

Selective Strictness and Parametricity in Structural Operational Semantics, Inequationally

Janis Voigtländer

Institut für Theoretische Informatik
Technische Universität Dresden
01062 Dresden, Germany
voigt@tcs.inf.tu-dresden.de

Patricia Johann*

Department of Computer Science
Rutgers University
Camden, NJ 08102 USA
pjohann@crab.rutgers.edu

Abstract

Parametric polymorphism constrains the behavior of pure functional programs in a way that allows the derivation of interesting theorems about them solely from their types, i.e., virtually for free. The formal background of such ‘free theorems’ is well developed for extensions of the Girard-Reynolds polymorphic lambda calculus by *algebraic datatypes* and *general recursion*, provided the resulting calculus is endowed with either a purely strict or a purely nonstrict semantics. But modern functional languages like Clean and Haskell, while using nonstrict evaluation by default, also provide means to enforce strict evaluation of subcomputations at will. The resulting *selective strictness* gives the advanced programmer explicit control over evaluation order, but is not without semantic consequences: it breaks standard parametricity results. This paper develops an operational semantics for a core calculus supporting all the language features emphasized above. Its main achievement is the characterization of observational approximation with respect to this operational semantics via a carefully constructed logical relation. This establishes the formal basis for new parametricity results, as illustrated by several example applications, including the first complete correctness proof for short cut fusion in the presence of selective strictness. The focus on observational approximation, rather than equivalence, allows a finer-grained analysis of computational behavior in the presence of selective strictness than would be possible with observational equivalence alone.

Keywords: Clean, Haskell, extensionality principles, fixpoint recursion, functional programming languages, identity extension, lambda calculus, logical relations, mixing strict and nonstrict semantics, parametric polymorphism, program transformations, *seq*, short cut fusion, theorems for free, types

*Supported in part by National Science Foundation grant CCF-0429072.

1 Introduction

To support the production of software that is rapidly prototyped, reliable, and maintainable, both programmers and programming language designers need to have at their disposal techniques for reasoning effectively about program semantics. One technique which is suitable for reasoning about programs in polymorphically typed functional languages is based on parametricity properties [Rey83] — more colorfully known as free theorems [Wad89] — associated with polymorphic functions.¹ A *parametricity property* formalizes the intuition that a polymorphic function must behave uniformly, i.e., must use the same algorithm to compute its result regardless of the concrete type at which it is instantiated. The parametricity property for a polymorphic function can be derived solely from the type of that function, with no knowledge whatsoever of the function’s actual definition; this is the sense in which parametricity properties are obtained ‘for free’. But a polymorphic function is only guaranteed to satisfy its parametricity property if a *parametricity theorem* [Wad89] — originally called *abstraction theorem* [Rey83] — holds for the underlying language of which it is part. Thus parametricity properties are not actually ‘free’ at all: the illusion of freeness is merely a reflection of the ease with which parametricity properties can be derived *once the considerable task of proving a parametricity theorem which guarantees that they hold* has been completed.

Some popular applications which are typically justified by cleverly instantiating parametricity properties actually require programming languages of interest to satisfy properties which are stronger than those guaranteed solely by their parametricity theorems. In particular, the past decade or so has seen the development of a number of parametricity-based techniques for automatically transforming modular-but-inefficient programs in nonstrict functional languages such as Haskell [Pey03] into efficient-but-monolithic equivalents [GLP93, TM95, Joh02, Sve02, Voi02, GJUV05]. Automatic transformation is essential to mitigating the inherent tension between the design and development of programs that are easy to reason about and the run-time efficiency of those programs. But the semantic correctness of parametricity-based program transformations with respect to some (typically equivalence or preorder) relation on programs in which we are interested does not always follow from parametricity theorems alone.

The key to stating and proving the parametricity theorem for a given language (or, more precisely, for a given model of a language) is to interpret the types of the language according to certain systematically-constructed relations, called *logical relations* [Plo73, Fri75, Rey83, Sta85]. A parametricity theorem then asserts that every closed term of closed type is related to itself by the relational interpretation of its type.² But proving correctness of parametricity-based program transformations

¹As is standard, we take the term ‘polymorphic’ to refer to parametric polymorphism, as opposed to *ad hoc* polymorphism of the kind supported, for example, by type classes [WB89].

²More precisely, this description applies only in an operational semantics framework. For a denotational model, the parametricity theorem asserts that *the interpretation of* every closed term of closed type is related to itself by the relational interpretation of its type.

in a given model often requires more than just reflexivity of the logical relation: the relational interpretation of each type — including \forall -types — must coincide exactly with the relation induced by the model. (See Theorem 7.19 for an example of a transformation whose correctness proof requires full coincidence.) When this is the case, i.e., when there is a logical relation which coincides with the relation on programs in which we are interested, we say that we have a *parametric model* of that relation. Constructing such a parametric model — and, therefore, proving correctness with respect to a relation of interest of certain parametricity-based program transformations — thus entails ensuring that a strengthening of the parametricity theorem holds. The required strengthening, given formally after Corollary 6.8 below, is closely related to Reynolds’ *identity extension lemma* [Rey83].

1.1 Parametric models and program transformations

The Girard-Reynolds polymorphic lambda calculus λ^\forall [Gir72, Rey74] — which is also known as ‘System F’ and provides the theoretical underpinning for many polymorphically typed functional languages — is well-known to admit parametric models [BFSS90, Has91, RR94]. Results derived from parametricity thus hold unconditionally in λ^\forall . But for calculi that more closely resemble modern functional languages the story is not so simple. For instance, adding a fixpoint primitive to a calculus, thus capturing general recursive definitions, weakens its parametricity properties by imposing strictness and continuity conditions on (some of the) functions appearing in those properties [Wad89, LP96]. The impact on identity extension and related strengthenings of parametricity theorems is also highly nontrivial. Moreover, to help programmers control the time and space behavior of programs, nonstrict languages often provide primitives for selectively forcing strict evaluation in computations. Such primitives can further compromise parametricity-based results, such as the correctness of program transformations, for calculi which include them; see, e.g., [JV04, JV06], Appendix B of [Voi02], and Section 2 below.

To study the impact of fixpoint recursion, Pitts [Pit00] constructed an operational model for the calculus PolyPCF obtained by adding a fixpoint primitive and an algebraic datatype to λ^\forall and endowing the result with a nonstrict semantics. Pitts proved that his model is parametric in the sense discussed above, paving the way for its use in [Joh02, Joh03, Joh05] to give the first fully satisfactory correctness proofs for *short cut fusion* [GLP93] and related program transformations. This suggests that an operational approach is also suitable for formal study of the (additional) impact of selective strictness on parametricity and program transformations based on it. Nevertheless, like recent work in the area [RS04, BMP06, Møg06], our studies of this issue [JV04, JV06] were first performed in a denotational setting. This is primarily because working in a denotational setting allowed a more intuitive initial approach to the problem. Concretely, to understand how free theorems and parametricity-based program transformations are affected by the presence of the polymorphic strict evaluation primitive *seq* in Haskell, we started with the standardly accepted, but naive, denotational model for Haskell and constructed a

logical relation for which a parametricity theorem could be proved. We then used this logical relation to derive appropriate variations of standard free theorems and to recover (partial and total) correctness of some well-known parametricity-based program transformations in the presence of *seq*. This work was significant in that it showed the decade-old conventional wisdom regarding the impact of *seq* on parametricity relative to the standardly accepted denotational model for Haskell to be in error. In particular, it was the first to correctly identify preconditions under which parametricity-based program transformations can safely be applied in the presence of *seq*. Indeed, while previous correctness arguments were for approximations of — and thus ultimately targeted — ‘Haskell minus *seq*’, our correctness arguments instead targeted ‘Haskell including *seq*’. The work in [JV04, JV06] thus helped move the focus of the discussion about the behavior of parametricity-based transformations more toward ‘real Haskell’.

But despite providing important insights, the work in [JV04, JV06] has two significant shortcomings. The first is that the standardly accepted denotational model in which we were working could not be shown to be parametric because the statement corresponding to Reynolds’ identity extension lemma for the logical relation which we offered as witness that this model is parametric could not be established. This statement remains a conjecture, which is unfortunate since correctness proofs for some program transformations of interest — such as short cut fusion — depend on it. The second shortcoming is that the status of the standardly accepted denotational model with respect to the observable behavior of Haskell programs has never been made precise. That is, the operational semantics that implementations of Haskell are, according to the language definition, expected to satisfy is not guaranteed to be in any way tied into this denotational semantics. As a result, equivalence in this denotational setting between a given Haskell program and the new program obtained by applying parametricity-based transformations to it is insufficient to allow us to draw any conclusions at all about the relationship between the observable behavior of the given program and the observable behavior of the transformed program.

Even with the above caveats, working in a denotational setting turned out to be an important first step in understanding the impact of *seq* on parametricity. As noted above, before [JV04, JV06], the impact of *seq* on parametricity, even with respect to the standardly accepted denotational model, was misunderstood, and [JV04, JV06] offer a complete correction to those long-standing misconceptions. Given the fruitfulness of our initial, intuitive approach to understanding the interplay between selective strictness and parametricity, it is only natural to now ask whether — and, if so, how — the insights developed in that denotational setting can be transferred to a more rigorous operational setting.

1.2 This paper

In this paper we study the calculus PolySeq which is obtained by adding a Haskell-like strictness primitive to Pitts’ PolyPCF. To provide a more fine-grained analysis

of program behavior than is possible by reasoning about observational equivalence alone, we instead focus on observational approximation. In other words, as in [JV04, JV06], we work in an inequational setting rather than an equational one. This is reasonable since, as noted in those earlier papers, the effect of selective strictness on an erstwhile (parametricity-based) program equivalence is to potentially make one side of the equivalence less defined than the other. The key contribution of this paper is thus the construction of a *parametric model of observational approximation* for PolySeq. This ports the results from [JV04, JV06] to an operational setting in which the outstanding issues mentioned in the previous subsection are resolved.

Our model of PolySeq observational approximation is similar to Pitts’ operational semantics-based parametric model of PolyPCF observational equivalence [Pit00], but his construction has been refined to accommodate inequational reasoning, as well as the extra constraints on relational interpretations of types imposed by selective strictness. While these constraints are in some sense ‘just’ operational analogues of denotational ones discussed in [JV04, JV06], their translation and incorporation into the operational semantics is delicate, and is accomplished rather differently than we anticipated in those earlier papers. Another difference between our construction and Pitts’ is that ours starts from a small-step semantics, rather than from a big-step semantics. This allows us to more explicitly model the operational behavior of *seq*, while at the same time providing some new insights into techniques for modularly constructing parametric models for extensions of λ^\forall that support multiple additional language features.

Constructing a parametric model of PolySeq observational approximation is a highly nontrivial undertaking precisely because the impact of *seq* on termination is subtle and complex. At first glance, it may appear that this impact can be accounted for by extending the operationally-based techniques of [Pit00] along the lines suggested for ‘Lazy PCF’ in the conclusion of that paper. But *seq* impacts the termination behavior of programs in ways that go beyond just making it possible to observe termination of whole programs at function types, or, indeed, at any type. In fact, *seq* can be used to force evaluation of *any* term of *any* type appearing at *any* place in a program, so that termination of both intermediate and ‘top-level’ computations of any type becomes observable. Since a parametric model of PolySeq observational approximation must relate every term of every type to every term of the same type whose observable behavior the original term approximates in every program context — in the sense that filling the program context with the original term yields a term whose termination implies termination of any term obtained by filling the same context with any term to which the original term is related — this must be taken into account when defining the logical relation which witnesses parametricity of the model. We accomplish this for the model constructed in this paper by enforcing a *convergence-preservation* property at all types, in addition to appropriately restricting the relations over which quantification is performed when defining the relational interpretation of \forall -types. While it is easy to see that convergence-preservation is necessary, it is not at all obvious that it combines with the restrictions on relations needed to accommodate fixpoints and algebraic data

types in a way that allows a Pitts-like construction to go through. Showing that it does is the major technical contribution of the paper, and the source from which all results about our parametric model of PolySeq observational approximation are ultimately derived.

An important secondary contribution of this paper is to show how the model we construct here can be used to prove the partial (or, indeed, total) correctness, with respect to observational approximation (resp., observational equivalence) of parametricity-based transformations on PolySeq programs. This will be exemplified for the classical short cut fusion technique [GLP93]. Although we focus in this paper on a calculus whose only algebraic datatypes are lists, all of the results here carry over to extensions of PolySeq with non-list algebraic datatypes; in particular, our results extend easily to prove correctness of short cut fusion for non-list algebraic datatypes. The techniques introduced in this paper can also be used to prove (partial or total) correctness of parametricity-based transformations which fuse consumers of algebraic data structures with producers parameterized over substitution values [Joh02], and which are category-theoretic duals of short cut fusion for algebraic datatypes [TM95, Sve02]. At present, it is not known how the logical relation techniques we use in this paper can be adapted for arbitrary recursive types.

The ultimate goal of the line of research advanced in this paper is the development of tools for reasoning about parametricity properties of, and parametricity-based transformations on programs in, real programming languages rather than toy calculi. This provides another point in favor of the operational approach taken in this paper. For while the Glasgow Haskell Compiler [GHC] uses a variant of λ^v as its intermediate language Core, a well-defined denotational semantics is not currently known even for relevant subsets of Core. It is thus unclear whether results derived relative to any particular denotational model of, say, PolySeq would eventually shed any light at all on parametricity properties of Core. On the other hand, we derive our results in this paper relative to an operational semantics very much like the one that implementations like GHC are expected to satisfy, so that the parametricity results we prove for PolySeq do indeed provide insights into those of Core and, by extension, full Haskell.

The remainder of this paper is structured as follows. Section 2 informally discusses selective strictness and how it can break parametricity. Our formal study of this issue begins in Section 3 with the syntax and semantics of PolySeq. In Section 4 we study PolySeq termination, introduce restrictions on relations to accommodate the fixpoint and selective strictness primitives, and examine their interplay. Based on the aforementioned restrictions, and indeed driven by them, we define our logical relation in Section 5. In Section 6 we prove our main technical result (Corollary 6.8), namely that the logical relation characterizes PolySeq observational approximation. In Section 7 we use this result to establish extensionality principles (Lemmas 7.6, 7.7, and 7.10), to enumerate terms up to observational equivalence (Lemmas 7.16 and 7.18), and to prove correctness of short cut fusion (Theorem 7.19). Section 8 concludes and explores the relationship of our work to other work in the area. Throughout, proofs that are either routine or essentially

repeated or elaborated from [Pit00] are omitted.

2 Selective strictness breaks parametricity

In PolyPCF, as in other nonstrict calculi, function arguments are evaluated only when required. But evaluation can be explicitly forced in the presence of a strict evaluation primitive such as Haskell’s *seq*. Although it was not originally given a polymorphic type, *seq* has for quite some time now been denotationally specified as follows in Haskell language definitions such as [Pey03]:³

$$\begin{aligned} seq &:: \forall \alpha \beta. \alpha \rightarrow \beta \rightarrow \beta \\ seq \perp b &= \perp \\ seq a b &= b \quad \text{if } a \neq \perp \end{aligned}$$

Here \perp is the undefined value corresponding to a nonterminating computation or a runtime error, such as might be obtained as the result of a failed pattern match. The operational behavior of *seq* is to evaluate its first argument before returning its second argument. Note that *seq* can be applied at all types.

A prototypical example of a function which uses *seq* — indeed, the one probably discussed most frequently on the Haskell Mailing List [HML] — is:

$$\begin{aligned} foldl' &:: \forall \alpha \beta. (\beta \rightarrow \alpha \rightarrow \beta) \rightarrow \beta \rightarrow [\alpha] \rightarrow \beta \\ foldl' f z [] &= z \\ foldl' f z (h : t) &= seq z' (foldl' f z' t) \\ &\quad \text{where } z' = f z h \end{aligned}$$

Here *seq* ensures that the accumulating parameter is computed immediately in each recursive step rather than constructing a complex closure, representing the overall accumulation, which would be computed only at the end of the call to *foldl'*. Thus, in many situations *foldl'* offers a useful and easily obtained efficiency improvement over the Haskell prelude function *foldl*. Further examples of programs which make use of selective strictness via *seq* can be found in [THLP98]. Note that other means of explicitly introducing strictness in Haskell programs — e.g., strict datatypes, the strict application function $\$!$, and the recently introduced bang patterns — are all definable in terms of *seq*. Similarly, language features for selective strictness in Clean [CLR] are interdefinable with *seq* [EM06].

The impact of selective strictness on the semantics of, and reasoning techniques for, languages like Clean and Haskell can be severe. This impact has been studied, for example, in [HK05] and [EM06], as well as in our own recent work [JV04, JV06]. It

³Actually, the type given there for *seq* is just $\alpha \rightarrow \beta \rightarrow \beta$. But the semantics of α and β occurring free in the type of *seq* is exactly an implicit universal quantification. For clarity, we prefer to make all quantification over type variables explicit. This is supported by most Haskell implementations via the keyword **forall**. Note that while adding or omitting outermost quantifications (as here for *seq*’s type) is just a matter of syntactic convenience, the positioning of ‘inner \forall s’ is crucial for functions like *build* in Figure 1 below (see also the footnote on the next page).

has also been noticed somewhat more in passing in [Voi02], [DJ04], and [DHJG06]. That the mixture of nonstrict and strict evaluation is currently a topic of great interest in programming languages research is further evidenced by recent work in the areas of program verification [ABB⁺05] and implementation [RMP06]. Our specific focus in this paper is on the impact of selective strictness on parametric polymorphism in nonstrict languages.

The classic example of a parametricity-based program transformation is the short cut fusion rule [GLP93]. This rule eliminates intermediate lists from compositions of list producers written in terms of *build* and list consumers written in terms of *foldr* using the following rule for appropriately typed g , c , and n :

$$\mathit{foldr} \ c \ n \ (\mathit{build} \ g) = g \ c \ n \quad (1)$$

Definitions of *foldr* and *build* are given in Figure 1. The function *foldr*, which is standard in the Haskell prelude, takes as input a function c , a value n , and a list l , and produces a value by replacing all occurrences of $(:)$ in l by c and any occurrence of $[]$ in l by n . For instance, $\mathit{foldr} \ (+) \ 0 \ l$ sums the (numeric) elements of the list l . The function *build*, on the other hand, takes as input a polymorphic function g providing a type-independent template for constructing ‘abstract’ lists, and applies it to the list constructors $(:)$ and $[]$ to get a corresponding ‘concrete’ list.⁴ For example, $\mathit{build} \ (\lambda c \ n \rightarrow c \ 4 \ (c \ 9 \ n))$ produces the list $[4,9]$. Applying rule (1) to a composition matching its left-hand side yields a corresponding instance of the right-hand side which avoids constructing intermediate lists produced by *build* g and immediately consumed by *foldr* $c \ n$. This is accomplished by applying g to the $(:)$ - and $[]$ -replacement functions c and n directly.

$$\begin{aligned} \mathit{foldr} &:: \forall \alpha \ \beta. (\alpha \rightarrow \beta \rightarrow \beta) \rightarrow \beta \rightarrow [\alpha] \rightarrow \beta \\ \mathit{foldr} \ c \ n &= f \ \mathbf{where} \ f \ [] = n \\ & \quad f \ (h : t) = c \ h \ (f \ t) \\ \\ \mathit{build} &:: \forall \alpha. (\forall \beta. (\alpha \rightarrow \beta \rightarrow \beta) \rightarrow \beta \rightarrow \beta) \rightarrow [\alpha] \\ \mathit{build} \ g &= g \ (:) \ [] \end{aligned}$$

Figure 1: Haskell functions for short cut fusion.

The short cut fusion rule is derived from the parametricity property for g , or, more accurately, for g ’s type. But the simple instantiation in which $g = \mathit{seq}$, $c = \perp$, and $n = 0$ shows that (1) is no longer an equivalence if seq is present. The intuitive reason for this breakdown of short cut fusion lies in the differences in definedness and strictness properties of the arguments supplied to g before and after applying the short cut fusion rule. The list constructors $(:)$ and $[]$ passed to g by *build* are

⁴Taking a polymorphic function as argument, *build* has a *rank-2 type* [Lei83]. While such higher-rank types are not covered by the current Haskell standard [Pey03], they are actually supported by most implementations. For recent work on type inference in this setting see [VWP06].

both non- \perp , and so is any value obtained by combining them. But since no such *a priori* guarantees exist for their replacement functions c and n , the same use of seq inside g might result in the undefined value \perp on the right-hand side of (1) and a non- \perp value on the left-hand side.

Before publication of [JV04, JV06], a folk theorem had long held that parametricity properties remain valid in the presence of seq if all of the functions appearing in them (where one is free to make a choice) are strict and total. But as shown there, this is not the case: although strictness and totality of the consumer ($foldr\ c\ n$) just happen to be sufficient for recovering correctness of the short cut fusion rule when seq is present, they are not enough to recover the parametricity properties of all polymorphic functions in this situation. In fact, it is probably this happenstance vis-a-vis the short cut fusion rule that is responsible for the failure of the folklore approach to parametricity in the presence of seq having gone unnoticed for so long. In [JV04, JV06] we gave constraints which guarantee that parametricity properties of polymorphic functions hold in the denotational setting considered there, even in the presence of seq . In this paper we translate these constraints to the operational setting and show how their operational counterparts can be used to recover parametricity properties of polymorphic functions — as well as correctness of program transformations based on parametricity — without the shortcomings of the denotationally-based development of [JV04, JV06] discussed in Section 1.1. Our results for short cut fusion in particular can be found in Section 7.5.

3 PolySeq

As a testbed for exploring the impact of selective strictness on parametricity results in an operational setting we use a concrete calculus, PolySeq, which extends the Girard-Reynolds calculus λ^\forall with an algebraic datatype of lists, as well as primitives for general recursion and selective strictness that can be applied at all types. As in λ^\forall and PolyPCF, and by contrast with Clean and Haskell, all typing is explicit in the syntax of terms. That is, lambda-bound variables always come with an attached type, and with regard to polymorphism, both type generalization and specialization are made explicit. Another, purely syntactic, difference from Clean and Haskell is that there is only one mechanism to perform pattern matching, namely by case expressions. And rather than using recursive function equations, recursion is made explicit using a fixpoint primitive in the standard way. As an example, consider the following ‘translation’ of the Haskell function $foldl'$ from the introduction:

$$\begin{aligned}
 foldl' = \Lambda\alpha.\Lambda\beta.\mathbf{fix}(\lambda g :: (\beta \rightarrow \alpha \rightarrow \beta) \rightarrow \beta \rightarrow \alpha\text{-list} \rightarrow \beta.\lambda f :: \beta \rightarrow \alpha \rightarrow \beta.\lambda z :: \beta. \\
 \lambda l :: \alpha\text{-list}.\mathbf{case}\ l\ \mathbf{of}\ \{\mathbf{nil} \Rightarrow z; \\
 h : t \Rightarrow \mathbf{seq}(\underline{f\ z\ h}, g\ f\ (\underline{f\ z\ h}\ t))\})
 \end{aligned}$$

Note that the sharing of the two underlined expressions, which was present in the Haskell version via the binding ‘**where** $z' = f\ z\ h$ ’, is lost in PolySeq, which does not provide any sharing construct. But this difference has no impact at all on any

semantic properties we are going to study. For while implementations of Haskell typically apply a lazy evaluation strategy, the language definition [Pey03] only mandates that the semantics be nonstrict (apart from where selective strictness is used, of course), without committing to either call-by-name or call-by-need. Clearly, such a commitment is unnecessary, as it would have no impact on the observable behavior of programs. Indeed, choosing between call-by-name and call-by-need can only impact program properties regarding time and space usage, neither of which is under study here or specified in [Pey03]. Put differently, the semantics of *foldl'* in Haskell is invariant under replacing its second defining equation from the introduction by

$$\text{foldl}' f z (h : t) = \text{seq} (\underline{f z h}) (\text{foldl}' f (\underline{f z h}) t),$$

so it makes no sense to complicate PolySeq by modeling a sharing construct.

3.1 Syntax and typing

The syntax of PolySeq types and terms is given in Figure 2, where α and x range over disjoint countably infinite sets of *type variables* and *term variables*, respectively. The only difference to Figure 1 in [Pit00], apart from the slightly different notation, is the addition of the new term former **seq**. Alongside τ and M , we also let σ and $A, B, C, F, G, H, L, N, R, T$, and V range over the syntactic categories of types and terms, respectively. Moreover, β and c, f, g, h, l, n, t , and y are used as additional type and term variables, respectively, and all the mentioned conventions apply to versions with indices or primes as well. To reduce the need for brackets, function types and function applications are read right- and left-associative, respectively, so that $\tau_1 \rightarrow \tau_2 \rightarrow \tau_3$ means $\tau_1 \rightarrow (\tau_2 \rightarrow \tau_3)$, while $F A B$ means $(F A) B$. The constructions $\forall\alpha.-$, $\lambda x :: \tau.-$, $\Lambda\alpha.-$, and **case** M **of** $\{\mathbf{nil} \Rightarrow M'; x : x' \Rightarrow -\}$ are binders for α , x , and x' . We identify types and terms up to renaming of bound (type and term) variables. The concept of free variables in a type or term is defined in the usual way. For example, α is a free variable in \mathbf{nil}_α , but not in $\Lambda\alpha.\mathbf{nil}_\alpha$. We write *Typ* for the set of closed types, that is, those having no free variables. The result of capture-avoiding substitution of a type τ' for all free occurrences of a type variable α in a type τ or a term M is denoted by $\tau[\tau'/\alpha]$ or $M[\tau'/\alpha]$, respectively. Similarly, $M[M'/x]$ denotes the result of capture-avoiding substitution of a term M' for all free occurrences of a term variable x in a term M . Additionally, we use substitution for lists of distinct variables (e.g., $M_2[H/h, T/t]$ and $\tau[\vec{\tau}/\vec{\alpha}]$) and substitution for type and term variables at once (e.g., $M[\vec{\sigma}/\vec{\alpha}, \vec{N}/\vec{x}]$).

Types are assigned to (some) terms according to the axioms and rules in Figure 3, where Γ ranges over *typing environments* of the form $\vec{\alpha}, x_1 :: \tau_1, \dots, x_m :: \tau_m$ for a finite list $\vec{\alpha}$ of distinct type variables, $m \in \mathbb{N}$, a list $\vec{x} = x_1, \dots, x_m$ of distinct term variables, and types τ_1, \dots, τ_m whose free variables are in $\vec{\alpha}$. The only difference of note to Figure 2 in [Pit00] is the addition of the typing rule for **seq**. In a typing judgment of the form $\Gamma \vdash M :: \tau$, with Γ as above, we require that M 's free variables are in $\vec{\alpha}, \vec{x}$ and that τ 's free variables are in $\vec{\alpha}$. As in [Pit00], the explicit type information in the syntax of function abstractions and empty lists ensures that

Types	$\tau ::= \alpha$	type variable
	$\tau \rightarrow \tau$	function type
	$\forall\alpha.\tau$	\forall -type
	$\tau\text{-list}$	list type
Terms	$M ::= x$	term variable
	$\lambda x :: \tau.M$	function abstraction
	$M M$	function application
	$\Lambda\alpha.M$	type generalization
	M_τ	type specialization
	\mathbf{nil}_τ	empty list
	$M : M$	non-empty list
	$\mathbf{case} M \mathbf{of} \{\mathbf{nil} \Rightarrow M; x : x \Rightarrow M\}$	case expression
	$\mathbf{fix}(M)$	fixpoint recursion
	$\mathbf{seq}(M, M)$	strictness primitive

Figure 2: Syntax of the PolySeq calculus.

for every Γ and M there is at most one τ with $\Gamma \vdash M :: \tau$. Given $\tau \in Typ$, we write $Term(\tau)$ for the set of terms M for which $\emptyset \vdash M :: \tau$ is derivable, where \emptyset is the empty typing environment. Further, we set $Term = \bigcup_{\tau \in Typ} Term(\tau)$.

3.2 Operational semantics

Our semantics for PolySeq follows Plotkin's style of structural operational semantics [Plo04]. In particular, and in contrast to [Pit00], we start from a small-step rather than from a big-step formulation. As pointed out below, the two approaches are equivalent in a precise sense, and indeed [Pit00] also makes essential use of a small-step semantics indirectly in the proof of Theorem 3.6 and (thus) in the structural termination relation \top . Working directly with a small-step semantics avoids such indirection, and it is interesting to see that an approach without any big-step overlay is itself sufficient for constructing the desired parametric model.

PolySeq *values* are given by the following grammar:

$$V ::= \lambda x :: \tau.M \mid \Lambda\alpha.M \mid \mathbf{nil}_\tau \mid M : M.$$

Note that adding **seq** to a calculus does not introduce any new values. The subset of $Term$ consisting of all elements that respect the above grammar is denoted by $Value$. Further, given $\tau \in Typ$, we set $Value(\tau) = Value \cap Term(\tau)$.

The remaining ingredients for setting up a small-step semantics are redex/reduct-pairs and a notion of reduction in context. The former are just as in the proof of Theorem 3.6 in [Pit00], except that an appropriate pair involving **seq** is added.

$$\begin{array}{c}
 \Gamma, x :: \tau \vdash x :: \tau \\
 \\
 \frac{\Gamma, x :: \tau \vdash M :: \tau'}{\Gamma \vdash (\lambda x :: \tau. M) :: \tau \rightarrow \tau'} \quad \frac{\Gamma \vdash F :: \tau \rightarrow \tau' \quad \Gamma \vdash A :: \tau}{\Gamma \vdash F A :: \tau'} \\
 \\
 \frac{\alpha, \Gamma \vdash M :: \tau}{\Gamma \vdash \Lambda \alpha. M :: \forall \alpha. \tau} \quad \frac{\Gamma \vdash G :: \forall \alpha. \tau}{\Gamma \vdash G_{\tau'} :: \tau[\tau'/\alpha]} \\
 \\
 \Gamma \vdash \mathbf{nil}_{\tau} :: \tau\text{-list} \quad \frac{\Gamma \vdash H :: \tau \quad \Gamma \vdash T :: \tau\text{-list}}{\Gamma \vdash (H : T) :: \tau\text{-list}} \\
 \\
 \frac{\Gamma \vdash L :: \tau\text{-list} \quad \Gamma \vdash M_1 :: \tau' \quad \Gamma, h :: \tau, t :: \tau\text{-list} \vdash M_2 :: \tau'}{\Gamma \vdash \mathbf{case } L \mathbf{ of } \{ \mathbf{nil} \Rightarrow M_1; h : t \Rightarrow M_2 \} :: \tau'} \\
 \\
 \frac{\Gamma \vdash F :: \tau \rightarrow \tau}{\Gamma \vdash \mathbf{fix}(F) :: \tau} \quad \frac{\Gamma \vdash A :: \tau \quad \Gamma \vdash B :: \tau'}{\Gamma \vdash \mathbf{seq}(A, B) :: \tau'}
 \end{array}$$

Figure 3: PolySeq type assignment relation.

Definition 3.1. Let $\tau \in Typ$ and $R, R' \in Term(\tau)$. We write $R \rightsquigarrow R'$ for the following pairs:

R	R'	if
$(\lambda x :: \tau'. N) A$	$N[A/x]$	$x :: \tau' \vdash N :: \tau$
$(\Lambda \alpha. N)_{\tau'}$	$N[\tau'/\alpha]$	$\alpha \vdash N :: \tau''$
$\mathbf{case } \mathbf{nil}_{\tau'} \mathbf{ of } \{ \mathbf{nil} \Rightarrow M; h : t \Rightarrow M' \}$	M	$h :: \tau', t :: \tau'\text{-list} \vdash M' :: \tau$
$\mathbf{case } H : T \mathbf{ of } \{ \mathbf{nil} \Rightarrow M; h : t \Rightarrow M' \}$	$M'[H/h, T/t]$	$h :: \tau', t :: \tau'\text{-list} \vdash M' :: \tau$
$\mathbf{fix}(F)$	$F \mathbf{fix}(F)$	
$\mathbf{seq}(V, M)$	M	$V \in Value,$

where x, h , and t are term variables, α is a type variable, $\tau' \in Typ$, $A, H \in Term(\tau')$, $M \in Term(\tau)$, $T \in Term(\tau'\text{-list})$, $F \in Term(\tau \rightarrow \tau)$, and the further types and terms that occur in the table are subject to the restrictions recorded on the right. \diamond

It is essential that V is a value in the last pair above, because otherwise one would not ensure the intended semantics of **seq**, which is to first evaluate its first argument before reducing to the second. Notice also that type instantiation requires an evaluation step in PolySeq, whereas in Haskell it does not.

To describe reduction in context, we use the machinery of evaluation contexts introduced in [FFKD87]. Following [HS97], we represent these contexts as stacks of evaluation frames. Compared to [Pit00], an appropriate additional kind of evaluation frame is introduced to account for **seq**.

Definition 3.2. The grammar for *evaluation frame stacks* is

$$S ::= Id \mid S \circ E,$$

where E ranges over *evaluation frames*:

$$E ::= (- M) \mid -_\tau \mid (\mathbf{case} - \mathbf{of} \{\mathbf{nil} \Rightarrow M; x : x \Rightarrow M\}) \mid \mathbf{seq}(-, M).$$

If a stack comprises a single evaluation frame E , then we denote it by E rather than $Id \circ E$. Moreover, given an evaluation frame E and a term M , we write $E\{M\}$ for the term that results from replacing ‘ $-$ ’ by M in E . \diamond

Argument and result types are assigned to (some) evaluation frame stacks according to the axiom and rules in Figure 4, where Γ again ranges over typing environments, with well-formedness conditions similar to those for term typing judgments. The only difference of note to Figure 6 in [Pit00] is the addition of the typing rule for the new evaluation frame. As in [Pit00], for every Γ , S , and τ there is at most one τ' with $\Gamma \vdash S :: \tau \multimap \tau'$; this satisfies all needs for type uniqueness we will encounter. Given $\tau, \tau' \in Typ$, we write $Stack(\tau, \tau')$ for the set of evaluation frame stacks S for which $\emptyset \vdash S :: \tau \multimap \tau'$ is derivable. Since we will later want to restrict our attention to evaluation frame stacks which return results of list type, we set, for every $\tau \in Typ$, $LStack(\tau) = \bigcup_{\tau' \in Typ} Stack(\tau, \tau'-list)$.

$$\begin{array}{c} \Gamma \vdash Id :: \tau \multimap \tau \\ \\ \frac{\Gamma \vdash S :: \tau' \multimap \tau'' \quad \Gamma \vdash A :: \tau}{\Gamma \vdash S \circ (- A) :: (\tau \rightarrow \tau') \multimap \tau''} \quad \frac{\Gamma \vdash S :: \tau[\tau'/\alpha] \multimap \tau''}{\Gamma \vdash S \circ -_{\tau'} :: (\forall \alpha. \tau) \multimap \tau''} \\ \\ \frac{\Gamma \vdash S :: \tau' \multimap \tau'' \quad \Gamma \vdash M_1 :: \tau' \quad \Gamma, h :: \tau, t :: \tau-list \vdash M_2 :: \tau'}{\Gamma \vdash S \circ (\mathbf{case} - \mathbf{of} \{\mathbf{nil} \Rightarrow M_1; h : t \Rightarrow M_2\}) :: \tau-list \multimap \tau''} \\ \\ \frac{\Gamma \vdash S :: \tau' \multimap \tau'' \quad \Gamma \vdash B :: \tau'}{\Gamma \vdash S \circ \mathbf{seq}(-, B) :: \tau \multimap \tau''} \end{array}$$

Figure 4: Typing evaluation frame stacks.

The (typed) operations of concatenating two evaluation frame stacks and of applying an evaluation frame stack to a term are given as follows.

Definition 3.3. Let $\tau \in Typ$. Given $S \in LStack(\tau)$, we define for every $\tau' \in Typ$ and $S' \in Stack(\tau', \tau)$ the concatenation $(S @ S') \in LStack(\tau')$ by induction on the structure of S' as follows:

$$\begin{aligned} S @ Id &= S \\ S @ (S'' \circ E) &= (S @ S'') \circ E. \end{aligned}$$

Moreover, we define for every $\tau' \in Typ$, $S \in Stack(\tau', \tau)$, and $M \in Term(\tau')$ the application $(S M) \in Term(\tau)$ by induction on the structure of S as follows:

$$\begin{aligned} Id M &= M \\ (S' \circ E) M &= S' (E\{M\}). \end{aligned} \quad \diamond$$

The transition relation induced by the choice of values, redex/reduct-pairs, and evaluation frames is defined in the following (standard) way.

Definition 3.4. Let $\tau_1, \tau_2, \tau' \in Typ$, $S_1 \in Stack(\tau_1, \tau')$, $M_1 \in Term(\tau_1)$, $S_2 \in Stack(\tau_2, \tau')$, and $M_2 \in Term(\tau_2)$. We write $(S_1, M_1) \mapsto (S_2, M_2)$ for the following pairs:

$$\frac{(S_1, M_1) \quad (S_2, M_2)}{(S, E\{N\}) \quad (S \circ E, N)} \Big| \text{if } N \notin Value$$

$$\frac{(S \circ E, V) \quad (S, E\{V\})}{(S, R) \quad (S, R')} \Big| V \in Value$$

$$(S, R) \quad (S, R') \Big| R \rightsquigarrow R',$$

where S is an evaluation frame stack, E is an evaluation frame, and the terms that occur in the table are subject to the restrictions recorded on the right. \diamond

Intuitively, the first two transition rules navigate a term to detect the next redex to be reduced, while the third rule performs a small-step reduction in a given evaluation context. Note that \mapsto is deterministic, but not terminating (due to **fix**). If we denote by \mapsto^* the reflexive, transitive closure of \mapsto , then evaluation of a term to a value can be captured as follows.

Definition 3.5. Given $M \in Term$ and $V \in Value$ of the same type, we write $M \Downarrow V$ if $(Id, M) \mapsto^* (Id, V)$. Given $M \in Term$, we write $M \Downarrow$ if there is some V with $M \Downarrow V$, and $M \Uparrow$ otherwise. In the former case we say that M *converges*, and in the latter we say that it *diverges*. Note that every value converges. \diamond

Thus, we have provided a stack-based abstract machine for PolySeq. The evaluation relation \Downarrow induced by this machine is the same as the one we would obtain via defining a big-step semantics by adding the rule

$$\frac{A \Downarrow V \quad B \Downarrow V'}{\mathbf{seq}(A, B) \Downarrow V'}$$

to Figure 3 in [Pit00]. The proof of this fact is very similar to that of (3) inside the proof of Theorem 3.6 in [Pit00]; we do not give it here. But since the small-step semantics is a bit more low-level than the big-step semantics, the small-step semantics more immediately reflects the operational behavior of **seq** in actual Haskell implementations. This is because the evaluation frame $\mathbf{seq}(-, M)$ and the reduction $\mathbf{seq}(V, M) \rightsquigarrow M$ make it explicit that **seq** first evaluates its first argument, before turning to the second one, while no such order is imposed in the above big-step rule.

Before moving on to the intended notion of operational approximation, we give three observations about termination issues and the existence of a ‘polymorphic bottom’. The first two observations follow easily from the definitions of \Downarrow and \mapsto , the third one arises by combination of the first two.

Observation 3.6. For every $\tau \in Typ$: $\mathbf{fix}(\lambda x :: \tau.x) \Uparrow$. ■

Observation 3.7. For every $R, R' \in Term$, if $R \rightsquigarrow R'$, then $R \Downarrow \Leftrightarrow R' \Downarrow$. ■

Observation 3.8. Let $\Omega = \Lambda\alpha.\mathbf{fix}(\lambda x :: \alpha.x) \in \mathit{Term}(\forall\alpha.\alpha)$. While $\Omega \Downarrow$, for every $\tau \in \mathit{Typ}$: $\Omega_\tau \Uparrow$. \blacksquare

Being able to describe to what value, if any, a term evaluates is usually not enough to reason about the operational behavior of a programming language. In particular, stipulating that two terms are to be considered equivalent (if and) only if both evaluate to the same value is a much too strong requirement. It would outlaw, for example, consideration of quicksort and heapsort implementations as semantically equivalent, on the grounds that their representations as function abstractions would necessarily be different. Such a situation would be highly undesirable, given that two different algorithms performing the same computational task (such as sorting a list) should clearly be equated by any reasonable semantics. For this reason, it is standard to allow only *particular observations* to be made about terms in the language (with comparison of function abstraction representations not being one of them), but to require that two terms are considered equivalent (if and) only if they lead to the same observations *in every possible context*.

Following [Pit00], we choose evaluation of terms of list type to the empty list as the only observation possible for PolySeq. Actually, due to the presence of **seq**, this allows the observation of termination at arbitrary types, in a sense later made precise in Corollary 4.16. So in contrast to the situation in the setting without **seq**, the initial choice to observe termination only at list types has no impact on the results here. The important point is that we can still observe only termination at arbitrary types, and not the full representation of any obtained value. To capture observations in context, we again follow the treatment in [Pit00]. That is, we specify a number of desirable properties our notion of semantic approximation for PolySeq should have, and then ask for the largest relation with those properties. Of course, reasoning about semantic approximation only makes sense if we have at least a preorder, i.e., a relation that is reflexive and transitive. But in addition to that, the intended notion of approximation should also be a congruence, i.e., should be compatible with all term formers and with substitution, in a sense made precise below. This is where the ‘context closure’ of observational approximation is ensured. To tie in the possible observations themselves, we impose an adequacy property, reflecting the above discussion about ‘**nil**-termination’. But since we are interested in approximation rather than equivalence, we trade the bidirectional implication in the conclusion of the definition of adequacy in [Pit00] for unidirectional implication in the first part of the next definition.

Definition 3.9. Let the relation \mathcal{E} comprise 4-tuples of the form (Γ, M, M', τ) with $\Gamma \vdash M :: \tau$ and $\Gamma \vdash M' :: \tau$. We write $\Gamma \vdash M \mathcal{E} M' :: \tau$ when the tuple (Γ, M, M', τ) is in \mathcal{E} , and we abbreviate this to $M \mathcal{E} M'$ if $\Gamma = \emptyset$ since τ is then uniquely determined as the closed type of both M and M' .

1. \mathcal{E} is *adequate* if for every $\tau \in \mathit{Typ}$ and $L, L' \in \mathit{Term}(\tau\text{-list})$:

$$L \mathcal{E} L' \Rightarrow (L \Downarrow \mathbf{nil}_\tau \Rightarrow L' \Downarrow \mathbf{nil}_\tau).$$

2. \mathcal{E} is *compatible* if it is closed under the axioms and rules in Figure 5, which differs from Figure 4 in [Pit00] only by the new rule for **seq**.
3. \mathcal{E} is *substitutive* if it is closed under the rules in Figure 6, where $\Gamma[\tau'/\alpha]$ is the typing environment obtained from Γ by replacing every $x :: \sigma$ therein by $x :: \sigma[\tau'/\alpha]$.
4. \mathcal{E} is *reflexive* if for every environment Γ , term M , and type τ with $\Gamma \vdash M :: \tau$:

$$\Gamma \vdash M \mathcal{E} M :: \tau.$$

5. \mathcal{E} is *transitive* if $\mathcal{E}; \mathcal{E} \subseteq \mathcal{E}$, where *relation composition* $\mathcal{E}_1; \mathcal{E}_2$ is defined by:

$$\Gamma \vdash M (\mathcal{E}_1; \mathcal{E}_2) M' :: \tau \Leftrightarrow \exists M''. \Gamma \vdash M \mathcal{E}_1 M'' :: \tau \wedge \Gamma \vdash M'' \mathcal{E}_2 M' :: \tau.$$

◇

$$\begin{array}{c}
 \Gamma, x :: \tau \vdash x \mathcal{E} x :: \tau \\
 \\
 \frac{\Gamma, x :: \tau \vdash M \mathcal{E} M' :: \tau'}{\Gamma \vdash (\lambda x :: \tau. M) \mathcal{E} (\lambda x :: \tau. M') :: \tau \rightarrow \tau'} \\
 \\
 \frac{\Gamma \vdash F \mathcal{E} F' :: \tau \rightarrow \tau' \quad \Gamma \vdash A \mathcal{E} A' :: \tau}{\Gamma \vdash (F A) \mathcal{E} (F' A') :: \tau'} \\
 \\
 \frac{\alpha, \Gamma \vdash M \mathcal{E} M' :: \tau}{\Gamma \vdash \Lambda \alpha. M \mathcal{E} \Lambda \alpha. M' :: \forall \alpha. \tau} \quad \frac{\Gamma \vdash G \mathcal{E} G' :: \forall \alpha. \tau}{\Gamma \vdash G_{\tau'} \mathcal{E} G'_{\tau'} :: \tau[\tau'/\alpha]} \\
 \\
 \Gamma \vdash \mathbf{nil}_\tau \mathcal{E} \mathbf{nil}_\tau :: \tau\text{-list} \quad \frac{\Gamma \vdash H \mathcal{E} H' :: \tau \quad \Gamma \vdash T \mathcal{E} T' :: \tau\text{-list}}{\Gamma \vdash (H : T) \mathcal{E} (H' : T') :: \tau\text{-list}} \\
 \\
 \frac{\Gamma \vdash L \mathcal{E} L' :: \tau\text{-list} \quad \Gamma \vdash M_1 \mathcal{E} M'_1 :: \tau' \quad \Gamma, h :: \tau, t :: \tau\text{-list} \vdash M_2 \mathcal{E} M'_2 :: \tau'}{\Gamma \vdash (\mathbf{case } L \mathbf{ of } \{ \mathbf{nil} \Rightarrow M_1; h : t \Rightarrow M_2 \}) \mathcal{E} (\mathbf{case } L' \mathbf{ of } \{ \mathbf{nil} \Rightarrow M'_1; h : t \Rightarrow M'_2 \}) :: \tau'} \\
 \\
 \frac{\Gamma \vdash F \mathcal{E} F' :: \tau \rightarrow \tau}{\Gamma \vdash \mathbf{fix}(F) \mathcal{E} \mathbf{fix}(F') :: \tau} \quad \frac{\Gamma \vdash A \mathcal{E} A' :: \tau \quad \Gamma \vdash B \mathcal{E} B' :: \tau'}{\Gamma \vdash \mathbf{seq}(A, B) \mathcal{E} \mathbf{seq}(A', B') :: \tau'}
 \end{array}$$

Figure 5: Compatibility properties.

It is easy to see that every compatible relation \mathcal{E} is also reflexive. Our intended notion of ‘contextual’ approximation is the largest relation satisfying all five properties from Definition 3.9. We call it *observational approximation* and write it as \sqsubseteq_{obs} . The existence of such a largest relation is, however, only stipulated for now. Using techniques of [Las98], existence could be proved in a direct manner here by

$$\frac{\alpha, \Gamma \vdash M \mathcal{E} M' :: \tau}{\Gamma[\tau'/\alpha] \vdash M[\tau'/\alpha] \mathcal{E} M'[\tau'/\alpha] :: \tau[\tau'/\alpha]}$$

$$\frac{\Gamma, x :: \tau \vdash M \mathcal{E} M' :: \tau' \quad \Gamma \vdash N \mathcal{E} N' :: \tau}{\Gamma \vdash M[N/x] \mathcal{E} M'[N'/x] :: \tau'}$$

Figure 6: Substitutivity properties.

characterizing \sqsubseteq_{obs} as the union of all adequate and compatible (and, therefore, substitutive) relations, from which coincidence of observational approximation and ‘contextual’ approximation follows straightforwardly. But since we are interested in a more constructive and ultimately more useful characterization, we defer the proof of the existence of \sqsubseteq_{obs} to Theorem 6.7, where it is characterized by a relation inductively derived based on the type structure of PolySeq. Of course, this means that up to that point in Section 6 we may not assume anything about \sqsubseteq_{obs} , and we will not do so. When we finally have returned to \sqsubseteq_{obs} , we will often use its reflexivity and transitivity without explicit mention.

4 PolySeq termination

In this section we study various aspects of termination in PolySeq. These are essential for characterizing observational approximation in the presence of **fix** and **seq** in the way we aim to do. Regarding **fix**, it has long been known [Wad89, LP96] that parametricity can only be achieved by restricting attention to relations that are admissible in a sense corresponding to the concepts of strictness and continuity in denotational semantics. The main technical contribution of [Pit00] was an account of such admissibility in an equational operational setting, based on a closure operator arising from **nil**-termination. As we will demonstrate, neither moving to an inequational setting nor adding **seq** breaks any of that machinery in principle. Indeed, the relevant Lemma 4.14 below can still be established in an inequational setting — whether or not **seq** is present — when adopting a directed counterpart to Pitts’ closure operator. But adding **seq** requires more. The reason is that with a strictness primitive that can be applied at all types, a new restriction must be imposed on relational interpretations of types. This was already observed in [LP96] and [PLST98] for the equational setting, and was both made more rigorous and extended to an inequational setting in [JV04, JV06], but all in denotational semantics only. Here we present (in Definition 4.17) an operational counterpart to the relevant new restriction on relations imposed in [JV04, JV06], show how it interacts with the aforementioned operational machinery for fixpoint admissibility (Lemma 4.18), and show how it guarantees the key property needed to achieve parametricity in the presence of **seq** (Lemma 4.19).

4.1 General properties of nil-termination

First, we fix a special notation for expressing that evaluating a particular term in a particular context described by an evaluation frame stack leads to the empty list.

Definition 4.1. Let $\tau, \tau' \in Typ$, $S \in Stack(\tau, \tau'\text{-list})$, and $M \in Term(\tau)$. We write $S \top M$ if $(S, M) \rightsquigarrow^* (Id, \mathbf{nil}_{\tau'})$. \diamond

Note that, rather than taking the described behavior as *definition* of \top , Pitts defines \top via a syntactic system of structural rules (see Figure 7 of [Pit00]) and then *proves* that $S \top M$ if and only if $(S, M) \rightsquigarrow^* (Id, \mathbf{nil}_{\tau'})$ inside his Theorem 3.6. Our approach is simpler in that it obviates the need for an extra set of syntactic rules and an attendant proof. The key point is, of course, that we still get all the properties of \top that we need. In particular, the structural properties present in Figure 7 of [Pit00], plus corresponding ones having to do with **seq**, are all embodied in Observations 4.2 and 4.3 below. Moreover, since these follow generically from Definitions 3.4, 3.5, and 4.1 (and from determinism of \rightsquigarrow), without considering the concrete sets of redex/reduct pairs and evaluation frames at hand, our approach promises to be more amenable to (further) extensions of the calculus.

Observation 4.2. For every $\tau \in Typ$, $L \in Term(\tau\text{-list})$: $Id \top L \Leftrightarrow L \Downarrow \mathbf{nil}_{\tau}$. \blacksquare

Observation 4.3. Let $\tau \in Typ$ and $S \in LStack(\tau)$.

1. For every $\tau' \in Typ$, $M \in Term(\tau')$, and evaluation frame E with $E\{M\} \in Term(\tau)$: $S \top E\{M\} \Leftrightarrow S \circ E \top M$.
2. For every $R, R' \in Term(\tau)$ with $R \rightsquigarrow R'$: $S \top R \Leftrightarrow S \top R'$. \blacksquare

By repeated applications of Observation 4.3(1), we also have the following corollary.

Corollary 4.4. For every $\tau, \tau' \in Typ$, $S \in Stack(\tau, \tau')$, $S' \in LStack(\tau')$, and $M \in Term(\tau)$:

$$(S' @ S) \top M \Leftrightarrow S' \top S M. \quad \blacksquare$$

The following lemma shows that **nil**-termination is respected in a certain sense by evaluation of the term put in context.

Lemma 4.5. For every $\tau \in Typ$, $S \in LStack(\tau)$, and $M \in Term(\tau)$:

$$S \top M \Leftrightarrow \exists V \in Value(\tau). M \Downarrow V \wedge S \top V. \quad \blacksquare$$

The proof, which proceeds by two inductions over \rightsquigarrow^* -sequences, is given in [VJ06]. An immediate consequence is the following ‘strictness of stacks’ result.

Corollary 4.6. For every $\tau \in Typ$ and $M \in Term(\tau)$, if $M \Uparrow$, then for every $S \in LStack(\tau)$, $S \top M$ does not hold. \blacksquare

4.2 Termination for **fix** and $\top\top$ -closedness

The key role of \top in [Pit00] is its use in defining an order-reversing Galois connection, and the use of the induced closure operator in characterizing a class of relations that admit a form of fixpoint induction. A similar construction, replacing bidirectional implication by unidirectional implication, is repeated in the following three definitions.

Definition 4.7. Given $\tau, \tau' \in Typ$, we define

$$Rel(\tau, \tau') = \mathcal{P}(Term(\tau) \times Term(\tau'))$$

and

$$StRel(\tau, \tau') = \mathcal{P}(LStack(\tau) \times LStack(\tau')).$$

Further, we set $Rel = \bigcup_{\tau, \tau' \in Typ} Rel(\tau, \tau')$. \diamond

Definition 4.8. Let $\tau, \tau' \in Typ$. Given $r \in Rel(\tau, \tau')$, we define $r^\top \in StRel(\tau, \tau')$ by

$$(S, S') \in r^\top \text{ iff } \forall(M, M') \in r. S \top M \Rightarrow S' \top M'.$$

Similarly, given $s \in StRel(\tau, \tau')$, we define $s^\top \in Rel(\tau, \tau')$ by

$$(M, M') \in s^\top \text{ iff } \forall(S, S') \in s. S \top M \Rightarrow S' \top M'. \quad \diamond$$

The following properties are standard, for every $\tau, \tau' \in Typ$ and $r, r_1, r_2 \in Rel(\tau, \tau')$:

$$r \subseteq r^{\top\top} \quad (2)$$

$$(r^{\top\top})^\top = r^\top \quad (3)$$

$$r_1 \subseteq r_2 \Rightarrow r_1^{\top\top} \subseteq r_2^{\top\top}. \quad (4)$$

Definition 4.9. A relation $r \in Rel$ is $\top\top$ -closed if $r^{\top\top} = r$. \diamond

Note that by (2) and (3), respectively, the condition $r^{\top\top} = r$ is equivalent to $r^{\top\top} \subseteq r$ and to the existence of some $r' \in Rel$ with $r = (r')^{\top\top}$.

Several important properties of $\top\top$ -closed relations can be established now. The first is that they respect adequate and compatible relations. The second is that they relate two appropriately typed terms if the first of them is diverging.

Lemma 4.10. Let \mathcal{E} be an adequate and compatible relation, let $\tau, \tau' \in Typ$, and $r \in Rel(\tau, \tau')$. If r is $\top\top$ -closed, then for every $M_1 \in Term(\tau)$ and $M_2, M_3 \in Term(\tau')$:

$$(M_1, M_2) \in r \wedge M_2 \mathcal{E} M_3 \Rightarrow (M_1, M_3) \in r.$$

Proof: The desired $(M_1, M_3) \in r$ follows from $\top\top$ -closedness of r and the following reasoning for every $(S, S') \in r^\top$:

$$\begin{aligned}
 S \top M_1 &\Rightarrow S' \top M_2 && \text{by } (S, S') \in r^\top \text{ and } (M_1, M_2) \in r \\
 &\Leftrightarrow Id \top S' M_2 && \text{by Corollary 4.4} \\
 &\Leftrightarrow S' M_2 \Downarrow \mathbf{nil}_{\tau''} && \text{by Observation 4.2} \\
 &\Rightarrow S' M_3 \Downarrow \mathbf{nil}_{\tau''} \\
 &\Leftrightarrow S' \top M_3 && \text{by Observation 4.2 and Corollary 4.4.}
 \end{aligned}$$

Here $\tau'' \in Typ$ is such that $S' \in Stack(\tau', \tau''\text{-list})$, and the second implication follows from $M_2 \mathcal{E} M_3$, because \mathcal{E} is compatible and adequate. \blacksquare

Lemma 4.11. For every $\tau, \tau' \in Typ$, $M \in Term(\tau)$, $M' \in Term(\tau')$, and $r \in Rel(\tau, \tau')$, if $M \uparrow$ and r is $\top\top$ -closed, then $(M, M') \in r$.

Proof: By Corollary 4.6, $M \uparrow$ implies that for every $(S, S') \in r^\top$: $S \top M \Rightarrow S' \top M'$. Thus, we have $(M, M') \in r^{\top\top} = r$. \blacksquare

The previous lemma provides a kind of base case for fixpoint induction. To establish this principle in full, we first need to look at the finite unwindings of a fixpoint.

Definition 4.12. Let $\tau \in Typ$ and $F \in Term(\tau \rightarrow \tau)$. By induction on $n \in \mathbb{N}$, we define $\mathbf{fix}^{(n)}(F) \in Term(\tau)$ as follows:

$$\mathbf{fix}^{(n)}(F) = \begin{cases} \mathbf{fix}(\lambda x :: \tau.x) & \text{if } n = 0 \\ F \mathbf{fix}^{(n-1)}(F) & \text{otherwise.} \end{cases} \quad \diamond$$

Then the following important termination property holds for \mathbf{fix} and its unwindings. The proof, which uses Observations 3.6 and 3.7 and Corollary 4.6, can be found in [VJ06].

Lemma 4.13. For every $\tau \in Typ$, $S \in LStack(\tau)$, and $F \in Term(\tau \rightarrow \tau)$:

$$S \top \mathbf{fix}(F) \Leftrightarrow \exists n \in \mathbb{N}. S \top \mathbf{fix}^{(n)}(F). \quad \blacksquare$$

This ‘unwinding lemma’ can now be used to establish a fixpoint induction principle for binary relations. It justifies the claim that $\top\top$ -closed relations have appropriate admissibility properties, and indeed corresponds to what is identified as the necessary parametricity property of \mathbf{fix} in Sections 7 and 5 of [Wad89] and [LP96], respectively. Its proof is essentially the same as that of Theorem 3.11 in [Pit00], and uses Observation 3.6 and Lemmas 4.11 and 4.13.

Lemma 4.14. Let $\tau, \tau' \in Typ$, $F \in Term(\tau \rightarrow \tau)$, $F' \in Term(\tau' \rightarrow \tau')$, and $r \in Rel(\tau, \tau')$. If

$$\forall (A, A') \in r. (F A, F' A') \in r$$

and r is $\top\top$ -closed, then $(\mathbf{fix}(F), \mathbf{fix}(F')) \in r$. \blacksquare

4.3 Termination for **seq** and convergence-preservation

To handle the strictness primitive, we ultimately need an analogue of Lemma 4.14 for **seq**. In Sections 5 of [JV04] and [JV06], the relevant parametricity property of **seq** was identified as saying that for appropriate relations r_1 and r_2 , terms A and A' related by r_1 , and terms B and B' related by r_2 , we should have that $\mathbf{seq}(A, B)$ and $\mathbf{seq}(A', B')$ are related by r_2 . It turns out that $\top\top$ -closedness is not a strong enough restriction on r_1 and r_2 to achieve this, just as strictness and continuity were not enough in the denotational setting. To see why, we first need the following key statement, playing a similar role for **seq** as Lemma 4.13 does for **fix**, as well as a corollary of this statement.

Lemma 4.15. Let $\tau \in Typ$, $S \in LStack(\tau)$, $A \in Term$, and $B \in Term(\tau)$. Then: $S \top \mathbf{seq}(A, B) \Leftrightarrow A \Downarrow \wedge S \top B$.

Proof: Let $\tau' \in Typ$ be such that $A \in Term(\tau')$. Then we reason as follows:

$$\begin{aligned}
 & S \top \mathbf{seq}(A, B) \\
 \Leftrightarrow & S \circ \mathbf{seq}(-, B) \top A && \text{by Observation 4.3(1)} \\
 \Leftrightarrow & \exists V \in Value(\tau'). A \Downarrow V \wedge S \circ \mathbf{seq}(-, B) \top V && \text{by Lemma 4.5} \\
 \Leftrightarrow & \exists V \in Value(\tau'). A \Downarrow V \wedge S \top B && \text{by Observation 4.3.} \quad \blacksquare
 \end{aligned}$$

Corollary 4.16. For every $M \in Term$ and $\tau \in Typ$: $M \Downarrow \Leftrightarrow \mathbf{seq}(-, \mathbf{nil}_\tau) \top M$.

Proof: We reason as follows:

$$\begin{aligned}
 M \Downarrow & \Leftrightarrow Id \top \mathbf{seq}(M, \mathbf{nil}_\tau) && \text{by Lemma 4.15 and } Id \top \mathbf{nil}_\tau \\
 & \Leftrightarrow \mathbf{seq}(-, \mathbf{nil}_\tau) \top M && \text{by Observation 4.3(1).} \quad \blacksquare
 \end{aligned}$$

The corollary essentially says that in the presence of **seq**, observing **nil**-termination of terms of list type suffices to observe general termination of arbitrary terms.

Now, we can give a counterexample to the parametricity property of **seq** mentioned at the start of this subsection. Let $r_1 = \{(\Omega, \Omega_{\forall\alpha.\alpha})\}^{\top\top}$ and $r_2 = \{(M, M) \mid M \in Term(\forall\alpha.\alpha)\}^{\top\top}$, both of which are in $Rel(\forall\alpha.\alpha, \forall\alpha.\alpha)$, and let $A' = \Omega_{\forall\alpha.\alpha}$ and $A = B = B' = \Omega$. Clearly, r_1 and r_2 are $\top\top$ -closed, and $(A, A') \in r_1$ and $(B, B') \in r_2$ hold by (2). So we would expect that $(\mathbf{seq}(A, B), \mathbf{seq}(A', B')) \in r_2$ also holds. But this does not hold, as is argued now. Note that for every $S \in LStack(\forall\alpha.\alpha)$ we have $(S, S) \in \{(M, M) \mid M \in Term(\forall\alpha.\alpha)\}^{\top}$. So to have $(\mathbf{seq}(\Omega, \Omega), \mathbf{seq}(\Omega_{\forall\alpha.\alpha}, \Omega)) \in \{(M, M) \mid M \in Term(\forall\alpha.\alpha)\}^{\top\top}$ would imply that for every such S we have $S \top \mathbf{seq}(\Omega, \Omega) \Rightarrow S \top \mathbf{seq}(\Omega_{\forall\alpha.\alpha}, \Omega)$. But this is contradicted by Observation 3.8, Lemma 4.15, and Corollary 4.16.

The reason for the failure described above is that the given r_1 relates a converging term to a diverging term. To repair this failure, we introduce a new restriction on relations, namely convergence-preservation. This restriction is a relatively direct translation of the totality restriction discussed in the denotational setting of [JV04, JV06]. In particular, in contrast to what was envisaged in those earlier papers, the

new restriction is not enforced by a variant of $\top\top$ -closure or by a completely new closure operator. But of course, the new restriction must at least nicely coexist with $\top\top$ -closure in a certain sense, to be discussed below the following definition.

Definition 4.17. We say that $r \in Rel$ is *convergence-preserving* if for every $(M, M') \in r$: $M \Downarrow \Rightarrow M' \Downarrow$. For given $\tau, \tau' \in Typ$, the restriction of $Rel(\tau, \tau')$ to convergence-preserving relations is denoted by $Rel^\Downarrow(\tau, \tau')$. We set $Rel^\Downarrow = \bigcup_{\tau, \tau' \in Typ} Rel^\Downarrow(\tau, \tau')$. \diamond

Note that the r_1 used in the counterexample above also provides evidence that not every $\top\top$ -closed relation is convergence-preserving. However, we can show that the notions of $\top\top$ -closure and convergence-preservation are compatible in the sense that the former preserves the latter. This is essential, since otherwise the modular way in which the two restrictions are introduced (and later preserved) would break down.

Lemma 4.18. For every $r \in Rel^\Downarrow$, $r^{\top\top} \in Rel^\Downarrow$ also holds.

Proof: For arbitrary $\tau \in Typ$ and every $(M, M') \in r$, we have:

$$\begin{aligned} \mathbf{seq}(-, \mathbf{nil}_\tau) \top M &\Leftrightarrow M \Downarrow && \text{by Corollary 4.16} \\ &\Rightarrow M' \Downarrow && \text{since } r \text{ is convergence-preserving} \\ &\Leftrightarrow \mathbf{seq}(-, \mathbf{nil}_\tau) \top M' && \text{by Corollary 4.16.} \end{aligned}$$

Thus, $(\mathbf{seq}(-, \mathbf{nil}_\tau), \mathbf{seq}(-, \mathbf{nil}_\tau)) \in r^\top$, and consequently for every $(N, N') \in r^{\top\top}$:

$$\mathbf{seq}(-, \mathbf{nil}_\tau) \top N \Rightarrow \mathbf{seq}(-, \mathbf{nil}_\tau) \top N',$$

which by Corollary 4.16, applied twice as above, is equivalent to $N \Downarrow \Rightarrow N' \Downarrow$. \blacksquare

A further connection between the concepts of convergence-preservation and $\top\top$ -closure is established in the following lemma, which provides the sought after parametricity property of \mathbf{seq} , and thus the analogue for \mathbf{seq} of Lemma 4.14.

Lemma 4.19. Let $r_1 \in Rel^\Downarrow$, $r_2 \in Rel$, $(A, A') \in r_1$, and $(B, B') \in r_2$. If r_2 is $\top\top$ -closed, then $(\mathbf{seq}(A, B), \mathbf{seq}(A', B')) \in r_2$.

Proof: The desired $(\mathbf{seq}(A, B), \mathbf{seq}(A', B')) \in r_2$ follows from $\top\top$ -closedness of r_2 and the following reasoning for every $(S, S') \in r_2^\top$:

$$\begin{aligned} S \top \mathbf{seq}(A, B) &\Leftrightarrow A \Downarrow \wedge S \top B && \text{by Lemma 4.15} \\ &\Rightarrow A' \Downarrow \wedge S \top B && \text{by } (A, A') \in r_1 \in Rel^\Downarrow \\ &\Rightarrow A' \Downarrow \wedge S' \top B' && \text{by } (S, S') \in r_2^\top \text{ and } (B, B') \in r_2 \\ &\Leftrightarrow S' \top \mathbf{seq}(A', B') && \text{by Lemma 4.15.} \end{aligned} \quad \blacksquare$$

Since in the previous lemma one relation must be convergence-preserving and the other one must be $\top\top$ -closed, we actually have to impose (and preserve) both restrictions on all relations in our type-based characterization of observational approximation. This is so because \mathbf{seq} can be applied at all types according to the

corresponding typing rule in Figure 3, in which no restrictions are imposed on τ or τ' . Hence, the relational interpretation of every type must be able to fulfill either of the two roles r_1 and r_2 in Lemma 4.19, and so must adhere to both restrictions. This is what will drive the development of the relational interpretation of PolySeq types in Section 5. But first we need to set up a few more auxiliary statements. The first is an analogue of Lemma 4.5 for general termination.

Lemma 4.20. For every $\tau, \tau' \in Typ$, $S \in Stack(\tau, \tau')$, and $M \in Term(\tau)$:

$$(S M)\Downarrow \Leftrightarrow \exists V \in Value(\tau). M \Downarrow V \wedge (S V)\Downarrow.$$

Proof: We reason as follows:

$$\begin{aligned} & (S M)\Downarrow \\ \Leftrightarrow & \mathbf{seq}(-, \mathbf{nil}_\tau) \top S M && \text{by Corollary 4.16} \\ \Leftrightarrow & (\mathbf{seq}(-, \mathbf{nil}_\tau) @ S) \top M && \text{by Corollary 4.4} \\ \Leftrightarrow & \exists V \in Value(\tau). M \Downarrow V \wedge (\mathbf{seq}(-, \mathbf{nil}_\tau) @ S) \top V && \text{by Lemma 4.5} \\ \Leftrightarrow & \exists V \in Value(\tau). M \Downarrow V \wedge (S V)\Downarrow, \end{aligned}$$

where the last equivalence is again by Corollaries 4.4 and 4.16. ■

Lemma 4.20 implies that for every $M \in Term$ and evaluation frame E with $E\{M\} \in Term$, if $E\{M\}\Downarrow$, then $M\Downarrow$. Letting $M = F$ and $E = (- A)$, we have the following corollary.

Corollary 4.21. For every $\tau, \tau' \in Typ$, $F \in Term(\tau \rightarrow \tau')$, and $A \in Term(\tau)$, if $(F A)\Downarrow$, then $F\Downarrow$. ■

5 The logical relation

The key to parametricity results, and to our characterization of observational approximation by a logical relation, is to build relational interpretations of types by induction on the type structure. Starting from an interpretation of type variables by relations (between terms), this requires defining a *relational action* for each of the ways of forming PolySeq types. Such an action takes an appropriate number of relations and produces a new one as the interpretation for the compound type. During this propagation of relations along the type structure it is essential that the restrictions needed to accommodate **fix** and **seq** — namely, $\top\top$ -closedness and convergence-preservation — are preserved. In our development convergence-preservation will always hold by construction, whereas preservation of $\top\top$ -closedness is established *a posteriori*.

The main characteristic of all logical relations, from the very beginning [Plo73, Fri75, Rey83, Sta85] up to newer accounts [Pit00, Pit05, Ahm06, DHJG06], is that for two functions to be related they must map related arguments to related results. In an inequational setting with **seq** present, the relational action for function types must additionally enforce the requirement that two function terms are only related

if convergence of the first implies convergence of second. This corresponds to the additional restriction in the relational action for function types in [JV04, JV06], transferred to the operational setting.

Definition 5.1. Given $\tau_1, \tau'_1, \tau_2, \tau'_2 \in Typ$, $r_1 \in Rel(\tau_1, \tau'_1)$, and $r_2 \in Rel(\tau_2, \tau'_2)$, we define $(r_1 \rightarrow r_2) \in Rel^\Downarrow(\tau_1 \rightarrow \tau_2, \tau'_1 \rightarrow \tau'_2)$ by

$$(F, F') \in (r_1 \rightarrow r_2) \text{ iff } (F \Downarrow \Rightarrow F' \Downarrow) \wedge \forall (A, A') \in r_1. (F A, F' A') \in r_2$$

for every $F \in Term(\tau_1 \rightarrow \tau_2)$ and $F' \in Term(\tau'_1 \rightarrow \tau'_2)$. \diamond

The relational action corresponding to \forall -types usually relates two polymorphic terms if respective instances, at arbitrary types, are related by the images, under a given relation-to-relation mapping, of certain relations between terms of the types at which instantiation occurs. The relation-to-relation mapping will be derived from the compound action of the body of the \forall -type on relations interpreting the type variable from its head. The range of relations over which the quantification occurs — i.e., the concretization of ‘certain’ above — depends on the primitives the calculus supports, at arbitrary types, in extension of the Girard-Reynolds calculus λ^\forall . As we have argued, **fix** and **seq** mandate relations to be $\top\top$ -closed and convergence-preserving. Since the later construction of the logical relation will ensure that the relation-to-relation mapping always first $\top\top$ -closes its argument relation (cf. clause (8) in Definition 5.4), only convergence-preservation needs to be explicitly enforced in the quantification of relations below. The aforementioned $\top\top$ -closure of the argument relation will not affect its convergence-preservation property due to Lemma 4.18. To ensure that also the result of the relational action is convergence-preserving, we explicitly enforce this property in a manner analogous to the way it was enforced in Definition 5.1. That this explicit enforcement is strictly necessary here is argued below the following definition.

Definition 5.2. Let τ_1 and τ'_1 be types with at most a single free variable, α say. Suppose R is a function that maps every $\tau_2, \tau'_2 \in Typ$ and $r \in Rel^\Downarrow(\tau_2, \tau'_2)$ to an $R_{\tau_2, \tau'_2}(r) \in Rel(\tau_1[\tau_2/\alpha], \tau'_1[\tau'_2/\alpha])$. Then we define $(\forall R) \in Rel^\Downarrow(\forall\alpha.\tau_1, \forall\alpha.\tau'_1)$ by

$$(G, G') \in (\forall R) \text{ iff } (G \Downarrow \Rightarrow G' \Downarrow) \\ \wedge \forall \tau_2, \tau'_2 \in Typ, r \in Rel^\Downarrow(\tau_2, \tau'_2). (G_{\tau_2}, G'_{\tau'_2}) \in R_{\tau_2, \tau'_2}(r)$$

for every $G \in Term(\forall\alpha.\tau_1)$ and $G' \in Term(\forall\alpha.\tau'_1)$. We also write $\forall R$ as $\forall r.R(r)$, suppressing reference to τ_2 and τ'_2 . \diamond

Assume the condition $G \Downarrow \Rightarrow G' \Downarrow$ were dropped from the previous definition. Then the relation $\forall R$ would no longer be guaranteed to be convergence-preserving. To see this, consider $\tau_1 = \tau'_1 = \alpha$ and let R be the function that maps every $\tau_2, \tau'_2 \in Typ$ and $r \in Rel^\Downarrow(\tau_2, \tau'_2)$ to $r^{\top\top}$. This function is well-behaved insofar as it maps every convergence-preserving relation to one that is both $\top\top$ -closed and convergence-preserving (cf. Lemma 4.18). Nevertheless, $\forall R$ would not be convergence-preserving,

because it would relate $G = \Omega$ and $G' = \Omega_{\forall\alpha.\alpha}$. This is because for every $\tau_2, \tau'_2 \in Typ$ and $r \in Rel^{\Downarrow}(\tau_2, \tau'_2)$, we have $(\Omega_{\tau_2}, (\Omega_{\forall\alpha.\alpha})_{\tau'_2}) \in r^{\top\top}$ by Observation 3.8 and Lemma 4.11. But $\Omega \Downarrow$ and $\Omega_{\forall\alpha.\alpha} \Uparrow$ by Observation 3.8.

The relational action for list types is a straightforward structural lifting, appropriately combined with $\top\top$ -closure. No special care is needed with respect to convergence-preservation, because it is satisfied automatically.

Definition 5.3. Given $\tau, \tau' \in Typ$ and $r \in Rel(\tau, \tau')$, we define $lift(r) \in Rel(\tau\text{-list}, \tau'\text{-list})$ as the greatest (post-)fixpoint (with respect to set inclusion) of the mapping $s \mapsto (1 + (r \times s))^{\top\top}$ for $s \in Rel(\tau\text{-list}, \tau'\text{-list})$, where

$$1 + (r \times s) = \{\mathbf{nil}_{\tau}, \mathbf{nil}_{\tau'}\} \cup \{(H : T, H' : T') \mid (H, H') \in r \wedge (T, T') \in s\}$$

for every such s . The existence of the greatest fixpoint is guaranteed by monotonicity of the mapping $s \mapsto (1 + (r \times s))^{\top\top}$ with respect to set inclusion, which in turn follows from (4). Note that the fixpoint property

$$lift(r) = (1 + (r \times lift(r)))^{\top\top}, \quad (5)$$

the observation that $1 + (r \times lift(r))$ is convergence-preserving (since it only relates values), and Lemma 4.18 imply that $lift(r)$ is convergence-preserving, and thus actually $lift(r) \in Rel^{\Downarrow}(\tau\text{-list}, \tau'\text{-list})$. \diamond

Combining the relational actions, the logical relation Δ is defined by induction on the structure of PolySeq types. It maps a type and a list containing convergence-preserving relations as interpretations for the type's free variables to a new relation. The new relation is convergence-preserving by construction, i.e., because the relational actions always deliver convergence-preserving relations.

Definition 5.4. For every type τ , $n \in \mathbb{N}$, list $\vec{\alpha} = \alpha_1, \dots, \alpha_n$ of distinct type variables containing the free variables of τ , lists $\vec{\tau} = \tau_1, \dots, \tau_n$ and $\vec{\tau}' = \tau'_1, \dots, \tau'_n$ of closed types, and list $\vec{r} = r_1, \dots, r_n$ with $r_i \in Rel^{\Downarrow}(\tau_i, \tau'_i)$ for every $1 \leq i \leq n$, we define $\Delta_{\tau}(\vec{r}/\vec{\alpha}) \in Rel^{\Downarrow}(\tau[\vec{\tau}/\vec{\alpha}], \tau[\vec{\tau}'/\vec{\alpha}])$ by induction on the structure of τ as follows:

$$\Delta_{\alpha_i}(\vec{r}/\vec{\alpha}) = r_i \quad (6)$$

$$\Delta_{\tau' \rightarrow \tau''}(\vec{r}/\vec{\alpha}) = \Delta_{\tau'}(\vec{r}/\vec{\alpha}) \rightarrow \Delta_{\tau''}(\vec{r}/\vec{\alpha}) \quad (7)$$

$$\Delta_{\forall\alpha.\tau'}(\vec{r}/\vec{\alpha}) = \forall r. \Delta_{\tau'}(\vec{r}, r^{\top\top}/\vec{\alpha}, \alpha) \quad (8)$$

$$\Delta_{\tau'\text{-list}}(\vec{r}/\vec{\alpha}) = lift(\Delta_{\tau'}(\vec{r}/\vec{\alpha})). \quad (9)$$

Note that without loss of generality the variable bound in the head of $\forall\alpha.\tau'$ in clause (8) can be assumed to not occur in $\vec{\alpha}$. Note also that the mapping $r \mapsto \Delta_{\tau'}(\vec{r}, r^{\top\top}/\vec{\alpha}, \alpha)$ in the right-hand side of that clause is well-defined since it will only be invoked for convergence-preserving relations r by Definition 5.2. But Lemma 4.18 ensures that $r^{\top\top}$ is also convergence-preserving for each such r , and can therefore be used as an argument for $\Delta_{\tau'}$. \diamond

To ultimately establish that $\Delta_\tau(\vec{r}/\vec{\alpha})$ is $\top\top$ -closed (in addition to being convergence-preserving) provided every relation in \vec{r} is, we have to show how $\top\top$ -closedness is pushed along the type structure by the relational actions for function and \forall -types. This is the task of the following two lemmas.

Lemma 5.5. For every $r_1, r_2 \in Rel$, if r_2 is $\top\top$ -closed, then so is $r_1 \rightarrow r_2$. \blacksquare

The proof of this lemma uses Observation 4.3(1) and Lemma 4.18. The main conceptual difference from the proof of Lemma 4.7(iii) in [Pit00] is the extra proof obligation arising from the additional convergence-preservation restriction in the definition of $r_1 \rightarrow r_2$. As in the equational setting of [VJ06], this extra obligation is met using Lemma 4.18. The proof of the following lemma is very similar.

Lemma 5.6. Let R be as in Definition 5.2. If $R_{\tau_2, \tau'_2}(r)$ is $\top\top$ -closed for every $\tau_2, \tau'_2 \in Typ$ and $r \in Rel^\Downarrow(\tau_2, \tau'_2)$, then $\forall R$ is also $\top\top$ -closed. \blacksquare

The following lemma gives the desired statement about propagation of $\top\top$ -closedness. It is the analogue of Lemma 4.11 in [Pit00], and is proved by induction on the structure of τ . The proof uses Lemmas 4.18, 5.5, and 5.6, and (5).

Lemma 5.7. Let τ , $\vec{\alpha}$, and \vec{r} be as in Definition 5.4. If every relation in \vec{r} is $\top\top$ -closed, then so is $\Delta_\tau(\vec{r}/\vec{\alpha})$. \blacksquare

By induction on the structure of types, we easily obtain the following results as well.

Observation 5.8. Let τ , $\vec{\alpha}$, and \vec{r} be as in Definition 5.4. Moreover, let $r' \in Rel^\Downarrow$ and let α' be a type variable not occurring in $\vec{\alpha}$ (and hence not occurring free in τ). Then:

$$\Delta_\tau(\vec{r}, r'/\vec{\alpha}, \alpha') = \Delta_\tau(\vec{r}/\vec{\alpha}). \quad \blacksquare$$

Observation 5.9. Let τ , $\vec{\alpha}$, and \vec{r} be as in Definition 5.4. Moreover, let α' be a type variable not occurring in $\vec{\alpha}$ and let τ' be a type with free variables in $\vec{\alpha}, \alpha'$. Then:

$$\Delta_{\tau'[\tau/\alpha']}(\vec{r}/\vec{\alpha}) = \Delta_{\tau'}(\vec{r}, \Delta_\tau(\vec{r}/\vec{\alpha})/\vec{\alpha}, \alpha'). \quad \blacksquare$$

6 Proving the main result

In this section we first prove the parametricity theorem and then prove its strengthening corresponding to identity extension. Proofs of parametricity theorems typically proceed by induction on typing derivations, which in our case would mean by induction on the axioms and rules in Figure 3. Instead, our proof will be based on the closely related system describing compatibility in Figure 5. In any case, considerations similar to those in [JV04, JV06] are necessary. In particular, since the relational action for function types was changed by adding a convergence-preservation restriction, a stronger statement than in the case of the ‘standard’ logical relation

must be proved for the rule in whose conclusion a function type appears, i.e., for the compatibility rule for function abstractions. The extra proof obligation thus arising in the following analogue of Lemma 4.7(i) in [Pit00] is met using the fact that every function abstraction is a value and thus trivially converges. The proof uses Observation 4.3(2) and proceeds as does the corresponding one for the equational setting, which can be found in its entirety in [VJ06].

Lemma 6.1. Let $\tau_1, \tau'_1, \tau_2, \tau'_2 \in Typ$, $r_1 \in Rel(\tau_1, \tau'_1)$, and $r_2 \in Rel(\tau_2, \tau'_2)$. Let x be a term variable and M and M' be terms such that $x :: \tau_1 \vdash M :: \tau_2$ and $x :: \tau'_1 \vdash M' :: \tau'_2$. If

$$\forall (A, A') \in r_1. (M[A/x], M'[A'/x]) \in r_2$$

and r_2 is $\top\top$ -closed, then $(\lambda x :: \tau_1. M, \lambda x :: \tau'_1. M') \in (r_1 \rightarrow r_2)$. \blacksquare

Similarly, we need an analogue of Lemma 4.8(i) in [Pit00], for type generalizations. In contrast to the situation in [JV04, JV06], we have now imposed an explicit convergence-preservation restriction on the relational action for \forall -types as well. Consequently, an extra proof obligation now also arises in the following ‘compatibility lemma’ for type generalizations, but can be met just as for function abstractions.

Lemma 6.2. Let τ_1, τ'_1, α , and R be as in Definition 5.2, and let M and M' be terms such that $\alpha \vdash M :: \tau_1$ and $\alpha \vdash M' :: \tau'_1$. If

$$\forall \tau_2, \tau'_2 \in Typ, r \in Rel^\Downarrow(\tau_2, \tau'_2). (M[\tau_2/\alpha], M'[\tau'_2/\alpha]) \in R_{\tau_2, \tau'_2}(r)$$

and $R_{\tau_2, \tau'_2}(r)$ is $\top\top$ -closed for every $\tau_2, \tau'_2 \in Typ$ and $r \in Rel^\Downarrow(\tau_2, \tau'_2)$, then $(\Lambda \alpha. M, \Lambda \alpha. M') \in (\forall R)$. \blacksquare

As is typical in proofs about logical relations, we have to generalize the statement we ultimately want to prove from terms and types without free variables to those with free variables. Hence, we also need to extend the logical relation itself to open terms, which is as usual done via closing substitutions.

Definition 6.3. Let $n, m \in \mathbb{N}$, let $\vec{\alpha}$ be a list of n type variables, $\vec{x} = x_1, \dots, x_m$ be a list of term variables, τ_1, \dots, τ_m be types, and $\Gamma = \vec{\alpha}, x_1 :: \tau_1, \dots, x_m :: \tau_m$. Given terms M and M' and a type τ with $\Gamma \vdash M :: \tau$ and $\Gamma \vdash M' :: \tau$, we write

$$\Gamma \vdash M \Delta M' :: \tau$$

if for every pair of lists $\vec{\sigma} = \sigma_1, \dots, \sigma_n$ and $\vec{\sigma}' = \sigma'_1, \dots, \sigma'_n$ of closed types and every list $\vec{r} = r_1, \dots, r_n$ of $\top\top$ -closed $r_i \in Rel^\Downarrow(\sigma_i, \sigma'_i)$, we have that for every pair of lists $\vec{N} = N_1, \dots, N_m$ and $\vec{N}' = N'_1, \dots, N'_m$ with $(N_j, N'_j) \in \Delta_{\tau_j}(\vec{r}/\vec{\alpha})$ for every $1 \leq j \leq m$, the following membership holds:

$$(M[\vec{\sigma}/\vec{\alpha}, \vec{N}/\vec{x}], M'[\vec{\sigma}'/\vec{\alpha}, \vec{N}'/\vec{x}]) \in \Delta_\tau(\vec{r}/\vec{\alpha}). \quad \diamond$$

Now, we can go about proving the fundamental property of Δ , namely that it is reflexive. That is, we prove the following parametricity theorem.

Theorem 6.4. The relation Δ is compatible, and thus reflexive. In particular, for every $\tau \in Typ$ and $M \in Term(\tau)$: $(M, M) \in \Delta_\tau()$.

Proof: To prove that Δ is compatible, we have to show that it is closed under each of the axioms and rules in Figure 5. The axiom $\Gamma, x :: \tau \vdash x \Delta x :: \tau$ is trivially satisfied due to the way Δ is defined. Also by that definition, to establish the rule

$$\frac{\Gamma \vdash A \Delta A' :: \tau \quad \Gamma \vdash B \Delta B' :: \tau'}{\Gamma \vdash \mathbf{seq}(A, B) \Delta \mathbf{seq}(A', B') :: \tau'}$$

it suffices to show that for Γ as in Definition 6.3, types τ and τ' , terms A, A', B , and B' with $\Gamma \vdash A :: \tau$, $\Gamma \vdash A' :: \tau$, $\Gamma \vdash B :: \tau'$, and $\Gamma \vdash B' :: \tau'$, lists $\vec{\sigma} = \sigma_1, \dots, \sigma_n$ and $\vec{\sigma}' = \sigma'_1, \dots, \sigma'_n$ of closed types, list $\vec{r} = r_1, \dots, r_n$ of $\top\top$ -closed $r_i \in Rel^{\downarrow}(\sigma_i, \sigma'_i)$, list $\vec{N} = N_1, \dots, N_m$ of $N_j \in Term(\tau_j[\vec{\sigma}/\vec{\alpha}])$, and list $\vec{N}' = N'_1, \dots, N'_m$ of $N'_j \in Term(\tau_j[\vec{\sigma}'/\vec{\alpha}])$,

$$(A[\vec{\sigma}/\vec{\alpha}, \vec{N}/\vec{x}], A'[\vec{\sigma}'/\vec{\alpha}, \vec{N}'/\vec{x}]) \in \Delta_\tau(\vec{r}/\vec{\alpha}) \quad (10)$$

and

$$(B[\vec{\sigma}/\vec{\alpha}, \vec{N}/\vec{x}], B'[\vec{\sigma}'/\vec{\alpha}, \vec{N}'/\vec{x}]) \in \Delta_{\tau'}(\vec{r}/\vec{\alpha}) \quad (11)$$

imply

$$(\mathbf{seq}(A, B)[\vec{\sigma}/\vec{\alpha}, \vec{N}/\vec{x}], \mathbf{seq}(A', B')[\vec{\sigma}'/\vec{\alpha}, \vec{N}'/\vec{x}]) \in \Delta_{\tau'}(\vec{r}/\vec{\alpha}).$$

By substitution, the latter is equivalent to

$$(\mathbf{seq}(A[\vec{\sigma}/\vec{\alpha}, \vec{N}/\vec{x}], B[\vec{\sigma}/\vec{\alpha}, \vec{N}/\vec{x}]), \mathbf{seq}(A'[\vec{\sigma}'/\vec{\alpha}, \vec{N}'/\vec{x}], B'[\vec{\sigma}'/\vec{\alpha}, \vec{N}'/\vec{x}])) \in \Delta_{\tau'}(\vec{r}/\vec{\alpha}).$$

But this is indeed implied by (10) and (11) due to Lemma 4.19, taking into account that $\Delta_\tau(\vec{r}/\vec{\alpha})$ is convergence-preserving (by Definition 5.4) and $\Delta_{\tau'}(\vec{r}/\vec{\alpha})$ is $\top\top$ -closed (by Lemma 5.7). The remaining axiom and rules from Figure 5 are established in a similar fashion using Lemmas 4.14, 4.18, 5.7, 6.1, and 6.2, Observation 5.9, (2), (5), and an analogue of Lemma 4.10(ii) in [Pit00], derived from (3), (5), and Observation 4.3. \blacksquare

We can use the technique of the preceding proof, together with Lemma 5.7 and Observation 5.9, to show that Δ is also closed under the rules in Figure 6. This gives the following lemma.

Lemma 6.5. The relation Δ is substitutive. \blacksquare

One further important property that Δ should have if it is to characterize observational approximation — indeed the very property tying in the observational aspect — is (inequational) adequacy with respect to **nil**-termination. The proof uses (5) and Observation 4.2 and is similar to that of the second part of Theorem 4.15 in [Pit00].

Lemma 6.6. The relation Δ is adequate. ■

Finally, we can show that Δ is not just any adequate, compatible, and substitutive relation, but is actually exactly the one we are looking for.

Theorem 6.7. The relation Δ is the largest adequate, compatible, and substitutive relation. It is also reflexive and transitive.

Proof: For the first statement, by Theorem 6.4 and Lemmas 6.5 and 6.6, it remains to prove that Δ subsumes every adequate, compatible, and substitutive relation. Let \mathcal{E} be such a relation, let Γ , M , M' , and τ be as in Definition 6.3, and assume $\Gamma \vdash M \mathcal{E} M' :: \tau$. Further, let $\vec{\sigma} = \sigma_1, \dots, \sigma_n$ and $\vec{\sigma}' = \sigma'_1, \dots, \sigma'_n$ be lists of closed types, $\vec{r} = r_1, \dots, r_n$ be a list of $\top\top$ -closed $r_i \in Rel^\downarrow(\sigma_i, \sigma'_i)$, and $\vec{N} = N_1, \dots, N_m$ and $\vec{N}' = N'_1, \dots, N'_m$ be lists with $(N_j, N'_j) \in \Delta_{\tau_j}(\vec{r}/\vec{\alpha})$ for every $1 \leq j \leq m$. Since Δ is reflexive (cf. Theorem 6.4), we have $\Gamma \vdash M \Delta M :: \tau$, which by Definition 6.3 implies that

$$(M[\vec{\sigma}/\vec{\alpha}, \vec{N}/\vec{x}], M[\vec{\sigma}'/\vec{\alpha}, \vec{N}'/\vec{x}]) \in \Delta_\tau(\vec{r}/\vec{\alpha}). \quad (12)$$

Moreover, $\Gamma \vdash M \mathcal{E} M' :: \tau$ and the substitutivity of \mathcal{E} imply that

$$M[\vec{\sigma}'/\vec{\alpha}, \vec{N}'/\vec{x}] \mathcal{E} M'[\vec{\sigma}'/\vec{\alpha}, \vec{N}'/\vec{x}]. \quad (13)$$

Since (12) and (13) combine into $(M[\vec{\sigma}/\vec{\alpha}, \vec{N}/\vec{x}], M'[\vec{\sigma}'/\vec{\alpha}, \vec{N}'/\vec{x}]) \in \Delta_\tau(\vec{r}/\vec{\alpha})$ by Lemmas 4.10 and 5.7, and since this is obtained independently of the choice of the lists $\vec{\sigma}$, $\vec{\sigma}'$, \vec{r} , \vec{N} , and \vec{N}' above, we indeed have the desired $\Gamma \vdash M \Delta M' :: \tau$ by Definition 6.3.

For the second statement, note that since Δ is compatible, it is also reflexive. Moreover, it is easy to see that the collection of adequate, compatible, and substitutive relations is closed under relation composition. This implies that $\Delta; \Delta$ is adequate, compatible, and substitutive, and is thus subsumed by the largest such relation, i.e., $\Delta; \Delta \subseteq \Delta$. But this means that Δ is also transitive. ■

With the above characterization, we have established that Δ coincides with our intended notion of approximation as discussed at the end of Section 3.2. So the following corollary can be read either as a definition of \sqsubseteq_{obs} in terms of Δ or as a coincidence statement between \sqsubseteq_{obs} and Δ . In the latter reading, we take \sqsubseteq_{obs} to be characterized as the union of all adequate, compatible, and substitutive relations and Δ to be characterized independently based on the type structure of PolySeq.

Corollary 6.8. Let Γ , M , M' , and τ be as in Definition 6.3. Then:

$$\Gamma \vdash M \sqsubseteq_{obs} M' :: \tau \Leftrightarrow \Gamma \vdash M \Delta M' :: \tau.$$

In particular, for every $\tau \in Typ$ and $M, M' \in Term(\tau)$:

$$M \sqsubseteq_{obs} M' \Leftrightarrow (M, M') \in \Delta_\tau(). \quad \blacksquare$$

Together with Observations 5.8 and 5.9, the statement about closed types in the previous corollary implies a statement about observational approximation corresponding to what is usually referred to as the identity extension lemma. This statement says that for every type τ with free variables in $\vec{\alpha} = \alpha_1, \dots, \alpha_n$, list $\vec{\tau} = \tau_1, \dots, \tau_n$ of closed types, and list $\vec{r} = r_1, \dots, r_n$ of relations with

$$r_i = \{(M, M') \mid M, M' \in \text{Term}(\tau_i) \wedge M \sqsubseteq_{\text{obs}} M'\},$$

we have:

$$\Delta_\tau(\vec{r}/\vec{\alpha}) = \{(M, M') \mid M, M' \in \text{Term}(\tau[\vec{r}/\vec{\alpha}]) \wedge M \sqsubseteq_{\text{obs}} M'\}.$$

7 Applications

In this section we show how our characterization of observational approximation by a logical relation can be used to study the semantics of PolySeq. Our study progresses from relatively basic statements about notions from the operational semantics in Section 3.2 to the correctness of parametricity-based program transformations.

7.1 Exploring observational approximation

Since for every $\tau \in \text{Typ}$, $\Delta_\tau()$ is convergence-preserving by Definition 5.4, we obtain the following simple consequence of Corollary 6.8.

Corollary 7.1. For every $M, M' \in \text{Term}$ of the same type, if $M \sqsubseteq_{\text{obs}} M'$, then $M \Downarrow \Rightarrow M' \Downarrow$. ■

Note that the above does not hold in the setting without **seq**. In particular, Examples 5.3 and 5.5 in [Pit00] provide PolyPCF observational equivalences — and hence approximations — that do not adhere to convergence-preservation. Namely, they yield that for every $\tau_1, \tau_2 \in \text{Typ}$, $\lambda x :: \tau_1. \Omega_{\tau_2} \sqsubseteq_{\text{obs}} \Omega_{\tau_1 \rightarrow \tau_2}$, and that for every type τ with at most a single free variable, α say, $\Lambda \alpha. \Omega_\tau \sqsubseteq_{\text{obs}} \Omega_{\forall \alpha. \tau}$. In light of Observation 3.8 and Corollary 7.1, neither of the two can be true in PolySeq. This motivates turning attention to extensionality principles, because that is where the two examples above come from. But before doing so, we look at three different ways of establishing observational approximations that work independently of whether or not **seq** is present in the calculus.

First, note that the converse implication does not hold in Corollary 7.1: the fact that convergence of one term implies convergence of another does not mean that the first term observationally approximates the second. The reason is that two terms that are both convergent can have otherwise completely different behavior. But if the first term is divergent, then it approximates any term of the same type, regardless of the convergence behavior of that second term, as established by the following lemma.

Lemma 7.2. For every $M, M' \in \text{Term}$ of the same type, if $M \Uparrow$, then $M \sqsubseteq_{\text{obs}} M'$.

Proof: By Corollary 6.8, it suffices to show that $(M, M') \in \Delta_\tau()$, where τ is the type of both M and M' . But this follows from Lemmas 4.11 and 5.7. ■

If we say that two terms M and M' are *observationally equivalent*, written $M =_{obs} M'$, if and only if $M \sqsubseteq_{obs} M'$ and $M' \sqsubseteq_{obs} M$ both hold, then Lemma 7.2 entails that all divergent terms of the same type are observationally equivalent.

To establish observational equivalence of two convergent terms, we may use one of the following two lemmas. Each states that a certain operational notion from Section 3.2 implies observational equivalence. Note that two separate observational approximations must be established for each lemma. The statements can also be regarded as instances of Corollary 4.16 in [Pit00].

Lemma 7.3. For every $M \in Term$ and $V \in Value$, if $M \Downarrow V$, then $M =_{obs} V$.

Proof: By Corollary 6.8, it suffices to show that $(M, V) \in \Delta_\tau()$ and $(V, M) \in \Delta_\tau()$, where τ is the type of both M and V . By Lemma 5.7, $\Delta_\tau()$ is $\top\top$ -closed, i.e., is equal to $(\Delta_\tau())^{\top\top}$. Then it suffices to reason for every $(S, S') \in (\Delta_\tau())^\top$ as follows:

$$\begin{aligned} S \top M &\Leftrightarrow S \top V && \text{by } M \Downarrow V \text{ and Lemma 4.5} \\ &\Rightarrow S' \top V && \text{by } (S, S') \in (\Delta_\tau())^\top \text{ and } (V, V) \in \Delta_\tau() \end{aligned}$$

and

$$\begin{aligned} S \top V &\Rightarrow S' \top V && \text{by } (S, S') \in (\Delta_\tau())^\top \text{ and } (V, V) \in \Delta_\tau() \\ &\Leftrightarrow S' \top M && \text{by } M \Downarrow V \text{ and Lemma 4.5,} \end{aligned}$$

where $(V, V) \in \Delta_\tau()$ holds by Theorem 6.4. ■

Lemma 7.4. For every $R, R' \in Term$, if $R \rightsquigarrow R'$, then $R =_{obs} R'$. ■

The proof of Lemma 7.4 is analogous to that of Lemma 7.3 above, but using Observation 4.3(2) instead of Lemma 4.5. Combining Lemmas 7.3 and 7.4 with the fact that \sqsubseteq_{obs} is compatible, we obtain the following corollary.

Corollary 7.5. For every $A, B \in Term$, if $A \Downarrow$, then $\mathbf{seq}(A, B) =_{obs} B$. ■

7.2 Extensionality principles

This subsection deals with extensionality principles. The standard one for function types — namely that two functions are ‘the same’ if (and only if) they return the same results for every possible argument — has to be revised for PolySeq, in addition to porting it to the inequational setting. Since \mathbf{seq} allows an additional observation to be made about function terms, namely checking them for convergence without applying them to any argument, it is necessary that convergence be preserved between the two function terms in question.

Lemma 7.6. Let $\tau_1, \tau_2 \in Typ$. Then for every $F, F' \in Term(\tau_1 \rightarrow \tau_2)$:

$$F \sqsubseteq_{obs} F' \text{ iff } (F \Downarrow \Rightarrow F' \Downarrow) \wedge \forall A \in Term(\tau_1). F A \sqsubseteq_{obs} F' A.$$

Proof: The left-to-right implication follows from Corollary 7.1 and the compatibility of \sqsubseteq_{obs} . For the right-to-left implication it suffices by Corollary 6.8 and clause (7) to prove that $(F, F') \in \Delta_{\tau_1}() \rightarrow \Delta_{\tau_2}()$. Since $F \Downarrow \Rightarrow F' \Downarrow$ is known, it remains by Definition 5.1 to show that for every $(M, M') \in \Delta_{\tau_1}()$: $(F M, F' M') \in \Delta_{\tau_2}()$. But this follows by Lemmas 4.10 (for the adequate and compatible relation \sqsubseteq_{obs}) and 5.7 from

$$(F M, F' M') \in \Delta_{\tau_2}() \wedge F M' \sqsubseteq_{obs} F' M'.$$

Here the first conjunct follows by Definition 5.1 from $(F, F) \in \Delta_{\tau_1}() \rightarrow \Delta_{\tau_2}()$, which holds due to Theorem 6.4, while the second one follows from the assumption that for every $A \in Term(\tau_1)$, $F A \sqsubseteq_{obs} F' A$. \blacksquare

The standard extensionality principle for \forall -types states that two polymorphic terms are observationally equivalent if and only if all their corresponding instances are. In PolySeq we must enrich the inequational version of this principle by an explicit convergence-preservation check, just as we did with the extensionality principle for function types above. The reason, of course, is that **seq** can also be used at \forall -types. The proof of the following lemma is completely analogous to that of Lemma 7.6 above.

Lemma 7.7. Let τ be a type with at most a single free variable, α say. Then for every $G, G' \in Term(\forall\alpha.\tau)$:

$$G \sqsubseteq_{obs} G' \text{ iff } (G \Downarrow \Rightarrow G' \Downarrow) \wedge \forall \tau' \in Typ. G_{\tau'} \sqsubseteq_{obs} G'_{\tau'}. \quad \blacksquare$$

One important consequence of Lemma 7.7 is that in PolySeq it is not true for every $G \in Term(\forall\alpha.\tau)$ that if for every $\tau' \in Typ$, $G_{\tau'} =_{obs} \Omega_{\tau[\tau'/\alpha]}$, then $G =_{obs} \Omega_{\forall\alpha.\tau}$. A counterexample is exactly the term $G = \Lambda\alpha.\Omega_{\tau}$ from Example 5.5 in [Pit00]. (Note that Example 5.5 in [Pit00] was shown above — below Corollary 7.1 — to fail in the presence of **seq**.) This might seem a bit surprising, given that the refuted implication is the direct translation into the operational setting of laws (2) and (5) in [JV04] and [JV06], respectively. The resolution of this apparent contradiction is that PolySeq combines selective strictness and impredicative polymorphism in such a way that type generalization and instantiation are observable, while they are not in Haskell (and while Pitts' PolyPCF has only impredicative polymorphism but no selective strictness). Ultimately, this is also the reason why we had to explicitly enforce $G \Downarrow \Rightarrow G' \Downarrow$ in the relational action for \forall -types (cf. Definition 5.2), whereas a corresponding enforcement was not necessary in [JV04, JV06].

The extensionality principle for list types is essentially a coinduction principle. To state it, we need the following notion.

Definition 7.8. Let $\tau \in Typ$ and $s \in Rel(\tau\text{-list}, \tau\text{-list})$. We say that s is a *simulation* if for every $(L, L') \in s$ the following implications hold:

1. if $L \Downarrow \mathbf{nil}_{\tau}$, then $L' \Downarrow \mathbf{nil}_{\tau}$,

2. if $L \Downarrow H : T$ (for some H and T), then there are H' and T' with $L' \Downarrow H' : T'$, $H \sqsubseteq_{obs} H'$, and $(T, T') \in s$. \diamond

Note that it is also possible to define the notion of simulation more generally, with an arbitrary relation to tie corresponding list elements together rather than doing so by \sqsubseteq_{obs} at some type τ . But the specialized definition above suffices for our purposes and we prefer not to add unnecessary complication. A particularly simple example of a simulation is given in the following lemma, whose proof is straightforward, using adequacy and compatibility of \sqsubseteq_{obs} , Observation 3.8, Corollary 7.1, and Lemmas 7.3 and 7.4.

Lemma 7.9. For every $\tau \in Typ$, the relation $\{(L, L') \mid L, L' \in Term(\tau\text{-list}) \wedge L \sqsubseteq_{obs} L'\}$ is a simulation. \blacksquare

The extensionality principle for list types then amounts to the statement that the simulation from the previous lemma is the greatest one at a given $\tau \in Typ$.

Lemma 7.10. Let $\tau \in Typ$. Then for every $L, L' \in Term(\tau\text{-list})$, $L \sqsubseteq_{obs} L'$ if and only if (L, L') is contained in some simulation. \blacksquare

The proof of the lemma is straightforward, proceeding essentially as in [Pit00] or, more simply, [VJ06]. That no changes (apart from replacing bidirectional implication by unidirectional implication) are necessary for list types — either here in the extensionality principle or prior to that in the development of the logical relation — makes sense intuitively. For inequationality at list types has been ‘built into’ the logical relation via the new notion of $\top\top$ -closure, rather than ‘added on’ *a posteriori* as in [JV04, JV06]. And inclusion of **seq** into the calculus has no real impact on what can happen at list types, given that **seq**(A, B) with A of list type can always be ‘mimicked’ by (**case** A **of** $\{\text{nil} \Rightarrow B; h : t \Rightarrow B\}$).

7.3 Manufacturing permissible relations

For applications of the logical relation that make more essential use of the power inherent in the quantification over relations in clause (8), such as those applications studied in the next two subsections, we need a source of appropriately restricted relations. In the equational setting without **seq**, Pitts identified a source of $\top\top$ -closed relations by considering graphs of evaluation frame stacks. We define two analogous notions of graphs suited to the inequational setting as follows.

Definition 7.11. For every $\tau, \tau' \in Typ$ and $S \in Stack(\tau, \tau')$, we define *left-graph* $_S \in Rel(\tau, \tau')$ by

$$(M, M') \in \text{left-graph}_S \text{ iff } S M \sqsubseteq_{obs} M'$$

and *right-graph* $_S \in Rel(\tau', \tau)$ by

$$(M, M') \in \text{right-graph}_S \text{ iff } M \sqsubseteq_{obs} S M'. \quad \diamond$$

Since in the inequational treatment of PolySeq we have to ensure that relational interpretations of types are not only $\top\top$ -closed but also convergence-preserving, we need to restrict attention to just particular stacks that give rise to such relations. As the following lemma shows, actually no restriction is necessary for right-graphs.

Lemma 7.12. For every $\tau, \tau' \in Typ$ and $S \in Stack(\tau, \tau')$, $right-graph_S$ is $\top\top$ -closed and convergence-preserving.

Proof: The proof of $\top\top$ -closure is essentially as that for Lemma 6.1 in [Pit00]. To see that $right-graph_S$ is convergence-preserving, let $(M, M') \in right-graph_S$, i.e., $M \sqsubseteq_{obs} S M'$. By Corollary 7.1, this implies that $M \Downarrow \Rightarrow (S M') \Downarrow$. So it remains to show that $(S M') \Downarrow \Rightarrow M' \Downarrow$. But this holds by Lemma 4.20. \blacksquare

For left-graphs, the above observation motivates the following definition.

Definition 7.13. Given $\tau, \tau' \in Typ$ and $S \in Stack(\tau, \tau')$, we say that S is *total* if for every $V \in Value(\tau)$: $(S V) \Downarrow$. \diamond

The analogue of Lemma 7.12 for left-graphs is now proved similarly to above (using the opposite direction of Lemma 4.20).

Lemma 7.14. For every $\tau, \tau' \in Typ$ and $S \in Stack(\tau, \tau')$, $left-graph_S$ is $\top\top$ -closed. Moreover, if S is total, then $left-graph_S$ is convergence-preserving as well. \blacksquare

7.4 Enumerating terms up to observational equivalence

One application of parametricity and extensionality results presented in [Pit00], and one which is also popular in the Haskell community (as various discussions on [HML] indicate), is to show that certain types have only a small number of inhabitants which differ with respect to observational equivalence. In particular, Pitts' Examples 2.6 and 2.7 state, respectively, that in the absence of **seq**, every element of $Term(\forall\alpha.\alpha)$ is observationally equivalent to Ω , and that every element of $Term(\forall\alpha.\alpha \rightarrow \alpha)$ is observationally equivalent either to $\Omega_{\forall\alpha.\alpha \rightarrow \alpha}$ or to $\Lambda\alpha.\lambda x :: \alpha.x$. We will revisit these examples for PolySeq and observe interesting differences. But first we need the following auxiliary lemma.

Lemma 7.15. For every $\tau \in Typ$ and $M, M' \in Term(\tau)$, if $M \Uparrow$ and $(M', M') \in \{(M, M)\}^{\top\top}$, then $M' \Uparrow$ as well.

Proof: From $(M', M') \in \{(M, M)\}^{\top\top}$ follows that for every $(S, S') \in \{(M, M)\}^{\top}$: $S \top M' \Rightarrow S' \top M'$. Choose $S = \mathbf{seq}(-, \mathbf{nil}_\tau)$ and $S' = \mathbf{seq}(-, \Omega_{\tau-list})$, which is a valid choice due to the following reasoning:

$$\begin{aligned}
 \mathbf{seq}(-, \mathbf{nil}_\tau) \top M &\Leftrightarrow M \Downarrow && \text{by Corollary 4.16} \\
 &\Leftrightarrow M \Downarrow \wedge Id \top \Omega_{\tau-list} && \text{by } M \Uparrow \\
 &\Leftrightarrow Id \top \mathbf{seq}(M, \Omega_{\tau-list}) && \text{by Lemma 4.15} \\
 &\Leftrightarrow \mathbf{seq}(-, \Omega_{\tau-list}) \top M && \text{by Observation 4.3(1).}
 \end{aligned}$$

Now, $\mathbf{seq}(-, \mathbf{nil}_\tau) \top M' \Rightarrow \mathbf{seq}(-, \Omega_{\tau\text{-list}}) \top M'$ implies

$$M' \Downarrow \Rightarrow (M' \Downarrow \wedge Id \top \Omega_{\tau\text{-list}})$$

by using Observation 4.3(1), Lemma 4.15, and Corollary 4.16 in a manner similar to that above. Since $Id \top \Omega_{\tau\text{-list}}$ does not hold (cf. Observations 3.8 and 4.2), this implies $M' \Uparrow$. \blacksquare

Now, we can show that although not all elements of $Term(\forall\alpha.\alpha)$ are observationally equivalent in PolySeq, they do separate into exactly two equivalence classes.

Lemma 7.16. For every $G \in Term(\forall\alpha.\alpha)$, either $G =_{obs} \Omega$ or $G =_{obs} \Omega_{\forall\alpha.\alpha}$.

Proof: By Theorem 6.4 and Definition 5.4, we have $(G, G) \in (\forall r.r^{\top\top})$. By Definition 5.2, this implies that for every $\tau \in Typ$ and $r \in Rel^\Downarrow(\tau, \tau)$: $(G_\tau, G_\tau) \in r^{\top\top}$. For fixed τ , choose $r = \{(\Omega_\tau, \Omega_\tau)\}$, which obviously is convergence-preserving. Then we get $G_\tau \Uparrow$ by Observation 3.8 and Lemma 7.15. Since this is so for every $\tau \in Typ$, the claim follows by Observation 3.8 and Lemmas 7.2 and 7.7. \blacksquare

To provide a similar account for the type $\forall\alpha.\alpha \rightarrow \alpha$, we first need a further auxiliary lemma. Its statement should also hold for Pitts' setting without \mathbf{seq} , but a proof would be much more complicated there, as it could not make use of the convergence-preservation of relational interpretations of types. In fact, there seems to be no way in PolyPCF to prove the statement solely based on the reflexivity of the logical relation as below.

Lemma 7.17. Let τ be a type with at most a single free variable, α say. Then for every $G \in Term(\forall\alpha.\tau)$:

$$(\exists\tau' \in Typ. G_{\tau'} \Downarrow) \Rightarrow (\forall\tau' \in Typ. G_{\tau'} \Downarrow).$$

Proof: By Theorem 6.4 and clause (8), we have $(G, G) \in (\forall r.\Delta_\tau(r^{\top\top}/\alpha))$. By Definition 5.2, this implies that for every $\tau_1, \tau_2 \in Typ$ and $r \in Rel^\Downarrow(\tau_1, \tau_2)$: $(G_{\tau_1}, G_{\tau_2}) \in \Delta_\tau(r^{\top\top}/\alpha)$. For fixed τ_1 and τ_2 , choose $r = \emptyset$, which obviously is convergence-preserving. By Lemma 4.18 and Definition 5.4, $\Delta_\tau(r^{\top\top}/\alpha)$ is then also convergence-preserving. This means that for every $\tau_1, \tau_2 \in Typ$: $G_{\tau_1} \Downarrow \Rightarrow G_{\tau_2} \Downarrow$. \blacksquare

Now, we can show that in PolySeq, $Term(\forall\alpha.\alpha \rightarrow \alpha)$ is factorized into four, rather than two, equivalence classes with respect to $=_{obs}$. The two additional ones arise exactly from the failures of Pitts' Examples 5.3 and 5.4 in the presence of \mathbf{seq} , as observed below Corollary 7.1.

Lemma 7.18. For every $G \in Term(\forall\alpha.\alpha \rightarrow \alpha)$, exactly one of the following alternatives holds:

1. $G =_{obs} \Omega_{\forall\alpha.\alpha \rightarrow \alpha}$,
2. $G =_{obs} \Lambda\alpha.\Omega_{\alpha \rightarrow \alpha}$,

3. $G =_{obs} \Lambda\alpha.\lambda x :: \alpha.\Omega_\alpha$, or

4. $G =_{obs} \Lambda\alpha.\lambda x :: \alpha.x$.

Proof: By Theorem 6.4 and Definition 5.4, we have $(G, G) \in (\forall r.r^{\top\top} \rightarrow r^{\top\top})$. By Definitions 5.1 and 5.2, this implies the following statement:

$$\begin{aligned} \forall \tau_1, \tau_2 \in Typ, r \in Rel^\downarrow(\tau_1, \tau_2), M_1 \in Term(\tau_1), M_2 \in Term(\tau_2). \\ (M_1, M_2) \in r^{\top\top} \Rightarrow (G_{\tau_1} M_1, G_{\tau_2} M_2) \in r^{\top\top}. \end{aligned} \quad (14)$$

From this, we get that

$$\forall \tau \in Typ, M \in Term(\tau). M \uparrow \Rightarrow (G_\tau M) \uparrow \quad (15)$$

by instantiating $\tau_1 = \tau_2 = \tau$, $r = \{(M, M)\}$, and $M_1 = M_2 = M$, and by using (2) and Lemma 7.15. Now, we prove that

$$\forall \tau \in Typ, M \in Term(\tau). G_\tau M \sqsubseteq_{obs} \Omega_\tau \quad (16)$$

or

$$\forall \tau \in Typ, M \in Term(\tau). G_\tau M =_{obs} M, \quad (17)$$

by contradiction. Assume that neither (16) nor (17) holds. Then there exist $\tau_1, \tau_2 \in Typ$, $M_1 \in Term(\tau_1)$, and $M_2 \in Term(\tau_2)$ with

$$\neg(G_{\tau_1} M_1 \sqsubseteq_{obs} \Omega_{\tau_1})$$

and

$$\neg(G_{\tau_2} M_2 =_{obs} M_2). \quad (18)$$

From these negations we get $(G_{\tau_1} M_1) \downarrow$, $M_1 \downarrow$, and $M_2 \downarrow$ by Lemma 7.2 and (15). Using (14) with $r = left-graph_{seq(-, M_2)}$, which is $\top\top$ -closed and convergence-preserving by $M_2 \downarrow$, Observation 3.7, and Lemma 7.14, we get that $(M_1, M_2) \in left-graph_{seq(-, M_2)}$ implies $(G_{\tau_1} M_1, G_{\tau_2} M_2) \in left-graph_{seq(-, M_2)}$, i.e.,

$$seq(M_1, M_2) \sqsubseteq_{obs} M_2 \Rightarrow seq(G_{\tau_1} M_1, M_2) \sqsubseteq_{obs} G_{\tau_2} M_2.$$

By $M_1 \downarrow$, $(G_{\tau_1} M_1) \downarrow$, and Corollary 7.5, the precondition of this implication is fulfilled, whereas its conclusion and (18) together imply that

$$\neg(G_{\tau_2} M_2 \sqsubseteq_{obs} M_2). \quad (19)$$

must hold. Now using (14) with $r = right-graph_{seq(-, M_2)}$, which is $\top\top$ -closed and convergence-preserving by Lemma 7.12, we get that $(M_2, M_1) \in right-graph_{seq(-, M_2)}$ implies $(G_{\tau_2} M_2, G_{\tau_1} M_1) \in right-graph_{seq(-, M_2)}$, i.e.,

$$M_2 \sqsubseteq_{obs} seq(M_1, M_2) \Rightarrow G_{\tau_2} M_2 \sqsubseteq_{obs} seq(G_{\tau_1} M_1, M_2).$$

But by $M_1 \downarrow$, $(G_{\tau_1} M_1) \downarrow$, and Corollary 7.5, the precondition of this implication is fulfilled, whereas its conclusion contradicts (19). Consequently, the assumption that neither (16) nor (17) holds was wrong. From this, the claim follows by Observation 3.8, Corollary 7.1, and Lemmas 7.2, 7.4, 7.6, 7.7, and 7.17. \blacksquare

7.5 Correctness of short cut fusion

The PolySeq equivalents of the Haskell functions *foldr* and *build* from Figure 1 in Section 2 are as follows:

$$\begin{aligned} \text{foldr} = \Lambda\alpha.\Lambda\beta.\lambda c :: \alpha \rightarrow \beta \rightarrow \beta.\lambda n :: \beta.\mathbf{fix}(\lambda f :: \alpha\text{-list} \rightarrow \beta.\lambda l :: \alpha\text{-list}. \\ \mathbf{case } l \mathbf{ of } \{\mathbf{nil} \Rightarrow n; h : t \Rightarrow c h (f t)\}) \end{aligned}$$

and

$$\text{build} = \Lambda\alpha.\lambda g :: \forall\beta.(\alpha \rightarrow \beta \rightarrow \beta) \rightarrow \beta \rightarrow \beta.g_{\alpha\text{-list}} (\lambda h :: \alpha.\lambda t :: \alpha\text{-list}.h : t) \mathbf{nil}_\alpha.$$

Due to the explicit syntactic representation of type specialization in PolySeq, the short cut fusion rule (1) becomes a bit more verbose now. It says that for every $\tau, \tau' \in \text{Typ}$, $G \in \text{Term}(\forall\beta.(\tau \rightarrow \beta \rightarrow \beta) \rightarrow \beta \rightarrow \beta)$, $C \in \text{Term}(\tau \rightarrow \tau' \rightarrow \tau')$, and $N \in \text{Term}(\tau')$:

$$(\text{foldr}_\tau)_{\tau'} C N (\text{build}_\tau G) =_{\text{obs}} G_{\tau'} C N \quad (20)$$

Of course, just as noted for ‘Haskell with *seq*’ in Section 2, the rule does not hold unconditionally in PolySeq. For example, it is easy to see that the above observational equivalence does not hold if G is $\Lambda\beta.\lambda c :: \tau \rightarrow \beta \rightarrow \beta.\lambda n :: \beta.\mathbf{seq}(c, n)$, C is $\Omega_{\tau \rightarrow \tau' \rightarrow \tau'}$, and N is any converging element of $\text{Term}(\tau')$.

But partial correctness of *foldr/build*-fusion does hold even in PolySeq with its selective strictness primitive, and total correctness can be recovered by imposing restrictions on N and C , as established in the following theorem.

Theorem 7.19. For every $\tau, \tau' \in \text{Typ}$, $G \in \text{Term}(\forall\beta.(\tau \rightarrow \beta \rightarrow \beta) \rightarrow \beta \rightarrow \beta)$, $C \in \text{Term}(\tau \rightarrow \tau' \rightarrow \tau')$, and $N \in \text{Term}(\tau')$:

$$G_{\tau'} C N \sqsubseteq_{\text{obs}} (\text{foldr}_\tau)_{\tau'} C N (\text{build}_\tau G).$$

Moreover, if $N \Downarrow$ and for every $A \in \text{Term}(\tau)$ and $B \in \text{Term}(\tau')$, $(C A B) \Downarrow$, then:

$$(\text{foldr}_\tau)_{\tau'} C N (\text{build}_\tau G) \sqsubseteq_{\text{obs}} G_{\tau'} C N.$$

Proof: By Theorem 6.4 and Definition 5.4, we have

$$(G, G) \in (\forall r.(\Delta_\tau(r^{\top\top}/\beta) \rightarrow (r^{\top\top} \rightarrow r^{\top\top})) \rightarrow (r^{\top\top} \rightarrow r^{\top\top})).$$

By Lemma 4.18, Definitions 5.1 and 5.2, and Observation 5.8, this implies that for every $\tau_1, \tau_2 \in \text{Typ}$, $r \in \text{Rel}^\Downarrow(\tau_1, \tau_2)$, $C_1 \in \text{Term}(\tau \rightarrow \tau_1 \rightarrow \tau_1)$, $C_2 \in \text{Term}(\tau \rightarrow \tau_2 \rightarrow \tau_2)$, $N_1 \in \text{Term}(\tau_1)$, and $N_2 \in \text{Term}(\tau_2)$:

$$\begin{aligned} & C_1 \Downarrow \Rightarrow C_2 \Downarrow \\ & \wedge (\forall(A_1, A_2) \in \Delta_\tau()). (C_1 A_1) \Downarrow \Rightarrow (C_2 A_2) \Downarrow \\ & \wedge \forall(B_1, B_2) \in r^{\top\top}. (C_1 A_1 B_1, C_2 A_2 B_2) \in r^{\top\top} \\ & \wedge (N_1, N_2) \in r^{\top\top} \\ & \Rightarrow (G_{\tau_1} C_1 N_1, G_{\tau_2} C_2 N_2) \in r^{\top\top}. \end{aligned}$$

We consider two instantiations of this.

First, we instantiate $\tau_1 = \tau'$, $\tau_2 = \tau\text{-list}$, $C_1 = C$, $C_2 = \lambda h :: \tau.\lambda t :: \tau\text{-list}.h : t$, $N_1 = N$, $N_2 = \mathbf{nil}_\tau$, and

$$r = \{(M, L) \mid M \in \text{Term}(\tau') \wedge L \in \text{Term}(\tau\text{-list}) \wedge M \sqsubseteq_{\text{obs}} (\text{foldr}_\tau)_{\tau'} C N L\}.$$

By the definition of *foldr*, Lemma 7.4, and the compatibility of \sqsubseteq_{obs} , this r equals *right-graph_S*, where

$$S = \mathbf{case} - \mathbf{of} \{ \mathbf{nil} \Rightarrow N; \\ h : t \Rightarrow C h (\mathbf{fix}(\lambda f :: \tau\text{-list} \rightarrow \tau'.\lambda l :: \tau\text{-list}. \\ \mathbf{case} l \mathbf{of} \{ \mathbf{nil} \Rightarrow N; h' : t' \Rightarrow C h' (f t') \}) t) \},$$

and thus is $\top\top$ -closed and convergence-preserving by Lemma 7.12. Hence, the choice of r was justified (i.e., r really is in $\text{Rel}^\downarrow(\tau', \tau\text{-list})$), and occurrences of $r^{\top\top}$ in the implication displayed above can be replaced by r itself. Then the first claim of the theorem follows from the definition of *build*, Lemma 7.4, the compatibility of \sqsubseteq_{obs} , and Observation 3.7, provided we can establish

$$\begin{aligned} \forall (A_1, A_2) \in \Delta_\tau(), B_1 \in \text{Term}(\tau'), B_2 \in \text{Term}(\tau\text{-list}). \\ B_1 \sqsubseteq_{\text{obs}} (\text{foldr}_\tau)_{\tau'} C N B_2 \\ \Rightarrow C A_1 B_1 \sqsubseteq_{\text{obs}} (\text{foldr}_\tau)_{\tau'} C N ((\lambda h :: \tau.\lambda t :: \tau\text{-list}.h : t) A_2 B_2) \end{aligned}$$

and

$$N \sqsubseteq_{\text{obs}} (\text{foldr}_\tau)_{\tau'} C N \mathbf{nil}_\tau.$$

But these two statements follow from Corollary 6.8, the definition of *foldr*, Lemma 7.4, and the compatibility of \sqsubseteq_{obs} .

Now, let $N \Downarrow$ and $(C A B) \Downarrow$ for every $A \in \text{Term}(\tau)$ and $B \in \text{Term}(\tau')$. Then we instantiate $\tau_1 = \tau\text{-list}$, $\tau_2 = \tau'$, $C_1 = \lambda h :: \tau.\lambda t :: \tau\text{-list}.h : t$, $C_2 = C$, $N_1 = \mathbf{nil}_\tau$, $N_2 = N$, and

$$r = \{(L, M) \mid L \in \text{Term}(\tau\text{-list}) \wedge M \in \text{Term}(\tau') \wedge (\text{foldr}_\tau)_{\tau'} C N L \sqsubseteq_{\text{obs}} M\}.$$

By the definition of *foldr*, Lemma 7.4, and the compatibility of \sqsubseteq_{obs} , this r equals *left-graph_S*, where S is as for the first instantiation. Since under the preconditions on N and C , S is total by Observation 3.7, the chosen r is again $\top\top$ -closed and convergence-preserving by Lemma 7.14. Then the second claim of the theorem follows from the definition of *build*, Lemma 7.4, and the compatibility of \sqsubseteq_{obs} , provided we can establish that $C \Downarrow$, that for every $A_2 \in \text{Term}(\tau)$, $(C A_2) \Downarrow$, that

$$\begin{aligned} \forall (A_1, A_2) \in \Delta_\tau(), B_1 \in \text{Term}(\tau\text{-list}), B_2 \in \text{Term}(\tau'). \\ (\text{foldr}_\tau)_{\tau'} C N B_1 \sqsubseteq_{\text{obs}} B_2 \\ \Rightarrow (\text{foldr}_\tau)_{\tau'} C N ((\lambda h :: \tau.\lambda t :: \tau\text{-list}.h : t) A_1 B_1) \sqsubseteq_{\text{obs}} C A_2 B_2, \end{aligned}$$

and that

$$(\text{foldr}_\tau)_{\tau'} C N \mathbf{nil}_\tau \sqsubseteq_{\text{obs}} N.$$

But the two ‘convergence conditions’ follow from the precondition on C by Corollary 4.21, whereas the other two statements follow from Corollary 6.8, the definition of *foldr*, Lemma 7.4, and the compatibility of \sqsubseteq_{obs} . \blacksquare

Note that the above proof relies on Corollary 6.8 rather than just on the weaker Theorem 6.4. This is exactly where the proofs in [JV04, JV06] had to resort to an unproved conjecture, because the statement corresponding to the identity extension lemma for the (denotational-style) logical relation offered there could not be shown to hold. This has now been remedied in the operational setting. The preconditions on N and C imposed for total correctness in Theorem 7.19 of course closely correspond to the denotational ones suggested for *foldr/build*-fusion in [JV04] and [JV06]. For pragmatic discussions of the implications such preconditions have for short cut fusion and related transformations in practice, consult the latter paper.

8 Conclusion and related work

We have constructed a parametric model of observational approximation for a non-strict polymorphic lambda calculus with general recursion, an algebraic datatype, and selective strictness. This puts earlier — more intuition-based than formally-derived — accounts [JV04, JV06] of free theorems and parametricity-based program transformations for functional languages mixing all these features, like Clean and Haskell, on a firm theoretical basis.

In retrospect, an interesting relationship can be observed between the restrictions introduced in [JV04, JV06] to accommodate selective strictness — or, more precisely, their operational-style incarnations coming forward in the present paper — and ideas from both Pitts’ original paper on PolyPCF [Pit00] and his extended study in [Pit05] of what is essentially ‘Call-by-value PolyPCF’. Firstly, Figure 9 of [Pit00] proposes using roughly the following as the relational interpretation of function types in a call-by-name calculus ‘Lazy PCF’ in which termination at function types is observable:

$$\{(\lambda x :: \tau_1.M, \lambda x :: \tau'_1.M') \mid \forall(A, A') \in r_1. (M[A/x], M'[A'/x]) \in r_2\}^{\top\top}$$

Since function abstractions are values, and thus converge, the relation of which Pitts’ (now bidirectional) $\top\top$ -closure is taken here is clearly convergence-reflecting, i.e., bidirectionally convergence-preserving. By the analogue of our Lemma 4.18 for the equational setting (which is known to hold for PolySeq [VJ06] and expected to hold for ‘Lazy PCF’ as well), the proposed relational interpretation of function types is then itself also convergence-reflecting. So the above is quite reminiscent of $r_1 \rightarrow r_2$ in our Definition 5.1. For a calculus incorporating impredicative polymorphic types and observable termination at those types — i.e., for ‘Lazy PolyPCF’ rather than ‘Lazy PCF’ — one would of course have to ensure convergence-preservation or -reflection of the relational interpretation of \forall -types as well, either explicitly as we do in Definition 5.2 or more indirectly in a manner similar to that above. But even then, one would not yet have a logical relation appropriate for PolySeq. For while the aforementioned adaptations account for observability of whole program termination at arbitrary types in ‘Lazy PolyPCF’, they do not fully capture the impact of selective strictness in PolySeq, namely that termination of arbitrary intermediate computations becomes observable as well. To account for that, we have

imposed convergence-preservation not only on the *result* of each relational action, but also on the relations over which we *quantify* in the relational action for \forall -types. And interestingly, something similar happens in Pitts' account for 'Call-by-value PolyPCF' in [Pit05]. The enforcement of convergence-reflection is again done quite indirectly rather than in the direct way we have preferred in the present paper (and which is indeed preferable when it comes to applying the logical relation as in the previous section), but the basic idea is the same. Nevertheless, as a whole, the logical relation from [Pit05] is not appropriate for PolySeq. In particular, the relational action for function types given there requires relatedness of function results only for related function arguments that are values. Regarding the extensionality principle for function types, this would imply that $A \in \text{Term}(\tau_1)$ in Pitts' (equational) analogue of our Lemma 7.6 is replaced by $A \in \text{Value}(\tau_1)$. And while this gives a correct principle for a purely strict calculus, it is wrong for PolySeq. That the logical relation characterizing PolySeq observational approximation — and thus equivalence — can ultimately be understood as a kind of blend of logical relations for 'Lazy PolyPCF' and 'Call-by-value PolyPCF' is not entirely surprising, since it corresponds to the fact that intuitively 'Haskell with *seq*' is situated somewhere between 'Haskell without *seq*' and a purely strict language like ML. But, of course, finding just the right blend, or equivalently, identifying the precise position selective strictness occupies between pure nonstrictness and pure strictness was the task we set out to solve with this paper. Interestingly, the picture drawn up above finds a complement, and is tied back to denotational semantics, in recent domain-theoretic work by Møgelberg [Møg06]. There, he sketches a program logic for a polymorphic call-by-value calculus with recursive types and terms and shows it to give rise to a computationally adequate model suitable for deriving consequences of parametricity along the lines of those in [Pit05]. And for the relational interpretations of types he arrives at something similar to our restrictions on relations in [JV04, JV06].

Pitts' study of parametric polymorphism for the purely strict version of PolyPCF was mainly motivated by concerns regarding existential types (\exists -types, see [MP88]) rather than \forall -types. A similar motivation underlies other recent work on operationally-based logical relations for purely strict calculi [Ahm06]. Hence, it would be interesting to see how things change when \exists -types are integrated into PolySeq as well. It is easy to see that the observationally isomorphic encoding of \exists -types by \forall -types provided for purely nonstrict PolyPCF in Section 7 of [Pit00] breaks down in PolySeq (just as the encoding of list types from Example 2.8 of that paper does). So if we want to reason about existential types, they must be explicitly added to the calculus, and their interplay with selective strictness must be studied very carefully. This is ongoing work.

Acknowledgments

We thank the anonymous reviewers for their suggestions for improving the presentation of the paper. We particularly thank the reviewer who encouraged us to move from the equational to the inequational setting in this paper.

References

- [ABB⁺05] A. Abel, M. Benke, A.L. Bove, R.J.M. Hughes, and U. Norell. Verifying Haskell programs using constructive type theory. In *Haskell Workshop, Proceedings*, pages 62–73. ACM Press, 2005.
- [Ahm06] A. Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *European Symposium on Programming, Proceedings*, volume 3924 of *LNCS*, pages 69–83. Springer-Verlag, 2006.
- [BFSS90] E.S. Bainbridge, P.J. Freyd, A. Scedrov, and P.J. Scott. Functorial polymorphism. *Theoretical Computer Science*, 70:35–64, 1990. Corrigendum *Ibid.*, 71:431, 1990.
- [BMP06] L. Birkedal, R.E. Møgelberg, and R.L. Petersen. Parametric domain-theoretic models of polymorphic intuitionistic / linear lambda calculus. In *Mathematical Foundations of Programming Semantics 2005, Proceedings*, volume 155 of *ENTCS*, pages 191–217. Elsevier B.V., 2006.
- [CLR] Clean Language Report 2.0 (<http://www.cs.ru.nl/~clean/CleanExtra/report20/index.html>).
- [DHJG06] N.A. Danielsson, R.J.M. Hughes, P. Jansson, and J. Gibbons. Fast and loose reasoning is morally correct. In *Principles of Programming Languages, Proceedings*, pages 206–217. ACM Press, 2006.
- [DJ04] N.A. Danielsson and P. Jansson. Chasing bottoms — A case study in program verification in the presence of partial and infinite values. In *Mathematics of Program Construction, Proceedings*, volume 3125 of *LNCS*, pages 85–109. Springer-Verlag, 2004.
- [EM06] M. van Eekelen and M. de Mol. Proof tool support for explicit strictness. In *Implementation and Application of Functional Languages 2005, Selected Papers*, volume 4015 of *LNCS*, pages 37–54. Springer-Verlag, 2006.
- [FFKD87] M. Felleisen, D.P. Friedman, E. Kohlbecker, and B. Duba. A syntactic theory of sequential control. *Theoretical Computer Science*, 52:205–237, 1987.
- [Fri75] H. Friedman. Equality between functionals. In *Logic Colloquium '72–73, Proceedings*, pages 22–37. Springer-Verlag, 1975.
- [GHC] The Glasgow Haskell Compiler (<http://www.haskell.org/ghc>).
- [Gir72] J.-Y. Girard. *Interprétation fonctionnelle et élimination des coupures dans l'arithmétique d'ordre supérieure*. PhD thesis, Université Paris VII, 1972.

- [GJUV05] N. Ghani, P. Johann, T. Uustalu, and V. Vene. Monadic augment and generalised short cut fusion. In *International Conference on Functional Programming, Proceedings*, pages 294–305. ACM Press, 2005.
- [GLP93] A. Gill, J. Launchbury, and S.L. Peyton Jones. A short cut to deforestation. In *Functional Programming Languages and Computer Architecture, Proceedings*, pages 223–232. ACM Press, 1993.
- [Has91] R. Hasegawa. Parametricity of extensionally collapsed term models of polymorphism and their categorical properties. In *Theoretical Aspects of Computer Software, Proceedings*, volume 526 of *LNCS*, pages 495–512. Springer-Verlag, 1991.
- [HK05] W.L. Harrison and R.B. Kieburtz. The logic of demand in Haskell. *Journal of Functional Programming*, 15:837–891, 2005.
- [HML] The Haskell Mailing List Archive (<http://www.mail-archive.com/haskell@haskell.org>).
- [HS97] R. Harper and C. Stone. An interpretation of Standard ML in type theory. Technical Report CMU-CS-97-147, Carnegie Mellon University, 1997.
- [Joh02] P. Johann. A generalization of short-cut fusion and its correctness proof. *Higher-Order and Symbolic Computation*, 15:273–300, 2002.
- [Joh03] P. Johann. Short cut fusion is correct. *Journal of Functional Programming*, 13:797–814, 2003.
- [Joh05] P. Johann. On proving the correctness of program transformations based on free theorems for higher-order polymorphic calculi. *Mathematical Structures in Computer Science*, 15:201–229, 2005.
- [JV04] P. Johann and J. Voigtländer. Free theorems in the presence of *seq*. In *Principles of Programming Languages, Proceedings*, pages 99–110. ACM Press, 2004.
- [JV06] P. Johann and J. Voigtländer. The impact of *seq* on free theorems-based program transformations. *Fundamenta Informaticae*, 69:63–102, 2006.
- [Las98] S.B. Lassen. *Relational Reasoning about Functions and Nondeterminism*. PhD thesis, University of Aarhus, 1998.
- [Lei83] D. Leivant. Polymorphic type inference. In *Principles of Programming Languages, Proceedings*, pages 88–98. ACM Press, 1983.
- [LP96] J. Launchbury and R. Paterson. Parametricity and unboxing with unpointed types. In *European Symposium on Programming, Proceedings*, pages 204–218. Springer-Verlag, 1996.

-
- [Møg06] R.E. Møgelberg. Interpreting polymorphic FPC into domain theoretic models of parametric polymorphism. In *International Colloquium on Automata, Languages and Programming, Proceedings*, volume 4052 of *LNCS*, pages 372–383. Springer-Verlag, 2006.
- [MP88] J.C. Mitchell and G.D. Plotkin. Abstract types have existential type. *ACM Transactions on Programming Languages and Systems*, 10:470–502, 1988.
- [Pey03] S.L. Peyton Jones, editor. *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, 2003.
- [Pit00] A.M. Pitts. Parametric polymorphism and operational equivalence. *Mathematical Structures in Computer Science*, 10:321–359, 2000.
- [Pit05] A.M. Pitts. Typed operational reasoning. In B.C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, pages 245–289. MIT Press, 2005.
- [Plo73] G.D. Plotkin. Lambda-definability and logical relations. Memorandum SAI-RM-4, University of Edinburgh, 1973.
- [Plo04] G.D. Plotkin. A structural approach to operational semantics. *Journal of Logic and Algebraic Programming*, 60-61:17–139, 2004.
- [PLST98] S.L. Peyton Jones, J. Launchbury, M.B. Shields, and A.P. Tolmach. Bridging the gulf: A common intermediate language for ML and Haskell. In *Principles of Programming Languages, Proceedings*, pages 49–61. ACM Press, 1998.
- [Rey74] J.C. Reynolds. Towards a theory of type structure. In *Colloque sur la Programmation, Proceedings*, pages 408–423. Springer-Verlag, 1974.
- [Rey83] J.C. Reynolds. Types, abstraction and parametric polymorphism. In *Information Processing, Proceedings*, pages 513–523. Elsevier Science Publishers B.V., 1983.
- [RMP06] B. Rudiak-Gould, A. Mycroft, and S.L. Peyton Jones. Haskell is not not ML. In *European Symposium on Programming, Proceedings*, volume 3924 of *LNCS*, pages 38–53. Springer-Verlag, 2006.
- [RR94] E.P. Robinson and G. Rosolini. Reflexive graphs and parametric polymorphism. In *Logic in Computer Science, Proceedings*, pages 364–371. IEEE Computer Society Press, 1994.
- [RS04] G. Rosolini and A. Simpson. Using synthetic domain theory to prove operational properties of a polymorphic programming language based on strictness. Manuscript, 2004.

- [Sta85] R. Statman. Logical relations and the typed lambda-calculus. *Information and Control*, 65:85–97, 1985.
- [Sve02] J. Svenningsson. Shortcut fusion for accumulating parameters & zip-like functions. In *International Conference on Functional Programming, Proceedings*, pages 124–132. ACM Press, 2002.
- [THLP98] P.W. Trinder, K. Hammond, H.-W. Loidl, and S.L. Peyton Jones. Algorithm + Strategy = Parallelism. *Journal of Functional Programming*, 8:23–60, 1998.
- [TM95] A. Takano and E. Meijer. Shortcut deforestation in calculational form. In *Functional Programming Languages and Computer Architecture, Proceedings*, pages 306–313. ACM Press, 1995.
- [VJ06] J. Voigtländer and P. Johann. Selective strictness and parametricity in structural operational semantics. Technical Report TUD-FI06-02, Technische Universität Dresden, 2006.
- [Voi02] J. Voigtländer. Concatenate, reverse and map vanish for free. In *International Conference on Functional Programming, Proceedings*, pages 14–25. ACM Press, 2002.
- [VWP06] D. Vytiniotis, S. Weirich, and S.L. Peyton Jones. Boxy types: inference for higher-rank types and impredicativity. In *International Conference on Functional Programming, Proceedings*, pages 251–262. ACM Press, 2006.
- [Wad89] P. Wadler. Theorems for free! In *Functional Programming Languages and Computer Architecture, Proceedings*, pages 347–359. ACM Press, 1989.
- [WB89] P. Wadler and S. Blott. How to make *ad-hoc* polymorphism less *ad hoc*. In *Principles of Programming Languages, Proceedings*, pages 60–76. ACM Press, 1989.