangle^l w$	values
$q ::= w \mid \uparrow l$	results
$F ::= [] s \mid w [] \mid \langle \{x:B \mid s\}, [], k \rangle^l$	evaluation contexts

Operational semantics for $\lambda_{\mathbf{H}}$

$(\lambda x:S_1. s_2) w_2 \longrightarrow_h s_2\{x := w_2\}$	F_BETA
$k w \longrightarrow_h \llbracket k \rrbracket(w)$	F_CONST
$\langle \{x:B \mid s_1\} \Rightarrow \{x:B \mid s_2\} \rangle^l k \longrightarrow_h \langle \{x:B \mid s_2\}, s_2\{x := k\}, k \rangle^l$	F_CCHECK
$\langle \{x:B \mid s\}, \text{true}, k \rangle^l \longrightarrow_h k$	F_OK
$\langle \{x:B \mid s\}, \text{false}, k \rangle^l \longrightarrow_h \uparrow l$	F_FAIL
$\langle \langle S_{11} \rightarrow S_{12} \Rightarrow S_{21} \rightarrow S_{22} \rangle^l w \rangle w' \longrightarrow_h \langle S_{12} \Rightarrow S_{22} \rangle^l (w (\langle S_{21} \Rightarrow S_{11} \rangle^l w'))$	F_CDECOMP
$F[\uparrow l] \longrightarrow_h \uparrow l$	F_BLAME
$F[s_1] \longrightarrow_h F[s_2]$	F_COMPAT
$\text{when } s_1 \longrightarrow_h s_2$	

Figure 2.4: Syntax and semantics for $\lambda_{\mathbf{H}}$

Typing rules for $\lambda_{\mathbf{H}}$

$$\begin{array}{c}
\boxed{\Delta \vdash s : S} \\
\\
\frac{x:S \in \Delta}{\Delta \vdash x : S} \text{ S_VAR} \qquad \frac{\vdash S_1 \quad \Delta, x:S_1 \vdash s_2 : S_2}{\Delta \vdash \lambda x:S_1. s_2 : S_1 \rightarrow S_2} \text{ S_LAM} \\
\\
\frac{}{\Delta \vdash k : \text{ty}_h(k)} \text{ S_CONST} \qquad \frac{\Delta \vdash s_1 : S_1 \rightarrow S_2 \quad \Delta \vdash s_2 : S_1}{\Delta \vdash s_1 s_2 : S_2} \text{ S_APP} \\
\\
\frac{\vdash S}{\Delta \vdash \uparrow l : S} \text{ S_BLAME} \qquad \frac{\vdash S_1 \quad \vdash S_2 \quad [S_1] = [S_2]}{\Delta \vdash \langle S_1 \Rightarrow S_2 \rangle^l : S_1 \rightarrow S_2} \text{ S_CAST} \\
\\
\frac{\Delta \vdash s : S_1 \quad \vdash S_2 \quad \vdash S_1 <: S_2}{\Delta \vdash s : S_2} \text{ S_SUB} \\
\\
\frac{\emptyset \vdash k : \{x:B \mid \text{true}\} \quad \emptyset \vdash s_2 : \{x:\text{Bool} \mid \text{true}\} \quad \vdash \{x:B \mid s_1\} \quad s_2 \rightarrow_h^* \text{true} \text{ implies } s_1 \{x := k\} \rightarrow_h^* \text{true}}{\emptyset \vdash \langle \{x:B \mid s_1\}, s_2, k \rangle^l : \{x:B \mid s_1\}} \text{ S_CHECKING} \\
\\
\boxed{\vdash S_1 <: S_2} \\
\\
\frac{\forall k \in \mathcal{K}_B. (s_1 \{x := k\} \rightarrow_h^* \text{true} \text{ implies } s_2 \{x := k\} \rightarrow_h^* \text{true})}{\vdash \{x:B \mid s_1\} <: \{x:B \mid s_2\}} \text{ SSUB_REFINE} \\
\\
\frac{\vdash S_{21} <: S_{11} \quad \vdash S_{12} <: S_{22}}{\vdash S_{11} \rightarrow S_{12} <: S_{21} \rightarrow S_{22}} \text{ SSUB_FUN} \\
\\
\boxed{\vdash S} \\
\\
\frac{}{\vdash \{x:B \mid \text{true}\}} \text{ SWF_RAW} \qquad \frac{\vdash S_1 \quad \vdash S_2}{\vdash S_1 \rightarrow S_2} \text{ SWF_FUN} \\
\\
\frac{x:\{x:B \mid \text{true}\} \vdash s : \{x:\text{Bool} \mid \text{true}\}}{\vdash \{x:B \mid s\}} \text{ SWF_REFINE}
\end{array}$$

Figure 2.5: Typing rules for $\lambda_{\mathbf{H}}$

Notice how the decomposition rules compare. In λ_C , the term $(\langle c_1 \mapsto c_2 \rangle^{l,l'} v) v'$ decomposes into two contract checks: c_1 checks the argument v' and c_2 checks the result of the application. In λ_H the term $(\langle S_{11} \rightarrow S_{12} \Rightarrow S_{21} \rightarrow S_{22} \rangle^l w) w'$ decomposes into two casts: a contravariant cast on the argument and a covariant cast on the result. The contravariant cast $\langle S_{21} \Rightarrow S_{11} \rangle^l w'$ makes w' a suitable input for w , while $\langle S_{12} \Rightarrow S_{22} \rangle^l$ casts the result from w applied to (the cast) w' . Suppose $S_{21} = \{x:\text{Int} \mid \text{pos } x\}$ and $S_{11} = \{x:B \mid \text{nonzero } x\}$. Then the check on the argument ensures that $\text{nonzero } x \rightarrow_h^* \text{true}$ —not, as one might expect, that $\text{pos } w' \rightarrow_h^* \text{true}$. While it is easy to read off from a λ_C contract exactly which checks will occur at runtime, a λ_H cast must be carefully inspected to see exactly which checks will take place. On the other hand, which label will be blamed is clearer with casts—there’s only one!

The typing rules for λ_H (Figure 2.5) are also similar to those of λ_C . Just as the λ_C rule `T_CONTRACT` checks to make sure that the contract has the right form, the λ_H rule `S_CAST` ensures that the two types in a cast are well-formed and have the same simple-type skeleton, defined as $[-] : S \rightarrow T$ (pronounced “erase S ”):

$$\begin{aligned} [\{x:B \mid s\}] &= B \\ [S_1 \rightarrow S_2] &= [S_1] \rightarrow [S_2] \end{aligned}$$

This prevents “stupid” casts, like $\langle [\text{Int}] \Rightarrow [\text{Bool}] \rangle^l$. I define a similar operator, $[-] : S \rightarrow S$ (pronounced “raw S ”), which trivializes all refinements:

$$\begin{aligned} [\{x:B \mid s\}] &= \{x:B \mid \text{true}\} \\ [S_1 \rightarrow S_2] &= [S_1] \rightarrow [S_2] \end{aligned}$$

These operations apply to λ_C contracts and types in the natural way. Type well-formedness in λ_H is similar to contract well-formedness in λ_C , though the `WF_RAW` case needs to be added to get things off the ground.

The active check rule `S_CHECKING` plays a role analogous to the `T_CHECKING` rule in λ_C , again using an implication to guarantee that we only have sensible terms in the predicate position. Note that we retain the target type in the active check, and that `S_CHECKING` gives active checks that type—technical moves necessary for preservation.

An important difference is that λ_H has subtyping. The `S_SUB` rule allows an expression to be promoted to any well-formed supertype. Refinement types are super-types if, for all constants of the base type, their condition evaluates to `true` whenever the subtype’s condition evaluates to `true`. For function types, we use the standard contravariant subtyping rule. I do not consider source programs with subtyping, since subtyping makes type checking undecidable³; subtyping is just a technical device for ensuring type preservation. Consider the following reduction:

$$\langle \{x:\text{Int} \mid \text{true}\} \Rightarrow \{x:\text{Int} \mid \text{pos } x\} \rangle^l 1 \rightarrow_h^* 1$$

³Flanagan [28] and Knowles and Flanagan [44] discuss trade-offs between static and dynamic checking that allow for decidable type systems and subtyping.

The source term is well-typed at $\{x:\text{Int} \mid \text{pos } x\}$. Since it evaluates to 1, we would like to have $\Delta \vdash 1 : \{x:\text{Int} \mid \text{pos } x\}$. To have type preservation in general, though, $\text{ty}_h(1)$ must be a subtype of $\{x:\text{Int} \mid s\}$ whenever $s\{x := 1\} \longrightarrow_h^* \text{true}$. That is, constants of base type must have “most-specific” types. One way to satisfy this requirement is to set $\text{ty}_h(k) = \{x:B \mid x = k\}$ for $k \in \mathcal{K}_B$; then if $s\{x := k\} \longrightarrow_h^* \text{true}$, we have $\vdash \text{ty}_h(k) <: \{x:B \mid s\}$. This approach is taken in Knowles and Flanagan [44] and Ou et al. [51].

Standard progress and preservation theorems hold for λ_H . We can also obtain a semantic type soundness theorem as a restriction of the one for dependent λ_H (Theorem 2.3.12).

2.2 The nondependent translations

The latent and manifest calculi differ in a few respects. Obviously, λ_C uses contract application and λ_H uses casts. Second, λ_C contracts have two labels—positive and negative—where λ_H contracts have a single label. Finally, λ_H has a much richer type system than λ_C . The translation ψ from λ_H to λ_C and Gronski and Flanagan’s ϕ from λ_C to λ_H must account for these differences while carefully mapping “feature for feature”.

Since I give the translations for the dependent languages in full in Section 2.3, I merely sketch the main ideas here.

The interesting parts of the translations deal with contracts and casts. Everything else is translated homomorphically, though the type annotation on lambdas must be chosen carefully. The full definitions of these translations are in Section 2.4; the nondependent definitions are a straightforward restriction, so I omit them.

For ψ , translating from λ_H ’s rich types to λ_C ’s simple types is easy: we just erase the types to their simple skeletons. The interesting case is translating the cast $\langle S_1 \Rightarrow S_2 \rangle^l$ to a contract by translating the pair of types together, $\langle \psi(S_1, S_2) \rangle^{l,l}$. I define ψ as two mutually recursive functions: $\psi(s)$ translates λ_H terms to λ_C terms; $\psi(S_1, S_2)$ translates a pair of λ_H types—effectively, a cast—to a λ_C contract. The latter function is defined as follows:

$$\begin{aligned} \psi(\{x:B \mid s_1\}, \{x:B \mid s_2\}) &= \{x:B \mid \psi(s_2)\} \\ \psi(S_{11} \rightarrow S_{12}, S_{21} \rightarrow S_{22}) &= \psi(S_{21}, S_{11}) \mapsto \psi(S_{12}, S_{22}) \end{aligned}$$

We use the single label on the cast in both the positive and negative positions of the resulting contract, i.e.:

$$\psi(\langle S_1 \Rightarrow S_2 \rangle^l) = \langle \psi(S_1, S_2) \rangle^{l,l}.$$

When we translate a pair of refinement types, we produce a contract that will check the predicate of the target type (like `F_CCHECK`); when translating a pair of function types, we translate the domain contravariantly (like `F_CDECOMP`). For example,

$$\langle \{x:\text{Int} \mid \text{nonzero } x\} \rightarrow [\text{Int}] \Rightarrow [\text{Int}] \rightarrow \{y:\text{Int} \mid \text{pos } y\} \rangle^l$$

translates to $\langle \{x:\text{Int} \mid \text{nonzero } x\} \mapsto \{y:\text{Int} \mid \text{pos } y\} \rangle^{l,l'}$.

Translating from λ_C to λ_H , we are moving from a simple type system to a rich one. The translation ϕ (essentially the same as Gronski and Flanagan’s) generates terms in λ_H with *raw* types— λ_H types with trivial refinements, corresponding to λ_C ’s simple types. Since the translation targets raw types, the type preservation theorem is stated as “if $\Gamma \vdash t : T$ then $[\Gamma] \vdash \phi(t) : [T]$ ” (see Section 2.6.1).

Whereas the difficulty with ψ is ensuring that the checks match up, the difficulty with ϕ is ensuring that the terms in λ_C and λ_H will blame the same labels. I deal with this problem by translating a single contract with two blame labels into two separate casts. Intuitively, the cast carrying the negative blame label will run all of the checks in negative positions in the contract, while the cast with the positive blame label will run the positive checks. Let

$$\phi(\langle c \rangle^{l,l'}) = \lambda x: [c]. \langle \phi(c) \Rightarrow [c] \rangle^{l'} (\langle [c] \Rightarrow \phi(c) \rangle^l x),$$

where the translation of contracts to refined types is:

$$\begin{aligned} \phi(\{x:B \mid t\}) &= \{x:B \mid \phi(t)\} \\ \phi(c_1 \mapsto c_2) &= \phi(c_1) \rightarrow \phi(c_2) \end{aligned}$$

The operation of casting into and out of raw types is a kind of “bulletproofing.” Bulletproofing maintains the raw-type invariant: the positive cast takes the argument out of $[c]$ and the negative cast returns it there. For example,

$$\langle \{x:\text{Int} \mid \text{nonzero } x\} \mapsto \{y:\text{Int} \mid \text{pos } y\} \rangle^{l,l'}$$

translates to the λ_H term

$$\begin{aligned} &\lambda f: [\text{Int} \rightarrow \text{Int}]. \\ &\langle \{x:\text{Int} \mid \text{nonzero } x\} \rightarrow \{y:\text{Int} \mid \text{pos } y\} \Rightarrow [\text{Int} \rightarrow \text{Int}] \rangle^{l'} \\ &(\langle [\text{Int} \rightarrow \text{Int}] \Rightarrow \{x:\text{Int} \mid \text{nonzero } x\} \rightarrow \{y:\text{Int} \mid \text{pos } y\} \rangle^l f). \end{aligned}$$

Unfolding the domain parts of the casts on f , the domain of the negative cast ensures that f ’s argument is nonzero with $\langle [\text{Int}] \Rightarrow \{x:\text{Int} \mid \text{nonzero } x\} \rangle^{l'}$; the domain of the positive cast does nothing, since $\langle \{x:\text{Int} \mid \text{nonzero } x\} \Rightarrow [\text{Int}] \rangle^l$ has no effect. Similarly, the codomain of the negative cast does nothing while the codomain of the positive cast checks that the result is positive. Separating the checks allows λ_H to keep track of blame labels, mimicking λ_C . Put more generally, in the positive cast, the positive positions may fail because they are “down casts”, whereas the negative positions are “up casts”, so they cannot fail. The opposite is true of the negative cast. This embodies the idea of contracts as pairs of projections [25]. Note that bulletproofing is “overkill” at base type: for example, $\langle \{x:\text{Int} \mid \text{nonzero } x\} \rangle^{l,l'}$ translates to

$$\begin{aligned} &\lambda x: [\text{Int}]. \\ &\langle \{x:\text{Int} \mid \text{nonzero } x\} \Rightarrow [\text{Int}] \rangle^{l'} \\ &(\langle [\text{Int}] \Rightarrow \{x:\text{Int} \mid \text{nonzero } x\} \rangle^l x). \end{aligned}$$

Only the positive cast does anything—the negative cast into $\llbracket \text{Int} \rrbracket$ always succeeds. This asymmetry is consistent with λ_{C} , where predicate contracts also ignore the negative label. In Section 2.3 I extend the bulletproofing translation to dependent contracts—one of my main contributions.

Both ϕ and ψ preserve behavior in a strong sense: if $\Gamma \vdash t : B$, then either t and $\phi(t)$ both evaluate to the same constant k or they both raise $\uparrow l$ for the same l ; and conversely for ψ . Interestingly, we need to set up this behavioral correspondence *before* we can prove that the translations preserve well-typedness, because of the T_CHECKING and S_CHECKING rules.

2.3 The dependent languages

We now extend λ_{C} to dependent function contracts and λ_{H} to dependent functions. Very little needs to be changed in λ_{C} , since contracts and types barely interact; the changes to E_CDECOMP and T_FUNC are the important ones. Adding dependency to λ_{H} is more involved. In particular, adding contexts to the subtyping judgment entails adding contexts to SSUB_REFINE. To avoid a dangerous circularity, I define closing substitutions in terms of a separate type semantics. Additionally, the new F_CDECOMP rule has a slightly tricky (but necessary) asymmetry, explained below.

2.3.1 Dependent λ_{C}

Dependent λ_{C} has been studied since Findler and Felleisen [26]; it received a very thorough treatment (with an untyped host language) in Blume and McAllester [11], was ported to Haskell by Hinze et al. [41] and Chitil and Huch [16], and was used as a specification language in Xu et al. [80]. Type soundness is not particularly difficult, since types and contracts are kept separate. My formulation follows Findler and Felleisen [26], with a few technical changes to make the proofs for ϕ easier.

We have marked the changed rules with a \bullet next to their names. The new T_REFINEC, T_FUNC, and E_CDECOMP rules in Figure 2.7 (typing) and Figure 2.6 (syntax and semantics) suffice to add dependency to λ_{C} . To help us work with the translations, we also make some small changes to the bindings in contexts, adding a new binding form to track the labels on a contract check throughout the contract well-formedness judgment. Note that T_FUNC adds $x:c_1^{l',l}$ to the context when checking the codomain of a function contract, swapping blame labels. I add a new variable rule, T_VARC, that treats $x:c^{l'}$ as if it were its skeleton, $x:[c]$. While unnecessary for λ_{C} 's metatheory, this new binding form helps ϕ preserve types when translating from λ_{H} to picky λ_{C} ; see Section 2.6.1.

Two different variants of the E_CDECOMP rule can be found in the literature: they are *lax* and *picky*. The original rule in Findler and Felleisen [26] is *lax* (like most other contract calculi): it does not recheck c_1 when substituting v' into c_2 . Blume and McAllester [11] used a *picky* semantics without observing their departure

Syntax for dependent λ_C

$T ::= B \mid T_1 \rightarrow T_2$	types
$c ::= \{x:B \mid t\} \mid x:c_1 \mapsto c_2$	contracts•
$\Gamma ::= \emptyset \mid \Gamma, x:T \mid \Gamma, x:c^{l,l'}$	typing contexts•
$t ::= x \mid k \mid \lambda x:T_1. t_2 \mid t_1 t_2 \mid \uparrow l \mid \langle c \rangle^{l,l'} \mid \langle \{x:B \mid t_1\}, t_2, k \rangle^l$	terms
$v ::= k \mid \lambda x:T_1. t_2 \mid \langle c \rangle^{l,l'} \mid \langle x:c_1 \mapsto c_2 \rangle^{l,l'} v$	values•
$r ::= v \mid \uparrow l$	results
$E ::= [] t \mid v [] \mid \langle \{x:B \mid t\}, [], k \rangle^l$	evaluation contexts

Operational semantics for dependent λ_C $t_1 \rightarrow_c t_2$

$(\lambda x:T_1. t_2) v \rightarrow_c t_2\{x := v\}$	E_BETA
$k v \rightarrow_c \llbracket k \rrbracket(v)$	E_CONST
$\langle \{x:B \mid t\} \rangle^{l,l'} k \rightarrow_c \langle \{x:B \mid t\}, t\{x := k\}, k \rangle^l$	E_CCHECK
$\langle \langle x:c_1 \mapsto c_2 \rangle^{l,l'} v \rangle v' \rightarrow_{lax} \langle c_2\{x := v'\} \rangle^{l,l'} (v \langle \langle c_1 \rangle^{l,l'} v' \rangle)$	E_CDECOMPLAX•
$\langle \langle x:c_1 \mapsto c_2 \rangle^{l,l'} v \rangle v' \rightarrow_{picky} \langle c_2\{x := \langle c_1 \rangle^{l,l'} v'\} \rangle^{l,l'} (v \langle \langle c_1 \rangle^{l,l'} v' \rangle)$	E_CDECOMPICKY•
$\langle \{x:B \mid t\}, \text{true}, k \rangle^l \rightarrow_c k$	E_OK
$\langle \{x:B \mid t\}, \text{false}, k \rangle^l \rightarrow_c \uparrow l$	E_FAIL
$E[\uparrow l] \rightarrow_c \uparrow l$	E_BLAME
$E[t_1] \rightarrow_c E[t_2] \text{ when } t_1 \rightarrow_c t_2$	E_COMPAT

Contract erasure

$$[\{x:B \mid t\}] = B \qquad [x:c_1 \mapsto c_2] = [c_1] \rightarrow [c_2]$$

Figure 2.6: Syntax and semantics for dependent λ_C

Typing rules for dependent λ_C

$$\begin{array}{c}
\boxed{\vdash \Gamma} \\
\\
\frac{}{\vdash \emptyset} \text{ T_EMPTY} \qquad \frac{\vdash \Gamma}{\vdash \Gamma, x:T} \text{ T_EXTVART}\bullet \qquad \frac{\vdash \Gamma}{\vdash \Gamma, x:T} \text{ T_EXTVART}\bullet \\
\qquad \qquad \qquad \boxed{\Gamma \vdash t : T} \\
\\
\frac{x:T \in \Gamma}{\Gamma \vdash x : T} \text{ T_VART} \qquad \frac{x:c^{l,l'} \in \Gamma}{\Gamma \vdash x : [c]} \text{ T_VARC}\bullet \qquad \frac{}{\Gamma \vdash k : \text{ty}_c(k)} \text{ T_CONST} \\
\\
\frac{\Gamma, x:T_1 \vdash t_2 : T_2}{\Gamma \vdash \lambda x:T_1. t_2 : T_1 \rightarrow T_2} \text{ T_LAM} \qquad \frac{\Gamma \vdash t_1 : T_1 \rightarrow T_2 \quad \Gamma \vdash t_2 : T_1}{\Gamma \vdash t_1 t_2 : T_2} \text{ T_APP} \\
\\
\frac{\Gamma \vdash^{l,l'} c : T}{\Gamma \vdash \langle c \rangle^{l,l'} : T \rightarrow T} \text{ T_CONTRACT}\bullet \qquad \frac{}{\Gamma \vdash \uparrow l : T} \text{ T_BLAME} \\
\\
\frac{\vdash \Gamma \quad \emptyset \vdash k : B \quad \emptyset \vdash t_2 : \text{Bool} \quad \emptyset \vdash^{l,l'} \{x:B \mid t_1\} : B \quad t_2 \rightarrow_c^* \text{true} \text{ implies } t_1 \{x := k\} \rightarrow_c^* \text{true}}{\Gamma \vdash \langle \{x:B \mid t_1\}, t_2, k \rangle^l : B} \text{ T_CHECKING} \\
\\
\frac{\Gamma, x:B \vdash t : \text{Bool}}{\Gamma \vdash^{l,l'} \{x:B \mid t\} : B} \text{ T_REFINEC} \qquad \frac{\boxed{\Gamma \vdash^{l,l'} c : T} \quad \Gamma \vdash^{l',l} c_1 : T_1 \quad \Gamma, x:c_1^{l',l} \vdash^{l,l'} c_2 : T_2}{\Gamma \vdash^{l,l'} x:c_1 \mapsto c_2 : T_1 \rightarrow T_2} \text{ T_FUNC}
\end{array}$$

Figure 2.7: Typing rules for dependent λ_C

from Findler and Felleisen; Hinze et al. [41] choose to be picky as well, substituting $\langle c_1 \rangle^{l',l} v'$ into c_2 because it makes their conjunction contract idempotent. We can show (straightforwardly) that both enjoy standard progress and preservation properties. Below, we consider translations to and from both dialects of λ_C : picky λ_C using only `E_CDECOMPICKY` in Sections 2.5.1 and 2.6.2, and lax λ_C using only `E_CDECOMPLAX` in Sections 2.5.2 and 2.6.1. Accordingly, I give two sets of evaluation rules: \longrightarrow_{lax} and \longrightarrow_{picky} . When I write \longrightarrow_c , the metavariable c ranges over *picky* and *lax*. I complete the type soundness proofs here generically, writing \longrightarrow_c for the evaluation relation. For the translations in Section 2.4, I specify which evaluation relation we use.

I make a standard assumption about constant denotations being well typed: if $\Gamma \vdash k \ v : T$ then $\Gamma \vdash \llbracket k \rrbracket(v) : T$.

2.3.1 Theorem [Progress]: If $\emptyset \vdash t : T$ then either $t \longrightarrow_c t'$ or $t = r$ (i.e., $t = v$ or $t = \uparrow l$).

Proof: By induction on the typing derivation. □

For preservation, we prove confluence and substitution lemmas. Note that our substitution lemma must now also cover contracts, since they are no longer closed.

2.3.2 Lemma [Determinacy]: Let \longrightarrow_c be either \longrightarrow_{picky} or \longrightarrow_{lax} . If $t \longrightarrow_c t'$ and $t \longrightarrow_c t''$, then $t' = t''$.

2.3.3 Corollary [Cotermination]: If Let \longrightarrow_c be either \longrightarrow_{picky} or \longrightarrow_{lax} . $t \longrightarrow_c^* r$ and $t \longrightarrow_c^* t'$, then $t' \longrightarrow_c^* r$.

2.3.4 Lemma [Term and contract substitution]: If $\emptyset \vdash v : T'$, then

1. if $\Gamma, x:T', \Gamma' \vdash t : T$ then $\Gamma, \Gamma'\{x := v\} \vdash t\{x := v\} : T$, and
2. if $\Gamma, x:T', \Gamma' \vdash^{l,l'} c : T$ then $\Gamma, \Gamma'\{x := v\} \vdash^{l,l'} c\{x := v\} : T$.

Proof: By mutual induction on the typing derivations for t and c . □

We omit the proof for $x:c^{l,l'}$ bindings, which is similar.

2.3.5 Theorem [Preservation]: If $\emptyset \vdash t : T$ and $t \longrightarrow_c t'$ then $\emptyset \vdash t' : T$.

Proof: By induction on the typing derivation. This proof is straightforward because typing and contracts hardly interact. □

Syntax for dependent λ_H

$S ::= \{x:B \mid s_1\} \mid x:S_1 \rightarrow S_2$	types/contracts•
$\Delta ::= \emptyset \mid \Delta, x:S$	typing contexts
$s ::= x \mid k \mid \lambda x:S_1. s_2 \mid s_1 s_2 \mid$	terms
$\uparrow l \mid \langle S_1 \Rightarrow S_2 \rangle^l \mid \langle \{x:B \mid s_1\}, s_2, k \rangle^l$	
$w ::= k \mid \lambda x:S_1. s_2 \mid \langle S_1 \Rightarrow S_2 \rangle^l \mid \langle x:S_{11} \rightarrow S_{12} \Rightarrow x:S_{21} \rightarrow S_{22} \rangle^l w$	values•
$q ::= w \mid \uparrow l$	results
$F ::= [] s \mid w [] \mid \langle \{x:B \mid s\}, [], k \rangle^l$	evaluation contexts

Operational semantics for dependent λ_H $s_1 \rightsquigarrow_h s_2$

$(\lambda x:S_1. s_2) w_2 \rightsquigarrow_h s_2\{x := w_2\}$	F_BETA
$k w \rightsquigarrow_h \llbracket k \rrbracket(w)$	F_CONST
$\langle \{x:B \mid s_1\} \Rightarrow \{x:B \mid s_2\} \rangle^l k \rightsquigarrow_h \langle \{x:B \mid s_2\}, s_2\{x := k\}, k \rangle^l$	F_CCHECK
$\langle \langle x:S_{11} \rightarrow S_{12} \Rightarrow x:S_{21} \rightarrow S_{22} \rangle^l w \rangle w' \rightsquigarrow_h \langle S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w'\} \Rightarrow S_{22}\{x := w'\} \rangle^l (w (\langle S_{21} \Rightarrow S_{11} \rangle^l w'))$	F_CDECOMP•
$\langle \{x:B \mid s\}, \text{true}, k \rangle^l \rightsquigarrow_h k$	F_OK
$\langle \{x:B \mid s\}, \text{false}, k \rangle^l \rightsquigarrow_h \uparrow l$	F_FAIL
$s_1 \longrightarrow_h s_2$	

$\frac{s_1 \rightsquigarrow_h s_2}{s_1 \longrightarrow_h s_2}$	F_REDUCE	$\frac{s_1 \longrightarrow_h s_2}{F[s_1] \longrightarrow_h F[s_2]}$	F_COMPAT	$\frac{}{F[\uparrow l] \longrightarrow_h \uparrow l}$	F_BLAZE
--	----------	---	----------	---	---------

Figure 2.8: Syntax and operational semantics for dependent λ_H

2.3.2 Dependent λ_H

Now we come to the challenging part: dependent λ_H and its proof of type soundness. These results require the most complex metatheory in this chapter—in fact, in the entirety of this dissertation—because we need some strong properties about call-by-value evaluation.⁴ The full definitions are in Figures 2.8 and 2.9. As before, I have marked the changed rules with a • next to their names.

Now we can enrich the type system with dependent function types, $x:S_1 \rightarrow S_2$, where x may appear in S_2 . The S_CAST rule and the proofs need a notion of type erasure, $\lfloor S \rfloor$; type height $|S|$ will also be used in the proofs.

$$\begin{array}{ll}
 \lfloor - \rfloor & : S \rightarrow T \\
 \lfloor \{x:B \mid s\} \rfloor & = B \\
 \lfloor x:S_1 \rightarrow S_2 \rfloor & = \lfloor S_1 \rfloor \rightarrow \lfloor S_2 \rfloor
 \end{array}
 \qquad
 \begin{array}{ll}
 | - | & : S \rightarrow \mathbb{N} \\
 |\{x:B \mid s\}| & = 1 \\
 |x:S_1 \rightarrow S_2| & = 1 + |S_1| + |S_2|
 \end{array}$$

⁴The benefit of a CBV semantics is a better treatment of blame; I would avoid the marriage of lazy evaluation and exceptions found in Haskell, as I find it very confusing. By contrast, Knowles and Flanagan [44] cannot treat failed casts as exceptions because that would destroy confluence. They treat them as stuck terms. Readers familiar with the soundness proof of Knowles and Flanagan [44] will notice that my proof is significantly different from theirs. I discuss this in related work, Chapter 5.

A new dependent application rule, `S_APP`, substitutes the argument into the result type of the application. I generalize `WF_REFINE` to allow refinement-type predicates that use variables from the enclosing context. `WF_FUN` adds the bound variable to the context when checking the codomain of function types. In `SSUB_FUN`, subtyping for dependent function types remains contravariant, but I also add the argument variable to the context with the smaller type. This is similar to the function subtyping rule of $F_{<}$: [14].

We need to be careful when implementing higher-order dependent casts in the rule `F_CDECOMP`. As the cast decomposes, the variables in the codomain types of such a cast must be replaced. However, this substitution is asymmetric; on one side, we cast the argument and on the other we do not. This behavior is required for type soundness. For suppose we have $\Delta \vdash x:S_{11} \rightarrow S_{12}$ and $\Delta \vdash x:S_{21} \rightarrow S_{22}$ with equal skeletons, and values $\Delta \vdash w : (x:S_{11} \rightarrow S_{12})$ and $\Delta \vdash w' : S_{21}$. Then $\Delta \vdash (\langle x:S_{11} \rightarrow S_{12} \Rightarrow x:S_{21} \rightarrow S_{22} \rangle^l w) w' : S_{22}\{x := w'\}$. When we decompose the cast, we must make *some* substitution into S_{12} and S_{22} , but which? It is clear that we must substitute w' into S_{22} , since the original application has type $S_{22}\{x := w'\}$. Decomposing the cast will produce the inner application $\Delta \vdash w (\langle S_{21} \Rightarrow S_{11} \rangle^l w') : S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w'\}$; in order to apply the codomain cast to this term, we must substitute $\langle S_{21} \Rightarrow S_{11} \rangle^l w'$ into S_{12} . This calculation determines the form of `F_CDECOMP`.

While the operational semantics changes only in `F_CDECOMP`, we have split the evaluation relation into two parts: reductions \rightsquigarrow_h and steps \longrightarrow_h . This is a technical change that allows me to factor the proofs more cleanly (particularly for the parallel reduction proofs).

The final change generalizes `SSUB_REFINE` to open terms. We must close these terms before we can compare their behavior, using closing substitutions σ and reading $\Delta \models \sigma$ as “ σ satisfies Δ ”.

Care is needed here to prevent the typing rules from becoming circular: the typing rule `S_SUB` references the subtyping judgment, the subtyping rule `SSUB_REFINE` references the implication judgment, and the single implication rule `S_IMP` has $\Delta \models \sigma$ in a negative position. This circularity would cause the typing rules to be non-monotonic, and so the existence of a least or greatest fixed-point would not be immediately obvious—the type system would not be well defined! To avoid this circularity, $\Delta \models \sigma$ must not refer back to the other judgments. (The reader may wonder why this was not a problem in λ_C , but notice that in λ_C , implication is only used in `T_CHECKING`—which has no (real) context. If we only needed implication in the `S_CHECKING` rule, we would not need contexts here, either—we can ensure that active checks only occur at the top-level, with an empty context. But the `SSUB_REFINE` subtyping rule refers to `S_IMP`, and subtyping may be used in arbitrary contexts.)

We can avoid the circularity and ensure that the type system is well defined by building the syntactic rules on top of a denotational semantics for λ_H 's types.⁵ The

⁵Knowles and Flanagan [44] also introduce a type semantics, but theirs differs from mine in two

Typing rules

$$\begin{array}{c}
\boxed{\vdash \Delta} \\
\frac{}{\vdash \emptyset} \text{S_EMPTY} \qquad \frac{\vdash \Delta \quad \Delta \vdash S}{\vdash \Delta, x:S} \text{S_EXTVAR} \\
\boxed{\Delta \vdash s : S} \\
\frac{x:S \in \Delta}{\Delta \vdash x : S} \text{S_VAR} \qquad \frac{\Delta \vdash S_1 \quad \Delta, x:S_1 \vdash s_2 : S_2}{\Delta \vdash \lambda x:S_1. s_2 : (x:S_1 \rightarrow S_2)} \text{S_LAM}\bullet \\
\frac{}{\Delta \vdash k : \text{ty}_h(k)} \text{S_CONST} \qquad \frac{\Delta \vdash s_1 : (x:S_1 \rightarrow S_2) \quad \Delta \vdash s_2 : S_1}{\Delta \vdash s_1 s_2 : S_2\{x := s_2\}} \text{S_APP}\bullet \\
\frac{\Delta \vdash S_1 \quad \Delta \vdash S_2 \quad [S_1] = [S_2]}{\Delta \vdash \langle S_1 \Rightarrow S_2 \rangle^l : S_1 \rightarrow S_2} \text{S_CAST} \qquad \frac{\Delta \vdash s : S_1 \quad \Delta \vdash S_2}{\Delta \vdash S_1 <: S_2} \text{S_SUB}\bullet \\
\frac{\vdash \Delta \quad \emptyset \vdash k : \{x:B \mid \text{true}\} \quad \emptyset \vdash s_2 : \{x:\text{Bool} \mid \text{true}\} \quad \emptyset \vdash \{x:B \mid s_1\} \quad \emptyset \vdash s_2 \supset s_1\{x := k\}}{\Delta \vdash \langle \{x:B \mid s_1\}, s_2, k \rangle^l : \{x:B \mid s_1\}} \text{S_CHECKING}\bullet \\
\boxed{\Delta \vdash S} \\
\frac{}{\Delta \vdash \{x:B \mid \text{true}\}} \text{SWF_RAW} \qquad \frac{\Delta \vdash S_1 \quad \Delta, x:S_1 \vdash S_2}{\Delta \vdash x:S_1 \rightarrow S_2} \text{SWF_FUN} \\
\frac{\Delta, x:\{x:B \mid \text{true}\} \vdash s : \{x:\text{Bool} \mid \text{true}\}}{\Delta \vdash \{x:B \mid s\}} \text{SWF_REFINE}\bullet \\
\boxed{\Delta \vdash S_1 <: S_2} \\
\frac{\Delta, x:\{x:B \mid \text{true}\} \vdash s_1 \supset s_2}{\Delta \vdash \{x:B \mid s_1\} <: \{x:B \mid s_2\}} \text{SSUB_REFINE}\bullet \\
\frac{\Delta \vdash S_{21} <: S_{11} \quad \Delta, x:S_{21} \vdash S_{12} <: S_{22}}{\Delta \vdash x:S_{11} \rightarrow S_{12} <: x:S_{21} \rightarrow S_{22}} \text{SSUB_FUN}\bullet \\
\boxed{\Delta \vdash s_1 \supset s_2} \\
\frac{\forall \sigma. (\Delta \models \sigma \wedge \sigma(s_1) \xrightarrow*_h \text{true}) \text{ implies } \sigma(s_2) \xrightarrow*_h \text{true}}{\Delta \vdash s_1 \supset s_2} \text{S_IMP}\bullet
\end{array}$$

$$\Delta \models \sigma \iff \forall x \in \text{dom}(\Delta). \sigma(x) \in \llbracket \sigma(\Delta(x)) \rrbracket$$

Figure 2.9: Typing rules for dependent λ_H

Denotations of types

$$\begin{aligned} s \in \llbracket \{x:B \mid s_0\} \rrbracket &\iff s \longrightarrow_h^* \uparrow l \vee (\exists k \in \mathcal{K}_B. s \longrightarrow_h^* k \wedge s_0\{x := k\} \longrightarrow_h^* \text{true}) \\ s \in \llbracket x:S_1 \rightarrow S_2 \rrbracket &\iff \forall q \in \llbracket S_1 \rrbracket. s \ q \in \llbracket S_2\{x := q\} \rrbracket \end{aligned}$$

Denotations of kinds

$$\begin{aligned} \{x:B \mid s\} \in \llbracket \star \rrbracket &\iff \forall k \in \mathcal{K}_B. s\{x := k\} \in \llbracket \{x:\text{Bool} \mid \text{true}\} \rrbracket \\ x:S_1 \rightarrow S_2 \in \llbracket \star \rrbracket &\iff S_1 \in \llbracket \star \rrbracket \wedge \forall q \in \llbracket S_1 \rrbracket. S_2\{x := q\} \in \llbracket \star \rrbracket \end{aligned}$$

Semantic judgments

$$\begin{aligned} \Delta \models S_1 <: S_2 &\iff \forall \sigma \text{ s.t. } \Delta \models \sigma, \llbracket \sigma(S_1) \rrbracket \subseteq \llbracket \sigma(S_2) \rrbracket \\ \Delta \models s : S &\iff \forall \sigma \text{ s.t. } \Delta \models \sigma, \sigma(s) \in \llbracket \sigma(S) \rrbracket \\ \Delta \models S &\iff \forall \sigma \text{ s.t. } \Delta \models \sigma, \sigma(S) \in \llbracket \star \rrbracket \end{aligned}$$

Figure 2.10: Type and kind semantics for dependent λ_H

idea is that the semantics of a type is a set of closed terms defined independently of the syntactic typing relation, but that turns out to contain all closed well-typed terms of that type. Thus, in the definition of $\Delta \models \sigma$, we quantify over a somewhat larger set than strictly necessary—not just the syntactically well-typed terms of appropriate type (which are all the ones that will ever appear in programs), but all semantically well-typed ones.

The type semantics appears in Figure 2.10. It is defined by induction on type skeletons. For refinement types, terms must either go to blame or produce a constant that satisfies (all instances of) the given predicate. For function types, well typed arguments must yield well-typed results. By construction, these sets include only terminating terms that do not get stuck. In order to show that casts inhabit the denotations of their types, we must also define a denotation of kinds. Since the only kind is \star , its denotation $\llbracket \star \rrbracket$ directly defines semantic well-formedness in terms of the denotations of types.

We must again make the assumption that constants have most-specific types: if $\llbracket \text{ty}_h(k) \rrbracket = B$ and $s\{x := k\} \longrightarrow_h^* \text{true}$ then $\emptyset \vdash \text{ty}_h(k) <: \{x:B \mid s\}$. We make some other, more standard assumptions, as well. Constants must have closed, well-formed types, and the types assigned must be well-formed. We require that constants are semantically well typed: $k \in \llbracket \text{ty}_h(k) \rrbracket$; this requirement is true by the “most-specific type” assumption at base types, but must be assumed at (first-

ways. First, because they cannot treat blame as an exception in their nondeterministic semantics, they must restrict the terms in the semantics to be those that only get stuck at failed casts. They do so by requiring the terms to be well-typed in the simply typed lambda calculus after all casts have been erased. Secondly, their type semantics does not require strong normalization. However, it is not clear whether their language actually admits nontermination—they include a `fix` constant, but their semantic type soundness proof appears to break down in that case. The problem is not insurmountable: either step indexing their semantics or a proof of unwinding as in Pitts [53] would resolve the issue.

order) functions types. Note that this rules out including `fix` as a constant, since my type semantics is inhabited only by strongly normalizing terms. We conjecture that expanding the denotation of refinement types to allow for divergence or a step-indexed logical relation [2] would allow us to consider nonterminating terms.

We introduce a few facts about the type semantics before proving semantic type soundness.

2.3.6 Lemma [Determinacy]: If $s \rightarrow_h s'$ and $s \rightarrow_h s''$, then $s' = s''$.

2.3.7 Corollary [Cotermination]: If $s \rightarrow_h^* s'$ and $s \rightarrow_h^* q$, then $s' \rightarrow_h^* q$.

2.3.8 Lemma [Expansion and contraction of $\llbracket S \rrbracket$]: If $s \rightarrow_h^* s'$, then $s' \in \llbracket S \rrbracket$ iff $s \in \llbracket S \rrbracket$.

Proof: By induction on $|S|$. □

2.3.9 Lemma [Blame inhabits all types]: For all S , $\uparrow l \in \llbracket S \rrbracket$.

Proof: By induction on $|S|$. □

2.3.10 Corollary [Nonemptiness]: For all S , there exists some q such that $q \in \llbracket S \rrbracket$.

The normal forms of \rightarrow_h^* are of the form $q = w$ or $\uparrow l$.

2.3.11 Lemma [Strong normalization]: If $s \in \llbracket S \rrbracket$, then there exists a q such that $s \rightarrow_h^* q$ —i.e., either $s \rightarrow_h^* w$ or $s \rightarrow_h^* \uparrow l$.

Proof: By induction on $|S|$.

$S = \{x:B \mid s_0\}$: Suppose $s \in \llbracket \{x:B \mid s_0\} \rrbracket$. By definition, either $s \rightarrow_h^* w$ or $s \rightarrow_h^* \uparrow l$, so s normalizes.

$S = x:S_1 \rightarrow S_2$: Suppose $s \in \llbracket x:S_1 \rightarrow S_2 \rrbracket$. We know that for any $q \in \llbracket S_1 \rrbracket$ that $s q \in \llbracket S_2\{x := q\} \rrbracket$. Since $\llbracket S_1 \rrbracket$ is nonempty (by Lemma 2.3.10), let $q \in \llbracket S_1 \rrbracket$. By the IH, $s q \rightarrow_h^* w$ or $s q \rightarrow_h^* \uparrow l$. By the definition of evaluation contexts and \rightarrow_h^* , the function position is evaluated first. If the application reduces to a value (i.e., $s q \rightarrow_h^* w$), then first $s q \rightarrow_h^* w' q$, and so $s \rightarrow_h^* w'$. Alternatively, the application could reduce to blame (i.e., $s q \rightarrow_h^* \uparrow l$). There are two ways for this to happen: either $s \rightarrow_h^* \uparrow l$, or $s \rightarrow_h^* w'$ and $q \rightarrow_h^* \uparrow l$. In both cases, s normalizes.

□ □

Unlike the rest of the chapter, I take a top-down approach to the rest of type soundness, to help motivate the steps. So: we are interested in relating the syntactic type system and the type semantics by *semantic type soundness*: if $\emptyset \vdash s : S$, then $s \in \llbracket S \rrbracket$. However, to prove this result, we must generalize it. In the bottom of Figure 2.10, I define three *semantic judgements* that correspond to each of the three typing judgments. (Note that the third one requires the definition of a *kind* semantics that picks out well-behaved types—those whose embedded terms belong to the type semantics.) We can then show that the typing judgments imply their semantic counterparts.

2.3.12 Theorem [Semantic type soundness]:

1. If $\Delta \vdash S_1 <: S_2$ then $\Delta \models S_1 <: S_2$.
2. If $\Delta \vdash s : S$ then $\Delta \models s : S$.
3. If $\Delta \vdash S$ then $\Delta \models S$.

Proof: Proof of (1) is in Lemma 2.3.14. Proof of (2) and (3) is in Lemma 2.3.21. \square

The first part follows by induction on the subtyping judgment.

2.3.13 Lemma [Trivial refinements of constants]: If $k \in \mathcal{K}_B$, then $k \in \llbracket \{x:B \mid \text{true}\} \rrbracket$.

2.3.14 Lemma [Semantic subtype soundness]: If $\Delta \vdash S_1 <: S_2$ then $\Delta \models S_1 <: S_2$.

Proof: By induction on the subtyping derivation.

SSUB_REFINE: We know $\Delta \vdash \{x:B \mid s_1\} <: \{x:B \mid s_2\}$, and must show the corresponding semantic subtyping. Inversion of this derivation gives us $\Delta, x:\{x:B \mid \text{true}\} \vdash s_1 \supset s_2$, which means:

$$\forall \sigma. ((\Delta, x:\{x:B \mid \text{true}\} \models \sigma \wedge \sigma(s_1) \longrightarrow_h^* \text{true}) \text{ implies } \sigma(s_2) \longrightarrow_h^* \text{true}) \quad (*)$$

We must show $\Delta \models \{x:B \mid s_1\} <: \{x:B \mid s_2\}$, i.e., that $\forall \sigma. \Delta \models \sigma$ implies $\llbracket \{x:B \mid s_1\} \rrbracket \subseteq \llbracket \{x:B \mid s_2\} \rrbracket$. Let σ be given such that $\Delta \models \sigma$. Suppose $s \in \llbracket \sigma(\{x:B \mid s_1\}) \rrbracket$. By definition, either s goes to $\uparrow l$, or it goes to $k \in \mathcal{K}_B$ such that $s_1\{x := k\} \longrightarrow_h^* \text{true}$. In the former case, $\uparrow l \in \llbracket \{x:B \mid s_2\} \rrbracket$ by definition. So consider the latter case, where $s \longrightarrow_h^* k$.

We already know that $k \in \mathcal{K}_B$, so it remains to see that: $\sigma(s_2)\{x := k\} \longrightarrow_h^* \text{true}$. We know by assumption that $\sigma(s_1)\{x := k\} \longrightarrow_h^* \text{true}$. By Lemma 2.3.13, $k \in \llbracket \{x:B \mid \text{true}\} \rrbracket$.

Now observe that $\Delta, x:\{x:B \mid \text{true}\} \models \sigma\{x := k\}$. Since $\sigma'(s_1) \longrightarrow_h^* \text{true}$, we can conclude that $\sigma'(s_2) \longrightarrow_h^* \text{true}$ by our assumption (*). This completes this case.

SSUB_FUN: $\Delta \vdash (x:S_{11} \rightarrow S_{12}) <: (x:S_{21} \rightarrow S_{22})$; by the IH, we have $\Delta \models S_{21} <: S_{11}$ and $\Delta, x:S_{21} \models S_{12} <: S_{22}$. We must show that $\Delta \models (x:S_{11} \rightarrow S_{12}) <: (x:S_{21} \rightarrow S_{22})$.

Let $\Delta \models \sigma$ and $s \in \llbracket \sigma(x:S_{11} \rightarrow S_{12}) \rrbracket$, for some σ . We must show, for all q , that if $q \in \llbracket \sigma(S_{21}) \rrbracket$ then $s \ q \in \llbracket \sigma(S_{22})\{x := q\} \rrbracket$.

Let $q \in \llbracket \sigma(S_{21}) \rrbracket$. Then $q \in \llbracket \sigma(S_{11}) \rrbracket$. Since $s \in \llbracket \sigma(x:S_{11} \rightarrow S_{12}) \rrbracket$, we know that $s \ q \in \llbracket \sigma(S_{12})\{x := q\} \rrbracket$. Finally, since $\Delta, x:S_{21} \models S_{12} <: S_{22}$ and $\Delta, x:S_{21} \models \sigma\{x := q\}$, we can conclude that $s \ q \in \llbracket \sigma(S_{22})\{x := q\} \rrbracket$, and so $s \in \llbracket \sigma(x:S_{21} \rightarrow S_{22}) \rrbracket$. \square

The proof semantic subtype soundness goes through easily, the first of the three parts of semantic soundness (Theorem 2.3.12). We run into some complications with semantic type and kind soundness, the second and third parts (which must be proven together). The crux of the difficulty lies with the S_APP rule. Suppose the application $s_1 \ s_2$ was well typed and $s_1 \in \llbracket x:S_1 \rightarrow S_2 \rrbracket$ and $s_2 \in \llbracket S_1 \rrbracket$. According to S_APP, the application's type is $S_2\{x := s_2\}$. By the type semantics defined in Figure 2.10, if $s_1 \in \llbracket x:S_1 \rightarrow S_2 \rrbracket$, then $s_1 \ q \in \llbracket S_2\{x := q\} \rrbracket$ for any $q \in \llbracket S_1 \rrbracket$. Sadly, s_2 is not necessarily a result! We do know, however, that $s_2 \in \llbracket S_1 \rrbracket$, so $s_2 \rightarrow_h^* q_2$ by strong normalization (Lemma 2.3.11). We need to ask, then, how the type semantics of $S_2\{x := s_2\}$ and $S_2\{x := q_2\}$ relate. (One might think that we can solve this by changing the type semantics to quantify over terms, not results. But this just pushes the problem to the S_LAM case.)

We can show that the two type semantics are in fact equal using a parallel reduction technique. I define a parallel reduction relation \Rightarrow on terms and types in Figure 2.11 that allows redexes in different parts of a term (or type) to be reduced in the same step, and then I prove that types that parallel-reduce to each other—like $S_2\{x := s_2\}$ and $S_2\{x := q_2\}$ —have the same semantics. The definition of parallel reduction is standard, though we need to be careful to make it respect call-by-value reduction order: the β -redex $(\lambda x:S_1. s_1) \ s_2$ should not be contracted unless s_2 is a value, since doing so can change the order of effects. (Other redexes within s_1 and s_2 can safely reduce.)⁶ The proof requires a longish sequence of technical lemmas that essentially show that \Rightarrow commutes with \rightarrow_h^* . Since the proofs require fussy symbol manipulation, we have done these proofs in Coq. My development is available at http://www.cis.upenn.edu/~mgree/papers/lambdah_parred.tgz. We restate the critical results here.

Lemma [Substitution of parallel-reducing terms, Lemma A3 in thy.v]:

If $w \Rightarrow w'$, then

1. if $s \Rightarrow s'$ then $s\{x := w\} \Rightarrow s'\{x := w'\}$, and
2. if $S \Rightarrow S'$ then $S\{x := w\} \Rightarrow S'\{x := w'\}$.

Lemma [Parallel reduction implies coevaluation, Lemma A20 in thy.v]:

If $s_1 \Rightarrow s_2$ then $s_1 \rightarrow_h^* k$ iff $s_2 \rightarrow_h^* k$. Similarly, $s_1 \rightarrow_h^* \uparrow l$ iff $s_2 \rightarrow_h^* \uparrow l$.

⁶I conjecture that the reflexive transitive closure of a similar “CBV-respecting” variant of full β -reduction could be used in place of parallel reduction. It is not clear whether it would lead to shorter proofs.

An alternative strategy would be to use \Rightarrow in the typing rules and \longrightarrow_h in the operational semantics. This would simplify some of the metatheory, but it would complicate the specification of the language. Using \longrightarrow_h in the typing rules gives a clearer intuition and keeps the core system small.

2.3.17 Lemma [Single parallel reduction preserves type semantics]:

If $S_1 \Rightarrow S_2$ then $\llbracket S_1 \rrbracket = \llbracket S_2 \rrbracket$.

Proof: By induction on $|S_1|$ (which is equal to $|S_2|$), with a case analysis on the final rule used to show $S_1 \Rightarrow S_2$. \square

2.3.18 Corollary [Parallel reduction preserves type semantics]: If $S_1 \Rightarrow^* S_2$ then $\llbracket S_1 \rrbracket = \llbracket S_2 \rrbracket$.

2.3.19 Lemma [Partial semantic substitution]: If $\Delta_1, x:S', \Delta_2 \models s : S$, and $\Delta_1, x:S', \Delta_2 \models S$, and $\Delta_1 \models s' : S'$ then $\Delta_1, \Delta_2\{x := s'\} \models s\{x := s'\} : S\{x := s'\}$ and $\Delta_1, \Delta_2\{x := s'\} \models S\{x := s'\}$.

Proof: By the definition of $\Delta \models \sigma$. \square

The semantic typing case for casts requires a separate induction.

2.3.20 Lemma [Semantic typing for casts]: If $\Delta \models S_1$ and $\Delta \models S_2$ and $\llbracket S_1 \rrbracket = \llbracket S_2 \rrbracket$, then $\Delta \models \langle S_1 \Rightarrow S_2 \rangle^l : x:S_1 \rightarrow S_2$ for fresh x .

Proof: By induction on $|S_1| = |S_2|$, going by cases on the shape of S_2 . Let $\Delta \models \sigma$; we show that $\sigma(\langle S_1 \Rightarrow S_2 \rangle^l) \in \llbracket \sigma(S_1 \rightarrow S_2) \rrbracket$.

$S_2 = \{x:B \mid \sigma(s_2)\}$: Let $q \in \llbracket \sigma(S_1) \rrbracket$. If $q = \uparrow l'$, then the applied cast goes to $\uparrow l'$, and we are done by Lemma 2.3.9. So $q = k \in \mathcal{K}_B$. By F_CCHECK $\langle S_1 \Rightarrow \{x:B \mid \sigma(s_2)\} \rangle^l k \longrightarrow_h \langle \{x:B \mid \sigma(s_2)\}, \sigma(s_2)\{x := k\}, k \rangle^l$. By the well-kinding of S_2 , we know that $\sigma(s_2)\{x := k\} \in \llbracket \{x:\text{Bool} \mid \text{true}\} \rrbracket$, so by strong normalization (Lemma 2.3.11), the predicate in the active check goes to blame or to a value. If it goes to blame, we are done. If it goes to a value, that value must be **true** or **false**. If it goes to **false**, then the whole term goes to blame and we are done. If it goes to **true**, then the check will step to k . But $\sigma(s_2)\{x := k\} \longrightarrow_h^* \text{true}$, so $k \in \llbracket \sigma(\{x:B \mid s_2\}) \rrbracket$ by definition. Expansion (Lemma 2.3.8) completes the proof.

$S_2 = x:S_{21} \rightarrow S_{22}$: We must have $S_1 = x:S_{11} \rightarrow S_{12}$. Let $q \in \llbracket \sigma(S_1) \rrbracket$; if it is blame we are done by Lemma 2.3.9, so let it be a value w . Let $q' \in \llbracket \sigma(S_{21}) \rrbracket$; if it is blame we are done, so let it be a value w' . By F_CDECOMP:

$$\langle \sigma(S_{12})\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w'\} \rangle \Rightarrow \sigma(S_{22})\{x := w'\} \rangle^l (w (\langle \sigma(S_{21}) \Rightarrow \sigma(S_{11}) \rangle^l w'))$$

By the IH, $\langle \sigma(S_{21}) \Rightarrow \sigma(S_{11}) \rangle^l$ is semantically well-typed, so we have $\langle \sigma(S_{21}) \Rightarrow \sigma(S_{11}) \rangle^l w' \in \llbracket \sigma(S_{11}) \rrbracket$. By strong normalization (Lemma 2.3.11), this term reduces (and therefore parallel reduces, by Lemma A4) to some q'' .

$$\begin{array}{c}
\boxed{s_1 \Rightarrow s_2} \\
\frac{}{s \Rightarrow s} \text{ FP_REFL} \qquad \frac{w \Rightarrow w'}{k \ w \Rightarrow \llbracket k \rrbracket(w')} \text{ FP_RCONST} \\
\frac{s_{12} \Rightarrow s'_{12} \quad w_2 \Rightarrow w'_2}{(\lambda x:S. s_{12}) \ w_2 \Rightarrow s'_{12}\{x := w'_2\}} \text{ FP_RBETA} \\
\frac{s_2 \Rightarrow s'_2}{\langle \{x:B \mid s_1\} \Rightarrow \{x:B \mid s_2\} \rangle^l \ k \Rightarrow \langle \{x:B \mid s'_2\}, s'_2\{x := k\}, k \rangle^l} \text{ FP_RCHECK} \\
\frac{}{\langle \{x:B \mid s\}, \text{true}, k \rangle^l \Rightarrow k} \text{ FP_ROK} \qquad \frac{}{\langle \{x:B \mid s\}, \text{false}, k \rangle^l \Rightarrow \uparrow l} \text{ FP_RFAIL} \\
\frac{S_{11} \Rightarrow S'_{11} \quad S_{12} \Rightarrow S'_{12} \quad S_{21} \Rightarrow S'_{21} \quad S_{22} \Rightarrow S'_{22} \quad w_1 \Rightarrow w'_1 \quad w_2 \Rightarrow w'_2}{\langle \langle x:S_{11} \rightarrow S_{12} \Rightarrow x:S_{21} \rightarrow S_{22} \rangle^l \ w_1 \rangle \ w_2 \Rightarrow \langle S'_{12}\{x := \langle S'_{21} \Rightarrow S'_{11} \rangle^l \ w'_2\} \Rightarrow S'_{22}\{x := w'_2\} \rangle^l \ (w'_1 \ (\langle S'_{21} \Rightarrow S'_{11} \rangle^l \ w'_2))} \text{ FP_RCDECOMP} \\
\frac{S_1 \Rightarrow S'_1 \quad s_{12} \Rightarrow s'_{12}}{\lambda x:S_1. s_{12} \Rightarrow \lambda x:S'_1. s'_{12}} \text{ FP_LAM} \qquad \frac{s_1 \Rightarrow s'_1 \quad s_2 \Rightarrow s'_2}{s_1 \ s_2 \Rightarrow s'_1 \ s'_2} \text{ FP_APP} \\
\frac{S_1 \Rightarrow S'_1 \quad S_2 \Rightarrow S'_2}{\langle S_1 \Rightarrow S_2 \rangle^l \Rightarrow \langle S'_1 \Rightarrow S'_2 \rangle^l} \text{ FP_CAST} \\
\frac{S \Rightarrow S' \quad s \Rightarrow s'}{\langle S, s, k \rangle^l \Rightarrow \langle S', s', k \rangle^l} \text{ FP_CHECK} \qquad \frac{}{F \ [\uparrow l] \Rightarrow \uparrow l} \text{ FP_BLAME} \\
\boxed{S_1 \Rightarrow S_2} \\
\frac{}{\overline{S} \Rightarrow \overline{S}} \text{ FP_SREFL} \qquad \frac{s \Rightarrow s'}{\{x:B \mid s\} \Rightarrow \{x:B \mid s'\}} \text{ FP_SREFINE} \\
\frac{S_1 \Rightarrow S'_1 \quad S_2 \Rightarrow S'_2}{x:S_1 \rightarrow S_2 \Rightarrow x:S'_1 \rightarrow S'_2} \text{ FP_SFUN}
\end{array}$$

Figure 2.11: Parallel reduction for dependent λ_H

We know that $w \ q'' \in \llbracket \sigma(S_{12})\{x := q''\} \rrbracket$ by assumption. Using parallel reduction (Corollary 2.3.18), we have $\llbracket \sigma(S_{12})\{x := q''\} \rrbracket = \llbracket \sigma(S_{12})\{x := \langle \sigma(S_{21}) \Rightarrow \sigma(S_{11}) \rangle^l w'\} \rrbracket$.

Before applying the IH, we note that $\Delta \models S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w'\}$ and $\Delta \models S_{22}\{x := w'\}$ by Lemma 2.3.19. Then, by the IH we see that $\langle \sigma(S_{12})\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w'\} \Rightarrow \sigma(S_{22})\{x := w'\} \rangle^l$ is semantically well-kinded, so

$$\langle \sigma(S_{12})\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w'\} \Rightarrow \sigma(S_{22})\{x := w'\} \rangle^l (w \ w'') \in \llbracket \sigma(S_{22})\{x := w'\} \rrbracket \quad \square$$

□

2.3.21 Lemma [Semantic type soundness]:

1. If $\Delta \vdash s : S$ then $\Delta \models s : S$.
2. If $\Delta \vdash S$ then $\Delta \models S$.

Proof: By induction on the typing and well-formedness derivations, using Corollary 2.3.18 in the S_APP case and Lemma 2.3.14 in the S_SUB case. □

Theorem 2.3.12 gives us type soundness, and it combines with Lemma 2.3.11 for an even stronger result: well-typed programs always evaluate to values of appropriate (semantic) type.

While one can prove progress and preservation theorems, I omit them: we already have type soundness. Later proofs will require standard weakening and substitution lemmas, though, so we can prove them now.

2.3.22 Lemma [Weakening]: If $\Delta \vdash s : S$ and $\Delta \vdash S$, and $\text{dom}(\Delta) \cap \text{dom}(\Delta') = \emptyset$ with $\vdash \Delta, \Delta'$, then $\Delta, \Delta' \vdash s : S$ and $\Delta, \Delta' \vdash S$.

Proof: By straightforward induction on s and $|S|$; we reuse the (critical) context well-formedness derivation in the S_CHECKING case. □

The substitution lemma has one complication: the operational judgment S_IMP requires the semantic type soundness theorem to show that a syntactically well-typed term can be used in a closing substitution. It is otherwise straightforward.

2.3.23 Lemma [Substitution (implication)]: If $\Delta_1, x:S, \Delta_2 \vdash s_1 \supset s_2$ and $\Delta_1 \vdash s : S$, then $\Delta_1, \Delta_2\{x := s\} \vdash s_1\{x := s\} \supset s_2\{x := s\}$.

Proof: Direct, unfolding the closing substitutions. □

2.3.24 Lemma [Substitution (subtyping)]: If $\Delta_1, x:S, \Delta_2 \vdash S_1 <: S_2$ and $\Delta_1 \vdash s : S$, then $\Delta_1, \Delta_2\{x := s\} \vdash S_1\{x := s\} <: S_2\{x := s\}$.

Proof: By induction on the subtyping derivation. □

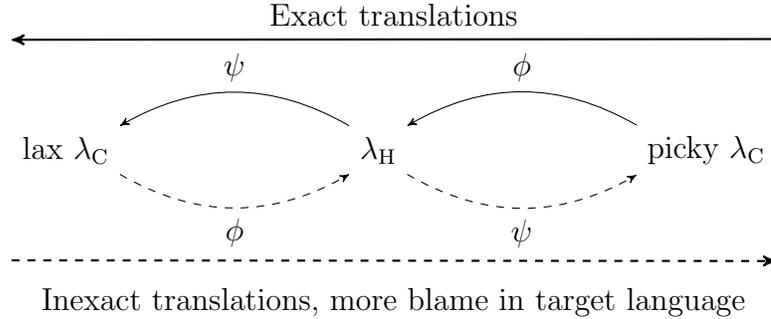
2.3.25 Lemma [Substitution (typing and well-formedness)]: If $\Delta_1 \vdash s : S$ then

1. if $\Delta_1, x:S, \Delta_2 \vdash s_1 : S_1$ then $\Delta_1, \Delta_2\{x := s\} \vdash s_1\{x := s\} : S_1\{x := s\}$,
2. if $\Delta_1, x:S, \Delta_2 \vdash S_1$ then $\Delta_1, \Delta_2\{x := s\} \vdash S_1\{x := s\}$, and
3. if $\vdash \Delta_1, x:S, \Delta_2$ then $\vdash \Delta_1, \Delta_2\{x := s\}$.

Proof: By mutual induction on the typing derivations. □

2.4 The translations

I divide treatment of the translations— ϕ from λ_H terms to λ_C terms and ψ from λ_C terms to λ_H terms—in two parts. I define the translations below. The axis of blame captures whether or not the translations exactly preserve semantics. The translations are *exact* when moving left on the axis of blame, but they are *inexact*—sometimes yielding more blame—moving right on the axis of blame.



I treat the exact translations in Section 2.5; the inexact translations are covered in Section 2.6.

Each translation proof follows the same basic schema. First, I define a logical relation between the two languages. Then we use the logical relation to prove a lemma relating the translation, contracts, and casts. Finally, we prove that the translation preserves evaluation behavior—that is, terms are logically related to their translations—and typing. All of the proofs make extensive use of expansion and contraction of evaluation and “cotermination” arguments. Every proof uses its own contract/cast logical relation. The proofs for the inexact translations of Section 2.6 demand custom term logical relations, too. I used σ to range over closing substitutions in λ_H ; I use δ to range over dual closing substitutions in the logical relations.

$$\begin{array}{l}
\text{Contexts} \quad \boxed{\phi : (\vdash \Gamma) \rightarrow \Delta} \\
\phi(\vdash \emptyset) = \emptyset \\
\phi(\vdash \Gamma, x:T) = \phi(\vdash \Gamma), x:[T] \\
\phi(\vdash \Gamma, x:c^{l,l'}) = \phi(\vdash \Gamma), x:\phi(\Gamma \vdash^{l,l'} c : [c]) \\
\\
\text{Terms} \quad \boxed{\phi : (\Gamma \vdash t : T) \rightarrow s} \\
\phi(\Gamma_1, x:T, \Gamma_2 \vdash x : T) = x \\
\phi(\Gamma_1, x:c^{l,l'}, \Gamma_2 \vdash x : [c]) = \langle \phi(\Gamma_1 \vdash^{l,l'} c : [c]) \Rightarrow [c] \rangle^{l'} x \\
\phi(\Gamma \vdash k : T) = k \\
\phi(\Gamma \vdash \lambda x:T_1. t_2 : T_1 \rightarrow T_2) = \lambda x:[T_1]. \phi(\Gamma, x:T_1 \vdash t_2 : T_2) \\
\phi(\Gamma \vdash t_1 t_2 : T_2) = \phi(\Gamma \vdash t_1 : T_1 \rightarrow T_2) \phi(\Gamma \vdash t_2 : T_1) \\
\phi(\Gamma \vdash \uparrow l : T) = \uparrow l \\
\phi(\emptyset \vdash \langle c, t, k \rangle^l : B) = \langle \phi(\emptyset \vdash^{l,l'} c : B), \phi(\emptyset \vdash t : \mathbf{Bool}), k \rangle^l \\
\phi(\Gamma \vdash \langle c \rangle^{l,l'} : T) = \lambda x:[c]. \langle \phi(\Gamma \vdash^{l,l} c : T) \Rightarrow [c] \rangle^{l'} \\
\qquad \qquad \qquad \langle \langle [c] \Rightarrow \phi(\Gamma \vdash^{l,l'} c : T) \rangle^l x \rangle \\
\qquad \qquad \qquad \text{where } x \text{ is fresh} \\
\\
\text{Types} \quad \boxed{\phi : (\Gamma \vdash^{l,l'} c : T) \rightarrow S} \\
\phi(\Gamma \vdash^{l,l'} \{x:B \mid t\} : B) = \{x:B \mid \phi(\Gamma, x:B \vdash t : \mathbf{Bool})\} \\
\phi(\Gamma \vdash^{l,l'} x:c_1 \mapsto c_2 : T_1 \rightarrow T_2) = x:\phi(\Gamma \vdash^{l,l} c_1 : T_1) \rightarrow \phi(\Gamma, x:c_1^{l',l} \vdash^{l,l'} c_2 : T_2)
\end{array}$$

Figure 2.12: The translation ϕ from dependent λ_C to dependent λ_H

2.4.1 Translating λ_C to λ_H : ϕ

We define the full ϕ for the dependent calculi in Figure 2.12. In the dependent case, we need to translate *derivations* of well-formedness and well-typing of λ_C contexts, terms, and contracts into λ_H contexts, terms, and types. I translate derivations to ensure type preservation, translating `T_VART` and `T_VARC` derivations differently: I leave variables of simple type alone, but I cast variables bound to contracts.

To see why we need this distinction, consider the function contract $f:(x:\{x:\mathbf{Int} \mid \mathbf{pos} \ x\} \mapsto \{y:\mathbf{Int} \mid \mathbf{true}\}) \mapsto \{z:\mathbf{Int} \mid f \ 0 = 0\}$. Note that this contract is well-formed in λ_C , but that the codomain “abuses” the bound variable. A naive translation will *not* be well-typed in λ_H . The term $f \ 0$ will not be typeable when f has type $x:\{x:\mathbf{Int} \mid \mathbf{pos} \ x\} \rightarrow [\mathbf{Int}]$, since f only accepts positive arguments. The problem is that `WF_FUN` can add a (possibly refined) type to the context when checking the codomain, so we need to restore the “variables have raw types” invariant—something we can’t always rely on subtyping to do, since types are not in general subtypes of their raw type. By tracking which variables were bound by contracts in λ_C , we can be sure to cast them to raw types when they’re referenced. We therefore translate the contract above to $f:S \rightarrow \{z:\mathbf{Int} \mid (\langle S \Rightarrow [\mathbf{Int} \rightarrow \mathbf{Int}] \rangle^{l'} f) \ 0 = 0\}$, where $S = x:\{x:\mathbf{Int} \mid \mathbf{pos} \ x\} \rightarrow [\mathbf{Int}]$. This (partially) motivates the $x:c^{l,l'}$ binding form in dependent λ_C .

Bulletproofing uses raw types, defined here for the dependent system.

$$\begin{aligned} \llbracket \{x:B \mid s\} \rrbracket &= \{x:B \mid \text{true}\} & \llbracket x:S_1 \rightarrow S_2 \rrbracket &= \llbracket S_1 \rrbracket \rightarrow \llbracket S_2 \rrbracket \\ \llbracket B \rrbracket &= \{x:B \mid \text{true}\} & \llbracket T_1 \rightarrow T_2 \rrbracket &= \llbracket T_1 \rrbracket \rightarrow \llbracket T_2 \rrbracket \\ \llbracket \{x:B \mid t\} \rrbracket &= \{x:B \mid \text{true}\} & \llbracket x:c_1 \mapsto c_2 \rrbracket &= \llbracket c_1 \rrbracket \rightarrow \llbracket c_2 \rrbracket \end{aligned}$$

Note that dependency is eliminated.

We could write the translation on terms instead of derivations, defining

$$\phi(x:c_1 \mapsto c_2) = x:\phi(c_1) \rightarrow \phi(c_2)\{x := \langle \phi(c_1) \Rightarrow \llbracket c_1 \rrbracket \rangle^l x\}$$

but the proofs are easier if we translate derivations.

Constants translate to themselves. One technical point: to maintain the raw type invariant, we need λ_H 's higher-order constants to have typings that can be seen as raw by the subtyping relation, i.e., $\Delta \vdash \text{ty}_h(k) <: \llbracket \text{ty}_c(k) \rrbracket$. This can be proven at base types (since we have already assumed that $\text{ty}_h(k)$ is the “most specific type” for each k), but must be assumed for first-order constant functions. This slightly restricts the types we might assign to constants, e.g., we cannot say $\text{ty}_h(\text{sqrt}) = x:\{x:\text{Float} \mid x \geq 0\} \rightarrow \{y:\text{Float} \mid (y * y) = x\}$, since it is not the case that $\Delta \vdash \text{ty}_h(\text{sqrt}) <: \llbracket \text{Float} \rightarrow \text{Float} \rrbracket$. Since its domain cannot be refined, $\llbracket \text{sqrt} \rrbracket$ must be defined for all $k \in \mathcal{K}_{\text{Float}}$, e.g., $\llbracket \text{sqrt} \rrbracket(-1)$ must be defined. We have already required that denotations be total over their simple types in λ_C , and λ_H uses the same denotation function $\llbracket - \rrbracket$, so this requirement does not seem too severe. In any case, we can define it to be equal to $\uparrow l_0$, for some l_0 . We could instead translate k to $\langle \text{ty}_h(k) \Rightarrow \llbracket \text{ty}_h(k) \rrbracket \rangle^{l_0} k$; however, in this case the nondependent fragments of the languages would no longer correspond exactly.

When translating a term with ϕ , the image behaves the same as the pre-image when we interpret the pre-image using picky λ_C semantics (proved in Section 2.5.1). When the pre-image is interpreted in the lax λ_C semantics, however, we find that the image can raise blame when the pre-image returns a value (proved in Section 2.6.1).

2.4.2 Translating λ_H to λ_C : ψ

In this section, I formally define ψ for the dependent versions of lax λ_C and λ_H . I prove that ψ is type preserving and induces behavioral correspondence.

The full definition of ψ is in Figure 2.13. Most terms are translated homomorphically. In abstractions, the annotation is translated by erasing the refined λ_H type to its skeleton. As we mentioned in Section 2.2, the trickiest part is the translation of casts between function types: when generating the codomain contract from a cast between two function types, we perform the same asymmetric substitution as F_CDECOMP . Since ψ inserts new casts, we need to pick a blame label: $\psi(\langle S_1 \Rightarrow S_2 \rangle^l)$ passes l as an index to $\psi^l(S_1, S_2)$.

I show that ψ preserves semantics exactly when its image is interpreted as lax λ_C in Section 2.5.2; I show that the picky λ_C interpretation of a term's ψ -image may blame more often than the original term in Section 2.6.2.

Term translation $\boxed{\psi : s \rightarrow t}$

$$\begin{array}{ll}
\psi(x) & = x & \psi(k) & = k \\
\psi(\lambda x:S. s) & = \lambda x:[S]. \psi(s) & \psi(s_1 s_2) & = \psi(s_1) \psi(s_2) \\
\psi(\langle S_1 \Rightarrow S_2 \rangle^l) & = \langle \psi^l(S_1, S_2) \rangle^{l,l} & \psi(\uparrow l) & = \uparrow l \\
\psi(\langle \{x:B \mid s_1\}, s_2, k \rangle^l) & = \langle \{x:B \mid \psi(s_1)\}, \psi(s_2), k \rangle^l & &
\end{array}$$

Cast translation $\boxed{\psi : S \times S \times l \rightarrow T}$

$$\begin{array}{ll}
\psi^l(\{x:B \mid s_1\}, \{x:B \mid s_2\}) & = \{x:B \mid \psi(s_2)\} \\
\psi^l(x:S_{11} \rightarrow S_{12}, x:S_{21} \rightarrow S_{22}) & = x:\psi^l(S_{21}, S_{11}) \mapsto \psi^l(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}, S_{22})
\end{array}$$

Figure 2.13: ψ mapping dependent λ_H to dependent λ_C

Result correspondence $r \approx q : T$

$$\begin{array}{ll}
k \approx k : B & \iff k \in \mathcal{K}_B \\
v \approx w : T_1 \rightarrow T_2 & \iff \forall t \sim s : T_1. v t \sim w s : T_2 \\
\uparrow l \approx \uparrow l : T &
\end{array}$$

Term correspondence $\boxed{t \sim s : T}$

$$t \sim s : T \iff t \longrightarrow_c^* r \wedge s \longrightarrow_h^* q \wedge r \approx q : T$$

Figure 2.14: A blame-exact result/term correspondence

2.5 Exact translations

Translations moving left on the axis of blame—from picky λ_C to λ_H , and from λ_H to lax λ_C —are exact. That is, we can show a tight behavioral correspondence between terms and their translations (see Figure 2.14). I read $t \sim s : T$ as “ t corresponds with s at type T ”.

The correspondence is a standard logical relation, defined in two intertwined parts: a relation on results, $r \approx q : T$ and its closure with respect to evaluation, $t \sim s : T$. The term correspondence is defined directly: terms correspond when they reduce to corresponding results. We write \longrightarrow_c in this single definition: in Section 2.5.1 I use this definition taking \longrightarrow_c to be \longrightarrow_{picky} ; in Section 2.5.2 I use this definition taking \longrightarrow_c to be \longrightarrow_{lax} . The result correspondence is defined inductively over λ_C ’s simple types. Blame corresponds to itself at any type. At B , constants in \mathcal{K}_B correspond to themselves; results at $T_1 \rightarrow T_2$ correspond when they applying them to corresponding *terms* yields corresponding *terms*. Stratifying the definition this way simplifies some of the proofs later. We call this correspondence *exact* because terms corresponding

Contract/type correspondence $c \sim^{l,l'} S : T$

$$\{x:B \mid t\} \sim^{l,l'} \{x:B \mid s\} : B \iff \forall k \in \mathcal{K}_B. t\{x := k\} \sim s\{x := k\} : \text{Bool}$$

$$x:c_1 \mapsto c_2 \sim^{l,l'} x:S_1 \rightarrow S_2 : T_1 \rightarrow T_2 \iff c_1 \sim^{l',l} S_1 : T_1 \wedge \forall t \sim s : T_1. c_2\{x := \langle c_1 \rangle^{l',l} t\} \sim^{l,l'} S_2\{x := \langle [S_1] \Rightarrow S_1 \rangle^{l'} s\} : T_2$$

Dual closing substitutions

$$\Gamma \models \delta \iff \begin{cases} \forall x:T \in \Gamma. \delta_1(x) \sim \delta_2(x) : T \\ \forall x:c^{l,l'} \in \Gamma. \delta_1(x) = \langle \delta_1(c) \rangle^{l,l'} t \wedge \delta_2(x) = \langle [c] \Rightarrow \delta_2(S) \rangle^{l'} s \\ \text{where } S = \phi(\Gamma \vdash^{l,l'} c : [c]) \wedge t \sim s : [c] \end{cases}$$

Lifted to open terms

$$\Gamma \vdash t \sim s : T \iff \forall \delta. (\Gamma \models \delta \text{ implies } \delta_1(t) \sim \delta_2(s) : T)$$

$$\Gamma \vdash c \sim^{l,l'} S : T \iff \forall \delta. (\Gamma \models \delta \text{ implies } \delta_1(c) \sim^{l,l'} \delta_2(S) : T)$$

Figure 2.15: Blame-exact correspondence for ϕ from picky λ_C

at base type yield identical results.

Note that I define the correspondence here on closed (or harmlessly open) terms. In the following two sections, I define translation specific extensions of the correspondence to open terms and contracts.

2.5.1 Translating picky λ_C to λ_H

We extend the term correspondence of Figure 2.14 to contracts and types, lifting the correspondences to open terms using dual closing substitutions. Recall that we interpret the term correspondence as using $\longrightarrow_{\text{picky}}$. For a binding $x:c^{l,l'} \in \Gamma$, we use ϕ to insert the negative cast (labelled with l') and closing substitutions (in Figure 2.15) to insert the positive cast (labelled with l). Do not be confused by the label used for function contract correspondence—this definition does, in fact, match up with closing substitutions. A binding $x:c^{l,l'} \in \Gamma$ must have come from the domain of an application of `T_FUNC`, so the labels on the binding are *already* swapped when ϕ or $\Gamma \models \delta$ sees them. In the definition of function contract correspondence, we swap manually—whence the l' on the inserted cast. It helps to think of polarity in terms of position rather than the presence or absence of a prime.

2.5.1 Lemma [Expansion and contraction]: If $t \longrightarrow_{\text{picky}}^* t'$, and $s \longrightarrow_h^* s'$ then $t \sim s : T$ iff $t' \sim s' : T$.

2.5.2 Lemma [Constants self-correspond]: For all k , $k \sim k : \text{ty}_c(k)$.

2.5.3 Lemma [Equivalence is closed under parallel reduction]: If $s \Rightarrow s'$, then $t \sim s : T$ iff $t \sim s' : T$. Similarly, if $S \Rightarrow S'$ then $c \sim^{l,l'} S : T$ iff $c \sim^{l,l'} S' : T$.

Proof: In both cases, by induction on T , using the first to prove the second. \square

2.5.4 Lemma [Trivial casts]: If $t \sim s : B$ and $\llbracket S \rrbracket = B$, then $t \sim \langle S \Rightarrow \llbracket B \rrbracket \rangle^l s : B$.

2.5.5 Lemma [Related base casts]: If $\{x:B \mid t\} \sim^{l_0, l_1} \{x:B \mid s\} : B$ and $t' \sim s' : B$ and $\llbracket S \rrbracket = B$, then $\langle \{x:B \mid t\} \rangle^{l, l'} t' \sim \langle S \Rightarrow \{x:B \mid s\} \rangle^l s' : B$.

Proof: Direct. Note that l_0 and l_1 are entirely irrelevant. \square

2.5.6 Lemma [Bulletproofing]: If $t \sim s : T$ and $c \sim^{l, l'} S : T$ then $\langle c \rangle^{l, l'} t \sim \langle S \Rightarrow \llbracket S \rrbracket \rangle^{l'} \langle \llbracket S \rrbracket \Rightarrow S \rangle^l s : T$.

Proof: By induction on T . First, observe that either both t and s go to $\uparrow l''$ or both t and s go to values related at T . In the former case, the outer terms also go to blame. So we only consider the case where $t \rightarrow_{\text{picky}^*} v$, $s \rightarrow_h^* w$, and $v \approx w : T$.

$T = B$: So $c = \{x:B \mid t_1\}$ and $S = \{x:B \mid s_1\}$ and $S' = \{x:B \mid s_2\}$. By Lemma 2.5.5 we have $\langle c \rangle^{l, l'} t \sim \langle \llbracket S \rrbracket \Rightarrow S \rangle^l s : B$. By Lemma 2.5.4 we can add the extra, trivial cast $\langle S \Rightarrow \llbracket S \rrbracket \rangle^{l'}$.

$T = T_1 \rightarrow T_2$: We know that $c = x:c_1 \mapsto c_2$ and $S = x:S_1 \rightarrow S_2$. Let $t' \sim s' : T_1$. We only need to consider the case where $t' \rightarrow_{\text{picky}^*} v'$ and $s' \rightarrow_h^* w'$ —if $t' \rightarrow_{\text{picky}^*} \uparrow l''$ and $s' \rightarrow_h^* \uparrow l''$ the outer terms correspond because both blame l'' .

On the λ_C side, $(\langle c \rangle^{l, l'} t) t' \rightarrow_{\text{picky}^*} \langle c_2 \{x := \langle c_1 \rangle^{l', l} v'\} \rangle^{l, l'} (v (\langle c_1 \rangle^{l', l} v'))$. In λ_H , we can see

$$\begin{aligned} & \langle \langle S \Rightarrow \llbracket S \rrbracket \rangle^{l'} \langle \llbracket S \rrbracket \Rightarrow S \rangle^l s \rangle s' \rightarrow_h^* \\ & \langle S_2 \{x := \langle \llbracket S_1 \rrbracket \Rightarrow S_1 \rangle^{l'} w'\} \Rightarrow \llbracket S_2 \rrbracket \rangle^{l'} ((\langle \llbracket S \rrbracket \Rightarrow S \rangle^l w) (\langle \llbracket S_1 \rrbracket \Rightarrow S_1 \rangle^{l'} w')) \end{aligned}$$

We cannot determine where the redex is until we know the shape of T_1 —does the negative argument cast step to an active check, or do we decompose the positive cast?

– $T_1 = B$. Since $v' \approx w' : B$, we must have $v' = w' = k \in \mathcal{K}_B$. By Lemma 2.5.5 and $c_1 \sim^{l', l} S_1 : B$, we know that $\langle c_1 \rangle^{l', l} v' \sim \langle \llbracket S_1 \rrbracket \Rightarrow S_1 \rangle^{l'} w' : B$. Both terms goes to blame or to the same value—which must be k , from inspection of the contract and cast evaluation rules.. The former case is immediate, since the outer terms then go to blame. So suppose $\langle c_1 \rangle^{l', l} k \rightarrow_{\text{picky}^*} k$ and $\langle \llbracket S_1 \rrbracket \Rightarrow S_1 \rangle^{l'} k \rightarrow_h^* k$. Now the terms evaluate like so:

$$\begin{aligned} & \langle c_2 \{x := \langle c_1 \rangle^{l', l} v'\} \rangle^{l, l'} (v (\langle c_1 \rangle^{l', l} v')) \rightarrow_{\text{picky}^*} \\ & \langle c_2 \{x := \langle c_1 \rangle^{l', l} k\} \rangle^{l, l'} (v k) \\ & \langle S_2 \{x := \langle \llbracket S_1 \rrbracket \Rightarrow S_1 \rangle^{l'} k\} \Rightarrow \llbracket S_2 \rrbracket \rangle^{l'} \\ & \quad ((\langle \llbracket S \rrbracket \Rightarrow S \rangle^l w) (\langle \llbracket S_1 \rrbracket \Rightarrow S_1 \rangle^{l'} k)) \rightarrow_h^* \\ & \langle S_2 \{x := \langle \llbracket S_1 \rrbracket \Rightarrow S_1 \rangle^{l'} k\} \Rightarrow \llbracket S_2 \rrbracket \rangle^{l'} \\ & \quad \langle \llbracket S_2 \rrbracket \Rightarrow S_2 \{x := k\} \rangle^l (w (\langle S_1 \Rightarrow \llbracket S_1 \rrbracket \rangle^l k)) \end{aligned}$$

By Lemma 2.5.4, $k \sim \langle S_1 \Rightarrow \lceil S_1 \rceil^l k : B$, so $v k \sim w (\langle S_1 \Rightarrow \lceil S_1 \rceil^l k) : T_2$.

We have by definition (and $k \sim k : B$) that $c_2\{x := \langle c_1 \rangle^{l,l} k\} \sim^{l,l'} S_2\{x := \langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} k\} : T_2$. Recall that $\langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} k \longrightarrow_h^* k$. This implies $\langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} k \Rightarrow^* k$ (Lemma A4 in the Coq). We can then see that $S_2\{x := \langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} k\} \Rightarrow^* S_2\{x := k\}$ by Lemma A1 in the Coq. By extension with the congruence rules:

$$\begin{aligned} & \langle S_2\{x := \langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} k\} \Rightarrow \lceil S_2 \rceil \rangle^{l'} \\ & \quad \langle \lceil S_2 \rceil \Rightarrow S_2\{x := \langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} k\} \rangle^l (w (\langle S_1 \Rightarrow \lceil S_1 \rceil^l k)) \Rightarrow \\ & \langle S_2\{x := \langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} k\} \Rightarrow \lceil S_2 \rceil \rangle^{l'} \\ & \quad \langle \lceil S_2 \rceil \Rightarrow S_2\{x := k\} \rangle^l (w (\langle S_1 \Rightarrow \lceil S_1 \rceil^l k)) \end{aligned}$$

By the IH $\langle c_2\{x := \langle c_1 \rangle^{l,l} k\} \rangle^{l,l'} (v k)$ corresponds to the former, which means it is related to the latter by Lemma 2.5.3. We conclude the case with expansion (Lemma 2.5.1).

– $T_1 = T_{11} \rightarrow T_{12}$. We continue with an application of F_CDECOMP in λ_H :

$$\begin{aligned} & \langle S_2\{x := \langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} w'\} \Rightarrow \lceil S_2 \rceil \rangle^{l'} \\ & \quad ((\langle \lceil S \rceil \Rightarrow S \rangle^l w) (\langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} w')) \longrightarrow_h^* \\ & \langle S_2\{x := \langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} w'\} \Rightarrow \lceil S_2 \rceil \rangle^{l'} \\ & \quad \langle \lceil S_2 \rceil \Rightarrow S_2\{x := \langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} w'\} \rangle^l \\ & \quad (w (\langle S_1 \Rightarrow \lceil S_1 \rceil^l (\langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} w'))) \end{aligned}$$

By the IH on $c_1 \sim^{l,l} S_1 : T_1$ and $v' \sim w' : T_1$, we can find what we need for the domain: $\langle c_1 \rangle^{l,l} v' \sim \langle S_1 \Rightarrow \lceil S_1 \rceil^l (\langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} w') : T_1$. By assumption, the results of applying v and w to these values correspond. (And they *are* values, since function contracts/casts applied to values are values.)

We have $c_2\{x := \langle c_1 \rangle^{l,l} v'\} \sim^{l,l'} S_2\{x := \langle \lceil S_1 \rceil \Rightarrow S_1 \rangle^{l'} w'\} : T_2$ by assumption, so the IH tells us that the codomain contract and bulletproofing correspond. We conclude by expansion (Lemma 2.5.1). $\square \square$

Having characterized how contracts and pairs of related casts relate, we show that terms correspond to their translation.

2.5.7 Theorem [Behavioral correspondence]: If $\vdash \Gamma$, then:

1. If $\phi(\Gamma \vdash t : T) = s$ then $\Gamma \vdash t \sim s : T$.
2. If $\phi(\Gamma \vdash^{l,l'} c : T) = S$ then $\Gamma \vdash c \sim^{l,l'} S : T$.

Proof: We simultaneously show both properties by induction on the depth of ϕ 's recursion. \square

We can now prove that ϕ preserves types, using Theorem 2.5.7 to show that ϕ preserves the implication judgment. As a preliminary, I use the behavioral correspondence to show that ϕ preserves the implication judgment.

2.5.8 Lemma: If $t_1 \longrightarrow_{picky}^* \text{true}$ implies $t_2 \longrightarrow_{picky}^* \text{true}$ then $\emptyset \vdash \phi(\emptyset \vdash t_1 : \text{Bool}) \supset \phi(\emptyset \vdash t_1 : \text{Bool})$.

Proof: By the logical relation. □

The type preservation proof is very similar to the correspondence proof of Theorem 2.5.7.

2.5.9 Theorem [Type preservation for ϕ]: If $\phi(\vdash \Gamma) = \Delta$, then:

1. $\vdash \Delta$.
2. If $\phi(\Gamma \vdash t : T) = s$ then $\Delta \vdash s : \lceil T \rceil$.
3. If $\phi(\Gamma \vdash^{l,l'} c : T) = S$ then $\Delta \vdash S$.

Proof: We prove all three properties simultaneously, by induction on the depth of ϕ 's recursion.

The proof is by cases on the λ_C context well-formedness/term typing/contract well-formedness derivations, which determine the branch of ϕ taken. □

2.5.2 Translating λ_H to lax λ_C

We reuse the term correspondence $t \sim s : T$ (Figure 2.14), interpreting it as using \longrightarrow_{lax} . and define a new contract/cast correspondence $c \sim S_1 \Rightarrow^l S_2 : T$ (Figure 2.16), relating contracts and pairs of λ_H types—effectively, casts. This correspondence uses the term correspondence in the base type case and follows the pattern of F_CDECOMP in the function case. Since it inserts a cast in the function case, I index the relation with a label, just like ψ . Note that the correspondence is blame-exact, relating λ_C and λ_H terms that either blame the same label or go to corresponding values. I define closing substitutions ignoring the contracts in the context; we lift the relation to open terms in the standard way.

We begin with some standard properties of the term correspondence relation.

2.5.10 Lemma [Expansion and contraction]: If $t \longrightarrow_{lax}^* t'$, and $s \longrightarrow_h^* s'$ then $t \sim s : T$ iff $t' \sim s' : T$.

2.5.11 Lemma [Blame corresponds to blame]: For all T , $\uparrow l \sim \uparrow l : T$.

2.5.12 Lemma [Constants self-correspond]: For all k , $k \approx k : \text{ty}_c(k)$.

As a corollary of Lemma 2.5.11 and Lemma 2.5.10, if two terms evaluate to blame, then they correspond. This will be used extensively in the proofs below, as it allows us to eliminate many cases.

Contract/type correspondence

$$\boxed{c \sim S_1 \Rightarrow^l S_2 : T}$$

$$\{x:B \mid t\} \sim \{x:B \mid s_1\} \Rightarrow^l \{x:B \mid s_2\} : B \iff \forall k \in \mathcal{K}_B. t\{x := k\} \sim s_2\{x := k\} : \mathbf{Bool}$$

$$x:c_1 \mapsto c_2 \sim x:S_{11} \rightarrow S_{12} \Rightarrow^l S_{21} \rightarrow S_{22} : T_1 \rightarrow T_2 \iff c_1 \sim S_{21} \Rightarrow^l S_{11} : T_1 \wedge \forall t \sim s : T_1. c_2\{x := t\} \sim S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l s\} \Rightarrow^l S_{22}\{x := s\} : T_2$$

Dual closing substitutions

$$\Gamma \models \delta \iff \begin{cases} \forall x:T \in \Gamma, & \delta_1(x) \sim \delta_2(x) : T \\ \forall x:c^{l,l'} \in \Gamma, & \delta_1(x) \sim \delta_2(x) : \lfloor c \rfloor \end{cases}$$

Lifted to open terms

$$\begin{aligned} \Gamma \vdash t \sim s : T &\iff \forall \delta. (\Gamma \models \delta \text{ implies } \delta_1(t) \sim \delta_2(s) : T) \\ \Gamma \vdash c \sim S_1 \Rightarrow^l S_2 : T &\iff \forall \delta. (\Gamma \models \delta \text{ implies } \delta_1(c) \sim \delta_2(S_1) \Rightarrow^l \delta_2(S_2) : T) \end{aligned}$$

Figure 2.16: Blame-exact correspondence for ψ into lax λ_C

2.5.13 Lemma [Corresponding terms coevaluate]: If $t \sim s : T$ then $t \longrightarrow_{lax}^* v \wedge s \longrightarrow_h^* w$ or $t \longrightarrow_{lax}^* \uparrow l \wedge s \longrightarrow_h^* \uparrow l'$; moreover, $t \longrightarrow_{lax}^* r$ and $s \longrightarrow_h^* q$ such that $r \approx q : T$.

2.5.14 Lemma [Contract/cast correspondence]: If $c \sim S_1 \Rightarrow^l S_2 : T$ and $t \sim s : T$ then $\langle c \rangle^{l,l} t \sim \langle S_1 \Rightarrow S_2 \rangle^l s : T$.

Proof: By induction on T . We reason via expansion (Lemma 2.5.10), showing that the initial terms reduce to corresponding terms.

$T = B$: So $c = \{x:B \mid t_1\}$, $S_1 = \{x:B \mid s_1\}$, and $S_2 = \{x:B \mid s_2\}$. Since $t \sim s : B$, we know that they either both reduce to $k \in \mathcal{K}_B$ or $\uparrow l'$. If the latter is the case, we are done. So suppose $t \longrightarrow_{lax}^* k$ along with $s \longrightarrow_h^* k$.

We can step our terms into active checks as follows, then:

$$\begin{aligned} \langle \{x:B \mid t_1\} \rangle^{l,l} t &\longrightarrow_{lax}^* \langle \{x:B \mid t_1\}, t_1\{x := k\}, k \rangle^l \\ \langle \{x:B \mid s_1\} \Rightarrow \{x:B \mid s_2\} \rangle^l s &\longrightarrow_h^* \langle \{x:B \mid s_2\}, s_2\{x := k\}, k \rangle^l \end{aligned}$$

By inversion of the contract/cast correspondence, we know that $t_1\{x := k\} \sim s_2\{x := k\} : \mathbf{Bool}$, so these terms go to blame or to a \mathbf{Bool} together. If they go to $\uparrow l'$, we are done. If they go to \mathbf{false} , then both the obligation and the cast will go to $\uparrow l$. Finally, if they both go to \mathbf{true} , then both terms will evaluate to k .

$T = T_1 \rightarrow T_2$: $c = x:c_1 \mapsto c_2$, $S_1 = x:S_{11} \rightarrow S_{12}$, and $S_2 = x:S_{21} \rightarrow S_{22}$. We know by inversion of the contract/cast relation that $c_1 \sim S_{21} \Rightarrow^l S_{11} : T_1$ and that for all $t \sim s : T_1$, $c_2\{x := t\} \sim S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l s\} \Rightarrow^l S_{22}\{x := s\} : T_2$. We

want to prove that $\langle c \rangle^{l,l} \sim \langle S_1 \Rightarrow S_2 \rangle^l s : T_1 \rightarrow T_2$. First, we can assume $t \rightarrow_{\text{Iax}}^* v$ and $s \rightarrow_h^* w$ where $v \sim w : T_1 \rightarrow T_2$ —if not, both cast and contracted terms go to blame and we are done.

We show that the decomposition of the contract and cast terms correspond for all inputs. Let $t' \sim s' : T_1$. Again, we can assume that they reduce to $v' \sim w' : T_1$, or else we are done by blame lifting. On the λ_C side, we have

$$(\langle c \rangle^{l,l} t) t' \rightarrow_{\text{Iax}}^* \langle c_2 \{x := v'\} \rangle^{l,l} (v (\langle c_1 \rangle^{l,l} v'))$$

In λ_H , we find

$$\begin{aligned} & (\langle S_1 \Rightarrow S_2 \rangle^l s) s' \rightarrow_h^* \\ & \langle S_{12} \{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w\} \Rightarrow S_{22} \{x := w'\} \rangle^l (w (\langle S_{21} \Rightarrow S_{11} \rangle^l w')) \end{aligned}$$

By the IH, we know that $\langle c_1 \rangle^{l,l} v' \sim \langle S_{21} \Rightarrow S_{11} \rangle^l w' : T_1$. Since $v \sim w : T_1 \rightarrow T_2$, we have $v (\langle c_1 \rangle^{l,l} v') \sim w (\langle S_{21} \Rightarrow S_{11} \rangle^l w') : T_2$. Again by the IH, we can see that $\langle c_2 \{x := v'\} \rangle^{l,l} (v (\langle c_1 \rangle^{l,l} v')) \sim \langle S_{12} \{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w'\} \Rightarrow S_{22} \{x := w'\} \rangle^l w (\langle S_{21} \Rightarrow S_{11} \rangle^l w') : T_2$. $\square \square$

We prove three more technical lemmas necessary for the behavioral and type correspondence.

2.5.15 Lemma [Skeletal equality of subtypes]: If $\Delta \vdash S_1 <: S_2$, then $\lfloor S_1 \rfloor = \lfloor S_2 \rfloor$.

2.5.16 Lemma: If $\lfloor S_1 \rfloor = \lfloor S_2 \rfloor = T$, then $\lfloor \psi^l(S_1, S_2) \rfloor = T$.

2.5.17 Lemma: If $\Delta_1 \vdash S_1$ and $\Delta_1 \vdash S_2$, where $\lfloor S_1 \rfloor = \lfloor S_2 \rfloor$ then

1. If $\Delta_1, x:S_1, \Delta_2 \vdash s : S$ then $\Delta_1, x:S_2, \Delta_2 \{x := \langle S_2 \Rightarrow S_1 \rangle^l x\} \vdash s \{x := \langle S_2 \Rightarrow S_1 \rangle^l x\} : S \{x := \langle S_2 \Rightarrow S_1 \rangle^l x\}$, and
2. If $\Delta_1, x:S_1, \Delta_2 \vdash S$ then $\Delta_1, x:S_2, \Delta_2 \{x := \langle S_2 \Rightarrow S_1 \rangle^l x\} \vdash S \{x := \langle S_2 \Rightarrow S_1 \rangle^l x\}$.

We use the correspondence relations to show that s and its translation $\psi(s)$ correspond—i.e., that ψ faithfully translates the λ_H semantics. We must choose the subject of induction carefully, however, to ensure that we can apply the IH in the case for function casts. An induction on the height of the well-formedness derivation is tricky because of the “extra” substitution that ψ does. Instead, we do induction on the depth of ψ ’s recursion, (and also derivation height, for the S_SUB case).

2.5.18 Theorem [Behavioral correspondence]:

1. If $\Delta \vdash s : S$ then $\lfloor \Delta \rfloor \vdash \psi(s) \sim s : \lfloor S \rfloor$.

2. If $\Delta \vdash S_1$ and $\Delta \vdash S_2$, where $\lfloor S_1 \rfloor = \lfloor S_2 \rfloor = \lfloor S \rfloor$, then $\lfloor \Delta \rfloor \vdash \psi^l(S_1, S_2) \sim S_1 \Rightarrow^l S_2 : \lfloor S \rfloor$ (for all l).

Proof: By induction on the lexicographically ordered pairs (m, n) , where m is the depth of the recursion of the translation $\psi(s)$ (for part 1) or $\psi^l(S_1, S_2)$ (for part 2) and n is either $|\Delta \vdash s : S|$ (for part 1) or $|\Delta \vdash S_1| + |\Delta \vdash S_2|$ (for part 2). The first component decreases in all uses of the IH except for the `S_SUB` case, where only the second component decreases. Part (1) of the proof proceeds by case analysis on the final rule used in the typing derivation $\Delta \vdash s : S$. Which rule was used determines the shape of $\psi(s)$ in all cases but `S_SUB`.

We give only the most interesting cases for the first part: `S_CAST`, `S_CHECKING`, and `S_SUB`.

`S_CAST`: $\Delta \vdash \langle S_1 \Rightarrow S_2 \rangle^l : S_1 \rightarrow S_2$ and $\psi(\langle S_1 \Rightarrow S_2 \rangle^l) = \langle \psi^l(S_1, S_2) \rangle^{l,l}$. By inversion, $\Delta \vdash S_1$ and $\Delta \vdash S_2$, where $\lfloor S_1 \rfloor = \lfloor S_2 \rfloor$.

By the IH for proposition (2), $\lfloor \Delta \rfloor \vdash \psi^l(S_1, S_2) \sim S_1 \Rightarrow^l S_2 : \lfloor S_2 \rfloor$.

Let $\lfloor \Delta \rfloor \models \delta$; we must show $\delta_1(\langle \psi^l(S_1, S_2) \rangle^{l,l}) \sim \delta_2(\langle S_1 \Rightarrow S_2 \rangle^l) : \lfloor S_1 \rfloor \rightarrow \lfloor S_2 \rfloor$. Let $t \sim s : \lfloor S_1 \rfloor$. We have $\delta_1(\langle \psi^l(S_1, S_2) \rangle^{l,l}) t \sim \delta_2(\langle S_1 \Rightarrow S_2 \rangle^l) s : \lfloor S_2 \rfloor$ by Lemma 2.5.14.

`S_CHECKING`: We have $\Delta \vdash \langle \{x:B \mid s_1\}, s_2, k \rangle^l : \{x:B \mid s_1\}$; translating yields $\psi(\langle \{x:B \mid s_1\}, s_2, k \rangle^l) = \langle \{x:B \mid \psi(s_1)\}, \psi(s_2), k \rangle^l$. Recall that the terms of the active check are closed. By inversion we have $\emptyset \vdash s_2 : \{x:\text{Bool} \mid \text{true}\}$ and $\emptyset \vdash k : \{x:B \mid \text{true}\}$, so $k \in \mathcal{K}_B$.

By the IH, $\psi(s_2) \sim s_2 : \text{Bool}$. These two terms coevaluate to blame or a boolean constant. There are three cases, all of which result in the active checks evaluating to \approx -corresponding values:

- If they go to $\uparrow l'$, then the checks do too, and $\uparrow l' \approx \uparrow l' : B$.
- If they go to **false**, then the checks go to $\uparrow l$, and $\uparrow l \approx \uparrow l : B$.
- If they go to **true**, then the checks go to $k \in \mathcal{K}_B$, and $k \approx k : B$.

`S_SUB`: $\Delta \vdash s : S$; we do not know anything about the shape of $\psi(s)$. By inversion, $\Delta \vdash s : S'$ and $\Delta \vdash S' <: S$. By Lemma 2.5.15, $\lfloor S' \rfloor = \lfloor S \rfloor$.

Since the sub-derivation $\Delta \vdash s : S'$ is smaller, by the IH $\lfloor \Delta \rfloor \vdash \psi(s) \sim s : \lfloor S' \rfloor$. But $\lfloor S' \rfloor = \lfloor S \rfloor$, so we are done.

Part (2) of this proof proceeds by cases on $\psi^l(S_1, S_2) = c$.

$\psi^l(S_1, \{x:B \mid s_2\}) = \{x:B \mid \psi(s_2)\}$: Note that $S_2 = \{x:B \mid s_2\}$. By inversion of $\Delta \vdash \{x:B \mid s_2\}$, we have $\Delta, x:\{x:B \mid \text{true}\} \vdash s_2 : \{x:\text{Bool} \mid \text{true}\}$.

By the IH for proposition (1), $\lfloor \Delta \rfloor, x:B \vdash \psi(s_2) \sim s_2 : \text{Bool}$.

We must show $\lfloor \Delta \rfloor \vdash \{x:B \mid \psi(s_2)\} \sim S_1 \Rightarrow^l \{x:B \mid s_2\} : B$. Let $\lfloor \Delta \rfloor \models \delta$; we prove that $\delta_1(\{x:B \mid \psi(s_2)\}) \sim \delta_2(S_1) \Rightarrow^l \delta_2(\{x:B \mid s_2\}) : B$, i.e., for all $k \in \mathcal{K}_B$, that $\delta_1(\psi(s_2))\{x := k\} \sim \delta_2(s_2)\{x := k\} : \text{Bool}$. Since $k \sim k : B$, we can see this last by the IH.

$\psi^l(x:S_{11} \rightarrow S_{12}, x:S_{21} \rightarrow S_{22}) = x:\psi^l(S_{21}, S_{11}) \mapsto \psi^l(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}, S_{22})$:
 We can see $S_2 = x:S_{21} \rightarrow S_{22}$ and so $S_1 = x:S_{11} \rightarrow S_{12}$, where $\lfloor S_{21} \rfloor = \lfloor S_{11} \rfloor$ and $\lfloor S_{22} \rfloor = \lfloor S_{12} \rfloor$. By inversion, we have the following well-formedness derivations:

$$\begin{array}{cc} \Delta \vdash S_{21} & \Delta \vdash S_{11} \\ \Delta, x:S_{21} \vdash S_{22} & \Delta, x:S_{22} \vdash S_{12} \end{array}$$

We can apply the IH (contravariantly) to see

$$\lfloor \Delta \rfloor \vdash \psi^l(S_{21}, S_{11}) \sim S_{21} \Rightarrow^l S_{11} : \lfloor S_{11} \rfloor \quad (*)$$

By weakening (Lemma 2.3.22), we can see $\Delta, x:S_{21} \vdash S_{21}$ and $\Delta, x:S_{21} \vdash S_{11}$. We can reapply the IH to show $\lfloor \Delta \rfloor, x:\lfloor S_{21} \rfloor \vdash \psi^l(S_{21}, S_{11}) \sim S_{21} \Rightarrow^l S_{11} : \lfloor S_{11} \rfloor$. Now $\Delta, x:S_{21} \vdash \langle S_{21} \Rightarrow S_{11} \rangle^l : S_{21} \rightarrow S_{11}$ and $\Delta, x:S_{21} \vdash \langle S_{21} \Rightarrow S_{11} \rangle^l x : S_{11}$. By Lemma 2.5.17, we can substitute this last into $\Delta, x:S_{11} \vdash S_{12}$, finding $\Delta, x:S_{21} \vdash S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}$.

We apply the IH for proposition (2) on $\Delta, x:S_{21} \vdash S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}$ and $\Delta, x:S_{21} \vdash S_{22}$, showing

$$\begin{array}{c} \lfloor \Delta \rfloor, x:\lfloor S_{21} \rfloor \vdash \psi^l(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}, S_{22}) \sim \\ S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\} \Rightarrow^l S_{22} : \lfloor S_{22} \rfloor \end{array} \quad (**)$$

We now combine (*) and (**) to show $\lfloor \Delta \rfloor \vdash \psi^l(x:S_{11} \rightarrow S_{12}, x:S_{21} \rightarrow S_{22}) \sim x:S_{11} \rightarrow S_{12} \Rightarrow^l x:S_{21} \rightarrow S_{22} : \lfloor S_2 \rfloor$. Let $\lfloor \Delta \rfloor \models \delta$. We can apply (*) to see $\delta_1(\psi^l(S_{21}, S_{11})) \sim \delta_2(S_{21}) \Rightarrow^l \delta_2(S_{11}) : \lfloor S_{11} \rfloor$. For the codomain we must show, for all $t \sim s : \lfloor S_{11} \rfloor$, that

$$\begin{array}{c} \delta_1(\psi^l(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}, S_{22}))\{x := t\} \sim \\ \delta_2(S_{12})\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l s\} \Rightarrow^l \delta_2(S_{22})\{x := s\} : \lfloor S_{22} \rfloor \end{array}$$

Let $t \sim s : \lfloor S_{11} \rfloor$. Recalling that $\lfloor S_{11} \rfloor = \lfloor S_{21} \rfloor$, observe $\lfloor \Delta \rfloor, x:\lfloor S_{21} \rfloor \models \delta\{x := t, s\}$. Call this δ' . By (**) we see

$$\begin{array}{c} \delta'_1(\psi^l(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}, S_{22})) \sim \\ \delta'_2(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}) \Rightarrow^l \delta'_2(S_{22}) : \lfloor S_{22} \rfloor \end{array}$$

which we can rewrite to

$$\begin{array}{c} \delta_1(\psi^l(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}, S_{22}))\{x := t\} \sim \\ \delta_2(S_{12})\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l s\} \Rightarrow^l \delta_2(S_{22})\{x := s\} : \lfloor S_{22} \rfloor \end{array}$$

This is exactly what we needed to finish the proof of correspondence. \square

As a preliminary to type-preservation, we can use behavioral correspondence to show that the implication judgment is preserved.

2.5.19 Lemma: If $\emptyset \vdash s_1 : \{x:\text{Bool} \mid \text{true}\}$ and $\emptyset \vdash s_2 : \{x:\text{Bool} \mid \text{true}\}$ and $\emptyset \vdash s_1 \supset s_2$, then $\psi(s_1) \longrightarrow_{\text{Iax}^*} \text{true}$ implies $\psi(s_2) \longrightarrow_{\text{Iax}^*} \text{true}$.

Proof: By the logical relation. □

The type preservation proof is very similar to the correspondence proof of Theorem 2.5.18, though the function case of the type/contract correspondence is intricate.

2.5.20 Theorem [Type preservation for ψ]:

1. If $\Delta \vdash s : S$ then $\lfloor \Delta \rfloor \vdash \psi(s) : \lfloor S \rfloor$.
2. If $\Delta \vdash S_1, \Delta \vdash S_2$, where $\lfloor S_1 \rfloor = \lfloor S_2 \rfloor = T$, then $\lfloor \Delta \rfloor \vdash^{l,l'} \psi^l(S_1, S_2) : T$.

Proof: By induction on the lexicographically ordered pair containing (a) the depth of the recursion of the translation ψ or $\psi(s)$, and (b) $|\Delta \vdash s : S|$ or $|\Delta \vdash S_1| + |\Delta \vdash S_2|$.

Part (1) of the proof proceeds by case analysis on the final rule of $\Delta \vdash s : S$, which determines the shape of $\psi(s) = t$ in all cases but S.SUB. Part (2) of the proof proceeds by case analysis on $\psi^l(S_1, S_2) = c$.

$\psi^l(x:S_{11} \rightarrow S_{12}, x:S_{21} \rightarrow S_{22}) = x:\psi^l(S_{21}, S_{11}) \mapsto \psi^l(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}, S_{22})$:

We must have $S_2 = x:S_{21} \rightarrow S_{22}$ and $S_1 = x:S_{11} \rightarrow S_{12}$, where $\lfloor S_{21} \rfloor = \lfloor S_{11} \rfloor$ and $\lfloor S_{22} \rfloor = \lfloor S_{12} \rfloor$. By inversion, we have the following well-formedness derivations:

$$\begin{array}{cc} \Delta \vdash S_{21} & \Delta \vdash S_{11} \\ \Delta, x:S_{21} \vdash S_{22} & \Delta, x:S_{22} \vdash S_{12} \end{array}$$

By the IH $\lfloor \Delta \rfloor \vdash^{l,l'} \psi^l(S_{21}, S_{11}) : \lfloor S_{11} \rfloor$. Note that $\lfloor \psi^l(S_{21}, S_{11}) \rfloor = \lfloor S_{21} \rfloor$.

By weakening, we can see $\Delta, x:S_{21} \vdash S_{21}$ and $\Delta, x:S_{21} \vdash S_{11}$. We can reapply the IH to show $\lfloor \Delta \rfloor, x:\lfloor S_{21} \rfloor \vdash^{l,l'} \psi^l(S_{21}, S_{11}) : \lfloor S_{11} \rfloor$. Now $\Delta, x:S_{21} \vdash \langle S_{21} \Rightarrow S_{11} \rangle^l : S_{21} \rightarrow S_{11}$. Next $\Delta, x:S_{21} \vdash \langle S_{21} \Rightarrow S_{11} \rangle^l x : S_{11}$. By Lemma 2.5.17, we can substitute this last into $\Delta, x:S_{11} \vdash S_{12}$, finding $\Delta, x:S_{21} \vdash S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}$.

By the IH for proposition (2) on $\Delta, x:S_{21} \vdash S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}$ and $\Delta, x:S_{21} \vdash S_{22}$,

$$\lfloor \Delta \rfloor, x:\lfloor S_{21} \rfloor \vdash^{l,l'} \psi^l(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}, S_{22}) : \lfloor S_{22} \rfloor$$

By Lemma 2.5.16, $\lfloor \psi^l(S_{21}, S_{11}) \rfloor = \lfloor S_{21} \rfloor$, so we can rewrite the above derivation to

$$\lfloor \Delta \rfloor, x:\psi^l(S_{21}, S_{11}) \vdash^{l,l'} \psi^l(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}, S_{22}) : \lfloor S_{22} \rfloor$$

Now by T.FUNC

$$\lfloor \Delta \rfloor \vdash^{l,l'} x:\psi^l(S_{21}, S_{11}) \mapsto \psi^l(S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l x\}, S_{22}) : \lfloor S_{21} \rfloor \rightarrow \lfloor S_{22} \rfloor \quad \square$$

□

2.6 Inexact translations

The same translations ϕ and ψ can be used to move right on the axis of blame (Figure 2.1). However, in this direction the images of these translations blame strictly more than their pre-images. We were able to use the same correspondence for both exact proofs in Section 2.5, but the following two proofs use custom correspondences: one where lax λ_C terms correspond to λ_H terms (with possibly more blame), and one where λ_H terms correspond to picky λ_C terms (with possibly more blame). In both cases, the λ_C terms will be on the left and the λ_H terms on the right.

2.6.1 Translating lax λ_C to λ_H

Translating with ϕ from terms in picky λ_C to exactly corresponding terms in λ_H was a relatively straightforward generalization of the nondependent case; things get more interesting when we consider the translation ϕ from lax λ_C to dependent λ_H . We can prove that it preserves types (for terms without active checks), but we can only show a weaker behavioral correspondence: sometimes lax λ_C terms terminate with values when their ϕ -images go to blame. This weaker property is a consequence of bulletproofing, the asymmetrically substituting `F_CDECOMP` rule, and the extra casts inserted for type preservation (i.e., for `T_VARC` derivations). This is not a weakness of my proof technique—I give a counterexample, a lax λ_C term $\emptyset \vdash t : T$ such that $t \longrightarrow_{lax}^* v$ and $\phi(\emptyset \vdash t : T) \longrightarrow_h^* \uparrow l$.

We can show the behavioral correspondence using a blame-inexact logical relation, defined in Figure 2.17. The behavioral correspondence here, though weaker than before, is still pretty strong: if $t \sim_{\prec} s : B$ (read “ t blames no more than s at type B ”), then either $s \longrightarrow_h^* \uparrow l$ or t and s both go to $k \in \mathcal{K}_B$. This correspondence differs slightly in construction from the earlier exact one—we define \approx_{\prec} as a relation on *values*, while \approx is a relation on *results*. Doing so simplifies the inexact treatment of blame—in particular, Lemma 2.6.2. We again use the term correspondence to relate contracts and λ_H types. We then lift the correspondences to open terms (Figure 2.17). Closing substitutions map variables to corresponding terms of appropriate type. Note that closing substitutions ignore the contract part of $x:c^{l,l'}$ bindings, treating them as if they were $x:[c]$.

2.6.1 Lemma [Expansion and contraction]: If $t \longrightarrow_{lax}^* t'$, and $s \longrightarrow_h^* s'$ then $t \sim_{\prec} s : T$ iff $t' \sim_{\prec} s' : T$.

Note that there are corresponding terms at every type. We can prove a much stronger lemma than we did for \sim in Lemma 2.5.11, since the correspondence here is much weaker.

2.6.2 Lemma [Everything corresponds to blame]: For all t and T , $t \sim_{\prec} \uparrow l' : T$.

Value correspondence $\boxed{v \approx_{\gamma} w : T}$

$$k \approx_{\gamma} k : B \iff k \in \mathcal{K}_B$$

$$v \approx_{\gamma} w : T_1 \rightarrow T_2 \iff \forall t \sim_{\gamma} s : T_1. v t \sim_{\gamma} w s : T_2$$

Term correspondence $\boxed{t \sim_{\gamma} s : T}$

$$t \sim_{\gamma} s : T \iff s \rightarrow_h^* \uparrow l \vee (t \rightarrow_{lax}^* v \wedge s \rightarrow_h^* w \wedge v \approx_{\gamma} w : T)$$

Contract/type correspondence $\boxed{c \sim_{\gamma} S : T}$

$$\{x:B \mid t\} \sim_{\gamma} \{x:B \mid s\} : B \iff \forall k \in \mathcal{K}_B. t\{x := k\} \sim_{\gamma} s\{x := k\} : \text{Bool}$$

$$x:c_1 \mapsto c_2 \sim_{\gamma} x:S_1 \rightarrow S_2 : T_1 \rightarrow T_2 \iff c_1 \sim_{\gamma} S_1 : T_1 \wedge \forall t \sim_{\gamma} s : T_1. c_2\{x := t\} \sim_{\gamma} S_2\{x := s\} : T_2$$

Dual closing substitutions

$$\Gamma \models_{\gamma} \delta \iff \forall x \in \text{dom}(\Gamma). \delta_1(x) \sim_{\gamma} \delta_2(x) : [\Gamma(x)]$$

Lifted to open terms

$$\begin{aligned} \Gamma \vdash t \sim_{\gamma} s : T &\iff \forall \delta. (\Gamma \models_{\gamma} \delta \text{ implies } \delta_1(t) \sim_{\gamma} \delta_2(s) : T) \\ \Gamma \vdash c \sim_{\gamma} S : T &\iff \forall \delta. (\Gamma \models_{\gamma} \delta \text{ implies } \delta_1(c) \sim_{\gamma} \delta_2(S) : T) \end{aligned}$$

Figure 2.17: Blame-inexact correspondence for ϕ from lax λ_C

2.6.3 Lemma [Constants self-correspond]: For all k , $k \approx_{\succ} k : \text{ty}_c(k)$.

Proof: By induction on $\text{ty}_c(k)$, recalling that constants are first order. \square

As a corollary of Lemma 2.6.2 and Lemma 2.6.1, if two terms evaluate to blame—or even just the λ_H side!—then they correspond. This will be used extensively in the proofs below, as it allows us to eliminate many cases.

We prove three lemmas about contracts and casts at base types. The first two characterize contracts and casts at base types.

2.6.4 Lemma [Trivial casts]: If $t \sim_{\succ} s : B$, then $t \sim_{\succ} \langle S \Rightarrow [B] \rangle^l s : B$ for all S .

2.6.5 Lemma [Related base casts]: If $\{x:B \mid t\} \sim_{\succ} \{x:B \mid s\} : B$ and $t' \sim_{\succ} s' : B$, then $\langle \{x:B \mid t\} \rangle^{l,l'} t' \sim_{\succ} \langle S \Rightarrow \{x:B \mid s\} \rangle^l s' : B$ for all S .

The third lemma shows that correspondence is closed under adding extra casts to the λ_H term, due to the inexactness of the behavioral correspondence. Since λ_H terms can go to blame more often than corresponding lax λ_C terms, we can add “extra” casts to λ_H terms. We formalize this in the following lemma, which captures the asymmetric treatment of blame by the \sim_{\succ} relation. We use it to show that the cast substituted in the codomain by `F_CDECOMP` does not affect behavioral correspondence. Note that the statement of the lemma requires that the types of the cast correspond to *some* contracts at the same type T , but we never use the contracts in the proof—they witness the well-formedness of the λ_H types.

2.6.6 Lemma [Extra casts]: If $t \sim_{\succ} s : T$ and $c_1 \sim_{\succ} S_1 : T$ and $c_2 \sim_{\succ} S_2 : T$, then $t \sim_{\succ} \langle S_1 \Rightarrow S_2 \rangle^l s : T$.

Proof: The proof is by induction on T . Note that we do not use c_1 or c_2 at all in the proof, but instead they are witnesses to the well-formedness of S_1 and S_2 .

$[S_1] = [S_2] = T$. Either $s \rightarrow_h^* \uparrow l'$ or t and s both go to corresponding values at T . If $s \rightarrow_h^* \uparrow l'$, then $\langle S_1 \Rightarrow S_2 \rangle^l s \rightarrow_h^* \uparrow l'$ and $t \sim_{\succ} \uparrow l' : T$ since everything is related to blame (Lemma 2.6.2).

Therefore, suppose that $t \rightarrow_{\text{Iax}}^* v$ and $s \rightarrow_h^* w$ and $v \approx_{\succ} w : T$ in each of the following cases of the induction.

$T = B$: So $S_2 = \{x:B \mid s_2\}$, and $c_2 = \{x:B \mid t_2\}$.

So $t \rightarrow_{\text{Iax}}^* k$ and $s \rightarrow_h^* k$ for $k \in \mathcal{K}_B$. If t and s both go to k , then $\langle S_1 \Rightarrow S_2 \rangle^l s \rightarrow_h^* \langle \{x:B \mid s_2\}, s_2\{x := k\}, k \rangle^l$. By $c_2 \sim_{\succ} S_2 : B$ we see (in particular) $t_2\{x := k\} \sim_{\succ} s_2\{x := k\} : \text{Bool}$. So $s_2\{x := k\}$ either goes to $\uparrow l'$ or $s_2\{x := k\}$ (and, irrelevantly, $t_2\{x := k\}$) go to some $k' \in \mathcal{K}_{\text{Bool}}$. In the former case, $\langle \{x:B \mid s_2\}, s_2\{x := k\}, k \rangle^l \rightarrow_h^* \uparrow l'$ and we are done (by Lemma 2.6.2). In the latter case, the λ_H term either goes to $\uparrow l$ (and everything is related to blame) or goes to k —but so does t , and $k \approx_{\succ} k : B$.

$T = T_1 \rightarrow T_2$: We have:

$$\begin{array}{ll} S_1 = x:S_{11} \rightarrow S_{12} & S_2 = x:S_{21} \rightarrow S_{22} \\ c_1 = x:c_{11} \mapsto c_{12} & c_2 = x:c_{21} \mapsto c_{22} \end{array}$$

We have $t \rightarrow_{\text{Iax}}^* v$ and $s \rightarrow_h^* w$, where $v \approx_{\succ} w : T_1 \rightarrow T_2$.

Let $t' \sim_{\succ} s' : T_1$; we wish to see that $v t' \sim_{\succ} (\langle S_1 \Rightarrow S_2 \rangle^l w) s' : T_2$. Either $s' \rightarrow_h^* \uparrow l'$ or both go to values. In the former case the whole cast goes to $\uparrow l'$ we are done by Lemma 2.6.2, so let $t' \rightarrow_{\text{Iax}}^* v'$ and $s' \rightarrow_h^* w'$.

Decomposing the cast in λ_H ,

$$\begin{array}{l} (\langle S_1 \Rightarrow S_2 \rangle^l w) s' \rightarrow_h^* \\ \langle S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w'\} \Rightarrow S_{22}\{x := w'\} \rangle^l (w (\langle S_{21} \Rightarrow S_{11} \rangle^l w')) \end{array}$$

We have $c_{21} \sim_{\succ} S_{21} : T_1$ and $c_{11} \sim_{\succ} S_{11} : T_1$, so $v' \sim_{\succ} \langle S_{21} \Rightarrow S_{11} \rangle^l w' : T_1$ by the IH. Since $v \approx_{\succ} w : T_1 \rightarrow T_2$, we can see that $v v' \sim_{\succ} w (\langle S_{21} \Rightarrow S_{11} \rangle^l w') : T_2$.

Furthermore, we know that for all $t'' \sim_{\succ} s'' : T_1$ that

- $c_{12}\{x := t''\} \sim_{\succ} S_{12}\{x := s''\} : T_2$ and
- $c_{22}\{x := t''\} \sim_{\succ} S_{22}\{x := s''\} : T_2$.

We know that $v' \sim_{\succ} w' : T_1$ and $v' \sim_{\succ} \langle S_{21} \Rightarrow S_{11} \rangle^l w' : T_1$, so we can see

- $c_{12}\{x := v'\} \sim_{\succ} S_{12}\{x := w'\} : T_2$ and
- $c_{22}\{x := v'\} \sim_{\succ} S_{22}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w'\} : T_2$.

So by the IH,

$$v v' \sim_{\succ} \langle S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w'\} \Rightarrow S_{22}\{x := w'\} \rangle^l (w (\langle S_{21} \Rightarrow S_{11} \rangle^l w')) : T_2$$

and we are done by expansion (Lemma 2.6.1). □ □

To apply the extra cast lemma, we'll need these “witness” contracts for raw types; to that end we define trivial contracts. These contracts are *lifted* from types, and are the λ_C correlate to λ_H 's raw types.

$$\begin{array}{l} B\uparrow = \{x:B \mid \text{true}\} \\ (T_1 \rightarrow T_2)\uparrow = (T_1\uparrow) \mapsto (T_2\uparrow) \end{array}$$

2.6.7 Lemma [Lifted types logically relate to raw types]: For all T , $T\uparrow \sim_{\succ} [T] : T$.

The “bulletproofing” lemma is the key to the behavioral correspondence proof. I show that a contract application corresponds to bulletproofing with related types. Note that I allow for different types in the two casts. This is necessary due to an asymmetric substitution that occurs when $T = B \rightarrow T_2$.

2.6.8 Lemma [Bulletproofing]: If $t \sim_{\succ} s : T$ and $c \sim_{\succ} S : T$ and $c \sim_{\succ} S' : T$, then $\langle c \rangle^{l,l'} t \sim_{\succ} \langle S' \Rightarrow [S'] \rangle^{l'} \langle [S] \Rightarrow S \rangle^l s : T$.

Proof: By induction on T . First, observe that either $s \rightarrow_h^* \uparrow l''$ or both t and s go to values related at T . In the former case, $\langle S' \Rightarrow [S'] \rangle^{l'} \langle [S] \Rightarrow S \rangle^l s \rightarrow_h^* \uparrow l''$, and everything is related to blame (Lemma 2.6.2). So $t \rightarrow_{lax}^* v$, $s \rightarrow_h^* w$, and $v \approx_{\succ} w : T$.

$T = B$: So $c = \{x:B \mid t_1\}$ and $S = \{x:B \mid s_1\}$ and $S' = \{x:B \mid s_2\}$. By Lemma 2.6.5 we have $\langle c \rangle^{l,l'} t \sim_{\succ} \langle [S] \Rightarrow S \rangle^l s : B$. By Lemma 2.6.4 we can add the extra, trivial cast.

$T = T_1 \rightarrow T_2$: We know that $c = x:c_1 \mapsto c_2$, $S = x:S_1 \rightarrow S_2$ and $S' = x:S'_1 \rightarrow S'_2$. Let $t' \sim_{\succ} s' : T_1$. By Lemma 2.6.2, we only need to consider the case where $t' \rightarrow_{lax}^* v'$ and $s' \rightarrow_h^* w'$ —if $s' \rightarrow_h^* \uparrow l''$ we are done.

On the λ_C side, $(\langle c \rangle^{l,l'} t) t' \rightarrow_{lax}^* \langle c_2\{x := v'\} \rangle^{l,l'} (v (\langle c_1 \rangle^{l,l'} v'))$. In λ_H , we can see

$$\begin{aligned} & (\langle S' \Rightarrow [S'] \rangle^{l'} \langle [S] \Rightarrow S \rangle^l s) s' \rightarrow_h^* \\ & \langle S'_2\{x := \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w'\} \Rightarrow [S'_2] \rangle^{l'} ((\langle [S] \Rightarrow S \rangle^l w) (\langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w')) \end{aligned}$$

We cannot determine where the redex is until we know the shape of T_1 —does the negative argument cast step to an active check, or do we decompose the positive cast?

– $T_1 = B$.

By Lemma 2.6.5 and $c_1 \sim_{\succ} S'_1 : B$, we know that $\langle c_1 \rangle^{l,l'} v' \sim_{\succ} \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w' : B$. The λ_H term goes to blame or both terms go to the same value, $v' = w' = k \in \mathcal{K}_B$. In the former case, the entire λ_H term goes to blame and we are done by Lemma 2.6.2. So suppose $\langle c_1 \rangle^{l,l'} k \rightarrow_{lax}^* k$ and $\langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w' \rightarrow_h^* k$. Now the terms evaluate like so:

$$\langle c_2\{x := v'\} \rangle^{l,l'} (v (\langle c_1 \rangle^{l,l'} v')) \rightarrow_{lax}^* \langle c_2\{x := k\} \rangle^{l,l'} (v k)$$

$$\begin{aligned} & \langle S'_2\{x := \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w'\} \Rightarrow [S'_2] \rangle^{l'} \\ & ((\langle [S] \Rightarrow S \rangle^l w) (\langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w')) \rightarrow_h^* \\ & \langle S'_2\{x := \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w'\} \Rightarrow [S'_2] \rangle^{l'} \\ & \langle [S_2] \Rightarrow S_2\{x := k\} \rangle^l (w (\langle S_1 \Rightarrow [S_1] \rangle^l k)) \end{aligned}$$

By Lemma 2.6.4, $k \sim_{\succ} \langle S_1 \Rightarrow [S_1] \rangle^l k : B$, so $v k \sim_{\succ} w (\langle S_1 \Rightarrow [S_1] \rangle^l k) : T_2$.

Noting that $k \sim_{\succ} k : B$ and $k \sim_{\succ} \langle [S_1] \Rightarrow S_1 \rangle^l k : B$, we can see that $c_2\{x := k\} \sim_{\succ} S_2\{x := k\} : T_2$ and $c_2\{x := k\} \sim_{\succ} S'_2\{x := \langle [S_1] \Rightarrow S_1 \rangle^l k\} : T_2$. Now the IH shows that $\langle c_2\{x := k\} \rangle^{l,l'} (v k) \sim_{\succ} \langle S'_2\{x := \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w'\} \Rightarrow [S'_2] \rangle^{l'} \langle [S_2] \Rightarrow S_2\{x := k\} \rangle^l (w (\langle S_1 \Rightarrow [S_1] \rangle^l k)) : T_2$, and we conclude the case with expansion (Lemma 2.6.1).

– $T_1 = T_{11} \rightarrow T_{12}$. We can continue with an application of `F_CDECOMP` in λ_H and find:

$$\begin{aligned} & \langle S'_2 \{x := \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w'\} \Rightarrow [S'_2] \rangle^{l'} \\ & \quad ((\langle [S] \Rightarrow S \rangle^l w) (\langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w')) \longrightarrow_h^* \\ & \langle S'_2 \{x := \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w'\} \Rightarrow [S'_2] \rangle^{l'} \\ & \quad \langle [S_2] \Rightarrow S_2 \{x := \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w'\} \rangle^l \\ & \quad (w (\langle S_1 \Rightarrow [S_1] \rangle^l \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w')) \end{aligned}$$

By the IH, $\langle c_1 \rangle^{l',l} v' \sim_{\succ} \langle S_1 \Rightarrow [S_1] \rangle^l \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w' : T_1$. By assumption, the results of applying v and w to these values correspond. (And they *are* values, since function contracts/casts applied to values are values.)

We know $c_1 \sim_{\succ} S'_1 : T_1$, and by Lemma 2.6.7 $T_1 \uparrow \sim_{\succ} [S'_1] : T_1$. Since $v' \sim_{\succ} w' : T_1$, Lemma 2.6.6 shows $v' \sim_{\succ} \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w' : T_1$. This lets us see that $c_2 \{x := v'\} \sim_{\succ} S'_2 \{x := \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w'\} : T_2$ and $c_2 \{x := v'\} \sim_{\succ} S_2 \{x := \langle [S'_1] \Rightarrow S'_1 \rangle^{l'} w'\} : T_2$. Now the IH and expansion (Lemma 2.6.1) complete the proof. \square \square

Having characterized how contracts and pairs of related casts relate, we show that translated terms correspond to their sources.

2.6.9 Theorem [Behavioral correspondence]: If $\vdash \Gamma$, then:

1. If $\phi(\Gamma \vdash t : T) = s$ then $\Gamma \vdash t \sim_{\succ} s : T$.
2. If $\phi(\Gamma \vdash^{l,l'} c : T) = S$ then $\Gamma \vdash c \sim_{\succ} S : T$.

Proof: We simultaneously show both properties by induction on the depth of ϕ 's recursion. To show $\Gamma \vdash t \sim_{\succ} s : T$, let $\Gamma \models \delta$ —we will show $\delta_1(t) \sim_{\succ} \delta_2(s) : T$.

The proof proceeds by case analysis on the final rule of the translated typing and well-formedness derivations. \square

We find a weak corollary: $\phi(\Gamma \vdash t : B) \longrightarrow_h^* k$ implies $t \longrightarrow_{\text{la}x}^* k$: if the λ_H term does *not* go to blame, then the original λ_C term must go to the same constant.

We can also show type preservation for terms not containing active checks. (We do not know that translated active checks are well typed, because Theorem 2.6.9 is not strong enough to preserve the implication judgment. We only expect these checks to occur at runtime, so this is good enough: ϕ preserves the types of source programs.)

2.6.10 Theorem [Type preservation for ϕ]: For programs without active checks, if $\phi(\vdash \Gamma) = \Delta$, then:

1. $\vdash \Delta$.
2. $\Delta \vdash \phi(\Gamma \vdash t : T) : [T]$.
3. $\Delta \vdash \phi(\Gamma \vdash^{l,l'} c : T)$.

Proof: We prove all three properties simultaneously, by induction on the depth of ϕ 's recursion.

The proof is by cases on the λ_C context well-formedness/term typing/contract well-formedness derivations, which determine the branch of ϕ taken. \square

To see that the ϕ in Figure 2.12 does not give us exact blame, let us look at two counterexamples; in both cases, a lax λ_C term goes to a value while its translation goes to blame. In the first example, blame is raised in λ_H due to bulletproofing. In the second, blame is raised due to the extra cast from the translation of `T_VARC`. In both examples, the contracts are *abusive*: higher-order contracts where the codomain places a contradictory requirement on the domain. For the first counterexample, let

$$\begin{aligned} c &= f:(x:\{x:\text{Int} \mid \text{true}\} \mapsto \{y:\text{Int} \mid \text{nonzero } y\}) \mapsto \{z:\text{Int} \mid f \ 0 = 0\} \\ S_1 &= x:\{x:\text{Int} \mid \text{true}\} \rightarrow \{y:\text{Int} \mid \text{nonzero } y\} \\ S &= \phi(\emptyset \vdash^{l,l} c : (\text{Int} \rightarrow \text{Int}) \rightarrow \text{Int}) \\ &= f:S_1 \rightarrow \{z:\text{Int} \mid (\langle S_1 \Rightarrow \lceil S_1 \rceil \rangle^l f) \ 0 = 0\}. \end{aligned}$$

Here, the contradiction comes when the codomain requires that $f \ 0$ yield 0, but f 's contract says it will return a non-zero value. We find $\langle c \rangle^{l,l} (\lambda f.0) (\lambda x.0) \longrightarrow_{\text{lax}}^* 0$ but

$$(\lambda x:\lceil c \rceil. \langle S \Rightarrow \lceil S \rceil \rangle^l (\langle \lceil S \rceil \Rightarrow S \rangle^l x)) (\lambda f.0) (\lambda x.0) \longrightarrow_h^* \uparrow l.$$

For the second counterexample, let

$$\begin{aligned} c' &= f:(x:\{x:\text{Int} \mid \text{nonzero } x\} \mapsto \{y:\text{Int} \mid \text{true}\}) \mapsto \{z:\text{Int} \mid f \ 0 = 0\} \\ S'_1 &= x:\{x:\text{Int} \mid \text{nonzero } x\} \rightarrow \{y:\text{Int} \mid \text{true}\} \\ S' &= \phi(\emptyset \vdash^{l,l} c' : (\text{Int} \rightarrow \text{Int}) \rightarrow \text{Int}) \\ &= f:S'_1 \rightarrow \{z:\text{Int} \mid (\langle S'_1 \Rightarrow \lceil S'_1 \rceil \rangle^l f) \ 0 = 0\}. \end{aligned}$$

This time, the contradiction comes from the codomain applying f to 0, while the domain contract requires that f 's input be non-zero. We find $\langle c' \rangle^{l,l} (\lambda f.0) (\lambda x.0) \longrightarrow_{\text{lax}}^* 0$ but

$$(\lambda x:\lceil c' \rceil. \langle S' \Rightarrow \lceil c' \rceil \rangle^l (\langle \lceil S' \rceil \Rightarrow \lceil c' \rceil \rangle^l x)) (\lambda f.0) (\lambda x.0) \longrightarrow_h^* \uparrow l.$$

The extra casts that ϕ inserts are all necessary—none can be removed. So while variations on this ϕ are possible, they can only add more casts, which won't resolve the problem that λ_H blames *more*.

2.6.2 Translating λ_H to picky λ_C

Terms in λ_H and their ψ -images in lax λ_C correspond exactly, as shown Section 2.5.2. When we change the operational semantics of λ_C to be *picky*, however, $\psi(s)$ blames (strictly) more often than s . Nevertheless, we can show an inexact correspondence, as we did for ϕ and lax λ_C in Section 2.6.1. I use a logical relation \sim_{\prec} for ψ into picky λ_C (Figure 2.18). Here I have reversed the asymmetry: picky λ_C may blame

Value correspondence $\boxed{v \approx_{\prec} w : T}$

$$k \approx_{\prec} k : B \iff k \in \mathcal{K}_B$$

$$v \approx_{\prec} w : T_1 \rightarrow T_2 \iff \forall t \sim_{\prec} s : T_1. v t \sim_{\prec} w s : T_2$$

Term correspondence $\boxed{t \sim_{\prec} s : T}$

$$t \sim_{\prec} s : T \iff t \rightarrow_{\text{picky}}^* \uparrow l \vee t \rightarrow_{\text{picky}}^* v \wedge s \rightarrow_h^* w \wedge v \approx_{\prec} w : T$$

Contract/type correspondence $\boxed{c \sim_{\prec} S_1 \Rightarrow S_2 : T}$

$$\{x:B \mid t\} \sim_{\prec} \{x:B \mid s_1\} \Rightarrow \{x:B \mid s_2\} : B \iff \forall k \in \mathcal{K}_B. t\{x := k\} \sim_{\prec} s_2\{x := k\} : \text{Bool}$$

$$x:c_1 \mapsto c_2 \sim_{\prec} x:S_{11} \rightarrow S_{12} \Rightarrow x:S_{21} \rightarrow S_{22} : T_1 \rightarrow T_2 \iff c_1 \sim_{\prec} S_{21} \Rightarrow S_{11} : T_1 \wedge \forall l. \forall t \sim_{\prec} s : T_1. c_2\{x := t\} \sim_{\prec} S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l s\} \Rightarrow S_{22}\{x := s\} : T_2$$

Dual closing substitutions

$$\Gamma \models \delta \iff \forall x \in \text{dom}(\Gamma). \delta_1(x) \sim_{\prec} \delta_2(x) : [\Gamma(x)]$$

Lifted to open terms

$$\begin{aligned} \Gamma \vdash t \sim_{\prec} s : T &\iff \forall \delta. (\Gamma \models \delta \text{ implies } \delta_1(t) \sim_{\prec} \delta_2(s) : T) \\ \Gamma \vdash c \sim_{\prec} S_1 \Rightarrow S_2 : T &\iff \forall \delta. (\Gamma \models \delta \text{ implies } \delta_1(c) \sim_{\prec} \delta_2(S_1) \Rightarrow \delta_2(S_2) : T) \end{aligned}$$

Figure 2.18: Blame-inexact correspondence for ψ into picky λ_C

more than λ_H . The proof follows the same general pattern: we first show that it is safe to add extra contract checks, then we show that contracts and casts correspond (inexactly), then the correspondence for well-typed terms. We can also show type preservation for source programs (excluding active checks).

2.6.11 Lemma [Expansion and contraction]: If $t \longrightarrow_{\text{picky}}^* t'$ and $s \longrightarrow_h^* s'$ then $t \sim_{\prec} s : T$ iff $t' \sim_{\prec} s' : T$.

2.6.12 Lemma [Blame corresponds to everything]: For all T , $\uparrow l \sim_{\prec} s : T$.

2.6.13 Lemma [Constants self-correspond]: For all k , $k \approx_{\prec} k : \text{ty}_c(k)$.

Proof: By induction on $\text{ty}_c(k)$, recalling that constants are first order. \square

As a corollary of Lemma 2.6.12 and Lemma 2.6.11, if a picky λ_C term evaluates to blame, then it corresponds to any λ_H term. This will be used extensively in the proofs below, as it allows us to eliminate many cases.

2.6.14 Lemma [Extra contracts]: If $t \sim_{\prec} s : T$ and $c \sim_{\prec} S_1 \Rightarrow S_2 : T$ then $\langle c \rangle^{l,l'} t \sim_{\prec} s : T$.

Proof: By induction on T . If $t \longrightarrow_{\text{picky}}^* \uparrow l''$ we are done, so let $t \longrightarrow_{\text{picky}}^* v$ and $s \longrightarrow_h^* w$ such that $v \approx_{\prec} w : T$.

$T = B$: So $c = \{x:B \mid t_2\}$ and $S_2 = \{x:B \mid s_2\}$. Moreover, $v = w = k \in \mathcal{K}_B$, since those are the only corresponding values at B .

We can step and see $\langle c \rangle^{l,l'} t \longrightarrow_{\text{picky}}^* \langle c, t_2\{x := k\}, k \rangle^l$. We know that $t_2\{x := k\} \sim_{\prec} s_2\{x := k\} : \text{Bool}$. There are two possibilities: either $t_2\{x := k\} \longrightarrow_{\text{picky}}^* \uparrow l''$ or both terms go to corresponding **Bools**. In the former case, the whole λ_C term goes to blame and we are done by Lemma 2.6.12. If both go to **false**, then the outer λ_C term evaluates to $\uparrow l$ and we are done by Lemma 2.6.12 again. If both go to **true**, then both outer terms go to k , and $k \approx_{\prec} k : B$.

$T = T_1 \rightarrow T_2$: So $c = x:c_1 \mapsto c_2$ and $S_1 = x:S_{11} \rightarrow S_{12}$ and $S_2 = x:S_{21} \rightarrow S_{22}$. Let $t' \sim_{\prec} s' : T_1$. If $t' \longrightarrow_{\text{picky}}^* \uparrow l''$ we are done by Lemm 2.6.12, so let $t' \longrightarrow_{\text{picky}}^* v'$ and $s' \longrightarrow_h^* w'$, where $v' \approx_{\prec} w' : T_1$. We want to prove $(\langle c \rangle^{l,l'} t) t' \sim_{\prec} s' : T_2$, which is true iff:

$$\langle c_2\{x := \langle c_1 \rangle^{l,l'} v'\} \rangle^{l,l'} (v (\langle c_1 \rangle^{l,l'} v')) \sim_{\prec} w w' : T_2$$

By the IH on $v \sim_{\prec} w' : T_1$ and $c_1 \sim_{\prec} S_{21} \Rightarrow S_{11} : T_1$, we have $\langle c_1 \rangle^{l,l'} v' \sim_{\prec} w' : T_1$. By definition, applying v and w yields related terms at T_2 . Since $\langle c_1 \rangle^{l,l'} v' \sim_{\prec} w' : T_1$, we have $c_2\{x := \langle c_1 \rangle^{l,l'} v'\} \sim_{\prec} S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^{l,l'} w'\} \Rightarrow S_{22}\{x := w'\} : T_2$. We can now apply the IH and see:

$$\langle c_2\{x := \langle c_1 \rangle^{l,l'} v'\} \rangle^{l,l'} (v (\langle c_1 \rangle^{l,l'} v')) \sim_{\prec} w w' : T_2 \quad \square$$

\square

2.6.15 Lemma [Contract/cast correspondence]: If $c \sim_{\prec} S_1 \Rightarrow S_2 : T$ and $t \sim_{\prec} s : T$ then $\langle c \rangle^{l,l'} t \sim_{\prec} \langle S_1 \Rightarrow S_2 \rangle^{l''} s : T$.

Proof: By induction on T . We reason via expansion (Lemma 2.6.11), showing that the initial terms reduce to corresponding terms.

$T = B$: So $c = \{x:B \mid t_1\}$, $S_1 = \{x:B \mid s_1\}$, and $S_2 = \{x:B \mid s_2\}$. Since $t \sim_{\prec} s : B$, we know that they either both reduce to $k \in \mathcal{K}_B$ or $t \rightarrow_{picky}^* \uparrow l'$. If the latter is the case, we are done. So suppose $t \rightarrow_{picky}^* k$ along with $s \rightarrow_h^* k$.

We can step our terms into active checks as follows, then:

$$\begin{array}{ccc} \langle \{x:B \mid t_1\} \rangle^{l,l'} t & \rightarrow_{picky}^* & \langle \{x:B \mid t_1\}, t_1 \{x := k\}, k \rangle^l \\ \langle \{x:B \mid s_1\} \Rightarrow \{x:B \mid s_2\} \rangle^l s & \rightarrow_h^* & \langle \{x:B \mid s_2\}, s_2 \{x := k\}, k \rangle^l \end{array}$$

By the contract/cast correspondence, we know that $t_1 \{x := k\} \sim_{\prec} s_2 \{x := k\} : \mathbf{Bool}$, so either $t_1 \{x := k\}$ goes to blame or both terms go to a \mathbf{Bool} together. In the former case, the outer λ_C term goes to blame and we are done by Lemma 2.6.12.. If they go to \mathbf{false} , then both the active check goes to $\uparrow l$ and we are done, again by Lemma 2.6.12. Finally, if they both go to \mathbf{true} , then both terms will evaluate to $k \in \mathcal{K}_B$, and $k \approx_{\prec} k : B$.

$T = T_1 \rightarrow T_2$: $c = x:c_1 \mapsto c_2$, $S_1 = x:S_{11} \rightarrow S_{12}$, and $S_2 = x:S_{21} \rightarrow S_{22}$. We know by inversion of the contract/cast relation that $c_1 \sim_{\prec} S_{21} \Rightarrow S_{11} : T_1$ and that for all l'' and $t \sim_{\prec} s : T_1$, $c_2 \{x := t\} \sim_{\prec} S_{12} \{x := \langle S_{21} \Rightarrow S_{11} \rangle^{l''} s\} \Rightarrow S_{22} \{x := s\} : T_2$. We want to prove that $\langle c \rangle^{l,l'} \sim_{\prec} \langle S_1 \Rightarrow S_2 \rangle^l s : T_1 \rightarrow T_2$. First, we can assume $t \rightarrow_{picky}^* v$ and $s \rightarrow_h^* w$ where $v \sim_{\prec} w : T_1 \rightarrow T_2$ —if not, both the contracted term goes to blame and we are done by Lemma 2.6.12.

We show that the decomposition of the contract and cast terms correspond for all inputs. Let $t' \sim_{\prec} s' : T_1$. Again, we can assume that they reduce to $v' \sim_{\prec} w' : T_1$, or else we are done by blame lifting in λ_C . On the λ_C side, we have

$$(\langle c \rangle^{l,l'} t) t' \rightarrow_{picky}^* \langle c_2 \{x := \langle c_1 \rangle^{l',l} v'\} \rangle^{l,l'} (v (\langle c_1 \rangle^{l',l} v'))$$

In λ_H , we find

$$\begin{array}{l} \langle S_1 \Rightarrow S_2 \rangle^{l''} s) s' \rightarrow_h^* \\ \langle S_{12} \{x := \langle S_{21} \Rightarrow S_{11} \rangle^{l''} w\} \Rightarrow S_{22} \{x := w'\} \rangle^{l''} (w (\langle S_{21} \Rightarrow S_{11} \rangle^{l''} w')) \end{array}$$

By the IH, we know that $\langle c_1 \rangle^{l',l} v' \sim_{\prec} \langle S_{21} \Rightarrow S_{11} \rangle^{l''} w' : T_1$. Since $v \sim_{\prec} w : T_1 \rightarrow T_2$, we have $v (\langle c_1 \rangle^{l',l} v') \sim_{\prec} w (\langle S_{21} \Rightarrow S_{11} \rangle^{l''} w') : T_2$. By Lemma 2.6.14, $\langle c_1 \rangle^{l',l} v' \sim_{\prec} w' : T_1$. We can then see that $c_2 \{x := \langle c_1 \rangle^{l',l} v'\} \sim_{\prec} S_{12} \{x := \langle S_{21} \Rightarrow S_{11} \rangle^{l''} w'\} \Rightarrow S_{22} \{x := w'\} : T_2$. By the IH, we therefore have

$$\begin{array}{l} \langle c_2 \{x := \langle c_1 \rangle^{l',l} v'\} \rangle^{l,l'} (v (\langle c_1 \rangle^{l',l} v')) \sim_{\prec} \\ \langle S_{12} \{x := \langle S_{21} \Rightarrow S_{11} \rangle^{l''} w'\} \Rightarrow S_{22} \{x := w'\} \rangle^{l''} w (\langle S_{21} \Rightarrow S_{11} \rangle^{l''} w') : T_2 \quad \square \end{array}$$

□

2.6.16 Theorem [Behavioral correspondence]:

1. If $\Delta \vdash s : S$ then $\lfloor \Delta \rfloor \vdash \psi(s) \sim_{\prec} s : \lfloor S \rfloor$.
2. If $\Delta \vdash S_1$ and $\Delta \vdash S_2$, where $\lfloor S_1 \rfloor = \lfloor S_2 \rfloor = \lfloor S \rfloor$, then $\lfloor \Delta \rfloor \vdash \psi^l(S_1, S_2) \sim_{\prec} S_1 \Rightarrow S_2 : \lfloor S \rfloor$.

Proof: By an induction similar to the proof Theorem 2.5.18. □ □

2.6.17 Theorem [Type preservation for ψ]: For programs with no active checks, if $\vdash \Delta$, then:

1. If $\Delta \vdash s : S$ then $\lfloor \Delta \rfloor \vdash \psi(s) : \lfloor S \rfloor$.
2. If $\Delta \vdash S_1, \Delta \vdash S_2$, where $\lfloor S_1 \rfloor = \lfloor S_2 \rfloor = T$, then $\lfloor \Delta \rfloor \vdash^{l,l'} \psi^l(S_1, S_2) : T$.

Proof: By an induction similar to the proof Theorem 2.5.20. □ □

Here is an example where a λ_H term reduces to a value while its ψ -image in picky λ_C term reduces to blame. As before, this counterexample uses an *abusive* contract: a higher-order contract where the codomain puts a contradictory requirement on the domain. Here, the contradiction is that f claims to return a non-zero value, but the codomain requires that it return 0.

$$\begin{aligned}
S_1 &= f:S_{11} \rightarrow S_{12} \\
&= f:(x:\lfloor \text{Int} \rfloor \rightarrow \{y:\text{Int} \mid \text{nonzero } y\}) \rightarrow \lfloor \text{Int} \rfloor \\
S_2 &= f:S_{21} \rightarrow S_{22} \\
&= f:(x:\lfloor \text{Int} \rfloor \rightarrow \lfloor \text{Int} \rfloor) \rightarrow \{z:\text{Int} \mid f \ 0 = 0\} \\
c &= \psi^l(S_1, S_2) \\
&= f:\psi^l(S_{21}, S_{11}) \mapsto \psi^l(S_{12}\{f := \langle S_{21} \Rightarrow S_{11} \rangle^l f\}, S_{22}) \\
&= f:(x:\{x:\text{Int} \mid \text{true}\} \mapsto \{y:\text{Int} \mid \text{nonzero } y\}) \mapsto \{z:\text{Int} \mid f \ 0 = 0\}
\end{aligned}$$

Let $w = (\lambda f:(x:\{x:\text{Int} \mid \text{true}\} \rightarrow \{y:\text{Int} \mid \text{nonzero } y\}). 0)$ and $w' = (\lambda x:\{x:\text{Int} \mid \text{true}\}. 0)$. The term is well typed: we can show $\emptyset \vdash w : S_1$ and $\emptyset \vdash w' : S_{21}$. Therefore $\emptyset \vdash (\langle S_1 \Rightarrow S_2 \rangle^l w) \ w' : S_{22}\{f := w'\}$. Translating, we find

$$\psi((\langle S_1 \Rightarrow S_2 \rangle^l w) \ w') = (\langle \psi^l(S_1, S_2) \rangle^{l,l} \psi(w)) \ \psi(w') = (\langle c \rangle^{l,l} \lambda f:\text{Int}. 0) \ \lambda x:\text{Int}. 0.$$

On the one hand $(\langle S_1 \Rightarrow S_2 \rangle^l w) \ w' \longrightarrow_h^* 0$, while $(\langle c \rangle^{l,l} \lambda f:\text{Int}. 0) \ \lambda x:\text{Int}. 0 \longrightarrow_{\text{picky}}^* \uparrow l$. This means we cannot hope to use ψ as an exact correspondence between λ_H and picky λ_C . (Removing the extra cast ψ inserts into S_{12} does not affect the example, since ψ ignores S_{12} here.) For example,

$$\psi^l(\{z:\text{Int} \mid \text{true}\}\{f := \langle S_{21} \Rightarrow S_{11} \rangle^l f\}, \{z:\text{Int} \mid f \ 0 = 0\}) = \{x:B \mid \psi(f \ 0 = 0)\}.$$

2.6.3 Alternative calculi

There are three alternative calculi I have not considered here: indy λ_C [22], superpicky λ_H , and nonterminating calculi. I describe them in detail below, but I leave them as future work.

Dimoulas et al. [22] add a third blame label to λ_C , representing the contract itself; I write it here as a subscript. They accordingly change the picky $E_CDECOMP$ rule:

$$(\langle x:c_1 \mapsto c_2 \rangle_{l''}^{l,l'} v_1) v_2 \longrightarrow_{indy} \langle c_2\{x := \langle c_1 \rangle_{l''}^{l,l'} v_2\} \rangle_{l''}^{l,l'} (v_1 (\langle c_1 \rangle_{l''}^{l,l'} v_2))$$

In the substitution in the codomain, note that the blame label on the domain contract uses the contract's blame label l'' . The intuition here is that any problem arising in c_2 is in the contract's context (label l''), not the original negative context (label l'). I conjecture (but have not proven) that indy λ_C is in the same position on the axis of blame as picky λ_C . We should only need to change the labels on the contracts ϕ inserts to have an exact correspondence; however, ψ will remain inexact.

Superpicky λ_H reworks the $F_CDECOMP$ rule in an attempt to harmonize λ_H and picky λ_C semantics:⁷

$$\begin{aligned} & (\langle x:S_{11} \rightarrow S_{12} \Rightarrow x:S_{21} \rightarrow S_{22} \rangle^l w_1) w_2 \rightsquigarrow_h \\ & \langle S_{12}\{x := \langle S_{21} \Rightarrow S_{11} \rangle^l w_2\} \Rightarrow S_{22}\{x := \langle S_{11} \Rightarrow S_{21} \rangle^l (\langle S_{21} \Rightarrow S_{11} \rangle^l w_2)\} \rangle^l \\ & (w_1 (\langle S_{21} \Rightarrow S_{11} \rangle^l w_2)) \end{aligned}$$

This seems to resolve the problem with ψ into picky λ_C , but it poses problems in the proof of semantic type soundness for λ_H : how do $S_{22}\{x := w_2\}$ and $S_{22}\{x := \langle S_{11} \Rightarrow S_{21} \rangle^l (\langle S_{21} \Rightarrow S_{11} \rangle^l w_2)\}$ relate?

Finally, I have been careful to ensure that all of the calculi are strongly normalizing. I do not believe this to be essential, though I would have to change the logical relations— λ_H 's type semantics and the correspondences—to account for nontermination. Step-indexing [2] should suffice.

2.7 Conclusion

The main contributions of the work in this chapter are (1) the dependent translations ϕ and ψ and their properties, and (2) the formulation and metatheory of dependent λ_H . (Dependent λ_C is not a contribution on its own: many similar systems have been studied, and in any case its properties are simple.) The nondependent part of our ϕ translation essentially coincides with the one studied by Gronski and Flanagan [36], and our behavioral correspondence theorem is essentially the same as theirs (though our proof is substantially different—and more detailed). Our ψ translation completes their story for the nondependent case, establishing a tight connection between the systems. The full dependent forms of ϕ and ψ studied in this chapter are novel, as

⁷This idea is due to Jeremy Siek (personal communication, January 2010).

is the observation that the correspondence between the latent and manifest worlds is more problematic in this setting.

We can faithfully encode dependent λ_H into λ_C —the behavioral correspondence is tight. λ_H 's `F_CDECOMP` rule forces us to accept a weaker behavioral correspondence when encoding λ_C into λ_H , so I conclude that the manifest and latent approaches are *not* equivalent in the dependent case. I do find, however, that the two approaches are entirely inter-encodable in the nondependent restriction.

Chapter 3

Polymorphic manifest contracts

And even the Abstract Entities
Circumambulate her charm;
But our lot crawls between dry ribs
To keep our metaphysics warm.

Whispers of Immortality
T.S. Eliot

Contracts are particularly useful at interfaces, allowing programmers to specify expectations and guarantees in code, rather than merely in comments. For example, consider an abstract datatype (ADT) modeling the natural numbers:

$$\mathbf{NAT} : \exists\alpha. (\mathbf{zero} : \alpha) \times (\mathbf{succ} : (\alpha \rightarrow \alpha)) \times (\mathbf{iszero} : (\alpha \rightarrow \mathbf{Bool})) \times (\mathbf{pred} : (\alpha \rightarrow \alpha))$$

The type **NAT** (with its existential quantification and products) can be encoded straightforwardly in System F. It is an *abstract* datatype because the actual representation of α is hidden: users of **NAT** interact with it through the constructors and operations provided. The **zero** constructor represents 0; the **succ** constructor takes a natural and produces its successor. The predicate **iszero** determines whether a given natural is zero. The **pred** operation takes a natural number and returns its predecessor. There are many ways to define an ADT matching **NAT**'s signature, but we can suppose here that **NAT** is implemented using the standard Church encoding

$$\alpha = \forall\beta. \beta \rightarrow (\beta \rightarrow \beta) \rightarrow \beta.$$

Under the standard definition, **pred zero** \longrightarrow^* **zero**, even though the mathematical natural-number predecessor operation isn't defined for zero. That is, the **pred** function is only *mostly* correct: **pred** n will return the same result as the mathematical natural-number predecessor so long as $n \neq 0$. Using contracts, we can specify this constraint explicitly, as we did in the introduction (Section 1.2):

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

In addition to requirements, we can also express guarantees. For example, the interface could indicate that successors are non-zero by giving `succ` the type $\alpha \rightarrow \{x:\alpha \mid \text{not}(\text{iszero } x)\}$. Both of these encodings rely on dependent types: we need to be able to refer back to `iszero` in the contracts for `pred` and `succ`. Fortunately, dependent functions can encode dependent sums:

$$\begin{aligned} (x : T_1) \times T_2 &= \forall \alpha. (x:T_1 \rightarrow T_2 \rightarrow \alpha) \rightarrow \alpha \\ (e_1, e_2) &= \Lambda \alpha. \lambda f:(x:T_1 \rightarrow T_2 \rightarrow \alpha). f e_1 e_2 \end{aligned}$$

Finally, note that to put a refinement type on the codomain of `succ`, we would have to rearrange the dependent sum above so `succ` came after `iszero`.

This chapter presents a new core calculus for contracts and polymorphism, F_H (so named since it extends the manifest-contract core calculus, λ_H , with the type abstractions/impredicative polymorphism of System F). The example I’ve just given puts contracts inside types, so it is manifest—a sensible choice for combining contracts and abstract datatypes. Abstract datatypes already use the type system to mediate access to abstractions; manifest contracts allow types to exercise a still finer grained control.

F_H is a polymorphic manifest contract calculus. In designing F_H , I take a new metatheoretical approach, since earlier manifest calculi don’t scale up to polymorphism easily—for two reasons.

First, the metatheory for existing calculi has been too complicated to extend easily [34, 44]. Previous approaches proved semantic type soundness, using denotational techniques. The denotational semantics used are harder to scale than standard syntactic methods (i.e., progress and preservation). I explain the complexity in detail in related work (Section 5.2), but, briefly, the problem is subtyping. In both Greenberg et al. and Knowles and Flanagan, subtyping between refinement types is what ultimately requires denotational semantics. The denotational semantics is used to avoid a dangerous circularity in the mutually recursive definition of the typing and subtyping judgments. The heart of my new technique is to replace subtyping with a simpler conversion relation, which allows for a simpler, more scalable, syntactic metatheory. Eliminating subtyping doesn’t lose useful reasoning principles: in Section 3.4, I define subtyping *post facto* and recover an “upcast” lemma from Knowles and Flanagan [44] showing that casts from subtypes to supertypes never fail.

Second, look again at our `NAT` abstract datatype. If we pare away the ADT’s type abstraction to the underlying Church encoding, we find

$$\{x:\alpha \mid \text{not}(\text{iszero } x)\} = \{x:\forall \beta. \beta \rightarrow (\beta \rightarrow \beta) \rightarrow \beta \mid \text{not}(\text{iszero } x)\}.$$

If we can put contracts on type variables, then the contracts on our ADTs will need to allow refinements on function and forall types. These so-called “general refinements” aren’t possible in existing manifest calculi, which restrict refinements to base types like `Bool` and `Int`. More specifically, earlier metatheory retains type preservation

by assigning constants most-specific types, but this approach doesn’t scale to refinements of functions. My new metatheory for F_H takes a different approach to type preservation that supports general refinements. (For more detail, see the discussion of `T_EXACT` in Section 3.2.)

In Section 3.1, I develop longer and more detailed examples. I describe F_H and prove type soundness in Section 3.2. I prove parametricity in Section 3.3 and the upcast lemma in Section 3.4.

3.1 Examples

In this section, I offer some longer examples, in order to develop an intuition for how contracts work in general and with polymorphism in particular. I introduce the mechanisms used to enforce contracts first; then I look at the `NAT` datatype in greater depth; I offer a “library as a language” as a final example.

Like other manifest calculi, F_H uses casts to dynamically enforce contracts. Applying a *cast* $\langle T_1 \Rightarrow T_2 \rangle^l$ to a value of type T_1 will ensure that the value behaves like a T_2 . I call T_1 the source type and T_2 the target type. The l superscript is a *blame label*, used to differentiate between different casts and identify the source of failures. These failures are indicated by *blame*, an uncatchable exception with a blame label attached; I write this exception $\uparrow l$ and pronounce it “blame l ”.

At base types, casts do one of two things: they either return the value they were applied to, or “raise blame”. For example, consider a cast from integers `Int` to positive integers, $\{x:\text{Int} \mid x > 0\}$. I write this cast $\langle \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rangle^l$, picking an arbitrary label l . If we apply this cast to 5, we expect to get 5 back, since $5 > 0$. That is,

$$\langle \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rangle^l 5 \longrightarrow^* 5.$$

On the other hand, suppose we apply the same cast to 0. This cast fails, since 0 is certainly not greater than itself. When the cast fails, it will raise blame with its label:

$$\langle \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rangle^l 0 \longrightarrow^* \uparrow l.$$

Checking predicate contracts with casts is easy. When we run $\langle T \Rightarrow \{x:T \mid e\} \rangle^l v$, we simply check if v satisfies the predicate, i.e., whether $e[v/x] \longrightarrow^* \text{true}$. If it does, then the entire application goes to v ; if not, then the program aborts, raising $\uparrow l$. When checking predicate contracts, only the target type matters—the type system will guarantee that whatever value we have is well typed at the source type, i.e., satisfies any predicates it has. For example, we’ll always have

$$\langle \{x:\text{Int} \mid x > 0\} \Rightarrow \text{Int} \rangle^l v \longrightarrow v$$

immediately, for all values v . The application will only be well typed if v actually *is* a positive number, but, operationally, the left-hand side doesn’t matter from the perspective of function refinements.

Unlike casts between base types and refinements, casts between function types aren't checked immediately. Instead, casts at function types wrap their argument up, decomposing the function cast into two parts: one cast for the domain and one for the codomain. In this way, checking is deferred until the cast function is called. Suppose we have a function f of type $\text{Int} \rightarrow \text{Int}$ and we want to ensure that maps positives to numbers greater than 5, casting it to $\{x:\text{Int} \mid x > 0\} \rightarrow \{y:\text{Int} \mid y > 5\}$. The cast decomposes as follows:

$$\langle \text{Int} \rightarrow \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rightarrow \{y:\text{Int} \mid y > 5\} \rangle^l f \quad \longrightarrow \\ \lambda x:\{x:\text{Int} \mid x > 0\}. (\langle \text{Int} \Rightarrow \{y:\text{Int} \mid y > 5\} \rangle^l (f (\langle \{x:\text{Int} \mid x > 0\} \Rightarrow \text{Int} \rangle^l x))).$$

Both casts in the wrapper have the same blame label as the original cast. Notice that the domains of the function types are treated contravariantly; note the inner term in the wrapped term: $f (\langle \{x:\text{Int} \mid x > 0\} \Rightarrow \text{Int} \rangle^l x)$. In this case, the domain cast $\langle \{x:\text{Int} \mid x > 0\} \Rightarrow \text{Int} \rangle^l x$ will never fail: every positive integer is also an integer. The codomain cast is covariant: $\langle \text{Int} \Rightarrow \{y:\text{Int} \mid y > 5\} \rangle^l$ checks that f returns a number greater than 5. Since not every integer is greater than 5, this cast will fail if f returns a number less than or equal to 5.

Let's consider a few concrete choices of the function f . First, we cast the identity function to $\{x:\text{Int} \mid x > 0\} \rightarrow \{y:\text{Int} \mid y > 5\}$. We can safely apply the cast function to six:

$$\langle \text{Int} \rightarrow \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rightarrow \{y:\text{Int} \mid y > 5\} \rangle^l (\lambda x:\text{Int}. x) 6 \longrightarrow^* 6.$$

In general, the identity function does not take positives to numbers greater than five. But when we apply the cast function to 6, which happens to satisfy the codomain contract, no blame is raised. However, when we apply the cast identity function to a value that doesn't satisfy the codomain contract, say 2, blame will be raised:

$$\langle \text{Int} \rightarrow \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rightarrow \{y:\text{Int} \mid y > 5\} \rangle^l (\lambda x:\text{Int}. x) 2 \longrightarrow^* \uparrow l.$$

Contrast this with the sorts of static checks offered by type systems: contract systems raise blame only when a violation is *detected*; type systems are usually conservative, signaling errors when a violation is *possible*.

Some functions will *never* work. If we cast the constant zero function to $\{x:\text{Int} \mid x > 0\} \rightarrow \{y:\text{Int} \mid y > 5\}$, it will raise blame for any value:

$$\langle \text{Int} \rightarrow \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rightarrow \{y:\text{Int} \mid y > 5\} \rangle^l (\lambda x:\text{Int}. 0) 6 \longrightarrow^* \uparrow l.$$

The constant zero function, which *never* returns a value satisfying the codomain contract $\{y:\text{Int} \mid y > 5\}$. No matter what value we apply the cast function to, it will always raise blame. Cast at function types are deferred, though: if we never call the cast function, it never has the opportunity to raise blame. Since casts check contracts dynamically, they only detect errors in parts of the program that are explored at runtime.

F_H 's type system rules out directly applying a function with domain type $\{x:\text{Int} \mid x > 0\}$ to 0. It is an important property of F_H that 0 doesn't have type $\{x:\text{Int} \mid x > 0\}$! One can try to cast 0 from Int to $\{x:\text{Int} \mid x > 0\}$, but this will always fail:

$$\langle \text{Int} \rightarrow \text{Int} \Rightarrow \{x:\text{Int} \mid x > 0\} \rightarrow \{y:\text{Int} \mid y > 5\} \rangle^l \quad \longrightarrow^* \uparrow l'$$

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

Finally, F_H supports dependent function types. For example, the type $x:\text{Int} \rightarrow \{y:\text{Int} \mid y > x\}$ is inhabited by functions over integers which produce results greater than their inputs. Dependent functions allow for very precise specifications. For example, $x:\text{Float} \rightarrow \{y:\text{Float} \mid |y^2 - x| < \epsilon\}$ specifies the square-root function. The exact unwinding rule for dependent functions is slightly subtle—see the discussion of `E_FUN` in Section 3.2.

With the basic operational ideas behind manifest contracts established, I offer two examples of contracts for abstract datatypes.

Contracts for abstract datatypes

The standard polymorphic encodings of existential and product types transfer over to F_H 's System F-style impredicative polymorphism without a problem. Indeed, 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.

$$(x : T_1) \times T_2 = \forall \alpha. (x:T_1 \rightarrow T_2 \rightarrow \alpha) \rightarrow \alpha$$

$$(e_1, e_2)_{(x:T_1) \times T_2} = \Lambda \alpha. \lambda f:(x:T_1 \rightarrow T_2 \rightarrow \alpha). f e_1 e_2$$

As in pure System F, we need to specify types on the encoding of pairs—I'll omit these types when they are clear from context. I also use record projection notation and field names, to make the example clearer.

Let's return to our simple example combining contracts and polymorphism—an abstract datatype of natural numbers.

$$\text{NAT} : \exists \alpha. (\text{zero} : \alpha) \times (\text{succ} : (\alpha \rightarrow \alpha)) \times (\text{iszero} : (\alpha \rightarrow \text{Bool})) \times$$

$$(\text{pred} : \{x:\alpha \mid \text{not } (\text{iszero } x)\} \rightarrow \alpha)$$

I omit the implementation, a standard Church encoding, where $\alpha = \forall \beta. \beta \rightarrow (\beta \rightarrow \beta) \rightarrow \beta$. The constructors `zero` and `succ` are standard; the operator `iszero` determines whether a natural is zero; the operator `pred` yields the predecessor. As we saw above, the standard representation the naturals is *inadequate* with respect to the mathematical natural numbers, in particular with respect to `pred`. In math, `pred zero` is undefined, but the implementation will return `zero`. The `NAT` interface hides our encoding of the naturals behind an existential type, but to ensure adequacy, we want

to ensure that `pred` is only ever applied to terms of type $\{x:\alpha \mid \text{not}(\text{iszero } x)\}$. With contracts, this is easy enough:

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

Recall that in the Church encoding, α will be instantiated with $\forall\beta.\beta \rightarrow (\beta \rightarrow \beta) \rightarrow \beta$. So the refinement $\{x:\alpha \mid \text{not}(\text{iszero } x)\}$ in the new type of `pred` is a refinement of a polymorphic function type. These general refinement types are available in F_H , but they were not in earlier manifest calculi.

To see why this more specific type for `pred` is useful, consider the following expression.

$$\text{unpack NAT} : \exists\alpha. I \text{ as } \alpha, n \text{ in } n.\text{iszero} (n.\text{pred} (n.\text{zero})) : \text{Bool}$$

Here I is the interface we specified for `NAT`. We’ve “unpacked” the ADT to make its type available as α ; its constructors and operators are in the dependent pair n . We then ask if the predecessor of 0 is 0, running $n.\text{iszero} (n.\text{pred} (n.\text{zero}))$. The inner application is *not well typed*! We have that $\text{zero} : \alpha$, but the domain type of `pred` is $\{x:\alpha \mid \text{not}(\text{iszero } x)\}$. In order to make the application well typed, we must insert a cast:

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

Naturally, this cast will ultimately raise $\uparrow l$, because $\text{not} (n.\text{iszero } n.\text{zero}) \longrightarrow^* \text{false}$.

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 datatype, consider the extension of the interface with a subtraction operation, `sub`, and a binary “less than or equal” operator, `leq`. Natural number subtraction `sub` x y is defined only when `leq` y x ; we can specify this pre-condition as before, by refining the type of `sub`’s second argument. But subtraction comes with a guarantee, as well: `sub` x y will always be less than or equal to x . That is, the result `sub` x y has the refined type $\{z:\alpha \mid \text{leq } z \ x\}$. We can specify both of these facts with 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’s contract requires that `sub`’s second argument is less than or equal to the first; the contract requires that `sub` returns a result that is less than or equal to the first argument.

How can we write an implementation to meet this interface? By putting casts in the implementations. We can impose the contracts on `pred` and `sub` when we “pack up” the implementation `NAT`. Writing `nat` for the type of the Church encoding $\forall\beta.\beta \rightarrow (\beta \rightarrow \beta) \rightarrow \beta$, we define the exported `pred` and `sub` in terms of the standard, unrefined implementations, `pred'` and `sub'`.

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

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, because we're casting out of a refinement. 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.

I can sum up the situation for contracts in abstract datatype interfaces as follows: the positive parts of the interface type are checked by the datatype's contract and can raise blame—these parts are the responsibility of the ADT's implementation; the negative parts of the interface type are not checked by the datatype's contract—these parts are the responsibility of the ADT's clients. Distributing obligations in this way recalls Findler and Felleisen's seminal idea of client and server blame [26].

Contracts as type systems

Inasmuch as abstract datatypes are little languages, contracts for abstract datatypes are like type systems for little languages. In this section, I develop a toy combinator language of string transducers, which specify mappings between regular languages.¹ A type system for the transducer combinators translates naturally to a contracted abstract datatype implementation.

Each transducer $t : \mathcal{L}_1 \hookrightarrow \mathcal{L}_2$ maps from a (regular) domain $\text{dom } t = \mathcal{L}_1$ to a (regular) range $\text{rng } t = \mathcal{L}_2$.

$$\begin{aligned} \mathcal{S} &::= \text{strings} \\ \mathcal{L} &::= \text{regular expressions} \\ t &::= \text{copy } \mathcal{L} \mid \text{delete } \mathcal{L} \mid \text{concat } t_1 t_2 \mid \text{seq } t_1 t_2 \end{aligned}$$

Let ϵ be the empty string, and let \cdot be string concatenation. The combinators compose to specify transducers. For this example, we only need two primitive combinators: `copy` \mathcal{L} , maps a string in the language \mathcal{L} to itself; and `delete` \mathcal{L} which maps a string in the language \mathcal{L} to the empty string, ϵ . We can combine transducers in two ways: `concat` $t_1 t_2$ runs t_1 on the first part of the input and t_2 on the second; `seq` $t_1 t_2$ runs t_2 on the output t_1 . We can define a semantics for transducers easily enough, with a function $\text{run } t$ that takes strings in $\text{dom } t$ to strings in $\text{rng } t$.

The `copy` \mathcal{L} transducer has the simplest semantics: it just copies strings in the given language, \mathcal{L} . Its typing rule is straightforward.

$$\text{run}(\text{copy } \mathcal{L}) \mathcal{S} = \mathcal{S} \qquad \text{copy } \mathcal{L} : \mathcal{L} \hookrightarrow \mathcal{L}$$

The `delete` \mathcal{L} transducer deletes a string in the language \mathcal{L} ; its typing rule indicates that its range is the language containing only the empty string, $\{\epsilon\}$.

¹This is effectively a unidirectional version of Boomerang [12].

$\text{run}(\text{delete } \mathcal{L}) \mathcal{S} = \epsilon$

$\text{delete } \mathcal{L} : \mathcal{L} \hookrightarrow \{\epsilon\}$

The $\text{concat } t_1 t_2$ combinator splits its input between its two sub-transducers. Splitting up the input makes the semantics somewhat subtle: in general, $\mathcal{S} \in \text{dom } t_1 \cdot \text{dom } t_2$ does not imply that there is a *unique* way to split \mathcal{S} . When two regular languages always split uniquely, I say they are *unambiguously splittable*, written $\mathcal{L}_1 \cdot^! \mathcal{L}_2$. Unambiguous splittability of regular languages is decidable [12]; if we only concatenate transducers with unambiguously splittable domains, then the run function will be unambiguous.

$$\begin{array}{l} \text{run}(\text{concat } t_1 t_2) (\mathcal{S}_1 \cdot \mathcal{S}_2) = \\ \quad \text{run } t_1 \mathcal{S}_1 \cdot \text{run } t_2 \mathcal{S}_2 \\ \text{where } \mathcal{S}_i \in \text{dom } t_i \end{array} \qquad \frac{t_1 : \mathcal{L}_{11} \hookrightarrow \mathcal{L}_{12} \quad t_2 : \mathcal{L}_{21} \hookrightarrow \mathcal{L}_{22} \quad \mathcal{L}_{11} \cdot^! \mathcal{L}_{21}}{\text{concat } t_1 t_2 : \mathcal{L}_{11} \cdot \mathcal{L}_{21} \hookrightarrow \mathcal{L}_{12} \cdot \mathcal{L}_{22}}$$

Finally, the $\text{seq } t_1 t_2$ combinator runs t_2 on the output of t_1 . The typing rule requires that the two sub-transducers match: $\text{rng } t_1$ and $\text{dom } t_2$ must be the same language.

$$\text{run}(\text{seq } t_1 t_2) \mathcal{S} = \text{run } t_2 (\text{run } t_1 \mathcal{S}) \qquad \frac{t_1 : \mathcal{L}_1 \hookrightarrow \mathcal{L}_2 \quad t_2 : \mathcal{L}_2 \hookrightarrow \mathcal{L}_3}{\text{seq } t_1 t_2 : \mathcal{L}_1 \hookrightarrow \mathcal{L}_3}$$

In order to facilitate programming with these transducers, we may try to embed these combinators inside of a functional programming. It is rather difficult to generalize this combinator language—typing a lambda calculus with string regular expression types is not easy [72, 9, 10]. It is easy, however, to define an abstract datatype offering these operations, assuming we have a type **String** of strings and **Regex** of regular expressions, with appropriate decision procedures for unambiguous splittability and equality.

$$\begin{aligned} \text{TRANS} : \exists \alpha. & (\text{dom} : \alpha \rightarrow \text{Regex}) \times (\text{rng} : \alpha \rightarrow \text{Regex}) \times \\ & (\text{run} : (\alpha \rightarrow \text{String} \rightarrow \text{String})) \times \\ & (\text{copy} : \text{Regex} \rightarrow \alpha) \times (\text{delete} : (\text{Regex} \rightarrow \alpha)) \times \\ & (\text{concat} : \alpha \rightarrow \alpha \rightarrow \alpha) \times (\text{seq} : (\alpha \rightarrow \alpha \rightarrow \alpha)) \end{aligned}$$

There is a problem, though: this datatype will let us write nonsensical combinators. In particular, we can give concat transducers that don't have unambiguously splittable domains, or we can give seq transducers which don't match up. For example, suppose $\epsilon \notin \mathcal{L}$ and $\mathcal{S} \in \mathcal{L}$. Let $t = \text{seq}(\text{delete } \mathcal{L})(\text{copy } \mathcal{L})$. Then:

$$\text{run}(\text{seq}(\text{delete } \mathcal{L})(\text{copy } \mathcal{L})) \mathcal{S} = \epsilon$$

Since $\epsilon \notin \mathcal{L}$, this means that run took a value in $\text{dom } t = \mathcal{L}$ and produced a value outside of $\text{rng } t = \mathcal{L}$.

We are in a difficult position: we have a little language and a type system. But scaling our transducer language's type system up to the lambda calculus increases

complexity, and distracts us from what we’d like to be doing—writing transducer programs!

Contracts offer a middle way. By putting contracts on the TRANS interface, we can ensure that all transducers are well formed.

$$\begin{aligned} \text{TRANS} : \exists \alpha. & (\text{dom} : \alpha \rightarrow \text{Regex}) \times (\text{rng} : \alpha \rightarrow \text{Regex}) \times \\ & (\text{run} : (t : \alpha \rightarrow \{x : \text{String} \mid x \in \text{dom } t\} \rightarrow \{x : \text{String} \mid x \in \text{rng } t\})) \times \\ & (\text{copy} : \text{Regex} \rightarrow \alpha) \times (\text{delete} : (\text{Regex} \rightarrow \alpha)) \times \\ & (\text{concat} : (t_1 : \alpha \rightarrow \{t_2 : \alpha \mid \text{splittable}(\text{dom } t_1) (\text{dom } t_2)\} \rightarrow \alpha)) \times \\ & (\text{seq} : (t_1 : \alpha \rightarrow \{t_2 : \alpha \mid \text{rng } t_1 = \text{dom } t_2\} \rightarrow \alpha)) \end{aligned}$$

The TRANS abstract datatype defines an embedded domain-specific language—with its own domain-specific type system. For example, `concat` will only accept transducers with splittable domains; `seq` will only sequence transducers that match up. The interface given here is only one of many: it checks inputs but not outputs. For example, we could ensure that `concat` has our intended behavior by giving its codomain the type $\{t_3 : \alpha \mid (\text{dom } t_3 = (\text{dom } t_1) \circ (\text{dom } t_2)) \wedge (\text{rng } t_3 = (\text{rng } t_1) \circ (\text{rng } t_2))\}$, where \circ denotes regular expression concatenation.

The checks on `run` and `seq` are easy enough to build into our TRANS implementation. But there is a distinct advantage to making the contracts explicit in the interface type: the type system will keep track of unambiguous splittability checks. Programmers can track relevant information in refinements in client modules, and we can statically eliminate redundant checks (see Section 3.4).

In general, contracting abstract datatype interfaces allows for library designers to extend the language’s type system with library-specific constraints. Clients then have two choices: propagate the library’s contracts through their code, possibly avoiding redundant checks; or ignore the contracts within their own code, allowing the checks to happen whenever they call into the library. Either way, the library’s users can rest assured that the contracts will guarantee the safety properties the library designers desired.

If programmers are careful to program in a “cover your ass” (CYA) style, wherein each library’s interface uses contracts that are strong enough to guarantee that other libraries’ contracts are satisfied, then error messages greatly improve. When libraries are stacked in a hierarchy several levels deep, CYA contracts in interfaces give programmers error messages earlier and at a higher level of abstraction.

As a final note before we begin the technical content: the foregoing is the current implementation strategy for Boomerang [12], a language of bidirectional string transducers called *lenses*. The semantic constraints on Boomerang combinators are decidable, but combining Boomerang’s typing rules with the lambda calculus would be cumbersome—a hard open problem. Boomerang is a complex language, and I believe there is no room in the “complexity budget” for a statically checked type system: lenses can already be difficult to program with and understand, and the complicated constraints necessary for type checking will only add more to the programmer’s burden. Instead, the Boomerang primitives have contracts that ensure

Terms and contexts

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

$$\Gamma ::= \emptyset \mid \Gamma, x:T \mid \Gamma, \alpha$$
Terms, values, results, and evaluation contexts

$$e ::= x \mid k \mid \text{op}(e_1, \dots, e_n) \mid \lambda x:T. e \mid \Lambda\alpha. e \mid e_1 e_2 \mid e T \mid \langle T_1 \Rightarrow T_2 \rangle^l \mid \uparrow l \mid \langle \{x:T \mid e_1\}, e_2, v \rangle^l$$

$$v ::= k \mid \lambda x:T. e \mid \Lambda\alpha. e \mid \langle T_1 \Rightarrow T_2 \rangle^l$$

$$r ::= v \mid \uparrow l$$

$$E ::= [] e_2 \mid v_1 [] \mid [] T \mid \langle \{x:T \mid e\}, [], v \rangle^l \mid \text{op}(v_1, \dots, v_{i-1}, [], e_{i+1}, \dots, e_n)$$
Figure 3.1: Syntax for F_H

that they produce sane bidirectional transformations. The Boomerang libraries built atop these primitives have contracts as well, in a CYA style. Even without optimizations to reduce the number of dynamic checks, this improvement in error handling has proved quite useful. Contracts are particularly suited to the phased nature of Boomerang, since the contracts on Boomerang’s lens combinators are “quasi-static”. Lenses are constructed only once and then run many times. Running a lens merely requires checking regular language membership, so higher cost one-time checks can be amortized over many lens runs.

3.2 Defining F_H

The syntax of F_H is given in Figure 3.1. 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. As usual, we write $T_1 \rightarrow T_2$ for $x:T_1 \rightarrow T_2$ when x does not appear free in T_2 . 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 , I fix a set \mathcal{K}_B of the constants in that type; I require the typing rules for constants and the typing and evaluation rules for operations to respect this set. I 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, several monomorphic first-order operations `op` (i.e., destructors of one or more base-type arguments), term and type abstractions, and term and type applications. The second line offers the standard constructs of a manifest contract calculus [28, 34, 44], 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 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. The type in the first position of an active check isn't necessary for the operational semantics, but I keep it around as a technical aid to type soundness. 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 my present scope.)

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 3.2.25 and 3.2.27 below—will show that evaluation produces a result, but not necessarily a value.) In Chapter 2, casts between function types applied to values were themselves considered values. I make the other choice here: excluding applications from the possible syntactic forms of values simplifies our inversion lemmas. Instead, casts between function types will η -expand their argument. This makes the notion of “function proxy” explicit: the cast semantics adds many new closures. I address this issue in Chapter 4, where I make manifest contracts space-efficient.

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 [28, 34, 36, 44, 51]); 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 for a variable.

Operational semantics

The call-by-value operational semantics in Figure 3.2 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 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_BETA, and E_TBETA rules should need no explanation—these are the standard call-by-value polymorphic lambda calculus reductions. The E_OP rule uses a denotation function $\llbracket - \rrbracket$ to give meaning to the first-order operations.

The E_REFL, E_FUN, and E_FORALL rules reduce casts structurally. E_REFL eliminates a cast from a type to itself; intuitively, such a cast should always succeed anyway. (I discuss this rule more in Section 3.3.) When a cast between function types is applied to a value v , the E_FUN rule produces a new lambda, wrapping v with

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

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

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

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

Figure 3.2: Operational semantics for F_H

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 [34] 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_REFL 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_REFL 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. We only

apply this rule 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, I offer a few examples. First, here is 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 is 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:

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

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

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

`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 3.3.)

I offer two higher level ways to understand the interactions of these complicated cast rules. First, we can see the reduction rules as an unfolding of a recursive function,

choosing the first clause in case of ambiguity. That is, the operational semantics unfolds a cast $\langle T_1 \Rightarrow T_2 \rangle^l v$ like $\mathcal{C}^l(T_1, T_2, v)$:

$$\begin{aligned}
\mathcal{C}^l(T, T, v) &= v \\
\mathcal{C}^l(\{x:T_1 \mid e\}, T_2, v) &= \mathcal{C}^l(T_1, T_2, v) \\
\mathcal{C}^l(T_1, \{x:T_2 \mid e\}, v) &= \text{let } x = \mathcal{C}^l(T_1, T_2, v) \text{ in } \langle \{x:T_2 \mid e\}, e, x \rangle^l \\
\mathcal{C}^l(\forall\alpha. T_1, \forall\alpha. T_2, v) &= \Lambda\alpha. \mathcal{C}^l(T_1, T_2, v) \\
\mathcal{C}^l(x:T_{11} \rightarrow T_{12}, x:T_{21} \rightarrow T_{22}, v) &= \\
&\lambda x:T_{21}. \mathcal{C}^l(T_{12}[\mathcal{C}^l(T_{21}, T_{11}, x)/x], T_{22}, v \mathcal{C}^l(T_{21}, T_{11}, x))
\end{aligned}$$

Alternatively, the rules firing during the evaluation of a cast in the small-step semantics obeys the following 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 $\langle T_1 \Rightarrow T_2 \rangle^l v$. To simplify the following discussion, I define $\text{unref}(T)$ as T without any outer refinements (though refinements on, e.g., the domain of a function would be unaffected); I write $\text{unref}_n(T)$ when we remove only the n outermost refinements:

$$\text{unref}(T) = \begin{cases} \text{unref}(T') & \text{if } T = \{x:T' \mid e\} \\ T & \text{otherwise} \end{cases}$$

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 :

$$\begin{aligned}
&\langle \text{unref}_1(T_2) \Rightarrow T_2 \rangle^l (\langle \text{unref}_2(T_2) \Rightarrow \text{unref}_1(T_2) \rangle^l \\
&\quad \dots (\langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l v) \dots)
\end{aligned}$$

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).

The E_REFL rule merits some more discussion. On the face of it, a cast $\langle T \Rightarrow T \rangle^l$ seems like it can't do anything: any value it applies must have already had type T , so what could go wrong during any checks? One might worry that adding such a cast will cause a different label to be blamed. In fact, we will find that such casts have no effect in Section 3.4, using the parametricity logical relation. But this safety proof is for a system where E_REFL was there in the first place! I haven't been able to prove parametricity for a system without E_REFL . Including it in the system, however, is a weaker assumption than including subsumption, as earlier systems did [34, 44]. It may be possible to further weaken the E_REFL by restricting it to base types.

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. We have no need for a value restriction because F_H doesn't have effects; any effectful extension, even nontermination, would force such a restriction. 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 3.4; in short, compatible types erase to identical simple type skeletons. I discuss type compatibility and type conversion (as used in `T_CONV`) below. Note that I assign casts a non-dependent function type. The `T_OP` rule uses the `ty` function to assign (possibly dependent) monomorphic first-order types to operations; I require that `ty(op)` and $\llbracket \text{op} \rrbracket$ agree.

Some of the typing rules—`T_CHECK`, `T_BLAKE`, `T_EXACT`, `T_FORGET`, and `T_CONV`—are “runtime only”. These rules aren't needed to type check source programs, but we need them to guarantee preservation. `T_CHECK`, `T_EXACT`, and `T_CONV` are excluded from source programs because I don't want the typing of source programs to rely on the evaluation relation; such an interaction is acceptable in this setting, but disrupts the phase distinction and is ultimately incompatible with nontermination and effects. I exclude `T_BLAKE` because programs shouldn't *start* with failures. Finally, I exclude `T_FORGET` because I imagine that source programs have all type changes explicitly managed by casts. In explicitly tagged calculus, as described in Section 3.5 and Chapter 4, we are able to abandon the `T_FORGET` rule in its entirety. 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. I 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. In Chapter 2, we have a similar but looser requirement: $e_2 \longrightarrow^* \text{true}$ implies $e_1[v/x] \longrightarrow^* \text{true}$. The proofs don't change radically, but I use the more syntactic requirement here because I aim to keep the system as syntactic as possible. The `T_EXACT` rule allows us to 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. Earlier systems used most-specific types and subtyping to show that the `E_OK` rule preserves typing. While the “most specific” requirement is abstract, a constant $k \in \mathcal{K}_B$ is typically given the *selfified* type $\text{ty}(k) = \{x:B \mid x = k\}$ [51]. But functions don't

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 $\boxed{\Gamma \vdash T}$

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

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

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

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

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

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

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

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

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

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

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

Figure 3.3: Typing rules for F_H

admit a decidable equality, so there isn't an obvious way to assign them most-specific types. `T_EXACT` is a suitably extensional, syntactic, and subtyping-free replacement for the earlier semantic requirement: constants and functions can be assigned less specific types, but we can use `T_EXACT` in the preservation proof to remember successful checks. 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). I define \equiv as an explicitly substitutive but otherwise structural type conversion relation in Figure 3.4. The `T_CONV` 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 $T_2[e'_2/x]$.

In the original ESOP 2011 work [8], we noted that $T_2[e_2/x]$ parallel reduces to $T_2[e'_2/x]$, so we took parallel reduction as type conversion. It turns out that parallel reduction doesn't quite have the properties we need. I discuss this further in Section 3.5. Here, I define a type conversion relation in Figure 3.4 that uses what I call *common subexpression reduction*, or CSR.. Belo et al. [8] also (falsely) claimed that symmetry wasn't necessary for type soundness or parametricity, but symmetry is in fact used in the proof of preservation (Lemma 3.2.27, when a term typed by `T_APP` steps by `E_REDUCE/E_REFL`).

The following proof of type soundness is (almost) entirely syntactic [79], offering a new approach to manifest calculi. In Chapter 2, we use subtyping instead of the \equiv relation; one of the contributions in this work is the insight that subtyping—with its accompanying metatheoretical complications that prevent a simple syntactic proof of type soundness—is not an essential component of manifest calculi.

I say *almost* entirely because both type soundness and parametricity rest on a conjecture about a semantic property which I call *cotermination*.

3.2.1 Conjecture [Cotermination at true]: If $\sigma \longrightarrow^* \sigma'$ then $\sigma(e) \longrightarrow^* \text{true}$ iff $\sigma'(e) \longrightarrow^* \text{true}$.

That is, if σ and σ' are common subexpression reductions, then terms behave equivalently when substituted with either. We discuss this conjecture further in Section 3.5.

I define type compatibility and type conversion in Figure 3.4. It is critical to understand that type variables are compatible only with themselves and refinements of themselves. To understand why this is, recall that compatibility determines when it is permissible to cast between two types. We can't allow a cast between α and a base type B , a function type $x:T_1 \rightarrow T_2$, or a quantified type $\forall\beta.T$ —we have no idea what type will fill in for α , and it's critical that compatibility is substitutive, i.e., that if $T_1 \parallel T_2$, then $T_1[e/x] \parallel T_2$ (Lemma 3.2.17). Moreover, we must avoid nontermination due to non-parametric operations (e.g., Girard's J operator); it's imperative that a term like

$$\text{let } \delta = \Lambda\alpha. \lambda x:\alpha. \langle \alpha \Rightarrow \forall\beta.\beta \rightarrow \beta \rangle^l \alpha x \text{ in } \delta \forall\beta.\beta \rightarrow \beta \delta$$

isn't well typed.

3.2.2 Lemma [Determinism]: If $e \longrightarrow e_1$ and $e \longrightarrow e_2$ then $e_1 = e_2$.

Type compatibility $\boxed{T_1 \parallel T_2}$

$$\begin{array}{c}
\frac{}{\alpha \parallel \alpha} \text{SIM_VAR} \qquad \frac{}{B \parallel B} \text{SIM_BASE} \\
\frac{T_1 \parallel T_2}{\{x:T_1 \mid e\} \parallel T_2} \text{SIM_REFINEL} \qquad \frac{T_1 \parallel T_2}{T_1 \parallel \{x:T_2 \mid e\}} \text{SIM_REFINER} \\
\frac{T_{11} \parallel T_{21} \quad T_{12} \parallel T_{22}}{x:T_{11} \rightarrow T_{12} \parallel x:T_{21} \rightarrow T_{22}} \text{SIM_FUN} \qquad \frac{T_1 \parallel T_2}{\forall \alpha. T_1 \parallel \forall \alpha. T_2} \text{SIM_FORALL}
\end{array}$$

Conversion $\boxed{\sigma_1 \longrightarrow^* \sigma_2}$ $\boxed{T_1 \equiv T_2}$

$$\begin{array}{c}
\sigma_1 \longrightarrow^* \sigma_2 \iff \begin{array}{l} \text{dom}(\sigma_1) = \text{dom}(\sigma_2) \wedge \\ \forall x \in \text{dom}(\sigma_1). \sigma_1(x) \longrightarrow^* \sigma_2(x) \wedge \\ \forall \alpha \in \text{dom}(\sigma_1). \sigma_1(\alpha) = \sigma_2(\alpha) \end{array} \\
\frac{}{\alpha \equiv \alpha} \text{C_VAR} \qquad \frac{}{B \equiv B} \text{C_BASE} \qquad \frac{\sigma_1 \longrightarrow^* \sigma_2 \quad T_1 \equiv T_2}{\{x:T_1 \mid \sigma_1(e)\} \equiv \{x:T_2 \mid \sigma_2(e)\}} \text{C_REFINE} \\
\frac{T_1 \equiv T'_1 \quad T_2 \equiv T'_2}{x:T_1 \rightarrow T_2 \equiv x:T'_1 \rightarrow T'_2} \text{C_FUN} \qquad \frac{T \equiv T'}{\forall \alpha. T \equiv \forall \alpha. T'} \text{C_FORALL} \\
\frac{T_2 \equiv T_1}{T_1 \equiv T_2} \text{C_SYM} \qquad \frac{T_1 \equiv T_2 \quad T_2 \equiv T_3}{T_1 \equiv T_3} \text{C_TRANS}
\end{array}$$

Figure 3.4: Type compatibility and conversion for F_H

3.2.3 Lemma [Reflexivity of conversion]:

$T \equiv T$ for all T .

Proof: By induction on T .

$(T = \alpha)$: By C_VAR.

$(T = B)$: By C_BASE.

$(T = \{x:T' \mid e\})$: By the IH on T' , we have $T' \equiv T'$. With an empty σ , we have $\sigma \longrightarrow^* \sigma$ trivially. We are done by C_REFINE.

$(T = x:T_1 \rightarrow T_2)$: By the IH on T_1 and T_2 and C_FUN.

$(T = \forall \alpha. T)$: By the IH on T and C_FORALL. □

3.2.4 Lemma [Like-type arrow conversion]:

If $x:T_{11} \rightarrow T_{12} \equiv T$ then $T = x:T_{21} \rightarrow T_{22}$.

Proof: By induction on the conversion relation. Only C_FUN applies, and C_SYM and C_TRANS are resolved by the IH. □

3.2.5 Lemma [Conversion arrow inversion]: If $x:T_{11} \rightarrow T_{12} \equiv x:T_{21} \rightarrow T_{22}$ then $T_{11} \equiv T_{21}$ and $T_{12} \equiv T_{22}$.

Proof: By induction on the conversion derivation.

(C_FUN): Immediate.

(C_SYM): By the IH and C_SYM.

(C_TRANS): It must be the case that T_2 is an arrow type (Lemma 3.2.4), so by the IH and C_TRANS. \square

3.2.6 Lemma [Like-type forall conversion]: If $\forall\alpha.T_1 \equiv T$ then $T = \forall\alpha.T_2$.

Proof: By induction on the conversion relation. Only C_FORALL applies, and C_SYM and C_TRANS are resolved by the IH. \square

3.2.7 Lemma [Conversion forall inversion]: If $\forall\alpha.T_1 \equiv \forall\alpha.T_2$ then $T_1 \equiv T_2$.

Proof: By induction on the conversion derivation.

(C_FORALL): Immediate.

(C_SYM): By the IH and C_SYM.

(C_TRANS): By Lemma 3.2.6, we know that the intermediate type is a forall type, so by the IH and C_TRANS. \square

3.2.8 Lemma [Term substitutivity of conversion]:

If $T_1 \equiv T_2$ and $e_1 \rightarrow^* e_2$ then $T_1[e_1/x] \equiv T_2[e_2/x]$.

3.2.9 Lemma [Type substitutivity of conversion]:

If $T_1 \equiv T_2$ then $T_1[T/\alpha] \equiv T_2[T/\alpha]$.

3.2.10 Lemma [Cotermination of refinement types]: If $\{x:T_1 \mid e_1\} \equiv \{x:T_2 \mid e_2\}$ then $T_1 \equiv T_2$ and $e_1[v/x] \rightarrow^* \text{true}$ iff $e_2[v/x]$, for all v .

Proof: By induction on the equivalence. There are three cases.

(C_REFINE): We have $T_1 \equiv T_2$ by assumption. We know that $e_1 = \sigma_1(e)$ and $e_2 = \sigma_2(e)$ for $\sigma_1 \rightarrow^* \sigma_2$. It is trivially true that $v \rightarrow^* v$, so $\sigma_1[v/x] \rightarrow^* \sigma_2[v/x]$. By cotermination (Lemma 3.2.1), we know that $\sigma_1(e)[v/x] \rightarrow^* \text{true}$ iff $\sigma_2(e)[v/x]$.

(C_SYM): By the IH.

(C_TRANS): By the IHs and transitivity of \equiv and cotermination. \square

3.2.11 Lemma [Value inversion]: If $\emptyset \vdash v : T$ and $\text{unref}_n(T) = \{x:T_n \mid e_n\}$ then $e_n[v/x] \rightarrow^* \text{true}$.

Proof: By induction on the height of the typing derivation; we list all the cases that could type values.

Proof:

(T_CONST): By assumption of valid typing of constants.

(T_ABS): Contradictory—the type is wrong.

(T_TABS): Contradictory—the type is wrong.

(T_CAST): Contradictory—the type is wrong.

(T_CONV): By applying Lemma 3.2.10 on the stack of refinements on T .

(T_FORGET): By the IH on $\emptyset \vdash v : \{x:T \mid e\}$, adjusting each of the n down by one to cover the stack of refinements on T .

(T_EXACT): By assumption for the outermost refinement; by the IH on $\emptyset \vdash v : T$ for the rest. \square

3.2.12 Lemma [Term weakening]: If x is fresh and $\Gamma \vdash T'$ then

1. $\Gamma, \Gamma' \vdash e : T$ implies $\Gamma, x:T', \Gamma' \vdash e : T$,
2. $\Gamma, \Gamma' \vdash T$ implies $\Gamma, x:T', \Gamma' \vdash T$, and
3. $\vdash \Gamma, \Gamma'$ implies $\vdash \Gamma, x:T', \Gamma'$.

Proof: By induction on e , T , and Γ' . The only interesting case is for terms where a runtime rule applies:

(T_CONV, T_EXACT, T_FORGET): The argument is the same for all terms, so: since $\vdash \Gamma, x:T', \Gamma'$, we can reapply T_CONV, T_EXACT, or T_FORGET, respectively. In the rest of this proof, we won't bother considering these rules. \square

3.2.13 Lemma [Type weakening]: If α is fresh then

1. $\Gamma, \Gamma' \vdash e : T$ implies $\Gamma, \alpha, \Gamma' \vdash e : T$,
2. $\Gamma, \Gamma' \vdash T$ implies $\Gamma, \alpha, \Gamma' \vdash T$, and
3. $\vdash \Gamma, \Gamma'$ implies $\vdash \Gamma, \alpha, \Gamma'$.

Proof: By induction on e , T , and Γ' . The proof is similar to term weakening, Lemma 3.2.12. \square

3.2.14 Lemma [Compatibility is symmetric]: $T_1 \parallel T_2$ iff $T_2 \parallel T_1$.

Proof: By induction on $T_1 \parallel T_2$.

(SIM_VAR): By SIM_VAR.

(SIM_BASE): By SIM_BASE.

(SIM_REFINEL): By SIM_REFINER and the IH.

(SIM_REFINER): By SIM_REFINEL and the IH.

(SIM_FUN): By SIM_FUN and the IHs.

(SIM_FORALL): By the IH and SIM_FORALL. □

3.2.15 Lemma [Substitution preserves compatibility]:

If $T_1 \parallel T_2$, then $T_1[e/x] \parallel T_2$.

Proof: By induction on the compatibility relation.

(SIM_VAR): By SIM_VAR.

(SIM_BASE): By SIM_BASE.

(SIM_REFINEL): By SIM_REFINEL and the IH.

(SIM_REFINER): By SIM_REFINER and the IH.

(SIM_FUN): By SIM_FUN and the IHs.

(SIM_FORALL): By SIM_FORALL and the IH. □

3.2.16 Lemma [Term substitution]: If $\Gamma \vdash e' : T'$, then

1. if $\Gamma, x:T', \Gamma' \vdash e : T$ then $\Gamma, \Gamma'[e'/x] \vdash e[e'/x] : T[e'/x]$,
2. if $\Gamma, x:T', \Gamma' \vdash T$ then $\Gamma, \Gamma'[e'/x] \vdash T[e'/x]$, and
3. if $\vdash \Gamma, x:T', \Gamma'$ then $\vdash \Gamma, \Gamma'[e'/x]$.

Proof: By induction on e , T , and Γ' . In the first two clauses, we are careful to leave Γ' as long as it is well formed. □

3.2.17 Lemma [Type substitution preserves compatibility]: If $T_1 \parallel T_2$ then $T_1[T'/\alpha] \parallel T_2$.

Proof: By induction on the compatibility relation.

(SIM_VAR): By SIM_VAR.

(SIM_BASE): By SIM_BASE.

(SIM_REFINEL): By SIM_REFINEL and the IH.

(SIM_REFINER): By SIM_REFINER and the IH.

(SIM_FUN): By SIM_FUN and the IHs.

(SIM_FORALL): By SIM_FORALL and the IH. □

3.2.18 Lemma [Type substitution]: If $\Gamma \vdash T'$ then

1. if $\Gamma, \alpha, \Gamma' \vdash e : T$, then $\Gamma, \Gamma'[T'/\alpha] \vdash e[T'/\alpha] : T[T'/\alpha]$,

2. if $\Gamma, \alpha, \Gamma' \vdash T$, then $\Gamma, \Gamma'[T'/\alpha] \vdash T[T'/\alpha]$, and
3. if $\vdash \Gamma, \alpha, \Gamma'$, then $\vdash \Gamma, \Gamma'[T'/\alpha]$.

Proof: By induction on e , T , and Γ' . □

As is standard for type systems with conversion rules, we must prove inversion lemmas in order to reason about typing derivations in a syntax-directed way.

3.2.19 Lemma [Lambda inversion]: If $\Gamma \vdash \lambda x:T_1. e_{12} : T$, then

1. $\Gamma \vdash T_1$,
2. $\Gamma, x:T_1 \vdash e_{12} : T_2$, and
3. $x:T_1 \rightarrow T_2 \equiv \text{unref}(T)$.

Proof: By induction on the typing derivation. Cases not mentioned only apply to terms which are not lambdas.

(T_ABS): By inversion, we have $\Gamma \vdash T_1$ and $\Gamma, x:T_1 \vdash e_{12} : T_2$. We find conversion immediately by reflexivity (Lemma 3.2.3), since $\text{unref}(T) = T = x:T_1 \rightarrow T_2$.

(T_CONV): We have $\Gamma \vdash \lambda x:T_1. e_{12} : T$; by inversion, $T \equiv T'$ and $\emptyset \vdash \lambda x:T_1. e_{12} : T'$. By the IH on this second derivation, we find $\emptyset \vdash T_1$ and $x:T_1 \vdash e_{12} : T_2$ where, $\text{unref}(T') \equiv x:T_1 \rightarrow T_2$. By weakening, we have $\Gamma \vdash T_1$ and $\Gamma, x:T_1 \vdash e_{12} : T_2$. Since $T' \equiv T$, we have $x:T_1 \rightarrow T_2 \equiv \text{unref}(T') \equiv \text{unref}(T)$ by C_TRANS.

(T_EXACT): $T = \{x:T' \mid e\}$, and we have $\Gamma \vdash \lambda x:T_1. e_{12} : \{x:T' \mid e\}$; by inversion, $\emptyset \vdash \lambda x:T_1. e_{12} : T'$. By the IH, $\emptyset \vdash T_1$ and $x:T_1 \vdash e_{12} : T_2$, where $x:T_1 \rightarrow T_2 \equiv \text{unref}(T')$. By weakening, $\Gamma \vdash T_1$ and $\Gamma, x:T_1 \vdash e_{12} : T_2$. Since $\text{unref}(T') = \text{unref}(\{x:T' \mid e\})$, we have the conversion by C_TRANS: $x:T_1 \rightarrow T_2 \equiv \text{unref}(T') = \text{unref}(\{x:T' \mid e\})$.

(T_FORGET): We have $\Gamma \vdash \lambda x:T_1. e_{12} : T$; by inversion, $\emptyset \vdash \lambda x:T_1. e_{12} : \{x:T \mid e\}$. By the IH on this latter derivation, we $\emptyset \vdash T_1$ and $x:T_1 \vdash e_{12} : T_2$, where $x:T_1 \rightarrow T_2 \equiv \text{unref}(\{x:T \mid e\})$. By weakening, $\Gamma \vdash T_1$ and $\Gamma, x:T_1 \vdash e_{12} : T_2$. Since $\text{unref}(\{x:T \mid e\}) = \text{unref}(T)$, we have by C_TRANS that $x:T_1 \rightarrow T_2 \equiv \text{unref}(\{x:T \mid e\}) = \text{unref}(T)$. □

3.2.20 Lemma [Cast inversion]: If $\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle^l : T$, then

1. $\Gamma \vdash T_1$,
2. $\Gamma \vdash T_2$,
3. $T_1 \parallel T_2$, and
4. $_ : T_1 \rightarrow T_2 \equiv \text{unref}(T)$ (i.e., T_2 does not mention the dependent variable).

Proof: By induction on the typing derivation. Cases not mentioned only apply to syntactically distinct terms.

(T_CAST): $T = _ : T_1 \rightarrow T_2$, and the derivation is by inversion. Conversion is by reflexivity (Lemma 3.2.3). T_2 does not mention the variable.

(T_CONV): $\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle^l : T$; by inversion, $\emptyset \vdash \langle T_1 \Rightarrow T_2 \rangle^l : T'$ and $T' \equiv T$. By the IH, we have $\emptyset \vdash T_1$ and $\emptyset \vdash T_2$ (which weaken to the derivations we want), as well as $T_1 \parallel T_2$ and $_ : T_1 \rightarrow T_2 \equiv \text{unref}(T')$. But $\text{unref}(T') \equiv \text{unref}(T)$, so we have the conversion we want by C_TRANS.

(T_EXACT): $\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle^l : \{x:T \mid e\}$; by inversion, $\emptyset \vdash \langle T_1 \Rightarrow T_2 \rangle^l : T$. By the IH, we have $\emptyset \vdash T_1$ and $\emptyset \vdash T_2$ (which weaken to the derivations we want). We also have $T_1 \parallel T_2$ and $_ : T_1 \rightarrow T_2 \equiv \text{unref}(T)$ — which is equal to $\text{unref}(\{x:T \mid e\})$, so we're done.

(T_FORGET): $\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle^l : T$; by inversion, $\emptyset \vdash \langle T_1 \Rightarrow T_2 \rangle^l : \{x:T \mid e\}$. By the IH, $\emptyset \vdash T_1$ and $\emptyset \vdash T_2$, so we find the derivations we want by weakening. We also have $T_1 \parallel T_2$ and $_ : T_1 \rightarrow T_2 \equiv \text{unref}(T)$, which is equal $\text{unref}(\{x:T \mid e\})$, so this case is complete. □

3.2.21 Lemma [Application inversion]: If $\Gamma \vdash e_1 e_2 : T$, then

1. $\Gamma \vdash e_1 : (x:T_1 \rightarrow T_2)$,
2. $\Gamma \vdash e_2 : T_1$, and
3. $T_2[e_2/x] \equiv T$.

Proof: By induction on the typing derivation. Cases not mentioned only apply to syntactically distinct terms.

(T_APP): We find the typings we want by inversion. Since $T = T_2[e_2/x]$, conversion is by reflexivity (Lemma 3.2.3).

(T_CONV): By the IH on $\emptyset \vdash e_1 e_2 : T'$, we find the typings we want (up to weakening), and $T_2[e_2/x] \equiv T'$. By inversion, $T' \equiv T$, so by C_TRANS we have $T_2[e_2/x] \equiv T$.

(T_EXACT): Only applies to values, and $e_1 e_2$ cannot be a value.

(T_FORGET): Only applies to values, and $e_1 e_2$ cannot be a value. □

3.2.22 Lemma [Type abstraction inversion]: If $\Gamma \vdash \Lambda\alpha. e : T$, then

1. $\Gamma, \alpha \vdash e : T'$ and
2. $\forall\alpha. T' \equiv \text{unref}(T)$.

Proof: By induction on the typing derivation. Cases not mentioned only apply to syntactically distinct terms.

(T_TABS): $T = \forall\alpha.T'$. Conversion is by reflexivity (Lemma 3.2.3).

(T_CONV): By the IH on $\emptyset \vdash \Lambda\alpha. e : T''$, using weakening to recover the typing derivation and using transitivity of conversion through unref to find $\text{unref}(T'') \equiv \text{unref}(T)$ by C_TRANS.

(T_EXACT): $T = \{x:T_0 \mid e_0\}$ for some T_0 and e_0 . By the IH on $\emptyset \vdash \Lambda\alpha. e : T_0$, using weakening for the derivation and the fact that $\text{unref}(\{x:T_0 \mid e_0\}) = \text{unref}(T_0)$ to find conversion.

(T_FORGET): By the IH on $\emptyset \vdash \Lambda\alpha. e : \{x:T \mid e_0\}$, using weakening for the derivation and the fact that $\text{unref}(\{x:T \mid e_0\}) = \text{unref}(T)$ to find conversion. \square

Inversion lemmas in hand, we prove a canonical forms lemma to support a proof of progress. The canonical forms proof is “modulo” the unref function: the shape of the values of type $\{x:T \mid e\}$ are determined by the inner type T .

3.2.23 Lemma [Conversion of unrefined types]: If $T_1 \equiv T_2$ then $\text{unref}(T_1) \equiv \text{unref}(T_2)$.

Proof: By induction on the derivation of $T_1 \equiv T_2$. \square

3.2.24 Lemma [Canonical forms]: If $\emptyset \vdash v : T$, then:

1. If $\text{unref}(T) = B$ then $v = k \in \mathcal{K}_B$ for some v
2. If $\text{unref}(T) = x:T_1 \rightarrow T_2$ then v is
 - (a) $\lambda x:T'_1. e_{12}$ and $T'_1 \equiv T_1$ for some x, T'_1 and e_{12} , or
 - (b) $\langle T'_1 \Rightarrow T'_2 \rangle^l$ and $T'_1 \equiv T_1$ and $T'_2 \equiv T_2$ for some T'_1, T'_2 , and l
3. If $\text{unref}(T) = \forall\alpha.T'$ then v is $\Lambda\alpha. e$ for some e .

Proof: By induction on the typing derivation.

(T_VAR): Contradictory: variables are not values.

(T_CONST): $\emptyset \vdash k : T$ and $\text{unref}(T) = B$; we are in case 1. By assumption, $k \in \mathcal{K}_B$.

(T_OP): Contradictory: $\text{op}(e_1, \dots, e_n)$ is not a value.

(T_ABS): $\emptyset \vdash \lambda x:T_1. e_{12} : T$ and $T = \text{unref}(T) = x:T_1 \rightarrow T_2$; we are in case 2a. Conversion is by reflexivity (Lemma 3.2.3).

(T_APP): Contradictory: $e_1 e_2$ is not a value.

(T_TABS): $\emptyset \vdash \Lambda\alpha. e : \forall\alpha.T$; we are in case 3. It is immediate that $v = \Lambda\alpha. e$, and conversion is by reflexivity (Lemma 3.2.3).

(T_TAPP): Contradictory: $e T$ is not a value.

(T_CAST): $\emptyset \vdash \langle T_1 \Rightarrow T_2 \rangle^l : _ : T_1 \rightarrow T_2$; we are in case 2b. It is immediate that $v = \langle T_1 \Rightarrow T_2 \rangle^l$. Conversion is by reflexivity (Lemma 3.2.3).

(T_CHECK): Contradictory: $\langle \{x:T \mid e_1\}, e_2, v \rangle^l$ is not a value.

(T_BLAZE): Contradictory: $\uparrow l$ is not a value.

(T_CONV): $\emptyset \vdash v : T$; by inversion, $\emptyset \vdash v : T'$ and $T' \equiv T$. We find an appropriate form for $\text{unref}(T')$ by the IH on $\emptyset \vdash v : T'$. We go by cases, in each case reproving whatever case was found in the IH and finding conversions by C_TRANS.

Case 1: $\text{unref}(T) = B$ and $v = k \in \mathcal{K}_B$. Since $\text{unref}(T') \equiv \text{unref}(T)$, we know that $\text{unref}(T') = B$, which is all we needed to show.

Case 2a: $\text{unref}(T) = x:T_1 \rightarrow T_2$ and $v = \lambda x:T_1''. e_{12}$ and $T_1'' \equiv T_1$. Since $T' \equiv T$, we have $\text{unref}(T') \equiv \text{unref}(T)$ (Lemma 3.2.23) and so $\text{unref}(T') = x:T_1' \rightarrow T_2'$ for some T_1' and T_2' such that $T_1' \equiv T_1$ (Lemma 3.2.5); by C_TRANS, we have $T_1'' \equiv T_1'$.

Case 2b: $\text{unref}(T) = x:T_1 \rightarrow T_2$ and $v = \langle T_1' \Rightarrow T_2' \rangle^l$ and $T_1' \equiv T_1$ and $T_2' \equiv T_2$. Since $T' \equiv T$, we have $\text{unref}(T') \equiv \text{unref}(T)$ (Lemma 3.2.23) and so $\text{unref}(T') = x:T_1'' \rightarrow T_2''$ for some T_1'' and T_2'' such that $T_1'' \equiv T_1$ and $T_2'' \equiv T_2$ (Lemma 3.2.5); by C_TRANS, we have $T_1' \equiv T_1''$ and $T_2' \equiv T_2''$ as required.

Case 3: $\text{unref}(T) = \forall \alpha. T_0$ and v is $\Lambda \alpha. e$. Since $T' \equiv T$, then $\text{unref}(T') \equiv \text{unref}(T)$ (Lemma 3.2.23).

(T_EXACT): $\emptyset \vdash v : \{x:T \mid e\}$; by inversion, $\emptyset \vdash v : T$. Noting that $\text{unref}(\{x:T \mid e\}) = \text{unref}(T)$, we apply the IH. Unlike the previous case, we need not change the conversion—it's in terms of the unrefined type.

(T_FORGET): $\emptyset \vdash v : T$; by inversion $\emptyset \vdash v : \{x:T \mid e\}$. By the IH (noting $\text{unref}(\{x:T \mid e\}) = \text{unref}(T)$), so we use the IH's conversion directly. \square

3.2.25 Theorem [Progress]: If $\emptyset \vdash e : T$, then either

1. $e \longrightarrow e'$, or
2. e is a result r , i.e., a value or blame.

Proof: By induction on the typing derivation.

(T_VAR): Contradictory: there is no derivation $\emptyset \vdash x : T$.

(T_CONST): $\emptyset \vdash k : \text{ty}(k)$. In this case, $e = k$ is a result.

(T_OP): $\emptyset \vdash \text{op}(e_1, \dots, e_n) : \sigma(T)$, where $\text{ty}(\text{op}) = x_1 : T_1 \rightarrow \dots \rightarrow x_n : T_n \rightarrow T$. By inversion, $\emptyset \vdash e_i : \sigma(T_i)$. Applying the IH from left to right, each of the e_i either steps or is a result.

Suppose everything to the left of e_i is a value. Then either e_i steps or is a result. If $e_i \rightarrow e'_i$, then $\text{op}(v_1, \dots, v_{i-1}, e_i, \dots, e_n) \rightarrow \text{op}(v_1, \dots, v_{i-1}, e'_i, \dots, e_n)$ by E_COMPAT. On the other hand, if e_i is a result, there are two cases. If $e_i = \uparrow l$, then the original expression steps to $\uparrow l$ by E_BLAZE. If e_i is a value, we can continue this process for each of the operation's arguments. Eventually, all of the operations arguments are values. By value inversion (Lemma 3.2.11), we know that we can type each of these values at the exact refinement types we need by T_EXACT. We assume that if $\text{op}(v_1, \dots, v_n)$ is well defined on values satisfying the refinements in its type, so E_OP applies.

(T_ABS): $\emptyset \vdash \lambda x:T_1. e_{12} : (x:T_1 \rightarrow T_2)$. In this case, $e = \lambda x:T_1. e_{12}$ is a result.

(T_APP): $\emptyset \vdash e_1 e_2 : T_2[e_2/x]$; by inversion, $\emptyset \vdash e_1 : (x:T_1 \rightarrow T_2)$ and $\emptyset \vdash e_2 : T_1$.

By the IH on the first derivation, e_1 steps or is a result. If e_1 steps, then the entire term steps by E_COMPAT. In the latter case, if e_1 is blame, we step by E_BLAZE. So e_1 is a value, v_1 .

By the IH on the second derivation, e_2 steps or is a result. If e_2 steps, then by E_COMPAT. Otherwise, if e_2 is blame, we step by E_BLAZE. So e_2 is a value, v_2 .

By canonical forms (Lemma 3.2.24) on $\emptyset \vdash e_1 : (x:T_1 \rightarrow T_2)$, there are two cases:

($e_1 = \lambda x:T'_1. e_{12}$ and $T'_1 \equiv T_1$): In this case, $(\lambda x:T'_1. e_{12}) v_2 \rightarrow e_{12}[v_2/x]$ by E_BETA.

($e_1 = \langle T'_1 \Rightarrow T'_2 \rangle^l$ and $T'_1 \equiv T_1$ and $T'_2 \equiv T_2$): We know that $T'_1 \parallel T'_2$ by cast inversion (Lemma 3.2.20). We determine which step is taken by cases on T'_1 and T'_2 .

($T'_1 = B$):

($T'_2 = B'$): It must be the case that $B = B'$, since $B \parallel B'$. By E_REFL, $\langle B \Rightarrow B \rangle^l v_2 \rightarrow v_2$.

($T'_2 = \alpha$ or $x:T_{21} \rightarrow T_{22}$ or $\forall \alpha. T_{22}$): Incompatible; contradictory.

($T'_2 = \{x:T''_2 \mid e\}$): If $T''_2 = B$, then by E_CHECK, $\langle B \Rightarrow \{x:B \mid e\} \rangle^l v_2 \rightarrow \langle \{x:B \mid e\}, e[v_2/x], v_2 \rangle^l$. Otherwise, by E_PRECHECK, we have:

$$\langle B \Rightarrow \{x:T''_2 \mid e\} \rangle^l v_2 \rightarrow \langle T''_2 \Rightarrow \{x:T''_2 \mid e\} \rangle^l (\langle B \Rightarrow T''_2 \rangle^l v_2)$$

($T'_1 = \alpha$):

($T'_2 = \alpha'$): It must be the case that $\alpha = \alpha'$, since $\alpha \parallel \alpha'$. By E_REFL, $\langle \alpha \Rightarrow \alpha \rangle^l v_2 \rightarrow v_2$.

($T'_2 = B$ or $x:T_{21} \rightarrow T_{22}$ or $\forall \alpha. T_{22}$): Incompatible; contradictory.

$(T'_2 = \{x:T''_2 \mid e\})$: If $T''_2 = \alpha$, then by E_CHECK, $\langle \alpha \Rightarrow \{x:\alpha \mid e\} \rangle^l v_2 \longrightarrow \langle \{x:\alpha \mid e\}, e[v_2/x], v_2 \rangle^l$. Otherwise,

$$\langle \alpha \Rightarrow \{x:T''_2 \mid e\} \rangle^l v_2 \longrightarrow \langle T''_2 \Rightarrow \{x:T''_2 \mid e\} \rangle^l (\langle \alpha \Rightarrow T''_2 \rangle^l v_2)$$

by E_PRECHECK.

$(T'_1 = x:T_{11} \rightarrow T_{12})$:

$(T'_2 = B \text{ or } \alpha \text{ or } \forall \alpha.T_{22})$: Incompatible; contradictory.

$(T'_2 = x:T_{21} \rightarrow T_{22})$: If $T'_1 = T'_2$, then $\langle T'_1 \Rightarrow T'_1 \rangle^l v_2 \longrightarrow v_2$ by E_REFL. If not, then

$$\begin{aligned} & \langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l v_2 \longrightarrow \\ & \lambda x:T_{21}. (\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l (v_2 (\langle T_{21} \Rightarrow T_{11} \rangle^l x))) \end{aligned}$$

by E_FUN.

$(T'_2 = \{x:T''_2 \mid e\})$: If $T'_1 = T''_2$, then $\langle T'_1 \Rightarrow \{x:T'_1 \mid e\} \rangle^l v_2 \longrightarrow \langle \{x:T'_1 \mid e\}, e[v_2/x], v_2 \rangle^l$ by E_CHECK. If not, then

$$\langle T'_1 \Rightarrow \{x:T''_2 \mid e\} \rangle^l v_2 \longrightarrow \langle T''_2 \Rightarrow \{x:T''_2 \mid e\} \rangle^l (\langle T'_1 \Rightarrow T''_2 \rangle^l v_2)$$

by E_PRECHECK.

$(T'_1 = \forall \alpha.T_{12})$:

$(T'_2 = B \text{ or } \alpha \text{ or } x:T_{21} \rightarrow T_{22})$: Incompatible; contradictory.

$(T'_2 = \forall \alpha.T_{22})$: If $T'_1 = T'_2$, then $\langle T'_1 \Rightarrow T'_1 \rangle^l v_2 \longrightarrow v_2$ by E_REFL. If not, then $\langle \forall \alpha.T_{11} \Rightarrow \forall \alpha.T_{22} \rangle^l v_2 \longrightarrow \Lambda \alpha. (\langle T_{11} \Rightarrow T_{22} \rangle^l (v_2 \alpha))$ by E_FORALL.

$(T'_2 = \{x:T''_2 \mid e\})$: If $T'_1 = T''_2$, then $\langle T'_1 \Rightarrow \{x:T'_1 \mid e\} \rangle^l v_2 \longrightarrow \langle \{x:T'_1 \mid e\}, e[v_2/x], v_2 \rangle^l$ by E_CHECK. If not, then $\langle T'_1 \Rightarrow \{x:T''_2 \mid e\} \rangle^l v_2 \longrightarrow \langle T''_2 \Rightarrow \{x:T''_2 \mid e\} \rangle^l (\langle T'_1 \Rightarrow T''_2 \rangle^l v_2)$ by E_PRECHECK.

$(T'_1 = \{x:T''_1 \mid e'_1\})$:

$(T'_2 = B \text{ or } \alpha \text{ or } x:T_{21} \rightarrow T_{22} \text{ or } \forall \alpha.T_{22})$: We see

$$\langle \{x:T''_1 \mid e'_1\} \Rightarrow T'_2 \rangle^l v_2 \longrightarrow \langle T''_1 \Rightarrow T'_2 \rangle^l v_2$$

by E_FORGET.

$(T'_2 = \{x:T''_2 \mid e'_2\})$: If $T'_1 = T'_2$, then we immediately have $\langle T'_1 \Rightarrow T'_1 \rangle^l v_2 \longrightarrow v_2$ by E_REFL. If $T'_1 = T''_2$, then

$$\langle T'_1 \Rightarrow \{x:T'_1 \mid e'_2\} \rangle^l v_2 \longrightarrow \langle \{x:T'_1 \mid e'_2\}, e'_2[v_2/x], v_2 \rangle^l$$

by E_CHECK. Otherwise,

$$\langle \{x:T''_1 \mid e'_1\} \Rightarrow \{x:T''_2 \mid e'_2\} \rangle^l v_2 \longrightarrow \langle T''_1 \Rightarrow \{x:T''_2 \mid e'_2\} \rangle^l v_2$$

by E_FORGET.

(T_TABS): $\emptyset \vdash \Lambda\alpha. e' : \forall\alpha. T$. In this case, $\Lambda\alpha. e'$ is a result.

(T_TAPP): $\emptyset \vdash e_1 T_2 : T_1[T_2/\alpha]$; by inversion, $\emptyset \vdash e_1 : \forall\alpha. T_1$ and $\emptyset \vdash T_2$. By the IH on the first derivation, e_1 steps or is a result. If $e_1 \longrightarrow e'_1$, then $e_1 T_2 \longrightarrow e'_1 T_2$ by E_COMPAT. If $e_1 = \uparrow l$, then $\uparrow l T_2 \longrightarrow \uparrow l$ by E_BLAKE.

If $e_1 = v_1$, then we know that $v_1 = \Lambda\alpha. e'_1$ by canonical forms (Lemma 3.2.24). We can see $(\Lambda\alpha. e'_1) T_2 \longrightarrow e'_1[T_2/\alpha]$ by E_TBETA.

(T_CAST): $\emptyset \vdash \langle T_1 \Rightarrow T_2 \rangle^l : T_1 \rightarrow T_2$. In this case, $\langle T_1 \Rightarrow T_2 \rangle^l$ is a result.

(T_CHECK): $\emptyset \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle^l : \{x:T \mid e_1\}$; by inversion, $\emptyset \vdash e_2 : \text{Bool}$. By the IH, either $e_2 \longrightarrow e'_2$ steps or $e_2 = r_2$. In the first case, $\langle \{x:T \mid e_1\}, e_2, v \rangle^l \longrightarrow \langle \{x:T \mid e_1\}, e'_2, v \rangle^l$ by E_COMPAT. In the second case, either $r_2 = \uparrow l$ or $r_2 = v_2$. If we have blame, then the entire term steps by E_BLAKE. If we have a value, then we know that v_2 is either `true` or `false`, since it's typed at `Bool`. If it's `true`, we step by E_OK. Otherwise we step by E_FAIL.

(T_BLAKE): $\emptyset \vdash \uparrow l : T$. In this case, $\uparrow l$ is a result.

(T_CONV): $\emptyset \vdash e : T'$; by inversion, $\emptyset \vdash e : T$. By the IH, we see that $e \longrightarrow e'$ or $e = r$.

(T_EXACT): $\emptyset \vdash v : \{x:T \mid e\}$. Here, v is a result by assumption.

(T_FORGET): $\emptyset \vdash v : T$. Again, v is a result by assumption. \square

\square

The following regularity property formalizes an important property of the type system: all contexts and types involved are well formed. This is critical for the proof of preservation: when a term raises blame, we must show that the blame is well typed. With regularity, we can immediately know that the original type is well formed.

3.2.26 Lemma [Context and type well formedness]: 1. If $\Gamma \vdash e : T$, then $\vdash \Gamma$ and $\Gamma \vdash T$.

2. If $\Gamma \vdash T$ then $\vdash \Gamma$.

Proof: By induction on the typing and well formedness derivations. \square

3.2.27 Theorem [Preservation]: If $\emptyset \vdash e : T$ and $e \longrightarrow e'$, then $\emptyset \vdash e' : T$.

Proof: By induction on the typing derivation.

(T_VAR): Contradictory—we can't have $\emptyset \vdash x : T$.

(T_CONST): $\emptyset \vdash k : \text{ty}(k)$. Contradictory—values don't step.

(T_OP): $\emptyset \vdash \text{op}(e_1, \dots, e_n) : \sigma(T)$. By cases on the step taken:

(E_REDUCE/E_OP): $\text{op}(v_1, \dots, v_n) \longrightarrow \llbracket \text{op} \rrbracket(v_1, \dots, v_n)$. This case is by assumption.

(E_BLAKE): $e_i = \uparrow l$, and everything to its left is a value. By context and type well formedness (Lemma 3.2.26), $\emptyset \vdash \sigma(T)$. So by T_BLAKE, $\emptyset \vdash \uparrow l : \sigma(T)$.

(E_COMPAT): Some $e_i \rightarrow e'_i$. By the IH and T_OP, using T_CONV to show that $\sigma(T) \equiv \sigma'(T)$ (Lemma 3.2.8).

(T_ABS): $\emptyset \vdash \lambda x:T_1. e_{12} : (x:T_1 \rightarrow T_2)$. Contradictory—values don't step.

(T_APP): $\emptyset \vdash e_1 e_2 : T'_2[e_2/x]$, with $\emptyset \vdash e_1 : (x:T'_1 \rightarrow T'_2)$ and $\emptyset \vdash e_2 : T'_1$, by inversion. By cases on the step taken.

(E_REDUCE/E_BETA): $(\lambda x:T_1. e_{12}) v_2 \rightarrow e_{12}[v_2/x]$. First, we have $\emptyset \vdash \lambda x:T_1. e_{12} : (x:T'_1 \rightarrow T'_2)$. By inversion for lambdas (Lemma 3.2.19), $x:T_1 \vdash e_{12} : T_2$. Moreover, $x:T_1 \rightarrow T_2 \equiv x:T'_1 \rightarrow T'_2$, which means $T_2 \equiv T'_2$ (Lemma 3.2.5).

By substitution, $\emptyset \vdash e_{12}[v_2/x] : T_2[v_2/x]$. We then see that $T_2[v_2/x] \equiv T'_2[v_2/x]$ (Lemma 3.2.8), so T_CONV completes this case.

(E_REDUCE/E_REFL): $\langle T \Rightarrow T \rangle^l v_2 \rightarrow v_2$. By cast inversion (Lemma 3.2.20), $_ : T \rightarrow T \equiv x:T'_1 \rightarrow T'_2$ and $\emptyset \vdash T$. In particular, we have $T \equiv T'_2$ and $T \equiv T'_1$ (Lemma 3.2.5). By substitutivity of conversion (Lemma 3.2.8), $T[v_2/x] \equiv T'_2[v_2/x]$. Since T is closed, we really know that $T \equiv T'_2[v_2/x]$.

By C_SYM and C_TRANS, we have $T'_1 \equiv T \equiv T'_2[v_2/x]$. By T_CONV on $\emptyset \vdash v_2 : T'_1$, we have $\emptyset \vdash v_2 : T'_2[v_2/x]$.

(E_REDUCE/E_FORGET): $\langle \{x:T_1 \mid e\} \Rightarrow T_2 \rangle^l v_2 \rightarrow \langle T_1 \Rightarrow T_2 \rangle^l v_2$. We restate the typing judgment and its inversion:

$$\begin{aligned} \emptyset \vdash \langle \{x:T_1 \mid e\} \Rightarrow T_2 \rangle^l v_2 &: T'_2[v_2/y] \\ \emptyset \vdash \langle \{x:T_1 \mid e\} \Rightarrow T_2 \rangle^l &: (y:T'_1 \rightarrow T'_2) \\ \emptyset \vdash v_2 &: T'_1 \end{aligned}$$

By cast inversion (Lemma 3.2.20), we know that $\emptyset \vdash T_1$ and $\emptyset \vdash T_2$ —as well as $_ : \{x:T_1 \mid e\} \rightarrow T_2 \equiv y:T'_1 \rightarrow T'_2$ and $\{x:T_1 \mid e\} \parallel T_2$. Inverting this conversion (Lemma 3.2.5), finding $\{x:T_1 \mid e\} \equiv T'_1$ and $T_2 \equiv T'_2$. Then by T_CONV and C_SYM, $\emptyset \vdash v_2 : \{x:T_1 \mid e\}$; by T_FORGET, $\emptyset \vdash v_2 : T_1$.

By T_CAST, we have $\emptyset \vdash \langle T_1 \Rightarrow T_2 \rangle^l : y:T_1 \rightarrow T_2$, with $T_1 \parallel T_2$ iff $\{x:T_1 \mid e\} \parallel T_2$. (Note, however, that y does not appear in T_2 —we write it to clarify the substitutions below.)

By T_APP, we find $\emptyset \vdash \langle T_1 \Rightarrow T_2 \rangle^l v_2 : T_2[v_2/y]$. Since $T_2 \equiv T'_2$, we have $T_2[v_2/y] \equiv T'_2[v_2/y]$ by Lemma 3.2.8. We are done by T_CONV.

(E_REDUCE/E_PRECHECK):

$$\begin{aligned} \langle T_1 \Rightarrow \{x:T_2 \mid e\} \rangle^l v_2 &\rightarrow \\ \langle T_2 \Rightarrow \{x:T_2 \mid e\} \rangle^l (\langle T_1 \Rightarrow T_2 \rangle^l v_2) & \end{aligned}$$

We restate the typing judgment and its inversion:

$$\begin{aligned} \emptyset \vdash \langle T_1 \Rightarrow \{x:T_2 \mid e\} \rangle^l v_2 &: T'_2[v_2/y] \\ \emptyset \vdash \langle T_1 \Rightarrow \{x:T_2 \mid e\} \rangle^l &: y:T'_1 \rightarrow T'_2 \\ \emptyset \vdash v_2 &: T'_1 \end{aligned}$$

By cast inversion (Lemma 3.2.20), $\emptyset \vdash T_1$ and $\emptyset \vdash \{x:T_2 \mid e\}$, and $y:T_1 \rightarrow \{x:T_2 \mid e\} \equiv y:T'_1 \rightarrow T'_2$. Also, $T_1 \parallel \{x:T_2 \mid e\}$.

By inversion on $\emptyset \vdash \{x:T_2 \mid e\}$, we find $\emptyset \vdash T_2$. Next, $T_1 \parallel T_2$ iff $T_1 \parallel \{x:T_2 \mid e\}$. Now by `T_CAST`, we find $\emptyset \vdash \langle T_1 \Rightarrow T_2 \rangle^l : y:T_1 \rightarrow T_2$. Note, however, that y doesn't occur in T_2 .

We can take the convertible function types and see that their parts are convertible: $T_1 \equiv T'_1$ and $\{x:T_2 \mid e\} \equiv T'_2$. Using the first conversion, we find $\emptyset \vdash v_2 : T_1$ by `T_CONV`. By `T_APP`, $\emptyset \vdash \langle T_1 \Rightarrow T_2 \rangle^l v_2 : T_2[v_2/y]$, where $T_2[v_2/y] = T_2$.

By `C_REFL` and `C_REFINER`, $T_2 \parallel \{x:T_2 \mid e\}$. We have well formedness derivations for both types, as well, so $\emptyset \vdash \langle T_2 \Rightarrow \{x:T_2 \mid e\} \rangle^l : y:T_2 \rightarrow \{x:T_2 \mid e\}$ by `T_CAST`. Again, y does not appear in e or T_2 . By `T_APP`, we have $\emptyset \vdash \langle T_2 \Rightarrow \{x:T_2 \mid e\} \rangle^l (\langle T_1 \Rightarrow T_2 \rangle^l v_2) : \{x:T_2 \mid e\}[\langle T_1 \Rightarrow T_2 \rangle^l v_2/y]$.

Since y isn't in $\{x:T_2 \mid e\}$, we can see:

$$\{x:T_2 \mid e\}[\langle T_1 \Rightarrow T_2 \rangle^l v_2/y] = \{x:T_2 \mid e\} = \{x:T_2 \mid e\}[v_2/y]$$

By substitutivity of conversion (Lemma 3.2.8), we have $\{x:T_2 \mid e\}[v_2/y] \equiv T'_2[v_2/y]$. We can now apply `T_CONV` to find $\emptyset \vdash \langle T_2 \Rightarrow \{x:T_2 \mid e\} \rangle^l (\langle T_1 \Rightarrow T_2 \rangle^l v_2) : T'_2[v_2/y]$.

(E_REDUCE/E_CHECK): $\langle T \Rightarrow \{x:T \mid e\} \rangle^l v_2 \longrightarrow \langle \{x:T \mid e\}, e[v_2/x], v_2 \rangle^l$.

We restate the typing judgment with its inversion:

$$\begin{aligned} \emptyset \vdash \langle T \Rightarrow \{x:T \mid e\} \rangle^l v_2 &: T'_2[v_2/y] \\ \emptyset \vdash \langle T \Rightarrow \{x:T \mid e\} \rangle^l &: y:T'_1 \rightarrow T'_2 \\ \emptyset \vdash v_2 &: T'_1 \end{aligned}$$

By cast inversion (Lemma 3.2.20), $\emptyset \vdash \{x:T \mid e\}$ and $\emptyset \vdash T$. Moreover, $y:T \rightarrow \{x:T \mid e\} \equiv y:T'_1 \rightarrow T'_2$, where y doesn't occur in $\{x:T \mid e\}$. This means that $T \equiv T'_1$ and $\{x:T \mid e\} \equiv T'_2$.

Using `T_CONV` and `C_SYM` with the first conversion shows $\emptyset \vdash v_2 : T$. By inversion on $\emptyset \vdash \{x:T \mid e\}$, we see $x:T \vdash e : \mathbf{Bool}$. By term substitution (Lemma 3.2.16), we find $\emptyset \vdash e[v_2/x] : \mathbf{Bool}$. Finally, $e[v_2/x] \longrightarrow^* e[v_2/x]$ by reflexivity (Lemma 3.2.3).

`T_CHECK` (with `WF_EMPTY`) shows $\emptyset \vdash \langle \{x:T \mid e\}, e[v_2/x], v_2 \rangle^l : \{x:T \mid e\}$. By substitutivity of conversion (Lemma 3.2.8), $\{x:T \mid e\}[v_2/y] \equiv T'_2[v_2/y]$.

Since y doesn't occur in $\{x:T \mid e\}$, we know that $\{x:T \mid e\}[v_2/y] = \{x:T \mid e\}$, so we can show that $\{x:T \mid e\} \equiv T_2[v_2/y]$ by C_SYM, and now $\emptyset \vdash \langle \{x:T \mid e\}, e[v_2/x], v_2 \rangle^l : T_2[v_2/y]$ by T_CONV.

(E_REDUCE/E_FUN):

$$\begin{aligned} & \langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l v_2 \longrightarrow \\ & \lambda x:T_{21}. (\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l (v_2 (\langle T_{21} \Rightarrow T_{11} \rangle^l x))) \end{aligned}$$

We restate the typing judgment with its inversion:

$$\begin{aligned} & \emptyset \vdash \langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l v_2 : T'_2[v_2/y] \\ & \emptyset \vdash \langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l : (y:T'_1 \rightarrow T'_2) \\ & \emptyset \vdash v_2 : T'_1 \end{aligned}$$

By cast inversion on the first derivation:

$$\begin{aligned} & \emptyset \vdash x:T_{11} \rightarrow T_{12} \quad \emptyset \vdash x:T_{21} \rightarrow T_{22} \\ & x:T_{11} \rightarrow T_{12} \parallel x:T_{21} \rightarrow T_{22} \\ & _:(x:T_{11} \rightarrow T_{12}) \rightarrow (x:T_{21} \rightarrow T_{22}) \equiv y:T'_1 \rightarrow T'_2 \end{aligned}$$

By inversion of this last (Lemma 3.2.5):

$$x:T_{11} \rightarrow T_{12} \equiv T'_1 \quad x:T_{21} \rightarrow T_{22} \equiv T'_2$$

So by T_CONV and C_SYM, we have $\emptyset \vdash v_2 : x:T_{11} \rightarrow T_{12}$. By weakening (Lemma 3.2.12), $x:T_{21} \vdash v_2 : x:T_{11} \rightarrow T_{12}$.

By inversion of the well formedness of the function types:

$$\emptyset \vdash T_{11} \quad x:T_{11} \vdash T_{12} \quad \emptyset \vdash T_{21} \quad x:T_{21} \vdash T_{22}$$

By weakening (Lemma 3.2.12), we find $x:T_{21} \vdash T_{11}$ and $x:T_{21} \vdash T_{21}$. By compatibility:

$$T_{11} \parallel T_{21} \quad T_{12} \parallel T_{22}$$

By T_CAST, $x:T_{21} \vdash \langle T_{21} \Rightarrow T_{11} \rangle^l : (_ : T_{21} \rightarrow T_{11})$ (notice that compatibility is symmetric, per Lemma 3.2.14). By T_APP and T_VAR, we can see $x:T_{21} \vdash \langle T_{21} \Rightarrow T_{11} \rangle^l x : T_{11}[x/_] = T_{11}$. Again by T_APP, we have $x:T_{21} \vdash v_2 (\langle T_{21} \Rightarrow T_{11} \rangle^l x) : T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x]$. By weakening (Lemma 3.2.12) and substitution (Lemma 3.2.16), we have the following two derivations:

$$x:T_{21} \vdash T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \quad x:T_{21} \vdash T_{22}$$

By T_CAST and Lemma 3.2.15:

$$x:T_{21} \vdash \langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l : (y:T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \rightarrow T_{22})$$

Noting that y is free here. By T_APP :

$$\begin{aligned} x:T_{21} \vdash & \langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l (v_2 (\langle T_{21} \Rightarrow T_{11} \rangle^l x)) \\ & : T_{22}[v_2 (\langle T_{21} \Rightarrow T_{11} \rangle^l x)/y] (= T_{22}) \end{aligned}$$

Finally, by T_ABS :

$$\emptyset \vdash \lambda x:T_{21}. \langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l (v_2 (\langle T_{21} \Rightarrow T_{11} \rangle^l x)) : x:T_{21} \rightarrow T_{22}$$

Since y isn't in $x:T_{21} \rightarrow T_{22}$, we can see that $x:T_{21} \rightarrow T_{22} = (x:T_{21} \rightarrow T_{22})[v_2/y]$. Using this fact with substitutivity of conversion (Lemma 3.2.8), we find $x:T_{21} \rightarrow T_{22} \equiv T'_2[v_2/y]$. So—finally—by T_CONV we have:

$$\emptyset \vdash \lambda x:T_{21}. \langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l (v_2 (\langle T_{21} \Rightarrow T_{11} \rangle^l x)) : T'_2[v_2/y]$$

(E_REDUCE/E_FORALL): $\langle \forall \alpha. T_1 \Rightarrow \forall \alpha. T_2 \rangle^l v_2 \longrightarrow (\Lambda \alpha. \langle T_1 \Rightarrow T_2 \rangle^l (v \alpha))$

We restate the typing and its inversion:

$$\begin{aligned} \emptyset \vdash & \langle \forall \alpha. T_1 \Rightarrow \forall \alpha. T_2 \rangle^l v_2 : T'_2[v_2/x] \\ \emptyset \vdash & \langle \forall \alpha. T_1 \Rightarrow \forall \alpha. T_2 \rangle^l : x:T'_1 \rightarrow T'_2 \\ \emptyset \vdash & v_2 : T'_1 \end{aligned}$$

By cast inversion (Lemma 3.2.20):

$$\begin{aligned} \emptyset \vdash \forall \alpha. T_1 & \quad \emptyset \vdash \forall \alpha. T_2 \\ \forall \alpha. T_1 \parallel \forall \alpha. T_2 & \\ -:(\forall \alpha. T_1) \rightarrow (\forall \alpha. T_2) \equiv x:T'_1 \rightarrow T'_2 & \end{aligned}$$

By inversion of this last $\forall \alpha. T_1 \equiv T'_1$ and $\forall \alpha. T_2 \equiv T'_2$ (Lemma 3.2.5). By T_CONV and C_SYM , $\emptyset \vdash v_2 : \forall \alpha. T_1$. By type variable weakening, WF_TVAR , and T_TAPP , we have:

$$\alpha \vdash v_2 \alpha : T_1[\alpha/\alpha] = T_1$$

By inversion of the universal types's well formedness and compatibility:

$$\alpha \vdash T_1 \quad \alpha \vdash T_2 \quad T_1 \parallel T_2$$

So by T_CAST , $\alpha \vdash \langle T_1 \Rightarrow T_2 \rangle^l : (x:T_1 \rightarrow T_2)$, noting that x doesn't occur in T_2 . By T_APP , $\alpha \vdash \langle T_1 \Rightarrow T_2 \rangle^l (v_2 \alpha) : T_2[v_2 \alpha/x] = T_2$. By T_TABS , $\emptyset \vdash \Lambda \alpha. (\langle T_1 \Rightarrow T_2 \rangle^l (v \alpha)) : \forall \alpha. T_2$.

We know that $\forall \alpha. T_2 \equiv T'_2$, so by type variable substitutivity of conversion (Lemma 3.2.9), $(\forall \alpha. T_2)[v_2/x] \equiv T'_2[v_2/x]$. Since x isn't in $\forall \alpha. T_2$, we know that $\forall \alpha. T_2 \equiv T'_2[v_2/x]$ (by way of Lemma 3.2.8).. Now we can see by T_CONV that $\emptyset \vdash \Lambda \alpha. (\langle T_1 \Rightarrow T_2 \rangle^l (v \alpha)) : T'_2[v_2/x]$.

(E_COMPAT): $E [e] \longrightarrow E [e']$ when $e \longrightarrow e'$ By cases on E:

($E = [] e_2, e_1 \longrightarrow e'_1$): By the IH and T_APP.

($E = v_1 [], e_2 \longrightarrow e'_2$): By the IH, T_APP, and T_CONV, since $T_2[e_2/x] \equiv T_2[e'_2/x]$ by reflexivity (Lemma 3.2.3) and substitutivity (Lemma 3.2.8).

(E_BLADE): $E [\uparrow l] \longrightarrow \uparrow l \emptyset \vdash E [\uparrow l] : T$ by assumption. By type well formedness (Lemma 3.2.26), we know that $\emptyset \vdash T$. We then have $\emptyset \vdash \uparrow l : T$ by T_BLADE.

(T_TABS): $\emptyset \vdash \Lambda \alpha. e : \forall \alpha. T$. This case is contradictory—values don't step.

(T_TAPP): $\emptyset \vdash e T : T'[T/\alpha]$. By cases on the step taken.

(E_REDUCE/E_TBETA): $(\Lambda \alpha. e') T \longrightarrow e'[T/\alpha]$ We restate the typing derivation and its inversion:

$$\emptyset \vdash (\Lambda \alpha. e') T : T'[T/\alpha] \quad \emptyset \vdash \Lambda \alpha. e' : \forall \alpha. T'' \quad \emptyset \vdash T$$

By type abstraction inversion (Lemma 3.2.22): $\alpha \vdash e' : T''$ and $\forall \alpha. T'' \equiv \forall \alpha. T'$; by inversion of this last (Lemma 3.2.7), $T'' \equiv T'$.

By type variable substitution (Lemma 3.2.18), $\emptyset \vdash e'[T/\alpha] : T''[T/\alpha]$. By type substitutivity of conversion (Lemma 3.2.9), $T''[T/\alpha] \equiv T'[T/\alpha]$. T_CONV gives us $\emptyset \vdash e'[T/\alpha] : T'[T/\alpha]$ as desired.

(E_COMPAT): $E [e] \longrightarrow E [e']$, where $E = [] T$. By the IH and T_TAPP.

(E_BLADE): $E [\uparrow l] \longrightarrow \uparrow l. \emptyset \vdash E [\uparrow l] : T$ by assumption. By type well formedness (Lemma 3.2.26), we know that $\emptyset \vdash T$. So we see $\emptyset \vdash \uparrow l : T$ by T_BLADE.

(T_CAST): $\emptyset \vdash \langle T_1 \Rightarrow T_2 \rangle^l : T_1 \rightarrow T_2$. This case is contradictory—values don't step.

(T_CHECK): $\emptyset \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle^l : \{x:T \mid e_1\}$. By cases on the step taken.

(E_REDUCE/E_OK): $\langle \{x:T \mid e_1\}, \text{true}, v \rangle^l \longrightarrow v$. By inversion, $\emptyset \vdash v : T$ and $\emptyset \vdash \{x:T \mid e\}$; we also have $e_1[v/x] \longrightarrow^* \text{true}$. By WF_EMPTY and the assumption that $e[v/x] \longrightarrow^* \text{true}$, we can find $\emptyset \vdash v : \{x:T \mid e\}$ by T_EXACT.

(E_REDUCE/E_FAIL): $\langle \{x:T \mid e_1\}, \text{false}, v \rangle^l \longrightarrow \uparrow l$ We have $\emptyset \vdash \{x:T \mid e\}$ by inversion. By WF_EMPTY and T_BLADE, $\emptyset \vdash \uparrow l : \{x:T \mid e\}$.

(E_COMPAT): $E [e] \longrightarrow E [e']$, where $E = \langle \{x:T \mid e_1\}, [], v \rangle^l$. By the IH on e , we know that $\emptyset \vdash e' : \text{Bool}$. We still have $\emptyset \vdash \{x:T \mid e_1\}$ and $\emptyset \vdash v : T$ from our original derivation. Since $e_1[v/x] \longrightarrow^* e$ and $e \longrightarrow e'$, then $e_1[v/x] \longrightarrow^* e'$. Therefore, $\emptyset \vdash \langle \{x:T \mid e_1\}, e', v \rangle^l : \{x:T \mid e_1\}$ by T_CHECK.

(E_BLAKE): $E[\uparrow l] \longrightarrow \uparrow l$. $\emptyset \vdash E[\uparrow l] : T$ by assumption. By type well formedness (Lemma 3.2.26), we know that $\emptyset \vdash T$. So $\emptyset \vdash \uparrow l : T$ by T_BLAKE.

(T_BLAKE): $\emptyset \vdash \uparrow l : T$. This case is contradictory—blame doesn't step.

(T_CONV): $\emptyset \vdash e : T'$; by inversion we have $\emptyset \vdash e : T$ and $T \equiv T'$ and $\emptyset \vdash T'$ (and, trivially, $\vdash \emptyset$). By the IH on the first derivation, we know that $\emptyset \vdash e' : T$. By T_CONV, we can see that $\emptyset \vdash e' : T'$.

(T_EXACT): $\emptyset \vdash v : \{x:T \mid e\}$. This case is contradictory—values don't step.

(T_FORGET): $\emptyset \vdash v : T$. This case is contradictory—values don't step. \square

\square

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. In fact, after cutting out universal types and restricting refinements to base types, it's possible to simplify the operational semantics and to do away with the T_FORGET rule. I don't give the proof here because it is subsumed by type soundness in F_H ; the CAST system in Chapter 4 follows this principle, though it uses explicit tags and also has the dynamic type Dyn.

3.3 Parametricity

I prove relational parametricity for two reasons: (1) it yields powerful reasoning techniques such as free theorems [76] and our subtyping proof in Section 3.4, and (2) it indicates that contracts don't interfere with type abstraction, i.e., that F_H supports polymorphism in the same way that System F does. The proof is mostly standard: I define a (syntactic) logical relation on terms and types, 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, I 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. I conjecture that logically related terms are *contextually equivalent* and upcasts can be eliminated without changing the meaning of a program.

We begin by defining two relations: $r_1 \sim r_2 : T; \theta; \delta$ relates closed results, defined by induction on types; $e_1 \simeq e_2 : T; \theta; \delta$ relates closed expressions which evaluate to results in the first relation. The definitions are shown in Figure 3.5.² 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 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 . A term substitution δ maps

²To save space, I write $\uparrow l \sim \uparrow l : T; \theta; \delta$ separately instead of manually adding such a clause for each type.

$$\begin{array}{l}
\text{Closed terms } \boxed{r_1 \sim r_2 : T; \theta; \delta} \quad \boxed{e_1 \simeq e_2 : T; \theta; \delta} \\
k \sim k : B; \theta; \delta \iff k \in \mathcal{K}_B \\
v_1 \sim v_2 : \alpha; \theta; \delta \iff \exists R T_1 T_2, \alpha \mapsto R, T_1, T_2 \in \theta \wedge v_1 R v_2 \\
v_1 \sim v_2 : (x:T_1 \rightarrow T_2); \theta; \delta \iff \forall v'_1 \sim v'_2 : T_1; \theta; \delta, v_1 v'_1 \simeq v_2 v'_2 : T_2; \theta; \delta[v'_1, v'_2/x] \\
v_1 \sim v_2 : \forall \alpha. T; \theta; \delta \iff \forall R T_1 T_2, v_1 T_1 \simeq v_2 T_2 : T; \theta[\alpha \mapsto R, T_1, T_2]; \delta \\
v_1 \sim v_2 : \{x:T \mid e\}; \theta; \delta \iff v_1 \sim v_2 : T; \theta; \delta \wedge \\
\theta_1(\delta_1(e))[v_1/x] \longrightarrow^* \text{true} \wedge \theta_2(\delta_2(e))[v_2/x] \longrightarrow^* \text{true} \\
\uparrow l \sim \uparrow l : T; \theta; \delta \\
e_1 \simeq e_2 : T; \theta; \delta \iff \exists r_1 r_2, e_1 \longrightarrow^* r_1 \wedge e_2 \longrightarrow^* r_2 \wedge r_1 \sim r_2 : T; \theta; \delta \\
R_{T, \theta, \delta} = \{(r_1, r_2) \mid r_1 \sim r_2 : T; \theta; \delta\}
\end{array}$$

$$\begin{array}{l}
\text{Types } \boxed{T_1 \simeq T_2 : *; \theta; \delta} \\
B \simeq B : *; \theta; \delta \\
\alpha \simeq \alpha : *; \theta; \delta \\
x:T_{11} \rightarrow T_{12} \simeq x:T_{21} \rightarrow T_{22} : *; \theta; \delta \iff T_{11} \simeq T_{21} : *; \theta; \delta \wedge \\
\forall v_1 \sim v_2 : T_{11}; \theta; \delta, \\
T_{12} \simeq T_{22} : *; \theta; \delta[v_1, v_2/x] \\
\forall \alpha. T_1 \simeq \forall \alpha. T_2 : *; \theta; \delta \iff \forall R T'_1 T'_2, T_1 \simeq T_2 : *; \theta[\alpha \mapsto R, T'_1, T'_2]; \delta \\
\{x:T_1 \mid e_1\} \simeq \{x:T_2 \mid e_2\} : *; \theta; \delta \iff T_1 \simeq T_2 : *; \theta; \delta \wedge \\
\forall v_1 \sim v_2 : T_1; \theta; \delta, \theta_1(\delta_1(e_1))[v_1/x] \simeq \theta_2(\delta_2(e_2))[v_2/x] : \text{Bool}; \theta; \delta
\end{array}$$

$$\begin{array}{l}
\text{Open terms and types } \boxed{\Gamma \vdash \theta; \delta} \quad \boxed{\Gamma \vdash e_1 \simeq e_2 : T} \quad \boxed{\Gamma \vdash T_1 \simeq T_2 : *} \\
\Gamma \vdash \theta; \delta \iff \forall x:T \in \Gamma, \theta_1(\delta_1(x)) \simeq \theta_2(\delta_2(x)) : T; \theta; \delta \wedge \\
\forall \alpha \in \Gamma, \exists R T_1 T_2, \alpha \mapsto R, T_1, T_2 \in \theta \\
\Gamma \vdash e_1 \simeq e_2 : T \iff \forall \Gamma \vdash \theta; \delta, \theta_1(\delta_1(e_1)) \simeq \theta_2(\delta_2(e_2)) : T; \theta; \delta \\
\Gamma \vdash T_1 \simeq T_2 : * \iff \forall \Gamma \vdash \theta; \delta, T_1 \simeq T_2 : *; \theta; \delta
\end{array}$$

Figure 3.5: The logical relation for parametricity

from variables to pairs of closed values. I write projections δ_i ($i = 1, 2$) to denote projections from this pair. I similarly write θ_i ($i = 1, 2$) for a substitution that maps a type variable α to T_i where $\theta(\alpha) = (R, T_1, T_2)$.

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. I require that both values satisfy their refinements rather than having the first satisfy the predicate iff the second does because I want to know that values related at refinement types *actually inhabit those types*, i.e., actually satisfy the predicates of the refinement. The \sim relation on results is extended to the relation \simeq on closed terms in a straightforward manner:

terms are related if and only if they both terminate at related results. Divergent terms are not related to each other—though we will discover that divergent terms do not exist in F_H . I 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 3.5).

To show that (well-typed) casts yield related results when applied to related inputs, we also need a relation on types $T_1 \simeq T_2 : *; \theta; \delta$; I define this relation in Figure 3.5. We can use the logical relation on terms to handle the arguments of function types and refinement types. Note that the T_1 and T_2 in this relation are not necessarily closed; terms in refinement types, which should be related at **Bool**, are closed by applying substitutions. In the function and 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, the second line of the function type case could change T_{11} to T_{21}). It does not really matter which side we choose, since they are related types. Here, we are “left-leaning”;³ in the proof, I justify this choice by proving a “type exchange” lemma (Lemma 3.4.2) 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 the type relation to open terms, writing $\Gamma \vdash T_1 \simeq T_2 : *$ when two types are equivalent for any Γ -respecting substitutions.

It is worth discussing two points peculiar to this formulation: terms in the logical relation are untyped, and the type indices are open.

I allow any relation on terms 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., [54]). I suspect that part of the reason this proof of parametricity works is its similarity to Girard’s untyped reducibility candidates. I have two motivations for leaving the relations untyped. First, functions of type $x:T_1 \rightarrow T_2$ must map 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 the type relation, terms that are well typed at one type won’t necessarily be well typed at the other, whether definitions are left- or right-leaning. Second, we prove in Section 3.4 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, I 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 $T_1 \rightarrow T_1$. There is one small hitch: $\lambda x:T_1. x$ has type $T_1 \rightarrow T_1$, not $T_1 \rightarrow T_2$! I therefore don’t demand that two expressions related at T be well typed at T , and I allow *any* relation to be chosen as R .

The type indices of the term relation are not necessarily closed. Instead, 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 $x:T_1 \rightarrow T_2$:

³I, too, lean quite far to the left.

the *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 $v_1 v'_1 \simeq v_2 v'_2$ at some type. The question here is which type index we should use. If we keep 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 [44] 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. I think this formulation is more uniform with regard to free term/type variables, since eager substitution is a non-starter for type variables, anyway.

As Atsushi and I developed the 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 decomposing 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 this syntactic setting. On the one hand, the evaluation of well-typed programs never encounters casts with uninstantiated type variables—a key property of our evaluation relation. On the other hand, by parametricity, we expect every value of type $\forall \alpha. \alpha \rightarrow \alpha$ 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 $\mathbf{Bool} \rightarrow \mathbf{Bool}$, and a function f of type $\mathbf{Bool} \rightarrow \mathbf{Bool}$ would return, by E_FUN, a wrapped version of f that is syntactically different from the f we passed in—that is, the function broke parametricity! We expect the returned value should behave the same as the input, though—the results are just *syntactically* different. 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.

Now we can set about proving parametricity (Lemma 3.3.7). We begin with compositionality theorems relating the closing substitutions θ and δ to substitution in terms (Lemma 3.3.1) and types (Lemma 3.3.3); convertibility shows that our logical relation relates terms at convertible types (Lemma 3.3.4); after some lemmas about casts and a separate induction relating casts between related types (Lemma 3.3.6), we prove parametricity.

3.3.1 Lemma [Term compositionality]: If $\delta_1(e) \longrightarrow^* e_1$ and $\delta_2(e) \longrightarrow^* e_2$ then $r_1 \sim r_2 : T; \theta; \delta[e_1, e_2/x]$ iff $r_1 \sim r_2 : T[e/x]; \theta; \delta$.

Proof: By induction on the (simple) structure of T , proving both directions simultaneously. We treat the case where $r_1 = r_2 = \uparrow l$ separately from the induction, since it’s the same easy proof in all cases: $\uparrow l \sim \uparrow l : T; \theta; \delta$ irrespective of T and δ . So for the rest of proof, we know $r_1 = v_1$ and $r_2 = v_2$. Only the refinement case is interesting.

$(T = \{y:T' \mid e'\})$: We show both directions simultaneously, where $x \neq y$, i.e., y is fresh. By the IH for T' , we know that

$$v_1 \sim v_2 : T'; \theta; \delta[e_1, e_2/x] \text{ iff } v_1 \sim v_2 : T'[e/x]; \theta; \delta.$$

It remains to show that the values satisfy their refinements.

That is, we must show:

$$\begin{aligned} \theta_1(\delta_1(e'[e_1/x][v_1/y])) &\longrightarrow^* \text{true} \text{ iff } \theta_1(\delta_1(e'[e/x][v_1/y])) \longrightarrow^* \text{true} \\ \theta_2(\delta_2(e'[e_2/x][v_2/y])) &\longrightarrow^* \text{true} \text{ iff } \theta_2(\delta_2(e'[e/x][v_2/y])) \longrightarrow^* \text{true} \end{aligned}$$

So let:

$$\begin{aligned} \sigma_1 &= \theta_1 \delta_1[\delta_1(e)/x, v_1/y] \longrightarrow^* \theta_1 \delta_1[e_1/x, v_1/y] = \sigma'_1 \\ \sigma_2 &= \theta_2 \delta_2[\delta_2(e)/x, v_2/y] \longrightarrow^* \theta_2 \delta_2[e_2/x, v_2/y] = \sigma'_2 \end{aligned}$$

We have $\sigma_1 \longrightarrow^* \sigma'_1$ by reflexivity except for $\delta_1(e) \longrightarrow^* e_1$, which we have by assumption; likewise, we have $\sigma_2 \longrightarrow^* \sigma'_2$. Then $\sigma_i(e')$ and $\sigma'_i(e')$ coterminate (Lemma 3.2.1), and we are done. \square

3.3.2 Lemma [Term Weakening/Strengthening]: If $x \notin T$, then $r_1 \sim r_2 : T; \theta; \delta[e_1, e_2/x]$ iff $r_1 \sim r_2 : T; \theta; \delta$.

Proof: Similar to Lemma 3.3.1. \square

3.3.3 Lemma [Type compositionality]:

$$r_1 \sim r_2 : T; \theta[\alpha \mapsto R_{T', \theta, \delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta \text{ iff } r_1 \sim r_2 : T[T'/\alpha]; \theta; \delta.$$

Proof: By induction on the (simple) structure of T , proving both directions simultaneously. As for Lemma 3.3.1, we treat the case where $r_1 = r_2 = \uparrow l$ separately from the induction, since it's the same easy proof in all cases: $\uparrow l \sim \uparrow l : T; \theta; \delta$ irrespective of T and δ . So for the rest of proof, we know $r_1 = v_1$ and $r_2 = v_2$. Here, the interesting case is for function types, where we must deal with some asymmetries in the definition of the logical relation. I also include the case for quantified types.

$(T = x:T_1 \rightarrow T_2)$: There are two cases:

(\Rightarrow) : Given $v_1 \sim v_2 : (x:T_1 \rightarrow T_2); \theta[\alpha \mapsto R_{T', \theta, \delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta$, we wish to show that $v_1 \sim v_2 : (x:T_1 \rightarrow T_2)[T'/\alpha]; \theta; \delta$. Let $v'_1 \sim v'_2 : T_1[T'/\alpha]; \theta; \delta$. We must show that $v_1 v'_1 \simeq v_2 v'_2 : T_2[T'/\alpha]; \theta; \delta[v'_1, v'_2/x]$. By the IH on T_1 , $v'_1 \sim v'_2 : T_1; \theta[\alpha \mapsto R_{T', \theta, \delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta$. By assumption,

$$v_1 v'_1 \simeq v_2 v'_2 : T_2; \theta[\alpha \mapsto R_{T', \theta, \delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta[v'_1, v'_2/x].$$

These normalize to $r'_1 \sim r'_2 : T_2; \theta[\alpha \mapsto R_{T', \theta, \delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta[v'_1, v'_2/x]$. Since $x \notin T'$, Lemma 3.3.2 gives $R_{T', \theta, \delta} = R_{T', \theta, \delta[v'_1, v'_2/x]}$ and so

$$r'_1 \sim r'_2 : T_2; \theta[\alpha \mapsto R_{T', \theta, \delta[v'_1, v'_2/x]}, \theta_1(\delta_1(T'[v'_1/x])), \theta_2(\delta_2(T'[v'_2/x]))]; \delta[v'_1, v'_2/x].$$

By the IH on T_2 , $r'_1 \sim r'_2 : T_2[T'/\alpha]; \theta; \delta[v'_1, v'_2/x]$. By expansion, $v_1 v'_1 \simeq v_2 v'_2 : T_2[T'/\alpha]; \theta; \delta[v'_1, v'_2/x]$.

(\Leftarrow): This case is similar: Given $v_1 \sim v_2 : (x:T_1 \rightarrow T_2)[T'/\alpha]; \theta; \delta$, we wish to show that $v_1 \sim v_2 : (x:T_1 \rightarrow T_2); \theta[\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta$. Let $v'_1 \sim v'_2 : T_1; \theta[\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta$. We must show that

$$v_1 v'_1 \simeq v_2 v'_2 : T_2; \theta[\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta[v'_1, v'_2/x].$$

By the IH on T_1 , $v'_1 \sim v'_2 : T_1[T'/\alpha]; \theta; \delta$. By assumption, $v_1 v'_1 \simeq v_2 v'_2 : T_2[T'/\alpha]; \theta; \delta[v'_1, v'_2/x]$. These normalize to $r'_1 \simeq r'_2 : T_2[T'/\alpha]; \theta; \delta[v'_1, v'_2/x]$. By the IH on T_2 ,

$$\begin{aligned} r'_1 \simeq r'_2 : & T_2[T'/\alpha]; \\ & \theta[\alpha \mapsto R_{T',\theta,\delta[v'_1, v'_2/x]}, \theta_1(\delta_1(T'[v'_1/x])), \theta_2(\delta_2(T'[v'_2/x]))]; \\ & \delta[v'_1, v'_2/x]. \end{aligned}$$

Since $x \notin T'$, Lemma 3.3.2 gives

$$r'_1 \simeq r'_2 : T_2[T'/\alpha]; \theta[\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta[v'_1, v'_2/x].$$

Finally, by expansion

$$\begin{aligned} v_1 v'_1 \simeq v_2 v'_2 : & T_2[T'/\alpha]; \\ & \theta[\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \\ & \delta[v'_1, v'_2/x]. \end{aligned}$$

($T = \forall \alpha'. T_0$): There are two cases:

(\Rightarrow): Given $v_1 \sim v_2 : \forall \alpha'. T_0; \theta[\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta$, we wish to show that $v_1 \sim v_2 : \forall \alpha'. (T_0[T'/\alpha]); \theta; \delta$. Let a relation R and closed types T_1 and T_2 be given. By assumption, we know that $v_1 T_1 \simeq v_2 T_2 : T_0; \theta[\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))][\alpha' \mapsto R, T_1, T_2]; \delta$. They normalize to $r'_1 \sim r'_2 : T_0; \theta[\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))][\alpha' \mapsto R, T_1, T_2]; \delta$. By the IH, $r'_1 \sim r'_2 : T_0[T'/\alpha]; \theta[\alpha' \mapsto R, T_1, T_2]; \delta$. By expansion, $v_1 T_1 \simeq v_2 T_2 : T_0[T'/\alpha]; \theta[\alpha' \mapsto R, T_1, T_2]; \delta$. Then, $v_1 \sim v_2 : \forall \alpha'. (T_0[T'/\alpha]); \theta; \delta$.

(\Leftarrow): This case is similar: given $v_1 \sim v_2 : \forall \alpha'. (T_0[T'/\alpha]); \theta; \delta$, we wish to show that $v_1 \sim v_2 : \forall \alpha'. T_0; \theta[\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta$. Let a relation R and closed types T_1 and T_2 be given. By assumption, we know that $v_1 T_1 \simeq v_2 T_2 : T_0[T'/\alpha]; \theta[\alpha' \mapsto R, T_1, T_2]; \delta$. They normalize to $r'_1 \sim r'_2 : T_0[T'/\alpha]; \theta[\alpha' \mapsto R, T_1, T_2]; \delta$. By the IH, $r'_1 \sim r'_2 : T_0; \theta[\alpha' \mapsto R, T_1, T_2][\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta$. By expansion, $v_1 T_1 \simeq v_2 T_2 : T_0; \theta[\alpha' \mapsto R, T_1, T_2][\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta$. Then, $v_1 \sim v_2 : \forall \alpha'. T_0; \theta[\alpha \mapsto R_{T',\theta,\delta}, \theta_1(\delta_1(T')), \theta_2(\delta_2(T'))]; \delta$.

□

□

Complexity of casts

$$\begin{aligned}
cc(\langle T \Rightarrow T \rangle^l) &= 1 \\
cc(\langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l) &= cc(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l) + \\
&\quad cc(\langle T_{21} \Rightarrow T_{11} \rangle^l) + 1 \\
cc(\langle \forall \alpha. T_1 \Rightarrow \forall \alpha. T_2 \rangle^l) &= cc(\langle T_1 \Rightarrow T_2 \rangle^l) + 1 \\
cc(\langle \{x:T_1 \mid e\} \Rightarrow T_2 \rangle^l) &= cc(\langle T_1 \Rightarrow T_2 \rangle^l) + 1 \\
&\quad (\text{if } T_2 \neq \{x:T_1 \mid e\} \text{ and } T_2 \neq \{y:\{x:T_1 \mid e\} \mid e'\}) \\
cc(\langle T_1 \Rightarrow \{x:T_1 \mid e\} \rangle^l) &= 1 \\
cc(\langle T_1 \Rightarrow \{x:T_2 \mid e\} \rangle^l) &= cc(\langle T_1 \Rightarrow T_2 \rangle^l) + 2 \\
&\quad (\text{if } T_1 \neq T_2 \text{ and } T_1 \text{ is not a refinement type})
\end{aligned}$$

Figure 3.6: Complexity of casts

3.3.4 Lemma [Convertibility]: If $T_1 \equiv T_2$ then $r_1 \sim r_2 : T_1; \theta; \delta$ iff $r_1 \sim r_2 : T_2; \theta; \delta$.

Proof: By induction on the conversion relation, leaving θ and δ general. The case where $r_1 = r_2 = \uparrow^l$ is immediate, so we only need to consider the case where $r_1 = v_1$ and $r_2 = v_2$.

(C_VAR): It must be that $T_1 = T_2 = \alpha$, so we are done immediately.

(C_BASE): It must be that $T_1 = T_2 = B$, so we are done immediately.

(C_REFINE): We have that $T_1 = \{x:T'_1 \mid \sigma_1(e)\}$ and $T_2 = \{x:T'_2 \mid \sigma_2(e)\}$, where $T'_1 \equiv T'_2$ and $\sigma_1 \longrightarrow^* \sigma_2$.

We have $\theta_1 \delta_1 \sigma_1 \longrightarrow^* \theta_1 \delta_1 \sigma_2$ and $\theta_2 \delta_2 \sigma_1 \longrightarrow^* \theta_2 \delta_2 \sigma_2$ (by reflexivity of \longrightarrow^* for θ_i and δ_i and assumption for $\sigma_1 \longrightarrow^* \sigma_2$), so by cotermination (Lemma 3.2.1):

$$\begin{aligned}
\theta_1(\delta_1(\sigma_1(e)))[v_1/x] &\longrightarrow^* \text{true iff } \theta_1(\delta_1(\sigma_2(e)))[v_1/x] \longrightarrow^* \text{true} \\
\theta_2(\delta_2(\sigma_1(e)))[v_2/x] &\longrightarrow^* \text{true iff } \theta_2(\delta_2(\sigma_2(e)))[v_2/x] \longrightarrow^* \text{true}
\end{aligned}$$

(C_FUN): We have that $T_1 = x:T_{11} \rightarrow T_{12} \equiv x:T_{21} \rightarrow T_{22} = T_2$.

Let $v'_1 \sim v'_2 : T_{21}; \theta; \delta$ be given; we must show that $v_1 v'_1 \simeq v_2 v'_2 : T_{22}; \theta; \delta[v'_1, v'_2/x]$.

By the IH, we know that $v'_1 \sim v'_2 : T_{11}; \theta; \delta$, so we know that $v_1 v'_1 \simeq v_2 v'_2 : T_{12}; \theta; \delta[v'_1, v'_2/x]$. We are done by another application of the IH.

The other direction is similar.

(C_FORALL): We have that $T_1 = \forall \alpha. T'_1 \equiv \forall \alpha. T'_2 = T_2$.

Let R, T , and T' be given. We must show that $v_1 T \simeq v_2 T' : T'_2; \theta[\alpha \mapsto R, T, T']; \delta$. We know that $v_1 T \simeq v_2 T' : T'_1; \theta[\alpha \mapsto R, T, T']; \delta$, so we are done by the IH.

The other direction is similar.

(C_SYM): By the IH.

(C_TRANS): By the IHs. □

Before we can show parametricity (Lemma 3.3.7), we prove in a separate induction that casts between related types are related (Lemma 3.3.6). Before we can prove *that* lemma, we need an auxiliary fact about casts and substitution.

3.3.5 Lemma [Cast substitution]: If $\vdash \Gamma, x:T_1$, $\Gamma, x:T_1 \vdash \langle T_1 \Rightarrow T_2 \rangle^l \simeq \langle T_1 \Rightarrow T_2 \rangle^l : T_2$, and $\Gamma, x:T_2 \vdash e_1 \simeq e_2 : T$ then $\Gamma_1, x:T_1, \Gamma_2 \vdash e_1[\langle T_1 \Rightarrow T_2 \rangle^l x/x] \simeq e_2[\langle T_1 \Rightarrow T_2 \rangle^l x/x] : T$.

Proof: Let $\Gamma \vdash \theta; \delta$. We must show that

$$\theta_1(\delta_1(e_1[\langle T_1 \Rightarrow T_2 \rangle^l x/x])) \simeq \theta_2(\delta_2(e_2[\langle T_1 \Rightarrow T_2 \rangle^l x/x])) : T; \theta; \delta.$$

Now $\theta_1(\delta_1(x)) \simeq \theta_2(\delta_2(x)) : T_1; \theta; \delta$ by definition. By assumption, $\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l x)) \simeq \theta_2(\delta_2(\langle T_1 \Rightarrow T_2 \rangle^l x)) : T_2; \theta; \delta$. So let δ' be $\delta[\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l x), \delta_2(\langle T_1 \Rightarrow T_2 \rangle^l x)/x]$. We have $\Gamma, x:T_2 \vdash \theta; \delta'$, so by assumption we have that $\theta_1(\delta'_1(e_1)) \simeq \theta_2(\delta'_2(e_2)) : T; \theta; \delta$, which is the same as $\theta_1(\delta_1(e_1[\langle T_1 \Rightarrow T_2 \rangle^l x/x])) \simeq \theta_2(\delta_2(e_2[\langle T_1 \Rightarrow T_2 \rangle^l x/x])) : T; \theta; \delta$, and we are done. \square

We show that (well typed) casts relate to themselves by induction a cast complexity metric, cc , defined in Figure 3.6. The complexity of a cast is the number of steps it and its subparts can take. This definition is carefully dependent on our definition of type compatibility and our cast reduction rules. Doing induction on this metric greatly simplifies the proof: we show that stepping casts at related types yields either related non-casts, or lower complexity casts between related types.

3.3.6 Lemma [Cast Reflexivity]: If $\vdash \Gamma$, $T_1 \parallel T_2$, $\Gamma \vdash T_1 \simeq T_1 : *$, and $\Gamma \vdash T_2 \simeq T_2 : *$, then $\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle^l \simeq \langle T_1 \Rightarrow T_2 \rangle^l : T_1 \rightarrow T_2$.

Proof: By induction on $cc(\langle T_1 \Rightarrow T_2 \rangle^l)$.

$(T_1 = T_2)$: Given $\Gamma \vdash \theta; \delta$, we wish to show that

$$\langle \theta_1(\delta_1(T_1)) \Rightarrow \theta_1(\delta_1(T_1)) \rangle^l \sim \langle \theta_2(\delta_2(T_1)) \Rightarrow \theta_2(\delta_2(T_1)) \rangle^l : T_1 \rightarrow T_1; \theta; \delta.$$

Let $v_1 \sim v_2 : T_1; \theta; \delta$. We must show that

$$\begin{aligned} & \langle \theta_1(\delta_1(T_1)) \Rightarrow \theta_1(\delta_1(T_1)) \rangle^l v_1 \simeq \\ & \langle \theta_2(\delta_2(T_1)) \Rightarrow \theta_2(\delta_2(T_1)) \rangle^l v_2 : T_1; \theta; \delta[v_1, v_2/z] \end{aligned}$$

for fresh z . By E_REFLECT, these normalize to $v_1 \sim v_2 : T_1; \theta; \delta[v_1, v_2/z]$. Lemma 3.3.2 finishes the case.

$(T_1 = x:T_{11} \rightarrow T_{12}$ and $T_2 = x:T_{21} \rightarrow T_{22}$ and $T_1 \neq T_2)$: Then, we have:

$$\begin{array}{ccc} T_{11} \parallel T_{21} & & T_{21} \parallel T_{22} \\ & \Gamma \vdash T_{11} & \\ \Gamma, x:T_{11} \vdash T_{12} & \Gamma, x:T_{21} \vdash T_{22}. & \end{array}$$

Given $\Gamma \vdash \theta; \delta$, we wish to show that

$$\theta_1(\delta_1(\langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l)) \sim \theta_1(\delta_1(\langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l)) \\ : (x:T_{11} \rightarrow T_{12}) \rightarrow (x:T_{21} \rightarrow T_{22}); \theta; \delta.$$

Let $v_1 \sim v_2 : x:T_{11} \rightarrow T_{12}; \theta; \delta$. We must show that

$$\theta_1(\delta_1(\langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l)) v_1 \simeq \\ \theta_1(\delta_1(\langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l)) v_2 : x:T_{21} \rightarrow T_{22}; \theta; \delta[v_1, v_2/z]$$

for fresh z . Let

$$v'_1 = \theta_1(\delta_1(\lambda x:T_{21}. (\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l (v_1 (\langle T_{21} \Rightarrow T_{11} \rangle^l x)))) \\ v'_2 = \theta_2(\delta_2(\lambda x:T_{21}. (\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l (v_2 (\langle T_{21} \Rightarrow T_{11} \rangle^l x))))).$$

Since

$$\theta_1(\delta_1(\langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l)) v_1 \longrightarrow v'_1 \\ \theta_2(\delta_2(\langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l)) v_2 \longrightarrow v'_2,$$

it suffices to show that

$$v'_1 \sim v'_2 : x:T_{21} \rightarrow T_{22}; \theta; \delta[v_1, v_2/z].$$

Let $v''_1 \sim v''_2 : T_{21}; \theta; \delta[v_1, v_2/z]$. We must show that

$$v'_1 v''_1 \simeq v'_2 v''_2 : T_{22}; \theta; \delta[v_1, v_2/z][v''_1, v''_2/x].$$

Since

$$v'_1 v''_1 \longrightarrow \theta_1(\delta_1(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v''_1/x] \Rightarrow T_{22} \rangle^l (v_1 (\langle T_{21} \Rightarrow T_{11} \rangle^l v''_1)))) \\ v'_2 v''_2 \longrightarrow \theta_2(\delta_2(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v''_2/x] \Rightarrow T_{22} \rangle^l (v_2 (\langle T_{21} \Rightarrow T_{11} \rangle^l v''_2))))),$$

it suffices to show that

$$\theta_1(\delta_1(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v''_1/x] \Rightarrow T_{22} \rangle^l (v_1 (\langle T_{21} \Rightarrow T_{11} \rangle^l v''_1)))) \\ \simeq \theta_2(\delta_2(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v''_2/x] \Rightarrow T_{22} \rangle^l (v_2 (\langle T_{21} \Rightarrow T_{11} \rangle^l v''_2)))) \\ : T_{22}; \theta; \delta[v_1, v_2/z][v''_1, v''_2/x].$$

Since $\Gamma \vdash T_1 \simeq T_1 : *$ and $\Gamma \vdash T_2 \simeq T_2 : *$, we have $\Gamma \vdash T_{11} \simeq T_{11} : *$ and $\Gamma \vdash T_{21} \simeq T_{21} : *$. Then, by the IH, we have

$$\theta_1(\delta_1(\langle T_{21} \Rightarrow T_{11} \rangle^l)) \sim \theta_2(\delta_2(\langle T_{21} \Rightarrow T_{11} \rangle^l)) : T_{21} \rightarrow T_{11}; \theta; \delta.$$

By Lemma 3.3.2 and assumption,

$$\theta_1(\delta_1(\langle T_{21} \Rightarrow T_{11} \rangle^l)) v''_1 \simeq \theta_2(\delta_2(\langle T_{21} \Rightarrow T_{11} \rangle^l)) v''_2 : T_{11}; \theta; \delta[v_1, v_2/z][v''_1, v''_2/z']$$

for fresh z' . These normalize to $r_1 \simeq r_2 : T_{11}; \theta; \delta[v_1, v_2/z][v_1'', v_2''/z']$. If $r_1 = r_2 = \uparrow l$ for some l , then we also have $v_1 v_1'' \longrightarrow^* \uparrow l$ and $v_2 v_2'' \longrightarrow^* \uparrow l$ and, by expansion,

$$v_1 v_1'' \simeq v_2 v_2'' : T_{22}; \theta; \delta[v_1, v_2/z][v_1'', v_2''/x].$$

Otherwise, let $v_1''' = r_1$ and $v_2''' = r_2$. By Lemma 3.3.2 and definition,

$$v_1 v_1''' \simeq v_2 v_2''' : T_{12}; \theta; \delta[v_1''', v_2'''/x].$$

These normalize to

$$r_1' \sim r_2' : T_{12}; \theta; \delta[v_1''', v_2'''/x].$$

If $r_1' = r_2' = \uparrow l'$ for some l' , then, again, we have $v_1 v_1'' \longrightarrow^* \uparrow l$ and $v_2 v_2'' \longrightarrow^* \uparrow l$ and, by expansion,

$$v_1 v_1'' \simeq v_2 v_2'' : T_{22}; \theta; \delta[v_1, v_2/z][v_1'', v_2''/x].$$

Otherwise, let $v_1'''' = r_1'$ and $v_2'''' = r_2'$. By Lemma 3.3.2 and 3.3.1 and the fact that

$$\begin{aligned} \langle T_{21} \Rightarrow T_{11} \rangle^l v_1'' &\longrightarrow^* v_1'''' \\ \langle T_{21} \Rightarrow T_{11} \rangle^l v_2'' &\longrightarrow^* v_2'''' \end{aligned}$$

we have

$$v_1'''' \sim v_2'''' : T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x]; \theta; \delta[v_1'', v_2''/x].$$

From $\Gamma \vdash T_1 \simeq T_1 : *$, we have $\Gamma, x:T_{11} \vdash T_{12} \simeq T_{12} : *$ by definition. Furthermore, we have $\Gamma, x:T_{21} \vdash T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \simeq T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] : *$ by Lemma 3.3.5.

Since $T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \parallel T_{22}$, by the IH,

$$\begin{aligned} \Gamma, x:T_{21} \vdash \langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l &\simeq \\ \langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \Rightarrow T_{22} \rangle^l : T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] &\rightarrow T_{22}. \end{aligned}$$

Since $\Gamma, x:T_{21} \vdash \theta; \delta[v_1'', v_2''/x]$ and $x \notin T_{21}, T_{11}$, we have

$$\begin{aligned} &\theta_1(\delta_1(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v_1''/x] \Rightarrow T_{22} \rangle^l)) \sim \\ &\theta_2(\delta_2(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v_2''/x] \Rightarrow T_{22} \rangle^l)) \\ &: T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l x/x] \rightarrow T_{22}; \theta; \delta[v_1'', v_2''/x]. \end{aligned}$$

By definition,

$$\begin{aligned} &(\theta_1(\delta_1(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v_1''/x] \Rightarrow T_{22} \rangle^l))) v_1'''' \\ &\simeq (\theta_2(\delta_2(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v_2''/x] \Rightarrow T_{22} \rangle^l))) v_2'''' \\ &: T_{22}; \theta; \delta[v_1'', v_2''/x][v_1''''', v_2'''''/z''] \end{aligned}$$

for fresh z'' . These normalize to $r_1'' \sim r_2'' : T_{22}; \theta; \delta[v_1'', v_2''/x][v_1''''', v_2'''''/z'']$. By expansion and Lemma 3.3.2, $v_1' v_1'' \simeq v_2' v_2'' : T_{22}; \theta; \delta[v_1, v_2/z][v_1'', v_2''/x]$.

$(T_1 = \forall\alpha.T'_1$ and $T_2 = \forall\alpha.T'_2$ and $T_1 \neq T_2$): Given $\Gamma \vdash \theta; \delta$, we wish to show that

$$\begin{aligned} \langle \forall\alpha.\theta_1(\delta_1(T'_1)) \Rightarrow \forall\alpha.\theta_1(\delta_1(T'_2)) \rangle^l &\sim \langle \forall\alpha.\theta_2(\delta_2(T'_1)) \Rightarrow \forall\alpha.\theta_2(\delta_2(T'_2)) \rangle^l \\ &: (\forall\alpha.T'_1) \rightarrow (\forall\alpha.T'_2); \theta; \delta. \end{aligned}$$

Let $v_1 \sim v_2 : \forall\alpha.T'_1; \theta; \delta$. We must show that

$$\begin{aligned} \langle \forall\alpha.\theta_1(\delta_1(T'_1)) \Rightarrow \forall\alpha.\theta_1(\delta_1(T'_2)) \rangle^l v_1 &\simeq \langle \forall\alpha.\theta_2(\delta_2(T'_1)) \Rightarrow \forall\alpha.\theta_2(\delta_2(T'_2)) \rangle^l v_2 \\ &: (\forall\alpha.T'_2); \theta; \delta[v_1, v_2/z] \end{aligned}$$

for fresh z . Since

$$\begin{aligned} \langle \forall\alpha.\theta_1(\delta_1(T'_1)) \Rightarrow \forall\alpha.\theta_1(\delta_1(T'_2)) \rangle^l v_1 &\longrightarrow \Lambda\alpha. \langle \theta_1(\delta_1(T'_1)) \Rightarrow \theta_1(\delta_1(T'_2)) \rangle^l (v_1 \alpha) \\ \langle \forall\alpha.\theta_2(\delta_2(T'_1)) \Rightarrow \forall\alpha.\theta_2(\delta_2(T'_2)) \rangle^l v_2 &\longrightarrow \Lambda\alpha. \langle \theta_2(\delta_2(T'_1)) \Rightarrow \theta_2(\delta_2(T'_2)) \rangle^l (v_2 \alpha), \end{aligned}$$

it suffices to show that

$$\begin{aligned} \Lambda\alpha. \langle \theta_1(\delta_1(T'_1)) \Rightarrow \theta_1(\delta_1(T'_2)) \rangle^l (v_1 \alpha) &\sim \\ \Lambda\alpha. \langle \theta_2(\delta_2(T'_1)) \Rightarrow \theta_2(\delta_2(T'_2)) \rangle^l (v_2 \alpha) & \\ &: (\forall\alpha.T'_2); \theta; \delta[v_1, v_2/z]. \end{aligned}$$

Let R, T''_1, T''_2 be given. We will show that

$$\begin{aligned} \Lambda\alpha. \langle \theta_1(\delta_1(T'_1)) \Rightarrow \theta_1(\delta_1(T'_2)) \rangle^l (v_1 \alpha) T''_1 &\sim \\ \Lambda\alpha. \langle \theta_2(\delta_2(T'_1)) \Rightarrow \theta_2(\delta_2(T'_2)) \rangle^l (v_2 \alpha) T''_2 & \\ &: T''_2; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta[v_1, v_2/z]. \end{aligned}$$

These terms normalize to

$$\begin{aligned} \theta_1(\delta_1(\langle T'_1 \Rightarrow T'_2 \rangle^l [T''_1/\alpha])) (v_1 T''_1) &\text{ and} \\ \theta_2(\delta_2(\langle T'_1 \Rightarrow T'_2 \rangle^l [T''_2/\alpha])) (v_2 T''_2), & \end{aligned}$$

and we will show

$$\begin{aligned} \theta_1(\delta_1(\langle T'_1 \Rightarrow T'_2 \rangle^l [T''_1/\alpha])) (v_1 T''_1) &\simeq \\ \theta_2(\delta_2(\langle T'_1 \Rightarrow T'_2 \rangle^l [T''_2/\alpha])) (v_2 T''_2) & \\ &: T''_2; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta[v_1, v_2/z]. \end{aligned}$$

By assumption,

$$v_1 T''_1 \simeq v_2 T''_2 : T''_1; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta.$$

These normalize to

$$r_1 \sim r_2 : T''_1; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta.$$

If $r_1 = r_2 = \uparrow l$ for some l , we have

$$\begin{aligned} \Lambda\alpha. \langle \theta_1(\delta_1(T'_1)) \Rightarrow \theta_1(\delta_1(T'_2)) \rangle^l (v_1 \alpha) T''_1 &\longrightarrow^* \uparrow l \\ \Lambda\alpha. \langle \theta_2(\delta_2(T'_1)) \Rightarrow \theta_2(\delta_2(T'_2)) \rangle^l (v_2 \alpha) T''_2 &\longrightarrow^* \uparrow l, \end{aligned}$$

finishing the case. Otherwise, we have $r_1 = v'_1$ and $r_2 = v'_2$ and

$$v'_1 \sim v'_2 : T'_1; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta.$$

It is easy to show that $\Gamma, \alpha \vdash T'_1 \simeq T''_1 : *$ and $\Gamma, \alpha \vdash T'_2 \simeq T''_2 : *$ from the assumptions $\Gamma, \alpha \vdash T_1 \simeq T'_1 : *$ and $\Gamma, \alpha \vdash T_2 \simeq T'_2 : *$. Then, by the IH, we have

$$\Gamma, \alpha \vdash \langle T'_1 \Rightarrow T'_2 \rangle^l \simeq \langle T'_1 \Rightarrow T'_2 \rangle^l : T'_1 \rightarrow T'_2.$$

Since $\Gamma, \alpha \vdash \theta[\alpha \mapsto R, T''_1, T''_2]; \delta$, we have

$$\begin{aligned} \theta_1(\delta_1(\langle T'_1[T''_1/\alpha] \Rightarrow T'_2[T''_1/\alpha] \rangle^l)) &\sim \theta_2(\delta_2(\langle T'_1[T''_1/\alpha] \Rightarrow T'_2[T''_1/\alpha] \rangle^l)) \\ &: T'_1 \rightarrow T'_2; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta. \end{aligned}$$

By definition,

$$\begin{aligned} \theta_1(\delta_1(\langle T'_1 \Rightarrow T'_2 \rangle^l[T''_1/\alpha])) v'_1 &\simeq \theta_2(\delta_2(\langle T'_1 \Rightarrow T'_2 \rangle^l[T''_1/\alpha])) v'_2 \\ &: T'_2; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta[v'_1, v'_2/z'] \end{aligned}$$

for fresh z' . By Lemma 3.3.2,

$$\theta_1(\delta_1(\langle T'_1 \Rightarrow T'_2 \rangle^l[T''_1/\alpha])) v'_1 \simeq \theta_2(\delta_2(\langle T'_1 \Rightarrow T'_2 \rangle^l[T''_2/\alpha])) v'_2 : T'_2; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta.$$

And now, by expansion,

$$\begin{aligned} \theta_1(\delta_1(\langle T'_1 \Rightarrow T'_2 \rangle^l[T''_1/\alpha])) (v_1 T''_1) &\longrightarrow^* \theta_1(\delta_1(\langle T'_1 \Rightarrow T'_2 \rangle^l[T''_1/\alpha])) v'_1 \\ \theta_2(\delta_2(\langle T'_1 \Rightarrow T'_2 \rangle^l[T''_2/\alpha])) (v_2 T''_2) &\longrightarrow^* \theta_2(\delta_2(\langle T'_1 \Rightarrow T'_2 \rangle^l[T''_2/\alpha])) v'_2 \end{aligned}$$

finish the case.

$(T_1 = \{x:T'_1 \mid e\}$ and $T_1 \neq T_2$ and $T_2 \neq \{y:T_1 \mid e'\})$: Given $\Gamma \vdash \theta; \delta$, we wish to show that

$$\theta_1(\delta_1(\langle \{x:T'_1 \mid e\} \Rightarrow T_2 \rangle^l)) \sim \theta_2(\delta_2(\langle \{x:T'_1 \mid e\} \Rightarrow T_2 \rangle^l)) : \{x:T'_1 \mid e\} \rightarrow T_2; \theta; \delta.$$

Let $v_1 \sim v_2 : \{x:T'_1 \mid e\}; \theta; \delta$, that is, $v_1 \sim v_2 : T'_1; \theta; \delta$ and

$$\begin{aligned} \theta_1(\delta_1(e))[v_1/x] &\longrightarrow^* \text{true} \\ \theta_2(\delta_2(e))[v_2/x] &\longrightarrow^* \text{true} \end{aligned}$$

We have to show that

$$\begin{aligned} \theta_1(\delta_1(\langle \{x:T'_1 \mid e\} \Rightarrow T_2 \rangle^l)) v_1 &\simeq \\ \theta_2(\delta_2(\langle \{x:T'_1 \mid e\} \Rightarrow T_2 \rangle^l)) v_2 &: T_2; \theta; \delta[v_1, v_2/z] \end{aligned}$$

for fresh z . Since

$$\begin{aligned} \theta_1(\delta_1(\langle \{x:T'_1 \mid e\} \Rightarrow T_2 \rangle^l)) v_1 &\longrightarrow \theta_1(\delta_1(\langle T'_1 \Rightarrow T_2 \rangle^l)) v_1 \\ \theta_2(\delta_2(\langle \{x:T'_1 \mid e\} \Rightarrow T_2 \rangle^l)) v_2 &\longrightarrow \theta_2(\delta_2(\langle T'_1 \Rightarrow T_2 \rangle^l)) v_2 \end{aligned}$$

by **E_FORGET**, it suffices to show that

$$\theta_1(\delta_1(\langle T'_1 \Rightarrow T_2 \rangle^l)) v_1 \simeq \theta_2(\delta_2(\langle T'_1 \Rightarrow T_2 \rangle^l)) v_2 : T_2; \theta; \delta[v_1, v_2/z].$$

Since $\Gamma \vdash T_1 \simeq T_1 : *$, we have $\Gamma \vdash T'_1 \simeq T'_1 : *$ by definition. We also have $T'_1 \parallel T_2$, by inversion. Then, by the IH

$$\theta_1(\delta_1(\langle T'_1 \Rightarrow T_2 \rangle^l)) \sim \theta_2(\delta_2(\langle T'_1 \Rightarrow T_2 \rangle^l)) : T'_1 \rightarrow T_2; \theta; \delta.$$

And finally, by assumption,

$$\theta_1(\delta_1(\langle T'_1 \Rightarrow T_2 \rangle^l)) v_1 \simeq \theta_2(\delta_2(\langle T'_1 \Rightarrow T_2 \rangle^l)) v_2 : T_2; \theta; \delta[v_1, v_2/z].$$

$(T_2 = \{x:T_1 \mid e\})$: Given $\Gamma \vdash \theta; \delta$, we wish to show that

$$\theta_1(\delta_1(\langle T_1 \Rightarrow \{x:T_1 \mid e\} \rangle^l)) \sim \theta_2(\delta_2(\langle T_1 \Rightarrow \{x:T_1 \mid e\} \rangle^l)) : T_1 \rightarrow \{x:T_1 \mid e\}; \theta; \delta.$$

Let $v_1 \sim v_2 : T_1; \theta; \delta$. We have to show that

$$\begin{aligned} \theta_1(\delta_1(\langle T_1 \Rightarrow \{x:T_1 \mid e\} \rangle^l)) v_1 &\simeq \\ \theta_2(\delta_2(\langle T_1 \Rightarrow \{x:T_1 \mid e\} \rangle^l)) v_2 &: \{x:T_1 \mid e\}; \theta; \delta[v_1, v_2/z] \end{aligned}$$

for fresh z . Since

$$\begin{aligned} \theta_1(\delta_1(\langle T_1 \Rightarrow \{x:T_1 \mid e\} \rangle^l)) v_1 &\longrightarrow \langle \{x:T_1 \mid e\}, e[v_1/x], v_1 \rangle^l \\ \theta_2(\delta_2(\langle T_1 \Rightarrow \{x:T_1 \mid e\} \rangle^l)) v_2 &\longrightarrow \langle \{x:T_1 \mid e\}, e[v_2/x], v_2 \rangle^l \end{aligned}$$

by **E_CHECK**, it suffices to show that

$$\langle \{x:T_1 \mid e\}, e[v_1/x], v_1 \rangle^l \simeq \langle \{x:T_1 \mid e\}, e[v_2/x], v_2 \rangle^l : \{x:T_1 \mid e\}; \theta; \delta[v_1, v_2/z].$$

By the assumption $\Gamma \vdash T_2 \simeq T_2 : *$, we have $T_2 \simeq T_2 : *; \theta; \delta$. By definition, $T_1 \simeq T_1 : *; \theta; \delta$ and

$$\theta_1(\delta_1(e))[v'_1/x] \simeq \theta_2(\delta_2(e))[v'_2/x] : \mathbf{Bool}; \theta; \delta$$

for any $v'_1 \sim v'_2 : T_1; \theta; \delta$. So, in particular,

$$\theta_1(\delta_1(e))[v_1/x] \simeq \theta_2(\delta_2(e))[v_2/x] : \mathbf{Bool}; \theta; \delta.$$

These reduce to $r_1 \sim r_2 : \mathbf{Bool}; \theta; \delta$.

We have three cases:

$(r_1 = r_2 = \uparrow l'$ for some l'): Then the case is immediate by:

$$\begin{aligned} \langle \{x:T_1 \mid e\}, e[v_1/x], v_1 \rangle^l &\longrightarrow^* \uparrow l' \\ \langle \{x:T_1 \mid e\}, e[v_2/x], v_2 \rangle^l &\longrightarrow^* \uparrow l'. \end{aligned}$$

$(r_1 = r_2 = \text{true})$: Then,

$$\begin{aligned} \langle \{x:T_1 \mid e\}, e[v_1/x], v_1 \rangle^l &\longrightarrow^* v_1 \\ \langle \{x:T_1 \mid e\}, e[v_2/x], v_2 \rangle^l &\longrightarrow^* v_2. \end{aligned}$$

By assumption and Lemma 3.3.2, $v_1 \sim v_2 : T_1; \theta; \delta[v_1, v_2/z]$, finishing the case.

$(r_1 = r_2 = \text{false})$: Then we conclude by,

$$\begin{aligned} \langle \{x:T_1 \mid e\}, e[v_1/x], v_1 \rangle^l &\longrightarrow^* \uparrow l \\ \langle \{x:T_1 \mid e\}, e[v_2/x], v_2 \rangle^l &\longrightarrow^* \uparrow l. \end{aligned}$$

$(T_1$ is not a refinement type and $T_2 = \{x:T'_2 \mid e\}$ and $T_1 \neq T'_2)$: Given $\Gamma \vdash \theta; \delta$, we wish to show that

$$\theta_1(\delta_1(\langle T_1 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) \sim \theta_2(\delta_2(\langle T_1 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) : T_1 \rightarrow \{x:T'_2 \mid e\}; \theta; \delta.$$

Let $v_1 \sim v_2 : T_1; \theta; \delta$. We have to show that

$$\begin{aligned} \theta_1(\delta_1(\langle T_1 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) v_1 &\simeq \\ \theta_2(\delta_2(\langle T_1 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) v_2 &: \{x:T'_2 \mid e\}; \theta; \delta[v_1, v_2/z] \end{aligned}$$

for fresh z . Since

$$\begin{aligned} \theta_1(\delta_1(\langle T_1 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) v_1 &\longrightarrow \\ \theta_1(\delta_1(\langle T'_2 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) (\theta_1(\delta_1(\langle T_1 \Rightarrow T'_2 \rangle^l)) v_1) & \\ \theta_2(\delta_2(\langle T_1 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) v_2 &\longrightarrow \\ \theta_2(\delta_2(\langle T'_2 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) (\theta_2(\delta_2(\langle T_1 \Rightarrow T'_2 \rangle^l)) v_2) & \end{aligned}$$

by E_PRECHECK, it suffices to show that

$$\begin{aligned} \theta_1(\delta_1(\langle T'_2 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) (\theta_1(\delta_1(\langle T_1 \Rightarrow T'_2 \rangle^l)) v_1) & \\ \simeq \theta_2(\delta_2(\langle T'_2 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) (\theta_2(\delta_2(\langle T_1 \Rightarrow T'_2 \rangle^l)) v_2) & \\ : \{x:T'_2 \mid e\}; \theta; \delta[v_1, v_2/z]. & \end{aligned}$$

Since $\Gamma \vdash T_2 \simeq T'_2 : *$, we have $\Gamma \vdash T'_2 \simeq T'_2 : *$ by definition. Noting that

$$\begin{aligned} cc(\langle T_1 \Rightarrow \{x:T'_2 \mid e\} \rangle^l) &= cc(\langle T_1 \Rightarrow T'_2 \rangle^l) + 2 > cc(\langle T_1 \Rightarrow T'_2 \rangle^l) \\ cc(\langle T_1 \Rightarrow \{x:T'_2 \mid e\} \rangle^l) &= cc(\langle T_1 \Rightarrow T'_2 \rangle^l) + 2 > 1 = cc(\langle T'_2 \Rightarrow \{x:T'_2 \mid e\} \rangle^l). \end{aligned}$$

we can apply the IH to see

$$\begin{aligned} \Gamma \vdash \langle T_1 \Rightarrow T'_2 \rangle^l &\simeq \langle T_1 \Rightarrow T'_2 \rangle^l : T_1 \rightarrow T'_2, \text{ and} \\ \Gamma \vdash \langle T'_2 \Rightarrow \{x:T'_2 \mid e\} \rangle^l &\simeq \langle T'_2 \Rightarrow \{x:T'_2 \mid e\} \rangle^l : T'_2 \rightarrow \{x:T'_2 \mid e\}. \end{aligned}$$

Then we can easily see

$$\begin{aligned} \theta_1(\delta_1(\langle T'_2 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) (\theta_1(\delta_1(\langle T_1 \Rightarrow T'_2 \rangle^l)) v_1) & \\ \simeq \theta_2(\delta_2(\langle T'_2 \Rightarrow \{x:T'_2 \mid e\} \rangle^l)) (\theta_2(\delta_2(\langle T_1 \Rightarrow T'_2 \rangle^l)) v_2) & \\ : \{x:T'_2 \mid e\}; \theta; \delta[v_1, v_2/z]. & \end{aligned}$$

□

Finally, we can prove relational parametricity—every well-typed term (under Γ) is related to itself for any Γ -respecting substitutions.

3.3.7 Theorem [Parametricity]: 1. If $\Gamma \vdash e : T$ then $\Gamma \vdash e \simeq e : T$, and

2. If $\Gamma \vdash T$ then $\Gamma \vdash T \simeq T : *$.

Proof: By simultaneous induction on the derivations with case analysis on the last rule used.

(T_VAR): Let $\Gamma \vdash \theta; \delta$. We wish to show that $\theta_1(\delta_1(x)) \simeq \theta_2(\delta_2(x)) : T; \theta; \delta$, which follows from the assumption.

(T_CONST): By the assumption that constants are assigned correct types.

(T_OP): By the assumption that operators are assigned correct types (and the IHs for the operator's arguments).

(T_ABS): We have $e = \lambda x:T_1. e_{12}$ and $T = x:T_1 \rightarrow T_2$ and $\Gamma, x:T_1 \vdash e_{12} : T_2$. Let $\Gamma \vdash \theta; \delta$. We wish to show that

$$\theta_1(\delta_1(\lambda x:T_1. e_{12})) \sim \theta_2(\delta_2(\lambda x:T_1. e_{12})) : (x:T_1 \rightarrow T_2); \theta; \delta.$$

Let $v_1 \sim v_2 : T_1; \theta; \delta$. We must show that

$$(\lambda x:\theta_1(\delta_1(T_1)). \theta_1(\delta_1(e_{12}))) v_1 \simeq (\lambda x:\theta_2(\delta_2(T_1)). \theta_2(\delta_2(e_{12}))) v_2 : T_2; \theta; \delta[v_1, v_2/x].$$

Since

$$\begin{aligned} (\lambda x:\theta_1(\delta_1(T_1)). \theta_1(\delta_1(e_{12}))) v_1 &\longrightarrow \theta_1(\delta_1(e_{12}))[v_1/x] \\ (\lambda x:\theta_2(\delta_2(T_1)). \theta_2(\delta_2(e_{12}))) v_2 &\longrightarrow \theta_2(\delta_2(e_{12}))[v_2/x], \end{aligned}$$

it suffices to show

$$\theta_1(\delta_1(e_{12}))[v_1/x] \simeq \theta_2(\delta_2(e_{12}))[v_2/x] : T_2; \theta; \delta[v_1, v_2/x].$$

By the IH, $\Gamma, x:T_1 \vdash e_{12} \simeq e_{12} : T_2$. The fact that $\Gamma, x:T_1 \vdash \theta; \delta[v_1, v_2/x]$ finishes the case.

(T_APP): We have $e = e_1 e_2$ and $\Gamma \vdash e_1 : x:T_1 \rightarrow T_2$ and $\Gamma \vdash e_2 : T_1$ and $T = T_2[e_2/x]$. Let $\Gamma \vdash \theta; \delta$. We wish to show that

$$\theta_1(\delta_1(e_1 e_2)) \simeq \theta_2(\delta_2(e_1 e_2)) : T_2[e_2/x]; \theta; \delta.$$

By the IH,

$$\begin{aligned} \theta_1(\delta_1(e_1)) &\simeq \theta_2(\delta_2(e_2)) : x:T_1 \rightarrow T_2; \theta; \delta, \text{ and} \\ \theta_1(\delta_1(e_2)) &\simeq \theta_2(\delta_2(e_2)) : T_1; \theta; \delta. \end{aligned}$$

These normalize to $r_{11} \sim r_{12} : x:T_1 \rightarrow T_2; \theta; \delta$ and $r_{21} \simeq r_{22} : T_1; \theta; \delta$, respectively. If $r_{11} = r_{12} = \uparrow l$ or $r_{21} = r_{22} = \uparrow l$ for some l , then we are done:

$$\begin{aligned}\theta_1(\delta_1(e_1 e_2)) &\longrightarrow^* \uparrow l \\ \theta_2(\delta_2(e_1 e_2)) &\longrightarrow^* \uparrow l.\end{aligned}$$

So let $r_{ij} = v_{ij}$. By definition,

$$v_{11} v_{21} \simeq v_{12} v_{22} : T_2; \theta; \delta[v_{21}, v_{22}/x].$$

These normalize to $r'_1 \sim r'_2 : T_2; \theta; \delta[v_{21}, v_{22}/x]$. By Lemma 3.3.1,

$$r'_1 \sim r'_2 : T_2[e_2/x]; \theta; \delta.$$

By expansion, we can then see

$$\theta_1(\delta_1(e_1 e_2)) \simeq \theta_2(\delta_2(e_1 e_2)) : T_2[e_2/x]; \theta; \delta.$$

(T_TABS): We have $e = \Lambda\alpha. e_0$ and $T = \forall\alpha. T_0$ and $\Gamma, \alpha \vdash e_0 : T_0$. Let $\Gamma \vdash \theta; \delta$. We wish to show that

$$\theta_1(\delta_1(\Lambda\alpha. e_0)) \sim \theta_2(\delta_2(\Lambda\alpha. e_0)) : \forall\alpha. T_0; \theta; \delta.$$

Let R, T_1, T_2 be given. We must show that

$$\theta_1(\delta_1(\Lambda\alpha. e_0)) T_1 \simeq \theta_2(\delta_2(\Lambda\alpha. e_0)) T_2 : T_0; \theta[\alpha \mapsto R, T_1, T_2]; \delta.$$

Since

$$\begin{aligned}\theta_1(\delta_1(\Lambda\alpha. e_0)) T_1 &\longrightarrow \theta_1(\delta_1(e_0))[T_1/\alpha] \\ \theta_2(\delta_2(\Lambda\alpha. e_0)) T_2 &\longrightarrow \theta_2(\delta_2(e_0))[T_2/\alpha]\end{aligned}$$

it suffices to show that

$$\theta_1(\delta_1(e_0))[T_1/\alpha] \simeq \theta_2(\delta_2(e_0))[T_2/\alpha] : T_0; \theta[\alpha \mapsto R, T_1, T_2]; \delta.$$

Since $\Gamma, \alpha \vdash \theta[\alpha \mapsto R, T_1, T_2]; \delta$, the IH finishes the case with $\Gamma, \alpha \vdash e_0 \simeq e_0 : T_0$.

(T_TAPP): We have $e = e_1 T_2$ and $\Gamma \vdash e_1 : \forall\alpha. T_0$ and $\Gamma \vdash T_2$ and $T = T_0[T_2/\alpha]$. Let $\Gamma \vdash \theta; \delta$. We wish to show that

$$\theta_1(\delta_1(e_1 T_2)) \simeq \theta_2(\delta_2(e_1 T_2)) : T_0[T_2/\alpha]; \theta; \delta.$$

By the IH,

$$\theta_1(\delta_1(e_1)) \simeq \theta_2(\delta_2(e_1)) : \forall\alpha. T_0; \theta; \delta.$$

These normalize to $r_1 \sim r_2 : \forall \alpha. T_0; \theta; \delta$. If both results are blame, $\theta_1(\delta_1(e_1 T_2))$ and $\theta_2(\delta_2(e_1 T_2))$ also normalize to blame, and we are done. So let $r_1 = v_1$ and $r_2 = v_2$. Then, by definition,

$$v_1 T'_1 \simeq v_2 T'_2 : T_0; \theta[\alpha \mapsto R, T'_1, T'_2]; \delta$$

for any R, T'_1, T'_2 . In particular,

$$v_1 \theta_1(\delta_1(T_2)) \simeq v_2 \theta_2(\delta_2(T_2)) : T_0; \theta[\alpha \mapsto R_{T_2, \theta, \delta}, \theta_1(\delta_1(T_2)), \theta_2(\delta_2(T_2))]; \delta.$$

These normalize to

$$r'_1 \sim r'_2 : T_0; \theta[\alpha \mapsto R_{T_2, \theta, \delta}, \theta_1(\delta_1(T_2)), \theta_2(\delta_2(T_2))]; \delta.$$

By Lemma 3.3.3, $r'_1 \sim r'_2 : T_0[T_2/\alpha]; \theta; \delta$. By expansion,

$$\theta_1(\delta_1(e_1 T_2)) \simeq \theta_2(\delta_2(e_1 T_2)) : T_0[T_2/\alpha]; \theta; \delta.$$

(T_CAST): We have $e = \langle T_1 \Rightarrow T_2 \rangle^l$ and $\vdash \Gamma$ and $T_1 \parallel T_2$ and $\Gamma \vdash T_1, \Gamma \vdash T_2$ and $\overline{T = T_1 \rightarrow T_2}$. By the IH, $\Gamma \vdash T_1 \simeq T_1 : *$ and $\Gamma \vdash T_2 \simeq T_2 : *$. By Lemma 3.3.6,

$$\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle^l \simeq \langle T_1 \Rightarrow T_2 \rangle^l : T_1 \rightarrow T_2.$$

(T_BLAKE): Immediate.

(T_CHECK): We have $e = \langle \{x:T_1 \mid e_1\}, e_2, v \rangle^l$ and $\emptyset \vdash v : T_1$ and $\emptyset \vdash e_2 : \mathbf{Bool}$, $\vdash \Gamma$ and $\emptyset \vdash \{x:T_1 \mid e_1\}$ and $e_1[v/x] \longrightarrow^* e_2$ and $T = \{x:T_1 \mid e_1\}$. Let $\Gamma \vdash \theta; \delta$. We wish to show that

$$\theta_1(\delta_1(\langle \{x:T_1 \mid e_1\}, e_2, v \rangle^l)) \simeq \theta_2(\delta_2(\langle \{x:T_1 \mid e_1\}, e_2, v \rangle^l)) : \{x:T_1 \mid e_1\}; \theta; \delta.$$

By the IH,

$$\theta_1(\delta_1(e_2)) \simeq \theta_2(\delta_2(e_2)) : \mathbf{Bool}; \theta; \delta$$

and these normalize to the same result. If the result is **false** or $\uparrow l'$ for some l' , then, for some l'' ,

$$\begin{aligned} \theta_1(\delta_1(\langle \{x:T_1 \mid e_1\}, e_2, v \rangle^l)) &\longrightarrow^* \uparrow l'' \\ \theta_2(\delta_2(\langle \{x:T_1 \mid e_1\}, e_2, v \rangle^l)) &\longrightarrow^* \uparrow l''. \end{aligned}$$

Otherwise, the result is **true**. Then, by the IH, $v \sim v : T_1; \theta; \delta$ and $\emptyset \vdash \{x:T_1 \mid e_1\} \simeq \{x:T_1 \mid e_1\} : *$. By definition,

$$\theta_1(\delta_1(e_1))[v/x] \simeq \theta_2(\delta_2(e_1))[v/x] : \mathbf{Bool}; \theta; \delta[v, v/x].$$

Then, we have

$$\begin{aligned} \theta_1(\delta_1(e_1))[v/x] &= e_1[v/x] \longrightarrow^* \mathbf{true} \\ \theta_2(\delta_2(e_1))[v/x] &= e_1[v/x] \longrightarrow^* \mathbf{true}. \end{aligned}$$

By definition, $v \simeq v : \{x:T_1 \mid e_1\}; \theta; \delta$. By expansion,

$$\theta_1(\delta_1(\langle \{x:T_1 \mid e_1\}, e_2, v \rangle^l)) \simeq \theta_2(\delta_2(\langle \{x:T_1 \mid e_1\}, e_2, v \rangle^l)) : \{x:T_1 \mid e_1\}; \theta; \delta.$$

(T_CONV): By Lemma 3.3.4.

(T_EXACT): We have $e = v$ and $\emptyset \vdash v : T$ and $\emptyset \vdash \{x:T_0 \mid e_0\}$ and $e_0[v/x] \longrightarrow^* \mathbf{true}$ and $T = \{x:T_0 \mid e_0\}$. Let $\Gamma \vdash \theta; \delta$. We wish to show that

$$v \sim v : \{x:T_0 \mid e_0\}; \theta; \delta.$$

By the IH, $v \sim v : T_0; \theta; \delta$. Since $\emptyset \vdash \{x:T_0 \mid e_0\}$, the only free variable in e_0 is x and

$$\begin{aligned} \theta_1(\delta_1(e_0))[v/x] &= e_0[v/x] \longrightarrow^* \mathbf{true} \\ \theta_2(\delta_2(e_0))[v/x] &= e_0[v/x] \longrightarrow^* \mathbf{true}. \end{aligned}$$

By definition, $v \sim v : \{x:T_0 \mid e_0\}; \theta; \delta$.

(T_FORGET): By the IH, $\emptyset \vdash v \simeq v : \{x:T \mid e\}$, which implies $\Gamma \vdash v \simeq v : T$.

(WF_BASE): Trivial.

(WF_TVAR): Trivial.

(WF_FUN): By the IH.

(WF_FORALL): By the IH.

(WF_REFINE): By the IH. □

I refer readers to [61] for a significantly expanded account of parametricity for F_H *with recursion*. There, they have proved that their logical relation based on $\top\top$ -closure [53] is sound⁴ with respect to contextual equivalence. Since the details of their technical developments are different from what I present here, I can only conjecture that my logical relation is also sound with respect to contextual equivalence.

We do have that logically related programs are by definition *behaviorally equivalent*: if $\emptyset \vdash e_1 \simeq e_2 : T$, then e_1 and e_2 coterminate at related results. When the results are constants or blame, the results are not only logically related, but equal.

3.4 Subtyping and Upcast Elimination

Knowles and Flanagan [44] 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. Upcast elimination is, at heart, an optimization: since the cast can never fail, there is no point in running it. I 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 behaviorally equivalent—to the identity.

⁴And also complete, under certain conditions.

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

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

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

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

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

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

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

Figure 3.7: Subtyping, implication, and closing substitutions

I define subtyping in Figure 3.7. My 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} [6]. 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! This extra cast is necessary in the `S_FUN` case when proving the upcast lemma, but the well typedness of the relation is necessary so we can apply parametricity (Lemma 3.3.7) in the `S_REFINE` case.

The 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 3.7 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 type $\text{unref}(T)$ 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. I 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 situation in the manifest calculi of Knowles and Flanagan [44] and Greenberg, Pierce, and Weirich [34]. They both define a separate denotational semantics for use in their refinement subtyping rule—a unary logical relation unrelated to the logical relations used in their respective proofs. But they *need* to do so, in order to avoid a circularity. F_H has no such issues, and I make the decision because it is expedient.

I formulate the implication judgment in terms of cotermination at values rather than cotermination at `true` (as in [34, 44]) 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.

3.4.1 Lemma [LR substitutivity]: If $\Gamma, x:T_1, \Gamma' \vdash T_2$ and $\Gamma \vdash e_1 \simeq e_2 : T_1$, then

$\Gamma, \Gamma'[e_1/x] \vdash T_2[e_1/x] \simeq T_2[e_2/x] : *$.

Proof: By induction on $\Gamma, x:T_1, \Gamma' \vdash T_2$. \square

3.4.2 Lemma [LR type exchange]: If $T_1 \simeq T_2 : *; \theta; \delta$ and $T_2 \simeq T_1 : *; \theta; \delta$, then $v_1 \sim v_2 : T_1; \theta; \delta$ if and only if $v_1 \sim v_2 : T_2; \theta; \delta$.

Proof: By induction on the size of T_1 and T_2 .

$(T_1 = T_2 = B)$: Trivial.

$(T_1 = T_2 = \alpha)$: Trivial.

$(T_1 = x:T_{11} \rightarrow T_{12}$ and $T_2 = x:T_{21} \rightarrow T_{22})$: By definition, $T_{11} \simeq T_{21} : *; \theta; \delta$ and $T_{21} \simeq T_{11} : *; \theta; \delta$ and

$$\begin{aligned} & \forall v'_1 \sim v'_2 : T_{11}; \theta; \delta. T_{12} \simeq T_{22} : *; \theta; \delta[v'_1, v'_2/x], \text{ and} \\ & \forall v'_1 \sim v'_2 : T_{21}; \theta; \delta. T_{22} \simeq T_{12} : *; \theta; \delta[v'_1, v'_2/x]. \end{aligned}$$

Then:

$$\begin{aligned} & v_1 \sim v_2 : T_1; \theta; \delta \\ \iff & \forall (v'_1 \sim v'_2 : T_{11}; \theta; \delta), v_1 v'_1 \simeq v_2 v'_2 : T_{12}; \theta; \delta[v'_1, v'_2/x] \\ \text{(by the IH)} \iff & \forall (v'_1 \sim v'_2 : T_{21}; \theta; \delta), v_1 v'_1 \simeq v_2 v'_2 : T_{22}; \theta; \delta[v'_1, v'_2/x] \\ \iff & v_1 \sim v_2 : T_2; \theta; \delta. \end{aligned}$$

$(T_1 = \forall \alpha. T'_1$ and $T_2 = \forall \alpha. T'_2)$: By definition,

$$\begin{aligned} \forall R T'_1 T''_1, \quad & T'_1 \simeq T'_2 : *; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta \\ \forall R T'_1 T''_2, \quad & T'_2 \simeq T'_1 : *; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta. \end{aligned}$$

Then:

$$\begin{aligned} & v_1 \sim v_2 : T_1; \theta; \delta \\ \iff & \forall R T''_1 T''_2, v_1 T''_1 \simeq v_2 T''_2 : T'_1; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta \\ \text{(by the IH)} \iff & \forall R T''_1 T''_2, v_1 T''_1 \simeq v_2 T''_2 : T'_2; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta \\ \iff & v_1 \sim v_2 : T_2; \theta; \delta. \end{aligned}$$

$(T_1 = \{x:T'_1 \mid e_1\}$ and $T_2 = \{x:T'_2 \mid e_2\})$: By definition, $T'_1 \simeq T'_2 : *; \theta; \delta$ and $T'_2 \simeq T'_1 : *; \theta; \delta$. And also:

$$\forall (v_1 \sim v_2 : T'_1; \theta; \delta), \theta_1(\delta_1(e_1))[v_1/x] \simeq \theta_2(\delta_2(e_2))[v_2/x] : \mathbf{Bool}; \theta; \delta \quad (3.1)$$

$$\forall (v_1 \sim v_2 : T'_2; \theta; \delta), \theta_1(\delta_1(e_2))[v_1/x] \simeq \theta_2(\delta_2(e_1))[v_2/x] : \mathbf{Bool}; \theta; \delta. \quad (3.2)$$

Then:

$$\begin{aligned} & v_1 \sim v_2 : \{x:T'_1 \mid e_1\}; \theta; \delta \\ \iff & \begin{cases} v_1 \sim v_2 : T'_1; \theta; \delta \\ \theta_1(\delta_1(e_1))[v_1/x] \longrightarrow^* \mathbf{true} \\ \theta_2(\delta_2(e_1))[v_2/x] \longrightarrow^* \mathbf{true} \end{cases} \\ \left(\begin{array}{l} \text{by the IH} \\ \text{by (3.1)} \\ \text{by (3.2)} \end{array} \right) \iff & \begin{cases} v_1 \sim v_2 : T'_2; \theta; \delta \\ \theta_2(\delta_2(e_2))[v_2/x] \longrightarrow^* \mathbf{true} \\ \theta_1(\delta_1(e_2))[v_1/x] \longrightarrow^* \mathbf{true} \end{cases} \\ \iff & v_1 \sim v_2 : \{x:T'_2 \mid e_2\}; \theta; \delta. \end{aligned}$$

□

3.4.3 Lemma [Upcast lemma]: If $\Gamma \vdash T_1 <: T_2$ and $\Gamma \vdash T_1$ and $\Gamma \vdash T_2$, then $\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle^l \simeq \lambda x:T_1. x : T_1 \rightarrow T_2$.

Proof: By induction on the subtyping derivation.

(S_BASE): Easy. We have $T_1 = B$ and $T_2 = B$. Let $\Gamma \vdash \theta; \delta$. We wish to show that

$$\theta_1(\delta_1(\langle B \Rightarrow B \rangle^l)) \simeq \theta_2(\delta_2(\lambda x:T_1. x)) : B \rightarrow B; \theta; \delta.$$

Let $v_1 \sim v_2 : B; \theta; \delta$. We must show that

$$\theta_1(\delta_1(\langle B \Rightarrow B \rangle^l)) v_1 \simeq \theta_2(\delta_2(\lambda x:T_1. x)) v_2 : B; \theta; \delta[v_1, v_2/z]$$

for fresh z , but these normalize to $v_1 \sim v_2 : B; \theta; \delta[v_1, v_2/z]$. Lemma 3.3.2 finishes the case.

(S_TVAR): Similar to the case for S_BASE.

(S_FUN): We have $T_1 = y:T_{11} \rightarrow T_{12}$ and $T_2 = y:T_{21} \rightarrow T_{22}$ and $\Gamma \vdash T_{21} <: T_{11}$ and $\Gamma, y:T_{21} \vdash T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y] <: T_{22}$.

Let $\Gamma \vdash \theta; \delta$. We wish to show that

$$\theta_1(\delta_1(\langle y:T_{11} \rightarrow T_{12} \Rightarrow y:T_{21} \rightarrow T_{22} \rangle^l)) \simeq \theta_2(\delta_2(\lambda x:T_1. x)) : (y:T_{11} \rightarrow T_{12}) \rightarrow (y:T_{21} \rightarrow T_{22}); \theta; \delta.$$

Let $v_1 \sim v_2 : y:T_{11} \rightarrow T_{12}; \theta; \delta$. We must show that

$$\theta_1(\delta_1(\langle y:T_{11} \rightarrow T_{12} \Rightarrow y:T_{21} \rightarrow T_{22} \rangle^l)) v_1 \simeq \theta_2(\delta_2(\lambda x:T_1. x)) v_2 : y:T_{21} \rightarrow T_{22}; \theta; \delta[v_1, v_2/z]$$

for fresh z .

If $T_1 = T_2$, then we are done, since

$$\begin{aligned} \theta_1(\delta_1(\langle y:T_{11} \rightarrow T_{12} \Rightarrow y:T_{21} \rightarrow T_{22} \rangle^l)) v_1 &\longrightarrow v_1 \\ \theta_2(\delta_2(\lambda x:T_1. x)) v_2 &\longrightarrow v_2. \end{aligned}$$

Otherwise,

$$\begin{aligned} \theta_1(\delta_1(\langle y:T_{11} \rightarrow T_{12} \Rightarrow y:T_{21} \rightarrow T_{22} \rangle^l)) v_1 &\longrightarrow \\ \theta_1(\delta_1(\lambda y:T_{21}. \langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y] \Rightarrow T_{21} \rangle^l (v_1 (\langle T_{21} \Rightarrow T_{11} \rangle^l y)))) & \end{aligned}$$

$$\theta_2(\delta_2(\lambda x:T_1. x)) v_2 \longrightarrow v_2.$$

So, it suffices to show that

$$\begin{aligned} \theta_1(\delta_1(\lambda y:T_{21}. \langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y] \Rightarrow T_{21} \rangle^l (v_1 (\langle T_{21} \Rightarrow T_{11} \rangle^l y)))) &\sim v_2 \\ : y:T_{21} \rightarrow T_{22}; \theta; \delta[v_1, v_2/z]. & \end{aligned}$$

Let $v'_1 \sim v'_2 : T_{21}; \theta; \delta[v_1, v_2/z]$. We will show that

$$\begin{aligned} \theta_1(\delta_1(\lambda y:T_{21}. \langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y] \Rightarrow T_{21} \rangle^l (v_1(\langle T_{21} \Rightarrow T_{11} \rangle^l y)))) v'_1 &\simeq v_2 v'_2 \\ &: T_{22}; \theta; \delta[v_1, v_2/z][v'_1, v'_2/y]. \end{aligned}$$

By the IH, $\Gamma \vdash \langle T_{21} \Rightarrow T_{11} \rangle^l \simeq \lambda x:T_{21}. x : T_{21} \rightarrow T_{11}$. So,

$$\theta_1(\delta_1(\langle T_{21} \Rightarrow T_{11} \rangle^l)) v'_1 \simeq (\lambda x:\theta_1(\delta_1(T_{21})). x) v'_2 : T_{11}; \theta; \delta[v'_1, v'_2/z']$$

for fresh z'' . Then, for some v''_1 , $\theta_1(\delta_1(\langle T_{21} \Rightarrow T_{11} \rangle^l)) v'_1 \rightarrow v''_1$ and $v''_1 \simeq v'_2 : T_{11}; \theta; \delta[v'_1, v'_2/z']$. By assumption,

$$v_1 v''_1 \simeq v_2 v'_2 : T_{12}; \theta; \delta[v''_1, v'_2/y],$$

which normalize to $v''_1 \simeq v''_2 : T_{12}; \theta; \delta[v''_1, v'_2/y]$. Now, we can show that

$$\begin{aligned} \Gamma, y:T_{21} \vdash \langle T_{21} \Rightarrow T_{11} \rangle^l y &\simeq y : T_{11}, \text{ and} \\ \Gamma \vdash \langle T_{21} \Rightarrow T_{11} \rangle^l &\simeq \lambda x:T_{21}. x : T_{21} \rightarrow T_{11}. \end{aligned}$$

Noting that $T_{12} = T_{12}[y/y]$, Lemma 3.4.1 obtains

$$\Gamma, y:T_{21} \vdash T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y] \simeq T_{12} : *.$$

Then, by Lemma 3.4.2 (note that it is easy to show that $\Gamma \vdash T_1 \simeq T_2 : *$ if and only if $\Gamma \vdash T_2 \simeq T_1 : *$), we have

$$v''_1 \simeq v''_2 : T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y]; \theta; \delta[v''_1, v'_2/y].$$

On the other hand, by the other IH,

$$\begin{aligned} \Gamma, y:T_{21} \vdash \langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y] \Rightarrow T_{22} \rangle^l &\simeq \lambda x:(T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y]). x \\ &: T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y] \rightarrow T_{22}. \end{aligned}$$

Since $\Gamma, y:T_{21} \vdash \theta; \delta[v'_1, v'_2/y]$,

$$\begin{aligned} \theta_1(\delta_1(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v'_1/y] \Rightarrow T_{22}[v'_1/y] \rangle^l)) &\sim \\ \theta_2(\delta_2(\lambda x:T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v'_2/y]. x)) & \\ : T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y] \rightarrow T_{22}; \theta; \delta[v'_1, v'_2/y]. & \end{aligned}$$

So,

$$\begin{aligned} \theta_1(\delta_1(\langle T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v'_1/y] \Rightarrow T_{22}[v'_1/y] \rangle^l)) v''_1 &\simeq \\ \theta_2(\delta_2((\lambda x:T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l v'_2/y]. x))) v''_2 & \\ : T_{22}; \theta; \delta[v'_1, v'_2/y][v''_1, v''_2/z'']. & \end{aligned}$$

for fresh z''' . They normalize to

$$v''_1 \sim v''_2 : T_{22}; \theta; \delta[v'_1, v'_2/y][v''_1, v''_2/z''].$$

Now, letting $T'_{12} = T_{12}[\langle T_{21} \Rightarrow T_{11} \rangle^l y/y]$, we have:

$$\begin{aligned}
& \theta_1(\delta_1(\lambda y: T_{21}. (\langle T'_{12} \Rightarrow T_{22} \rangle^l (v_1 (\langle T_{21} \Rightarrow T_{11} \rangle^l y)))) v'_1 \\
\longrightarrow & \theta_1(\delta_1(\langle T'_{12}[v'_1/y] \Rightarrow T_{22}[v'_1/y] \rangle^l (v_1 (\langle T_{21} \Rightarrow T_{11} \rangle^l v'_1)))) \\
\longrightarrow^* & \theta_1(\delta_1(\langle T'_{12}[v'_1/y] \Rightarrow T_{22}[v'_1/y] \rangle^l)) (v_1 v''_1) \\
\longrightarrow^* & \theta_1(\delta_1(\langle T'_{12}[v'_1/y] \Rightarrow T_{22}[v'_1/y] \rangle^l)) v'''_1 \\
\longrightarrow^* & v''''_1
\end{aligned}$$

and $v_2 v'_2 \longrightarrow^* v''''_2$ and

$$v''''_1 \sim v''''_2 : T_{22}; \theta; \delta[v'_1, v'_2/y][v''''_1, v''''_2/z''''],$$

which concludes this case by expansion.

(S_FORALL): We have $T_1 = \forall \alpha. T'_1$ and $T_2 = \forall \alpha. T'_2$ and $\Gamma, \alpha \vdash T'_1 <: T'_2$.

Let $\Gamma \vdash \theta; \delta$. We wish to show that

$$\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) \simeq \theta_2(\delta_2(\lambda x: T_1. x)) : T_1 \rightarrow T_2; \theta; \delta.$$

Let $v_1 \sim v_2 : T_1; \theta; \delta$. We must show that

$$\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) v_1 \simeq \theta_2(\delta_2(\lambda x: T_1. x)) v_2 : T_2; \theta; \delta[v_1, v_2/z]$$

for fresh z . If $T_1 = T_2$, then we are done:

$$\begin{aligned}
\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) v_1 & \longrightarrow v_1 \\
\theta_2(\delta_2(\lambda x: T_1. x)) v_2 & \longrightarrow v_2.
\end{aligned}$$

Otherwise,

$$\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) v_1 \longrightarrow \Lambda \alpha. (\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) (v_1 \alpha)).$$

So, it suffices to show that

$$\Lambda \alpha. (\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) (v_1 \alpha)) \simeq v_2 : T_2; \theta; \delta[v_1, v_2/z].$$

Let R, T''_1, T''_2 be given. We must show that

$$\Lambda \alpha. (\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) (v_1 \alpha)) T''_1 \simeq v_2 T''_2 : T'_2; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta[v_1, v_2/z].$$

Since

$$\Lambda \alpha. (\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) (v_1 \alpha)) T''_1 \longrightarrow \theta_1(\delta_1(\langle T_1[T''_1/\alpha] \Rightarrow T_2[T''_1/\alpha] \rangle^l)) (v_1 T''_1),$$

it suffices to show that

$$\theta_1(\delta_1(\langle T_1[T''_1/\alpha] \Rightarrow T_2[T''_1/\alpha] \rangle^l)) (v_1 T''_1) \simeq v_2 T''_2 : T'_2; \theta[\alpha \mapsto R, T''_1, T''_2]; \delta[v_1, v_2/z].$$

By assumption

$$v_1 T_1'' \simeq v_2 T_2'' : T_1; \theta[\alpha \mapsto R, T_1'', T_2'']; \delta,$$

which normalize to

$$v_1' \sim v_2' : T_1; \theta[\alpha \mapsto R, T_1'', T_2'']; \delta.$$

By the IH,

$$\Gamma, \alpha \vdash \langle T_1' \Rightarrow T_2' \rangle^l \simeq \lambda x: T_1'. x : T_1' \rightarrow T_2'$$

and so,

$$\theta_1(\delta_1(\langle T_1[T_1''/\alpha] \Rightarrow T_2[T_1''/\alpha] \rangle^l)) \sim \lambda x: T_1'[T_2''/\alpha]. x : T_1' \rightarrow T_2'; \theta[\alpha \mapsto R, T_1'', T_2'']; \delta.$$

Then, by definition,

$$\begin{aligned} \theta_1(\delta_1(\langle T_1[T_1''/\alpha] \Rightarrow T_2[T_1''/\alpha] \rangle^l)) v_1' &\simeq (\lambda x: T_1'[T_2''/\alpha]. x) v_2' \\ &: T_2'; \theta[\alpha \mapsto R, T_1'', T_2'']; \delta[v_1', v_2'/z'] \end{aligned}$$

for fresh z' . They normalize to

$$v_1'' \sim v_2' : T_2'; \theta[\alpha \mapsto R, T_1'', T_2'']; \delta[v_1', v_2'/z'],$$

and by expansion we are done:

$$\theta_1(\delta_1(\langle T_1[T_1''/\alpha] \Rightarrow T_2[T_1''/\alpha] \rangle^l)) (v_1 T_1'') \simeq v_2 T_2'' : T_2'; \theta[\alpha \mapsto R, T_1'', T_2'']; \delta[v_1, v_2/z].$$

(S_REFINE): We have $\Gamma \vdash \text{unref}(T_1) <: \text{unref}(T_2)$ and

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

Let $\Gamma \vdash \theta; \delta$. We wish to show that

$$\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) \simeq \theta_2(\delta_2(\lambda x: T_1. x)) : T_1 \rightarrow T_2; \theta; \delta.$$

Let $v_1 \sim v_2 : T_1; \theta; \delta$. We must show that

$$\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) v_1 \simeq \theta_2(\delta_2(\lambda x: T_1. x)) v_2 : T_2; \theta; \delta[v_1, v_2/z]$$

for fresh z .

We have three cases according to how the LHS reduces.

($T_2 = \text{unref}_1^j(T_1)$ for some j): Then, we have

$$\begin{aligned} &\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) v_1 \\ \text{(by E_FORGET)} &\longrightarrow^* \theta_1(\delta_1(\langle T_2 \Rightarrow T_2 \rangle^l)) v_1 \\ \text{(by E_REFL)} &\longrightarrow v_1. \end{aligned}$$

Since $\theta_2(\delta_2(\lambda x: T_1. x)) v_2 \longrightarrow v_2$, we have $v_1 \sim v_2 : T_1; \theta; \delta[v_1, v_2/z]$. Since $\text{unref}_1^j(T_1) = T_2$, we also have $v_1 \sim v_2 : T_2; \theta; \delta[v_1, v_2/z]$. So, by expansion,

$$\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) v_1 \simeq \theta_2(\delta_2(\lambda x: T_1. x)) v_2 : T_2; \theta; \delta[v_1, v_2/z].$$

($T'_1 = \text{unref}_1^j(T_1)$ for some j and $T_2 = \{x:T'_1 \mid e_2\}$): Then, we have

$$\begin{aligned} & \theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) v_1 \\ \text{(by E_FORGET)} & \longrightarrow^* \theta_1(\delta_1(\langle T'_1 \Rightarrow \{x:T'_1 \mid e_2\} \rangle^l)) v_1 \\ \text{(by E_CHECK)} & \longrightarrow \langle \theta_1(\delta_1(\{x:T'_1 \mid e_2\})), \theta_1(\delta_1(e_2))[v_1/x], v_1 \rangle^l. \end{aligned}$$

By assumption, $v_1 \simeq v_2 : \text{unref}(T_1); \theta; \delta$ and $\theta_1(\delta_1(\text{casts}(T_1))) v_1 \longrightarrow^* v_1$. By IMP, for some v'_1 we have,

$$\theta_1(\delta_1(\text{casts}(T_2))) (\theta_1(\delta_1(\langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l)) v_1) \longrightarrow^* v'_1.$$

By inspecting the reduction sequence, we have

$$\begin{aligned} & \text{casts}(T_2) (\langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l v_1) \\ \text{(by E_REFL)} & \longrightarrow \text{casts}(T_2) v_1 \\ & \longrightarrow^* \langle T'_1 \Rightarrow \{x:T'_1 \mid e_2\} \rangle^l v_1 \\ \text{(by E_CHECK)} & \longrightarrow \langle \theta_1(\delta_1(\{x:T'_1 \mid e_2\})), \theta_1(\delta_1(e_2))[v_1/x], v_1 \rangle^l \\ & \longrightarrow^* \langle \theta_1(\delta_1(\{x:T'_1 \mid e_2\})), \text{true}, v_1 \rangle^l \\ & \longrightarrow v_1. \end{aligned}$$

Since $\theta_2(\delta_2(\lambda x:T_1. x)) v_2 \longrightarrow v_2$, we have

$$v_1 \sim v_2 : T_1; \theta; \delta[v_1, v_2/z].$$

Since $\text{unref}_1^j(T_1) = T'_1$, we also have

$$v_1 \sim v_2 : T'_1; \theta; \delta[v_1, v_2/z].$$

Finally, $\theta_1(\delta_1(e_2))[v_1/x] \longrightarrow^* \text{true}$ gives $v_1 \sim v_2 : \{x:T'_1 \mid e_2\}; \theta; \delta[v_1, v_2/z]$.
Finally, by expansion

$$\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) v_1 \simeq \theta_2(\delta_2(\lambda x:T_1. x)) v_2 : T_2; \theta; \delta[v_1, v_2/z].$$

(Otherwise): We have

$$\begin{aligned} & \theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) v_1 \\ \text{(by E_FORGET)} & \longrightarrow^* \theta_1(\delta_1(\langle \text{unref}(T_1) \Rightarrow T_2 \rangle^l)) v_1 \\ \text{(by E_PRECHECK)} & \longrightarrow^* e \end{aligned}$$

where

$$e = \theta_1(\delta_1(\langle \text{unref}_1(T_2) \Rightarrow T_2 \rangle^l (\langle \text{unref}_2(T_2) \Rightarrow \text{unref}_1(T_2) \rangle^l (\dots (\langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l v_1)))))).$$

By the IH,

$$\Gamma \vdash \langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l \simeq \lambda x: \text{unref}(T_1). x : (\text{unref}(T_1) \rightarrow \text{unref}(T_2)).$$

That is,

$$\theta_1(\delta_1(\langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l)) v_1 \simeq \theta_1(\delta_1(\lambda x: \text{unref}(T_1). x)) v_2 \\ : \text{unref}(T_2); \theta; \delta[v_1, v_2/z].$$

They normalize to $v'_1 \sim v_2 : \text{unref}(T_2); \theta; \delta[v_1, v_2/z]$.

By assumption, $v_1 \simeq v_2 : \text{unref}(T_1); \theta; \delta$ and $\theta_1(\delta_1(\text{casts}(T_1))) v_1 \longrightarrow^* v_1$. By IMP, for some v''_1 we have:

$$\theta_1(\delta_1(\text{casts}(T_2) (\langle \text{unref}(T_1) \Rightarrow \text{unref}(T_2) \rangle^l v_1))) \longrightarrow \\ \theta_1(\delta_1(\text{casts}(T_2))) v'_1 \longrightarrow^* v''_1.$$

Inspecting this reduction sequence gives $v''_1 = v'_1$.

By Lemma 3.3.7, $\Gamma \vdash T_2 \simeq T_2 : *$, which shows $\theta_2(\delta_2(\text{casts}(T_2))) v_2 \longrightarrow^* v'_2$ for some v'_2 —that is, running all of the refinements doesn't produce blame, which means the checks must go to true. We can then conclude that $v'_1 \sim v_2 : T_2; \theta; \delta[v_1, v_2/z]$, and so, by expansion,

$$\theta_1(\delta_1(\langle T_1 \Rightarrow T_2 \rangle^l)) v_1 \simeq \theta_2(\delta_2(\lambda x: T_1. x)) v_2 : T_2; \theta; \delta[v_1, v_2/z].$$

□

□

Other optimizations

We can use subtyping to prove other optimizations correct. For example, we can show that $\{y:\{x:T \mid e_1\} \mid e_2\}$ and $\{x:T \mid e_1 \wedge e_2[x/y]\}$ are equivalent. To do so, we must first show that subtyping is reflexive. In order to show reflexivity, we must show that subtyping respects contextual equivalence of types—two types are contextually equivalent if their type parts are identical and the matching term parts are contextually equivalent. (This implies that contextually equivalent types have contextually equivalent sub-parts.) We lift this to contexts in the natural way.

3.4.4 Lemma [Subtyping respects contextual equivalence]: If T_1 is contextually equivalent to T'_1 , then (1) $\Gamma \vdash T_1 <: T_2$ iff $\Gamma \vdash T'_1 <: T_2$ and (2) $\Gamma \vdash T_2 <: T_1$ iff $\Gamma \vdash T_2 <: T'_1$.

Proof: By induction on the subtyping derivations. □

3.4.5 Lemma [Reflexivity of subtyping]: If $\Gamma \vdash T$, then $\Gamma \vdash T <: T$.

Proof: By induction on the height of $\Gamma \vdash T$, using Lemma 3.4.4 to account for the cast in the domain of S_FUN. □

3.4.6 Lemma: If $\Gamma \vdash \{y:\{x:T \mid e_1\} \mid e_2\}$, then $\Gamma \vdash \{y:\{x:T \mid e_1\} \mid e_2\} <: \{x:T \mid e_1 \wedge e_2[x/y]\}$ and $\Gamma \vdash \{x:T \mid e_1 \wedge e_2[x/y]\} <: \{y:\{x:T \mid e_1\} \mid e_2\}$.

Proof: By inversion, we know that the inner types are well formed ($\Gamma \vdash \{x:T \mid e_1\}$ and $\Gamma \vdash T$) and that the predicates are well formed ($\Gamma, x:T \vdash e_1 : \mathbf{Bool}$ and $\Gamma, y:\{x:T \mid e_1\} \vdash e_2 : \mathbf{Bool}$).

We want to apply `S_REFINE` to prove both subtypings. By reflexivity of subtyping (Lemma 3.4.5), we know that $\Gamma \vdash \text{unref}(\{y:\{x:T \mid e_1\} \mid e_2\}) <: \text{unref}(\{x:T \mid e_1 \wedge e_2[x/y]\})$ and vice versa. We must then show, forall $\Gamma, x:T \vdash \theta; \delta$, that if

$$\theta_1(\delta_1(\langle T \Rightarrow \{x:T \mid e_1\} \rangle^l (\langle \{x:T \mid e_1\} \Rightarrow \{y:\{x:T \mid e_1\} \mid e_2\} \rangle^l x)))$$

reduces to a value iff

$$\theta_1(\delta_1(\langle T \Rightarrow \{x:T \mid e_1 \wedge e_2[x/y]\} \rangle^l x))$$

does, too.

Let $\theta_1(\delta_1(x)) = v$. The former steps to a value when $e_1[v/x] \longrightarrow^* \mathbf{true}$ and $e_2[v/y] \longrightarrow^* \mathbf{true}$. The latter steps to a value $(e_1 \wedge e_2[x/y])[v/x] \longrightarrow^* \mathbf{true}$, i.e., when $e_1[v/x] \longrightarrow^* \mathbf{true}$ and $e_2[v/y] \longrightarrow^* \mathbf{true}$. \square

3.5 Type conversion: parallel reduction vs. common subexpression reduction

Belo et al. [8] used parallel reduction (defined in Figure 3.8) as the conversion relation between types, while here I've used a structural conversion relation with common subexpression reduction (CSR). In this section, I explain the need for a conversion relation, why I've made the change, and how the two approaches differ.

First, why do we need a type conversion relation at all? Suppose we are trying to prove preservation, in order to find syntactic type soundness. Consider the term $v_1 e_2$, where v_1 has the type $x:T_1 \rightarrow T_2$ and e_2 has the type T_1 . According to `T_APP`, the type of $v_1 e_2$ is $T_2[e_2/x]$. What happens when $e_2 \longrightarrow e_2'$? The syntactic type system gives us $v_1 e_2 : T_2[e_2'/x]$. To finish such a preservation proof, we must know how $T_2[e_2/x]$ and $T_2[e_2'/x]$ relate. Intuitively, they ought to be inhabited by the same values: since our evaluation semantics is deterministic, any value that satisfies the checks in $T_2[e_2/x]$ should also satisfy the checks in $T_2[e_2'/x]$, since the latter type is just a few extra steps along. We must modify the definition of our syntactic type system to make it respect this equivalence in a formal way.

In λ_H in Chapter 2, we observe that while $T_2[e_2/x]$ may not reduce to $T_2[e_2'/x]$ —types don't reduce at all, in fact—we can relate them as subtypes of each other. For this reason (among others), the λ_H system introduces a subtyping relation. But it turns out that subtyping in that language introduces a vicious cycle (see Section 5.2.2), forcing us to adopt a semantic approach to types soundness. I end up

Parallel term reduction $\boxed{e_1 \Rightarrow e_2}$

$$\begin{array}{c}
\frac{v_i \Rightarrow v'_i}{\text{op}(v_1, \dots, v_n) \Rightarrow \llbracket \text{op} \rrbracket(v'_1, \dots, v'_n)} \quad \text{EP_ROP} \qquad \frac{e_{12} \Rightarrow e'_{12} \quad v_2 \Rightarrow v'_2}{(\lambda x:T. e_{12}) v_2 \Rightarrow e'_{12}[v'_2/x]} \quad \text{EP_RBETA} \\
\\
\frac{e \Rightarrow e' \quad T_2 \Rightarrow T'_2}{(\Lambda \alpha. e) T_2 \Rightarrow e'[T'_2/\alpha]} \quad \text{EP_RTBETA} \qquad \frac{v \Rightarrow v'}{\langle T \Rightarrow T \rangle^l v \Rightarrow v'} \quad \text{EP_RREFL} \\
\\
\frac{T_2 \neq \{x:T_1 \mid e\} \quad T_2 \neq \{y:\{x:T_1 \mid e\} \mid e_2\} \quad T_1 \Rightarrow T'_1 \quad T_2 \Rightarrow T'_2 \quad v \Rightarrow v'}{\langle \{x:T_1 \mid e\} \Rightarrow T_2 \rangle^l v \Rightarrow \langle T'_1 \Rightarrow T'_2 \rangle^l v'} \quad \text{EP_RFORGET} \\
\\
\frac{T_1 \neq T_2 \quad T_1 \neq \{x:T \mid e\} \quad T_1 \Rightarrow T'_1 \quad T_2 \Rightarrow T'_2 \quad e \Rightarrow e' \quad v \Rightarrow v'}{\langle T_1 \Rightarrow \{x:T_2 \mid e\} \rangle^l v \Rightarrow \langle T'_2 \Rightarrow \{x:T'_2 \mid e'\} \rangle^l (\langle T'_1 \Rightarrow T'_2 \rangle^l v')} \quad \text{EP_RPRECHECK} \\
\\
\frac{T \Rightarrow T' \quad e \Rightarrow e' \quad v \Rightarrow v'}{\langle T \Rightarrow \{x:T \mid e\} \rangle^l v \Rightarrow \langle \{x:T' \mid e'\} \rangle^l (e'[v'/x], v')^l} \quad \text{EP_RCHECK} \\
\\
\frac{v \Rightarrow v'}{\langle \{x:T \mid e_1\}, \text{true}, v \rangle^l \Rightarrow v'} \quad \text{EP_ROK} \qquad \frac{}{\langle \{x:T \mid e_1\}, \text{false}, v \rangle^l \Rightarrow \uparrow^l} \quad \text{EP_RFAIL} \\
\\
\frac{x:T_{11} \rightarrow T_{12} \neq x:T_{21} \rightarrow T_{22} \quad T_{11} \Rightarrow T'_{11} \quad T_{12} \Rightarrow T'_{12} \quad T_{21} \Rightarrow T'_{21} \quad T_{22} \Rightarrow T'_{22} \quad v \Rightarrow v'}{\langle x:T_{11} \rightarrow T_{12} \Rightarrow x:T_{21} \rightarrow T_{22} \rangle^l v \Rightarrow \lambda x:T'_{21}. (\langle T'_{12} \langle T'_{21} \Rightarrow T'_{11} \rangle^l x/x \rangle \Rightarrow T'_{22})^l (v' (\langle T'_{21} \Rightarrow T'_{11} \rangle^l x))} \quad \text{EP_RFUN} \\
\\
\frac{\forall \alpha. T_1 \neq \forall \alpha. T_2 \quad T_1 \Rightarrow T'_1 \quad T_2 \Rightarrow T'_2 \quad v \Rightarrow v'}{\langle \forall \alpha. T_1 \Rightarrow \forall \alpha. T_2 \rangle^l v \Rightarrow \Lambda \alpha. (\langle T'_1 \Rightarrow T'_2 \rangle^l (v' \alpha))} \quad \text{EP_RFORALL} \\
\\
\frac{}{e \Rightarrow e} \quad \text{EP_REFL} \quad \frac{T_1 \Rightarrow T'_1 \quad e_{12} \Rightarrow e'_{12}}{\lambda x:T_1. e_{12} \Rightarrow \lambda x:T'_1. e'_{12}} \quad \text{EP_ABS} \quad \frac{e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2}{e_1 e_2 \Rightarrow e'_1 e'_2} \quad \text{EP_APP} \\
\\
\frac{e \Rightarrow e'}{\Lambda \alpha. e \Rightarrow \Lambda \alpha. e'} \quad \text{EP_TABS} \quad \frac{e_1 \Rightarrow e'_1 \quad T_2 \Rightarrow T'_2}{e_1 T_2 \Rightarrow e'_1 T'_2} \quad \text{EP_TAPP} \\
\\
\frac{e_i \Rightarrow e'_i}{\text{op}(e_1, \dots, e_n) \Rightarrow \text{op}(e'_1, \dots, e'_n)} \quad \text{EP_OP} \quad \frac{T_1 \Rightarrow T'_1 \quad T_2 \Rightarrow T'_2}{\langle T_1 \Rightarrow T_2 \rangle^l \Rightarrow \langle T'_1 \Rightarrow T'_2 \rangle^l} \quad \text{EP_CAST} \\
\\
\frac{T \Rightarrow T' \quad e \Rightarrow e'}{\langle T, e, k \rangle^l \Rightarrow \langle T', e', k \rangle^l} \quad \text{EP_CHECK} \quad \frac{}{E \llbracket \uparrow^l \rrbracket \Rightarrow \uparrow^l} \quad \text{EP_BLAME}
\end{array}$$

Parallel type reduction $\boxed{T_1 \Rightarrow T_2}$

$$\begin{array}{c}
\frac{}{T \Rightarrow T} \quad \text{EP_TREFL} \quad \frac{\sigma_1 \longrightarrow^* \sigma_2 \quad T_1 \Rightarrow T_2}{\{x:T_1 \mid \sigma_1(e)\} \Rightarrow \{x:T_2 \mid \sigma_2(e)\}} \quad \text{EP_TREFINE} \\
\\
\frac{T_1 \Rightarrow T'_1 \quad T_2 \Rightarrow T'_2}{x:T_1 \rightarrow T_2 \Rightarrow x:T'_1 \rightarrow T'_2} \quad \text{EP_TFUN} \quad \frac{T \Rightarrow T'}{\forall \alpha. T \Rightarrow \forall \alpha. T'} \quad \text{EP_TFORALL}
\end{array}$$

Figure 3.8: Parallel reduction

Conversion $\boxed{\sigma_1 \longrightarrow^* \sigma_2}$ $\boxed{T_1 \equiv T_2}$

$$\begin{array}{c}
\sigma_1 \longrightarrow^* \sigma_2 \iff \begin{array}{l} \text{dom}(\sigma_1) = \text{dom}(\sigma_2) \wedge \\ \forall x \in \text{dom}(\sigma_1). \sigma_1(x) \longrightarrow^* \sigma_2(x) \wedge \\ \forall \alpha \in \text{dom}(\sigma_1). \sigma_1(\alpha) = \sigma_2(\alpha) \end{array} \\
\\
\frac{}{\alpha \equiv \alpha} \text{C_VAR} \quad \frac{}{B \equiv B} \text{C_BASE} \quad \frac{\sigma_1 \longrightarrow^* \sigma_2 \quad T_1 \equiv T_2}{\{x:T_1 \mid \sigma_1(e)\} \equiv \{x:T_2 \mid \sigma_2(e)\}} \text{C_REFINE} \\
\\
\frac{T_1 \equiv T'_1 \quad T_2 \equiv T'_2}{x:T_1 \rightarrow T_2 \equiv x:T'_1 \rightarrow T'_2} \text{C_FUN} \quad \frac{T \equiv T'}{\forall \alpha. T \equiv \forall \alpha. T'} \text{C_FORALL} \\
\\
\frac{T_2 \equiv T_1}{T_1 \equiv T_2} \text{C_SYM} \quad \frac{T_1 \equiv T_2 \quad T_2 \equiv T_3}{T_1 \equiv T_3} \text{C_TRANS}
\end{array}$$

Figure 3.9: Type conversion via common subexpression reduction

showing that that $T_2[e_2/x]$ *parallel reduces* to $T_2[e'_2/x]$. That is, we can take some number of reduction steps in parallel (one for each free occurrence of x) from $T_2[e_2/x]$ to $T_2[e'_2/x]$. I then show that the denotations of such types are equal (Lemma 2.3.17), and then prove *semantic* type soundness for λ_H with respect to those denotations. Note that Lemma 2.3.17 depends on a long Coq development showing *cotermination*: if $e_1 \Rightarrow e_2$, then $e_1 \longrightarrow^* v_1$ iff $e_2 \longrightarrow^* v_2$ such that $v_1 \Rightarrow v_2$ (Lemma A20 in `thy.v`)—more on this later.

The situation for λ_H is somewhat unsatisfying. We set out to prove syntactic type soundness and ended up proving semantic type soundness along the way. While not a serious burden for a language as small as λ_H , having to use semantic techniques throughout makes adding some features—polymorphism, state and other effects, concurrency—difficult. For example, a semantic proof of type soundness for F_H is very close to a proof of parametricity—must we prove parametricity while proving type soundness?

In originally doing the work in this chapter (Belo et al. [8]), we observed that we could get rid of subtyping and explicitly use the symmetric, transitive closure of parallel reduction as the conversion relation. (Parallel reduction is reflexive by definition.) Between that (and a few other changes), we found subtyping no longer necessary. We still, however, needed cotermination for parallel reduction. We also explicitly depended on *substitutivity*: if $e_1 \Rightarrow e_2$ and $e'_1 \Rightarrow e'_2$ then $e_1[e'_1/x] \Rightarrow e_2[e'_2/x]$. It turns out that the proof of cotermination for λ_H also needs substitutivity, Lemma A3 in `thy.v`, but we needed it for the substitution in types discussed above. Unfortunately, parallel reduction in F_H is *not* substitutive [60]. There are two counterexamples, both in Figure 3.10.

Why doesn't substitutivity hold in F_H , when it did (so easily) in λ_H ? There are two reasons. First, the cast semantics of F_H is much more complicated than that of λ_H (six rules, as opposed to two). The F_H rules depend on upon certain (syntactic) equalities between types—both counterexamples in Figure 3.10 take advantage these

Counterexample 1

Let T be a type with a free variable x .

$$\begin{aligned} e_1 &= \langle T \Rightarrow \{y: T[5/x] \mid \text{true}\} \rangle^l 0 \\ e_2 &= \langle T[5/x] \Rightarrow \{y: T[5/x] \mid \text{true}\} \rangle^l (\langle T \Rightarrow T[5/x] \rangle^l 0) \\ e'_1 = e'_2 &= 5 \end{aligned}$$

Observe that $e'_1 \Rightarrow e'_2$ (by EP_REFL) and $e_1 \Rightarrow e_2$ (by EP_RPRECHECK) but $e_1[5/x] = \langle T[5/x] \Rightarrow \{y: T[5/x] \mid \text{true}\} \rangle^l 0 \Rightarrow \langle \{y: T[5/x] \mid \text{true}\}, \text{true}, 0 \rangle^l$ by EP_RCHECK, not $e_2[5/x]$.

Counterexample 2

Let T_2 be a type with a free variable x .

$$\begin{aligned} e_1 &= \langle T_1 \rightarrow T_2 \Rightarrow T_1 \rightarrow T_2[5/x] \rangle^l v \\ e_2 &= \lambda y: T_1. \langle T_2 \Rightarrow T_2[5/x] \rangle^l (v (\langle T_1 \Rightarrow T_1 \rangle^l y)) \\ e'_1 = e'_2 &= 5 \end{aligned}$$

Observe that $e'_1 \Rightarrow e'_2$ (by EP_REFL) and $e_1 \Rightarrow e_2$ (by EP_RFUN). We have $e_1[5/x] = \langle T_1 \rightarrow T_2[5/x] \Rightarrow T_1 \rightarrow T_2[5/x] \rangle^l v \Rightarrow v[5/x]$ by EP_RREFL, not $e_2[5/x]$.

Figure 3.10: Counterexamples to substitutivity of parallel reduction in F_H

equalities to break substitutivity. Second, λ_H treats $\langle x: T_{11} \rightarrow T_{12} \Rightarrow x: T_{21} \rightarrow T_{22} \rangle^l v$ as a value, while it is a redex in F_H . This syntactic coincidence makes an exact cotermination lemma possible. But as the second counterexample shows, in F_H it's possible to have a substitution introduce a function proxy in e_2 but not e_1 . While I conjecture that e_1 and e_2 are *contextually* equivalent, they won't yield values that parallel reduce to each other. The rules that are the source of the problem for substitutivity of parallel reduction are the EP_R... rules, where a reduction in the outer term happens at the same time as parallel reductions deep inside the term. Note that both counterexamples make use of such rules.

The semantics of F_H as published in ESOP 2011 are wrong. To fix them, I introduced a simpler conversion relation, defined in Figure 3.9. Instead of allowing full parallel reduction, I restrict convertible types to the CSR $\sigma_1 \longrightarrow^* \sigma_2$, i.e., substitutions over the same set of term and type variables where (a) every term binding in σ_1 reduces to its corresponding binding in σ_2 , and (b) the type bindings are identical. My conversion relation is essentially the symmetric, transitive closure⁵ of parallel reduction—*without* these reducing rules.

Phrasing the conversion relation in terms of CSR gives us substitutivity nearly automatically, but cotermination remains an issue. In Conjecture 3.2.1, I suggest that terms related by CSR coterminate at **true**; this is enough to prove type soundness and parametricity of F_H .

It is unclear if cotermination of parallel reduction holds in F_H despite the absence

⁵I prove that my relation is reflexive in Lemma 3.2.3.

of substitutivity, i.e., whether if $e_1 \Rightarrow e_2$ then $e_1 \longrightarrow^* v_1$ iff $e_2 \longrightarrow v_2$ such that $v_1 \Rightarrow v_2$. But it turns out that we can get by with a simpler property: cotermination at `true`, rather than at arbitrary values. This property is a corollary of cotermination, since `true` \Rightarrow `true`, but it is less likely to be interfered with by function proxies which may be introduced as in the second counterexample. The intuition that leads me to believe that cotermination at `true` holds for CSR is that in a well typed program, any extra checks or function proxies introduced due to differing substitutions must eventually disappear if the type of the final expression is `Bool`.

I believe that weak bisimulation is a promising proof technique: it’s syntactic enough to avoid issues of circularity with typing, but semantic enough to relate terms that are related by \longrightarrow^* reductions. I think *weak* bisimulations in particular are appropriate, because of the need for \longrightarrow^* reductions on both sides of the relation. The system as defined may make such a proof slightly difficult. Recalling the first counterexample to substitutivity, we will need to have $e_1[e'_1/x]$ and $e_2[e'_2/x]$ in the relation, but how can the relation “remember” that any checks that occur on one side but not the other inevitably succeed? Introducing explicit tagging, as I do in Chapter 4, is an attractive approach to solving this technical problem. In an explicitly tagged manifest contract system, the only values inhabiting refinement types are tagged as such, e.g., $(v, \{x:T \mid e\})$; the operational semantics then manages tags on values, tagging in `E_CHECKOK` and untagging in `E_FORGET`. Explicit tagging has several advantages: it clarifies the staging of the operational semantics; it eliminates the need for a `T_FORGET` rule; it gives value inversion directly (Lemma 3.2.11). Finally, any proof of cotermination at `true` (Conjecture 3.2.1) must be careful to *not* rely on type soundness, preservation, or substitution properties. Those theorems in F_H rely on Conjecture 3.2.1, so we can’t use them in its proof.

Finally: what kind of calculus *wouldn’t* have cotermination at `true`? In a non-deterministic language, CSR may make one choice with σ_1 and another with σ_2 . Fortunately, F_H is deterministic. In a deterministic language, cotermination at `true` may not hold for CSR if the evaluation relation abuses equalities that are violated by reduction. F_H ’s semantics *does* use equalities that are violated by term reduction; I believe that “abuse” means using an equality on part of a term to determine which step to take, but then ignoring that part of the term later in evaluation. Since F_H *doesn’t* do that, I am confident enough to conjecture that cotermination at `true` holds.

3.6 Conclusion

This chapter presented a simpler approach to manifest contract calculi, which I applied to defining F_H , a parametrically polymorphic manifest contract calculus. When I say “parametrically” polymorphic, I mean in particular that the relation R used to related terms at type variables in the logical relation is a *parameter* of the logical re-

lation, which admits any instantiation of R .⁶ I offered the first operational semantics for *general* refinements, where refinements can apply to any type, not just base types. Finally, I defined a *post facto* subtyping relation, proving that “upcasts” from subtypes to supertypes always succeed in F_H , i.e., that subtyping is sound. This recovers the reasoning principles lost when we left subtyping out of the language definition.

⁶Earlier versions [8] only admit relations that respect parallel reduction, but that restriction has been relaxed.

Chapter 4

Space-efficient manifest contracts

Space and Time! now I see it is true, what I guessd at,
What I guessd when I loafd on the grass,
What I guessd while I lay alone in my bed,
And again as I walkd the beach under the paling stars of the morning.

Song of Myself
Walt Whitman

There has been a great deal of investigation into so-called *full-spectrum* programming languages, combining dynamic types, simple types, and contract-style refinements. The promise of a single language admitting the full development cycle—from a small script to a more manageable, statically typed program to a robust, verified system—has held great allure for some time. Prior attempts to fulfill this promise have attacked the problem piecemeal: script to program, and program to verified program. This dissertation so far is no exception: we have studied contracts as an extension on top of simple types (Chapter 2) and System F (Chapter 3).

I find it useful to think of one axis of the design space as a spectrum of expressivity, ranging from dynamic types on the left to dependent refinement types on the right (illustrated in Figure 4.1). I call it the “dyn/refine spectrum”. On the one hand, work in the script-to-program category goes back at least to Abadi et al.’s work with type **Dynamic**, with more recent work falling under Siek and Taha’s rubric of “gradual typing” [1, 73, 67, 46, 74, 4, 65, 23, 18, 39, 68]. On the other hand, program verification is an enormous field in its own right; in this dissertation, I focus on dynamic enforcement methods in general and contracts in particular. We have seen that some more recent work takes a type-oriented, or *manifest*, approach to contracts, allowing so-called *refinement types* of the form $\{x:T \mid e\}$, inhabited by values v that satisfy the *predicate* e , i.e. $e[v/x] \longrightarrow^* \text{true}$.

Over the last decade, the state of the art combining these two paradigms—in gradual types and manifest contracts in particular—has steadily progressed [10, 45, 78, 51]. Starting with Sage [45] and continuing with Wadler and Findler [78], many

dynamic types \longleftrightarrow simple types \longleftrightarrow refinement types \longleftrightarrow dependency

Figure 4.1: The dyn/refine spectrum of cast expressiveness

languages have expressed the interactions between dynamic and simple static types and between simple static types and refinement types using a single syntactic form: the cast. The previous two chapters are no exception.

Casts are promising. They offer a unified view of changes in type information, have straightforward operational semantics, and enjoy a fruitful relationship with subtyping (see [78, 65, 44, 34, 68, 8]). The space-efficient full-spectrum language I develop in this chapter derives its semantics from a full-spectrum language with casts. Before we can continue, I must explain how casts work with **Dyn**, the dynamic type.

The first consequence of introducing **Dyn** to our language is the introduction of nontermination: in earlier chapters, the typing regime prevented divergence. But it is easy to write a nonterminating program using **Dyn**—we must simply deal with it. In any case, this is a theoretical difficulty; pragmatically, I *want* languages which have general recursion (and, so, the potential for nontermination).

As for casts into and out of **Dyn**, a cast of the form $\langle \text{Int} \Rightarrow \text{Dyn} \rangle 5$ asks for the number 5, which has type **Int**, to be treated as a value of type **Dyn**, the dynamic type. The operational semantics of such a cast will mark the value with a tag, as in $5_{\text{Int!}}$. Similarly, a cast of the form $\langle \text{Dyn} \Rightarrow \text{Int} \rangle 5_{\text{Int!}}$ will project the tagged integer out of type dynamic, yielding the original value 5. In the case where the cast’s argument isn’t tagged correctly, the cast must raise an error. Consider the term $\langle \text{Dyn} \Rightarrow \text{Int} \rangle \text{true}_{\text{Bool!}}$. It tries to project an **Int** out of type dynamic, but the dynamic value is really a **Bool**—a type error. All we can reasonably do is abort the program, evaluating to the uncatchable exception `fail`.¹

We have already seen that casts on functions are the most interesting casts: values with functional casts on them, like $\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle v_1$, are themselves values; they are *wrapped* with a *function proxy*. When such wrapped values are applied to a value v_2 , the cast unfolds:

$$(\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle v_1) v_2 \longrightarrow \langle T_{12} \Rightarrow T_{22} \rangle (v_1 (\langle T_{21} \Rightarrow T_{11} \rangle v_2))$$

Note that this rule is contravariant in the domain. I use a different semantics in Chapter 3, where I introduce new lambdas. Since our subject in this chapter is space efficiency, introducing extra closures for explicit function proxies is a non-starter.

One problem common to calculi with casts is the problem of *space efficiency*. In particular, casts can accumulate in an unbounded way: redundant casts can grow the stack arbitrarily; casting functions can introduce an arbitrary number of function

¹I forgo blame in this chapter, leaving it as future work. Since my final result is an inexact relationship between the naïve and space-efficient calculi, tracking blame doesn’t seem particularly useful—it won’t match up.

proxies. This unbounded growth of casts can, in the extreme, change the asymptotic complexity of programs. To see why the naïve treatment isn’t space efficient, recall the mutually recursive definition of `even` and `odd` from Section 1.3, adapted from Herman et al. [39]:

$$\begin{aligned} \text{even} & : \text{Dyn} \rightarrow \text{Dyn} = \langle \text{Int} \rightarrow \text{Bool} \Rightarrow \text{Dyn} \rightarrow \text{Dyn} \rangle \lambda x : \text{Int}. \\ & \quad \text{if } x = 0 \text{ then true else odd } (x - 1) \\ \text{odd} & : \text{Int} \rightarrow \text{Bool} = \lambda x : \text{Int}. \\ & \quad \text{if } x = 0 \text{ then false else } (\langle \text{Dyn} \rightarrow \text{Dyn} \Rightarrow \text{Int} \rightarrow \text{Bool} \rangle \text{even}) (x - 1) \end{aligned}$$

In this example, `even` is written in a more dynamically typed style than `odd`. (While this example is contrived, it is easy to imagine mutually recursive modules with dynamic typing and refinement types. For example, in PLT Racket [55], Typed Racket [74] modules interoperate with a number of untyped Racket modules. In those cases, it’s possible to accumulate an unbounded number of redundant checks on the stack or as function proxies, e.g., on continuations or callbacks.) The cast $\langle \text{Int} \rightarrow \text{Bool} \Rightarrow \text{Dyn} \rightarrow \text{Dyn} \rangle (\lambda x : \text{Int}. \dots)$ in the definition of `even` will (a) check that `even`’s dynamically typed arguments are in fact `Int`s and (b) tag the resulting booleans into the dynamic type. The symmetric cast on `even` in the definition of `odd` serves to make the function `even` behave as if it were typed. This cast will ultimately cast the integer value $n - 1$ into the dynamic type, `Dyn`, as well as projecting `even`’s result out of type `Dyn` and into type `Bool`. Now consider the reduction sequence in Figure 4.2, observing how the number of coercions grows (redexes are highlighted).

While the operational semantics doesn’t have an explicit stack, we can still see the accumulation of pending casts. What *were* tail-recursive calls are accumulating extra work in the continuation.² What’s more, the work is redundant: we must tag and untag `true` twice. In short, casts have taken an algorithm that should use a constant amount of stack space and turned it into an algorithm that uses $O(n)$ stack space. Such a large space overhead is impractical: casts aren’t space efficient.

The same holds true for casts into and out of refinement types. Consider a library of drawing primitives based around painters, functions of type `Canvas`→`Canvas`. An underlying graphics library offers basic functions for manipulating canvases and functions over canvases, e.g., `primFlipH` : $(\text{Canvas} \rightarrow \text{Canvas}) \rightarrow (\text{Canvas} \rightarrow \text{Canvas})$ flips an image on its horizontal axis. A wrapper library may add derived functions while re-exporting the underlying functions with refinement types specifying a square canvas dimensions, where `SquareCanvas` = $\{x : \text{Canvas} \mid \text{width}(x) = \text{height}(x)\}$:

$$\begin{aligned} \text{flipH } p & = \langle \text{Canvas} \rightarrow \text{Canvas} \Rightarrow \text{SquareCanvas} \rightarrow \text{SquareCanvas} \rangle \\ & \quad (\text{primFlipH } (\langle \text{SquareCanvas} \rightarrow \text{SquareCanvas} \Rightarrow \text{Canvas} \rightarrow \text{Canvas} \rangle p)) \end{aligned}$$

The wrapper library only accepts painters with appropriately refined types, but must strip away these refinements before calling the underlying implementation—which de-

²The call to `even` in the `else` branch of `odd` doesn’t look like a tail call, but if the coercions are inserted automatically—as in Herman et al.—then it can be very difficult to tell what is and isn’t a tail call. See related work (Chapter 5) for a discussion of cast insertion.

```

odd 3
→ (⟨Dyn→Dyn ⇒ Int→Bool⟩ even) 2
→ ⟨Dyn ⇒ Bool⟩ (even (⟨Int ⇒ Dyn⟩ 2))
→ ⟨Dyn ⇒ Bool⟩ (even 2Int!)
→ ⟨Dyn ⇒ Bool⟩ ((⟨Int→Bool ⇒ Dyn→Dyn⟩ (λx:Int. ...)) 2Int!)
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ ((λx:Int. ...) (⟨Dyn ⇒ Int⟩ 2Int!)))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ ((λx:Int. ...) 2))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ (odd 1))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ ((⟨Dyn→Dyn ⇒ Int→Bool⟩ even) 0))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ (⟨Dyn ⇒ Bool⟩ (even (⟨Int ⇒ Dyn⟩ 0))))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ (⟨Dyn ⇒ Bool⟩ (even 0Int!)))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ (⟨Dyn ⇒ Bool⟩
  ((⟨Int→Bool ⇒ Dyn→Dyn⟩ (λx:Int. ...)) 0Int!)))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ (⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩
  ((λx:Int. ...) (⟨Dyn ⇒ Int⟩ 0Int!))))))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ (⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩
  ((λx:Int. ...) 0))))))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ (⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ true)))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ (⟨Dyn ⇒ Bool⟩ trueBool!))
→ ⟨Dyn ⇒ Bool⟩ (⟨Bool ⇒ Dyn⟩ true)
→ ⟨Dyn ⇒ Bool⟩ trueBool!
→ true

```

Figure 4.2: Space-inefficient reduction

mands $\text{Canvas} \rightarrow \text{Canvas}$ painters. The wrapper library then has to cast these modified functions *back* to the refined types. Calling `flipH` (`flipH p`) will yield:

```

⟨Canvas→Canvas ⇒ SquareCanvas→SquareCanvas⟩
  (primFlipH
    (⟨SquareCanvas→SquareCanvas ⇒ Canvas→Canvas⟩
      (⟨Canvas→Canvas ⇒ SquareCanvas→SquareCanvas⟩ (primFlipH p')))))

```

where $\langle \text{SquareCanvas} \rightarrow \text{SquareCanvas} \Rightarrow \text{Canvas} \rightarrow \text{Canvas} \rangle p \rightarrow^* p'$

That is, we first cast p to a plain painter and return a new painter p' . We then cast p' into and then immediately out of the refined type, before continuing on to flip p' . All the while, we are accumulating many function proxies beyond the wrapping that the primitive library is doing. A space-efficient scheme for manifest contracts avoids accumulating these extra function proxies on p' .

Two of the existing approaches to space efficiency in the world of gradual typing [39, 65, 68] factor casts into their constituent *coercions*, adapting Henglein’s system [38]. I take a similar approach. For example, the cast $\langle \text{Dyn} \Rightarrow \text{Bool} \rangle$, which checks that a dynamic value is a boolean and then untags it, is written as the coercion `Bool?`; the cast $\langle \text{Bool} \Rightarrow \text{Dyn} \rangle$, which tags a boolean into type dynamic, is written `Bool!`. Most importantly, coercions can be composed, so $\langle \text{Bool}? \rangle (\langle \text{Bool}! \rangle e) \rightarrow \langle \text{Bool}!; \text{Bool}? \rangle e$.

Herman et al. normalize the coercion $\text{Bool!}; \text{Bool?}$ into the no-op coercion Id . This normalization process is how they achieve space efficiency. For example:

$$\langle \text{Bool?} \rangle (\langle \text{Bool!} \rangle ((\lambda x:\text{Int}. \dots) (\langle \text{Int?} \rangle 2_{\text{Int!}}))) \longrightarrow \langle \text{Id} \rangle ((\lambda x:\text{Int}. \dots) (\langle \text{Int?} \rangle 2_{\text{Int!}}))$$

This composition and normalization of pending coercions allows them to prove a bound on the size of any coercion that occurs during the run of a given program. This bound effectively restores the possibility of tail-call optimization.

However, it isn't obvious how to extend Herman et al.'s [39] coercion system to refinement types. When do we test that values satisfy predicates? How do refinement type coercions normalize? I show that refinement types should have a checking coercion $\{x:T \mid e\}?$ and an (un)tagging coercion $\{x:T \mid e\}!$; the key insight for space efficiency is that the composition $\{x:T \mid e\}?.\{x:T \mid e\}!$ should normalize to Id , i.e., checks that are immediately forgotten should be thrown away. Throwing away checks sounds dangerous, but the calculus is still sound—values typed at refinement types must satisfy their refinements. On the one hand, this is great news—space-efficiency is not only more practical, but there are fewer errors! On the other hand, the space-efficient semantics aren't *exactly* equivalent to the naïve semantics. Whether or not this is good news, these dropped checks are part and parcel of space efficiency. I develop this idea in detail in Section 4.5; I show that this means that space-efficient programs fail less often than their naïve counterparts in Section 4.6.

In this chapter, I make several contributions, extending the existing solutions in a number of dimensions. In particular, I:

- Present a new approach to coercions that extends the earlier work to refinement types in a novel formulation (Section 4.5);
- Introduce what is, at present, the most expressive full-spectrum contract language, offering type dynamic as well as refinements of both base types and the dynamic type (Section 4.3 and Section 4.5); this new language narrowly edges out Wadler and Findler [78] by including refinements of type dynamic;
- Show that this language is space efficient, i.e., there are a bounded number of coercions in any program, and those coercions are bounded in size (Section 4.7); and
- Offer proof my new language is sound with respect to the naïve, inefficient semantics: if the naïve semantics reaches a value, so does the space-efficient one, but occasionally, the naïve semantics will fail when the space-efficient one succeeds (Section 4.6).

Outline

In this long chapter, I begin by briefly overviewing my design philosophy in Section 4.1. I then define a cast calculus, CAST , with dynamic and refinement types (Section 4.2), followed by a naïve, space-inefficient coercion calculus, NAIVE , with

dynamic and refinement types (Section 4.3). I translate CAST terms into behaviorally equivalent NAIVE terms in Section 4.4. This proof justifies my eventual claim that I’ve made *manifest contracts* space-efficient: since CAST terms have exactly corresponding NAIVE terms, which are simulated by EFFICIENT terms, we know that EFFICIENT terms simulate CAST terms. If I didn’t relate CAST and NAIVE, I would have invented a system that enjoys space efficiency—and looks like manifest contracts, if you squint. In Section 4.5 I make the naïve calculus into a space-efficient one, which I call EFFICIENT. In Section 4.6, I relate the two coercion calculi, showing that the two are *mostly* observationally equivalent: if the naïve semantics reduces a term to a value, so will the space efficient one; sometimes the naïve semantics will diverge or produce a failure when the space-efficient one succeeds. I give formal justification for my claim of space efficiency in Section 4.7.

4.1 Design philosophy

My design philosophy has three principles:

1. Base values have simple types; e.g., all integers are typed at `Int`.
2. We give operations types precise enough to guarantee totality; e.g., division has a type at least as precise as $\text{Int} \rightarrow \{x:\text{Int} \mid x \neq 0\} \rightarrow \text{Int}$.
3. Values satisfy their refinement types; e.g., if $\emptyset \vdash v : \{x:T \mid e\}$, then $e[v/x] \rightarrow^* \text{true}$.

Giving base values simple types by avoids some of the technical problems seen in previous refinement type systems that gave constants most specific types [51, 28, 45, 44, 34]. These other systems set $\text{ty}(k)$ such that if $e[k/x] \rightarrow^* \text{true}$, then $\text{ty}(k) <: \{x:T \mid e\}$. We wish to avoid subtyping—and its concomitant circularities [34, 44, 8]. Beyond technical issues, simple types are a de facto default in functional programming; Hindley-Milner is a sweet spot in the design space. I believe that taking this sweet spot as the default programming paradigm will increase the usability and practicality of manifest contracts.

There is one exception to assigning values simple types: it makes sense to give `null` or `undefined` a dynamic type; attempting to coerce such a value to a simple type would result in an error. I omit such values—it’s not clear that they’re a good idea [42]. One alternative, due to Stephanie Weirich, is to have these dummy values inhabit *all* types, in which case they would never need to be tagged or checked at all, and elimination forms would raise errors.

To me, keeping operations total is the main goal of refinement types. In fact, the original use of refinement types by Freeman and Pfenning [30] was to make partial pattern matching total. The calculi here don’t have algebraic datatypes, but partial operations are just as undesirable as partial pattern matches. This point of view is nice for implementations: primitive partial operations can be exposed at types strong

enough to guarantee safety; the coercion algorithm will insert appropriate checks which may or may not be optimized away. (This is *hybrid type checking* [28].)

The flipside of using refinements to protect partial operations is that refinement checks that don't end up protecting partial operations aren't as important as those that do. In EFFICIENT in Section 4.5, we occasionally skip checking whether or not a value satisfies a refinement. We only do this when the program's next step would be to untag the value as having satisfied the refinement—why bother checking if we don't care about the result? This lenient approach to refinement type checking means that the naïve and space-efficient calculi don't behave exactly the same: some programs raise errors in the former but not the latter, precisely because we skip checks.

One might ask: if it's okay to have a ψ rule skip refinement type checks, why not do the same for checks when moving into and out of Dyn? The simplest explanation is that we *must* have a ψ rule for refinement types—without it, we won't have space efficiency. But the (safe) ϕ rule suffices when we consider Dyn. Since we *can* have the gradual typing parts of NAIVE and EFFICIENT behave the same, we do. As a philosophical difference, I hold the line at the structure of values, i.e., simple types. If simple types are the default paradigm of the language, we want to avoid errors in this fragment as much as possible. As Siek and Wadler [68] say:

[the] boundary between static and dynamic typing regions require[s] certain run-time checks to maintain the integrity of the static region.

The refined world has the structure of the simply typed world plus a predicate—it has more type information. On the other hand, dynamic types have less information, so we must be stricter with them. Put another way, simple types are about the broad structure of values, while refinements are about the safety of operations. If we never end up running the operation whose safety is ensured by a refinement type, no problem; the programmer's specification may be too tight. Polymorphism aside, there is no such thing as a simple type that is “too tight”.

My third principle is a compromise between stringency and lenience. If a value is typed at a refinement type, then it satisfies its predicate. All well-typed refinement calculi have value inversion—it is a consequence of subject reduction. Having only value inversion instead of stronger reasoning principles (as in the semantics of Xu et al. [80]) is a compromise because some checks can be skipped (lenience), but once the program reduces to a value it must actually inhabit its type (stringency). This *value inversion* principle is valuable, and one that we have used in practice when programming with refinement types. In call-by-value (CBV) languages (including all of the calculi in this thesis), value inversion lets programmers reason easily about functions which take refined inputs.

Value inversion is the *least* that a refinement type can mean in a sound system: any less and type soundness won't hold. While every sound refinement type system has value inversion, I find it desirable to be able to directly invert the typing derivation to find that a value satisfies the predicate of its refinement type, rather than applying a

Types and base types
 $T ::= B \mid T_1 \rightarrow T_2 \mid \{x:B \mid e\} \mid \text{Dyn} \mid \{x:\text{Dyn} \mid e\}$
 $B ::= \text{Bool} \mid \text{Int} \mid \dots$

Terms, results, values, and pre-values
 $e ::= x \mid r \mid \text{op}(e_1, \dots, e_n) \mid e_1 e_2 \mid \langle T_1 \Rightarrow T_2 \rangle e \mid \langle \{x:T \mid e_1\}, e_2, v \rangle$
 $r ::= v \mid \text{fail}$
 $v ::= u_{\text{d}} \mid v_{\text{B!}} \mid v_{\text{Fun!}} \mid v_{\{x:T \mid e\}^?} \mid v_{\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle}$
 $u ::= k \mid \lambda x:T. e$

Typing contexts
 $\Gamma ::= \emptyset \mid \Gamma, x:T$

Figure 4.3: CAST syntax

more complicated or general reasoning principle. I believe that having value inversion directly is easier for programmers to understand and use.

4.2 A cast calculus

In this section, we define CAST, a cast calculus with dynamic and refinement types. It is a small extension of Wadler and Findler’s system [78].

I define the syntax and static semantics in Section 4.2.1; the operational semantics are in Section 4.2.2. I prove type soundness in Section 4.2.3.

Rule naming conventions

Before I begin my technical work in earnest, a word about conventions. We are defining three calculi: one with casts (CAST) and two with coercions (NAIVE and EFFICIENT). The latter two largely share syntax and typing rules; all typing rules will be of the form T_NAME. In general, we will rely on context to differentiate terms. The evaluation rules for CAST, in this section, are named G_NAME, using a subscripted arrow \longrightarrow_c . The evaluation rules for the naïve calculus, NAIVE, in Section 4.3 are named F_NAME, using a subscripted arrow \longrightarrow_n for the reduction relation; the space-efficient evaluation rules for EFFICIENT are in Section 4.5 are named E_NAME using a plain arrow \longrightarrow .

4.2.1 Syntax and typing

I give the syntax of CAST in Figure 4.3. The syntax is much like those of Chapters 2 and 3: I extend the simply typed lambda calculus with Dyn, refinements of base types and dynamic, casts, active checks, and failures. I sometimes write $\{x:T \mid e\}$ when it doesn’t matter whether the underlying type is Dyn or B. Notice that

Well formed contexts and types

$\boxed{\vdash \Gamma}$ $\boxed{\vdash T}$

$$\begin{array}{c}
\overline{\vdash \emptyset} \quad \text{WF_EMPTY} \qquad \qquad \qquad \frac{\vdash \Gamma \quad \vdash T}{\vdash \Gamma, x:T} \quad \text{WF_EXTEND} \\
\\
\overline{\vdash B} \quad \text{WF_BASE} \qquad \overline{\vdash \text{Dyn}} \quad \text{WF_DYN} \qquad \frac{\vdash T_1 \quad \vdash T_2}{\vdash T_1 \rightarrow T_2} \quad \text{WF_FUN} \\
\\
\frac{x:T \vdash e : \text{Bool} \quad T = B \text{ or } T = \text{Dyn}}{\vdash \{x:T \mid e\}} \quad \text{WF_REFINE}
\end{array}$$

Similar types

$\boxed{\vdash T_1 \parallel T_2}$

$$\begin{array}{c}
\frac{T \neq T_1 \rightarrow T_2}{\vdash T \parallel T} \quad \text{P_ID} \qquad \overline{\vdash \text{Dyn} \parallel T} \quad \text{P_DYNL} \qquad \overline{\vdash T \parallel \text{Dyn}} \quad \text{P_DYNR} \\
\\
\frac{\vdash T_{11} \parallel T_{21} \quad \vdash T_{12} \parallel T_{22}}{\vdash T_{11} \rightarrow T_{12} \parallel T_{21} \rightarrow T_{22}} \quad \text{P_FUN} \\
\\
\frac{\vdash T_1 \parallel T_2}{\vdash \{x:T_1 \mid e\} \parallel T_2} \quad \text{P_REFINEL} \qquad \frac{\vdash T_1 \parallel T_2}{\vdash T_1 \parallel \{x:T_2 \mid e\}} \quad \text{P_REFINER}
\end{array}$$

Figure 4.4: Typing for CAST, part 1

refinements of dynamic indirectly include refinements of functions. At the cost of having even more canonical coercions in Section 4.5.1, we could add refinements of functions. I omit them because they would have brought complexity without new insights. Going beyond refinements of functions to the general refinements of F_H in Chapter 3, however, is challenging future work (see Chapter 6). The values in CAST are structured slightly differently from what I have presented so far—I separate *values* v from *pre-values* u . In particular, values have tags while pre-values are constants and lambdas.

The typing rules for CAST are in Figures 4.4 and 4.5. Most of the rule are standard. I follow Wadler and Findler [78] in tracking completed casts with explicit tags—this approach plays particularly well with the coercion-based approach in the sequel.

First, the well formedness rules for contexts and types are straightforward. In WF_REFINE, we ensure that we only ever refine base types and Dyn; as for other refinement type rules, we must make sure that the refinement predicate is well typed. As we did for F_H (in Figure 3.4), we must define type similarity. Note that Dyn is similar to every type. We restrict the P_ID rule so that the notion of compatibility

Well typed terms and values

$$\begin{array}{c}
\boxed{\Gamma \vdash u : T} \quad \boxed{\Gamma \vdash e : T} \\
\\
\frac{\vdash \Gamma}{\Gamma \vdash k : \text{ty}(k)} \quad \text{T_CONST} \qquad \frac{\vdash T_1 \quad \Gamma, x:T_1 \vdash e_{12} : T_2}{\Gamma \vdash \lambda x:T_1. e_{12} : T_1 \rightarrow T_2} \quad \text{T_ABS} \\
\\
\frac{\Gamma \vdash u : T}{\Gamma \vdash u_{\text{d}} : T} \quad \text{T_PREVAL} \quad \frac{\Gamma \vdash v : B}{\Gamma \vdash v_{B!} : \text{Dyn}} \quad \text{T_TAGB} \quad \frac{\Gamma \vdash v : \text{Dyn} \rightarrow \text{Dyn}}{\Gamma \vdash v_{\text{Fun!}} : \text{Dyn}} \quad \text{T_TAGFUN} \\
\\
\frac{\vdash \Gamma \quad \emptyset \vdash v : T \quad \vdash \{x:T \mid e\} \quad e[v/x] \longrightarrow_c^* \text{true}_{\text{id}}}{\Gamma \vdash v_{\{x:T \mid e\} ?} : \{x:T \mid e\}} \quad \text{T_TAGREFINE} \\
\\
\frac{\Gamma \vdash v : T_{11} \rightarrow T_{12} \quad \vdash T_{21} \rightarrow T_{22} \quad \vdash T_{11} \rightarrow T_{12} \parallel T_{21} \rightarrow T_{22}}{\Gamma \vdash v_{\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle} : T_{21} \rightarrow T_{22}} \quad \text{T_WRAP} \\
\\
\frac{\vdash \Gamma \quad x:T \in \Gamma}{\Gamma \vdash x : T} \quad \text{T_VAR} \qquad \frac{\vdash T \quad \vdash \Gamma}{\Gamma \vdash \text{fail} : T} \quad \text{T_FAIL} \\
\\
\frac{\Gamma \vdash e : T_1 \quad \vdash T_2 \quad \vdash T_1 \parallel T_2}{\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle e : T_2} \quad \text{T_CAST} \quad \frac{\Gamma \vdash e_1 : (T_1 \rightarrow T_2) \quad \Gamma \vdash e_2 : T_1}{\Gamma \vdash e_1 e_2 : T_2} \quad \text{T_APP} \\
\\
\frac{\text{ty}(op) = T_1 \rightarrow \dots \rightarrow T_n \rightarrow T \quad \Gamma \vdash e_i : T_i}{\Gamma \vdash op(e_1, \dots, e_n) : T} \quad \text{T_OP} \\
\\
\frac{\vdash \Gamma \quad \vdash \{x:T \mid e_1\} \quad \emptyset \vdash v : T \quad \emptyset \vdash e_2 : \text{Bool} \quad e_1[v/x] \longrightarrow_c^* e_2}{\Gamma \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle : \{x:T \mid e_1\}} \quad \text{T_CHECKCAST}
\end{array}$$

Figure 4.5: Typing for CAST, part 2

lines up more neatly with the operational semantics for casts.

The typing rules for terms are mostly standard. We have a specialized treatment of values to account for tags. `T_PREVAL` types pre-values tagged with `ld`. `T_TAGB` and `T_TAGFUN` type values that have been injected into the dynamic type, `Dyn`. `T_TAGREFINE` types values that have been checked as satisfying refinement types. Note that we structure this rule to apply only to closed values, but with a context to allow for weakening—as in Chapter 3. `T_WRAP` types values with function proxies. Otherwise, the rules should be familiar: `T_CAST` and `T_CHECKCAST` are standard, based on what we have seen so far.

It is worth taking a moment to comment on the type assignment functions $\text{ty}(k)$ and $\text{ty}(op)$ used in the `T_CONST` and `T_OP` rules. In line with my philosophy (Section 4.1), the rule for constants gives them base types: $\text{ty}(k) = B$. I require that no constants have, by default, type dynamic or a refinement type. By the same philosophy, the operator type assignment function $\text{ty}(op)$ takes operations to first-order types that ensure totality. If $\text{ty}(op) = T_1 \rightarrow \dots \rightarrow T_n \rightarrow T$, then the operation’s denotation $\llbracket op \rrbracket$ is a total function from (T_1, \dots, T_n) to T . If, for example, division is expressed as the operator `div`, then $\text{ty}(\text{div}) = \text{Int} \rightarrow \{x : \text{Int} \mid x \neq 0_{\text{ld}}\} \rightarrow \text{Int}$ or some similarly exact type. This property of operator types is critical: I believe that refinement types are meant to help programmers avoid failures in general, and failures of (fundamentally partial) primitive operations in particular.

To be clear, constants must have simple types, but operations on dynamic values are permitted—so long as they only look at the tag, not the underlying value. In particular, the proofs work for tag testing operations like `isFun` and `isB`, but not for operations that extract and then apply functions tagged with `Fun!`.

4.2.2 Operational semantics

The operational semantics are defined in Figure 4.6 (for core rules) and Figure 4.7 (for cast rules). The rules implement straightforward tagging and checking operations.

The core rules should be unsurprising. `G_FUN` captures the behavior of function proxies in application positions. Much more interesting are the cast rules of Figure 4.7. Even though we don’t have general refinements, I still adopt the cast staging techniques we used for `FH` in Chapter 3. That is, `G_CASTCHECK`, `G_CASTPRECHECK`, and `G_CASTPREDPRED` (which could just as easily be named `G_CASTFORGET`) should be familiar. `G_CASTID` is the reflexivity rule for casts, but limited to *not* apply for functions. This is done to mirror how coercions behave on function types; addressing this issue is future work (and discussed in Chapter 6).

`G_CASTFUN` and `G_CASTB` handle injections into `Dyn`: the simply attach a type tag to the value. Note that `G_CASTFUN` only applies to functions of type `Dyn`→`Dyn`. We use `G_CASTFUNDYN` to inject functions of other types into `Dyn` by stages.

`G_CASTFUNFUN` and `G_CASTBB` handle successful projections back out of `Dyn`; failed projections—mismatches between the target type of the cast and the tag on

$$\boxed{e_1 \longrightarrow_c e_2}$$

$$\begin{array}{c}
\frac{}{(\lambda x:T. e_{12})_{\text{ld}} v_2 \longrightarrow_c e_{12}[v_2/x]} \quad \text{G_BETA} \qquad \frac{}{op(v_1, \dots, v_n) \longrightarrow_c \llbracket op \rrbracket (v_1, \dots, v_n)} \quad \text{G_OP} \\
\\
\frac{}{v_1 \langle T_{11} \Rightarrow T_{12} \Rightarrow T_{21} \Rightarrow T_{22} \rangle v_2 \longrightarrow_c \langle T_{12} \Rightarrow T_{22} \rangle (v_1 \langle \langle T_{21} \Rightarrow T_{11} \rangle v_2 \rangle)} \quad \text{G_FUN} \\
\\
\frac{e_1 \longrightarrow_c e'_1}{e_1 e_2 \longrightarrow_c e'_1 e_2} \quad \text{G_APPL} \qquad \frac{e_2 \longrightarrow_c e'_2}{v_1 e_2 \longrightarrow_c v_1 e'_2} \quad \text{G_APPR} \\
\\
\frac{e_i \longrightarrow_c e'_i}{op(v_1, \dots, v_{i-1}, e_i, \dots, e_n) \longrightarrow_c op(v_1, \dots, v_{i-1}, e'_i, \dots, e_n)} \quad \text{G_OPINNER} \\
\\
\frac{e \longrightarrow_c e'}{\langle T_1 \Rightarrow T_2 \rangle e \longrightarrow_c \langle T_1 \Rightarrow T_2 \rangle e'} \quad \text{G_CASTINNER} \\
\\
\frac{e_2 \longrightarrow_c e'_2}{\langle \{x:T \mid e_1\}, e_2, v \rangle \longrightarrow_c \langle \{x:T \mid e_1\}, e'_2, v \rangle} \quad \text{G_CHECKINNER} \\
\\
\frac{}{\text{fail } e_2 \longrightarrow_c \text{fail}} \quad \text{G_APPRAISEL} \qquad \frac{}{v_1 \text{fail} \longrightarrow_c \text{fail}} \quad \text{G_APPRAISER} \\
\\
\frac{}{op(v_1, \dots, v_{i-1}, \text{fail}, \dots, e_n) \longrightarrow_c \text{fail}} \quad \text{G_OPRAISE} \\
\\
\frac{}{\langle T_1 \Rightarrow T_2 \rangle \text{fail} \longrightarrow_c \text{fail}} \quad \text{G_CASTRAISE} \qquad \frac{}{\langle \{x:T \mid e\}, \text{fail}, v \rangle \longrightarrow_c \text{fail}} \quad \text{G_CHECKRAISE}
\end{array}$$

Figure 4.6: CAST operational semantics (core rules)

$$\boxed{e_1 \longrightarrow_c e_2}$$

$$\begin{array}{c}
\frac{}{\langle T \Rightarrow \{x:T \mid e\} \rangle v \longrightarrow_c \langle \{x:T \mid e\}, e[v/x], v \rangle} \text{G_CASTCHECK} \\
\frac{}{\langle \{x:T \mid e\}, \text{true}_{\text{id}}, v \rangle \longrightarrow_c v_{\{x:T \mid e\}?}} \text{G_CHECKOK} \\
\frac{}{\langle \{x:T \mid e\}, \text{false}_{\text{id}}, v \rangle \longrightarrow_c \text{fail}} \text{G_CHECKFAIL} \qquad \frac{T \neq T_1 \rightarrow T_2}{\langle T \Rightarrow T \rangle v \longrightarrow_c v} \text{G_CASTID} \\
\frac{}{\langle \text{Dyn} \Rightarrow T_1 \rightarrow T_2 \rangle v_{B!} \longrightarrow_c \text{fail}} \text{G_CASTFUNFAILB} \\
\frac{}{\langle \text{Dyn} \Rightarrow T_1 \rightarrow T_2 \rangle v_{\text{Fun}!} \longrightarrow_c \langle \text{Dyn} \rightarrow \text{Dyn} \Rightarrow T_1 \rightarrow T_2 \rangle v} \text{G_CASTFUNFUN} \\
\frac{}{\langle B \Rightarrow \text{Dyn} \rangle v \longrightarrow_c v_{B!}} \text{G_CASTB} \qquad \frac{}{\langle \text{Dyn} \rightarrow \text{Dyn} \Rightarrow \text{Dyn} \rangle v \longrightarrow_c v_{\text{Fun}!}} \text{G_CASTFUN} \\
\frac{}{\langle \text{Dyn} \Rightarrow B \rangle v_{B!} \longrightarrow_c v} \text{G_CASTBB} \qquad \frac{B \neq B'}{\langle \text{Dyn} \Rightarrow B \rangle v_{B'!} \longrightarrow_c \text{fail}} \text{G_CASTBFAILB} \\
\frac{}{\langle \text{Dyn} \Rightarrow B \rangle v_{\text{Fun}!} \longrightarrow_c \text{fail}} \text{G_CASTBFAILFUN} \\
\frac{}{\langle T_{11} \rightarrow T_{12} \Rightarrow \text{Dyn} \rangle v \longrightarrow_c \langle \text{Dyn} \rightarrow \text{Dyn} \Rightarrow \text{Dyn} \rangle (\langle T_{11} \rightarrow T_{12} \Rightarrow \text{Dyn} \rightarrow \text{Dyn} \rangle v)} \text{G_CASTFUNDYN} \\
\frac{}{\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle v \longrightarrow_c v_{\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle}} \text{G_CASTFUNWRAP} \\
\frac{T_2 \neq \{x:T_1 \mid e\}}{\langle \{x:T_1 \mid e\} \Rightarrow T_2 \rangle v_{\{x:T_1 \mid e\}?} \longrightarrow_c \langle T_1 \Rightarrow T_2 \rangle v} \text{G_CASTPREDPRED} \\
\frac{T_1 \neq T_2 \quad T_1 \neq \{x:T'_1 \mid e'\}}{\langle T_1 \Rightarrow \{x:T_2 \mid e\} \rangle v \longrightarrow_c \langle T_2 \Rightarrow \{x:T_2 \mid e\} \rangle (\langle T_1 \Rightarrow T_2 \rangle v)} \text{G_CASTPRECHECK}
\end{array}$$

Figure 4.7: CAST operational semantics (cast rules)

		Target type			
Source type		$\{x:\text{Dyn} \mid e\}$	Dyn	Dyn→Dyn	$T_1 \rightarrow T_2$
	$\{x:\text{Dyn} \mid e\}$	I (PP, PC, C)	PP	PP, FFB FF, I	PP, FFB (FF, FW)
	Dyn	C	I	FFB (FF, I)	FFB (FF, FW)
	Dyn→Dyn	PC, F, C	F	FD, FW, F	FW
	$T_1 \rightarrow T_2$	PC, FD, FW, F, C	FD, FW, F	FW	FW
	B	PC, B, C	B	not well typed	
	$\{x:B \mid e\}$	PP, PC, B, C	PP, B	not well typed	

		Target type		
Source type		B	$\{x:B \mid e\}$	
	$\{x:\text{Dyn} \mid e\}$	PP, FFB BB	PP, PC, BFF BFB (BB, C)	
	Dyn	BB BFB BFF	PC, BFF BFB (BB, C)	
	Dyn→Dyn	not well typed		
	$T_1 \rightarrow T_2$	not well typed		
	B	I	C	
	$\{x:B \mid e\}$	PP	I (PP, PC, C)	

The vertical axis is source types, the horizontal axis is target types. Rule names are abbreviated, where I means G_CASTID, C means G_CASTCHECK, PC means G_CASTPRECHECK, etc. I use commas for sequencing and | for disjunction.

Table 4.1: CAST cast reductions, by type

the value—are handled by G_CASTBFAILB, G_CASTBFAILFUN (for casts into base types B), and G_CASTFUNFAILB (for casts into Dyn→Dyn).

The G_CASTFUNWRAP rule wraps functions with function proxies. This is the same as treating $\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle v$ as a value, as we did in Chapter 2. We use a separate tag for symmetry’s sake, since function proxies on values must be represented similarly in the sequel.

There are enough cast rules that it may not be immediately clear which rules apply when. In Table 4.1, we show how casts between each relevant type reduce. (I split out Dyn→Dyn because it is treated specially.) In order to fit the table easily on a single page, I use abbreviated rule names: a rule G_CASTRULENAME is written RN: I drop the G_CAST prefix and then take the initials of the rule. I use a regular expressions-like syntax: commas are sequential concatenation and | is disjunction. For example, a cast from $\{x:\text{Dyn} \mid e\}$ to $\{x:B \mid e\}$ will run PP, PC, BFF | BFB | (BB, C): that is, it first runs G_CASTPREDPRED, then G_CASTPRECHECK. We then check the tag, running G_CASTBFAILFUN or G_CASTBFAILB if the tag is wrong. If the type in the case matches the tag on the value being cast, we run G_CASTBB and then G_CASTCHECK.

4.2.3 Proofs

The proofs are entirely standard, culminating in a syntactic type soundness proof (Theorem 4.2.7) by progress and preservation (Lemmas 4.2.3 and 4.2.6).

4.2.1 Lemma [Determinism]: If $e \longrightarrow_c e_1$ and $e \longrightarrow_c e_2$ then $e_1 = e_2$.

Proof: By induction on $e \longrightarrow_c e_1$, observing that in each case the same rule must have applied to find $e \longrightarrow_c e_2$. \square

4.2.2 Lemma [Canonical forms]: If $\Gamma \vdash v : T$, then:

- $T = \text{Dyn}$ implies that $v = v'_{B!}$ or $v = v'_{\text{Fun!}}$ for some v' .
- $T = B$ implies that $v = k_{\text{Id}}$.
- $T = T_1 \rightarrow T_2$ implies that $v = \lambda x : T_1. e_{\text{Id}}$ or $v = v'_{\langle T'_1 \rightarrow T'_2 \Rightarrow T_1 \rightarrow T_2 \rangle}$ for some v' .
- $T = \{x : T' \mid e\}$ implies that $v = v'_{\{x : T \mid e\} ?}$ for some v' .

Proof: By case analysis on the typing derivation. \square

4.2.3 Lemma [Progress]: If $\emptyset \vdash e : T$ then there exists an e' such that $e \longrightarrow_c e'$ or e is a result.

Proof: By induction on the typing derivation.

(T_PREVAL) u_{Id} is a result.

(T_TAGB) $v_{B!}$ is a result.

(T_TAGFUN) $v_{\text{Fun!}}$ is a result.

(T_TAGREFINE) $v_{\{x : T \mid e\} ?}$ is a result.

(T_WRAP) $v_{\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle}$ is a result.

(T_VAR) Contradictory—variables aren't well typed in the empty context.

(T_FAIL) fail is a result.

(T_CAST) By the IH on $\emptyset \vdash e : T_1$, either $e \longrightarrow_c e'$ or e is a result. In the former case, we step by G_CASTINNER. If e is a result, then it is either fail or a value v . In the former case, we step by G_CASTRAISE. Otherwise, we go by cases on T_1 and T_2 .

($T_1 = \text{Dyn}$) By cases on T_2 .

($T_2 = \text{Dyn}$) We step by G_CASTID.

($T_2 = B$) By canonical forms (Lemma 4.2.2), v is either $v'_{B'}$ or $v'_{\mathbf{Fun}}$. If $B = B'$, we step by $\mathbf{G_CASTBB}$. If not, we step by $\mathbf{G_CASTBFAILB}$ or $\mathbf{G_CASTBFAILFUN}$.

($T_2 = T_{21} \rightarrow T_{22}$) By canonical forms (Lemma 4.2.2), $v = v'_{B'}$ or $v = v'_{\mathbf{Fun}}$. We step by $\mathbf{G_CASTFUNFAILB}$ or $\mathbf{G_CASTFUNFUN}$.

($T_2 = \{x:T \mid e\}$) If $T = \mathbf{Dyn}$, we step by $\mathbf{G_CASTCHECK}$. If $T = B$ (the only other option), we step by $\mathbf{G_CASTPRECHECK}$.

($T_1 = B$) By cases on T_2 .

($T_2 = \mathbf{Dyn}$) We step by $\mathbf{G_CASTB}$.

($T_2 = B'$) By inversion of $\vdash B \parallel B'$, we have $B = B'$. We step by $\mathbf{G_CASTID}$.

($T_2 = T_{21} \rightarrow T_{22}$) Contradictory, since it is not the case that $\vdash B \parallel T_{21} \rightarrow T_{22}$.

($T_2 = \{x:T \mid e\}$) If $T = B'$ (and so $B' = B$), we step by $\mathbf{G_CASTCHECK}$. Otherwise we step by $\mathbf{G_CASTPRECHECK}$.

($T_1 = T_{11} \rightarrow T_{12}$) By cases on T_2 .

($T_2 = \mathbf{Dyn}$) We step by $\mathbf{G_CASTFUN}$ or $\mathbf{G_CASTFUNDYN}$.

($T_2 = B$) Contradictory, since it is not the case that $\vdash T_{21} \rightarrow T_{22} \parallel B$.

($T_2 = T_{21} \rightarrow T_{22}$) We step by $\mathbf{G_CASTFUNWRAP}$.

($T_2 = \{x:T \mid e\}$) It must be that $T = \mathbf{Dyn}$, since it is not the case that $\vdash T_{21} \rightarrow T_{22} \parallel B$. We step by $\mathbf{G_CASTPRECHECK}$.

($T_1 = \{x:T \mid e\}$) By cases on T_2 .

($T_2 = \mathbf{Dyn}$) We step by $\mathbf{G_CASTPREDPRED}$.

($T_2 = B$) We step by $\mathbf{G_CASTPREDPRED}$.

($T_2 = T_{21} \rightarrow T_{22}$) We step by $\mathbf{G_CASTPREDPRED}$.

($T_2 = \{x:T' \mid e'\}$) We step by $\mathbf{G_CASTPREDPRED}$ or $\mathbf{G_CASTID}$.

($\mathbf{T_APP}$) We have $\emptyset \vdash e_1 e_2 : T_2$. By the IH on $\emptyset \vdash e_1 : T_1 \rightarrow T_2$, either e_1 steps, or it is a result. In the former case, we go by $\mathbf{G_APPL}$. In the latter, e_1 is either \mathbf{fail} (and we step by $\mathbf{G_APPRAISEL}$) or e_1 is a value.

Similarly, by the IH on $\emptyset \vdash e_2 : T_1$, either e_2 steps or is a result. We can apply $\mathbf{G_APPR}$ or $\mathbf{G_APPRAISER}$ (using that e_1 is a value), unless e_2 is some value v_2 . In that case, we use canonical forms to see that e_1 is either $\lambda x:T_1. e'_{1d}$ or $v_1 \langle T'_1 \rightarrow T'_2 \Rightarrow T_1 \rightarrow T_2 \rangle$. We step by $\mathbf{G_BETA}$ and $\mathbf{G_FUN}$, respectively.

($\mathbf{T_OP}$) By induction on n , applying the IH to step by either $\mathbf{G_OPINNER}$ or $\mathbf{G_OPFAIL}$. If all of the arguments are values, we step by $\mathbf{G_OP}$.

($\mathbf{T_CHECKCAST}$) By the IH, we can step the active check by $\mathbf{G_CASTINNER}$ or $\mathbf{G_CASTRAISE}$. If it's a value, we have $\emptyset \vdash v_2 : \mathbf{Bool}$, so v_2 is either \mathbf{true}_{1d} or \mathbf{false}_{1d} . We step by $\mathbf{G_CHECKKOK}$ and $\mathbf{G_CHECKFAIL}$, respectively.

□

4.2.4 Lemma [Regularity]: • If $\Gamma \vdash e : T$ then $\vdash \Gamma$ and $\vdash T$.

- If $\Gamma \vdash u : T$ then $\vdash \Gamma$ and $\vdash T$.

Proof: By mutual induction on the derivations.

(T_VAR) $\vdash \Gamma$ by assumption, which gives us $\vdash T$.

(T_CONST) $\vdash \Gamma$ By assumption, and we assume that $\vdash \text{ty}(k)$.

(T_ABS) We have $\vdash T_1$ and $\Gamma, x:T_1 \vdash e_{12} : T_2$ by assumption. By the IH, $\vdash \Gamma, x:T_1$ and $\vdash T_2$. By inversion, $\vdash \Gamma$. We have $\vdash T_1 \rightarrow T_2$ by WF_FUN.

(T_PREVAL) By the IH.

(T_TAGB) By the IH and WF_DYN.

(T_TAGFUN) By the IH and WF_DYN.

(T_TAGREFINE) By the IH and assumption.

(T_WRAP) By the IH and assumption.

(T_CAST) By the IH and assumption.

(T_FAIL) By assumption.

(T_APP) By the IH.

(T_OP) By the IH and the assumption that operators have well formed types.

(T_CHECKCAST) By assumption.

□

4.2.5 Lemma [Substitution]: If $\emptyset \vdash v : T$ then:

- If $\Gamma_1, x:T, \Gamma_2 \vdash e : T'$ then $\Gamma_1, \Gamma_2 \vdash e[v/x] : T'$.
- If $\Gamma_1, x:T, \Gamma_2 \vdash u : T'$ then $\Gamma_1, \Gamma_2 \vdash u[v/x] : T'$.

Proof: By mutual induction on the typing derivations for terms and pre-values, leaving Γ_2 general.

(T_CONST) Immediate by T_CONST.

(T_ABS) By T_ABS, using the IH on $\Gamma_1, x:T, \Gamma_2, y:T_1 \vdash e : T_2$.

(T_VAR) If x is the variable in question, then by weakening. If not, then by T_VAR.

(T_PREVAL) By the IH and T_PREVAL.

(T_TAGB) By the IH and T_TAGB.

(T_TAGFUN) By the IH and T_TAGFUN.

(T_TAGREFINE) Immediate by T_TAGREFINE, since the terms themselves are actually closed.

(T_WRAP) By the IH and T_WRAP.

(T_CAST) By the IH and T_CAST.

(T_FAIL) Immediate by T_FAIL.

(T_APP) By the IH and T_APP.

(T_OP) By the IH and T_OP.

(T_CHECKCAST) Immediate by T_CHECKCAST, since the terms themselves are actually closed.

□

4.2.6 Lemma [Preservation]: If $\emptyset \vdash e : T$ and $e \longrightarrow_c e'$, then $\emptyset \vdash e' : T$.

Proof: By induction on the evaluation derivation.

(G_BETA) By inversion, $x:T_1 \vdash e_1 : T_2$ and $\emptyset \vdash v_2 : T_1$. By substitution (Lemma 4.2.5).

(G_FUN) By inversion of T_WRAP, we know that $\vdash T_{11} \rightarrow T_{12} \parallel T_{21} \rightarrow T_{22}$. By T_CAST, T_APP, and T_CAST; we find the similarities necessary for the T_CAST rules by inversion, since only P_FUN could have applied.

(G_OP) By assumption.

(G_CASTID) Immediate.

(G_CASTFUNFAILB) We have $\vdash T_1 \rightarrow T_2$ by inversion; we are done by T_FAIL.

(G_CASTFUNFUN) We have $\vdash \text{Dyn} \parallel T_i$ by P_DYNL; by P_FUN, T_CAST, and assumption.

(G_CASTB) By assumption and T_TAGB.

(G_CASTFUN) By assumption and T_TAGFUN.

(G_CASTBB) By assumption.

(G_CASTBFAILB) We have $\vdash B$ immediately; by T_FAIL.

(G_CASTBFAILFUN) We have $\vdash B$ immediately; by T_FAIL.

(G_CASTFUNDDYN) We use P_DYNR, P_FUN in two applications of T_CAST.

(G_CASTFUNWRAP) By T_WRAP.

(G_CASTPREDPRED) By assumption and T_CAST, using $\vdash T_1 \parallel T_2$ from the inversion of $\vdash \{x:T_1 \mid e\} \parallel T_2$.

(G_CASTPRECHECK) By assumption and T_CAST. We use $\vdash T_1 \parallel T_2$ from the inversion of $\vdash T_1 \parallel \{x:T_2 \mid e\}$ in the first case and P_ID (it can't be a function type!) with P_REFINER in the second.

(G_CASTCHECK) By T_CHECKCAST, using $e[v/x] \longrightarrow_c^* e[v/x]$.

(G_CHECKOK) By T_TAGREFINE, using $e[v/x] \longrightarrow_c^* \text{true}_{\text{Id}}$.

(G_CHECKFAIL) By T_FAIL, using the assumption that $\vdash \{x:T \mid e\}$.

(G_APPL) By T_APP and the IH.

(G_APPR) By T_APP and the IH.

(G_OPINNER) By T_OP and the IH.

(G_CASTINNER) By T_CAST and the IH.

(G_CHECKINNER) By T_CHECKCAST and the IH, extending $e[v/x] \longrightarrow_c^* e_2 \longrightarrow_c e_2'$.

(G_APPRAISEL) By regularity (Lemma 4.2.4) and T_FAIL.

(G_APPRAISER) By regularity (Lemma 4.2.4) and T_FAIL.

(G_OPRAISE) By regularity (Lemma 4.2.4) and T_FAIL.

(G_CASTRAISE) By inversion, $\vdash T_2$; then, by T_FAIL.

(G_CHECKRAISE) By inversion, $\vdash \{x:T \mid e\}$; then, by T_FAIL.

□

4.2.7 Theorem [Type soundness]: If $\emptyset \vdash e : T$, then either $e \longrightarrow^* r$ such that $\emptyset \vdash r : T$ or e diverges.

Proof: Using progress (Lemma 4.2.3) and preservation (Lemma 4.2.6). Unsurprisingly, this is not a constructive proof. □

Types and base types
 $T ::= B \mid T_1 \rightarrow T_2 \mid \{x:B \mid e\} \mid \text{Dyn} \mid \{x:\text{Dyn} \mid e\}$
 $B ::= \text{Bool} \mid \text{Int} \mid \dots$

Coercions, primitive coercions, and type tags
 $c ::= d_1; \dots; d_n$
 $d ::= D! \mid D? \mid c_1 \mapsto c_2 \mid \text{Fail}$
 $D ::= B \mid \mathbf{Fun} \mid \{x:B \mid e\} \mid \{x:\text{Dyn} \mid e\}$

Terms, results, values, and pre-values
 $e ::= x \mid r \mid \text{op}(e_1, \dots, e_n) \mid e_1 e_2 \mid \langle c \rangle e \mid$
 $\langle \{x:T \mid e_1\}, e_2, v \rangle$
 $r ::= v \mid \text{fail}$
 $v ::= u_{\text{id}} \mid v_{B!} \mid v_{\mathbf{Fun}!} \mid v_{\{x:T \mid e\}?) \mid v_{c_1 \mapsto c_2}$
 $u ::= k \mid \lambda x:T. e$

Typing contexts
 $\Gamma ::= \emptyset \mid \Gamma, x:T$

Figure 4.8: NAIVE syntax

4.3 A naïve coercion calculus

In this section, we define a naïve coercion calculus, NAIVE. Its syntax is in Figure 4.8, its typing rules are in Figure 4.9 (both in Section 4.3.1), and its operational semantics are in Figure 4.12. The type soundness proof appears in Section 4.3.3; we show how to translate CAST terms into behaviorally equivalent NAIVE terms in Section 4.4.

This language adheres to a design philosophy of “simple types by default, dynamic and refinement types by coercion”. I am the first to articulate a typing philosophy for full-spectrum languages. My design philosophy, elaborated more fully in Section 4.1, has three principles. First, base values have simple types; e.g., all integers are typed at `Int`. Second, we give operations types precise enough to guarantee totality; e.g., division has a type at least as precise as $\text{Int} \rightarrow \{x:\text{Int} \mid x \neq 0\} \rightarrow \text{Int}$. I understand refinement types as being designed for protecting partial operations (the original name is due to a method for protecting partial pattern matches [30]); giving operations types that make them total means that any reasoning about runtime errors can entirely revolve around cast (here, coercion) failures. And third, values satisfy their refinement types; e.g., if $\emptyset \vdash v : \{x:T \mid e\}$, then $e[v/x] \longrightarrow^* \text{true}$. Every sound refinement type calculus has this, but ours will have it as an inversion of the typing rule. Since we are defining call-by-value (CBV) languages, this means that functions can assume their inputs actually satisfy their refinement types.

4.3.1 Syntax and typing

Most of the terms here are standard parts of the lambda calculus. The most pertinent extension here is the coercion term, $\langle c \rangle e$; I describe my language of coercions in greater detail below. Evaluation returns *results*: either a value or a failure `fail`. The term `fail` represents coercion failure. Coercion failures can occur when the predicate fails—i.e., $e_1[v/x] \longrightarrow_n^* \text{false}_{\text{id}}$ (see `F_CHECKFAIL`)—or when dynamically typed values don’t match their type—e.g., $\langle \text{Bool?} \rangle 5_{\text{Int!}}$ (see `F_TAGFAIL`). I treat `fail` as an uncatchable exception. Just like for `CAST`, values in `NAIVE` are split in two parts: *pre-values* u are the typical values of other languages: constants and lambdas; *values* v are pre-values with a stack of primitive coercions. (See below for an explanation of the different kinds of coercions.) Technically, a value is either a pre-value tagged with the identity coercion, u_{id} , or an inner value *tagged* with an extra coercion, v_d . That is, in this language every value has a list of coercions. Values are introduced in source programs with the identity coercion, u_{id} . Keeping a coercion on *every* value is a slight departure from prior formulations. Doing so is technically expedient—simplifying the structure of the language and clearly differentiating terms with pending coercions and values with tags.

The terms of the calculus are otherwise fairly unremarkable: we have variables, application, and a fixed set of built-in operations. We have two additional runtime terms. The *active check* $\langle \{x:T \mid e_1\}, e_2, v \rangle$ represents an ongoing check that the value v satisfies the predicate e_1 ; it is invariant that $e_1[v/x] \longrightarrow_n^* e_2$. The second term, `fail`, represents the uncatchable exception thrown when a check fails. Adding exception handling facilities would completely destroy any obvious formal equivalence between `NAIVE` and `EFFICIENT`, though I conjecture that an adequately instrumented semantics would be able to determine the precise point at which a `NAIVE` program diverged from its `EFFICIENT` translation.

The calculus has: simple types, where B is a base type and $T_1 \rightarrow T_2$ is the standard function type; the dynamic type `Dyn`; and refinements of both base types and type dynamic. The base refinement $\{x:B \mid e\}$ includes all constants k of type B such that $e[k_{\text{id}}/x] \longrightarrow_n^* \text{true}_{\text{id}}$. Similarly, the dynamic refinement $\{x:\text{Dyn} \mid e\}$ includes all *values* v such that v has type `Dyn` and $e[v/x] \longrightarrow_n^* \text{true}_{\text{id}}$. When defining inference rules in this fashion (e.g., `T_TAGVALREFINE`), I treat such a rule as a rule schema that expands into two separate rules. I find it useful to think of three different “domains” of types: the base domain with the types B and $\{x:B \mid e\}$; the functional domain with the types of the form $T_1 \rightarrow T_2$, with special attention paid to functions on dynamic values, of type `Dyn`→`Dyn`; and the dynamic domain with the types `Dyn` and $\{x:\text{Dyn} \mid e\}$.

Well formedness judgments of types, and contexts are defined in Figure 4.9. It is worth noting, however, that well formedness of refinements refers back to the term typing judgment. Coercion typing is in defined in Figure 4.10.

Term typing (also defined in Figure 4.9) is mostly standard. Readers should find the rules for constants (`T_CONST`), variables (`T_VAR`), functions (`T_ABS`), failure

Well formed contexts and types

$\boxed{\vdash \Gamma}$ $\boxed{\vdash T}$

$$\begin{array}{c}
\frac{}{\vdash \emptyset} \text{WF_EMPTY} \qquad \frac{\vdash \Gamma \quad \vdash T}{\vdash \Gamma, x:T} \text{WF_EXTEND} \\
\\
\frac{}{\vdash B} \text{WF_BASE} \qquad \frac{}{\vdash \text{Dyn}} \text{WF_DYN} \qquad \frac{\vdash T_1 \quad \vdash T_2}{\vdash T_1 \rightarrow T_2} \text{WF_FUN} \\
\\
\frac{x:T \vdash e : \text{Bool} \quad T = B \text{ or } T = \text{Dyn}}{\vdash \{x:T \mid e\}} \text{WF_REFINE}
\end{array}$$

Well typed terms and values

$\boxed{\Gamma \vdash u : T}$

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

$$\begin{array}{c}
\frac{\vdash \Gamma}{\Gamma \vdash k : \text{ty}(k)} \text{T_CONST} \qquad \frac{\vdash T_1 \quad \Gamma, x:T_1 \vdash e_{12} : T_2}{\Gamma \vdash \lambda x:T_1. e_{12} : T_1 \rightarrow T_2} \text{T_ABS} \\
\\
\frac{\Gamma \vdash u : T}{\Gamma \vdash u_d : T} \text{T_PREVAL} \qquad \frac{d \neq \{x:T \mid e\}? \quad d \neq \text{Fail}}{\Gamma \vdash v : T_1 \quad \vdash d : T_1 \rightsquigarrow T_2} \text{T_TAGVAL} \\
\\
\frac{\vdash \Gamma \quad \emptyset \vdash v : T \quad \vdash \{x:T \mid e\} \quad e[v/x] \rightarrow_n^* \text{true}_{\text{Id}}}{\Gamma \vdash v_{\{x:T \mid e\}?} : T} \text{T_TAGVALREFINE} \\
\\
\frac{\vdash \Gamma \quad x:T \in \Gamma}{\Gamma \vdash x : T} \text{T_VAR} \quad \frac{\vdash T \quad \vdash \Gamma}{\Gamma \vdash \text{fail} : T} \text{T_FAIL} \quad \frac{\vdash c : T_1 \rightsquigarrow T_2 \quad \Gamma \vdash e : T_1}{\Gamma \vdash \langle c \rangle e : T_2} \text{T_COERCE} \\
\\
\frac{\text{ty}(op) = T_1 \rightarrow \dots \rightarrow T_n \rightarrow T \quad \Gamma \vdash e_i : T_i}{\Gamma \vdash op(e_1, \dots, e_n) : T} \text{T_OP} \\
\\
\frac{\Gamma \vdash e_1 : (T_1 \rightarrow T_2) \quad \Gamma \vdash e_2 : T_1}{\Gamma \vdash e_1 e_2 : T_2} \text{T_APP} \\
\\
\frac{\vdash \Gamma \quad \vdash \{x:T \mid e_1\} \quad \emptyset \vdash v : T \quad \emptyset \vdash e_2 : \text{Bool} \quad e_1[v/x] \rightarrow_n^* e_2}{\Gamma \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle : \{x:T \mid e_1\}} \text{T_CHECKNAIVE}
\end{array}$$

Figure 4.9: Typing for NAIVE

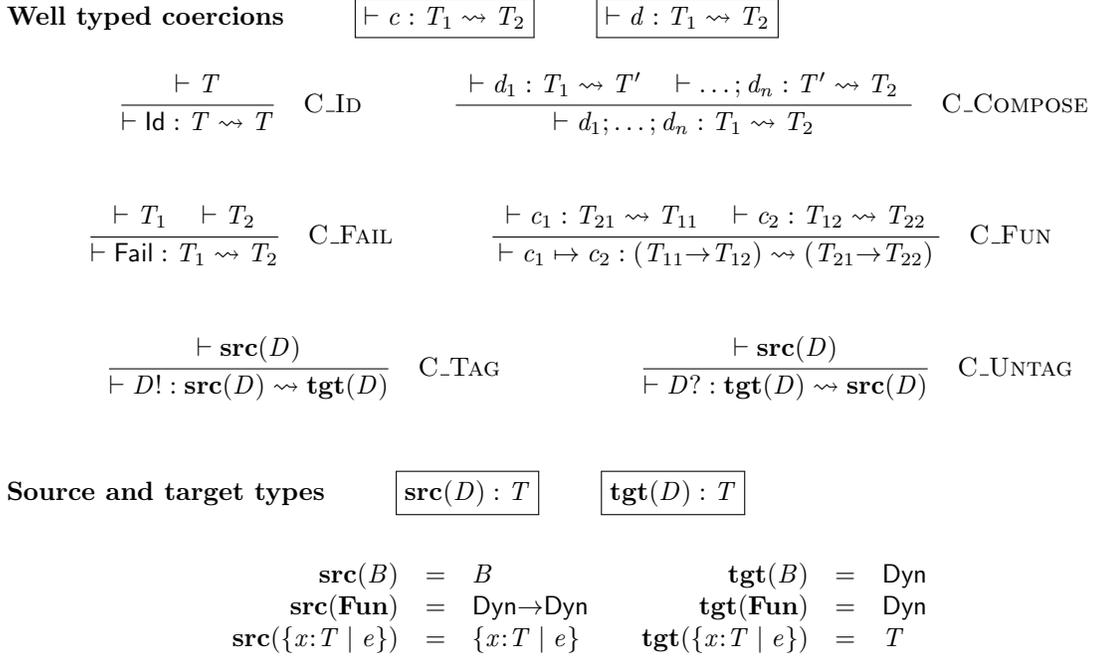


Figure 4.10: Coercion typing

(T_FAIL), application (T_APP), and built-in operations (T_OP) familiar.

Pre-values are typed by T_CONST and T_ABS; pre-values tagged with the `ld` coercion are typed as values by T_PREVAL. T_TAGVAL types values that are tagged with anything but a refinement type, for which we use a separate rule. We want all values at a refined type to satisfy their refinement—a key property and part of my philosophy of refinement types. The T_TAGVALREFINE rule ensures that values typed at a refinement type actually satisfy their refinement. In the metatheory, the typing rule for active check forms, T_CHECKNAIVE, holds onto a trace of the evaluation of the predicate. If the check succeeds, the trace can then be put directly into a T_TAGVALREFINE derivation. Naturally, none of these rules with premises concerning evaluation (T_TAGVALREFINE, T_CHECKNAIVE) are necessary for source programs—they are technicalities for the proofs of preservation (Lemma 4.3.8 below) and equivalence (Section 4.6).

The coercions are the essence of this calculus: they represent the step-by-step checks that are done to move values between dynamic, simple, and refinement types. The syntax of coercions in Figure 4.8 splits coercions into three parts: composite coercions c , primitive coercions d , and tags D . The typing rules for coercions are written in Figure 4.9; the types of primitive coercions are shown graphically in Figure 4.11. When it is clear from context whether I mean a composite or a primitive coercion, I will simply call either a “coercion”. A composite coercion c is a list of primitive coercions. I write the empty coercion—the composite coercion comprising zero primitive coercions—as `ld`. When I write c ; d or d ; c in a rule and it matches against a coercion

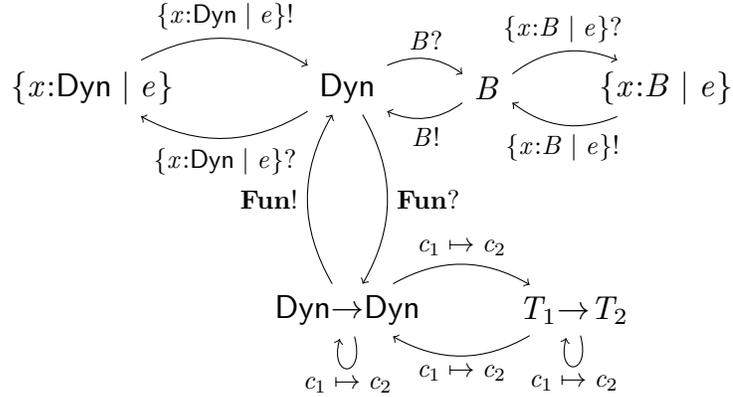


Figure 4.11: Primitive coercions

with a single primitive coercion—that is, when we match $c; B!$ against $B!$ —we let $c = \text{Id}$. This is a slight departure from earlier coercion systems; this construction avoids messing around too much with re-association of coercion composition. I compare my coercions to other formulations in related work (Chapter 5). There are four kinds of primitive coercions: failures **Fail**, tag coercions $D!$, checking coercions $D?$, and functional coercions $c_1 \mapsto c_2$. (Note that c_1 and c_2 are composite.) Finally, the tags D are a flattening of the type space: each base type B has a corresponding tag (which I also write B); functions have a single tag **Fun**. Intuitively, these are the type tags that are commonly used in dynamically typed languages. We also have refinement tags for both types of refinement, which I write the same as the corresponding types: $\{x:B \mid e\}$ for refinements of base types and $\{x:\text{Dyn} \mid e\}$ for refinements of type dynamic.

Failure coercions are present only for showing the equivalence with the space-efficient calculus; rule **F_FAIL** gives a semantics for **Fail**. (I discuss the operational semantics more fully below, in Section 4.3.2.) It is worthwhile to contrast my treatment of failure with that of Herman et al. [39]. Whereas Henglein treats mismatched tag/untag operations, such as $B!; \text{Fun?}$, as stuck, I follow Herman et al. [39] in having an explicit failure coercion, **Fail**, which leads to an uncatchable program failure, **fail**. The **F_FAIL** rule causes the program to fail when a failure coercion appears. (I carefully keep **Fail** out of the tags placed on pre-values and values.) In fact, **F_FAIL** will never apply when evaluating sensible source programs—no sane program will start with **Fail** in it, and no **NAIVE** evaluation rule generates **Fail**. Instead, failures arise in the **NAIVE** when the other rules with **FAIL** in their name fire. We include **F_FAIL** as a technicality for the soundness proof (Theorem 4.6.10) in Section 4.6. In Herman et al.’s calculus, $\langle \text{Fail} \rangle v$ is a value—the program won’t actually fail until this value reaches an elimination form. While systems with lazy error detection have been proposed [43], here $\langle \text{Fail} \rangle e$ raises a program-terminating exception immediately, before stepping e further. Eager error detection is more in line with standard error behavior, particularly other calculi where failed casts result in

blame [26, 28, 41, 67, 16, 36, 74, 78, 65, 34, 68, 8, 4, 22].

The *tagging* coercions $D!$ and *checking* coercions $D?$ fall into two groups: those which move values into type dynamic and those which deal with refinements (of either base types or type dynamic). They are typed by C_TAG and C_UNTAG ; note that $\vdash \mathbf{tgt}(D)$ for all D , but we need the $\vdash \mathbf{src}(D)$ premise for $D = \{x:T \mid e\}$ in particular, to ensure that the predicate is well typed.

The tags B and **Fun** are used to move values to and from the dynamic type \mathbf{Dyn} . The tagging coercions $B!$ and **Fun!** mark a base value (typed B) or functional value (typed $\mathbf{Dyn} \rightarrow \mathbf{Dyn}$) as having the dynamic type \mathbf{Dyn} . The checking coercions $B?$ and **Fun?** are the corresponding *untagging* coercions, taking a dynamic value and checking its tag. If the tags match, the original typed value is returned: $\langle \mathbf{Bool} \rangle \mathbf{true}_{B!} \rightarrow_n \mathbf{true}_{\mathbf{Id}}$. If the tags don't match, the program fails: $\langle \mathbf{Bool} \rangle \mathbf{5}_{\mathbf{Int}!} \rightarrow_n^* \mathbf{fail}$.

The tags $\{x:B \mid e\}$ and $\{x:\mathbf{Dyn} \mid e\}$ are used for refinements. The checking coercion $\{x:T \mid e\}?$ checks that a value v satisfies the predicate e , i.e., that $e[v/x] \rightarrow_n \mathbf{true}_{\mathbf{Id}}$; see $F_CHECKOK$ and $F_CHECKFAIL$ below. The coercion $\{x:T \mid e\}!$ is correspondingly used to ‘forget’ refinement checks.

Note that in both the dynamic and the refinement cases, the checking coercions are the ones that might fail. Given my philosophy of values starting out simply typed, the two types of coercions differ in that tagging coercions are applied first when moving from simple types to dynamic typing, but checking coercions are applied first when moving to refinements. Any simply typed value is just fine as an appropriately tagged dynamic value, but a simply typed value must be checked to see if it satisfies a refinement.

Finally, while the **Fun!** and **Fun?** coercions injecting and project functions on $\mathbf{Dyn} \rightarrow \mathbf{Dyn}$ into type \mathbf{Dyn} , there is a separate structural coercion that works on typed functions: $c_1 \mapsto c_2$, typed by the rule C_FUN . Note that the C_FUN rule is contravariant; see the F_FUN rule below.

Only certain so-called *value coercions* can appear on values: \mathbf{Id} appears on *all* values; $B!$ and **Fun!** tag values into the dynamic type; $\{x:T \mid e\}?$ marks successful refinement checks; and $c_1 \mapsto c_2$ wraps a value with pending checks, creating a function proxy. I revisit the notion of value coercions below, in the space efficient calculus of Section 4.5.

4.3.2 Operational semantics

The rules are adapted from the evaluation contexts used in Herman et al. [39]. The core rules in Figure 4.12 are standard CBV rules (F_BETA and F_OP), as are most of the congruence and exception raising rules (F_APPL , F_APPR , $F_OPINNER$, $F_APPRAISEL$, $F_APPRAISER$, $F_OPRAISE$). F_FUN and F_MERGE manage function proxies and pending coercions on the stack; the tag management rules appear in Figure 4.13. Before discussing the coercion evaluation rules, it is worth taking a moment to talk about the denotation of operators. In particular, $\llbracket op \rrbracket (v_1, \dots, v_n)$ must (a) be total when applied to correctly typed values and (b) ignore the tags on its

$$\boxed{e_1 \longrightarrow_n e_2}$$

$$\begin{array}{c}
\frac{}{(\lambda x:T. e_{12})_{\text{Id}} v_2 \longrightarrow_n e_{12}[v_2/x]} \text{ F_BETA} \qquad \frac{}{v_1 \langle c_1 \mapsto c_2 \rangle v_2 \longrightarrow_n \langle c_2 \rangle (v_1 \langle c_1 \rangle v_2)} \text{ F_FUN} \\
\\
\frac{}{op(v_1, \dots, v_n) \longrightarrow_n [[op]](v_1, \dots, v_n)} \text{ F_OP} \\
\\
\frac{e \neq \langle c' \rangle e'}{\langle \text{Fail}; c \rangle e \longrightarrow_n \text{fail}} \text{ F_FAIL} \qquad \frac{e_1 \longrightarrow_n e'_1}{e_1 e_2 \longrightarrow_n e'_1 e_2} \text{ F_APPL} \qquad \frac{e_2 \longrightarrow_n e'_2}{v_1 e_2 \longrightarrow_n v_1 e'_2} \text{ F_APPR} \\
\\
\frac{e_i \longrightarrow_n e'_i}{op(v_1, \dots, v_{i-1}, e_i, \dots, e_n) \longrightarrow_n op(v_1, \dots, v_{i-1}, e'_i, \dots, e_n)} \text{ F_OPINNER} \\
\\
\frac{e \neq \langle c' \rangle e'' \quad e \longrightarrow_n e'}{\langle c \rangle e \longrightarrow_n \langle c \rangle e'} \text{ F_COERCEINNER} \\
\\
\frac{}{\langle c_1 \rangle \langle c_2 \rangle e \longrightarrow_n \langle c_2; c_1 \rangle e} \text{ F_MERGE} \\
\\
\frac{e_2 \longrightarrow_n e'_2}{\langle \{x:T \mid e_1\}, e_2, v \rangle \longrightarrow_n \langle \{x:T \mid e_1\}, e'_2, v \rangle} \text{ F_CHECKINNER} \\
\\
\frac{}{\langle c \rangle \text{fail} \longrightarrow_n \text{fail}} \text{ F_COERCERAISE} \qquad \frac{}{op(v_1, \dots, v_{i-1}, \text{fail}, \dots, e_n) \longrightarrow_n \text{fail}} \text{ F_OPRAISE} \\
\\
\frac{}{\text{fail } e_2 \longrightarrow_n \text{fail}} \text{ F_APPRAISEL} \qquad \frac{}{v_1 \text{fail} \longrightarrow_n \text{fail}} \text{ F_APPRAISER} \\
\\
\frac{}{\langle \{x:T \mid e\}, \text{fail}, v \rangle \longrightarrow_n \text{fail}} \text{ F_CHECKRAISE}
\end{array}$$

Figure 4.12: NAIVE operational semantics (core rules)

$$\begin{array}{c}
\frac{}{\langle \text{ld} \rangle v \longrightarrow_n v} \text{ F_TAGID} \\
\\
\frac{}{\langle \{x:T \mid e\}^?; c \rangle v \longrightarrow_n \langle c \rangle \langle \{x:T \mid e\}, e[v/x], v \rangle} \text{ F_CHECK} \\
\\
\frac{}{\langle \{x:T \mid e\}, \text{true}_{\text{ld}}, v \rangle \longrightarrow_n v_{\{x:T \mid e\}^?}} \text{ F_CHECKOK} \\
\\
\frac{}{\langle \{x:T \mid e\}, \text{false}_{\text{ld}}, v \rangle \longrightarrow_n \text{fail}} \text{ F_CHECKFAIL} \\
\\
\frac{}{\langle B!; c \rangle v \longrightarrow_n \langle c \rangle v_{B!}} \text{ F_TAGB} \qquad \frac{}{\langle \mathbf{Fun}!; c \rangle v \longrightarrow_n \langle c \rangle v_{\mathbf{Fun}!}} \text{ F_TAGFUN} \\
\\
\frac{}{\langle B^?; c \rangle v_{B!} \longrightarrow_n \langle c \rangle v} \text{ F_TAGBB} \qquad \frac{}{\langle \mathbf{Fun}^?; c \rangle v_{\mathbf{Fun}!} \longrightarrow_n \langle c \rangle v} \text{ F_TAGFUNFUN} \\
\\
\frac{}{\langle \mathbf{Fun}^?; c \rangle v_{B!} \longrightarrow_n \text{fail}} \text{ F_TAGFUNFAILB} \\
\\
\frac{B \neq B'}{\langle B^?; c \rangle v_{B!} \longrightarrow_n \text{fail}} \text{ F_TAGBFAILB} \qquad \frac{}{\langle B^?; c \rangle v_{\mathbf{Fun}!} \longrightarrow_n \text{fail}} \text{ F_TAGBFAILFUN} \\
\\
\frac{}{\langle (c_1 \mapsto c_2); c \rangle v \longrightarrow_n \langle c \rangle v_{c_1 \mapsto c_2}} \text{ F_TAGFUNWRAP} \\
\\
\frac{}{\langle \{x:T \mid e\}!; c \rangle v_{\{x:T \mid e\}^?} \longrightarrow_n \langle c \rangle v} \text{ F_TAGPREDPRED}
\end{array}$$

Figure 4.13: NAIVE operational semantics (coercion rules)

inputs. This disallows some potentially useful operators—e.g., testing or projecting the tag from a dynamic value—but greatly simplifies the technicalities relating the two calculi in Section 4.6. I don’t believe that adding such tag-dependent operators would break anything in a deep way, but I omit them for simplicity’s sake. While on the subject of tags, I want to stress that the dynamic type tags are *not* erasable, since the evaluation rules depend on them. I conjecture that the refinement tags $\{x:T \mid e\}?$ are in fact erasable, but have not proven so.

Most of the rules for coercions take a term of the form $\langle d; c \rangle v$ and somehow apply the primitive coercion d to v . The rest cover more structural uses of coercions. I cover these structural rules first and then explain the “tagging” rules. `F_TAGID` (in Figure 4.13) applies when we have used up all of the primitive coercions, in which case we simply drop the coercion form.

The `F_MERGE` and `F_COERCEINNER` rules coordinate coercion merging and congruence. The `F_MERGE` rule simply concatenates two adjacent coercions. This concatenation isn’t space efficient—in the space-efficient calculus, we normalize the concatenation to a canonical coercion of bounded size. `F_COERCEINNER` steps congruently inside a coerced term—we are careful to ensure that it can only apply after `F_MERGE` has fired. Carefully staging `F_COERCEINNER` after `F_MERGE` helps maintain determinism (Lemma 4.3.4)—if we didn’t force `F_MERGE` to apply first, the number of coercions might grow out of control.³ Note that in `F_MERGE` and `F_FAIL`, the innermost term need not be a value. If I formulated the semantics as an abstract machine, we could have an explicit stack of coercions; instead, we combine them as they collide in the term.

The remaining coercion rules have `TAG` in their name and work on a term $\langle d; c \rangle v$ by combining d and v . `F_TAGB` and `F_TAGFUN` tag base values and functions (of type `Dyn`→`Dyn`) into type dynamic, using the tagging coercions $B!$ and $\mathbf{Fun}!$, respectively. `F_TAGBB` and `F_TAGFUNFUN` apply $B?$ and $\mathbf{Fun}?$ to values that have matching $B!$ and $\mathbf{Fun}!$ tags; the effect is to simply strip the tag off the value. The `F_TAGFUNFAILB`, `F_TAGBFAILFUN`, and `F_TAGBFAILB` rules cause the program to fail when it tries to strip a tag off with a checking coercion that doesn’t match.

`F_CHECK` starts the *active check* for a refinement check. An active check $\langle \{x:T \mid e_1\}, e_2, v \rangle$ is a special kind of condition: if $e_2 \rightarrow_n^* \mathbf{true}_{\text{id}}$, then it returns $v_{\{x:T \mid e_1\}?$ (rule `F_CHECKOK`); if, on the other hand, $e_2 \rightarrow_n^* \mathbf{false}_{\text{id}}$, then the active check returns `fail` (rule `F_CHECKFAIL`). Note that the typing rules for active checks make sure that $e_1[v/x] \rightarrow_n^* e_2$, i.e., that the active check is actually checking whether or not the value satisfies the predicate. The `F_TAGPREDPRED` rule is similar the untagging rules `F_TAGBB` and `F_TAGFUNFUN`, though there is no chance of failure here. Having $\{x:T \mid e\}!$ eliminate the tag $\{x:T \mid e\}?$ is reminiscent of the coercion normalization rule given in the introduction—which we said occasionally skips checks. But here in the `NAIVE`, `F_TAGPREDPRED` only applies when removing a tag from a

³Siek and Wadler [68] pointed out that Herman et al.’s evaluation contexts introduce nondeterminism; explicit congruence rules avoid this problem.

value, i.e., when the check has already been done. In the space-efficient semantics in Section 4.5, coercion normalization will actually skip checks.

It is worth emphasizing here: `F_TAGPREDPRED` is not an optimization. It is simply the natural translation of `G_CASTPREDPRED` into coercions:

$$\langle \{x:T_1 \mid e\} \Rightarrow T_2 \rangle v_{\{x:T_1|e\}?) \longrightarrow_c \langle T_1 \Rightarrow T_2 \rangle v$$

When a value is tagged as having been checked but must be cast out of the refinement type, we simply drop the tag on the floor. Rephrasing this in coercions, we find:

$$\langle \{x:T \mid e\}!; c \rangle v_{\{x:T|e\}?) \longrightarrow_n \langle c \rangle v$$

Without a rule like this, there would never be a way to untag a value;⁴ suppose $x:\{x:\text{Int} \mid x \neq 0_{\text{Id}}\}$. We need coercions c_1 and c_2 such that

$$(\lambda x:\{x:\text{Int} \mid x \neq 0_{\text{Id}}\}. \text{if } \dots \text{ then } \text{div}(5_{\text{Id}}, x) \text{ else } \text{plus}(1_{\text{Id}}, \langle c_1 \rangle x))_{\text{Id}} (\langle c_2 \rangle \dots)$$

is well typed. Whether we write them with question marks or exclamation points is irrelevant: we need *some* coercions such that c_2 and c_1 cancel each other out. I choose to have $?$ represent checking operations that may fail (testing a dynamic type tag, adding a refinement tag) and $!$ to represent operations that won't fail (adding a dynamic type tag, dropping a refinement tag).

The `F_TAGFUNWRAP` rule wraps a value in a functional coercion. The `F_FUN` rule unwinds applications of wrapped values, coercing the wrapped function's argument and result.

Readers particularly familiar with contracts will recognize that these coercions are a lower level account of the steps taken in running casts (see Belo et al. [8] for an account); alternatively, these coercions are a lower level formulation of the projections underlying contracts [25]. See Greenberg et al. [34] for a comparison.

In Figure 4.14, I translate the cast example from the introduction (Figure 4.2). It is easy to see that this calculus isn't space efficient, either: coercions can consume an unbounded amount of space. As the function evaluates, a stack of coercions builds up—here, proportional to the size of the input. Again, I highlight redexes; note that the whole term is highlighted when the outermost coercions merge. The casts of the earlier example match the coercions here, e.g. the cast $\langle \text{Dyn} \Rightarrow \text{Bool} \rangle e$ is just like the coercion $\langle \text{Bool} \rangle e$.

4.3.3 Proofs

`NAIVE` is type sound, i.e, it enjoys progress and preservation.

4.3.1 Lemma: For all D , $\vdash \mathbf{tgt}(D)$.

⁴Short of accumulating refinements of values, which is obviously not space efficient. We could leave tags on, adding a rule like `T_FORGET` in Chapter 3, at the expense of implicitly dropping a tag in other rules. That wins nothing: we still have to drop the tag, and the rules are more confusing.

```

odd 3Id
→n evenInt!→Bool? 2Id
→n ⟨Bool?⟩ (even ⟨Int!⟩ 2Id)
→n ⟨Bool?⟩ (((λx:Int. ...) Id)Int?→Bool! (2Id)Int!)
→n ⟨Bool?⟩ (⟨Bool!⟩ ((λx:Int. ...) Id (⟨Int?⟩ (2Id)Int!)))
→n ⟨Bool!; Bool?⟩ ((λx:Int. ...) Id (⟨Int?⟩ (2Id)Int!))
→n ⟨Bool!; Bool?⟩ ((λx:Int. ...) Id 2Id)
→n ⟨Bool!; Bool?⟩ odd 1Id
→n ⟨Bool!; Bool?⟩ evenInt!→Bool? 0Id
→n ⟨Bool!; Bool?⟩ (⟨Bool?⟩ (even ⟨Int!⟩ 0Id))
→n ⟨Bool?; Bool!; Bool?⟩ (even ⟨Int!⟩ 0Id)
→n ⟨Bool?; Bool!; Bool?⟩ (((λx:Int. ...) Id)Int?→Bool! (0Id)Int!)
→n ⟨Bool?; Bool!; Bool?⟩ (⟨Bool!⟩
  ((λx:Int. ...) Id (⟨Int?⟩ (0Id)Int!)))
→n ⟨Bool!; Bool?; Bool!; Bool?⟩
  ((λx:Int. ...) Id (⟨Int?⟩ (0Id)Int!))
→n ⟨Bool!; Bool?; Bool!; Bool?⟩ ((λx:Int. ...) Id 0Id)
→n ⟨Bool!; Bool?; Bool!; Bool?⟩ trueId
→n ⟨Bool?; Bool!; Bool?⟩ (trueId)Bool!
→n ⟨Bool!; Bool?⟩ trueId
→n ⟨Bool?⟩ (trueId)Bool!
→n ⟨Id⟩ trueId
→n trueId

```

Figure 4.14: NAIVE reduction

Proof: By cases on D , we use either `WF_DYN` or `WF_BASE`. □

4.3.2 Lemma [Regularity of coercion typing]: • If $\vdash c : T_1 \rightsquigarrow T_2$ then $\vdash T_1$ and $\vdash T_2$.

- If $\vdash d : T_1 \rightsquigarrow T_2$ then $\vdash T_1$ and $\vdash T_2$.

Proof: By mutual induction on the typing derivations.

(C_ID) By inversion.

(C_FAIL) By inversion.

(C_COMPOSE) By the IH.

(C_UNTAG) By assumption on the left; by Lemma 4.3.1 on the right.

(C_TAG) By Lemma 4.3.1 on the left; by assumption on the right.

(C_FUN) By the IH and `WF_FUN`.

□

4.3.3 Lemma [Regularity]: If $\Gamma \vdash e : T$ then $\vdash T$, and if $\Gamma \vdash u : T$ then $\vdash T$.

Proof: By induction on the typing derivation.

(T_CONST) By assumption.

(T_ABS) By the assumption, the IH, and `WF_FUN`.

(T_PREVAL) By the IH.

(T_TAGVAL) By regularity of coercion typing (Lemma 4.3.2).

(T_TAGVALREFINE) By assumption.

(T_VAR) By induction on $\vdash \Gamma$.

(T_FAIL) By assumption.

(T_COERCE) By regularity of coercion typing (Lemma 4.3.2).

(T_OP) By assumption on operation typing.

(T_APP) By the IH

(T_CHECKNAIVE) By assumption.

□

4.3.4 Lemma [Determinism]: If $e \longrightarrow_n e_1$ and $e \longrightarrow_n e_2$ then $e_1 = e_2$.

Proof: By induction on $e \longrightarrow_n e_1$, observing that in each case the same rule must have applied to find $e \longrightarrow_n e_2$. \square

4.3.5 Lemma [Canonical forms]: If $\emptyset \vdash v : T$, then:

- If $T = \text{Bool}$, then v is either true_{ld} or false_{ld} .
- If $T = T_1 \rightarrow T_2$, then v is either $\lambda x : T_1. e_{12\text{ld}}$ or $v'_{c_1 \mapsto c_2}$.
- If $T = \text{Dyn}$, then v is either $v'_{B!}$ or $v'_{\text{Fun}!}$.
- If $T = \{x : T' \mid e\}$, then v is $v'_{\{x : T' \mid e\}?$.

Proof: By induction on the typing derivation $\emptyset \vdash v : T$.

Proof:

(T_PREVAL) Constants have base types, so we must only consider the case where $T = \text{Bool}$; if $\text{ty}(k) = \text{Bool}$ then k is either true or false , and we are done.

(T_TAGVAL) $\emptyset \vdash v'_d : T_2$ and $\vdash d : T_1 \rightsquigarrow T_2$. It must be that d is one of $B!$ or $\text{Fun}!$ or $c_1 \mapsto c_2$. In the first two cases, we fulfill the $T_2 = \text{Dyn}$ case. In the latter case, we fulfill the arrow case.

(T_TAGVALREFINE) $\emptyset \vdash v'_{\{x : T \mid e\}?$: $\{x : T \mid e\}$, which fulfills the refinement type case. \square

4.3.6 Lemma [Progress]: If $\emptyset \vdash e : T$ then either e is a result or $e \longrightarrow_n e'$.

Proof: By induction on the typing derivation.

(T_PREVAL) u_{ld} is a result.

(T_TAGVAL) v_d is a result.

(T_TAGVALREFINE) $v_{\{x : T \mid e\}?$ is a result.

(T_VAR) Contradictory—no such typing derivation.

(T_FAIL) fail is a result.

(T_COERCE) $\emptyset \vdash \langle c \rangle e : T_2$, where $\vdash c : T_1 \rightsquigarrow T_2$ and $\emptyset \vdash e : T_1$. We go by cases on the IH for e :

($e = v$) We go by cases on c .

($c = \text{ld}$) We step by F_TAGID.

($c = \text{Fail}; c'$) We step by F_FAIL.

$(c = B!; c')$ We step by `F_TAGB`.

$(c = B?; c')$ By canonical forms (Lemma 4.3.5), v is either $v'_{B!}$ or $v'_{\mathbf{Fun}!}$. In the former case, we step by either `F_TAGBB` or `F_TAGBFAILB`, depending on whether $B = B'$. In the latter case, we step by `F_TAGBFAILFUN`.

$(c = \mathbf{Fun}!; c')$ We step by `F_TAGFUN`.

$(c = \mathbf{Fun}?; c')$ By canonical forms (Lemma 4.3.5), v is either $v'_{B!}$ or $v'_{\mathbf{Fun}!}$. In the former case, we step by `F_TAGFUNFAILB`. In the latter case, we step by `F_TAGFUNFUN`.

$(c = \{x:T \mid e\}?.; c')$ We step by `F_CHECK`.

$(c = \{x:T \mid e\}!; c')$ By canonical forms (Lemma 4.3.5), v must be of the form $v'_{\{x:T \mid e\}?.}$, so we step by `F_TAGPREDPRED`.

$(c = (c_1 \mapsto c_2); c')$ We step by `F_TAGFUNWRAP`.

$(e = \mathbf{fail})$ We step by `F_COERCEFAIL`.

$(e \longrightarrow_n e')$ If $e = \langle c' \rangle e'$, we ignore the step from the IH and instead step by `F_MERGE`. Otherwise, we step by `F_COERCEINNER`.

$(\mathbf{T_OP}) \emptyset \vdash \mathit{op}(e_1, \dots, e_n) : T$. We go by cases on the IH of each e_i , from left to right. If any of the e_i steps, we step by `F_OPINNER`. If any of e_i is `fail`, we step by `F_OPRAISE`. Finally, if all of the e_i are values, we step by `F_OP`.

$(\mathbf{T_APP})$ We have $\emptyset \vdash e_1 e_2 : T_2$, where $\emptyset \vdash e_1 : T_1 \rightarrow T_2$ and $\emptyset \vdash e_2 : T_1$. We go by cases on IH for e_1 .

$(e_1 = v_1)$ We go by cases on the IH for $\emptyset \vdash e_2 : T_1$.

$(e_2 = v_2)$ By canonical forms (Lemma 4.3.5), v_1 is either $\lambda x:T_1. e_{12\text{id}}$ or $v'_{1c_1 \mapsto c_2}$. We step by `F_BETA` or `F_FUN`, respectively.

$(e_2 = \mathbf{fail})$ We step by `F_APPRAISER`.

$(e_2 \longrightarrow_n e'_2)$ We step by `F_APPR`.

$(e_1 = \mathbf{fail})$ We step by `F_APPRAISEL`.

(otherwise) By the IH on $\emptyset \vdash e_1 : T_1 \rightarrow T_2$, we have $e_1 \longrightarrow_n e'_1$, and we step by `F_APPL`.

$(\mathbf{T_CHECKNAIVE}) \emptyset \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle : \{x:T \mid e_1\}$. We go by cases on the IH for $\emptyset \vdash e_2 : \mathbf{Bool}$.

$(e_2 = v_2)$ By canonical forms (Lemma 4.3.5), v_2 is either `trueid` or `falseid`. We step by `F_CHECKOK` or `F_CHECKFAIL`, respectively.

$(e_2 = \mathbf{fail})$ We step by `F_CHECKRAISE`.

$(e_2 \longrightarrow_n e'_2)$ We step by `F_CHECKINNER`.

□

4.3.7 Lemma [Substitution]: If $\Gamma_1, x:T, \Gamma_2 \vdash e : T'$ and $\emptyset \vdash v : T$ then $\Gamma_1, \Gamma_2 \vdash e[v/x] : T'$. Similarly, if $\Gamma_1, x:T, \Gamma_2 \vdash u : T'$ and $\emptyset \vdash v : T$ then $\Gamma_1, \Gamma_2 \vdash u[v/x] : T'$.

Proof: By induction on the typing derivation, leaving Γ_2 general.

(T_CONST) By WF_CONST.

(T_ABS) By WF_ABS and the IH.

(T_PREVAL) By T_PREVAL and the IH.

(T_TAGVAL) By T_TAGVAL and the IH.

(T_TAGVALREFINE) By T_TAGVALREFINE and the IH, noting that all terms involved are actually closed.

(T_VAR) Either by assumption (if the two variables are the same) or by T_VAR and weakening.

(T_FAIL) By T_FAIL.

(T_COERCE) By T_COERCE and the IH.

(T_OP) By T_OP and the IH.

(T_APP) By T_APP and the IH.

(T_CHECKNAIVE) By T_CHECKNAIVE, noting that all terms involved are actually closed.

□

4.3.8 Lemma [Preservation]: If $\emptyset \vdash e : T$ and $e \longrightarrow_n e'$ then $\emptyset \vdash e' : T$.

Proof: By induction on the typing derivation $\emptyset \vdash e : T$.

(T_PREVAL) Contradictory—doesn't step.

(T_TAGVAL) Contradictory—doesn't step.

(T_TAGVALREFINE) Contradictory—doesn't step.

(T_VAR) Contradictory—there is no such derivation.

(T_FAIL) Contradictory—doesn't step.

(T_COERCE) By cases on the step taken.

(F_CHECK) $\emptyset \vdash \langle \{x:T_1 \mid e\}^?; c \rangle v : T_2$ and $\langle \{x:T_1 \mid e\}^?; c \rangle v \longrightarrow_n \langle c \rangle \langle \{x:T_1 \mid e\}, e[v/x], v \rangle$. By inversion, $\vdash c : \{x:T_1 \mid e\} \rightsquigarrow T_2$. We are done by T_CHECKNAIVE and T_COERCE.

(F_TAGID) $\emptyset \vdash \langle \text{Id} \rangle v : T$ and $\langle \text{Id} \rangle v \longrightarrow_n v$. By assumption.

(F_TAGFUNFAILB) $\emptyset \vdash \langle \mathbf{Fun}^?; c \rangle v_{B!} : T$ and $\langle \mathbf{Fun}^?; c \rangle v_{B!} \longrightarrow_n \text{fail}$. By regularity of coercion typing (Lemma 4.3.2) we can apply T_FAIL.

(F_TAGFUNFUN) $\emptyset \vdash \langle \mathbf{Fun}^?; c \rangle v_{\mathbf{Fun}!} : T$ and $\langle \mathbf{Fun}^?; c \rangle v_{\mathbf{Fun}!} \longrightarrow_n \langle c \rangle v$. By inversion $\vdash c : (\text{Dyn} \rightarrow \text{Dyn}) \rightsquigarrow T$, and we are done by assumption and T_COERCE.

(F_TAGB) $\emptyset \vdash \langle B!; c \rangle v : T$ and $\langle B!; c \rangle v \longrightarrow_n \langle c \rangle v_{B!}$. By inversion, $\vdash c : \text{Dyn} \rightsquigarrow T$ and $\emptyset \vdash v : B$. We finish by T_TAGVAL and T_COERCE.

(F_TAGFUN) $\emptyset \vdash \langle \mathbf{Fun}!; c \rangle v : T$ and $\langle \mathbf{Fun}!; c \rangle v \longrightarrow_n \langle c \rangle v_{\mathbf{Fun}!}$. By inversion, $\vdash c : \text{Dyn} \rightsquigarrow T$ and $\emptyset \vdash v : \text{Dyn} \rightarrow \text{Dyn}$. We finish by T_TAGVAL and T_COERCE.

(F_TAGBB) $\emptyset \vdash \langle B^?; c \rangle v_{B!} : T$ and $\langle B^?; c \rangle v_{B!} \longrightarrow_n \langle c \rangle v$. By inversion, $\vdash c : B \rightsquigarrow T$ and $\emptyset \vdash v : B$. We finish by T_COERCE.

(F_TAGBFAILB) $\emptyset \vdash \langle B^?; c \rangle v_{B!} : T$ and $\langle B^?; c \rangle v_{B!} \longrightarrow_n \text{fail}$. By regularity of coercion typing (Lemma 4.3.2) and T_FAIL.

(F_TAGBFAILFUN) $\emptyset \vdash \langle B^?; c \rangle v_{\mathbf{Fun}!} : T$ and $\langle B^?; c \rangle v_{\mathbf{Fun}!} \longrightarrow_n \text{fail}$. By regularity of coercion typing (Lemma 4.3.2) and T_FAIL.

(F_TAGFUNWRAP) $\emptyset \vdash \langle (c_1 \mapsto c_2); c \rangle v : T$ and $\langle (c_1 \mapsto c_2); c \rangle v \longrightarrow_n \langle c \rangle v_{c_1 \mapsto c_2}$. By inversion, $\vdash c : (T_{21} \rightarrow T_{22}) \rightsquigarrow T$ and $\emptyset \vdash v : T_{11} \rightarrow T_{12}$. By T_TAGVAL and T_COERCE.

(F_TAGPREDPRED) We have:

$$\begin{aligned} \emptyset \vdash \langle \{x:T_1 \mid e\}!; c \rangle v_{\{x:T_1 \mid e\}^?} : T_2 \\ \langle \{x:T_1 \mid e\}!; c \rangle v_{\{x:T_1 \mid e\}^?} \longrightarrow_n \langle c \rangle v \end{aligned}$$

By inversion, $\vdash c : T_1 \rightsquigarrow T_2$ and $\emptyset \vdash v : T_1$. By T_COERCE.

(F_FAIL) $\emptyset \vdash \langle \text{Fail}; c \rangle e : T$ and $\langle \text{Fail}; c \rangle e \longrightarrow_n \text{fail}$. By regularity of coercion typing (Lemma 4.3.2), $\vdash T$, so we are done by T_FAIL.

(F_COERCEINNER) $\emptyset \vdash \langle c \rangle e : T_2$ and $\langle c \rangle e \longrightarrow_n \langle c \rangle e'$, where $e \longrightarrow_n e'$. By inversion, $\vdash c : T_1 \rightsquigarrow T_2$ and $\emptyset \vdash e : T_1$. By the IH on $\emptyset \vdash e : T_1$, $\emptyset \vdash e' : T_1$. We are done by T_COERCE.

(F_MERGE) $\emptyset \vdash \langle c_1 \rangle (\langle c_2 \rangle e) : T_3$ and $\langle c_1 \rangle (\langle c_2 \rangle e) \longrightarrow_n \langle c_2; c_1 \rangle e$. By inversion, $\vdash c_1 : T_2 \rightsquigarrow T_3$ and $\vdash c_2 : T_1 \rightsquigarrow T_2$ and $\emptyset \vdash e : T_1$. By C_COMPOSE, $\vdash c_2; c_1 : T_1 \rightsquigarrow T_3$, so we are done by T_COERCE.

(F_COERCERAISE) $\emptyset \vdash \langle c \rangle \text{fail} : T$ and $\langle c \rangle \text{fail} \longrightarrow_n \text{fail}$. By regularity of coercion typing (Lemma 4.3.2), we have $\vdash T$. We are done by T_FAIL.

(T_OP) We go by cases on the step taken.

(F_OP) By assumption on the denotations of operations.

(F_OPINNER) By the IH and T_OPINNER.

(F_OPRAISE) By regularity (Lemma 4.3.3) and T_FAIL.

(T_APP) We go by cases on the step taken.

(F_BETA) $\emptyset \vdash \lambda x:T_1. e_{12\text{Id}} v_2 : T_2$ and $\lambda x:T_1. e_{12\text{Id}} v_2 \longrightarrow_n e_{12}[v_2/x]$. By inversion, $\emptyset \vdash \lambda x:T_1. e_{12} : T_1 \rightarrow T_2$; by further inversion, $x:T_1 \vdash e_1 : T_2$. By substitution (Lemma 4.3.7), $\emptyset \vdash e_1[v_2/x] : T_{22}$.

(F_FUN) $\emptyset \vdash v_{1c_1 \rightarrow c_2} v_2 : T_{22}$ and $v_{1c_1 \rightarrow c_2} v_2 \longrightarrow_n \langle c_2 \rangle (v_1 (\langle c_1 \rangle v_2))$. By inversion, $\emptyset \vdash v_{1c_1 \rightarrow c_2} : T_{21} \rightarrow T_{22}$ and $\emptyset \vdash v_2 : T_{21}$; by further inversion, $\emptyset \vdash v_1 : T_{11} \rightarrow T_{12}$ and $\vdash c_1 : T_{21} \rightsquigarrow T_{11}$ and $\vdash c_2 : T_{12} \rightsquigarrow T_{22}$.

By T_COERCE, $\emptyset \vdash \langle c_1 \rangle v_2 : T_{11}$. By T_APP, $\emptyset \vdash v_1 (\langle c_1 \rangle v_2) : T_{12}$. By T_COERCE, $\emptyset \vdash \langle c_2 \rangle (v_1 (\langle c_1 \rangle v_2)) : T_{22}$, and we are done.

(F_APPL) $\emptyset \vdash e_1 e_2 : T_2$ and $e_1 e_2 \longrightarrow_n e'_1 e_2$ where $e_1 \longrightarrow_n e'_1$. By inversion, $\emptyset \vdash e_1 : T_1 \rightarrow T_2$ and $\emptyset \vdash e_2 : T_2$. By the IH on the former derivation, $\emptyset \vdash e'_1 : T_1 \rightarrow T_2$, and we are done by T_APP.

(F_APPR) $\emptyset \vdash v_1 e_2 : T_2$ and $v_1 e_2 \longrightarrow_n v_1 e'_2$ where $e_2 \longrightarrow_n e'_2$. By inversion, $\emptyset \vdash v_1 : T_1 \rightarrow T_2$ and $\emptyset \vdash e_2 : T_2$. By the IH on the latter derivation, $\emptyset \vdash e'_2 : T_2$, and we are done by T_APP.

(F_APPRAISEL) $\emptyset \vdash \text{fail } e_2 : T$. By regularity (Lemma 4.3.3), $\vdash T$, so we can apply T_FAIL and be done.

(F_APPRAISER) $\emptyset \vdash v_1 \text{fail} : T$. By regularity (Lemma 4.3.3), $\vdash T$, so we can apply T_FAIL and be done.

(T_CHECKNAIVE) We go by cases on the step taken.

(F_CHECKOK) We have:

$$\begin{aligned} \emptyset \vdash \langle \{x:T \mid e\}, \text{true}_{\text{Id}}, v \rangle &: \{x:T \mid e\} \\ \langle \{x:T \mid e\}, \text{true}_{\text{Id}}, v \rangle &\longrightarrow_n v_{\{x:T \mid e\}}? \end{aligned}$$

By inversion, $\emptyset \vdash v : T$ and $e[v/x] \longrightarrow_n^* \text{true}_{\text{Id}}$, so we are done by T_TAGVALREFINE.

(F_CHECKFAIL) We have:

$$\begin{aligned} \emptyset \vdash \langle \{x:T \mid e\}, \text{false}_{\text{Id}}, v \rangle &: \{x:T \mid e\} \\ \langle \{x:T \mid e\}, \text{false}_{\text{Id}}, v \rangle &\longrightarrow_n \text{fail} \end{aligned}$$

By inversion, $\vdash \{x:T \mid e\}$, and we are done by T_FAIL.

(F_CHECKINNER) We have:

$$\begin{aligned} \emptyset \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle &: \{x:T \mid e_1\} \\ \langle \{x:T \mid e_1\}, e_2, v \rangle &\longrightarrow_n \langle \{x:T \mid e_1\}, e'_2, v \rangle \end{aligned}$$

where $e_2 \rightarrow_n e'_2$. By inversion, we have $\vdash \{x:T \mid e_1\}$ and $\emptyset \vdash v : T$ and $\emptyset \vdash e_2 : \mathbf{Bool}$ and $e_1[v/x] \rightarrow_n^* e_2$. We now have $e_1[v/x] \rightarrow_n^* e'_2$, and $\emptyset \vdash e'_2 : \mathbf{Bool}$ by the IH, so we are done by `T_CHECKNAIVE`.

(`F_CHECKRAISE`) $\emptyset \vdash \langle \{x:T \mid e\}, \mathbf{fail}, v \rangle : \{x:T \mid e\}$ and $\langle \{x:T \mid e\}, \mathbf{fail}, v \rangle \rightarrow_n \mathbf{fail}$. By inversion, $\vdash \{x:T \mid e\}$, so we can apply `T_FAIL` and be done. □

4.3.9 Theorem [Type soundness]: If $\emptyset \vdash e : T$ then either $e \rightarrow_n^* r$ such that $\emptyset \vdash r : T$ or e diverges.

Proof: Using progress (Lemma 4.3.6) and preservation (Lemma 4.3.8). Naturally the proof is not constructive! □

Before concluding the section, I establish *coercion congruence*, a property of evaluation that will be critical in relating `NAIVE` and `EFFICIENT` in Section 4.6. We say e_1 and e_2 coterminate when (a) e_1 diverges iff e_2 diverges, and (b) $e_1 \rightarrow_n^* v$ iff $e_2 \rightarrow_n^* v$. Cotermination is obviously reflexive, symmetric, and transitive.

When I say that e diverges, we mean that for all e' such that $e \rightarrow_n^* e'$, there exists an e'' such that $e' \rightarrow_n e''$. Combined with the fact that values don't step, we can see that this definition coincides with another standard definition of divergence: e never reduces to a value, i.e., for all e' such that $e \rightarrow_n^* e'$, it is never the case that e' is a value.

4.3.10 Lemma: If $e \rightarrow_n e'$ then e and e' coterminate.

Proof: By soundness (Theorem 4.3.9), either e diverges or reduces a result.

By assumption, e takes at least one step, so it isn't a value. So if $e \rightarrow_n^* e''$ to diverge or reduce to a value, then $e \rightarrow_n e' \rightarrow_n^* e''$ by determinism (Lemma 4.3.4). So e' behaves the same way. □

4.3.11 Corollary: $\langle c_1 \rangle (\langle c_2 \rangle e)$ and $\langle c_2; c_1 \rangle e$ coterminate.

Proof: By Lemma 4.3.10 and `F_MERGE`. □

4.3.12 Lemma [Coercion divergence congruence]: If e diverges then $\langle c \rangle e$ diverges.

Proof: It is never the case that $e \rightarrow_n^* v$, and $\langle c \rangle e$ is never a value for any c or e . □

4.3.13 Lemma [Coercion congruence]: If $e \rightarrow_n^* r$ then $\langle c \rangle r$ and $\langle c \rangle e$ coterminate.

Proof: We instead prove that there exists an e' such that $\langle c \rangle r \rightarrow^* e'$ and $\langle c \rangle e \rightarrow^* e'$, which implies cotermination by way of Lemma 4.3.10. We go by induction on $e \rightarrow_n^* r$.

$(r \longrightarrow_n^0 r)$ Immediate.

$(e \longrightarrow_n e' \longrightarrow_n^* r)$ By cases on $e \longrightarrow_n e'$. Rules without coercions (the core rules `F_BETA`, `F_FUN`, `F_OP`, `F_CHECKOK`, `F_CHECKFAIL`; the congruence rules `F_APPL`, `F_APPR`, `F_OPINNER`, `F_CHECKINNER`; and the exception raising rules `F_APPRAISEL`, `F_APPRAISER`, `F_OPRAISE`, and `F_CHECKRAISE`) are relatively easy: there is no merging. These cases work by first applying `F_COERCEINNER` and the original rule to find $\langle c \rangle e \longrightarrow_n \langle c \rangle e'$, and then we are done by the IH, since $e \longrightarrow_n e'$ and $e \longrightarrow_n^* r$ implies $e' \longrightarrow_n^* r$ by determinism (Lemma 4.3.4).

The remaining cases must confront some degree of merging between coercions in e and the coercion c . The general thrust is to observe that if $e = \langle c' \rangle e \longrightarrow_n \langle c'' \rangle e'$, then $\langle c \rangle (\langle c' \rangle e) \longrightarrow_n \langle c'; c \rangle e \longrightarrow_n \langle c''; c \rangle e''$. We can then apply the IH on $\langle c'' \rangle e' \longrightarrow_n^* r$ to find an e'' such that $\langle c \rangle (\langle c'' \rangle e') \longrightarrow_n^* e''$ and $\langle c \rangle r \longrightarrow_n^* e''$. Finally, the former term steps by `F_MERGE` to $\langle c''; c \rangle e'$, which must in turn reduce to e'' —so we are done.

`(F_TAGID)` $\langle \text{ld} \rangle e \longrightarrow_n e$, so $\langle c \rangle (\langle \text{ld} \rangle e) \longrightarrow_n \langle \text{ld}; c \rangle e = \langle c \rangle e$. By the IH on $e \longrightarrow_n^* r$, we know that there exists an e' such that $\langle c \rangle e \longrightarrow_n^* e'$ and $\langle c \rangle r \longrightarrow_n^* e'$.

`(F_TAGFUNFAILB)` $\langle \mathbf{Fun?}; c' \rangle v_{0B!} \longrightarrow_n \text{fail}$, so

$$\langle c \rangle (\langle \mathbf{Fun?}; c' \rangle v_{0B!}) \longrightarrow_n \langle \mathbf{Fun?}; c'; c \rangle v_{0B!} \longrightarrow_n \text{fail}$$

by `F_MERGE` and `F_TAGFUNFAILB`.

So $e' = \text{fail}$, since $\langle c \rangle \text{fail} \longrightarrow_n \text{fail}$ by `F_COERCERAISE`.

`(F_TAGFUNFUN)` $\langle \mathbf{Fun?}; c' \rangle v_{0\mathbf{Fun}!} \longrightarrow_n \langle c' \rangle v_0$, so

$$\langle c \rangle (\langle \mathbf{Fun?}; c' \rangle v_{0\mathbf{Fun}!}) \longrightarrow_n \langle \mathbf{Fun?}; c'; c \rangle v_{0\mathbf{Fun}!} \longrightarrow_n \langle c'; c \rangle v_0$$

by `F_MERGE` and `F_TAGFUNFUN`. We know that $\langle c' \rangle v_0 \longrightarrow_n^* v$, and by the IH there exists an e' such that $\langle c \rangle (\langle c' \rangle v_0) \longrightarrow_n^* e'$ and $\langle c \rangle r \longrightarrow_n^* e'$. We conclude by applying `F_MERGE` on the former term along with Corollary 4.3.11.

`(F_TAGB)` $\langle B!; c' \rangle v_0 \longrightarrow_n \langle c \rangle v_{0B!}$, so $\langle c \rangle (\langle B!; c' \rangle v_0) \longrightarrow_n \langle B!; c'; c \rangle v_0 \longrightarrow_n \langle c'; c \rangle v_{0B!}$. Since $\langle c' \rangle v_{0B!} \longrightarrow_n^* v$, we know that there exists an e' such that $\langle c \rangle (\langle c' \rangle v_{0B!}) \longrightarrow_n^* e'$ and $\langle c \rangle r \longrightarrow_n^* e'$ by the IH. We conclude by applying `F_MERGE` on the former term along with Corollary 4.3.11.

`(F_TAGFUN)` As for `F_TAGB`.

`(F_TAGBB)` As for `F_TAGFUNFUN`.

`(F_TAGBFAILB)` As for `F_TAGFUNFAILB`.

`(F_TAGBFAILFUN)` As for `F_TAGBFAILFUN`.

`(F_TAGFUNWRAP)` $\langle (c_1 \mapsto c_2); c' \rangle v_0 \longrightarrow_n \langle c' \rangle v_{0_{c_1 \mapsto c_2}}$, so

$$\langle c \rangle (\langle (c_1 \mapsto c_2); c' \rangle v_0) \longrightarrow_n \langle (c_1 \mapsto c_2); c'; c \rangle v_0 \longrightarrow_n \langle c'; c \rangle v_{0_{c_1 \mapsto c_2}}$$

Since $\langle c' \rangle v_{0_{c_1 \mapsto c_2}} \longrightarrow_n^* v$, we know by the IH that there exists an e' such that $\langle c \rangle (\langle c' \rangle v_{0_{c_1 \mapsto c_2}}) \longrightarrow_n^* e'$ and $\langle c \rangle r \longrightarrow_n^* e'$. We conclude by applying `F_MERGE` on the former term along with Corollary 4.3.11.

(`F_TAGPREDPRED`) As for `F_TAGFUNFUN` and `F_TAGBB`.

(`F_CHECK`) $\langle \{x:T \mid e\}^?; c' \rangle v_0 \longrightarrow_n \langle c' \rangle \langle \{x:T \mid e\}, e[v_0/x], v_0 \rangle$, so

$\langle c \rangle (\langle \{x:T \mid e\}^?; c' \rangle v_0) \longrightarrow_n \langle \{x:T \mid e\}^?; c'; c \rangle v_0 \longrightarrow_n \langle c'; c \rangle \langle \{x:T \mid e\}, e[v_0/x], v_0 \rangle$

by `F_MERGE` and `F_CHECK`. Since $\langle c' \rangle \langle \{x:T \mid e\}, e[v_0/x], v_0 \rangle \longrightarrow_n^* r$, the IH gives us an e' such that $\langle c \rangle (\langle c' \rangle \langle \{x:T \mid e\}, e[v_0/x], v_0 \rangle) \longrightarrow_n^* e'$ and $\langle c \rangle r \longrightarrow_n^* e'$; we are done by applying `F_MERGE` and Corollary 4.3.11 on the left.

(`F_COERCEINNER`) $\langle c' \rangle e \longrightarrow_n \langle c' \rangle e'$, so $\langle c \rangle (\langle c' \rangle e) \longrightarrow_n \langle c'; c \rangle e \longrightarrow_n \langle c'; c \rangle e'$. Since $\langle c' \rangle e' \longrightarrow_n^* r$, the IH gives us an e'' such that $\langle c \rangle (\langle c' \rangle e') \longrightarrow_n^* e''$. This last steps immediately by `F_MERGE` to $\langle c'; c \rangle e'$, so we know that term goes to e''' , too.

(`F_MERGE`) $\langle c_1 \rangle (\langle c_2 \rangle e) \longrightarrow_n \langle c_2; c_1 \rangle e$, so

$\langle c \rangle (\langle c_1 \rangle (\langle c_2 \rangle e)) \longrightarrow_n \langle c_1; c \rangle (\langle c_2 \rangle e) \longrightarrow_n \langle c_2; c_1; c \rangle e$

We know that $\langle c_2; c_1 \rangle e \longrightarrow_n^* r$, so the IH gives us an e' such that $\langle c \rangle (\langle c_2; c_1 \rangle e) \longrightarrow_n^* e'$ and $\langle c \rangle r \longrightarrow_n^* e'$. But the former term steps by `F_MERGE` to $\langle c_2; c_1; c \rangle e$, so we are done by Corollary 4.3.11.

(`F_FAIL`) $\langle \text{Fail}; c' \rangle e \longrightarrow_n \text{fail}$, so $\langle c \rangle (\langle \text{Fail}; c' \rangle e) \longrightarrow_n \langle \text{Fail}; c'; c \rangle e \longrightarrow_n \text{fail}$. We have $e' = r = \text{fail}$, since $\langle c \rangle \text{fail} \longrightarrow_n^* \text{fail}$ by `F_COERCERAISE`.

(`F_COERCERAISE`) $\langle c' \rangle \text{fail} \longrightarrow_n \text{fail}$, so $\langle c \rangle (\langle c' \rangle \text{fail}) \longrightarrow_n \langle c'; c \rangle \text{fail} \longrightarrow_n \text{fail}$, and again $e' = r = \text{fail}$ as in `F_FAIL`, so we are done.

□

4.4 Soundness of NAIVE with regard to CAST

I use a step-indexed logical relation (defined in Figure 4.16) to show that the terms translated from `CAST` to `NAIVE` by way of the function `coerce` (defined in Figure 4.15) are behaviorally equivalent. We are forced to use step indices to account for the type structure of dynamic types, i.e., that the “universal” type `Dyn` includes (tagged values from) the types `Dyn`→`Dyn` and `B`.

The proof strategy is the same as that of Greenberg et al. [34] and Belo et al. [8]: we separately relate the contract portions of the two languages (Figure 4.17). After showing that related contract checks (here, casts and coercions) behave in related ways on related values, we can show that `CAST` terms are related to their `NAIVE` translations.

This proof is a simpler version of the proof relating `NAIVE` and `EFFICIENT` in Section 4.6.

Translating casts

$\boxed{\text{coerce}(T_1, T_2) : T}$

$$\begin{aligned}
\text{coerce}(T, T) &= \text{Id} \\
&\text{when } T \neq T_1 \rightarrow T_2 \\
\text{coerce}(\text{Dyn}, T_1 \rightarrow T_2) &= \mathbf{Fun}^?; \text{coerce}(\text{Dyn} \rightarrow \text{Dyn}, T_1 \rightarrow T_2) \\
\text{coerce}(\text{Dyn}, B) &= B^? \\
\text{coerce}(B, \text{Dyn}) &= B! \\
\text{coerce}(T_1 \rightarrow T_2, \text{Dyn}) &= \text{coerce}(T_1 \rightarrow T_2, \text{Dyn} \rightarrow \text{Dyn}); \mathbf{Fun}! \\
\text{coerce}(T_{11} \rightarrow T_{12}, T_{21} \rightarrow T_{22}) &= \text{coerce}(T_{21}, T_{11}) \mapsto \text{coerce}(T_{12}, T_{22}) \\
\text{coerce}(\{x: T_1 \mid e\}, T_2) &= \{x: T_1 \mid \text{coerce}(e)\}!; \text{coerce}(T_1, T_2) \\
&\text{when } T_2 \neq \{x: T_1 \mid e\} \\
\text{coerce}(T_1, \{x: T_2 \mid e\}) &= \text{coerce}(T_1, T_2); \{x: T_2 \mid \text{coerce}(e)\}^? \\
&\text{when } T_1 \neq T_2 \text{ and } T_1 \neq \{x: T_1' \mid e'\}
\end{aligned}$$

Translating pre-values and terms

$\boxed{\text{coerce}(u) : u}$

$\boxed{\text{coerce}(e) : e}$

$$\begin{aligned}
\text{coerce}(k) &= k \\
\text{coerce}(\lambda x: T. e) &= \lambda x: \text{coerce}(T). \text{coerce}(e) \\
\text{coerce}(x) &= x \\
\text{coerce}(u_{\text{Id}}) &= \text{coerce}(u)_{\text{Id}} \\
\text{coerce}(v_{B!}) &= \text{coerce}(v)_{B!} \\
\text{coerce}(v_{\mathbf{Fun}!}) &= \text{coerce}(v)_{\mathbf{Fun}!} \\
\text{coerce}(v_{\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle}) &= \text{coerce}(v)_{(\text{coerce}(T_{21}, T_{11}) \mapsto \text{coerce}(T_{12}, T_{22}))} \\
\text{coerce}(v_{\{x: T \mid e\}^?}) &= \text{coerce}(v)_{\{x: T \mid \text{coerce}(e)\}^?} \\
\text{coerce}(\langle T_1 \Rightarrow T_2 \rangle e) &= \langle \text{coerce}(T_1, T_2) \rangle \text{coerce}(e) \\
\text{coerce}(e_1 e_2) &= \text{coerce}(e_1) \text{coerce}(e_2) \\
\text{coerce}(op(e_1, \dots, e_n)) &= op(\text{coerce}(e_1), \dots, \text{coerce}(e_n)) \\
\text{coerce}(\langle \{x: T \mid e_1\}, e_2, v \rangle) &= \langle \text{coerce}(\{x: T \mid e_1\}), \text{coerce}(e_2), \text{coerce}(v) \rangle
\end{aligned}$$

Translating types and contexts

$\boxed{\text{coerce}(T) : T}$

$\boxed{\text{coerce}(\Gamma) : \Gamma}$

$$\begin{aligned}
\text{coerce}(\text{Dyn}) &= \text{Dyn} \\
\text{coerce}(B) &= B \\
\text{coerce}(T_1 \rightarrow T_2) &= \text{coerce}(T_1) \rightarrow \text{coerce}(T_2) \\
\text{coerce}(\{x: T \mid e\}) &= \{x: T \mid \text{coerce}(e)\} \\
\text{coerce}(\emptyset) &= \emptyset \\
\text{coerce}(\Gamma, x: T) &= \text{coerce}(\Gamma), x: \text{coerce}(T)
\end{aligned}$$

Figure 4.15: Translating from CAST to NAIVE

Value rules

$$\boxed{v_1 \sim^j v_2 : T}$$

$$\begin{aligned} \forall j. k_{\text{id}} \sim^j k_{\text{id}} : B &\iff \text{ty}(k) = B \\ v_{11} \sim^j v_{21} : T_1 \rightarrow T_2 &\iff \\ \forall m < j. \forall v_{12} \sim^m v_{22} : T_1. v_{11} v_{12} \simeq^m v_{21} v_{22} : T_2 & \\ \\ v_{1B!} \sim^j v_{2B!} : \text{Dyn} &\iff v_1 \sim^j v_2 : B \\ v_{1\text{Fun}!} \sim^j v_{2\text{Fun}!} : \text{Dyn} &\iff v_1 \sim^j v_2 : \text{Dyn} \rightarrow \text{Dyn} \\ \\ v_{1\{x:T|e_1\}^?} \sim^j v_{2\{x:T|e_2\}^?} : \{x:T \mid e_1\} & \\ \iff & \\ \forall m < j. v_1 \sim^m v_2 : T \wedge \{x:T \mid e_1\} \sim^m \{x:T \mid e_2\} & \end{aligned}$$

Term rules

$$\boxed{e_1 \simeq^j e_2 : T}$$

$$\begin{aligned} e_1 \simeq^j e_2 : T &\iff \\ e_1 \text{ diverges} \vee & \\ \forall m < j. e_1 \rightarrow_c^m \text{fail} \implies e_2 \rightarrow_c^* \text{fail} & \\ \wedge e_1 \rightarrow_c^m v_1 \implies e_2 \rightarrow_c^* v_2 \wedge v_1 \sim^{(j-m)} v_2 : T & \end{aligned}$$

Type rules

$$\boxed{T_1 \sim^j T_2}$$

$$\begin{aligned} B \sim^j B & \quad \text{Dyn} \sim^j \text{Dyn} \\ T_{11} \rightarrow T_{12} \sim^j T_{21} \rightarrow T_{22} & \iff T_{11} \sim^j T_{21} \wedge T_{12} \sim^j T_{22} \\ \{x:T \mid e_1\} \sim^j \{x:T \mid e_2\} & \iff \\ \forall m < j. \forall v_1 \sim^m v_2 : T. e_1[v_1/x] \simeq^m e_2[v_2/x] : \text{Bool} & \end{aligned}$$

Closing substitutions and open terms

$$\boxed{\Gamma \models^j \delta}$$

$$\boxed{\Gamma \vdash e_1 \simeq e_2 : T}$$

$$\begin{aligned} \Gamma \models^j \delta &\iff \forall x \in \text{dom}(\Gamma). \delta_1(x) \sim^j \delta_2(x) : T \\ \Gamma \vdash e_1 \simeq e_2 : T &\iff \forall j \geq 0. \forall \Gamma \models^j \delta. \delta_1(e_1) \simeq^j \delta_2(e_2) : T \end{aligned}$$

Figure 4.16: Relating CAST and NAIVE

4.4.1 Lemma [Preservation for coerce]: Assuming that no refinement tags or active checking forms are present:

1. If $\Gamma \vdash e : T$ then $\text{coerce}(\Gamma) \vdash \text{coerce}(e) : \text{coerce}(T)$.
2. If $\Gamma \vdash u : T$ then $\text{coerce}(\Gamma) \vdash \text{coerce}(u) : \text{coerce}(T)$.
3. If $\vdash T_1$ and $\vdash T_2$ and $\vdash T_1 \parallel T_2$ then $\vdash \text{coerce}(T_1, T_2) : \text{coerce}(T_1) \rightsquigarrow \text{coerce}(T_2)$.
4. If $\vdash T$ then $\vdash \text{coerce}(T)$.
5. If $\vdash \Gamma$ then $\vdash \text{coerce}(\Gamma)$.

Proof: By simultaneous induction on the typing/well formedness/similarity derivations.

(WF_EMPTY) By WF_EMPTY.

(WF_EXTEND) By WF_EXTEND, using IHs (4) and (5).

(WF_DYN) By WF_DYN.

(WF_BASE) By WF_BASE.

(WF_FUN) By WF_FUN and IH (4).

(WF_REFINE) By WF_REFINE and IH (1).

(T_CONST) By WF_CONST and IH (5).

(T_ABS) By WF_ABS and IHs (4) and (1).

(T_VAR) By WF_VAR and IH (5).

(T_PREVAL) By WF_PREVAL and IH (2).

(T_TAGB) By IH (1), $\text{coerce}(\Gamma) \vdash \text{coerce}(v) : B$. By C_TAG, $\vdash B! : B \rightsquigarrow \text{Dyn}$. We are done by T_TAGVAL.

(T_TAGFUN) By IH (1), $\text{coerce}(\Gamma) \vdash \text{coerce}(v) : \text{Dyn} \rightarrow \text{Dyn}$. By C_TAG, we have $\vdash \mathbf{Fun}! : (\text{Dyn} \rightarrow \text{Dyn}) \rightsquigarrow \text{Dyn}$. We are done by T_TAGVAL.

(T_TAGREFINE) Contradiction—we assumed that refinement tags do not appear.

(T_WRAP) By IH (1), $\text{coerce}(\Gamma) \vdash \text{coerce}(v) : \text{coerce}(T_{11} \rightarrow T_{12})$. By NAIVE regularity (Lemma 4.3.3) and IH (4), we have $\vdash \text{coerce}(T_{11} \rightarrow T_{12})$ and $\vdash \text{coerce}(T_{21} \rightarrow T_{22})$. Since we also have $\vdash T_{11} \rightarrow T_{12} \parallel T_{21} \rightarrow T_{22}$, we can apply IH (3) to find that

$$\vdash \text{coerce}(T_{11} \rightarrow T_{12}, T_{21} \rightarrow T_{22}) : \text{coerce}(T_{11} \rightarrow T_{12}) \rightsquigarrow \text{coerce}(T_{21} \rightarrow T_{22})$$

Recall that $\text{coerce}(T_{11} \rightarrow T_{12}, T_{21} \rightarrow T_{22}) = \text{coerce}(T_{21}, T_{11}) \mapsto \text{coerce}(T_{12}, T_{22})$. So by WF_TAGVAL, we are done.

(T_FAIL) By WF_FAIL and IHs (4) and (5).

(T_CAST) By WF_COERCE, NAIVE regularity (Lemma 4.3.3), and IHs (1), (4), and (3).

(T_APP) By WF_APP and IH (1).

(T_OP) By WF_OP and IH (1).

(T_CHECKCAST) Contradictory—we assumed that there were no active checks.

(P_ID) Since T can't be a function type, $\text{coerce}(T, T) = \text{ld}$, and we are done by C_ID.

(P_DYNL) If T_2 is a function, then by C_COMPOSE with C_UNTAG and IH (3). If it is B , then by C_UNTAG. Finally, if it is a refinement, then by C_UNTAG and IH (4) along with C_COMPOSE and IH (3).

(P_DYNR) By C_TAG and IH (4), C_COMPOSE and IH (3).

(P_FUN) By C_FUN and IH (3).

(P_REFINEL) By C_COMPOSE and IH (3), with C_TAG and IH (4).

(P_REFINER) By C_COMPOSE and IH (3), with C_UNTAG and IH (4).

□

4.4.2 Lemma: If $\vdash T_1 \Rightarrow T_2 \equiv c$, then for all $v_1 \sim^j v_2 : T_1$, we have $\langle T_1 \Rightarrow T_2 \rangle v_1 \simeq^j \langle \text{coerce}(T_1, T_2) \rangle v_2 : T_2$.

Proof: By induction on the derivation of $\vdash T_1 \Rightarrow T_2 \equiv c$. In all cases, we begin by letting $m < j$ be given such that $\langle T_1 \Rightarrow T_2 \rangle v_1 \xrightarrow{m}_c r_1$. (If the left-hand side diverges, we are done.)

(RC_ID) We must show that $\langle T \Rightarrow T \rangle v_1 \simeq^m \langle \text{ld} \rangle v_2 : T$, where T is not a function type. The left-hand side must step by G_CASTID to v_1 . The right-hand side steps by F_TAGID to v_2 . Since $v_1 \sim^j v_2 : T$ and $m < j$, we are done.

$$\boxed{\vdash T_1 \Rightarrow T_2 \equiv c}$$

$$\frac{T \neq T_1 \rightarrow T_2}{\vdash T \Rightarrow T \equiv \text{Id}} \text{RC_ID} \quad \frac{}{\vdash \text{Dyn} \Rightarrow B \equiv B?} \text{RC_DYNB} \quad \frac{}{\vdash B \Rightarrow \text{Dyn} \equiv B!} \text{RC_BDYN}$$

$$\frac{\vdash \text{Dyn} \rightarrow \text{Dyn} \Rightarrow T_1 \rightarrow T_2 \equiv c}{\vdash \text{Dyn} \Rightarrow T_1 \rightarrow T_2 \equiv \mathbf{Fun?}; c} \text{RC_DYNFUN} \quad \frac{\vdash T_1 \rightarrow T_2 \Rightarrow \text{Dyn} \rightarrow \text{Dyn} \equiv c}{\vdash T_1 \rightarrow T_2 \Rightarrow \text{Dyn} \equiv c; \mathbf{Fun!}} \text{RC_FUNDDYN}$$

$$\frac{\vdash T_{21} \Rightarrow T_{11} \equiv c_1 \quad \vdash T_{12} \Rightarrow T_{22} \equiv c_2}{\vdash T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \equiv c_1 \mapsto c_2} \text{RC_FUN}$$

$$\frac{T_2 \neq \{x:T_1 \mid e_1\} \quad x:T_1 \vdash e_1 \simeq e_2 : \mathbf{Bool} \quad \vdash T_1 \Rightarrow T_2 \equiv c}{\vdash \{x:T_1 \mid e_1\} \Rightarrow T_2 \equiv \{x:T_1 \mid e_2\}!; c} \text{RC_PREDPRED}$$

$$\frac{T_1 \neq \{x:T_1' \mid e'\} \quad \vdash T_1 \Rightarrow T_2 \equiv c \quad x:T_2 \vdash e_1 \simeq e_2 : \mathbf{Bool}}{\vdash T_1 \Rightarrow \{x:T_2 \mid e_1\} \equiv c; \{x:T_2 \mid e_2\}?) \text{RC_PRECHECK}$$

Figure 4.17: Relating casts and NAIVE coercions

(RC_DYNB) We must show that $\langle \text{Dyn} \Rightarrow B \rangle v_1 \simeq^m \langle B? \rangle v_2 : B$. There are two ways to have $v_1 \simeq^j v_2 : \text{Dyn}$: either they are both tagged with base type $B!$ or with as functions with $\mathbf{Fun!}$.

If they are tagged as functions, then both sides step to fail. (The left-hand side steps by G_CASTBFAILFUN , the right by F_TAGBFAILFUN .) If they are tagged with $B' \neq B$, then both sides step to fail, on the left by G_CASTBFAILB and by F_TAGBFAILB on the right.

If they are tagged with $B' = B$, then they will reduce to v_1 and v_2 respectively, by G_CASTBB on the left and F_TAGBB and F_TAGID on the right. Since we have $v_1 \simeq^j v_2 : B$, we can find $v_1 \simeq^{(j-m)} v_2 : B$ and be done.

(RC_BDYN) We must show that $\langle B \Rightarrow \text{Dyn} \rangle v_1 \simeq^m \langle B! \rangle v_2 : \text{Dyn}$. The left-hand side steps by G_CASTB and the right-hand side by F_TAGB and F_TAGID , leaving us to prove $v_{1B!} \simeq^{(j-m)} v_{2B!} : \text{Dyn}$, which we have immediately by assumption (since $m < j$).

(RC_DYNFUN) We must show that $\langle \text{Dyn} \Rightarrow T_1 \rightarrow T_2 \rangle v_1 \simeq^m \langle \mathbf{Fun?}; c \rangle v_2 : T_1 \rightarrow T_2$. First, we go by cases as for the RC_DYNB case above. Easily solving the cases when v_1 and v_2 are tagged with $B!$. So $v_1 = v'_{1\mathbf{Fun!}}$ and $v_2 = v'_{2\mathbf{Fun!}}$. The left-hand side steps by G_CASTFUNFUN and the right hand side steps by F_TAGFUNFUN , giving us $\langle \text{Dyn} \rightarrow \text{Dyn} \Rightarrow T_1 \rightarrow T_2 \rangle v'_1$ on the left and $\langle c \rangle v'_2$ on the right. We are done by the IH.

(RC_FUNDDYN) We must show that $\langle T_1 \rightarrow T_2 \Rightarrow \text{Dyn} \rangle v_1 \simeq^m \langle c; \mathbf{Fun!} \rangle v_2 : \text{Dyn}$,

where $\vdash \text{Dyn} \rightarrow \text{Dyn} \Rightarrow T_1 \rightarrow T_2 \equiv c$.

The left-hand side must step by `G_CASTFUNDDYN`, giving us $\langle \text{Dyn} \rightarrow \text{Dyn} \Rightarrow \text{Dyn} \rangle (\langle T_1 \rightarrow T_2 \Rightarrow \text{Dyn} \rightarrow \text{Dyn} \rangle v_1)$. We know that $\langle T_1 \rightarrow T_2 \Rightarrow \text{Dyn} \rightarrow \text{Dyn} \rangle v_1 \simeq^j \langle c \rangle v_2 : \text{Dyn} \rightarrow \text{Dyn}$ by the IH.

Now we consider for a moment $\langle T_1 \rightarrow T_2 \Rightarrow \text{Dyn} \rightarrow \text{Dyn} \rangle v_1 \simeq^j \langle c \rangle v_2 : \text{Dyn} \rightarrow \text{Dyn}$. If the left-hand side diverges, we are done. Similarly, if both sides reduce to `fail`, then so do our original terms, and we are done. So suppose they reduce to $v'_1 \sim^{(j-m')} v'_2 : \text{Dyn} \rightarrow \text{Dyn}$ for some $m' < j$. By coercion congruence (Lemma 4.3.13), we know that $\langle c; \mathbf{Fun}! \rangle v_2 \longrightarrow_n^* \langle \mathbf{Fun}! \rangle v'_2$

We can expand the right-hand side to $\langle \mathbf{Fun}! \rangle (\langle c \rangle v_2) \longrightarrow_n \langle c; \mathbf{Fun}! \rangle v_2$, so we can now reduce on the left and right to find $\langle \text{Dyn} \rightarrow \text{Dyn} \Rightarrow \text{Dyn} \rangle v'_1$ on the left and $\langle \mathbf{Fun}! \rangle v'_2$ on the right. The left steps by `G_CASTFUN` and the right by `F_TAGFUN` and `F_TAGID`. We then have $v'_{1\mathbf{Fun}!} \sim^{(j-m'-1)} v'_{2\mathbf{Fun}!} : \text{Dyn}$. If $j - m' - 1$ is below m , we are done with a smaller evaluation derivation. If it's larger, we are done trivially, as determinism (Lemma 4.2.1) guarantees that this is the only evaluation derivation.

(`RC_FUN`) We must show that $\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle v_1 \simeq^m \langle c_1 \mapsto c_2 \rangle v_2 : T_{21} \rightarrow T_{22}$, given that $\vdash T_{21} \Rightarrow T_{11} \equiv c_1$ and $\vdash T_{12} \Rightarrow T_{22} \equiv c_2$.

The left-hand side steps by `G_CASTFUNWRAP` to $v_1 \langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle$; the right-hand side steps by `F_TAGFUNWRAP` and `F_TAGID` to $v_{1c_1 \mapsto c_2}$.

Let $m' < m$ such that $v'_1 \sim^{m'} v'_2 : T_{21}$. We must show that

$$v_1 \langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle v'_1 \simeq^{m'} v_{2c_1 \mapsto c_2} v'_2 : T_{21} \rightarrow T_{22}$$

The left-hand side steps by `G_FUN` to $\langle T_{12} \Rightarrow T_{22} \rangle (v_1 (\langle T_{21} \Rightarrow T_{11} \rangle v'_1))$; the right-hand side steps by `F_FUN` to $\langle c_2 \rangle (v_2 (\langle c_1 \rangle v'_2))$. By the IH on $\vdash T_{21} \Rightarrow T_{11} \equiv c_1$, we can find either (a) divergence on the left, (b) failure on the left and right, or (c) a value $v''_1 \sim^j v''_2 : T_{11}$.

In this last case, we know that $v_1 v''_1 \simeq^{m''} v_2 v''_2 : T_{12}$ for all $m'' < j$, so we again find either (a) divergence on the left, (b) failure on the left and right, or (c) a value $v'''_1 \sim^j v'''_2 : T_{12}$.

In this last case, we know by the IH on $\vdash T_{12} \Rightarrow T_{22} \equiv c_2$ that $\langle T_{12} \Rightarrow T_{22} \rangle v'''_1 \simeq^j \langle c_2 \rangle v'''_2 : T_{22}$, and we are done.

(`RC_PREDPRED`) We must show $\langle \{x:T_1 \mid e_1\} \Rightarrow T_2 \rangle v_1 \simeq^m \langle \{x:T_1 \mid e_2\}!; c \rangle v_2 : T_2$, given $x:T_1 \vdash e_1 \simeq e_2 : \mathbf{Bool}$ and $\vdash T_1 \Rightarrow T_2 \equiv c$.

The only way to have $v_1 \sim^j v_2 : \{x:T_1 \mid e_1\}$ is to have both values tagged as $v'_{1\{x:T_1|e_1\}?$ and $v'_{2\{x:T_1|e_2\}?$. (Anything else would be ill typed.)

So the left-hand side steps by `G_CASTPREDPRED` and the right-hand steps by `F_TAGPREDPRED`, giving us $\langle T_1 \Rightarrow T_2 \rangle v'_1$ and $\langle c \rangle v'_2$, respectively. By assumption we have $v'_1 \sim^j v'_2 : T_1$, so we are done by the IH.

(RC_PRECHECK) We must show $\langle T_1 \Rightarrow \{x:T_2 \mid e_1\} \rangle v_1 \simeq^m \langle c; \{x:T_2 \mid e_2\} \rangle v_2 : \{x:T_2 \mid e_1\}$ given $\vdash T_1 \Rightarrow T_2 \equiv c$ and $x:T_2 \vdash e_1 \simeq e_2 : \mathbf{Bool}$.

In CAST we step by G_CASTPRECHECK, yielding $\langle T_2 \Rightarrow \{x:T_2 \mid e_1\} \rangle (\langle T_1 \Rightarrow T_2 \rangle v_1)$.

As in RC_FUNDYN, we consider the inner term for a moment. We know by the IH that $\langle T_1 \Rightarrow T_2 \rangle v_1 \simeq^j \langle c \rangle v_2 : T_2$. If the left-hand side diverges or both go to fail, we are done—our whole term will behave the same way, in the same number of steps. So we only need to consider the case where both sides reduce to $v'_1 \sim^{(j-m')} v'_2 : T_2$ for some $m' < j$. By coercion congruence (Lemma 4.3.13), we know that $\langle c; \{x:T \mid e_2\} \rangle v_2 \xrightarrow*_n \langle \{x:T \mid e_2\} \rangle v'_2$.

Now we can see that both sides reduce again: on the left, by G_CASTCHECK; on the right, by F_TAGCHECK.

Since we have $x:T_2 \vdash e_1 \simeq e_2 : \mathbf{Bool}$, we know that $e_1[v'_1/x] \simeq^j e_2[v'_2/x] : \mathbf{Bool}$, so the checks behave the same: either the left-hand side diverges, both return fail, or they return related values at type Bool—that is, they both return $\mathbf{true}_{\text{Id}}$ or $\mathbf{false}_{\text{Id}}$.

In this last case, both sides step: to $v'_{1\{x:T|e_1\}?$ and $v'_{2\{x:T|e_2\}?$ (by G_CHECKOK or F_CHECKOK and then F_TAGID), or to fail (by G_CHECKFAIL on the left and by F_CHECKFAIL and F_COERCERAISE on the right). In either case, we may or may not have enough steps left in our index. Either way we are done: if we do have enough steps, we have found a derivation to related values. If not, we are done trivially, as determinism (Lemma 4.2.1) guarantees that this is the only evaluation derivation.

□

4.4.3 Theorem [Soundness]: 1. If $\Gamma \vdash u : T$ then $\Gamma \vdash u_{\text{Id}} \simeq \mathbf{coerce}(u_{\text{Id}}) : T$.

2. If $\Gamma \vdash v : T$ then $\Gamma \vdash v \simeq \mathbf{coerce}(v) : T$.

3. If $\Gamma \vdash e : T$ then $\Gamma \vdash e \simeq \mathbf{coerce}(e) : T$.

4. If $\vdash T_1$ and $\vdash T_2$ and $\vdash T_1 \parallel T_2$ then $\vdash T_1 \Rightarrow T_2 \equiv \mathbf{coerce}(T_1, T_2)$.

5. If $\vdash T$ then $\forall j. T \sim^j \mathbf{coerce}(T)$.

Proof: By simultaneous induction on the typing/well formedness/similarity derivations.

(T_CONST) Immediate, by definition: $k \sim^j k : B$ for all j .

(T_ABS) We must show that $\Gamma \vdash (\lambda x:T_1. e)_{\text{Id}} \simeq (\lambda x:\mathbf{coerce}(T_1). \mathbf{coerce}(e))_{\text{Id}} : T_1 \rightarrow T_2$, given that $\Gamma, x:T_1 \vdash e \simeq \mathbf{coerce}(e) : T_2$ by IH (3).

Let j be given such that $\Gamma \models^j \delta$. Unfolding some definitions, we must show that $\delta_1(\lambda x:T_1. e)_{\text{id}} \sim^j \delta_2(\lambda x:\text{coerce}(T_1). \text{coerce}(e))_{\text{id}} : T_1 \rightarrow T_2$. Let $m < j$ be given such that $v_1 \sim^m v_2 : T$. We must show that

$$(\lambda x:T_1. \delta_1(e))_{\text{id}} v_1 \simeq^m (\lambda x:\text{coerce}(T_1). \delta_2(\text{coerce}(e)))_{\text{id}} v_2 : T_2$$

Stepping each side, we see we must show that $\delta_1(e[v_1/x]) \simeq^{(m-1)} \delta_2(\text{coerce}(e)[v_2/x]) : T_2$. We have $\Gamma \models^{(m-1)} \delta[v_1, v_2/x]$, so we are done by instantiating IH (3) at $m-1$.

(T_PREVAL) By IH (1).

(T_TAGB) We must show that $\Gamma \vdash v_{B!} \simeq \text{coerce}(v)_{B!} : \text{Dyn}$. Let j be given such that $\Gamma \models^j \delta$; we must show that $\delta_1(v_{B!}) \sim^j \delta_2(\text{coerce}(v)_{B!}) : \text{Dyn}$.

By IH (2), we know that $\Gamma \vdash v \simeq \text{coerce}(v) : B$; instantiating this proposition with j and δ , we are done.

(T_TAGFUN) We must show that $\Gamma \vdash v_{\text{Fun!}} \simeq \text{coerce}(v)_{\text{Fun!}} : \text{Dyn}$. Let j be given such that $\Gamma \models^j \delta$; we must show that $\delta_1(v_{\text{Fun!}}) \sim^j \delta_2(\text{coerce}(v)_{\text{Fun!}}) : \text{Dyn}$.

By IH (2), we know that $\Gamma \vdash v \simeq \text{coerce}(v) : \text{Dyn} \rightarrow \text{Dyn}$; instantiating this proposition with j and δ , we are done.

(T_TAGREFINE) We must show that $\Gamma \vdash v_{\{x:T|e\}^?} \simeq \text{coerce}(v)_{\{x:T|\text{coerce}(e)\}^?} : \{x:T | e\}$. Let j be given such that $\Gamma \models^j \delta$; we must show that

$$\delta_1(v_{\{x:T|e\}^?}) \sim^j \delta_2(\text{coerce}(v)_{\{x:T|\text{coerce}(e)\}^?}) : \{x:T | e\}$$

Let $m < j$ be given; we must show that $v \sim^m \text{coerce}(v) : T$ and $\{x:T | e\} \sim^m \{x:T | \text{coerce}(e)\}$. We find this by instantiating IH (2) and IH (5) at m and δ .

(T_WRAP) We must show that

$$\Gamma \vdash \text{coerce}(v)_{\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle} \simeq \text{coerce}(v)_{\text{coerce}(T_{21}, T_{11}) \mapsto \text{coerce}(T_{12}, T_{22})} : T_{21} \rightarrow T_{22}$$

Let j be given such that $\Gamma \models^j \delta$; we must show that

$$\delta_1(v)_{\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle} \simeq^j \delta_2(\text{coerce}(v))_{\text{coerce}(T_{21}, T_{11}) \mapsto \text{coerce}(T_{12}, T_{22})} : T_{21} \rightarrow T_{22}$$

Let $m < j$ be given such that $v_1 \sim^m v_2 : T_{21}$. We must show (after stepping each side) that

$$\begin{aligned} & \langle T_{12} \Rightarrow T_{22} \rangle (\delta_1(v) (\langle T_{21} \Rightarrow T_{11} \rangle v_1)) \simeq^{(m-1)} \\ & \langle \text{coerce}(T_{12}, T_{22}) \rangle (\delta_2(\text{coerce}(v)) (\langle \text{coerce}(T_{21}, T_{11}) \rangle v_2)) : T_{22} \end{aligned}$$

By IH (4) and Lemma 4.4.2, we know that

$$\langle T_{21} \Rightarrow T_{11} \rangle v_1 \simeq^{j'} \langle \text{coerce}(T_{21}, T_{11}) \rangle v_2 : T_{11}$$

for all j' . Instantiating this at $m - 1$, we see that either: (a) the left-hand side diverges, and so therefore does the whole term; (b) both sides reduce to **fail**, and so therefore does the whole term; or (c) both sides reduce to $v'_1 \sim^{(m-1-m')} v'_2 : T_{11}$.

In all but the last case, we are done. In the last case, we know by the IH that $\Gamma \vdash v \simeq \text{coerce}(v) : T_{11} \rightarrow T_{12}$, so we can instantiate this at $m - 1 - m'$ to find a similar set of cases. Again, we only need to consider the case where v , v'_1 and $\text{coerce}(v)$, v'_2 reduce to values $v''_1 \sim^{(m-1-m'-m'')} v''_2 : T_{12}$ in m'' steps.

We now reason as in the domain cast, using IH (4) and Lemma 4.4.2 to finally conclude that $\langle T_{12} \Rightarrow T_{22} \rangle v''_1 \simeq^{(m-1-m'-m'')} \langle \text{coerce}(T_{12}, T_{22}) \rangle v''_2 : T_{22}$.

(T_VAR) If $\Gamma \vdash x \simeq x : T$, then $\delta_1(x) \simeq^j \delta_2(x) : T$ for all j by definition.

(T_FAIL) We immediately have **fail** \simeq^j **fail** : T for all j and T .

(T_CAST) We must show that $\Gamma \vdash \langle T_1 \Rightarrow T_2 \rangle e \simeq \langle \text{coerce}(T_1, T_2) \rangle \text{coerce}(e) : T_2$. Let j be given such that $\Gamma \models^j \delta$.

By IH (3), we know that $\Gamma \vdash e \simeq \text{coerce}(e) : T_1$. We can instantiate this at j and δ to find that either: (a) the left side diverges, and so does the cast term; (b) both sides reduce to **fail**, and so do the cast/coercion terms in one more step; or (c) both sides reduce in m steps to $v_1 \sim^{(j-m)} v_2 : T_1$. By coercion congruence (Lemma 4.3.13), we know that $\langle \text{coerce}(T_1, T_2) \rangle \delta_2(\text{coerce}(e)) \rightarrow_n^* \langle \text{coerce}(T_1, T_2) \rangle v_2$.

By IH (4), we know that $\vdash T_1 \Rightarrow T_2 \equiv \text{coerce}(T_1, T_2)$. By Lemma 4.4.2, we know that $\langle T_1 \Rightarrow T_2 \rangle v_1 \simeq^j \langle \text{coerce}(T_1, T_2) \rangle v_2 : T_2$ for all j —and for $m < j$ in particular, so we are done.

(T_APP) We must show that $\Gamma \vdash e_1 e_2 \simeq \text{coerce}(e_1) \text{coerce}(e_2) : T_2$. Let j be given such that $\Gamma \models^j \delta$. We must show $\delta_1(e_1) \delta_1(e_2) \simeq^j \delta_2(\text{coerce}(e_1)) \delta_2(\text{coerce}(e_2)) : T_2$.

By IH (3), we know that $\Gamma \vdash e_1 \simeq \text{coerce}(e_1) : T_1 \rightarrow T_2$; by instantiating at j and δ , we have either: a left-side divergence (and are done); **fail** in $m < j$ steps on both sides (and are done in $m + 1$ steps); or $v_1 \sim^{(j-m)} v_2 : T_1 \rightarrow T_2$.

Again by IH (3), $\Gamma \vdash e_2 \simeq \text{coerce}(e_2) : T_1$. Instantiating at $j - m$, we again have three possibilities: left-side divergence, **fail** on both sides, or reduction to $v'_1 \sim^{(j-m-m')} v'_2 : T_1$. In the former two cases we are done. In the latter, we know by definition that $v_1 v'_1 \simeq^{(j-m-m')} v_2 v'_2 : T_2$, and we are done.

(T_OP) Using IH (3), we can reduce all of the arguments of the operation to values. Since operations are first-order and don't use dynamic types, we know that the arguments must all be base values or refinements of base values, so the related values are *equal* constants. The G_OP/F_OP will then produce identical outputs, which must also be related.

(T_CHECKCAST) We must show that

$$\Gamma \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle \simeq \langle \{x:T \mid \text{coerce}(e_1)\}, \text{coerce}(e_2), \text{coerce}(v) \rangle : \{x:T \mid e_1\}$$

Let j and $\Gamma \models^j \delta$ be given. By IH (3) instantiated at j , we know that $e_2 \simeq^j \text{coerce}(e_2) : \text{Bool}$. If this diverges, so does the whole term, and we are done. If both sides go to fail in m steps, we are done in $m + 1$ steps. If both sides go to $v_1 \simeq^{j-m} v_2 : \text{Bool}$ in m steps, then there are two possibilities.

If $v_1 = v_2 = \text{true}_{\text{ld}}$, then both sides step by G_CHECKOK/F_CHECKOK to $\delta_1(v)_{\{x:T \mid e_1\}?$ and $\delta_2(\text{coerce}(v))_{\{x:T \mid \text{coerce}(e_1)\}?$. Let $m' < j - m - 1$ be given; We already have $\delta_1(v) \sim^{m'} \delta_2(\text{coerce}(v)) : T$ by IH (2), and we are done by IH (5), finding:

$$\{x:T \mid e_1\} \sim^{m'} \{x:T \mid \text{coerce}(e_1)\}$$

(P_ID) We have $\vdash T \Rightarrow T \equiv \text{ld}$ by RC_ID immediately.

(P_DYNL) We have $\vdash \text{Dyn} \parallel T$. We go by cases on T :

($T = \text{Dyn}$) Immediate, by RC_ID.

($T = B$) By RC_DYNB.

($T = T_1 \rightarrow T_2$) We can find $\vdash \text{Dyn} \rightarrow \text{Dyn} \parallel T_1 \rightarrow T_2$ by an application of P_FUN and two applications of P_DYNL. Then by IH (4), we have $\vdash \text{Dyn} \rightarrow \text{Dyn} \Rightarrow T_1 \rightarrow T_2 \equiv \text{coerce}(\text{Dyn} \rightarrow \text{Dyn}, T_1 \rightarrow T_2)$, and then by RC_DYNFUN we have $\vdash \text{Dyn} \Rightarrow T_1 \rightarrow T_2 \equiv \text{Fun?}; \text{coerce}(\text{Dyn} \rightarrow \text{Dyn}, T_1 \rightarrow T_2)$.

($T = \{x:T' \mid e\}$) We have $\vdash \text{Dyn} \parallel T'$ by P_DYNL immediately, so by IH (4) we have $\vdash \text{Dyn} \Rightarrow T' \equiv \text{coerce}(\text{Dyn}, T')$. By IH (5), $\{x:T' \mid e\} \sim^j \{x:T' \mid \text{coerce}(e)\}$ for all j ; unfolding this definition, we can find $x:T' \vdash e \simeq \text{coerce}(e) : \text{Bool}$. So now we have $\vdash \text{Dyn} \Rightarrow \{x:T' \mid e\} \equiv \text{coerce}(\text{Dyn}, T'); \{x:B \mid \text{coerce}(e)\}?$ by RC_PRECHECK.

(P_DYNR) We have $\vdash T \parallel \text{Dyn}$. We go by cases on T :

($T = \text{Dyn}$) Immediate, by RC_ID.

($T = B$) By RC_BDYN.

($T = T_1 \rightarrow T_2$) We can find $\vdash T_1 \rightarrow T_2 \parallel \text{Dyn} \rightarrow \text{Dyn}$ by an application of P_FUN and two applications of P_DYNR. Then by IH (4), we have $\vdash T_1 \rightarrow T_2 \Rightarrow \text{Dyn} \rightarrow \text{Dyn} \equiv \text{coerce}(T_1 \rightarrow T_2, \text{Dyn} \rightarrow \text{Dyn})$, and then by RC_DYNFUN we have $\vdash T_1 \rightarrow T_2 \Rightarrow \text{Dyn} \equiv \text{coerce}(T_1 \rightarrow T_2, \text{Dyn} \rightarrow \text{Dyn}); \text{Fun!}$.

($T = \{x:T' \mid e\}$) We have $\vdash T' \parallel \text{Dyn}$ by P_DYNL immediately, so by IH (4) we have $\vdash T' \Rightarrow \text{Dyn} \equiv \text{coerce}(T', \text{Dyn})$. By IH (5), we can unfold $\{x:T' \mid e\} \sim^j \{x:T' \mid \text{coerce}(e)\}$ into $x:T' \vdash e \simeq \text{coerce}(e) : \text{Bool}$. We are now done by RC_PREDPRED.

(P_FUN) We have $\vdash T_{11} \rightarrow T_{12} \parallel T_{21} \rightarrow T_{22}$. By IH (4), we find:

$$\vdash T_{21} \Rightarrow T_{11} \equiv \text{coerce}(T_{21}, T_{11}) \vdash T_{12} \Rightarrow T_{22} \equiv \text{coerce}(T_{12}, T_{22})$$

We find $\vdash T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \equiv \text{coerce}(T_{21}, T_{11}) \mapsto \text{coerce}(T_{12}, T_{22})$ by RC_FUN.

(P_REFINEL) We have $\vdash \{x:T_1 \mid e\} \parallel T_2$; by inversion, $\vdash T_1 \parallel T_2$. If $T_2 = \{x:T_1 \mid e\}$, then we are done by RC_ID. Otherwise, we use RC_PREDPRED and IHS (4) and (5).

(P_REFINER) We have $\vdash T_1 \parallel \{x:T_2 \mid e\}$; by inversion, $\vdash T_1 \parallel T_2$. If $T_1 = \{x:T_2 \mid e\}$, then we are done by RC_ID. Otherwise, we use RC_PRECHECK and IHS (4) and (5).

(WF_DYN) By definition.

(WF_BASE) By definition.

(WF_FUN) By IH (5).

(WF_REFINE) By IH (3) and rearranging quantifiers, recalling that e_1 (and so $\text{coerce}(e_1)$) are both closed.

□

4.5 A space-efficient coercion calculus

Having developed the naïve semantics in NAIVE, I now turn to space efficiency. In this section, I define EFFICIENT, a space-efficient coercion calculus.

There are two loci of inefficiency: coercion merges and function proxies (functional coercions). When F_MERGE applies, it merely concatenates two coercions: $\langle c_1 \rangle (\langle c_2 \rangle e) \rightarrow_n \langle c_2; c_1 \rangle e$. EFFICIENT’s semantics combines c_2 and c_1 to eliminate redundant checks. To bound the number of function proxies, I ensure that coercion merging combines adjacent functional coercions, and I change the tagging scheme on values—while NAIVE allows an arbitrary stack of tags values, EFFICIENT will keep the size of value tags bounded. The solution to both of these problems lies in *canonical coercions* and the *merge* algorithm. Before I describe them below in Section 4.5.1, I discuss changes to the syntax of values and to the typing rules.

I make changes to both the syntax (Figure 4.18) and typing rules (Figure 4.19) of NAIVE from Section 4.3; I define an entirely new operational semantics for EFFICIENT in Section 4.5.2 (see Figure 4.22). Changes are marked with highlighting and/or a bullet marker, •. The proof of type soundness of this development is in Section 4.5.3; it relies on my development of canonical coercions in Section 4.5.1. Throughout the new typing rules, I assume that coercions are canonical (see below).

Types

$T ::= B \mid T_1 \rightarrow T_2 \mid \{x:B \mid e\} \mid \text{Dyn} \mid \{x:\text{Dyn} \mid e\}$
 $B ::= \text{Bool} \mid \text{Int} \mid \dots$

Coercions, primitive coercions, and type tags

$\bullet c ::= d_1; \dots; d_n \mid \text{Fail}$
 $\bullet d ::= D! \mid D? \mid c_1 \mapsto c_2$
 $D ::= B \mid \text{Fun} \mid \{x:B \mid e\} \mid \{x:\text{Dyn} \mid e\}$

Terms, values, pre-values, and results

$e ::= x \mid r \mid \text{op}(e_1, \dots, e_n) \mid e_1 e_2 \mid \langle c \rangle e \mid$
 $\quad \langle \{x:T \mid e_1\}, e_2, v \rangle$
 $r ::= v \mid \text{fail}$
 $\bullet v ::= u_c$
 $u ::= k \mid \lambda x:T. e$

Typing contexts

$\Gamma ::= \emptyset \mid \Gamma, x:T$

Figure 4.18: Updated syntax for EFFICIENT

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

$$\frac{c \neq \text{Fail} \quad c \neq c'; \{x:T \mid e\}? \quad c \text{ is canonical} \quad \Gamma \vdash u : T_1 \quad \vdash c : T_1 \rightsquigarrow T_2}{\Gamma \vdash u_c : T_2} \quad \text{T_VAL}$$
$$\frac{\emptyset \vdash u_c : T \quad \vdash \{x:T \mid e\}? : T \rightsquigarrow \{x:T \mid e\} \quad c; \{x:T \mid e\}? \text{ is canonical} \quad \vdash \Gamma \quad e[u_c/x] \longrightarrow^* \text{true}_{\text{id}}}{\Gamma \vdash u_{c; \{x:T \mid e\}?} : \{x:T \mid e\}} \quad \text{T_VALREFINE}$$
$$\frac{\vdash \Gamma \quad \vdash \{x:T \mid e_1\} \quad \emptyset \vdash v : T \quad \emptyset \vdash e_2 : \text{Bool} \quad e_1[v/x] \longrightarrow^* e_2}{\Gamma \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle : \{x:T \mid e_1\}} \quad \text{T_CHECK}$$
$$\frac{\vdash c : T_1 \rightsquigarrow T_2 \quad \Gamma \vdash e : T_1 \quad c \text{ is canonical}}{\Gamma \vdash \langle c \rangle e : T_2} \quad \text{T_COERCE}$$

Figure 4.19: Updated typing rules for EFFICIENT

In NAIVE, tagging is stacked: a value is either a pre-value tagged with `ld` or a value tagged with a single primitive coercion. In EFFICIENT, I collapse this stack: values are pre-values tagged with a composite coercion, u_c . While NAIVE typed values with stacks of coercions using `T_PREVAL`, `T_TAGVAL`, and `T_TAGVALREFINE`, we now use rules `T_VAL` and `T_VALREFINE` to type values with a single, composite coercion on them.

I also change the structure of coercions slightly, treating `Fail` as a composite coercion. I do this because coercion normalization doesn't allow composite coercions with `Fail` at the top level—such coercions are normalized to just `Fail` itself.

The changes to the typing rules aren't major: `T_VAL` and `T_VALREFINE` account for flattened values: `T_VAL` will apply to u_c unless the coercion c ends in $\{x:T \mid e\}?$, in which case the typing derivation for the value will be `T_VALREFINE` around `T_VAL`. I separate the two rules to make sure we have value inversion, as outlined in my philosophy (Section 4.1). That is, I want to ensure that if a value has a refinement check tag on it, it satisfies that refinement. The `T_CHECK` rule changes to use the space-efficient semantics, but it remains a technical rule for supporting the evaluation of programs. Finally, I change `T_COERCE` to require that coercions appearing in the program source are *canonical*. Before discussing the new evaluation rules, I discuss my space-efficient coercions and what I mean by a canonical coercion.

4.5.1 Space-efficient coercions

I define a set of *canonical coercions*, further subdivided into value coercions for constants and for functions. I list these coercions in Table 4.2 below; I prove that they are in fact *the* normal coercions for a standard set of rewrite rules given in Figure 4.20 in Lemma 4.5.5 and Lemma 4.5.7. Next, I define a set of rules for merging coercions, proving that merging two well typed canonical coercions yields a well typed canonical coercion—no bigger than the previous two combined. This is how I show space efficiency: where a naïve implementation would accumulate and discharge all checks, EFFICIENT will keep its coercions in canonical form for which we have a bounded size (see Table 4.2).

Henglein and Herman et al. define coercions with slightly different structure: for me, `ld` is notation for the empty composite coercion, but for them it is a coercion in its own right. Henglein and Herman et al.'s systems work by taking a single rewrite rule, the so-called ϕ rule:

$$D!; D? \longrightarrow \text{ld} \quad (\phi)$$

The ϕ rule stands in distinction to the ψ rule:

$$D?; D! \longrightarrow \text{ld} \quad (\psi)$$

Henglein observed that for ϕ and ψ rules, the right-hand side is “more efficient”. The ϕ rules are *efficient* and *harmless*: nothing can fail (if a value is already a D -typed value, then tagging as D into `Dyn` and then immediately untagging can't fail, so you

$$\boxed{c \longrightarrow c}$$

$$\begin{array}{l}
\text{Fail}; c \longrightarrow \text{Fail} \quad (\text{F_FAIL}) \\
c; \text{Fail} \longrightarrow \text{Fail} \\
(c_{11} \mapsto c_{12}); (c_{21} \mapsto c_{22}) \longrightarrow (c_{21}; c_{11}) \mapsto (c_{12}; c_{22}) \\
B!; B? \longrightarrow \text{Id} \quad (\text{F_TAGBB}) \quad (\phi) \\
B!; B'? \longrightarrow \text{Fail when } B \neq B' \\
\quad (\text{F_TAGBFAILB}) \\
\mathbf{Fun!}; \mathbf{Fun}? \longrightarrow \text{Id} \quad (\text{F_TAGFUNFUN}) \quad (\phi) \\
B!; \mathbf{Fun}? \longrightarrow \text{Fail} \quad (\text{F_TAGBFAILFUN}) \\
\mathbf{Fun!}; B? \longrightarrow \text{Fail} \quad (\text{F_TAGFUNFAILB}) \\
\{x:T \mid e\}?, \{x:T \mid e\}! \longrightarrow \text{Id} \quad (\text{F_TAGPREDPRED}) \quad (\psi) \\
\{x:T \mid e\}!, \{x:T \mid e\}? \longrightarrow \text{Id} \quad (\phi)
\end{array}$$

$$\frac{d_{i-1}; d_i \longrightarrow c}{d_1; \dots; d_{i-1}; d_i; \dots; d_n \longrightarrow d_1; \dots; c; \dots; d_n} \text{ COMPAT}$$

$$\frac{c_1 \longrightarrow c'_1}{d_1; \dots; (c_1 \mapsto c_2); \dots; d_n \longrightarrow d_1; \dots; (c'_1 \mapsto c_2); \dots; d_n} \text{ FUNDOM}$$

$$\frac{c_2 \longrightarrow c'_2}{d_1; \dots; (c_1 \mapsto c_2); \dots; d_n \longrightarrow d_1; \dots; (c_1 \mapsto c'_2); \dots; d_n} \text{ FUNCOD}$$

Figure 4.20: Coercion rewriting rules

might as well just run `ld`). But the right-hand side of a ψ rule is “safer”—the left-hand side of a ψ rule may fail at runtime, but the right-hand side won’t. Checking that a value of type `Dyn` holds a D -typed value and then immediately retagging it might fail, since the value may actually have a different type. Put another way, ϕ rules don’t affect the extensional behavior of the program, just the (intensional) cost of running it; ψ rules affect extensional behavior, possibly hiding errors.

Henglein and Herman et al. define term rewriting systems with ϕ rules modulo an equational theory obtained by completing the following rules with reflexivity, symmetry, transitivity, and compatibility:

$$\begin{aligned} \text{ld}; c &= c \\ c; \text{ld} &= c \\ c_{11} \mapsto c_{12}; c_{21} \mapsto c_{22} &= (c_{21}; c_{11}) \mapsto (c_{12}; c_{22}) \\ (c_1; c_2); c_3 &= c_1; (c_2; c_3) \end{aligned}$$

My system has ϕ rules for dynamic types and for manifest contracts, but it also has a ψ rule for manifest contracts. To make the technicalities simpler, I directly define a rewrite system in Figure 4.20, where I lift the two-coercion rewrite rules to composite coercions with the rules `COMPAT`, `FUNDOM`, and `FUNCOD`. Rules that are instances of ϕ or ψ rules are marked as such. When reading the rules, recall that $c = \text{ld}; c = c; \text{ld}$. I expand out the ϕ rule for clarity here (and to ease some technicalities below).

Note that many of these rules correspond to reduction rules in `NAIVE`’s operational semantics; I’ve written the reduction rule names next to such rewrite rules. The rules for predicates over base types and type dynamic are new. Just as $B?$ takes a less specific type, `Dyn`, to a more specific type B while performing a check, we have $\{x:B \mid e\}?$ take a less specific type, B to a more specific type $\{x:B \mid e\}$ while performing a check. That is, $\text{src}(D)$ is more specific than $\text{tgt}(D)$. By the same analogy, $\{x:B \mid e\}!$ takes a more specific type to a less specific one. The rules for refinements don’t have just a ϕ rule—they must have a ψ rule, too:

$$\{x:T \mid e\}?, \{x:T \mid e\}! \longrightarrow \text{ld}$$

Note that this is a *syntactic* equivalence on refinement types and their predicate terms. This may lead to some implementation issues; see Chapter 6.

To see why we need a ψ rule, consider the rule `F_TAGPREDPRED` from the naïve operational semantics (Figure 4.12): we must have a ψ rule if $\{x:T \mid e\}!$ is going to `untag` $v_{\{x:T \mid e\}?$. But if $\{x:T \mid e\}?, \{x:T \mid e\}! \longrightarrow \text{ld}$ on tags, it must also hold for coercions on the stack, if we want space efficiency. That is, the coercion

$$\{x:T \mid e_1\}?, \{x:T \mid e_1\}!, \{x:T \mid e_2\}?, \{x:T \mid e_2\}!, \dots; \{x:T \mid e_n\}?, \{x:T \mid e_n\}!$$

will accumulate checks unless each pair of e_i can be annihilated.

Space-efficient refinement checking *must* drop some checks on the floor. The ϕ rule for refinements is an optimization, unnecessary for soundness and space efficiency. Removing this rule would add one extra canonical coercion. Suppose

$\{x:T \mid e\}!; \{x:T \mid e\}?$; c was canonical. What can c be that is (a) well typed, and (b) doesn't reduce? There is no such coercion; the only new canonical coercion would be $\{x:T \mid e\}!; \{x:T \mid e\}?$. Dropping the ϕ rule doesn't affect soundness, either. In fact, having it complicates matters a tiny bit:

$$\begin{aligned}
& \langle \{x:T \mid e\} \rangle (\langle \{x:T \mid e\} \rangle v_{\{x:T|e\}?) \\
\longrightarrow_n & \langle \{x:T \mid e\}!; \{x:T \mid e\} \rangle v_{\{x:T|e\}?) \\
\longrightarrow_n & \langle \{x:T \mid e\} \rangle v \\
\longrightarrow_n & \langle \text{Id} \rangle \langle \{x:T \mid e\}, e[v/x], v \rangle \\
\longrightarrow_n^* & \langle \text{Id} \rangle \langle \{x:T \mid e\}, \text{true}_{\text{Id}}, v \rangle \quad (\text{by typing of } v_{\{x:T|e\}?) \\
\longrightarrow_n & \langle \text{Id} \rangle v_{\{x:T|e\}?) \\
\longrightarrow_n & v_{\{x:T|e\}?)
\end{aligned}$$

And in the space-efficient calculus with N_PHI for $D = \{x:T \mid e\}$, the merge equivalent of the ϕ rule:

$$\begin{aligned}
& \langle \{x:T \mid e\} \rangle (\langle \{x:T \mid e\} \rangle u_{c;\{x:T|e\}?) \\
\longrightarrow & \langle \{x:T \mid e\}!; \{x:T \mid e\} \rangle u_{c;\{x:T|e\}?) \\
\longrightarrow & \langle \text{Id} \rangle u_{c;\{x:T|e\}?) \\
\longrightarrow & u_{c;\{x:T|e\}?)
\end{aligned}$$

And without the ϕ rule:

$$\begin{aligned}
& \langle \{x:T \mid e\} \rangle (\langle \{x:T \mid e\} \rangle u_{c;\{x:T|e\}?) \\
\longrightarrow & \langle \{x:T \mid e\}!; \{x:T \mid e\} \rangle u_{c;\{x:T|e\}?) \\
\longrightarrow & \langle \{x:T \mid e\} \rangle u_c \\
\longrightarrow & \langle \text{Id} \rangle \langle \{x:T \mid e\}, e[u_c/x], u_c \rangle \\
\longrightarrow^* & \langle \text{Id} \rangle \langle \{x:T \mid e\}, e[u_c/x], u_c \rangle \quad (\text{by typing of } u_{c;\{x:T|e\}?) \\
\longrightarrow & \langle \text{Id} \rangle u_{c;\{x:T|e\}?) \\
\longrightarrow & u_{c;\{x:T|e\}?)
\end{aligned}$$

The evaluation derivation is closer to that of the naive system without the ϕ rule for refinements—including the ϕ rule merely skips rechecking the predicate. I do not do so here, but I conjecture that EFFICIENT terms are behaviorally equivalent with or without the ϕ rule, though the runtime costs will be different.

Finally, one may wonder: we have ψ rule for refinements, but no such rule for tagging into and checking from Dyn. Why not? Henglein pointed out that ψ rules don't preserve the behavior of the original program exactly, possibly hiding errors. It turns out that I *must* have a ψ rule to make manifest contracts space efficient. But gradual types don't need the ψ rule to be space efficient. Since I'd like the behavior of NAIVE and EFFICIENT to be as close as possible, I minimize the mismatch by only using ψ rules for manifest contracts, not dynamic contracts—even if it doesn't give us a stronger theorem.

The situation is neatly dual: manifest contracts must have ψ rules, but need not have ϕ rules; dynamic types must have ϕ rules, but need not have ψ rules.

4.5.1 Lemma [Preservation for rewriting]: If $\vdash c : T_1 \rightsquigarrow T_2$ and $c \longrightarrow c'$ then $\vdash c' : T_1 \rightsquigarrow T_2$.

Proof: By case analysis on the reduction step taken. \square

4.5.2 Lemma [Confluence]: Given a well typed coercion c_1 , if $c_1 \longrightarrow c_2$ and $c_1 \longrightarrow c'_2$ then $c_2 \longrightarrow^* c_3$ and $c'_2 \longrightarrow^* c_3$.

Proof: By case analysis on the critical pairs. If the reductions take place in non-overlapping parts of the term, then we simply make the appropriate, other, reduction step in both terms.

If the reduction takes place in an overlapping critical pair, we must reason more carefully. If one of the steps is any of the **ld** or **Fail** cases, we are done: simply take the appropriate step in the other term.

The remaining critical pairs are multiple functional coercions in a row and overlapping refinement coercions.

Suppose we have $c_1 = c; (c_{11} \mapsto c_{12}); (c_{21} \mapsto c_{22}); (c_{31} \mapsto c_{32}); c'$ and

$$\begin{aligned} c_1 &\longrightarrow c_2 = c; ((c_{21}; c_{11}) \mapsto (c_{12}; c_{22})); (c_{31} \mapsto c_{32}); c' \\ c_1 &\longrightarrow c'_2 = c; (c_{11} \mapsto c_{12}); ((c_{31}; c_{21}) \mapsto (c_{22}; c_{32})); c'. \end{aligned}$$

We can simply reduce each side once more to find $c; (c_{31}; c_{21}; c_{11}) \mapsto (c_{12}; c_{22}; c_{32}); c'$ on both sides.

Suppose we have $c_1 = c; \{x:T \mid e\}^?; \{x:T \mid e\}^!; \{x:T \mid e\}^?; c'$ and

$$\begin{aligned} c_1 &\longrightarrow c_2 = c; \mathbf{ld}; \{x:T \mid e\}^?; c' \\ c_1 &\longrightarrow c'_2 = c; \{x:T \mid e\}^?; \mathbf{ld}; c'. \end{aligned}$$

Note that the refinements must be the same to have a well typed critical pair where the either reduction rule could fire. In this case, we can simply reduce the **ld** on either side to find confluence. The case for $c_1 = c; \{x:T \mid e\}^!; \{x:T \mid e\}^?; \{x:T \mid e\}^!; c'$ is symmetric. \square

4.5.3 Lemma: The rewrite system of Figure 4.20 is normalizing.

Proof: We define coercion size straightforwardly, as follows:

$$\begin{aligned} \text{size}(\mathbf{ld}) = \text{size}(\mathbf{Fail}) &= 1 \\ \text{size}(D!) = \text{size}(D?) &= 1 \\ \text{size}(d_1; \dots; d_n) &= \sum \text{size}(d_i) \\ \text{size}(c_1 \mapsto c_2) &= 1 + \text{size}(c_1) + \text{size}(c_2) \end{aligned}$$

We show that if $c_1 \longrightarrow c_2$, then $\text{size}(c_1) > \text{size}(c_2)$ (by inspection). The only subtle case is $(c_{11} \mapsto c_{12}); (c_{21} \mapsto c_{22}) \longrightarrow (c_{21}; c_{11}) \mapsto (c_{12}; c_{22})$, wherein the size reduces by exactly 1, by eliminating one of the functional coercion constructors. \square

4.5.4 Lemma: The rewrite system of Figure 4.20 is strongly normalizing on well typed coercions.

Proof: This follows directly from the fact that the rewrite system is normalizing (Lemma 4.5.3) and confluent (Lemma 4.5.2). \square

4.5.5 Lemma: The canonical coercions are normal.

Proof: By inspection. \square

4.5.6 Lemma: If $d_1; \dots; d_n$ is canonical, then any prefix $d_1; \dots$ or suffix $\dots; d_n$ is also canonical.

Proof: By inspection of Figure 4.11. \square

4.5.7 Lemma: If $\vdash c : T_1 \rightsquigarrow T_2$ and c is normal, then c is canonical.

Proof: By induction on the derivation of $\vdash c : T_1 \rightsquigarrow T_2$. Consulting Table 4.2 will speed up the case analysis, since the possible typings quickly circumscribe the possibilities.

(C_ID) **ld** is canonical.

(C_FAIL) **Fail** is canonical.

(C_COMPOSE) We have $\vdash d; c : T_1 \rightsquigarrow T_2$; by inversion we have $\vdash d : T_1 \rightsquigarrow T'$ and $\vdash c : T' \rightsquigarrow T_2$. Moreover, if $d; c$ is normal, then so is c —so by the IH, we know that c is canonical.

We will consider various possibilities for c , but in all cases we may rule out **ld**—because each of the primitive coercions is canonical—and **Fail**—because $d; \text{Fail}$ is not normal. So in the following cases, we are looking at a coercion $d; c$ where c is non-empty. We go by cases on the former derivation:

(C_UNTAG) By cases on d :

($d = B?$) It must be the case that $T' = B$, so c is one of $B!$ or $B!; \{x:\text{Dyn} \mid e\}?$ or $\{x:B \mid e\}?$. In all cases, we have immediately that $d; c$ is canonical.

($d = \text{Fun}?$) It must be the case that $T' = (\text{Dyn} \rightarrow \text{Dyn})$, so c is one of **Fun!** or **Fun!**; $\{x:\text{Dyn} \mid e\}?$ or one of the $(c_1 \mapsto c_2); c'$ coercions. In all cases, $d; c$ is canonical.

($d = \{x:T \mid e\}?$) The only possibility is that $c = \{x:T \mid e\}!; c'$, but this would not be normal—a contradiction. (That is, the only canonical coercion beginning $\{x:T \mid e\}?$ is exactly $\{x:T \mid e\}?$.)

(C_TAG) By cases on d :

($d = B!$) It must be the case that $T' = \text{Dyn}$. It cannot be the case that $c = \text{Fun}!; c'$, for then $d; c$ wouldn't be normal. By the same token, we can't have that $c = B?; c'$. The only remaining possibility is that $c = \{x:\text{Dyn} \mid e\}?$, and $B!; \{x:\text{Dyn} \mid e\}?$ is canonical.

$$\boxed{c_1 * c_2 \Rightarrow c_3}$$

$$\begin{array}{c}
\frac{c_1; c_2 \text{ is canonical}}{c_1 * c_2 \Rightarrow (c_1; c_2)} \quad \text{N_CANONICAL} \\
\\
\frac{c_{21} * c_{11} \Rightarrow c_{31} \quad c_{12} * c_{22} \Rightarrow c_{32} \quad c_1 * (c_{31} \mapsto c_{32}); c_2 \Rightarrow c}{c_1; (c_{11} \mapsto c_{12}) * (c_{21} \mapsto c_{22}); c_2 \Rightarrow c} \quad \text{N_FUN} \\
\\
\frac{}{\text{Fail} * c \Rightarrow \text{Fail}} \quad \text{N_FAILL} \qquad \frac{c \neq \text{Fail}}{c * \text{Fail} \Rightarrow \text{Fail}} \quad \text{N_FAILR} \qquad \frac{c_1 * c_2 \Rightarrow c}{c_1; D! * D?; c_2 \Rightarrow c} \quad \text{N_PHI} \\
\\
\frac{B \neq B'}{c_1; B! * B'?; c_2 \Rightarrow \text{Fail}} \quad \text{N_BFAILB} \qquad \frac{}{c_1; B! * \mathbf{Fun}?; c_2 \Rightarrow c} \quad \text{N_BFAILFUN} \\
\\
\frac{}{c_1; \mathbf{Fun}! * B?; c_2 \Rightarrow \text{Fail}} \quad \text{N_FUNFAILB} \qquad \frac{c_1 * c_2 \Rightarrow c}{c_1; \{x:T \mid e\}? * \{x:T \mid e'\}; c_2 \Rightarrow c} \quad \text{N_PREDPRED}
\end{array}$$

Figure 4.21: Merging coercions

($d = \mathbf{Fun}!$) It must be the case that $T' = \text{Dyn}$. It cannot be the case that $c = \mathbf{Fun}?; c'$ or $c = B!; c'$, because then $d; c$ wouldn't be normal. We can therefore conclude that $c = \{x:\text{Dyn} \mid e\}?$, and $\mathbf{Fun}!; \{x:\text{Dyn} \mid e\}?$ is canonical.

($d = \{x:B \mid e\}!$) c is either $B!$ or $B!; \{x:\text{Dyn} \mid e'\}?$ or $\{x:B \mid e'\}?$. For the first two, we have that $d; c$ is canonical. In the last case, $d; c$ is normal iff $e \neq e'$, in which case $d; c$ is canonical.

($d = \{x:\text{Dyn} \mid e\}!$) c must come from Dyn . If it is $\{x:\text{Dyn} \mid e'\}?$, we must have $e \neq e'$ for $d; c$ to be normal, but then $d; c$ is canonical. If $c = B?; c'$, then $\{x:\text{Dyn} \mid e\}!; B?; c'$ is canonical. The same is true when $c = \mathbf{Fun}?; c'$.

(C_FUN) $d = c_1 \mapsto c_2$. It must be the case that c_1 and c_2 are normal, so by the IH they are also canonical. It can't be the case that $c = (c'_1 \mapsto c'_2); c'$ —then $d; c$ wouldn't be normal.

So the only possibility is that (depending on the types of c_1 and c_2) the coercion c could be $\mathbf{Fun}!$ or $\mathbf{Fun}!; \{x:\text{Dyn} \mid e\}?$. In both of these cases, $d; c$ is canonical. □

4.5.8 Corollary: All well typed coercions rewrite to a canonical coercion.

Proof: By strong normalization (Lemma 4.5.4), preservation (Lemma 4.5.1), and the fact well typed normal coercions are canonical (Lemma 4.5.7). □

Having developed the rewrite rules for my (somewhat relaxed) coercions, I define a merge algorithm in Figure 4.21 that takes two coercions and merges them from the

“inside out”: when merging $d_{11}; \dots; d_{1n}$ and $d_{21}; \dots; d_{2i}$, we first merge d_{1n} and d_{21} , and then d_{1n-1} and d_{22} , and so on. We will only ever merge *canonical* coercions, greatly simplifying the algorithm. We are justified in doing this by Corollary 4.5.8—every coercion reduces (i.e., is equivalent to) some normalized, canonical form. My merging relation implements the ϕ rule for all tags in `N_PHI`; since we have ψ rule only for refinements, we implement it separately in `N_PREDPRED`. We implement failure rules in `N_BFAILB`, `N_BFAILFUN`, and `N_FUNFAILB`. I relate my coercion system to others in Chapter 5.

Looking at Table 4.2, we can see why `T_VALREFINE` only needs to check the last coercion on a tagged pre-value: if $\{x:T \mid e\}?$ appears in a canonical coercion, it appears at the end. Similarly, the observation that certain coercions are *value coercions*, i.e., are the only coercions that can be applied to values, can be made based on typing: if constants are typed at simple types and lambdas are assigned functional types, then all value coercions must come from `B` or `T1→T2`.

I write $c_1 \Downarrow c_2$ (read “merge c_1 and c_2 ”) for canonical coercions c_1 and c_2 to mean the coercion c such that $c_1 * c_2 \Rightarrow c$. This notation is justified by Lemma 4.5.11, which shows that merging is an operator on canonical coercions.

4.5.9 Lemma [Preservation for merge]: If $\vdash c_1 : T_1 \rightsquigarrow T_2$ and $\vdash c_2 : T_2 \rightsquigarrow T_3$ and $c_1 * c_2 \Rightarrow c_3$ then $\vdash c_3 : T_1 \rightsquigarrow T_3$.

Proof: By induction on the derivation of $c_1 * c_2 \Rightarrow c_3$.

(`N_CANONICAL`) By induction on the length of c_1 , using `C_COMPOSE`.

(`N_FAILL`) We have `Fail` * $c_2 \Rightarrow$ `Fail`. By regularity (Lemma 4.3.2), $\vdash T_3$. By inversion we have $\vdash T_1$, so by `C_FAIL` we have \vdash `Fail` : $T_1 \rightsquigarrow T_3$.

(`N_FAILR`) We have c_1 * `Fail` \Rightarrow `Fail`. By regularity (Lemma 4.3.2), $\vdash T_1$. By inversion we have $\vdash T_3$, so by `C_FAIL` we have \vdash `Fail` : $T_1 \rightsquigarrow T_3$.

(`N_PHI`) We have $\vdash c_1; D! : T_1 \rightsquigarrow \mathbf{tgt}(D)$ and $\vdash D?; c_2 : \mathbf{tgt}(D) \rightsquigarrow T_3$. By inversion, $\vdash c_1 : T_1 \rightsquigarrow \mathbf{src}(D)$ and $\vdash c_2 : \mathbf{src}(D) \rightsquigarrow T_3$. We can now apply the IH to find that $\vdash c : T_1 \rightsquigarrow T_3$.

(`N_BFAILB`) We have $\vdash c_1; B! : T_1 \rightsquigarrow \mathbf{Dyn}$ and $\vdash B'?; c_2 : \mathbf{Dyn} \rightsquigarrow T_3$. By regularity (Lemma 4.3.2), we know that $\vdash T_1$ and $\vdash T_3$. So we can conclude by `C_FAIL` that \vdash `Fail` : $T_1 \rightsquigarrow T_3$.

(`N_BFAILFUN`) We have $\vdash c_1; B! : T_1 \rightsquigarrow \mathbf{Dyn}$ and $\vdash \mathbf{Fun}?; c_2 : \mathbf{Dyn} \rightsquigarrow T_3$. By regularity (Lemma 4.3.2), we know that $\vdash T_1$ and $\vdash T_3$. So we can conclude by `C_FAIL` that \vdash `Fail` : $T_1 \rightsquigarrow T_3$.

(`N_FUNFAILB`) We have $\vdash c_1; \mathbf{Fun}! : T_1 \rightsquigarrow \mathbf{Dyn}$ and $\vdash B?; c_2 : \mathbf{Dyn} \rightsquigarrow T_3$. By regularity (Lemma 4.3.2), we know that $\vdash T_1$ and $\vdash T_3$. So we can conclude by `C_FAIL` that \vdash `Fail` : $T_1 \rightsquigarrow T_3$.

Coercion	Type
$\text{Id} :$	$T \rightsquigarrow T$
$\text{Fail} :$	$T \rightsquigarrow T'$
$\{x:\text{Dyn} \mid e\}?$	$\text{Dyn} \rightsquigarrow \{x:\text{Dyn} \mid e\}$
$B?$	$\text{Dyn} \rightsquigarrow B$
$B?; B!$	$\text{Dyn} \rightsquigarrow \text{Dyn}$
$B?; B!; \{x:\text{Dyn} \mid e\}?$	$\text{Dyn} \rightsquigarrow \{x:\text{Dyn} \mid e\}$
$B?; \{x:B \mid e\}?$	$\text{Dyn} \rightsquigarrow \{x:B \mid e\}$
$\text{Fun}?$	$\text{Dyn} \rightsquigarrow \text{Dyn} \rightarrow \text{Dyn}$
$\text{Fun}?$; $\text{Fun}!$	$\text{Dyn} \rightsquigarrow \text{Dyn}$
$\text{Fun}?$; $\text{Fun}!$; $\{x:\text{Dyn} \mid e\}?$	$\text{Dyn} \rightsquigarrow \{x:\text{Dyn} \mid e\}$
$\text{Fun}?$; $c_1 \mapsto c_2$	$\text{Dyn} \rightsquigarrow T_{21} \rightarrow T_{22}$
$\text{Fun}?$; $c_1 \mapsto c_2$; $\text{Fun}!$	$\text{Dyn} \rightsquigarrow \text{Dyn}$
$\text{Fun}?$; $c_1 \mapsto c_2$; $\text{Fun}!$; $\{x:\text{Dyn} \mid e\}?$	$\text{Dyn} \rightsquigarrow \{x:\text{Dyn} \mid e\}$
$B!$	$B \rightsquigarrow \text{Dyn}$
$B!; \{x:\text{Dyn} \mid e\}?$	$B \rightsquigarrow \{x:\text{Dyn} \mid e\}$
$\{x:B \mid e\}?$	$B \rightsquigarrow \{x:B \mid e\}$
$\text{Fun}!$	$(\text{Dyn} \rightarrow \text{Dyn}) \rightsquigarrow \text{Dyn}$
$\text{Fun}!$; $\{x:\text{Dyn} \mid e\}?$	$(\text{Dyn} \rightarrow \text{Dyn}) \rightsquigarrow \{x:\text{Dyn} \mid e\}$
$c_1 \mapsto c_2$	$(T_{11} \rightarrow T_{12}) \rightsquigarrow (T_{21} \rightarrow T_{22})$
$c_1 \mapsto c_2$; $\text{Fun}!$	$(T_{11} \rightarrow T_{12}) \rightsquigarrow \text{Dyn}$
$c_1 \mapsto c_2$; $\text{Fun}!$; $\{x:\text{Dyn} \mid e\}?$	$(T_{11} \rightarrow T_{12}) \rightsquigarrow \{x:\text{Dyn} \mid e\}$
$\{x:\text{Dyn} \mid e\}!$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow \text{Dyn}$
$\{x:\text{Dyn} \mid e\}!; \{x:\text{Dyn} \mid e'\}?$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow \text{Dyn where } e \neq e'$
$\{x:\text{Dyn} \mid e\}!; B?$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow B$
$\{x:\text{Dyn} \mid e\}!; B?; B!$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow \text{Dyn}$
$\{x:\text{Dyn} \mid e\}!; B?; B!; \{x:\text{Dyn} \mid e'\}?$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow \{x:\text{Dyn} \mid e'\}$
$\{x:\text{Dyn} \mid e\}!; B?; \{x:B \mid e'\}?$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow \{x:B \mid e'\}$
$\{x:\text{Dyn} \mid e\}!; \text{Fun}?$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow (\text{Dyn} \rightarrow \text{Dyn})$
$\{x:\text{Dyn} \mid e\}!; \text{Fun}?$; $\text{Fun}!$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow \text{Dyn}$
$\{x:\text{Dyn} \mid e\}!; \text{Fun}?$; $\text{Fun}!$; $\{x:\text{Dyn} \mid e'\}?$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow \{x:\text{Dyn} \mid e'\}$
$\{x:\text{Dyn} \mid e\}!; \text{Fun}?$; $c_1 \mapsto c_2$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow T_{21} \rightarrow T_{22}$
$\{x:\text{Dyn} \mid e\}!; \text{Fun}?$; $c_1 \mapsto c_2$; $\text{Fun}!$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow \text{Dyn}$
$\{x:\text{Dyn} \mid e\}!; \text{Fun}?$; $c_1 \mapsto c_2$; $\text{Fun}!$; $\{x:\text{Dyn} \mid e'\}?$	$\{x:\text{Dyn} \mid e\} \rightsquigarrow \{x:\text{Dyn} \mid e'\}$
$\{x:B \mid e\}!$	$\{x:B \mid e\} \rightsquigarrow B$
$\{x:B \mid e\}!; B!$	$\{x:B \mid e\} \rightsquigarrow \text{Dyn}$
$\{x:B \mid e\}!; B!; \{x:\text{Dyn} \mid e'\}?$	$\{x:B \mid e\} \rightsquigarrow \{x:\text{Dyn} \mid e'\}$
$\{x:B \mid e\}!; \{x:B \mid e'\}?$	$\{x:B \mid e\} \rightsquigarrow \{x:B \mid e'\} \text{ where } e \neq e'$

Rows with a blue background are *value coercions*, and are the only coercions that can appear as tags on pre-values. Horizontal rules mark a change of initial primitive coercion.

Table 4.2: Canonical coercions

(N_PREDPRED) We have $\vdash c_1; \{x:T \mid e\}^? : T_1 \rightsquigarrow \{x:T \mid e\}$ and $\vdash \{x:T \mid e\}!; c_2 : \{x:T \mid e\} \rightsquigarrow T_3$, where T is either B or Dyn . By inversion, $\vdash c_1 : T_1 \rightsquigarrow T$ and $\vdash c_2 : T \rightsquigarrow T_3$. We can now apply the IH to find that $\vdash c : T_1 \rightsquigarrow T_3$.

(N_FUN) We have:

$$\begin{aligned} &\vdash c_1; (c_{11} \mapsto c_{12}) : T_1 \rightsquigarrow (T_{21} \rightarrow T_{22}) \\ &\vdash (c_{21} \mapsto c_{22}); c_2 : (T_{21} \rightarrow T_{22}) \rightsquigarrow T_3 \end{aligned}$$

By inversion, we know that:

$$\begin{aligned} \vdash c_1 : & T_1 \rightsquigarrow (T_{11} \rightarrow T_{12}) \\ \vdash c_{11} : & T_{21} \rightsquigarrow T_{11} \\ \vdash c_{12} : & T_{12} \rightsquigarrow T_{22} \\ \vdash c_{21} : & T_{31} \rightsquigarrow T_{21} \\ \vdash c_{22} : & T_{22} \rightsquigarrow T_{32} \\ \vdash c_2 : & (T_{31} \rightarrow T_{32}) \rightsquigarrow T_3 \end{aligned}$$

We know $c_{21} * c_{11} \Rightarrow c_{31}$ and $c_{12} * c_{22} \Rightarrow c_{32}$, so by the IH we have that $\vdash c_{31} : T_{31} \rightsquigarrow T_{11}$ and $\vdash c_{32} : T_{12} \rightsquigarrow T_{32}$. By C_FUN, $\vdash c_{31} \mapsto c_{32} : (T_{11} \rightarrow T_{12}) \rightsquigarrow (T_{31} \rightarrow T_{32})$. We now have enough typing to apply the IH on $c_1 * (c_{31} \mapsto c_{32}); c_2 \Rightarrow c$ and find that $\vdash c : T_1 \rightsquigarrow T_3$.

□

4.5.10 Lemma [Merge is a function]: Given canonical coercions c_1 and c_2 , if $c_1 * c_2 \Rightarrow c_3$ and $c_1 * c_2 \Rightarrow c'_3$, then $c_3 = c'_3$.

Proof: By induction on the derivation of $c_1 * c_2 \Rightarrow c_3$, observing in each case that whatever rule applied to form the derivation must be the one used to form $c_1 * c_2 \Rightarrow c'_3$. The only tricky case is N_CANONICAL. But if $c_1; c_2$ is canonical, then none of the N... rules can apply. □

4.5.11 Lemma [Merge is an operator]: Given canonical coercions $\vdash c_1 : T_1 \rightsquigarrow T_2$ and $\vdash c_2 : T_2 \rightsquigarrow T_3$, then there exists a unique canonical coercion c such that $c_1 * c_2 \Rightarrow c$.

Proof: By induction on $\text{size}(c_1) + \text{size}(c_2)$, with a long case analysis. Uniqueness is by Lemma 4.5.10; we use Lemma 4.7.1 to apply the IH when merging two functional coercions. We can use Table 4.2 and types to narrow the search.

First, we can rule out cases where either coercion is **ld** (we just get the other coercion, which is already canonical) or **Fail** (we just get **Fail**, which is canonical). Now we go by analysis on the left coercion c_1 . In many cases we will reduce out a few primitive coercions and then use the fact that prefixes and suffixes of canonical coercions are canonical (Lemma 4.5.6) to apply the IH.

$(\{x:\text{Dyn} \mid e\}?)$ The only coercions that can apply are of the form $\{x:\text{Dyn} \mid e\}!; c'_2$. In all cases, we will first apply `N_PREDPRED`, leaving us with `ld` on the left and c'_2 on the right. We will have just c'_2 , which is canonical since it is a suffix of a canonical coercion (Lemma 4.5.6).

$(B?)$ The only canonical coercions that can apply are:

1. $B!$, where $B?; B!$ is canonical;
2. $B!; \{x:\text{Dyn} \mid e\}?$, where $B?; B!; \{x:\text{Dyn} \mid e\}?$ is canonical; and
3. $\{x:B \mid e\}?$, where $B?; \{x:B \mid e\}?$ is canonical.

$(B?; B!)$ Here c_2 can be any canonical coercion from `Dyn`. The possibilities are:

1. $\{x:\text{Dyn} \mid e\}?$, where $B?; B!; \{x:\text{Dyn} \mid e\}?$ is canonical;
2. $B?; c'_2$, with $B? * c'_2 \Rightarrow c_3$ by the IH and Lemma 4.5.6, and we can then apply `N_PREDPRED`;
3. $B'?; c'_2$, where $B \neq B'$ and we have `Fail`;
4. **Fun?**; c'_2 , where we have `Fail`.

$(B?; B!; \{x:\text{Dyn} \mid e\}?)$ Here c_2 can be any canonical coercion from $\{x:\text{Dyn} \mid e\}?$. All of these are of the form $\{x:\text{Dyn} \mid e\}!; c'_2$. By the IH, $B?; B! * c'_2 \Rightarrow c$ for some canonical c ; then we can apply `N_PREDPRED`.

$(B?; \{x:B \mid e\}?)$ Here c_2 can be any canonical coercion from $\{x:B \mid e\}$. All of these are of the form $\{x:B \mid e\}!; c'_2$. By the IH, $B? * c'_2 \Rightarrow c$ for some canonical c ; then we can apply `N_PREDPRED`.

(Fun?) Here c_2 can be any canonical coercion from `Dyn`→`Dyn`, to which is either of the form **Fun!**; c'_2 or $c_1 \mapsto c_2; c'_2$. In either case, **Fun?**; c_2 is already canonical, so by `N_CANONICAL` we are done.

(Fun?; Fun!) Here c_2 can be any canonical coercion from `Dyn`. The possibilities are:

1. $\{x:\text{Dyn} \mid e\}?$, where **Fun?**; **Fun!**; $\{x:\text{Dyn} \mid e\}?$ is canonical;
2. $B?; c'_2$, where we have `Fail`;
3. **Fun?**; c'_2 , where we can step once by `N_PREDPRED` after observing that **Fun?** * $c'_2 \Rightarrow c$ for some canonical c by the IH and Lemma 4.5.6.

(Fun?; Fun!; \{x:\text{Dyn} \mid e\}?) Here c_2 comes from $\{x:\text{Dyn} \mid e\}$. These are all of the form $\{x:\text{Dyn} \mid e\}!; c'_2$. If we are to step by `N_PREDPRED`, we must consider **Fun?**; **Fun!** and c'_2 (canonical by Lemma 4.5.6). Since their combined size is smaller than our original coercions, we know by the IH that **Fun?**; **Fun!** * $c'_2 \Rightarrow c$ for some canonical c , and can step.

(**Fun?**; $c_{11} \mapsto c_{12}$) The coercion c_2 must come from a functional type that agrees with the type of $\vdash c_{11} \mapsto c_{12} : (T_{11} \rightarrow T_{12}) \rightsquigarrow (T_{21} \rightarrow T_{22})$, i.e., $\vdash c_2 : (T_{21} \rightarrow T_{22}) \rightsquigarrow T_3$. These coercions either begin **Fun!**; c'_2 or $c_{21} \mapsto c_{22}$; c'_2 . In the first case, **Fun?**; $c_{11} \mapsto c_{12}$; **Fun!**; c'_2 is canonical if **Fun!**; c'_2 is canonical. In the second case, we will step by N_PHI: we first apply the IH on the smaller coercions **Fun?** and $c_{31} \mapsto c_{32}$; c'_2 (using Lemma 4.7.1) to find a canonical c .

(**Fun?**; $c_{11} \mapsto c_{12}$; **Fun!**) Here c_2 comes from type **Dyn**. If it is of the form $B?$; c'_2 , we have **Fail**, which is canonical. If $c_2 = \{x:\text{Dyn} \mid e\}?$, then their concatenation, **Fun?**; $c_{11} \mapsto c_{12}$; **Fun!**; $\{x:\text{Dyn} \mid e\}?$, is canonical. If it is of the form **Fun?**; c'_2 , then we can step by N_PHI: we must show how **Fun?**; $c_{11} \mapsto c_{12}$ and c'_2 merge. The latter is canonical by Lemma 4.5.6. Since both are canonical and their combined size is smaller than the original pair of coercions (by Lemma 4.7.1), we can apply the IH to find a canonical c that they merge to.

(**Fun?**; $c_{11} \mapsto c_{12}$; **Fun!**; $\{x:\text{Dyn} \mid e\}?$) Here c_2 comes from $\{x:\text{Dyn} \mid e\}?$, so it must be a canonical coercion of the form $\{x:\text{Dyn} \mid e\}!$; c'_2 . We can combine them by N_PREDPRED, using Lemma 4.5.6 and the IH.

($B!$) In this case, c_2 comes from **Dyn**. If it is $\{x:\text{Dyn} \mid e\}?$, we are done— $B!$; $\{x:\text{Dyn} \mid e\}?$ is canonical. If it begins **Fun!**; c'_2 , we get **Fail**, which is canonical. If it begins $B?$; c'_2 , we get c'_2 , which is canonical by Lemma 4.5.6.

($B!$; $\{x:\text{Dyn} \mid e\}?$) Here $c_2 = \{x:\text{Dyn} \mid e\}!$; c'_2 , for canonical c'_2 (Lemma 4.5.6). We can apply the IH to find $B! * c'_2 \Rightarrow c$ for some canonical c , stepping by N_PREDPRED.

($\{x:B \mid e\}?$) Here $c_2 = \{x:B \mid e\}!$; c'_2 , for canonical c'_2 (Lemma 4.5.6). By N_PREDPRED and N_CANONICAL, we find that $\{x:B \mid e\}! * \{x:B \mid e\}!; c'_2 \Rightarrow c'_2$.

(**Fun!**) In this case, c_2 comes from **Dyn**. If it is $\{x:\text{Dyn} \mid e\}?$, we are done; **Fun!**; $\{x:\text{Dyn} \mid e\}?$ is canonical. If it begins $B!$; c'_2 , we get **Fail**, which is canonical. If it begins **Fun?**; c'_2 , we get c'_2 , which is canonical by Lemma 4.5.6.

(**Fun!**; $\{x:\text{Dyn} \mid e\}?$) Here $c_2 = \{x:\text{Dyn} \mid e\}!$; c'_2 , with c'_2 canonical (Lemma 4.5.6). We can apply the IH to find **Fun!** * $c'_2 \Rightarrow c$ for some canonical c .

($c_{11} \mapsto c_{12}$) The coercion c_2 must come from a functional type that agrees with the type of $\vdash c_{11} \mapsto c_{12} : (T_{11} \rightarrow T_{12}) \rightsquigarrow (T_{21} \rightarrow T_{22})$, i.e., $\vdash c_2 : (T_{21} \rightarrow T_{22}) \rightsquigarrow T_3$. These coercions either begin **Fun!**; c'_2 or $c_{21} \mapsto c_{22}$; c'_2 . In the first case, $c_{11} \mapsto c_{12}$; **Fun!**; c'_2 is canonical if **Fun!**; c'_2 is canonical.

In the second case, we use the IH to find that $c_{21} * c_{11} \Rightarrow c_{31}$ and $c_{12} * c_{22} \Rightarrow c_{32}$. If $c_{21} \mapsto c_{22}$; c'_2 was canonical, so must be $c_{31} \mapsto c_{32}$; c'_2 , which is what we are left with after stepping by N_FUN and by N_CANONICAL.

$(c_{11} \mapsto c_{12}; \mathbf{Fun}!)$ In this case, c_2 comes from **Dyn**. If it is $\{x:\mathbf{Dyn} \mid e\}?$, we are done; $\mathbf{Fun}!; \{x:\mathbf{Dyn} \mid e\}?$ is canonical. If it begins $B!; c'_2$, we get **Fail**, which is canonical. If it begins $\mathbf{Fun}?; c'_2$, we can apply the IH (since c'_2 is canonical by Lemma 4.5.6) to show that $c_{11} \mapsto c_{12} * c'_2 \Rightarrow c$ for some canonical c , stepping by **N_PHI**.

$(c_{11} \mapsto c_{12}; \mathbf{Fun}!; \{x:\mathbf{Dyn} \mid e\}?)$ Here $c_2 = \{x:\mathbf{Dyn} \mid e\}!; c'_2$, where c'_2 is canonical by Lemma 4.5.6. We can apply the IH to find $c_{11} \mapsto c_{12}; \mathbf{Fun}! * c'_2 \Rightarrow c$ for some canonical c , stepping by **N_PREDPRED**.

$(\{x:\mathbf{Dyn} \mid e\}!)$ Here c_2 must come from **Dyn**.

If it is $\{x:\mathbf{Dyn} \mid e\}?$, then we step by **N_PHI** and are left with **Id**. If it is $\{x:\mathbf{Dyn} \mid e'\}?$ for $e \neq e'$, then $\{x:\mathbf{Dyn} \mid e\}!; \{x:\mathbf{Dyn} \mid e'\}?$ is canonical.

If it is $B!; c'_2$ or $\mathbf{Fun}!; c'_2$, then again $\{x:\mathbf{Dyn} \mid e\}!; c_2$ is canonical.

$(\{x:\mathbf{Dyn} \mid e\}!; \{x:\mathbf{Dyn} \mid e'\}?)$ The coercion c_2 must be of the form $\{x:\mathbf{Dyn} \mid e''\}!; c'_2$, where c'_2 is canonical by Lemma 4.5.6. So by the IH, $\{x:\mathbf{Dyn} \mid e\}! * c'_2 \Rightarrow c$ for some canonical c , and we can step by **N_PREDPRED**.

$(\{x:\mathbf{Dyn} \mid e\}!; B?)$ Here c_2 must be either $B!; c'_2$ or $\{x:B \mid e'\}?$. In either case, the concatenation of the two is canonical.

$(\{x:\mathbf{Dyn} \mid e\}!; B?; B!)$ In this case, c_2 comes from **Dyn**. If it is $\{x:\mathbf{Dyn} \mid e'\}?$, we are done— $\{x:\mathbf{Dyn} \mid e\}!; B?; B!; \{x:\mathbf{Dyn} \mid e'\}?$ is canonical. If it begins $\mathbf{Fun}!; c'_2$, we get **Fail**, which is canonical. If it begins $B?; c'_2$, we can apply the IH to find $\{x:\mathbf{Dyn} \mid e\}!; B? * c'_2 \Rightarrow c$, since c'_2 is canonical by Lemma 4.5.6. We can then apply **N_PHI**.

$(\{x:\mathbf{Dyn} \mid e\}!; B?; B!; \{x:\mathbf{Dyn} \mid e'\}?)$ Here $c_2 = \{x:\mathbf{Dyn} \mid e'\}!; c'_2$, where c'_2 is canonical by Lemma 4.5.6. We can apply the IH to find $\{x:\mathbf{Dyn} \mid e\}!; B?; B! * c'_2 \Rightarrow c$ for some canonical c , stepping by **N_PREDPRED**.

$(\{x:\mathbf{Dyn} \mid e\}!; B?; \{x:B \mid e'\}?)$ Here $c_2 = \{x:B \mid e'\}!; c'_2$, where c'_2 is canonical by Lemma 4.5.6. We can apply the IH to find $\{x:\mathbf{Dyn} \mid e\}!; B? * c'_2 \Rightarrow c$ for some canonical c , stepping by **N_PREDPRED**.

$(\{x:\mathbf{Dyn} \mid e\}!; \mathbf{Fun}?)$ Here c_2 can be any canonical coercion from **Dyn**→**Dyn**, i.e., of the form $\mathbf{Fun}!; c'_2$ or $c_1 \mapsto c_2; c'_2$ where c'_2 is canonical. In either case, $\{x:\mathbf{Dyn} \mid e\}!; \mathbf{Fun}?; c_2$ is already canonical, so by **N_CANONICAL** we are done.

$(\{x:\mathbf{Dyn} \mid e\}!; \mathbf{Fun}?; \mathbf{Fun}!)$ In this case, c_2 comes from **Dyn**. If it is $\{x:\mathbf{Dyn} \mid e\}?$, we are done— $\{x:\mathbf{Dyn} \mid e\}!; \mathbf{Fun}?; \mathbf{Fun}!; \{x:\mathbf{Dyn} \mid e\}?$ is canonical. If it begins $B!; c'_2$, we get **Fail**, which is canonical. If it begins $\mathbf{Fun}?; c'_2$, we can apply the IH (since c'_2 is canonical by Lemma 4.5.6) to show that $\{x:\mathbf{Dyn} \mid e\}!; \mathbf{Fun}? * c'_2 \Rightarrow c$ for some canonical c , stepping by **N_PHI**.

$(\{x:\text{Dyn} \mid e\}!; \mathbf{Fun}^?; \mathbf{Fun}!; \{x:\text{Dyn} \mid e'\}?)$ Here $c_2 = \{x:\text{Dyn} \mid e'\}!; c'_2$, where c'_2 is canonical by Lemma 4.5.6. We can apply the IH to find $\{x:\text{Dyn} \mid e\}!; \mathbf{Fun}^?; \mathbf{Fun}! * c'_2 \Rightarrow c$ for some canonical c , stepping by `N_PREDPRED`.

$(\{x:\text{Dyn} \mid e\}!; \mathbf{Fun}^?; c_{11} \mapsto c_{12})$ The coercion c_2 must come from a functional type that agrees with the type of $\vdash c_{11} \mapsto c_{12} : (T_{11} \rightarrow T_{12}) \rightsquigarrow (T_{21} \rightarrow T_{22})$, i.e., $\vdash c_2 : (T_{21} \rightarrow T_{22}) \rightsquigarrow T_3$. These coercions either begin $\mathbf{Fun}!; c'_2$ or $c_{21} \mapsto c_{22}; c'_2$. In the first case, $c_{11} \mapsto c_{12}; \mathbf{Fun}!; c'_2$ is canonical if $\mathbf{Fun}!; c'_2$ is canonical.

In the second case, the IH gives us that $c_{21} * c_{11} \Rightarrow c_{31}$ and $c_{12} * c_{22} \Rightarrow c_{32}$ (using Lemma 4.7.1 for the size argument). If $c_{21} \mapsto c_{22}; c'_2$ was canonical, so must be $(c_{31} \mapsto c_{32}); c'_2$. By the IH again, $\{x:\text{Dyn} \mid e\}!; \mathbf{Fun}^? * (c_{31} \mapsto c_{32}); c'_2 \Rightarrow c$ for some canonical c , so we can step by `N_FUN`.

$(\{x:\text{Dyn} \mid e\}!; \mathbf{Fun}^?; c_{11} \mapsto c_{12}; \mathbf{Fun}!)$ In this case, c_2 comes from `Dyn`. If it is $\{x:\text{Dyn} \mid e'\}?$, we are done— $\{x:\text{Dyn} \mid e\}!; \mathbf{Fun}^?; c_{11} \mapsto c_{12}; \mathbf{Fun}!; \{x:\text{Dyn} \mid e'\}?$ is canonical. If it begins $B!; c'_2$, we get `Fail`, which is canonical. If it begins $\mathbf{Fun}^?; c'_2$, we can apply the IH (since c'_2 is canonical by Lemma 4.5.6) to show that $\{x:\text{Dyn} \mid e\}!; \mathbf{Fun}^?; c_{11} \mapsto c_{12} * c'_2 \Rightarrow c$ for some canonical c , stepping by `N_PHI`.

$(\{x:\text{Dyn} \mid e\}!; \mathbf{Fun}^?; c_{11} \mapsto c_{12}; \mathbf{Fun}!; \{x:\text{Dyn} \mid e'\}?)$ Here $c_2 = \{x:\text{Dyn} \mid e'\}!; c'_2$, where c'_2 is canonical by Lemma 4.5.6. We can apply the IH to find $\{x:\text{Dyn} \mid e\}!; \mathbf{Fun}^?; c_{11} \mapsto c_{12}; \mathbf{Fun}! * c'_2 \Rightarrow c$ for some canonical c , stepping by `N_PREDPRED`.

$(\{x:B \mid e\}!)$ Here $c_2 = \{x:B \mid e\}^?; c'_2$, with c'_2 canonical by Lemma 4.5.6. We merge to c'_2 by `N_PHI` and `N_CANONICAL`.

$(\{x:B \mid e\}!; B!)$ Here c_2 comes from `Dyn`. If c_2 is $\{x:\text{Dyn} \mid e'\}?$, then we are done, since $\{x:B \mid e\}!; B!; \{x:\text{Dyn} \mid e'\}?$ is canonical. If it begins $\mathbf{Fun}!; c'_2$, we get `Fail`, which is canonical. If it begins $B^?; c'_2$, we can use the IH to show that $\{x:B \mid e\}! * c'_2 \Rightarrow c$ for a canonical c (since c'_2 is canonical by Lemma 4.5.6). Then we can step by `N_PHI`.

$(\{x:B \mid e\}!; B!; \{x:\text{Dyn} \mid e'\}?)$ Here $c_2 = \{x:\text{Dyn} \mid e'\}!; c'_2$. We know that c'_2 is canonical (Lemma 4.5.6), so by the IH we have $\{x:B \mid e\}!; B! * c'_2 \Rightarrow c$ for some canonical c . We step by `N_PREDPRED`.

$(\{x:B \mid e\}!; \{x:B \mid e'\}?)$ Here $c_2 = \{x:B \mid e'\}!; c'_2$. We know that c'_2 is canonical (Lemma 4.5.6), so by the IH we have $\{x:B \mid e\}! * c'_2 \Rightarrow c$ for some canonical c . We step by `N_PREDPRED`.

□

4.5.12 Lemma: For all canonical coercions c_1 and c_2 , if $c_1 * c_2 \Rightarrow c_3$ then $c_1; c_2 \longrightarrow^* c_3$.

Proof: By induction on the derivation of $c_1 * c_2 \Rightarrow c_3$.

(N_CANONICAL) Since $c_3 = c_1; c_2$, we are done immediately by reflexivity.

(N_FAILL) We have $\text{Fail} * c_2 \Rightarrow \text{Fail}$. By induction on the length of c_2 , we have $\text{Fail}; c_2 \longrightarrow^* \text{Fail}$.

(N_FAILR) We have $c_1 * \text{Fail} \Rightarrow \text{Fail}$. By induction on the length of c_1 , we have $c_1; \text{Fail} \longrightarrow^* \text{Fail}$.

(N_PHI) We have $c_1 * c_2 \Rightarrow c_3$. We have $c_1; D!; D?; c_2 \longrightarrow c_1; c_2$ for all D , so we are done by the IH.

(N_BFAILB) We have $c_1; B!; B'?; c_2 \longrightarrow c_1; \text{Fail}; c_2$. By induction on the length of c_1 and c_2 , we can conclude that $c_1; \text{Fail}; c_2 \longrightarrow^* \text{Fail}$.

(N_BFAILFUN) We have $c_1; B!; \text{Fun}?; c_2 \longrightarrow c_1; \text{Fail}; c_2$. By induction on the length of c_1 and c_2 , we know that $c_1; \text{Fail}; c_2 \longrightarrow^* \text{Fail}$.

(N_FUNFAILB) We have $c_1; \text{Fun}!; B?; c_2 \longrightarrow c_1; \text{Fail}; c_2$. By induction on the length of c_1 and c_2 , we know that $c_1; \text{Fail}; c_2 \longrightarrow^* \text{Fail}$.

(N_PREDPRED) We have $c_1 * c_2 \Rightarrow c_3$ and $c_1; \{x:T \mid e\}?; \{x:T \mid e\}!; c_2 \longrightarrow c_1; c_2$, so we have $c_1; c_2 \longrightarrow^* c_3$ by the IH.

(N_FUN) We know that $c_{21} * c_{11} \Rightarrow c_{31}$ and $c_{12} * c_{22} \Rightarrow c_{32}$ and $c_1 * (c_{31} \mapsto c_{32}); c_2 \Rightarrow c_3$.

We can rewrite $c_1; (c_{11} \mapsto c_{12}); (c_{21} \mapsto c_{22}); c_2 \longrightarrow c_1; ((c_{21}; c_{11}) \mapsto (c_{12}; c_{22})); c_2$. By the IH, we know that $c_{21}; c_{11} \longrightarrow^* c_{31}$ and $c_{12}; c_{22} \longrightarrow^* c_{32}$, so we can rewrite to $c_1; (c_{31} \mapsto c_{32}); c_2$. But we know by the IH that this rewrites to c_3 , so we are done.

□

4.5.13 Lemma [Merge is associative]: $c_1 \Downarrow (c_2 \Downarrow c_3) = (c_1 \Downarrow c_2) \Downarrow c_3$ for all canonical coercions c_1 , c_2 , and c_3 .

Proof: Consider the term $c_1; c_2; c_3$. On the one hand, we can say that $c_1; c_2; c_3 \longrightarrow^* c_1; c_2 \Downarrow c_3 \longrightarrow^* c_1 \Downarrow c_2 \Downarrow c_3$ by Lemma 4.5.12. On the other hand, we also have $c_1; c_2; c_3 \longrightarrow^* c_1 \Downarrow c_2; c_3 \longrightarrow^* c_1 \Downarrow c_2 \Downarrow c_3$.

Recall that results of merges are canonical forms (Lemma 4.5.11), which are normal (Lemma 4.5.5). Since rewriting is confluent (Lemma 4.5.2), it must be that case that $c_1 \Downarrow (c_2 \Downarrow c_3) = (c_1 \Downarrow c_2) \Downarrow c_3$. □

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

$$\begin{array}{c}
\frac{}{\langle \lambda x:T. e_{12} \rangle_{\text{ld}} v_2 \longrightarrow e_{12}[v_2/x]} \text{E_BETA} \qquad \frac{}{u_1 \langle c_1 \mapsto c_2 \rangle v_2 \longrightarrow \langle c_2 \rangle (u_1 \text{ld} (\langle c_1 \rangle v_2))} \text{E_FUN} \\
\\
\frac{}{op(v_1, \dots, v_n) \longrightarrow \llbracket op \rrbracket (v_1, \dots, v_n)} \text{E_OP} \\
\\
\frac{}{\langle \{x:T \mid e\}?\ ; c \rangle v \longrightarrow \langle c \rangle \langle \{x:T \mid e\}, e[v/x], v \rangle} \text{E_CHECK} \\
\\
\frac{}{\langle \{x:T \mid e\}, \text{true}_{\text{ld}}, u_c \rangle \longrightarrow u_{c \downarrow \{x:T \mid e\} ?}} \text{E_CHECKOK}\bullet \\
\\
\frac{}{\langle \{x:T \mid e\}, \text{false}_{\text{ld}}, v \rangle \longrightarrow \text{fail}} \text{E_CHECKFAIL} \quad \frac{d_1 \neq \{x:T \mid e\}?\ \quad c * d_1 \Rightarrow \text{Fail}}{\langle d_1; c_2 \rangle u_c \longrightarrow \text{fail}} \text{E_TAGFAIL}\bullet \\
\\
\frac{d_1 \neq \{x:T \mid e\}?\ \quad c * d_1 \Rightarrow c' \quad c' \neq \text{Fail}}{\langle d_1; c_2 \rangle u_c \longrightarrow \langle c_2 \rangle u_{c'}} \text{E_TAG}\bullet \\
\\
\frac{}{\langle \text{ld} \rangle v \longrightarrow v} \text{E_TAGID} \qquad \frac{e_1 \longrightarrow e'_1}{e_1 e_2 \longrightarrow e'_1 e_2} \text{E_APPL} \qquad \frac{e_2 \longrightarrow e'_2}{v_1 e_2 \longrightarrow v_1 e'_2} \text{E_APPR} \\
\\
\frac{e_i \longrightarrow e'_i}{op(v_1, \dots, v_{i-1}, e_i, \dots, e_n) \longrightarrow op(v_1, \dots, v_{i-1}, e'_i, \dots, e_n)} \text{E_OPINNER} \\
\\
\frac{}{\langle c_1 \rangle (\langle c_2 \rangle e) \longrightarrow \langle c_2 \downarrow c_1 \rangle e} \text{E_MERGE}\bullet \quad \frac{e \neq \langle c' \rangle e'' \quad c \neq \text{Fail} \quad e \longrightarrow e'}{\langle c \rangle e \longrightarrow \langle c \rangle e'} \text{E_COERCEINNER} \\
\\
\frac{e_2 \longrightarrow e'_2}{\langle \{x:T \mid e_1\}, e_2, v \rangle \longrightarrow \langle \{x:T \mid e_1\}, e'_2, v \rangle} \text{E_CHECKINNER} \\
\\
\frac{e \neq \langle c \rangle e'}{\langle \text{Fail} \rangle e \longrightarrow \text{fail}} \text{E_FAIL} \qquad \frac{}{\langle c \rangle \text{fail} \longrightarrow \text{fail}} \text{E_COERCERAISE} \\
\\
\frac{}{\text{fail } e_2 \longrightarrow \text{fail}} \text{E_APPRAISEL} \qquad \frac{}{v_1 \text{fail} \longrightarrow \text{fail}} \text{E_APPRAISER} \\
\\
\frac{}{op(v_1, \dots, v_{i-1}, \text{fail}, \dots, e_n) \longrightarrow \text{fail}} \text{E_OPRAISE} \quad \frac{}{\langle \{x:T \mid e\}, \text{fail}, v \rangle \longrightarrow \text{fail}} \text{E_CHECKRAISE}
\end{array}$$

Figure 4.22: EFFICIENT operational semantics

```

odd 3id
→ evenInt!→Bool? 2id
→ ⟨Bool?⟩ (even (⟨Int!⟩ 2id))
→ ⟨Bool?⟩ (((λx:Int. ...) Int?→Bool! 2Int!)
→ ⟨Bool!⟩ (⟨Bool!⟩ ((λx:Int. ...) Id (⟨Int?⟩ 2Int!)))
→ ⟨Id⟩ ((λx:Int. ...) Id (⟨Int?⟩ 2Int!))
→ ⟨Id⟩ ((λx:Int. ...) Id 2id)
→ ⟨Id⟩ (odd 1id)
→ ⟨Id⟩ (evenInt!→Bool? 0id)
→ ⟨Id⟩ (⟨Bool?⟩ (even (⟨Int!⟩ 0id)))
→ ⟨Bool?⟩ (even (⟨Int!⟩ 0id))
→ ⟨Bool?⟩ ((λx:Int. ...) Int?→Bool! 0Int!)
→ ⟨Bool!⟩ (⟨Bool!⟩ (λx:Int. ...) Id (⟨Int?⟩ 0Int!))
→ ⟨Id⟩ ((λx:Int. ...) Id (⟨Int?⟩ 0Int!))
→ ⟨Id⟩ ((λx:Int. ...) Id 0id)
→ ⟨Id⟩ trueid
→ trueid

```

Figure 4.23: Space-efficient reduction

4.5.2 Operational semantics

I give the changed operational semantics in Figure 4.22. The biggest change to the operational semantics is that `E_MERGE` explicitly merges the two coercions. Herman et al. simply say that they keep their coercions in normal form—that is, we should interpret normalization happening automatically when `E_MERGE` applies, even though they write `E_MERGE` as directly concatenating the two coercions into $c_2; c_1$. My semantics explicitly normalizes the coercions (rule `E_MERGE`), possibly stopping the program on the next step (the no-longer-useless rule `E_FAIL`).

Otherwise, the rules are largely the same as the naïve semantics, though we’re now able to use merges to distill the tag rules into a few possibilities: `E_TAG` replaces all of the successful `F_TAG*` rules; `E_TAGFAIL` replaces all of the failing `F_TAG*` rules. `E_CHECKOK` is essentially `F_CHECKOK`, though it uses a merge instead of concatenation (though the typing rules mean that the merge will apply `N_CANONICAL` every time).

We can finally observe that our reduction is space efficient: the coercions in Figure 4.23 don’t grow with the size of the input like the coercions in Figure 4.14 or the casts in Figure 4.2. I discuss this claim in more detail in Section 4.7.

4.5.3 Proofs

`EFFICIENT` enjoys type soundness; we can show as much using standard syntactic methods. We must assume that $\text{ty}(k)$ and $\text{ty}(op)$ are always well formed.

4.5.14 Lemma [Regularity]: 1. If $\Gamma \vdash e : T$, then $\vdash T$.

2. If $\Gamma \vdash u : T$, then $\vdash T$.
3. If $\vdash \Gamma$, then $\vdash T$ for all $x:T \in \Gamma$.

Proof: By mutual induction on the typing derivations. □

4.5.15 Lemma [Determinism]: If $e \longrightarrow e_1$ and $e \longrightarrow e_2$, then $e_1 = e_2$.

Proof: By induction on $e \longrightarrow e_1$, observing that in each case the same rule must have applied to find $e \longrightarrow e_2$. In the E_MERGE case, we rely on the fact that merge is an operator on canonical coercions (Lemma 4.5.11). □

4.5.16 Lemma [Canonical forms]: If $\emptyset \vdash v : T$ then:

- If $T = \text{Bool}$, then $v = \text{true}_{\text{ld}}$ or $v = \text{false}_{\text{ld}}$.
- If $T = T_1 \rightarrow T_2$, then $v = \lambda x:T_1'. e_c$ where $c = \text{ld}$ or $c = c_1 \mapsto c_2$.

Proof: By cases on the typing derivation. We observe that the only coercions from B to B are **ld** and **Fail**, and $\text{true}_{\text{Fail}}$ isn't well formed.

In the function case, the only coercions from $T_{11} \rightarrow T_{12}$ to $T_{21} \rightarrow T_{22}$ are **ld** and **Fail** and $c_1 \mapsto c_2$. Since u_{Fail} isn't well formed, c must be either the identity or a functional coercion. □

4.5.17 Lemma [Progress]: If $\emptyset \vdash e : T$ then either e is a result, or there exists an e' such that $e \longrightarrow e'$.

Proof: By induction on the typing derivation $\emptyset \vdash e : T$. I give only the interesting cases.

(T_COERCE) We have $\emptyset \vdash \langle c \rangle e : T_2$, where $\vdash c : T_1 \rightsquigarrow T_2$ and $\emptyset \vdash e : T_1$. If the inner term is of the form $\langle c' \rangle e'$, we step by E_MERGE (in which case we are well typed by Lemma 4.5.9). If $c = \text{Fail}$, we step E_FAIL (in which case we are well typed by Lemma 4.5.14 and T_FAIL). If not, it either steps by E_COERCEINNER or e is a result. If $e = \text{fail}$, we step by E_COERCERAISE. If $e = u_{c'}$, we step by E_TAG, E_TAGFAIL, or E_CHECK, depending on what the leftmost coercion in c is. If there is no leftmost coercion, i.e., we have $\langle \text{ld} \rangle v$, then we step by E_TAGID.

By inversion, $\emptyset \vdash e_2 : \text{Bool}$. By canonical forms (Lemma 4.5.16), if e_2 is a value, then it is either true_{ld} or false_{ld} . In the former case, we step by E_CHECKOK; in the latter case, we step by E_CHECKFAIL. □

4.5.18 Lemma [Weakening]: If $\Gamma_1, \Gamma_2 \vdash e : T$ and $x \notin \text{dom}(\Gamma)$ and $\vdash T'$, then $\Gamma_1, x:T', \Gamma_2 \vdash e : T$.

Proof: By induction on e . □

4.5.19 Lemma [Substitution]: If $\emptyset \vdash v : T'$:

- if $\Gamma_1, x:T', \Gamma_2 \vdash e : T$ then $\Gamma_1, \Gamma_2 \vdash e[v/x] : T$.
- if $\Gamma_1, x:T', \Gamma_2 \vdash u : T$ then $\Gamma_1, \Gamma_2 \vdash u[v/x] : T$.

Proof: By induction on the typing derivation of e . □

4.5.20 Lemma [Preservation]: If $\emptyset \vdash e : T$ and $e \longrightarrow e'$, then $\emptyset \vdash e' : T$.

Proof: By induction on the typing derivation. I give only the interesting cases.

(T_COERCE) We go by cases on the step taken:

(E_COERCEINNER) By the IH.

(E_MERGE) By T_COERCE and Lemma 4.5.9.

(E_FAIL) By T_FAIL and regularity (Lemma 4.5.14).

(E_COERCERAISE) By T_FAIL and regularity (Lemma 4.5.14).

(E_TAG) By T_COERCE, T_VAL, and C_COMPOSE, knowing that the result of the merge isn't a fail or $\{x:T \mid e\}$? for T_VAL.

(E_TAGFAIL) By T_FAIL and regularity (Lemma 4.5.14).

(E_TAGID) By the typing assumption on the value.

(E_CHECK) By T_COERCE and T_CHECK.

(T_CHECK) We go by cases on the step taken:

(E_CHECKINNER) By the IH.

(E_CHECKOK) By T_VALREFINE and C_COMPOSE, knowing that the merge can't be a fail—predicates never introduce failures. We can take the evaluation we need directly from the T_CHECK derivation.

(E_CHECKFAIL) By T_FAIL and regularity (Lemma 4.5.14).

(E_CHECKRAISE) By T_FAIL and regularity (Lemma 4.5.14). □

4.5.21 Theorem [Type soundness]: If $\emptyset \vdash e : T$ then either $e \longrightarrow^* r$ such that $\emptyset \vdash r : T$ or e diverges.

Proof: Using progress (Lemma 4.5.17) and preservation (Lemma 4.5.20). Naturally the proof is not constructive! □

In Section 4.3, I define cotermination and show that if $e \longrightarrow_n^* r$ then $\langle c \rangle e$ and $\langle c \rangle r$ coterminate, and if e diverges so does $\langle c \rangle e$ (Lemma 4.3.13). Unfortunately, neither of these exactly hold in EFFICIENT. For the former, consider $e = \langle \{x:\text{Bool} \mid \text{false}_{\text{Id}}\} \rangle \text{true}_{\text{Id}} \longrightarrow^* \text{fail}$. The term $\langle \{x:\text{Bool} \mid \text{false}_{\text{Id}}\} \rangle e \longrightarrow^* \text{true}_{\text{Id}}$. For the latter, suppose that `diverge` is a closed, divergent term, such as

$$(\lambda x:(\text{Dyn} \rightarrow \text{Dyn}). x (\langle \mathbf{Fun!} \rangle x))_{\text{Id}} (\langle \mathbf{Fun?} \mapsto \text{Id} \rangle (\lambda x:(\text{Dyn} \rightarrow \text{Dyn}). x (\langle \mathbf{Fun!} \rangle x))_{\text{Id}})$$

On the one hand

$$\begin{aligned} & \langle \{x:\text{Bool} \mid \text{diverge}\} \rangle (\langle \{x:\text{Bool} \mid \text{diverge}\} \rangle \text{true}_{\text{Id}}) \\ \longrightarrow & \langle \text{Id} \rangle \text{true}_{\text{Id}} \\ \longrightarrow & \text{true}_{\text{Id}} \end{aligned}$$

but $\langle \{x:\text{Bool} \mid \text{diverge}\} \rangle \text{true}_{\text{Id}}$ diverges. The theorem doesn't hold for when e diverges or evaluates to `fail`—but it *does* hold for values.

4.5.22 Lemma: If $e \longrightarrow^* v$ and $\langle c \rangle v \longrightarrow^* v'$ then $\langle c \rangle e \longrightarrow^* v'$.

Proof: By induction on $e \longrightarrow^* v$, and then by cases on the step taken. (The base case ($e = v$) is immediate.)

If the step doesn't have an exposed coercion, we can just reapply the step taken. The core rules (E_BETA, E_FUN, E_OP, E_CHECKOK, E_CHECKFAIL), the congruence rules (E_APPL, E_APPR, E_OPINNER, E_CHECKINNER), and the exception raising rules (E_APPRAISEL, E_APPRAISER, E_OPRAISE, E_CHECKRAISE) all fit this rubric. The remaining cases must merge coercions in e with c . The general thrust is as for Lemma 4.3.13: we will merge and normalize, take some small tag manipulation step (determined by analyzing how the canonical merge went), and then reproduce that for a stepped e .

(E_TAGID) We step by E_MERGE (using N_CANONICAL) and apply the IH.

(E_TAG) We have $e_2 = \langle d_1; c_2 \rangle u_{c'_2} \longrightarrow \langle c_2 \rangle u_{c'_2 \downarrow d_1}$ with $d_1 \neq \{x:T \mid e\}$?

Instead we step to $\langle d_1; c_2 \downarrow c \rangle u_{c'_2}$ by E_MERGE.

If d_1 remains unaffected by the merge, we are done easily: we step by E_TAG and can apply the IH on $\langle c \rangle (\langle c_2 \rangle u_{c'})$.

If d_1 is affected, then all of $d_1; c_2$ must have disappeared. If $d_1; c_2 \downarrow c = c$, then we are done by assumption. So instead it must be the case that $d_1; c_2 \downarrow c$ is some suffix of c .

By the IH, $\langle c \rangle (\langle c_2 \rangle u_{c'_2 \downarrow d_1}) \longrightarrow^* v'$, which steps to $\langle c_2 \downarrow c \rangle u_{c'_2 \downarrow d_1}$. Whatever coercion was left in $c_2 \downarrow c$ that eliminated d_1 must now be exposed, so we can step by E_TAG to find $\langle d_1; c_2 \downarrow c \rangle u_{c'_2} \longrightarrow^* v'$.

(E_TAGFAIL) Contradictory—`fail` isn't a value.

(E_CHECK) We have $e_2 = \langle \{x:T \mid e\}^?; c' \rangle v'_2 \longrightarrow \langle c' \rangle \langle \{x:T \mid e\}, e[v'_2/x], v'_2 \rangle$; since this steps to some value v' , it must be the case that $e[v'_2/x] \longrightarrow^* \text{true}_{\text{Id}}$.

The combined term steps by E_MERGE to $\langle \{x:T \mid e\}^?; c' \Downarrow c \rangle v'_2$.

Now, the result of $\{x:T \mid e\}^?; c' \Downarrow c$ either has the same refinement checking coercion on the front or it doesn't. If it does, we can use the evaluation we found above to step to $\langle c' \Downarrow c \rangle v'_2$ (using Lemma 4.5.13).

If it doesn't, we're already at $\langle c' \Downarrow c \rangle v'_2$. In either case, that term is equivalent to $\langle c \rangle (\langle c' \rangle v'_2)$, and we can apply the IH on $\langle c' \rangle v'_2 \longrightarrow^* v_2$.

(E_COERCEINNER) We have $e_2 = \langle c_1 \rangle (\langle c_2 \rangle e'_2) \longrightarrow \langle c_2 \Downarrow c_1 \rangle e'_2$. We'll step by E_MERGE twice to find $\langle c_2 \Downarrow c_1 \Downarrow c \rangle e'_2$. By associativity of merges, that term is equal to $\langle c_1 \Downarrow c_2 \Downarrow c \rangle e_2$.

By the IH we know that $\langle c \rangle (\langle c_2 \Downarrow c_1 \rangle e'_2) \longrightarrow^* v'$, so we are done.

(E_FAIL) Contradictory—fail isn't a value.

(E_MERGE) We step by E_MERGE and then by E_MERGE again to find the same result, and we are done.

□

4.6 Soundness of EFFICIENT with regard to NAIVE

We will never get an exact semantic match between the naïve and space-efficient semantics: the ψ rule for refinements in the space-efficient semantics means that some checks will happen in the naïve semantics that won't happen in the space-efficient semantics. If those checks fail or diverge, then the NAIVE term won't behave the same as its translation into EFFICIENT. All we can hope for is that *if* NAIVE produces a value, then EFFICIENT will produce a similar one.

I adapt the asymmetric logical relations from Chapter 2 to show that the two calculi behave mostly the same, with NAIVE diverging and failing more often. My proof method largely follows the early one, as well: I augment the logical relation with a separate relation $\vdash c_1 \equiv c_2$ (read “ c_1 is equivalent to c_2 ”) on coercions (Figure 4.26). Just as the $\vdash T_1 \Rightarrow T_2 \equiv c$ relation from Section 4.4 was an invariant relating casts to coercions, $\vdash c_1 \equiv c_2$ relates casts to their normal forms. With this relation in hand, we first prove that applying related coercions to related values yields related values (Lemma 4.6.8). Then we show that well typed terms are related to their translations and that coercions are related to their translations (Theorem 4.6.10).

We are forced into a *step-indexed* logical relation because of the type **Dyn**. In particular, the more common, fixpoint-on-types definition wouldn't work for the function tag case, since we need to take values at **Dyn** and relate them at the type **Dyn**→**Dyn**. I define the step-indexed logical relation in Figure 4.24; when I say $e \longrightarrow_n^m e'$, I mean

Value rules

$$\boxed{v_1 \lesssim^j v_2 : T}$$

$$\begin{aligned} \forall j. k_{\text{ld}} \lesssim^j k_{\text{ld}} : B &\iff \text{ty}(k) = B \\ v_{11} \lesssim^j v_{21} : T_1 \rightarrow T_2 &\iff \\ \forall m < j. \forall v_{12} \lesssim^m v_{22} : T_1. v_{11} v_{12} \approx^m v_{21} v_{22} : T_2 & \end{aligned}$$

$$\begin{aligned} v_{1B!} \lesssim^j u_{2c\Downarrow B!} : \text{Dyn} &\iff v_1 \lesssim^j u_{2c} : B \\ v_{1\text{Fun}!} \lesssim^j u_{2c\Downarrow \text{Fun}!} : \text{Dyn} &\iff v_1 \lesssim^j u_{2c} : \text{Dyn} \rightarrow \text{Dyn} \end{aligned}$$

$$\begin{aligned} v_{1\{x:T|e_1\}^?} \lesssim^j u_{c\Downarrow\{x:T|e_2\}^?} : \{x:T \mid e_1\} &\iff \\ \forall m < j. v_1 \lesssim^m u_{2c} : T \wedge \{x:T \mid e_1\} \lesssim^m \{x:T \mid e_2\} & \end{aligned}$$

Term rules

$$\boxed{e_1 \approx^j e_2 : T}$$

$$\begin{aligned} e_1 \approx^j e_2 : T &\iff \\ e_1 \text{ diverges} \vee & \\ \forall m < j. e_1 \rightarrow_n^m \text{fail} & \\ \vee e_1 \rightarrow_n^m v_1 \implies e_2 \rightarrow^* v_2 \wedge v_1 \lesssim^{(j-m)} v_2 : T & \end{aligned}$$

Type rules

$$\boxed{T_1 \lesssim^j T_2}$$

$$\begin{aligned} B \lesssim^j B &\iff \text{Dyn} \lesssim^j \text{Dyn} \\ T_{11} \rightarrow T_{12} \lesssim^j T_{21} \rightarrow T_{22} &\iff T_{11} \lesssim^j T_{21} \wedge T_{12} \lesssim^j T_{22} \\ \{x:T \mid e_1\} \lesssim^j \{x:T \mid e_2\} &\iff \\ \forall m < j. \forall v_1 \lesssim^m v_2 : T. e_1[v_1/x] \approx^m e_2[v_2/x] : \text{Bool} & \end{aligned}$$

Closing substitutions and open terms

$$\boxed{\Gamma \Vdash_{\lesssim}^j \delta}$$

$$\boxed{\Gamma \vdash e_1 \approx e_2 : T}$$

$$\begin{aligned} \Gamma \Vdash_{\lesssim}^j \delta &\iff \forall x \in \text{dom}(\Gamma). \delta_1(x) \lesssim^j \delta_2(x) : T \\ \Gamma \vdash e_1 \approx e_2 : T &\iff \forall j \geq 0. \forall \Gamma \Vdash_{\lesssim}^j \delta. \delta_1(e_1) \approx^j \delta_2(e_2) : T \end{aligned}$$

Figure 4.24: Relating NAIVE and EFFICIENT

that e steps to e' in *exactly* m steps. The definitions begin by defining a relation $v_1 \lesssim^j v_2 : T$ for closed values and a relation $e_1 \lesssim^j e_2 : T$ for closed terms as a fixpoint on the index j .⁵ I lift the definitions to open terms by defining *dual closing value substitutions* δ ; if $\Gamma \Vdash_{\lesssim}^j \delta$ and $x:T \in \Gamma$, then $\delta_1(x) \lesssim^j \delta_2(x) : T$.

NAIVE terms are on the left of the relation, while EFFICIENT terms are on the right. We require that both sides be well typed. We obtain the asymmetry we seek by saying that *when* the naïve semantics yields a value, then the space-efficient yields a similar one—but otherwise, the naïve semantics will fail or diverge. This definition still allows EFFICIENT to diverge or to fail, but then the naïve semantics must also diverge or fail—but note that it’s possible for the left-hand side of the relation to diverge and the right-hand side to fail, and vice versa. This is possible because the naïve semantics could run a check that diverges, while the space-efficient semantics skips that check and instead runs a failing one.

Step-indexed logical relations are commonly asymmetric, for a separate reason: only one side needs the index; for us, it is particularly convenient to put the step-index on the naïve side, allowing us to skip reasoning about how step indices and merges interact.

We could try to be more specific in the relation: either you get the same thing, or you get divergence or a failure—that can be traced back exactly to a check that happened in NAIVE but not in EFFICIENT. I omit this more precise tracking for simplicity’s sake.

The value relation $v_1 \lesssim^j v_2 : T$ is subtler than usual for this logical relation: the definitions at Dyn and $\{x:T \mid e\}$ must shuffle some tags around. In particular, the rule for type Dyn is split into cases by the underlying tag of values. The case for refinements $\{x:T \mid e\}$ requires that the values be related at the underlying type T (recalling that $T = B$ or $T = \text{Dyn}$) and also that the values be tagged as satisfying the predicate (or a related predicate in EFFICIENT). For well typed terms, this is enough to ensure that the underlying values satisfy their refinement predicates.

My proof works by showing that a well-typed naïve term e is related to its translation into the space-efficient term $\text{canonical}(e)$. I define canonical in Figure 4.25, omitting most of the cases since they are homomorphic. In particular, $\text{canonical}(v)$ must unfold the stacked tags on a NAIVE value and merge them into a single coercion.

Before we continue, I must justify that canonical is a function—in particular, is the case for composite coercions valid?

4.6.1 Lemma: canonical is a well defined function that produces canonical coercions.

Proof: By induction on c , using Lemma 4.5.11 to show that $c_1 \Downarrow c_2$ is a unique canonical coercion when c_1 and c_2 are canonical. \square

⁵The definition for terms is just fancy notation for a proposition involving the relation on values; this stratified definition conveniently separates the definition of related values and evaluation, while avoiding the need for $\top\top$ -closure [54].

Composite coercions	$\boxed{\text{canonical}(c) : c}$
	$\text{canonical}(\text{Id}) = \text{Id}$
	$\text{canonical}(d_1; \dots; d_n) = \text{canonical}(d_1) \Downarrow \text{canonical}(d_2; \dots; d_n)$
Primitive coercions	$\boxed{\text{canonical}(d) : d}$
	$\text{canonical}(\text{Fail}) = \text{Fail}$
	$\text{canonical}(D!) = \text{canonical}(D)!$
	$\text{canonical}(D?) = \text{canonical}(D)?$
	$\text{canonical}(c_1 \mapsto c_2) = \text{canonical}(c_1) \mapsto \text{canonical}(c_2)$
Tags	$\boxed{\text{canonical}(D) : D}$
	$\text{canonical}(B) = B$
	$\text{canonical}(\mathbf{Fun}) = \mathbf{Fun}$
	$\text{canonical}(\{x:T \mid e\}) = \{x:T \mid \text{canonical}(e)\}$
Pre-values	$\boxed{\text{canonical}(u) : u}$
	$\text{canonical}(k) = k$
	$\text{canonical}(\lambda x:T. e) = \lambda x:\text{canonical}(T). \text{canonical}(e)$
Expressions	$\boxed{\text{canonical}(e) : e}$
	$\text{canonical}(x) = x$
	$\text{canonical}(u_{\text{Id}}) = \text{canonical}(u)_{\text{Id}}$
	$\text{canonical}(v_d) = u_{c \Downarrow d}$ where $\text{canonical}(v) = u_c$
	$\text{canonical}(\text{fail}) = \text{fail}$
	$\text{canonical}(op(e_1, \dots, e_n)) = op(\text{canonical}(e_1), \dots, \text{canonical}(e_n))$
	$\text{canonical}(e_1 e_2) = \text{canonical}(e_1) \text{canonical}(e_2)$
	$\text{canonical}(\langle c \rangle e) = \langle \text{canonical}(c) \rangle \text{canonical}(e)$
	$\text{canonical}(\langle \{x:T \mid e_1\}, e_2, v \rangle) = \langle \text{canonical}(\{x:T \mid e_1\}), \text{canonical}(e_2), \text{canonical}(v) \rangle$
Types	$\boxed{\text{canonical}(T) : T}$
	$\text{canonical}(B) = B$
	$\text{canonical}(T_1 \rightarrow T_2) = \text{canonical}(T_1) \rightarrow \text{canonical}(T_2)$
	$\text{canonical}(\text{Dyn}) = \text{Dyn}$
	$\text{canonical}(\{x:T \mid e\}) = \{x:T \mid \text{canonical}(e)\}$
Contexts	$\boxed{\text{canonical}(\Gamma) : \Gamma}$
	$\text{canonical}(\emptyset) = \emptyset$
	$\text{canonical}(\Gamma, x:T) = \text{canonical}(\Gamma), x:\text{canonical}(T)$

Figure 4.25: Canonicalizing NAIVE terms

In order to put $\text{canonical}(e)$ in the relation, we must know that it is well typed. Since some of the runtime typing terms require facts about derivations that will be hard to translate—in particular, T_TAGVALREFINE and T_CHECKNAIVE —we’ll exclude them from the proof. Normal “source” terms shouldn’t have any of these, anyway. Later on, when we’ve proved that naïve terms are logically related to their canonical translations, we’ll have the evaluation derivations after all.

4.6.2 Lemma [Preservation for canonical]: Assuming that no refinement tags or active checking forms are present:

1. If $\vdash c : T_1 \rightsquigarrow T_2$ then $\vdash \text{canonical}(c) : \text{canonical}(T_1) \rightsquigarrow \text{canonical}(T_2)$.
2. If $\vdash d : T_1 \rightsquigarrow T_2$ then $\vdash \text{canonical}(d) : \text{canonical}(T_1) \rightsquigarrow \text{canonical}(T_2)$.
3. If $\Gamma \vdash u : T$ then $\text{canonical}(\Gamma) \vdash \text{canonical}(u) : \text{canonical}(T)$.
4. If $\Gamma \vdash e : T$ then $\text{canonical}(\Gamma) \vdash \text{canonical}(e) : \text{canonical}(T)$.
5. If $\vdash T$ then $\vdash \text{canonical}(T)$.
6. If $\vdash \Gamma$ then $\vdash \text{canonical}(\Gamma)$.

Proof: By simultaneous induction on the derivations, using Lemma 4.5.9 when merging coercions.

(C_ID) By C_ID.

(C_COMPOSE) By IH (2) and IH (1) and Lemma 4.5.9.

(C_FAIL) By C_FAIL.

(C_UNTAG) By C_UNTAG and the IH.

(C_TAG) By C_TAG and the IH.

(C_FUN) By IH (1) and C_FUN.

(T_CONST) By T_CONST and IH (6).

(T_ABS) By IH (5) and IH (4), we can reapply T_ABS.

(T_VAR) By IH (6) and T_VAR.

(T_PREVAL) By IH (3) and T_VAL.

(T_TAGVAL) By IH (4), Lemma 4.5.9, and T_VAL.

(T_TAGVALREFINE) Contradictory—we assumed these weren’t present.

(T_OP) By IH (4) and T_OP.

(T_APP) By IH (4) and T_APP.

(T_COERCE) By IH (1) and IH (4) and T_COERCE.

(T_FAIL) By IH (6) and IH (5), we can reapply T_FAIL.

(T_CHECKNAIVE) Contradictory—we assumed these weren't present.

(WF_BASE) By WF_BASE.

(WF_DYN) By WF_DYN.

(WF_FUN) By IH (5) and WF_FUN.

(WF_REFINE) By IH (4) and WF_REFINE, noting that T is either B or Dyn , so $\text{canonical}(T) = T$.

(WF_EMPTY) Immediate.

(WF_EXTEND) By IH (6) and IH (5), we can reapply WF_EXTEND.

□

The ultimate goal is soundness: if $\Gamma \vdash e : T$ then $\Gamma \vdash e \approx \text{canonical}(e) : T$. The proof works in a few stages: first I define relations $\vdash c$ **ignorable** (coercions which are equivalent to `Id` or `Fail`) and $\vdash c$ **failable** (coercions which are equivalent to `Fail`). I define these relations in Figure 4.26. We then prove lemmas that allow us to easily work with ignorable and failable coercions (Lemma 4.6.6 and Lemma 4.6.7, respectively). Then we relate non-canonical coercions to canonical ones (using a separate inductive relation $\vdash c_1 \equiv c_2$, defined in Figure 4.26). We show that such related coercions take related values to related values (Lemma 4.6.8). We then prove a separate lemma showing that related coercions are logically related on related terms (Lemma 4.6.9). This is not a trivial extension of the similar lemma for values, due to coercion merges. With those lemmas to hand, we can finally prove soundness: that terms are related to their translations (Theorem 4.6.10). Don't worry—I'll explain the proof less tersely as we go.

We begin by establishing standard properties of the logical relation. We will use closure under evaluation (Lemmas 4.6.3 and 4.6.4), fast-path failure (Lemma 4.6.5), and determinism (Lemmas 4.3.4 and 4.5.15) extensively. If you've already read Chapter 2, the proof techniques here should be old hat.

4.6.3 Lemma [Expansion]: If $e_1 \longrightarrow_n^* e'_1$ and $e_2 \longrightarrow^* e'_2$ then $e'_1 \approx^j e_2 : T$ implies $e_1 \approx^j e'_2 : T$.

Proof: Let $m < j$. If e'_1 diverges, so does e_1 by determinism (Lemma 4.3.4) and we are done.

Otherwise, we have $e'_1 \rightarrow_n^m r$. So there exists an m' such that $e_1 \rightarrow_n^{m'} r$. If $m' > j$, then we are done vacuously; if $m' < j$ then we are done by assumption. \square

4.6.4 Lemma [Contraction]: If $e_1 \rightarrow_n^* e'_1$ and $e_2 \rightarrow_n^* e'_2$ then $e_1 \approx^j e_2 : T$ implies $e'_1 \approx^j e'_2 : T$.

Proof: Let $m < j$. If e_1 diverges, so does e'_1 by determinism (Lemma 4.3.4), and we are done.

Otherwise, we have $e_1 \rightarrow_n^m r$. But since $e_1 \rightarrow_n^* e'_1$, we know that $e'_1 \rightarrow_n^{m'} r$ for $m' < m$ by determinism (Lemma 4.3.4), so we are done by assumption. \square

4.6.5 Lemma: For all indices j , if $e_1 \rightarrow_n^m \text{fail}$ then $e_1 \approx^j e_2 : T$.

Proof: If e_1 goes to fail in $m \geq j$ steps, then we are done vacuously. If $m < j$, then we are done by definition. \square

Ignorable coercions can be freely added or removed to naïve terms while preserving logical relation to space-efficient terms. The `I_ID` rule obviously fits this bill; `I_FAIL` is also acceptable, when we realize that `NAIVE` can fail more often than `EFFICIENT`. `I_BB` captures the case of an injection of a base-value into type dynamic, with some possible extra coercions in the middle. `I_FUNFUN` and `I_PREDSAME` are similar. Note that `I_PREDPRED` can't have extra coercions in the middle—the coercion typing rules ensure that the only non-Fail coercion that can come after $\{x:T \mid e\}?$ is $\{x:T \mid e\}!$. The `I_*` rules try to capture the logic of similarly named `N_*` rules.

4.6.6 Lemma: If $\langle c_1; c_2 \rangle v_1 \approx^j e_2 : T$ and $\vdash c_1$ **ignorable** then $\langle c_2 \rangle v_1 \approx^j e_2 : T$.

Proof: By induction on $\vdash c_1$ **ignorable**.

Let $m < j$.

(`I_ID`) By definition, `ld`; $c_2 = c_2$, so we are done immediately.

(`I_FAIL`) $\langle \text{Fail}; c_2 \rangle v_1 \rightarrow_n \text{fail}$, and $\text{fail} \approx^j e_2 : T$ by Lemma 4.6.5.

(`I_BB`) $\langle B!; c'_1; B?; c_2 \rangle v_1 \rightarrow_n \langle c'_1; B?; c_2 \rangle v_{1B!}$ by `F_TAGB`; the reduced term is then related to e_2 at T (Lemma 4.6.4). Since $\vdash c'_1$ **ignorable**, we can apply the IH and find that $\langle B?; c_2 \rangle v_{1B!} \approx^j e_2 : T$. By `F_TAGBB` and Lemma 4.6.4, we have $\langle c_2 \rangle v_1 \approx^j e_2 : T$.

(`I_FUNFUN`) $\langle \text{Fun!}; c'_1; \text{Fun?}; c_2 \rangle v_1 \rightarrow_n \langle c'_1; \text{Fun?}; c_2 \rangle v_{1\text{Fun!}}$ by `F_TAGFUN`; the reduced term is then related to e_2 at T (Lemma 4.6.4). Since $\vdash c'_1$ **ignorable**, we can apply the IH and find that $\langle \text{Fun?}; c_2 \rangle v_{1\text{Fun!}} \approx^j e_2 : T$. By `F_TAGFUNFUN` and Lemma 4.6.4, we have $\langle c_2 \rangle v_1 \approx^j e_2 : T$.

Ignorable coercions

$\boxed{\vdash c \text{ ignorable}}$

$$\begin{array}{c}
\frac{}{\vdash \text{ld ignorable}} \text{I_ID} \qquad \frac{}{\vdash \text{Fail ignorable}} \text{I_FAIL} \qquad \frac{\vdash c \text{ ignorable}}{\vdash B!; c; B? \text{ ignorable}} \text{I_BB} \\
\\
\frac{\vdash c \text{ ignorable}}{\vdash \text{Fun!}; c; \text{Fun? ignorable}} \text{I_FUNFUN} \qquad \frac{}{\vdash \{x:T \mid e\}?, \{x:T \mid e\}! \text{ ignorable}} \text{I_PREDPRED} \\
\\
\frac{\vdash c \text{ ignorable}}{\vdash \{x:T \mid e\}!, c; \{x:T \mid e\}? \text{ ignorable}} \text{I_PREDSAME} \\
\\
\frac{\vdash c_1 \text{ ignorable} \quad \vdash c_2 \text{ ignorable}}{\vdash c_1; c_2 \text{ ignorable}} \text{I_CONCAT}
\end{array}$$

Failable coercions

$\boxed{\vdash c \text{ failable}}$

$$\begin{array}{c}
\frac{}{\vdash \text{Fail failable}} \text{L_FAIL} \qquad \frac{B \neq B' \quad \vdash c \text{ ignorable}}{\vdash B!; c; B'? \text{ failable}} \text{L_BB} \\
\\
\frac{\vdash c \text{ ignorable}}{\vdash B!; c; \text{Fun? failable}} \text{L_BFUN} \qquad \frac{\vdash c \text{ ignorable}}{\vdash \text{Fun!}; c; B? \text{ failable}} \text{L_FUNB}
\end{array}$$

Relating coercions

$\boxed{\vdash c_1 \equiv c_2}$

$\boxed{\vdash d_1 \equiv d_2}$

$\boxed{\vdash D_1 \equiv D_2}$

$$\begin{array}{c}
\frac{\vdash c'_i \text{ ignorable} \quad \vdash c_i \equiv d_i}{\vdash c'_0; c_1; c'_1; c_2; \dots; c'_{n-1}; c_n; c'_n \equiv d_1; \dots; d_n} \text{R_COMPOSITE} \qquad \frac{\vdash c \text{ ignorable}}{\vdash c \equiv \text{ld}} \text{R_ID} \\
\\
\frac{\vdash c'_i \text{ ignorable} \quad \vdash (c_{1 \ n \ 1}; \dots; c_{1 \ 1 \ 1}) \mapsto (c_{1 \ 1 \ 2}; \dots; c_{1 \ n \ 2}) \equiv c_{21} \mapsto c_{22}}{\vdash (c_{1 \ 1 \ 1} \mapsto c_{1 \ 1 \ 2}); c'_1; \dots; c'_n; (c_{1 \ n \ 1} \mapsto c_{1 \ n \ 2}) \equiv c_{21} \mapsto c_{22}} \text{R_FUN} \qquad \frac{\vdash c_1 \text{ failable}}{\vdash c'_1; c_1; c'_2 \equiv c_2} \text{R_FAIL} \\
\\
\frac{\vdash D \equiv D}{\vdash D! \equiv D!} \text{R_TAG} \qquad \frac{\vdash D \equiv D}{\vdash D? \equiv D?} \text{R_UNTAG} \\
\\
\frac{}{\vdash B \equiv B} \text{R_DB} \qquad \frac{}{\vdash \text{Fun} \equiv \text{Fun}} \text{R_DFUN} \qquad \frac{x:\text{Bool} \vdash e_1 \approx e_2 : \text{Bool}}{\vdash \{x:T \mid e_1\} \equiv \{x:T \mid e_2\}} \text{R_DPRED}
\end{array}$$

Figure 4.26: Relating NAIVE coercions to canonical EFFICIENT coercions

(I_PREDPRED) $\langle \{x:T \mid e\}?, \{x:T \mid e\}!, c_2 \rangle v_1 \longrightarrow_n \langle \{x:T \mid e\}!, c_2 \rangle \langle \{x:T \mid e\}, e[v_1/x], v_1 \rangle$ by F_CHECK. By type soundness (Theorem 4.3.9), we know that $e[v_1/x]$ either reduces to a value, reduces to **fail**, or diverges. In either of the last two cases, we are done by the definition of the logical relation or Lemma 4.6.5, respectively.

Suppose that $e[v_1/x] \longrightarrow_n^* v$; the value v must be either $\mathbf{true}_{\text{id}}$ or $\mathbf{false}_{\text{id}}$, since $\emptyset \vdash e[v_1/x] : \mathbf{Bool}$. In the latter case we are done by Lemma 4.6.5, just like for when the whole term reduced to **fail**.

So then $\langle \{x:T \mid e\}!, c_2 \rangle \langle \{x:T \mid e\}, \mathbf{true}_{\text{id}}, v_1 \rangle \longrightarrow_n \langle \{x:T \mid e\}!, c_2 \rangle v_{1_{\{x:T \mid e\}}?}$ by F_CHECKOK. We can then step by F_TAGPREDPRED and apply Lemma 4.6.4 to find that $\langle c_2 \rangle v_1 \approx^j e_2 : T$.

(I_PREDSAME) $\langle \{x:T \mid e\}!, c'_1; \{x:T \mid e\}?, c_2 \rangle v_1 \approx^j e_2 : T$. By inversion of the typing of v_1 , we know that it must be of the form $v'_{1_{\{x:T \mid e\}}?}$, and furthermore $e[v_1/x] \longrightarrow_n^* \mathbf{true}_{\text{id}}$.

Lemma 4.6.4 and F_PREDPRED give us $\langle c'_1; \{x:T \mid e\}?, c_2 \rangle v'_1 \approx^j e_2 : T$, so we can apply the IH on $\vdash c'_1$ **ignorable** to find $\langle \{x:T \mid e\}?, c_2 \rangle v'_1 \approx^j e_2 : T$. We then step by F_TAGPREDPRED, then by F_CHECK. Since we know that $e[v'_1/x] \longrightarrow_n^* \mathbf{true}_{\text{id}}$, we can step the whole term using that relation and F_CHECKOK to find $\langle c_2 \rangle v'_{1_{\{x:T \mid e\}}?} = \langle c_2 \rangle v_1 \approx^j e_2 : T$ by Lemma 4.6.4.

(I_CONCAT) $\langle c_{11}; c_{12}; c_2 \rangle v_1 \approx^j e_2 : T$. By the IH on $\vdash c_{11}$ **ignorable**, we find that

$$\langle c_{12}; c_2 \rangle v_1 \approx^j e_2 : T$$

By the IH on $\vdash c_{12}$ **ignorable**, we find that $\langle c_2 \rangle v_1 \approx^j e_2 : T$.

□

We prove a similar lemma that *failable* coercions always fail. The L_* rules try to capture the logic of similarly named N_FAIL^* rules.

4.6.7 Lemma: If $\vdash c_1$ **failable**, then $\langle c'_1; c_1; c'_2 \rangle v_1 \approx^j e_2 : T$.

Proof: We begin by using type soundness (Theorem 4.3.9) to find that either $\langle c'_1; c_1; c'_2 \rangle v_1 \longrightarrow_n^* \langle c_1; c'_2 \rangle v'_1$, the whole left-hand side reduces to **fail** (and we are done by Lemma 4.6.5), or the whole left-hand side diverges (and we are done by definition). In the last two cases we are done, so we consider the first case.

We proceed by induction on $\vdash c_1$ **failable**.

(L_FAIL) We step by F_FAIL, and have **fail** $\approx^j e_2 : T$ by Lemma 4.6.5.

(L_BB) We step by F_TAGB followed by F_TAGBFAILB, and then we have **fail** $\approx^j e_2 : T$ by definition; we finish by expansion (Lemma 4.6.3).

(L_BFUN) We step by F_TAGB followed by F_TAGBFAILFUN, and we have fail $\approx^j e_2 : T$ by Lemma 4.6.5; we finish by expansion (Lemma 4.6.3).

(L_FUNB) We step by F_TAGFUN followed by F_TAGFUNFAILB, and we have fail $\approx^j e_2 : T$ by Lemma 4.6.5; we finish by expansion (Lemma 4.6.3)

□

With ignorable and failable coercions, we can characterize *all* non-canonical coercions, relating them to canonical coercions. The relation $\vdash c_1 \equiv c_2$ relates a non-canonical coercion c_1 to a canonical coercion c_2 . Note that this inductively defined relation isn't the same thing as the logical relation—it's a separate invariant relation, used to relate coercions to their canonical forms. Only when relating refinements does this invariant relation dip into the logical relation, to relate the predicates. First we show that well typed coercions c in NAIVE are related to $\mathbf{canonical}(c)$ in EFFICIENT. Then we'll use this general relation to relate *in the logical relation* how coercion forms work on logically related values.

We now show that *any* coercions $\vdash c_1 \equiv c_2$ yield related results when applied to related values. I defined the relation $\vdash c_1 \equiv c_2$ because this lemma is easier to prove on the relation than on the $\mathbf{canonical}$ function itself.

4.6.8 Lemma [Relating canonical coercions]:

If $v_1 \approx^j v_2 : T_1$ and $\vdash c_1 : T_1 \rightsquigarrow T_2$ and $\vdash c_1 \equiv c_2$, then $\langle c_1 \rangle v_1 \approx^j \langle c_2 \rangle v_2 : T_2$.

Proof: By induction on $\vdash c_1 \equiv c_2$ using ignorability (Lemma 4.6.6) and failability (Lemma 4.6.7) extensively.

(R_ID) Observe that $T_1 = T_2$. By ignorability (Lemma 4.6.6), $\langle c_1 \rangle v_1 \longrightarrow_n^* \langle \mathbf{ld} \rangle v_1$; by F_TAGID, $\langle \mathbf{ld} \rangle v_1 \longrightarrow_n v_1$. On the right-hand side, E_TAGID steps $\langle \mathbf{ld} \rangle v_2 \longrightarrow v_2$, and we have $v_1 \approx^j v_2 : T_1$ by assumption. We are done by expansion (Lemma 4.6.3).

(R_FAIL) By failability (Lemma 4.6.7).

(R_COMPOSITE) We want to show

$$\langle c'_0; c_1; c'_1; c_2; \dots; c'_{n-1}; c_n; c'_n \rangle v_1 \approx^j \langle d_1; \dots; d_n \rangle v_2 : T_2$$

given that $\vdash c'_i$ **ignorable** and $\vdash c_i \equiv d_i$.

We go by induction on n . The $n = 0$ case is already covered by R_ID above.

We first use ignorability (Lemma 4.6.6), and we need to show:

$$\langle c_1; c'_1; c_2; \dots; c'_{n-1}; c_n; c'_n \rangle v_1 \approx^j \langle d_1; \dots; d_n \rangle v_2 : T_2$$

We now go by cases on $\vdash c_1 \equiv d_1$ for the first hypothesis.

(R_FAIL) By failability (Lemma 4.6.7).

(R_TAG) By well typing of the LR, the expressions must have the same type $T_1 = B$ (if the tag is $B!$), $T_1 = \text{Dyn} \rightarrow \text{Dyn}$ (if the tag is **Fun!**), or $T_1 = \{x:T_1 \mid e\}$ (if the tag is $\{x:T_1 \mid e\}!$).

Whatever the case, we know by $v_1 \approx^j v_2 : T_1$ that v_1 and v_2 are similarly tagged, so we can step each side by F_TAGB, F_TAGFUN, or F_PREDPRED on the left (and E_TAG on the right). We then finish by the first IH and expansion (Lemma 4.6.3).

(R_UNTAG) By well typing of the LR, the expressions must have the same type $T_1 = \text{Dyn}$ (if the tag is $B?$ or **Fun?**), $T_1 = B$ (if the tag is $\{x:B \mid e\}?$), or $T_1 = \text{Dyn}$ (if the tag is $\{x:\text{Dyn} \mid e\}?$).

In the first case, we step by one of the following pairs on each side:

$$\begin{array}{c} \text{F_TAGBB/E_TAG} \\ \text{F_TAGFUNFUN/E_TAG} \\ \text{F_TAGBFAILB/E_TAGFAIL} \\ \text{F_TAGBFAILFUN/E_TAGFAIL} \\ \text{F_TAGFUNFAILB/E_TAGFAIL} \end{array}$$

For the first two, we are done by value relation we have and the first IH and expansion (Lemma 4.6.3); for the latter three, we are done by having **fail** on the left (Lemma 4.6.5).

In the last two cases, both sides step to checking forms, which we know coterminate at *all* indices, and we finish by the first IH.

(R_FUN) We must show that if:

- * $v_1 \approx^j v_2 : T_1 \rightarrow T_2$,
- * $\vdash ((c_{1n1}; \dots; c_{111}) \mapsto (c_{112}; \dots; c_{1n2})) : (T_1 \rightarrow T_2) \rightsquigarrow (T'_1 \rightarrow T'_2)$, and
- * $\vdash ((c_{1n1}; \dots; c_{111}) \mapsto (c_{112}; \dots; c_{1n2})) \equiv c_{21} \mapsto c_{22}$

then $\langle (c_{111} \mapsto c_{112}); \dots; (c_{1n1} \mapsto c_{1n2}) \rangle v_1 \approx^j \langle c_{21} \mapsto c_{22} \rangle v_2 : T'_1 \rightarrow T'_2$. In order to reduce the already heavy bookkeeping, I'll skip over the ignorable coercions, implicitly applying Lemma 4.6.6.

Let $m < j$ be given such that $v'_1 \approx^m v'_2 : T'_1$. We repeatedly unwrap the tags on the left-hand side, applying F_MERGE. On the right-hand side, we have $v_2 = u_{2c}$, where c is either **ld** or $c'_{21} \mapsto c'_{22}$. In either case, we run E_MERGE, yielding either $u_{c_{21} \mapsto c_{22}}$ or $u_{(c'_{21} \downarrow c_{21} \mapsto c'_{22} \downarrow c_{22})}$.

In the first case, where $c = \text{ld}$, we must show (after unwrapping the left n times):

$$\langle c_{112}; \dots; c_{1n2} \rangle (v_1 (\langle c_{1n1}; \dots; c_{111} \rangle v'_1)) \approx^m \langle c_{22} \rangle (v_2 (\langle c_{21} \rangle v'_2)) : T'_2$$

This is accomplished straightforwardly by deconstructing our relation on the (co-)domain coercions and using the IH on the inner and outer cases.

In the second case, where $c = c'_{21} \mapsto c_{22}$, we must show (after unwrapping the left n times):

$$\langle c_{112}; \dots; c_{1n2} \rangle (v_1 (\langle c_{1n1}; \dots; c_{111} \rangle v'_1)) \approx^m \langle c'_{22} \Downarrow c_{22} \rangle (u_{21d} (\langle c_{21} \Downarrow c'_{21} \rangle v'_2)) : T'_2$$

The right-hand side expands to $\langle c_{22} \rangle (\langle c'_{22} \rangle (u_{21d} (\langle c'_{21} \rangle (\langle c_{21} \rangle v'_2))))$. We know that $\langle c_{1n1}; \dots; c_{111} \rangle v'_1 \approx^m \langle c_{21} \rangle v'_2 : T_1$ by the IH, so they reduce to $v''_1 \approx^m v''_2 : T_1$ (or the left-hand side raises blame or diverges, in which case we are done.)

By assumption, $v_1 v''_1 \approx^m v_2 v''_2 : T_2$; in particular, this means that $v_1 v''_1 \approx^m \langle c'_{22} \rangle (u_{21d} (\langle c'_{21} \rangle v''_2)) : T_2$. By Lemma 4.3.13 and Lemma 4.5.22, we can see that $v_1 (\langle c_{1n1}; \dots; c_{111} \rangle v'_1)$ either:

- * diverges, and we are done;
- * raises blame, and we are done; or
- * reduces to $v'''_1 \approx^m v'''_2 : T_2$ such that $(\langle c'_{22} \rangle (u_{21d} (\langle c'_{21} \rangle (\langle c_{21} \rangle v'_2)))) \longrightarrow^* v'''_2$.

By the deconstructing our relationship on the codomains of the function contracts, we know that $\langle c_{112}; \dots; c_{1n2} \rangle v'''_1 \approx^m \langle c_{22} \rangle v'''_2 : T'_2$. Since we got here through reduction, we can finish by expansion (Lemma 4.6.3).

□

In Chapter 2, a similar characterization of casts is sufficient: once we'd related λ_C contracts and λ_H casts, we had what we needed to handle the corresponding cases of the final proof. But that strategy won't work here: congruent reductions under coercions may introduce coercions on the outside; these extra coercions will merge into c_1 and c_2 (by F_MERGE or E_MERGE), possibly disrupting $\vdash c_1 \equiv c_2$. Consider the T_MERGE case of the proof of soundness. Even if we have $\langle c_1 \rangle v_1 \approx^j \langle c_2 \rangle v_2 : T$ and $e_1 \longrightarrow_n^* v_1$ and $e_2 \longrightarrow^* v_2$, we can't just put the derivations in and be done: what if e_1 or e_2 produce terms like $\langle c'_1 \rangle e'_1$ or $\langle c'_2 \rangle e'_2$ as they evaluate? Then F_MERGE or E_MERGE will fire, and we won't know anything about the related coercions, nor do we know how many extra steps may have been taken. Accounting for the steps, it turns out, is not particularly hard: if too many new steps are added, the terms are vacuously in the relation; if not, then we merely need to account for the extra merged coercions.

In short: having Lemma 4.6.8 on values doesn't immediately tell us anything about how coercions work on arbitrary terms; we must prove that separately.

4.6.9 Lemma [Relating coercions with merges]: If $e_1 \approx^j e_2 : T_1$ and $\vdash c_1 : T_1 \rightsquigarrow T_2$ and $\vdash c_1 \equiv c_2$, then $\langle c_1 \rangle e_1 \approx^j \langle c_2 \rangle e_2 : T_2$.

Proof: First, we can ignore the cases where $e_1 \longrightarrow_n^m$ fail or e_1 diverges—those are immediately related, since $\langle c_1 \rangle e_1$ also diverges or goes to fail (by Lemma 4.3.13).

So $e_1 \longrightarrow_n^m v_1$ and $e_2 \longrightarrow^* v_2$, and by definition, $v_1 \approx^j v_2 : T_1$. By coercion congruence (Lemma 4.3.13 and Lemma 4.5.22), there exist e'_1 and e'_2 such that (a)

$\langle c_1 \rangle e_1 \longrightarrow_n^* e'_1$ and $\langle c_1 \rangle v_1 \longrightarrow_n^* e'_1$, and (b) $\langle c_2 \rangle e_2 \longrightarrow^* e'_2$ and $\langle c_2 \rangle v_2 \longrightarrow^* e'_2$. But we know that $\langle c_1 \rangle v_1 \approx^j \langle c_2 \rangle v_2 : T_2$ by Lemma 4.6.8, so we are done by contraction and expansion (Lemma 4.6.4 and Lemma 4.6.3). \square

4.6.10 Theorem [Soundness]: • If $\Gamma \vdash u : T$ then $\Gamma \vdash u_{\text{ld}} \approx \text{canonical}(u_{\text{ld}}) : T$.

- If $\Gamma \vdash e : T$ then $\Gamma \vdash e \approx \text{canonical}(e) : T$.
- If $\vdash c : T_1 \rightsquigarrow T_2$ then $\vdash c \equiv \text{canonical}(c)$.
- If $\vdash T$ then $\forall j. T \succ^j \text{canonical}(T)$.

Proof: By lexicographic induction on the typing derivation and the size of the term (v or e , respectively), using Lemma 4.6.9 in the `T_COERCE` case. We use Lemma 4.6.8 in the `T_TAGVAL`, `T_TAGVALREFINE`, and coercion cases.

Let a j be given. In all cases, we begin by letting $\Gamma \models_{\prec}^j \delta$, so we must show that $\delta_1(e) \approx^j \delta_2(\text{canonical}(e)) : T$.

(`T_CONST`) Immediate— $k_{\text{ld}} \succ^j k_{\text{ld}} : B$ for any j .

(`T_ABS`) We must show that

$$\delta_1(\lambda x : T_1. e_{1\text{ld}}) \succ^j \delta_2(\lambda x : \text{canonical}(T_1). \text{canonical}(e_1)_{\text{ld}}) : T_1 \rightarrow T_2$$

given that $\Gamma, x : T_1 \vdash e_1 : T_2$. Let $m < j$, and let $v_1 \succ^m v_2 : T_1$. We have

$$\begin{aligned} \delta_1(\lambda x : T_1. e_{1\text{ld}}) v_1 &\longrightarrow_n \delta_1(e_1)[v_1/x] \\ \delta_2(\lambda x : \text{canonical}(T_1). \text{canonical}(e_1)_{\text{ld}}) v_2 &\longrightarrow \delta_2(\text{canonical}(e_1))[v_2/x] \end{aligned}$$

We must now show that these two terms are related at m . We can apply IH (4.6.10) on $\Gamma, x : T_1 \vdash e_1 : T_2$ at the index m using the closing substitution $\Gamma, x : T_1 \models_{\prec}^m \delta[v_1, v_2/x]$.

(`T_PREVAL`) By IH (4.6.10).

(`T_TAGVAL`) We must show that

$$\delta_1(v_{1d}) \succ^j \delta_2(u_{1c} \Downarrow \text{canonical}(d)) : T_2$$

where $\text{canonical}(v_1) = u_{1c}$ given that $\delta_1(v_1) \approx^j \delta_2(\text{canonical}(u_1)_c) : T_1$, knowing that $d \neq \{x : T_1 \mid e\}$? and $\vdash d : T_1 \rightsquigarrow T_2$

By IH (4.6.10), we find that $\vdash d \equiv \text{canonical}(d)$. We can then apply Lemma 4.6.8 to find that $\langle d \rangle \delta_1(v_1) \approx^j \langle \text{canonical}(d) \rangle \delta_2(\text{canonical}(u_1)_c) : T_2$. Since we originally had v_{1d} , we know that d must be a value tag, we can step both sides (by `F_TAGB`, `F_TAGFUN`, `F_TAGFUNWRAP` or `F_TAGPREDPRED` on the left; `E_TAG` or `E_FUN` on the right) to find

$$\begin{aligned} \langle d \rangle \delta_1(v_1) &\longrightarrow_n^* \delta_1(v_{1d}) \\ \langle \text{canonical}(d) \rangle \delta_2(\text{canonical}(u_1)_c) &\longrightarrow^* \delta_2(\text{canonical}(u_1)_{c \Downarrow \text{canonical}(d)}) \end{aligned}$$

We are then done by expansion (Lemma 4.6.3).

(T_TAGVALREFINE) As for the previous case, stepping through related checking forms to find related results. (Or, possibly, by finding fail on the left and ignoring the right entirely, by Lemma 4.6.5.)

(T_VAR) By definition, $\delta_1(x) \approx^j \delta_2(x) : T$.

(T_OP) By IH (4.6.10) we can either find divergence or failure on the left, or each argument reduces to a value. In this case, we know that operations are first-order and don't take dynamic values, so we get the exact same output on both sides—which must then be related.

(T_APP) By IH (4.6.10) and the definition of the logical relation at function types.

(T_COERCE) By Lemma 4.6.9 and IH (4.6.10).

(T_FAIL) Immediate by definition.

(T_CHECK) By IH (4.6.10).

(C_ID) By R_ID.

(C_FAIL) By R_FAIL, L_FAIL, and I_ID.

(C_COMPOSE) We have

$$\begin{aligned} \vdash d_1 &\equiv \text{canonical}(d_1) \\ \vdash d_2; \dots; d_n &\equiv \text{canonical}(d_2; \dots; d_n) \end{aligned}$$

It remains to show that concatenation on the left is related to merging on the right.

We go by cases on d_1 . Throughout the analysis, we will examine the rule used to find $\vdash d_2; \dots; d_n \equiv \text{canonical}(d_2; \dots; d_n)$ —when the rule was R_FAIL, we are immediately done, since the merge will produce Fail on the right and the left will always satisfy R_FAIL.

($d_1 = B!$) Either $\text{canonical}(d_1) \Downarrow \text{canonical}(d_2; \dots; d_n)$ begins with $B!$ or it doesn't. If it does, then we invert $\vdash d_2; \dots; d_n \equiv \text{canonical}(d_2; \dots; d_n)$:

(R_ID) We are done by R_COMPOSITE.

(R_FAIL) We are done by R_FAIL.

(R_COMPOSITE) We are done by R_COMPOSITE, with $c'_0 = \text{Id}$ and $c_1 = B!$ and $c'_1; \dots; c'_n = d_2; \dots; d_n$.

If it doesn't, then we go by inversion of the derivation of

$$\text{canonical}(d_1) * \text{canonical}(d_2; \dots; d_n) \Rightarrow \dots$$

We exclude obviously contradictory cases (N_CANONICAL, N_FAILL, N_PHI with $D = \mathbf{Fun}$ or $D = \{x:T \mid e\}$, N_FUNFAILB, N_PREDPRED, N_FUN).

(N_FAILR) The only relation rule that could have applied is R_FAIL, so we are done again by R_FAIL.

(N_PHI with $D = B$) By R_COMPOSITE, since $\vdash B!$; c ; $B?$ **ignorable**.

(N_BFAILB) By R_FAIL and L_BB.

(N_BFAILFUN) By R_FAIL and L_BFUN.

($d_1 = \mathbf{Fun}!$) As for $B!$.

($d_1 = \{x:T \mid e\}!$) If d_1 isn't at the beginning of the coercion, then it must be that $\{x:T \mid e\}?$ begins $\mathbf{canonical}(d_2; \dots; d_n)$.

If d_1 remains at the beginning, then we go by R_COMPOSITE or R_FAIL, depending on how we found the derivation of $\vdash d_2; \dots; d_n \equiv \mathbf{canonical}(d_2; \dots; d_n)$: we use the former if R_ID or R_COMPOSITE was used, the latter if R_FAIL was used.

We go by cases on the derivation of $\vdash d_2; \dots; d_n \equiv \mathbf{canonical}(d_2; \dots; d_n)$.

(R_ID) Contradictory—we assumed that the final merge result began with $\{x:T \mid e\}?$.

(R_FAIL) By R_FAIL.

(R_COMPOSITE) By R_COMPOSITE, noting that

$$\vdash \{x:T \mid e\}!; c'_0; \{x:T \mid e\}?$$
 ignorable

by I_PREDSAME.

($d_1 = B?$) It must be the case that the merged coercion starts with $B?$, since no merge eliminates $B?$ on the left.

If R_ID or R_COMPOSITE derived $\vdash d_2; \dots; d_n \equiv \mathbf{canonical}(d_2; \dots; d_n)$, we are done by R_COMPOSITE. If R_FAIL was used, we're done immediately.

($d_1 = \mathbf{Fun}?$) As for $B?$.

($d_1 = \{x:T \mid e\}?$) The outermost rule applying to find $\{x:T \mid \mathbf{canonical}(e)\}?$ \Downarrow $\mathbf{canonical}(d_2; \dots; d_n)$ must be either N_FAILR, N_CANONICAL, or N_PREDPRD. In the first two cases, we are done (by R_COMPOSITE, R_TAG, and R_DPRED—or by R_FAIL). In the third case, we are done by the IH with R_COMPOSITE and R_TAG, with $c'_0 = \text{ld}$. In the fourth case, we are done by the IH (since $\vdash \{x:T \mid e\}?$; $\{x:T \mid e\}!$ **ignorable**, and we can use either R_ID or R_COMPOSITE with I_CONCAT).

($d_1 = c_{11} \mapsto c_{12}$) The outermost rule applying in $\mathbf{canonical}(c_{11}) \mapsto \mathbf{canonical}(c_{12}) \Downarrow \mathbf{canonical}(d_2; \dots; d_n)$ must be one of N_FAILR, N_CANONICAL, or N_FUN.

In the first two cases we are done, either by R_COMPOSITE with R_FUN or by R_FAIL. In the third case, we are done by the IH with R_COMPOSITE and R_FUN. In the fourth case, we are done by the IH, noticing that the relation on the right comes from R_COMPOSITE with an initial R_FUN, so we can fold $c_{11} \mapsto c_{12}$ into that R_FUN derivation and reconstruct a new R_COMPOSITE derivation.

(WF_DYN) Immediate

(WF_BASE) Immediate

(WF_FUN) By IH (4.6.10).

(WF_REFINE) We have $T \succ^j T$ immediately. As IH (4.6.10) on $x:T \vdash e : \mathbf{Bool}$, we have $x:T \vdash e \approx \mathbf{canonical}(e) : \mathbf{Bool}$. Let $m < j$ and $v_1 \succ^m v_2 : T$ be given; we can then find $e[v_1/x] \approx^m \mathbf{canonical}(e)[v_2/x] : \mathbf{Bool}$ by instantiating at m and δ .

□

The definition of $\emptyset \vdash e_1 \approx e_2 : T$ gives us approximate observational equivalence: either e diverges, e reduces to **fail**, or $e \rightarrow_n^* v_1$ and $\mathbf{canonical}(e) \rightarrow^* v_2$ such that $v_1 \succ^j v_2 : T$ for arbitrary j . (I write this result without indices because the definition of $\Gamma \vdash e_1 \approx e_2 : T$ quantifies over all indices.) Note that for base values, we have exactly the same result on both sides.

4.7 Space efficiency

The structure of my space-efficiency proof is largely the same as in prior work. Coercion size is broken down by the order of the types involved; the maximum size of any coercion is $|\mathbf{largest\ coercion}| \cdot 2^{\mathbf{tallest\ type}}$. Inspecting the canonical coercions, the largest is $\{x:\mathbf{Dyn} \mid e\}!; \mathbf{Fun}?; c_1 \mapsto c_2; \mathbf{Fun}!; \{x:\mathbf{Dyn} \mid e'\}?$, with a size of 5. The largest possible canonical coercion therefore has size $M = 5 \cdot 2^h$.

Formally, observe that merging canonical coercions c_1 and c_2 either produces a smaller coercion or $c_1; c_2$ is canonical (and has size $\mathbf{size}(c_1) + \mathbf{size}(c_2)$).

4.7.1 Lemma [Merge reduces size]: If $c_1 * c_2 \Rightarrow c_3$, then either:

- $\mathbf{size}(c_1) + \mathbf{size}(c_2) > \mathbf{size}(c_3)$, or
- $c_3 = c_1; c_2$ is canonical.

Proof: By induction on the derivation of $c_1 * c_2 \Rightarrow c_3$. Note that in either case, $\mathbf{size}(c_1) + \mathbf{size}(c_2) \geq \mathbf{size}(c_3)$.

(N_CANONICAL) $c_1; c_2$ is canonical.

(N_FAILL) We have $\mathbf{Fail} * c_2 \Rightarrow \mathbf{Fail}$, with $1 + \mathbf{size}(c_2) > 1$.

(N_FAILR) We have $c_1 * \mathbf{Fail} \Rightarrow \mathbf{Fail}$, with $\mathbf{size}(c_1) + 1 > 1$.

(N_PHI) We have $\vdash c_1; D! : T_1 \rightsquigarrow \mathbf{tgt}(D)$ and $\vdash D?; c_2 : \mathbf{tgt}(D) \rightsquigarrow T_3$. By the IH, $\mathbf{size}(c_1) + \mathbf{size}(c_2) \geq \mathbf{size}(c_3)$, so we immediately have $\mathbf{size}(c_1) + 1 + \mathbf{size}(c_2) + 1 > \mathbf{size}(c_3)$.

(N_BFAILB) We have $\vdash c_1; B! : T_1 \rightsquigarrow \text{Dyn}$ and $\vdash B'?; c_2 : \text{Dyn} \rightsquigarrow T_3$. It is immediate that $\text{size}(\text{Fail})$ is smaller.

(N_BFAILFUN) We have $\vdash c_1; B! : T_1 \rightsquigarrow \text{Dyn}$ and $\vdash \mathbf{Fun}?; c_2 : \text{Dyn} \rightsquigarrow T_3$. It is immediate that $\text{size}(\text{Fail})$ is smaller.

(N_FUNFAILB) We have $\vdash c_1; \mathbf{Fun}! : T_1 \rightsquigarrow \text{Dyn}$ and $\vdash B'?; c_2 : \text{Dyn} \rightsquigarrow T_3$. It is immediate that $\text{size}(\text{Fail})$ is smaller.

(N_PREDPRED) We have $\vdash c_1; \{x:T \mid e\}? : T_1 \rightsquigarrow \{x:T \mid e\}$ and $\vdash \{x:T \mid e\}!; c_2 : \{x:T \mid e\} \rightsquigarrow T_3$, where T is either B or Dyn . By the IH, $\text{size}(c_1) + \text{size}(c_2) \geq \text{size}(c_3)$, so we immediately have $\text{size}(c_1) + 1 + 1 + \text{size}(c_2) \geq \text{size}(c_3)$.

(N_FUN) We have:

$$\begin{aligned} \vdash c_1; (c_{11} \mapsto c_{12}) : T_1 \rightsquigarrow (T_{21} \rightarrow T_{22}) \\ \vdash (c_{21} \mapsto c_{22}); c_2 : (T_{21} \rightarrow T_{22}) \rightsquigarrow T_3 \end{aligned}$$

By the IH:

$$\begin{aligned} \text{size}(c_{21}) + \text{size}(c_{11}) &\geq \text{size}(c_{31}) \\ \text{size}(c_{12}) + \text{size}(c_{22}) &\geq \text{size}(c_{33}) \\ \text{size}(c_1) + \text{size}((c_{31} \mapsto c_{32}); c_2) &\geq \text{size}(c_3) \end{aligned}$$

Now we can see that:

$$\begin{aligned} \text{size}(c_{11} \mapsto c_{12}) + \text{size}(c_{21} \mapsto c_{22}) &= 1 + \text{size}(c_{11}) + \text{size}(c_{12}) + \\ &\quad 1 + \text{size}(c_{21}) + \text{size}(c_{22}) \\ &> 1 + \text{size}(c_{31}) + \text{size}(c_{32}) \end{aligned}$$

Now we can finally conclude that:

$$\begin{aligned} \text{size}(c_1; (c_{11} \mapsto c_{12})) + \text{size}((c_{21} \mapsto c_{22}); c_2) &> \text{size}(c_1) + \text{size}((c_{31} \mapsto c_{32}); c_2) \\ &> \text{size}(c_3) \end{aligned}$$

□

Rules with merges (and E_MERGE in particular) don't increase the size of the largest coercion in the program. Applying this lemma across an evaluation $e \longrightarrow^* e'$, we can see that no coercion ever exceeds the size of the largest coercion in e . If M is the size of the largest coercion, then there is at most an M -fold space overhead of coercions. But this size bound is galactic; I find it hard to believe that this overhead is observable in practice. A much more interesting notion of space efficiency—not studied here—is to determine implementation schemes for space-efficient layout of coercions in memory and time-efficient merges of coercions. I believe that explicitly enumerating the canonical coercions is a step towards this goal: the canonical coercions in Table 4.2 are exactly those which must be represented.

4.8 Conclusion

Space-efficiency is attainable for contract languages that cover the whole spectrum, from dynamic types to refinement types. While canonical coercions may skip checks—and behave slightly differently from naïve, inefficient implementations—we still have the guarantee that when the naïve implementation produces a value, the space-efficient one will produce a behaviorally equivalent one. Resolving the space inefficiency of naïve contract systems is an important first step in making contracts amenable to pervasive use.

Chapter 5

Related work

Literature does not exist in a vacuum.

ABC of Reading
Ezra Pound

We begin by surveying the field in Section 5.1, with a focus on λ_H in Section 5.1.2. We then focus more closely on work related to F_H and space-efficiency in Sections 5.2 and 5.3, respectively.

5.1 Contracts: a survey

Conferences in recent years have seen a profusion of papers on higher-order contracts and related features. This is all to the good, but for newcomers to the area it can be a bit overwhelming, especially given the great variety of technical approaches. To help reduce the level of confusion, in Table 5.1 I summarize the important points of comparison between a number of systems that are closely related to ours. This table is an updated version of those in Greenberg et al. [34, 35]. A similar comparison for gradual typed systems—which necessarily has some overlap with the material in this table—is in Table 5.2.

The largest difference is between latent and manifest treatments of contracts—i.e., whether contract checking (under whatever name) is a completely dynamic matter or whether it leaves a “trace” that the type system can track.

Another major distinction (labeled “dep” in the figure) is the presence of dependent contracts or, in manifest systems, dependent function types. Latent systems with dependent contracts also vary in whether their semantics is lax or picky.

Next, most contract calculi use a standard call-by-value order of evaluation (“eval order” in the figure). Notable exceptions include those of Hinze et al. [41], which is embedded in Haskell, Flanagan [28], which uses a variant of call-by-name, and Knowles and Flanagan [44], which uses full β -reduction (more on this below).

Latent systems						
	FF02	HJL06	GF07 λ_C	BM06	DFFF11	λ_C
	(1)	(2)	(3)	(4)	(5)	(Ch. 2)
dep (6)	✓ lax	✓ picky	×	(7)	✓ indy	✓ either
eval order	CBV	lazy	CBV	CBV	CBV	CBV
blame (8)	$\uparrow l$	$\uparrow l$	$\uparrow l$	$\uparrow l$ or \perp	$\uparrow l$	$\uparrow l$
checking (9)	if	if	\bigcirc	active	active	active
typing (10)	✓	✓	✓	n/a	✓	✓
any con (11)	✓	✓	✓	✓	✓	✓

Manifest systems									
	GF07 λ_H	F06	KF10	WF09	OTMW04	BGIP11	λ_H	F_H	EFF.
	(3)	(12)	(13)	(14)	(15)	(16)	(Ch. 2)	(Ch. 3)	(Ch. 4)
dep (6)	×	✓	✓	×	✓	✓	✓	✓	×
eval order	CBV	CBN(17)	full β	CBV	CBV	CBV	CBV	CBV	CBV
blame (8)	$\uparrow l$	stuck	stuck	$\uparrow l$	\uparrow	$\uparrow l$	$\uparrow l$	$\uparrow l$	\uparrow
checking (9)	\bigcirc	\bigcirc	active	active	if	active	active	active	active
typing (10)	×	×	✓	✓	✓	✓	✓	✓	✓
any con (11)	✓	✓	✓	✓	×	✓	✓	✓	✓
poly (18)	×	×	×	×	×	✓	×	✓	×
space (19)	×	×	×	×	×	×	×	×	✓

(1) Findler and Felleisen [26]. (2) Hinze et al. [41]. (3) Gronski and Flanagan [36]. (4) Blume and McAllester [11]. (5) Dimoulas et al. [22]. (12) Flanagan [28]. (13) Knowles and Flanagan [44]. (14) Wadler and Findler [78]. (15) Ou et al. [51]. (16) Belo et al. [8]. This is not the same as the F_H offered in Chapter 3. I omit CAST and NAIVE from this analysis. (6) Does the system include dependent contracts or function types (✓) or not (×) and, for latent systems, is the semantics lax or picky? (See below for more on “indy” checking.) (7) An “unusual” form of dependency, where negative blame in the codomain results in nontermination. (17) A nondeterministic variant of CBN. (8) Do failed contracts raise labeled blame ($\uparrow l$), raise blame without a label (\uparrow), get stuck, or sometimes raise blame and sometimes diverge (\perp)? (9) Is contract or cast checking performed using an “active check” syntactic form (active), an “if” construct with a refined typing rule (if), or “inlined” by making the operational semantics refer to its own reflexive and transitive closure (\bigcirc)? (10) Is the typing relation monotonic, i.e., is the typing relation known to be uniquely defined? (11) Are arbitrary user-defined boolean functions allowed as contracts or refinements (✓), or only built-in ones (×)? (18) Does the type system support polymorphism? (19) Is the operational semantics space efficient?

Table 5.1: Comparison between contract systems

Another point of variation (“blame” in the figure) is how contract violations or cast failures are reported—by raising an exception or by getting stuck. I return to this below.

The next two rows in the table (“checking” and “typing”) concern more technical points in the papers most closely related to the work in Chapter 2. In both Gronski and Flanagan [36] and Flanagan [28], the operational semantics checks casts “all in one go”:

$$\frac{s_2\{x := k\} \rightarrow_h^* \mathbf{true}}{\langle \{x:B \mid s_1\} \Rightarrow \{x:B \mid s_2\} \rangle^l k \rightarrow_h k}$$

Such rules are formally awkward, and in any case they violate the spirit of a small-step semantics. Also, the formal definitions of λ_H in both Gronski and Flanagan [36] and Flanagan [28] involve a circularity between the typing, subtyping, and implication relations. Knowles and Flanagan [44] improve the technical presentation of λ_H in both respects. In particular, they avoid circularity (as I do) by introducing a denotational interpretation of types and maintain small-step evaluation by using a new syntactic form of “partially evaluated casts” (like most of the other systems).

5.1.1 Refinement types and contracts

Refinement types, in fact, first referred to refinements of datatypes in particular [30], though the term has been appropriated to mean types of the form $\{x:T \mid e\}$ in general. We discuss the connection between our manifest contracts and datatypes more in Section 5.3.

Findler and Felleisen’s contracts—and those in this dissertation—have runtime effects and are *not* erasable in general. That is, the default semantics is to check all of these specifications as the program runs. Runtime checking contrasts markedly with refinement type theories that use SMT (satisfiability modulo theories) solvers, where there is little or no checking at runtime. The work on liquid types is one example [58, 59]. Others go further, including type Dyn: F* [71, 29] and DJS [19, 17]. All of these languages use SMT solvers to resolve refinement types written as propositions in a logic—while I use code. The logics of many of these languages allow first-order quantification, which rules out run-time checking as a possibility. In general, languages with propositional refinement types begin to push into the territory of program verification and program logics. This is a broad field in its own right, and I do not attempt to survey it here.

I should contrast the SMT-solver approach with the Curry–Howard approach taken in, e.g., Coq [21]. There are many differences, but the chief one is in the nature of evidence: refinement types and contracts use SMT solvers or other extensional “proof on the side” methods, while Coq and its brethren have intensional proof languages where the programmer constructs proofs like programs (and vice versa). Introducing impurities to the Curry–Howard style is an active area of research, but the current state of the art favors the extensional approach (e.g., F*).

Some work has been done on optimizing away unnecessary contract checks. In Chapter 3, I prove the soundness of *upcast elimination* (Lemma 3.4.3); other such optimizations have been the subject of some study already [44, 8]. In Chapter 4, I optimize the space usage of contracts in a way that requires us to (soundly) skip some checks. Findler et al. [27] (discussed more in Section 6.1.2) reduce re-checking of contracts on data structures.

5.1.2 Situating λ_H

The formulation of λ_H in Chapter 2 is most comparable to that of Knowles and Flanagan [44], but there are some significant differences. First, my cast-checking constructs are equipped with labels, and failed casts go to explicit blame—i.e., they raise labeled exceptions. In the λ_H of Knowles and Flanagan (though not the earlier one of Gronski and Flanagan), failed casts are simply stuck terms—their progress theorem says “If a well-typed term cannot step, then either it is a value or it contains a stuck cast.” Second, their operational semantics uses full, non-deterministic β -reduction, rather than specifying a particular order of reduction, as I have done. This significantly simplifies parts of the metatheory by allowing them to avoid introducing parallel reduction. I prefer standard call-by-value reduction because I consider blame to be an exception—a computational effect—and I want to be able to reason about *which* blame will be raised by expressions involving many casts. At first glance, it might seem that my theorems follow directly from the results for Knowles and Flanagan’s language, since CBV is a restriction of full β -reduction. However, the reduction relation is used in the type system (in rule S_IMP), so the type systems for the two languages are not the same. For example, suppose the term *bad* contains a cast that fails. In my system $\{y:B \mid \text{true}\}$ is not a subtype of $\{y:B \mid (\lambda x:S. \text{true}) \text{ bad}\}$ because the contract evaluates to blame. However, the subtyping does hold in the Knowles and Flanagan system because the predicate reduces to **true**.

The system studied by Ou et al. [51] is also close in spirit to my λ_H . The main difference is that, because their system includes general recursion, they restrict the terms that can appear in contracts to just applications involving predefined constants: only “pure” terms can be substituted into types, and these do not include lambda-abstractions. My system (like all of the others in Table 5.1—see the row labeled “any con”) allows arbitrary user-defined boolean functions to be used as contracts.

My description of λ_C is ultimately based on λ_{CON} [26], though my presentation is slightly different in its use of checks. Hinze et al. [41] adapted Findler and Felleisen-style contracts to a location-passing implementation in Haskell, using picky dependent function contracts.

My λ_H type semantics in Section 2.3.2 is effectively a semantics of contracts. Blume and McAllester [11] offers a semantics of contracts that is slightly different—my semantics includes blame at every type, while theirs explicitly excludes it. Xu et al. [80] is also similar, though their “contracts” have no dynamic semantics at all: they are simply specifications.

Dimoulas et al. [22] introduce a new dialect of picky λ_C , where contract checks in the codomain are given a distinct negative label. If labels represent “contexts” for values, then this treats the contract as an independent context. “Indy” λ_C and picky λ_C will raise exactly the same *amount* of blame, but they will blame different labels.

Chapter 3, a corrected extension of Belo et al. [8], at once simplifies and extends the CBV λ_H given in Chapter 2. The type system is redesigned to avoid subtyping and closing substitutions, so type soundness is proved with easy syntactic methods [79]. The language also allows *general refinements*—refinements of any type, not just base types—and extends the type system to polymorphism. This can be seen as completing some of the future work of Greenberg et al. [34].

I have discussed only a small sample of the many papers on contracts and related ideas. I refer the reader to Knowles and Flanagan [44] for a more comprehensive survey. Another useful resource is Wadler and Findler [77] (technically superseded by Wadler and Findler [78], but with a longer related work section), which surveys work combining contracts with type Dyn and related features.

There are also *many* other systems that employ various kinds of precise types, but in a completely static manner. One notable example is the work of Xu et al. [80], which uses user-defined boolean predicates to classify values (justifying their use of the term “contracts”) but checks statically that these predicates hold.

Sage [45] and Knowles and Flanagan [44] both support mixed static and dynamic checking of contracts, using, e.g., a theorem prover. I have not addressed this aspect of their work, since I have chosen to work directly with the core calculus λ_H , which for them was the target of an elaboration function.

5.2 F_H : polymorphism and manifest metatheory

I discuss work related to F_H (Chapter 3) in two parts. First, I distinguish my work from the untyped contract systems that enforce parametric polymorphism *dynamically*, rather than statically as F_H does. Then I discuss how F_H differs from existing manifest contract calculi in greater detail.

5.2.1 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 [52]).

Guha et al. [37] define contracts with polymorphic signatures, maintaining abstraction with sealed “coffers”; they do not prove parametricity. Matthews and Ahmed [46] prove parametricity for a polymorphic multi-language system with a similar policy. Neis et al. [50] use dynamic type generation to restore parametricity in the presence of intensional type analysis. F_H ’s contracts are subordinate to the type system, so

the parametricity result does not require dynamic type generation. Ahmed et al. [3] prove parametricity for a gradual typing [67] calculus which enforces polymorphism with a set of global runtime seals. Ahmed et al. [4] define a polymorphic calculus for gradual typing, using local syntactic “barriers” instead of global seals. I believe that it is possible to combine F_H with the barrier calculus of Ahmed et al., yielding a polymorphic blame calculus [78]. I leave this to future work.

5.2.2 F_H and other manifest calculi

For a coarse comparison, please refer back to Table 5.1. In this section, I give a more technical comparison with the closest related work. Four existing manifest calculi with dependent function types ([28, 34, 44, 51]) use subtyping and theorem provers as part of the definition of their type systems. All four of these calculi have complicated metatheory. Ou et al. [51] 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 [44] as well as Greenberg, Pierce, and Weirich [34] use denotational semantics to give a firm foundation to earlier work [28]. I consider three systems in more detail: Knowles and Flanagan’s λ_H (KF); Chapter 2’s λ_H (which is the same as Greenberg, Pierce, and Weirich’s λ_H , which I write here as GPW); and F_H . The rest of this subsection addresses the differences between KF, GPW, and F_H .

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 use subtyping to resolve these issues. 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 $\{x:\text{Int} \mid x > 0\}$. KF and GPW use a rule like the following for refinement subtyping:¹

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

Combined with a “constants get most specific types” requirement—for example, assigning n the type $\{x:\text{Int} \mid x = n\}$ —subtyping allows n to be typed at any predicate contract it satisfies. Second, KF and GPW use subtyping to 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? (An important question when proving preservation!) Both KF and GPW relate reduction and subtyping, showing

¹Readers familiar with the systems will recognize that I’ve folded the implication judgment into the relevant subtyping rule.

that types that reduce to each other are mutual subtypes. KF use full beta reduction throughout their system. GPW use call-by-value reduction in their operational semantics, showing that parallel reducing types are mutual subtypes, separately relating CBV and 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 the 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. Instead, I resolve it with a runtime rule that allows us to type values with any refinement they satisfy:

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

Adding this rule eliminates one use of subtyping as well as 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 at the beginning of this chapter 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 $e_2 \rightarrow e'_2$, 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.

Other consequences of subtyping

The F_H operational semantics is essentially a superset of λ_H from Chapter 2, barring some slight differences in the function cast decomposition rule. The type system, however, is not a superset: λ_H types some programs F_H does not. In particular, λ_H builds in subsumption, while F_H only has a subsumption principle *post facto*. We can, however, take a λ_H typing derivation and eliminate every occurrence of subsumption: by the upcast lemma, the two programs are equivalent, even if one of them is not well typed. That is, I have taken subsumption out of the type system and proved subsumption safe as an optimization—and, in doing so, greatly simplified the type system.

5.3 Space efficiency and gradual types

There are two threads of work related to the development of Chapter 4: a more recent line of work on gradual types, refinement types, and full-spectrum programming languages; and an older, more general line of work on coercions, which may or may not have runtime semantics. Space efficiency and representation have been studied in both settings.

5.3.1 Space efficiency, gradual typing, and refinement types

We give an overview of the field of gradual typing in 5.2. The commonality between all systems is the presence of a type like `Dyn`.

In Siek and Taha’s seminal work on gradual typing [67], space efficiency is already a concern—they point out that the canonical forms lemma has implications for which values can be unboxed (the typed ones). Herman, Tomb, and Flanagan [39] compiled a language like Siek and Taha’s into a calculus with Henglein’s coercions [38], proving a space-efficiency result with a galactic bound similar to mine. Herman et al. stop at proving that their compilation is type preserving without proving soundness of their compilation. (I compare my system to Herman et al.’s in greater detail below.) Siek, Garcia, and Taha [65] explore the design space around Herman et al.’s result,

	ACPP89	T90	CF91	H94	KTGFF06	ST06	HTF07/10	SGT09	WF09
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
blame (10)	total	\uparrow	\uparrow	stuck	\uparrow	\uparrow	\uparrow	\uparrow^l	\uparrow^l
con (11)	\times	\times	\times	\times	$\checkmark?$	\times	\times	\times	\checkmark
poly (12)	\times	\times	\checkmark	\times	\checkmark	\times	\times	\times	\times
space (13)	\times	\times	\times	\checkmark	\times	\times	\times	$\checkmark+$	\times

	BGHL10	SW10	AFSW11	SG12	G13	EFF.
	(14)	(15)	(16)	(17)	(18)	(Ch. 4)
blame (10)	\uparrow	\uparrow^l	\uparrow^l	\uparrow^l	\uparrow^l	\uparrow
con (11)	$\checkmark(\text{FO})$	\times	\times	\times	\times	\checkmark
poly (12)	\times	\times	\checkmark	\times	\times	\times
space (13)	\times	\checkmark	\times	\checkmark	$\checkmark+$	$\checkmark+$

(1) Abadi et al. [1]. (2) Thatte [73] (3) Cartwright and Fagan [15] (4) Henglein [38] (5) Knowles et al. [45] (6) Siek and Taha [67] (7) Herman et al. [39, 40] (8) Siek et al. [65] (9) Wadler and Findler [78] (14) Bierman et al. [10] (15) Siek and Wadler [68] (16) Ahmed et al. [4] (17) Siek and Garcia [66] (18) Garcia [31] (10) Do failed contracts raise labeled blame (\uparrow^l), raise blame without a label (\uparrow), or get stuck? NB that (1) requires that typecases be total, so there are no errors. (11) Are there contracts or refinements types? NB that (5) lacks a soundness proof, and (14) is a first-order language. (12) Does the type system support polymorphism? (13) Is the operational semantics space efficient? If so (\checkmark), is there a proof of soundness relating the space-efficient calculus to a naïve semantics ($\checkmark+$)?

Table 5.2: Comparison between gradual typing systems

this time with an observational equivalence theorem exactly relating two coercion semantics.

Siek and Wadler [68] study an alternative, cast-based formulation of space efficiency, proving tighter bounds than Herman et al. [39] and an exact observational equivalence. Their insight is that casts can be factored not merely as a “twosome” $\langle S \Rightarrow T \rangle$, but rather as a threesome: $\langle S \xrightarrow{R} T \rangle$. They maintain the invariant that S downcasts to R , and R upcasts to T ; merging casts amounts to calculating a greatest lower bound. They come up with an elegant theory of merging casts, with a detailed accounting for blame. While the mathematics is beautiful, I believe that their algorithm is overkill: Herman et al.’s journal article [40] cleanly enumerates the recursive structure of the canonical coercions for dynamic and simple types, with only 17 possible structures at the top level. Siek and Wadler’s theory is the theory of these 17 structures. Many of the solutions can be simply pre-computed and looked up in a table at runtime. I have 37 canonical coercions. I don’t study the question here, but I believe that a careful analysis would allow for very compact representations with very fast merges—by pointer comparison and table lookup when functional coercions aren’t involved. I discuss this issue further in future work (Chapter 6).

Before considering other full-spectrum languages, I compare this work to the most closely related work: Herman, Tomb, Flanagan [39, 40], Garcia [31], Siek and Garcia [66], and Henglein [38]. Henglein is trying to reason carefully about programs written in a dynamic style, rather than thinking about multi-paradigm programming (though it is clear that he knows that his work applies to “dynamic typing in a static language”). His theory of coercions has no `Fail` coercion and treats `Id` slightly differently at function types. Herman et al. adapt his calculus to match the setting of gradual types, though they never rebuild his theory. Henglein develops a general theory *characterizing* canonical coercions, but I *enumerate* them, as in Herman et al. [40].

Perhaps the biggest difference is that Henglein and Herman et al. formulate coercions as having arbitrary composition: $c_1; c_2$ is a coercion that can be used freely. As a consequence, it is somewhat difficult to reason directly about coercions in the calculus: what should $\langle (c_1 \mapsto c_2; \mathbf{Fun!}); \{x:\mathbf{Dyn} \mid e\} \rangle v$ do? Their solution is to work with coercions up to an equivalence relation that includes associativity of coercion composition; coercions normalize in a term rewriting system modulo this equivalence relation. Henglein studies some algorithmic rewriting systems. But Herman et al. don’t develop the rewriting system at all, never showing that their rewriting system is strongly normalizing, and even when they enumerate canonical coercions in their journal version [40], they do so without proof. I feel that term rewriting modulo equational theories is insufficient for guiding an implementation of a coercion calculus: the compiler needs a concrete representation for coercions and a concrete algorithm for merging them. I accordingly adopt a constrained form of coercion composition out of a desire to aid implementation, but also out of expedience: I don’t need to worry about associativity at all. I don’t believe that free composition buys anything,

anyway: I don't expect programmers to be writing coercions by hand, so ease of expression in the coercion language isn't particularly important.

Herman et al.'s calculus is a little odd: the value $\langle \text{id} \rangle v$ takes a step to v . My approach makes a clearer distinction between terms and the results that they produce. Siek and Wadler noticed a separate problem with nondeterminism at cast merges, most likely due to a mistake in defining evaluation contexts.

Siek and Garcia [66] study various interpretations and implementations of gradual typing. Their `seq-lazy` function is very similar to my merge algorithm, though our treatments of associativity are different.

Garcia's work [31] is remarkably similar to mine, though restricted to gradual types: starting with a gradually typed calculus with casts, he develops coercions and then threesomes as a series of derivations. His supercoercions are remarkably similar to my canonical coercions, though mine are complicated by the presence of refinement types. Our approaches differ, though; he says:

One might try to devise an ad hoc reassociation scheme or represent a sequence of coercions as lists, but it would necessarily involve pairwise comparisons, bidirectional search, and splicing into the middle of complex coercion expressions.

My representation is very nearly a "sequence of coercions as lists", though I abandon the standard linked-list structure. My merge algorithm roughly fits his description: I do pairwise comparisons from the inside out and concatenation, though I do not splice things in the middle nor do I do bidirectional search. Coercions and threesomes are two representations of the same idea. Which of the two implementations is best in practice is an open question. Coercion merges are amenable to table lookup, but the threesome merge operator can be memoized. Which is faster, which—if any—has a more compact representation? Which is easier to implement?

The work discussed so far consisted of calculi devised expressly for space-efficient gradual typing. Findler et al. [27] discuss space efficiency from the perspective of an implementation in PLT Racket (then PLT Scheme). Their setting—latent contracts, no type system—is rather different from the foregoing systems; they address datatypes, while the foundational calculi omit datatypes.

Considering the wider world of full-spectrum programming languages, we summarize existing solutions. None of the following are space efficient; I am the first to combine space efficiency, gradual types, and refinement types. Ou et al. [51] cover the spectrum and include dependent types, but allow only a constrained set of refinement predicates; Sage [45] covers the entire spectrum and also includes dependent types, but lacks a soundness proof; Wadler and Findler's [78] development covers dynamic types through refinements of base types; Bierman et al. [10] cover the whole spectrum but (also with dependency) only for first-order types.

5.3.2 Coercions

There are many other systems that use coercions to other ends. Henglein gives an excellent summary of work up to 1994 in the related work section of his article [38]. One of the classic uses of coercions is subtyping [13, 48]; more recent work relates subtyping and polymorphism [20]. Work on unboxing [63, 49] confronts similar issues of space efficiency. Many of these works carefully ensure that coercions are erasable, while my coercions are definitely not.

Swamy, Hicks, and Bierman [70] study coercion insertion in general, showing that their framework can encode gradual types. I haven't studied coercion insertion at all, though Swamy et al.'s framework would be a natural one to use. I am not aware of work on how coercion insertion algorithms affect space consumption, though experience with the implementation of Boomerang [12] shows that small changes in coercion insertion can affect the efficiency of checking.

Chapter 6

Conclusion and future work

Why mince words, anyway, since you are not completely real Dreyfuses, Edisons, and Napoleons? You have assumed these names vicariously, for lack of anything better. Now you will swell the numbers of many of your predecessors, those anonymous Garibaldiis, Bismarcks, and MacMahons who wander in their thousands, unacknowledged, all over the world.

Sanatorium under the Sign of the Hourglass
Bruno Schulz

Now at last without fantasies or self-deception, cut off from the mistakes and confusion of the past, grave and simple, carrying a small suitcase, getting on a bus, like girls in movies leaving home, convents, lovers, I supposed I would get started on my real life.

Lives of Girls and Women
Alice Munro

Manifest contracts, as opposed to latent contracts, are the way forward for investigating strong specifications in general purpose programming languages. The existing uses of latent contracts have been lackluster: PLT Racket uses contracts to recover simple types. I suspect that the lack of a type discipline has prevented pervasive use of strong specifications. Taking programming in a type discipline as a baseline may avoid the “contracts for types are good enough” under-specification problem found in PLT Racket.

In this dissertation, I have developed the field of manifest contracts with three contributions:

- Characterized manifest contracts in relation to latent contracts (Chapter 2);
- Developed F_H , a manifest contract calculus with a powerful reasoning principle: relational parametricity (Chapter 3); and

- Shown how to resolve the unbounded space consumption in function proxies and on the stack that can accrue with contract checking (Chapter 4).

Each of these contributions serves language designers in different ways, but all of this goes to support a firm theoretical basis for manifest contracts. Obtaining reasoning principles and eliminating the gross inefficiencies in naïve formulations are necessary first steps in designing and implementing higher-order languages with pervasive manifest contracts.

The characterization of latent and manifest contracts outlines the field: what options does the language designer have, and how do they relate? This allows for informed choices—and, as in the example of Dimoulas et al. [22], further study. The proof technique for translation correctness—logical relations extended with an inductively defined invariant for contracts—offers theoreticians a powerful tool for metatheoretical work with contracts; I use it again in Chapter 4.

The F_H metatheory of Chapter 3 offers several things. First, its metatheory is much simpler than that of Chapter 2’s λ_H , as it eschews denotational techniques.¹ The F_H metatheory offers a cleaner and easier way to design core calculi for manifest contracts, a natural first step in language design. The F_H parametricity relation offers powerful reasoning principles to implementors and programmers alike: the upcast lemma of Section 3.4, for example, could drive an upcast elimination optimization in a compiler; the parametricity relation could be used to reason about different representations of abstract types.

Finally, the space efficiency work is necessary for a real implementation of designs that use contracts pervasively, rather than just at module boundaries like in PLT Racket. Moreover, my work in Chapter 4 reveals fundamental trade-offs in the design space: space-efficiency can be had for manifest contracts, but at the price of skipping some checks. Finally, I offer a different perspective on and new algorithms for coercions and threesomes.

I hope that language designers can, after reading this dissertation, begin to think about what it would mean to include manifest contracts in their systems: how the operational semantics might ensure space efficiency; how the type system might be kept small enough to admit easy analysis, while offering good reasoning principles; what equalities the reasoning principles permit. I view my work as digging—not even pouring—foundation for languages with manifest contracts. I have surveyed the field, and I offer workably efficient semantics and powerful principles for reasoning and type abstraction.

¹The original, flawed conception of Belo et al. [8] was still simpler—it didn’t “need” a cotermination theorem; however, it did need correction.

6.1 Future work

The work in this dissertation aims to put manifest contracts on a firm type-theoretical footing. There are several avenues of future work on the road to putting manifest contracts on a firm *pragmatic* footing. I conclude by discussing future work.

6.1.1 State and effects

The biggest subject missing from this dissertation is *state* and other forms of effects, beyond uncatchable exceptions (blame) and nontermination. Adding state to a manifest calculus is problematic. Consider the type $\{x:\text{Ref Int} \mid !x > !y\}$. When a value v has this type, what does the programmer know? When must the contract be checked? If either x or y update, we may need to recheck the contract. But what if x and y need to *both* be updated? What if there are circular dependencies between refined references? The folklore consensus is that *some* sort of limits are necessary on stateful contracts. Some work has been done on this question, in particular Shinnar [64]. Shinnar uses a notion of delimited (transactional memory) transactions to (a) check contracts at the end of transactions, and (b) roll back changes contracts make to state. Unfortunately for our purposes, he emphasizes the delimited transactions in favor of a detailed study of stateful contracts. Dimoulas et al. [22] offer a latent semantics without a study of what the semantics actually *means*. The temporal contracts due to Disney et al. [24] may seem unrelated, but there is no real difference between a temporal contract and a stateful contract. Phrased as stateful contracts, their temporal contracts are essentially limited to advancing the state of a state machine—they forbid contracts from reading or writing to state, or even calling functions.

6.1.2 Datatypes

In Chapter 3, I investigated how contracts interact with abstract types and polymorphism; Sekiyama and Igarashi [61] have extended that work to include a language with fixpoints. It remains to address datatypes; very little study has been done on how contracts and concrete implementations of algebraic datatypes interact. I am aware only of Findler et al. [27]; Rehof [57] studies gradual typing calculi with sums, but within the framework of dynamic types, and without an eye to implementation. While the programming language community normally contents itself with Church encodings (as in Belo et al. [8]), they are insufficient here—as Findler et al. demonstrate, we must finely control abstract datatype representations in order to check contracts without changing asymptotic time complexity. Supporting space-efficient datatypes in addition to space-efficient functions and stacks would complement the work in Chapter 4.

I believe it would be very interesting future work to try to combine my refinement types, protecting partial operations (space efficiently), with classic refinements of datatypes [30] protecting partial matches (also space efficiently). As a first step,

adding pairs and sums presents a design choice: is tag merging deep, as it is for functions, or is it shallow and deferred upon projection? The deep choice is more obviously space efficient, but both options should be explored.

6.1.3 Coercion insertion

While some of the early gradual typing papers give coercion insertion algorithm, very little has been done in terms of comparative study; I am aware only of Swamy et al. [70] and Allende et al. [5]. Can different coercion insertion strategies affect the space and time efficiency of a program?

I am able to answer an emphatic “yes” to those questions based on experience in Boomerang [12]. Consider a partial binary operation $\text{op} : (x:T) \rightarrow \{y:T \mid \text{P } x \ y\} \rightarrow T$ (Recall `seq` and `concat` from Chapter 3.) Suppose we have a nested application $\text{op} (\text{op } e_1 \ e_2) \ e_3$ where each of the e_i is typed at T . It is natural to insert a cast $\langle T \Rightarrow \{y:T \mid \text{P } x \ y\} \rangle$ on e_2 . Well, almost—we actually need to cast to $\{y:T \mid \text{P } e_1 \ y\}$, to account for the application of the dependent function type. So the new inner term is $e_{\text{inner}} = (\text{op } e_1 \ (\langle T \Rightarrow \{y:T \mid \text{P } e_1 \ y\} \rangle e_2))$. If we apply the coercion insertion similarly on the outer term, we need to cast e_3 to $\{y:T \mid \text{P } e_{\text{inner}} \ y\}$ —which will force us to completely rerun the check on e_2 and the application of `op`. Even with parser tricks to keep expression trees balanced, this is so inefficient as to be a non-starter.

Our solution in Boomerang was to evaluate terms with inserted checks in an *aliased* call-by-value fashion: terms are evaluated in a call-by-value order, but there is some sharing in the expression tree. This seriously convoluted the definition of our interpreter—we implemented it with reference cells in OCaml—but led to an enormous speedup. While this problem arises in every language with a dependent application rule, it is particular dire for us because the terms in types are sometimes evaluated!

6.1.4 Theoretical curiosities

I am curious to see what free theorems and reasoning principles come derive from relational parametricity for manifest calculi. I also wonder whether or not we can soundly perform type erasure in F_H —the careful treatment of compatibility at type variables seems to indicate yes, but it remains to be shown.

With the introduction of abstract types, there is room to draw connections between the client/server blame from the ADTs of Section 3.1 and the classic Findler and Felleisen-style client/server blame. On a similar note, I believe that an interesting connection can be drawn between blame labels and stack traces.

On the one hand, Typed Racket née Scheme [74] has a “static analysis” feel in its flow-sensitive type system; later work by Tobin-Hochstadt and Van Horn [75] develops a static analysis. On the other hand, most of the work on manifest contracts uses subtyping, relying on SMT solvers to resolve implications between predicates [44, 45,

28, 10, 59, 58, 51]. How do static analysis and subtyping compare at eliminating unnecessary checks and/or detecting failures?

Knowles and Flanagan [44] give a cast insertion algorithm with a three-valued theorem prover to compile source-level programs into λ_H terms with casts: if the prover says “yes”, then subtyping holds and no cast is necessary; if the prover says “maybe”, then subtyping may or may not hold, so a cast is inserted; if the prover says “no”, then the program is rejected. Rejecting such programs is tempting, but too conservative: a cast that isn’t an upcast doesn’t *always* fail, it just *may* fail. A program should be rejected not when a non-upcast is needed, but when attempting to cast between types that don’t share any values.

6.1.5 Extensions for EFFICIENT

The obvious next step for the space-efficiency work is to add blame [26]. Phil Wadler berated me about the omission, and insists that I prove a blame theorem [78], which would allow me to claim that blame comes from less specific types. Siek and Wadler [68] were the first space-efficient calculus to have blame, which they obtain with some effort—their threesome merging takes place outside in, making it hard to compute which label to blame. Later work [66, 31] simplifies the treatment of blame somewhat. I conjecture that my inside-out coercion merge algorithm offers an easy way to compute blame: blame comes from left to right. The `seq-lazy` definition from Siek and Garcia [66] bolsters my confidence in this conjecture.

In `CAST`, the rule `G_CASTID` doesn’t apply to casts between function types, which instead use `G_CASTFUNWRAP`. Given the cast $\langle T_1 \rightarrow T_2 \Rightarrow T_1 \rightarrow T_2 \rangle v$, we will wrap v with redundant checks: when the T_i aren’t arrow types, the casts immediately disappear. That is, we η -expand all function contracts, even trivial ones. To make matters worse, `G_CASTFUNDYN` and `G_CASTFUNFUN` introduce function casts. In `EFFICIENT`, each function will only ever have a single function proxy on it, so the cost is not so great. Even so, an implementation would want to avoid this unnecessary η -expansion. The corresponding rewrite rule would be $(\text{ld} \mapsto \text{ld}) \longrightarrow \text{ld}$. I would need a merge rule similar to `FUNDOM` to handle this case. Henglein includes the equality $(\text{ld} \mapsto \text{ld}) = \text{ld}$ along with associativity. Relatedly, I could consider eager semantics, where $(\text{Fail} \mapsto c_2) = (c_1 \mapsto \text{Fail}) = \text{Fail}$, as discussed in the literature [65, 66, 31].

If checks are expensive, predicting when checks happen could be important. When will the `N_PREDPRED` merge rule apply, and how can programmers predict performance? There are further implementation concerns: the coercion merging algorithm needs to compare refinement types, which, as Greg Morrisett pointed out, amounts to comparing closures. Comparing closures is dangerous business: optimizers may disrupt programmer expectations. If I were to introduce dependencies, the comparison of closures would include the comparison of environments containing functions; extensionally equivalent functions may not be intensionally equal, leading to still more unexpected behavior. A more nominal approach may serve here. Finally, there is an open question about calling conventions, since tagging introduces a second kind of

closure: when calling a function, do we need to run coercions or not? Jeremy Siek suggested a “smart closure” which holds the logic for branching inside its own code; this may support better branch prediction than an indirect jump or branching at call sites.

Extending the calculus to general refinements, where any type T can be refined to $\{x:T \mid e\}$, would be a challenging but important step towards adding polymorphism. (You can’t allow refinement of type variables unless *any* type can be refined, since there’s no way to know what type will be substituted in for the variable.) It wouldn’t be too difficult to add function refinements $\{x:(T_1 \rightarrow T_2) \mid e\}$ to this calculus, but refinements of refinements seem to break space efficiency: if $\{x:B \mid e\}?$ is canonical, so is $\{x:B \mid e\}?$; $\{x:\{x:B \mid e\} \mid e'\}?$ —there are an infinite number of canonical coercions. In a monomorphic calculus, the number of canonical coercions can be bounded by the types in the original program, but not so in a polymorphic calculus. Prior work relating dynamic types and polymorphism will apply here [57, 46, 4]—though attaining a relational parametricity proof remains a hard open problem.

Other systems have treated failure more eagerly, e.g., $\text{Fail} \mapsto \text{Id} * c \Rightarrow \text{Fail}$. This would further disrupt the connection between the naïve and space-efficient semantics.

Adding dependent functions to the coercion calculus above would complicate matters significantly, but would also add a great deal of expressiveness. Adding the type $(x : T_1) \rightarrow T_2$ is straightforward enough: we should be able to prove type soundness using entirely syntactic techniques, adapting work on FH, a polymorphic calculus with manifest contracts and general refinements [8]. Designing the coercions is a challenge, though. Dependency means that in the coercion $(x:c_1) \mapsto c_2$, the variable x is bound in c_2 . Coercion well formedness now needs a context to keep track of such bound variables:

$$\frac{\Gamma \vdash c_1 : T_{11} \rightsquigarrow T_{21} \quad \Gamma, x:T_{11} \vdash c_2 : T_{12} \rightsquigarrow T_{22}}{\Gamma \vdash (x:c_1) \mapsto c_2 : ((x:T_{21}) \rightarrow T_{12}) \rightsquigarrow ((x:T_{11}) \rightarrow T_{22})}$$

How do dependent functions and their corresponding coercions affect coercion normalization? The following structural equivalence rule, derived from the for dependent functions looks improbable, though its asymmetry echoes the asymmetry of the dependent function cast rule in other manifest calculi, e.g., F_H and λ_H .

$$\begin{aligned} & ((x:c_1) \mapsto c_2); ((y:c'_1) \mapsto c'_2) = \\ & (y:(c'_1; c_1)) \mapsto (c_2\{\langle c'_1 \rangle y/x\}; c'_2) \end{aligned}$$

The metatheory surrounding dependent functions in coercion calculi will be difficult. In fact, the rule above isn’t obvious: trying to use $x:T_{21}$ as the binding for x will raise difficulties in typing the equational rule.

Programs with dependent types have a potentially infinite set of types (and, so, coercions) which may appear as the program evaluates. In the type $(x : \text{Real}) \rightarrow \{y:\text{Real} \mid |x - y^2| < \epsilon\}$, there are potentially infinitely many different codomain types: one for each **Real** value of x . But even with an infinite number of possible coercions,

the set of canonical coercions shouldn't change (beyond the addition of dependency to functions).

Set semantics for refinement types extend refinement types to have a set of predicates, rather than a single one. The `N_PHI` merge rule with $D = \{x:T \mid e\}$ skips a check when we would have projected out of and then back into a refinement type. If refinements were sets, we could broaden this optimization to allow the space-efficient calculus to avoid even more redundant checks. (This was, in fact, the original motivation for this work.)

In an extension of the system of Section 4.5 with so-called “general” refinements [8], the coercion $\{x:\text{Int} \mid x > 0\}?$; $\{x:\{x:\text{Int} \mid x > 0\} \mid x > 5\}?$ is well typed but $\{x:\text{Int} \mid x > 0\}?$; $\{x:\text{Int} \mid x > 5\}?$ isn't. Furthermore, $\{x:\{x:\text{Int} \mid x > 0\} \mid x > 5\}?$ and $\{x:\{x:\text{Int} \mid x > 5\} \mid x > 0\}?$ are totally different coercions, even though the underlying predicates in the refinements are the same. Could we convert $\{x:T \mid e_1\}?$; $\{x:\{x:T \mid e_1\} \mid e_2\}?$ to a single coercion? Simply changing the coercion to $\{x:T \mid e_1 \wedge e_2\}?$ disrupts typing, but the upcast lemma (Lemma 3.4.3 in Chapter 3) shows that this is still behaviorally equivalent. Treating a refinement as being flat with a single set of predicates, rather than a tower of separate refinements, would resolve these ordering issues. That is, we could have a single type that represents both orderings of refinement: $\{x:\text{Int} \mid \{x > 5, x > 0\}\}$. Casts into and out of refinement types could be simplified to adding and removing refinements from the set. When a value is coerced into a refinement type with a set of predicates, the type system remembers all of the predicates equally, acting as a cache of multiple satisfied contracts. The utility of the set semantics is that helps address the library problem: when writing a list library, what refinements are important—emptiness, sortedness, length? When refinements are treated as sets of predicates, libraries can deal only with their own predicates, treating extra client predicates parametrically.

For example, remembering that a list is both non-empty and sorted might be useful for a sorted-list representation of sets. When the set code needs to take the head of a list (which happens to be a minimal member of the set), it can do so directly. Similarly, when calling the `insertSorted` function to add an element to the set, it knows that both its original representation and the extended one are still valid, sorted representations.

Another example where the set semantics allows multiple predicates to interact is a value $v : \{x:\text{Int} \mid \{x \neq 0, \text{prime } x\}\}$. Since $v \neq 0$, we can use it as the divisor with `div : Int → {x: Int | x ≠ 0} → Int`; since `prime v`, we can use it as half of a private key.

Set semantics was my initial motivation for revisiting space-efficient coercions: I was interested in ways of remembering contract checks on values to reduce redundant checking.

Bibliography

- [1] M. Abadi, L. Cardelli, B. Pierce, and G. Plotkin. Dynamic typing in a statically-typed language. In *Principles of Programming Languages (POPL)*, 1989.
- [2] Amal Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *European Symposium on Programming (ESOP)*, 2006.
- [3] Amal Ahmed, Robert Bruce Findler, Jacob Matthews, and Philip Wadler. Blame for all. In *Workshop on Script-to-Program Evolution (STOP)*, 2009.
- [4] Amal Ahmed, Robert Bruce Findler, Jeremy Siek, and Philip Wadler. Blame for all. In *Principles of Programming Languages (POPL)*, 2011.
- [5] Esteban Allende, Johan Fabry, and Eric Tanter. Cast insertion strategies for gradually-typed objects. In *Dynamic Languages Symposium (DLS)*, 2013.
- [6] David Aspinall and Adriana Compagnoni. Subtyping dependent types. *IFIP Conference on Theoretical Computer Science (TCS)*, September 2001.
- [7] Brian Aydemir and Stephanie Weirich. LNgen: Tool support for locally nameless representations. Technical Report MS-CIS-10-24, University of Pennsylvania, June 2010.
- [8] João Filipe Belo, Michael Greenberg, Atsushi Igarashi, and Benjamin C. Pierce. Polymorphic contracts. In *European Symposium on Programming (ESOP)*, 2011.
- [9] Véronique Benzaken, Giuseppe Castagna, and Alain Frisch. CDuce: an XML-centric general-purpose language. In *International Conference on Functional Programming (ICFP)*, pages 51–63, New York, NY, USA, 2003. ACM. ISBN 1-58113-756-7. doi:10.1145/944705.944711.
- [10] Gavin M. Bierman, Andrew D. Gordon, Cătălin Hrițcu, and David Langworthy. Semantic subtyping with an SMT solver. In *International Conference on Functional Programming (ICFP)*, 2010.
- [11] Matthias Blume and David A. McAllester. Sound and complete models of contracts. *Journal of Functional Programming (JFP)*, 2006.

- [12] Aaron Bohannon, J. Nathan Foster, Benjamin C. Pierce, Alexandre Pilkiewicz, and Alan Schmitt. Boomerang: resourceful lenses for string data. In *Principles of Programming Languages (POPL)*, 2008. doi:10.1145/1328438.1328487.
- [13] Val Breazu-Tannen, Thierry Coquand, Carl A. Gunter, and Andre Scedrov. Inheritance as implicit coercion. *Information and Computation*, 93(1):172 – 221, 1991. ISSN 0890-5401. doi:http://dx.doi.org/10.1016/0890-5401(91)90055-7. Selections from 1989 {IEEE} Symposium on Logic in Computer Science.
- [14] Luca Cardelli, Simone Martini, John C. Mitchell, and Andre Scedrov. An extension of system F with subtyping. In *Information and Computation*, 1991.
- [15] Robert Cartwright and Mike Fagan. Soft typing. In *Programming Language Design and Implementation (PLDI)*, pages 278–292, New York, NY, USA, 1991. ACM. ISBN 0-89791-428-7. doi:10.1145/113445.113469.
- [16] Olaf Chitil and Frank Huch. Monadic, prompt lazy assertions in haskell. In *Asian Symposium on Programming Languages and Systems (APLAS)*, 2007.
- [17] Ravi Chugh, David Herman, and Ranjit Jhala. Dependent types for javascript. In *OOPSLA*, pages 587–606, 2012.
- [18] Ravi Chugh, Patrick M. Rondon, and Ranjit Jhala. Nested refinements: a logic for duck typing. In *Principles of Programming Languages (POPL)*, POPL '12, pages 231–244, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1083-3. doi:10.1145/2103656.2103686.
- [19] Ravi Chugh, Patrick Maxim Rondon, and Ranjit Jhala. Nested refinements: a logic for duck typing. In *Principles of Programming Languages (POPL)*, pages 231–244, 2012.
- [20] Julien Cretin and Didier Rémy. On the power of coercion abstraction. In *Principles of Programming Languages (POPL)*, POPL '12, pages 361–372, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1083-3. doi:10.1145/2103656.2103699.
- [21] Coq development team. The Coq proof assistant reference manual, version 8.2, August 2009. <http://coq.inria.fr/>.
- [22] Christos Dimoulas, Robert Bruce Findler, Cormac Flanagan, and Matthias Felleisen. Correct blame for contracts: no more scapegoating. In *Principles of Programming Languages (POPL)*, 2011. doi:10.1145/1926385.1926410.
- [23] Tim Disney and Cormac Flanagan. Gradual information flow typing. In *Workshop on Script-to-Program Evolution (STOP)*, 2011.

- [24] Tim Disney, Cormac Flanagan, and Jay McCarthy. Temporal higher-order contracts. In *Proceedings of the 16th ACM SIGPLAN international conference on Functional programming*, ICFP '11, pages 176–188, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0865-6. doi:10.1145/2034773.2034800.
- [25] Robert Bruce Findler. Contracts as pairs of projections. In *Symposium on Logic Programming*, 2006.
- [26] Robert Bruce Findler and Matthias Felleisen. Contracts for higher-order functions. In *International Conference on Functional Programming (ICFP)*, 2002.
- [27] Robert Bruce Findler, Shu-Yu Guo, and Anne Rogers. Lazy contract checking for immutable data structures. In Olaf Chitil, Zoltán Horváth, and Viktória Zsóka, editors, *Implementation and Application of Functional Languages*, pages 111–128. Springer-Verlag, Berlin, Heidelberg, 2008. ISBN 978-3-540-85372-5. doi:10.1007/978-3-540-85373-2_7.
- [28] Cormac Flanagan. Hybrid type checking. In *Principles of Programming Languages (POPL)*, 2006.
- [29] Cédric Fournet, Nikhil Swamy, Juan Chen, Pierre-Évariste Dagand, Pierre-Yves Strub, and Benjamin Livshits. Fully abstract compilation to javascript. In *Principles of Programming Languages (POPL)*, pages 371–384, 2013.
- [30] Tim Freeman and Frank Pfenning. Refinement types for ML. In *Programming Language Design and Implementation (PLDI)*, June 1991.
- [31] Ronald Garcia. Calculating threesomes, with blame. In *International Conference on Functional Programming (ICFP)*, 2013.
- [32] Mike Gordon. In Gordon Plotkin, Colin Stirling, and Mads Tofte, editors, *Proof, language, and interaction*, chapter From LCF to HOL: a short history, pages 169–185. MIT Press, Cambridge, MA, USA, 2000. ISBN 0-262-16188-5.
- [33] Michael Greenberg. Space-efficient manifest contracts. Rejected from POPL, 2013.
- [34] Michael Greenberg, Benjamin C. Pierce, and Stephanie Weirich. Contracts made manifest. In *Principles of Programming Languages (POPL)*, 2010.
- [35] Michael Greenberg, Benjamin C. Pierce, and Stephanie Weirich. Contracts made manifest. *JFP*, 22(3):225–274, May 2012.
- [36] Jessica Gronski and Cormac Flanagan. Unifying hybrid types and contracts. In *Trends in Functional Programming (TFP)*, 2007.

- [37] Arjun Guha, Jacob Matthews, Robert Bruce Findler, and Shriram Krishnamurthi. Relationally-parametric polymorphic contracts. In *Dynamic Languages Symposium (DLS)*, 2007.
- [38] Fritz Henglein. Dynamic typing: Syntax and proof theory. *Sci. Comput. Program.*, 22(3):197–230, 1994.
- [39] David Herman, Aaron Tomb, and Cormac Flanagan. Space-efficient gradual typing. In *Trends in Functional Programming (TFP)*, pages 404–419, April 2007.
- [40] David Herman, Aaron Tomb, and Cormac Flanagan. Space-efficient gradual typing. *Higher Order Symbol. Comput.*, 23(2):167–189, June 2010. ISSN 1388-3690. doi:10.1007/s10990-011-9066-z.
- [41] Ralf Hinze, Johan Jeuring, and Andres Löf. Typed contracts for functional programming. In *Functional and Logic Programming (FLOPS)*, 2006.
- [42] Tony Hoare. Null references: The billion dollar mistake. QCon talk, 2009.
- [43] Catalin Hritcu, Michael Greenberg, Ben Karel, Benjamin C. Pierce, and Greg Morrisett. All your IFCEException are belong to us. In *Security and Privacy (SP), 2013 IEEE Symposium on*, pages 3–17, 2013. doi:10.1109/SP.2013.10. The author is deeply embarassed by the title of this paper.
- [44] Kenneth Knowles and Cormac Flanagan. Hybrid type checking. *ACM Transactions on Programming Languages and Systems*, 32:6:1–6:34, 2010.
- [45] Kenneth Knowles, Aaron Tomb, Jessica Gronski, Stephen N. Freund, and Cormac Flanagan. Sage: Hybrid checking for flexible specifications. In *Scheme and Functional Programming Workshop*, 2006.
- [46] Jacob Matthews and Amal Ahmed. Parametric polymorphism through run-time sealing or, theorems for low, low prices! In *European Symposium on Programming (ESOP)*, 2008.
- [47] Bertrand Meyer. *Eiffel: the language*. Prentice-Hall, Inc., 1992. ISBN 0-13-247925-7.
- [48] Yasuhiko Minamide. Runtime behavior of conversion interpretation of subtyping. In *Selected Papers from the 13th International Workshop on Implementation of Functional Languages, IFL '02*, pages 155–167, London, UK, UK, 2002. Springer-Verlag. ISBN 3-540-43537-9.
- [49] Yasuhiko Minamide and Jacques Garrigue. On the runtime complexity of type-directed unboxing. In *International Conference on Functional Programming (ICFP)*, ICFP '98, pages 1–12, New York, NY, USA, 1998. ACM. ISBN 1-58113-024-4. doi:10.1145/289423.289424.

- [50] Georg Neis, Derek Dreyer, and Andreas Rossberg. Non-parametric parametricity. In *International Conference on Functional Programming (ICFP)*, 2009.
- [51] Xinming Ou, Gang Tan, Yitzhak Mandelbaum, and David Walker. Dynamic typing with dependent types. In *IFIP Conference on Theoretical Computer Science (TCS)*, 2004.
- [52] Benjamin Pierce and Eijiro Sumii. Relating cryptography and polymorphism, July 2000.
- [53] A. M. Pitts. Typed operational reasoning. In B. C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 7. The MIT Press, 2005.
- [54] A. M. Pitts. Typed operational reasoning. In B. C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 7, pages 245–289. The MIT Press, 2005. ISBN 0-262-16228-8.
- [55] PLT. PLT Racket, 2013. URL <http://racket-lang.org>.
- [56] PLT. PLT Racket contract system, 2013. URL <http://pre.plt-scheme.org/docs/html/guide/contracts.html>.
- [57] Jakob Rehof. Polymorphic dynamic typing: Aspects of proof theory and inference. Master’s thesis, DIKU, 1995. DIKU Technical Report D-249.
- [58] Patrick M. Rondon, Ming Kawaguchi, and Ranjit Jhala. Liquid types. In *Programming Language Design and Implementation (PLDI)*, 2008.
- [59] Patrick Maxim Rondon, Ming Kawaguchi, and Ranjit Jhala. Low-level liquid types. In *Principles of Programming Languages (POPL)*, 2010.
- [60] Taro Sekiyama and Atsushi Igarashi. Personal communication, November 2013.
- [61] Taro Sekiyama and Atsushi Igarashi. Logical relations for a manifest calculus, fixed. In preparation for submission., 2013.
- [62] Peter Sewell, Francesco Zappa Nardelli, Scott Owens, Gilles Peskine, Tom Ridge, Susmit Sarkar, and Rok Strnisa. Ott: effective tool support for the working semanticist. In *International Conference on Functional Programming (ICFP)*, 2007.
- [63] Zhong Shao. Flexible representation analysis. In *International Conference on Functional Programming (ICFP)*, pages 85–98, Amsterdam, The Netherlands, June 1997.
- [64] Avraham Shinnar. *Safe and Effective Contracts*. PhD thesis, Harvard University, May 2011.

- [65] Jeremy Siek, Ronald Garcia, and Walid Taha. Exploring the design space of higher-order casts. In Giuseppe Castagna, editor, *Programming Languages and Systems*, volume 5502 of *Lecture Notes in Computer Science*, pages 17–31. Springer Berlin Heidelberg, 2009. ISBN 978-3-642-00589-3. doi:10.1007/978-3-642-00590-9_2.
- [66] Jeremy G Siek and Ronald Garcia. Interpretations of the gradually-typed lambda calculus. In *Scheme and Functional Programming (SFP)*, 2012.
- [67] Jeremy G. Siek and Walid Taha. Gradual typing for functional languages. In *Scheme and Functional Programming Workshop*, September 2006.
- [68] Jeremy G. Siek and Philip Wadler. Threesomes, with and without blame. In *Principles of Programming Languages (POPL)*, pages 365–376, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-479-9. doi:10.1145/1706299.1706342.
- [69] Michael Sperber, R. Kent Dybvig, Matthew Flatt, Anton van Straaten, Robby Findler, and Jacob Matthews, editors. *Revised [6] Report on the Algorithmic Language Scheme*. Cambridge University Press, Cambridge, UK, June 2010. ISBN 9780521193993.
- [70] Nikhil Swamy, Michael Hicks, and Gavin M. Bierman. A theory of typed coercions and its applications. In *International Conference on Functional Programming (ICFP)*, ICFP '09, pages 329–340, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-332-7. doi:10.1145/1596550.1596598. URL <http://doi.acm.org/10.1145/1596550.1596598>.
- [71] Nikhil Swamy, Joel Weinberger, Cole Schlesinger, Juan Chen, and Benjamin Livshits. Verifying higher-order programs with the dijkstra monad. In *Programming Language Design and Implementation (PLDI)*, pages 387–398, 2013.
- [72] Naoshi Tabuchi, Eijiro Sumii, and Akinori Yonezawa. Regular expression types for strings in a text processing language. *Electronic Notes in Theoretical Computer Science*, 2003. doi:10.1016/S1571-0661(04)80781-3. International Workshop in Types in Programming.
- [73] Satish Thatte. Quasi-static typing. In *Principles of Programming Languages (POPL)*, 1990.
- [74] Sam Tobin-Hochstadt and Matthias Felleisen. The design and implementation of typed scheme. In *Principles of Programming Languages (POPL)*, 2008.
- [75] Sam Tobin-Hochstadt and David Van Horn. Higher-order symbolic execution via contracts. In *OOPSLA*, OOPSLA '12, pages 537–554, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1561-6. doi:10.1145/2384616.2384655.

- [76] Philip Wadler. Theorems for free! In *Conference on Functional Programming and Computer Architecture (FPCA)*, 1989.
- [77] Philip Wadler and Robert Bruce Findler. Well-typed programs can't be blamed. In *Scheme and Functional Programming Workshop*, 2007.
- [78] Philip Wadler and Robert Bruce Findler. Well-typed programs can't be blamed. In *European Symposium on Programming (ESOP)*, 2009.
- [79] Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115:38–94, 1994.
- [80] Dana N. Xu, Simon Peyton Jones, and Koen Claessen. Static contract checking for haskell. In *Principles of Programming Languages (POPL)*, 2009.