

Number 948



**UNIVERSITY OF
CAMBRIDGE**

Computer Laboratory

Reasoning about effectful programs and evaluation order

Dylan McDermott

June 2020

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500
<https://www.cl.cam.ac.uk/>

© 2020 Dylan McDermott

This technical report is based on a dissertation submitted October 2019 by the author for the degree of Doctor of Philosophy to the University of Cambridge, Homerton College.

Technical reports published by the University of Cambridge Computer Laboratory are freely available via the Internet:

<https://www.cl.cam.ac.uk/techreports/>

ISSN 1476-2986

Abstract

Program transformations have various applications, such as in compiler optimizations. These transformations are often effect-dependent: replacing one program with another relies on some restriction on the side-effects of subprograms. For example, we cannot eliminate a dead computation that raises an exception, or a duplicated computation that prints to the screen. Effect-dependent program transformations can be described formally using effect systems, which annotate types with information about the side-effects of expressions.

In this thesis, we extend previous work on effect systems and correctness of effect-dependent transformations in two related directions.

First, we consider evaluation order. Effect systems for call-by-value languages are well-known, but are not sound for other evaluation orders. We describe sound and precise effect systems for various evaluation orders, including call-by-name. We also describe an effect system for Levy's call-by-push-value, and show that this subsumes those for call-by-value and call-by-name. This naturally leads us to consider effect-dependent transformations that replace one evaluation order with another. We show how to use the call-by-push-value effect system to prove the correctness of transformations that replace call-by-value with call-by-name, using an argument based on logical relations. Finally, we extend call-by-push-value to additionally capture call-by-need. We use our extension to show a classic example of a relationship between evaluation orders: if the side-effects are restricted to (at most) nontermination, then call-by-name is equivalent to call-by-need.

The second direction we consider is non-invertible transformations. A program transformation is non-invertible if only one direction is correct. Such transformations arise, for example, when considering undefined behaviour, nondeterminism, or concurrency. We present a general framework for verifying noninvertible effect-dependent transformations, based on our effect system for call-by-push-value. The framework includes a non-symmetric notion of correctness for effect-dependent transformations, and a denotational semantics based on order-enriched category theory that can be used to prove correctness.

Acknowledgements

Most importantly, I thank Alan Mycroft for his enthusiastic supervision during my PhD. We had many useful discussions about a wide range of topics, and he provided a lot of encouragement, and advice on how to write papers.

I also thank all of the people who I had helpful technical discussions with, including Paul Downen, Paul Blain Levy, Martin Hyland, Dominic Orchard, Tomas Petricek, Sam Staton, and Tarmo Uustalu. I'm particularly grateful to Ohad Kammar for collaboration and many conversations about category theory and denotational semantics.

My friends and colleagues at the Computer Lab deserve thanks for their help and support. My office-mate Ian Orton provided many entertaining (and occasionally useful) distractions after 12:30pm, along with Sam Ainsworth, Aurore Alcolei, Matthew Daggitt, Hugo Paquet and Philip Saville, Nathanael Arkor, Adam Ó Conghaile, Andrej Ivašković, Ben Simner, Shaun Steenkamp and Dima Szamozvancev. And probably some others I forgot to mention. I also thank the admin staff at the CL, especially Megan Sammons for her tireless work organising supervisions.

I'm grateful to Matthew Parkinson, Dimitrios Vytiniotis and others at MSR Cambridge for an interesting internship project, discussions about PhDs, and a surprisingly large amount of fun debugging garbage collectors.

Finally, I thank the EPSRC for funding this work.

Contents

1	Introduction	7
1.1	Approach	8
1.2	Contributions	10
2	Effect systems and evaluation orders	11
2.1	Effect algebras	12
2.2	Simply-typed lambda calculus	13
2.3	Call-by-value	15
2.4	Moggi-style call-by-name	17
2.5	Graded monadic metalanguage	18
2.6	Levy-style call-by-name	25
2.7	Graded call-by-push-value	27
2.8	Related work	42
2.9	Summary	42
3	Call-by-value and call-by-name	43
3.1	Logical relations for graded call-by-push-value	44
3.2	Restricting side-effects in call-by-value and call-by-name	51
3.3	A Galois connection between call-by-value and call-by-name	51
3.4	Reasoning principle for call-by-value and call-by-name	56
3.5	Related work	59
3.6	Summary	60
4	Noninvertible program transformations	63
4.1	Examples of noninvertible transformations	64
4.2	Order-enriched semantics of GCBPV	68
4.3	Relating syntax and semantics	88
4.4	Relating call-by-value and call-by-name, semantically	95
4.5	Related work	98
4.6	Summary	99
5	Call-by-need and extended call-by-push-value	101
5.1	Extended call-by-push-value	102
5.2	Call-by-need translation	109
5.3	Equivalence between call-by-name and call-by-need	112
5.4	An effect system for extended call-by-push-value	116
5.5	Related work	118
5.6	Summary	119

6 Conclusions	121
6.1 Future work	122
6.2 Final remarks	123
Bibliography	125
A Order-enriched category theory	133
B Additional proofs	137
B.1 Erasing coercions in GCBPV terms	137
B.2 Logical relations and the free lifting	142
B.3 Call-by-name and call-by-need	144

Chapter 1

Introduction

At a high level, the question considered in this thesis is:

When is it correct to replace one program phrase with another?

Suppose we are writing an optimizing compiler. The optimizations are applied locally, replacing expressions e_1 (which form part of the program) with other expressions e_2 . Setting aside the question of whether replacing e_1 with e_2 actually improves the program (we do not consider this), we want to know whether it is *valid* to do so. The output of the compiler should not do anything the original program could not, so it is valid to replace e_1 with e_2 when *every observable behaviour of e_2 is an observable behaviour of e_1* . We develop techniques for formal reasoning about programs that can be used to prove the validity of program transformations.

Of course, formal reasoning about programs is not a new area. To motivate the specific problems we consider, we give three examples of common program transformations:

	Old expression	Replacement
Dead code elimination (where x not free in e')	let $x = e$ in e'	e'
Inlining	$(\lambda x. e) e'$	$e'[x \mapsto e]$
Common subexpression elimination	let $x_1 = e$ in let $x_2 = e$ in $e' (x_1, x_2)$	let $x = e$ in $e' (x, x)$

These examples highlight three important aspects of program transformations:

- *They are often effect-dependent.* The first example is valid if e has no side-effects, but not e.g. if e changes some state that e' then reads from, or if e raises an exception. (For these examples, **let** represents eager evaluation, so first e is evaluated, and then e' is.) The first example is therefore *effect-dependent*, which means applying it requires static knowledge of the side-effects of the expressions involved. The wide array of side-effects available in practical languages means that a majority of the transformations we would like to perform are effect-dependent. We therefore focus on verifying effect-dependent transformations throughout this thesis, continuing a long line of work started by Tolmach [95] and Benton et al. [12].
- *They often depend on evaluation order.* Consider the second example. This is clearly valid in call-by-name languages (it is just β -reduction), even with arbitrary side-effects. It is not in general valid in call-by-value languages. We cannot consider program transformations

without thinking about evaluation order. Even worse, the side-effects of expressions depend on evaluation order, so we cannot validate *any* effect-dependent transformations without considering evaluation order. Previous work on effect-dependent transformations has mostly considered only a single evaluation order (usually call-by-value). We focus on capturing a range of evaluation orders. This aspect also suggests another useful class of program transformations: those that replace one evaluation order with another, such as call-by-value with call-by-name.

- *They are often noninvertible.* One way to justify the validity of the third example (or in general, any program transformation) is to show that both sides have the *same* semantics. This is the method normally used, but it is unsuitable for many examples. Consider the third example, and suppose that the only side-effect of e is to make nondeterministic choices internally. We could allow the compiler to make nondeterministic choices statically, by replacing the expression $(e_1 \text{ or } e_2)$ with e_1 or with e_2 . If this is the case, then we can replace the expression on the left of the example, which makes the choices in e twice, with the one on the right, even though it does not have the same semantics. Program transformations are in general *noninvertible*: only one direction is valid.

This thesis constructs a general framework for proving the validity of program transformations with all three of these aspects.

1.1 Approach

We have stated our aim, now we explain the approach that we take. Throughout, we express program transformations inside some *intermediate language* that we construct for this purpose. Given an expression that we want to transform, the first step is to *translate* it into our intermediate language of choice. We then reason about the translation of the expression. The intermediate languages we use in this thesis (GCBPV in Section 2.7, ECBPV in Section 5.1) are based on Levy's [56] call-by-push-value calculus (CBPV). (We use e for expressions in source languages, and M for terms in intermediate languages, so each expression e can be translated into a term M .)

Using intermediate languages in this way has several advantages. They can be designed specifically for validating program transformations, ignoring considerations that apply to source languages (we do not need to write programs directly in intermediate languages, because intermediate-language programs are constructed by the translations). In particular, the intermediate language can be close to the denotational semantics. (For CBPV, several side-effects have simpler models than e.g. for monadic intermediate languages [78]. This is one of the reasons why CBPV is useful for verifying program transformations.)

However, the main advantage of intermediate languages is that we can design them so that evaluation order is irrelevant (for a given intermediate-language program, every evaluation order gives the same behaviour). The evaluation order of the source language is encoded in the intermediate language *programs* (in the syntax) by the translation. We use intermediate languages that capture several evaluation orders. CBPV and the variants we use capture both call-by-value and call-by-name, so for every source-language expression e we have two intermediate-language terms: a call-by-value translation $\llbracket e \rrbracket^v$ and a call-by-name translation $\llbracket e \rrbracket^n$. (We also have other translations in Chapters 2 and 5.) Encoding evaluation orders inside programs in this way allows us to consider transformations that replace one evaluation order with another inside a single expression, since they really can just transform the syntax of the program. There is no change in the semantics of the intermediate language itself. For

call-by-value and call-by-name, we relate the behaviours of $(e)^v$ and $(e)^n$, which are terms in the same language.

We use *effect systems* [63] to enable *effect-dependent* transformations, following previous work. Effect systems refine type systems so that they statically (over-)estimate the side-effects of expressions. (We do not attempt to define¹ *side-effect* precisely, but we use a broad interpretation, which includes at least nontermination, nondeterministic and probabilistic choice, and exceptions.) Effect systems are the traditional way of working with effect-dependent transformations. They are a lightweight technique (requiring only minor modifications to type systems), and are expressive (we can use them to restrict the order and the number of times that side-effects are used, not just *whether* particular side-effects are used). They are also amenable to mechanization (for example, type inference algorithms can be extended to *effect* inference algorithms). We design several effect systems to account for the fact that the side-effects of expressions depend on evaluation order.

We formally define *validity* of program transformations in terms of *contextual preorders* \leq_{ctx} . It is correct to replace M with N when $M \leq_{\text{ctx}} N$, which means that every observable behaviour of N is an observable behaviour of M . For example, if the only side-effect is nondeterministic choice, this might mean that every possible result of a program containing N is a possible result of the same program but containing M instead: the transformation does not make the set of possible results of a program larger. If we want to validate a transformation that replaces e with e' inside a call-by-value source language, we show that $(e)^v \leq_{\text{ctx}} (e')^v$. Contextual preorders might not be symmetric in examples; this allows us to consider noninvertible transformations.

Comparing evaluation orders requires more work than just asking whether the contextual preorder relates them. Everything we do is typed, and it only makes sense to ask if $M \leq_{\text{ctx}} N$ if M and N have the same type, but in general this is not the case if M and N are two different translations of the same expression. Call-by-name functions take (thunked) computations as arguments, call-by-value functions take values as arguments (the distinction is important in CBPV), so if e is a function, it does not make sense to replace the call-by-value translation $(e)^v$ with the call-by-name translation $(e)^n$. To fix this, we need to add some extra glue to the transformations. To pass a call-by-name argument to a call-by-value function it needs to be evaluated first; to pass a call-by-value argument to a call-by-name function we need to convert it into a trivial computation. To compare call-by-value and call-by-name, we therefore construct maps between the two evaluation orders, and compose these with the two translations (this is similar to the technique used by Reynolds [89] to relate direct and continuation semantics). Other comparisons between evaluation orders require similar considerations.

We define contextual preorders in terms of *inequational theories*, which are specified by axioms characterizing the constructs in the intermediate language. These axioms include β - and η -laws, and also laws capturing the behaviour of the side-effects, for example, associativity of nondeterministic choice, or that two reads from the same location can be merged. Different axioms are used for different side-effects. We prefer inequational theories to e.g. operational semantics for two reasons: theories allow us to apply inequations under lambdas (which is important because we mainly consider open terms), and they allow uniformly specifying various side-effects (e.g. the usual way of specifying operational semantics for global state is to carry around an extra state parameter; this is not required for inequational theories).

To prove instances of contextual preorders, and hence to prove the validity of program transformations, we employ two different techniques. The first is to use logical relations on

¹The best we can do is cheat by saying that a side-effect is roughly anything that we can capture in the languages we use. (This includes the usual examples.) A proper definition could be along the lines of Sabry [91]: a side-effect is anything that breaks equivalences between evaluation orders. (Though unlike Sabry we do not treat divergence and errors as special cases.)

the syntax (type-indexed families of binary relations on terms). These lie somewhere between inequational theories and contextual preorders. We construct them so that they relate more pairs of terms than the inequational theory (in most cases, the program transformations we consider are not instances of the inequational theory), and also so that they only relate M to N when $M \preceq_{\text{ctx}} N$ (which means we can actually use them to validate program transformations). We use logical relations to relate call-by-value and call-by-name executions of expressions in Chapter 3. However, they turn out to be quite cumbersome to use. We therefore prefer the second technique, which is to use denotational semantics.

We define a *categorical* semantics of GCBPV in Chapter 4, and use it to verify noninvertible program transformations. The semantics allows for models in various categories, which allows us to model as wide a range of side-effects as possible. The categories we use are *order-enriched*, which enables us to reason about noninvertible transformations: to show $M \preceq_{\text{ctx}} N$ holds, we show $\llbracket M \rrbracket \sqsubseteq \llbracket N \rrbracket$, where \sqsubseteq is part of the data of the model.

When we relate call-by-name and call-by-need evaluation in Chapter 5 we are forced to go back to logical relations on the syntax because we do not have a general call-by-need denotational semantics. We actually need to use a more powerful notion of logical relation than previously (*Kripke logical relations of varying arity* [39]), because of the extra difficulties that call-by-need introduces.

Our framework for validating program transformations therefore has several components. We have various source-language effect systems, one for each evaluation order we wish to consider. We also have intermediate languages that capture all of these evaluation orders via translations from source languages ($(_)^\vee$, $(_)^\eta$, etc.). We define a notion of inequational theory for each intermediate language, and use these to define contextual preorders. Finally, we also include machinery (logical relations and denotational semantics) for proving instances of contextual preorders. Together, these components enable us to validate all of the kinds mentioned above.

1.2 Contributions

This thesis develops a framework for proving the correctness of effect-dependent program transformations:

- We develop tools for formulating effect-dependent transformations for various evaluation orders, including two forms of call-by-name (Chapter 2). This includes source-language effect systems for each of the evaluation orders we consider, and an intermediate language (GCBPV) we use for reasoning throughout most of this thesis.
- We establish a novel reasoning principle for relating call-by-value and call-by-name (Chapter 3), which enables us to prove the validity of program transformations that replace call-by-value with call-by-name, and vice versa. The technique we use to do this is new, and we expect it to work more generally (to relate other evaluation orders).
- We develop machinery for proving the validity of noninvertible program transformations (Chapter 4), including an order-enriched denotational semantics for GCBPV. We apply this to examples including undefined behaviour, nondeterminism, and shared global state.
- We extend GCBPV to additionally capture call-by-need evaluation (Chapter 5). This extension enables us to give a new proof technique for relating evaluation orders. We exemplify the technique by proving the classic result that call-by-name and call-by-need are equivalent if the only side-effect is nontermination.

Chapter 2

Effect systems and evaluation orders

We formalize effect-dependent program transformations using *effect systems*. The goal of an effect system is to statically associate computations with *effects* ε , which abstractly represent the side-effects the computation has. Hence by requiring computations to be associated with a chosen effect ε , we can view effect systems as a way to restrict the side-effects of computations.

As noted in the introduction, different expressions may have different side-effects depending on the evaluation order used. The effect system used must therefore reflect the evaluation order. Using a call-by-value effect system to analyse a call-by-name language could lead to imprecise results (leading us to miss valid effect-dependent transformations), or worse, might be unsound (leading us to incorrectly believe certain transformations are valid).

To reason about transformations that change the evaluation order it is useful to have a single *intermediate* language that allows programs to express their evaluation order, rather than one language for each evaluation order. Two such languages are Moggi's *monadic metalanguage* [78] and Levy's *call-by-push-value* [56]. We can use these to reason about *source* languages with a single evaluation order by translating from source to intermediate. Since valid changes in evaluation order are also effect-dependent, we therefore also wish to define effect systems for intermediate languages.

This chapter serves two purposes. First, it provides a general introduction to effect systems. In particular we explain *effect algebras* (Section 2.1), which provide the abstract notion of *effect* we track, the well-known call-by-value effect system (Section 2.3), and the less well-known effect system for the monadic metalanguage (Section 2.5). Second, it also makes two key contributions:

- We describe effect systems for two different forms of call-by-name (Sections 2.4 and 2.6). These allow us to do effect-dependent reasoning in call-by-name languages. We contrast the two forms of call-by-name in Section 2.6.
- We describe an effect system for call-by-push-value (Section 2.7), and show that translations from source languages into call-by-push-value respect the effect systems. We use this call-by-push-value effect system as the basis for work in later chapters of this thesis.

Source and intermediate languages Before continuing, we clarify what we mean by *source* and *intermediate* languages. If a language is intended primarily to be used for reasoning (e.g. proving correctness of program transformations), then we call it an intermediate language. In this chapter, the intermediate languages are the monadic metalanguage (Section 2.5) and call-by-push-value (Section 2.7). We focus on making them easy to reason about (for example, by ensuring that their effect systems are syntax-directed, and by making sequencing of side-effecting computations explicit). Source languages are intended to more closely resemble

programming languages (in particular, sequencing of computations is implicit as it is in ML), and formal reasoning about source languages is intended to be done by translating source programs into an intermediate language, and then reasoning in the intermediate language. In this thesis, we concentrate mainly on the intermediate languages, and hence we keep the source languages as small as possible (we only include booleans, unit type and function types).

2.1 Effect algebras

An *effect* ε abstractly represents some collection of side-effects. Effect systems statically assign effects to program fragments; when the effect ε is assigned, this means the possible runtime side-effects of the program fragment are restricted to ε . To do this, we require the effects to have some extra structure, called an *effect algebra*¹. There are various forms of effect algebra. Commonly effects are assumed to form semilattices, but other structures, such as *effectors* [94] or *effect quantales* [31] are also used. Here we follow Katsumata [45] and use *preordered monoids*.

Definition 2.1.1 (Preordered monoid) A *preordered monoid* $(\mathcal{E}, \leq, \cdot, 1)$ consists of a monoid $(\mathcal{E}, \cdot, 1)$ and a preorder \leq on \mathcal{E} , such that the binary operation \cdot is monotone in each argument separately. ◀

Preordered monoids have several advantages: they are simple, and they generalize the other forms of effect algebra we have just mentioned (Katsumata postulates that effects *always* form a preordered monoid). They can also be easily extended with additional data (such as iteration operators for recursion).

The effects ε are elements of the set \mathcal{E} . The *subeffecting relation* \leq provides a notion of approximation of effects: $\varepsilon \leq \varepsilon'$ means ε is more restrictive than ε' . The *multiplication* \cdot represents sequencing of effects: $\varepsilon \cdot \varepsilon'$ is the effect of running a computation with effect ε followed by a computation with effect ε' . Finally, the *unit* 1 is used for computations with no side-effects. Each effect system is parameterized by the effect algebra, so we can instantiate an effect system with different effect algebras for different use cases.

The simplest effect algebra is the trivial preordered monoid, in which \mathcal{E} is a singleton $\{\star\}$. When instantiated with the trivial preordered monoid, effect systems do not track effect information, and are the same as ordinary type systems. More interesting examples are *Gifford-style effect algebras*, which are used for example by Lucassen and Gifford [63].

Example 2.1.2 *Gifford-style* effect algebras have the form $(\mathcal{P}\Sigma, \subseteq, \cup, \emptyset)$, where Σ is some set of operations (e.g. $\Sigma := \{\text{raise, get, put}\}$), and $\mathcal{P}\Sigma$ is the powerset. In this case, effects $\varepsilon \subseteq \Sigma$ give the set of operations that a computation *may* use during its execution. For example, a computation with effect $\{\text{raise, get}\}$ may raise an exception or get the value of the state, but does not change the value of the state. The subeffecting relation \subseteq allows additional operations (that are not used) to be included in the effect (this is useful to balance the two branches of *if*-expressions). We can use a Gifford-style effect system to express effect-dependent program transformations. For example, consider

$\begin{array}{l} \text{let } x = e \text{ in} \\ \text{let } y = e \text{ in} \\ (x, y) \end{array}$	$\begin{array}{l} \text{let } x = e \text{ in} \\ (x, x) \end{array}$
---	---

¹For the avoidance of doubt, these have nothing to do with effect algebras in quantum theory [22].

If the effect of e is $\{\text{raise, get}\}$, then it is correct to replace the expression on the left with the one on the right. However, if e uses `put`, this might not be the case. ◀

This example gives a *may* analysis: it allows overapproximation of effects. In general, if a preordered monoid $(\mathcal{E}, \leq, \cdot, 1)$ describes a may analysis, then the preordered monoid $(\mathcal{E}, \geq, \cdot, 1)$ (which is the same but with the opposite order) describes the corresponding must analysis that allows under- instead of over-approximation.

Example 2.1.3 Consider a language with an operation `flip` that nondeterministically chooses either `true` or `false`. We can use the preordered monoid $(\mathbb{N}_+, \leq, \cdot, 1)$, where \mathbb{N}_+ is the set of positive integers, \cdot is the usual integer multiplication, and the unit is the integer 1, to statically (over-)estimate the number of potential results of computations. The effect of `flip` would be the integer 2. ◀

Example 2.1.4 Consider a language with global state, including operations `get` and `put` that read from and write to the state. In general, to evaluate an expression we might have to supply an initial state, but if `put` must be used before `get`, then we do not. We can determine whether this is the case statically using a preordered monoid with underlying preorder $\{\text{pf} \leq 1 \leq \text{gf}\}$. The effect `pf` means must set the value of the state, and does not get the value before setting it; the unit 1 means does not get before setting the value, but may do neither; and `gf` means may get the value first. The multiplication is defined by:

$$\text{pf} \cdot \varepsilon = \text{pf} \quad 1 \cdot \varepsilon = \varepsilon \quad \text{gf} \cdot \varepsilon = \text{gf}$$

The effect of applying `get` is `gf`, and the effect of applying `put` is `pf`. Programs that have the effect `pf` or the effect 1 do not require an initial state. For example, if e' has effect `pf`, then it is correct to replace $(e; e')$ with e' . Even if e changes the state, it will be overwritten by e' , and hence will not change the behaviour of e' . ◀

In some cases, it is useful to assume effect algebras with more structure than just a preordered monoid. For example, for languages that include fixed points, we might assume an operator $(-)^*$ that assigns to each effect $\varepsilon \in \mathcal{E}$ an effect $\varepsilon^* \in \mathcal{E}$ of a recursive computation. Mycroft et al. [80] discuss several cases in which we want additional structure. In this thesis, to keep the discussion general, we mostly assume only a preordered monoid. We will occasionally require \mathcal{E} to have *bounded binary joins*:

Definition 2.1.5 A *partially ordered monoid* is a preordered monoid $(\mathcal{E}, \leq, \cdot, 1)$ in which \leq is antisymmetric. An element $\varepsilon' \in \mathcal{E}$ is an *upper bound* of $\varepsilon_1, \varepsilon_2 \in \mathcal{E}$ if $\varepsilon_1 \leq \varepsilon'$ and $\varepsilon_2 \leq \varepsilon'$. It is the *join* of ε_1 and ε_2 if additionally $\varepsilon' \leq \varepsilon''$ for all upper bounds ε'' . A partially ordered monoid has *bounded binary joins* if for each pair of elements $\varepsilon_1, \varepsilon_2 \in \mathcal{E}$, existence of an upper bound implies the existence of a join $\varepsilon_1 \vee \varepsilon_2$. ◀

In this definition, we restrict to partial orders so that joins are unique. All of the above examples are partially ordered monoids with bounded binary joins. Joins are also useful e.g. for effect inference, but we do not discuss inference in this thesis.

2.2 Simply-typed lambda calculus

The source language we consider is based on the simply-typed lambda calculus. In later sections of this chapter we consider various evaluation orders, and give effect systems that correspond to

$$\begin{array}{c}
\frac{}{\Gamma \vdash x : \tau} \text{ if } (x : \tau) \in \Gamma \qquad \frac{\Gamma \vdash e : \text{car}_{\text{op}}}{\Gamma \vdash \text{op } e : \text{ar}_{\text{op}}} \qquad \frac{}{\Gamma \vdash () : \text{unit}} \\
\\
\frac{}{\Gamma \vdash \text{true} : \text{bool}} \qquad \frac{}{\Gamma \vdash \text{false} : \text{bool}} \qquad \frac{\Gamma \vdash e_1 : \text{bool} \quad \Gamma \vdash e_2 : \tau \quad \Gamma \vdash e_3 : \tau}{\Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau} \\
\\
\frac{\Gamma, x : \tau \vdash e : \tau'}{\Gamma \vdash \lambda x : \tau. e : \tau \rightarrow \tau'} \qquad \frac{\Gamma \vdash e_1 : \tau \rightarrow \tau' \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash e_1 e_2 : \tau'}
\end{array}$$

Figure 2.1: Source language type system

them. Here we give the syntax and type system of the source language without yet considering how to track effects.

The syntax of source-language types τ and expressions e is as follows:

$$\begin{aligned}
\tau &::= \text{unit} \mid \text{bool} \mid \tau \rightarrow \tau' \\
e &::= x \mid \text{op } e \mid () \mid \text{true} \mid \text{false} \mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \mid \lambda x : \tau. e \mid e_1 e_2
\end{aligned}$$

Most of the syntax is standard. The particular choice of types is unimportant; we choose the unit type, booleans and function types because we use them in examples.

We assume some set Σ of *operations*, ranged over by op . The purpose of these is to provide side-effects. We could for example have operations `get` and `put` for interacting with global state, or `raise` for raising exceptions. Each operation takes exactly one argument; hence the syntax of expressions includes `op e`, which means apply the operation op to the argument e . We can include nullary operations by using `unit` for the argument. If we had included products in the syntax we could also have operations with more than one argument (we cannot get multiple arguments by currying because of the restriction to ground types below). We do not give any examples involving operations of more than one argument, so do not need to consider product types here.

The operations op are *generic effects* [85]. Including operations in this way, rather than for example just assuming an arbitrary collection of constants c , is useful when we consider equivalences between evaluation orders (see Chapter 3).

To specify the type system, we assume an assignment of types for the operations op . The *ground types* are `unit` and `bool` (i.e. types that do not include functions). We restrict the operations to ground types (again this is useful for equivalences between evaluation orders). We collect the data required to specify the type system into a notion of *signature*. Signatures also include an effect algebra. This effect algebra is not used in this section, but is used in each of our effect systems.

Definition 2.2.1 (Signature) A (*source-language*) *signature* consists of the following data:

- A preordered monoid $(\mathcal{E}, \leq, \cdot, 1)$ of *effects*.
- A set Σ of *operations*.
- For each operation $\text{op} \in \Sigma$, ground types car_{op} and ar_{op} , respectively called the *coarity* and *arity* of op , and an effect $\text{eff}_{\text{op}} \in \mathcal{E}$. ◀

$$\begin{array}{c}
\frac{}{\Gamma \vdash_v x : \tau \& 1} \text{if } (x : \tau) \in \Gamma \qquad \frac{\Gamma \vdash_v e : \text{car}_{\text{op}} \& \varepsilon}{\Gamma \vdash_v \text{op } e : \text{ar}_{\text{op}} \& \varepsilon \cdot \text{eff}_{\text{op}}} \qquad \frac{}{\Gamma \vdash_v () : \text{unit} \& 1} \\
\\
\frac{}{\Gamma \vdash_v \text{true} : \text{bool} \& 1} \qquad \frac{}{\Gamma \vdash_v \text{false} : \text{bool} \& 1} \\
\\
\frac{\Gamma \vdash_v e_1 : \text{bool} \& \varepsilon \qquad \Gamma \vdash_v e_2 : \tau \& \varepsilon' \qquad \Gamma \vdash_v e_3 : \tau \& \varepsilon'}{\Gamma \vdash_v \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau \& \varepsilon \cdot \varepsilon'} \\
\\
\frac{\Gamma, x : \tau \vdash_v e : \tau' \& \varepsilon}{\Gamma \vdash_v \lambda x : \tau. e : \tau \xrightarrow{\varepsilon} \tau' \& 1} \qquad \frac{\Gamma \vdash_v e_1 : \tau \xrightarrow{\varepsilon_3} \tau' \& \varepsilon_1 \qquad \Gamma \vdash_v e_2 : \tau \& \varepsilon_2}{\Gamma \vdash_v e_1 e_2 : \tau' \& \varepsilon_1 \cdot \varepsilon_2 \cdot \varepsilon_3} \\
\\
\frac{\Gamma \vdash_v e : \tau \& \varepsilon}{\Gamma \vdash_v e : \tau \& \varepsilon'} \text{if } \varepsilon \leq \varepsilon'
\end{array}$$

Figure 2.2: Call-by-value effect system

The coarity of an operation is the type of its argument; the arity is the type of its result.² For example, for global state storing a value of type **bool**, we would have an operation `get` with coarity $\text{car}_{\text{get}} := \text{unit}$ and arity $\text{ar}_{\text{get}} := \text{bool}$, and an operation `put` with coarity **bool** and arity **unit**. For binary nondeterministic choice we would have an operation `flip` with coarity **unit** and arity **bool**, which nondeterministically chooses between **true** and **false**. (We can use this to make one form of nondeterministic choice between two expressions e_1, e_2 of an arbitrary type by writing `if flip () then e_1 else e_2` .)

Each operation `op` is also associated with the effect eff_{op} of using it. One option would be to use a Gifford-style effect algebra (Section 2.1), so that $\text{eff}_{\text{op}} := \{\text{op}\}$ for each $\text{op} \in \Sigma$, and the preordered monoid is $(\mathcal{P}\Sigma, \subseteq, \cup, \emptyset)$. However, we allow other effect algebras to be used.

We assume a fixed signature to specify the type system. A *typing context* Γ is an ordered list of (variable, type) pairs such that no variable appears more than once. We write \diamond for the empty typing context. The typing judgment $\Gamma \vdash e : \tau$ is defined inductively by the rules in Figure 2.1. Rules that add variables to typing contexts implicitly assume that those variables are fresh. (This is the case for all of the typing rules in this thesis.)

2.3 Call-by-value

We augment the source language type system so that it tracks the effects of expressions. The first effect system we present is well-known: it is for the traditional example of a call-by-value source language. The call-by-value effect system consists of a typing judgment of the form $\Gamma \vdash_v e : \tau \& \varepsilon$, which assigns to each expression e an effect ε in addition to the type τ .

²The reader may be concerned that these are the wrong way around, since the arity should specify the type of the argument and the coarity the type of the result. However, *generic effects* with argument type τ and result type τ' are in bijection with *algebraic operations* [85], which are τ'' -indexed families of maps from $(\tau' \rightarrow \tau'')$ to $(\tau \rightarrow \tau'')$ satisfying certain conditions. When taking the algebraic operations view, the arity and coarity are the correct way around. Our usage of the terms arity and coarity matches e.g. Katsumata [44].

$$\begin{array}{c}
\frac{e \overset{v}{\rightsquigarrow} e'}{\text{op } e \overset{v}{\rightsquigarrow} \text{op } e'} \quad \text{if true then } e_2 \text{ else } e_3 \overset{v}{\rightsquigarrow} e_2 \quad (\lambda x:\tau. e) v \overset{v}{\rightsquigarrow} e[x \mapsto v] \\
\text{if false then } e_2 \text{ else } e_3 \overset{v}{\rightsquigarrow} e_3 \\
\hline
\frac{e_1 \overset{v}{\rightsquigarrow} e'_1}{\text{if } e_1 \text{ then } e_2 \text{ else } e_3 \overset{v}{\rightsquigarrow} \text{if } e'_1 \text{ then } e_2 \text{ else } e_3} \quad \frac{e_1 \overset{v}{\rightsquigarrow} e'_1}{e_1 e_2 \overset{v}{\rightsquigarrow} e'_1 e_2} \quad \frac{e_2 \overset{v}{\rightsquigarrow} e'_2}{(\lambda x:\tau. e) e_2 \overset{v}{\rightsquigarrow} (\lambda x:\tau. e) e'_2}
\end{array}$$

Figure 2.3: Call-by-value reduction relation

To add effects, we replace the syntax of types and expressions with the following:

$$\begin{aligned}
\tau, \tau' &::= \mathbf{unit} \mid \mathbf{bool} \mid \tau \xrightarrow{\varepsilon} \tau' \\
e &::= x \mid \text{op } e \mid () \mid \mathbf{true} \mid \mathbf{false} \mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \mid \lambda x:\tau. e \mid e_1 e_2
\end{aligned}$$

The only difference with the previous syntax is that each function type is annotated with a *latent* effect ε . The latent effect is the effect of applying the function.

We again assume a fixed source-language signature (Definition 2.2.1) to specify the type system. Since arities and coarities are required to be ground types, it does not matter that the syntax of types has changed (the ground types are still **unit** and **bool**). Typing contexts Γ are ordered lists of (variable, type) pairs without repetitions, as before. The typing judgment $\Gamma \vdash_v e : \tau \ \& \ \varepsilon$ is defined inductively by the rules in Figure 2.2. For operations, we multiply the effect of the argument by the effect of the operation itself (the argument is evaluated before the operation is performed). Function application evaluates the function, then the argument, and then the body of the function; hence the effect in the conclusion of the typing rule for application is $\varepsilon_1 \cdot \varepsilon_2 \cdot \varepsilon_3$. A different evaluation order would require a different effect here. The final rule is subeffecting, which allows effects to be overapproximated. The overapproximation is necessary for programs that use **if**: we do not attempt to determine which branch will be taken, so we require both branches to have the same effect.

We give a call-by-value operational semantics for the source language. It consists of a small-step reduction relation $\overset{v}{\rightsquigarrow}$. The definition is straightforward: it is the smallest relation closed under the rules in Figure 2.3. To specify β -reduction of functions, we use the following syntax of *values* v , which form a subset of expressions:

$$v ::= () \mid \mathbf{true} \mid \mathbf{false} \mid \lambda x:\tau. e$$

Values are pure: if $\Gamma \vdash_v v : \tau \ \& \ \varepsilon$ for any effect ε , then $\Gamma \vdash_v v : \tau \ \& \ 1$. (The proof is a trivial induction on the typing derivation.) We also use capture-avoiding substitution $e[x \mapsto v]$. There are no rules for reducing $\text{op } v$. To give a complete operational semantics to an instantiation of the source language, one would augment the rules in the figure to characterize the behaviour of the operations, but we do not consider this here. (Though we do consider this for call-by-push-value; see Definition 2.7.4.)

Call-by-value reductions preserve the effect (and type) assigned by the call-by-value effect system. (Recall that \diamond is the empty typing context.)

Theorem 2.3.1 (Subject reduction for call-by-value) If $\diamond \vdash_v e : \tau \ \& \ \varepsilon$ and $e \overset{v}{\rightsquigarrow} e'$ then $\diamond \vdash_v e' : \tau \ \& \ \varepsilon$. ◀

$$\begin{array}{c}
\frac{}{\Gamma \vdash_{\text{moggi}} x : \tau \& \varepsilon} \text{ if } (x : \tau \& \varepsilon) \in \Gamma \quad \frac{\Gamma \vdash_{\text{moggi}} e : \text{car}_{\text{op}} \& \varepsilon}{\Gamma \vdash_{\text{moggi}} \text{op } e : \text{ar}_{\text{op}} \& \varepsilon \cdot \text{eff}_{\text{op}}} \quad \frac{}{\Gamma \vdash_{\text{moggi}} () : \text{unit} \& 1} \\
\\
\frac{}{\Gamma \vdash_{\text{moggi}} \text{true} : \text{bool} \& 1} \quad \frac{}{\Gamma \vdash_{\text{moggi}} \text{false} : \text{bool} \& 1} \\
\\
\frac{\Gamma \vdash_{\text{moggi}} e_1 : \text{bool} \& \varepsilon \quad \Gamma \vdash_{\text{moggi}} e_2 : \tau \& \varepsilon' \quad \Gamma \vdash_{\text{moggi}} e_3 : \tau \& \varepsilon'}{\Gamma \vdash_{\text{moggi}} \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau \& \varepsilon \cdot \varepsilon'} \\
\\
\frac{\Gamma, x : \tau \& \varepsilon \vdash_{\text{moggi}} e : \tau' \& \varepsilon'}{\Gamma \vdash_{\text{moggi}} \lambda x : \tau \& \varepsilon. e : \tau \xrightarrow{\varepsilon, \varepsilon'} \tau' \& 1} \quad \frac{\Gamma \vdash_{\text{moggi}} e_1 : \tau \xrightarrow{\varepsilon_2, \varepsilon_3} \tau' \& \varepsilon_1 \quad \Gamma \vdash_{\text{moggi}} e_2 : \tau \& \varepsilon_2}{\Gamma \vdash_{\text{moggi}} e_1 e_2 : \tau' \& \varepsilon_1 \cdot \varepsilon_3} \\
\\
\frac{\Gamma \vdash_{\text{moggi}} e : \tau \& \varepsilon}{\Gamma \vdash_{\text{moggi}} e : \tau \& \varepsilon'} \text{ if } \varepsilon \leq \varepsilon'
\end{array}$$

Figure 2.4: Call-by-name effect system

2.4 Moggi-style call-by-name

The effect system given in the previous section is not suitable for call-by-name evaluation. This is because it assigns the effect 1 to each variable, but in call-by-name, variables may have side-effects. Subject reduction does not hold in general if we use call-by-name reduction instead of call-by-value.

In this section, we give an effect system for call-by-name. We refer to this as *Moggi-style* call-by-name here because the effect system is partly based on Moggi's monadic semantics for call-by-name [78]. It is distinguished from *Levy-style* call-by-name (Section 2.6) in that side-effects can occur at any type, rather than just at base types.³

Unlike the call-by-value effect system, a call-by-name effect system should not assign the effect 1 to every variable. To assign the correct effects to variables, we change the notion of typing context to also include effects. For Moggi-style call-by-name, a typing context is an ordered list of triples of the form $x : \tau \& \varepsilon$, where ε is the effect associated with the variable x . The syntax is changed so that function types are annotated with the effect ε of the argument as well as the latent effect ε' of the function, and we also annotate lambdas with the effect of the argument:

$$\begin{aligned}
\tau, \tau' &::= \text{unit} \mid \text{bool} \mid \tau \xrightarrow{\varepsilon, \varepsilon'} B \\
e &::= c \mid x \mid () \mid \text{true} \mid \text{false} \mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \mid \lambda x : \tau \& \varepsilon. e \mid e_1 e_2
\end{aligned}$$

A practical effect system would probably add effect polymorphism to this, so that each function can be applied to arguments of various effects. The syntax of expressions e is the same as for call-by-value.

The call-by-name effect system is parameterized by the same notion of signature (Definition 2.2.1) as the other source-language type systems. The typing judgment $\Gamma \vdash_{\text{moggi}} e : \tau \& \varepsilon$ is defined by the rules in Figure 2.4. Compared to the call-by-value effect system, only the rules

³Benton et al. [7] refer to Levy-style call-by-name as *Algol-like*.

$$\begin{array}{c}
\frac{e \overset{\text{moggi}}{\rightsquigarrow} e'}{\text{op } e \overset{\text{moggi}}{\rightsquigarrow} \text{op } e'} \quad \text{if true then } e_2 \text{ else } e_3 \overset{\text{moggi}}{\rightsquigarrow} e_2 \quad \text{if false then } e_2 \text{ else } e_3 \overset{\text{moggi}}{\rightsquigarrow} e_3 \quad (\lambda x : \tau \& \varepsilon. e) e' \overset{\text{moggi}}{\rightsquigarrow} e[x \mapsto e'] \\
\frac{e_1 \overset{\text{moggi}}{\rightsquigarrow} e'_1}{\text{if } e_1 \text{ then } e_2 \text{ else } e_3 \overset{\text{moggi}}{\rightsquigarrow} \text{if } e'_1 \text{ then } e_2 \text{ else } e_3} \quad \frac{e_1 \overset{\text{moggi}}{\rightsquigarrow} e'_1}{e_1 e_2 \overset{\text{moggi}}{\rightsquigarrow} e'_1 e_2}
\end{array}$$

Figure 2.5: Call-by-name reduction relation

for variables, lambda abstraction and application change. The rule for application requires the effect ε_2 of the argument to be the same as the annotation on the function type. The conclusion of this rule does not mention the effect ε_2 , since the expression e_2 is only evaluated if the function uses its argument. If it does, the latent effect ε_3 will reflect this. The typing rule for operations has not changed: operations still evaluate their arguments eagerly.

We also define a call-by-name operational semantics, consisting of a small-step reduction relation $\overset{\text{moggi}}{\rightsquigarrow}$. Again the definition is straightforward; it is given in Figure 2.5. Call-by-name reductions preserve effects.

Theorem 2.4.1 (Subject reduction for Moggi-style call-by-name) If $\diamond \vdash_{\text{moggi}} e : \tau \& \varepsilon$ and $e \overset{\text{moggi}}{\rightsquigarrow} e'$ then $\diamond \vdash_{\text{moggi}} e' : \tau \& \varepsilon$. \blacktriangleleft

2.5 Graded monadic metalanguage

So far, this chapter has given two source-language effect systems: one for call-by-value and one for Moggi-style call-by-name. We now consider an intermediate language that captures both of these evaluation orders, by making (side-effecting) computations into a first-class notion.

The intermediate language we describe is the *graded monadic metalanguage* (GMM), which is essentially the monadic metalanguage [78]: it consists of a pure fragment together with a type constructor that forms types of effectful computations. The primary difference is that the type constructor is *graded* by the effect algebra. GMM is similar to languages described by Katsumata [45] and by Gaboardi et al. [29].

The syntax of GMM types A, B and terms M, N is as follows:

$$\begin{array}{ll}
A, B ::= b & M, N ::= c \mid x \\
\mid \mathbf{unit} & \mid () \\
\mid A_1 \times A_2 & \mid (M_1, M_2) \mid \mathbf{fst } M \mid \mathbf{snd } M \\
\mid \mathbf{empty} & \mid \mathbf{case}_A M \mathbf{ of } \{ \} \\
\mid A_1 + A_2 & \mid \mathbf{inl}_{A_2} M \mid \mathbf{inr}_{A_1} M \mid \mathbf{case } M \mathbf{ of } \{ \mathbf{inl } x_1. N_1, \mathbf{inr } x_2. N_2 \} \\
\mid A \rightarrow B & \mid \lambda x : A. M \mid M N \\
\mid \langle \varepsilon \rangle A & \mid \mathbf{op } M \mid \langle M \rangle \mid \mathbf{let } \langle x \rangle = M \mathbf{ in } N \mid \mathbf{coerce}_{\varepsilon \leq \varepsilon'} M
\end{array}$$

We include a richer syntax of types than we do for source languages. The non-highlighted part of the syntax is just the simply-typed lambda calculus, with base types b , unit type, products, empty type, sum types, and function types. GMM does not have general recursion built in (but

it can be added). The terms include constants c , and the eliminator $\text{case}_A M \text{ of } \{\}$ of the empty type. The constants c are intended to introduce elements of base types; for example, we could have a base type nat of natural numbers and constants $\text{zero} : \text{nat}$ and $\text{suc} : \text{nat} \rightarrow \text{nat}$ for zero and successor. They are not intended to provide side-effects (this is what the operations op are for). Modulo constants, there are no closed terms of type empty . (This does not mean the empty type is useless: there may still be closed terms of type $\langle \varepsilon \rangle \text{empty}$.) We sometimes omit the type ascriptions (e.g. writing $\text{inl } M$ instead of $\text{inl}_{A_2} M$). We also write $_$ for an arbitrary fresh variable.

The highlighted part of the syntax concerns effectful computations. The type $\langle \varepsilon \rangle A$ contains computations that return results of type A , and have side-effects restricted to $\varepsilon \in \mathcal{E}$. For example, using a Gifford-style effect algebra (Example 2.1.2), computations $M : \langle \{\text{raise}, \text{get}\} \rangle A$ may use raise and get , but do not use the put operation. The type $\langle \varepsilon \rangle A$ is the same as the monadic type TA of Moggi's language, except for the restriction on side-effects. We refer to the family of type constructors $\langle - \rangle$ as the *graded monad*. The intention is that *all* side-effects are encapsulated in the graded monad, including e.g. divergence.

The term $\text{op } M$ is the application of the operation op to the argument M . The term $\langle M \rangle$ is a computation that immediately returns M , causing no side-effects (this is also commonly written $\text{return } M$). Computations are sequenced using $\text{let } \langle x \rangle = M \text{ in } N$, which runs M , binds x to the result (if any) and then runs N . (This is the same as $M \gg= \lambda x. N$ in Haskell.) The variable x is bound inside N but not M . We emphasize that the evaluation order of $\text{let } \langle X \rangle = M \text{ in } N$ is fixed: M is always evaluated eagerly. It is possible to express various evaluation orders in GMM because computations are first-class. For example, it is possible to duplicate a computation by passing it to a function. Computations with effect ε can be regarded as computations with a less restrictive effect ε' using $\text{coerce}_{\varepsilon \leq \varepsilon'}$. This has the same purpose as the subeffecting rule in the above source-language effect systems. Subeffecting in GMM is explicit so that the effect system is syntax-directed.

As usual, we instantiate GMM with different side-effects by choosing different operations. For example, defining $\text{bool} := \text{unit} + \text{unit}$, we could have the following (recall that the coarity is the type of the argument of the operation, and the arity is the type of the output).

Operation op	Coarity car_{op}	Arity ar_{op}	Effect eff_{op}
raise	unit	empty	$\{\text{raise}\}$
get	unit	bool	$\{\text{get}\}$
put	bool	unit	$\{\text{put}\}$

Again the arity and coarity of each operation are required to be ground types. In GMM, ground types G are types that do not contain function types or the graded monad:

$$G ::= b \mid \text{unit} \mid G_1 \times G_2 \mid \text{empty} \mid G_1 + G_2$$

GMM is parameterized by the base types b , operations, effect algebra and constants c .

Definition 2.5.1 A GMM signature consists of the following data:

- A preordered monoid $(\mathcal{E}, \leq, \cdot, 1)$ of *effects*.
- A set \mathcal{B} of *base types*.
- A type-indexed family of pairwise disjoint sets \mathcal{K}_A of *constants of type A*.
- A set Σ of *operations*.
- For each operation $\text{op} \in \Sigma$, ground types car_{op} and ar_{op} , respectively called the *coarity* and *arity* of op , and an effect $\text{eff}_{\text{op}} \in \mathcal{E}$. ◀

$$\begin{array}{c}
\frac{}{\Gamma \vdash_{\text{mon}} c : A} \text{ if } c \in \mathcal{K}_A \qquad \frac{}{\Gamma \vdash_{\text{mon}} x : A} \text{ if } (x : A) \in \Gamma \qquad \frac{}{\Gamma \vdash_{\text{mon}} () : \text{unit}} \\
\frac{\Gamma \vdash_{\text{mon}} M_1 : A_1 \quad \Gamma \vdash_{\text{mon}} M_2 : A_2}{\Gamma \vdash_{\text{mon}} (M_1, M_2) : A_1 \times A_2} \qquad \frac{\Gamma \vdash_{\text{mon}} M : A_1 \times A_2}{\Gamma \vdash_{\text{mon}} \text{fst } M : A_1} \qquad \frac{\Gamma \vdash_{\text{mon}} M : A_1 \times A_2}{\Gamma \vdash_{\text{mon}} \text{snd } M : A_2} \\
\frac{\Gamma \vdash_{\text{mon}} M : \text{empty}}{\Gamma \vdash_{\text{mon}} \text{case}_A M \text{ of } \{ \} : A} \\
\frac{\Gamma \vdash_{\text{mon}} M : A_1}{\Gamma \vdash_{\text{mon}} \text{inl}_{A_2} M : A_1 + A_2} \qquad \frac{\Gamma \vdash_{\text{mon}} M : A_2}{\Gamma \vdash_{\text{mon}} \text{inr}_{A_1} M : A_1 + A_2} \\
\frac{\Gamma \vdash_{\text{mon}} M : A_1 + A_2 \quad \Gamma, x_1 : A_1 \vdash_{\text{mon}} N_1 : B \quad \Gamma, x_2 : A_2 \vdash_{\text{mon}} N_2 : B}{\Gamma \vdash_{\text{mon}} \text{case } M \text{ of } \{ \text{inl } x_1. N_1, \text{inr } x_2. N_2 \} : B} \\
\frac{\Gamma, x : A \vdash_{\text{mon}} M : B}{\Gamma \vdash_{\text{mon}} \lambda x : A. M : A \rightarrow B} \qquad \frac{\Gamma \vdash_{\text{mon}} M : A \rightarrow B \quad \Gamma \vdash_{\text{mon}} N : A}{\Gamma \vdash_{\text{mon}} MN : B} \\
\frac{\Gamma \vdash_{\text{mon}} M : \text{car}_{\text{op}}}{\Gamma \vdash_{\text{mon}} \text{op } M : \langle \varepsilon \rangle \text{ar}_{\text{op}}} \qquad \frac{\Gamma \vdash_{\text{mon}} M : A}{\Gamma \vdash_{\text{mon}} \langle M \rangle : \langle 1 \rangle A} \\
\frac{\Gamma \vdash_{\text{mon}} M : \langle \varepsilon \rangle A \quad \Gamma, x : A \vdash_{\text{mon}} N : \langle \varepsilon' \rangle B}{\Gamma \vdash_{\text{mon}} \text{let } \langle x \rangle = M \text{ in } N : \langle \varepsilon \cdot \varepsilon' \rangle B} \qquad \frac{\Gamma \vdash_{\text{mon}} M : \langle \varepsilon \rangle A}{\Gamma \vdash_{\text{mon}} \text{coerce}_{\varepsilon \leq \varepsilon'} M : \langle \varepsilon' \rangle A} \text{ if } \varepsilon \leq \varepsilon'
\end{array}$$

Figure 2.6: GMM typing rules

We assume a fixed GMM signature for the rest of this section.

Typing contexts Γ are ordered lists of pairs of variable names and types such that each variable appears at most once. The typing judgment has the form $\Gamma \vdash_{\text{mon}} M : A$. By contrast with the above effect systems, it is not annotated with an extra effect: in GMM, effects only appear inside types. The typing rules are given in Figure 2.6. Each of the rules that add variables to the typing context implicitly assume the variable is fresh. The rules are syntax-directed, so given a context Γ and term M , there is at most one type A such that $\Gamma \vdash_{\text{mon}} M : A$, and if $\Gamma \vdash_{\text{mon}} M : A$ holds then it has exactly one derivation.

The non-highlighted part consists of the standard typing rules of the simply-typed lambda calculus. The rules for the graded monad are highlighted. For operations, we assume a pure term M of type car_{op} . In the source language effect systems, the arguments of operations are computations, and the effect system reflects the fact that they are eagerly evaluated. Here, if we wish to use a computation M as an argument to an operation, it must be evaluated first. If M has type $\langle \varepsilon \rangle \text{car}_{\text{op}}$, we can do this by writing

$$\text{let } \langle x \rangle = M \text{ in op } x$$

which has type $\langle \varepsilon \cdot \text{eff}_{\text{op}} \rangle \text{ar}_{\text{op}}$.

$$\begin{array}{c}
\text{fst } (M_1, M_2) \equiv M_1 \\
\left(\text{case } \text{inl}_{A_2} M \text{ of } \begin{array}{l} \{\text{inl } x_1. N_1 \\ , \text{inr } x_2. N_2\} \end{array} \right) \equiv N_1[x_1 \mapsto M] \\
(\lambda x : A. M) N \equiv M[x \mapsto N]
\end{array}
\qquad
\begin{array}{c}
\text{snd } (M_1, M_2) \equiv M_2 \\
\left(\text{case } \text{inr}_{A_1} M \text{ of } \begin{array}{l} \{\text{inl } x_1. N_1 \\ , \text{inr } x_2. N_2\} \end{array} \right) \equiv N_2[x_2 \mapsto M] \\
\text{let } \langle x \rangle = \langle M \rangle \text{ in } N \equiv N[x \mapsto M]
\end{array}$$

$$\begin{array}{c}
M \equiv () \\
M \equiv \text{case}_A N \text{ of } \{\} \\
M \equiv \lambda x : A. M x \\
M \equiv \text{let } \langle x \rangle = M \text{ in } \langle x \rangle
\end{array}
\qquad
\begin{array}{c}
M \equiv (\text{fst } M, \text{snd } M) \\
M[x \mapsto N] \equiv \left(\text{case } N \text{ of } \begin{array}{l} \{\text{inl } y_1. M[x \mapsto \text{inl}_{A_2} y_1] \\ , \text{inr } y_2. M[x \mapsto \text{inr}_{A_1} y_2]\} \end{array} \right)
\end{array}$$

$$\begin{array}{c}
\text{let } \langle y \rangle = (\text{let } \langle x \rangle = M_1 \text{ in } M_2) \text{ in } M_3 \equiv \text{let } \langle x \rangle = M_1 \text{ in } \text{let } \langle y \rangle = M_2 \text{ in } M_3 \\
\text{coerce}_{\varepsilon_1 \cdot \varepsilon_2 \leq \varepsilon'_1 \cdot \varepsilon'_2} (\text{let } \langle x \rangle = M \text{ in } N) \equiv \text{let } \langle x \rangle = \text{coerce}_{\varepsilon_1 \leq \varepsilon'_1} M \text{ in } \text{coerce}_{\varepsilon_2 \leq \varepsilon'_2} N \\
\text{coerce}_{\varepsilon \leq \varepsilon'} M \equiv M \\
\text{coerce}_{\varepsilon' \leq \varepsilon''} (\text{coerce}_{\varepsilon \leq \varepsilon'} M) \equiv \text{coerce}_{\varepsilon \leq \varepsilon''} M
\end{array}$$

Figure 2.7: Axioms of the GMM equational theory

Each of our other three highlighted rules corresponds to one part of the structure of the effect algebra: we use 1 for returning a value, the binary operation \cdot for sequencing two computations, and the order \leq for coercions from a more restrictive effect to a less restrictive effect.

2.5.1 Equational theory

GMM is designed to prove validity of program transformations. We therefore require a notion of *contextual equivalence*, which specifies when it is correct to replace one program with another. This is defined in terms of an *equational theory* (rather than an operational semantics). The equational theory for GMM consists of a judgment of the form $\Gamma \vdash_{\text{mon}} M \equiv N : A$, which means M and N both have type A in context Γ , and they have the same semantics. We sometimes write this just as $M \equiv N$.

The axioms of the equational theory are listed in Figure 2.7. Since we relate well-typed terms, each of the axioms should be read as assuming that both sides have the same type in the same typing context (but not that they are closed). For example, the axiom $M \equiv ()$ holds when $\Gamma \vdash_{\text{mon}} M : \text{unit}$. The necessary constraints can be derived by looking at the typing rules (because they are syntax-directed). As usual, the non-highlighted part of the figure is standard from the simply-typed lambda calculus (it consists of the usual β - and η - laws), and the highlighted part describes the behaviour of computations. Each of the axioms of the highlighted part corresponds to one of the requirements in the definition of preordered monoid.

The first group of axioms consists of β -laws. The highlighted β -law states that sequencing can be eliminated if the first computation immediately returns a value. The constraints necessary to ensure that both sides of the axiom have the same type are $\Gamma \vdash_{\text{mon}} M : A$ and

$\Gamma, x : A \vdash_{\text{mon}} N : \langle \varepsilon \rangle B$. When these hold both sides have the same type because of the left unit law of the preordered monoid: $1 \cdot \varepsilon = \varepsilon$.

The second group consists of one η -law for each of the type formers of GMM (excluding base types). The highlighted axiom states that sequencing of computations can be added without changing behaviour by returning the value of the variable x . It corresponds to the right unit law $\varepsilon \cdot 1 = \varepsilon$.

The third group contains the remaining axioms for computations. The first axiom states that sequencing of computations is associative, and corresponds to associativity of \cdot . The second states that coercion commutes with sequencing of computations. The two terms have the same type because \cdot is monotone. (A similar equation can be stated for commutativity of **coerce** with the eliminator for sum types. This equation follows from the other axioms in the figure, so we do not need it as an additional axiom.) The final two axioms of the figure allow us to remove coercions that do not change the effect, and to combine adjacent coercions. They correspond to reflexivity and transitivity of the preorder \leq .

In addition to the core axioms, we also close the equational theory under congruence. To define this, we use a notion of *term context*. A GMM term context $C[\]$ consists of a term with a single *hole* (which we write as \square):

$$\begin{aligned} C[\] ::= & \square \mid (C[\], M_2) \mid (M_1, C[\]) \mid \text{fst } C[\] \mid \text{snd } C[\] \mid \text{case}_A C[\] \text{ of } \{ \} \\ & \mid \text{inl}_{A_2} C[\] \mid \text{inr}_{A_1} C[\] \mid \text{case } C[\] \text{ of } \{ \text{inl } x_1. N_1, \text{inr } x_2. N_2 \} \\ & \mid \text{case } M \text{ of } \{ \text{inl } x_1. C[\], \text{inr } x_2. N_2 \} \mid \text{case } M \text{ of } \{ \text{inl } x_1. N_1, \text{inr } x_2. C[\] \} \\ & \mid \lambda x : A. C[\] \mid C[\] N \mid M C[\] \\ & \mid \text{op } C[\] \mid \langle C[\] \rangle \mid \text{let } \langle x \rangle = C[\] \text{ in } N \mid \text{let } \langle x \rangle = M \text{ in } C[\] \mid \text{coerce}_{\varepsilon \leq \varepsilon'} C[\] \end{aligned}$$

We write $C[M]$ for the term formed by replacing \square with M . The term context may capture some of the free variables of M : for example, if $C[\]$ is $\text{let } \langle x \rangle = M \text{ in } \langle \square \rangle$ then $C[(x, y)]$ is $\text{let } \langle x \rangle = M \text{ in } \langle (x, y) \rangle$; in the latter term, only y is free.

This notion of term context is used in the definition of the GMM equational theory.

Definition 2.5.2 We define $\Gamma \vdash_{\text{mon}} M \equiv N : A$ inductively by:

- Preorder: if $\Gamma \vdash_{\text{mon}} M : A$ then $\Gamma \vdash_{\text{mon}} M \equiv M : A$, and if $\Gamma \vdash_{\text{mon}} M_1 \equiv M_2 : A$ and $\Gamma \vdash_{\text{mon}} M_2 \equiv M_3 : A$ then $\Gamma \vdash_{\text{mon}} M_1 \equiv M_3 : A$.
- Congruence: if $\Gamma \vdash_{\text{mon}} M \equiv N : A$ and both $\Gamma' \vdash_{\text{mon}} C[M] : B$ and $\Gamma' \vdash_{\text{mon}} C[N] : B$ then $\Gamma' \vdash_{\text{mon}} C[M] \equiv C[N] : B$.
- Axioms: If $\Gamma \vdash_{\text{mon}} M : A$ and $\Gamma \vdash_{\text{mon}} N : A$, and $M \equiv N$ is an instance of an axiom in Figure 2.7, then $\Gamma \vdash_{\text{mon}} M \equiv N : A$ and $\Gamma \vdash_{\text{mon}} N \equiv M : A$. \blacktriangleleft

For a fixed typing context Γ and type A , the equational theory is an equivalence relation on terms of type A in context Γ . It is also closed under substitution: if $x_1 : A_1, \dots, x_n : A_n \vdash_{\text{mon}} N \equiv N' : B$, and $\Gamma \vdash_{\text{mon}} M_i \equiv M'_i : A_i$ for each i , then

$$\Gamma \vdash_{\text{mon}} N[x_1 \mapsto M_1, \dots, x_n \mapsto M_n] \equiv N'[x_1 \mapsto M'_1, \dots, x_n \mapsto M'_n] : B$$

Two terms that are syntactically equal up to the placement of coercions are also usually related by the equational theory. Given a term M , define $\lfloor M \rfloor$ to be the same as M but with all uses of **coerce** deleted (for example, $\lfloor \lambda x : A. \text{coerce}_{1 \leq \varepsilon} \langle x \rangle \rfloor = \lambda x : A. \langle x \rangle$).

Lemma 2.5.3 Suppose that the effect algebra is a partially ordered monoid with bounded binary joins. Given two terms M_1, M_2 such that $\Gamma \vdash_{\text{mon}} M_1 : A$ and $\Gamma \vdash_{\text{mon}} M_2 : A$, if $\lfloor M_1 \rfloor = \lfloor M_2 \rfloor$ then $\Gamma \vdash_{\text{mon}} M_1 \equiv M_2 : A$. \blacktriangleleft

$$\begin{array}{c}
\langle \mathbf{unit} \rangle := \mathbf{unit} \\
\langle \mathbf{bool} \rangle := \mathbf{unit} + \mathbf{unit} \\
\langle \tau \xrightarrow{\varepsilon} \tau' \rangle := \langle \tau \rangle \rightarrow \langle \varepsilon \rangle \langle \tau' \rangle \\
\langle \diamond \rangle := \diamond \\
\langle \Gamma, x : \tau \rangle := \langle \Gamma \rangle, x : \langle \tau \rangle \\
\frac{}{\langle \Gamma \vdash_v x : \tau \& 1 \rangle := \langle x \rangle} \quad \frac{\langle \Gamma \vdash_v e : \mathbf{car}_{\text{op}} \& \varepsilon \rangle = M}{\langle \Gamma \vdash_v \text{op } e : \mathbf{ar}_{\text{op}} \& \varepsilon \cdot \mathbf{eff}_{\text{op}} \rangle = \mathbf{let} \langle x \rangle = M \mathbf{in} \text{op } x} \\
\frac{}{\langle \Gamma \vdash_v () : \mathbf{unit} \& 1 \rangle := \langle () \rangle} \\
\frac{}{\langle \Gamma \vdash_v \mathbf{true} : \mathbf{bool} \& 1 \rangle := \langle \mathbf{inl}_{\mathbf{unit}} () \rangle} \quad \frac{}{\langle \Gamma \vdash_v \mathbf{false} : \mathbf{bool} \& 1 \rangle := \langle \mathbf{inr}_{\mathbf{unit}} () \rangle} \\
\frac{\langle \Gamma \vdash_v e_1 : \mathbf{bool} \& \varepsilon \rangle = M_1 \quad \langle \Gamma \vdash_v e_2 : \tau \& \varepsilon' \rangle = M_2 \quad \langle \Gamma \vdash_v e_3 : \tau \& \varepsilon' \rangle = M_3}{\langle \Gamma \vdash_v \mathbf{if } e_1 \mathbf{ then } e_2 \mathbf{ else } e_3 : \tau \& \varepsilon \cdot \varepsilon' \rangle := \mathbf{let} \langle x \rangle = M_1 \mathbf{ in case } x \mathbf{ of } \{ \mathbf{inl} _ . M_2, \mathbf{inr} _ . M_3 \}} \\
\frac{\langle \Gamma, x : \tau \vdash_v e : \tau' \& \varepsilon \rangle = M}{\langle \Gamma \vdash_v \lambda x : \tau . e : \tau \xrightarrow{\varepsilon} \tau' \& 1 \rangle := \langle \lambda x : \langle \tau \rangle . M \rangle} \\
\frac{\langle \Gamma \vdash_v e_1 : \tau \xrightarrow{\varepsilon_3} \tau' \& \varepsilon_1 \rangle = M_1 \quad \langle \Gamma \vdash_v e_2 : \tau \& \varepsilon_2 \rangle = M_2}{\langle \Gamma \vdash_v e_1 e_2 : \tau' \& \varepsilon_1 \cdot \varepsilon_2 \cdot \varepsilon_3 \rangle := \mathbf{let} \langle f \rangle = M_1 \mathbf{ in let } \langle x \rangle = M_2 \mathbf{ in } f x} \\
\frac{\langle \Gamma \vdash_v e : \tau \& \varepsilon \rangle = M}{\langle \Gamma \vdash_v e : \tau \& \varepsilon' \rangle := \mathbf{coerce}_{\varepsilon \leq \varepsilon'} M}
\end{array}$$

Figure 2.8: Translation of call-by-value types (top left), contexts (top right), and typing derivations (bottom) into GMM.

A proof of a similar fact for our call-by-push-value effect system (Lemma 2.7.5) is given in Appendix B.1, where we also conjecture a counterexample for an effect algebra that does not have bounded binary joins. The proof for GMM is similar to the proof for call-by-push-value.

2.5.2 Translations into GMM

GMM captures both call-by-value and Moggi-style call-by-name via two translations (one for each evaluation order) from source-language expressions e into GMM terms. They are based on similar translations given by Moggi [77, 78], Wadler [98] and Benton et al. [7]. Translations like these allow us to use intermediate languages to reason about source languages. Having multiple evaluation-order-directed translations enables reasoning about changes in evaluation order in the source language using the intermediate language.

In this section, we assume a source-language signature and GMM signature that are compatible. Recall that the source-language ground types are \mathbf{unit} and \mathbf{bool} . We translate these into GMM ground types:

$$\langle \mathbf{unit} \rangle := \mathbf{unit} \quad \langle \mathbf{bool} \rangle := \mathbf{unit} + \mathbf{unit}$$

(The two translations we give below agree on ground types with this definition.) We require that both signatures have the same effect algebra, and that for each operation op in the source-language signature: op is also included in the GMM signature; both signatures agree on the effect of op ; and if τ and τ' are respectively the coarity and arity assigned to op by the source-language signature, then $\langle\tau\rangle$ and $\langle\tau'\rangle$ are the coarity and arity assigned to op by the GMM signature. We allow the GMM signature to contain additional operations, and place no restrictions on base types or constants.

The first translation maps call-by-value source-language types τ to GMM types $\langle\tau\rangle_{\text{mon}}^{\text{v}}$ and call-by-value contexts Γ to GMM contexts $\langle\Gamma\rangle_{\text{mon}}^{\text{v}}$. It maps *derivations* of $\Gamma \vdash_{\text{v}} e : \tau \ \& \ \varepsilon$ to GMM *terms* $\langle\Gamma \vdash_{\text{v}} e : \tau \ \& \ \varepsilon\rangle_{\text{mon}}^{\text{v}}$. Since the call-by-value effect system is not syntax-directed, derivations are not unique (there is a choice of where to use the subeffecting rule), and different derivations are translated into syntactically different GMM terms. The term $\langle\Gamma \vdash_{\text{v}} e : \tau \ \& \ \varepsilon\rangle_{\text{mon}}^{\text{v}}$ is a computation with effect ε : in context $\langle\Gamma\rangle_{\text{mon}}^{\text{v}}$, it has type $\langle\varepsilon\rangle\langle\tau\rangle_{\text{mon}}^{\text{v}}$. The definition of the call-by-value translation is given in Figure 2.8, where we omit the sub- and superscripts. The important part is the translation of functions. Source-language function types are translated into GMM function types where the result is a computation (so may have side-effects), but the argument is not. Function applications first evaluate the function itself, then the argument, and then the function body (this matches the operational semantics in Figure 2.3).

Call-by-name types τ and contexts Γ are translated into GMM types $\langle\tau\rangle_{\text{mon}}^{\text{moggi}}$ and contexts $\langle\Gamma\rangle_{\text{mon}}^{\text{moggi}}$ respectively. The translation of derivations of $\Gamma \vdash_{\text{moggi}} e : \tau \ \& \ \varepsilon$ has similar typing to the call-by-value translation: $\langle\Gamma \vdash_{\text{moggi}} e : \tau \ \& \ \varepsilon\rangle_{\text{mon}}^{\text{moggi}}$ is a GMM term of type $\langle\varepsilon\rangle\langle\tau\rangle_{\text{mon}}^{\text{moggi}}$ in context $\langle\Gamma\rangle_{\text{mon}}^{\text{moggi}}$. The definition of the call-by-name translation is in Figure 2.9. Many of the cases are the same as for call-by-value. The key differences are for typing contexts and functions. The translation of typing contexts uses the graded monad because variables have side-effects. Similarly, arguments of functions have effects, and therefore are encapsulated in the graded monad. The translation of a function application does not evaluate the argument immediately; instead, the computation is passed to the function, and is evaluated each time the function uses its argument (matching Figure 2.5).

The call-by-value and Moggi-style call-by-name translations have the desired typing:

Lemma 2.5.4

1. If M is the call-by-value translation of a derivation of $\Gamma \vdash_{\text{v}} e : \tau \ \& \ \varepsilon$, then

$$\langle\Gamma\rangle_{\text{mon}}^{\text{v}} \vdash_{\text{mon}} M : \langle\varepsilon\rangle\langle\tau\rangle_{\text{mon}}^{\text{v}}$$

2. If M is the call-by-name translation of a derivation of $\Gamma \vdash_{\text{moggi}} e : \tau \ \& \ \varepsilon$, then

$$\langle\Gamma\rangle_{\text{mon}}^{\text{moggi}} \vdash_{\text{mon}} M : \langle\varepsilon\rangle\langle\tau\rangle_{\text{mon}}^{\text{moggi}} \quad \blacktriangleleft$$

Although it is important that we translate derivations rather than just well-typed expressions, for suitable effect algebras it does not matter which derivation is chosen up to the equational theory:

Lemma 2.5.5 Suppose that the effect algebra $(\mathcal{E}, \leq, \cdot, 1)$ is a partially ordered monoid with bounded binary joins.

1. If M and M' are call-by-value translations of derivations of $\Gamma \vdash_{\text{v}} e : \tau \ \& \ \varepsilon$, then $M \equiv M'$.
2. If M and M' are call-by-name translations of derivations of $\Gamma \vdash_{\text{moggi}} e : \tau \ \& \ \varepsilon$, then $M \equiv M'$.

This follows immediately from Lemma 2.5.3. Hence when considering GMM terms up to \equiv , we can refer to *the* translation of a well-typed expression. We do this in the following theorem.

$$\begin{array}{c}
\langle \mathbf{unit} \rangle := \mathbf{unit} \\
\langle \mathbf{bool} \rangle := \mathbf{unit} + \mathbf{unit} \\
\langle \tau \xrightarrow{\varepsilon, \varepsilon'} \tau' \rangle := \langle \varepsilon \rangle \langle \tau \rangle \rightarrow \langle \varepsilon' \rangle \langle \tau' \rangle \\
\hline
\langle \Gamma \vdash_{\text{moggi}} x : \tau \ \& \ \varepsilon \rangle := x \\
\langle \Gamma \vdash_{\text{moggi}} \text{op } e : \text{ar}_{\text{op}} \ \& \ \varepsilon \cdot \text{eff}_{\text{op}} \rangle := \mathbf{let} \langle x \rangle = M \mathbf{in} \text{op } x \\
\hline
\langle \Gamma \vdash_{\text{moggi}} () : \mathbf{unit} \ \& \ 1 \rangle := \langle () \rangle \\
\hline
\langle \Gamma \vdash_{\text{moggi}} \mathbf{true} : \mathbf{bool} \ \& \ 1 \rangle := \langle \mathbf{inl}_{\mathbf{unit}} () \rangle \quad \langle \Gamma \vdash_{\text{moggi}} \mathbf{false} : \mathbf{bool} \ \& \ 1 \rangle := \langle \mathbf{inr}_{\mathbf{unit}} () \rangle \\
\hline
\langle \Gamma \vdash_{\text{moggi}} e_1 : \mathbf{bool} \ \& \ \varepsilon \rangle = M_1 \quad \langle \Gamma \vdash_{\text{moggi}} e_2 : \tau \ \& \ \varepsilon' \rangle = M_2 \quad \langle \Gamma \vdash_{\text{moggi}} e_3 : \tau \ \& \ \varepsilon' \rangle = M_3 \\
\hline
\langle \Gamma \vdash_{\text{moggi}} \mathbf{if} \ e_1 \ \mathbf{then} \ e_2 \ \mathbf{else} \ e_3 : \tau \ \& \ \varepsilon \cdot \varepsilon' \rangle := \mathbf{let} \langle x \rangle = M_1 \mathbf{in} \ \mathbf{case} \ x \ \mathbf{of} \ \{ \mathbf{inl} _ . M_2, \ \mathbf{inr} _ . M_3 \} \\
\hline
\langle \Gamma, x : \tau \ \& \ \varepsilon \vdash_{\text{moggi}} e : \tau' \ \& \ \varepsilon' \rangle = M \\
\hline
\langle \Gamma \vdash_{\text{moggi}} \lambda x : \tau \ \& \ \varepsilon . e : \tau \xrightarrow{\varepsilon, \varepsilon'} \tau' \ \& \ 1 \rangle := \langle \lambda x : \langle \varepsilon \rangle \langle \tau \rangle . M \rangle \\
\hline
\langle \Gamma \vdash_{\text{moggi}} e_1 : \tau \xrightarrow{\varepsilon_2, \varepsilon_3} \tau' \ \& \ \varepsilon_1 \rangle = M_1 \quad \langle \Gamma \vdash_{\text{moggi}} e_2 : \tau \ \& \ \varepsilon_2 \rangle = M_2 \\
\hline
\langle \Gamma \vdash_{\text{moggi}} e_1 \ e_2 : \tau' \ \& \ \varepsilon_1 \cdot \varepsilon_3 \rangle := \mathbf{let} \langle f \rangle = M_1 \mathbf{in} \ f \ M_2 \\
\hline
\langle \Gamma \vdash_{\text{moggi}} e : \tau \ \& \ \varepsilon \rangle = M \\
\hline
\langle \Gamma \vdash_{\text{moggi}} e : \tau \ \& \ \varepsilon' \rangle := \mathbf{coerce}_{\varepsilon \leq \varepsilon'} M
\end{array}$$

Figure 2.9: Translation of Moggi-style call-by-name types (top left), contexts (top right), and typing derivations (bottom) into GMM.

Theorem 2.5.6 (Soundness) Suppose that the effect algebra $(\mathcal{E}, \leq, \cdot, 1)$ is a partially ordered monoid with bounded binary joins.

1. If $\diamond \vdash_{\text{v}} e : \tau \ \& \ \varepsilon$ and $e \rightsquigarrow^{\text{v}} e'$, then $\langle e \rangle_{\text{mon}}^{\text{v}} \equiv \langle e' \rangle_{\text{mon}}^{\text{v}}$.
2. If $\diamond \vdash_{\text{moggi}} e : \tau \ \& \ \varepsilon$ and $e \rightsquigarrow^{\text{moggi}} e'$, then $\langle e \rangle_{\text{mon}}^{\text{moggi}} \equiv \langle e' \rangle_{\text{mon}}^{\text{moggi}}$. ◀

Both of the translations described in this section are *compositional*: the translation of a typing derivation depends only on the translations of its subderivations (rather than the actual structure of the derivations). This property is crucial for reasoning, because we look at parts of programs rather than complete programs. Compositionality implies it is valid to do so.

2.6 Levy-style call-by-name

As we mentioned above, Moggi-style call-by-name is not the only form of call-by-name. We also discuss *Levy-style* call-by-name.

First we explain why there are two forms of call-by-name. Suppose we are in a language (such as PCF) that has divergence as the only side-effect, and that Ω_τ is a closed diverging term of type τ . Each program M of type **bool** will either diverge or return a boolean. In a domain-theoretic denotational semantics, the denotation of M is an element of 2_\perp , where $2 = \{\text{true}, \text{false}\}$ is discrete, and the operation $(-)_\perp$ freely adds a least element. What should the denotation of a closed term of type **bool** \rightarrow **bool** be? It could be an element of either of the following (where \Rightarrow means function space):

$$(\mathbf{bool}_\perp \Rightarrow \mathbf{bool}_\perp)_\perp \qquad \mathbf{bool}_\perp \Rightarrow \mathbf{bool}_\perp$$

The left distinguishes between $\Omega_{\mathbf{bool} \rightarrow \mathbf{bool}}$ (which has interpretation \perp) and $\lambda x. \Omega_{\mathbf{bool}}$. On the right, there is no separate least element \perp for diverging terms of function type, so $\llbracket \Omega_{\mathbf{bool} \rightarrow \mathbf{bool}} \rrbracket = \llbracket \lambda x. \Omega_{\mathbf{bool}} \rrbracket$.

In PCF there is no *observable* difference between $\Omega_{\mathbf{bool} \rightarrow \mathbf{bool}}$ and $\lambda x. \Omega_{\mathbf{bool}}$, because programs can only use functions by applying them. Hence a denotational semantics does not need to distinguish between these terms, and both interpretations of **bool** \rightarrow **bool** are correct. This breaks down if we add other ways to use functions: in Haskell, $M \text{ `seq` } N$ roughly means evaluate M , discard the result, and then evaluate N ; by passing these two terms as M we can distinguish between them. In languages with `seq`, only the interpretation on the left is valid. However, in languages without `seq`, the interpretation on the right more closely reflects the observable behaviours of terms. In particular, η -expansion preserves the interpretation of functions. The interpretation on the right also has other useful properties: for example, the evident function $\text{swap} : (\tau_1 \rightarrow \tau_2 \rightarrow \tau') \rightarrow (\tau_2 \rightarrow \tau_1 \rightarrow \tau')$ is an isomorphism $(\tau_1 \rightarrow \tau_2 \rightarrow \tau') \cong (\tau_2 \rightarrow \tau_1 \rightarrow \tau')$ (by which we mean $\llbracket \text{swap} (\text{swap } e) \rrbracket = \llbracket e \rrbracket$). Both of the two interpretations of call-by-name have advantages.

The Moggi-style call-by-name translation we define in Figure 2.9 takes the option on the left (think of $\langle \varepsilon \rangle$ as being $(-)_\perp$). It assumes that it might be possible to observe side-effects at any type. *Levy-style* call-by-name corresponds to the option on the right. It is the form of call-by-name primarily considered by Levy [56].⁴ For Levy-style, we assume that side-effects can only be observed at base types. This is the case for the source language we consider (and so Levy-style is suitable) because we have not included any `seq`-like constructs.

To describe a source-language effect system for Levy-style call-by-name, we attach the effects ε to base types (rather than having an extra parameter in the typing judgment). The syntax of types is therefore

$$\tau, \tau' ::= \langle \varepsilon \rangle \mathbf{unit} \mid \langle \varepsilon \rangle \mathbf{bool} \mid \tau \rightarrow \tau'$$

The similarity with the graded monad $\langle - \rangle$ of GMM (Section 2.5) is not accidental. For example, $\langle \varepsilon \rangle \mathbf{bool}$ should be thought of as the type of computations of effect ε that return booleans. Typing contexts Γ are lists of pairs of the form $x : \tau$. The syntax of expressions is similar to the other source languages:

$$e ::= x \mid \text{op } e \mid () \mid \mathbf{true} \mid \mathbf{false} \mid \mathbf{if } e_1 \mathbf{ then } e_2 \mathbf{ else } e_3 \mid \lambda x : \tau. e \mid e_1 e_2$$

The effect system is again parameterized by a source-language signature (Definition 2.2.1). The coarity and arity of each operation is either **unit** or **bool**; these are not annotated with effects, and hence are not Levy-style call-by-name types (but $\langle \varepsilon \rangle \tau$ is for $\tau \in \{\mathbf{unit}, \mathbf{bool}\}$). The

⁴And this is why we call it *Levy-style*, though it has been considered previously. For example, it is used for PCF and idealized Algol [90], and is discussed by Benton et al. [7].

$$\begin{array}{c}
\frac{}{\Gamma \vdash_n x : \tau} \text{ if } (x : \tau) \in \Gamma \qquad \frac{\Gamma \vdash_n e : \langle \varepsilon \rangle \text{car}_{\text{op}}}{\Gamma \vdash_n \text{op } e : \langle \varepsilon \cdot \text{eff}_{\text{op}} \rangle \text{ar}_{\text{op}}} \qquad \frac{}{\Gamma \vdash_n () : \langle 1 \rangle \text{unit}} \\
\frac{}{\Gamma \vdash_n \text{true} : \langle 1 \rangle \text{bool}} \qquad \frac{}{\Gamma \vdash_n \text{false} : \langle 1 \rangle \text{bool}} \qquad \frac{\Gamma \vdash_n e_1 : \langle \varepsilon \rangle \text{bool} \quad \Gamma \vdash_n e_2 : \tau \quad \Gamma \vdash_n e_3 : \tau}{\Gamma \vdash_n \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \langle \varepsilon \rangle \tau} \\
\frac{\Gamma, x : \tau \vdash_n e : \tau'}{\Gamma \vdash_n \lambda x : A. e : \tau \rightarrow \tau'} \qquad \frac{\Gamma \vdash_n e_1 : \tau \rightarrow \tau' \quad \Gamma \vdash_n e_2 : \tau}{\Gamma \vdash_n e_1 e_2 : \tau'} \qquad \frac{\Gamma \vdash_n e : \langle \varepsilon \rangle \tau}{\Gamma \vdash_n e : \langle \varepsilon' \rangle \tau} \text{ if } \varepsilon \leq \varepsilon'
\end{array}$$

Figure 2.10: Effect system for Levy-style call-by-name

typing judgment has the form $\Gamma \vdash_n e : \tau$. To define it, we use a function $\langle \varepsilon \rangle$ that adds the effect ε to types. The type $\langle \varepsilon \rangle \tau$ is given by:

$$\langle \varepsilon \rangle (\langle \varepsilon' \rangle \text{unit}) := \langle \varepsilon \cdot \varepsilon' \rangle \text{unit} \qquad \langle \varepsilon \rangle (\langle \varepsilon' \rangle \text{bool}) := \langle \varepsilon \cdot \varepsilon' \rangle \text{bool} \qquad \langle \varepsilon \rangle (\tau \rightarrow \tau') := \tau \rightarrow \langle \varepsilon \rangle \tau'$$

The typing judgment is inductively defined by the rules of Figure 2.10. Most of the rules are standard; in particular, the rules for λ -abstraction and application are the same as for the simply-typed λ -calculus. Hence η -expansion preserves types: if $\Gamma \vdash_n e : \tau \rightarrow \tau'$ then $\Gamma \vdash_n \lambda x : \tau. e x : \tau \rightarrow \tau'$. (This is the key property that we would expect if we can only observe effects at base types. It is not true for the Moggi-style effect system.) In the rule for **if** e_1 **then** e_2 **else** e_3 , the operation $\langle \varepsilon \rangle$ is used because the expression e_1 (with effect ε) is evaluated before evaluating either e_2 or e_3 . The subeffecting rule is similar to those in our other effect systems.

The Levy-style call-by-name operational semantics consists of a small-step reduction relation $\overset{n}{\rightsquigarrow}$. The rules defining this are exactly the same as for Moggi-style call-by-name (Figure 2.5), except for the slightly different λ -abstraction syntax. The Levy-style call-by-name effect system also satisfies the following subject reduction theorem. (Recall that \diamond is the empty typing context.)

Theorem 2.6.1 (Subject reduction for Levy-style call-by-name) If $\diamond \vdash_n e : \tau$ and $e \overset{n}{\rightsquigarrow} e'$ then $\diamond \vdash_n e' : \tau$. \blacktriangleleft

A natural question to ask is whether we can add Levy-style call-by-name to the translations into the graded monadic metalanguage described in Section 2.5.2. Filinski [24] defines a *generalized let* construct that can be used to do this. However, as explained by Levy [56], the result is difficult to reason about because the translation of **if** depends on the type τ . This motivates *call-by-push-value*, which is designed to admit a more suitable translation of Levy-style call-by-name.

2.7 Graded call-by-push-value

The language we use for the remainder of this thesis is Levy's *call-by-push-value* (CBPV) [56]. CBPV has several advantages over other intermediate languages such as the monadic metalanguage. The main advantage for our purposes is that it allows us to express more evaluation

orders. There is a compositional translation from Moggi’s monadic metalanguage (and therefore from call-by-value and Moggi-style call-by-name), and a compositional translation from Levy-style call-by-name. As for the other languages in this chapter, we wish to use an effect system to enable effect-dependent reasoning. We therefore present *graded call-by-push-value* (GCBPV), which is similar to CBPV except that the type system tracks effect information. Before doing this, we take a closer look at how CBPV allows programs to control their evaluation order.

CBPV stratifies terms into *values*, which do not have side-effects, and *computations*, which might.⁵ Evaluation order is irrelevant for values (because of the lack of side-effects), so we only need to be careful about how computations are sequenced. There is exactly one CBPV primitive that causes the evaluation of more than one computation. The computation M to x . N runs the computation M , binds x to the result, and then runs the computation N . (It is similar to $\text{let } \langle x \rangle = M \text{ in } N$ in GMM.) The evaluation order is fixed: M is always evaluated eagerly. To allow more control over evaluation order, CBPV allows computations to be thunked. The term $\text{thunk } M$ is a value that contains the thunk of the computation M . Thunks can be duplicated (to allow a single computation to be evaluated more than once), and discarded (so that the computation is not evaluated). If V is a thunk, it can be converted back into a computation with $\text{force } V$.

2.7.1 Syntax

The syntax of graded call-by-push-value is as follows. It is similar to the syntax of ordinary call-by-push-value, except that it tracks effect information.⁶

$A, B ::= b$ $ \text{unit}$ $ A_1 \times A_2$ $ \text{empty}$ $ A_1 + A_2$ $ \underline{U} \underline{C}$	$V, W ::= c \mid x$ $ ()$ $ (V_1, V_2) \mid \text{fst } V \mid \text{snd } V$ $ \text{case}_A V \text{ of } \{ \}$ $ \text{inl}_{A_2} V \mid \text{inr}_{A_1} V \mid \text{case } V \text{ of } \{ \text{inl } x_1. W_1, \text{inr } x_2. W_2 \}$ $ \text{thunk } M$
$\underline{C}, \underline{D} ::= \underline{\text{unit}}$ $ \underline{C}_1 \times \underline{C}_2$ $ A \rightarrow \underline{C}$ $ \langle \varepsilon \rangle A$	$M, N ::= \lambda \{ \}$ $ \lambda \{ 1. M_1, 2. M_2 \} \mid 1' M \mid 2' M$ $ \lambda x : A. M \mid V' M$ $ \text{op } V \mid \langle V \rangle \mid M \text{ to } x. N \mid \text{coerce}_{\varepsilon \leq \varepsilon'} M$ $ \text{force } V$

Types are stratified into *value types* A, B and *computation types* $\underline{C}, \underline{D}$; terms are stratified into *value terms* V, W and *computation terms* M, N . Value types include base types b , the unit type, product types, the empty type, and sum types. As for GMM, we include a collection of constants

⁵There are several calculi that similarly capture various evaluation orders by stratifying their syntax into two parts, such as *polarized* calculi [102]. These have various viewpoints on the distinction between the two parts, and these affect the design of the language. We return to this issue when we add call-by-need in Chapter 5.

⁶Other than tracking effect information, the only difference with ordinary CBPV is that eliminators of product and sum types are value terms rather than computation terms (which makes value terms slightly more general). Levy [56] calls this CBPV with *complex values*. This influences how we think about the split between values and computations: with complex values, value terms do not have side-effects when they are executed; without complex values, values do not reduce at all.

c. All GCBPV constants are value terms; we do not need constants on the computation level (this is discussed further below). Values also include *thunks*. The value type \mathbf{UC} contains thunks of computations of type \underline{C} . The value term **thunk** V suspends a computation M ; the computation term **force** V runs a suspended computation V . Computation types include a unit type **unit** (introduced by the computation term $\lambda\{\}$), and *lazy* binary products $\underline{C}_1 \times \underline{C}_2$. Pairs of computation terms are written $\lambda\{1. M_1, 2. M_2\}$, and the first and second projections are $1'M$ and $2'M$. Laziness means $i'M$ evaluates only the i th component of the pair; the other component (and its side-effects) are ignored. Binary products of value types give us strict pairing. Functions send values to computations, and are computations themselves. Application is written $V'M$, where V is the argument and M is the function to apply.

The highlighted computation type $\langle \varepsilon \rangle A$ is a *returner type*. It contains computations which have side-effects restricted to the effect ε , before potentially returning a value (which must be an element of A). The side-effects could include e.g. divergence (and hence such a computation might never return a value). The corresponding type in ordinary CBPV is usually written $\mathbf{F}A$; we add grading by the effect ε . Application of an operation op to the argument V is written $\text{op } V$. As for GMM, the argument is not a computation; to apply an operation to the result of a computation M we can use $(M \text{ to } x. \text{op } x)$. Returners (i.e. elements of returner types) are introduced by $\langle V \rangle$, which is a computation that immediately returns V (with no side-effects). The eliminator is $M \text{ to } x. N$ (which requires M to have returner type). As mentioned above, this runs M , then binds x to the result and runs N . Notationally, **to** is right-associative (so $M_1 \text{ to } x. M_2 \text{ to } y. M_3$ means $M_1 \text{ to } x. (M_2 \text{ to } y. M_3)$), and variable bindings extend as far to the right as possible (so $\lambda x : A. M_1 \text{ to } y. M_2$ means $\lambda x : A. (M_1 \text{ to } y. M_2)$). Finally, we can coerce computations of returner type from effect ε to a less restrictive effect ε' with **coerce** $_{\varepsilon \leq \varepsilon'}$.

The eliminators of the empty and sum types are value terms. We also have syntactic sugar for eliminators of value terms on the computation level:

$$\underline{\text{case}}_{\underline{C}} V \text{ of } \{\} := \text{force} (\text{case}_{\mathbf{UC}} V \text{ of } \{\})$$

$$\underline{\text{case}} V \text{ of } \{\text{inl } x_1. M_1, \text{inr } x_2. M_2\} := \text{force} (\text{case } V \text{ of } \{\text{inl } x_1. \text{thunk } M_1, \text{inr } x_2. \text{thunk } M_2\})$$

In GCBPV, all side-effects occur at returner types. This is similar to Levy-style call-by-name, where side-effects occur at base types. Lambda abstraction therefore does not thunk effects, so η -expansion is valid for function types. Similarly, products of computation types are *lazy*. The computation $1'\lambda\{1. M_1, 2. M_2\}$ has the same semantics as M_1 (it does not evaluate M_2), and η -expansion is valid (if M is a product of computations, it has the same semantics as $\lambda\{1. 1'M, 2. 2'M\}$).

Ground types G in GCBPV are value types that do not contain thunks.

$$G ::= b \mid \mathbf{unit} \mid G_1 \times G_2 \mid \mathbf{empty} \mid G_1 + G_2$$

The notion of signature for graded call-by-push-value is almost identical to that of the graded monadic metalanguage (Definition 2.5.1).

Definition 2.7.1 A GCBPV signature consists of the following data:

- A preordered monoid $(\mathcal{E}, \leq, \cdot, 1)$ of *effects*.
- A set \mathcal{B} of *base types*.
- A family of pairwise disjoint sets \mathcal{K}_A of *constants of type* A , indexed by value types A .
- A set Σ of *operations*.
- For each operation $\text{op} \in \Sigma$, ground types car_{op} and ar_{op} , respectively called the *coarity* and *arity* of op , and an effect $\text{eff}_{\text{op}} \in \mathcal{E}$. ◀

$$\begin{array}{c}
\boxed{\Gamma \vdash V : A} \\
\frac{}{\Gamma \vdash c : A} \text{ if } c \in \mathcal{K}_A \qquad \frac{}{\Gamma \vdash x : A} \text{ if } (x : A) \in \Gamma \qquad \frac{}{\Gamma \vdash () : \mathbf{unit}} \\
\frac{\Gamma \vdash V_1 : A_1 \quad \Gamma \vdash V_2 : A_2}{\Gamma \vdash (V_1, V_2) : A_1 \times A_2} \qquad \frac{\Gamma \vdash V : A_1 \times A_2}{\Gamma \vdash \mathbf{fst} V : A_1} \qquad \frac{\Gamma \vdash V : A_1 \times A_2}{\Gamma \vdash \mathbf{snd} V : A_2} \qquad \frac{\Gamma \vdash V : \mathbf{empty}}{\Gamma \vdash \mathbf{case}_A V \text{ of } \{\} : A} \\
\frac{\Gamma \vdash V : A_1}{\Gamma \vdash \mathbf{inl}_{A_2} V : A_1 + A_2} \qquad \frac{\Gamma \vdash V : A_2}{\Gamma \vdash \mathbf{inr}_{A_1} V : A_1 + A_2} \\
\frac{\Gamma \vdash V : A_1 + A_2 \quad \Gamma, x_1 : A_1 \vdash W_1 : B \quad \Gamma, x_2 : A_2 \vdash W_2 : B}{\Gamma \vdash \mathbf{case} V \text{ of } \{\mathbf{inl} x_1. W_1, \mathbf{inr} x_2. W_2\} : B} \\
\frac{\Gamma \perp M : \underline{C}}{\Gamma \vdash \mathbf{thunk} M : \mathbf{UC} \underline{C}} \\
\boxed{\Gamma \perp M : \underline{C}} \\
\frac{}{\Gamma \perp \lambda\{\} : \mathbf{unit}} \qquad \frac{\Gamma \perp M_1 : \underline{C}_1 \quad \Gamma \perp M_2 : \underline{C}_2}{\Gamma \perp \lambda\{1. M_1, 2. M_2\} : \underline{C}_1 \times \underline{C}_2} \qquad \frac{\Gamma \perp M : \underline{C}_1 \times \underline{C}_2}{\Gamma \perp 1'M : \underline{C}_1} \qquad \frac{\Gamma \perp M : \underline{C}_1 \times \underline{C}_2}{\Gamma \perp 2'M : \underline{C}_2} \\
\frac{\Gamma, x : A \perp M : \underline{C}}{\Gamma \perp \lambda x : A. M : A \rightarrow \underline{C}} \qquad \frac{\Gamma \vdash V : A \quad \Gamma \perp M : A \rightarrow \underline{C}}{\Gamma \perp V'M : \underline{C}} \\
\frac{\Gamma \vdash V : \mathbf{car}_{\text{op}}}{\Gamma \perp \text{op} V : \langle \mathbf{eff}_{\text{op}} \rangle \mathbf{ar}_{\text{op}}} \qquad \frac{\Gamma \vdash V : A}{\Gamma \perp \langle V \rangle : \langle 1 \rangle A} \qquad \frac{\Gamma \perp M : \langle \varepsilon \rangle A \quad \Gamma, x : A \perp N : \underline{C}}{\Gamma \perp M \mathbf{to} x. N : \langle \varepsilon \rangle \underline{C}} \\
\frac{\Gamma \perp M : \langle \varepsilon \rangle A}{\Gamma \perp \mathbf{coerce}_{\varepsilon \leq \varepsilon'} M : \langle \varepsilon' \rangle A} \text{ if } \varepsilon \leq \varepsilon' \qquad \frac{\Gamma \vdash V : \mathbf{UC} \underline{C}}{\Gamma \perp \mathbf{force} V : \underline{C}}
\end{array}$$

Figure 2.11: Graded call-by-push-value typing rules

Typing contexts Γ are ordered lists mapping variable names to value types A ; as usual, we assume no variable appears more than once. Variables do not have computation types (but they may be thought of as computations). There are two typing judgments: $\Gamma \vdash V : A$ means the value term V has value type A in context Γ , and similarly $\Gamma \vdash M : \underline{C}$ means the computation term M has computation type \underline{C} in context Γ . To define them, we use a function $\langle\langle \varepsilon \rangle\rangle$ that adds the effect ε to computation types, similar to the function on types used for Levy-style call-by-name. The computation type $\langle\langle \varepsilon \rangle\rangle \underline{C}$ is defined by recursion on the structure of \underline{C} .

$$\begin{aligned} \langle\langle \varepsilon \rangle\rangle \mathbf{unit} &:= \mathbf{unit} & \langle\langle \varepsilon \rangle\rangle (\underline{C}_1 \times \underline{C}_2) &:= \langle\langle \varepsilon \rangle\rangle \underline{C}_1 \times \langle\langle \varepsilon \rangle\rangle \underline{C}_2 \\ \langle\langle \varepsilon \rangle\rangle (A \rightarrow \underline{C}) &:= A \rightarrow \langle\langle \varepsilon \rangle\rangle \underline{C} & \langle\langle \varepsilon \rangle\rangle (\langle \varepsilon' \rangle A) &:= \langle \varepsilon \cdot \varepsilon' \rangle A \end{aligned}$$

The base case of the recursion is returner types, where we collect the two effects ε and ε' together by multiplying them. We do not look at the structure of A in this case. The definition of the typing judgments is given inductively by the rules in Figure 2.11. As for our other type systems, rules that extend typing contexts implicitly assume the extra variable is fresh. We use $\langle\langle \varepsilon \rangle\rangle$ in the typing rule for **to** because it evaluates more than one computation.

We also have a substitution lemma for GCBPV. Since we need it later, we first define a notion of well-typed substitution. *Substitutions* are given by the following grammar:

$$\sigma ::= \diamond \mid \sigma, x \mapsto V$$

where \diamond is the empty substitution. Variables are associated with value types (not computation types) in typing contexts, and are therefore mapped to value terms by substitutions. We have a typing judgment $\Gamma \vdash \sigma : \Delta$ for substitutions, meaning in the context Γ the terms in σ have the types given in the context Δ . This is defined as follows:

$$\frac{}{\Gamma \vdash \diamond : \diamond} \qquad \frac{\Gamma \vdash \sigma : \Delta \quad \Gamma \vdash V : A}{\Gamma \vdash (\sigma, x \mapsto V) : (\Delta, x : A)}$$

We write $V[\sigma]$ and $M[\sigma]$ for the applications of the substitution σ to the value term V and to the computation term M . These are defined by induction on the structure of the terms, and are capture-avoiding. The key property of the substitution typing judgment is the substitution lemma:

Lemma 2.7.2 (Substitution) Suppose that $\Gamma \vdash \sigma : \Delta$.

1. (Values) If $\Delta \vdash V : B$ then $\Gamma \vdash V[\sigma] : B$.
2. (Computations) If $\Delta \vdash M : \underline{C}$ then $\Gamma \vdash M[\sigma] : \underline{C}$.

Proof sketch. There is a weakening lemma on value terms, which states that if $\Gamma, \Gamma' \vdash W : A$ then $\Gamma, x : A', \Gamma' \vdash W : A$. This is proved by induction on the typing derivation mutually with a similar weakening lemma for computation terms. The weakening lemma for value terms can then be extended to substitutions. Finally, the substitution lemma is proved by induction on the typing derivation for V (or M), using weakening. \square

GCBPV allows subeffecting on returner types using **coerce**. We could have instead included a general form of *subtyping*, by including terms of the form **coerce** $_{\underline{C} <: \underline{D}}$ M in the syntax, but this turns out to be unnecessary. Subtyping $\underline{C} <: \underline{D}$ for computation types and $A <: B$ for value types is defined inductively by the rules on the left of Figure 2.12. Both subtyping relations are provably reflexive and transitive. Subtyping on terms is then syntactic sugar: value terms **coerce** $_{A <: B}$ V and computation terms **coerce** $_{\underline{C} <: \underline{D}}$ M are given by induction on the derivation

Subtyping rule	Syntactic sugar
$\boxed{A <: B}$	$\boxed{\Gamma \vdash \mathbf{coerce}_{A <: B} V : B}$
$\frac{}{b <: b}$	$\mathbf{coerce}_{b <: b} V := V$
$\frac{}{\mathbf{unit} <: \mathbf{unit}}$	$\mathbf{coerce}_{\mathbf{unit} <: \mathbf{unit}} V := V$
$\frac{A_1 <: B_1 \quad A_2 <: B_2}{A_1 \times A_2 <: B_1 \times B_2}$	$\mathbf{coerce}_{A_1 \times A_2 <: B_1 \times B_2} V :=$ $(\mathbf{coerce}_{A_1 <: B_1} (\mathbf{fst} V), \mathbf{coerce}_{A_2 <: B_2} (\mathbf{snd} V))$
$\frac{}{\mathbf{empty} <: \mathbf{empty}}$	$\mathbf{coerce}_{\mathbf{empty} <: \mathbf{empty}} V := V$
$\frac{A_1 <: B_1 \quad A_2 <: B_2}{A_1 + A_2 <: B_1 + B_2}$	$\mathbf{coerce}_{A_1 + A_2 <: B_1 + B_2} V :=$ $\mathbf{case} V \mathbf{of}$ $\quad \{ \mathbf{inl} x_1. \mathbf{inl}_{B_2} (\mathbf{coerce}_{A_1 <: B_1} x_1)$ $\quad , \mathbf{inr} x_2. \mathbf{inr}_{B_1} (\mathbf{coerce}_{A_2 <: B_2} x_2) \}$
$\frac{\underline{C} <: \underline{D}}{\underline{U} \underline{C} <: \underline{U} \underline{D}}$	$\mathbf{coerce}_{\underline{U} \underline{C} <: \underline{U} \underline{D}} V := \mathbf{thunk} (\mathbf{coerce}_{\underline{C} <: \underline{D}} (\mathbf{force} V))$
$\boxed{\underline{C} <: \underline{D}}$	$\boxed{\Gamma \vdash \mathbf{coerce}_{\underline{C} <: \underline{D}} M : \underline{D}}$
$\frac{}{\mathbf{unit} <: \mathbf{unit}}$	$\mathbf{coerce}_{\mathbf{unit} <: \mathbf{unit}} M := M$
$\frac{\underline{C}_1 <: \underline{D}_1 \quad \underline{C}_2 <: \underline{D}_2}{\underline{C}_1 \times \underline{C}_2 <: \underline{D}_1 \times \underline{D}_2}$	$\mathbf{coerce}_{\underline{C}_1 \times \underline{C}_2 <: \underline{D}_1 \times \underline{D}_2} M :=$ $\lambda \{ 1. \mathbf{coerce}_{\underline{C}_1 <: \underline{D}_1} (1' M), 2. \mathbf{coerce}_{\underline{C}_2 <: \underline{D}_2} (2' M) \}$
$\frac{B <: A \quad \underline{C} <: \underline{D}}{(A \rightarrow \underline{C}) <: (B \rightarrow \underline{D})}$	$\mathbf{coerce}_{(A \rightarrow \underline{C}) <: (B \rightarrow \underline{D})} M :=$ $\lambda x : B. \mathbf{coerce}_{\underline{C} <: \underline{D}} ((\mathbf{coerce}_{B <: A} x) ' M)$
$\frac{\varepsilon \leq \varepsilon' \quad A <: B}{\langle \varepsilon \rangle A <: \langle \varepsilon' \rangle B}$	$\mathbf{coerce}_{\langle \varepsilon \rangle A <: \langle \varepsilon' \rangle B} M := \mathbf{coerce}_{\varepsilon \leq \varepsilon'} (M \mathbf{to} x. \langle \mathbf{coerce}_{A <: B} x \rangle)$

Figure 2.12: Subtyping in graded call-by-push-value

of subtyping on the right of Figure 2.12. In general, the definition η -expands the terms (for example, coercions between function types introduce lambdas) and propagates coercions down to returner types. Recall that η -expansion preserves the semantics of GCBPV terms. Coercions have the expected typing: if $A \prec B$ then $\Gamma \vdash V : A$ implies $\Gamma \vdash \mathbf{coerce}_{A \prec B} V : B$, and if $\underline{C} \prec \underline{D}$ then $\Gamma \vdash M : \underline{C}$ implies $\Gamma \vdash \mathbf{coerce}_{\underline{C} \prec \underline{D}} M : \underline{D}$.

A useful property is that the functions $\langle\langle \varepsilon \rangle\rangle$ on computation types form an *action* of the preordered monoid of effects on the preorder of computation types (ordered by subtyping).

Definition 2.7.3 An *action* of a preordered monoid $(\mathcal{E}, \leq, \cdot, 1)$ on a preordered set (X, \sqsubseteq) is a function $*$: $\mathcal{E} \times X \rightarrow X$ that is monotone in both arguments and satisfies

$$1 * x = x \quad \varepsilon * (\varepsilon' * x) = (\varepsilon \cdot \varepsilon') * x$$

for all $\varepsilon, \varepsilon' \in \mathcal{E}$ and $x \in X$. ◀

2.7.2 (In)equational theories

As for GMM (Section 2.5) we give an equational theory for determining when it is correct to replace one program with another. However, there are two key differences in this section. The first is that we define a general notion of *theory* for GCBPV, consisting of some core axioms possibly augmented with additional axioms, rather than constraining ourselves to one particular equational theory. The reason for this is that we add extra axioms to capture the behaviour of side-effects included in the signature. The second difference is that we do not constrain ourselves to *symmetric* theories (hence we prefer the term *inequational* theory rather than equational theory). We give examples of non-symmetric theories in Chapter 4.

For GCBPV an inequational theory consists of two judgment forms, one for values and one for computations:

$$\Gamma \vdash V \leq W : A \quad \Gamma \vdash M \leq N : \underline{C}$$

These mean that the terms are well-typed, and every behaviour of V (or M) is a possible behaviour of W (or N). (For this chapter, we can use the same intuition as for \equiv in GMM. The lack of symmetry only becomes important in Chapter 4.) As before, we omit the typing context and type when they are obvious or unimportant.

Each inequational theory must satisfy a core set of symmetric axioms, which are listed in Figure 2.13. We use the symbol \equiv to indicate that they should be read symmetrically (we define \equiv to be the intersection of \leq and its opposite \geq). They are just the usual axioms for call-by-push-value (with complex values), plus some additional axioms for coercions. Each axiom is subject to suitable restrictions on free variables and typing (for example, $M \equiv \lambda\{ \}$ can hold only if M has type unit).

The first group consists of β laws. The second consists of η laws; there is an η law for every GCBPV type (excluding base types). In the third group, the first two equations state that sequencing of computations commutes with formation of binary products and functions. A consequence of these two axioms and the η laws is that it is not necessary to allow arbitrary computation types on the right-hand side of **to**; returner types suffice. This is the case because:

$$\begin{aligned} M \mathbf{to} x. N &\equiv \lambda\{ \} && \text{if } N : \mathbf{unit} \\ M \mathbf{to} x. N &\equiv \lambda\{1. M \mathbf{to} x. 1'N, 2. M \mathbf{to} x. 2'N\} && \text{if } N : \underline{C}_1 \times \underline{C}_2 \\ M \mathbf{to} x. N &\equiv \lambda y:A. M \mathbf{to} x. y'N && \text{if } N : A \rightarrow \underline{C} \end{aligned}$$

where y is not free in N . (These are instances of the equational theory defined below, but not axioms.) The third axiom of the third group is associativity of **to**. The fourth is commutativity

$$\begin{array}{l}
\text{fst } (V_1, V_2) \equiv V_1 \\
\text{case } \text{inl}_{A_2} V \text{ of} \\
\quad \{\text{inl } x_1. W_1 \equiv W_1[x_1 \mapsto V] \\
\quad , \text{inr } x_2. W_2\} \\
1' \lambda\{1. M_1, 2. M_2\} \equiv M_1 \\
V' \lambda x:A. M \equiv M[x \mapsto V] \\
\text{force } (\text{thunk } M) \equiv M \\
\hline
V \equiv () \\
V \equiv \text{case}_A W \text{ of } \{\} \\
V \equiv \text{thunk } (\text{force } V) \\
M \equiv \lambda\{\} \\
M \equiv \lambda x:A. x' M \\
\hline
\text{snd } (V_1, V_2) \equiv V_2 \\
\text{case } \text{inr}_{A_1} V \text{ of} \\
\quad \{\text{inl } x_1. W_1 \equiv W_2[x_2 \mapsto V] \\
\quad , \text{inr } x_2. W_2\} \\
2' \lambda\{1. M_1, 2. M_2\} \equiv M_2 \\
\langle V \rangle \text{ to } x. M \equiv M[x \mapsto V] \\
\hline
V \equiv (\text{fst } V, \text{snd } V) \\
\text{case } W \text{ of} \\
V[x \mapsto W] \equiv \{\text{inl } y_1. V[x \mapsto \text{inl}_{A_2} y_1] \\
\quad , \text{inr } y_2. V[x \mapsto \text{inr}_{A_1} y_2]\} \\
M \equiv (1' M, 2' M) \\
M \equiv M \text{ to } x. \langle x \rangle \\
\hline
\lambda\{1. M \text{ to } x. N_1, 2. M \text{ to } x. N_2\} \equiv M \text{ to } x. \lambda\{1. N_1, 2. N_2\} \\
\lambda y:A. M \text{ to } x. N \equiv M \text{ to } x. \lambda y:A. N \\
(M_1 \text{ to } x. M_2) \text{ to } y. M_3 \equiv M_1 \text{ to } x. M_2 \text{ to } y. M_3 \\
\text{coerce}_{\varepsilon_1 \cdot \varepsilon_2 \leq \varepsilon'_1 \cdot \varepsilon'_2} (M \text{ to } x. N) \equiv (\text{coerce}_{\varepsilon_1 \leq \varepsilon'_1} M) \text{ to } x. \text{coerce}_{\varepsilon_2 \leq \varepsilon'_2} N \\
\text{coerce}_{\varepsilon \leq \varepsilon'} M \equiv M \\
\text{coerce}_{\varepsilon' \leq \varepsilon''} (\text{coerce}_{\varepsilon \leq \varepsilon'} M) \equiv \text{coerce}_{\varepsilon \leq \varepsilon''} M
\end{array}$$

Figure 2.13: Core axioms of GCBPV inequational theories

of **coerce** with **to**. This axiom can only be applied if N has returner type; this is not an issue because we can write all uses of **to** in terms of those that only use returner types. The equation

$$\text{coerce}_{\langle \varepsilon \rangle \underline{C} <: \langle \varepsilon' \rangle \underline{C}'} (M \text{ to } x. N) \equiv (\text{coerce}_{\varepsilon \leq \varepsilon'} M) \text{ to } x. \text{coerce}_{\underline{C} <: \underline{C}'} N$$

follows. Finally, there are two equations corresponding to reflexivity and transitivity of the order \leq on effects. The equations

$$\text{coerce}_{\underline{C} <: \underline{C}} M \equiv M \quad \text{coerce}_{\underline{C}' <: \underline{C}''} (\text{coerce}_{\underline{C} <: \underline{C}'} M) \equiv \text{coerce}_{\underline{C} <: \underline{C}''} M$$

follow from these.

Each inequational theory is also required to be closed under congruence. Congruence involves two kinds of *term context*: we have $C[\]$ for value terms with a single hole, and $\underline{C}[\]$ for computation terms with a single hole. We write \square for a hole where a value term is expected,

and \square for a hole where a computation term is expected.

$$\begin{aligned}
C[] &::= \square \mid (C[], V_2) \mid (V_1, C[]) \mid \mathbf{fst} C[] \mid \mathbf{snd} C[] \mid \mathbf{case}_A C[] \mathbf{of} \{ \} \\
&\mid \mathbf{inl}_{A_2} C[] \mid \mathbf{inr}_{A_1} C[] \mid \mathbf{case} C[] \mathbf{of} \{ \mathbf{inl} x_1. W_1, \mathbf{inr} x_2. W_2 \} \\
&\mid \mathbf{case} V \mathbf{of} \{ \mathbf{inl} x_1. C[], \mathbf{inr} x_2. W_2 \} \mid \mathbf{case} V \mathbf{of} \{ \mathbf{inl} x_1. W_1, \mathbf{inr} x_2. C[] \} \\
&\mid \mathbf{thunk} C[] \\
\underline{C}[] &::= \square \mid \lambda\{1. M_1, 2. \underline{C}[]\} \mid \lambda\{1. \underline{C}[], 2. M_2\} \mid 1' \underline{C}[] \mid 2' \underline{C}[] \mid \lambda x:A. \underline{C}[] \mid C[]'M \mid V' \underline{C}[] \\
&\mid \mathbf{op} C[] \mid \langle C[] \rangle \mid \underline{C}[] \mathbf{to} x. N \mid M \mathbf{to} x. \underline{C}[] \mid \mathbf{coerce}_{\varepsilon \leq \varepsilon'} \underline{C}[] \mid \mathbf{force} C[]
\end{aligned}$$

If $C[]$ is a term context that expects a value term, we write $C[V]$ for the term formed by replacing \square with the value term V , and similarly for the other kinds of term context. The term context may capture some of the free variables of V .

Both judgments of an inequational theory are also required to be closed under substitution (see also Lemma 2.7.2). To state closure under substitution, we extend the value judgment

$$\Gamma \vdash V \leq W : A$$

of each inequational theory to substitutions componentwise: given a typing context Γ , as well as another typing context and two substitutions

$$\Delta = x_1 : A_1, \dots, x_n : A_n \quad \sigma = x_1 \mapsto V_1, \dots, x_n \mapsto V_n \quad \sigma' = x_1 \mapsto V'_1, \dots, x_n \mapsto V'_n$$

we write $\Gamma \vdash \sigma \leq \sigma' : \Delta$ if $\Gamma \vdash \sigma : \Delta$, $\Gamma \vdash \sigma' : \Delta$, and for all i ,

$$\Gamma \vdash V_i \leq V'_i : A_i$$

We now collect together our requirements on inequational theories into a single definition:

Definition 2.7.4 (Inequational theory) An *inequational theory* consists of a GCBPV signature and two judgments

$$\Gamma \vdash V \leq W : A \quad \Gamma \vdash M \leq N : \underline{C}$$

such that:

- Preorder:
 - Values: if $\Gamma \vdash V : A$ then $\Gamma \vdash V \leq V : A$, and if $\Gamma \vdash V_1 \leq V_2 : A$ and $\Gamma \vdash V_2 \leq V_3 : A$ then $\Gamma \vdash V_1 \leq V_3 : A$.
 - Computations: if $\Gamma \vdash M : \underline{C}$ then $\Gamma \vdash M \leq M : \underline{C}$, and if $\Gamma \vdash M_1 \leq M_2 : \underline{C}$ and $\Gamma \vdash M_2 \leq M_3 : \underline{C}$ then $\Gamma \vdash M_1 \leq M_3 : \underline{C}$.
- Congruence: if $\Gamma \vdash V \leq W : A$ then for term contexts with hole \square

$$\begin{aligned}
\Gamma' \vdash C[V] : B \wedge \Gamma' \vdash C[W] : B &\Rightarrow \Gamma' \vdash C[V] \leq C[W] : B \\
\Gamma' \vdash \underline{C}[V] : \underline{D} \wedge \Gamma' \vdash \underline{C}[W] : \underline{D} &\Rightarrow \Gamma' \vdash \underline{C}[V] \leq \underline{C}[W] : \underline{D}
\end{aligned}$$

If $\Gamma \vdash M \leq N : \underline{C}$ then for term contexts with hole \square

$$\begin{aligned}
\Gamma' \vdash C[M] : B \wedge \Gamma' \vdash C[N] : B &\Rightarrow \Gamma' \vdash C[M] \leq C[N] : B \\
\Gamma' \vdash \underline{C}[M] : \underline{D} \wedge \Gamma' \vdash \underline{C}[N] : \underline{D} &\Rightarrow \Gamma' \vdash \underline{C}[M] \leq \underline{C}[N] : \underline{D}
\end{aligned}$$

- Substitution: If $\Gamma \vdash \sigma \leq \sigma' : \Delta$ then

$$\begin{aligned}
\Delta \vdash V \leq W : B &\Rightarrow \Gamma \vdash V[\sigma] \leq W[\sigma'] : B \\
\Delta \vdash M \leq N : \underline{C} &\Rightarrow \Gamma \vdash M[\sigma] \leq N[\sigma'] : \underline{C}
\end{aligned}$$

$$\begin{aligned}
\text{get } () \text{ to } x. (\text{put } x; \langle x \rangle) &\equiv \text{coerce}_{\{\text{get}\} \leq \{\text{get}, \text{put}\}} (\text{get } ()) \\
\text{put } V; \text{get } () &\equiv \text{put } V; \text{coerce}_{\emptyset \leq \{\text{get}\}} \langle V \rangle \\
\text{put } V_1; \text{put } V_2 &\equiv \text{put } V_2 \\
\text{get } () \text{ to } x. \text{get } () \text{ to } y. \langle (x, y) \rangle &\equiv \text{get } () \text{ to } x. \langle (x, x) \rangle \\
\text{get } (); \langle () \rangle &\equiv \text{coerce}_{\emptyset \leq \{\text{get}\}} \langle () \rangle
\end{aligned}$$

Figure 2.14: Signature axioms for global state

- Core axioms:

- Values: if $\Gamma \vdash V : A$ and $\Gamma \vdash W : A$, and $V \equiv W$ is an instance of an axiom in Figure 2.13, then $\Gamma \vdash V \leq W : A$ and $\Gamma \vdash W \leq V : A$.
- Computations: if $\Gamma \vdash M : \underline{C}$ and $\Gamma \vdash N : \underline{C}$, and $M \equiv N$ is an instance of an axiom in Figure 2.13, then $\Gamma \vdash M \leq N : \underline{C}$ and $\Gamma \vdash N \leq M : \underline{C}$. ◀

For a fixed typing context Γ and type A , each inequational theory is a preorder on terms of type A in context Γ , and similarly for computations. We use the symbol \equiv for the symmetric part of the inequational theory (e.g. on values, $\Gamma \vdash V \equiv W : A$ means both $\Gamma \vdash V \leq W : A$ and $\Gamma \vdash W \leq V : A$ hold).

The smallest possible inequational theory (i.e. the one that relates the fewest terms), is given by taking the core axioms, and closing them under reflexivity, transitivity and congruence. (Closure under substitution holds automatically in this case.) Since all of the core axioms are symmetric, the smallest inequational theory is also symmetric.

In general, to specify an inequational theory, we give a list of (possibly non-symmetric) *signature axioms* (which characterize the side-effects in the signature) to be taken in addition to the core axioms. The \leq judgments are then given by closing under congruence, reflexivity and transitivity. We do not assume that inequational theories are given by a finite set of axioms (or axiom schemas). Nor do we forbid infinitary rules in the definition. This allows us to consider inequational theories for recursion (using an infinitary *rational continuity* rule [19]).

We give one example here, for global state (further examples are given in Chapter 4). For this example, we use the Gifford-style effect algebra with two operations: `get` with coarity **unit** and arity **bool**, and `put` with coarity **unit** and arity **bool**. Effects are subsets of $\{\text{get}, \text{put}\}$. We write $M; N$ as syntactic sugar for $M \text{ to } x. N$, where x is fresh. The signature axioms for this example are listed in Figure 2.14; they characterize the behaviour of `get` and `put`. The signature axioms for global state imply more general equations. For example, the fourth axiom, which allows us to merge two adjacent reads, implies the following general equation:

$$\text{get } () \text{ to } x. \text{get } () \text{ to } y. M \equiv \text{get } () \text{ to } z. M[x \mapsto z, y \mapsto z]$$

As an aside, we mention that these are not the same as the axioms given by Plotkin and Power [88]. If we forget about the effect system (more formally, if we use the trivial effect algebra), then our axioms are equivalent to Plotkin and Power's. However, the proof that our fifth axiom (which is just $\text{get } (); \langle () \rangle \equiv \langle () \rangle$ for the trivial effect algebra) follows from theirs is a sequence of equalities involving `put`. We cannot use that proof with a Gifford-style effect algebra because the effect for this axiom does not contain $\{\text{put}\}$. In general, when adding an effect system it may be necessary to specify additional axioms.

In the rest of this chapter we consider arbitrary inequational theories \leq . Given a value term V we define $\lfloor V \rfloor$ to be V but with all uses of `coerce` deleted, and similarly for computation terms. We have:

Lemma 2.7.5 Suppose that the effect algebra $(\mathcal{E}, \leq, \cdot, 1)$ is a partially ordered monoid with bounded binary joins.

1. Given two value terms V_1, V_2 such that $\Gamma \vdash V_1 : A$ and $\Gamma \vdash V_2 : A$, if $\lfloor V_1 \rfloor = \lfloor V_2 \rfloor$ then $\Gamma \vdash V_1 \equiv V_2 : A$.
2. Given two computation terms M_1, M_2 such that $\Gamma \vdash M_1 : \underline{C}$ and $\Gamma \vdash M_2 : \underline{C}$, if $\lfloor M_1 \rfloor = \lfloor M_2 \rfloor$ then $\Gamma \vdash M_1 \equiv M_2 : \underline{C}$. \blacktriangleleft

We use this fact when considering the translations from source-language effect systems. A proof is given in Appendix B.1. We conjecture that this lemma does not hold when the assumption about the effect algebra is dropped. A conjectured counterexample is given along with the proof.

When we consider equivalences between terms in graded call-by-push-value, we ask for the terms to have the same observable behaviour. We therefore define a *contextual preorder* for GCBPV, by considering closed computations that return elements of ground types.

Definition 2.7.6 (Contextual preorder) Given an inequational theory, the *contextual preorder* consists of two judgment forms.

1. Between value terms: $\Gamma \vdash V \leq_{\text{ctx}} W : A$ if $\Gamma \vdash V : A$, $\Gamma \vdash W : A$, and for all ground types G , effects ε , and term contexts $\underline{C}[\]$ with hole \square such that $\diamond \vdash \underline{C}[V] : \langle \varepsilon \rangle G$ and $\diamond \vdash \underline{C}[W] : \langle \varepsilon \rangle G$ we have

$$\diamond \vdash \underline{C}[V] \leq \underline{C}[W] : \langle \varepsilon \rangle G$$

2. Between computation terms: $\Gamma \vdash M \leq_{\text{ctx}} N : \underline{C}$ if $\Gamma \vdash M : \underline{C}$, $\Gamma \vdash N : \underline{C}$, and for all ground types G , effects ε , and term contexts $\underline{C}[\]$ with hole \square such that $\diamond \vdash \underline{C}[M] : \langle \varepsilon \rangle G$ and $\diamond \vdash \underline{C}[N] : \langle \varepsilon \rangle G$ we have

$$\diamond \vdash \underline{C}[M] \leq \underline{C}[N] : \langle \varepsilon \rangle G \quad \blacktriangleleft$$

As for inequational theories, we often omit the typing context and type. The inequational theory is stronger than contextual equivalence: if $M \leq N$ then $M \leq_{\text{ctx}} N$ (and similarly for value terms), because \leq is a congruence. For closed terms of $\langle \varepsilon \rangle G$, where G is a ground type, the converse also holds: if $\diamond \vdash M \leq_{\text{ctx}} N : \langle \varepsilon \rangle G$ then $\diamond \vdash M \leq N : \langle \varepsilon \rangle G$. This is not generally the case for open terms or for other types. For example, with global state

$$x : \langle \{\text{get}\} \rangle \mathbf{unit} \vdash x \leq \mathbf{coerce}_{\emptyset \leq \{\text{get}\}} \langle () \rangle : \langle \{\text{get}\} \rangle \mathbf{unit}$$

does not hold, but the two computations have the same observable behaviour.

2.7.3 Translations into GCBPV

To reason about evaluation orders using graded call-by-push-value, we give compositional translations from our source-language effect systems, similar to the translations into the graded monadic metalanguage described in Section 2.5.2. In addition to call-by-value and Moggi-style call-by-name we give a translation from Levy-style call-by-name into GCBPV.⁷ We could not do this satisfactorily for GMM. The call-by-value and Levy-style call-by-name translations we give here are based on those of Levy [56], with some minor modifications to deal with effects.

⁷We can also translate most of GMM into GCBPV by decomposing the graded monad $\langle \varepsilon \rangle -$ as $U(\langle \varepsilon \rangle -)$. (Function types cannot be translated.)

$$\begin{array}{c}
\langle \mathbf{unit} \rangle := \mathbf{unit} \\
\langle \mathbf{bool} \rangle := \mathbf{unit} + \mathbf{unit} \\
\langle \tau \xrightarrow{\varepsilon} \tau' \rangle := \mathbf{U} (\langle \tau \rangle \rightarrow \langle \varepsilon \rangle \langle \tau' \rangle) \\
\hline
\langle \Gamma \vdash_v x : \tau \& 1 \rangle := \langle x \rangle \\
\hline
\langle \Gamma \vdash_v () : \mathbf{unit} \& 1 \rangle := \langle () \rangle \\
\hline
\langle \Gamma \vdash_v \mathbf{true} : \mathbf{bool} \& 1 \rangle := \langle \mathbf{inl}_{\mathbf{unit}} () \rangle \quad \langle \Gamma \vdash_v \mathbf{false} : \mathbf{bool} \& 1 \rangle := \langle \mathbf{inr}_{\mathbf{unit}} () \rangle \\
\hline
\langle \Gamma \vdash_v e_1 : \mathbf{bool} \& \varepsilon \rangle = M_1 \quad \langle \Gamma \vdash_v e_2 : \tau \& \varepsilon' \rangle = M_2 \quad \langle \Gamma \vdash_v e_3 : \tau \& \varepsilon' \rangle = M_3 \\
\hline
\langle \Gamma \vdash_v \mathbf{if} e_1 \mathbf{then} e_2 \mathbf{else} e_3 : \tau \& \varepsilon \cdot \varepsilon' \rangle := M_1 \mathbf{to} x. \mathbf{case} x \mathbf{of} \{ \mathbf{inl} _ . M_2, \mathbf{inr} _ . M_3 \} \\
\hline
\langle \Gamma, x : \tau \vdash_v e : \tau' \& \varepsilon \rangle = M \\
\hline
\langle \Gamma \vdash_v \lambda x : \tau. e : \tau \xrightarrow{\varepsilon} \tau' \& 1 \rangle := \langle \mathbf{think} (\lambda x : \langle \tau \rangle. M) \rangle \\
\hline
\langle \Gamma \vdash_v e_1 : \tau \xrightarrow{\varepsilon_3} \tau' \& \varepsilon_1 \rangle = M_1 \quad \langle \Gamma \vdash_v e_2 : \tau \& \varepsilon_2 \rangle = M_2 \\
\hline
\langle \Gamma \vdash_v e_1 e_2 : \tau' \& \varepsilon_1 \cdot \varepsilon_2 \cdot \varepsilon_3 \rangle := M_1 \mathbf{to} f. M_2 \mathbf{to} x. x \mathbf{force} f \\
\hline
\langle \Gamma \vdash_v e : \tau \& \varepsilon \rangle = M \\
\hline
\langle \Gamma \vdash_v e : \tau \& \varepsilon' \rangle := \mathbf{coerce}_{\varepsilon \leq \varepsilon'} M
\end{array}$$

Figure 2.15: Translation of call-by-value types (top left), contexts (top right), and typing derivations (bottom) into GCBPV.

The translations we give in this section assume a fixed source-language signature (Definition 2.2.1) that is used to instantiate the source-language effect systems, and a fixed GCBPV signature (Definition 2.7.1). The two signatures are required to be compatible. We translate source-language ground types $\tau \in \{\mathbf{unit}, \mathbf{bool}\}$ into GCBPV as follows:

$$\langle \mathbf{unit} \rangle := \mathbf{unit} \quad \langle \mathbf{bool} \rangle := \mathbf{unit} + \mathbf{unit}$$

For compatibility we assume that the two signatures have the same effect algebra, and that for each source-language operation op : the GCBPV signature also includes op ; the two signatures assign the same effect eff_{op} to op ; and if τ and τ' are respectively the coarity and arity of op in the source-language signature, then $\langle \tau \rangle$ and $\langle \tau' \rangle$ are the coarity and arity of op in the GCBPV signature. The GCBPV signature is allowed to contain operations that are not in the source-language signature, and we place no requirements on base types or constants.

Call-by-value types τ are translated into GCBPV value types $\langle \tau \rangle^v$, which represent values of type τ (with no side-effects), and call-by-value contexts Γ are translated into GCBPV contexts $\langle \Gamma \rangle^v$. Call-by-value *derivations* $\Gamma \vdash_v e : \tau \& \varepsilon$ are translated into GCBPV *computations* $\langle \Gamma \rangle^v \vdash e : \tau \& \varepsilon^v$ that return elements of $\langle \tau \rangle^v$, so $\langle \Gamma \rangle^v \vdash \langle e \rangle^v : \langle \varepsilon \rangle \langle \tau \rangle^v$, where we omit the context and type in $\langle e \rangle^v$. (Recall that we have to translate *derivations* rather than just well-typed

$$\begin{array}{c}
\langle \mathbf{unit} \rangle := \mathbf{unit} \\
\langle \mathbf{bool} \rangle := \mathbf{unit} + \mathbf{unit} \\
\langle \tau \xrightarrow{\varepsilon, \varepsilon'} \tau' \rangle := \mathbf{U} (\mathbf{U} (\langle \varepsilon \rangle \langle \tau \rangle) \rightarrow \langle \varepsilon' \rangle \langle \tau' \rangle) \\
\langle \diamond \rangle := \diamond \\
\langle \Gamma, x : \tau \& \varepsilon \rangle := \langle \Gamma \rangle, x : \mathbf{U} (\langle \varepsilon \rangle \langle \tau \rangle) \\
\frac{}{\langle \Gamma \vdash_{\text{moggi}} x : \tau \& \varepsilon \rangle := \mathbf{force} x} \qquad \frac{\langle \Gamma \vdash_{\text{moggi}} e : \mathbf{car}_{\text{op}} \& \varepsilon \rangle = M}{\langle \Gamma \vdash_{\text{moggi}} \mathbf{op} e : \mathbf{ar}_{\text{op}} \& \varepsilon \cdot \mathbf{eff}_{\text{op}} \rangle = M \mathbf{to} x. \mathbf{op} x} \\
\frac{}{\langle \Gamma \vdash_{\text{moggi}} () : \mathbf{unit} \& 1 \rangle := \langle () \rangle} \\
\frac{}{\langle \Gamma \vdash_{\text{moggi}} \mathbf{true} : \mathbf{bool} \& 1 \rangle := \langle \mathbf{inl}_{\mathbf{unit}} () \rangle} \qquad \frac{}{\langle \Gamma \vdash_{\text{moggi}} \mathbf{false} : \mathbf{bool} \& 1 \rangle := \langle \mathbf{inr}_{\mathbf{unit}} () \rangle} \\
\frac{\langle \Gamma \vdash_{\text{moggi}} e_1 : \mathbf{bool} \& \varepsilon \rangle = M_1 \quad \langle \Gamma \vdash_{\text{moggi}} e_2 : \tau \& \varepsilon' \rangle = M_2 \quad \langle \Gamma \vdash_{\text{moggi}} e_3 : \tau \& \varepsilon' \rangle = M_3}{\langle \Gamma \vdash_{\text{moggi}} \mathbf{if} e_1 \mathbf{then} e_2 \mathbf{else} e_3 : \tau \& \varepsilon \cdot \varepsilon' \rangle := M_1 \mathbf{to} x. \mathbf{case} x \mathbf{of} \{ \mathbf{inl} _ . M_2, \mathbf{inr} _ . M_3 \}} \\
\frac{\langle \Gamma, x : \tau \& \varepsilon \vdash_{\text{moggi}} e : \tau' \& \varepsilon' \rangle = M}{\langle \Gamma \vdash_{\text{moggi}} \lambda x : \tau \& \varepsilon. e : \tau \xrightarrow{\varepsilon, \varepsilon'} \tau' \& 1 \rangle := \langle \mathbf{thunk} (\lambda x : \mathbf{U} (\langle \varepsilon \rangle \langle \tau \rangle). M) \rangle} \\
\frac{\langle \Gamma \vdash_{\text{moggi}} e_1 : \tau \xrightarrow{\varepsilon_2, \varepsilon_3} \tau' \& \varepsilon_1 \rangle = M_1 \quad \langle \Gamma \vdash_{\text{moggi}} e_2 : \tau \& \varepsilon_2 \rangle = M_2}{\langle \Gamma \vdash_{\text{moggi}} e_1 e_2 : \tau' \& \varepsilon_1 \cdot \varepsilon_3 \rangle := M_1 \mathbf{to} f. (\mathbf{thunk} M_2) \mathbf{' (force} f)} \\
\frac{\langle \Gamma \vdash_{\text{moggi}} e : \tau \& \varepsilon \rangle = M}{\langle \Gamma \vdash_{\text{moggi}} e : \tau \& \varepsilon' \rangle := \mathbf{coerce}_{\varepsilon \leq \varepsilon'} M}
\end{array}$$

Figure 2.16: Translation of Moggi-style call-by-name types (top left), contexts (top right), and typing derivations (bottom) into GCBPV.

expressions because the call-by-value effect system is not syntax directed.) The side-effects of e are encapsulated in the returner type.

The definition of the call-by-value translation is in Figure 2.15, where we omit the superscript on $\langle - \rangle^v$. For function types we use computations of returner type with effect ε (for the latent effect of the function), and the side-effects of the function are thunked by wrapping it in \mathbf{U} . This means that call-by-value expressions of function type are translated into computation terms that return thunks of GCBPV functions. For operations \mathbf{op} , the expression e is first eagerly evaluated, and then \mathbf{op} is applied to the result. In the translation of \mathbf{if} -expressions we use the eliminator of GCBPV sum types on the computation level (recall that this is defined in terms of the eliminator on the value level). In the translation of lambda abstraction and function application we explicitly specify where side-effects are thunked, and where thunks are forced.

Most of the translation of Moggi-style call-by-name is similar to the call-by-value translation. Types τ are translated into GCBPV value types $\langle \tau \rangle^{\text{moggi}}$, which represent values of type τ with no side-effects. Contexts Γ become GCBPV contexts $\langle \Gamma \rangle^{\text{moggi}}$. Derivations of $\Gamma \vdash_{\text{moggi}} e : \tau \& \varepsilon$ become computation terms $\langle \Gamma \vdash_{\text{moggi}} e : \tau \& \varepsilon \rangle^{\text{moggi}}$ of type $\langle \varepsilon \rangle \langle \tau \rangle^{\text{moggi}}$.

The definition is in Figure 2.16. The primary difference is in the translation of function types,

$$\begin{array}{c}
\langle \varepsilon \rangle \mathbf{unit} := \langle \varepsilon \rangle \mathbf{unit} \\
\langle \varepsilon \rangle \mathbf{bool} := \langle \varepsilon \rangle (\mathbf{unit} + \mathbf{unit}) \\
\langle \tau \rightarrow \tau' \rangle := \mathbf{U}(\langle \tau \rangle \rightarrow \langle \tau' \rangle) \\
\hline
\langle \Gamma \vdash_n x : \tau \rangle := \mathbf{force} \ x \\
\hline
\langle \Gamma \vdash_n () : \langle 1 \rangle \mathbf{unit} \rangle := \langle () \rangle \\
\hline
\langle \Gamma \vdash_n \mathbf{true} : \langle 1 \rangle \mathbf{bool} \rangle := \langle \mathbf{inl}_{\mathbf{unit}} () \rangle \qquad \langle \Gamma \vdash_n \mathbf{false} : \langle 1 \rangle \mathbf{bool} \rangle := \langle \mathbf{inr}_{\mathbf{unit}} () \rangle \\
\hline
\langle \Gamma \vdash_n e_1 : \langle \varepsilon \rangle \mathbf{bool} \rangle = M_1 \quad \langle \Gamma \vdash_n e_2 : \tau \rangle = M_2 \quad \langle \Gamma \vdash_n e_3 : \tau \rangle = M_3 \\
\langle \Gamma \vdash_n \mathbf{if} \ e_1 \ \mathbf{then} \ e_2 \ \mathbf{else} \ e_3 : \langle \varepsilon \rangle \tau \rangle := M_1 \ \mathbf{to} \ x. \ \mathbf{case} \ x \ \mathbf{of} \ \{ \mathbf{inl} _ . M_2, \ \mathbf{inr} _ . M_3 \} \\
\hline
\langle \Gamma, x : \tau \vdash_n e : \tau' \rangle = M \qquad \langle \Gamma \vdash_n e_1 : \tau \rightarrow \tau' \rangle = M_1 \quad \langle \Gamma \vdash_n e_2 : \tau \rangle = M_2 \\
\langle \Gamma \vdash_n \lambda x : \tau. e : \tau \rightarrow \tau' \rangle := \lambda x : \mathbf{U}(\langle \tau \rangle). M \qquad \langle \Gamma \vdash_n e_1 \ e_2 : \tau' \rangle := (\mathbf{thunk} \ M_2) \ ' M_1 \\
\hline
\langle \Gamma \vdash_n e : \langle \varepsilon \rangle \tau \rangle = M \\
\langle \Gamma \vdash_n e : \langle \varepsilon' \rangle \tau \rangle := \mathbf{coerce}_{\langle \varepsilon \rangle \langle \tau \rangle <: \langle \varepsilon' \rangle \langle \tau \rangle} M
\end{array}$$

Figure 2.17: Translation of Levy-style call-by-name types (top left), contexts (top right), and typing derivations (bottom) into GCBPV.

where arguments are thunked computations rather than just values. Similarly, typing contexts contain thunks. Variables x force the corresponding thunk; hence we reevaluate a computation each time x is used. Operations eagerly evaluate their argument. For function application, the argument is not evaluated immediately; instead a thunk is passed to the function.

Both the call-by-value and Moggi-style call-by-name translations into GCBPV are somewhat similar to the corresponding translations into GMM. In this section, we add a translation for Levy-style call-by-name. Unlike the previous two, Levy-style call-by-name types τ are translated into GCBPV *computation* types $(\tau)^n$. The elements of $(\tau)^n$ are therefore computations, and might have side-effects. A derivation $\Gamma \vdash_n e : \tau$ is translated into a computation term $(\Gamma \vdash_n e : \tau)^n$ of type $(\tau)^n$. We do not introduce an additional returner type (in contrast to the previous translations), because the side-effects occur at returner types that appear in $(\tau)^n$. Typing contexts Γ are translated into GCBPV typing contexts $(\Gamma)^n$ as before.

The definition of the translation from Levy-style call-by-name into GCBPV is given in Figure 2.17. Returner types are attached to **unit** and **bool** because these are the types at which side-effects occur for Levy-style call-by-name. As for Moggi-style call-by-name, functions take thunks of computations as arguments. Otherwise, they just use the GCBPV function types directly. The translation of types satisfies $(\langle \varepsilon \rangle \tau)^n = \langle \varepsilon \rangle (\tau)^n$. Contexts are translated using thunks. Most of the cases in the translation of typing derivations are similar to the Moggi-style call-by-name translation. The case for function application is simpler because the side-effects associated with computing the function are not thunked. Like the other two translations in this section, the translation from Levy-style call-by-name into GCBPV is compositional.

Each of the three translations in this section satisfies some useful properties. First, the terms that arise from translating typing derivations into graded call-by-push-value have the desired types.

Lemma 2.7.7

1. If M is the translation of a derivation of $\Gamma \vdash_v e : \tau \ \& \ \varepsilon$ in call-by-value then

$$\langle \Gamma \rangle^v \vdash M : \langle \varepsilon \rangle \langle \tau \rangle^v$$

2. If M is the translation of a derivation of $\Gamma \vdash_{\text{moggi}} e : \tau \ \& \ \varepsilon$ in Moggi-style call-by-name then

$$\langle \Gamma \rangle^{\text{moggi}} \vdash M : \langle \varepsilon \rangle \langle \tau \rangle^{\text{moggi}}$$

3. If M is the translation of a derivation of $\Gamma \vdash_n e : \tau$ in Levy-style call-by-name then

$$\langle \Gamma \rangle^n \vdash M : \langle \tau \rangle^n \quad \blacktriangleleft$$

Next, we note that any two typing derivations with the same conclusion in our source-language effect systems can differ only in where subeffecting is used. Hence by Lemma 2.7.5, for typical effect algebras, translations of derivations with the same conclusion are the same up to the GCBPV inequational theory.

Lemma 2.7.8 Suppose that the effect algebra $(\mathcal{E}, \leq, \cdot, 1)$ is a partially ordered monoid with bounded binary joins.

1. If M and M' are call-by-value translations of derivations of $\Gamma \vdash_v e : \tau \ \& \ \varepsilon$ then $M \equiv M'$.
2. If M and M' are Moggi-style call-by-name translations of derivations of $\Gamma \vdash_{\text{moggi}} e : \tau \ \& \ \varepsilon$ then $M \equiv M'$.
3. If M and M' are Levy-style call-by-name translations of derivations of $\Gamma \vdash_n e : \tau$ then $M \equiv M'$. \(\blacktriangleleft\)

In particular, when considering GCBPV terms up to the (symmetric part of) the inequational theory, we can consider translations of well-typed expressions; the choice of derivation is irrelevant. (For example, for each Γ , τ and ε , the call-by-value translation is a function from expressions e such that $\Gamma \vdash_v e : \tau \ \& \ \varepsilon$ to equivalence classes of computation terms up to \equiv .) We use this in the statement of soundness for the three translations. Recall that \diamond is the empty typing context. In the statement of soundness we omit the typing context, type and effect when writing translations of well-typed expressions.

Lemma 2.7.9 Suppose that the effect algebra $(\mathcal{E}, \leq, \cdot, 1)$ is a partially ordered monoid with bounded binary joins.

1. If $\diamond \vdash_v e : \tau \ \& \ \varepsilon$ and $e \overset{v}{\rightsquigarrow} e'$ then $\langle e \rangle^v \equiv \langle e' \rangle^v$.
2. If $\diamond \vdash_{\text{moggi}} e : \tau \ \& \ \varepsilon$ and $e \overset{\text{moggi}}{\rightsquigarrow} e'$ then $\langle e \rangle^{\text{moggi}} \equiv \langle e' \rangle^{\text{moggi}}$.
3. If $\diamond \vdash_n e : \tau$ and $e \overset{n}{\rightsquigarrow} e'$ then $\langle e \rangle^n \equiv \langle e' \rangle^n$. \(\blacktriangleleft\)

2.8 Related work

Previous work on effect systems has focused on call-by-value languages. Historically these used Gifford-style effect algebras [63], which are well-known. The more general form of effect algebra we use here (preordered monoids) was first described by Katsumata [45]. Effect systems can be extended, for example by adding operations to the effect algebras to support more language constructs [80], or by adding support for regions [81]. The Moggi-style call-by-name effect presented here was first described by McDermott and Mycroft [69], but is somewhat similar to a system described by Abadi et al. [1]. The Levy-style effect system is new here. GMM is similar to languages described by Katsumata [45] and Gaboardi et al. [29].

A version of the call-by-push-value effect system we present here was first described by McDermott and Mycroft [70]. Kammar and Plotkin [42] describe a Gifford-style effect system for call-by-push-value, called MAIL (multi-adjunctive intermediate language). Unlike GCBPV, in MAIL thunk types, returner types, and the typing judgment are all annotated with effects. MAIL does not generalize to other effect algebras: if the effect algebra is allowed to be an arbitrary preordered monoid then it is unsound, and some computations that are typable in GCBPV are not typable in MAIL. It also overapproximates effects more than ours, even for Gifford-style. Hence our effect system has several advantages, though in many cases Kammar and Plotkin's suffices.

Our translations into GMM and GCBPV are partly based on translations given by Moggi [78], Wadler [98] and Levy [56, 59]. We could have presented these translations differently. For example, the translation of Moggi-style call-by-name into GMM can be presented by first translating into call-by-value [35, 34]. We choose to translate directly from source languages into GMM and GCBPV for simplicity. We could also have considered languages other than GMM or GCBPV, such as Filinski's [24].

2.9 Summary

One of the themes of this thesis is reasoning about programs using knowledge of *local* restrictions on side-effects. That is, by supposing that parts of programs are restricted to particular side-effects, rather than every part of every program. A standard method for doing this is to use effect systems. We take this approach here. This chapter provides the effect systems we use to do effect-dependent reasoning.

In later chapters, we reason about different evaluation orders for the source language, in the presence of side-effects. The method we use relies on having an intermediate language together with translations, one for each evaluation order of interest, from the source language. This chapter supplies these. The intermediate language we use is graded call-by-push-value (Section 2.7), and we have effect systems for call-by-value (Section 2.3), Moggi-style call-by-name (Section 2.4) and Levy-style call-by-name (Section 2.6), each with a translation (Section 2.7.3) into GCBPV. We are therefore in a position to consider equivalences between evaluation orders in the next chapter.

Chapter 3

Call-by-value and call-by-name

One of the goals of this thesis is to relate evaluation orders. Specifically, suppose we have some program M (for example in GCBPV) that uses a mix of several evaluation orders and we construct a new program M' by replacing one evaluation order with another for some subterm. We want to reason formally about the relationship between M and M' . These relationships generally depend on restrictions on side-effects. For example, for replacing call-by-value with call-by-name:

- If there are no side-effects at all (in particular, all programs are strongly normalizing), the choice of call-by-value or call-by-name does not affect the semantics of the program: M and M' terminate with the same result.
- If there are diverging terms (for instance, via recursion), then the behaviour may change: a program might diverge under call-by-value and return a result under call-by-name. However, we can say something about how the behaviour changes: if M terminates with some result, then M' terminates with the same result.
- If nondeterminism is the only side-effect, every possible result of M is a possible result of M' .

These properties are intuitively obvious, and can be proved for specific examples by reasoning at the meta-level.

In this chapter, we show how to relate call-by-value (Section 2.3) and Levy-style call-by-name (Section 2.6). We obtain a reasoning principle (Theorem 3.4.2) for relating these two evaluation orders. This reasoning principle can be applied to various effects. As an example, we show that, in GCBPV instantiated with global state, call-by-value and Levy-style call-by-name have the same observable behaviour for subterms restricted to no side-effects. (Of course, is not the case if we drop the restriction.)

Rather than reasoning at the meta-level, we use a common intermediate language (here GCBPV) that captures both evaluation orders via translations from source-language effect systems. We relate the observable behaviours of the call-by-value translation $\llbracket e \rrbracket^v$ and the Levy-style call-by-name translation $\llbracket e \rrbracket^n$ of each source-language expression e (the two translations are the subterms we refer to above). Intuitively, we want to show $\llbracket e \rrbracket^v \leq_{\text{ctx}} \llbracket e \rrbracket^n$ (recall that \leq_{ctx} is the contextual preorder defined in Definition 2.7.6). However, there are two subtleties we need to deal with.

The first is that the two translations have different types, and hence we cannot relate them directly using the contextual preorder. The way around this is inspired by work relating direct and continuation interpretations of languages [89]: we identify maps between the call-by-value and call-by-name interpretations¹, and then compose these with the translations of expressions to arrive at two terms that can be compared directly. We show that, under certain conditions

on the side-effects, the maps between call-by-value and call-by-name form a *Galois connection*, and this fact allows us to derive our reasoning principle.

The second subtlety arises when we consider effect systems, which we do so that we only need to restrict subterms to particular side-effects, rather than the entire source language. We have various effect systems, each for a different evaluation order, and these assign different effects to each source language expression e . We deal with this by only considering effect algebras in which, if e is typable without tracking effects (i.e. our first source-language type system, defined in Section 2.2), then it is typable under both the call-by-value and call-by-name effect systems. In this chapter, we consider only Gifford-style effect algebras, for which this holds.

The technique we propose is intended to be *general*, in that it should work for other pairs of evaluation orders, such as call-by-need and call-by-name (even though we only concentrate on one pair here). We also work abstractly and identify properties of side-effects that enable us to relate call-by-value and call-by-name, rather than just considering some fixed collection of side-effects. An advantage of the technique of using Galois connections is that the properties required are derived from the structure of the two maps between evaluation orders.

To derive our reasoning principle, we need a method of proving instances of the GCBPV contextual preorder. We use *logical relations* to do this. We therefore discuss logical relations for graded call-by-push-value in this chapter before considering call-by-value and call-by-name.

This chapter has two main contributions:

- We describe a notion of logical relation for graded call-by-push-value (Section 3.1), based on previous work on logical relations for other languages. We also provide a method of constructing such a logical relation, called the *free lifting* (Section 3.1.1).
- We derive a reasoning principle for relating call-by-value and Levy-style call-by-name (Section 3.4). This is formulated in terms of a Galois connection between call-by-value and call-by-name computations (Section 3.3).

In this chapter, we work exclusively with syntax. In particular, we do not use denotational semantics. Chapter 4 shows (amongst other things) how to derive a similar reasoning principle using denotational semantics, and argues that the denotational approach has several advantages over the approach we use in this chapter. However, Chapter 4 requires significantly more machinery than this one.

3.1 Logical relations for graded call-by-push-value

To prove instances of GCBPV contextual preorders, we use *logical relations*. We consider binary logical relations on GCBPV terms in general, before considering evaluation orders in later sections of this chapter.

Recall that an inequational theory for GCBPV consists of a signature (Definition 2.7.1) and two judgments

$$\Gamma \vdash V \leqslant W : A \quad \Gamma \Vdash M \leqslant N : \underline{C}$$

satisfying certain properties (see Definition 2.7.4). Each inequational theory induces a contextual preorder \leqslant_{ctx} (Definition 2.7.6). As usual, we write

$$\Gamma \vdash V \equiv W : A \quad \Gamma \Vdash M \equiv N : \underline{C}$$

¹Specifically, these map between computations of type $\langle \varepsilon \rangle (\tau_\varepsilon^v)^v$ and computations of type $(\tau_\varepsilon^n)^n$, where τ_ε^v and τ_ε^n are source-language types annotated with the effect ε .

$$\begin{aligned}
\mathcal{R}[\mathbf{unit}] &= \text{Term}_{\mathbf{unit}} \times \text{Term}_{\mathbf{unit}} \\
\mathcal{R}[A_1 \times A_2] &= \{(V, V') \mid (\text{fst } V, \text{fst } V') \in \mathcal{R}[A_1] \wedge (\text{snd } V, \text{snd } V') \in \mathcal{R}[A_2]\} \\
\mathcal{R}[\mathbf{empty}] &= \emptyset \\
\mathcal{R}[A_1 + A_2] &= \{(\text{inl}_{A_2} V, \text{inl}_{A_2} V') \mid (V, V') \in \mathcal{R}[A_1]\} \cup \{(\text{inr}_{A_1} V, \text{inr}_{A_1} V') \mid (V, V') \in \mathcal{R}[A_2]\} \\
\mathcal{R}[\mathbf{U} \underline{C}] &= \{(V, V') \mid (\text{force } V, \text{force } V') \in \mathcal{R}[\underline{C}]\} \\
\mathcal{R}[\underline{\mathbf{unit}}] &= \underline{\text{Term}}_{\underline{\mathbf{unit}}} \times \underline{\text{Term}}_{\underline{\mathbf{unit}}} \\
\mathcal{R}[\underline{C}_1 \times \underline{C}_2] &= \{(M, M') \mid (1'M, 1'M') \in \mathcal{R}[\underline{C}_1] \wedge (2'M, 2'M') \in \mathcal{R}[\underline{C}_2]\} \\
\mathcal{R}[A \rightarrow \underline{C}] &= \{(M, M') \mid \forall (V, V') \in \mathcal{R}[A]. (V'M, V'M') \in \mathcal{R}[\underline{C}]\}
\end{aligned}$$

Figure 3.1: Logical relations for graded call-by-push-value

for the symmetric part of the equational theory (i.e. $V \equiv W$ if $V \leq W$ and $W \leq V$). For each value type A we define Term_A as the set of \equiv -equivalence classes of closed terms of type A , and similarly define $\underline{\text{Term}}_{\underline{C}}$ for computation types:

$$\text{Term}_A := \{[V]_{\equiv} \mid \diamond \vdash V : A\} \quad \underline{\text{Term}}_{\underline{C}} := \{[M]_{\equiv} \mid \diamond \vdash M : \underline{C}\}$$

Throughout this chapter, we consider terms up to \equiv . All of the definitions we use are invariant under \equiv (including for example the contextual preorder \leq_{ctx}). We omit the square brackets when writing equivalence classes.

In GCBPV, a logical relation consists of two families of types

$$A \mapsto \mathcal{R}[A] \subseteq \text{Term}_A \times \text{Term}_A \quad \underline{C} \mapsto \mathcal{R}[\underline{C}] \subseteq \underline{\text{Term}}_{\underline{C}} \times \underline{\text{Term}}_{\underline{C}}$$

indexed by value types A and by computation types \underline{C} . These are usually defined by induction on the structure of the types. They are *logical* in the sense that they respect the equations in Figure 3.1. These equations uniquely determine the logical relation on every type former, except for base types and returner types. For ground types these are standard. Note that for \mathbf{unit} the set $\text{Term}_{\mathbf{unit}}$ contains only one equivalence class. Since the only way to use a thunk is to force it, the definition on thunk types just requires the two forced computations to be related. For products of computation types the definition is similar to products of value types: we require that each of the projections are related. For function types, we require that related arguments are sent to related results.

The only type formers that are omitted from the figure are base types and returner types. For these, the logical relation must be chosen based on the constants and side-effects included in the language. In particular, the definition of the logical relation on returner types depends on the choice of operations $\text{op} \in \Sigma$. In the definition of logical relations, we impose some additional requirements on these:

Definition 3.1.1 (Logical relation) A *logical relation* consists of a relation $\mathcal{R}[A] \subseteq \text{Term}_A \times \text{Term}_A$ for each value type A and a relation $\mathcal{R}[\underline{C}] \subseteq \underline{\text{Term}}_{\underline{C}} \times \underline{\text{Term}}_{\underline{C}}$ for each computation type \underline{C} , such that:

- The relations satisfy the equations in Figure 3.1.
- Closure under operations: for each operation $\text{op} \in \Sigma$, if $(V, V') \in \mathcal{R}[\langle \text{car}_{\text{op}} \rangle]$ then $(\text{op } V, \text{op } V') \in \mathcal{R}[\langle \text{eff}_{\text{op}} \rangle \text{ar}_{\text{op}}]$.

- Closure under pure computations: if $(V, V') \in \mathcal{R}[[A]]$ then $(\langle V \rangle, \langle V' \rangle) \in \mathcal{R}[\langle 1 \rangle A]$.
- Closure under **to**: if $(M, M') \in \mathcal{R}[\langle \varepsilon \rangle A]$, and $x : A \perp N : \langle \varepsilon' \rangle B$ and $x : A \perp N' : \langle \varepsilon' \rangle B$ satisfy

$$\forall (V, V') \in \mathcal{R}[[A]]. (N[x \mapsto V], N'[x \mapsto V']) \in \mathcal{R}[\langle \varepsilon' \rangle B]$$

then

$$(M \text{ to } x. N, M' \text{ to } x. N') \in \mathcal{R}[\langle \varepsilon \cdot \varepsilon' \rangle B]$$

- Closure under **coerce**: if $(M, M') \in \mathcal{R}[\langle \varepsilon \rangle A]$ and $\varepsilon \leq \varepsilon'$ then

$$(\text{coerce}_{\varepsilon \leq \varepsilon'} M, \text{coerce}_{\varepsilon \leq \varepsilon'} M') \in \mathcal{R}[\langle \varepsilon' \rangle A]$$

- For each constant $c \in \mathcal{K}_A$ we have $(c, c) \in \mathcal{R}[[A]]$. ◀

In the definition, we require closure under **to** only when N and N' have returner type. This suffices to get closure under **to** for arbitrary computation types \underline{C} , because, as we noted in Section 2.7.2, having **to** only for returner types on the right-hand side suffices.

Lemma 3.1.2 Suppose that $\mathcal{R}[-]$ is a logical relation. If $(M, M') \in \mathcal{R}[\langle \varepsilon \rangle A]$, and $x : A \perp N : \underline{C}$ and $x : A \perp N' : \underline{C}$ are well-typed computations such that

$$\forall (V, V') \in \mathcal{R}[[A]]. (N[x \mapsto V], N'[x \mapsto V']) \in \mathcal{R}[\underline{C}]$$

then

$$(M \text{ to } x. N, M' \text{ to } x. N') \in \mathcal{R}[\langle \varepsilon \rangle \underline{C}]$$

Proof. By induction on the structure of \underline{C} . For **unit** this holds by the equation for $\mathcal{R}[\underline{\text{unit}}]$. For $\underline{C}_1 \times \underline{C}_2$, it suffices to show that

$$(i^i(M \text{ to } x. N), i^i(M' \text{ to } x. N')) \in \mathcal{R}[\langle \varepsilon \rangle \underline{C}_i]$$

for each $i \in \{1, 2\}$. We have

$$\begin{aligned} i^i(M \text{ to } x. N) &\equiv i^i(\lambda\{1. M \text{ to } x. 1^i N, 2. M \text{ to } x. 2^i N\}) \\ &\equiv M \text{ to } x. i^i N \end{aligned}$$

and similarly for $M' \text{ to } x. N'$. The result then follows from the inductive hypothesis, using the fact that

$$\forall (V, V') \in \mathcal{R}[[A]]. ((i^i N)[x \mapsto V], (i^i N')[x \mapsto V']) \in \mathcal{R}[\underline{C}_i]$$

For $B \rightarrow \underline{C}$ we show that

$$(W^i(M \text{ to } x. N), W^i(M' \text{ to } x. N')) \in \mathcal{R}[\langle \varepsilon \rangle \underline{C}]$$

for each $(W, W') \in \mathcal{R}[[B]]$. This is the case because

$$\begin{aligned} W^i(M \text{ to } x. N) &\equiv W^i(\lambda y:A. M \text{ to } x. y^i N) \\ &\equiv M \text{ to } x. W^i N \end{aligned}$$

and similarly for $M' \text{ to } x. N'$, so we can apply the inductive hypothesis using

$$\forall (V, V') \in \mathcal{R}[[A]]. ((W^i N)[x \mapsto V], (W^i N')[x \mapsto V']) \in \mathcal{R}[\underline{C}]$$

Finally, for returner types we have closure under **to** by assumption. □

$$\begin{aligned}
\mathcal{R}[\langle \emptyset \rangle A] &:= \{(\langle V \rangle, \langle V' \rangle) \mid (V, V') \in \mathcal{R}[A]\} \\
\mathcal{R}[\langle \{\text{get}\} \rangle A] &:= \left\{ \left(\begin{array}{l} \text{get } () \text{ to } b. \text{ if } b \text{ then } \langle W_1 \rangle \text{ else } \langle W_2 \rangle, \\ \text{get } () \text{ to } b. \text{ if } b \text{ then } \langle W'_1 \rangle \text{ else } \langle W'_2 \rangle \end{array} \right) \mid (W_1, W'_1), (W_2, W'_2) \in \mathcal{R}[A] \right\} \\
\mathcal{R}[\langle \{\text{put}\} \rangle A] &:= \{(\text{coerce}_{\emptyset \leq \{\text{put}\}} \langle W \rangle, \text{coerce}_{\emptyset \leq \{\text{put}\}} \langle W' \rangle) \mid (W, W') \in \mathcal{R}[A]\} \cup \text{MustPut } A \\
\mathcal{R}[\langle \{\text{get}, \text{put}\} \rangle A] &:= \left\{ \left(\begin{array}{l} \text{get } () \text{ to } b. \text{ if } b \text{ then } M_1 \text{ else } M_2, \\ \text{get } () \text{ to } b. \text{ if } b \text{ then } M'_1 \text{ else } M'_2 \end{array} \right) \mid (M_1, M'_1), (M_2, M'_2) \in \text{MustPut } A \right\}
\end{aligned}$$

where $(\text{MustPut } A) \subseteq \underline{\text{Term}}_{\langle \{\text{put}\} \rangle A} \times \underline{\text{Term}}_{\langle \{\text{put}\} \rangle A}$ is given by

$$\text{MustPut } A := \{(\text{put } V; \langle W \rangle), (\text{put } V; \langle W' \rangle) \mid V \in \{\mathbf{true}, \mathbf{false}\} \wedge (W, W') \in \mathcal{R}[A]\}$$

Figure 3.2: Logical relation for global state

For many concrete examples, it is difficult to show that closure under **to** holds for all computation types directly. We instead show that it holds for returner types.

Example 3.1.3 As an example, consider global state. We have no base types and a Gifford-style effect algebra with two operations: **get** with coarity **unit** and arity **bool** = **unit** + **unit**, and **put** with coarity **bool** and arity **unit**. We gave a (symmetric) inequational theory for this example in Figure 2.14. To define a logical relation, we only need to give $\mathcal{R}[\langle \varepsilon \rangle A]$ in terms of $\mathcal{R}[A]$ for each $\varepsilon \subseteq \{\text{get}, \text{put}\}$. We do this in Figure 3.2. The idea behind the definition is to think about how computations with each effect could behave. For example, a computation with effect $\{\text{put}\}$ can either just return a result (without setting the value of the state), or set the state once and then return. Computations that use **put** more than once are equivalent to computations that use **put** exactly once because $\text{put } V_1; \text{put } V_2 \equiv \text{put } V_2$ (lack of recursion implies that computations are finite).

The proof that the requirements in the definition of logical relation are met in this case uses the inequational theory. In particular, to show that the logical relations are closed under operations, we need to show that $(\text{get } (), \text{get } ()) \in \mathcal{R}[\langle \{\text{get}\} \rangle \mathbf{bool}]$ and if $V \in \{\mathbf{true}, \mathbf{false}\}$ then $(\text{put } V, \text{put } V) \in \mathcal{R}[\langle \{\text{put}\} \rangle \mathbf{unit}]$. These follow from the η -laws:

$$\begin{aligned}
\text{get } () &\equiv \text{get } () \text{ to } b. \langle b \rangle \equiv \text{get } () \text{ to } b. \text{ if } b \text{ then } \langle \mathbf{true} \rangle \text{ else } \langle \mathbf{false} \rangle \\
\text{put } V &\equiv \text{put } V \text{ to } x. \langle x \rangle \equiv \text{put } V \text{ to } x. \langle () \rangle \equiv \text{put } V; \langle () \rangle
\end{aligned}$$

So the terms have the correct form to be related by the logical relation. The signature axioms for global state come into play when showing that the logical relation is closed under pure computations and under **to**. To show closure under **to** for $\varepsilon = \varepsilon' = \{\text{get}\}$, we use the axiom that merges adjacent uses of **get**:

$$\begin{aligned}
&(\text{get } () \text{ to } b. \text{ if } b \text{ then } \langle V_1 \rangle \text{ else } \langle V_2 \rangle) \text{ to } x. \text{get } () \text{ to } b'. \text{ if } b' \text{ then } \langle W_1 \rangle \text{ else } \langle W_2 \rangle \\
&\equiv \text{get } () \text{ to } b. \text{get } () \text{ to } b'. \text{ if } b \text{ then if } b' \text{ then } \langle W_1[x \mapsto V_1] \rangle \text{ else } \langle W_2[x \mapsto V_1] \rangle \\
&\quad \text{else if } b' \text{ then } \langle W_1[x \mapsto V_2] \rangle \text{ else } \langle W_2[x \mapsto V_2] \rangle \\
&\equiv \text{get } () \text{ to } b. \text{ if } b \text{ then if } b \text{ then } \langle W_1[x \mapsto V_1] \rangle \text{ else } \langle W_2[x \mapsto V_1] \rangle \\
&\quad \text{else if } b \text{ then } \langle W_1[x \mapsto V_2] \rangle \text{ else } \langle W_2[x \mapsto V_2] \rangle \\
&\equiv \text{get } () \text{ to } b. \text{ if } b \text{ then } \langle W_1[x \mapsto V_1] \rangle \text{ else } \langle W_2[x \mapsto V_2] \rangle
\end{aligned}$$

We return to our global state example in Section 3.1.1. ◀

We extend logical relations to typing contexts Γ , for which we relate closed substitutions. Recall from Section 2.7 that substitutions σ satisfying $\diamond \vdash \sigma : \Gamma$ map variables $(x : A) \in \Gamma$ to closed terms of type A . Also recall that we extend the inequational theory to substitutions componentwise. The set of equivalence classes of closed substitutions is

$$\text{Subst}_\Gamma := \{[\sigma]_\equiv \mid \diamond \vdash \sigma : \Gamma\}$$

In this chapter we identify substitutions that are related by \equiv , and omit the square brackets on equivalence classes.

We extend logical relations $\mathcal{R}[-]$ to typing contexts Γ by defining a relation $\mathcal{R}[\Gamma] \subseteq \text{Subst}_\Gamma \times \text{Subst}_\Gamma$ on closed substitutions componentwise:

$$\mathcal{R}[\diamond] := \{(\diamond, \diamond)\} \quad \mathcal{R}[\Gamma, x : A] := \{((\sigma, x \mapsto V), (\sigma', x \mapsto V')) \mid (\sigma, \sigma') \in \mathcal{R}[\Gamma] \wedge (V, V') \in \mathcal{R}[A]\}$$

The key property that follows from the requirements in the definition of logical relation is the *fundamental lemma*.

Lemma 3.1.4 (Fundamental) Suppose that $\mathcal{R}[-]$ is a logical relation and $(\sigma, \sigma') \in \mathcal{R}[\Gamma]$.

1. (Values) If $\Gamma \vdash V : A$ then $(V[\sigma], V[\sigma']) \in \mathcal{R}[A]$.
2. (Computations) If $\Gamma \vdash M : \underline{C}$ then $(M[\sigma], M[\sigma']) \in \mathcal{R}[\underline{C}]$.

Proof sketch. By induction on the derivations of $\Gamma \vdash V : A$ and $\Gamma \vdash M : \underline{C}$. Most of the cases are standard, so we do not give them. For $\text{op } V$ and $\langle V \rangle$ we use the assumptions that the logical relation is closed under operations and pure computations. For $M \text{ to } x.N$ we use closure under **to** for arbitrary computation types (Lemma 3.1.2): by the inductive hypothesis we have $(M[\sigma], M[\sigma']) \in \mathcal{R}[\langle \varepsilon \rangle A]$, and

$$\forall (V, V') \in \mathcal{R}[A]. (N[\sigma, x \mapsto V], N[\sigma', x \mapsto V'])$$

Hence

$$((M[\sigma] \text{ to } x.N[\sigma]), (M[\sigma'] \text{ to } x.N'[\sigma'])) \in \mathcal{R}[\langle \varepsilon \rangle \underline{C}]$$

For **coerce** $_{\varepsilon \leq \varepsilon'}$ M we use closure under **coerce**. For constants we use the assumption about constants. \square

A consequence of the fundamental lemma is that each relation is reflexive (using the empty context for Γ and the empty substitution for σ and σ'). Hence if $M \equiv M'$ for closed, well-typed terms $M, M' : \underline{C}$, then $(M, M') \in \mathcal{R}[\underline{C}]$ (recall that we relate \equiv -equivalence classes of terms). In general the converse $(M, M') \in \mathcal{R}[\underline{C}] \Rightarrow M \equiv M'$ does not hold. Nor does $M \leq M' \Rightarrow (M, M') \in \mathcal{R}[\underline{C}]$.

In the fundamental lemma the logical relation is applied to open terms by quantifying over pairs of related substitutions. We do this frequently, and so define abbreviations for this notion. For values we write $(V, V') : \mathcal{R}[\Gamma] \rightarrow \mathcal{R}[A]$ if $\Gamma \vdash V : A$ and $\Gamma \vdash V' : A$, and $(V[\sigma], V'[\sigma']) \in \mathcal{R}[A]$ for all $(\sigma, \sigma') \in \mathcal{R}[\Gamma]$. Similarly, for computations we write $(M, M') : \mathcal{R}[\Gamma] \rightarrow \mathcal{R}[\underline{C}]$ if $\Gamma \vdash M : \underline{C}$ and $\Gamma \vdash M' : \underline{C}$, and $(M[\sigma], M'[\sigma']) \in \mathcal{R}[\underline{C}]$ for all $(\sigma, \sigma') \in \mathcal{R}[\Gamma]$.

Recall that the contextual preorder \leq_{ctx} is defined by considering closed term-contexts of type $\langle \varepsilon \rangle G$, where G is a ground type. In this chapter, we use logical relations to prove instances of \leq_{ctx} . We show that we can do this for logical relations that satisfy an additional assumption at the types $\langle \varepsilon \rangle G$. (This assumption is satisfied for our global state example.)

Lemma 3.1.5 Suppose that $\mathcal{R}[-]$ is a logical relation, and that for each ground type G , effect ε , and pair of computations N, N' we have:

$$(N, N') \in \mathcal{R}[\langle \varepsilon \rangle G] \quad \Rightarrow \quad N \leq N'$$

1. If $(V, V') : \mathcal{R}[\Gamma] \rightarrow \mathcal{R}[A]$ then $\Gamma \vdash V \leq_{\text{ctx}} V' : A$.
2. If $(M, M') : \mathcal{R}[\Gamma] \rightarrow \mathcal{R}[\underline{C}]$ then $\Gamma \vdash M \leq_{\text{ctx}} M' : \underline{C}$.

Proof sketch. For (1) we show for all typing contexts Γ' , types B and term contexts $C[\]$ with hole \square that if $\Gamma' \vdash C[V] : B$ and $\Gamma' \vdash C[V'] : B$ then $(C[V], C[V']) : \mathcal{R}[\Gamma'] \rightarrow \mathcal{R}[B]$, and similarly for computation-typed term contexts $\underline{C}[\]$. These are by induction on the typing derivations, and are similar to the proof of the fundamental lemma. For $C[\] = \square$ we use the assumption about (V, V') .

Then given a ground type G , effect ε and term context $\underline{C}[\]$ with hole \square such that $\diamond \vdash \underline{C}[V] : \langle \varepsilon \rangle G$ and $\diamond \vdash \underline{C}[V'] : \langle \varepsilon \rangle G$, we have $(C[V], C[V']) \in \mathcal{R}[\langle \varepsilon \rangle G]$ (using the above with the empty context for Γ'). By the assumption about ground types, it follows that $\underline{C}[V] \leq \underline{C}[V']$.

The proof of (2) is similar. □

3.1.1 Free lifting

Logical relations must be compatible with the side-effects and effect system that GCBPV is instantiated with. The main difficulty in constructing a logical relation is *lifting* a relation $\mathcal{R}[A]$ for the value type A , to form relations $\mathcal{R}[\langle \varepsilon \rangle A]$ on computations for each effect ε . There are several existing techniques for lifting relations to computations, such as $\top\top$ -lifting [62, 43] and codensity lifting [46]. Here we give a further lifting technique, called the *free lifting*. This technique is folklore, and is described for category-theoretic monads by Kammar and McDermott [41]. The free lifting has the important property that it is *initial*: it relates only the computations that must be related in order to form a logical relation (see Lemma 3.1.7). This does not mean it is the best possible lifting for all applications (a logical relation can relate too few terms), but it is for some (e.g. Example 3.4.5). The example logical relation we give above (for global state) is an instance of the free lifting.

The insight used to define the free lifting is that every closed computation of returner type, when evaluated, will execute a sequence of operations (get, put, raise, etc.), before possibly returning a result. This means that, up to \equiv , every closed computation has one of three forms: it is either a return $\langle V \rangle$, a coercion $\text{coerce}_{\varepsilon \leq \varepsilon'} M$ where M is a closed computation of returner type, or a sequencing $\text{op } V \text{ to } x. M$ of an operation op followed by a closed computation M . The free lifting is constructed as the smallest relation that relates computations of all three forms.

Given a relation $\mathcal{R}_{\text{free}}[b] \in \text{Term}_b \times \text{Term}_b$ for each base type b , we form a logical relation $\mathcal{R}_{\text{free}}[-]$, by defining $\mathcal{R}_{\text{free}}[A]$ on non-base types A by induction on the structure of A . We use the equations in Figure 3.1 as the definition on the rest of the type formers except returner types. It remains to provide, given a type A for which the logical relation is defined, a family of relations

$$\mathcal{R}_{\text{free}}[\langle \varepsilon \rangle A] \subseteq \underline{\text{Term}}_{\langle \varepsilon \rangle A} \times \underline{\text{Term}}_{\langle \varepsilon \rangle A}$$

indexed by $\varepsilon \in \mathcal{E}$. We define this family inductively by the following rules:

- If $(V, V') \in \mathcal{R}_{\text{free}}[A]$ then $(\langle V \rangle, \langle V' \rangle) \in \mathcal{R}_{\text{free}}[\langle 1 \rangle A]$.
- If $(M, M') \in \mathcal{R}_{\text{free}}[\langle \varepsilon \rangle A]$ and $\varepsilon \leq \varepsilon'$ then $(\text{coerce}_{\varepsilon \leq \varepsilon'} M, \text{coerce}_{\varepsilon \leq \varepsilon'} M') \in \mathcal{R}_{\text{free}}[\langle \varepsilon' \rangle A]$.

- For each $\text{op} \in \Sigma$, if $(M, M') : \mathcal{R}_{\text{free}}[[x : \text{ar}_{\text{op}}]] \dot{\rightarrow} \mathcal{R}_{\text{free}}[[\langle \varepsilon \rangle A]]$ and $(V, V') \in \mathcal{R}_{\text{free}}[[\text{car}_{\text{op}}]]$ then

$$((\text{op } V \text{ to } x. M), (\text{op } V' \text{ to } x. M')) \in \mathcal{R}_{\text{free}}[[\langle \text{eff}_{\text{op}} \cdot \varepsilon \rangle A]]$$

In this definition we use the logical relation on the arity and coarity of operations. The arity and coarity are ground types (in particular do not contain returner types), so the definition is well-founded.

Lemma 3.1.6 If $(c, c) \in \mathcal{R}_{\text{free}}[[A]]$ for each constant $c \in \mathcal{K}_A$, then $\mathcal{R}_{\text{free}}[[-]]$ is a logical relation.

Proof. See Appendix B.2. □

Hence the free lifting does give us a way of constructing logical relations. Logical relations constructed in this way are *initial* in the following sense:

Lemma 3.1.7 Suppose that $\mathcal{R}[[-]]$ is a logical relation and $\mathcal{R}_{\text{free}}[[b]] = \mathcal{R}[[b]]$ for all base types b . Then $\mathcal{R}_{\text{free}}[[A]] \subseteq \mathcal{R}[[A]]$ implies $\mathcal{R}_{\text{free}}[[\langle \varepsilon \rangle A]] \subseteq \mathcal{R}[[\langle \varepsilon \rangle A]]$.

Proof. See Appendix B.2. □

This lemma does not imply that $\mathcal{R}_{\text{free}}[[-]]$ is included in $\mathcal{R}[[-]]$ at *every* type, because of the contravariance of the logical relation at function types. However, it does imply that $\mathcal{R}_{\text{free}}[[\langle \varepsilon \rangle G]] \subseteq \mathcal{R}[[\langle \varepsilon \rangle G]]$ for ground types G . In particular, this means if any logical relation satisfies the assumptions in Lemma 3.1.5 (which we use to prove contextual equivalence), then a logical relation based on the free lifting does.

The definition of the free lifting slightly simplifies when we restrict to Gifford-style effect algebras. In particular, in the above definition of the free lifting we define the entire ε -indexed family of relations $\mathcal{R}[[\langle \varepsilon \rangle A]]$ by an induction. For Gifford-style effect algebras, we can instead define each relation $\mathcal{R}[[\langle \varepsilon \rangle A]]$ separately by induction.

Suppose that the effect algebra is the preordered monoid

$$(\mathcal{P}\Sigma, \subseteq, \cup, \emptyset)$$

so that each effect ε is just a set of operations, and that for each operation $\text{op} \in \Sigma$ we have $\text{eff}_{\text{op}} = \{\text{op}\}$. Given $\mathcal{R}_{\text{free}}[[A]]$, we define relations $\mathcal{R}'_{\text{free}}[[\langle \varepsilon \rangle A]]$ (our alternative definition). We show that the alternative definition coincides with our general definition below.

For each effect $\varepsilon \subseteq \Sigma$, the relation $\mathcal{R}'_{\text{free}}[[\langle \varepsilon \rangle A]] \subseteq \text{Term}_{\langle \varepsilon \rangle A} \times \text{Term}_{\langle \varepsilon \rangle A}$ is defined inductively by the following rules:

- If $(V, V') \in \mathcal{R}_{\text{free}}[[A]]$ then $(\text{coerce}_{\emptyset \leq \varepsilon} \langle V \rangle, \text{coerce}_{\emptyset \leq \varepsilon} \langle V' \rangle) \in \mathcal{R}'_{\text{free}}[[\langle \varepsilon \rangle A]]$.
- For each $\text{op} \in \varepsilon$, if $(M, M') : \mathcal{R}_{\text{free}}[[x : \text{ar}_{\text{op}}]] \dot{\rightarrow} \mathcal{R}'_{\text{free}}[[\langle \varepsilon \rangle A]]$ and $(V, V') \in \mathcal{R}_{\text{free}}[[\text{car}_{\text{op}}]]$ then

$$((\text{op } V \text{ to } x. M), (\text{op } V' \text{ to } x. M')) \in \mathcal{R}'_{\text{free}}[[\langle \varepsilon \rangle A]]$$

Lemma 3.1.8 For each value type A and effect $\varepsilon \subseteq \Sigma$,

$$\mathcal{R}'_{\text{free}}[[\langle \varepsilon \rangle A]] = \mathcal{R}_{\text{free}}[[\langle \varepsilon \rangle A]]$$

Proof. See Appendix B.2. □

This alternative definition is closer to the one given by Kammar and McDermott [41].

3.2 Restricting side-effects in call-by-value and call-by-name

As we mentioned in the introduction to this chapter, there are two major difficulties in relating the call-by-value and call-by-name translations of expressions when restricting side-effects. The first is that the side-effects of expressions depend on the evaluation order. This is why different evaluation orders have different effect systems (each of which can be instantiated with the same effect algebra). In particular, we have one effect system for call-by-value (Section 2.3) and another for Levy-style call-by-name. This means we cannot in general use the same restriction on side-effects for *both* call-by-value and call-by-name.

For the rest of this chapter, we therefore restrict to Gifford-style effect algebras (Example 2.1.2). Hence we take the preordered monoid to be $(\mathcal{P}\Sigma, \subseteq, \cup, \emptyset)$, where Σ is the set of operations from the signature. The two key properties that Gifford-style effect algebras satisfy are that the unit \emptyset is the least element and that the multiplication \cup is idempotent. Hence it does not matter that call-by-name might discard or duplicate side-effects; the call-by-value and call-by-name effect systems assign the same effects.

To impose restrictions on side-effects for both call-by-value and call-by-name we do the following. Recall that the syntax of source-language types τ does not have any effect annotations. For call-by-value and Levy-style call-by-name, we use types that do have effect annotations. Given any source-language type τ and effect ε , we write τ_ε^v and τ_ε^n respectively for the call-by-value and call-by-name types constructed by annotating τ with ε everywhere:

$$\begin{array}{l} \boxed{\tau_\varepsilon^n} \\ \boxed{\tau_\varepsilon^v} \end{array} \quad \begin{array}{l} \mathbf{unit}_\varepsilon^n := \langle \varepsilon \rangle \mathbf{unit} \\ \mathbf{unit}_\varepsilon^v := \mathbf{unit} \end{array} \quad \begin{array}{l} \mathbf{bool}_\varepsilon^n := \langle \varepsilon \rangle \mathbf{bool} \\ \mathbf{bool}_\varepsilon^v := \mathbf{bool} \end{array} \quad \begin{array}{l} (\tau \rightarrow \tau')_\varepsilon^n := \tau_\varepsilon^n \rightarrow \tau'_\varepsilon^n \\ (\tau \rightarrow \tau')_\varepsilon^v := \tau_\varepsilon^v \xrightarrow{\varepsilon} \tau'_\varepsilon^v \end{array}$$

We annotate source-language typing contexts Γ pointwise to get Γ_ε^v for call-by-value and Γ_ε^n for call-by-name. (Since expressions contain types we similarly have expressions e_ε^v and e_ε^n , although we write both of these just as e to avoid cluttering the notation. It can always be determined which we mean from context.) Define $\mathit{ops} \, e \subseteq \Sigma$ to be the set of operations that appear syntactically in the (unannotated) expression e . If we annotate types with an effect $\varepsilon \supseteq \mathit{ops} \, e$, then e is well-typed in both the call-by-value and call-by-name effect systems:

Lemma 3.2.1 If $\Gamma \vdash e : \tau$ and $\mathit{ops} \, e \subseteq \varepsilon$ then $\Gamma_\varepsilon^v \vdash_v e : \tau_\varepsilon^v \ \& \ \varepsilon$ and $\Gamma_\varepsilon^n \vdash_n e : \tau_\varepsilon^n$. ◀

Hence we can consider expressions restricted to ε in both evaluation orders simultaneously. Even though we translate derivations of expressions in this lemma, it does not matter which derivations we choose. This is because we identify GCBPV terms up to \equiv throughout this chapter, and can apply Lemma 2.7.8.

This way of restricting side-effects of source-language expressions, where we just annotate everything with the same effect, is slightly ad hoc. There may be better ways of restricting side-effects across evaluation orders. In particular, effect polymorphism may enable a more principled way of restricting side-effects.

3.3 A Galois connection between call-by-value and call-by-name

The second problem with relating the call-by-value and call-by-name translations of expressions is that they have different types. Given an unannotated source-language type τ and effect ε

we write $(\tau)_\varepsilon^v$ for $(\tau_\varepsilon^v)^v$ and $(\tau)_\varepsilon^n$ for $(\tau_\varepsilon^n)^n$, i.e. the call-by-value and call-by-name translations of τ into GCBPV, where all of the effect annotations are ε . Similarly for typing contexts. If $\Gamma \vdash e : \tau$ and $\text{ops } e \subseteq \varepsilon$ then we also write $(e)_\varepsilon^v$ and $(e)_\varepsilon^n$ for the call-by-value and Levy-style call-by-name translations of e with ε for the effect annotations. These are typable in the two effect systems by Lemma 3.2.1. Pictorially, we have:

$$\begin{array}{c} \Gamma \vdash e : \tau \\ \swarrow \quad \searrow \\ \Gamma_\varepsilon^v \vdash_v e : \tau_\varepsilon^v \ \& \ \varepsilon \mapsto (\Gamma)_\varepsilon^v \vdash (e)_\varepsilon^v : \langle \varepsilon \rangle (\tau)_\varepsilon^v \\ \Gamma_\varepsilon^n \vdash_n e : \tau_\varepsilon^n \longmapsto (\Gamma)_\varepsilon^n \vdash (e)_\varepsilon^n : (\tau)_\varepsilon^n \end{array}$$

The two computations on the right are the ones that we want to relate using the contextual preorder. However, we cannot ask whether $(e)_\varepsilon^v \ll_{\text{ctx}} (e)_\varepsilon^n$ because one computation has a call-by-value type and one has a call-by-name type. It does not make sense to replace $(e)_\varepsilon^v$ with $(e)_\varepsilon^n$ inside a GCBPV program, because the result would not be well-typed.

A similar problem arises when comparing two different denotational semantics of the same language. When comparing direct and continuation semantics of the lambda calculus, Reynolds [89], solves this problem by defining maps between the two semantics, so that a denotation in the direct semantics can be viewed as a denotation in the continuation semantics and vice versa. We use a similar idea here.

Specifically, we define maps Φ from call-by-value computations to call-by-name computations, and Ψ from call-by-name to call-by-value:

$$\Gamma \vdash M : \langle \varepsilon \rangle (\tau)_\varepsilon^v \mapsto \Gamma \vdash \Phi_{\tau,\varepsilon} M : (\tau)_\varepsilon^n \qquad \Gamma \vdash N : (\tau)_\varepsilon^n \mapsto \Gamma \vdash \Psi_{\tau,\varepsilon} N : \langle \varepsilon \rangle (\tau)_\varepsilon^v$$

Then instead of replacing $(e)_\varepsilon^v$ with $(e)_\varepsilon^n$ directly, we use Φ and Ψ to convert $(e)_\varepsilon^n$ to a computation of the correct type (defined formally in Section 3.4). The result looks something like the following composition of the translation of the expression e with two maps:

$$(\Gamma)_\varepsilon^v \longrightarrow (\Gamma)_\varepsilon^n \xrightarrow{(e)_\varepsilon^n} (\tau)_\varepsilon^n \longrightarrow \langle \varepsilon \rangle (\tau)_\varepsilon^v$$

This behaves like a call-by-name computation, but has the same type as a call-by-value computation. We could instead have chosen to convert $(e)_\varepsilon^v$ into a computation of the same type as $(e)_\varepsilon^n$. This choice is arbitrary, because of the properties of Galois connections.

We do not want just *any* maps between call-by-value and call-by-name. We show that under certain conditions (which is where the choice of side-effects becomes important) the maps we define form *Galois connections* [73]. This is crucial for the correctness of our reasoning principle. It also intuitively implies that it does not matter where we use Φ and Ψ inside programs.

The maps $\Phi_{\tau,\varepsilon}$ and $\Psi_{\tau,\varepsilon}$ are defined by induction on the structure of τ in Figure 3.3. We use some extra variables in the definition, which are assumed to be fresh. To go from a call-by-value computation to a call-by-name computation we first evaluate the call-by-value computation, and then map the result to call-by-name using $\hat{\Phi}_{\tau,\varepsilon}$, which has the following typing:

$$\Gamma \vdash V : (\tau)_\varepsilon^v \mapsto \Gamma \vdash \hat{\Phi}_{\tau,\varepsilon} V : (\tau)_\varepsilon^n$$

On the base types **bool** and **unit**, the maps Φ and Ψ do nothing. On function types, they convert the argument to the other evaluation order, apply the function, and then convert back. (Due to the contravariance of arguments the two maps are mutually defined.) These have the types stated because the unit effect \emptyset is the least element of the effect algebra, and the multiplication \cup is idempotent.

$$\begin{aligned}
\Phi_{\tau,\varepsilon}M &:= M \text{ to } x. \hat{\Phi}_{\tau,\varepsilon}x \\
\hat{\Phi}_{\text{unit},\varepsilon}V &:= \text{coerce}_{0 \leq \varepsilon} \langle V \rangle \\
\hat{\Phi}_{\text{bool},\varepsilon}V &:= \text{coerce}_{0 \leq \varepsilon} \langle V \rangle \\
\hat{\Phi}_{\tau \rightarrow \tau',\varepsilon}V &:= \lambda x : \mathbf{U} \langle \tau \rangle_{\varepsilon}^n. \Psi_{\tau,\varepsilon}(\text{force } x) \text{ to } z. \Phi_{\tau',\varepsilon}(z \text{ 'force } V) \\
\Psi_{\text{unit},\varepsilon}N &:= N \\
\Psi_{\text{bool},\varepsilon}N &:= N \\
\Psi_{\tau \rightarrow \tau',\varepsilon}N &:= \text{coerce}_{0 \leq \varepsilon} \langle \text{thunk } \lambda x : \langle \tau \rangle_{\varepsilon}^v. \Psi_{\tau',\varepsilon}((\text{thunk } (\hat{\Phi}_{\tau,\varepsilon}x)) \text{ ' } N) \rangle
\end{aligned}$$

Figure 3.3: Syntactic maps Φ from call-by-value to Levy-style call-by-name and Ψ from Levy-style call-by-name to call-by-value

In the rest of this section, we show that for side-effects satisfying certain conditions, the two maps form *Galois connections*. We do this with respect to some given logical relation $\mathcal{R}[\![-]\!]$, with the aim that we can use the logical relation to show instances of the contextual preorder (by applying Lemma 3.1.5). In addition to the usual requirements on logical relations (Definition 3.1.1), we assume that the relations $\mathcal{R}[\![A]\!]$ and $\mathcal{R}[\![C]\!]$ are transitive for each A and C . This implies (using Lemma 3.1.4) that each relation is a preorder.² We emphasize that the assumptions we have made about the logical relation so far are weak: it should be possible to define a useful logical relation satisfying these assumptions for every collection of side-effects (in particular, they hold for our global state example). We add further constraints below that allow us to derive our reasoning principle.

Definition 3.3.1 (Galois connection) A Galois connection (Φ, Ψ) from \underline{C} to \underline{D} is pair of maps

$$M : \underline{C} \mapsto \Phi M : \underline{D} \qquad N : \underline{D} \mapsto \Psi N : \underline{C}$$

such that for all closed $M : \underline{C}$ and $N : \underline{D}$,

$$(\Phi M, N) \in \mathcal{R}[\![\underline{D}]\!] \quad \Leftrightarrow \quad (M, \Psi N) \in \mathcal{R}[\![\underline{C}]\!] \quad \blacktriangleleft$$

In this definition we restrict to closed terms. Although $\Phi_{\tau,\varepsilon}$ and $\Psi_{\tau,\varepsilon}$ are defined for open terms, we do not need to consider open terms when showing that they form a Galois connection, because $\mathcal{R}[\![\Gamma]\!] \rightarrow \mathcal{R}[\![\underline{C}]\!]$ is defined in terms of $\mathcal{R}[\![\underline{C}]\!]$ on closed computations, and this is what we need to prove instances of.

For $(\Phi_{\tau,\varepsilon}, \Psi_{\tau,\varepsilon})$ to be a Galois connection the following must hold:

$$(\Phi_{\tau,\varepsilon}(\Psi_{\tau,\varepsilon}M), M) \in \mathcal{R}[\![\langle \tau \rangle_{\varepsilon}^n]\!] \qquad (N, \Psi_{\tau,\varepsilon}(\Phi_{\tau,\varepsilon}N)) \in \mathcal{R}[\![\langle \varepsilon \rangle \langle \tau \rangle_{\varepsilon}^v]\!]$$

The right cannot hold in general when τ is a function type $\tau' \rightarrow \tau''$, because the term $\Psi_{\tau,\varepsilon}(\Phi_{\tau,\varepsilon}N)$ has no side-effects, even if N does. If we convert a computation N from call-by-value to call-by-name and then back using these maps then the side-effects of N are *thunked*. The side-effects do not occur until the function is applied. For $\tau = \text{unit} \rightarrow \text{unit}$ we get

$$\Psi_{\text{unit} \rightarrow \text{unit},\varepsilon}(\Phi_{\text{unit} \rightarrow \text{unit},\varepsilon}N) \equiv \text{coerce}_{0 \leq \varepsilon} \langle \text{thunk } \lambda x : \text{unit}. N \text{ to } z. x \text{ ' force } z \rangle$$

²Galois connections are normally defined for partial orders. We do not have antisymmetry, and hence need to generalize to preorders, but this is unimportant.

Based on this observation, we require a further constraint on the logical relation to ensure that we get Galois connections. This is where it becomes important which side-effects we have: it is only satisfied in certain cases (see Example 3.3.5). The constraint is that computations are *thinkable*:

Definition 3.3.2 (Thinkable) A closed computation $M : \langle \varepsilon \rangle A$ is *thinkable* if

$$(M \text{ to } x. \langle \mathbf{think} \text{ coerce}_{0 \leq \varepsilon} \langle x \rangle \rangle, \text{coerce}_{0 \leq \varepsilon} \langle \mathbf{think} M \rangle) \in \mathcal{R}[\langle \varepsilon \rangle \mathbf{U} \langle \varepsilon \rangle A]$$

An effect ε is *thinkable* if for every value type A , every closed computation $M : \langle \varepsilon \rangle A$ is thinkable. \blacktriangleleft

This property was first defined in a symmetric setting (non-enriched Kleisli categories) by Führmann [27]. Our setting is not in general symmetric (even though we only give a symmetric example in this chapter), so we give a directed version (an alternative name might be *lax thinkable*). We only need one direction to hold for the proofs.

For intuition, suppose we have a computation $y : \mathbf{U} \langle \varepsilon \rangle A \vdash N : \underline{C}$. Then

$$M \text{ to } x. N[y \mapsto \mathbf{think} \text{ coerce}_{0 \leq \varepsilon} \langle x \rangle] \quad \text{coerce}_{\underline{C} <: \langle \varepsilon \rangle \underline{C}} (N[y \mapsto \mathbf{think} M])$$

are two closed computations of type $\langle \varepsilon \rangle \underline{C}$. The left computation evaluates M exactly once (like function arguments in call-by-value); the right potentially evaluates M several times (like call-by-name). If M is thinkable, then the left is related to the right by $\mathcal{R}[\langle \varepsilon \rangle \underline{C}]$. In other words, it is correct to replace a computation that evaluates M once with one that evaluates M zero or more times.

A key example is thinking the side-effects of call-by-value functions. Every computation $M : \langle \varepsilon \rangle (\tau' \rightarrow \tau'')_{\varepsilon}^{\vee}$ has two opportunities to have side-effects: when evaluating M to a function (immediate), and after applying the function (latent). If M is thinkable, the immediate side-effects can be turned into latent side-effects, because M is related to the following computation by $\mathcal{R}[\langle \varepsilon \rangle (\tau' \rightarrow \tau'')_{\varepsilon}^{\vee}]$:

$$\text{coerce}_{0 \leq \varepsilon} \langle \mathbf{think} \lambda x : (\tau')_{\varepsilon}^{\vee}. M \text{ to } f. x \text{ 'force } f \rangle$$

where the M is now inside the lambda.

To solve the issue with $\Psi_{\mathbf{unit} \rightarrow \mathbf{unit}, \varepsilon} (\Phi_{\mathbf{unit} \rightarrow \mathbf{unit}, \varepsilon} M)$ above it suffices for M to be thinkable. In fact, if enough computations are thinkable then the maps defined above do in fact form Galois connections:

Theorem 3.3.3 Suppose that the effect ε is thinkable. For every source-language type τ the pair $(\Phi_{\tau, \varepsilon}, \Psi_{\tau, \varepsilon})$ is a Galois connection from $\langle \varepsilon \rangle (\tau)_{\varepsilon}^{\vee}$ to $(\tau)_{\varepsilon}^{\text{n}}$.

Proof. By induction on the structure on τ . This is trivial for **unit** and **bool**, because the maps are just identities.

The interesting case is when τ is a function type $\tau' \rightarrow \tau''$. We first show that $(M, \Psi_{\tau, \varepsilon} N) \in \mathcal{R}[\langle \varepsilon \rangle (\tau)_{\varepsilon}^{\vee}]$ implies $(\Phi_{\tau, \varepsilon} M, N) \in \mathcal{R}[(\tau)_{\varepsilon}^{\text{n}}]$. It suffices to show for arbitrary $(P, P') \in \mathcal{R}[(\tau')_{\varepsilon}^{\text{n}}]$ that we have

$$(\mathbf{think} P) \text{ ' } \Phi_{\tau, \varepsilon} M \quad \mathcal{R} \quad (\mathbf{think} P') \text{ ' } N$$

We reason as follows:

$$\begin{aligned} & M \text{ to } x. \Psi_{\tau', \varepsilon} P \text{ to } z. z \text{ 'force } x \\ \mathcal{R} & \Psi_{\tau, \varepsilon} N \text{ to } x. \Psi_{\tau', \varepsilon} P \text{ to } z. z \text{ 'force } x && \text{(assumption)} \\ \equiv & \Psi_{\tau', \varepsilon} P \text{ to } z. \Psi_{\tau'', \varepsilon} ((\mathbf{think} \Phi_{\tau', \varepsilon} (\text{coerce}_{0 \leq \varepsilon} \langle z \rangle)) \text{ ' } N) \\ \mathcal{R} & \Psi_{\tau'', \varepsilon} ((\mathbf{think} \Phi_{\tau', \varepsilon} (\Psi_{\tau', \varepsilon} P)) \text{ ' } N) && (\varepsilon \text{ is thinkable}) \\ \mathcal{R} & \Psi_{\tau'', \varepsilon} ((\mathbf{think} P') \text{ ' } N) && \text{(Galois connection on } \tau') \end{aligned}$$

This and the fact that we have a Galois connection on τ'' then imply

$$\mathbf{thunk} P \text{ ' } \Phi_{\tau,\varepsilon} M \quad \equiv \quad \Phi_{\tau'',\varepsilon} (M \text{ to } x. \Psi_{\tau',\varepsilon} P \text{ to } z. z \text{ ' force } x) \quad \mathcal{R} \quad (\mathbf{thunk} P') \text{ ' } N$$

as required.

Second, we show that $(\Phi_{\tau,\varepsilon} M, N) \in \mathcal{R}[\langle \varepsilon \rangle (\tau)_{\varepsilon}^{\mathbb{N}}]$ implies $(M, \Psi_{\tau,\varepsilon} N) \in \mathcal{R}[\langle \varepsilon \rangle (\tau)_{\varepsilon}^{\mathbb{V}}]$. We have already mentioned that M is related to

$$\mathbf{coerce}_{\emptyset \leq \varepsilon} \langle \mathbf{thunk} \lambda x : (\tau')_{\varepsilon}^{\mathbb{V}}. M \text{ to } f. x \text{ ' force } f \rangle$$

(using the fact that M is thunkable) so it suffices to show that this computation is related to $\Psi_{\tau,\varepsilon} N$ by $\mathcal{R}[\langle \varepsilon \rangle (\tau)_{\varepsilon}^{\mathbb{V}}]$. By looking at the definitions of $\Psi_{\tau,\varepsilon}$ and this relation, and using the fact that logical relations are closed under **coerce**, it is therefore enough to show that, for all $(V, V') \in \mathcal{R}[\langle \varepsilon \rangle (\tau')_{\varepsilon}^{\mathbb{V}}]$, the following computations are related by $\mathcal{R}[\langle \varepsilon \rangle (\tau')_{\varepsilon}^{\mathbb{V}}]$:

$$M \text{ to } f. V \text{ ' force } f \quad \Psi_{\tau',\varepsilon} ((\mathbf{thunk} (\hat{\Phi}_{\tau',\varepsilon} V')) \text{ ' } N)$$

We prove this by first reasoning as follows:

$$\begin{aligned} \Phi_{\tau'',\varepsilon} (M \text{ to } f. V \text{ ' force } f) &\equiv \Phi_{\tau'',\varepsilon} (M \text{ to } f. \mathbf{coerce}_{\emptyset \leq \varepsilon} \langle V \rangle \text{ to } z. z \text{ ' force } f) \\ &\mathcal{R} \Phi_{\tau'',\varepsilon} (M \text{ to } f. \Psi_{\tau',\varepsilon} (\Phi_{\tau',\varepsilon} (\mathbf{coerce}_{\emptyset \leq \varepsilon} \langle V' \rangle)) \text{ to } z. z \text{ ' force } f) \\ &\quad \text{(Galois connection on } \tau' \text{)} \\ &\equiv (\mathbf{thunk} (\hat{\Phi}_{\tau',\varepsilon} V')) \text{ ' } \Phi_{\tau,\varepsilon} M \\ &\mathcal{R} (\mathbf{thunk} (\hat{\Phi}_{\tau',\varepsilon} V')) \text{ ' } N \quad \text{(assumption)} \end{aligned}$$

and then using the fact that we have a Galois connection for τ'' . \square

This theorem also has a partial converse: if the maps form Galois connections, then computations that return elements of base types are thunkable:

Lemma 3.3.4 If the pair $(\Phi_{\mathbf{unit} \rightarrow \mathbf{unit}, \varepsilon}, \Psi_{\mathbf{unit} \rightarrow \mathbf{unit}, \varepsilon})$ is a Galois connection then every closed computation $M : \langle \varepsilon \rangle \tau$, where $\tau \in \{\mathbf{unit}, \mathbf{bool}\}$, is thunkable.

Proof. Define the closed computation $M' : \langle \varepsilon \rangle (\mathbf{unit} \rightarrow \tau)_{\varepsilon}^{\mathbb{V}}$ by

$$M' := M \text{ to } x. \langle \mathbf{thunk} \lambda _ : \mathbf{unit}. \mathbf{coerce}_{\emptyset \leq \varepsilon} \langle x \rangle \rangle$$

By a standard property of Galois connections, M' is related by $\mathcal{R}[\langle \varepsilon \rangle (\mathbf{unit} \rightarrow \tau)_{\varepsilon}^{\mathbb{V}}]$ to the computation $\Psi_{\mathbf{unit} \rightarrow \tau, \varepsilon} (\Phi_{\mathbf{unit} \rightarrow \tau, \varepsilon} M')$. Hence:

$$\begin{aligned} M \text{ to } x. \langle \mathbf{thunk} \mathbf{coerce}_{\emptyset \leq \varepsilon} \langle x \rangle \rangle &\equiv M' \text{ to } t. \langle \mathbf{thunk} (()) \text{ ' force } t \rangle \\ &\mathcal{R} \Psi_{\mathbf{unit} \rightarrow \tau, \varepsilon} (\Phi_{\mathbf{unit} \rightarrow \tau, \varepsilon} M') \text{ to } t. \langle \mathbf{thunk} (()) \text{ ' force } t \rangle \\ &\equiv \mathbf{coerce}_{\emptyset \leq \varepsilon} \langle \mathbf{thunk} M \rangle \quad \square \end{aligned}$$

This proof works for base types τ because the two maps $\Phi_{\tau,\varepsilon}$ and $\Psi_{\tau,\varepsilon}$ are identities.

We end this section by returning to our main example.

Example 3.3.5 Recall our global state example. There are four possible effects: \emptyset , $\{\text{get}\}$, $\{\text{put}\}$ and $\{\text{get}, \text{put}\}$. As expected, the effect \emptyset is thunkable. This is the case because if $M : \langle \emptyset \rangle A$ then by the fundamental lemma (Lemma 3.1.4) there is some $V : A$ such that $M \equiv \langle V \rangle$. We therefore have:

$$M \text{ to } x. \langle \mathbf{thunk} \langle x \rangle \rangle \equiv \langle \mathbf{thunk} \langle V \rangle \rangle \equiv \mathbf{coerce}_{\emptyset \leq \emptyset} \langle \mathbf{thunk} M \rangle$$

and these terms are related by $\mathcal{R}[\langle \emptyset \rangle A]$.

The effect $\{\text{put}\}$ is not thinkable, because the computation put true is not. This is also what we would expect: if put true were thinkable then we would have

$$\text{put true to } x. \langle \mathbf{thunk} \text{coerce}_{\emptyset \leq \{\text{put}\}} \langle x \rangle \rangle \cong_{\text{ctx}} \text{coerce}_{\emptyset \leq \{\text{put}\}} \langle \mathbf{thunk} (\text{put true}) \rangle$$

which implies (using the computation context \square to t . put false ; $\mathbf{force } t$) that $\text{put false} \leq \text{put true}$. Clearly we do not want this to be the case (and it is not). The effect $\{\text{get}, \text{put}\}$ similarly is not thinkable.

The most interesting case is $\{\text{get}\}$. This effect is not thinkable because the computation $\text{get } ()$ is not (this can be verified trivially by looking at the definition of the logical relation). We would not expect get to be thinkable because the state is not read-only: if get were thinkable then we would have

$$\text{get } () \text{ to } x. \langle \mathbf{thunk} \text{coerce}_{\emptyset \leq \{\text{get}\}} \langle x \rangle \rangle \cong_{\text{ctx}} \text{coerce}_{\emptyset \leq \{\text{get}\}} \langle \mathbf{thunk} (\text{get } ()) \rangle \quad (3.1)$$

and using the context put false ; \square to t . (put true ; $\mathbf{force } t$) this implies

$$\text{coerce}_{\{\text{put}\} \leq \{\text{get}, \text{put}\}} (\text{put true}; \langle \mathbf{false} \rangle) \equiv \text{coerce}_{\{\text{put}\} \leq \{\text{get}, \text{put}\}} (\text{put true}; \langle \mathbf{true} \rangle)$$

which is not the case because the results of the two computations differ.³

To summarize: the only effect that is thinkable for this example is \emptyset , and hence \emptyset is the only effect for which we can apply our reasoning principle (more interesting examples are discussed in Section 4.4). In other words, we can only replace call-by-value with call-by-name in our global state example for subterms that have no side-effects. (Though the rest of the program might still have side-effects.) \blacktriangleleft

3.4 Reasoning principle for call-by-value and call-by-name

We now use the Galois connections defined in the previous section to relate the call-by-value and Levy-style call-by-name translations of expressions, and arrive at our reasoning principle.

Recall that we compose the call-by-name translation of each expression e with the maps Φ and Ψ defined above, to arrive at a GCBPV computation of the same type as the call-by-value translation⁴:

$$(\Gamma)_{\varepsilon}^{\vee} \longrightarrow (\Gamma)_{\varepsilon}^{\text{n}} \xrightarrow{\langle e \rangle_{\varepsilon}^{\text{n}}} (\tau)_{\varepsilon}^{\text{n}} \longrightarrow \langle \varepsilon \rangle (\tau)_{\varepsilon}^{\vee}$$

We first give a precise definition of this composition. The arrow on the right is just given by applying $\Psi_{\tau, \varepsilon}$. The arrow on the left is a substitution $\hat{\Phi}_{\Gamma, \varepsilon}$ from terms in call-by-name contexts $(\Gamma)_{\varepsilon}^{\text{n}}$ to terms in call-by-value contexts $(\Gamma)_{\varepsilon}^{\vee}$. We define $\hat{\Phi}_{\Gamma, \varepsilon}$ for the source-language typing context $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$ as

$$\hat{\Phi}_{\Gamma, \varepsilon} := x_1 \mapsto \mathbf{thunk} (\hat{\Phi}_{\tau_1, \varepsilon} x_1), \dots, x_n \mapsto \mathbf{thunk} (\hat{\Phi}_{\tau_n, \varepsilon} x_n)$$

³However, even if we removed put from the language, and used the same definition of the logical relation (but omitting the $\{\text{put}\}$ and $\{\text{get}, \text{put}\}$ cases, the computation $\text{get } ()$ would still not be thinkable with respect to the logical relation. This suggests that the definition of $\mathcal{R}[\langle \text{get} \rangle]$ above (and hence the free lifting) might not be suitable in that case. We might be able to use a logical relation that relates more computations.

⁴Also recall that we identify GCBPV terms up to \equiv throughout this chapter, and therefore, because the effect algebra is Gifford-style, it does not matter which typing derivation for e we choose to translate. We use this fact implicitly in this section.

The composition above is then the term $\Psi_{\tau,\varepsilon}((e)_\varepsilon^n[\hat{\Phi}_{\Gamma,\varepsilon}])$. If $\Gamma \vdash e : \tau$ and $ops\ e \subseteq \varepsilon$ then $(\Gamma)_\varepsilon^v \vdash \Psi_{\tau,\varepsilon}((e)_\varepsilon^n[\hat{\Phi}_{\Gamma,\varepsilon}]) : \langle \varepsilon \rangle (\tau)_\varepsilon^v$, so this computation has the same typing as $(e)_\varepsilon^v$.

The reasoning principle we derive proves instances of the contextual preorder that have the form

$$(e)_\varepsilon^v \leq_{\text{ctx}} \Psi_{\tau,\varepsilon}((e)_\varepsilon^n[\hat{\Phi}_{\Gamma,\varepsilon}])$$

This works by first showing the two terms are related by the logical relation:

$$((e)_\varepsilon^v, \Psi_{\tau,\varepsilon}((e)_\varepsilon^n[\hat{\Phi}_{\Gamma,\varepsilon}])) : \mathcal{R}[(\Gamma)_\varepsilon^v] \dot{\rightarrow} \mathcal{R}[\langle \varepsilon \rangle (\tau)_\varepsilon^v]$$

and then applying Lemma 3.1.5, which allows us to promote this to an instance of the contextual preorder. We prove that the two terms are related using the properties of Galois connections, which allow us to push composition with $\Psi_{\tau,\varepsilon}$ into the structure of terms.

Lemma 3.4.1 Suppose that $(\Phi_{\tau',\varepsilon}, \Psi_{\tau',\varepsilon})$ is a Galois connection from $\langle \varepsilon \rangle (\tau')_\varepsilon^v$ to $(\tau')_\varepsilon^n$ for every source language type τ' . If $\Gamma \vdash e : \tau$ and $ops\ e \subseteq \varepsilon$ then

$$((e)_\varepsilon^v, \Psi_{\tau,\varepsilon}((e)_\varepsilon^n[\hat{\Phi}_{\Gamma,\varepsilon}])) : \mathcal{R}[(\Gamma)_\varepsilon^v] \dot{\rightarrow} \mathcal{R}[\langle \varepsilon \rangle (\tau)_\varepsilon^v]$$

Proof. By induction on the derivation of $\Gamma \vdash e : \tau$. In each case of the induction we consider an arbitrary pair $(\sigma, \sigma') \in \mathcal{R}[(\Gamma)_\varepsilon^v]$ (recall that this is how $\dot{\rightarrow}$ is defined). We define σ'' to be the composition of $\hat{\Phi}_{\Gamma,\varepsilon}$ with σ' . In each case we are required to show $((e)_\varepsilon^v[\sigma], \Psi_{\tau,\varepsilon}((e)_\varepsilon^n[\sigma''])) \in \mathcal{R}[\langle \varepsilon \rangle (\tau)_\varepsilon^v]$. We only give three representative cases.

- If e is a variable x then, writing V for $\sigma\ x$ and V' for $\sigma'\ x$, we have:

$$(x)_\varepsilon^v[\sigma] \equiv \mathbf{coerce}_{0 \leq \varepsilon} \langle V \rangle \quad \mathcal{R} \quad \Psi_{\tau,\varepsilon}(\Phi_{\tau,\varepsilon'}(\mathbf{coerce}_{0 \leq \varepsilon} \langle V' \rangle)) \equiv \Psi_{\tau,\varepsilon}((x)_\varepsilon^n[\sigma''])$$

- If e is a λ -abstraction $\lambda x : \tau'. e$ and $\tau = \tau' \rightarrow \tau''$ then (expanding the definition of $\Psi_{\tau,\varepsilon}$) it suffices to show that $(\lambda x : (\tau')_\varepsilon^v. (e)_\varepsilon^v)[\sigma]$ is related by $\mathcal{R}[(\tau')_\varepsilon^v \rightarrow \langle \varepsilon \rangle (\tau'')_\varepsilon^v]$ to

$$(\lambda x : (\tau')_\varepsilon^v. \Psi_{\tau'',\varepsilon}((e)_\varepsilon^n[x \mapsto \mathbf{thunk}(\hat{\Phi}_{\tau',\varepsilon}x)]))[\sigma'']$$

Consider arbitrary $(W, W') \in \mathcal{R}[(\tau')_\varepsilon^v]$. We have:

$$\begin{aligned} W \text{ ' } (\lambda x : (\tau')_\varepsilon^v. (e)_\varepsilon^v)[\sigma] &\equiv (e)_\varepsilon^v[\sigma, x \mapsto W] \\ &\mathcal{R} \quad \Psi_{\tau'',\varepsilon}((e)_\varepsilon^n[\sigma'', x \mapsto \mathbf{thunk}(\hat{\Phi}_{\tau',\varepsilon}W)]) \\ &\equiv W' \text{ ' } (\lambda x : (\tau')_\varepsilon^v. \Psi_{\tau'',\varepsilon}((e)_\varepsilon^n[x \mapsto \mathbf{thunk}(\hat{\Phi}_{\tau',\varepsilon}x)]))[\sigma''] \end{aligned}$$

- If e is a function application $e_1\ e_2$, where e_1 has type $\tau' \rightarrow \tau$, then:

$$\begin{aligned} (e_1\ e_2)_\varepsilon^v[\sigma] &\equiv (e_1)_\varepsilon^v \text{ to } f. (e_2)_\varepsilon^v \text{ to } x. x \text{ ' force } f \\ &\mathcal{R} \quad \Psi_{\tau' \rightarrow \tau, \varepsilon}((e_1)_\varepsilon^n[\sigma'']) \text{ to } f. \Psi_{\tau', \varepsilon}((e_2)_\varepsilon^n[\sigma'']) \text{ to } x. x \text{ ' force } f \\ &\mathcal{R} \quad \Psi_{\tau, \varepsilon}(\Phi_{\tau, \varepsilon}(\Psi_{\tau' \rightarrow \tau, \varepsilon}((e_1)_\varepsilon^n[\sigma'']))) \text{ to } f. \Psi_{\tau', \varepsilon}((e_2)_\varepsilon^n[\sigma'']) \text{ to } x. x \text{ ' force } f \\ &\hspace{15em} (\text{Galois connection on } \tau) \\ &\equiv \Psi_{\tau, \varepsilon}((\mathbf{thunk}((e_2)_\varepsilon^n[\sigma'']))) \text{ ' } \Phi_{\tau' \rightarrow \tau, \varepsilon}(\Psi_{\tau' \rightarrow \tau, \varepsilon}((e_1)_\varepsilon^n[\sigma'']))) \\ &\mathcal{R} \quad \Psi_{\tau, \varepsilon}((\mathbf{thunk}((e_2)_\varepsilon^n[\sigma'']))) \text{ ' } (e_1)_\varepsilon^n[\sigma''] \\ &\hspace{15em} (\text{Galois connection on } \tau' \rightarrow \tau) \\ &\equiv \Psi_{\tau, \varepsilon}((e_1\ e_2)_\varepsilon^n[\sigma'']) \quad \square \end{aligned}$$

We showed in the previous section (Theorem 3.3.3) that the maps between the two evaluation orders form Galois connections if the effect ε is thunkable. Hence we arrive at our reasoning principle, which we state formally as Theorem 3.4.2. Recall that each inequational theory \leq (Definition 2.7.4) induces a contextual preorder \leq_{ctx} (Definition 2.7.6). Given any such inequational theory, to show that the call-by-value and call-by-name translations of source-language expressions restricted to the effect ε are related by \leq_{ctx} it is enough to define a suitable logical relation, and show that the effect ε is thunkable.

Theorem 3.4.2 (Call-by-value and call-by-name) Suppose we given some GCBPV inequational theory and a logical relation such that:

- Each relation $\mathcal{R}[[A]]$ and each relation $\mathcal{R}[[C]]$ is transitive.
- For each ground type G and effect ε' we have:

$$(N, N') \in \mathcal{R}[[\langle \varepsilon \rangle G]] \quad \Rightarrow \quad N \leq N'$$

for all N, N' .

- The effect ε is thunkable.

If $\Gamma \vdash e : \tau$ and $\text{ops } e \subseteq \varepsilon$ then

$$(e)_\varepsilon^v \leq_{\text{ctx}} \Psi_{\tau, \varepsilon}((e)_\varepsilon^n[\hat{\Phi}_{\Gamma, \varepsilon}])$$

Proof. By Theorem 3.3.3, the maps between call-by-value and call-by-name computations form Galois connections. Hence we can apply Lemma 3.4.1, which tells us that

$$((e)_\varepsilon^v, \Psi_{\tau, \varepsilon}((e)_\varepsilon^n[\hat{\Phi}_{\Gamma, \varepsilon}])) : \mathcal{R}[[\langle \Gamma \rangle_\varepsilon^v]] \rightarrow \mathcal{R}[[\langle \varepsilon \rangle(\tau)_\varepsilon^v]]$$

Finally, Lemma 3.1.5 implies the corresponding instance of the contextual preorder, which is the result we want. \square

The generality of this theorem comes from two sources. First, we consider arbitrary inequational theories \leq in which the effect algebra is Gifford-style. (It may be possible to prove a similar reasoning principle for other effect algebras on a case-by-case basis.) The only other requirement is the existence of a suitable logical relation in which enough computations are thunkable. Second, this theorem applies to terms that are open and have higher types, using the maps between the two evaluation orders. We obtain a result about source-language *programs* (closed expressions of base types) as a corollary. The corollary is closer to the standard results that are proved for specific side-effects, because the maps are trivial.

Corollary 3.4.3 Suppose that the assumptions of Theorem 3.4.2 hold. If e is a closed source-language expression of type $\tau \in \{\mathbf{unit}, \mathbf{bool}\}$ and $\text{ops } e \subseteq \varepsilon$ then $(e)_\varepsilon^v \leq (e)_\varepsilon^n$. \blacktriangleleft

Like our result about Galois connections, the reasoning principle has a partial converse. If call-by-value can be replaced with call-by-name, then computations of certain types are thunkable:

Lemma 3.4.4 Suppose that for each $\Gamma \vdash e : \tau$ with $\text{ops } e \subseteq \varepsilon$ we have

$$((e)_\varepsilon^v, \Psi_{\tau, \varepsilon}((e)_\varepsilon^n[\hat{\Phi}_{\Gamma, \varepsilon}])) : \mathcal{R}[[\langle \Gamma \rangle_\varepsilon^v]] \rightarrow \mathcal{R}[[\langle \varepsilon \rangle(\tau)_\varepsilon^v]]$$

For every base type $A \in \{\mathbf{unit}, \mathbf{bool}\}$, all closed computations $M : \langle \varepsilon \rangle A$ are thunkable.

Proof. First we show that every closed computation $N : \langle \varepsilon \rangle (\tau)_\varepsilon^v$ is related to $\Psi_{\tau,\varepsilon}(\Phi_{\tau,\varepsilon}N)$ (this is one of the properties of Galois connections). To do this, define the following source-language expression e and GCBPV value V :

$$\begin{aligned} e &:= x () & (x : \mathbf{unit} \rightarrow \tau \vdash e : \tau) \\ V &:= \mathbf{thunk} (\lambda y : \mathbf{unit}. N) & (\diamond \vdash V : \mathbf{U} (\mathbf{unit} \rightarrow \langle \varepsilon \rangle (\tau)_\varepsilon^v)) \end{aligned}$$

Now we apply the assumption to e :

$$N \equiv \langle e \rangle_\varepsilon^v [x \mapsto V] \quad \mathcal{R} \quad \Psi_{\tau,\varepsilon} (\langle e \rangle_\varepsilon^n [x \mapsto \mathbf{thunk} (\hat{\Phi}_{\mathbf{unit} \rightarrow \tau, \varepsilon} V)]) \equiv \Psi_{\tau,\varepsilon} (\Phi_{\tau,\varepsilon} N)$$

We use this to prove the result as follows. Let τ' be either \mathbf{unit} or \mathbf{bool} , so that $(\tau')_\varepsilon^v = A$. Define the closed computation $N : \langle \varepsilon \rangle (\mathbf{unit} \rightarrow \tau')_\varepsilon^v$ by:

$$N := M \text{ to } x. \langle \mathbf{thunk} \lambda y : \mathbf{unit}. \mathbf{coerce}_{\emptyset \leq \varepsilon} \langle x \rangle \rangle$$

Using the above with $\tau = \mathbf{unit} \rightarrow \tau'$ we have:

$$\begin{aligned} M \text{ to } x. \langle \mathbf{thunk} \mathbf{coerce}_{\emptyset \leq \varepsilon} \langle x \rangle \rangle &\equiv N \text{ to } f. \langle \mathbf{thunk} (()) \text{force } f \rangle \\ &\mathcal{R} \quad \Psi_{\mathbf{unit} \rightarrow \tau', \varepsilon} (\Phi_{\mathbf{unit} \rightarrow \tau', \varepsilon} N) \text{ to } f. \langle \mathbf{thunk} (()) \text{force } f \rangle \\ &\equiv \mathbf{coerce}_{\emptyset \leq \varepsilon} \langle \mathbf{thunk } M \rangle \end{aligned}$$

so M is thinkable. □

We again return to our main example, and apply our reasoning principle to it.

Example 3.4.5 Recall that for our global state example the effect \emptyset is thinkable, but the other three effects $\{\text{get}\}$, $\{\text{put}\}$, and $\{\text{get}, \text{put}\}$ are not. (So the maps $\Phi_{\tau,\varepsilon}$ and $\Psi_{\tau,\varepsilon}$ are Galois connections for $\varepsilon = \emptyset$.) We have already mentioned that the logical relation for global state satisfies the assumption about ground types that allows us to use it to prove instances of the contextual preorder (this is not difficult to show). We can therefore apply our reasoning principle to expressions e that do not use either of the operations `get` and `put`. (This restriction also applies to the free variables of e . Their types are annotated with the effect \emptyset in both the call-by-value and call-by-name effect systems, and hence we cannot bind them to expressions that have side-effects.)

If we have any program with a call-by-value subterm that does not have side-effects then the reasoning principle shows that we can replace the evaluation order of the subterm with call-by-name. For this example, the contextual preorder is symmetric, so we could also replace call-by-name with call-by-value: the reasoning principle shows they have identical behaviour. In both cases, the program itself may have side-effects. The restriction on side-effects applies only on the subterm. The corollary shows that if we have an entire program e that does not have any side-effects, then call-by-value and call-by-name are equivalent: we have $\langle e \rangle_\emptyset^v \equiv \langle e \rangle_\emptyset^n$. ◀

3.5 Related work

Comparing evaluation orders Plotkin [84] and many others relate call-by-value and call-by-name. Crucially, they consider lambda-calculi with no side-effects other than divergence. This makes a significant difference to the techniques that can be used, in particular because in this case the equational theory for call-by-name is strictly weaker than for call-by-value. This is not necessarily true for other side-effects. Other evaluation orders (such as call-by-need) have also been compared in similarly restricted settings [67, 70, 33].

It might also be possible to recast some of our work in terms of the *duality* between call-by-value and call-by-name [23, 18, 99].

Relating semantics of languages The technique we use here to relate call-by-value and call-by-name is based on the idea used first by Reynolds [89] to relate direct and continuation semantics of the lambda calculus, and later used by others (e.g. [75, 52, 16, 24]). There are several differences with our approach. Reynolds first constructs a logical relation *between the two semantics*, and uses this to establish a relationship between direct and continuation semantics using the two maps. Our logical relations do not relate call-by-value and call-by-name, they instead relate arbitrary GCBPV terms. Reynolds also relies on continuations with a large-enough domain of answers (e.g. a solution to a particular recursive domain equation) to construct the maps between the two semantics. In our case we do not need to impose such a restriction for the maps to exist.

There has been previous work [92, 54, 93] on soundness and completeness properties of translations (similar to the translations into GCBPV). Galois connections (and similar structures) in which the order is reduction of programs play a significant role in these. A key difference in our case is that we consider the observable behaviour of programs. We cannot consider reductions in our case, because moving from call-by-value to call-by-name can add, remove, and reorder reductions.

Axiomatic properties of side-effects In this chapter, we emphasize the use of axiomatic properties of side-effects [27] for formal reasoning about programs. Many of these properties have been used for this and similar purposes, and have been studied for particular side-effects [14, 42, 60]. We contribute a new use for them.

3.6 Summary

This chapter has two primary goals. The first is to show that we can use GCBPV for formal reasoning about effect-dependent program transformations, including those that involve source languages. We do this primarily by defining a notion of logical relation for GCBPV (Definition 3.1.1) and relating it to the contextual preorder (Lemma 3.1.5). We also describe a technique for constructing these logical relations (the free lifting in Section 3.1.1).

The second is to derive a reasoning principle (Theorem 3.4.2) that relates call-by-value and call-by-name. The principle shows that it is correct to replace call-by-value with call-by-name for subterms restricted to *thinkable* effects (Definition 3.3.2). It is about open expressions, and allows us to change evaluation order *within* programs. We obtain a result about call-by-value and call-by-name evaluations of *programs* as a corollary (Corollary 3.4.3). The reasoning principle is not restricted to a particular collection of side-effects; we instead identify the axiomatic property of side-effects (thinkable) that gives rise to a relationship between the two evaluation orders.

We expect that the technique we use can be applied to other evaluation orders. Two evaluation orders can be related by giving translations into some common intermediate language (here we use GCBPV), constructing maps between the two translations, and showing that (for some models) these maps form Galois connections. In Chapter 5 we give a specific example of a relationship between call-by-name and call-by-need (with nontermination as the only side-effect), but not a *general* reasoning principle.

In this chapter we work exclusively at the level of syntax. For example, the notion of logical relation we use relates pairs of GCBPV terms, and we define thinkable in terms of the inequational theory. The advantage of using syntax rather than a denotational semantics is that the syntax requires less machinery. However, it also has disadvantages. Sticking to syntax makes it more difficult to work with some side-effects (such as recursion) syntactically,

and many of the proofs (e.g. Theorem 3.3.3 and Lemma 3.4.1) are more complex. This partly motivates the next chapter, where we switch to semantics. We redevelop the relationship between call-by-value and call-by-name in terms of the semantics, and give more examples of side-effects that it can be applied to. At a high-level, the principle does not change much, and the technique used to derive it is the same. Both versions have merits.

Chapter 4

Noninvertible program transformations

Most of the previous work on proving the correctness of effect-dependent transformations has been restricted to cases where the transformations are *invertible* in the sense that both directions are valid, since they rely on the expressions having the same behaviour. However, there are many situations in which only one direction of a program transformation is valid. Consider the following two expressions:

let $x = e$ in	let $x = e$ in
let $y = e$ in	
(x, y)	(x, x)

In some cases these have the same behaviour, e.g. if the only side-effect of e is to write to some state that is not shared between threads. When this is true and the expression on the left appears inside some program, then it can safely be replaced with the expression on the right, perhaps as part of a compiler optimization.

Suppose that e above is allowed to write to state that is shared between threads. In this case, the two expressions are not equivalent, because some other thread may concurrently use the state between the two executions of e on the left. The observable behaviours of the expression on the right are an (in general proper) superset of those on the left. Replacing the left with the right would be a valid transformation, but the reverse is not. We cannot verify all effect-dependent transformations by proving only equivalences between expressions.

There are also other examples. Consider a language with boxing (i.e. conversion of primitive values into immutable references), and an equality operator on such references. We might wish to validate a transformation that merges boxes of the same value. This transformation would change uses of equality that could have evaluated to either true or false into ones that always give true. The reverse transformation is not valid because it could replace a program that is guaranteed to give true with one that might give false. Another example is undefined behaviour (e.g. in the C language). The C standard allows each expression that has undefined behaviour to be replaced with any other (even one that causes arbitrary effects), but the reverse is not true. Including undefined behaviour when considering program transformations is useful even if the source language does not have undefined behaviour [55].

These examples explain why we defined *inequational* theories in Section 2.7.2. This chapter develops machinery that we can use for reasoning about these *noninvertible* effect-dependent transformations.

Contributions We show that previous work on equational reasoning about effects can be adapted to *noninvertible* effect-dependent transformations:

- We describe a general *order-enriched* categorical semantics for GCBPV that allows us to state and prove noninvertible transformations semantically (Section 4.2).
- We *relate the syntax to the semantics* by describing how to prove adequacy of models of GCBPV (Section 4.3). We do this by developing an abstract notion of logical relation that generalizes the one given in Section 3.1.
- We apply the framework to examples by verifying effect-dependent transformations involving undefined behaviour, nondeterminism, and mutable state shared between threads. The examples are introduced in Section 4.1.
- We use the denotational semantics to redevelop our reasoning principle for relating call-by-value and call-by-name (Section 4.4).

Much of this chapter simply shows that techniques that have previously been applied to reasoning about effect-dependent transformations can be adapted to non-invertible cases, allowing support for previously neglected applications. To do this we have to redevelop some of the definitions that appear in the previous chapter. There we worked exclusively with syntax (in particular, the notion of logical relation we used relates terms to other terms). Many of the definitions we give here generalize those in the previous chapter.

4.1 Examples of noninvertible transformations

We first instantiate GCBPV (by choosing suitable inequational theories) to characterize three different collections of side-effects that we use as examples. For each example, we also state instances of the contextual preorder \leq_{ctx} (and its symmetric counterpart \cong_{ctx}) that represent effect-dependent program transformations. We defer the proofs that these instances hold to later sections of this chapter.

We use the following syntactic sugar for booleans (with eliminators for values and for computations), let bindings, and sequencing of computations throughout.

$$\begin{aligned}
\text{bool} &:= \text{unit} + \text{unit} & \text{true} &:= \text{inl } () & \text{false} &:= \text{inr } () \\
\text{if } V \text{ then } W_1 \text{ else } W_2 &:= \text{case } V \text{ of } \{\text{inl } x_1. W_1, \text{inr } x_2. W_2\} & (x_i \text{ not free in } W_i) \\
\underline{\text{if}} V \text{ then } M_1 \text{ else } M_2 &:= \underline{\text{case}} V \text{ of } \{\text{inl } x_1. M_1, \text{inr } x_2. M_2\} & (x_i \text{ not free in } M_i) \\
\text{let } x = V \text{ in } M &:= \langle V \rangle \text{ to } x. M \\
M; N &:= M \text{ to } y. N & (y \text{ not free in } N)
\end{aligned}$$

The syntactic sugar for let-bindings can be written equivalently using lambdas or substitution:

$$\text{let } x = V \text{ in } M \equiv V \text{ ' } \lambda x. M \equiv M[x \mapsto V]$$

4.1.1 Undefined behaviour

As a simple example, we consider C-style *undefined behaviour*. Computations with undefined behaviour are allowed to do anything (including arbitrary side-effects) at runtime. Here we consider undefined behaviour in isolation, without any other side-effects in the language. We contribute a denotational semantics for undefined behaviour.

Recall that to instantiate GCBPV we need to choose a signature (see Definition 2.7.1), which consists of an effect algebra, base types, constants, and operations (with coarities, arities and effects). For this example, the signature consists of the following data:

- The effect algebra is the trivial preordered monoid $\{\star\}$. (So we do not actually track any effect information, and none of the transformations we consider are effect-dependent. This example is intended to be as simple as possible.)
- The only base type is `int`, which contains signed 32-bit integers.
- We have a constant $\underline{n} : \text{int}$ for each n in the set $\llbracket \text{int} \rrbracket := \{-2^{31}, \dots, 2^{31} - 1\}$, and two constants

$$\text{geq} : \text{int} \times \text{int} \rightarrow \text{bool} \quad \text{add} : \text{int} \times \text{int} \rightarrow \text{int}$$

The first determines if its first argument is greater than or equal to its second argument, and the second adds two integers, with the result defined to wrap around on overflow.

- There are two operations. The first is `undef`, which has coarity `unit`, arity `empty` and effect \star . Hence `undef ()` is a computation of type $\langle \star \rangle \text{empty}$. Semantically, it has undefined behaviour (and the inequational theory we define below captures this property.) As syntactic sugar we have an undefined computation $\text{undef}_{\underline{C}}$ of type \underline{C} for every computation type \underline{C} :

$$\text{undef}_{\underline{C}} := \text{undef} () \text{ to } x. \text{case}_{\underline{C}} x \text{ of } \{ \}$$

(We have $\text{undef} () \equiv \text{undef}_{\langle \star \rangle \text{empty}}$.) The second operation is `addnsw` (no signed wrap), which adds two integers with undefined behaviour on overflow. Its coarity is `int × int`, its arity is `int`, and its effect is \star , so if $V_1, V_2 : \text{int}$ then `addnsw (V1, V2)` : $\langle \star \rangle \text{int}$ adds V_1 and V_2 .

As we explain in Section 2.7.2, we can give an inequational theory by specifying a list of signature axioms, which characterize the behaviour of the operations and constants.

The main axiom we have is that undefined behaviour can be replaced with any computation (for all \underline{C} and M):

$$\text{undef}_{\underline{C}} \leqslant M$$

(Recall our usual convention on typing of equations: the axiom holds only if both sides have the same type in the same typing context.)

We also have axioms that specify the behaviour of the three functions on constant integers. These are (for all $m, n \in \llbracket \text{int} \rrbracket$):

$$\begin{aligned} \text{add}(m, n) &\equiv \frac{m + n + 2^{32}}{2^{32}} && \text{if } (m + n) < -2^{31} \\ \text{add}(m, n) &\equiv m + n && \text{if } (m + n) \in \llbracket \text{int} \rrbracket \\ \text{add}(m, n) &\equiv \frac{m + n - 2^{32}}{2^{32}} && \text{if } (m + n) \geq 2^{31} \\ \text{geq}(m, n) &\equiv \text{true} && \text{if } m \geq n \\ \text{geq}(m, n) &\equiv \text{false} && \text{if } m < n \\ \text{add}_{\text{nsw}}(m, n) &\equiv \langle m + n \rangle && \text{if } (m + n) \in \llbracket \text{int} \rrbracket \\ \text{add}_{\text{nsw}}(m, n) &\equiv \text{undef}_{\text{int}} && \text{if } (m + n) \notin \llbracket \text{int} \rrbracket \end{aligned}$$

where $+$ is the usual addition of integers in \mathbb{Z} . (Recall that we write \equiv to mean both directions are taken as axioms.) This completes the definition of the inequational theory.

A consequence of these signature axioms is that if undefined behaviour occurs as part of a larger computation, then the whole computation is undefined:

$$\text{undef}_{\underline{C}} \text{ to } x. M \equiv \text{undef}_{\underline{D}}$$

We give some valid program transformations for the inequational theory for undefined behaviour. The first is replacing add_{nsw} with **add** inside programs:

$$\text{add}_{\text{nsw}}(v_1, v_2) \leq_{\text{ctx}} \langle \mathbf{add}(v_1, v_2) \rangle$$

Validity of this transformation implies it is correct to compile add_{nsw} as **add**. It is not invertible in general (since the reverse can introduce undefined behaviour).

A common optimization is to simplify comparisons of integers by using the fact that some operations on signed integers have undefined behaviour. One example is the following transformation:

$$\begin{array}{ccc} \text{add}_{\text{nsw}}(V_1, V_2) \text{ to } x. & & \text{add}_{\text{nsw}}(V_1, V_2) \text{ to } x. \\ \text{let } b = \text{geq}(x, V_1) \text{ in } & \leq_{\text{ctx}} & \text{let } b = \text{geq}(V_2, \underline{0}) \text{ in } \\ N & & N \end{array}$$

(This is especially useful when V_2 is a constant, because the comparison on the right can be evaluated statically.) Perhaps surprisingly, the reverse direction is also valid. This is the case because if the addition overflows, both terms have undefined behaviour, so the result of the comparison is irrelevant. It would not be correct to apply either direction of this transformation if **add** was used instead of add_{nsw} .

4.1.2 Nondeterminism

We also consider binary nondeterminism [11], so that programs choose from finite nonempty sets of results via a binary choice operation flip (that “flips a coin”). Unlike previous work (such as [11]), we allow nondeterministic choices to be made statically by program transformations, reducing the amount of runtime nondeterminism. These transformations are correct because they restrict the behaviour of programs; they do not add new behaviours. This is not the case for the reverse direction, and hence these transformations are not invertible. Finite nondeterminism is superficially similar to undefined behaviour, but not the same since there is no equation that corresponds to $\mathbf{undef}_{\underline{C}} \text{ to } x. M \equiv \mathbf{undef}_{\underline{D}}$.

The GCBPV signature for this example consists of the following data:

- The effect algebra tracks nondeterminism. It is the preordered monoid $\{1 \leq +\}$ with least upper bound as the multiplication and 1 as the unit. In this case, 1 means deterministic, and + means potentially nondeterministic.
- There are no base types or constants.
- There is a single operation flip with coarity **unit**, arity **bool** and effect +. The computation $\text{flip}()$ has type $\langle + \rangle \mathbf{bool}$, and nondeterministically chooses either **true** or **false**.

We use the operation flip to choose between arbitrary computations M_1, M_2 by defining

$$M_1 \text{ or } M_2 := \text{flip}() \text{ to } x. \mathbf{if } x \text{ then } M_1 \text{ else } M_2$$

If $\Gamma \vdash M_1 : \underline{C}$ and $\Gamma \vdash M_2 : \underline{C}$ then $\Gamma \vdash M_1 \text{ or } M_2 : \langle + \rangle \underline{C}$. The η -law for **bool** implies that $\text{flip}() \equiv \langle \mathbf{true} \rangle \text{ or } \langle \mathbf{false} \rangle$.

Again we specify the inequational theory by giving signature axioms. These are listed in Figure 4.1. The first is the crucial one: it allows us to replace $\text{flip}()$ with **true** or **false** statically (a similar axiom with **false** instead of **true** is derivable). The other direction of this axiom (with \geq rather than \leq) is not derivable, so we cannot replace constant booleans with $\text{flip}()$. The final three axioms in this group ensure that nondeterministic choice is idempotent, commutative and associative. Although these three axioms are non-symmetric (we only assume \leq rather than \equiv) the other direction of each is derivable.

$$\begin{aligned}
& \text{flip } () \preceq \text{coerce}_{1 \leq +} \langle \text{true} \rangle \\
& \text{coerce}_{1 \leq +} \langle () \rangle \preceq \text{flip } (); \langle () \rangle \\
& \text{flip } () \preceq \text{flip } () \text{ to } x. \langle \text{if } x \text{ then false else true} \rangle \\
& \left(\begin{array}{l} \text{flip } () \text{ to } x. \\ \text{if } x \text{ then flip } () \text{ to } y. \langle \text{inl}_{\text{unit}} () \rangle \\ \text{else coerce}_{1 \leq +} \langle \text{inr}_{\text{bool}} () \rangle \end{array} \right) \preceq \left(\begin{array}{l} \text{flip } () \text{ to } x. \\ \text{if } x \text{ then coerce}_{1 \leq +} \langle \text{inl}_{\text{unit}} \text{true} \rangle \\ \text{else flip } () \text{ to } y. \\ \langle \text{if } y \text{ then inl}_{\text{unit}} \text{false else inr}_{\text{bool}} () \rangle \end{array} \right)
\end{aligned}$$

Figure 4.1: Signature axioms for nondeterminism

Lemma 4.1.1

1. If $\Gamma \vdash M_1 : \underline{C}$ and $\Gamma \vdash M_2 : \underline{C}$ then $M_1 \text{ or } M_2 \equiv M_2 \text{ or } M_1$.
2. If $\Gamma \vdash M : \underline{C}$ then $M \text{ or } M \equiv \text{coerce}_{\underline{C} < \langle + \rangle \underline{C}} M$.
3. If $\Gamma \vdash M_1 : \underline{C}$ and $\Gamma \vdash M_2 : \underline{C}$ then

$$M_1 \text{ or } M_2 \preceq \text{coerce}_{\underline{C} < \langle + \rangle \underline{C}} M_1 \quad M_1 \text{ or } M_2 \preceq \text{coerce}_{\underline{C} < \langle + \rangle \underline{C}} M_2$$

4. If $\Gamma \vdash M_1 : \underline{C}$, $\Gamma \vdash M_2 : \underline{C}$ and $\Gamma \vdash M_3 : \underline{C}$ then $(M_1 \text{ or } M_2) \text{ or } M_3 \equiv M_1 \text{ or } (M_2 \text{ or } M_3)$.
5. If $\Gamma \vdash M_1 : \langle \varepsilon \rangle A$, $\Gamma \vdash M_2 : \langle \varepsilon \rangle A$ and $\Gamma, x : A \vdash N : \underline{C}$ then

$$(M_1 \text{ or } M_2) \text{ to } x. N \equiv (M_1 \text{ to } x. N) \text{ or } (M_2 \text{ to } x. N) \quad \blacktriangleleft$$

This completes the inequational theory we use for our nondeterminism example.

We consider a transformation that reuses the result of a duplicated computation:

$$M \text{ to } x. M \text{ to } y. N \preceq_{\text{ctx}} M \text{ to } x. N[y \mapsto x]$$

Validity of this transformation is shown in later sections of this chapter. Benton et al. [11] consider a similar transformation when M is deterministic; in this case, the reverse direction also holds (if M has effect 1, then \cong_{ctx} holds). Their equational framework cannot validate this duplicated computation transformation when nondeterminism is allowed because it is not invertible (the right-to-left direction is not correct in general). Our framework can. We can also validate other transformations involving nondeterminism, for example, dead code elimination:

$$M; N \cong_{\text{ctx}} N$$

4.1.3 Shared global state

We also consider a concurrency example. We again use mutable global state as a source of side-effects, but in this case assume the state can be accessed by multiple threads, which requires us to use directed versions of the signature axioms for global state. We do not attempt to fully characterize a concurrent language in this example. Instead, we focus only on a single thread (assuming there might be other threads that can arbitrarily access and mutate the state), and do not include any concurrency primitives (e.g. creation of new threads). The idea is that we can validate transformations on subprograms that use the shared state but do not use any concurrency primitives. For simplicity, the state is the only source of side-effects in this example.

The signature for this example consists of the following data:

$$\begin{aligned}
\text{get } () \text{ to } x. (\text{put } x; \langle x \rangle) &\leq \mathbf{coerce}_{\{\text{get}\} \leq \{\text{get}, \text{put}\}} (\text{get } ()) \\
\text{put } V; \text{get } () &\leq \text{put } V; \mathbf{coerce}_{\emptyset \leq \{\text{get}\}} \langle V \rangle \\
\text{put } V_1; \text{put } V_2 &\leq \text{put } V_2 \\
\text{get } () \text{ to } x. \text{get } () \text{ to } y. \langle (x, y) \rangle &\leq \text{get } () \text{ to } x. \langle (x, x) \rangle \\
\text{get } (); \langle () \rangle &\equiv \mathbf{coerce}_{\emptyset \leq \{\text{get}\}} \langle () \rangle
\end{aligned}$$

Figure 4.2: Signature axioms for shared global state

- We use a Gifford-style effect algebra (Example 2.1.2) with set of operations $\Sigma := \{\text{get}, \text{put}\}$, so that effects $\varepsilon \subseteq \Sigma$ specify the operations that may be used.
- There are no base types or constants.
- There are two operations:

Operation op	Coarity car_{op}	Arity ar_{op}	Effect eff_{op}
get	unit	bool	$\{\text{get}\}$
put	bool	unit	$\{\text{put}\}$

The signature axioms of the inequational theory are given in Figure 4.2. The axioms are the same as those for non-shared global state (Figure 2.14), except that four of them are directed. The reverse directions introduce data races if another thread is accessing the state concurrently and hence they are not included.

We validate similar program transformations to the previous example: subject to suitable restrictions on effects, it is sound to reuse the result of a duplicated computation, and eliminate dead computations. Specifically, if M has type $\langle \varepsilon \rangle A$ where $\varepsilon \subseteq \{\text{put}\}$, and M' has type $\langle \varepsilon' \rangle A'$ where $\varepsilon' \subseteq \{\text{get}\}$, then the following hold:

$$M \text{ to } x. M \text{ to } y. N \leq_{\text{ctx}} M \text{ to } x. N[y \mapsto x] \qquad M'; N \cong_{\text{ctx}} N$$

The instance on the left does not hold when the effect of M is $\{\text{get}\}$, because we do not allow reordering of consecutive uses of `get`. We could instead capture a weaker memory model by including a signature axiom that allows this reordering; in this case, the program transformation would be valid. The instance on the left is also not symmetric (so the transformation is not invertible). This is unlike non-shared state, in which it is symmetric.

4.2 Order-enriched semantics of GCBPV

The goal of this chapter is to prove instances of contextual preorders \leq_{ctx} . It is difficult to prove these directly so, as in the previous chapter, we develop some extra machinery. Unlike the previous chapter, here we focus on denotational semantics. This is primarily because denotational models are easier to reason about than syntactic logical relations.

A common way of proving contextual equivalences is to give an adequate model of the chosen side-effects, so that it suffices to show that denotations of terms are equal in the model. In our case, \leq_{ctx} is not symmetric, so equality between denotations is not suitable. We therefore use *order-enriched* models, which come with partial orders \sqsubseteq between denotations.

We have three examples, which need to be modelled in different settings. There are also other examples of side-effects (e.g. local state [88] and probability [37]) that require other

settings. We aim to provide a general *categorical* semantics for GCBPV that can be applied to as many of these as possible. Our semantics is based on *adjunction models* for CBPV [58], which we extend with grading and with order-enrichment.¹ The order \sqsubseteq on morphisms is the semantic counterpart of the contextual preorder.

We use some order-enriched category theory to describe the semantics (see Appendix A for background).

4.2.1 Graded adjunctions

The data required for a model of GCBPV closely mirrors the syntax. There are two kinds of types, and therefore we use two **Poset**-categories: a *value category* \mathbf{C} and a *computation category* \mathbf{D} . Value types are interpreted as objects of \mathbf{C} , and computation types are interpreted as objects of \mathbf{D} . We use X, Y, \dots to refer to objects of \mathbf{C} and $\underline{X}, \underline{Y}, \dots$ for objects of \mathbf{D} .

The key component in each model is the data required to interpret thunk types and returner types. In ordinary CBPV, there is a single type former F for returner types (because there is no grading), and in the semantics this is left adjoint to the type former U for thunk types:

$$\begin{array}{ccc} & F & \\ \mathbf{C} & \begin{array}{c} \curvearrowright \\ \perp \\ \curvearrowleft \end{array} & \mathbf{D} \\ & U & \end{array}$$

(Of course, we actually need a *strong* adjunction; we add strength below.) We have an adjunction because for each value of type A we can form a computation of type $F A$ (by returning it), and any computation $\Gamma, x : A \vdash M : \underline{C}$ can be *extended* to a computation

$$\Gamma, y : U F A \vdash \text{force } y \text{ to } x. M : \underline{C}$$

that satisfies certain properties. To add grading, we note that this is still the case with the returner type $\langle 1 \rangle A$ (i.e. with the unit effect). We therefore use the left adjoint F as the interpretation of $\langle 1 \rangle A$. For returner types annotated with other effects, we note that they can be written using $\langle 1 \rangle$ and the *action* of the effect algebra on computation types (see Definition 2.7.3), as $\langle \varepsilon \rangle A = \langle \varepsilon \rangle \langle 1 \rangle A$. Hence the remaining structure we need to interpret the other returner types is an action $\varepsilon \otimes -$ of the preordered monoid of effects on the computation category \mathbf{D} .

Definition 4.2.1 Suppose that $(\mathcal{E}, \leq, \cdot, 1)$ is a preordered monoid, and \mathbf{D} is a **Poset**-category. A *strict Poset-action* of \mathcal{E} on \mathbf{D} is a **Poset**-functor $\otimes : \mathcal{E} \times \mathbf{D} \rightarrow \mathbf{D}$ that respects the unit and multiplication:

$$1 \otimes - = Id_{\mathbf{D}} \quad (\varepsilon \cdot \varepsilon') \otimes - = \varepsilon \otimes (\varepsilon' \otimes -) \quad \blacktriangleleft$$

Here, and throughout this chapter, we treat the preordered monoid as a **Poset**-category with set of objects \mathcal{E} , and a single morphism from ε to ε' if $\varepsilon \leq \varepsilon'$ (we write this morphism as $\varepsilon \leq \varepsilon'$). The ordering on morphisms is equality.

Each strict **Poset**-action comes in particular with morphisms $(\varepsilon \leq \varepsilon') \otimes \underline{X}$ from $\varepsilon \otimes \underline{X}$ to $\varepsilon' \otimes \underline{X}$ whenever $\varepsilon \leq \varepsilon'$. This corresponds to the subtyping $\langle \varepsilon \rangle \underline{C} <: \langle \varepsilon' \rangle \underline{C}$.

We call the combination of an adjunction and an action a *graded Poset-adjunction*. This definition is implicit in [28] (which considers *resolutions* of graded monads).

Definition 4.2.2 Suppose that $(\mathcal{E}, \leq, \cdot, 1)$ is a preordered monoid, and that \mathbf{C} and \mathbf{D} are **Poset**-categories. An $(\mathcal{E}$ -)graded **Poset-adjunction** is a triple (F, U, \otimes) where:

¹Though unlike Levy, we do not define adjunction models in terms of *locally indexed* categories, instead we use more a more elementary (but less elegant) description.

- $F : \mathbf{C} \rightarrow \mathbf{D}$ and $U : \mathbf{D} \rightarrow \mathbf{C}$ form a **Poset**-adjunction $F \dashv U$, with unit $\eta : Id_{\mathbf{C}} \rightarrow U \circ F$ and counit $\varpi : F \circ U \rightarrow Id_{\mathbf{D}}$.
- $\otimes : \mathcal{E} \times \mathbf{D} \rightarrow \mathbf{D}$ is a strict \mathcal{E} -action on \mathbf{D} . ◀

The counit induces an extension operator: every \mathbf{C} -morphism $f : X \rightarrow U\underline{Y}$ can be extended to a \mathbf{C} -morphism

$$U(\varepsilon \otimes FX) \xrightarrow{U(\varepsilon \otimes Ff)} U(\varepsilon \otimes F(U\underline{Y})) \xrightarrow{U(\varepsilon \otimes \varpi)} U\underline{Y}$$

This is similar to the extension of computations in context $\Gamma, x : A$ to computations in context $\Gamma, x : \mathbf{U}FA$ mentioned above, except for the extra context Γ . This extension operator is enough to interpret *closed* computations M to $x.N$, but because there is no Γ , it cannot interpret this computation if it is open. We require a more general extension operator than this, and to define it we need a strong adjunction. (Strength appears for exactly the same reason as it does for models of the monadic metalanguage [78].)

Definition 4.2.3 (Graded strong Poset-adjunction) Suppose that $(\mathcal{E}, \leq, \cdot, 1)$ is a preordered monoid, \mathbf{C} is a cartesian **Poset**-category, and \mathbf{D} is a **Poset**-category. An *graded strong Poset-adjunction* consists of a graded **Poset**-adjunction (F, U, \otimes) and a family of morphisms

$$str_{\varepsilon, X, Y} : X \times U(\varepsilon \otimes FY) \rightarrow U(\varepsilon \otimes F(X \times Y))$$

indexed by $\varepsilon \in \mathcal{E}$ and $X, Y \in \mathbf{C}$, such that:

- For each ε , the pair $(U(\varepsilon \otimes F-), str_{\varepsilon})$ is a strong **Poset**-functor $\mathbf{C} \rightarrow \mathbf{C}$.
- The unit η is a strong natural transformation $Id_{\mathbf{C}} \rightarrow (U \circ F, str_1)$.
- For each $\varepsilon, \varepsilon'$ the family of morphisms $U(\varepsilon \otimes \varpi_{(\varepsilon' \otimes F-)})$ is a strong natural transformation

$$(U(\varepsilon \otimes F-), str_{\varepsilon}) \circ (U(\varepsilon' \otimes F-), str_{\varepsilon'}) \rightarrow (U((\varepsilon \cdot \varepsilon') \otimes F-), str_{\varepsilon \cdot \varepsilon'})$$

- The strength is natural in ε : if $\varepsilon \leq \varepsilon'$ then

$$\begin{array}{ccc} X \times U(\varepsilon \otimes FY) & \xrightarrow{str_{\varepsilon, X, Y}} & U(\varepsilon \otimes F(X \times Y)) \\ \downarrow X \times U((\varepsilon \leq \varepsilon') \otimes FY) & & \downarrow U((\varepsilon \leq \varepsilon') \otimes F(X \times Y)) \\ X \times U(\varepsilon' \otimes FY) & \xrightarrow{str_{\varepsilon', X, Y}} & U(\varepsilon' \otimes F(X \times Y)) \end{array}$$

commutes for all X, Y . ◀

The strength enables us to define a more general version of the extension operator above. Given a \mathbf{C} -morphism $f : Z \times X \rightarrow U\underline{Y}$, define its extension $f^{\dagger} : Z \times U(\varepsilon \otimes FX) \rightarrow U(\varepsilon \otimes \underline{Y})$ as

$$Z \times U(\varepsilon \otimes FX) \xrightarrow{str_{\varepsilon, Z, X}} U(\varepsilon \otimes F(Z \times X)) \xrightarrow{U(\varepsilon \otimes Ff)} U(\varepsilon \otimes F(U\underline{Y})) \xrightarrow{U(\varepsilon \otimes \varpi)} U(\varepsilon \otimes \underline{Y})$$

In the semantics, we use this extension operator to interpret open computations M to $x.N$ in typing context Γ . The object Z is the interpretation of Γ .

Each graded **Poset**-adjunction with $\mathbf{C} = \mathbf{Set}$ or $\mathbf{C} = \mathbf{Poset}$ forms a graded strong **Poset**-adjunction in exactly one way. The strength in both of these cases is given by:

$$str_{\varepsilon, X, Y}(x, t) = U(\varepsilon \otimes (F(y \mapsto (x, y)))) t$$

For $\mathbf{C} = \omega\mathbf{Cpo}$, if a strength exists it is also given by this formula, and this formula defines a strength when the functions $U(\varepsilon \otimes F-) : \mathbf{C}(X, Y) \rightarrow \mathbf{D}(U(\varepsilon \otimes FX), U(\varepsilon \otimes FY))$ preserve least upper bounds.² These facts are useful when defining models based on **Set**, **Poset** or $\omega\mathbf{Cpo}$.

²This characterization of strengths for **Set**, **Poset** and $\omega\mathbf{Cpo}$ follows from the fact that tensorial strengths correspond to enrichments of functors [50]. One can also show that if \mathbf{C} is *well-pointed*, then strengths are uniquely determined by the formula above.

4.2.2 Models of GCBPV

We are almost ready to give the definition of GCBPV *model*. We do this formally below (Definition 4.2.7), but first give an informal description. In addition to the graded adjunction, we need data to interpret the rest of the type formers. On the value level we have product and sum types, so we ask for the value category \mathbf{C} to be bicartesian, and require distributivity. For computations we have product types, and so we ask for the computation category \mathbf{D} to be cartesian. Since for types we have $\langle\!\langle \varepsilon \rangle\!\rangle \underline{\mathbf{unit}} = \underline{\mathbf{unit}}$ and $\langle\!\langle \varepsilon \rangle\!\rangle (\underline{C}_1 \times \underline{C}_2) = \langle\!\langle \varepsilon \rangle\!\rangle \underline{C}_1 \times \langle\!\langle \varepsilon \rangle\!\rangle \underline{C}_2$ we also require the action \otimes to *strictly preserve* the cartesian structure of \mathbf{D} . In particular, we require

$$\varepsilon \otimes \underline{1} = \underline{1} \quad \varepsilon \otimes (\underline{X}_1 \times \underline{X}_2) = (\varepsilon \otimes \underline{X}_1) \times (\varepsilon \otimes \underline{X}_2)$$

where we write the cartesian structure of \mathbf{D} with an underline.

Functions in GCBPV map values to computations, and are computations themselves. To interpret these we use a **Poset**-functor $X \Rightarrow - : \mathbf{D} \rightarrow \mathbf{D}$ for each object $X \in \mathbf{C}$, with *currying* functions Λ and *uncurrying* functions Λ^{-1} :

$$\Lambda : \mathbf{C}(Z \times X, U\underline{Y}) \cong \mathbf{C}(Z, U(X \Rightarrow \underline{Y})) : \Lambda^{-1}$$

As for products, we require exponentials to strictly preserve the action, in particular,

$$\varepsilon \otimes (X \Rightarrow \underline{Y}) = X \Rightarrow (\varepsilon \otimes \underline{Y})$$

We also impose an additional requirement on the exponentials $X \Rightarrow -$. To explain this we define a notion of *linearity*³ for morphisms of \mathbf{C} . The definition uses the *right* strength str^r (which is the *left* strength str with the products reversed):

$$str_{\varepsilon, X, Y}^r := U(\varepsilon \otimes F\langle \pi_2, \pi_1 \rangle) \circ str_{\varepsilon, Y, X} \circ \langle \pi_2, \pi_1 \rangle : U(\varepsilon \otimes FX) \times Y \rightarrow U(\varepsilon \otimes F(X \times Y))$$

Definition 4.2.4 A natural transformation $\alpha_\varepsilon : U(\varepsilon \otimes \underline{X}) \rightarrow U(\varepsilon \otimes \underline{Y})$ is *linear* if the following diagram commutes

$$\begin{array}{ccc} U(\varepsilon \otimes FU(\varepsilon' \otimes \underline{X})) & \xrightarrow{U(\varepsilon \otimes F\alpha_{\varepsilon'})} & U(\varepsilon \otimes FU(\varepsilon' \otimes \underline{Y})) \\ U(\varepsilon \otimes \omega_{\varepsilon' \otimes \underline{X}}) \downarrow & & \downarrow U(\varepsilon \otimes \omega_{\varepsilon' \otimes \underline{Y}}) \\ U((\varepsilon \cdot \varepsilon') \otimes \underline{X}) & \xrightarrow{\alpha_{\varepsilon \cdot \varepsilon'}} & U((\varepsilon \cdot \varepsilon') \otimes \underline{Y}) \end{array}$$

for all $\varepsilon, \varepsilon' \in \mathcal{E}$.

A natural transformation $\alpha_\varepsilon : U(\varepsilon \otimes \underline{X}) \times Z \rightarrow U(\varepsilon \otimes \underline{Y})$ is *left-linear* if the following diagram commutes

$$\begin{array}{ccc} U(\varepsilon \otimes FU(\varepsilon' \otimes \underline{X})) \times Z & \xrightarrow{str^r} & U(\varepsilon \otimes F(U(\varepsilon' \otimes \underline{X}) \times Z)) & \xrightarrow{U(\varepsilon \otimes F\alpha_{\varepsilon'})} & U(\varepsilon \otimes FU(\varepsilon' \otimes \underline{Y})) \\ U(\varepsilon \otimes \omega_{\varepsilon' \otimes \underline{X}}) \times Z \downarrow & & & & \downarrow U(\varepsilon \otimes \omega_{\varepsilon' \otimes \underline{Y}}) \\ U((\varepsilon \cdot \varepsilon') \otimes \underline{X}) \times Z & \xrightarrow{\alpha_{\varepsilon \cdot \varepsilon'}} & & & U((\varepsilon \cdot \varepsilon') \otimes \underline{Y}) \end{array}$$

for all $\varepsilon, \varepsilon' \in \mathcal{E}$. ◀

³The word *linear* comes from Kock [49]. A similar property for syntax (see Section 4.2.3.5) is also called *linearity* by Munch-Maccagnoni [79] (and by Levy [60] for call-by-push-value).

$\llbracket A \rrbracket \in \mathbf{C}$	$\llbracket \underline{C} \rrbracket \in \mathbf{D}$	$\llbracket \Gamma \rrbracket \in \mathbf{C}$
$\llbracket \mathbf{unit} \rrbracket := 1$	$\llbracket \underline{\mathbf{unit}} \rrbracket := \underline{1}$	$\llbracket \diamond \rrbracket := 1$
$\llbracket A_1 \times A_2 \rrbracket := \llbracket A_1 \rrbracket \times \llbracket A_2 \rrbracket$	$\llbracket \underline{C}_1 \times \underline{C}_2 \rrbracket := \llbracket \underline{C}_1 \rrbracket \times \llbracket \underline{C}_2 \rrbracket$	$\llbracket \Gamma, x : A \rrbracket := \llbracket \Gamma \rrbracket \times \llbracket A \rrbracket$
$\llbracket \mathbf{empty} \rrbracket := 0$	$\llbracket A \rightarrow \underline{C} \rrbracket := \llbracket A \rrbracket \Rightarrow \llbracket \underline{C} \rrbracket$	
$\llbracket A_1 + A_2 \rrbracket := \llbracket A_1 \rrbracket + \llbracket A_2 \rrbracket$	$\llbracket \langle \varepsilon \rangle A \rrbracket := \varepsilon \otimes F \llbracket A \rrbracket$	
$\llbracket \mathbf{UC} \rrbracket := U \llbracket \underline{C} \rrbracket$		
$\llbracket \Gamma \vdash V : A \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket$	$\llbracket \Gamma \vdash M : \underline{C} \rrbracket : \llbracket \Gamma \rrbracket \rightarrow U \llbracket \underline{C} \rrbracket$	
$\llbracket c \rrbracket := \llbracket c \rrbracket \circ \langle \rangle_{\llbracket \Gamma \rrbracket}$	$\llbracket \lambda \{ \} \rrbracket := \cong \circ \langle \rangle_{\llbracket \Gamma \rrbracket}$	
$\llbracket x \rrbracket := \pi_x$	$\llbracket \lambda \{ 1. M_1, 2. M_2 \} \rrbracket := \cong \circ \langle \llbracket M_1 \rrbracket, \llbracket M_2 \rrbracket \rangle$	
$\llbracket () \rrbracket := \langle \rangle_{\llbracket \Gamma \rrbracket}$	$\llbracket i' M \rrbracket := U \pi_i \circ \llbracket M \rrbracket \quad (i \in \{1, 2\})$	
$\llbracket (V_1, V_2) \rrbracket := \langle \llbracket V_1 \rrbracket, \llbracket V_2 \rrbracket \rangle$	$\llbracket \lambda x : A. M \rrbracket := \Lambda \llbracket M \rrbracket$	
$\llbracket \mathbf{fst} V \rrbracket := \pi_1 \circ \llbracket V \rrbracket$	$\llbracket V' M \rrbracket := \Lambda^{-1} \llbracket M \rrbracket \circ \langle id, \llbracket V \rrbracket \rangle$	
$\llbracket \mathbf{snd} V \rrbracket := \pi_2 \circ \llbracket V \rrbracket$	$\llbracket \mathbf{op} V \rrbracket := \llbracket \mathbf{op} \rrbracket \circ \llbracket V \rrbracket$	
$\llbracket \mathbf{case}_A V \mathbf{of} \{ \} \rrbracket := \llbracket \cdot \rrbracket_{\llbracket A \rrbracket} \circ \llbracket V \rrbracket$	$\llbracket \langle V \rangle \rrbracket := \eta \circ \llbracket V \rrbracket$	
$\llbracket \mathbf{inl} V \rrbracket := inl \circ \llbracket V \rrbracket$	$\llbracket M \mathbf{to} x. N \rrbracket := \llbracket N \rrbracket^\dagger \circ \langle id, \llbracket M \rrbracket \rangle$	
$\llbracket \mathbf{inr} V \rrbracket := inr \circ \llbracket V \rrbracket$	$\llbracket \mathbf{coerce}_{\varepsilon \leq \varepsilon'} M \rrbracket := U((\varepsilon \leq \varepsilon') \otimes F \llbracket A \rrbracket) \circ \llbracket M \rrbracket$	
$\llbracket \mathbf{case} V \mathbf{of} \left. \begin{array}{l} \{ \mathbf{inl} x_1. W_1 \\ , \mathbf{inr} x_2. W_2 \} \end{array} \right\} \rrbracket := \left(\begin{array}{l} \llbracket W_1 \rrbracket, \llbracket W_2 \rrbracket \\ \circ dist \\ \circ \langle id, \llbracket V \rrbracket \rangle \end{array} \right)$	$\llbracket \mathbf{force} V \rrbracket := \llbracket V \rrbracket$	
$\llbracket \mathbf{thunk} M \rrbracket := \llbracket M \rrbracket$		

Figure 4.3: Denotational semantics of GCBPV

The intuition is that (left-)linearity means α uses the $U(\varepsilon \otimes X)$ argument exactly once, and does not add any other side-effects.

Using the data for exponentials above we define a natural transformation ev for function application:

$$ev_{\varepsilon, X, Y} := \Lambda^{-1} id : U(\varepsilon \otimes (X \Rightarrow Y)) \times X \rightarrow U(\varepsilon \otimes Y)$$

The extra requirement is that this natural transformation is left-linear. This is equivalent to requiring $\Lambda^{-1} \alpha_\varepsilon$ to be left-linear for all linear α .

The remaining data required to interpret GCBPV is an interpretation $\llbracket b \rrbracket \in \mathbf{C}$ of each base type b , an interpretation $\llbracket c \rrbracket : 1 \rightarrow \llbracket A \rrbracket$ of each constant c of type A , and an interpretation $\llbracket \mathbf{op} \rrbracket : \mathbf{car}_{\mathbf{op}} \rightarrow U(\mathbf{eff}_{\mathbf{op}} \otimes \mathbf{Far}_{\mathbf{op}})$ of each operation \mathbf{op} .

Using this data, we give the denotational semantics of GCBPV in Figure 4.3. Each value type A is interpreted as an object $\llbracket A \rrbracket$ in the value category \mathbf{C} , and each computation type \underline{C} is interpreted as an object $\llbracket \underline{C} \rrbracket$ of the computation category \mathbf{D} . The strict preservation requirements we impose imply that $\llbracket \langle \varepsilon \rangle \underline{C} \rrbracket = \varepsilon \otimes \llbracket \underline{C} \rrbracket$, and so $\llbracket M \mathbf{to} x. N \rrbracket$ is well-defined. Typing contexts Γ are interpreted as objects $\llbracket \Gamma \rrbracket \in \mathbf{C}$ using the terminal object and binary products. If $(x : A) \in \Gamma$ then we write $\pi_x : \llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket$ for the corresponding projection.

The interpretation of terms is given on the bottom of Figure 4.3. Values $\Gamma \vdash V : A$ are interpreted as morphisms $\llbracket \Gamma \vdash V : A \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket$ in the value category \mathbf{C} . Computations $\Gamma \vdash M : \underline{C}$ are also interpreted in the value category, as morphisms $\llbracket \Gamma \vdash M : \underline{C} \rrbracket : \llbracket \Gamma \rrbracket \rightarrow U\llbracket \underline{C} \rrbracket$. (Intuitively, we interpret computations as morphisms of \mathbf{C} because morphisms in the computation category \mathbf{D} are effect-preserving, while computations are not. The interpretation of a computation M is also not in general linear.) We often omit the typing context and type from denotations of terms. The two isomorphisms \cong in Figure 4.3 are the inverses of the following morphisms:

$$\langle \rangle_{U\underline{1}} : U\underline{1} \rightarrow 1 \quad \langle U\pi_1, U\pi_2 \rangle : U(\underline{X}_1 \times \underline{X}_2) \rightarrow U\underline{X}_1 \times U\underline{X}_2$$

These inverses exist because U is a right adjoint. In adjunction models, **think** and **force** are invisible because computations are interpreted in \mathbf{C} . The interpretation of **to** uses the general extension operator $(-)^{\dagger}$ defined in Section 4.2.1. Computations of returner type are coerced using $(\varepsilon \leq \varepsilon')^{\otimes}(-)$.

We collect together all of the data required for the semantics into the notion of GCBPV *structure*.

Definition 4.2.5 (GCBPV structure) Given some GCBPV signature, a *structure* consists of

- A distributive **Poset**-category \mathbf{C} and a cartesian **Poset**-category \mathbf{D} .
- A **Poset**-functor $X \Rightarrow - : \mathbf{D} \rightarrow \mathbf{D}$ for each $X \in \mathbf{C}$, together with a family of bijections

$$\Lambda : \mathbf{C}(Z \times X, U\underline{Y}) \cong \mathbf{C}(Z, U(X \Rightarrow \underline{Y})) : \Lambda^{-1}$$

natural in $Z \in \mathbf{C}$ and $\underline{Y} \in \mathbf{D}$.

- A graded strong **Poset**-adjunction (F, U, \otimes) , where $F : \mathbf{C} \rightarrow \mathbf{D}$ and $U : \mathbf{D} \rightarrow \mathbf{C}$, such that

$$\begin{aligned} \varepsilon^{\otimes} \underline{1} &= \underline{1} & \langle \varepsilon^{\otimes} \pi_1, \varepsilon^{\otimes} \pi_2 \rangle &= id_{\varepsilon^{\otimes}(\underline{X}_1 \times \underline{X}_2)} \\ \varepsilon^{\otimes}(\underline{X}_1 \times \underline{X}_2) &= (\varepsilon^{\otimes} \underline{X}_1) \times (\varepsilon^{\otimes} \underline{X}_2) & ev & \text{ is left-linear} \\ \varepsilon^{\otimes}(X \Rightarrow \underline{Y}) &= X \Rightarrow (\varepsilon^{\otimes} \underline{Y}) \end{aligned}$$

- An object $\llbracket b \rrbracket \in \mathbf{C}$ for each base type b .
- A morphism $\llbracket c \rrbracket : 1 \rightarrow \llbracket A \rrbracket$ for each constant $c \in \mathcal{K}_A$, and a morphism $\llbracket \text{op} \rrbracket : \llbracket \text{car}_{\text{op}} \rrbracket \rightarrow U(\text{eff}_{\text{op}}^{\otimes} F \llbracket \text{ar}_{\text{op}} \rrbracket)$ for each operation $\text{op} \in \Sigma$. \blacktriangleleft

Given any such structure, the denotational semantics defined in Figure 4.3 satisfies some useful properties. Recall that by definition, each GCBPV inequational theory is required to satisfy congruence and substitution properties, and also to respect the core axioms given in Figure 2.13. The denotational semantics satisfies similar properties. In the following lemma, we interpret well-typed substitutions $\Gamma \vdash \sigma : \Delta$, as \mathbf{C} -morphisms $\llbracket \Gamma \vdash \sigma : \Delta \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket \Delta \rrbracket$, given by tupling their components.

Lemma 4.2.6 Given any GCBPV structure, the following properties hold:

1. Compositionality:

- If $\llbracket \Gamma \vdash V : A \rrbracket \sqsubseteq \llbracket \Gamma \vdash W : A \rrbracket$ then for term contexts with hole \square ,

$$\begin{aligned} \Gamma' \vdash C[V] : B \wedge \Gamma' \vdash C[W] : B &\Rightarrow \llbracket \Gamma' \vdash C[V] : B \rrbracket \sqsubseteq \llbracket \Gamma' \vdash C[W] : B \rrbracket \\ \Gamma' \vdash \underline{C}[V] : \underline{D} \wedge \Gamma' \vdash \underline{C}[W] : \underline{D} &\Rightarrow \llbracket \Gamma' \vdash \underline{C}[V] : \underline{D} \rrbracket \sqsubseteq \llbracket \Gamma' \vdash \underline{C}[W] : \underline{D} \rrbracket \end{aligned}$$

- If $\llbracket \Gamma \vdash M : \underline{C} \rrbracket \sqsubseteq \llbracket \Gamma \vdash N : \underline{C} \rrbracket$ then for term contexts with hole \square ,

$$\begin{aligned} \Gamma' \vdash C[M] : B \wedge \Gamma' \vdash C[N] : B &\Rightarrow \llbracket \Gamma' \vdash C[M] : B \rrbracket \sqsubseteq \llbracket \Gamma' \vdash C[N] : B \rrbracket \\ \Gamma' \vdash \underline{C}[M] : \underline{D} \wedge \Gamma' \vdash \underline{C}[N] : \underline{D} &\Rightarrow \llbracket \Gamma' \vdash \underline{C}[M] : \underline{D} \rrbracket \sqsubseteq \llbracket \Gamma' \vdash \underline{C}[N] : \underline{D} \rrbracket \end{aligned}$$

2. Substitution: if $\Gamma \vdash \sigma : \Delta$ then:

$$\begin{aligned} \Delta \vdash V : A &\Rightarrow \llbracket \Gamma \vdash V[\sigma] : A \rrbracket = \llbracket \Delta \vdash V : A \rrbracket \circ \llbracket \Gamma \vdash \sigma : \Delta \rrbracket \\ \Delta \vdash M : \underline{C} &\Rightarrow \llbracket \Gamma \vdash M[\sigma] : \underline{C} \rrbracket = \llbracket \Delta \vdash M : \underline{C} \rrbracket \circ \llbracket \Gamma \vdash \sigma : \Delta \rrbracket \end{aligned}$$

Hence if $\llbracket \Gamma \vdash \sigma : \Delta \rrbracket \sqsubseteq \llbracket \Gamma \vdash \sigma' : \Delta \rrbracket$ then

$$\begin{aligned} \llbracket \Delta \vdash V : A \rrbracket \sqsubseteq \llbracket \Delta \vdash W : A \rrbracket &\Rightarrow \llbracket \Gamma \vdash V[\sigma] : A \rrbracket \sqsubseteq \llbracket \Gamma \vdash W[\sigma'] : A \rrbracket \\ \llbracket \Delta \vdash M : \underline{C} \rrbracket \sqsubseteq \llbracket \Delta \vdash N : \underline{C} \rrbracket &\Rightarrow \llbracket \Gamma \vdash M[\sigma] : \underline{C} \rrbracket \sqsubseteq \llbracket \Gamma \vdash N[\sigma'] : \underline{C} \rrbracket \end{aligned}$$

3. Core axioms:

- Values: if $\Gamma \vdash V : A$ and $\Gamma \vdash W : A$, and $V \equiv W$ is an instance of an axiom in Figure 2.13, then $\llbracket \Gamma \vdash V : A \rrbracket = \llbracket \Gamma \vdash W : A \rrbracket$.
- Computations: if $\Gamma \vdash M : \underline{C}$ and $\Gamma \vdash N : \underline{C}$, and $M \equiv N$ is an instance of an axiom in Figure 2.13, then $\llbracket \Gamma \vdash M : \underline{C} \rrbracket = \llbracket \Gamma \vdash N : \underline{C} \rrbracket$.

Proof sketch. Each part of this lemma is proved separately.

For compositionality, the proof is by induction on the structure of the term contexts $C[\]$ and $\underline{C}[\]$. Each case follows immediately from monotonicity (for example of the copairing operations $\langle -, - \rangle$).

For substitution, we first show a weakening lemma: there is an evident morphism

$$\llbracket \Gamma, x : A, \Gamma' \rrbracket \rightarrow \llbracket \Gamma, \Gamma' \rrbracket$$

and if $\Gamma, \Gamma' \vdash M : \underline{C}$ then the interpretation of $\Gamma, x : A, \Gamma' \vdash M : \underline{C}$ is equal to

$$\llbracket \Gamma, x : A, \Gamma' \rrbracket \longrightarrow \llbracket \Gamma, \Gamma' \rrbracket \xrightarrow{\llbracket \Gamma, \Gamma' \vdash M : \underline{C} \rrbracket} \llbracket \underline{C} \rrbracket$$

There is a similar fact for values. We then show the substitution lemma by induction on V and M , weakening substitutions where necessary. The second part of substitution follows immediately from the first part and monotonicity of \circ .

For the core axioms we go case-by-case. Most follow from the universal properties of the structure used to interpret each type former. For the axioms that involve substitution we use the previous part of this lemma. For the axiom $\lambda y : A. M \text{ to } x. N \equiv M \text{ to } x. \lambda y : A. N$ we use left-linearity of ev . \square

This lemma implies that the denotational semantics is sound with respect to the smallest inequational theory, in which the core axioms are the only axioms. It is not in general sound with respect to an arbitrary inequational theory. For our three examples in Section 4.1, we give signature axioms in addition to the core axioms, and define the inequational theory by closing under reflexivity, transitivity and congruence. The above lemma implies that to check soundness with respect to an inequational theory defined in this way, we need only check the signature axioms.

We call a structure a *model* if it is sound with respect to some given inequational theory.

Definition 4.2.7 (GCBPV model) A *model* of some GCBPV inequational theory \leq is a structure such that the interpretation of terms is *sound*: if $\Gamma \vdash V \leq W : A$ then $\llbracket \Gamma \vdash V : A \rrbracket \sqsubseteq \llbracket \Gamma \vdash W : A \rrbracket$, and if $\Gamma \vdash M \leq N : \underline{C}$ then $\llbracket \Gamma \vdash M : \underline{C} \rrbracket \sqsubseteq \llbracket \Gamma \vdash N : \underline{C} \rrbracket$. ◀

Each model induces a semantic notion of validity of program transformations: a transformation is valid in a model if it only replaces M with N when $\llbracket M \rrbracket \sqsubseteq \llbracket N \rrbracket$.

4.2.3 Constructing models

We give various models of GCBPV. First we present the three models we use to verify the example program transformations in Section 4.1, one each for undefined behaviour, nondeterminism and shared global state. We then describe how to construct *monadic* models of GCBPV. Finally we describe the term model, which we use for adequacy proofs in Section 4.3.

4.2.3.1 Undefined behaviour

Our undefined behaviour example has the simplest model. We interpret value types as posets (in the value category $\mathbf{C} = \mathbf{Poset}$), and computation types as pointed posets (in the computation category $\mathbf{D} = \mathbf{Poset}_\perp$). The least element of each object of \mathbf{Poset}_\perp represents undefined behaviour. We use the usual bicartesian structure of \mathbf{C} and cartesian structure of \mathbf{D} .

For the graded adjunction, the left adjoint sends each poset X to the pointed poset X_\perp obtained by freely adding a least element, and sends each monotone function $f : X \rightarrow Y$ to its strict extension $f_\perp : X_\perp \rightarrow Y_\perp$ (so $f_\perp \perp := \perp$ and $f_\perp x := x$ for $x \in X$). The right adjoint is the forgetful functor, which sends each pointed poset to itself. The unit $\eta_X : X \rightarrow X_\perp$ maps each element of X to itself, and the counit $\omega_X : (X_\perp)_\perp \rightarrow X_\perp$ merges the two bottom elements. For this example, the effect algebra is the trivial preordered monoid with underlying set $\{\star\}$. Hence the action $\star \otimes -$ is just the identity. Since the value category is \mathbf{Poset} , strength is also trivial in this case: the strength is given as above, and all natural transformations are automatically strong. With this graded adjunction, we can show that the computation M to x . N eagerly evaluates M in the semantics:

$$\llbracket M \text{ to } x. N \rrbracket \rho = \begin{cases} \perp & \text{if } \llbracket M \rrbracket \rho = \perp \\ \llbracket N \rrbracket (\rho, \llbracket M \rrbracket \rho) & \text{otherwise} \end{cases}$$

In this case, a natural transformation $\alpha_\varepsilon : U(\varepsilon \otimes \underline{X}) \rightarrow U(\varepsilon \otimes \underline{Y})$ is just a monotone function $\alpha_\star : \underline{X} \rightarrow \underline{Y}$ between the pointed posets \underline{X} and \underline{Y} . Such a natural transformation is linear if it is strict (i.e. is a morphism in \mathbf{Poset}_\perp). Similarly, left-linearity just means strict in the left argument.

The function space $X \Rightarrow Y$ is the set of monotone functions, ordered pointwise. The least element of the function space is the function that maps every element of X to the least element of Y . Currying and uncurrying are the usual operations on functions. The evaluation map is given by

$$ev_{\star, X, Y}(f, x) = f x$$

and is clearly left-linear.

The interpretation of the only base type **int** is just $\llbracket \mathbf{int} \rrbracket := \{-2^{31}, \dots, 2^{31} - 1\}$ (ordered by equality). Recall that constants $c \in \mathcal{K}_A$ are interpreted as morphisms $\llbracket c \rrbracket : 1 \rightarrow \llbracket A \rrbracket$. On \mathbf{Poset} , these are just elements $\llbracket c \rrbracket \in \llbracket A \rrbracket$ by abuse of notation. For undefined behaviour, **geq** and **add** have obvious interpretations (as functions $\llbracket \mathbf{int} \rrbracket \rightarrow \llbracket \mathbf{int} \rrbracket \rightarrow 2$ and $\llbracket \mathbf{int} \rrbracket \rightarrow \llbracket \mathbf{int} \rrbracket \rightarrow \llbracket \mathbf{int} \rrbracket$ respectively). The constant integers are interpreted as $\llbracket \underline{n} \rrbracket := n \in \llbracket \mathbf{int} \rrbracket$.

The operation `undef` is interpreted as a morphism $\llbracket \text{undef} \rrbracket : 1 \rightarrow 0_{\perp}$ (recall that 0 is the empty poset). For this we take the least element \perp . Finally, the operation `addnsw` is interpreted as a morphism $\llbracket \text{add}_{\text{nsw}} \rrbracket : \llbracket \text{int} \rrbracket \times \llbracket \text{int} \rrbracket \rightarrow \llbracket \text{int} \rrbracket_{\perp}$, which is given by:

$$\llbracket \text{add}_{\text{nsw}} \rrbracket (m, n) := \begin{cases} m + n & \text{if } (m + n) \in \llbracket \text{int} \rrbracket \\ \perp & \text{otherwise} \end{cases}$$

It is easy to see that this interpretation is sound, and hence a model in the sense of Definition 4.2.7, by checking the signature axioms. In particular, we have $\llbracket \text{undef}_{\underline{c}} \rrbracket \sqsubseteq \llbracket M \rrbracket$ because the left-hand side is the least element.

We validate each program transformation $M \leq_{\text{ctx}} N$ by showing that $\llbracket M \rrbracket \sqsubseteq \llbracket N \rrbracket$. This is sufficient because the model of undefined behaviour is *adequate*; we prove adequacy in Section 4.3.2.1. Hence, to validate the example transformations we show

$$\begin{aligned} & \llbracket \text{add}_{\text{nsw}}(V_1, V_2) \rrbracket \sqsubseteq \llbracket \langle \text{add}(V_1, V_2) \rangle \rrbracket \\ & \left[\begin{array}{l} \text{add}_{\text{nsw}}(V_1, V_2) \text{ to } x. \\ \text{let } b = \text{geq}(x, V_1) \text{ in} \\ N \end{array} \right] = \left[\begin{array}{l} \text{add}_{\text{nsw}}(V_1, V_2) \text{ to } x. \\ \text{let } b = \text{geq}(V_2, \underline{0}) \text{ in} \\ N \end{array} \right] \end{aligned}$$

To do this we just use cases on whether the addition overflows. For the first transformation if the addition does not overflow then the two sides are equal by definition, otherwise the left-hand side is the least element. The second is similar.

4.2.3.2 Nondeterminism

For our nondeterminism example, we again interpret value types as posets (in the value category **Poset**). For computation types we define a new **Poset**-category **D**. Recall that a meet-semilattice is a poset X such that each pair of elements $x_1, x_2 \in X$ has a *meet* (greatest lower bound) $x_1 \sqcap x_2$, i.e. an element of $x_1 \sqcap x_2 \in X$ that satisfies

$$\forall y \in X. y \sqsubseteq x_1 \sqcap x_2 \Leftrightarrow y \sqsubseteq x_1 \wedge y \sqsubseteq x_2$$

Objects of **D** are pairs of a meet-semilattice A_+ and a subset A_1 (which is not required to be closed under meets). We write such an object as $(A_1 \subseteq A_+)$. Morphisms $f : (A_1 \subseteq A_+) \rightarrow (B_1 \subseteq B_+)$ in **D** are monotone functions $f : A_+ \rightarrow B_+$ that preserve meets ($f(x_1 \sqcap x_2) = f x_1 \sqcap f x_2$), and such that if $x \in A_1$ then $f x \in B_1$.

The idea is that the elements of A_+ are nondeterministic, and elements of A_1 are deterministic. We use the usual bicartesian structure of **Poset**. The terminal object of **D** is $\{\star\} \subseteq \{\star\}$ and the binary product $(A_1 \subseteq A_+) \times (B_1 \subseteq B_+)$ is $(A_1 \times B_1) \subseteq (A_+ \times B_+)$. The exponential $X \Rightarrow (A_1 \subseteq A_+)$ consists of the meet-semilattice of monotone functions $f : X \rightarrow A_+$. The subset contains the monotone functions whose images are subsets of A_1 .

We next define the graded adjunction. The right adjoint $U : \mathbf{D} \rightarrow \mathbf{Poset}$ sends $(A_1 \subseteq A_+)$ to A_1 (the order on A_1 is the restriction of the order on A_+), and sends each morphism $f : (A_1 \subseteq A_+) \rightarrow (B_1 \subseteq B_+)$ to its restriction to $A_1 \rightarrow B_1$. The monoid action is defined for the effect $+$ by:

$$(+) \otimes (A_1 \subseteq A_+) := (A_+ \subseteq A_+) \quad (+) \otimes f := f$$

and on the effect 1 it is the identity.

For the left adjoint, recall that binary nondeterminism is traditionally modelled using the monad $\mathcal{P}_{\text{fin}}^+$ on **Set**, where $\mathcal{P}_{\text{fin}}^+ X$ is the set of nonempty finite subsets of X . A subset $S \in \mathcal{P}_{\text{fin}}^+ X$

is the set of possible results of a computation. Here we also have to take into account the ordering on elements of X . If $x \in X$ is a possible result of a computation and $x \sqsubseteq x'$ then x' should also be viewed as a possible result. Hence instead of using subsets $S \in \mathcal{P}_{\text{fin}}^+ X$ we use their *upwards closures* $\uparrow S := \{x' \in X \mid \exists x \in S. x \sqsubseteq x'\}$.

Formally, the left adjoint $F : \mathbf{Poset} \rightarrow \mathbf{D}$ sends each poset X to $(A_1 \subseteq A_+)$, where

$$A_1 := \{\uparrow\{x\} \mid x \in X\} \quad A_+ := \{\uparrow S \mid S \in \mathcal{P}_{\text{fin}}^+ X\}$$

and A_+ has superset as the order.⁴ The meet of $S_1, S_2 \in A_+$ is the union $S_1 \cup S_2$. On monotone functions $f : X \rightarrow Y$, the **Poset**-functor F is given by

$$Ff := \lambda S. \uparrow\{f x \mid x \in S\}$$

The unit of the adjunction is $\eta_X x := \uparrow\{x\}$. For the counit we note that for any $(A_1 \subseteq A_+) \in \mathbf{D}$ each set $S' \in \mathcal{P}_{\text{fin}}^+ A_+$ has a meet $\prod S' \in A_+$, and hence if $S \subseteq A_+$ is the upwards closure of some $S' \in \mathcal{P}_{\text{fin}}^+ A_+$ then S has a meet $\prod S = \prod S'$. The counit $\omega_{(A_1 \subseteq A_+)} : FA_1 \rightarrow (A_1 \subseteq A_+)$ is given by

$$\omega_{(A_1 \subseteq A_+)} S := \prod S$$

For this graded adjunction, sequencing of computations takes the meet of all possible results:

$$\llbracket M \text{ to } x. N \rrbracket = \prod_{a \in \llbracket M \rrbracket \rho} \llbracket N \rrbracket(\rho, a)$$

If M is deterministic (has type $\langle 1 \rangle A$ for some A), then $\llbracket M \rrbracket \rho = \uparrow\{a\}$ for some a , and so $\llbracket M \text{ to } x. N \rrbracket = \llbracket N \rrbracket(\rho, a)$ by monotonicity of $\llbracket N \rrbracket$.

To complete the model we give the interpretation of flip, which is a function $\llbracket \text{flip} \rrbracket : 1 \rightarrow \{\uparrow S \mid S \in \mathcal{P}_{\text{fin}}^+ \{\text{true}, \text{false}\}\}$. This maps the unique element of 1 to the set $\{\text{true}, \text{false}\}$, so that the two possible results of flip () are true and false. Under this interpretation we can show that binary nondeterministic choice takes the meet:

$$\llbracket M_1 \text{ or } M_2 \rrbracket \rho = \llbracket M_1 \rrbracket \rho \sqcap \llbracket M_2 \rrbracket \rho$$

It is not difficult to check soundness of the signature axioms (Figure 4.1), and hence that this data forms a model of the inequational theory for nondeterminism. For example, for the axiom $\text{flip} () \preceq \text{coerce}_{1 \leq +} \langle \text{true} \rangle$ we have

$$\llbracket \text{flip} () \rrbracket \rho = \{\text{true}, \text{false}\} \supseteq \{\text{true}\} = \llbracket \text{coerce}_{1 \leq +} \langle \text{true} \rangle \rrbracket \rho$$

We show that this model is adequate in Section 4.3.2.2.

We verify that both of the program transformations given in Section 4.1.2 hold. For reusing the result of a duplicated computation this means showing

$$\llbracket M \text{ to } x. M \text{ to } y. N \rrbracket \sqsubseteq \llbracket M \text{ to } x. N[y \mapsto x] \rrbracket$$

Using the interpretation of **to** above this inequality becomes

$$\prod_{x \in \llbracket M \rrbracket \rho} \prod_{y \in \llbracket M \rrbracket \rho} \llbracket N \rrbracket((\rho, x), y) \sqsubseteq \prod_{x \in \llbracket M \rrbracket \rho} \llbracket N \rrbracket((\rho, x), x)$$

⁴For the left adjoint, A_+ is the *free meet-semilattice* on the poset X . (Similarly, $\mathcal{P}_{\text{fin}}^+ Y$ ordered by \supseteq is the free meet-semilattice on the set Y .)

and this holds, so the left-to-right transformation is valid. If M is deterministic (has type $\langle 1 \rangle A$ for some A), then $\llbracket M \rrbracket \rho = \uparrow\{a\}$ for some a , and the two sides of this inequality are just $\llbracket N \rrbracket((\rho, a), a)$. So in this case, we have an equality, and the right-to-left direction of the transformation is also valid.

The other example we give is eliminating a dead computation, which is valid because

$$\llbracket M; N \rrbracket \rho = \bigsqcap_{x \in \llbracket M \rrbracket \rho} \llbracket N \rrbracket \rho = \llbracket N \rrbracket \rho$$

4.2.3.3 Shared global state

We give a model of shared global state on **Poset**, in the style of *algebraic effects* [88]. We do not give a general description of algebraic effects on **Poset**; instead we consider only `get` and `put`. A general description of algebraic models of GCBPV on **Poset** would be similar to [42, 64, 47], but adapted to include grading.

We interpret value types as posets (in the value category **Poset**). The computation category is more complicated, and somewhat similar to the computation category we use for nondeterminism. Objects are roughly posets, together with interpretations of the `get` and `put` operations.

Recall that the set of operations is $\Sigma := \{\text{get}, \text{put}\}$, and effects are subsets $\varepsilon \subseteq \Sigma$.

Definition 4.2.8 A *graded shared mnemoid*⁵ (A, g, p) consists of a family of sets $(A_\varepsilon)_{\varepsilon \subseteq \Sigma}$ such that $\varepsilon \subseteq \varepsilon'$ implies $A_\varepsilon \subseteq A_{\varepsilon'}$, a partial order \sqsubseteq on A_Σ , and monotone functions

$$g : A_\Sigma \times A_\Sigma \rightarrow A_\Sigma \quad p_{\text{true}} : A_\Sigma \rightarrow A_\Sigma \quad p_{\text{false}} : A_\Sigma \rightarrow A_\Sigma$$

such that if $a_{\text{true}}, a_{\text{false}} \in A_\varepsilon$ then $g(a_{\text{true}}, a_{\text{false}}) \in A_{\varepsilon \cup \{\text{get}\}}$, if $a \in A_\varepsilon$ then $p_b a \in A_{\varepsilon \cup \{\text{put}\}}$, and the following axioms hold:

$$\begin{aligned} g(p_{\text{true}} a_{\text{true}}, p_{\text{false}} a_{\text{false}}) &\sqsubseteq g(a_{\text{true}}, a_{\text{false}}) & p_b(g(a_{\text{true}}, a_{\text{false}})) &\sqsubseteq p_b a_b \\ g(g(a_{\text{true}, \text{true}}, a_{\text{true}, \text{false}}), g(a_{\text{false}, \text{true}}, a_{\text{false}, \text{false}})) &\sqsubseteq g(a_{\text{true}, \text{true}}, a_{\text{false}, \text{false}}) \\ p_{b_1}(p_{b_2} a) &\sqsubseteq p_{b_2} a & g(a, a) &= a \end{aligned}$$

The computation category **GSMnem** has graded shared mnemoids as objects. Morphisms $f : (A, g, p) \rightarrow (A', g', p')$ are monotone functions $f : A_\Sigma \rightarrow A'_\Sigma$ such that

$$f(g(a_{\text{true}}, a_{\text{false}})) = g'(f a_{\text{true}}, f a_{\text{false}}) \quad f(p_b a) = p'_b(f a)$$

and $a \in A_\varepsilon$ implies $f a \in A'_\varepsilon$. Morphisms are ordered pointwise. \blacktriangleleft

The set A_ε contains the interpretations of computations with effect ε . The functions provide the operations: $g(a_{\text{true}}, a_{\text{false}})$ means get the value b of the state and then run a_b , and $p_b a$ means put the value b and then run a . The axioms are similar to those of the inequational theory in Figure 4.2.

We use the usual bicartesian structure on **Poset**. The cartesian structure on **GSMnem**, and exponentials, are straightforward. In particular, the family of sets in the product $(A, g, p) \times (A', g', p')$ is $(A_\varepsilon \times A'_\varepsilon)_{\varepsilon \subseteq \Sigma}$, and the functions are given componentwise. For the exponential $X \Rightarrow (A, g, p)$ we take the set of monotone functions $X \rightarrow A_\Sigma$.

⁵These are based on Melliès's [72] *mnemoids*. We adapt the definition to *shared* global state, and also add grading (given by the subsets). The usual global state monad on a fixed set forms a graded shared mnemoid (just as it forms a mnemoid); this implies consistency of our axioms.

For the graded adjunction the right adjoint $U : \mathbf{GSMnem} \rightarrow \mathbf{Poset}$ sends (A, g, p) to A_\emptyset , and morphisms f to their restrictions to A_\emptyset . The monoid action is given by:

$$\varepsilon \otimes (A, g, p) := ((A_{\varepsilon \cup \varepsilon'})_{\varepsilon' \subseteq \Sigma}, g, p) \quad \varepsilon \otimes f := f$$

and $(\varepsilon \subseteq \varepsilon') \otimes -$ is the identity on A_Σ .

For the left adjoint, given a poset (X, \sqsubseteq_X) , terms t over X are given by the following formal grammar, where x ranges over elements of X and b over elements of $2 = \{\text{true}, \text{false}\}$:

$$t ::= x \mid \text{get}(t_{\text{true}}, t_{\text{false}}) \mid \text{put}_b t$$

The term x just returns x ; the term $\text{get}(t_{\text{true}}, t_{\text{false}})$ gets the value b of the state and then runs the term t_b ; and $\text{put}_b t$ puts b and then runs the term t . For example, $\text{get}(\text{put}_{\text{false}} x, \text{put}_{\text{true}} x)$ is a term that gets the value of the state, puts its negation, and then returns x . We capture the required behaviour of these terms using a (non-antisymmetric) preorder \preceq on terms. It is defined as the smallest preorder that is closed under congruence

$$\frac{x \sqsubseteq_X x'}{x \preceq x'} \quad \frac{t_{\text{true}} \preceq t'_{\text{true}} \quad t_{\text{false}} \preceq t'_{\text{false}}}{\text{get}(t_{\text{true}}, t_{\text{false}}) \preceq \text{get}(t'_{\text{true}}, t'_{\text{false}})} \quad \frac{t \preceq t' \quad b \in 2}{\text{put}_b t \preceq \text{put}_b t'}$$

and under the following six axioms (which should be compared with the shared mnemoid axioms above):

$$\begin{aligned} \text{get}(\text{put}_{\text{true}} t_{\text{true}}, \text{put}_{\text{false}} t_{\text{false}}) &\preceq \text{get}(t_{\text{true}}, t_{\text{false}}) & \text{put}_b(\text{get}(t_{\text{true}}, t_{\text{false}})) &\preceq \text{put}_b t_b \\ \text{get}(\text{get}(t_{\text{true}, \text{true}}, t_{\text{true}, \text{false}}), \text{get}(t_{\text{false}, \text{true}}, t_{\text{false}, \text{false}})) &\preceq \text{get}(t_{\text{true}, \text{true}}, t_{\text{false}, \text{false}}) \\ \text{put}_{b_1}(\text{put}_{b_2} t) &\preceq \text{put}_{b_2} t & \text{get}(t, t) &\preceq t & t &\preceq \text{get}(t, t) \end{aligned}$$

The intersection of \preceq and its converse \succeq is an equivalence relation on computations. We quotient by this equivalence relation, so that \preceq induces a partial order on equivalence classes $[t]$ of computations t . (We do not know of a description that does not involve a quotient.) From this point onwards, we always consider equivalence classes, and suppress the square brackets in the notation.

We construct a graded shared mnemoid $FX := (A, g, p)$ as follows: A_ε is the set of equivalence classes of terms t over X that contain only the operations in ε . The functions g and p_b are straightforward:

$$g(t_{\text{true}}, t_{\text{false}}) := \text{get}(t_{\text{true}}, t_{\text{false}}) \quad p_b t := \text{put}_b t$$

This defines the left adjoint on objects. To define it on morphisms, we first define the *interpretations* of terms in graded shared mnemoids (B, g, p) . Given a term t over X , and a monotone function $\rho : X \rightarrow B_\Sigma$, we define $\mathcal{I}[[t]]\rho \in B_\Sigma$ by:

$$\mathcal{I}[[x]]\rho := \rho x \quad \mathcal{I}[[\text{get}(t_{\text{true}}, t_{\text{false}})]]\rho := g(\mathcal{I}[[t_{\text{true}}]]\rho, \mathcal{I}[[t_{\text{false}}]]\rho) \quad \mathcal{I}[[\text{put}_b t]]\rho := p_b(\mathcal{I}[[t]]\rho)$$

(This is well-defined and monotone by the axioms required for graded shared mnemoids.) Any monotone function $f : X \rightarrow Y$ can be viewed as a monotone function $f : X \rightarrow B_\Sigma$ where B_Σ is the set of equivalence classes of terms over Y , and so we define $Ff := \mathcal{I}[[-]]f$. The unit of the graded adjunction maps each x to the trivial computation x . The counit $\omega_{(A, g, p)} : FA_\emptyset \rightarrow (A, g, p)$ is given by $\omega_{(A, g, p)} := \mathcal{I}[[-]]id$.

The constants are interpreted as follows:

$$[[\text{get}]] \star := (\text{get}(\text{true}, \text{false})) \quad [[\text{put}]] b := (\text{put}_b \star)$$

This defines a model of the inequational theory for shared global state, in which

$$\llbracket M \text{ to } x. N \rrbracket \rho := \mathcal{I}[\llbracket M \rrbracket \rho](a \mapsto \llbracket N \rrbracket(\rho, a))$$

We show adequacy in Section 4.3.2.3.

Finally, we verify that the example transformations in Section 4.1.3 are valid in this model. For reusing the result of a duplicated computation we show

$$\llbracket M \text{ to } x. M \text{ to } y. N \rrbracket \sqsubseteq \llbracket M \text{ to } x. N[y \mapsto x] \rrbracket$$

for computations M of type $\langle \varepsilon \rangle A$ with $\varepsilon \subseteq \{\text{put}\}$. Given any $\rho \in \llbracket \Gamma \rrbracket$, the term $\llbracket M \rrbracket \rho$ has the form $\text{put}_{b_1}(\text{put}_{b_2}(\cdots(\text{put}_{b_n} a) \cdots))$ for some $b_1, \dots, b_n \in 2$ and $a \in \llbracket A \rrbracket$, because of the restriction on the effect ε . By expanding the interpretations of the two computations, the goal becomes

$$p_{b_1}(\cdots(p_{b_n}(p_{b_1}(\cdots(p_{b_n}(\llbracket N \rrbracket(\rho, (a, a)))) \cdots))) \cdots) \sqsubseteq p_{b_1}(\cdots(p_{b_n}(\llbracket N \rrbracket(\rho, (a, a)))) \cdots)$$

If $n = 0$ then this holds automatically; if $n > 0$ then we can show by induction on n that

$$p_{b_1}(\cdots(p_{b_n}(p_b a)) \cdots) \sqsubseteq p_b a$$

for all a, b , using the axiom $p_{b'}(p_{b''} a') \sqsubseteq p_{b''} a'$, and instantiate this to get the inequality we need.

For eliminating a dead computation we show

$$\llbracket M'; N' \rrbracket = \llbracket N' \rrbracket$$

for computations M' of type $\langle \varepsilon' \rangle A'$ where $\varepsilon' \subseteq \{\text{get}\}$. This uses a similar induction, except that we use the axiom $g(a, a) = a$.

4.2.3.4 Monadic models

The three models above suffice for our examples. We additionally give a *general* source of models. It is well-known that monads can be used to model various side-effects, and strong monads are used for models of the monadic metalanguage [78]. Monads can also be used as models for CBPV by taking the adjunction to be the *Eilenberg-Moore resolution*, as shown by Levy [57, Chapter 12]. Here we generalize the notion of monadic model of CBPV to the graded, **Poset**-enriched case. This allows to use previous monadic models of side-effects as models of GCBPV. The construction in this chapter is a relatively straightforward generalization, and the proofs are almost identical.

Monadic models of GCBPV are based on Katsumata's [45] *graded monads*. We again need tensorial strengths (to interpret open terms) and, in our case, need **Poset**-enrichment. Hence we use *graded strong Poset-monads*. In the following definition, \mathbf{C} should be thought of as the value category of a model of GCBPV.

Definition 4.2.9 Suppose that $(\mathcal{E}, \leq, \cdot, 1)$ is a preordered monoid and \mathbf{C} is a cartesian **Poset**-category. A (\mathcal{E}) -*graded strong Poset-monad*⁶ $(T, \text{str}, \eta, \mu)$ on \mathbf{C} consists of:

- A strong **Poset**-functor $(T\varepsilon, \text{str}_\varepsilon)$ on \mathbf{C} for each $\varepsilon \in \mathcal{E}$.
- A strong natural transformation $T(\varepsilon \leq \varepsilon') : (T\varepsilon, \text{str}_\varepsilon) \rightarrow (T\varepsilon', \text{str}_{\varepsilon'})$ for each $\varepsilon \leq \varepsilon' \in \mathcal{E}$.
- A strong natural transformation $\eta : \text{Id}_{\mathbf{C}} \rightarrow (T1, \text{str}_1)$ from the identity strong **Poset**-functor.

- A strong natural transformation $\mu_{\varepsilon, \varepsilon'} : (T\varepsilon, str_\varepsilon) \circ (T\varepsilon', str_{\varepsilon'}) \rightarrow (T(\varepsilon \cdot \varepsilon'), str_{\varepsilon \cdot \varepsilon'})$ for each $\varepsilon, \varepsilon' \in \mathcal{E}$.

Such that:

- T is functorial: $T(\varepsilon \leq \varepsilon) = id_{T\varepsilon}$ and if $\varepsilon \leq \varepsilon' \leq \varepsilon''$ then $T(\varepsilon \leq \varepsilon'') = T(\varepsilon' \leq \varepsilon'') \circ T(\varepsilon \leq \varepsilon')$.
- μ is natural in ε and ε' : if $\varepsilon_1 \leq \varepsilon_2$ and $\varepsilon'_1 \leq \varepsilon'_2$ then the following diagrams commute

$$\begin{array}{ccc}
T\varepsilon_1 \circ T\varepsilon' & \xrightarrow{(T(\varepsilon_1 \leq \varepsilon_2))_{T\varepsilon'}} & T\varepsilon_2 \circ T\varepsilon' & & T\varepsilon \circ T\varepsilon'_1 & \xrightarrow{T\varepsilon(T(\varepsilon'_1 \leq \varepsilon'_2))} & T\varepsilon \circ T\varepsilon'_2 \\
\mu_{\varepsilon_1, \varepsilon'} \downarrow & & \downarrow \mu_{\varepsilon_2, \varepsilon'} & & \mu_{\varepsilon, \varepsilon'_1} \downarrow & & \downarrow \mu_{\varepsilon, \varepsilon'_2} \\
T(\varepsilon_1 \cdot \varepsilon') & \xrightarrow{T(\varepsilon_1 \cdot \varepsilon' \leq \varepsilon_2 \cdot \varepsilon')} & T(\varepsilon_2 \cdot \varepsilon') & & T(\varepsilon \cdot \varepsilon'_1) & \xrightarrow{T(\varepsilon \cdot \varepsilon'_1 \leq \varepsilon \cdot \varepsilon'_2)} & T(\varepsilon \cdot \varepsilon'_2)
\end{array}$$

- The (graded) monad laws hold: the following diagrams commute

$$\begin{array}{ccc}
\eta_{T\varepsilon} \curvearrowright & T\varepsilon & \curvearrowleft T\varepsilon\eta \\
\parallel & & \\
T1 \circ T\varepsilon & \xrightarrow{\mu_{1, \varepsilon}} & T\varepsilon \xleftarrow{\mu_{\varepsilon, 1}} T\varepsilon \circ T1
\end{array}
\quad
\begin{array}{ccc}
T\varepsilon \circ T\varepsilon' \circ T\varepsilon'' & \xrightarrow{\mu_{\varepsilon, \varepsilon', T\varepsilon''}} & T(\varepsilon \cdot \varepsilon') \circ T\varepsilon'' \\
T\varepsilon\mu_{\varepsilon', \varepsilon''} \downarrow & & \downarrow \mu_{\varepsilon \cdot \varepsilon', \varepsilon''} \\
T\varepsilon \circ T(\varepsilon' \cdot \varepsilon'') & \xrightarrow{\mu_{\varepsilon, \varepsilon' \cdot \varepsilon''}} & T(\varepsilon \cdot \varepsilon' \cdot \varepsilon'')
\end{array}
\quad \blacktriangleleft$$

Graded strong **Poset**-functors can be used as models of our graded version of the monadic metalanguage (GMM in Section 2.5). The **Poset**-functors $T\varepsilon$ are used to interpret the type constructors $\langle \varepsilon \rangle$. The strong natural transformations $T(\varepsilon \leq \varepsilon')$ are used for coercions. The *unit* η is used to interpret pure computations, and the *multiplication* μ is used for sequencing of computations.

If T is a graded strong **Poset**-monad then it forms a functor $\mathcal{E} \rightarrow [\mathbf{C}, \mathbf{C}]$ (where as usual we view \mathcal{E} as a category). The strength is trivial in many cases in the same way that it is for adjunctions. If $\mathbf{C} = \mathbf{Set}$ or $\mathbf{C} = \mathbf{Poset}$, there is always a unique strength, and every natural transformation is strong. If $\mathbf{C} = \omega\mathbf{Cpo}$ the strength is unique if it exists, and it does exist if $T\varepsilon$ preserves least upper bounds of morphisms for each ε .

It is well-known that every adjunction induces a monad, and there is an analogous fact in our situation:

Lemma 4.2.10 Suppose that (F, U, \otimes) is a graded strong **Poset**-adjunction with $F : \mathbf{C} \rightarrow \mathbf{D}$. Define $T : \mathcal{E} \rightarrow [\mathbf{C}, \mathbf{C}]$ by:

$$T\varepsilon := U(\varepsilon \otimes F-) \quad T(\varepsilon \leq \varepsilon') := U((\varepsilon \leq \varepsilon') \otimes F-)$$

Then T forms a graded strong **Poset**-monad. Each strength str_ε is the strength of the adjunction, the unit η is the unit of the adjunction, and the multiplication is given by:

$$\mu_{\varepsilon, \varepsilon', X} := U(\varepsilon \otimes \omega_{\varepsilon' \otimes X}) : U(\varepsilon \otimes F(U(\varepsilon' \otimes FX))) \rightarrow U((\varepsilon \cdot \varepsilon') \otimes FX) \quad \blacktriangleleft$$

It is possible to continue in this direction and show that each GCBPV structure induces most of a GMM structure (the data required to interpret the graded monadic metalanguage (Section 2.5), which includes a graded strong **Poset**-monad; the only missing part is that we do not quite

⁶Similar to the non-enriched case, graded strong **Poset**-monads are just *lax monoidal functors* $\mathcal{E} \rightarrow [\mathbf{C}, \mathbf{C}]_s$, where the preordered monoid \mathcal{E} is viewed as a thin monoidal category, and $[\mathbf{C}, \mathbf{C}]_s$ is the monoidal category of strong **Poset**-functors on \mathbf{C} , with composition as the tensor product. Graded strong **Poset**-monads can be equivalently presented as *Kleisli triples*, which have morphisms $(\gg) : (X \Rightarrow T\varepsilon'Y) \rightarrow (T\varepsilon X \Rightarrow T(\varepsilon \cdot \varepsilon')Y)$.

have exponentials on \mathbf{C}). This should not be surprising: we mentioned in Section 2.7.3 that we can translate most of GMM into GCBPV (by interpreting the graded monad $\langle \varepsilon \rangle -$ as $\mathbf{U} \langle \varepsilon \rangle -$).

In this section we go in the other direction, and construct models of GCBPV from models of GMM. Given a graded strong **Poset**-monad on \mathbf{C} , we construct a graded strong **Poset**-adjunction. We use \mathbf{C} as the value category of the model, so the first step is to construct the computation category \mathbf{D} .

Definition 4.2.11 (Graded algebra [76, 28]) Suppose that $T : \mathcal{E} \rightarrow [\mathbf{C}, \mathbf{C}]$ is a graded strong **Poset**-monad. A T -algebra is a pair (A, a) of a functor $A : \mathcal{E} \rightarrow \mathbf{C}$ and a natural transformation

$$a_{\varepsilon, \varepsilon'} : T\varepsilon(A\varepsilon') \rightarrow A(\varepsilon \cdot \varepsilon')$$

such that the following diagrams commute for all $\varepsilon, \varepsilon', \varepsilon''$:

$$\begin{array}{ccc} A\varepsilon & \xrightarrow{\eta} & T1(A\varepsilon) & & T\varepsilon(T\varepsilon'(A\varepsilon'')) & \xrightarrow{T\varepsilon a} & T\varepsilon(A(\varepsilon' \cdot \varepsilon'')) \\ & \searrow & \downarrow a_{1, \varepsilon} & & \mu \downarrow & & \downarrow a \\ & & A\varepsilon & & T(\varepsilon \cdot \varepsilon')(A\varepsilon'') & \xrightarrow{a} & A(\varepsilon \cdot \varepsilon' \cdot \varepsilon'') \end{array}$$

We call the functor A the *carrier* of the algebra.

A *homomorphism* from (A, a) to (A', a') is a natural transformation $h : A \rightarrow A'$ such that

$$\begin{array}{ccc} T\varepsilon(A\varepsilon') & \xrightarrow{T\varepsilon h} & T\varepsilon(A'\varepsilon') \\ a \downarrow & & \downarrow a' \\ A(\varepsilon \cdot \varepsilon') & \xrightarrow{h} & A'(\varepsilon \cdot \varepsilon') \end{array}$$

commutes.

The *Eilenberg-Moore category* $T\text{-Alg}$ of the graded monad T has T -algebras as objects and graded algebra homomorphisms as morphisms. It is a **Poset**-category: $h \sqsubseteq h'$ if $h_\varepsilon \sqsubseteq h'_\varepsilon$ for all ε . \blacktriangleleft

(This definition does not actually require strength, but we assume strength throughout this section for consistency.)

We take $T\text{-Alg}$ as the computation category \mathbf{D} , and construct a graded adjunction

$$\begin{array}{ccc} \mathbf{C} & \begin{array}{c} \xrightarrow{F} \\ \perp \\ \xleftarrow{U} \end{array} & T\text{-Alg} \end{array} \quad \varepsilon \otimes -$$

The right adjoint U is the *forgetful Poset*-functor, which is given by

$$U(A, a) := A1 \quad Uh := h_1$$

The left adjoint F maps each object X to the *free T -algebra on X* :

$$FX := (T(-)X, \mu) \quad (Ff)_\varepsilon := T\varepsilon f$$

So the carrier of FX on the effect ε is $T\varepsilon X$. The action $\varepsilon \otimes -$ multiplies by the effect ε : given a T -algebra (A, a) , the T -algebra $\varepsilon \otimes (A, a)$ is (B, b) , where:

$$B\varepsilon' := A(\varepsilon \cdot \varepsilon') \quad B(\varepsilon' \leq \varepsilon'') := A(\varepsilon \cdot \varepsilon' \leq \varepsilon \cdot \varepsilon'') \quad b_{\varepsilon', \varepsilon''} := a_{\varepsilon', \varepsilon \cdot \varepsilon''}$$

This defines the action on objects. On homomorphisms $h : (A, a) \rightarrow (A', a')$, the homomorphism $\varepsilon \otimes h : \varepsilon \otimes (A, a) \rightarrow \varepsilon \otimes (A', a')$ is given by

$$(\varepsilon \otimes h)_{\varepsilon'} := h_{\varepsilon \cdot \varepsilon'}$$

and if $\varepsilon \leq \varepsilon'$, then the homomorphism $(\varepsilon \leq \varepsilon') \otimes (A, a) : \varepsilon \otimes (A, a) \rightarrow \varepsilon' \otimes (A, a)$ is given by:

$$((\varepsilon \leq \varepsilon') \otimes (A, a))_{\varepsilon''} := A(\varepsilon \cdot \varepsilon'' \leq \varepsilon' \cdot \varepsilon'')$$

Finally, the unit of the graded adjunction is just the unit η of the monad. The counit is given by

$$(\varpi_{(A,a)})_{\varepsilon} := a_{1,\varepsilon} : T1(A\varepsilon) \rightarrow A\varepsilon$$

and the strength is just the strength of the monad.

The data that we construct here satisfies $U(\varepsilon \otimes F-) = T\varepsilon$. If we apply Lemma 4.2.10 to it then we recover the graded monad we started with. Showing this construction does give a graded strong **Poset**-adjunction is straightforward.

Lemma 4.2.12 Given any graded strong **Poset**-monad, the construction above defines a graded strong **Poset**-adjunction. ◀

We still have to construct the rest of the data required for a model, in particular, we need the computation category $T\text{-Alg}$ to be cartesian, and also require exponentials. We can construct these out of similar structure on \mathbf{C} .

Lemma 4.2.13 Suppose that $T : \mathcal{E} \rightarrow [\mathbf{C}, \mathbf{C}]$ is a graded strong **Poset**-monad on a cartesian closed **Poset**-category \mathbf{C} . Then:

1. The terminal object $\underline{1}$ of $T\text{-Alg}$ is the unique T -algebra with the constantly-1 functor as carrier, and satisfies $\varepsilon \otimes \underline{1} = \underline{1}$.
2. The binary product of (A_1, a_1) and (A_2, a_2) is the T -algebra with carrier $A_1 \times A_2$ and natural transformation

$$T\varepsilon(A_1\varepsilon' \times A_2\varepsilon') \xrightarrow{\langle T\varepsilon\pi_1, T\varepsilon\pi_2 \rangle} T\varepsilon(A_1\varepsilon') \times T\varepsilon(A_2\varepsilon') \xrightarrow{a_{1,\varepsilon,\varepsilon'} \times a_{2,\varepsilon,\varepsilon'}} A_1(\varepsilon \cdot \varepsilon') \times A_2(\varepsilon \cdot \varepsilon')$$

together with $\pi_i : A_1\varepsilon \times A_2\varepsilon \rightarrow A_i\varepsilon$ as the i th projection. It satisfies

$$\varepsilon \otimes ((A_1, a_1) \times (A_2, a_2)) = (\varepsilon \otimes (A_1, a_1)) \times (\varepsilon \otimes (A_2, a_2)) \quad \langle \varepsilon \otimes \pi_1, \varepsilon \otimes \pi_2 \rangle = id_{\varepsilon \otimes ((A_1, a_1) \times (A_2, a_2))}$$

3. The exponential $X \Rightarrow (A, a)$ is the T -algebra with carrier $X \Rightarrow A$ and natural transformation

$$T\varepsilon(X \Rightarrow A\varepsilon') \xrightarrow{\Lambda(\text{ev} \circ \text{str}^r)} X \Rightarrow T\varepsilon(A\varepsilon') \xrightarrow{X \Rightarrow a_{\varepsilon,\varepsilon'}} X \Rightarrow A(\varepsilon \cdot \varepsilon')$$

with currying given by currying for exponentials in \mathbf{C} . It satisfies

$$\varepsilon \otimes (X \Rightarrow (A, a)) = X \Rightarrow (\varepsilon \otimes (A, a)) \quad \text{ev is left-linear} \quad \blacktriangleleft$$

The proof is essentially the same as for the non-graded, non-enriched case.

To give a GCBPV structure, it therefore suffices to give:

- A bicartesian closed **Poset**-category \mathbf{C} .
- A graded strong **Poset**-monad T on \mathbf{C} .
- An object $\llbracket b \rrbracket \in \mathbf{C}$ for each base type b .

- A morphism $\llbracket c \rrbracket : 1 \rightarrow \llbracket A \rrbracket$ for each constant $c \in \mathcal{K}_A$ and a morphism $\llbracket \text{op} \rrbracket : \llbracket \text{ar}_{\text{op}} \rrbracket \rightarrow T\varepsilon \llbracket \text{car}_{\text{op}} \rrbracket$ for each operation $\text{op} \in \Sigma$.

The amount of data required and number of diagrams to check in this definition can be large (for Gifford-style effect algebras, it is exponential in the number of operations). We end this section by giving a lemma that allows us to reduce the amount of effort required to specify monadic models.

A graded strong monad *structure* consists of the same data as a graded strong **Poset**-monad, but none of the laws.

Definition 4.2.14 Suppose that $(\mathcal{E}, \leq, \cdot, 1)$ is a preordered monoid and \mathbf{C} is a cartesian **Poset**-category. A *graded strong monad structure* on \mathbf{C} consists of:

- For each $\varepsilon \in \mathcal{E}$: an object $T\varepsilon X \in \mathbf{C}$ for each $X \in \mathbf{C}$, a morphism $T\varepsilon f : T\varepsilon X \rightarrow T\varepsilon Y$ for each $f : X \rightarrow Y$, and a morphism $\text{str}_{\varepsilon, X, Y} : X \times T\varepsilon Y \rightarrow T\varepsilon(X \times Y)$ for each $X, Y \in \mathbf{C}$.
- A morphism $(T(\varepsilon \leq \varepsilon'))_X : T\varepsilon X \rightarrow T\varepsilon' X$ for each $\varepsilon \leq \varepsilon' \in \mathcal{E}$ and $X \in \mathbf{C}$.
- A morphism $\eta_X : X \rightarrow T1X$ for each $X \in \mathbf{C}$.
- A morphism $\mu_{\varepsilon, \varepsilon', X} : T\varepsilon(T\varepsilon' X) \rightarrow T(\varepsilon \cdot \varepsilon')X$ for each $\varepsilon, \varepsilon' \in \mathcal{E}$ and $X \in \mathbf{C}$. ◀

Each graded strong monad structure either is or is not a graded strong **Poset**-monad. We show that such a structure is a graded strong **Poset**-monad if it forms a *grading* of a strong **Poset**-monad S , which means that it comes with morphisms $m_{\varepsilon, X} : T\varepsilon X \rightarrow SX$. These morphisms allow us to view computations in $T\varepsilon X$ as computations in S . They are required to preserve the structure (multiplication, etc.), and also to be *full monomorphisms*:

Definition 4.2.15 (Meseguer [74]) A morphism $m : X \rightarrow Y$ in a **Poset**-category is a *full monomorphism* if $m \circ f \sqsubseteq m \circ g$ implies $f \sqsubseteq g$ for all $f, g : Y \rightarrow Z$. ◀

Each full monomorphism is in particular a monomorphism ($m \circ f = m \circ g \Rightarrow f = g$). In **Set** the full monomorphisms are the injections, in **Poset** they are the monotone functions such that $m x \sqsubseteq m y \Rightarrow x \sqsubseteq y$.

Definition 4.2.16 Suppose that $(\mathcal{E}, \leq, \cdot, 1)$ is a preordered monoid, \mathbf{C} is a cartesian **Poset**-category and that S is a strong **Poset**-monad on \mathbf{C} . A *grading* of S consists of a graded strong monad structure T on \mathbf{C} together with a full monomorphism

$$m_{\varepsilon, X} : T\varepsilon X \rightarrow SX$$

for each $\varepsilon \in \mathcal{E}$ and $X \in \mathbf{C}$, such that the following diagrams commute:

$$\begin{array}{ccc}
 \begin{array}{ccc}
 T\varepsilon X & \xrightarrow{T\varepsilon f} & T\varepsilon Y \\
 m_{\varepsilon, X} \downarrow & & \downarrow m_{\varepsilon, Y} \\
 SX & \xrightarrow{Sf} & SY
 \end{array} &
 \begin{array}{ccc}
 X \times T\varepsilon Y & \xrightarrow{\text{str}_{\varepsilon, X, Y}} & T\varepsilon(X \times Y) \\
 X \times m_{\varepsilon, Y} \downarrow & & \downarrow m_{\varepsilon, X \times Y} \\
 X \times SY & \xrightarrow{\text{str}_{X, Y}} & S(X \times Y)
 \end{array} &
 \begin{array}{ccc}
 T\varepsilon X & & \\
 (T(\varepsilon \leq \varepsilon'))_X \downarrow & \searrow m_{\varepsilon, X} & \\
 T\varepsilon' X & \xrightarrow{m_{\varepsilon', X}} & SX
 \end{array} \\
 \\
 \begin{array}{ccc}
 X & & \\
 \eta_X \downarrow & \searrow \eta_X & \\
 T1X & \xrightarrow{m_{1, X}} & SX
 \end{array} &
 \begin{array}{ccc}
 T\varepsilon(T\varepsilon' X) & \xrightarrow{m_{\varepsilon, T\varepsilon' X}} & S(T\varepsilon' X) \xrightarrow{Sm_{\varepsilon', X}} & S(SX) \\
 \mu_{\varepsilon, \varepsilon', X} \downarrow & & & \downarrow \mu_X \\
 T(\varepsilon \cdot \varepsilon')X & \xrightarrow{m_{\varepsilon \cdot \varepsilon', X}} & SX &
 \end{array}
 \end{array}$$

◀

The lemma we use to reduce the amount of effort required to specify graded strong **Poset**-monads is as follows.

Lemma 4.2.17 Suppose that $(\mathcal{E}, \leq, \cdot, 1)$ is a preordered monoid, \mathbf{C} is a cartesian **Poset**-category, and S is a strong **Poset**-monad on \mathbf{C} . The graded monad structure T of any grading of S is a graded strong **Poset**-monad.

Proof. We reduce each of the requirements in the definition of graded strong **Poset**-monad to the corresponding requirement on S by composing with m , and then using the diagrams in the definition of grading. This suffices because m is a full monomorphism. We do not give all of the cases since they are all similar to each other, we just give a few representative examples.

We must show that each $T\mathcal{E}$ forms a **Poset**-functor. Monotonicity holds because if $f \sqsubseteq g$ then:

$$m_{\varepsilon, Y} \circ T\mathcal{E}f = Sf \circ m_{\varepsilon, X} \sqsubseteq Sg \circ m_{\varepsilon, X} = m_{\varepsilon, Y} \circ T\mathcal{E}g$$

Composition is preserved because:

$$m_{\varepsilon, Z} \circ T\mathcal{E}g \circ T\mathcal{E}f = Sg \circ Sf \circ m_{\varepsilon, X} = S(g \circ f) \circ m_{\varepsilon, X} = m_{\varepsilon, Z} \circ T\mathcal{E}(g \circ f)$$

The associativity law for the strength holds because of the following diagram chase

$$\begin{array}{ccccc}
(X \times Y) \times T\mathcal{E}Z & \xrightarrow{\quad str_{\varepsilon, X \times Y, Z} \quad} & T\mathcal{E}((X \times Y) \times Z) & & \\
\downarrow \text{assoc} & \searrow^{(X \times Y) \times m_{\varepsilon, Z}} & \downarrow m_{\varepsilon, (X \times Y) \times Z} & & \downarrow T\mathcal{E} \text{assoc} \\
& & (X \times Y) \times SZ & \xrightarrow{\quad str_{X \times Y, Z} \quad} & S((X \times Y) \times Z) & \xrightarrow{\quad S \text{assoc} \quad} & T\mathcal{E}(X \times (Y \times Z)) \\
& & \downarrow \text{assoc} & & \downarrow m_{\varepsilon, X \times (Y \times Z)} \\
X \times (Y \times T\mathcal{E}Z) & \xrightarrow{\quad X \times (Y \times m_{\varepsilon, Z}) \quad} & X \times (Y \times SZ) & \xrightarrow{\quad X \times str_{Y, Z} \quad} & X \times S(Y \times Z) & \xrightarrow{\quad str_{X, Y \times Z} \quad} & S(X \times (Y \times Z)) \\
& \searrow^{X \times str_{\varepsilon, Y, Z}} & \downarrow X \times m_{\varepsilon, Y \times Z} & & \downarrow m_{\varepsilon, X \times (Y \times Z)} \\
& & X \times T\mathcal{E}(Y \times Z) & \xrightarrow{\quad str_{\varepsilon, X, Y \times Z} \quad} & T\mathcal{E}(X \times (Y \times Z)) & \xrightarrow{\quad m_{\varepsilon, X \times (Y \times Z)} \quad} & S(X \times (Y \times Z))
\end{array}$$

where the centre is the associativity law for the strength of S , and other parts of the diagram are from the definition of grading.

The left- and right-unit laws hold because of the following diagram chase

$$\begin{array}{ccccccc}
T\mathcal{E}X & \xrightarrow{\quad m_{\varepsilon, X} \quad} & SX & \xleftarrow{\quad m_{\varepsilon, X} \quad} & T\mathcal{E}X & & \\
\downarrow \eta_{T\mathcal{E}X} & \searrow \eta_{T\mathcal{E}X} & \downarrow \eta_{SX} & \searrow S\eta_X & \downarrow T\mathcal{E}\eta_X & & \\
& & S(T\mathcal{E}X) & \xrightarrow{\quad Sm_{\varepsilon, X} \quad} & S(SX) & \xleftarrow{\quad m_{\varepsilon, SX} \quad} & T\mathcal{E}(SX) \\
& \searrow m_{\varepsilon, T\mathcal{E}X} & \downarrow \mu_X & \searrow \mu_X & \downarrow T\mathcal{E}m_{1, X} & & \\
T1(T\mathcal{E}X) & \xrightarrow{\quad \mu_{1, \varepsilon, X} \quad} & T\mathcal{E}X & \xrightarrow{\quad m_{\varepsilon, X} \quad} & SX & \xleftarrow{\quad m_{\varepsilon, X} \quad} & T\mathcal{E}X & \xleftarrow{\quad \mu_{\varepsilon, 1, X} \quad} & T\mathcal{E}(T1X)
\end{array}$$

The centre is the left- and right-unit laws for S ; the other parts are diagrams from the definition of grading (other than the pentagon on the bottom-right, which combines two of them). \square

To further reduce the amount of work required to construct graded monads, we could also attempt to construct gradings from monads. One method of doing this is given for Gifford-style effect algebras by Kammar and McDermott [41].

4.2.3.5 Term model

The last model we describe is the *term model* of GCBPV, which is a purely syntactic model, for example, the morphisms of the value category are (equivalence classes of) GCBPV values. Terms are interpreted essentially as themselves. We do not use this to model any of our examples, and it is not itself useful for proving the validity of program transformations. However, it is useful for showing that models (such as the examples we give above) are *adequate*. In Section 4.3 we use the term model to prove adequacy.

Recall that for any given inequational theory \leq we write \equiv for the intersection of \leq and its converse. For each typing context Γ and value type A , this is an equivalence relation on closed values of type A in context Γ , and similarly for computation types.

We first define the value category \mathbf{C}_λ . Objects are value types A . Morphisms from A to B are equivalence classes of values $x : A \vdash V : B$. We write $[x : A \vdash V : B]$, or just $[V]$, for such a morphism. The ordering on morphisms is \leq . The identity on A is $[x]$, and the composition of $[V]$ and $[W]$ is $[V[x \mapsto W]]$.

The computation category \mathbf{D}_λ is more difficult because contexts contain only value types, and because we need to be able to define the action \otimes . First we define a notion of *linearity* for computations, similar to linearity for natural transformations (Definition 4.2.4).

Definition 4.2.18 A family $(M^\varepsilon)_{\varepsilon \in \mathcal{E}}$ of computations

$$y : \mathbf{U}(\langle \varepsilon \rangle \underline{C}) \vdash M^\varepsilon : \langle \varepsilon \rangle \underline{D}$$

is *natural* if

$$y : \mathbf{U}(\langle \varepsilon \rangle \underline{C}) \vdash \mathbf{coerce}_{\langle \varepsilon \rangle \underline{C} <: \langle \varepsilon' \rangle \underline{C}} M^\varepsilon \equiv M^{\varepsilon'} [y \mapsto \mathbf{coerce}_{\mathbf{U} \langle \varepsilon \rangle \underline{C} <: \mathbf{U} \langle \varepsilon' \rangle \underline{C}} y] : \langle \varepsilon' \rangle \underline{D}$$

for each $\varepsilon \leq \varepsilon' \in \mathcal{E}$, and is *linear* if

$$z : \mathbf{U}(\langle \varepsilon \rangle (\mathbf{U}(\langle \varepsilon' \rangle \underline{C}))) \vdash \mathbf{force} z \text{ to } y. M^{\varepsilon'} \equiv M^{\varepsilon \cdot \varepsilon'} [y \mapsto \mathbf{think}(\mathbf{force} z \text{ to } w. \mathbf{force} w)] : \langle \varepsilon \cdot \varepsilon' \rangle \underline{D}$$

for each $\varepsilon, \varepsilon' \in \mathcal{E}$. ◀

(This definition of linearity is based on the one given by Munch-Maccagnoni [79] and Levy [60], but here we need natural *families* of computations because of the effect system.) Objects of the computation category \mathbf{D}_λ are computation types \underline{C} . A morphism from \underline{C} to \underline{D} is a family $([M^\varepsilon])_{\varepsilon \in \mathcal{E}}$ of equivalence classes of computations

$$y : \mathbf{U}(\langle \varepsilon \rangle \underline{C}) \vdash M^\varepsilon : \langle \varepsilon \rangle \underline{D}$$

such that $(M^\varepsilon)_{\varepsilon \in \mathcal{E}}$ is natural and linear. The order on morphisms is

$$([M^\varepsilon])_{\varepsilon \in \mathcal{E}} \sqsubseteq ([N^\varepsilon])_{\varepsilon \in \mathcal{E}} \iff \forall \varepsilon \in \mathcal{E}. M^\varepsilon \leq N^\varepsilon$$

We sometimes leave the effect ε implicit and write $[y : \mathbf{U} \underline{C} \vdash M : \underline{D}]$, or just $[M]$, for these morphisms. The identity on \underline{C} at ε is $[y : \mathbf{U}(\langle \varepsilon \rangle \underline{C}) \vdash \mathbf{force} y : \langle \varepsilon \rangle \underline{C}]$ and the composition of $[y : \mathbf{U} \underline{C}_2 \vdash M : \underline{C}_3]$ and $[y : \mathbf{U} \underline{C}_1 \vdash N : \underline{C}_2]$, for each effect, is $[y : \mathbf{U} \underline{C}_1 \vdash M[y \mapsto \mathbf{think} N] : \underline{C}_3]$.

For every morphism in the computation category that we define, each of the equivalence classes of terms is the same up to typing annotations. They can therefore be thought of as single terms that satisfy $y : \mathbf{U} \underline{C} \vdash M : \underline{D}$ and, by replacing typing annotations, can be typed as $y : \mathbf{U}(\langle \varepsilon \rangle \underline{C}) \vdash M^\varepsilon : \langle \varepsilon \rangle \underline{D}$. They are essentially polymorphic in the effect ε . Of course, this is not the case for every computation (just as not every computation is linear), but this is not a problem because computations are interpreted as morphisms in the value category.

Bicartesian structure of \mathbf{C}_λ :

Terminal object	unit	$\langle \rangle_A := [()]$
Binary products	$A_1 \times A_2$	$\pi_1 := [\text{fst } x] \quad \pi_2 := [\text{snd } x]$
Initial object	empty	$[]_A := [\text{case}_A x \text{ of } \{\}]$
Binary coproducts	$A_1 + A_2$	$\text{inl} := [\text{inl}_{A_2} x] \quad \text{inr} := [\text{inr}_{A_1} x]$

Cartesian structure of \mathbf{D}_λ :

Terminal object	<u>unit</u>	$\langle \rangle_{\underline{C}} := [\lambda \{\}]$
Binary products	$\underline{C}_1 \times \underline{C}_2$	$\pi_1 := [1' \text{force } y] \quad \pi_2 := [2' \text{force } y]$

Exponentials:

$$(A \Rightarrow \underline{C}) := (A \rightarrow \underline{C}) \quad (A \Rightarrow [M]) := [\lambda z : A. M[y \mapsto \text{thunk } (z' \text{force } y)]]$$

$$\Lambda[V] := [\text{thunk } (\lambda z : A. \text{force } (V[x \mapsto (x, z)]))]$$

Graded strong **Poset**-adjunction:

$$U : \mathbf{D}_\lambda \rightarrow \mathbf{C}_\lambda \quad U\underline{C} := \underline{C} \quad U[y : \underline{C} \perp M : \underline{C}'] := [x : \underline{C} \vdash \text{thunk } M[y \mapsto x] : \underline{C}']$$

$$F : \mathbf{C}_\lambda \rightarrow \mathbf{D}_\lambda \quad FA := \langle 1 \rangle A \quad F[x : A \vdash V : B] := [y : U \langle \varepsilon \rangle A \perp \text{force } y \text{ to } x. \langle V \rangle : \langle \varepsilon \rangle B]$$

$$\varepsilon' \otimes \underline{C} := \langle \varepsilon' \rangle \underline{C} \quad (\varepsilon' \leq \varepsilon'') \otimes \underline{C} := [\text{coerce}_{\langle \varepsilon' \rangle \underline{C} <: \langle \varepsilon'' \rangle \underline{C}} (\text{force } y)] \quad \varepsilon' \otimes ([M^\varepsilon])_{\varepsilon \in \mathcal{E}} := ([M^{\varepsilon \cdot \varepsilon'}])_{\varepsilon \in \mathcal{E}}$$

$$\eta_A := [\text{thunk } \langle x \rangle] \quad \omega_{\underline{C}} := [\text{force } y \text{ to } z. \text{force } z]$$

$$\text{str}_{\varepsilon, A, B} := [\text{thunk } (\text{force } (\text{fst } x) \text{ to } z. \langle (z, \text{snd } x) \rangle)]$$

Base types, constants and operations:

$$\llbracket b \rrbracket := b \quad \llbracket c \rrbracket := [c] \quad \llbracket \text{op} \rrbracket := [\text{thunk } (\text{op } x)]$$

Figure 4.4: The GCBPV term model

The remaining structure of the term model is defined in Figure 4.4. The interesting part is the definition of the graded strong **Poset**-adjunction. In the definition of the right adjoint U on morphisms, only the equivalence class for the effect 1 is used (similar to the right adjoint in monadic models (Section 4.2.3.4)). The definition of the action \otimes is where it is important that we have natural families of computations as morphisms of \mathbf{D}_λ rather than just individual computations. In particular, the definition of $\varepsilon \otimes$ – on morphisms reindexes the family given (this is why it is the only definition in which we make the indices explicit). Naturality of the counit ω follows from linearity of morphisms in \mathbf{D}_λ .

The term model is in fact a model, as the following lemma shows.

Lemma 4.2.19 Given any inequational theory, the data given above forms a GCBPV structure in which the interpretation of terms satisfies

$$\Gamma \vdash V \leq W : A \Leftrightarrow \llbracket V \rrbracket \sqsubseteq \llbracket W \rrbracket \quad \Gamma \perp M \leq N : \underline{C} \Leftrightarrow \llbracket M \rrbracket \sqsubseteq \llbracket N \rrbracket$$

Hence this GCBPV structure is a GCBPV model.

Proof sketch. Showing that this data forms a GCBPV structure just requires checking each of the requirements in turn. All are easy to check using the core axioms of inequational theories.

The fact about the interpretations of terms holds because terms are interpreted essentially as themselves. Formally, for each typing context Γ , there is a substitution σ_Γ such that

$$x : \llbracket \Gamma \rrbracket \vdash \sigma_\Gamma : \Gamma \quad \llbracket \Gamma \vdash V : A \rrbracket = [V[\sigma_\Gamma]] \quad \llbracket \Gamma \vdash M : \underline{C} \rrbracket = [\mathbf{thunk} M[\sigma_\Gamma]]$$

There is also a substitution σ_Γ^{-1} that behaves as an inverse to σ_Γ . It satisfies

$$\Gamma \vdash \sigma_\Gamma^{-1} : (x : \llbracket \Gamma \rrbracket) \quad V[\sigma_\Gamma][\sigma_\Gamma^{-1}] \equiv V \quad M[\sigma_\Gamma][\sigma_\Gamma^{-1}] \equiv M$$

Hence on computations M we have:

$$\begin{array}{ll} M \leqslant M' & \llbracket M \rrbracket \sqsubseteq \llbracket M' \rrbracket \\ \Rightarrow \mathbf{thunk} M[\sigma_\Gamma] \leqslant \mathbf{thunk} M[\sigma_\Gamma] & \Rightarrow \mathbf{thunk} M[\sigma_\Gamma] \leqslant \mathbf{thunk} M[\sigma_\Gamma] \\ \Rightarrow \llbracket M \rrbracket \sqsubseteq \llbracket M' \rrbracket & \Rightarrow \mathbf{force} \mathbf{thunk} M[\sigma_\Gamma][\sigma_\Gamma^{-1}] \leqslant \mathbf{force} \mathbf{thunk} M[\sigma_\Gamma][\sigma_\Gamma^{-1}] \\ & \Rightarrow M \sqsubseteq M' \end{array}$$

Values are similar. □

This lemma actually shows that the term model is *complete* as well as sound. Completeness is the crucial property of the term model: it allows us to use term model to show adequacy of our other models (Section 4.3 shows how to do this). Moreover, the fact that the syntax of GCBPV forms a model shows that we have not included too many requirements in the definition of model. The definition exactly matches the inequational theory in the following sense:

Corollary 4.2.20 Given any inequational theory:

1. $\llbracket \Gamma \vdash V : A \rrbracket \sqsubseteq \llbracket \Gamma \vdash W : A \rrbracket$ for every model of \leqslant if and only if $\Gamma \vdash V \leqslant W : A$.
2. $\llbracket \Gamma \vdash M : \underline{C} \rrbracket \sqsubseteq \llbracket \Gamma \vdash N : \underline{C} \rrbracket$ for every model of \leqslant if and only if $\Gamma \vdash M \leqslant N : \underline{C}$. ◀

4.3 Relating syntax and semantics

The previous sections describe examples of program transformations we would like to perform by conjecturing instances $M \leqslant_{\text{ctx}} N$ of the contextual preorder, and also show for each of these instances that $\llbracket M \rrbracket \sqsubseteq \llbracket N \rrbracket$. To prove the validity of the program transformations, we show that our example models are adequate:

Definition 4.3.1 (Adequacy) A model of some GCBPV inequational theory is *adequate* if both of the following hold:

- (Values) $\llbracket \Gamma \vdash V : A \rrbracket \sqsubseteq \llbracket \Gamma \vdash W : A \rrbracket$ implies $\Gamma \vdash V \leqslant_{\text{ctx}} W : A$.
- (Computations) $\llbracket \Gamma \vdash M : \underline{C} \rrbracket \sqsubseteq \llbracket \Gamma \vdash N : \underline{C} \rrbracket$ implies $\Gamma \vdash M \leqslant_{\text{ctx}} N : \underline{C}$. ◀

Adequacy can be formulated in terms of the more general problem of *abstraction* between models of GCBPV.

Definition 4.3.2 (Abstraction) Suppose that \mathcal{M}_1 and \mathcal{M}_2 are models of the same inequational theory. We say that \mathcal{M}_2 *abstracts* \mathcal{M}_1 if for all ground types G , effects ε and closed computations $M, N : \langle \varepsilon \rangle G$ we have

$$\mathcal{M}_1 \llbracket M \rrbracket \sqsubseteq \mathcal{M}_1 \llbracket N \rrbracket \quad \Rightarrow \quad \mathcal{M}_2 \llbracket M \rrbracket \sqsubseteq \mathcal{M}_2 \llbracket N \rrbracket \quad \leftarrow$$

Adequacy is a special case of abstraction because of the following lemma.

Lemma 4.3.3 A model \mathcal{M} of some inequational theory is adequate if and only if the term model (Section 4.2.3.5) abstracts \mathcal{M} .

Proof. Since the term model is complete (Lemma 4.2.19), it abstracts \mathcal{M} if and only if for all ground types G , effects ε and closed computations $M, N : \langle \varepsilon \rangle G$ we have

$$\mathcal{M}[\![M]\!] \sqsubseteq \mathcal{M}[\![N]\!] \quad \Rightarrow \quad M \leqslant N$$

so it suffices to show that this property is equivalent to adequacy.

Suppose that this property holds and that $\mathcal{M}[\![V]\!] \sqsubseteq \mathcal{M}[\![W]\!]$ for well-typed values V, W . By compositionality (Lemma 4.2.6 (1)) we have $\mathcal{M}[\![\underline{C}[V]]\!] \sqsubseteq \mathcal{M}[\![\underline{C}[W]]\!]$. So if $\underline{C}[V]$ and $\underline{C}[W]$ satisfy the conditions on M and N we have $\underline{C}[V] \leqslant_{\text{ctx}} \underline{C}[W]$. Hence $V \leqslant_{\text{ctx}} W$. Similarly for computations.

If \mathcal{M} is adequate and $\mathcal{M}[\![M]\!] \sqsubseteq \mathcal{M}[\![N]\!]$ then $M \leqslant_{\text{ctx}} N$, which implies $M \leqslant N$ by applying the definition of the contextual preorder to the term context $\underline{\square}$. \square

We describe a method for proving instances of abstraction, and then show that our example models are adequate using Lemma 4.3.3.

4.3.1 Logical relations and abstraction

To prove abstraction between models, we use logical relations. We give a general categorical description of logical relations for GCBPV that can be used to state sufficient conditions for abstraction to hold. This is based on the work of Hermida [36] and Katsumata [44], adapted to our **Poset**-enriched models. This description of logical relations is more general than the one we gave in Section 3.1: instead of covering just binary relations between terms, our new definition covers relations of other arities (e.g. unary relations, and relations of varying arity), and are not restricted to relating terms (they can relate denotations).

These relations are formulated in terms of functors $p : \mathbf{R}_C \rightarrow \mathbf{C}$ from an (ordinary) category \mathbf{R}_C into (the underlying ordinary category of) a **Poset**-category \mathbf{C} . Here \mathbf{C} is either the value or computation category of a model of GCBPV (we actually require two such functors: one for the value category and one for the computation category). An object $P \in \mathbf{R}_C$ should be thought of as a unary predicate on $pP \in \mathbf{C}$, and a morphism $\hat{f} : P \rightarrow Q$ as a witness that $p\hat{f} : pP \rightarrow pQ$ is truth-preserving. We assume that the functor p is faithful (i.e. for all $\hat{f}, \hat{g} : P \rightarrow Q$ if $p\hat{f} = p\hat{g}$ then $\hat{f} = \hat{g}$), which implies such witnesses are unique. We say that an object $P \in \mathbf{R}_C$ is *above* $X \in \mathbf{C}$ if $pP = X$, and for $f : pP \rightarrow pQ$ we write $f : P \dot{\rightarrow} Q$ to mean there is some \hat{f} such that $p\hat{f} = f$ (i.e. f is truth-preserving).

The general idea is to *lift* the structure of some GCBPV model (in **Poset**-categories \mathbf{C} and \mathbf{D}) to categories \mathbf{R}_C and \mathbf{R}_D of relations, and hence construct a new *logical-relations* (LR) model (we define LR model below). In the LR model, value types A are interpreted as relations $\mathcal{R}[\![A]\!] \in \mathbf{R}_C$ and computation types \underline{C} are interpreted as relations $\mathcal{R}[\![\underline{C}]\!] \in \mathbf{R}_D$.

The next definition captures much of the required structure. It is essentially the same as a definition given by Ma and Reynolds [65].

Definition 4.3.4 Suppose that \mathbf{C} is a cartesian **Poset**-category and $p : \mathbf{R}_C \rightarrow \mathbf{C}$ is a functor. A *lifting* of the cartesian structure of \mathbf{C} consists of:

- An \mathbf{R}_C -object $\dot{1}$ above 1, such that for each P , we have $\langle \rangle_{pP} : P \dot{\rightarrow} \dot{1}$.
- For each P_1, P_2 , an \mathbf{R}_C -object $P_1 \dot{\times} P_2$ above $pP_1 \times pP_2$ such that $\pi_1 : P_1 \dot{\times} P_2 \dot{\rightarrow} P_1$, $\pi_2 : P_1 \dot{\times} P_2 \dot{\rightarrow} P_2$, and for each P, f_1, f_2 if $f_1 : Q \dot{\rightarrow} P_1$ and $f_2 : Q \dot{\rightarrow} P_2$ then $\langle f_1, f_2 \rangle : Q \dot{\rightarrow} P_1 \dot{\times} P_2$.

If \mathbf{C} is bicartesian, a *lifting* of the bicartesian structure additionally consists of:

- An $\mathbf{R}_\mathbf{C}$ -object $\dot{0}$ above 0 , such that for each P , we have $[\]_{pP} : \dot{0} \dot{\rightarrow} P$.
- For each P_1, P_2 , an $\mathbf{R}_\mathbf{C}$ -object $P_1 \dot{+} P_2$ above $pP_1 + pP_2$ such that $inl : P_1 \dot{\rightarrow} P_1 \dot{+} P_2$, $inr : P_2 \dot{\rightarrow} P_1 \dot{+} P_2$, and for each P, f_1, f_2 if $f_1 : P_1 \dot{\rightarrow} Q$ and $f_2 : P_2 \dot{\rightarrow} Q$ then $[f_1, f_2] : P_1 \dot{+} P_2 \dot{\rightarrow} Q$.

Finally, if \mathbf{C} is distributive, a *lifting* of the distributive structure is a lifting of the bicartesian structure such that $dist_{pP, pQ, pR} : pP \dot{\times} (pQ \dot{+} pR) \dot{\rightarrow} (pP \dot{\times} pQ) \dot{+} (pP \dot{\times} pR)$ for all P, Q, R . ◀

Example 4.3.5 For **Poset**, we define the category **Pred** as follows. Objects $P \subseteq X$ are pairs of a poset X and an arbitrary set P of elements of X ; we call these *predicates*. The predicate $P \subseteq X$ is considered to be true on elements of P , and false on the other elements of X . Morphisms $f : (P \subseteq X) \rightarrow (Q \subseteq Y)$ are monotone functions $f : X \rightarrow Y$ such that $f x \in Q$ for all $x \in P$. There is a functor $p : \mathbf{Pred} \rightarrow \mathbf{Poset}$ that sends $(P \subseteq X)$ to X and f to itself. In this case, $f : (P \subseteq X) \dot{\rightarrow} (Q \subseteq Y)$ holds iff $x \in P$ implies $f x \in Q$ for all x . We lift the bicartesian structure of **Poset** as follows:

$$\begin{aligned} \dot{1} &:= \{\star\} \subseteq 1 & (P_1 \subseteq X_1) \dot{\times} (P_2 \subseteq X_2) &:= (P_1 \times P_2) \subseteq (X_1 \times X_2) \\ \dot{0} &:= \emptyset \subseteq 0 & (P_1 \subseteq X_1) \dot{+} (P_2 \subseteq X_2) &:= (P_1 + P_2) \subseteq (X_1 + X_2) \end{aligned} \quad \blacktriangleleft$$

Given a faithful functor $p : \mathbf{R}_\mathbf{C} \rightarrow \mathbf{C}$ we **Poset**-enrich $\mathbf{R}_\mathbf{C}$ in the following way: if $\dot{f}, \dot{g} : P \rightarrow Q$ are morphisms in $\mathbf{R}_\mathbf{C}$ then $\dot{f} \sqsubseteq \dot{g}$ if $p\dot{f} \sqsubseteq p\dot{g}$ (this is a partial order because f is faithful). Any lifting of the cartesian (or bicartesian) structure on \mathbf{C} induces a cartesian (bicartesian) structure on $\mathbf{R}_\mathbf{C}$: we use the objects and witnesses required in the definition of lifting (e.g. the product of P_1 and P_2 is $P_1 \dot{\times} P_2$, where each projection $\dot{\pi}_i : P_1 \dot{\times} P_2 \rightarrow P_i$ is the witness of $\pi_i : P_1 \dot{\times} P_2 \dot{\rightarrow} P_i$).

We also require liftings of certain **Poset**-functors.

Definition 4.3.6 Suppose that $p : \mathbf{R}_\mathbf{C} \rightarrow \mathbf{C}$ and $q : \mathbf{R}_\mathbf{D} \rightarrow \mathbf{D}$ are functors. A *lifting* of a **Poset**-functor $G : \mathbf{C} \rightarrow \mathbf{D}$ consists of an object $\dot{G}P \in \mathbf{R}_\mathbf{D}$ that satisfies $q(\dot{G}P) = G(pP)$ for each $P \in \mathbf{R}_\mathbf{C}$, such that for each f if $f : P \dot{\rightarrow} Q$ then $Gf : \dot{G}P \dot{\rightarrow} \dot{G}Q$. ◀

The data required for a lifting of a **Poset**-functor $G : \mathbf{C} \rightarrow \mathbf{D}$ forms a **Poset**-functor $\dot{G} : \mathbf{R}_\mathbf{C} \rightarrow \mathbf{R}_\mathbf{D}$ such that $G \circ p = q \circ \dot{G}$.

In the following definitions, we suppose we are given some model \mathcal{M} of GCBPV with value category \mathbf{C} and computation category \mathbf{D} . To form an LR model we need to be able to lift the graded adjunction.

Definition 4.3.7 Suppose that $p : \mathbf{R}_\mathbf{C} \rightarrow \mathbf{C}$ and $q : \mathbf{R}_\mathbf{D} \rightarrow \mathbf{D}$ are functors, and that $\mathbf{R}_\mathbf{C}$ has an object $P_1 \dot{\times} P_2$ such that $p(P_1 \dot{\times} P_2) = pP_1 \times pP_2$ for each pair of objects $P_1, P_2 \in \mathbf{R}_\mathbf{C}$. A *lifting* of the graded strong **Poset**-adjunction (F, U, \otimes) consists of liftings of the **Poset**-functors $F : \mathbf{C} \rightarrow \mathbf{D}, U : \mathbf{D} \rightarrow \mathbf{C}$ and $\varepsilon \otimes - : \mathbf{D} \rightarrow \mathbf{D}$ for each $\varepsilon \in \mathcal{E}$, such that

- For each $\varepsilon, \varepsilon'$ we have $(\varepsilon \cdot \varepsilon') \otimes - = \varepsilon \otimes (\varepsilon' \otimes -)$, and $1 \otimes -$ is the identity.
- For each ε, P, Q we have $str_{\varepsilon, pP, pQ} : P \dot{\times} \dot{U}(\varepsilon \otimes \dot{F}Q) \dot{\rightarrow} \dot{U}(\varepsilon \otimes \dot{F}(P \dot{\times} Q))$.
- For each $\varepsilon, \varepsilon', Q$, if $\varepsilon \leq \varepsilon'$ then $(\varepsilon \leq \varepsilon') \otimes qQ : (\varepsilon \otimes Q) \dot{\rightarrow} (\varepsilon' \otimes Q)$.
- For each P we have $\eta_{pP} : P \dot{\rightarrow} \dot{U}\dot{F}P$.
- For each Q we have $\omega_{qQ} : \dot{F}\dot{U}Q \dot{\rightarrow} Q$. ◀

Each lifting of (F, U, \otimes) determines a graded strong **Poset**-adjunction $(\dot{F}, \dot{U}, \otimes)$ on \mathbf{D} , where each of the required morphisms is given by a witness. In general, the correct choice of lifting depends on the specific effects being modelled, though some general techniques for constructing

liftings of monads (in the non-enriched case) exist [32, 43, 46]. We also expect that a version of the *free lifting* (Section 3.1.1) could be developed for graded adjunctions (this would be similar to the monadic version given by Kammar and McDermott [41]). We do not use these general techniques, but instead construct liftings specific to our examples (see Section 4.3.2).

Recall from Section 3.1 that the crucial part of the definition of a logical relation is its definition on returner types. Since left adjoints are unique, the definition on returner types is uniquely determined by the right adjoint in this case. The crucial part of the definition here is the choice of \mathbf{R}_D .

We collect the data required to form an LR model \mathcal{R} in the following definition. (We write $\mathcal{M}[\![-]\!]$ for interpretations in the base model \mathcal{M} , and $\mathcal{R}[\![-]\!]$ for interpretations in the logical-relations model \mathcal{R} .)

Definition 4.3.8 A *logical-relations* (LR) model above the model \mathcal{M} consists of

- Categories \mathbf{R}_C and \mathbf{R}_D and functors $p : \mathbf{R}_C \rightarrow \mathbf{C}$ and $q : \mathbf{R}_D \rightarrow \mathbf{D}$, together with liftings of the distributive structure of \mathbf{C} and the cartesian structure of \mathbf{D} .
- For each $P \in \mathbf{R}_C$, a lifting $P \dot{\Rightarrow} - : \mathbf{R}_D \rightarrow \mathbf{R}_D$ of the **Poset**-functor $pP \Rightarrow - : \mathbf{D} \rightarrow \mathbf{D}$, such that for all f, f' if $f : R \dot{\times} P \dot{\rightarrow} \dot{U}Q$ then $\Lambda f : R \dot{\rightarrow} \dot{U}(P \dot{\Rightarrow} Q)$ and if $f' : R \dot{\rightarrow} \dot{U}(P \dot{\Rightarrow} Q)$ then $\Lambda^{-1} f' : R \dot{\times} P \dot{\rightarrow} \dot{U}Q$.
- A lifting of the graded strong **Poset**-adjunction.
- A \mathbf{R}_C -object $\mathcal{R}[\![b]\!]$ above $\mathcal{M}[\![b]\!]$ for each base type b .

such that $\mathcal{M}[\![c]\!] : \dot{1} \dot{\rightarrow} \mathcal{R}[\![A]\!]$ for each constant $c \in \mathcal{K}_A$ and

$$\mathcal{M}[\![\text{op}]\!] : \mathcal{R}[\![\text{car}_{\text{op}}]\!] \dot{\rightarrow} \dot{U}(\text{eff}_{\text{op}} \dot{\otimes} \dot{F}(\mathcal{R}[\![\text{ar}_{\text{op}}]\!]))$$

for each operation $\text{op} \in \Sigma$. ◀

Soundness of any such model follows from the definition of the order on morphisms in \mathbf{R}_C and \mathbf{R}_D and soundness of the base model \mathcal{M} .

As usual, the logical relation satisfies a *fundamental lemma*, which states that interpretations of terms preserve relations.

Lemma 4.3.9 (Fundamental) If \mathcal{R} is an LR model over \mathcal{M} then

1. If $\Gamma \vdash V : A$ then $\mathcal{M}[\![V]\!] : \mathcal{R}[\![\Gamma]\!] \dot{\rightarrow} \mathcal{R}[\![A]\!]$.
2. If $\Gamma \vdash M : \underline{C}$ then $\mathcal{M}[\![M]\!] : \mathcal{R}[\![\Gamma]\!] \dot{\rightarrow} \dot{U}(\mathcal{R}[\![\underline{C}]\!])$.

Proof. By a trivial induction, the interpretations $\mathcal{R}[\![V]\!]$ and $\mathcal{R}[\![M]\!]$ in the LR model satisfy $p(\mathcal{R}[\![V]\!]) = \mathcal{M}[\![V]\!]$ and $p(\mathcal{R}[\![M]\!]) = \mathcal{M}[\![M]\!]$, and hence are the required witnesses. □

To prove abstraction, we need to relate two models. Given two **Poset**-categories \mathbf{C}_1 and \mathbf{C}_2 , define $\mathbf{C}_1 \times \mathbf{C}_2$ as the **Poset**-category in which objects are pairs (X_1, X_2) with $X_1 \in \mathbf{C}_1$ and $X_2 \in \mathbf{C}_2$, and morphisms (f_1, f_2) are similarly pairs of morphisms (ordered componentwise). Given models \mathcal{M}_1 (with value category \mathbf{C}_1 and computation category \mathbf{D}_1) and \mathcal{M}_2 (with value category \mathbf{C}_2 and computation category \mathbf{D}_2), we define their *product* $\mathcal{M}_1 \times \mathcal{M}_2$ as the model with value category $\mathbf{C}_1 \times \mathbf{C}_2$ and computation category $\mathbf{D}_1 \times \mathbf{D}_2$, where each part of the structure is given by pairing the structure in \mathcal{M}_1 and \mathcal{M}_2 . The product model satisfies $(\mathcal{M}_1 \times \mathcal{M}_2)[\![-]\!] = (\mathcal{M}_1[\![-]\!], \mathcal{M}_2[\![-]\!])$.

To relate \mathcal{M}_1 to \mathcal{M}_2 , we consider *binary* functors $p : \mathbf{R}_{\mathbf{C}_1 \times \mathbf{C}_2} \rightarrow \mathbf{C}_1 \times \mathbf{C}_2$, in which an object $R \in \mathbf{R}_{\mathbf{C}_1 \times \mathbf{C}_2}$ above (X_1, X_2) relates X_1 to X_2 . A key part of the proof of abstraction is to show that certain relations are *monotone*:

Definition 4.3.10 Suppose that \mathbf{C}_1 and \mathbf{C}_2 are **Poset**-categories and $p : \mathbf{R}_{\mathbf{C}_1 \times \mathbf{C}_2} \rightarrow \mathbf{C}_1 \times \mathbf{C}_2$ is a faithful functor. An object $R \in \mathbf{D}$ is *monotone* if whenever $(f_1, f_2) : \dot{1} \dot{\rightarrow} R$ and $(g_1, g_2) : \dot{1} \dot{\rightarrow} R$,

$$f_1 \sqsubseteq g_1 \quad \Rightarrow \quad f_2 \sqsubseteq g_2 \quad \blacktriangleleft$$

Note the similarity with the implication in the definition of abstraction (Definition 4.3.2). The relations $\dot{1}$ and $\dot{0}$ are monotone (the latter because we assume distributivity⁷), and if R_1, R_2 are monotone then so is $R_1 \dot{\times} R_2$ (if these relations exist). We say that monotone relations are *closed under coproducts* if $R_1 \dot{+} R_2$ is monotone for monotone R_1, R_2 (this does not hold in general).

Example 4.3.11 The functor $\mathbf{BRel} \rightarrow \mathbf{Poset} \times \mathbf{Poset}$, where \mathbf{BRel} is similar to \mathbf{Pred} except that objects are binary relations $R \subseteq X_1 \times X_2$, could be used to relate two models on the value category \mathbf{Poset} . In this case, monotonicity of R just means that if $(a_1, a_2), (b_1, b_2) \in R$ then $a_1 \sqsubseteq b_1$ implies $a_2 \sqsubseteq b_2$. \blacktriangleleft

Using this definition of monotonicity of relations, we give sufficient conditions for one model of GCBPV to abstract another. This is the key lemma that allows us to prove abstraction. We then use it to prove adequacy of our example models.

Lemma 4.3.12 (Abstraction) Suppose that \mathcal{M}_1 and \mathcal{M}_2 are models of the same inequational theory. If there exists an LR model \mathcal{R} above $\mathcal{M}_1 \times \mathcal{M}_2$ such that each $\mathcal{R}[\![b]\!]$ is monotone, $\dot{U}(\varepsilon \dot{\otimes} \dot{F}R)$ is monotone for each effect ε and monotone R , and monotone relations in $\mathbf{R}_{\mathbf{C}}$ are closed under coproducts, then \mathcal{M}_2 abstracts \mathcal{M}_1 .

Proof. An easy induction on ground types G shows that each $\mathcal{R}[\![G]\!]$ is monotone, and hence so is $\dot{U}(\mathcal{R}[\![\langle \varepsilon \rangle G]\!]) = \dot{U}(\varepsilon \dot{\otimes} \mathcal{R}[\![G]\!])$. By the fundamental lemma (Lemma 4.3.9), if $M, N : \langle \varepsilon \rangle G$ are closed computations then

$$(\mathcal{M}_1[\![M]\!], \mathcal{M}_2[\![M]\!]), (\mathcal{M}_1[\![N]\!], \mathcal{M}_2[\![N]\!]) : \dot{1} \dot{\rightarrow} \mathcal{R}[\![\langle \varepsilon \rangle G]\!]$$

and the result follows from monotonicity. \square

4.3.2 Examples

We now show that the models we use for each of our three examples are adequate, by using Lemma 4.3.12 to show that each example model \mathcal{M} abstracts the term model.

We therefore need three LR models, one for each of our example models. Recall that each example model uses the value category \mathbf{Poset} , and the value category of the term model is \mathbf{C}_λ . Also recall that we write \mathbf{Term}_A and \mathbf{Term}_C for the sets of equivalence classes of closed values and computations (we also omit square brackets around equivalence classes in this section). We use a common construction for $\mathbf{R}_{\mathbf{Poset} \times \mathbf{C}_\lambda}$:

- Objects are triples (X, A, R) where X is a poset, A is a value type, and $R \subseteq X \times \mathbf{Term}_A$ is a relation.
- Morphisms $(f, V) : (X, A, R) \rightarrow (Y, B, R')$ are pairs of a morphism $f : X \rightarrow Y$ in \mathbf{Poset} and a morphism $W : A \rightarrow B$ in \mathbf{C}_λ , such that if $(x, N) \in R$ then $(fx, W[x \mapsto V]) \in R'$.

⁷To show that $\dot{0}$ is monotone, we use the fact that distributivity of binary coproducts implies [15, Proposition 3.4] the initial object is *strict* (all morphisms into 0 are isomorphisms), so any two morphisms $1 \rightarrow 0$ are equal.

We have a functor $\mathbf{R}_{\mathbf{Poset} \times \mathbf{C}_\lambda} \rightarrow \mathbf{Poset} \times \mathbf{C}_\lambda$ that sends (X, A, R) to (X, A) , and is the identity on morphisms. For the lifting of the bicartesian structure, we take the bicartesian structure of \mathbf{Poset} and \mathbf{C}_λ , together with the following relations:

$$\begin{aligned} \dot{1} &:= \{(\star, ())\} & R_1 \dot{\times} R_2 &:= \{(x, V) \mid (\pi_1 x, \mathbf{fst} V) \in R_1 \wedge (\pi_2 x, \mathbf{snd} V) \in R_2\} \\ \dot{0} &:= \emptyset & R_1 \dot{+} R_2 &:= \{(\mathbf{inl} x, \mathbf{inl} V) \mid (x, V) \in R_1\} \cup \{(\mathbf{inr} x, \mathbf{inr} V) \mid (x, V) \in R_2\} \end{aligned}$$

Note in particular that $\mathcal{R}[\mathbf{bool}] = \{(\mathbf{true}, \mathbf{true}), (\mathbf{false}, \mathbf{false})\}$, and that (X, A, R) is monotone (Definition 4.3.10) if $x \sqsubseteq x'$ implies $V \leq V'$ for all $(x, V), (x', V') \in R$. Monotone relations are closed under coproducts in this case, so we have one of the requirements of the abstraction lemma (Lemma 4.3.12).

Each of our examples has a different computation category, and hence we use different categories of relations on computations in each case. We go through each example in turn.

4.3.2.1 Undefined behaviour

For undefined behaviour, recall that we use the trivial effect algebra, and hence morphisms in the computation category \mathbf{D}_λ are linear (Definition 4.2.18) equivalence classes of computations $y : \mathbf{U} \underline{C} \perp N : \underline{D}$. We define $\mathbf{R}_{\mathbf{Poset}_\perp \times \mathbf{D}_\lambda}$ by:

- Objects (X, \underline{C}, R) consist of a pointed poset X , a computation type \underline{C} , and a relation $R \subseteq X \times \mathbf{Term}_{\underline{C}}$, such that $(\perp, \mathbf{undef}_{\underline{C}}) \in R$.
- Morphisms $(f, N) : (X, \underline{C}, R) \rightarrow (Y, \underline{D}, R')$ are pairs of a morphism $f : X \rightarrow Y$ in \mathbf{Poset}_\perp and a morphism $N : \underline{C} \rightarrow \underline{D}$ in \mathbf{D}_λ , such that if $(x, M) \in R$ then $(f x, N[y \mapsto \mathbf{thunk} M]) \in R'$.

The crucial part of this definition is that \perp is required to be related to $\mathbf{undef}_{\underline{C}}$. This ensures the counit lifts.

The functor $\mathbf{R}_{\mathbf{Poset}_\perp \times \mathbf{D}_\lambda} \rightarrow \mathbf{Poset}_\perp \times \mathbf{D}_\lambda$ forgets the relation (maps (X, \underline{C}, R) to (X, \underline{C}) and (f, N) to itself). The cartesian structure on $\mathbf{R}_{\mathbf{Poset}_\perp \times \mathbf{D}_\lambda}$, and exponentials, are given by the corresponding structures in the underlying models and the relations

$$\begin{aligned} R_1 \dot{\perp} R_2 &:= \{(\star, \lambda\{\})\} & R_1 \dot{\times} R_2 &:= \{(x, M) \mid (\pi_1 x, 1'M) \in R_1 \wedge (\pi_2 x, 2'M) \in R_2\} \\ R \dot{\Rightarrow} R' &:= \{(f, M) \mid \forall (x, V) \in R. (f x, V'M) \in R'\} \end{aligned}$$

To lift the graded adjunction, we define $\dot{U}(X, \underline{C}, R_1) := (X, \mathbf{U} \underline{C}, R'_1)$ and $\dot{F}(X, A, R_2) := (X_\perp, \langle \star \rangle A, R'_2)$ where

$$R'_1 := \{(x, \mathbf{thunk} M) \mid (x, M) \in R_1\} \quad R'_2 := \{(x, \langle V \rangle) \mid (x, V) \in R_2\} \cup \{(\perp, \mathbf{undef}_{\langle \star \rangle A})\}$$

The idea behind the definition of the left adjoint \dot{F} is that each computation of returner type, when executed, will either return some value V , or have undefined behaviour. The lifting of the action \otimes is trivial. To complete the LR model we define the interpretation of the base type \mathbf{int} to have relation $\{(n, \underline{n}) \mid n \in \llbracket \mathbf{int} \rrbracket\}$. The constants \mathbf{add} and \mathbf{geq} lift because of the axioms of the inequational theory that specify their behaviour. Similarly for the operation $\mathbf{add}_{\mathbf{nsw}}$. The operation \mathbf{undef} lifts because the computation type $\langle \star \rangle \mathbf{empty}$ is interpreted as the relation $\{(\perp, \mathbf{undef}())\}$. We therefore have a LR model that we can use to relate our model of undefined behaviour to the term model.

To apply the abstraction lemma (Lemma 4.3.12), we just need to check that $\mathcal{R}[\mathbf{int}]$ is monotone, which is trivial, and that $\dot{U}(\dot{F}R)$ is monotone for monotone R . For the latter, suppose that (x, M) and (x', M') are in

$$\{(x, \langle V \rangle) \mid (x, V) \in R\} \cup \{(\perp, \mathbf{undef}_{\underline{C}})\}$$

and $x \sqsubseteq x'$. Then if $(x, M) = (x, \langle V \rangle)$ for $(x, V) \in R$ we have $(x', M) = (x', \langle V' \rangle)$ with $(x', V') \in R$ (because $x' \neq \perp$), so $M = \langle V \rangle \leq \langle V' \rangle = M'$ because R is monotone. If $(x, M) = (\perp, \mathbf{undef}_{\underline{C}})$ then we have $M \leq M'$ as one of the axioms of the inequational theory. So the abstraction lemma tells us the our model of undefined behaviour is adequate, and hence that the example program transformations we give in Section 4.1.1 are valid instances of the contextual preorder.

4.3.2.2 Nondeterminism

Recall that for nondeterminism we use the effects $1 \leq +$. We define the computation category $\mathbf{R}_{\mathbf{D} \times \mathbf{D}_\lambda}$ of the logical relations by:

- Objects (X, \underline{C}, R) consist of an object $X = (X_1 \sqsubseteq X_+)$ in \mathbf{D} (i.e. a meet-semilattice X_+ with a subset X_1), a computation type \underline{C} , and a family of relations $(R_\varepsilon \subseteq X_\varepsilon \times \mathbf{Term}_{\langle \varepsilon \rangle \underline{C}})_{\varepsilon \in \{1, +\}}$, such that $(x, M) \in R_1$ implies $(x, \mathbf{coerce}_{\underline{C} \leftarrow \langle + \rangle \underline{C}} M) \in R_+$, and $(x_1, M_1), (x_2, M_2) \in R_+$ implies $(x_1 \sqcap x_2, M_1 \mathbf{or} M_2) \in R_+$.
- Morphisms $(f, N) : (X, \underline{C}, R) \rightarrow (Y, \underline{D}, R')$ are pairs of a morphism $f : X \rightarrow Y$ in \mathbf{D} and a morphism $N : \underline{C} \rightarrow \underline{D}$ in \mathbf{D}_λ , such that $(x, M) \in R_\varepsilon$ implies $(fx, N^\varepsilon[y \mapsto \mathbf{thunk} M]) \in R'_\varepsilon$ for $\varepsilon \in \{1, +\}$.

The relation R_1 relates deterministic computations (of type \underline{C}), and R_+ relates nondeterministic computations (of type $\langle + \rangle \underline{C}$). The condition that meets are related ensures that the counit of the graded adjunction lifts. The functor $\mathbf{R}_{\mathbf{D} \times \mathbf{D}_\lambda} \rightarrow \mathbf{D} \times \mathbf{D}_\lambda$ forgets the relation (maps (X, \underline{C}, R) to (X, \underline{C}) and (f, N) to itself). The cartesian structure on $\mathbf{R}_{\mathbf{D} \times \mathbf{D}_\lambda}$, and exponentials, are defined in the same way as for undefined behaviour (on each effect ε).

To lift the graded adjunction, we define $\dot{U}(X, \underline{C}, R) := (X_1, \mathbf{UC}, R')$ where R' is defined in terms of the relation on deterministic computations:

$$R' := \{(x, \mathbf{thunk} M) \mid (n, M) \in R_1\}$$

The lifting of the left adjoint is $\dot{F}(X, A, T) := (FX, \langle 1 \rangle A, T')$, where

$$T'_1 := \{(\uparrow\{x\}, \langle V \rangle) \mid (x, V) \in T\}$$

$$T'_+ := \{(\uparrow\{x_1, \dots, x_n\}, \mathbf{coerce}_{1 \leq +} \langle V_1 \rangle \mathbf{or} \dots \mathbf{or} \mathbf{coerce}_{1 \leq +} \langle V_n \rangle) \mid \forall i. (x_i, V_i) \in T\}$$

In this definition, the order in which the computations V_i appear in the nondeterministic choice is irrelevant because \mathbf{or} is commutative (and we consider computations up to \equiv). Similarly, duplicates and bracketing are irrelevant because of idempotence and associativity. For the lifting of the monoid action, $1 \otimes -$ is the identity, and for $(+) \otimes -$ we map the family $(\mathcal{R}_\varepsilon)_{\varepsilon \in \{1, +\}}$ to $(\mathcal{R}_+)_{\varepsilon \in \{1, +\}}$ (forgetting the relation on deterministic computations). This data forms an LR model.

To apply the abstraction lemma (Lemma 4.3.12) it suffices to show that the left adjoint preserves monotonicity of relations. We sketch the key part of this proof, which is showing that if T above is monotone, and $(S, M), (S', M') \in T'_+$ and $S \supseteq S'$, then $M \leq M'$. Because of the definition of T'_+ , there are lists of pairs (x_i, V_i) and (x'_i, V'_i) in T such that:

$$\begin{aligned} S &= \uparrow\{x_1, \dots, x_m\} & M &\equiv \mathbf{coerce}_{1 \leq +} \langle V_1 \rangle \mathbf{or} \dots \mathbf{or} \mathbf{coerce}_{1 \leq +} \langle V_m \rangle \\ S' &= \uparrow\{x'_1, \dots, x'_n\} & M' &\equiv \mathbf{coerce}_{1 \leq +} \langle V'_1 \rangle \mathbf{or} \dots \mathbf{or} \mathbf{coerce}_{1 \leq +} \langle V'_n \rangle \end{aligned}$$

Now $S \supseteq S'$ implies that for each i there is some ϕ_i such that $x_{\phi_i} \sqsubseteq x'_i$. Monotonicity of T implies $V_{\phi_i} \leq V'_i$. Hence we have the following, where the first instance of \leq follows from $M_1 \mathbf{or} M_2 \leq M_j$ (as well as commutativity, associativity and idempotence of \mathbf{or}).

$$M \leq \mathbf{coerce}_{1 \leq +} \langle V_{\phi_1} \rangle \mathbf{or} \dots \mathbf{or} \mathbf{coerce}_{1 \leq +} \langle V_{\phi_n} \rangle \leq M'$$

4.3.2.3 Shared global state

For shared global state we use a Gifford-style effect algebra, so that effects are subsets $\varepsilon \subseteq \{\text{get}, \text{put}\}$, and objects of the computation category \mathbf{GSMnem} are graded shared mnemoids. We define the computation category $\mathbf{R}_{\mathbf{GSMnem} \times \mathbf{D}_\lambda}$ of the LR model by:

- Objects (X, \underline{C}, R) consist of a graded shared mnemoid X (we write X_ε for the individual sets, and g and p for the get and put functions), a computation type \underline{C} , and a family of relations $(R_\varepsilon \subseteq X_\varepsilon \times \underline{\text{Term}}_{\langle\langle\varepsilon\rangle\rangle \underline{C}})_{\varepsilon \subseteq \{\text{get}, \text{put}\}}$, such that:
 - If $(x, M) \in R_\varepsilon$ and $\varepsilon \subseteq \varepsilon'$ then $(x, \text{coerce}_{\langle\langle\varepsilon\rangle\rangle \underline{C} : \langle\langle\varepsilon'\rangle\rangle \underline{C}} M) \in R_{\varepsilon'}$.
 - If $(x_{\text{true}}, M_{\text{true}}), (x_{\text{false}}, M_{\text{false}}) \in R_\varepsilon$ then

$$(g(x_{\text{true}}, x_{\text{false}}), \text{get } () \text{ to } b. \underline{\text{if}} \ b \ \text{then } M_{\text{true}} \ \text{else } M_{\text{false}}) \in R_{\varepsilon \cup \{\text{get}\}}$$
 - If $(x, M) \in R_\varepsilon$ and $b \in \{\text{true}, \text{false}\}$ then $(p_b x, \text{put } b \ \text{to } _ . M) \in R_{\varepsilon \cup \{\text{put}\}}$.
- Morphisms $(f, N) : (X, \underline{C}, R) \rightarrow (Y, \underline{D}, R')$ are pairs of a morphism $f : X \rightarrow Y$ in \mathbf{GSMnem} and a morphism $N : \underline{C} \rightarrow \underline{D}$ in \mathbf{D}_λ , such that $(x, M) \in R_\varepsilon$ implies $(fx, N^\varepsilon[y \mapsto \text{thunk } M]) \in R'_\varepsilon$ for $\varepsilon \subseteq \{\text{get}, \text{put}\}$.

This is very similar to the corresponding category for our nondeterminism example; much of the rest of the structure is also very similar, including the cartesian structure of $\mathbf{R}_{\mathbf{GSMnem} \times \mathbf{D}_\lambda}$ and exponentials. The relation R_ε relates computations with effect ε , and we require relations to be closed under the get and put operations so that the counit lifts. The functor $\mathbf{R}_{\mathbf{GSMnem} \times \mathbf{D}_\lambda} \rightarrow \mathbf{GSMnem} \times \mathbf{D}_\lambda$ maps each object (X, \underline{C}, R) to (X, \underline{C}) and each morphism (f, N) to itself.

We lift the right adjoint by defining $\dot{U}(X, \underline{C}, R) := (X_\emptyset, \underline{U} \underline{C}, R')$ where

$$R' := \{(x, \text{thunk } M) \mid (x, M) \in R_\emptyset\}$$

The lifting of the monoid action is $\varepsilon \dot{\circ} (X, \underline{C}, R) := ((X_{\varepsilon' \cup \varepsilon})_{\varepsilon' \subseteq \{\text{get}, \text{put}\}}, \langle\langle\varepsilon\rangle\rangle \underline{C}, (R_{\varepsilon' \cup \varepsilon})_{\varepsilon' \subseteq \{\text{get}, \text{put}\}})$. Finally, for the lifting of the left adjoint suppose that X is a poset, A is a value type, and $R \subseteq X \times \text{Term}_A$. We take $\dot{F}(X, A, R) = (FX, \langle\emptyset\rangle A, R')$, where each relation R'_ε is defined inductively by the following rules:

$$\frac{(x, V) \in R}{(x, \text{coerce}_{\emptyset \leq \varepsilon} \langle V \rangle) \in R'_\varepsilon} \quad \frac{(t, M) \in R'_\varepsilon \quad (b, V) \in \mathcal{R}[\![\text{bool}]\!]}{(\text{put}_b t, \text{put } V; M) \in R'_\varepsilon} \text{ if } \text{put} \in \varepsilon$$

$$\frac{(t_{\text{true}}, M_{\text{false}}), (t_{\text{false}}, M_{\text{false}}) \in R'_\varepsilon}{(\text{get}(t_{\text{true}}, t_{\text{false}}), \text{get } () \text{ to } b. \underline{\text{if}} \ b \ \text{then } M_{\text{true}} \ \text{else } M_{\text{false}}) \in R'_\varepsilon} \text{ if } \text{get} \in \varepsilon$$

This definition is based on the free lifting from Section 3.1.1.

This defines an LR model that satisfies all of the requirements of the abstraction lemma. Hence our model of shared global state is adequate, and we have verified all of our example program transformations.

4.4 Relating call-by-value and call-by-name, semantically

The final contribution of this chapter is to redevelop our syntax-based reasoning principle for relating call-by-value and call-by-name evaluation (Chapter 3) using the denotational semantics. Working inside the semantics makes it easier to apply our reasoning principle.

$$\begin{array}{l}
\boxed{\phi_{\tau,\varepsilon} : F_\varepsilon \llbracket \tau \rrbracket_\varepsilon^v \rightarrow \llbracket \tau \rrbracket_\varepsilon^n} \\
\phi_{\mathbf{unit},\varepsilon} := id : F_\varepsilon 1 \rightarrow F_\varepsilon 1 \\
\phi_{\mathbf{bool},\varepsilon} := id : F_\varepsilon 2 \rightarrow F_\varepsilon 2 \\
\phi_{\tau \rightarrow \tau',\varepsilon} := \omega \circ F_\varepsilon(\Lambda(U\phi_{\tau',\varepsilon} \circ ev^\dagger \circ (id \times \psi_{\tau,\varepsilon}))) : F_\varepsilon(U(\llbracket \tau \rrbracket_\varepsilon^v \Rightarrow F_\varepsilon \llbracket \tau' \rrbracket_\varepsilon^v)) \rightarrow U \llbracket \tau \rrbracket_\varepsilon^n \Rightarrow \llbracket \tau' \rrbracket_\varepsilon^n \\
\\
\boxed{\psi_{\tau,\varepsilon} : U \llbracket \tau \rrbracket_\varepsilon^n \rightarrow U(F_\varepsilon \llbracket \tau \rrbracket_\varepsilon^v)} \\
\psi_{\mathbf{unit},\varepsilon} := id : U(F_\varepsilon 1) \rightarrow U(F_\varepsilon 1) \\
\psi_{\mathbf{bool},\varepsilon} := id : U(F_\varepsilon 2) \rightarrow U(F_\varepsilon 2) \\
\psi_{\tau \rightarrow \tau',\varepsilon} := \eta_\varepsilon \circ \Lambda(\psi_{\tau',\varepsilon} \circ ev \circ (id \times (U\phi_{\tau,\varepsilon} \circ \eta_\varepsilon))) \\
\quad : U(U \llbracket \tau \rrbracket_\varepsilon^n \Rightarrow \llbracket \tau' \rrbracket_\varepsilon^n) \rightarrow U(F_\varepsilon(U(\llbracket \tau \rrbracket_\varepsilon^v \Rightarrow F_\varepsilon \llbracket \tau' \rrbracket_\varepsilon^v)))
\end{array}$$

Figure 4.5: Semantic morphisms ϕ from call-by-value to Levy-style call-by-name and ψ from Levy-style call-by-name to call-by-value

Recall that in Chapter 3 we restrict to Gifford-style effect algebras, so that effects are subsets ε of some fixed set Σ of operations. We impose the same restriction in this section. Section 3.3 defines maps

$$\Gamma \vDash M : \langle \varepsilon \rangle \langle \tau \rangle_\varepsilon^v \mapsto \Gamma \vDash \Phi_{\tau,\varepsilon} M : \langle \tau \rangle_\varepsilon^n \quad \Gamma \vDash N : \langle \tau \rangle_\varepsilon^n \mapsto \Gamma \vDash \Psi_{\tau,\varepsilon} N : \langle \varepsilon \rangle \langle \tau \rangle_\varepsilon^v$$

between call-by-value and Levy-style call-by-name computations (with effect ε), and shows that they form Galois connections under certain circumstances. In this section we show that the *interpretations* of these maps in models satisfying certain conditions (which we state later) form Galois connections.

Definition 4.4.1 Suppose that X and Y are posets. A *Galois connection* (ϕ, ψ) from X to Y is a pair of monotone functions

$$\phi : X \rightarrow Y \quad \psi : Y \rightarrow X$$

such that for all $x \in X$ and $y \in Y$,

$$\phi x \sqsubseteq y \quad \Leftrightarrow \quad x \sqsubseteq \psi y \quad \blacktriangleleft$$

Figure 4.5 defines *semantic* analogues of $\Phi_{\tau,\varepsilon}$ and $\Psi_{\tau,\varepsilon}$. We use some extra notation: $\llbracket \tau \rrbracket_\varepsilon^v$ is $\llbracket \langle \tau \rangle_\varepsilon^v \rrbracket$ (i.e. the interpretation of the call-by-value translation of the type τ , with the effect annotation ε), and similarly $\llbracket \tau \rrbracket_\varepsilon^n$ is $\llbracket \langle \tau \rangle_\varepsilon^n \rrbracket$. We also define F_ε to be $\varepsilon \otimes F-$, so that $\llbracket \langle \varepsilon \rangle A \rrbracket = F_\varepsilon \llbracket A \rrbracket$, and define $\eta_{\varepsilon,X} : X \rightarrow U(F_\varepsilon X)$ as $U((\emptyset \leq \varepsilon) \otimes X) \circ \eta_X$ (using $\emptyset \subseteq \varepsilon$). The figure uses the extension operator $(-)^{\dagger}$ defined in Section 4.2.1, and the natural transformation ev from Section 4.2.2. The maps between call-by-value and call-by-name in the semantics are morphisms $\phi_{\tau,\varepsilon} : F_\varepsilon \llbracket \tau \rrbracket_\varepsilon^v \rightarrow \llbracket \tau \rrbracket_\varepsilon^n$ in the computation category \mathbf{D} and $\psi_{\tau,\varepsilon} : U \llbracket \tau \rrbracket_\varepsilon^n \rightarrow U(F_\varepsilon \llbracket \tau \rrbracket_\varepsilon^v)$ in the value category \mathbf{C} . The definition relies on our assumption that the effect algebra is Gifford-style (which implies the multiplication is idempotent and $\varepsilon \otimes \llbracket \tau \rrbracket_\varepsilon^n = \llbracket \tau \rrbracket_\varepsilon^n$).

Using composition, the morphisms $\phi_{\tau,\varepsilon}$ and $\psi_{\tau,\varepsilon}$ induce functions between posets of morphisms (for each $X \in \mathbf{C}$):

$$U\phi_{\tau,\varepsilon} \circ - : \mathbf{C}(X, U(F_\varepsilon \llbracket \tau \rrbracket_\varepsilon^v)) \rightarrow \mathbf{C}(X, U \llbracket \tau \rrbracket_\varepsilon^n) \quad \psi_{\tau,\varepsilon} \circ - : \mathbf{C}(X, U \llbracket \tau \rrbracket_\varepsilon^n) \rightarrow \mathbf{C}(X, U(F_\varepsilon \llbracket \tau \rrbracket_\varepsilon^v))$$

These are the interpretations of Φ and Ψ in the following sense.

Lemma 4.4.2 If $\Gamma \vdash M : \langle \varepsilon \rangle (\tau)_\varepsilon^v$ then $\llbracket \Phi_{\tau,\varepsilon} M \rrbracket = U\phi_{\tau,\varepsilon} \circ \llbracket M \rrbracket$, and if $\Gamma \vdash N : (\tau)_\varepsilon^n$ then $\llbracket \Psi_{\tau,\varepsilon} N \rrbracket = \psi_{\tau,\varepsilon} \circ \llbracket N \rrbracket$. ◀

Our goal is to show that the pair $(U\phi_{\tau,\varepsilon} \circ -, \psi_{\tau,\varepsilon} \circ -)$ is a Galois connection for each τ . For the same reason as in Section 3.3, we cannot expect this to hold in every case; it will only hold if enough computations are thunkable. We defined thunkable for syntax in Definition 3.3.2; we now define thunkable for semantics.

Definition 4.4.3 (Thunkable) A morphism $f : X \rightarrow U(F_\varepsilon Y)$ in \mathbf{C} is *thunkable* if

$$U(F_\varepsilon \eta_{\varepsilon,Y}) \circ f \sqsubseteq \eta_{\varepsilon,UF_\varepsilon Y} \circ f$$

An effect ε is *thunkable* if $U(F_\varepsilon \eta_{\varepsilon,Z}) \sqsubseteq \eta_{\varepsilon,UF_\varepsilon Z}$ for each $Z \in \mathbf{C}$. ◀

Note that an effect ε is thunkable if and only if for all X, Y , all morphisms $f : X \rightarrow U(F_\varepsilon Y)$ are thunkable. This definition is based on the one given by Führmann [27], adapted to order-enriched models.⁸

An example of a side-effect that is thunkable under this definition is undefined behaviour (and the reasoning principle we give below can be used to show that, if side-effects are restricted to undefined behaviour, call-by-value can be replaced with call-by-name). Nondeterministic choice as we present it above (in which we can statically make nondeterministic choices) is not thunkable, but it is if we reverse all of the orders on morphisms. This means we can use the reasoning principle to replace call-by-name with call-by-value, but not the reverse, if we allow nondeterministic choices to be made statically.

Assuming the effect ε is thunkable is enough to show we have a Galois connection in the semantics (this theorem is a semantic version of Theorem 3.3.3).

Theorem 4.4.4 If the effect ε is thunkable, then for every source-language type τ and object $X \in \mathbf{C}$, the pair $(U\phi_{\tau,\varepsilon} \circ -, \psi_{\tau,\varepsilon} \circ -)$ is a Galois connection from $\mathbf{C}(X, U(F_\varepsilon \llbracket \tau \rrbracket_\varepsilon^v))$ to $\mathbf{C}(X, U \llbracket \tau \rrbracket_\varepsilon^n)$. ◀

Recall that in Section 3.4 we state our reasoning principle that relates call-by-value and call-by-name on open terms by composing the call-by-name translation with our maps between evaluation orders, to obtain a computation with the same typing as the call-by-value translation:

$$(\Gamma)_\varepsilon^v \longrightarrow (\Gamma)_\varepsilon^n \xrightarrow{(\hat{e})_\varepsilon^n} (\tau)_\varepsilon^n \longrightarrow \langle \varepsilon \rangle (\tau)_\varepsilon^v$$

We defined this formally by giving a substitution $\hat{\Phi}_{\Gamma,\varepsilon}$ from terms in call-by-name contexts $(\Gamma)_\varepsilon^n$ to terms in call-by-value contexts $(\Gamma)_\varepsilon^v$. The composition is then the term $\Psi_{\tau,\varepsilon}((\hat{e})_\varepsilon^n[\hat{\Phi}_{\Gamma,\varepsilon}])$.

Since in this section we reason using the denotational semantics, we can now say what this means inside models of GCBPV. The interpretation of the substitution $\hat{\Phi}_{\Gamma,\varepsilon}$ is a morphism $\hat{\phi}_{\Gamma,\varepsilon} : \llbracket \Gamma \rrbracket_\varepsilon^v \rightarrow \llbracket \Gamma \rrbracket_\varepsilon^n$, which is defined as follows (recall that \diamond is the empty typing context).

$$\hat{\phi}_{\diamond,\varepsilon} := id : 1 \rightarrow 1 \quad \hat{\phi}_{(\Gamma,x:\tau),\varepsilon} := \hat{\phi}_{\Gamma,\varepsilon} \times (U\phi_{\tau,\varepsilon} \circ \eta_\varepsilon) : \llbracket \Gamma \rrbracket_\varepsilon^v \times \llbracket \tau \rrbracket_\varepsilon^v \rightarrow \llbracket \Gamma \rrbracket_\varepsilon^n \times U \llbracket \tau \rrbracket_\varepsilon^n$$

The interpretations of the computations we wish to relate are the two sides of the inequality given as the following diagram. Assuming the model is adequate, to show that the two terms

⁸The (directed) definition of thunkable also appears in another form: an effect ε is thunkable if and only if $U \circ F_\varepsilon$ forms a *Kock-Zöberlein* monad [51].

are related by the contextual preorder it therefore suffices to show that this inequality holds.

$$\begin{array}{ccc} \llbracket \Gamma \rrbracket_{\varepsilon}^{\vee} & \xrightarrow{\hat{\phi}_{\Gamma, \varepsilon}} & \llbracket \Gamma \rrbracket_{\varepsilon}^{\text{n}} \\ \llbracket e \rrbracket_{\varepsilon}^{\vee} \downarrow & \sqsubseteq & \downarrow \llbracket e \rrbracket_{\varepsilon}^{\text{n}} \\ U(F_{\varepsilon} \llbracket \tau \rrbracket_{\varepsilon}^{\vee}) & \xleftarrow{\psi_{\tau, \varepsilon}} & U \llbracket \tau \rrbracket_{\varepsilon}^{\text{n}} \end{array}$$

Again we can show this is the case directly using the properties of Galois connections (see also Lemma 3.4.1).

Lemma 4.4.5 Suppose that $(U\phi_{\tau', \varepsilon} \circ -, \psi_{\tau', \varepsilon} \circ -)$ is a Galois connection from $\mathbf{C}(X, U(F_{\varepsilon} \llbracket \tau' \rrbracket_{\varepsilon}^{\vee}))$ to $\mathbf{C}(X, U \llbracket \tau' \rrbracket_{\varepsilon}^{\text{n}})$ for each τ' . If $\Gamma \vdash e : \tau$ and $\text{ops } e \subseteq \varepsilon$ then

$$\llbracket e \rrbracket_{\varepsilon}^{\vee} \sqsubseteq \psi_{\tau, \varepsilon} \circ \llbracket e \rrbracket_{\varepsilon}^{\text{n}} \circ \hat{\phi}_{\Gamma, \varepsilon} \quad \blacktriangleleft$$

We use this lemma to show the semantic version of our reasoning principle for relating call-by-value and Levy-style call-by-name (see also Theorem 3.4.2):

Theorem 4.4.6 (Call-by-value and call-by-name) Suppose we given some GCBPV inequational theory containing a Gifford-style effect algebra, and an adequate model of it in which the effect ε is thunkable. If $\Gamma \vdash e : \tau$ and $\text{ops } e \subseteq \varepsilon$ then

$$\llbracket e \rrbracket_{\varepsilon}^{\vee} \leq_{\text{ctx}} \Psi_{\tau, \varepsilon}(\llbracket e \rrbracket_{\varepsilon}^{\text{n}}[\hat{\Phi}_{\Gamma, \varepsilon}])$$

Proof. By Theorem 4.4.4, the maps between call-by-value and call-by-name computations form Galois connections. Hence we can apply Lemma 4.4.5, which tells us that

$$\llbracket e \rrbracket_{\varepsilon}^{\vee} \sqsubseteq \psi_{\tau, \varepsilon} \circ \llbracket e \rrbracket_{\varepsilon}^{\text{n}} \circ \hat{\phi}_{\Gamma, \varepsilon}$$

The result follows from adequacy. \square

4.5 Related work

Effect-dependent transformations There is a long line of work on using denotational semantics to verify the correctness of effect-dependent program transformations, for example by Benton and Buchlovsky [5], Benton and Kennedy [8], Benton et al. [10, 9], Birkedal et al. [13]. Here we differ in two major ways. We focus on giving a general framework (somewhat similar to Kammar and Plotkin [42]) that incorporates several side-effects with little effort, and instantiate it for examples. Second, we focus on noninvertible transformations, which are often ignored in previous work. (For example, they are not mentioned at all in [11] but are briefly considered in [6].)

Inequations and order-enriched semantics The inequations and orders we use are induced by considering the possible behaviours of computations. Hoare and He [38] emphasize a similar idea to this in the setting of process calculus.

Previous work on order-enriched semantics (such as Fiore [26]) has concentrated mainly on fixed points and recursive types, using partiality as the order (the least element represents nontermination). Goncharov and Schröder [30] also incorporate other side-effects, but again use partiality for the order on morphisms. They also use a more general notion of order-enrichment, where only certain sets of morphisms are required to be partially ordered, allowing e.g. $\mathbf{C} = \mathbf{Set}$ to be used with a non-trivial order. We do not use this approach, because it is unclear how to extend it to a higher-order language (Goncharov and Schröder consider only first-order).

Relating syntax and semantics To prove adequacy we rely heavily on the fibrational view of logical relations started by Hermida [36]. We also based our presentation somewhat on Katsumata [44], who relates models of a monadic language similar to GMM. We obtain adequacy as a direct consequence of the fundamental lemma. Plotkin and Power [87] and Kammar et al. [40] instead prove termination first, then obtain adequacy as a corollary. We expect our technique for proving adequacy to work even for languages with nontermination.

4.6 Summary

This chapter presents a general denotational semantics for GCBPV, and uses it to prove the validity effect-dependent program transformations. This improves on Chapter 3, in which we reasoned inside the syntax. Denotational models makes the proofs less cumbersome. We include a discussion of various sources of models. In particular, we discuss monadic models to allow us use previous monadic models of side-effects, such as local state [88, 40] and probability [37], as the basis of models of GCBPV. We also include a general logical-relations-based method for relating models of GCBPV, including proofs of adequacy.

We are careful to ensure that we allow noninvertible transformations, by using contextual *preorders* and order-enriched models. This is important because there are various applications that give rise to noninvertible transformations, as our examples show.

Chapter 5

Call-by-need and extended call-by-push-value

We have developed a framework for reasoning about program transformations that supports call-by-value, Moggi-style call-by-name and Levy-style call-by-name. Another common evaluation order is call-by-need (which is sometimes also referred to as “lazy evaluation” when data constructors defer evaluation of arguments until the data structure is traversed). So far we have said nothing about call-by-need, and in fact CBPV (and GCBPV) do not enable us to reason about call-by-need.

An intuitive reason is that call-by-need has “action at a distance” in that reduction of one subterm causes reduction of all other subterms that originated as copies during variable substitution. Indeed call-by-need is often framed using mutable stores (graph reduction [100], or reducing a thunk which is accessed by multiple pointers [53]). CBPV does not allow these to be encoded.

This chapter presents *extended call-by-push-value* (ECBPV), a calculus similar to CBPV, but which can capture call-by-need reduction in addition to call-by-value and call-by-name. Specifically, ECBPV adds an extra primitive $M \text{ need } \underline{x}. N$ which runs N , with M being evaluated the first time \underline{x} is used. On subsequent uses of \underline{x} , the result of the first run is returned immediately. The term M is evaluated at most once. We give the syntax and type system of ECBPV, together with an equational theory that expresses when terms are considered equal. We then give a compositional translation from a lambda calculus (specifically our source language in Section 2.2) into ECBPV for call-by-need reduction.

Initially we extend *ordinary* CBPV, which does not track the effects of computations in the type system. As we will see, support for call-by-need (and action at a distance more generally) makes tracking effects inside the type system more difficult, so we postpone the discussion of grading until the end of the chapter.

ECBPV can be used to reason about call-by-need and all of the other evaluation orders captured by CBPV. Hence we can consider equivalences between call-by-need and other evaluation orders. For example, if there are no side-effects at all in the source language¹ then call-by-need, call-by-value and call-by-name should be semantically equivalent. If the only effect is nondeterminism, then need and value (but not name) are equivalent. If the only effect is nontermination then need and name (but not value) are equivalent. We show that ECBPV can be used to prove such equivalences by proving the latter. We do this using a technique similar to the one we used in Chapter 3 for call-by-value and call-by-name. The main difference

¹Without an effect system we have to restrict the side-effects of the entire source language. With an effect system we can restrict individual expressions.

is that we rely on *Kripke logical relations of varying arity* [39], which generalize ordinary logical relations.

This chapter makes the following contributions:

- We describe *extended call-by-push-value*, a version of CBPV containing an extra construct that adds support for call-by-need. We give its syntax, type system, and equational theory (Section 5.1).
- We describe a call-by-need translation from our source language into ECBPV (Section 5.2).
- We show that, if the source language has nontermination as the only side-effect, then call-by-need and Moggi-style call-by-name are equivalent (Section 5.3).
- We refine the type system of ECBPV so that its types also carry effect information (Section 5.4). This allows equivalences between evaluation orders to be exploited, both at ECBPV and at source level, without a whole-language restriction on side-effects.

5.1 Extended call-by-push-value

We describe an extension to call-by-push-value with support for call-by-need. The primary difference between ordinary CBPV and ECBPV is the addition of a primitive that allows *computations* to be added to the environment so that they are evaluated only the first time they are used. Before describing this change, we take a closer look at how CBPV supports call-by-value and call-by-name.

Recall that CBPV stratifies terms into *values*, which do not have side-effects, and *computations*, which might. Evaluation order is irrelevant for values, so we are only concerned with how computations are sequenced. There is exactly one primitive that causes the evaluation of more than one computation, which is $M \text{ to } x. N$. This behaves in the same way as it does in GCBPV. The evaluation order is fixed: M is always eagerly evaluated. This construct can be used to implement call-by-value as in Figure 2.15: to apply a function, eagerly evaluate the argument and then evaluate the body of the function.

Thanks let us implement call-by-name in CBPV: arguments to functions are thunked computations. Arguments are used by forcing them, so that the computation is evaluated every time the argument is used. Effectively, we simulate a construct $M \text{ name } \underline{x}. N$, which evaluates M each time the variable \underline{x} is used by N , rather than eagerly evaluating. (The variable \underline{x} is underlined here to indicate that it refers to a computation rather than a value: uses of it may have side-effects.)

To support call-by-need, extended call-by-push-value adds another construct $M \text{ need } \underline{x}. N$. This runs the computation N , with the computation M being evaluated the first time \underline{x} is used. On subsequent uses of \underline{x} , the result of the first run is returned immediately. The computation M is evaluated at most once. This new construct adds the “action at a distance” missing from ordinary CBPV.

We observe that adding general mutable references to call-by-push-value would allow call-by-need to be encoded. However, reasoning about evaluation order would be difficult, and so we do not take this option.

5.1.1 Syntax

The syntax of extended call-by-push-value is as follows. The **highlighted** parts are new in this chapter. The rest of the syntax is similar to CBPV and GCBPV. The type constructor F

corresponds to $\langle \varepsilon \rangle$ in GCBPV (without the grading by the effect ε). There are no coercions because of the lack of grading.

$$\begin{array}{ll}
A, B ::= b & V, W ::= c \mid x \\
\quad \mid \mathbf{unit} & \quad \mid () \\
\quad \mid A_1 \times A_2 & \quad \mid (V_1, V_2) \mid \mathbf{fst} V \mid \mathbf{snd} V \\
\quad \mid \mathbf{empty} & \quad \mid \mathbf{case}_A V \mathbf{of} \{ \} \\
\quad \mid A_1 + A_2 & \quad \mid \mathbf{inl}_{A_2} V \mid \mathbf{inr}_{A_1} V \mid \mathbf{case} V \mathbf{of} \{ \mathbf{inl} x_1. W_1, \mathbf{inr} x_2. W_2 \} \\
\quad \mid \underline{U} \underline{C} & \quad \mid \mathbf{thunk} M \\
\\
\underline{C}, \underline{D} ::= \mathbf{unit} & M, N ::= \lambda \{ \} \\
\quad \mid \underline{C}_1 \times \underline{C}_2 & \quad \mid \lambda \{ 1. M_1, 2. M_2 \} \mid 1' M \mid 2' M \\
\quad \mid A \rightarrow \underline{C} & \quad \mid \lambda x : A. M \mid V' M \\
\quad \mid \mathbf{F} A & \quad \mid \mathbf{op} V \mid \langle V \rangle \mid M \mathbf{to} x. N \\
& \quad \mid \mathbf{force} V \mid \underline{x} \mid M \mathbf{need} \underline{x}. N
\end{array}$$

We assume two sets of variables: *value variables* x, y, \dots (which are the same as the variables in GCBPV) and *computation variables* $\underline{x}, \underline{y}, \dots$. While ordinary CBPV does not include computation variables, they do not of themselves add any expressive power to it. The ability to express call-by-need in ECBPV comes from the **need** construct used to bind them.²

The primary new construct is $M \mathbf{need} \underline{x}. N$. This term evaluates N . The first time \underline{x} is evaluated (due to a use of \underline{x} inside N) it behaves the same as the computation M . If M returns a value V then subsequent uses of \underline{x} behave the same as $\langle V \rangle$. Hence only the first use of \underline{x} will evaluate M . If \underline{x} is not used then M is not evaluated at all. The computation variable \underline{x} bound inside the term is primarily used by eagerly sequencing it with other computations. For example,

$$M \mathbf{need} \underline{x}. \underline{x} \mathbf{to} y. \underline{x} \mathbf{to} z. \langle (y, z) \rangle$$

uses \underline{x} twice: once where the result is bound to y , and once where the result is bound to z . Only the first of these uses will evaluate M , so this term has the same semantics as $M \mathbf{to} x. \langle (x, x) \rangle$. The term $M \mathbf{need} \underline{x}. \langle () \rangle$ does not evaluate M at all, and has the same semantics as $\langle () \rangle$.

With the addition of **need** it is not in general possible to statically determine the order in which computations are executed. Uses of computation variables are given statically, but not all of these actually evaluate the corresponding computation dynamically. The set of uses of computation variables that actually cause effects depends on run-time behaviour. This will be important when describing the effect system in Section 5.4.

Typing contexts Γ are ordered lists mapping value variables to value types A , and computation variables to computation types of the form $\mathbf{F} A$. The restriction to $\mathbf{F} A$ is due to the fact that the only construct that binds computation variables is **need**, and this only sequences computations of returner type. Allowing computation variables to be associated with other forms of computation type in typing contexts is therefore unnecessary. Ground types are value types that do not contain thunks:

$$G ::= b \mid \mathbf{unit} \mid G_1 \times G_2 \mid \mathbf{empty} \mid G_1 + G_2$$

The syntax is parameterized by a notion of signature identical to the one for GCBPV (Definition 2.7.1), except for the omission of the effects ε .

²Computation variables are not strictly required to support call-by-need (since we can use $x : \mathbf{U}(\mathbf{F} A)$ instead of $\underline{x} : \mathbf{F} A$), but they simplify reasoning about evaluation order, and therefore we choose to include them.

Definition 5.1.1 A ECBPV signature consists of the following data:

- A set \mathcal{B} of base types.
- A family of pairwise disjoint sets \mathcal{K}_A of constants of type A , indexed by value types A .
- A set Σ of operations.
- For each operation $\text{op} \in \Sigma$, ground types car_{op} and ar_{op} , respectively called the *coarity* and *arity* of op . ◀

(Recall from Definition 2.7.1 that the coarity is the type of the argument of the operation, and the arity is the type of the result, not the other way around.)

The type system of extended call-by-push-value is a minor extension of the type system of ordinary call-by-push-value: we add one typing rule for each of the two new constructs. The rules are given in Figure 5.1. We again write the value typing judgment as $\Gamma \vdash V : A$ and the computation typing judgment as $\Gamma \vdash M : \underline{C}$, and follow our usual convention: rules that add a new variable to the typing context implicitly require freshness.

ECBPV admits a substitution lemma similar to the one for GCBPV (Lemma 2.7.2). *Substitutions* σ are given by the following grammar:

$$\sigma ::= \diamond \mid \sigma, x \mapsto V \mid \sigma, \underline{x} \mapsto M$$

where \diamond is the empty substitution. The typing judgment $\Gamma \vdash \sigma : \Delta$ for substitutions means in the context Γ the terms in σ have the types given in the context Δ . It is defined as follows:

$$\frac{}{\Gamma \vdash \diamond : \diamond} \quad \frac{\Gamma \vdash \sigma : \Delta \quad \Gamma \vdash V : A}{\Gamma \vdash (\sigma, x \mapsto V) : (\Delta, x : A)} \quad \frac{\Gamma \vdash \sigma : \Delta \quad \Gamma \vdash M : \mathbf{F}A}{\Gamma \vdash (\sigma, \underline{x} \mapsto M) : (\Delta, \underline{x} : \mathbf{F}A)}$$

We write $V[\sigma]$ and $M[\sigma]$ for the (capture-avoiding) applications of the substitution σ to the value term V and to the computation term M , and $x \mapsto V$ and $\underline{x} \mapsto M$ for the obvious substitutions that are identities on all other variables. The substitution lemma for ECBPV is:

Lemma 5.1.2 (Substitution) Suppose that $\Gamma \vdash \sigma : \Delta$.

1. (Values) If $\Delta \vdash V : B$ then $\Gamma \vdash V[\sigma] : B$.
2. (Computations) If $\Delta \vdash M : \underline{C}$ then $\Gamma \vdash M[\sigma] : \underline{C}$. ◀

We define the call-by-name construct mentioned above as syntactic sugar for other CBPV primitives:

$$M \text{ name } \underline{x}. N := \mathbf{thunk} M \text{ ' } \lambda y. N[\underline{x} \mapsto \mathbf{force} y]$$

where y is not free in N .

5.1.2 Inequational theories for ECBPV

The behaviour of **need** is captured by the core axioms for ECBPV inequational theories, which are listed in Figure 5.2. Each of these axioms is symmetric and holds when both sides of the equation have suitable types. For example, the second axiom of the third group ($M \text{ need } \underline{x}. N \equiv N$) holds only if \underline{x} is not free in N .

For the core axioms to capture call-by-need, we might expect computation terms that are not of the form $\langle V \rangle$ to never be duplicated, since they should not be evaluated more than once. There are two exceptions to this rule. Such terms can be duplicated in the axioms that duplicate value terms (such as the β -laws for sum types). In this case, the syntax ensures such terms are thunked.

$$\begin{array}{c}
\boxed{\Gamma \vdash V : A} \\
\frac{}{\Gamma \vdash c : A} \text{ if } c \in \mathcal{K}_A \qquad \frac{}{\Gamma \vdash x : A} \text{ if } (x : A) \in \Gamma \qquad \frac{}{\Gamma \vdash () : \mathbf{unit}} \\
\frac{\Gamma \vdash V_1 : A_1 \quad \Gamma \vdash V_2 : A_2}{\Gamma \vdash (V_1, V_2) : A_1 \times A_2} \qquad \frac{\Gamma \vdash V : A_1 \times A_2}{\Gamma \vdash \mathbf{fst} V : A_1} \qquad \frac{\Gamma \vdash V : A_1 \times A_2}{\Gamma \vdash \mathbf{snd} V : A_2} \qquad \frac{\Gamma \vdash V : \mathbf{empty}}{\Gamma \vdash \mathbf{case}_A V \text{ of } \{\} : A} \\
\frac{\Gamma \vdash V : A_1}{\Gamma \vdash \mathbf{inl}_{A_2} V : A_1 + A_2} \qquad \frac{\Gamma \vdash V : A_2}{\Gamma \vdash \mathbf{inr}_{A_1} V : A_1 + A_2} \\
\frac{\Gamma \vdash V : A_1 + A_2 \quad \Gamma, x_1 : A_1 \vdash W_1 : B \quad \Gamma, x_2 : A_2 \vdash W_2 : B}{\Gamma \vdash \mathbf{case} V \text{ of } \{\mathbf{inl} x_1. W_1, \mathbf{inr} x_2. W_2\} : B} \\
\frac{\Gamma \Vdash M : \underline{C}}{\Gamma \vdash \mathbf{thunk} M : \mathbf{UC}} \\
\boxed{\Gamma \Vdash M : \underline{C}} \\
\frac{}{\Gamma \Vdash \lambda\{\} : \mathbf{unit}} \qquad \frac{\Gamma \Vdash M_1 : \underline{C}_1 \quad \Gamma \Vdash M_2 : \underline{C}_2}{\Gamma \Vdash \lambda\{1. M_1, 2. M_2\} : \underline{C}_1 \times \underline{C}_2} \qquad \frac{\Gamma \Vdash M : \underline{C}_1 \times \underline{C}_2}{\Gamma \Vdash 1'M : \underline{C}_1} \qquad \frac{\Gamma \Vdash M : \underline{C}_1 \times \underline{C}_2}{\Gamma \Vdash 2'M : \underline{C}_2} \\
\frac{\Gamma, x : A \Vdash M : \underline{C}}{\Gamma \Vdash \lambda x : A. M : A \rightarrow \underline{C}} \qquad \frac{\Gamma \vdash V : A \quad \Gamma \Vdash M : A \rightarrow \underline{C}}{\Gamma \Vdash V'M : \underline{C}} \\
\frac{\Gamma \vdash V : \mathbf{car}_{\text{op}}}{\Gamma \Vdash \text{op} V : \mathbf{F ar}_{\text{op}}} \qquad \frac{\Gamma \vdash V : A}{\Gamma \Vdash \langle V \rangle : \mathbf{FA}} \qquad \frac{\Gamma \Vdash M : \mathbf{FA} \quad \Gamma, x : A \Vdash N : \underline{C}}{\Gamma \Vdash M \text{ to } x. N : \underline{C}} \qquad \frac{\Gamma \vdash V : \mathbf{UC}}{\Gamma \Vdash \mathbf{force} V : \underline{C}} \\
\frac{}{\Gamma \Vdash \underline{x} : \mathbf{FA}} \text{ if } (\underline{x} : \mathbf{FA}) \in \Gamma \qquad \frac{\Gamma \Vdash M : \mathbf{FA} \quad \Gamma, \underline{x} : \mathbf{FA} \Vdash N : \underline{C}}{\Gamma \Vdash M \mathbf{need} \underline{x}. N : \underline{C}}
\end{array}$$

Figure 5.1: Extended call-by-push-value typing rules

$\text{fst } (V_1, V_2) \equiv V_1$ $\text{case inl}_{A_2} V \text{ of}$ $\{ \text{inl } x_1. W_1 \equiv W_1[x_1 \mapsto V]$ $, \text{inr } x_2. W_2 \}$ $1' \lambda\{1. M_1, 2. M_2\} \equiv M_1$ $V' \lambda x:A. M \equiv M[x \mapsto V]$ $\text{force } (\text{thunk } M) \equiv M$	$\text{snd } (V_1, V_2) \equiv V_2$ $\text{case inr}_{A_1} V \text{ of}$ $\{ \text{inl } x_1. W_1 \equiv W_2[x_2 \mapsto V]$ $, \text{inr } x_2. W_2 \}$ $2' \lambda\{1. M_1, 2. M_2\} \equiv M_2$ $\langle V \rangle \text{ to } x. M \equiv M[x \mapsto V]$ $\langle V \rangle \text{ need } \underline{x}. M \equiv M[\underline{x} \mapsto \langle V \rangle]$	} β -laws
$V \equiv ()$ $V \equiv \text{case}_A W \text{ of } \{ \}$ $V \equiv \text{thunk } (\text{force } V)$ $M \equiv \lambda\{ \}$ $M \equiv \lambda x:A. x' M$	$V \equiv (\text{fst } V, \text{snd } V)$ $\text{case } W \text{ of}$ $\{ \text{inl } y_1. V[x \mapsto \text{inl}_{A_2} y_1]$ $, \text{inr } y_2. V[x \mapsto \text{inr}_{A_1} y_2] \}$ $M \equiv \lambda\{1. 1' M, 2. 2' M\}$ $M \equiv M \text{ to } x. \langle x \rangle$	} η -laws
$M \text{ need } \underline{x}. x \text{ to } y. N \equiv M \text{ to } y. N[\underline{x} \mapsto \langle y \rangle]$ $M \text{ need } \underline{x}. N \equiv N$		
$\lambda\{1. M \text{ to } x. N_1, 2. M \text{ to } x. N_2\} \equiv M \text{ to } x. \lambda\{1. N_1, 2. N_2\}$ $\lambda y:A. M \text{ to } x. N \equiv M \text{ to } x. \lambda y:A. N$ $\lambda\{1. M \text{ need } \underline{x}. N_1, 2. M \text{ need } \underline{x}. N_2\} \equiv M \text{ need } \underline{x}. \lambda\{1. N_1, 2. N_2\}$ $\lambda y:A. M \text{ need } \underline{x}. N \equiv M \text{ need } \underline{x}. \lambda y:A. N$ $M_1 \text{ to } x. M_2 \text{ need } \underline{y}. M_3 \equiv M_2 \text{ need } \underline{y}. M_1 \text{ to } x. M_3$ $(M_1 \text{ to } x. M_2) \text{ to } y. M_3 \equiv M_1 \text{ to } x. M_2 \text{ to } y. M_3$ $(M_1 \text{ need } \underline{x}. M_2) \text{ to } y. M_3 \equiv M_1 \text{ need } \underline{x}. M_2 \text{ to } y. M_3$ $(M_1 \text{ need } \underline{x}. M_2) \text{ need } \underline{y}. M_3 \equiv M_1 \text{ need } \underline{x}. M_2 \text{ need } \underline{y}. M_3$		} Sequencing laws

Figure 5.2: Core axioms of ECBPV inequational theories

The duplication is acceptable because we should allow these terms to be executed once in each separate execution of a computation (and separate executions arise from duplication of thunks). Only duplication within a *single* execution of a computation is problematic. Computations can also be duplicated across both elements of a pair $\lambda\{1. M_1, 2. M_2\}$. This is also correct, because only one component of a pair can be used within a single computation (without thinking), so the side-effects will still not happen twice. (There is a similar consideration for functions, which can only be applied once.) The other axioms never duplicate **need**-bound computations that might have effects.

We have seen the majority of the core axioms before: they are the same as the ones for GCBPV and CBPV (see Section 2.7.2). Only the axioms involving **need** are new; these are highlighted. The β -law for **need** parallels the usual β -law for **to**: it gives the behaviour of computation terms that return values without having any side-effects. The interesting axioms are those in the third group of Figure 5.2, which we call *sequencing axioms*.

The first sequencing axiom is the crucial one. It states that if a computation will next evaluate \underline{x} , where \underline{x} is a computation variable bound to M , then this is the same as evaluating M , and then using the result for subsequent uses of \underline{x} . It implies that

$$\begin{aligned} M \text{ need } \underline{x}. \underline{x} &\equiv M \text{ need } \underline{x}. \underline{x} \text{ to } y. \langle y \rangle && (\eta \text{ law for returner types}) \\ &\equiv M \text{ to } y. \langle y \rangle \\ &\equiv M && (\eta \text{ law for returner types}) \end{aligned}$$

The second sequencing axiom does *garbage collection* [68]: if a computation bound by **need** is not used (because the variable does not appear), then the binding can be dropped. This equation implies, for example, that

$$M_1 \text{ need } \underline{x}_1. M_2 \text{ need } \underline{x}_2. \dots M_n \text{ need } \underline{x}_n. \langle () \rangle \equiv \langle () \rangle$$

The next five sequencing axioms (two from CBPV and three new) state that binding a computation with **to** or **need** commutes with the remaining forms of computation terms. These allow **to** and **need** to be moved to the outside of other constructs *except* **thunk**s. The final three axioms (one from CBPV and two new) capture associativity involving **need** and **to**. Note that associativity between different evaluation orders is not necessarily valid. In particular, we do not have

$$(M_1 \text{ to } x. M_2) \text{ need } y. M_3 \equiv M_1 \text{ to } x. (M_2 \text{ need } x. M_3)$$

(The first term might not evaluate M_1 , the second always does.) This is usually the case when evaluation orders are mixed [79].

These final eight axioms allow computation terms to be placed in normal forms where bindings of computations are on the outside. (Compare this with the translation of source-language *answers* in Section 5.2.)

An inequational theory again consists of two judgment forms, one for values and one for computations:

$$\Gamma \vdash V \leqslant W : A \quad \Gamma \vdash M \leqslant N : \underline{C}$$

We require inequational theories to be closed under congruence, so define term contexts $C[]$ (value terms with a single hole) and $\underline{C}[]$ (computation terms with a single hole):

$$\begin{aligned} C[] &::= \square \mid (C[], V_2) \mid (V_1, C[]) \mid \text{fst } C[] \mid \text{snd } C[] \mid \text{case}_A C[] \text{ of } \{ \\ &\mid \text{inl}_{A_2} C[] \mid \text{inr}_{A_1} C[] \mid \text{case } C[] \text{ of } \{\text{inl } x_1. W_1, \text{inr } x_2. W_2\} \\ &\mid \text{case } V \text{ of } \{\text{inl } x_1. C[], \text{inr } x_2. W_2\} \mid \text{case } V \text{ of } \{\text{inl } x_1. W_1, \text{inr } x_2. C[]\} \\ &\mid \text{thunk } \underline{C}[] \\ \underline{C}[] &::= \square \mid \lambda\{1. M_1, 2. \underline{C}[]\} \mid \lambda\{1. \underline{C}[], 2. M_2\} \mid 1' \underline{C}[] \mid 2' \underline{C}[] \mid \lambda x:A. \underline{C}[] \mid C[]'M \mid V' \underline{C}[] \\ &\mid \text{op } C[] \mid \langle C[] \rangle \mid \underline{C}[] \text{ to } x. N \mid M \text{ to } x. \underline{C}[] \mid \text{force } C[] \\ &\mid \underline{C}[] \text{ need } \underline{x}. N \mid M \text{ need } \underline{x}. \underline{C}[] \end{aligned}$$

We also require closure under substitution, so define a judgment

$$\Gamma \vdash \sigma \leqslant \sigma' : \Delta$$

on well-typed substitutions componentwise, using both judgments of the inequational theory. The definition of ECBPV inequational theory is similar to the definition for GCBPV (Definition 2.7.4):

Definition 5.1.3 (Inequational theory) An *inequational theory* consists of a ECBPV signature and two judgments

$$\Gamma \vdash V \leq W : A \quad \Gamma \perp M \leq N : \underline{C}$$

such that:

- Preorder:
 - Values: if $\Gamma \vdash V : A$ then $\Gamma \vdash V \leq V : A$, and if $\Gamma \vdash V_1 \leq V_2 : A$ and $\Gamma \vdash V_2 \leq V_3 : A$ then $\Gamma \vdash V_1 \leq V_3 : A$.
 - Computations: if $\Gamma \perp M : \underline{C}$ then $\Gamma \perp M \leq M : \underline{C}$, and if $\Gamma \perp M_1 \leq M_2 : \underline{C}$ and $\Gamma \perp M_2 \leq M_3 : \underline{C}$ then $\Gamma \perp M_1 \leq M_3 : \underline{C}$.
- Congruence: if $\Gamma \vdash V \leq W : A$ then for term contexts with hole \square

$$\begin{aligned} \Gamma' \vdash C[V] : B \wedge \Gamma' \vdash C[W] : B &\Rightarrow \Gamma' \vdash C[V] \leq C[W] : B \\ \Gamma' \perp \underline{C}[V] : \underline{D} \wedge \Gamma' \perp \underline{C}[W] : \underline{D} &\Rightarrow \Gamma' \perp \underline{C}[V] \leq \underline{C}[W] : \underline{D} \end{aligned}$$

If $\Gamma \perp M \leq N : \underline{C}$ then for term contexts with hole \square

$$\begin{aligned} \Gamma' \vdash C[M] : B \wedge \Gamma' \vdash C[N] : B &\Rightarrow \Gamma' \vdash C[M] \leq C[N] : B \\ \Gamma' \perp \underline{C}[M] : \underline{D} \wedge \Gamma' \perp \underline{C}[N] : \underline{D} &\Rightarrow \Gamma' \perp \underline{C}[M] \leq \underline{C}[N] : \underline{D} \end{aligned}$$

- Substitution: If $\Gamma \vdash \sigma \leq \sigma' : \Delta$ then

$$\begin{aligned} \Delta \vdash V \leq W : B &\Rightarrow \Gamma \vdash V[\sigma] \leq W[\sigma'] : B \\ \Delta \perp M \leq N : \underline{C} &\Rightarrow \Gamma \perp M[\sigma] \leq N[\sigma'] : \underline{C} \end{aligned}$$

- Core axioms:
 - Values: if $\Gamma \vdash V : A$ and $\Gamma \vdash W : A$, and $V \equiv W$ is an instance of an axiom in Figure 5.2, then $\Gamma \vdash V \leq W : A$ and $\Gamma \vdash W \leq V : A$.
 - Computations: if $\Gamma \perp M : \underline{C}$ and $\Gamma \perp N : \underline{C}$, and $M \equiv N$ is an instance of an axiom in Figure 5.2, then $\Gamma \perp M \leq N : \underline{C}$ and $\Gamma \perp N \leq M : \underline{C}$. ◀

To specify an inequational theory, one gives a collection of signature axioms and then closes under the core axioms, reflexivity, transitivity and congruence. The signature axioms may involve **to** (e.g. Figure 2.14), but we do not expect any examples to require signature axioms involving **need**, because the behaviour of **need** is completely characterized by the sequencing axioms.

We define the contextual preorder in the usual way (recall that \diamond is the empty typing context).

Definition 5.1.4 (Contextual preorder) The *contextual preorder* consists of two judgment forms.

1. Between value terms: $\Gamma \vdash V \leq_{\text{ctx}} W : A$ if $\Gamma \vdash V : A$, $\Gamma \vdash W : A$, and for all ground types G and term contexts $\underline{C}[\]$ with hole \square such that $\diamond \perp \underline{C}[V] : FG$ and $\diamond \perp \underline{C}[W] : FG$ we have

$$\diamond \perp \underline{C}[V] \leq \underline{C}[W] : FG$$

2. Between computation terms: $\Gamma \perp M \leq_{\text{ctx}} N : \underline{C}$ if $\Gamma \perp M : \underline{C}$, $\Gamma \perp N : \underline{C}$, and for all ground types G and term contexts $\underline{C}[\]$ with hole \square such that $\diamond \perp \underline{C}[M] : FG$ and $\diamond \perp \underline{C}[N] : FG$ we have

$$\diamond \perp \underline{C}[M] \leq \underline{C}[N] : FG \quad \blacktriangleleft$$

$$\begin{array}{l}
\text{Evaluation contexts } \mathcal{E}[] ::= \square \mid \text{op } \mathcal{E}[] \mid \text{if } \mathcal{E}[] \text{ then } e_2 \text{ else } e_3 \mid \mathcal{E}[] e_2 \\
\quad \mid (\lambda x:\tau. \mathcal{E}[x]) \mathcal{E}'[] \mid (\lambda x:\tau. \mathcal{E}[]) e_2 \\
\text{Values } v ::= \text{true} \mid \text{false} \mid \lambda x:\tau. e \\
\text{Answers } a ::= v \mid (\lambda x:\tau. a) e \\
\\
\begin{array}{c}
\text{if true then } e_2 \text{ else } e_3 \xrightarrow{\text{need}} e_2 \\
\text{if false then } e_2 \text{ else } e_3 \xrightarrow{\text{need}} e_3
\end{array}
\quad
\frac{e \xrightarrow{\text{need}} e'}{\mathcal{E}[e] \xrightarrow{\text{need}} \mathcal{E}[e']} \\
\\
(\lambda x:\tau. \mathcal{E}[x]) v \xrightarrow{\text{need}} (\lambda x:\tau. \mathcal{E}[v]) v
\quad
(\lambda x:\tau. a) e_1 e_2 \xrightarrow{\text{need}} (\lambda x:\tau. a e_2) e_1 \\
(\lambda x:\tau. \mathcal{E}[x]) ((\lambda y:\tau'. a) e) \xrightarrow{\text{need}} (\lambda y:\tau'. (\lambda x:\tau. \mathcal{E}[x]) a) e
\end{array}$$

Figure 5.3: Call-by-need operational semantics

5.2 Call-by-need translation

Our goal with ECBPV is to use it to reason about call-by-need evaluation. To do this, we give a call-by-need translation from our source language (Section 2.2) into ECBPV.

Most of the translation is similar to the Moggi-style call-by-name translation in Figure 2.16. The critical difference is how functions are dealt with. We encode call-by-need functions as terms of the form

$$\lambda x': \text{UFA. (force } x') \text{ need } \underline{x}. M$$

where x' is not free in M . This is an ECBPV function that accepts a thunk as an argument. The thunk is added to the context, and the body of the function is executed. The first time the argument is used (via \underline{x}), the computation inside the thunk is evaluated. Subsequent uses do not evaluate the computation again.

Before defining the translation we give a call-by-need operational semantics for our source language, based on Ariola and Felleisen's [3]. The only differences between our source language and Ariola and Felleisen's calculus are the addition of booleans, operations, and a type system. It is likely that we can translate other call-by-need calculi, such as those of Launchbury [53] and Maraist et al. [68]. Call-by-need small-step reductions are written $e \xrightarrow{\text{need}} e'$; this is defined in Figure 5.3.

The call-by-need semantics needs some auxiliary definitions. An *evaluation context* $\mathcal{E}[]$ is a source-language expression with a single hole \square , picked from the grammar given in the figure. (There is only one kind of reduction context and one kind of hole because the syntax of the source language is not stratified into values and computations.) The hole in an evaluation context indicates where reduction is currently taking place: it says which part of the expression is currently *needed*. We write $\mathcal{E}[e]$ for the expression in which the hole is replaced with e . A (source-language) *value* is the result of a computation (the word value should not be confused with the value terms of ECBPV). An *answer* is a value in some *environment*, which maps variables to expressions. Answers can be thought of as *closures*. The environment is encoded in an answer using application and lambda abstraction: the answer $(\lambda x:\tau. a) e$ means the answer a where the environment maps x to e . Encoding environments in this way makes the translation slightly simpler than if we had used a Launchbury-style [53] call-by-need language with explicit environments. In the latter case, the translation would need to encode

$$\begin{array}{c}
\langle \mathbf{unit} \rangle := \mathbf{unit} \\
\langle \mathbf{bool} \rangle := \mathbf{unit} + \mathbf{unit} \\
\langle \tau \rightarrow \tau' \rangle := \mathbf{U} (\mathbf{U} (\mathbf{F} (\langle \tau \rangle)) \rightarrow \mathbf{F} (\langle \tau' \rangle)) \\
\langle \diamond \rangle := \diamond \\
\langle \Gamma, x : \tau \rangle := \langle \Gamma \rangle, \underline{x} : \mathbf{F} (\langle \tau \rangle) \\
\hline
\langle \Gamma \vdash x : \tau \rangle := \underline{x} \qquad \langle \Gamma \vdash e : \mathbf{car}_{\text{op}} \rangle = M \qquad \langle \Gamma \vdash () : \mathbf{unit} \rangle := \langle () \rangle \\
\hline
\langle \Gamma \vdash \mathbf{true} : \mathbf{bool} \rangle := \langle \mathbf{inl}_{\mathbf{unit}} () \rangle \qquad \langle \Gamma \vdash \mathbf{false} : \mathbf{bool} \rangle := \langle \mathbf{inr}_{\mathbf{unit}} () \rangle \\
\hline
\langle \Gamma \vdash e_1 : \mathbf{bool} \rangle = M_1 \qquad \langle \Gamma \vdash e_2 : \tau \rangle = M_2 \qquad \langle \Gamma \vdash e_3 : \tau \rangle = M_3 \\
\langle \Gamma \vdash \mathbf{if} \ e_1 \ \mathbf{then} \ e_2 \ \mathbf{else} \ e_3 : \tau \rangle := M_1 \ \mathbf{to} \ x. \ \underline{\mathbf{case}} \ x \ \mathbf{of} \ \{ \mathbf{inl} \ _ . M_2, \ \mathbf{inr} \ _ . M_3 \} \\
\hline
\langle \Gamma, x : \tau \vdash e : \tau' \rangle = M \\
\langle \Gamma \vdash \lambda x : \tau. e : \tau \rightarrow \tau' \rangle := \langle \mathbf{thunk} (\lambda x' : \mathbf{U} (\mathbf{F} (\langle \tau \rangle)). \ \mathbf{force} \ x' \ \mathbf{need} \ \underline{x}. M) \rangle \\
\hline
\langle \Gamma \vdash e_1 : \tau \rightarrow \tau' \rangle = M_1 \qquad \langle \Gamma \vdash e_2 : \tau \rangle = M_2 \\
\langle \Gamma \vdash e_1 \ e_2 : \tau' \rangle := M_1 \ \mathbf{to} \ f. \ (\mathbf{thunk} \ M_2) \ ' \ (\mathbf{force} \ f)
\end{array}$$

Figure 5.4: Call-by-need translation of call-by-need types (top left), contexts (top right), and expressions into ECBPV.

the environments; here they are already encoded inside expressions. Answers are terminal computations: they do not reduce.

The two axioms for reducing **if**-expressions are obvious. The congruence rule only allows reducing the expression that is needed by the computation. The axiom on the bottom left of Figure 5.3 is the most important: it states that if the subexpression currently being evaluated is a variable x , and the environment maps x to a source-language value v , then that use of x can be replaced with v . Note that $\mathcal{E}[v]$ may contain other uses of x ; the replacement only occurs when the value is actually needed. This axiom roughly corresponds to the first sequencing axiom in Figure 5.2. The two axioms on the bottom right of Figure 5.3 rearrange the environment into a standard form. Both have a syntactic restriction to answers so that each expression has at most one reduct (this restriction is not needed to ensure that $\overset{\text{need}}{\rightsquigarrow}$ captures call-by-need).

The call-by-need translation from the source language to ECBPV is defined in Figure 5.4. We assume that each operation op in the source-language signature is also included in the ECBPV signature, and that if τ and τ' are respectively the coarity and arity of op in the source-language signature then $(\tau)^{\text{need}}$ and $(\tau')^{\text{need}}$ are the coarity and arity in the ECBPV signature.

The translation of types is the same as the Moggi-style call-by-name translation in Figure 2.16 (ignoring the grading), so types τ are translated into value types $(\tau)^{\text{need}}$. (We omit the superscript in the figure.) On function types $\tau \rightarrow \tau'$, we use ECBPV function types that receive **thunks** of computations as arguments (the argument is not evaluated before the function is applied). The call-by-need translation differs only on typing contexts and expressions (since there is no subeffecting, we do not have to be careful about the distinction between typing derivations and well-typed expressions).

Source-language typing contexts Γ are translated into ECBPV typing contexts $(\Gamma)^{\text{need}}$,

containing computations that return values. The computations in the context are all bound using **need**. An expression $\Gamma \vdash e : \tau$ is translated to a computation $\llbracket e \rrbracket^{\text{need}}$ that returns $\llbracket \tau \rrbracket^{\text{need}}$ in the context $\llbracket \Gamma \rrbracket^{\text{need}}$. The key case is the translation of lambdas. These become computations that immediately return a thunk of a function. The function places the computation given as an argument onto the context using **need**, so that it is evaluated at most once, before executing the body. The remainder of the cases are similar to the Moggi-style call-by-name translation.

Under the call-by-need translation, the expression $(\lambda x : \tau. e_1) e_2$ is translated into a term that executes the computation $\llbracket e_1 \rrbracket^{\text{need}}$, and executes $\llbracket e_2 \rrbracket^{\text{need}}$ only when needed. This is the case because, by the β rules for thunks, functions, and returner types:

$$\llbracket (\lambda x : \tau. e_1) e_2 \rrbracket^{\text{need}} \equiv \llbracket e_2 \rrbracket^{\text{need}} \mathbf{need} \underline{x}. \llbracket e_1 \rrbracket^{\text{need}}$$

As a consequence, translations of answers are particularly simple: they have the following form (up to \equiv):

$$M_1 \mathbf{need} \underline{x}_1. M_2 \mathbf{need} \underline{x}_2. \cdots M_n \mathbf{need} \underline{x}_n. \langle V \rangle$$

which intuitively means the value V in the environment mapping each \underline{x}_i to M_i .

The call-by-need translation produces ECBPV computations of the correct type.

Lemma 5.2.1 If $\Gamma \vdash e : \tau$ then $\llbracket \Gamma \rrbracket^{\text{need}} \vdash \llbracket e \rrbracket^{\text{need}} : \mathbf{F} \llbracket \tau \rrbracket^{\text{need}}$. ◀

We prove that the call-by-need translation is *sound*: if $e \rightsquigarrow^{\text{need}} e'$ then $\llbracket e \rrbracket^{\text{need}} \equiv \llbracket e' \rrbracket^{\text{need}}$. To do this, we first look at translations of evaluation contexts. The following lemma says the translation captures the idea that the hole in an evaluation context corresponds to the term being evaluated.

Lemma 5.2.2 Define, for each evaluation context $\mathcal{E}[]$, the term context $\llbracket \mathcal{E}[] \rrbracket^{\text{need}}$ by:

$$\begin{aligned} \llbracket \square \rrbracket^{\text{need}} &:= \langle \square \rangle \\ \llbracket \text{op } \mathcal{E}[] \rrbracket^{\text{need}} &:= \llbracket \mathcal{E}[] \rrbracket^{\text{need}} \mathbf{to} x. \text{op } x \\ \llbracket \text{if } \mathcal{E}[] \text{ then } e_2 \text{ else } e_3 \rrbracket^{\text{need}} &:= \llbracket \mathcal{E}[] \rrbracket^{\text{need}} \mathbf{to} x. \mathbf{case} x \text{ of} \\ &\quad \{ \mathbf{inl} _ . \mathbf{thunk} \llbracket e_2 \rrbracket^{\text{need}} \\ &\quad \quad , \mathbf{inr} _ . \mathbf{thunk} \llbracket e_3 \rrbracket^{\text{need}} \} \\ \llbracket \mathcal{E}[] e_2 \rrbracket^{\text{need}} &:= \llbracket \mathcal{E}[] \rrbracket^{\text{need}} \mathbf{to} z. \mathbf{thunk} \llbracket e_2 \rrbracket^{\text{need}} \text{ force } z \\ \llbracket (\lambda x : \tau. \mathcal{E}[x]) \mathcal{E}'[] \rrbracket^{\text{need}} &:= \llbracket \mathcal{E}'[] \rrbracket^{\text{need}} \mathbf{need} \underline{x}. \llbracket \mathcal{E}[x] \rrbracket^{\text{need}} \\ \llbracket (\lambda x. \mathcal{E}[]) e_2 \rrbracket^{\text{need}} &:= \llbracket e_2 \rrbracket^{\text{need}} \mathbf{need} \underline{x}. \llbracket \mathcal{E}[] \rrbracket^{\text{need}} \end{aligned}$$

For each expression e we have:

$$\llbracket \mathcal{E}[e] \rrbracket^{\text{need}} \equiv \llbracket e \rrbracket^{\text{need}} \mathbf{to} y. \llbracket \mathcal{E}[] \rrbracket^{\text{need}} [y]$$

(where y is fresh). ◀

This lemma omits the typing of expressions for presentational purposes. Soundness is now easy to show:

Theorem 5.2.3 (Soundness) If e and e' are closed, well-typed source-language expressions then $e \rightsquigarrow^{\text{need}} e'$ implies $\llbracket e \rrbracket^{\text{need}} \equiv \llbracket e' \rrbracket^{\text{need}}$. ◀

5.3 Equivalence between call-by-name and call-by-need

Extended call-by-push-value can be used to prove equivalences between evaluation orders. In this section we prove a classic example: if the only side-effect is nontermination, then call-by-name is equivalent to call-by-need. We do this in two stages.

First, we show that call-by-name is equivalent to call-by-need *within* ECBPV. Specifically, we show that

$$M \text{ name } \underline{x}. N \cong_{\text{ctx}} M \text{ need } \underline{x}. N$$

(Recall that $M \text{ name } \underline{x}. N$ is syntactic sugar for $\mathbf{thunk } M \text{ ' } \lambda y. N[\underline{x} \mapsto \mathbf{force } y]$.)

Second, an important corollary is that the meta-level reduction strategies are equivalent. We show that the call-by-need translation in Section 5.2 and the Moggi-style call-by-name translation in Figure 2.16 (ignoring grading) give contextually equivalent ECBPV terms. We expect the Levy-style translation in Figure 2.17 to work equally well, except that we would have to define a more complicated Galois connection between the two evaluation orders (see Chapter 3).

To model nontermination being the sole source-language effect, we choose the ECBPV signature with no constants or base types, and a single operation:

Operation op	Coarity car_{op}	Arity ar_{op}
diverge	unit	empty

We expect our proofs to work with general fixed-point operators, but for simplicity we do not consider this here. The operation enables us to define a diverging computation $\Omega_{\underline{C}}$ of each computation type \underline{C} , using the eliminator of the empty type:

$$\Omega_{\underline{C}} := \text{diverge } () \text{ to } x. \text{case}_{\underline{C}} x \text{ of } \{\}$$

We use the smallest inequational theory with this signature. This inequational theory is symmetric. We do not need any signature axioms to characterize nontermination; the core axioms are sufficient. For example, associativity of **to** and the η -law for **empty** imply

$$\Omega_{FA} \text{ to } x. M \equiv \Omega_{\underline{C}}$$

So diverging as part of a larger computation causes the entire computation to diverge.

We first show that

$$M \text{ name } \underline{x}. N \cong_{\text{ctx}} M \text{ need } \underline{x}. N$$

As we did in Chapter 3 we use logical relations to get a strong enough inductive hypothesis for the proof to go through. However, unlike the usual case, it does not suffice to relate *closed* terms. To see why, consider a closed term M of the form

$$\Omega_{FA} \text{ need } \underline{x}. N_1 \text{ to } y. N_2$$

If we relate only closed terms, then we do not learn anything about N_1 itself (since \underline{x} may be free in it). We could attempt to proceed by considering the closed term $\Omega_{FA} \text{ need } \underline{x}. N_1$. For example, if this returns a value V then \underline{x} cannot have been evaluated and M should have the same behaviour as $\Omega_{FA} \text{ need } \underline{x}. N_2[y \mapsto V]$. However, we get stuck when proving the last step. This is only a problem because Ω_{FA} is a nonterminating computation: every closed, terminating computation of returner type has the form $\langle V \rangle$ (up to \equiv), and when these are bound using **need** we can eliminate the binding using the equation

$$\langle V \rangle \text{ need } \underline{x}. N_1 \text{ to } y. N_2 \equiv (N_1 \text{ to } y. N_2)[\underline{x} \mapsto \langle V \rangle]$$

The solution is to relate terms that may have free computation variables (we do not have to consider free value variables). The free computation variables should be thought of as referring to nonterminating computations, because we can remove the bindings of variables that refer to terminating computations. We relate open terms using *Kripke logical relations of varying arity*, which were introduced by Jung and Tiuryn [39] to study lambda definability.

We need a number of definitions first. We define Term_A^Γ as the set of equivalence classes (up to \equiv) of terms of value type A in context Γ , and similarly define $\underline{\text{Term}}_D^\Gamma$ for computation types:

$$\text{Term}_A^\Gamma := \{[V]_{\equiv} \mid \Gamma \vdash V : A\} \quad \underline{\text{Term}}_D^\Gamma := \{[M]_{\equiv} \mid \Gamma \vdash M : D\}$$

A computation-type context Δ is an ECBPV typing context that contains only computation variables \underline{x} . In this section Δ ranges over computation-type contexts only, Γ ranges over arbitrary typing contexts. Variables in computation-type contexts refer to nonterminating computations for the proof of contextual equivalence. A computation-type context Δ' *weakens* another context Δ , written $\Delta' \triangleright \Delta$, whenever Δ is a sublist of Δ' . Formally, this relation is generated by the following rules:

$$\frac{}{\diamond \triangleright \diamond} \quad \frac{\Delta' \triangleright \Delta}{\Delta', \underline{x} : \text{FA} \triangleright \Delta, \underline{x} : \text{FA}} \quad \frac{\Delta' \triangleright \Delta}{\Delta', \underline{x} : \text{FA} \triangleright \Delta}$$

The type system of ECBPV admits a weakening lemma, so if $\Delta' \triangleright \Delta$ we can weaken any $\Delta \vdash V : A$ to get $\Delta' \vdash V : A$, and similarly for computations.

A *Kripke relation* is a family of binary relations indexed by computation-type contexts that respects weakening of terms:

Definition 5.3.1 (Kripke relation) A *Kripke relation* R over a value type A (respectively a computation type D) is a family of relations $R^\Delta \subseteq \text{Term}_A^\Delta \times \text{Term}_A^\Delta$ (respectively $R^\Delta \subseteq \underline{\text{Term}}_D^\Delta \times \underline{\text{Term}}_D^\Delta$) indexed by computation-type contexts Δ such that whenever $\Delta' \triangleright \Delta$, we have $M \in R^\Delta$ implies $M \in R^{\Delta'}$ (where M is weakened). ◀

We consider binary relations on equivalence classes of terms because we want to relate pairs of terms up to \equiv , as we did in Chapter 3 (to prove contextual equivalence).

First we observe that, because computation variables are bound to nonterminating computations, some computations can easily be seen to diverge:

Definition 5.3.2 A computation $\Delta \vdash M : \underline{C}$ is *trivially diverging* if

$$M \equiv N_1 \text{ to } x. N_2$$

for some value type A , computation $N_1 \in \{\underline{x} \mid (\underline{x} : \text{FA}) \in \Delta\} \cup \{\Omega_{\text{FA}}\}$ and computation $\Gamma, \underline{x} : \text{FA} \vdash N_2 : \underline{C}$. ◀

If M is trivially diverging then we cannot necessarily show $M \equiv \Omega_{\underline{C}}$ because N_1 might be \underline{x} . However, we can show this once we bind all of the variables in the computation context Δ to nonterminating computations. We need the Kripke relations we define over computation terms to be *closed under divergence*. (For the rest of this section, we omit the square brackets around equivalence classes.)

Definition 5.3.3 A Kripke relation R over a computation type \underline{C} is *closed under divergence* if it is transitive and each of the following holds:

1. If $\Delta \vdash M : \underline{C}$ and $\Delta \vdash M' : \underline{C}$ are trivially diverging, then $(M, M') \in R^\Delta$.

$$\begin{aligned}
\mathcal{R}[\mathbf{unit}]^\Delta &:= \text{Term}_{\mathbf{unit}} \times \text{Term}_{\mathbf{unit}} \\
\mathcal{R}[A_1 \times A_2]^\Delta &:= \{(V, V') \mid (\mathbf{fst} V, \mathbf{fst} V') \in \mathcal{R}[A_1]^\Delta \wedge (\mathbf{snd} V, \mathbf{snd} V') \in \mathcal{R}[A_2]^\Delta\} \\
\mathcal{R}[\mathbf{empty}]^\Delta &:= \emptyset \\
\mathcal{R}[A_1 + A_2]^\Delta &:= \{(\mathbf{inl}_{A_2} V, \mathbf{inl}_{A_2} V') \mid (V, V') \in \mathcal{R}[A_1]^\Delta\} \cup \\
&\quad \{(\mathbf{inr}_{A_1} V, \mathbf{inr}_{A_1} V') \mid (V, V') \in \mathcal{R}[A_2]^\Delta\} \\
\mathcal{R}[\mathbf{U} \underline{C}]^\Delta &:= \{(V, V') \mid (\mathbf{force} V, \mathbf{force} V') \in \mathcal{R}[\underline{C}]^\Delta\} \\
\mathcal{R}[\mathbf{F} A] &:= \text{the smallest closed-under-divergence Kripke relation such that} \\
&\quad (V, V') \in \mathcal{R}[A]^\Delta \Rightarrow (\langle V \rangle, \langle V' \rangle) \in \mathcal{R}[\mathbf{F} A]^\Delta \\
\mathcal{R}[\mathbf{unit}]^\Delta &:= \underline{\text{Term}}_{\mathbf{unit}} \times \underline{\text{Term}}_{\mathbf{unit}} \\
\mathcal{R}[\underline{C}_1 \times \underline{C}_2]^\Delta &:= \{(M, M') \mid (1'M, 1'M') \in \mathcal{R}[\underline{C}_1]^\Delta \wedge (2'M, 2'M') \in \mathcal{R}[\underline{C}_2]^\Delta\} \\
\mathcal{R}[A \rightarrow \underline{C}]^\Delta &:= \{(M, M') \mid \forall \Delta' \triangleright \Delta, (V, V') \in \mathcal{R}[A]^\Delta, (V'M, V'M') \in \mathcal{R}[\underline{C}]^\Delta\}
\end{aligned}$$

Figure 5.5: Kripke logical relation for ECBPV with nontermination

2. If $\Delta \perp M : \mathbf{F} A$ and $\Delta \perp M' : \mathbf{F} A$ are trivially diverging, and $(N, N') \in R^{\Delta, y:\mathbf{F} A}$, then all four of the following pairs are in R^Δ :

$$\begin{array}{ll}
(M \mathbf{need} y. N, M' \mathbf{need} y. N') & (N[y \mapsto M], N'[y \mapsto M']) \\
(N[y \mapsto M], M' \mathbf{need} y. N') & (M \mathbf{need} y. N, N'[y \mapsto M']) \quad \blacktriangleleft
\end{array}$$

This definition works because all of the trivially diverging computations are interchangeable in the semantics. The second part of the definition is the most important. Since we are showing that **need** is contextually equivalent to substitution, we want these to be related; this justifies the two pairs at the bottom of the definition. We have to consider computation variables in the definition (inside trivially diverging computations) only because of our use of Kripke logical relations. For ordinary logical relations, there would be no free variables to consider.

The key part of the proof of contextual equivalence is the definition of the Kripke logical relation $\mathcal{R}[-]$, which is a family of relations indexed by value and computation types. It is defined in Figure 5.5 by induction on the structure of the types. In the figure, we again omit square brackets around equivalence classes.

The definition on most of the type formers is standard, and similar to the equations required to hold for GCBPV logical relations (Figure 3.1). For returner types, we want any pair of computations that return related values to be related. We also want the relation to be closed under divergence, in order to show the fundamental lemma (below) for **to** and **need**. We therefore define $\mathcal{R}[\mathbf{F} A]$ as the smallest such Kripke relation. For function types, we require as usual that related arguments are sent to related results. For this to define a Kripke relation, we have to quantify over all computation-type contexts Δ' that weaken Δ , because of the contravariance of the argument. The computations M and M' are weakened so that they are in the context Δ' .

The relations we define are Kripke relations. Using the sequencing axioms of ECBPV inequational theories, and the β - and η -laws for computation types, we can show that $\mathcal{R}[\underline{C}]$ is closed under divergence for each computation type \underline{C} . These facts are important for the proof of the fundamental lemma.

We define sets $\text{Subst}_\Gamma^\Delta$ of equivalence classes of substitutions (where \equiv is given component-wise on substitutions), and extend the logical relation by defining $\mathcal{R}[\Gamma]^\Delta \subseteq \text{Subst}_\Gamma^\Delta \times \text{Subst}_\Gamma^\Delta$:

$$\begin{aligned} \text{Subst}_\Gamma^\Delta &:= \{[\sigma]_\equiv \mid \Delta \vdash \sigma : \Gamma\} \\ \mathcal{R}[\diamond]^\Delta &:= \{(\diamond, \diamond)\} \\ \mathcal{R}[\Gamma, x : A]^\Delta &:= \{((\sigma, x \mapsto V), (\sigma', x \mapsto V')) \mid (\sigma, \sigma') \in \mathcal{R}[\Gamma]^\Delta \wedge (V, V') \in \mathcal{R}[A]^\Delta\} \\ \mathcal{R}[\Gamma, \underline{x} : \mathbf{F}A]^\Delta &:= \{((\sigma, \underline{x} \mapsto M), (\sigma', \underline{x} \mapsto M')) \mid (\sigma, \sigma') \in \mathcal{R}[\Gamma]^\Delta \wedge (M, M') \in \mathcal{R}[\mathbf{F}A]^\Delta\} \end{aligned}$$

As usual, the logical relations satisfy a *fundamental lemma*.

Lemma 5.3.4 (Fundamental) Suppose that $(\sigma, \sigma') \in \mathcal{R}[\Gamma]^\Delta$.

1. (Values) If $\Gamma \vdash V : A$ then $(V[\sigma], V[\sigma']) \in \mathcal{R}[A]^\Delta$.
2. (Computations) If $\Gamma \vdash M : \underline{C}$ then $(M[\sigma], M[\sigma']) \in \mathcal{R}[\underline{C}]^\Delta$.

Proof. By induction on the structure of the terms. For two of the cases we use lemmas given in the appendix: Lemma B.3.3 for **need** and Lemma B.3.4 for **to**. For computation variables we use the assumption about σ and σ' . For the operation `undef`, we use the fact that `undef ()` is trivially diverging. The rest of the cases are standard. \square

We also have the following two facts about the logical relation. The first roughly is that **name** is related to **need** by the logical relation, and is true because of the additional pairs that are related in the definition of closed-under-divergence (Definition 5.3.3).

Lemma 5.3.5 For all computation terms $\Gamma \vdash M : \mathbf{F}A$ and $\Gamma, \underline{x} : \mathbf{F}A \vdash N : \underline{C}$, if $(\sigma, \sigma') \in \mathcal{R}[\Gamma]^\Delta$ then

$$((N[\underline{x} \mapsto M])[\sigma], (M \text{ \textbf{need} } \underline{x}. N)[\sigma']) \in \mathcal{R}[\underline{C}]^\Delta$$

Proof. Apply the fundamental lemma (Lemma 5.3.4) to M , and then use Lemma B.3.5. \square

The second fact is that related terms are contextually equivalent (similar to Lemma 3.1.5).

Lemma 5.3.6

1. For all value terms $\Gamma \vdash V : A$ and $\Gamma \vdash V' : A$, if $(V[\sigma], V'[\sigma']) \in \mathcal{R}[A]^\Delta$ for all $(\sigma, \sigma') \in \mathcal{R}[\Gamma]^\Delta$ then

$$\Gamma \vdash V \cong_{\text{ctx}} V' : A$$

2. For all computation terms $\Gamma \vdash M : \underline{C}$ and $\Gamma \vdash M' : \underline{C}$, if $(M[\sigma], M'[\sigma']) \in \mathcal{R}[\underline{C}]^\Delta$ for all $(\sigma, \sigma') \in \mathcal{R}[\Gamma]^\Delta$ then

$$\Gamma \vdash M \cong_{\text{ctx}} M' : \underline{C}$$

Proof. The logical relation is closed under congruence (placing terms in contexts $C[\]$ and $\underline{C}[\]$) by a similar proof to the fundamental lemma (Lemma 5.3.4). Hence it suffices to show that if $N, N' : \mathbf{F}G$ are closed computations, with G a ground type, then $(N, N') \in \mathcal{R}[\mathbf{F}G]^\Delta$ implies $N \equiv N'$. This is proved in the appendix as Corollary B.3.9. \square

This gives us enough to prove the internal equivalence between call-by-name and call-by-need:

Theorem 5.3.7 For all computation terms $\Gamma \vdash M : \mathbf{F}A$ and $\Gamma, \underline{x} : \mathbf{F}A \vdash N : \underline{C}$, we have

$$\Gamma \vdash M \text{ \textbf{name} } \underline{x}. N \cong_{\text{ctx}} M \text{ \textbf{need} } \underline{x}. N : \underline{C}$$

Proof. Apply Lemma 5.3.5 to show that the two computations are related to each other, and then apply Lemma 5.3.6 to obtain the contextual equivalence. \square

We now move on to the meta-level equivalence. Suppose we are given a (possibly open) source-language expression $\Gamma \vdash e : \tau$. Recall that the call-by-need translation uses a context containing computation variables (i.e. $(\Gamma)^{\text{need}}$) and the call-by-name translation uses a context containing value variables, which map to thunks of computations. We have two ECBPV computation terms of type $F(\tau)$ in context $(\Gamma)^{\text{need}}$: one is just $(e)^{\text{need}}$, the other is the call-by-name translation $(e)^{\text{moggi}}$ (defined in Section 2.4) with all of its variables substituted with thunked computations. The theorem then states that these are contextually equivalent.

Theorem 5.3.8 (Equivalence between call-by-name and call-by-need) If e is a source-language expression that satisfies $\underline{x}_1 : \tau_1, \dots, \underline{x}_n : \tau_n \vdash e : \tau'$ then

$$(e)^{\text{moggi}}[x_1 \mapsto \mathbf{thunk} \underline{x}_1, \dots, x_n \mapsto \mathbf{thunk} \underline{x}_n] \cong_{\text{ctx}} (e)^{\text{need}}$$

Proof. By induction on the typing derivation of e . The interesting case is lambda abstraction, where we use the internal equivalence between call-by-name and call-by-need above (Theorem 5.3.7). \square

5.4 An effect system for extended call-by-push-value

The equivalence between call-by-name and call-by-need in the previous section is predicated on the only side-effect in the language being nontermination. We discuss how to relax this restriction so that only subterms need to be restricted to nontermination and the language itself have other side-effects.

Call-by-need makes it difficult to statically estimate effects. Computation variables bound using **need** might have effects on their first use, but on subsequent uses do not. Hence to precisely determine the effects of a term, we must track which variables have been previously used.

McDermott and Mycroft [69] show how to achieve this for a call-by-need effect system, and we expect that this technique can be adapted to ECBPV. Here we take a simpler approach. By slightly restricting the effect algebras we consider, we remove the need to track variable usage information, while still ensuring the effect information is not an underestimate (an underestimate would enable incorrect transformations). This can reduce the precision of the effect information obtained, but for our use case (determining equivalences between evaluation orders) this is not an issue, since we primarily care about which side-effects are used (rather than e.g. how many times they are used).

If we do not track usage information and assume any use of a variable may have side-effects, we might misestimate the effect of a call-by-need computation variable evaluated for a second time (whose true effect is 1). To ensure this misestimate is an overestimate, we require that the effect algebra is *pointed*.

Definition 5.4.1 (Pointed preordered monoid) A preordered monoid $(\mathcal{E}, \leq, \cdot, 1)$ is *pointed* if for all $\varepsilon \in \mathcal{E}$ we have $1 \leq \varepsilon$. \blacktriangleleft

Many of our examples of effect algebras are pointed, including Gifford-style effect algebras (Example 2.1.2).

In our example, where we wish to establish whether the effects of an expression are restricted to nontermination, we use the two-element preorder $\{\Omega \leq \top\}$ with join for sequencing and Ω

$$\begin{array}{c}
\frac{\Gamma \vdash V : A}{\Gamma \vdash \langle V \rangle : \langle 1 \rangle A} \qquad \frac{\Gamma \vdash M : \langle \varepsilon \rangle A \quad \Gamma, x : A \vdash N : \underline{C}}{\Gamma \vdash M \text{ to } x. N : \langle \varepsilon \rangle \underline{C}} \\
\\
\frac{\Gamma \vdash M : \langle \varepsilon \rangle A \quad \Gamma, \underline{x} : \langle \varepsilon \rangle A \vdash N : \underline{C}}{\Gamma \vdash M \text{ need } \underline{x}. N : \underline{C}} \qquad \frac{}{\Gamma \vdash \underline{x} : \langle \varepsilon \rangle A} \text{ if } (\underline{x} : \langle \varepsilon \rangle A) \in \Gamma \\
\\
\frac{\Gamma \vdash M : \langle \varepsilon \rangle A}{\Gamma \vdash \text{coerce}_{\varepsilon \leq \varepsilon'} M : \langle \varepsilon' \rangle A} \text{ if } \varepsilon \leq \varepsilon'
\end{array}$$

Figure 5.6: Effect system modifications to ECBPV

as the unit 1. The effect Ω means side-effects restricted to (at most) nontermination, and \top means unrestricted side-effects. Thus we would enable the equivalence between call-by-name and call-by-need when the effect is Ω , and not when it is \top . This effect algebra is pointed.

The effect system is similar to GCBPV. The syntax is a minor modification of the syntax for ECBPV in Section 5.1. We grade the returner types $\mathbf{F}A$ so that they have the form $\langle \varepsilon \rangle A$, where $\varepsilon \in \mathcal{E}$. The grammar of types is therefore identical to the one for GCBPV:

$$\underline{C}, \underline{D} ::= \mathbf{unit} \mid \underline{C}_1 \times \underline{C}_2 \mid A \rightarrow \underline{C} \mid \langle \varepsilon \rangle A$$

Typing contexts still assign returner types to computation variables, but they now also have effects:

$$\Gamma ::= \diamond \mid \Gamma, x : A \mid \Gamma, \underline{x} : \langle \varepsilon \rangle \underline{C}$$

We need to be able to overapproximate effects, and therefore add explicit subtyping to the grammar of ECBPV terms:

$$\begin{aligned}
M, N ::= & \lambda\{\} \mid \lambda\{1. M_1, 2. M_2\} \mid 1'M \mid 2'M \mid \lambda x:A. M \mid V'M \\
& \mid \text{op } V \mid \langle V \rangle \mid M \text{ to } x. N \mid \text{coerce}_{\varepsilon \leq \varepsilon'} M \mid \text{force } V \mid \underline{x} \mid M \text{ need } x. N
\end{aligned}$$

The syntax of value types and value terms is generated by the same grammar as before (in Section 5.1). The ECBPV effect system is also parameterized by the same notion of signature as for GCBPV (Definition 2.7.1).

We define the action of the preordered monoid of effects on computation types as for GCBPV:

$$\begin{aligned}
\langle \varepsilon \rangle \mathbf{unit} &::= \mathbf{unit} & \langle \varepsilon \rangle (\underline{C}_1 \times \underline{C}_2) &::= \langle \varepsilon \rangle \underline{C}_1 \times \langle \varepsilon \rangle \underline{C}_2 \\
\langle \varepsilon \rangle (A \rightarrow \underline{C}) &::= A \rightarrow \langle \varepsilon \rangle \underline{C} & \langle \varepsilon \rangle (\langle \varepsilon' \rangle A) &::= \langle \varepsilon \cdot \varepsilon' \rangle A
\end{aligned}$$

The typing judgments have exactly the same form as before (except for the new syntax of types). The majority of the typing rules, including all of the rules for values, are also unchanged.

The only rules we change are those for computation variables, returning values, **to** and **need**, which are replaced with the first four rules in Figure 5.6. We also add a subeffecting rule, which is the last rule of the figure. The rules for return, **to** and **coerce** are the same as the GCBPV rules.

We also have to change the inequational theory to add axioms for **coerce** (e.g. transitivity and commutativity with **to**), and also need to add coercions to the axiom

$$M \text{ need } \underline{x}. \underline{x} \text{ to } y. N \equiv M \text{ to } y. N[\underline{x} \mapsto \langle y \rangle]$$

replacing it with

$$M \text{ need } \underline{x}. \underline{x} \text{ to } y. N \equiv M \text{ to } y. N[\underline{x} \mapsto \text{coerce}_{1 \leq \varepsilon} \langle y \rangle]$$

where ε is the effect of M . This works because the effect algebra is pointed, so we have $1 \leq \varepsilon$.

5.4.1 Exploiting effect-dependent equivalences

Our primary goal in adding an effect system to ECBPV is to exploit (local, effect-justified) equivalences between evaluation orders even without a whole-language restriction on effects. We sketch how to do this for our example.

When proving the equivalence between call-by-name and call-by-need in Section 5.3 we assumed that the only operation in the language was diverge. To relax this restriction, we use the effect algebra with preorder $\{\Omega \leq \top\}$ described above, and assign the effect Ω to diverge. We can include other effectful operations (e.g. raise, put), and give them the effect \top . The statement of the internal (object-level) equivalence becomes:

If $\Gamma \vdash M : \langle \Omega \rangle A$ and $\Gamma, \underline{x} : \langle \Omega \rangle A \vdash N : \underline{C}$ then

$$\Gamma \vdash M \text{ name } \underline{x}. N \cong_{\text{ctx}} M \text{ need } \underline{x}. N : \underline{C}$$

The premise restricts the effect of M to Ω so that nontermination is its only possible side-effect, but N is allowed to have *any* effect.

To prove this equivalence, we need a logical relation for the effect system, which means we have to define a Kripke relation $\mathcal{R}[\langle \varepsilon \rangle A]$ for each effect ε . For $\mathcal{R}[\langle \Omega \rangle A]$ we use the same definition as before (the definition of $\mathcal{R}[\langle \text{F} \rangle A]$). The definition of $\mathcal{R}[\langle \top \rangle A]$ depends on the specific other effects included (we could adapt the free lifting in Section 3.1.1).

To state and prove a meta-level equivalence for a source language that includes other side-effects, we need to define a call-by-need effect system for the source language. This would just be the same as the Moggi-style call-by-name effect system in Section 2.4, which is sound for call-by-need for pointed effect algebras. The call-by-need translation of types is then the same as the call-by-name translation of types. Just as for the object-level equivalence, the statement of the meta-level equivalence requires the source-language expression to have the effect Ω . We omit the details here.

5.5 Related work

Reasoning about call-by-need The majority of work on reasoning about call-by-need source languages has concentrated on operational semantics based on environments [53], graphs [100, 96], and answers [4, 3, 17, 68]. However, these do not compare call-by-need with other evaluation orders. The only type-based analysis of a lazy source language we know of apart from McDermott and Mycroft's effect system [69] is Turner et al.'s [97] (and its extension [101]).

Polarized type theories (e.g. Zeilberger [102]) stratify types into several kinds (like values and computations in CBPV) to capture multiple evaluation orders. Downen and Ariola [20, 21] recently described how to capture call-by-need using polarity. They take a different approach to ours, by stratifying the syntax according to evaluation order, rather than whether terms might have effects. This means they have three kinds of type (call-by-value, call-by-need, and call-by-name), resulting in a more complex language than ours. In their language, information about evaluation order can be deduced from types, in ours it cannot (the call-by-need and

call-by-name translations of types are identical). They also do not apply their language to reasoning about the differences between evaluation orders, which was the primary motivation for ECBPV. It is not clear whether their language can also be used for this purpose.

Multiple evaluation orders can also be captured in a Moggi-style language by using *joinads* instead of monads [83]. There may be some joinad structure implicit in extended call-by-push-value.

Logical relations Kripke logical relations have previously been applied to the problems of lambda definability [39] and normalization [2, 25]. Previous proofs of contextual equivalence relate only closed terms. We were forced to relate open terms because of the **need** construct.

Reasoning about effects using logical relations often runs into a difficulty in ensuring the relations are closed under sequencing of computations. We are able to work around this due to our specific choice of effects. It is possible that considering other effects would require a technique such as Lindley and Stark’s *leapfrog method* [62, 61].

5.6 Summary

We have described extended call-by-push-value, a calculus that can be used for reasoning about several evaluation orders. In particular, ECBPV supports call-by-need via the addition of the construct $M \text{ need } \underline{x}. N$. This allows us to prove that call-by-name and call-by-need reduction are equivalent if nontermination is the only effect in the source language, both inside the language itself, and on the meta-level. We proved the latter by giving a translation from our source language into ECBPV that captures call-by-need reduction. We also defined an effect system for ECBPV. The effect system statically bounds the side-effects of terms, allowing us to validate equivalences between evaluation orders without restricting the entire language to particular effects. We end this chapter with a description of some possible future work.

Other equivalences between evaluation orders We have proved one example of an equivalence between evaluation orders using ECBPV, but there are others that we might also expect to hold. For example, we would expect call-by-need and call-by-value to be equivalent if the effects are restricted to nondeterminism, allocating state, and reading from state (but not writing). It should be possible to use ECBPV to prove these by defining suitable logical relations. More generally, it might be possible to characterize when particular equivalences hold in terms of the algebraic properties of the effects we restrict to, and give a reasoning principle similar to the one we derived in Chapter 3.

Denotational semantics In Chapter 4 we emphasized the use of denotational semantics for proving the validity of program transformations. Developing an ECBPV denotational semantics would enable us to reason about ECBPV program transformations in this way. Composing the denotational semantics with the call-by-need translation would also result in a call-by-need denotational semantics for the source language.

Some potential approaches to describing the denotational semantics of ECBPV are Maraist et al.’s [67] translation into an affine calculus, combined with a semantics of linear logic [71], and also continuation-passing-style translations [82]. None of these consider side-effects however.

Chapter 6

Conclusions

This thesis develops a general framework for proving the validity of effect-dependent program transformations. We include an intermediate language in which such transformations can be stated (graded call-by-push-value, defined in Section 2.7), a way to use this intermediate language to reason about a source language (translations into GCBPV in Section 2.7.3), and machinery for proving the instances of contextual preorders (logical relations and order-enriched categorical semantics in Section 3.1 and Chapter 4).

Our framework supports various evaluation orders. The side-effects of programs depend on the evaluation order used. Hence, to discuss effect-dependent transformations for some evaluation order it is desirable to construct an effect system that is bespoke to that evaluation order. It is well-known that we can do this for call-by-value; Chapter 2, we show how to do this for two forms of call-by-name (Moggi-style call-by-name and Levy-style). We also describe GCBPV, which augments Levy’s call-by-push-value with an effect system that generalises all of these. We can therefore use the framework to prove the validity of effect-dependent program transformations for source languages with any of these evaluation orders. This extends previous work, which primarily supports only call-by-value.

We then applied the framework to program transformations that change the evaluation order used inside programs. We proved two versions of a novel reasoning principle that relates call-by-value and call-by-name evaluation in the presence of side-effects: one syntactic (Chapter 3) and one semantic (Section 4.4). The reasoning principle arises because of the existence of certain Galois connections between call-by-value and call-by-name computations. We expect that our approach to relating evaluation orders will work more generally: we can relate two evaluation orders by translating into a single intermediate language and then defining similar Galois connections.

Chapter 4 deals with *noninvertible* program transformations. These have mostly been neglected in previous work, which tends to consider only symmetric transformations. We gave several examples, including for undefined behaviour, nondeterminism and concurrency, and showed how to use our framework to validate them. These justify our use of order-enrichment in the denotational semantics of GCBPV. We argue that order-enrichment is crucial for a *general* framework for formal reasoning about program transformations.

Finally, we added call-by-need to our framework. Call-by-need is more difficult to reason about than call-by-value or call-by-name because it involves *action at a distance*. In Chapter 5 we developed ECBPV, which extends CBPV with primitives that allow it to capture call-by-need evaluation. We then discussed how to use ECBPV to reason about call-by-need. In particular, we showed that call-by-name and call-by-need are equivalent when side-effects are limited to nontermination. It turns out that ordinary logical relations do not suffice to do this, and so we used Kripke logical relations of varying arity.

6.1 Future work

There are many possible directions for future work. Some are obvious extensions of the contributions of this thesis, others are open questions that should be solved to continue with the aims of this thesis. We highlight just a few here.

Relating evaluation orders We have given two relationships between evaluation orders: a general reasoning principle for call-by-value and Levy-style call-by-name in Chapter 3 and Section 4.4, and one instance of a relationship between call-by-need and Moggi-style call-by-name in Section 5.3. An obvious question is whether we can develop the latter into a more general reasoning principle that relies only on axiomatic properties of side-effects. It would also be interesting to formulate reasoning principles for other pairs of evaluation orders. In particular, it should be possible to give a reasoning principle that relates the two forms of call-by-name.

Other side-effects We have considered a small number of example side-effects, such as global state, nondeterminism, nontermination and undefined behaviour. There are many other examples of side-effects that it should be possible to apply our framework to, such as local state [88, 40] and probabilistic choice. Incorporating effect *handlers* [86] into our framework is another obvious next step.

Semantics of call-by-need Throughout Chapter 4 we emphasized the (well-known) advantages of denotational semantics for formal reasoning about program transformations, but when adding call-by-need in Chapter 5 we did not consider denotational semantics at all. Denotational semantics for call-by-need (in the presence of arbitrary side-effects) is still an open problem, and it is not clear what a model of ECBPV would look like. Our use of Kripke logical relations of varying arity suggests we might be able to model call-by-need in (Poset-)categories of the form $[\text{Ctx}, \mathbf{C}]$, where Ctx is the category of typing contexts and weakenings, but we have not pursued this.

How should models be graded? In Chapter 4 we construct a number of graded adjunctions that we use as models of GCBPV. A question we might ask is *which* graded adjunctions are useful as models. For example, we might expect each of the subeffecting morphisms $(\varepsilon \leq \varepsilon') \otimes (-) : \varepsilon \otimes (-) \rightarrow \varepsilon' \otimes (-)$ to be a full monomorphism (Definition 4.2.15), which corresponds to $\text{coerce}_{\varepsilon \leq \varepsilon'} M \leq \text{coerce}_{\varepsilon \leq \varepsilon'} N \Rightarrow M \leq N$ in the syntax. (This is the case for all of our examples.) There are graded adjunctions that do not have this property, but it is not clear whether any of them are useful for models. Are there any other properties we would expect to hold for all side-effects? Constructing graded adjunctions also takes a lot of work, so it would be useful to find a general technique for doing this. Kammar and McDermott [41] consider such a technique for constructing graded monads for Gifford-style effect algebras.

Intermediate languages for program reasoning We have presented GCBPV as an intermediate language that is useful for formal reasoning about programs. In particular, we chose to base our intermediate language on CBPV partly because we can easily verify the correctness of program transformations in CBPV. However, we do not claim that GCBPV is the best possible intermediate language. More work is needed to settle the question of what the best language is. If we are proving the validity of compiler optimizations then we might want to

find an intermediate language that can be used both for formal reasoning and for applying optimizations inside a compiler. There may be a tension between these two applications.

Practical effect systems There are very few instances of effect systems being used in practice, and this thesis does not consider *practical* use of grading. One way to use effect systems as a way to do effect-dependent optimizations inside a compiler would be to have the compiler infer the effects of terms (rather than exposing the effect system to the programmer). This would require effect inference, and to do inference we would probably want to include effect polymorphism. Extending GCBPV (and related machinery, such as the denotational semantics), with effect polymorphism, and then describing how to do effect inference on it, would be interesting future work.

6.2 Final remarks

There are several key aspects of the approach used throughout this thesis that are worth highlighting here.

The first is our use of a common *intermediate* language for stating and proving the validity of program transformations, inspired by compiler design. Intermediate languages are useful because they allow us to prove *general* results about program transformations. When relating call-by-value and call-by-name in Chapter 3 and Section 4.4 we obtained a reasoning principle that is not specific to the choice of side-effects in the language. This was aided partly by our use of inequational theories for specifying the behaviour of side-effects (rather than e.g. operational semantics). The structure of the translations into the intermediate language guided us in doing this. Our results about the relationship between call-by-value and call-by-name, and between call-by-name and call-by-need (Section 5.3) are also about open terms. We can relate open terms because we use intermediate languages that capture all of these evaluation orders.

The intermediate languages approach also allows us to ignore considerations about the design of source languages. We chose to base our intermediate languages on CBPV. Whether CBPV is good for writing programs is irrelevant, as long as there are suitable translations from source languages. On the other hand, the design of CBPV makes it is easy to reason about. Crucially, there are no choices to make about evaluation order in CBPV. We therefore get a separation of concerns: the translations from the source language specify the evaluation order but are independent of the choice of side-effects, and the semantics of the intermediate language specifies the behaviour of the side-effects, but is independent of the evaluation order. This was particularly helpful in Chapter 4, where we barely had to consider evaluation orders.

Finally, the intermediate language can include features that may not be useful in source languages. Call-by-name might be an example of this: even if we do not want to use call-by-name in an effectful *source* language, we might still want to replace call-by-value with call-by-name as an optimization inside the intermediate language.

Another important aspect of our approach is the use of *grading* to specify restrictions on side-effects. This is crucial for considering *effect-dependent* program transformations without restricting the side-effects that can appear in the language itself. It allows us to consider languages with a variety of effects (I/O, mutable state, nontermination, nondeterminism, probabilistic choice, etc.) even though most transformations are only correct for subterms that do not use most of them. Another advantage of grading is that we can use it to obtain *stable* semantics [16], where side-effects can be added without changing the denotations of existing programs.

We also emphasize the use of axiomatic properties of side-effects (such as *thunkable* effects in Chapter 3), which we used to give general reasoning principles for program transformations.

These aspects of our approach, especially our use of intermediate languages, may turn out to be helpful more widely in program semantics.

Bibliography

- [1] Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon G. Riecke. 1999. A Core Calculus of Dependency. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 147–160. <https://doi.org/10.1145/292540.292555>
- [2] Thorsten Altenkirch, Martin Hofmann, and Thomas Streicher. 1995. Categorical Reconstruction of a Reduction Free Normalization Proof. *Lecture Notes in Computer Science* 953, 182–199. https://doi.org/10.1007/3-540-60164-3_27
- [3] Zena M. Ariola and Matthias Felleisen. 1997. The call-by-need lambda calculus. *Journal of functional programming* 7, 3, 265–301.
- [4] Zena M. Ariola, John Maraist, Martin Odersky, Matthias Felleisen, and Philip Wadler. 1995. A Call-by-need Lambda Calculus. In *Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 233–246. <http://doi.acm.org/10.1145/199448.199507>
- [5] Nick Benton and Peter Buchlovsky. 2007. Semantics of an Effect Analysis for Exceptions. In *Proceedings of the 2007 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation*. ACM, 15–26. <https://doi.org/10.1145/1190315.1190320>
- [6] Nick Benton, Martin Hofmann, and Vivek Nigam. 2016. Effect-dependent Transformations for Concurrent Programs. In *Proceedings of the 18th International Symposium on Principles and Practice of Declarative Programming*. ACM, 188–201. <https://doi.org/10.1145/2967973.2968602>
- [7] Nick Benton, John Hughes, and Eugenio Moggi. 2002. Monads and Effects. In *Applied Semantics*. Springer, 42–122. <http://dl.acm.org/citation.cfm?id=647424.725798>
- [8] Nick Benton and Andrew Kennedy. 1999. Monads, Effects and Transformations. *Electronic Notes in Theoretical Computer Science* 26, 3–20. [https://doi.org/10.1016/S1571-0661\(05\)80280-4](https://doi.org/10.1016/S1571-0661(05)80280-4)
- [9] Nick Benton, Andrew Kennedy, Lennart Beringer, and Martin Hofmann. 2009. Relational Semantics for Effect-based Program Transformations: Higher-order Store. In *Proceedings of the 11th ACM SIGPLAN Conference on Principles and Practice of Declarative Programming*. ACM, 301–312. <https://doi.org/10.1145/1599410.1599447>
- [10] Nick Benton, Andrew Kennedy, Martin Hofmann, and Lennart Beringer. 2006. Reading, Writing and Relations: Towards Extensional Semantics for Effect Analyses. In *Proceedings of the 4th Asian Conference on Programming Languages and Systems*. Springer, 114–130. https://doi.org/10.1007/11924661_7

- [11] Nick Benton, Andrew Kennedy, Martin Hofmann, and Vivek Nigam. 2016. Counting Successes: Effects and Transformations for Non-deterministic Programs. In *A List of Successes That Can Change the World: Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*, Sam Lindley, Conor McBride, Phil Trinder, and Don Sannella (Eds.). Springer, 56–72. https://doi.org/10.1007/978-3-319-30936-1_3
- [12] Nick Benton, Andrew Kennedy, and George Russell. 1998. Compiling Standard ML to Java Bytecodes. In *Proceedings of the Third ACM SIGPLAN International Conference on Functional Programming*. ACM, 129–140. <https://doi.org/10.1145/289423.289435>
- [13] Lars Birkedal, Filip Sieczkowski, and Jacob Thamsborg. 2012. A Concurrent Logical Relation. In *26th International Workshop/21st Annual Conference of the EACSL (Leibniz International Proceedings in Informatics)*, Patrick Cégielski and Arnaud Durand (Eds.), Vol. 16. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 107–121. <https://doi.org/10.4230/LIPIcs.CSL.2012.107>
- [14] Anna Bucalo, Carsten Führmann, and Alex Simpson. 2003. An equational notion of lifting monad. *Theoretical Computer Science* 294, 1, 31–60. [https://doi.org/10.1016/S0304-3975\(01\)00243-2](https://doi.org/10.1016/S0304-3975(01)00243-2)
- [15] Aurelio Carboni, Stephen Lack, and R.F.C. Walters. 1993. Introduction to extensive and distributive categories. *Journal of Pure and Applied Algebra* 84, 2, 145–158. [https://doi.org/10.1016/0022-4049\(93\)90035-R](https://doi.org/10.1016/0022-4049(93)90035-R)
- [16] Robert Cartwright and Matthias Felleisen. 1994. Extensible Denotational Language Specifications. In *Proceedings of the International Conference on Theoretical Aspects of Computer Software*. Springer, 244–272. <http://dl.acm.org/citation.cfm?id=645868.668496>
- [17] Stephen Chang and Matthias Felleisen. 2012. The call-by-need lambda calculus, revisited. In *Proceedings of the 21st European Conference on Programming Languages and Systems*. Springer, 128–147. http://dx.doi.org/10.1007/978-3-642-28869-2_7
- [18] Pierre-Louis Curien and Hugo Herbelin. 2000. The Duality of Computation. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming*. ACM, 233–243. <https://doi.org/10.1145/351240.351262>
- [19] Marco Devesas Campos and Paul Blain Levy. 2018. A Syntactic View of Computational Adequacy. In *Foundations of Software Science and Computation Structures*, Christel Baier and Ugo Dal Lago (Eds.). Springer, 71–87.
- [20] Paul Downen and Zena M. Ariola. 2018. Beyond Polarity: Towards a Multi-Discipline Intermediate Language with Sharing. In *27th EACSL Annual Conference on Computer Science Logic*. 21:1–21:23. <https://doi.org/10.4230/LIPIcs.CSL.2018.21>
- [21] Paul Downen and Zena M. Ariola. 2019. Compiling With Classical Connectives. *arXiv e-prints*, Article 1907.13227.
- [22] Anatolij Dvurečenskij and Sylvia Pulmannová. 2000. *New trends in quantum structures*. Mathematics and its Applications, Vol. 516. Kluwer Academic Publishers.
- [23] Andrzej Filinski. 1989. *Declarative Continuations and Categorical Duality*. Master’s thesis. University of Copenhagen.

- [24] Andrzej Filinski. 1996. *Controlling Effects*. Ph.D. Dissertation. Carnegie Mellon University.
- [25] Marcelo Fiore. 2002. Semantic Analysis of Normalisation by Evaluation for Typed Lambda Calculus. In *Proceedings of the 4th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*. ACM, 26–37. <https://doi.org/10.1145/571157.571161>
- [26] Marcelo Pablo Fiore. 1996. *Axiomatic Domain Theory in Categories of Partial Maps*. Cambridge University Press.
- [27] Carsten Führmann. 1999. Direct models of the computational lambda-calculus. *Electronic Notes in Theoretical Computer Science* 20, 245–292.
- [28] Soichiro Fujii, Shin-ya Katsumata, and Paul-André Melliès. 2016. Towards a Formal Theory of Graded Monads. In *Foundations of Software Science and Computation Structures*, Bart Jacobs and Christof Löding (Eds.). Springer, 513–530.
- [29] Marco Gaboardi, Shin-ya Katsumata, Dominic Orchard, Flavien Breuvert, and Tarmo Uustalu. 2016. Combining Effects and Coeffects via Grading. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming*. ACM, 476–489. <https://doi.org/10.1145/2951913.2951939>
- [30] Sergey Goncharov and Lutz Schröder. 2013. A Relatively Complete Generic Hoare Logic for Order-Enriched Effects. In *Proceedings of the 28th Annual ACM/IEEE Symposium on Logic in Computer Science*. IEEE Computer Society, 273–282. <https://doi.org/10.1109/LICS.2013.33>
- [31] Colin S. Gordon. 2017. A Generic Approach to Flow-Sensitive Polymorphic Effects. In *31st European Conference on Object-Oriented Programming*, Peter Müller (Ed.), Vol. 74. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 13:1–13:31. <https://doi.org/10.4230/LIPIcs.ECOOP.2017.13>
- [32] Jean Goubault-Larrecq, Slawomir Lasota, and David Nowak. 2002. Logical Relations for Monadic Types. In *Computer Science Logic*, Julian Bradfield (Ed.). Springer, 553–568.
- [33] Jennifer Hackett and Graham Hutton. 2019. Call-by-need is Clairvoyant Call-by-value. *Proceedings of the ACM Programming Languages* 3, Article 114, 23 pages. <https://doi.org/10.1145/3341718>
- [34] John Hatcliff and Olivier Danvy. 1997. Thunks and the λ -calculus. *Journal of Functional Programming* 7, 3, 303–319. <https://doi.org/10.1017/S0956796897002748>
- [35] John Mark Hatcliff. 1995. *The structure of continuation-passing styles*. Ph.D. Dissertation. Kansas State University.
- [36] Claudio Hermida. 1993. *Fibrations, Logical Predicates and Indeterminates*. Ph.D. Dissertation. University of Edinburgh.
- [37] Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. 2017. A convenient category for higher-order probability theory. In *Proceedings of the 32nd Annual IEEE Symposium on Logic in Computer Science*. IEEE Computer Society Press, 1–12. <https://doi.org/10.1109/LICS.2017.8005137>
- [38] C.A.R. Hoare and He Jifeng. 1998. *Unifying Theories of Programming*. Prentice Hall.

- [39] Achim Jung and Jerzy Tiuryn. 1993. A New Characterization of Lambda Definability. In *Proceedings of the International Conference on Typed Lambda Calculi and Applications*. Springer, 245–257. <http://dl.acm.org/citation.cfm?id=645891.671429>
- [40] Ohad Kammar, Paul B. Levy, Sean K. Moss, and Sam Staton. 2017. A monad for full ground reference cells. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science*. IEEE Press, 49:1–49:12. <http://dl.acm.org/citation.cfm?id=3329995.3330044>
- [41] Ohad Kammar and Dylan McDermott. 2018. Factorisation Systems for Logical Relations and Monadic Lifting in Type-and-effect System Semantics. *Electronic Notes in Theoretical Computer Science* 341, 239 – 260. <https://doi.org/10.1016/j.entcs.2018.11.012> Proceedings of the Thirty-Fourth Conference on the Mathematical Foundations of Programming Semantics (MFPS XXXIV).
- [42] Ohad Kammar and Gordon D. Plotkin. 2012. Algebraic Foundations for Effect-dependent Optimisations. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 349–360. <https://doi.org/10.1145/2103656.2103698>
- [43] Shin-ya Katsumata. 2005. A Semantic Formulation of $\top\top$ -lifting and Logical Predicates for Computational Metalanguage. In *Proceedings of the 19th International Conference on Computer Science Logic*. Springer, 87–102. https://doi.org/10.1007/11538363_8
- [44] Shin-ya Katsumata. 2013. Relating Computational Effects by $\top\top$ -lifting. *Inf. Comput.* 222, 228–246. <https://doi.org/10.1016/j.ic.2012.10.014>
- [45] Shin-ya Katsumata. 2014. Parametric Effect Monads and Semantics of Effect Systems. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 633–645. <https://doi.org/10.1145/2535838.2535846>
- [46] Shin-ya Katsumata and Tetsuya Sato. 2015. Codensity Liftings of Monads. In *6th Conference on Algebra and Coalgebra in Computer Science*, Lawrence S. Moss and Pawel Sobocinski (Eds.), Vol. 35. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 156–170. <https://doi.org/10.4230/LIPIcs.CALCO.2015.156>
- [47] G.M. Kelly and A.J. Power. 1993. Adjunctions whose counits are coequalizers, and presentations of finitary enriched monads. *Journal of Pure and Applied Algebra* 89, 1, 163 – 179. [https://doi.org/10.1016/0022-4049\(93\)90092-8](https://doi.org/10.1016/0022-4049(93)90092-8)
- [48] Max Kelly. 1982. *Basic concepts of enriched category theory*. Cambridge University Press.
- [49] Anders Kock. 1971. Bilinearity and cartesian closed monads. *Math. Scand.* 29, 2, 161–174. <https://doi.org/10.7146/math.scand.a-11042>
- [50] Anders Kock. 1972. Strong functors and monoidal monads. *Archiv der Mathematik* 23, 113–120.
- [51] Anders Kock. 1995. Monads for which structures are adjoint to units. *Journal of Pure and Applied Algebra* 104, 1, 41–59.
- [52] Jakov Kučan. 1998. Retraction Approach to CPS Transform. *Higher-Order and Symbolic Computation* 11, 2, 145–175. <https://doi.org/10.1023/A:1010012532463>

- [53] John Launchbury. 1993. A Natural Semantics for Lazy Evaluation. In *Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 144–154. <https://doi.org/10.1145/158511.158618>
- [54] Julia L. Lawall and Olivier Danvy. 1993. Separating Stages in the Continuation-passing Style Transformation. In *Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 124–136. <https://doi.org/10.1145/158511.158613>
- [55] Juneyoung Lee, Yoonseung Kim, Youngju Song, Chung-Kil Hur, Sanjoy Das, David Majnemer, John Regehr, and Nuno P. Lopes. 2017. Taming Undefined Behavior in LLVM. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM, 633–647. <https://doi.org/10.1145/3062341.3062343>
- [56] Paul Blain Levy. 1999. Call-by-Push-Value: A Subsuming Paradigm. In *Typed Lambda Calculi and Applications*, Jean-Yves Girard (Ed.). Springer, 228–243. https://doi.org/10.1007/3-540-48959-2_17
- [57] Paul Blain Levy. 2001. *Call-by-push-value*. Ph.D. Dissertation. Queen Mary, University of London, UK.
- [58] Paul Blain Levy. 2003. Adjunction Models For Call-By-Push-Value With Stacks. *Electronic Notes in Theoretical Computer Science* 69, 248–271. [https://doi.org/10.1016/S1571-0661\(04\)80568-1](https://doi.org/10.1016/S1571-0661(04)80568-1) CTCS’02, Category Theory and Computer Science.
- [59] Paul Blain Levy. 2006. Call-by-push-value: Decomposing call-by-value and call-by-name. *Higher-Order and Symbolic Computation* 19, 4, 377–414. <https://doi.org/10.1007/s10990-006-0480-6>
- [60] Paul Blain Levy. 2017. Contextual Isomorphisms. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages*. ACM, 400–414. <https://doi.org/10.1145/3009837.3009898>
- [61] Sam Lindley. 2005. *Normalisation by Evaluation in the Compilation of Typed Functional Programming Languages*. Ph.D. Dissertation. University of Edinburgh, UK.
- [62] Sam Lindley and Ian Stark. 2005. Reducibility and $\top\top$ -lifting for Computation Types. In *Proceedings of the 7th International Conference on Typed Lambda Calculi and Applications*. Springer, 262–277. https://doi.org/10.1007/11417170_20
- [63] J. M. Lucassen and D. K. Gifford. 1988. Polymorphic Effect Systems. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 47–57. <https://doi.org/10.1145/73560.73564>
- [64] Christoph Lüth and Neil Ghani. 1997. Monads and modular term rewriting. In *Category Theory and Computer Science*, Eugenio Moggi and Giuseppe Rosolini (Eds.). Springer, 69–86.
- [65] QingMing Ma and John C. Reynolds. 1992. Types, abstraction, and parametric polymorphism, part 2. In *Mathematical Foundations of Programming Semantics*, Stephen Brookes, Michael Main, Austin Melton, Michael Mislove, and David Schmidt (Eds.). Springer, 1–40.

- [66] Saunders Mac Lane. 1998. *Categories for the working mathematician* (second ed.). Springer.
- [67] John Maraist, Martin Odersky, David N. Turner, and Philip Wadler. 1995. Call-by-Name, Call-by-Value, Call-by-Need, and the Linear Lambda Calculus, In Proceedings of the Eleventh Annual Mathematical Foundations of Programming Semantics Conference. *Electronic Notes in Theoretical Computer Science*, 370–392.
- [68] John Maraist, Martin Odersky, and Philip Wadler. 1998. The call-by-need lambda calculus. *Journal of Functional Programming* 8, 3, 275–317. <https://doi.org/10.1017/S0956796898003037>
- [69] Dylan McDermott and Alan Mycroft. 2018. Call-by-need effects via coeffects. *Open Computer Science* 8, 93–108. <https://doi.org/10.1515/comp-2018-0009>
- [70] Dylan McDermott and Alan Mycroft. 2019. Extended Call-by-Push-Value: Reasoning About Effectful Programs and Evaluation Order. In *Programming Languages and Systems*, Luís Caires (Ed.). Springer, 235–262.
- [71] Paul-André Melliès. 2009. Categorical semantics of linear logic. In *Interactive Models of Computation and Program Behaviour, Panoramas et Synthèses 27, Société Mathématique de France*.
- [72] Paul-André Melliès. 2010. Segal Condition Meets Computational Effects. In *25th Annual IEEE Symposium on Logic in Computer Science*. 150–159. <https://doi.org/10.1109/LICS.2010.46>
- [73] A Melton, D A Schmidt, and G E Strecker. 1986. Galois Connections and Computer Science Applications. In *Proceedings of a Tutorial and Workshop on Category Theory and Computer Programming*. Springer, 299–312. <http://dl.acm.org/citation.cfm?id=20081.20099>
- [74] José Meseguer. 1980. Varieties of chain-complete algebras. *Journal of Pure and Applied Algebra* 19, 347–383. [https://doi.org/10.1016/0022-4049\(80\)90106-1](https://doi.org/10.1016/0022-4049(80)90106-1)
- [75] Albert R. Meyer and Mitchell Wand. 1985. Continuation semantics in typed lambda-calculi. In *Logics of Programs*, Rohit Parikh (Ed.). Springer, 219–224.
- [76] Stefan Milius, Dirk Pattinson, and Lutz Schröder. 2015. Generic Trace Semantics and Graded Monads. In *6th Conference on Algebra and Coalgebra in Computer Science*, Lawrence S. Moss and Pawel Sobocinski (Eds.), Vol. 35. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 253–269. <https://doi.org/10.4230/LIPIcs.CALCO.2015.253>
- [77] Eugenio Moggi. 1989. Computational lambda-calculus and monads. In *Proceedings of the Fourth Annual Symposium on Logic in Computer Science*. IEEE Press, 14–23. <http://dl.acm.org/citation.cfm?id=77350.77353>
- [78] Eugenio Moggi. 1991. Notions of Computation and Monads. *Inf. Comput.* 93, 1, 55–92. [https://doi.org/10.1016/0890-5401\(91\)90052-4](https://doi.org/10.1016/0890-5401(91)90052-4)
- [79] Guillaume Munch-Maccagnoni. 2014. Models of a Non-associative Composition. In *Foundations of Software Science and Computation Structures*, Anca Muscholl (Ed.). Springer, 396–410.

- [80] Alan Mycroft, Dominic Orchard, and Tomas Petricek. 2016. Effect Systems Revisited—Control-Flow Algebra and Semantics. In *Semantics, Logics, and Calculi: Essays Dedicated to Hanne Riis Nielson and Flemming Nielson on the Occasion of Their 60th Birthdays*. Springer, 1–32. https://doi.org/10.1007/978-3-319-27810-0_1
- [81] Flemming Nielson and Hanne Riis Nielson. 1999. Type and Effect Systems. In *Correct System Design, Recent Insight and Advances*. Springer, 114–136. <http://dl.acm.org/citation.cfm?id=646005.673740>
- [82] Chris Okasaki, Peter Lee, and David Tarditi. 1994. Call-by-need and continuation-passing style. *LISP and Symbolic Computation* 7, 57–81. <https://doi.org/10.1007/BF01019945>
- [83] Tomas Petricek and Don Syme. 2011. Joinads: A Retargetable Control-flow Construct for Reactive, Parallel and Concurrent Programming. In *Proceedings of the 13th International Conference on Practical Aspects of Declarative Languages*. Springer, 205–219. <http://dl.acm.org/citation.cfm?id=1946313.1946336>
- [84] G.D. Plotkin. 1975. Call-by-name, call-by-value and the λ -calculus. *Theoretical Computer Science* 1, 2, 125–159. [https://doi.org/10.1016/0304-3975\(75\)90017-1](https://doi.org/10.1016/0304-3975(75)90017-1)
- [85] Gordon Plotkin and John Power. 2003. Algebraic Operations and Generic Effects. *Applied Categorical Structures* 11, 69–94. <https://doi.org/10.1023/A:1023064908962>
- [86] Gordon Plotkin and Matija Pretnar. 2009. Handlers of Algebraic Effects. In *Proceedings of the 18th European Symposium on Programming Languages and Systems: Held As Part of the Joint European Conferences on Theory and Practice of Software*. Springer, 80–94. https://doi.org/10.1007/978-3-642-00590-9_7
- [87] Gordon D. Plotkin and John Power. 2001. Adequacy for Algebraic Effects. In *Proceedings of the 4th International Conference on Foundations of Software Science and Computation Structures*. Springer, 1–24. <http://dl.acm.org/citation.cfm?id=646793.704708>
- [88] Gordon D. Plotkin and John Power. 2002. Notions of Computation Determine Monads. In *Proceedings of the 5th International Conference on Foundations of Software Science and Computation Structures*. Springer, 342–356. <http://dl.acm.org/citation.cfm?id=646794.704856>
- [89] John C. Reynolds. 1974. On the Relation Between Direct and Continuation Semantics. In *Proceedings of the 2nd Colloquium on Automata, Languages and Programming*. Springer, 141–156. <http://dl.acm.org/citation.cfm?id=646230.681878>
- [90] John C. Reynolds. 1997. The Essence of Algol. In *Algol-like Languages*, Peter W. O’Hearn and Robert D. Tennent (Eds.). Birkhäuser Boston, 67–88. https://doi.org/10.1007/978-1-4612-4118-8_4
- [91] Amr Sabry. 1998. What is a Purely Functional Language? *Journal of Functional Programming* 8, 1, 1–22. <https://doi.org/10.1017/S0956796897002943>
- [92] Amr Sabry and Matthias Felleisen. 1992. Reasoning About Programs in Continuation-passing Style. In *Proceedings of the 1992 ACM Conference on LISP and Functional Programming*. ACM, 288–298. <https://doi.org/10.1145/141471.141563>

- [93] Amr Sabry and Philip Wadler. 1996. A Reflection on Call-by-value. In *Proceedings of the First ACM SIGPLAN International Conference on Functional Programming*. ACM, 13–24. <https://doi.org/10.1145/232627.232631>
- [94] Ross Tate. 2013. The Sequential Semantics of Producer Effect Systems. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 15–26. <https://doi.org/10.1145/2429069.2429074>
- [95] Andrew Tolmach. 1998. Optimizing ML using a hierarchy of monadic types. In *Types in Compilation*, Xavier Leroy and Atsushi Ohori (Eds.). Springer, 97–115.
- [96] D. A. Turner. 1979. A new implementation technique for applicative languages. *Software: Practice and Experience* 9, 1, 31–49. <https://doi.org/10.1002/spe.4380090105>
- [97] David N. Turner, Philip Wadler, and Christian Mossin. 1995. Once Upon a Type. In *Proceedings of the Seventh International Conference on Functional Programming Languages and Computer Architecture*. ACM, 1–11. <http://doi.acm.org/10.1145/224164.224168>
- [98] Philip Wadler. 1998. The Marriage of Effects and Monads. In *Proceedings of the Third ACM SIGPLAN International Conference on Functional Programming*. ACM, 63–74. <https://doi.org/10.1145/289423.289429>
- [99] Philip Wadler. 2003. Call-by-value is Dual to Call-by-name. In *Proceedings of the Eighth ACM SIGPLAN International Conference on Functional Programming*. ACM, 189–201. <https://doi.org/10.1145/944705.944723>
- [100] C.P. Wadsworth. 1971. *Semantics and Pragmatics of the Lambda-calculus*. University of Oxford.
- [101] Keith Wansbrough and Simon Peyton Jones. 1999. Once Upon a Polymorphic Type. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 15–28. <http://doi.acm.org/10.1145/292540.292545>
- [102] Noam Zeilberger. 2009. *The Logical Basis of Evaluation Order and Pattern-matching*. Ph.D. Dissertation. Carnegie Mellon University.

Appendix A

Order-enriched category theory

We give a brief introduction to the concepts of order-enriched (specifically **Poset**-enriched) category theory that we use for the denotational semantics of GCBPV in Chapter 4. A more comprehensive introduction to enriched category theory in general is given by Kelly [48] (though note that restricting to **Poset**-enrichment simplifies the definitions significantly). We assume some basic (ordinary) category theory (e.g. Mac Lane [66]).

A **Poset**-category is like an ordinary category, except that morphisms form partially ordered sets and composition of morphisms is monotone.

Definition A.0.1 (Poset-category) A **Poset**-category \mathbf{C} consists of a collection of objects and, for each pair of objects X, Y , a poset $\mathbf{C}(X, Y)$ of morphisms, together with:

- For each triple of objects X, Y, Z , a composition function $\circ : \mathbf{C}(Y, Z) \times \mathbf{C}(X, Y) \rightarrow \mathbf{C}(X, Z)$. Composition is required to be monotone (in both arguments) and associative.
- For each object X , an identity morphism $id_X \in \mathbf{C}(X, X)$ that is the unit for composition.

◀

We identify each poset (including objects of **Poset** below) with its underlying set, and always use the symbol \sqsubseteq for the order on morphisms in a **Poset**-category. Every **Poset**-category has an underlying ordinary category, by forgetting the order on morphisms. We give several examples, all of which are useful for constructing models of GCBPV.

Example A.0.2 Every locally small category (in particular, the category **Set** of sets and functions) forms a **Poset**-category by using equality for each of the partial orders. In these cases, all of the **Poset**-enriched definitions coincide with the ordinary ones.

◀

Example A.0.3 The prototypical example of a **Poset**-category is **Poset** itself, which has partially ordered sets (X, \sqsubseteq) as objects and monotone functions as morphisms. Morphisms are ordered pointwise (i.e. $f \sqsubseteq g$ if $f x \sqsubseteq g x$ for all x). When we use **Poset** as a **Poset**-category, this is the order on morphisms we use.

◀

Example A.0.4 We also restrict **Poset** to a smaller category. A poset X is *pointed* if it has a least element \perp , and a monotone function $f : X \rightarrow Y$ between pointed posets is *strict* if $f \perp = \perp$. The **Poset**-category \mathbf{Poset}_\perp is formed by restricting **Poset** to pointed posets and strict monotone functions.

◀

Example A.0.5 A poset (X, \sqsubseteq) is an ω -complete partial order if every ω -chain $x_0 \sqsubseteq x_1 \sqsubseteq \dots$ has a least upper bound $\bigsqcup x$, and a monotone function between two ω cpos is *continuous* if it preserves least upper bounds of ω -chains. The **Poset**-category $\omega\mathbf{Cpo}$ has ω cpos as objects and continuous functions as morphisms. The order on morphisms is pointwise.

◀

As these examples show, each ordinary category may have more than one possible order-enrichment. When using **Poset**-categories as models of GCBPV, the order on morphisms must be chosen based on the side-effects being modelled. The order on morphisms is the only extra structure required in **Poset**-category theory compared to ordinary category theory. The remaining definitions are the same as the ordinary ones, except that wherever there is a function on morphisms it is required to be monotone.

To model various type formers, we require **Poset**-categories \mathbf{C} that come with some given extra structure. A *terminal object* 1 of a **Poset**-category is just a terminal object in the underlying ordinary category. We write $\langle \rangle_X : X \rightarrow 1$ for the universal morphisms. Binary *products* are also the same as in the underlying category, except that the pairing operations $\langle -, - \rangle$ on morphisms are required to be monotone (if $f_1 \sqsubseteq f'_1 : X \rightarrow Y_1$ and $f_2 \sqsubseteq f'_2 : X \rightarrow Y_2$ then $\langle f_1, f_2 \rangle \sqsubseteq \langle f'_1, f'_2 \rangle : X \rightarrow Y_1 \times Y_2$). For each X, Y, Z , there is a canonical associativity morphism $\text{assoc} : (X \times Y) \times Z \rightarrow X \times (Y \times Z)$. Similarly, binary *coproducts* (with copairing morphisms $[f_1, f_2] : X_1 + X_2 \rightarrow Y$) are the same as in the underlying ordinary category, except that the copairing operations are required to be monotone. *Exponentials* in **Poset**-categories are just exponentials in the underlying category, except that the currying operations $\Lambda : \mathbf{C}(Z \times X, Y) \rightarrow \mathbf{C}(Z, X \Rightarrow Y)$ are required to be monotone. This implies that uncurrying $\Lambda^{-1} : \mathbf{C}(Z, X \Rightarrow Y) \rightarrow \mathbf{C}(Z \times X, Y)$ is also monotone.

We say that a **Poset**-category \mathbf{C} is *cartesian* if it has chosen finite products (terminal object and binary products). It is *bicartesian* if it is cartesian and has chosen finite coproducts (initial object and binary coproducts). It is *distributive* if it is bicartesian and, for all $X, Y, Z \in \mathbf{C}$, the morphism

$$(X \times Y) + (X \times Z) \xrightarrow{[\langle \pi_1, \text{inl} \circ \pi_2 \rangle, \langle \pi_1, \text{inr} \circ \pi_2 \rangle]} X \times (Y + Z)$$

has an inverse $\text{dist}_{X,Y,Z} : X \times (Y + Z) \rightarrow (X \times Y) + (X \times Z)$. A **Poset**-category \mathbf{C} is *(bi)cartesian closed* if it is (bi)cartesian and has chosen exponentials. Every bicartesian closed **Poset**-category is distributive.

Example A.0.6 The **Poset**-categories **Set**, **Poset** and $\omega\mathbf{Cpo}$ are bicartesian closed. For **Set**, the terminal object is the singleton $\{\star\}$, and the binary product $X_1 \times X_2$ is the usual cartesian product. The initial object is the empty set, and the coproduct $X_1 + X_2$ is the disjoint union. For **Poset** and $\omega\mathbf{Cpo}$ each of these is the same as in **Set**, with the obvious orderings. In **Set** exponentials are sets of functions, in **Poset** they are sets of monotone functions, ordered pointwise, and in $\omega\mathbf{Cpo}$ they are sets of continuous functions, ordered pointwise.

The **Poset**-category \mathbf{Poset}_\perp is cartesian; finite products are the same as in **Poset**. (It is also bicartesian, but not distributive or closed.) \blacktriangleleft

For models of GCBPV we also need **Poset**-adjunctions. To define these we first consider **Poset**-enrichment of functors. This is property of functors, not some additional structure that must be chosen: each ordinary functor either does or does not **Poset**-enrich.

Definition A.0.7 (Poset-functor) If \mathbf{C} and \mathbf{D} are **Poset**-categories, a **Poset**-functor $F : \mathbf{C} \rightarrow \mathbf{D}$ is a functor between the underlying ordinary categories, such that each function $F : \mathbf{C}(X, Y) \rightarrow \mathbf{D}(FX, FY)$ is monotone. *Natural transformations* between **Poset**-functors are just natural transformations in the ordinary sense. \blacktriangleleft

The identity $\text{Id}_{\mathbf{C}} : \mathbf{C} \rightarrow \mathbf{C}$ is a **Poset**-functor for every **Poset**-category \mathbf{C} , and the composition $G \circ F$ of two **Poset**-functors is also a **Poset**-functor.

A **Poset**-adjunction is then the same as an ordinary adjunction, except that the two functors are required to **Poset**-enrich.¹

¹Throughout we use the name ω for counits of adjunctions, reserving ε for effects.

Definition A.0.8 (Poset-adjunction) Suppose that \mathbf{C} and \mathbf{D} are **Poset**-categories. A **Poset-adjunction** $F \dashv U$ consists of two **Poset**-functors: the *left adjoint* $F : \mathbf{C} \rightarrow \mathbf{D}$ and the *right adjoint* $U : \mathbf{D} \rightarrow \mathbf{C}$, together with two natural transformations

$$\eta : Id_{\mathbf{C}} \rightarrow U \circ F \quad \varpi : F \circ U \rightarrow Id_{\mathbf{D}}$$

such that the two triangle laws $U \varpi \circ \eta_U = id_U$ and $\varpi_F \circ F \eta = id_F$ hold. ◀

Example A.0.9 There is a **Poset**-adjunction between posets and pointed posets:

$$\begin{array}{ccc} & \xrightarrow{(-)_{\perp}} & \\ \mathbf{Poset} & \xleftrightarrow{\perp} & \mathbf{Poset}_{\perp} \\ & \xleftarrow{U} & \end{array}$$

The left adjoint $(-)_{\perp}$ sends each poset X to the pointed poset X_{\perp} formed by freely adding a least element, and each monotone function f to its strict extension f_{\perp} (so $f_{\perp} \perp = \perp$ and $f_{\perp} x = f x$ for $x \neq \perp$). The right adjoint is the forgetful functor, which sends each pointed poset to itself, and each strict monotone function to itself. The unit $\eta_X : X \rightarrow X_{\perp}$ is the monotone function that maps $x \in X$ to itself, and the counit $\varpi_X : X_{\perp} \rightarrow X$ is the strict monotone function that merges the least elements. ◀

As in the ordinary case, there is an equivalent definition of **Poset**-adjunction in terms of bijections between hom-posets:

Definition A.0.10 (Poset-adjunction (alternative definition)) Suppose that \mathbf{C} and \mathbf{D} are **Poset**-categories. A **Poset-adjunction** $F \dashv U$ consists of two **Poset**-functors: the *left adjoint* $F : \mathbf{C} \rightarrow \mathbf{D}$ and the *right adjoint* $U : \mathbf{D} \rightarrow \mathbf{C}$, together with a family of bijections

$$\theta : \mathbf{C}(FX, Y) \rightarrow \mathbf{D}(X, UY)$$

natural in $X \in \mathbf{C}$ and $Y \in \mathbf{D}$. ◀

This definition does not explicitly require the functions θ (or their inverses) to be monotone, but this follows from naturality: given $f \sqsubseteq f' : FX \rightarrow Y$ we have

$$\theta f = Uf \circ \theta id_X \sqsubseteq Uf' \circ \theta id_X = \theta f'$$

We also need *tensorial strengths* [50] for denotational semantics. The following definition is exactly the same as the ordinary one, except we additionally require the functor to enrich.

Definition A.0.11 (Strong Poset-functor) Suppose that \mathbf{C} is a cartesian **Poset**-category. A *strong Poset-functor* (F, str^F) on \mathbf{C} consists of a **Poset**-functor $F : \mathbf{C} \rightarrow \mathbf{C}$ and a natural transformation $str_{X,Y}^F : X \times FY \rightarrow F(X \times Y)$ called the (*tensorial*) *strength*, such that the following diagrams commute:

$$\begin{array}{ccc} 1 \times FX & \xrightarrow{str_{1,X}^F} & F(1 \times X) \\ & \searrow \pi_2 & \downarrow F\pi_2 \\ & & FX \\ (X \times Y) \times FZ & \xrightarrow{str_{X \times Y, Z}^F} & F((X \times Y) \times Z) \\ \downarrow assoc & & \downarrow Fassoc \\ X \times (Y \times FZ) & \xrightarrow{X \times str_{Y,Z}^F} & X \times F(Y \times Z) \xrightarrow{str_{X, Y \times Z}^F} F(X \times (Y \times Z)) \end{array}$$

A strong natural transformation $\alpha : (F, str^F) \rightarrow (G, str^G)$ between strong **Poset**-functors is a natural transformation $\alpha : F \rightarrow G$ such that

$$\begin{array}{ccc} X \times FY & \xrightarrow{str^F_{X,Y}} & F(X \times Y) \\ X \times \alpha_Y \downarrow & & \downarrow \alpha_{X \times Y} \\ X \times GY & \xrightarrow{str^G_{X,Y}} & G(X \times Y) \end{array}$$

commutes. ◀

Examples include the identity strong **Poset**-functor (Id_C, id) on C (which we sometimes write as Id_C), and the composition $(G, str^G) \circ (F, str^F) := (G \circ F, G str^F \circ str^G)$ of strong **Poset**-functors. Strong functors are required in models of programming languages when we interpret open terms. For example, models of the monadic metalanguage [78] use strong monads. We can use the definitions above to give a succinct definition of order-enriched strong monads:

Definition A.0.12 (Strong Poset-monad) Suppose that C is a cartesian **Poset**-category. A strong **Poset**-monad on C consists of a strong **Poset**-functor (T, str) and two strong natural transformations

$$\eta : Id_C \rightarrow (T, str) \quad \mu : (T, str) \circ (T, str) \rightarrow (T, str)$$

such that the monad laws hold:

$$\begin{array}{ccc} & \eta_T \curvearrowright & T & \curvearrowleft T\eta \\ & \downarrow & \parallel & \downarrow \\ T \circ T & \xrightarrow{\mu} & T & \xleftarrow{\mu} T \circ T \end{array} \qquad \begin{array}{ccc} T \circ T \circ T & \xrightarrow{\mu_T} & T \circ T \\ T\mu \downarrow & & \downarrow \mu \\ T \circ T & \xrightarrow{\mu} & T \end{array} \quad \blacktriangleleft$$

Appendix B

Additional proofs

B.1 Erasing coercions in GCBPV terms

Recall the following lemma about graded call-by-push-value.

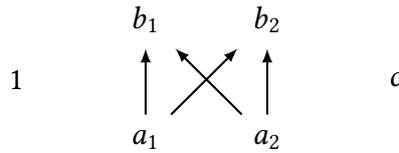
Lemma 2.7.5 Suppose that the effect algebra $(\mathcal{E}, \leq, \cdot, 1)$ is a partially ordered monoid with bounded binary joins.

1. Given two value terms V_1, V_2 such that $\Gamma \vdash V_1 : A$ and $\Gamma \vdash V_2 : A$, if $\lfloor V_1 \rfloor = \lfloor V_2 \rfloor$ then $\Gamma \vdash V_1 \equiv V_2 : A$.
2. Given two computation terms M_1, M_2 such that $\Gamma \vdash M_1 : \underline{C}$ and $\Gamma \vdash M_2 : \underline{C}$, if $\lfloor M_1 \rfloor = \lfloor M_2 \rfloor$ then $\Gamma \vdash M_1 \equiv M_2 : \underline{C}$. \blacktriangleleft

We give the proof of this lemma in this section.

B.1.1 Conjectured counterexample

Before giving the proof, we note that Lemma 2.7.5 does not appear to hold when the condition that the effect algebra has bounded binary joins is dropped. Consider the partial order (\mathcal{E}, \leq) with Hasse diagram



This does not have bounded binary joins: the pair (a_1, a_2) has two upper bounds $(b_1$ and $b_2)$, but neither is less than the other. We define the multiplication so that 1 is the unit and $\varepsilon \cdot \varepsilon' = c$ for $\varepsilon, \varepsilon' \neq 1$. Using this as an effect algebra, we define a typing context Γ and two computation terms M_1 and M_2 such that $\Gamma \vdash M_i : \langle c \rangle \mathbf{unit}$ for $i \in \{1, 2\}$:

$$\begin{aligned}
 \Gamma &:= x : \langle c \rangle (\mathbf{unit} + \mathbf{unit}), y_1 : \langle a_1 \rangle \mathbf{unit}, y_2 : \langle a_1 \rangle \mathbf{unit} \\
 M_i &:= y \text{ to } z. \text{ case } z \text{ of } \{ \mathbf{inl } _ . \mathbf{coerce}_{a_1 \leq b_1} y_1, \mathbf{inr } _ . \mathbf{coerce}_{a_2 \leq b_1} y_2 \}
 \end{aligned}$$

The terms M_i are identical apart from the coercions ($\lfloor M_1 \rfloor = \lfloor M_2 \rfloor$), but we conjecture that $\Gamma \vdash M_1 \not\equiv M_2 : \langle c \rangle \mathbf{unit}$. The situation would be different if a_1 and a_2 had a join $a_1 \vee a_2$: we would be able to reason as follows:

$$\begin{aligned}
 M_i &\equiv y \text{ to } z. \mathbf{coerce}_{a_1 \vee a_2 \leq b_1} \text{ case } z \text{ of } \{ \mathbf{inl } _ . \mathbf{coerce}_{a_1 \leq a_1 \vee a_2} y_1, \mathbf{inr } _ . \mathbf{coerce}_{a_2 \leq a_1 \vee a_2} y_2 \} \\
 &\equiv y \text{ to } z. \text{ case } z \text{ of } \{ \mathbf{inl } _ . \mathbf{coerce}_{a_1 \leq a_1 \vee a_2} y_1, \mathbf{inr } _ . \mathbf{coerce}_{a_2 \leq a_1 \vee a_2} y_2 \} \\
 &\quad \text{(to commutes with coerce)}
 \end{aligned}$$

and hence $M_1 \equiv M_2$. We use similar reasoning in the proof of Lemma 2.7.5.

B.1.2 Proof of Lemma 2.7.5

To prove Lemma 2.7.5, we first define a stronger notion of subtyping for GCBPV in which contravariant subtyping is not allowed. The judgments $A <:' B$ and $\underline{C} <:' \underline{D}$ are given by the rules on the left of Figure 2.12 (which generate the usual notion of subtyping), except that the rule for function types is replaced with

$$\frac{\underline{C} <:' \underline{D}}{(A \rightarrow \underline{C}) <:' (A \rightarrow \underline{D})}$$

Clearly $A <:' B$ implies $A <: B$, and we have coercions $\text{coerce}_{A <: B}$ from type A to type B if $A <:' B$ holds, as in Figure 2.12.

We have partial binary operations \vee on value and computation types. When it is defined $A \vee B$ is the join of the value types A and B with respect to $<:'$, and similarly for computation types. In the definition, $\varepsilon \vee \varepsilon'$ denotes the join of the effects $\varepsilon, \varepsilon'$ if it exists. The cases are not exhaustive; the join is undefined if no case applies. The left-hand side of a case is undefined if any part of the right-hand side is.

$$\begin{aligned} b \vee b &:= b & \text{unit} \vee \text{unit} &:= \text{unit} & (A_1 \times A_2) \vee (B_1 \times B_2) &:= (A_1 \vee B_1) \times (A_2 \vee B_2) \\ \text{empty} \vee \text{empty} &:= \text{empty} & (A_1 + A_2) \vee (B_1 + B_2) &:= (A_1 \vee B_1) + (A_2 \vee B_2) \\ \underline{U} \underline{C} \vee \underline{U} \underline{D} &:= \underline{U} (\underline{C} \vee \underline{D}) \end{aligned}$$

$$\begin{aligned} \underline{\text{unit}} \vee \underline{\text{unit}} &:= \underline{\text{unit}} & (\underline{C}_1 \times \underline{C}_2) \vee (\underline{D}_1 \times \underline{D}_2) &:= (\underline{C}_1 \vee \underline{D}_1) \times (\underline{C}_2 \vee \underline{D}_2) \\ (A \rightarrow \underline{C}) \vee (A \rightarrow \underline{D}) &:= A \rightarrow (\underline{C} \vee \underline{D}) & \langle \varepsilon \rangle A \vee \langle \varepsilon' \rangle B &:= \langle \varepsilon \vee \varepsilon' \rangle (A \vee B) \end{aligned}$$

Not allowing contravariant subtyping means we do not have to consider meets. Bounded binary joins lift from the effect algebra to types:

Lemma B.1.1 Suppose that the effect algebra is a partially ordered monoid with bounded binary joins.

1. If there is some B such that $A <:' B$ and $A' <:' B$ then $A \vee A'$ is defined and is the join of A and A' .
2. If there is some \underline{D} such that $\underline{C} <:' \underline{D}$ and $\underline{C}' <:' \underline{D}$ then $\underline{C} \vee \underline{C}'$ is defined and is the join of \underline{C} and \underline{C}' .

Proof. The proof is by induction on the structure of B and \underline{D} . If $B = b$ then $A = A' = b$ and $b \vee b = b$ is clearly the join. Similarly for $B = \text{unit}$, for $B = \text{empty}$ and for $\underline{D} = \underline{\text{unit}}$. If $B = B_1 \times B_2$ then $A = A_1 \times A_2$ and $A' = A'_1 \times A'_2$ for some $A_1, A'_1 <: B_1$ and $A_2, A'_2 <: B_2$. Then $A_1 \vee A'_1$ and $A_2 \vee A'_2$ are defined, and are joins. Hence $A \times A' = (A_1 \times A_2) \vee (A'_1 \times A'_2) := (A_1 \vee A'_1) \times (A_2 \vee A'_2)$, which is an upper bound. It is the least because if B' is an upper bound then $B' = B'_1 \times B'_2$ for some B'_1, B'_2 , such that B'_i is an upper bound of A_i and A'_i . Then because they are joins, $A_i \vee A'_i <: B'_i$. Similar reasoning can be applied for $B = B_1 + B_2$, for $B = \underline{U} B'$, for $\underline{D} = \underline{D}_1 \times \underline{D}_2$ and for $\underline{D} = A \rightarrow \underline{D}'$. The latter uses the fact that subtyping is not allowed on A . Finally, if $\underline{D} = \langle \varepsilon'' \rangle B$ then $\underline{C} = \langle \varepsilon \rangle A$ and $\underline{C}' = \langle \varepsilon' \rangle A'$ for some $\varepsilon, \varepsilon' \leq \varepsilon''$ and $A, A' <: B$. Then $\varepsilon \vee \varepsilon'$ is defined because we assume bounded binary joins, and so is $A \vee A'$. Hence $\langle \varepsilon \rangle A \vee \langle \varepsilon' \rangle A' = \langle \varepsilon \vee \varepsilon' \rangle (A \vee A')$ is defined, and is the join. \square

We use this notion of subtyping to give typing judgments $\Gamma \vdash_e V : A$ and $\Gamma \vdash_e M : \underline{C}$ for terms that do not have explicit coercions. These are generated by the same rules as the usual typing judgments (Figure 2.11), except that there is no rule for **coerce**, and the rules for case expressions on sum types and for function application are replaced with:

$$\frac{\Gamma \vdash_e V : A_1 + A_2 \quad \Gamma, x_1 : A_1 \vdash_e W_1 : B_1 \quad \Gamma, x_2 : A_2 \vdash_e W_2 : B_2 \quad B_1 \vee B_2 \text{ defined}}{\Gamma \vdash_e \text{case } V \text{ of } \{\text{inl } x_1. W_1, \text{inr } x_2. W_2\} : B_1 \vee B_2}$$

$$\frac{\Gamma \vdash_e V : A \quad \Gamma \vdash_e M : B \rightarrow \underline{C} \quad A <:' B}{\Gamma \vdash_e V' M : \underline{C}}$$

Note that these are syntax directed. Given Γ and V , if there is some A such that $\Gamma \vdash_e V : A$ then it is unique, and derivations are unique. Similarly for computations. For each term that is well typed under these judgments, we can add explicit coercions to get terms that are well typed under the original judgments. Given a typing context Γ and value term V such that $\Gamma \vdash_e V : A$ for some A , we have a value term $\llbracket V \rrbracket_\Gamma$ such that $\Gamma \vdash \llbracket V \rrbracket_\Gamma : A$. Similarly, given Γ and M such that $\Gamma \vdash_e M : \underline{C}$ for some \underline{C} we have $\llbracket M \rrbracket_\Gamma$ such that $\Gamma \vdash \llbracket M \rrbracket_\Gamma : \underline{C}$. These are defined by induction on the structure of M and V . Most of the cases just apply $\llbracket - \rrbracket$ to subterms. The interesting cases correspond to the new typing rules:

$$\left[\begin{array}{l} \text{case } V \text{ of} \\ \{ \text{inl } x_1. W_1 \\ , \text{inr } x_2. W_2 \} \end{array} \right]_\Gamma := \begin{array}{l} \text{case } \llbracket V \rrbracket_\Gamma \text{ of} \\ \{ \text{inl } x_1. \text{coerce}_{B_1 <: B_1 \vee B_2} \llbracket W_1 \rrbracket_{\Gamma, x_1 : A_1} \\ , \text{inr } x_2. \text{coerce}_{B_2 <: B_1 \vee B_2} \llbracket W_2 \rrbracket_{\Gamma, x_2 : A_2} \} \end{array} \quad \text{if } \begin{cases} \Gamma \vdash_e V : A_1 + A_2 \\ \Gamma, x_1 : A_1 \vdash_e W_1 : B_1 \\ \Gamma, x_2 : A_2 \vdash_e W_2 : B_2 \end{cases}$$

$$\llbracket V' M \rrbracket_\Gamma := (\text{coerce}_{A <: B} \llbracket V \rrbracket_\Gamma)' \llbracket M \rrbracket_\Gamma \quad \text{if } \begin{cases} \Gamma \vdash_e V : A \\ \Gamma \vdash_e M : B \rightarrow \underline{C} \end{cases}$$

These terms clearly have the correct types.

We also extend subtyping to typing contexts. Given two typing contexts $\Gamma' = x_1 : A'_1, \dots, x_n : A'_n$ and $\Gamma = x_1 : A_1, \dots, x_n : A_n$ of the same length and with the same variable names, we write $\Gamma' <:' \Gamma$ if $A'_i <:' A_i$ for each i , and also write $\text{coerce}_{\Gamma' <: \Gamma}$ for the substitution

$$x_1 \mapsto \text{coerce}_{A'_1 <: A_1} x_1, \dots, x_n \mapsto \text{coerce}_{A'_n <: A_n} x_n$$

If $\Gamma \vdash V : A$ then $\Gamma' \vdash V[\text{coerce}_{\Gamma' <: \Gamma}] : A$, and if $\Gamma \vdash_e M : \underline{C}$ then $\Gamma' \vdash_e M[\text{coerce}_{\Gamma' <: \Gamma}] : \underline{C}$.

Lemma B.1.2 Suppose that the effect algebra is a partially ordered monoid with bounded binary joins.

1. If $\Gamma \vdash V : A$ and $\Gamma' <:' \Gamma$ then there is some $A' <:' A$ such that $\Gamma' \vdash_e \llbracket V \rrbracket : A'$ and

$$\Gamma' \vdash \text{coerce}_{A' <: A} \llbracket \llbracket V \rrbracket \rrbracket_{\Gamma'} \equiv \llbracket V[\text{coerce}_{\Gamma' <: \Gamma}] \rrbracket : A$$

2. If $\Gamma \vdash_e M : \underline{C}$ and $\Gamma' <:' \Gamma$ then there is some $\underline{C}' <:' \underline{C}$ such that $\Gamma' \vdash_e \llbracket M \rrbracket : \underline{C}'$ and

$$\Gamma' \vdash \text{coerce}_{\underline{C}' <: \underline{C}} \llbracket \llbracket M \rrbracket \rrbracket_{\Gamma'} \equiv \llbracket M[\text{coerce}_{\Gamma' <: \Gamma}] \rrbracket : \underline{C}$$

Proof. By induction the derivations of $\Gamma \vdash V : A$ and $\Gamma \vdash_e M : \underline{C}$.

- Constants c , and the unit terms $()$ and $\lambda\{\}$ are trivial.
- For variables x , we have $(x : A) \in \Gamma$, and $(x : A') \in \Gamma'$ for some $A' <:' A$. Then we have $\llbracket x \rrbracket = x$ and $\text{coerce}_{A' <: A} \llbracket \llbracket x \rrbracket \rrbracket_{\Gamma'} = \text{coerce}_{A' <: A} x = x[\text{coerce}_{\Gamma' <: \Gamma}]$.

- For pairs $\Gamma \vdash (V_1, V_2) : A_1 \times A_2$ by the inductive hypothesis we have

$$\Gamma' \vdash_e \llbracket V_i \rrbracket : A'_i \quad \Gamma' \vdash \mathbf{coerce}_{A'_i <: A_i} \llbracket \llbracket V_i \rrbracket \rrbracket_{\Gamma'} \equiv V_i[\mathbf{coerce}_{\Gamma' <: \Gamma}] : A_i$$

for some $A'_i <: A_i$. Hence $A'_1 \times A'_2 <: A_1 \times A_2$, $\Gamma' \vdash_e \llbracket (V_1, V_2) \rrbracket : A'_1 \times A'_2$, and

$$\begin{aligned} & \mathbf{coerce}_{A'_1 \times A'_2 <: A_1 \times A_2} \llbracket \llbracket (V_1, V_2) \rrbracket \rrbracket_{\Gamma'} \\ \equiv & (\mathbf{coerce}_{A'_1 <: A_1} \llbracket \llbracket V_1 \rrbracket \rrbracket_{\Gamma'}, \mathbf{coerce}_{A'_2 <: A_2} \llbracket \llbracket V_2 \rrbracket \rrbracket_{\Gamma'}) && (\beta \text{ laws for products}) \\ \equiv & (V_1[\mathbf{coerce}_{\Gamma' <: \Gamma}], V_2[\mathbf{coerce}_{\Gamma' <: \Gamma}]) \\ = & (V_1, V_2)[\mathbf{coerce}_{\Gamma' <: \Gamma}] \end{aligned}$$

The cases for **inl**, **inr**, **thunk**, pairs of computation terms, and $\langle - \rangle$ are similar.

- For $\Gamma \vdash \mathbf{fst} V : A_1$ we have $\Gamma \vdash V : A_1 \times A_2$. By the inductive hypothesis there is some subtype of $A_1 \times A_2$, which necessarily has the form $A'_1 \times A'_2$ for $A'_i <: A_i$, such that

$$\Gamma' \vdash_e \llbracket V \rrbracket : A'_1 \times A'_2 \quad \Gamma' \vdash \mathbf{coerce}_{A'_1 \times A'_2 <: A_1 \times A_2} \llbracket \llbracket V \rrbracket \rrbracket_{\Gamma'} \equiv V[\mathbf{coerce}_{\Gamma' <: \Gamma}] : A_1 \times A_2$$

We have $\Gamma' \vdash_e \llbracket \mathbf{fst} V \rrbracket : A'_1$ and

$$\begin{aligned} \mathbf{coerce}_{A'_1 <: A_1} \llbracket \llbracket \mathbf{fst} V \rrbracket \rrbracket_{\Gamma'} & \equiv \mathbf{fst} (\mathbf{coerce}_{A'_1 \times A'_2 <: A_1 \times A_2} \llbracket \llbracket V \rrbracket \rrbracket_{\Gamma'}) && (\beta \text{ for products}) \\ & \equiv \mathbf{fst} (V[\mathbf{coerce}_{\Gamma' <: \Gamma}]) \\ & = (\mathbf{fst} V)[\mathbf{coerce}_{\Gamma' <: \Gamma}] \end{aligned}$$

The cases for **snd**, projections for products of computations and **force** are similar.

- For $\Gamma \vdash \mathbf{case}_A V \text{ of } \{\} : A$ we have $\Gamma \vdash V : \mathbf{empty}$. By the inductive hypothesis there is some subtype of **empty**, which can only be **empty** itself, such that $\Gamma' \vdash_e \llbracket V \rrbracket : \mathbf{empty}$. We have $A <: A$, $\Gamma' \vdash_e \llbracket \mathbf{case}_A V \text{ of } \{\} \rrbracket : A$, and

$$\begin{aligned} \mathbf{coerce}_{A <: A} \llbracket \llbracket \mathbf{case}_A V \text{ of } \{\} \rrbracket \rrbracket_{\Gamma'} & \equiv \mathbf{case}_A V[\mathbf{coerce}_{\Gamma' <: \Gamma}] \text{ of } \{\} && (\eta \text{ law for empty}) \\ & = (\mathbf{case}_A V \text{ of } \{\})[\mathbf{coerce}_{\Gamma' <: \Gamma}] \end{aligned}$$

- For $\Gamma \vdash \lambda x : A. M : A \rightarrow \underline{C}$ we have $\Gamma, x : A \vdash M : \underline{C}$. Applying the inductive hypothesis to $\Gamma', x : A <: \Gamma, x : A$, there is some $\underline{C}' <: \underline{C}$ such that

$$\Gamma', x : A \vdash_e \llbracket M \rrbracket : \underline{C}' \quad \Gamma', x : A \vdash \mathbf{coerce}_{\underline{C}' <: \underline{C}} \llbracket \llbracket M \rrbracket \rrbracket_{\Gamma', x : A} \equiv M[\mathbf{coerce}_{\Gamma', x : A <: \Gamma, x : A}] : \underline{C}$$

We have $(A \rightarrow \underline{C}') <: (A \rightarrow \underline{C})$, $\Gamma' \vdash_e \llbracket \lambda x : A. M \rrbracket : A \rightarrow \underline{C}'$, and

$$\begin{aligned} & \mathbf{coerce}_{(A \rightarrow \underline{C}') <: (A \rightarrow \underline{C})} \llbracket \llbracket \lambda x : A. M \rrbracket \rrbracket_{\Gamma'} \\ \equiv & \lambda y : A. \mathbf{coerce}_{\underline{C}' <: \underline{C}} (y \text{ ' } \lambda x : A. \llbracket \llbracket M \rrbracket \rrbracket_{\Gamma', x : A}) && (\mathbf{coerce}_{A <: A} W \equiv W) \\ \equiv & \lambda x : A. \mathbf{coerce}_{\underline{C}' <: \underline{C}} \llbracket \llbracket M \rrbracket \rrbracket_{\Gamma', x : A} && (\beta \text{ law for functions}) \\ \equiv & \lambda x : A. M[\mathbf{coerce}_{\Gamma', x : A <: \Gamma, x : A}] \\ \equiv & (\lambda x : A. M)[\mathbf{coerce}_{\Gamma' <: \Gamma}] && (\mathbf{coerce}_{A <: A} W \equiv W) \end{aligned}$$

- For $\Gamma \vdash \text{op } V : \langle \text{eff}_{\text{op}} \rangle \text{ar}_{\text{op}}$ we have $\Gamma \vdash V : \text{car}_{\text{op}}$. Since car_{op} is a ground type, its only subtype is car_{op} itself. Hence by the inductive hypothesis we have

$$\Gamma' \vdash_e \llbracket V \rrbracket : \text{car}_{\text{op}} \quad \Gamma' \vdash \mathbf{coerce}_{\text{car}_{\text{op}} <: \text{car}_{\text{op}}} \llbracket \llbracket V \rrbracket \rrbracket_{\Gamma'} \equiv V[\mathbf{coerce}_{\Gamma' <: \Gamma}] : \text{car}_{\text{op}}$$

We have $\langle \text{eff}_{\text{op}} \rangle_{\text{ar}_{\text{op}}} <:' \langle \text{eff}_{\text{op}} \rangle_{\text{ar}_{\text{op}}}$, $\Gamma' \vdash_e \text{op } V : \langle \text{eff}_{\text{op}} \rangle_{\text{ar}_{\text{op}}}$, and

$$\begin{aligned} & \text{coerce}_{\langle \text{eff}_{\text{op}} \rangle_{\text{ar}_{\text{op}}} <: \langle \text{eff}_{\text{op}} \rangle_{\text{ar}_{\text{op}}}} \llbracket \text{op } V \rrbracket_{\Gamma'} \\ & \equiv \text{op} \llbracket V \rrbracket_{\Gamma'} && (\text{coerce}_{\underline{C} <: \underline{C}} N \equiv N) \\ & \equiv \text{op} (\text{coerce}_{\text{car}_{\text{op}} <: \text{car}_{\text{op}}} \llbracket V \rrbracket_{\Gamma'}) && (\text{coerce}_{A <: A} W \equiv W) \\ & \equiv \text{op} (V[\text{coerce}_{\Gamma' <: \Gamma}]) \\ & = (\text{op } V)[\text{coerce}_{\Gamma' <: \Gamma}] \end{aligned}$$

- For $\Gamma \vdash M \text{ to } x. N : \langle \varepsilon \rangle \underline{C}$ we have $\Gamma \vdash M : \langle \varepsilon \rangle A$ and $\Gamma, x : A \vdash N : \underline{C}$. There is some subtype of $\langle \varepsilon \rangle A$, which necessarily has the form $\langle \varepsilon' \rangle A'$ for $\varepsilon' \leq \varepsilon$ and $A' <:' A$ such that

$$\Gamma \vdash_e \llbracket M \rrbracket : \langle \varepsilon' \rangle A' \quad \Gamma' \vdash \text{coerce}_{\langle \varepsilon' \rangle A' <: \langle \varepsilon \rangle A} \llbracket \llbracket M \rrbracket \rrbracket_{\Gamma'} \equiv M[\text{coerce}_{\Gamma' <: \Gamma} : \langle \varepsilon \rangle A]$$

Now $\Gamma', x : A' <:' \Gamma, x : A$, so there is some $\underline{C}' <: \underline{C}$ such that

$$\Gamma', x : A' \vdash_e \llbracket N \rrbracket : \underline{C}' \quad \Gamma', x : A' \vdash \text{coerce}_{\underline{C}' <: \underline{C}} \llbracket \llbracket N \rrbracket \rrbracket_{\Gamma', x : A'} \equiv N[\text{coerce}_{\Gamma', x : A' <: \Gamma, x : A} : \underline{C}]$$

We have $\langle \varepsilon' \rangle \underline{C}' <:' \langle \varepsilon \rangle \underline{C}$, $\Gamma' \vdash_e \llbracket M \text{ to } x. N \rrbracket : \langle \varepsilon' \rangle \underline{C}'$, and

$$\begin{aligned} & \text{coerce}_{\langle \varepsilon' \rangle \underline{C}' <: \langle \varepsilon \rangle \underline{C}} \llbracket \llbracket M \text{ to } x. N \rrbracket \rrbracket_{\Gamma'} \\ & \equiv (\text{coerce}_{\varepsilon' \leq \varepsilon} \llbracket \llbracket M \rrbracket \rrbracket_{\Gamma'}) \text{ to } x. \text{coerce}_{\underline{C}' <: \underline{C}} \llbracket \llbracket N \rrbracket \rrbracket_{\Gamma', x : A'} \\ & \equiv (\text{coerce}_{\varepsilon' \leq \varepsilon} \llbracket \llbracket M \rrbracket \rrbracket_{\Gamma'}) \text{ to } x. N[\text{coerce}_{\Gamma', x : A' <: \Gamma, x : A}] \\ & \equiv (\text{coerce}_{\langle \varepsilon' \rangle A' <: \langle \varepsilon \rangle A} \llbracket \llbracket M \rrbracket \rrbracket_{\Gamma'}) \text{ to } x. N[\text{coerce}_{\Gamma' <: \Gamma}] \\ & \equiv M[\text{coerce}_{\Gamma' <: \Gamma}] \text{ to } x. N[\text{coerce}_{\Gamma' <: \Gamma}] \\ & = (M \text{ to } x. N)[\text{coerce}_{\Gamma' <: \Gamma}] \end{aligned}$$

- For $\Gamma \vdash \text{coerce}_{\varepsilon' \leq \varepsilon} M : \langle \varepsilon \rangle A$ we have $\Gamma \vdash M : \langle \varepsilon' \rangle A$ and $\varepsilon' \leq \varepsilon$. By the inductive hypothesis there is some subtype of $\langle \varepsilon' \rangle A$, which necessarily has the form $\langle \varepsilon'' \rangle A'$ for some $\varepsilon'' \leq \varepsilon'$ and $A' <:' A$, such that

$$\Gamma' \vdash_e \llbracket M \rrbracket : \langle \varepsilon'' \rangle A' \quad \Gamma' \vdash \text{coerce}_{\langle \varepsilon'' \rangle A' <: \langle \varepsilon' \rangle A} \llbracket \llbracket M \rrbracket \rrbracket_{\Gamma'} \equiv M[\text{coerce}_{\Gamma' <: \Gamma} : \langle \varepsilon' \rangle A]$$

We have $\langle \varepsilon'' \rangle A' <:' \langle \varepsilon \rangle A$, and $\Gamma' \vdash_e \llbracket \text{coerce}_{\varepsilon' \leq \varepsilon} M \rrbracket : \langle \varepsilon'' \rangle A'$ because $\llbracket \text{coerce}_{\varepsilon' \leq \varepsilon} M \rrbracket = \llbracket M \rrbracket$. The required equation also holds:

$$\begin{aligned} & \text{coerce}_{\langle \varepsilon'' \rangle A' <: \langle \varepsilon \rangle A} \llbracket \llbracket \text{coerce}_{\varepsilon' \leq \varepsilon} M \rrbracket \rrbracket_{\Gamma'} \\ & \equiv \text{coerce}_{\langle \varepsilon'' \rangle A' <: \langle \varepsilon \rangle A} (\text{coerce}_{\langle \varepsilon'' \rangle A' <: \langle \varepsilon' \rangle A} \llbracket \llbracket M \rrbracket \rrbracket_{\Gamma'}) \\ & \quad (\text{coerce}_{\underline{D}'' <: \underline{D}} N \equiv \text{coerce}_{\underline{D}' <: \underline{D}} (\text{coerce}_{\underline{D}'' <: \underline{D}'} N)) \\ & \equiv \text{coerce}_{\langle \varepsilon'' \rangle A' <: \langle \varepsilon \rangle A} (M[\text{coerce}_{\Gamma' <: \Gamma}]) \\ & \equiv (\text{coerce}_{\varepsilon' \leq \varepsilon} M)[\text{coerce}_{\Gamma' <: \Gamma}] \end{aligned}$$

- For $\Gamma \vdash \text{case } V \text{ of } \{\text{inl } x_1. W_1, \text{inr } x_2. W_2\} : B$ we have $\Gamma \vdash V : A_1 + A_2$. By the inductive hypothesis there is a subtype of $A_1 + A_2$, which necessarily has the form $A'_1 + A'_2$ for $A'_i <:' A_i$, such that

$$\Gamma' \vdash_e \llbracket V \rrbracket : A'_1 + A'_2 \quad \Gamma' \vdash \text{coerce}_{A'_1 + A'_2 <: A_1 + A_2} \llbracket \llbracket V \rrbracket \rrbracket_{\Gamma'} \equiv V[\text{coerce}_{\Gamma' <: \Gamma} : A_1 + A_2]$$

For $i \in \{1, 2\}$ we have $\Gamma, x_i : A_i \vdash W_i : B$ and $\Gamma', x_i : A'_i <:' \Gamma, x_i : A_i$, so there exists $B'_i <:' B$ such that

$$\begin{aligned} & \Gamma', x_i : A'_i \vdash_e \llbracket W_i \rrbracket : B'_i \\ & \Gamma', x_i : A'_i \vdash \text{coerce}_{B'_i <: B} \llbracket \llbracket W_i \rrbracket \rrbracket_{\Gamma', x_i : A'_i} \equiv W_i[\text{coerce}_{\Gamma', x_i : A'_i <: \Gamma, x_i : A_i} : B] \end{aligned}$$

Since B'_1 and B'_2 have an upper bound, by Lemma B.1.1 (1) their join $B'_1 \vee B'_2 <: B$ is defined. We therefore have $\Gamma' \vdash_e \llbracket \text{case } V \text{ of } \{\text{inl } x_1. W_1, \text{inr } x_2. W_2\} \rrbracket_{\Gamma'} : B'_1 \vee B'_2$ and

$$\begin{aligned}
& \text{coerce}_{B'_1 \vee B'_2 <: B} \llbracket \llbracket \text{case } V \text{ of } \{\text{inl } x_1. W_1, \text{inr } x_2. W_2\} \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} \\
& \text{case } \llbracket \llbracket V \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} \text{ of} \\
\equiv & \{ \text{inl } x_1. \text{coerce}_{B'_1 \vee B'_2 <: B} (\text{coerce}_{B'_1 <: B'_1 \vee B'_2} \llbracket \llbracket W_1 \rrbracket_{\Gamma', x_1: A'_1} \rrbracket_{\Gamma'} \rrbracket_{\Gamma'}) \quad (\beta \text{ and } \eta \text{ laws for sums}) \\
& \quad , \text{inr } x_2. \text{coerce}_{B'_2 \vee B'_2 <: B} (\text{coerce}_{B'_2 <: B'_2 \vee B'_2} \llbracket \llbracket W_2 \rrbracket_{\Gamma', x_2: A'_2} \rrbracket_{\Gamma'} \rrbracket_{\Gamma'}) \} \\
\equiv & \text{case } \llbracket \llbracket V \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} \text{ of } \{ \text{inl } x_1. \text{coerce}_{B'_1 <: B} \llbracket \llbracket W_1 \rrbracket_{\Gamma', x_1: A'_1} \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} , \text{inr } x_2. \text{coerce}_{B'_2 <: B} \llbracket \llbracket W_2 \rrbracket_{\Gamma', x_2: A'_2} \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} \} \\
& \quad (\text{coerce}_{A'' <: A} V \equiv \text{coerce}_{A' <: A} (\text{coerce}_{A'' <: A'} V)) \\
\equiv & \text{case } \llbracket \llbracket V \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} \text{ of } \{ \text{inl } x_1. W_1 [\text{coerce}_{\Gamma', x_1: A'_1 <: \Gamma, x_1: A_1}] , \text{inr } x_2. W_2 [\text{coerce}_{\Gamma', x_2: A'_2 <: \Gamma, x_2: A_2}] \} \\
\equiv & \text{case } \text{coerce}_{A'_1 + A'_2 <: A_1 + A_2} \llbracket \llbracket V \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} \text{ of } \{ \text{inl } x_1. W_1 [\text{coerce}_{\Gamma' <: \Gamma}] , \text{inr } x_2. W_2 [\text{coerce}_{\Gamma' <: \Gamma}] \} \\
& \quad (\beta \text{ and } \eta \text{ laws for sums}) \\
\equiv & (\text{case } V \text{ of } \{ \text{inl } x_1. W_1, \text{inr } x_2. W_2 \}) [\text{coerce}_{\Gamma' <: \Gamma}]
\end{aligned}$$

- For $\Gamma \vdash V' M : \underline{C}$ we have $\Gamma \vdash V : A$ and $\Gamma \vdash M : A \rightarrow \underline{C}$. By the inductive hypothesis, there exist $A' <: A$ and a subtype of $A \rightarrow \underline{C}$ that necessarily has the form $A \rightarrow \underline{C}'$ for $\underline{C}' <: \underline{C}$ (recall that we do not allow subtyping on arguments), such that

$$\begin{aligned}
\Gamma' \vdash_e \llbracket V \rrbracket : A' & & \Gamma' \vdash \text{coerce}_{A' <: A} \llbracket \llbracket V \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} & \equiv V [\text{coerce}_{\Gamma' <: \Gamma}] : A \\
\Gamma' \vdash_e \llbracket M \rrbracket : A \rightarrow \underline{C}' & & \Gamma' \vdash \text{coerce}_{A \rightarrow \underline{C}' <: A \rightarrow \underline{C}} \llbracket \llbracket M \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} & \equiv M [\text{coerce}_{\Gamma' <: \Gamma}] : A \rightarrow \underline{C}
\end{aligned}$$

We therefore have $\Gamma \vdash_e \llbracket V' M \rrbracket : \underline{C}'$ and

$$\begin{aligned}
\text{coerce}_{\underline{C}' <: \underline{C}} \llbracket \llbracket V' M \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} & \equiv \text{coerce}_{\underline{C}' <: \underline{C}} ((V [\text{coerce}_{\Gamma' <: \Gamma}]) ' \llbracket \llbracket M \rrbracket_{\Gamma'} \rrbracket_{\Gamma'}) \\
& \equiv (V [\text{coerce}_{\Gamma' <: \Gamma}]) ' \text{coerce}_{A \rightarrow \underline{C}' <: A \rightarrow \underline{C}} \llbracket \llbracket M \rrbracket_{\Gamma'} \rrbracket_{\Gamma'} \\
& \quad (\beta \text{ law for function types}) \\
& \equiv (V' M) [\text{coerce}_{\Gamma' <: \Gamma}] \quad \square
\end{aligned}$$

We obtain Lemma 2.7.5 as a corollary.

Proof (of Lemma 2.7.5). We give the proof for the first part of the lemma (values). The proof of the second part (computations) is similar.

Suppose that $\Gamma \vdash V_1 : A$ and $\Gamma \vdash V_2 : A$. We have $\Gamma <: \Gamma$, so by Lemma B.1.2 (1) there are value types $A'_i <: A$ such that $\Gamma \vdash_e \llbracket V_i \rrbracket : A'_i$ and

$$\Gamma \vdash \text{coerce}_{A'_i <: A} \llbracket \llbracket V_i \rrbracket_{\Gamma} \rrbracket_{\Gamma} \equiv V_i [\text{coerce}_{\Gamma <: \Gamma}] : A$$

for $i \in \{1, 2\}$. If $\llbracket V_1 \rrbracket = \llbracket V_2 \rrbracket$ then $A'_1 = A'_2$, and we have:

$$V_1 \equiv V_1 [\text{coerce}_{\Gamma <: \Gamma}] \equiv \text{coerce}_{A'_1 <: A} \llbracket \llbracket V_1 \rrbracket_{\Gamma} \rrbracket_{\Gamma} = \text{coerce}_{A'_2 <: A} \llbracket \llbracket V_2 \rrbracket_{\Gamma} \rrbracket_{\Gamma} \equiv V_2 [\text{coerce}_{\Gamma <: \Gamma}] \equiv V_2$$

where we use the fact that the inequational theory is closed under substitution. \square

B.2 Logical relations and the free lifting

This section contains proofs related to the free lifting defined in Section 3.1.1.

Lemma 3.1.6 If $(c, c) \in \mathcal{R}_{\text{free}} \llbracket A \rrbracket$ for each constant $c \in \mathcal{K}_A$, then $\mathcal{R}_{\text{free}} \llbracket - \rrbracket$ is a logical relation.

Proof. We show that the free lifting meets each requirement in the definition of logical relation.

- The equations in Figure 3.1 hold by definition.
- For closure under operations we note that

$$\begin{aligned} \text{op } V &\equiv \text{op } V \text{ to } x. \langle x \rangle & (\langle x \rangle, \langle x \rangle) &: \mathcal{R}_{\text{free}} \llbracket x : \text{ar}_{\text{op}} \rrbracket \dot{\rightarrow} \mathcal{R}_{\text{free}} \llbracket \langle 1 \rangle \text{ar}_{\text{op}} \rrbracket \\ \text{op } V' &\equiv \text{op } V' \text{ to } x. \langle x \rangle \end{aligned}$$

and apply the final case in the definition of $\mathcal{R}_{\text{free}} \llbracket \langle 1 \rangle \text{ar}_{\text{op}} \rrbracket$.

- Closure under pure computations is one of the cases in the definition of $\mathcal{R}_{\text{free}} \llbracket \langle \varepsilon \rangle A \rrbracket$.
- For closure under **to** suppose that $(M, M') \in \mathcal{R}_{\text{free}} \llbracket \langle \varepsilon \rangle A \rrbracket$ and $(N, N') : \mathcal{R}_{\text{free}} \llbracket x : A \rrbracket \dot{\rightarrow} \mathcal{R}_{\text{free}} \llbracket \langle \varepsilon' \rangle B \rrbracket$. We show that

$$(M \text{ to } x. N, M' \text{ to } x. N') \in \mathcal{R} \llbracket \langle \varepsilon \cdot \varepsilon' \rangle B \rrbracket$$

by induction on $(M, M') \in \mathcal{R}_{\text{free}} \llbracket \langle \varepsilon \rangle A \rrbracket$. For the $(\langle V \rangle, \langle V' \rangle)$ case use the assumption about (N, N') and

$$\langle V \rangle \text{ to } x. N \equiv N[x \mapsto V] \quad \langle V' \rangle \text{ to } x. N' \equiv N'[x \mapsto V']$$

For $(\text{coerce}_{\varepsilon'' \leq \varepsilon} M, \text{coerce}_{\varepsilon'' \leq \varepsilon} M')$ use the inductive hypothesis and

$$\begin{aligned} (\text{coerce}_{\varepsilon'' \leq \varepsilon} M) \text{ to } x. N &\equiv \text{coerce}_{\varepsilon'' \cdot \varepsilon' \leq \varepsilon \cdot \varepsilon'} (M \text{ to } x. N) \\ (\text{coerce}_{\varepsilon'' \leq \varepsilon} M') \text{ to } x. N' &\equiv \text{coerce}_{\varepsilon'' \cdot \varepsilon' \leq \varepsilon \cdot \varepsilon'} (M' \text{ to } x. N') \end{aligned}$$

For $(\text{op } V \text{ to } y. M, \text{op } V' \text{ to } y. M')$, note that the inductive hypothesis implies

$$(M \text{ to } x. N, M' \text{ to } x. N') : \mathcal{R}_{\text{free}} \llbracket y : \text{ar}_{\text{op}} \rrbracket \dot{\rightarrow} \mathcal{R}_{\text{free}} \llbracket \langle \varepsilon \cdot \varepsilon' \rangle B \rrbracket$$

and then use

$$\begin{aligned} (\text{op } V \text{ to } y. M) \text{ to } x. N &\equiv \text{op } V \text{ to } y. (M \text{ to } x. N) \\ (\text{op } V' \text{ to } y. M') \text{ to } x. N' &\equiv \text{op } V' \text{ to } y. (M' \text{ to } x. N') \end{aligned}$$

- Closure under **coerce** is one of the cases in the definition of $\mathcal{R}_{\text{free}} \llbracket \langle \varepsilon \rangle A \rrbracket$.
- Each constant is related to itself by assumption. □

Lemma 3.1.7 Suppose that $\mathcal{R} \llbracket - \rrbracket$ is a logical relation and $\mathcal{R}_{\text{free}} \llbracket b \rrbracket = \mathcal{R} \llbracket b \rrbracket$ for all base types b . Then $\mathcal{R}_{\text{free}} \llbracket A \rrbracket \subseteq \mathcal{R} \llbracket A \rrbracket$ implies $\mathcal{R}_{\text{free}} \llbracket \langle \varepsilon \rangle A \rrbracket \subseteq \mathcal{R} \llbracket \langle \varepsilon \rangle A \rrbracket$.

Proof. By induction on the definition of $\mathcal{R}_{\text{free}} \llbracket \langle \varepsilon \rangle A \rrbracket$. Each case uses the closure properties of the logical relation $\mathcal{R} \llbracket - \rrbracket$. For the $(\langle V \rangle, \langle V' \rangle)$ case use $\mathcal{R}_{\text{free}} \llbracket A \rrbracket \subseteq \mathcal{R} \llbracket A \rrbracket$ and closure under pure computations. For $(\text{coerce}_{\varepsilon \leq \varepsilon'} M, \text{coerce}_{\varepsilon \leq \varepsilon'} M')$ use the inductive hypothesis and closure under **coerce**. The $(\text{op } V \text{ to } x. M, \text{op } V' \text{ to } x. M')$ case is the most difficult. Suppose that $(V, V') \in \mathcal{R}_{\text{free}} \llbracket \text{car}_{\text{op}} \rrbracket$ and $(M, M') : \mathcal{R}_{\text{free}} \llbracket x : \text{ar}_{\text{op}} \rrbracket \dot{\rightarrow} \mathcal{R}_{\text{free}} \llbracket \langle \varepsilon \rangle A \rrbracket$. By the assumption on base types, we have $\mathcal{R}_{\text{free}} \llbracket G \rrbracket = \mathcal{R} \llbracket G \rrbracket$ for all ground types G , in particular for the coarity and arity of op . Hence the inductive hypothesis implies $(V, V') \in \mathcal{R} \llbracket \text{car}_{\text{op}} \rrbracket$ and $(M, M') : \mathcal{R} \llbracket x : \text{ar}_{\text{op}} \rrbracket \dot{\rightarrow} \mathcal{R} \llbracket \langle \varepsilon \rangle A \rrbracket$, and the result follows by closure of $\mathcal{R} \llbracket - \rrbracket$ under operations and under **to**. □

Lemma 3.1.8 For each value type A and effect $\varepsilon \subseteq \Sigma$,

$$\mathcal{R}'_{\text{free}}[\langle \varepsilon \rangle A] = \mathcal{R}_{\text{free}}[\langle \varepsilon \rangle A]$$

Proof. (\subseteq): For a fixed effect ε , we have $\mathcal{R}'_{\text{free}}[\langle \varepsilon \rangle A] \subseteq \mathcal{R}_{\text{free}}[\langle \varepsilon \rangle A]$ by a trivial induction.

(\supseteq): First, a trivial induction shows that if $\varepsilon \subseteq \varepsilon'$ then

$$(M, M') \in \mathcal{R}'_{\text{free}}[\langle \varepsilon \rangle A] \quad \Rightarrow \quad (\text{coerce}_{\varepsilon \leq \varepsilon'} M, \text{coerce}_{\varepsilon \leq \varepsilon'} M') \in \mathcal{R}'_{\text{free}}[\langle \varepsilon' \rangle A]$$

We then show by another induction that for all ε and $(M, M') \in \mathcal{R}_{\text{free}}[\langle \varepsilon \rangle A]$ we have $(M, M') \in \mathcal{R}'_{\text{free}}[\langle \varepsilon \rangle A]$. For the $(\langle V \rangle, \langle V' \rangle) \in \mathcal{R}_{\text{free}}[\langle 1 \rangle A]$ case use

$$\langle V \rangle \equiv \text{coerce}_{\emptyset \leq \emptyset} \langle V \rangle \quad \langle V' \rangle \equiv \text{coerce}_{\emptyset \leq \emptyset} \langle V' \rangle$$

For $(\text{coerce}_{\varepsilon \leq \varepsilon'} M, \text{coerce}_{\varepsilon \leq \varepsilon'} M) \in \mathcal{R}_{\text{free}}[\langle \varepsilon' \rangle A]$ use the above implication. Finally, for

$$((\text{op } V \text{ to } x. M), (\text{op } V' \text{ to } x. M')) \in \mathcal{R}_{\text{free}}[\langle \{\text{op}\} \cup \varepsilon \rangle A]$$

we assume that

$$(M, M') : \mathcal{R}_{\text{free}}[x : \text{ar}_{\text{op}}] \dot{\rightarrow} \mathcal{R}_{\text{free}}[\langle \varepsilon \rangle A] \quad (V, V') \in \mathcal{R}_{\text{free}}[\text{car}_{\text{op}}]$$

By the inductive hypothesis, the above implication, and the fact that $\mathcal{R}_{\text{free}}$ and $\mathcal{R}'_{\text{free}}$ coincide on ground types, we have

$$(\text{coerce}_{\varepsilon \leq \{\text{op}\} \cup \varepsilon} M, \text{coerce}_{\varepsilon \leq \{\text{op}\} \cup \varepsilon} M') : \mathcal{R}_{\text{free}}[x : \text{ar}_{\text{op}}] \dot{\rightarrow} \mathcal{R}'_{\text{free}}[\langle \{\text{op}\} \cup \varepsilon \rangle A]$$

so by the definition of $\mathcal{R}'_{\text{free}}[\langle \{\text{op}\} \cup \varepsilon \rangle A]$,

$$((\text{op } V \text{ to } x. \text{coerce}_{\varepsilon \leq \{\text{op}\} \cup \varepsilon} M), (\text{op } V' \text{ to } x. \text{coerce}_{\varepsilon \leq \{\text{op}\} \cup \varepsilon} M')) \in \mathcal{R}'_{\text{free}}[\langle \{\text{op}\} \cup \varepsilon \rangle A]$$

The result follows because

$$\text{op } V \text{ to } x. \text{coerce}_{\varepsilon \leq \{\text{op}\} \cup \varepsilon} M \equiv \text{op } V \text{ to } x. M$$

and similarly for $\text{op } V' \text{ to } x. M'$. □

B.3 Call-by-name and call-by-need

This section contains lemmas that are used to relate call-by-name and call-by-need evaluation Section 5.3.

First, we give two lemmas that are useful in later proofs in this section.

Lemma B.3.1 The relations $\mathcal{R}[\underline{A}]^\Delta$ and $\mathcal{R}[\underline{C}]^\Delta$ are partial equivalence relations (symmetric and transitive).

Proof. By induction on A and \underline{C} . Most of the cases are trivial. For returner types, transitivity comes from the definition of closed under divergence. The proof that $(M, M') \in \mathcal{R}[\underline{F}A]^\Delta$ implies $(M', M) \in \mathcal{R}[\underline{F}A]^\Delta$ is a simple induction on $(M, M') \in \mathcal{R}[\underline{F}A]^\Delta$. □

Lemma B.3.2 The Kripke relation $\mathcal{R}[\underline{C}]$ is closed under divergence for every computation type \underline{C} .

Proof. By induction on \underline{C} . We already have transitivity (by Lemma B.3.1), so in each case we only consider the other requirements.

- For returner types this holds by definition.
- For $\underline{C} = \mathbf{unit}$, this is trivial because $\mathcal{R}[\llbracket \mathbf{unit} \rrbracket^\Delta]$ is the total relation.
- For $\underline{C} = \underline{C}_1 \times \underline{C}_2$ we check each of the requirements individually.
 - If $\Delta \vdash M : \underline{C}_1 \times \underline{C}_2$ and $\Delta \vdash M' : \underline{C}_1 \times \underline{C}_2$ are trivially diverging then we have

$$M \equiv P_1 \text{ to } x.P_2 \quad M' \equiv P'_1 \text{ to } x'.P'_2$$

for some

$$P_1 \in \{\underline{z} \mid (\underline{z} : \mathbf{F}B) \in \Delta\} \cup \{\Omega_{\mathbf{F}B}\} \quad P'_1 \in \{\underline{z} \mid (\underline{z} : \mathbf{F}B') \in \Delta\} \cup \{\Omega_{\mathbf{F}B'}\}$$

and P_2, P'_2 . For each $i \in \{1, 2\}$ we have $i'M \equiv P_1 \text{ to } x.i'P_2$, and similarly for M' . Hence $i'M$ and $i'M'$ are both trivially diverging, so are related. This implies that $(M, M') \in \mathcal{R}[\llbracket \underline{C}_1 \times \underline{C}_2 \rrbracket^\Delta]$.

- For the second case of the definition of closed under divergence (Definition 5.3.3) we give the proof for the pair

$$(M \text{ need } y. N, N'[y \mapsto M'])$$

where M and M' are trivially diverging and $(N, N') \in \mathcal{R}[\llbracket \underline{C}_1 \times \underline{C}_2 \rrbracket^{\Delta, y: \mathbf{F}A}]$ (the other three pairs are similar). For each $i \in \{1, 2\}$, we have

$$i'(M \text{ need } y. N) \equiv M \text{ need } y. i'N \quad i'(N'[y \mapsto M']) = (i'N')[y \mapsto M']$$

So it suffices to show that $(M \text{ need } y. i'N, (i'N')[y \mapsto M']) \in \mathcal{R}[\llbracket \underline{C}_i \rrbracket^\Delta]$. This follows from the inductive hypothesis.

- For $\underline{C} = A \rightarrow \underline{D}$ we again check each requirement individually.
 - If $\Delta \vdash M : A \rightarrow \underline{D}$ and $\Delta \vdash M' : A \rightarrow \underline{D}$ are trivially diverging then we have

$$M \equiv P_1 \text{ to } x.P_2 \quad M' \equiv P'_1 \text{ to } x'.P'_2$$

for some

$$P_1 \in \{\underline{z} \mid (\underline{z} : \mathbf{F}B) \in \Delta\} \cup \{\Omega_{\mathbf{F}B}\} \quad P'_1 \in \{\underline{z} \mid (\underline{z} : \mathbf{F}B') \in \Delta\} \cup \{\Omega_{\mathbf{F}B'}\}$$

and P_2, P'_2 . For each $\Delta' \triangleright \Delta$ and $(V, V') \in \mathcal{R}[\llbracket A \rrbracket^{\Delta'}]$ we have $V'M \equiv P_1 \text{ to } x.V'P_2$, and similarly for M' . Hence $V'M$ and $V'M'$ are both trivially diverging, so are related. This implies that $(M, M') \in \mathcal{R}[\llbracket A \rightarrow \underline{D} \rrbracket^\Delta]$.

- For the second case of the definition of closed under divergence (Definition 5.3.3) we give the proof for the pair

$$(M \text{ need } y. N, N'[y \mapsto M'])$$

where M and M' are trivially diverging and $(N, N') \in \mathcal{R}[\llbracket A \rightarrow \underline{D} \rrbracket^{\Delta, y: \mathbf{F}A}]$ (the other three pairs are similar). For each $\Delta' \triangleright \Delta$ and $(V, V') \in \mathcal{R}[\llbracket A \rrbracket^{\Delta'}]$, we have

$$V'(M \text{ need } y. N) \equiv M \text{ need } y. V'N \quad V''(N'[y \mapsto M']) = (V''N')[y \mapsto M']$$

So it suffices to show that $(M \text{ need } y. V'N, (V''N')[y \mapsto M']) \in \mathcal{R}[\llbracket \underline{D} \rrbracket^{\Delta'}]$. This follows from the inductive hypothesis, using the fact that $\mathcal{R}[\llbracket \underline{D} \rrbracket]$ is a Kripke relation to weaken N and N' . \square

We next turn to the proof of the fundamental lemma. As usual, this is by induction on the structure of the terms. Most of the cases are completely standard, so in this section we only give the cases for **need** and **to**, which involve the definition of the logical relation on returner types. These are the next two lemmas.

Lemma B.3.3 If $\Gamma, \underline{x} : FA \vdash M : \underline{C}$ satisfies

$$(M[\sigma], M[\sigma']) \in \mathcal{R}[\underline{C}]^\Delta$$

for all Δ and $(\sigma, \sigma') \in \mathcal{R}[\Gamma, \underline{x} : FA]^\Delta$, then for all $(N, N') \in \mathcal{R}[FA]^\Delta$ and $(\sigma, \sigma') \in \mathcal{R}[\Gamma]^\Delta$, we have

$$(N \text{ need } \underline{x}. M[\sigma], N' \text{ need } \underline{x}. M[\sigma']) \in \mathcal{R}[\underline{C}]^\Delta$$

Proof. By induction on the derivation of $(N, N') \in \mathcal{R}[FA]^\Delta$.

- For transitivity, use the inductive hypothesis and transitivity of $\mathcal{R}[\underline{C}]^\Delta$ (Lemma B.3.1).
- If $N \equiv \langle V \rangle$ and $N' \equiv \langle V' \rangle$ with $(V, V') \in \mathcal{R}[A]^\Delta$ then the pair we are considering is the same up to \equiv as

$$(M[\sigma, \underline{x} \mapsto \langle V \rangle], M[\sigma', \underline{x} \mapsto \langle V' \rangle])$$

The result therefore follows from the assumption about M .

- If N and N' are trivially diverging, the result follows from the fact that $\mathcal{R}[\underline{C}]$ is closed under divergence (Lemma B.3.2).
- For the second case of the definition of closed under divergence (Definition 5.3.3), the pair (N, N') can have one of four forms. We give the proof for one of them (the other three are similar). Suppose that

$$N \equiv N_1 \text{ need } \underline{y}. N_2 \quad N' \equiv N'_2[\underline{y} \mapsto N'_1]$$

where N_1 and N'_1 are trivially diverging and $(N_2, N'_2) \in \mathcal{R}[FA]^\Delta, \underline{y}:FB$. We have

$$\begin{aligned} N \text{ need } \underline{x}. M[\sigma] &\equiv N_1 \text{ need } \underline{y}. N_2 \text{ need } \underline{x}. M[\sigma] \\ N' \text{ need } \underline{x}. M[\sigma'] &\equiv (N'_2 \text{ need } \underline{x}. M[\sigma'])[\underline{y} \mapsto N'_1] \end{aligned}$$

Since $\mathcal{R}[\underline{C}]$ is closed under divergence (Lemma B.3.2), it therefore suffices to show that

$$(N_2 \text{ need } \underline{x}. M[\sigma], N'_2 \text{ need } \underline{x}. M[\sigma']) \in \mathcal{R}[\underline{C}]^\Delta, \underline{y}:FB$$

This follows from the inductive hypothesis, using the fact that Kripke relations respect weakening (σ and σ' are weakened). \square

Lemma B.3.4 If $\Gamma, x : A \vdash M : \underline{C}$ satisfies

$$(M[\sigma], M[\sigma']) \in \mathcal{R}[\underline{C}]^\Delta$$

for all Δ and $(\sigma, \sigma') \in \mathcal{R}[\Gamma, x : A]^\Delta$, then for all $(N, N') \in \mathcal{R}[FA]^\Delta$ and $(\sigma, \sigma') \in \mathcal{R}[\Gamma]^\Delta$ we have

$$(N \text{ to } x. M[\sigma], N' \text{ to } x. M[\sigma']) \in \mathcal{R}[\underline{C}]^\Delta$$

Proof. By induction on the derivation of $(N, N') \in \mathcal{R}[FA]^\Delta$.

- For transitivity, use the inductive hypothesis and transitivity of $\mathcal{R}[\underline{C}]^\Delta$ (Lemma B.3.1).

- If $N \equiv \langle V \rangle$ and $N' \equiv \langle V' \rangle$ with $(V, V') \in \mathcal{R}[[A]]^\Delta$ then the pair we are considering is the same up to \equiv as

$$(M[\sigma, x \mapsto \langle V \rangle], M[\sigma', x \mapsto \langle V' \rangle])$$

The result therefore follows from the assumption about M .

- If N and N' are trivially diverging then for some

$$P_1 \in \{\underline{y} \mid (\underline{y} : \mathbf{F} B) \in \Delta\} \cup \{\Omega_{\mathbf{F} B}\} \quad P'_1 \in \{\underline{y} \mid (\underline{y} : \mathbf{F} B') \in \Delta\} \cup \{\Omega_{\mathbf{F} B'}\}$$

and P_2, P'_2 we have

$$N \equiv P_1 \mathbf{to} z. P_2 \quad N' \equiv P'_1 \mathbf{to} z'. P'_2$$

Hence

$$N \mathbf{to} x. M[\sigma] \equiv P_1 \mathbf{to} z. P_2 \mathbf{to} x. M[\sigma] \quad N' \mathbf{to} x. M[\sigma'] \equiv P'_1 \mathbf{to} z'. P'_2 \mathbf{to} x. M[\sigma']$$

so the two computations we wish to relate are both trivially diverging, and we can use the fact that $\mathcal{R}[[\underline{C}]]$ is closed under divergence (Lemma B.3.2).

- For the second case of the definition of closed under divergence (Definition 5.3.3), the pair (N, N') can have one of four forms. We give the proof for one of them (the other three are similar). Suppose that

$$N \equiv N_1 \mathbf{need} \underline{y}. N_2 \quad N' \equiv N'_2[\underline{y} \mapsto N'_1]$$

where N_1 and N'_1 are trivially diverging and $(N_2, N'_2) \in \mathcal{R}[[\mathbf{F} A]]^{\Delta, \underline{y} : \mathbf{F} B}$. We have

$$\begin{aligned} N \mathbf{to} x. M[\sigma] &\equiv N_1 \mathbf{need} \underline{y}. N_2 \mathbf{to} x. M[\sigma] \\ N' \mathbf{to} x. M[\sigma'] &\equiv (N'_2 \mathbf{to} x. M[\sigma'])[\underline{y} \mapsto N'_1] \end{aligned}$$

Since $\mathcal{R}[[\underline{C}]]$ is closed under divergence (Lemma B.3.2), it therefore suffices to show that

$$(N_2 \mathbf{to} x. M[\sigma], N'_2 \mathbf{to} x. M[\sigma']) \in \mathcal{R}[[\underline{C}]]^{\Delta, \underline{y} : \mathbf{F} B}$$

This follows from the inductive hypothesis, using the fact that Kripke relations respect weakening (σ and σ' are weakened). \square

At this point, we know that the fundamental lemma holds. We next show that call-by-name and call-by-need are related by the logical relation:

Lemma B.3.5 If $\Gamma, \underline{x} : \mathbf{F} A \vdash M : \underline{C}$ then for all $(N, N') \in \mathcal{R}[[\mathbf{F} A]]^\Delta$ and $(\sigma, \sigma') \in \mathcal{R}[[\Gamma]]^\Delta$, each of the following pairs is in $\mathcal{R}[[\underline{C}]]^\Delta$:

$$\begin{aligned} (N \mathbf{need} \underline{x}. M[\sigma], N' \mathbf{need} \underline{x}. M[\sigma']) &\quad (M[\sigma][\underline{x} \mapsto N], M[\sigma'][\underline{x} \mapsto N']) \\ (M[\sigma][\underline{x} \mapsto N], N' \mathbf{need} \underline{x}. M[\sigma']) &\quad (N \mathbf{need} \underline{x}. M[\sigma], M[\sigma'][\underline{x} \mapsto N']) \end{aligned}$$

Proof. For the two pairs on the top row, we apply Lemma B.3.3 and the fundamental lemma (Lemma 5.3.4). The remaining two require more effort. We give only the proof for

$$(M[\sigma][\underline{x} \mapsto N], N' \mathbf{need} \underline{x}. M[\sigma'])$$

The other is similar. The proof is by induction on the derivation of $(N, N') \in \mathcal{R}[[\mathbf{F} A]]^\Delta$.

- For transitivity, use the inductive hypothesis and transitivity of $\mathcal{R}[[\underline{C}]]^\Delta$ (Lemma B.3.1).

- If $N \equiv \langle V \rangle$ and $N' \equiv \langle V' \rangle$ with $(V, V') \in \mathcal{R}[\![A]\!]^\Delta$ then the pair we are considering is the same up to \equiv as

$$(M[\sigma, \underline{x} \mapsto \langle V \rangle], M[\sigma', \underline{x} \mapsto \langle V' \rangle])$$

The result therefore follows from the fundamental lemma (Lemma 5.3.4).

- If N and N' are trivially diverging, the result follows from the fact that $\mathcal{R}[\![\underline{C}]\!]$ is closed under divergence (Lemma B.3.2).
- For the second case of the definition of closed under divergence (Definition 5.3.3), suppose that $\Delta \vdash N_1 : \mathbf{F} B$ and $\Delta \vdash N'_1 : \mathbf{F} B$ are trivially diverging, and that $(N_2, N'_2) \in \mathcal{R}[\![\mathbf{F} A]\!]^{\Delta, y : \mathbf{F} B}$. We consider each of the four forms of pair (N, N') in turn. In each case, we use the fact that the logical relation is a partial equivalence relation (Lemma B.3.1), and that σ is related to itself (which follows from the fundamental lemma). We also weaken terms, and therefore use the fact that we have Kripke relations.

- If $(N, N') = (N_1 \text{ need } y. N_2, N'_1 \text{ need } y. N'_2)$ then

$$\begin{aligned} M[\sigma][\underline{x} \mapsto N] &\equiv M[\sigma, \underline{x} \mapsto (N_1 \text{ need } y. N_2)] \\ &\mathcal{R} M[\sigma, \underline{x} \mapsto N_2[\underline{y} \mapsto N_1]] && (5.3.4, \text{ closure under divergence}) \\ &\equiv M[\sigma, \underline{x} \mapsto N_2][\underline{y} \mapsto N_1] \\ &\mathcal{R} N'_1 \text{ need } y. N'_2 \text{ need } \underline{x}. M[\sigma'] && (\text{IH, closure under divergence}) \\ &\equiv N' \text{ need } \underline{x}. M[\sigma'] \end{aligned}$$

- If $(N, N') = (N_2[\underline{y} \mapsto N_1], N'_2[\underline{y} \mapsto N'_1])$ then

$$\begin{aligned} M[\sigma][\underline{x} \mapsto N] &\equiv M[\sigma, \underline{x} \mapsto N_2][\underline{y} \mapsto N_1] \\ &\mathcal{R} (N'_2 \text{ need } \underline{x}. M[\sigma'])[\underline{y} \mapsto N'_1] && (\text{IH, closure under divergence}) \\ &\equiv N' \text{ need } \underline{x}. M[\sigma'] \end{aligned}$$

- If $(N, N') = (N_2[\underline{y} \mapsto N_1], N'_1 \text{ need } y. N'_2)$ then

$$\begin{aligned} M[\sigma][\underline{x} \mapsto N] &\equiv M[\sigma, \underline{x} \mapsto N_2][\underline{y} \mapsto N_1] \\ &\mathcal{R} N'_1 \text{ need } y. N'_2 \text{ need } \underline{x}. M[\sigma'] && (\text{IH, closure under divergence}) \\ &\equiv N' \text{ need } \underline{x}. M[\sigma'] \end{aligned}$$

- If $(N, N') = (N_1 \text{ need } y. N_2, N'_2[\underline{y} \mapsto N'_1])$ then

$$\begin{aligned} M[\sigma][\underline{x} \mapsto N] &\equiv M[\sigma, \underline{x} \mapsto (N_1 \text{ need } y. N_2)] \\ &\mathcal{R} M[\sigma, \underline{x} \mapsto N_2[\underline{y} \mapsto N_1]] && (5.3.4, \text{ closure under divergence}) \\ &\equiv M[\sigma, \underline{x} \mapsto N_2][\underline{y} \mapsto N_1] \\ &\mathcal{R} (N'_2 \text{ need } \underline{x}. M[\sigma'])[\underline{y} \mapsto N'_1] && (\text{IH, closure under divergence}) \\ &\equiv N' \text{ need } \underline{x}. M[\sigma'] \quad \square \end{aligned}$$

Finally, we show that we can use the logical relation to prove contextual equivalences. First we note that closed, trivially diverging computations actually diverge:

Lemma B.3.6 If $\diamond \vdash M : \underline{C}$ is trivially diverging then $M \equiv \Omega_{\underline{C}}$.

Proof. There are no free computation variables, so $M \equiv \Omega_{\mathbf{F} A} \text{ to } x. N$ for some N . The result then follows from the η -law for the empty type. \square

Next, we note that on ground types the logical relation matches the equational theory:

Lemma B.3.7 If G is a ground type and $(V, V') \in \mathcal{R}[[G]]^\Delta$, then $V \equiv W \equiv V'$ for some closed value W .

Proof. By induction on G .

- If $G = \mathbf{unit}$ then $V \equiv () \equiv V'$ by the η -law for \mathbf{unit} .
- If $G = G_1 \times G_2$ then

$$V \equiv (\mathbf{fst} V_1, \mathbf{snd} V_2) \equiv (W_1, W_2) \equiv (\mathbf{fst} V'_1, \mathbf{snd} V'_2) \equiv V'$$

where W_1 and W_2 come from applying the inductive hypothesis.

- If $G = \mathbf{empty}$ then we have a contradiction ($\mathcal{R}[[\mathbf{empty}]]^\Delta$ is empty).
- If $G = G_1 + G_2$ then there are two cases to consider. If $V \equiv \mathbf{inl}_{G_2} V_1$ and $V' \equiv \mathbf{inl}_{G_2} V'_1$ with $(V_1, V'_1) \in \mathcal{R}[[G_1]]^\Delta$, then

$$V \equiv \mathbf{inl}_{G_2} V_1 \equiv \mathbf{inl}_{G_2} W \equiv \mathbf{inl}_{G_2} V'_1 \equiv V'$$

where W comes from applying the inductive hypothesis. The \mathbf{inr} case is similar. \square

We also show that closed returners of ground type either return a result or diverge. To do this, we also need to consider computations with free computation variables. The next lemma is where it is important that computation variables are only bound inside the definition of the logical relation to trivially diverging computations.

Lemma B.3.8 If G is a ground type and $(M, M') \in \mathcal{R}[[FG]]^\Delta$ then one of the following holds:

1. $M \equiv \langle V \rangle \equiv M'$ for some closed value V .
2. M and M' are both trivially diverging.

Proof. By definition, $\mathcal{R}[[FA]]$ is the smallest closed-under-divergence relation such that

$$(V, V') \in \mathcal{R}[[A]]^\Delta \Rightarrow (\langle V \rangle, \langle V' \rangle) \in \mathcal{R}[[FA]]^\Delta \quad (\text{B.1})$$

so we can proceed by induction on the derivation of $(M, M') \in \mathcal{R}[[FA]]^\Delta$. If this holds because of Equation (B.1) then the result follows immediately from Lemma B.3.7. Transitivity is trivial (using transitivity of \equiv). Otherwise, we consider each of the cases in the definition of closed under divergence (Definition 5.3.3).

- In the first case of the definition, both computations are required to be trivially diverging.
- For the second case we have two trivially diverging computations, which necessarily have the form

$$M \equiv P_1 \mathbf{to} x. P_2 \quad M' \equiv P'_1 \mathbf{to} x'. P'_2$$

for some

$$P_1 \in \{\underline{z} \mid (z : \mathbf{FB}) \in \Delta\} \cup \{\Omega_{\mathbf{FB}}\} \quad P'_1 \in \{\underline{z} \mid (z : \mathbf{FB}') \in \Delta\} \cup \{\Omega_{\mathbf{FB}'}\}$$

and also have some $(N_2, N'_2) \in \mathcal{R}[[FA]]^{\Delta, y:FA}$. By the inductive hypothesis, there are two cases to consider:

- We have $N_2 \equiv \langle V \rangle \equiv N'_2$ for some closed V . Then for each of the pairs in the second case of Definition 5.3.3 we have

$$M \equiv \langle V \rangle \equiv M'$$

where we use the garbage collection rule and the fact that V is closed to remove the **need** bindings.

- Both N_2 and N'_2 are trivially diverging, which implies that

$$N_2 \equiv Q_1 \text{ to } z. Q_2 \quad N'_2 \equiv Q'_1 \text{ to } z'. Q'_2$$

for some

$$Q_1 \in \{\underline{z} \mid (z : \mathbf{F} B_1) \in \Delta\} \cup \{\Omega_{\mathbf{F} B_1}\} \quad Q'_1 \in \{\underline{z} \mid (z : \mathbf{F} B'_1) \in \Delta\} \cup \{\Omega_{\mathbf{F} B'_1}\}$$

We show that M is trivially diverging (M' is trivially diverging by an identical proof). By looking at the four pairs in the second case of Definition 5.3.3, we know that M has one of two forms. The first is

$$M \equiv (P_1 \text{ to } x. P_2) \text{ need } \underline{y}. Q_1 \text{ to } z. Q_2$$

In this case, if $Q_1 \neq \underline{y}$ then

$$M \equiv Q_1 \text{ to } z. (P_1 \text{ to } x. P_2) \text{ need } \underline{y}. Q_2$$

and if $Q_1 = \underline{y}$ then

$$M \equiv P_1 \text{ to } x. P_2 \text{ to } z. Q_2[\underline{y} \mapsto \langle z \rangle]$$

both of which are trivially diverging. The second form of M is

$$M \equiv (Q_1 \text{ to } z. Q_2)[\underline{y} \mapsto P_1 \text{ to } x. P_2]$$

In this case, if $Q_1 \neq \underline{y}$ then

$$M \equiv Q_1 \text{ to } z. (Q_2[\underline{y} \mapsto (P_1 \text{ to } x. P_2)])$$

and if $Q_1 = \underline{y}$ then

$$M \equiv P_1 \text{ to } x. P_2 \text{ to } z. Q_2[\underline{y} \mapsto (P_1 \text{ to } x. P_2)]$$

both of which are again trivially diverging. □

Finally, we can put the last few lemmas together to show that the logical relation matches the equational theory on closed returners of ground type.

Corollary B.3.9 If G is a ground type and $(M, M') \in \mathcal{R}[[\mathbf{F} G]]^\diamond$ then $M \equiv M'$.

Proof. Applying Lemma B.3.8 gives us two cases to consider:

- If we have $M \equiv \langle V \rangle \equiv M'$, then we are done.
- If both M and M' are both trivially diverging, then we apply Lemma B.3.6. □