

# Two Parametricities versus Three Universal Types\*

DOMINIQUE DEVRIESE, Vrije Universiteit Brussel, Belgium

MARCO PATRIGNANI, Stanford University & CISPA Helmholtz Center for Information Security, USA & Germany

FRANK PIESENS, KU Leuven, Belgium

The formal study of programming languages makes heavy use of formal calculi: small but expressive language models which capture the essence of important programming language features. One of the most important such calculi is System F, co-discovered by Girard and Reynolds, which models the essence of polymorphism and abstract data types, features that are present in many programming languages. The core property of System F is parametricity: a formal theorem expressing the abstractions offered by the language and validating important principles like information hiding and modularity.

When System F is combined with other features such as general recursion, recursive types, mutable state or control flow primitives (e.g., continuations and exceptions), it is well known that the formulation of parametricity needs to be adapted to follow suit, for example using techniques like step-indexing, Kripke world-indexing or biorthogonality. However, it is less well understood what should happen with parametricity when System F is combined with untyped languages, gradual types, dynamic sealing and runtime type analysis (typecase) alongside type generation. Extensions of System F with these features have been proven to have a form of parametricity (e.g., parametricity with type worlds). However little is known about why exactly this modified parametricity works and why it is needed. Moreover, its relative power with respect to traditional parametricity and the relative expressiveness of System F with and without these extensions are also unknown.

This paper sheds light on these unknowns. We explain that the aforementioned different settings have a common idea: they combine System F with a *universal type*. We explain why exactly standard Reynolds parametricity is incompatible with such a type and how parametricity with type worlds is compatible with it. Building on these insights, we answer two open conjectures from the literature, negatively, and we point out a deficiency in current proposals for combining System F with gradual types.

CCS Concepts: • **Security and privacy** → **Formal security models**; *Logic and verification*; • **Theory of computation** → *Logic and verification*.

Additional Key Words and Phrases: Fully abstract compilation, System F, sealing, parametricity, universal type

## ACM Reference Format:

Dominique Devriese, Marco Patrignani, and Frank Piessens. 2019. Two Parametricities versus Three Universal Types. *J. ACM* 0, 0, Article 0 (2019), 36 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

\*This paper extends the previous version by Devriese et al. [2018] by (1) explaining the difference between Reynolds-style parametricity and parametricity with type worlds, (2) proving the degeneracy of Univ in a simpler way and (3) by disproving the Neis-Dreyer-Rossberg conjecture about fully abstractly embedding System F into a language with a dynamic type-case [Neis et al. 2009] (Section 5).

---

Authors' addresses: Dominique Devriese, Vrije Universiteit Brussel, Brussels, Belgium, name.surname@vub.be; Marco Patrignani, cs, Stanford University & CISPA Helmholtz Center for Information Security, USA & Germany, mp@cs.stanford.edu; Frank Piessens, imec-DistriNet, KU Leuven, Leuven, Belgium, name.surname@cs.kuleuven.be;

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2019 Association for Computing Machinery.

0004-5411/2019/0-ART0 \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

*To better explain and clarify notions, this paper uses colours. Please print or view it in colours.*

## 1 INTRODUCTION

System F is a widely influential type system, originally defined by Reynolds [1974] and Girard [1972], featuring parametric polymorphism and an impredicative universe. The core property of System F is *parametricity*. Parametricity guarantees that polymorphic functions in System F cannot behave differently when invoked at different types. Parametricity is formalised using a logical relation: an (often relational) property about values and terms of a certain type [Reynolds 1983]. Parametricity then states that the LR properties are automatically satisfied by any term, without the need to verify their code. For this reason, Wadler [1989] has described them as free theorems. A canonical example is the fact that any value of type  $\forall X. X \rightarrow X$  must behave as the identity function.

It is well known that the meaning of parametricity must be adapted when additional features are added to System F. For example, consider a value  $f$  of type  $\forall X. X \rightarrow X$ . In vanilla System F, it must behave as the identity function, i.e., return its argument in every invocation. If we add general recursion, then there is another possibility:  $f$  can also be the function that diverges on every invocation. Additionally adding mutable state variables changes the situation again: it is now possible that  $f$  returns its argument in some invocations and diverges in others, even though the choice can still not depend on the argument. And if we add continuations, then it may be possible for  $f$  to return more than once, but still: on every return, the return value must be the one it received as an argument. These intuitive differences are reflected formally in a different definition of the logical relation used to define parametricity. In other words, parametricity is a property that does not merely express the absence of certain undesired behaviour, but instead, it captures semantic guarantees about the language as a whole.

In many cases like the above, it is well understood why adding a certain feature to System F requires a particular change in the formulation of parametricity. However, there are also cases where a certain formulation of parametricity is used because it works, but it is not well understood why that formulation is needed. Specifically, in this paper, we look at the work by Sumii and Pierce [2003] on logical relations for encryption, the work by Neis et al. [2009, 2011] on parametricity in the presence of a dynamic type inspection primitive and runtime type generation, and the work by Ahmed et al. [2017] and Toro et al. [2019] on parametricity in gradually-typed variants of System F. Although they work in different contexts, these all prove (a form of) parametricity with a logical relation indexed by System F types (or a superset of these types). Additionally, the logical relation used by these parametricity results share a particular, non-standard structure: they are indexed by a *type world*. In traditional, Reynolds-style logical relations (RLRs), semantic interpretations for opaque type variables are represented in a type environment (see Section 2.3). However, the aforementioned logical relations use type worlds to carry these interpretations (we refer to them as type-world logical relations, TWLRs).

In this paper, we clarify that the reason these three different domains require the same kind of logical relation is the fact that they all extend System F with some form of universal type. By this, we mean a type which all other types can be embedded into and extracted from again. In Section 3, we will show that RLRs are incompatible with such universal types and that TWLRs remove this incompatibility.

The poor understanding of the need for TWLRs can be illustrated in two ways. First, if we look at the history of one of the aforementioned results [Ahmed et al. 2017], then we see that a RLR was actually used in a precursor of the work [Matthews and Ahmed 2008]. However, a flaw was later discovered in this proof Ahmed et al. [2017]. Few details are available about this flaw but we suspect it was a consequence of the fundamental incompatibility of RLR with universal types, that we will explain later. This flaw was then resolved by the authors in an unpublished draft by moving

to a TWLR [Ahmed et al. 2011c] and the whole effort culminated in the mature TWLR of Ahmed et al. [2017].

The second consequence of the poor understanding of the link between TWLRs and universal types is that researchers had wrong expectations of how program equivalence changes when adding a universal type to System F. Concretely, our insights allow us to disprove two conjectures made by experts in published literature. These conjectures are respectively related to *secure compilation* of System F using dynamic sealing (read, idealised encryption) and to the *enforcement of parametricity* in the presence of dynamic type analysis and runtime type generation. Additionally, we also identify a concern in the *interaction between gradually typed languages and polymorphic types*. Let us take a closer look at these three topics.

*Secure compilation.* The field of secure compilation studies high-level programming languages that are compiled to low-level target languages where they may interact with untrusted target-level components. The goal of secure compilation is to ensure that those target-level components can only interact with the compiled code in ways that high-level components can interact with the original code. This constitutes a powerful security property, as it effectively excludes a wide variety of low-level attacks like improper stack manipulation, breaking control flow guarantees, reading from or writing to private memory of other components, inspecting or modifying the implementation of a function etc.

Formally, this formulation of secure compilation has been expressed as full abstraction: given two source-level contextually equivalent programs, their target-level compilation are also contextually equivalent (and vice versa) [Abadi 1998].<sup>1</sup> Compiler full abstraction has been proven for compilers that rely on address-space layout randomisation [Abadi and Plotkin 2012; Jagadeesan et al. 2011], or secure enclaves [Agten et al. 2012; Larmuseau et al. 2015, 2016; Patrignani et al. 2015, 2016], tagged architectures [Juglaret et al. 2015, 2016], JavaScript [Fournet et al. 2013], typed closure conversion [Ahmed and Blume 2008], cryptographic primitives [Abadi et al. 1998, 1999, 2000; Bugliesi and Giunti 2007] etc, we refer the interested reader to the survey of Patrignani et al. [2019].

In a paper from the year 2000, Pierce and Sumii [2000] proposed a compiler from System F to a cryptographic lambda calculus, which enforces parametricity using a form of idealised encryption primitives (sealing) called lambda-seal ( $\lambda^\sigma$ ). They conjectured that this compiler was fully abstract, a conjecture that has received further research attention, but remains open to this day [Siek and Wadler 2016; Sumii and Pierce 2004]. In other work, the same authors proposed TWLR for the cryptographic lambda calculus [Sumii and Pierce 2003] but they lacked the main insight from this paper: namely that they need this TWLR to accommodate the existence of universal types which is contradicted by regular RLRs.

*Non-Parametric Parametricity.* Some modern programming languages include a way to perform *intensional type analysis* through a type cast operator [Abadi et al. 1995; Rossberg 2003]. This appears to be in direct conflict with parametric polymorphism, possibly violating parametricity and representation independence guarantees [Mitchell 1986]. Researchers have proposed runtime type generation as a way to regain parametricity for languages with a type cast operator. Ideally, when an abstract type is defined, a fresh type name should also be generated at runtime. Such a name should then be used as a runtime representative of the abstract type for type analysis purposes.

Runtime type generation has been proven to indeed provide parametricity guarantees for System F terms that interact with terms that can perform type casts [Neis et al. 2009]. This guarantee has been dubbed *non-parametric parametricity*.

<sup>1</sup>Other formal characterisations of secure compilation have been proposed more recently [Abate et al. 2018, 2019; Patrignani and Garg 2017, 2019]. We discuss their relation to our work in Section 8.

Neis et al. [2009] also conjectured that their way of using runtime type generation preserves *any* System F abstraction in their type case language. Albeit they do not explain the reasons behind this conjecture, we think we understand the intuition behind it. Parametricity is the most powerful abstraction programmers think that System F has, so it seems logical to believe that once that has been preserved, all other abstractions will follow. These authors also use a TWLR for stating parametricity in their system, but they did not see a subtlety that we highlight here: the resulting parametricity is (1) weaker than a more traditional RLR-based version and (2) RLR-based parametricity conflicts with the universal type implied by their type-cast operator.

*Gradual typing.* The final field where we contribute new insights is gradual typing. In order to allow programmers to incrementally migrate large, untyped code bases to a typed programming language, gradual programming languages allow for typed and untyped code to interact. Such languages generally strive to preserve the benefits of the statically-typed components of an application, even when interacting with untyped components. Such benefits include performance benefits, absence of runtime type errors but also benefits for reasoning. While the literature mentions several ways to formalise the former two properties, the latter has received less attention.

A natural way to formally express that a gradual language preserves the reasoning principles of the typed language is to reuse the same notion of fully abstract compilation that we mentioned above. Specifically, if the embedding of the typed language into the gradual language is fully abstract, then similarly to secure compilation, this expresses that untyped code can only interact with typed code in ways that are also possible using just typed code.<sup>2</sup>

Also in the field of gradual typing, System F's parametric polymorphism presents a formidable challenge. The observation that sealing could be useful to combine parametric polymorphism with dynamic typing was already made by Pierce and Sumii [2000]. This idea was further developed with the definition of a gradually-typed language based on this idea [Ahmed et al. 2011c; Matthews and Ahmed 2008], the addition of blame in the polymorphic blame calculus [Ahmed et al. 2011b], further developed by Igarashi et al. [2017], Ahmed et al. [2017] and Toro et al. [2019]. The latter two papers also prove parametricity by relying on TWLR in the polymorphic blame calculus. Also in these cases there is no clear mention of the implications that relying on TWLR has as opposed to relying on RLR for the kind of parametricity that these work obtain.

*Three questions, one answer.* In this paper, thanks to our understanding of parametricity with universal types, we solve these open problems, answering the questions negatively. We prove that Sumii and Pierce's compiler is not fully abstract, and that System F does not embed fully abstractly into either Neis et al.'s language with a type-case, or any of the published polymorphic blame calculi.

Our counterexamples are essentially based on the observation that System F parametricity forbids the existence of a universal type: a type which any other type can be embedded into and extracted from.<sup>3</sup> More precisely, consider the following type:

$$\mathbf{Univ} \stackrel{\text{def}}{=} \exists Y. \forall X. (X \rightarrow Y) \times (Y \rightarrow X)$$

<sup>2</sup>This is related to the notion that fully abstract compilation can also be used to reason about language expressiveness in general, beyond just language security (which is what is done in the case of secure compilation) [Felleisen 1991; Mitchell 1993; Parrow 2008].

<sup>3</sup>Note that by this definition, the term *universal type* is broader than Abadi's *dynamic type* [Abadi et al. 1991] or Siek and Taha's *gradual type* [Siek and Taha 2006]: the former includes *any* type that arbitrary values can be embedded into and extracted from, while the latter are specific primitive types, specifically intended for representing untyped values in a typed language.

**Univ** can be read as expressing the existence of a universal type: a type **Y** such that for any other type **X**, there is a mapping from **X** into **Y** and vice versa. In a non-terminating variant of System F, there exist inhabitants of **Univ**, but we prove that they are all degenerate in the sense that mapping a value into the universal type and back must necessarily diverge. This degeneracy follows from the RLR, illustrating the incompatibility of RLRs with universal types.

However, Sumii and Pierce’s compiler fails to enforce this degeneracy of **Univ**. In fact, their target language really does contain a universal type: since it is untyped we can think about it as being “uni-typed”, citing Dana Scott [Statman 1991]. As a consequence, their fully abstract compilation conjecture is false, as we will formally show by constructing two System F terms  $t_u$  and  $t_\omega$  whose contextual equivalence relies on the degeneracy of **Univ**. We can then falsify Sumii and Pierce’s conjecture by showing that these two terms are mapped to non-equivalent terms by their proposed compiler.

In Neis et al.’s language with non-parametric parametricity, we show that type cast operators also break the degeneracy of **Univ**, despite the presence of runtime type generation primitives. The type  $\forall Z. Z$ , which is normally only inhabited by diverging terms, becomes a universal type in the presence of a dynamic type-case, as any value can be embedded in it and extracted from it. By relying on this type, System F does not embed fully abstractly into this language, contradicting Neis et al.’s conjecture.

Finally, in the field of gradual typing, we demonstrate that existing polymorphic blame calculi also break the degeneracy of **Univ**. Like Sumii and Pierce’s target language, they also provide a universal type: the type of untyped values  $\star$ , common to most gradual languages (also often indicated as  $?$ ). Exploiting the existence of this type, we demonstrate that the polymorphic blame calculus does not embed System F in a fully abstract way and as a result, they do not preserve System F’s parametricity.

To close off, we discuss a number of consequences and perspectives that follow from our results. First, we discuss some thoughts on how Sumii and Pierce’s compiler might be fixed (so that it does enforce full abstraction), and what could be modified in the polymorphic blame calculus to make it preserve all of System F’s contextual equivalences. However, in neither case there appears to be a panacea solution: potential fixes all seem to come with certain downsides. Because of this, we also discuss whether we should not instead adjust our expectations and find a way to formalise the guarantees that we do get from both Sumii and Pierce’s compiler, Neis et al.’s sealing wrappers and the polymorphic blame calculus.

*Outline.* We start our discussion by repeating the definition of System F and defining two logical relations for it, which we will use in subsequent proofs (Section 2). Then we present type **Univ** and two key terms:  $t_u$  and  $t_\omega$ . We prove that **Univ** is degenerate and that these terms are contextually equivalent using the previously-defined logical relations (Section 3). Next, we disprove Sumii and Pierce’s conjecture by explaining their compiler and how it treats  $t_\omega$  and  $t_u$  and fails to preserve their equivalence (Section 4). We then present **G**, an extension of System F with type casts and runtime type generation and demonstrate how embedding  $t_u$  and  $t_\omega$  into **G** breaks contextual equivalence (Section 5). Next, we turn to gradual typing, introducing polymorphic blame calculi and demonstrating how  $t_u$  and  $t_\omega$ ’s contextual equivalence is also lost in the presence of these calculi’s universal type (Section 6). Finally, we discuss perspectives and consequences of our results (Section 7), related work (Section 8) and we conclude (Section 9). For space constraints we omit most proofs and formalisation, which can be found in the supplementary material.

To make the distinction between languages visually apparent, we typeset elements of System F in a **blue, bold** font, elements of  $\lambda^\sigma$  in a **red, sans-serif** font, elements of **G** in a **emerald, verbatim**

font and elements of the blame calculus in an *orange, italic* font. This kind of syntax highlighting has been proven effective for colourblind and black&white readers too.

## 2 SYSTEM F AND THE TYPE Univ

Let us first consider System F itself, its parametricity and our claim that parametricity excludes the existence of a universal type.

### 2.1 The Source Language $\lambda^F$

Figure 1 presents the variant of System F that we will be using in this paper, which we indicate with  $\lambda^F$ . In addition to standard polymorphic functions ( $\forall X. \tau$ ) and existential packages ( $\exists X. \tau$ ), the variant includes recursive types  $\mu X. \tau$  (so there are no termination guarantees), as well as **Unit** and **Bool** and product  $\tau_1 \times \tau_2$  types. In the figure, we present terms  $t$ , values  $v$  and types  $\tau$ . We show the most important typing rules  $\Delta; \Gamma \vdash t : \tau$  in terms of term and type variable contexts  $\Gamma$  and  $\Delta$  (but we omit context and type well-formedness judgements  $\Delta; \Gamma \vdash \diamond$  and  $\Delta \vdash \tau$ ). Finally, we define call-by-value evaluation rules in terms of evaluation contexts  $E$ .

Program contexts  $C$  are defined as terms with exactly one subterm replaced by a hole  $[\cdot]$ . An omitted well-typedness judgement for program contexts  $C : \Delta; \Gamma; \tau \rightarrow \Delta'; \Gamma'; \tau'$  guarantees that plugging a well-typed term  $\Delta; \Gamma \vdash t : \tau$  in the hole produces the well-typed resulting term  $\Delta'; \Gamma' \vdash C[t] : \tau'$ .

*Definition 2.1 ( $\lambda^F$  Contextual equivalence).* For two terms  $t_1, t_2$  that have the same type  $\tau$  in the same context  $\Delta; \Gamma$ , we define that they are contextually equivalent ( $\Delta; \Gamma \vdash t_1 \simeq t_2 : \tau$ ) iff for all  $C$  such that  $\vdash C : \Delta; \Gamma, \tau \rightarrow \emptyset; \emptyset, \tau'$ , we have that  $C[t_1] \uparrow$  iff  $C[t_2] \uparrow$ , where  $\uparrow$  indicates divergence [Plotkin 1977].

### 2.2 Parametricity

The main information hiding mechanism in  $\lambda^F$  is parametric polymorphism: the representation type of an existentially quantified package is invisible outside the package, and hence clients of the package cannot depend on that representation type.

*Example 2.2 ( $\mathbb{Z}_n$  implementation in  $\lambda^F$ ).* We could for instance represent the type  $\mathbb{Z}_n$  of integers modulo  $n$  as a tuple  $\langle \langle \text{zero}, \text{succ} \rangle, \text{zero?} \rangle$  of type  $\exists X. X \times (X \rightarrow X) \times (X \rightarrow \text{Bool})$ , and then either implement this type as a  $k$ -tuple of booleans, or by means of a type of natural numbers defined using recursive types.

*Example 2.3 (Example polymorphic function type in  $\lambda^F$ ).* Dually, the type system ensures that code cannot depend on parameters of universally quantified types, for instance the only thing a function of type  $\forall X. X \times X \rightarrow X$  can do is return one of its two arguments or diverge.

Reynolds formalised (relational) parametricity in the form of a theorem that all  $\lambda^F$  terms of a certain type satisfy a property that can be derived from their type [Reynolds 1983]. For example, if we assume a value  $f$  of type  $\forall X. X \rightarrow X$ , then parametricity states that for any closed types  $\tau_1, \tau_2$ , any relation  $R$  between values of types  $\tau_1$  and  $\tau_2$ , and any two closed values  $v_1, v_2$  of type  $\tau_1$  and  $\tau_2$  respectively, if  $(v_1, v_1)$  is in  $R$ , then  $f \tau_1 v_1$  and  $f \tau_2 v_2$  will either both diverge or reduce to values  $(v'_1, v'_2) \in R$ . The relational property is derived from the type using what is known as a logical relation.

## Syntax:

$$\begin{aligned}
t &::= v \mid x \mid tt \mid t.1 \mid t.2 \mid \langle t, t \rangle \mid t \tau \mid \text{if } t \text{ then } t \text{ else } t \\
&\quad \mid \text{pack } \langle \tau, t \rangle \text{ as } \exists X. \tau \mid \text{unpack } t \text{ as } \langle X, x \rangle \text{ in } t \mid \text{roll } t \mid \text{unroll } t \\
\text{Val } \ni v &::= \text{unit} \mid \text{true} \mid \text{false} \mid \lambda x : \tau. t \mid \langle v, v \rangle \mid \Lambda X. t \mid \text{pack } \langle \tau, v \rangle \text{ as } \exists X. \tau \mid \text{roll } v \\
\tau &::= \text{Unit} \mid \text{Bool} \mid \tau \rightarrow \tau \mid \tau \times \tau \mid X \mid \forall X. \tau \mid \exists X. \tau \mid \mu X. \tau \\
\Gamma &::= \emptyset \mid \Gamma, (x : \tau) \quad \Delta ::= \emptyset \mid \Delta, X \\
E &::= [\cdot] \mid E t \mid v E \mid E.1 \mid E.2 \mid \langle E, t \rangle \mid \langle v, E \rangle \mid E \tau \mid E; t \\
&\quad \mid \text{if } E \text{ then } t \text{ else } t \mid \text{pack } \langle \tau, E \rangle \text{ as } \exists X. \tau \\
&\quad \mid \text{unpack } E \text{ as } \langle X, x \rangle \text{ in } t \mid \text{unpack } v \text{ as } \langle X, x \rangle \text{ in } E \mid \text{roll } E \mid \text{unroll } E
\end{aligned}$$

## Typing rules (excerpts):

$$\begin{array}{c}
\frac{\Delta; \Gamma \vdash \diamond}{(\mathbf{x} : \tau) \in \Gamma} \quad \frac{\Delta; \Gamma, (\mathbf{x} : \tau) \vdash t : \tau'}{\Delta; \Gamma \vdash \lambda \mathbf{x} : \tau. t : \tau \rightarrow \tau'} \quad \frac{\Delta, X; \Gamma \vdash t : \tau}{\Delta; \Gamma \vdash \Lambda X. t : \forall X. \tau} \quad \frac{\Delta; \Gamma \vdash t : \tau' \rightarrow \tau}{\Delta; \Gamma \vdash t' : \tau'} \\
\frac{\Delta \vdash \tau'}{\Delta; \Gamma \vdash t \tau' : \tau[\tau'/X]} \quad \frac{\Delta; \Gamma \vdash t : \forall X. \tau}{\Delta; \Gamma \vdash \text{roll } t : \mu X. \tau} \quad \frac{\Delta; \Gamma \vdash t : \tau[\mu X. \tau/X]}{\Delta; \Gamma \vdash \text{unroll } t : \tau[\mu X. \tau/X]} \\
\frac{\Delta \vdash \tau}{\Delta; \Gamma \vdash \text{pack } \langle \tau', t \rangle \text{ as } \exists X. \tau : \exists X. \tau} \quad \frac{\Delta; \Gamma \vdash t : \exists X. \tau \quad \Delta \vdash \tau'}{\Delta; \Gamma \vdash \text{unpack } t \text{ as } \langle X, x \rangle \text{ in } t_1 : \tau'}
\end{array}$$

## Evaluation rules (excerpts):

$$\begin{array}{c}
\frac{t \hookrightarrow_0 t'}{E[t] \hookrightarrow E[t']} \quad \frac{}{(\lambda \mathbf{x} : \tau. t) v \hookrightarrow_0 t[v/x]} \quad \frac{}{(\Lambda X. t) \tau \hookrightarrow_0 t[\tau/X]} \\
\frac{}{\text{unpack } (\text{pack } \langle \tau', v \rangle \text{ as } \exists X. \tau) \text{ as } \langle X', x \rangle \text{ in } t \hookrightarrow_0 t[v/x][\tau'/X']} \\
\frac{}{\text{unroll } (\text{roll } v) \hookrightarrow_0 v}
\end{array}$$

Fig. 1. System F syntax, typing rules and evaluation rules (excerpts). The semantics relation  $\hookrightarrow$  relies on the primitive reductions indicated as  $\hookrightarrow_0$ .

### 2.3 Reynolds-style Logical Relation

As we explained in the introduction, this logical relation can be defined in various ways. For our purposes, it is instructive to first consider Reynolds' original definition. For simplicity and because it suffices for our purposes, we present a unary variant and omit recursive types.<sup>4</sup>

The Reynolds-style Logical Relation (RLR) defines a relation on values  $\mathcal{V}[\cdot]$  and on terms  $\mathcal{E}[\cdot]$ . Both are indexed with a type environment  $\rho$ , which maps type variables to a closed type  $\tau$  and a relation  $\mathbf{R}$  on values of type  $\tau$ . The value relation  $\mathcal{V}[\cdot]^\rho$  for type variables  $X$  defers to the

<sup>4</sup>We will continue to use words like relation, related etc. despite the fact that they are unary.

appropriate entry in  $\rho$  for type variables  $\mathbf{X}$ . Otherwise, the value relation accepts boolean and unit values when they are of canonical form and pairs when both components are in the appropriate relation themselves. For function types, the value relation accepts appropriately typed lambdas that map related values to related terms. Polymorphic functions are accepted when they can be applied to an arbitrary type  $\tau'$  and an arbitrary relation  $\mathbf{R}$  on  $\tau'$  to obtain a term in the appropriate term relation, with  $\rho$  extended with  $\mathbf{R}$ . Finally, existential packages must contain a type  $\tau'$  and a value related at the appropriate type, extending  $\rho$  with some relation  $\mathbf{R}$  on  $\tau'$ . The term relation accepts terms that diverge or produce a suitable value.

$$\begin{aligned}
\rho &\stackrel{\text{def}}{=} \{\mathbf{X} \mapsto (\tau, \mathbf{R}) \mid \mathbf{R} \in \text{Re1}(\tau)\} \\
\text{Re1}(\tau) &= \mathcal{P}(\{\mathbf{v} \mid \emptyset \vdash \mathbf{v} : \tau\}) \\
\mathcal{V}[\mathbf{X}]^\rho &= \rho(\mathbf{X}).\mathbf{R} \\
\mathcal{V}[\mathbf{Unit}]^\rho &= \{\mathbf{unit}\} \\
\mathcal{V}[\mathbf{Bool}]^\rho &= \{\mathbf{true}, \mathbf{false}\} \\
\mathcal{V}[\tau \rightarrow \tau']^\rho &= \{\lambda \mathbf{x} : \tau. \mathbf{t} \mid \forall \mathbf{v}. \text{ if } \mathbf{v} \in \mathcal{V}[\tau]^\rho \text{ then } \mathbf{t}[\mathbf{v}/\mathbf{x}] \in \mathcal{E}[\tau']^\rho\} \\
\mathcal{V}[\tau \times \tau']^\rho &= \{\langle \mathbf{v}, \mathbf{v}' \rangle \mid \mathbf{v} \in \mathcal{V}[\tau]^\rho \text{ and } \mathbf{v}' \in \mathcal{V}[\tau']^\rho\} \\
\mathcal{V}[\forall \mathbf{X}. \tau]^\rho &= \{\lambda \mathbf{X}. \mathbf{v} \mid \forall \mathbf{R} \in \text{Re1}(\tau'). \mathbf{v}[\tau'/\mathbf{X}] \in \mathcal{E}[\tau]^\rho, \mathbf{X} \mapsto (\tau', \mathbf{R})\} \\
\mathcal{V}[\exists \mathbf{X}. \tau]^\rho &= \{\text{pack } \langle \tau', \mathbf{v} \rangle \text{ as } \exists \mathbf{X}. \tau \mid \exists \mathbf{R} \in \text{Re1}(\tau'). \mathbf{v} \in \mathcal{V}[\tau]^\rho, \mathbf{X} \mapsto (\tau', \mathbf{R})\} \\
\mathcal{E}[\tau]^\rho &= \{\mathbf{t} \mid \text{ if } \mathbf{t} \hookrightarrow^* \mathbf{v} \text{ then } \mathbf{v} \in \mathcal{V}[\tau]^\rho\}
\end{aligned}$$

We can then define relations on environments  $\mathcal{G}[\cdot]$  and on type environments  $\mathcal{D}[\cdot]$ . The former accepts environments which map free variables to values in the appropriate value relation. The latter requires a type and a relation on that type for every free type variable.

$$\begin{aligned}
\mathcal{G}[\emptyset]^\rho &= \{\emptyset\} \\
\mathcal{G}[\Gamma, (\mathbf{x} : \tau)]^\rho &= \{\gamma; [\mathbf{v}/\mathbf{x}] \mid \gamma \in \mathcal{G}[\Gamma]^\rho \text{ and } \mathbf{v} \in \mathcal{V}[\tau]^\rho\} \\
\mathcal{D}[\emptyset] &= \{\emptyset\} \\
\mathcal{D}[\Delta; \alpha] &= \{\rho; \alpha \mapsto (\tau, \mathbf{R}) \mid \mathbf{R} = \text{Re1}(\tau)\}
\end{aligned}$$

With these ingredients, we can define the relation  $\Delta; \Gamma \Vdash \mathbf{t} : \tau$  on open terms. This relation accepts terms  $\mathbf{t}$  which are in the term relation after appropriately closing their free variables and type variables (Definition 2.4).

*Definition 2.4 (Reynolds-style Logical Relation).*

$$\Delta; \Gamma \Vdash \mathbf{t} : \tau \stackrel{\text{def}}{=} \forall \rho \in \mathcal{D}[\Delta], \forall \gamma \in \mathcal{G}[\Gamma]^\rho, \mathbf{t}\gamma \in \mathcal{E}[\tau]^\rho$$

We rely on a few results for the RLR, which are listed below. The fundamental property (Theorem 2.5) states that syntactically well-typed terms are in the relation (i.e., they are *semantically* well-typed).

**THEOREM 2.5 (FUNDAMENTAL PROPERTY FOR RLR).** *if  $\Delta; \Gamma \vdash \mathbf{v} : \tau$  then  $\Delta; \Gamma \Vdash \mathbf{v} : \tau$*

Proving the fundamental property relies on a number of standard lemmas. We only mention a few which we will need in the rest of this paper.

First, we have an antireduction lemma (Lemma 2.6) which states that a term  $t$  is in the term relation if a term  $t'$  that it reduces to is.

LEMMA 2.6 (ANTIREDUCTION). *If  $t \hookrightarrow^* t'$  and  $t' \in \mathcal{E} \llbracket \tau \rrbracket^\rho$ , then  $t \in \mathcal{E} \llbracket \tau \rrbracket^\rho$ .*

Next, we mention two compatibility lemmas: one for functions (Lemma 2.7) and one for applications (Lemma 2.8). Essentially, they state that lambdas and applications are in the logical relation if their subterms are (at appropriate types).

LEMMA 2.7 (COMPATIBILITY FOR FUNCTIONS).

$$\text{If } \Delta; \Gamma, x : \tau' \Vdash t : \tau \text{ then } \Delta; \Gamma \Vdash \lambda x : \tau'. t : \tau' \rightarrow \tau$$

LEMMA 2.8 (COMPATIBILITY FOR APPLICATIONS).

$$\text{If } \Delta; \Gamma \Vdash t : \tau' \rightarrow \tau \text{ and } \Delta; \Gamma \Vdash t' : \tau' \text{ then } \Delta; \Gamma \Vdash t t' : \tau$$

Finally, we mention the Boring lemma (Lemma 2.9)<sup>5</sup>, which states that the semantic type relation  $\rho$  can be altered freely as long as the type variables mentioned in the type  $\tau$  are left untouched.

LEMMA 2.9 (BORING LEMMA). *If  $\rho_1$  and  $\rho_2$  agree on the free type variables of  $\tau$ , then*

$$\mathcal{E} \llbracket \tau \rrbracket^{\rho_1} = \mathcal{E} \llbracket \tau \rrbracket^{\rho_2}$$

## 2.4 Logical Relation with Type Worlds

Not all logical relations follow the Reynolds-style, researchers have also used a different style of logical relation that uses type worlds (TWLR) [Ahmed et al. 2017; Neis et al. 2009; Sumii and Pierce 2003; Toro et al. 2019]. The crucial difference between this TWLR and the RLR of Section 2.3 is that TWLR uses *Kripke worlds*  $\mathbf{W}$  instead of semantic type relations  $\rho$  to represent the interpretations of type variables. Interestingly, this different treatment of type environments makes the logical relation compatible with universal types, as we will explain in Section 3.

Here, we define a unary logical relation with type worlds inspired by Ahmed et al.. As a notation convention, we typeset elements of the TWLR with a grey background, to distinguish them from elements of the RLR (which have no background).

Logical relations with type worlds are a special case of Kripke logical relations. Such logical relations use a world  $\mathbf{W}$  to carry assumptions, e.g., regarding a heap if the language has one. In a TWLR, the worlds are used to store the type  $\tau$  and predicate  $R$  on values of type  $\tau$  that are bound to a type variable, i.e., the information that was stored in  $\rho$  in the RLR (see Section 2.3). However, unlike  $\rho$ , Kripke LR, and TWLRs specifically, treat worlds *monotonically*. Intuitively, this means that once a binding is added to a world, it can never be removed. This monotonicity is formalised using a *future world relation* ( $\mathbf{W}' \sqsupseteq \mathbf{W}$ ), which defines that a world  $\mathbf{W}'$  is a future world of  $\mathbf{W}$  if the former extends the latter.

$$\mathbf{World} \ni \mathbf{W} = \emptyset \mid (\mathbf{W}; (X, \tau, R))$$

$$\mathbf{W}' \sqsupseteq \mathbf{W} = \mathbf{W}' \supseteq \mathbf{W}$$

$$\mathbf{W} + (X, \tau, R) = (\mathbf{W}; (X, \tau, R)) \quad \text{if } X \notin \text{dom}(\mathbf{W})$$

$$R \in \text{Rel}(\tau)$$

<sup>5</sup>For lack of a better name, we call this lemma as in Derek Dreyer's lecture notes [Dreyer et al. 2018]. Neis et al. [2011] call this the *irrelevance* lemma.

$$\text{Re1} [\tau] = \{\mathbf{R} \in \mathcal{P}(\mathbf{World} \times \mathbf{Val}) \mid \forall (\mathbf{W}, \mathbf{v}) \in \mathbf{R}. \forall \mathbf{W}' \supseteq \mathbf{W}. (\mathbf{W}', \mathbf{v}) \in \mathbf{R} \text{ and } \emptyset; \emptyset \vdash \mathbf{v} : \tau\}$$

The attentive reader may notice that our definition of worlds is actually cyclic: worlds map type variables to types and relations, and the relations are themselves world-indexed. As a result, the worlds presented here are not actually well-defined.

TODO: check how sumii and pierce solved the cyclicity.

Fortunately, this is a well-known problem and a standard solution exists: step-indexing [Ahmed et al. 2009a; Appel and McAllester 2001; Dreyer et al. 2011].

Essentially, the idea is to solve the cyclicity by indexing worlds with a number of steps which indicates up to which level they are defined. By carefully (and often tediously) ensuring that all world-indexed definitions only depend on the world up to a suitable number of steps, cyclic reasoning can be ruled out without otherwise changing the argumentation. Because step-indexing tends to complicate the technicalities while otherwise contributing little insight to the reasoning, we choose not to use it here. Instead, we simply ignore the problem in this text and stick to our illegal but comprehensible definition. To readers who wish to understand how the cyclicity can be solved without breaking the basic rules of mathematics, we recommend consulting Ahmed et al. [2017].

Having defined worlds, we can present the TWLR, which follows the same intuition of the RLR save for replacing  $\rho$  with worlds  $\mathbf{W}$ .

$$\mathcal{V} [\mathbf{X}] = \{(\mathbf{W}, \mathbf{v}) \mid (\mathbf{W}, \mathbf{v}) \in \mathbf{W}.\kappa(\mathbf{X})\}$$

$$\mathcal{V} [\mathbf{Unit}] = \{(\mathbf{W}, \mathbf{unit})\}$$

$$\mathcal{V} [\mathbf{Bool}] = \{(\mathbf{W}, \mathbf{true}), (\mathbf{W}, \mathbf{false})\}$$

$$\mathcal{V} [\tau \rightarrow \tau'] = \left\{ (\mathbf{W}, \lambda x : \tau. t) \mid \begin{array}{l} \forall \mathbf{W}' \supseteq \mathbf{W}. \forall \mathbf{v}. \\ \text{if } (\mathbf{W}', \mathbf{v}) \in \mathcal{V} [\tau] \text{ then } (\mathbf{W}', (\lambda x : \tau. t) \mathbf{v}) \in \mathcal{E} [\tau'] \end{array} \right\}$$

$$\mathcal{V} [\tau \times \tau'] = \{(\mathbf{W}, \langle \mathbf{v}, \mathbf{v}' \rangle) \mid (\mathbf{W}, \mathbf{v}) \in \mathcal{V} [\tau] \text{ and } (\mathbf{W}, \mathbf{v}') \in \mathcal{V} [\tau']\}$$

$$\mathcal{V} [\forall \mathbf{X}. \tau] = \left\{ (\mathbf{W}, \Lambda \mathbf{X}. \mathbf{v}) \mid \begin{array}{l} \forall \mathbf{W}' \supseteq \mathbf{W}. \forall \tau', \mathbf{R} \in \text{Re1} [\tau']. \\ (\mathbf{W}' + (\mathbf{X}, \tau', \mathbf{R}), \mathbf{v}[\tau'/\mathbf{X}]) \in \mathcal{E} [\tau] \end{array} \right\}$$

$$\mathcal{V} [\exists \mathbf{X}. \tau] = \left\{ (\mathbf{W}, \text{pack } \langle \tau', \mathbf{v} \rangle \text{ as } \exists \mathbf{X}. \tau) \mid \begin{array}{l} \exists \mathbf{R} \in \text{Re1} [\tau']. \exists (\mathbf{X}, \tau', \mathbf{R}) \in \mathbf{W}. \\ (\mathbf{W}, \mathbf{v}) \in \mathcal{V} [\tau] \end{array} \right\}$$

$$\mathcal{E} [\tau] = \{(\mathbf{W}, t) \mid \forall \mathbf{v}. \text{if } t \leftrightarrow^* \mathbf{v} \text{ then } \exists \mathbf{W}' \supseteq \mathbf{W}. (\mathbf{W}', \mathbf{v}) \in \mathcal{V} [\tau] \}$$

$$\mathcal{G} [\emptyset]^\rho = \{\emptyset\}$$

$$\mathcal{G} [\Gamma, (\mathbf{x} : \tau)]^\rho = \left\{ (\mathbf{W}, \gamma; [\mathbf{v}/\mathbf{x}]) \mid (\mathbf{W}, \gamma) \in \mathcal{G} [\Gamma]^\rho \text{ and } (\mathbf{W}, \mathbf{v}) \in \mathcal{V} [\tau]^\rho \right\}$$

The last piece we need to formalise is what it means for a world  $\mathbf{W}$  to agree with a type environment  $\Delta$ , which we denote with  $\mathbf{W} \vdash \Delta$ . This intuitively replaces the type environment relation  $\mathcal{D} [\cdot]$  present in RLR. A world agrees with a type environment when it maps all the type variables of the latter to valid relations, formally:

$$\emptyset \vdash \emptyset$$

$$\mathbf{W}, (\mathbf{X}, \tau, \mathbf{R}) \vdash \Delta, \mathbf{X} \text{ if } \mathbf{W} \vdash \Delta \text{ and } \mathbf{R} \in \text{Rel} [\tau]$$

With these ingredients we can define our Type-World Logical Relation.

*Definition 2.10 (Type-World Logical Relation).*

$$\Delta; \Gamma \Vdash t : \tau \stackrel{\text{def}}{=} \forall \mathbf{W} \vdash \Delta, \forall \gamma \in \mathcal{G} [\Gamma]^\rho, (\mathbf{W}, \mathbf{ty}) \in \mathcal{E} [\tau]^\rho$$

The fundamental property of TWLR is analogous to that for RLR.

**THEOREM 2.11 (FUNDAMENTAL PROPERTY FOR TWLR).** *if  $\Delta; \Gamma \vdash v : \tau$  then  $\Delta; \Gamma \Vdash v : \tau$*

The main result we need from this logical relation is an analogous to Lemma 2.9 ([Boring lemma](#)). If we try to naively state such a result for this TWLR too we obtain the wrong statement below (Lemma 2.12).

**LEMMA 2.12 (WRONG BORING LEMMA FOR TWLR).** *If  $\mathbf{W}_1$  and  $\mathbf{W}_2$  agree on the free type variables of  $\tau$ , then*

$$(\mathbf{W}_1, \mathbf{t}) \in \mathcal{V} [\tau] \iff (\mathbf{W}_2, \mathbf{t}) \in \mathcal{V} [\tau]$$

Contrary to the RLR, this lemma does not hold for our TWLR. If we were to try and prove it, we would quickly get stuck in the cases for lambdas and big lambdas where we get a future world  $W'$  of, for example,  $W_1$  but we have no way to relate it to world  $W_2$ .

The correct statement of Lemma 2.9 ([Boring lemma](#)) applied to TWLR is the following one, which respects world monotonicity.

**LEMMA 2.13 (BORING LEMMA FOR TWLR).** *If  $\mathbf{W}_2 \sqsupseteq \mathbf{W}_1$ , then,  $\forall \mathbf{t}, v$*

$$\text{if } (\mathbf{W}_1, v) \in \mathcal{V} [\tau] \text{ then } (\mathbf{W}_2, v) \in \mathcal{V} [\tau]$$

The premise has changed in this lemma, in fact we are only allowed to *extend* a world, but not remove bindings from it.

The difference between the RLR of Section 2.3 and the TWLR of this section may appear technical. However, we will see in the next section that this technical change from RLR to TWLR renders parametricity compatible with universal types. This fact is witnessed by several TWLR-based parametricity proofs for languages with such types [[Ahmed et al. 2011c, 2017](#); [Neis et al. 2009, 2011](#); [Sumii and Pierce 2003](#); [Toro et al. 2019](#)]. However, none of those papers point out that using a TWLR is not just technically convenient, but *required* in the presence of universal types.

In the next Section, we provide further insight into how the use of a TWLR enables compatibility with universal types, by demonstrating an RLR-based proof of the absence of universal types and explaining why this proof fails using a TWLR.

### 3 THE TYPE [Univ](#)

To construct this proof, we rely heavily on the type [Univ](#):

$$\mathbf{Univ} \stackrel{\text{def}}{=} \exists Y. \forall X. (X \rightarrow Y) \times (Y \rightarrow X)$$

The type can be read as stating the existence of a universal type  $Y$ : a type that all other types can be embedded into and extracted from.

In our non-terminating variant of  $\lambda^F$ , [Univ](#) is clearly inhabited, for example by this value:

$$\text{pack } \langle \mathbf{Unit}, \Lambda X. \langle \lambda_ : X. \mathbf{unit}, \lambda_ : \mathbf{Unit}. \omega_X \rangle \rangle \text{ as } \mathbf{Univ}$$

Note that we write  $\omega_X$  for a diverging term of type  $X$ , which can be constructed in the standard way using recursive types. However, this value is *degenerate* in the sense that injecting a value into the packaged  $Y$  and extracting it again diverges.

A crucial observation for this paper is that *all* System F values of type  $\mathbf{Univ}$  are degenerate in this sense. Intuitively, this is because a single type  $Y$  needs to be chosen, independently of the types  $X$  that will be embedded into it. Because nothing is known upfront about these  $X$ s and nothing can be learnt about them after invocation (because  $X$  must be treated parametrically), no viable choice for  $Y$  can be made.<sup>6</sup>

### 3.1 Two Contextually Equivalent Terms

This degeneracy of  $\mathbf{Univ}$  implies the contextual equivalence of the following two terms of type  $\mathbf{Univ} \rightarrow \mathbf{Unit}$ .

$$\begin{aligned} t_u &\stackrel{\text{def}}{=} \lambda x : \mathbf{Univ}. \text{unpack } x \text{ as } \langle Y, x' \rangle \text{ in} \\ &\quad \text{let } x'' : (\mathbf{Unit} \rightarrow Y) \times (Y \rightarrow \mathbf{Unit}) = x' \text{ Unit in } x''.2 (x''.1 \text{ unit}) \\ t_\omega &\stackrel{\text{def}}{=} \lambda x : \mathbf{Univ}. \omega_{\mathbf{Unit}} \end{aligned}$$

The reason that  $t_u$  and  $t_\omega$  are contextually equivalent is that both will diverge when applied to any argument of type  $\mathbf{Univ}$ . For  $t_u$ , this follows from the degeneracy of the type  $\mathbf{Univ}$ , as we demonstrate below, while for  $t_\omega$ , the term  $\omega_{\mathbf{Unit}}$  in the body ensures divergence. Note that the degeneracy of  $\mathbf{Univ}$  is essential: if the context were able to produce a non-degenerate value of type  $\mathbf{Univ}$ , then  $t_u$  would not diverge when applied to it, so that the context could distinguish  $t_u$  from  $t_\omega$ .

Thus, we have the following theorem.

**THEOREM 3.1** ( $t_u$  AND  $t_\omega$  ARE CONTEXTUALLY EQUIVALENT IN  $\lambda^F$ ).  $\emptyset; \emptyset \vdash t_u \simeq t_\omega : \mathbf{Univ} \rightarrow \mathbf{Unit}$ .

The first step of this proof is captured by Lemma 3.2 (*Diverging functions are contextually equivalent to an omega function*), which states that a function that diverges for every argument is equivalent to a function whose body is the omega term.

**LEMMA 3.2** (DIVERGING FUNCTIONS ARE CONTEXTUALLY EQUIVALENT TO AN OMEGA FUNCTION).

$$\begin{aligned} &\text{If } \emptyset; \emptyset \vdash \lambda x : \tau'. t : \tau' \rightarrow \mathbf{Unit} \text{ and (for all } \emptyset; \emptyset \vdash v : \tau' \text{ we have that } (\lambda x : \tau'. t) v \uparrow) \\ &\text{Then } \emptyset; \emptyset \vdash \lambda x : \tau'. t \simeq \lambda x : \tau'. \omega_{\mathbf{Unit}} : \tau' \rightarrow \mathbf{Unit} \end{aligned}$$

We do not think this lemma is very hard to believe. It can be proven using standard techniques, either using an ad hoc simulation argument or by relying on existing binary logical relations for System F, like the one by Dreyer et al. [2011]. To avoid distracting from the main point of our paper, we do not offer a proof here.

With Lemma 3.2, it suffices to prove that  $t_u$  always diverges when supplied with a value of type  $\mathbf{Univ}$  in order to conclude Theorem 3.1. This result we derive from the “degeneracy” of  $\mathbf{Univ}$ , i.e., from the fact that  $\mathbf{Univ}$  is only inhabited by diverging terms. We show this fact with the aid of the RLR of Section 2.3 (Lemma 3.3).

**LEMMA 3.3** ( $\mathbf{Univ}$  IS DEGENERATE). *For all*  $\emptyset; \emptyset \vdash v : \mathbf{Univ}$ , *we have that*  $t_u v \uparrow$ .

**PROOF.** The proof proceeds largely in a standard way: we unfold the definition of value relation for the value whose type is known and we rely on antireduction and compatibility lemmas in order to reason about terms after they evaluate. The goal is to show divergence of term  $t_u v$ . This can be achieved by showing that that term belongs to the term relation for a type variable whose set of

<sup>6</sup>Note that the situation is entirely different if we swap the quantifications in the type:  $\mathbf{Triv} \stackrel{\text{def}}{=} \forall X. \exists Y. (X \rightarrow Y) \times (Y \rightarrow X)$ .

inhabitants is empty. Specifically, we will prove that it belongs to  $\mathcal{E} \llbracket \mathbf{X} \rrbracket^{\dots, X \mapsto (\tau_X, \emptyset)}$ . By definition this means that the term must diverge or reduce to a value in  $\mathcal{V} \llbracket \mathbf{X} \rrbracket^{\dots, X \mapsto (\tau_X, \emptyset)}$ . Since that value relation is empty, the latter is not possible, so  $t_u v$  must diverge.

To prove that  $t_u v$  is in  $\mathcal{E} \llbracket \mathbf{X} \rrbracket^{\dots, X \mapsto (\tau_X, \emptyset)}$ , the main trick we use relies on having two relations for  $\mathbf{X}$  to inhabit, for  $\tau_X = \mathbf{Unit}$ . We will then instantiate the quantification over  $\mathbf{X}$  with two different semantic interpretations.

- (1) the first one  $\mathbf{RU} = \{\mathbf{unit}\}$ ;
- (2) the second one  $\mathbf{RE} = \emptyset$ .

Now take  $\emptyset; \emptyset \vdash v : \mathbf{Univ}$ . By Theorem 2.5 (Fundamental property for RLR) with  $\emptyset; \emptyset \vdash v : \mathbf{Univ}$ , we have (HLR)  $\emptyset; \emptyset \Vdash v : \mathbf{Univ}$ . By Definition 2.4 (Reynolds-style Logical Relation) with HLR (taking  $\rho = \emptyset$  and  $\gamma = \emptyset$ ), we have  $v \in \mathcal{E} \llbracket \mathbf{Univ} \rrbracket^{\emptyset}$ . Because  $v$  is a value, we have  $v \in \mathcal{V} \llbracket \mathbf{Univ} \rrbracket^{\emptyset}$ .

By unfolding the definition of  $\mathbf{Univ}$  and the value relation of the related type(s), it follows for some  $\tau_Y, v'$  and  $\mathbf{R}_Y \in \text{Rel}(\tau_Y)$ , that

- $v = \mathbf{pack} \langle \tau_Y, v' \rangle$  as  $\mathbf{Univ}$
- (HPVP)  $v' \in \mathcal{V} \llbracket \forall X. (X \rightarrow Y) \times (Y \rightarrow X) \rrbracket^{\emptyset, Y \mapsto (\tau_Y, \mathbf{R}_Y)}$ .

Let us now consider the reductions of  $t_u v$ :

$$\begin{aligned} t_u v & \\ & \equiv (\lambda x : \mathbf{Univ}. \mathbf{unpack} \ x \ \text{as} \ \langle Y, x' \rangle \ \text{in} \ (x' \ [\mathbf{Unit}]).2 \ ((x' \ [\mathbf{Unit}]).1 \ \mathbf{unit})) \ v \\ & \hookrightarrow \mathbf{unpack} \ v \ \text{as} \ \langle Y, x' \rangle \ \text{in} \ (x' \ [\mathbf{Unit}]).2 \ ((x' \ [\mathbf{Unit}]).1 \ \mathbf{unit}) \\ & \equiv \mathbf{unpack} \ (\mathbf{pack} \ \langle \tau_Y, v' \rangle \ \text{as} \ \mathbf{Univ}) \ \text{as} \ \langle Y, x' \rangle \ \text{in} \ (x' \ [\mathbf{Unit}]).2 \ ((x' \ [\mathbf{Unit}]).1 \ \mathbf{unit}) \\ & \hookrightarrow (v' \ [\mathbf{Unit}]).2 \ ((v' \ [\mathbf{Unit}]).1 \ \mathbf{unit}) \end{aligned}$$

By Lemma 2.6 (Antireduction) and Lemma 2.9 (Boring lemma), to conclude our thesis ( $t_u v \in \mathcal{E} \llbracket \mathbf{X} \rrbracket^{\emptyset, X \mapsto (\tau_X, \emptyset)}$ ) it is sufficient to show that:

$$(v' \ [\mathbf{Unit}]).2 \ ((v' \ [\mathbf{Unit}]).1 \ \mathbf{unit}) \in \mathcal{E} \llbracket \mathbf{X} \rrbracket^{\emptyset, Y \mapsto (\tau_Y, \mathbf{R}_Y), X \mapsto (\tau_X, \emptyset)}$$

We can use compatibility lemmas to reason about this term and the term relation that it inhabits. Specifically, we can see that this term is a top-level application, so by Lemma 2.8 (Compatibility for applications), it suffices to prove the following (recall that  $\mathbf{RE} = \emptyset$ ):

- (1)  $(v' \ [\mathbf{Unit}]).2 \in \mathcal{E} \llbracket Y \rightarrow X \rrbracket^{\emptyset, Y \mapsto (\tau_Y, \mathbf{R}_Y), X \mapsto (\tau_X, \mathbf{RE})}$
- (2)  $(v' \ [\mathbf{Unit}]).1 \ \mathbf{unit} \in \mathcal{E} \llbracket Y \rrbracket^{\emptyset, Y \mapsto (\tau_Y, \mathbf{R}_Y), X \mapsto (\tau_X, \mathbf{RE})}$

The former follows easily from HPVP by unfolding the definition of value relation for universals and pairs.

To prove Item 2 we apply our main trick. By Lemma 2.9 (Boring lemma), we can replace the relation with an equivalent one. We first drop the first relation for  $\mathbf{X}$  in the term relation for  $\mathbf{Y}$  (since variable  $\mathbf{X}$  is not mentioned in type  $\mathbf{Y}$ ), then add the second relation and all the term relations are equivalent:

$$\mathcal{E} \llbracket Y \rrbracket^{\emptyset, Y \mapsto (\tau_Y, \mathbf{R}_Y), X \mapsto (\tau_X, \mathbf{RE})} = \mathcal{E} \llbracket Y \rrbracket^{\emptyset, Y \mapsto (\tau_Y, \mathbf{R}_Y)} = \mathcal{E} \llbracket Y \rrbracket^{\emptyset, Y \mapsto (\tau_Y, \mathbf{R}_Y), X \mapsto (\tau_X, \mathbf{RU})}$$

So instead of proving:

$$(v' \ [\mathbf{Unit}]).1 \ \mathbf{unit} \in \mathcal{E} \llbracket Y \rrbracket^{\emptyset, Y \mapsto (\tau_Y, \mathbf{R}_Y), X \mapsto (\tau_X, \mathbf{RE})}$$

we can just prove (notice the change in the semantic interpretation for  $\mathbf{X}$ ):

$$(v' \ [\mathbf{Unit}]).1 \ \mathbf{unit} \in \mathcal{E} \llbracket Y \rrbracket^{\emptyset, Y \mapsto (\tau_Y, \mathbf{R}_Y), X \mapsto (\tau_X, \mathbf{RU})}$$

This is necessary in order to reason about the function application within this term, as we explain below.

Again, we apply Lemma 2.8 (Compatibility for applications) and it suffices to prove the following:

- (4)  $(\mathbf{v}' \text{ [Unit]}.1) \in \mathcal{E} \llbracket \mathbf{X} \rightarrow \mathbf{Y} \rrbracket^{\emptyset, Y \mapsto (\tau_Y, R_Y), X \mapsto (\tau_X, RU)}$   
 (5)  $\mathbf{unit} \in \mathcal{E} \llbracket \mathbf{X} \rrbracket^{\emptyset, Y \mapsto (\tau_Y, R_Y), X \mapsto (\tau_X, RU)}$

The former is again straightforward from HPVP by unfolding the definition of the value relation for universals and pairs.

Proving Item 5 is not hard either. The term relation includes the value relation, so it suffices to prove:

$$\mathbf{unit} \in \mathcal{V} \llbracket \mathbf{X} \rrbracket^{\emptyset, Y \mapsto (\tau_Y, R_Y), X \mapsto (\tau_X, RU)}$$

By definition of the value relation for type variables, this holds if  $\mathbf{unit}$  inhabits the semantic interpretation for  $\mathbf{X}$ , i.e.  $RU = \{\mathbf{unit}\}$ .  $\square$

If we had not performed our main trick, all other proof obligations would hold, but proving Item 5 would not be possible. In fact, there we would have had to prove that:

$$\mathbf{unit} \in \mathcal{V} \llbracket \mathbf{X} \rrbracket^{\emptyset, Y \mapsto (\tau_Y, R_Y), X \mapsto (\tau_X, RE)}$$

i.e., (according to the value relation for type variables) that  $\mathbf{unit}$  is in the current semantic interpretation for  $\mathbf{X}$ , i.e.,  $\emptyset$ , which is not possible.

### 3.2 Why Universal Types Require Type Worlds

The counterexample in the previous section demonstrates that Reynolds-style logical relations are incompatible with a universal type. But to thoroughly understand what's going on, it is useful to take a closer look at the counterexample.

Recall the definition of the type  $\underline{\mathbf{Univ}}$ .

$$\underline{\mathbf{Univ}} \stackrel{\text{def}}{=} \exists Y. \forall X. (\mathbf{X} \rightarrow \mathbf{Y}) \times (\mathbf{Y} \rightarrow \mathbf{X})$$

According to the RLR, any value of this type is of the form  $\text{pack } \langle \tau_Y, \mathbf{v} \rangle \text{ as } \underline{\mathbf{Univ}}$  and comes with a predicate  $R_Y \in \text{Re1 } (\tau_Y)$ , i.e., a predicate on values of type  $\tau'$ . Importantly, this predicate needs to be chosen independently of choices that are made later, particularly the choice for the type  $\tau_X$  and the predicate  $R_X \in \text{Re1 } (\tau_X)$  which the universal quantification will be instantiated with.

It is this independence (of the choice of  $R_Y$ ) that is fundamentally incompatible with the existence of a universal type. Imagine there is a universal type  $\mathbf{U}$  in System F with total injection and extraction functions  $\text{in}_Z : \mathbf{Z} \rightarrow \mathbf{U}$  and  $\text{out}_Z : \mathbf{U} \rightarrow \mathbf{Z}$  for arbitrary types  $\mathbf{Z}$ . Then we could define a value of type  $\underline{\mathbf{Univ}}$  by taking  $\mathbf{Y} = \mathbf{U}$  and constructing a pair with  $\text{in}_X$  and  $\text{out}_X$ :

$$\text{pack } \langle \mathbf{U}, \lambda X. \langle \text{in}_X, \text{out}_X \rangle \rangle \text{ as } \underline{\mathbf{Univ}}$$

Now, to make a choice for the predicate  $R_U$ , one thing to decide is whether or not the predicate should accept the result of  $\text{in}_X \mathbf{unit}$  when  $\mathbf{X}$  is later instantiated to  $\mathbf{Unit}$  (as in the counterexample). However, whether or not this value can be legally embedded in  $\mathbf{U}$  fundamentally depends on the choice for  $R_X \in \text{Re1 } (\mathbf{Unit})$ . Particularly, if  $R_X = \emptyset$ , then it should be rejected, but if  $R_X = \{\mathbf{unit}\}$  then it should be accepted. Clearly, this is impossible if we need to choose  $R_Y$  before we know what  $R_X$  will be instantiated with.

So how do type-worlds fit into this picture? In a TWLR like that of Section 2.4, we still need to make a choice for  $R_Y$  before a choice is made for  $R_X$ . However, the type of  $R_Y$  is different now:

$R_Y \in \mathcal{P}(\text{World} \times \text{Val})$  with

$$\forall (W, v) \in R. \forall W' \sqsupseteq W. (W', v) \in R \text{ and } \emptyset; \emptyset \vdash v : \tau_Y$$

Another way to think of this set is as  $\text{World} \xrightarrow{\text{mon}} \mathcal{P}(\text{Val})$ : the set of monotone functions from **World** to sets of values. In other words,  $R_Y$  is now a family of relations, and can decide which values to accept based on the current type world  $W$ . This world  $W$  will contain choices of predicates for other type variables, particularly the choice for  $R_X$ . In each of the TWLRs where a universal type is at play, the logical relation will accept different values depending on the world  $W$ , thus solving the conundrum outlined above.

It is also interesting to consider what this means in practice, when proving theorems using parametricity. Many proofs go through with TWLRs as they do with RLRs, as demonstrated, for example, by Ahmed et al. [2017]. However, this is not the case for all proofs and Lemma 3.3 (**Univ is degenerate**) from Section 3.1 is a perfect example. To see this, let us try to adapt the proof to use the TWLR from Section 2.4 instead of the RLR from Section 2.3, and see where this fails. Interestingly, the place we get stuck is at the application of Lemma 2.9 (**Boring lemma**), suggesting the theorem is not as boring as the name suggests.

When we look at how Lemma 2.9 (**Boring lemma**) was applied in the proof of Lemma 3.3, we see that we were able to prove the following:

$$(v' [\text{Unit}]).1 \text{ unit} \in \mathcal{E} \llbracket Y \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\tau_X, RU)}$$

Then, we used Lemma 2.9 (**Boring lemma**) twice, the first time to forget a binding for  $X$ , and the second time to add a different binding for the same  $X$ :

$$\mathcal{E} \llbracket Y \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\tau_X, RU)} = \mathcal{E} \llbracket Y \rrbracket^{0, Y \mapsto (\tau_Y, R_Y)} = \mathcal{E} \llbracket Y \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\tau_X, RE)}$$

We could then conclude that

$$(v' [\text{Unit}]).1 \text{ unit} \in \mathcal{E} \llbracket Y \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\tau_X, RE)}$$

With TWLR, these applications of the lemma cannot be replicated. Consider the following worlds:

- $W_{yx1} = ((Y, \tau_Y, R_Y); (X, \tau_X, RU))$
- $W_{yx2} = ((Y, \tau_Y, R_Y); (X, \tau_X, RE))$

Clearly,  $W_{yx2} \not\sqsupseteq W_{yx1}$  (HPNF).

We can replicate the first steps of the proof until we have the following relation:

$$(W_{yx1}, (v' [\text{Unit}]).1 \text{ unit}) \in \mathcal{E} \llbracket Y \rrbracket$$

but the step to  $W_{yx2}$  no longer holds:

$$(W_{yx2}, (v' [\text{Unit}]).1 \text{ unit}) \notin \mathcal{E} \llbracket Y \rrbracket$$

We cannot use Lemma 2.13 (**Boring lemma for TWLR**) because of (HPNF).

In other words, the use of a TWLR means that different terms may be valid at type  $Y$  in world  $W_{yx1}$  than in  $W_{yx2}$ . This dependence of the relation for  $Y$  on the world is precisely what a TWLR purposefully allows. As a result, it is impossible to transfer the result about  $v' [\text{Unit}]$  from the world  $W_{yx1}$  (with the permissive predicate  $RU$  for  $X$ ) to the world  $W_{yx2}$  (with the more restrictive predicate  $RE$  for  $X$ ).

In addition to clarifying the relation between TWLRs and universal types, this paper also discusses some consequences that follow from these insights. Particularly, the example from Section 3.1 can be used to disprove two open conjectures and expose a previously unmentioned deficiency of polymorphic blame calculi. In the next three sections, we discuss these three topics in turn, starting with the secure compilation of System F in the cryptographic  $\lambda$ -calculus, a conjecture by [Pierce and Sumii \[2000\]](#); [Sumii and Pierce \[2004\]](#).

## 4 ENFORCING PARAMETRICITY IN AN UNTYPED TARGET LANGUAGE

Sumii and Pierce’s conjecture is about a compiler from  $\lambda^F$  to an untyped lambda calculus with sealing (idealised encryption) called  $\lambda^\sigma$ . In this section, we first introduce the target language  $\lambda^\sigma$  (Section 4.1) and the compiler from  $\lambda^F$  to  $\lambda^\sigma$  (Section 4.2). We then prove that the compiler is not fully-abstract: there exist two terms that are contextually-equivalent in  $\lambda^F$  but whose compilation is inequivalent in  $\lambda^\sigma$  (Section 4.3).

*Remark.* Sumii and Pierce have presented both a typed [[Pierce and Sumii 2000](#)] as well as an untyped [[Sumii and Pierce 2004](#)] version of  $\lambda^\sigma$ . We use the untyped version because it is quite a bit simpler.<sup>7</sup> However, the typed version suffers from the same problem, as we show in detail in the technical report. Essentially, both settings break degeneracy of [Univ](#) because they feature a universal type: the untype of all values in the untyped target language and the type `bits` of ciphertexts produced by encryption (sealing) in the typed target language.

### 4.1 The Cryptographic Lambda Calculus $\lambda^\sigma$

$\lambda^\sigma$  (Figure 2) is an untyped  $\lambda$ -calculus, extended with *sealing*, which models a *dynamic* protection mechanism such as (idealized) symmetric encryption [[Sumii and Pierce 2004](#)].

Most syntactic constructs are standard for an untyped  $\lambda$ -calculus. The term `wrong` models run-time errors and is a stuck term. Sealing introduces four new syntactic constructs in the calculus: `vx.t` creates a fresh seal (symmetric encryption key) and then evaluates `t` with `x` bound to the newly created seal. There is no surface syntax for seals, but the internal syntax  $\sigma$  represents run-time seal values created with `vx.t`. The construct `{t1}t2` first evaluates `t1` and `t2` to values `v1` and `v2` and then creates the sealed value `{v1}v2` (or leads to a run-time error if `v2` is not a seal). One can think of such a sealed value as `v1` encrypted under `v2`. The final construct `let {x}t1 = t2 in t3 else t4` is for unsealing or decrypting. It first evaluates `t1` to a seal  `$\sigma_1$`  and `t2` to `{v2} $\sigma_2$`  (or produces a run-time error if either result is not of that form). If  `$\sigma_1$`  and  `$\sigma_2$`  are equal, `t3` is evaluated with `x` bound to the decrypted value `v2`, otherwise `t4` is evaluated.

Program contexts in  $\lambda^\sigma$  are defined as for  $\lambda^F$  and are denoted with `C`. Contextual equivalence in  $\lambda^\sigma$ , indicated with  $\approx$  is defined analogously to Definition 2.1. Note that the quantified contexts are not allowed to contain literal seals (as they are internal syntax), but they are allowed to allocate and use fresh seals of their own.

Sealing is the main information hiding mechanism in  $\lambda^\sigma$ : by creating a new seal  $\sigma$  and making sure it does not leak to the context, a term can create values `{v} $\sigma$`  that are opaque to the context. [Pierce and Sumii \[2000\]](#) explain how one can use this information hiding mechanism to implement a protection similar to that offered by parametric polymorphism. For instance, the following implementation of  $\mathbb{Z}_3$  from Example 2.2 in terms of pairs of booleans, uses sealing to protect the abstraction.

<sup>7</sup>The extra complexity in the typed version comes from the difficulty of erasing a polymorphic function to a simply typed language, and from the fact that the compiler protects all type variables in a separate pass rather than using a single pass for all type variables.

## Syntax:

$$\begin{aligned}
t ::= & v \mid x \mid t \ t \mid t.1 \mid t.2 \mid \langle t, t \rangle \mid \text{if } t \text{ then } t \text{ else } t \\
& \mid vx.t \mid \{t\}_t \mid \sigma \mid \text{let } \{x\}_t = t \text{ in } t \text{ else } t \mid \text{roll } t \mid \text{unroll } t \mid \text{wrong} \\
v ::= & \text{unit} \mid \text{true} \mid \text{false} \mid \lambda x. t \mid \langle v, v \rangle \mid \{v\}_\sigma \mid \sigma \mid \text{roll } v
\end{aligned}$$

## Evaluation rules (excerpts):

$$\begin{array}{c}
\frac{\sigma \notin \text{dom}(h)}{(h, vx.t) \hookrightarrow_0 (h; \sigma, t[\sigma/x])} \qquad \frac{\sigma \equiv \sigma'}{\text{let } \{x\}_\sigma = \{v\}_{\sigma'} \text{ in } t \text{ else } t' \hookrightarrow_0 t[v/x]} \\
\frac{\sigma \not\equiv \sigma'}{\text{let } \{x\}_\sigma = \{v\}_{\sigma'} \text{ in } t \text{ else } t' \hookrightarrow_0 t'} \qquad \frac{\not\equiv v', \sigma'. v \equiv \{v'\}_{\sigma'}}{\text{let } \{x\}_\sigma = v \text{ in } t \text{ else } t' \hookrightarrow_0 \text{wrong}} \\
\frac{\not\equiv \sigma. v \equiv \sigma}{\text{let } \{x\}_v = v' \text{ in } t \text{ else } t' \hookrightarrow_0 \text{wrong}} \qquad \frac{\not\equiv \sigma. v' \equiv \sigma}{\{v\}_v \hookrightarrow_0 \text{wrong}} \qquad \frac{t \hookrightarrow_0 t'}{(h, E[t]) \hookrightarrow_0 (h, E[t'])}
\end{array}$$

Fig. 2.  $\lambda^\sigma$  syntax and evaluation rules (excerpts). A  $\lambda^\sigma$  program state is a pair  $h; t$  where  $h$  is the list of allocated seals. The semantics relation  $\hookrightarrow$  relies on the beta reductions indicated as  $\hookrightarrow_0$ , which do not require a list of allocated seals to reduce.

*Example 4.1* ( $\mathbb{Z}_3$  in  $\lambda^\sigma$ ). First, we introduce two helper functions:

$$\text{seal}_\sigma \stackrel{\text{def}}{=} \lambda y. \{y\}_\sigma \qquad \text{unseal}_\sigma \stackrel{\text{def}}{=} \lambda y. \text{let } \{x\}_\sigma = y \text{ in } x \text{ else wrong}$$

Then, we can define a  $\lambda^\sigma$  correspondent of **z3** as:

$$z3 = vs. \langle \langle \text{zero}, \text{succ} \rangle, \text{zero?} \rangle$$

$$\text{where} \begin{cases} \text{zero} = \text{seal}_s \langle \text{false}, \text{false} \rangle \\ \text{succ} = \lambda p. \text{if } (\text{unseal}_s p).2 \text{ then } \text{seal}_s \langle \text{false}, \text{false} \rangle \text{ else} \\ \qquad \qquad \text{if } (\text{unseal}_s p).1 \text{ then } \text{seal}_s \langle \text{false}, \text{true} \rangle \text{ else } \text{seal}_s \langle \text{true}, \text{false} \rangle \\ \text{zero?} = \lambda p. \text{if } (\text{unseal}_s p).1 \text{ then false else if } (\text{unseal}_s p).2 \text{ then false else true} \end{cases}$$

Whenever values of the abstract type leave the scope of the abstract type definition, they are encrypted, and when they are passed back in they are decrypted before use. Intuitively, one can see that this protects against a context looking into or tampering with representation values, similarly to the protection offered by type checking of parametric polymorphism.

Also the protection required for the dual case, where a term calls a universally quantified function provided by the context, can be implemented in  $\lambda^\sigma$ . Consider for instance the  $\lambda^F$  term:

$$\lambda f : \forall X. X \times X \rightarrow X. \text{if } f \text{ Bool } \langle \text{true}, \text{true} \rangle \text{ then true else false}$$

The polymorphic function  $f$  that will be passed in by the context, can only return one of its arguments (or diverge), and hence the invocation in the term above will necessarily return **true** (or diverge).

We can implement a similar protection in  $\lambda^\sigma$ . To enforce parametric behaviour of a polymorphic function received from the context, we create a new seal for every invocation, and we encrypt all parameters of the quantified type with that seal. For instance, the term above could be implemented



$$\begin{aligned}
\text{protect}_{\eta;\text{Unit}} x &\stackrel{\text{def}}{=} x \\
\text{protect}_{\eta;\text{Bool}} x &\stackrel{\text{def}}{=} x \\
\text{protect}_{\eta;\tau_1 \times \tau_2} x &\stackrel{\text{def}}{=} \text{let } x_1 = x.1 \text{ in let } x_2 = x.2 \text{ in } \langle \text{protect}_{\eta;\tau_1} x_1, \text{protect}_{\eta;\tau_2} x_2 \rangle \\
\text{protect}_{\eta;\tau_1 \rightarrow \tau_2} x &\stackrel{\text{def}}{=} \lambda y. \text{let } z = x(\text{confine}_{\eta;\tau_1} y) \text{ in } \text{protect}_{\eta;\tau_2} z \\
\text{protect}_{\eta;\forall X. \tau} x &\stackrel{\text{def}}{=} \lambda \_ . \text{let } y = x \text{ unit in } \text{protect}_{\eta, X \mapsto (\lambda x.x, \lambda x.x); \tau} y \\
\text{protect}_{\eta;\exists X. \tau} x &\stackrel{\text{def}}{=} \text{vs. } \text{protect}_{\eta, X \mapsto (\text{seals}, \text{unseals}); \tau} x \\
\text{protect}_{\eta;\mu X. \tau} x &\stackrel{\text{def}}{=} (\text{fix}_2 (\lambda \text{rec.} \langle \lambda y. \text{protect}_{\eta, X \mapsto (\text{rec}.1, \text{rec}.2); \tau} y, \lambda y. \text{confine}_{\eta, X \mapsto (\text{rec}.1, \text{rec}.2); \tau} y \rangle)).1 x \\
\text{protect}_{\eta;X} x &\stackrel{\text{def}}{=} t_p x \quad \text{where } \eta(X) = (t_p, \_) \\
\text{confine}_{\eta;\text{Unit}} x &\stackrel{\text{def}}{=} x; \text{unit} \\
\text{confine}_{\eta;\text{Bool}} x &\stackrel{\text{def}}{=} \text{if } x \text{ then true else false} \\
\text{confine}_{\eta;\tau_1 \times \tau_2} x &\stackrel{\text{def}}{=} \text{let } x_1 = x.1 \text{ in let } x_2 = x.2 \text{ in } \langle \text{confine}_{\eta;\tau_1} x_1, \text{confine}_{\eta;\tau_2} x_2 \rangle \\
\text{confine}_{\eta;\tau_1 \rightarrow \tau_2} x &\stackrel{\text{def}}{=} \lambda y. \text{let } z = x(\text{protect}_{\eta;\tau_1} y) \text{ in } \text{confine}_{\eta;\tau_2} z \\
\text{confine}_{\eta;\forall X. \tau} x &\stackrel{\text{def}}{=} \lambda \_ . \text{vs. let } x' = x \text{ unit in } \text{confine}_{\eta, X \mapsto (\text{seals}, \text{unseals}); \tau} x' \\
\text{confine}_{\eta;\exists X. \tau} x &\stackrel{\text{def}}{=} \text{confine}_{\eta, X \mapsto (\lambda x.x, \lambda x.x); \tau} x \\
\text{confine}_{\eta;\mu X. \tau} x &\stackrel{\text{def}}{=} (\text{fix}_2 (\lambda \text{rec.} \langle \lambda y. \text{protect}_{\eta, X \mapsto (\text{rec}.1, \text{rec}.2); \tau} y, \lambda y. \text{confine}_{\eta, X \mapsto (\text{rec}.1, \text{rec}.2); \tau} y \rangle)).2 x \\
\text{confine}_{\eta;X} x &\stackrel{\text{def}}{=} t_c x \quad \text{where } \eta(X) = (\_, t_c)
\end{aligned}$$

Fig. 3. The dynamic wrappers of Sumii and Pierce’s compiler.  $\text{fix}_2$  is presented in Eq. (2) p. 20.

that the confined value is indeed of the expected type (**unit** or **true/false**). If any of these checks fail, the term reduces to **wrong**. Protecting at a ground type does nothing, because there is no way for the context to use such values that is not allowed by the type.

Protecting at type  $\forall X. \tau$  does nothing but forward the dummy **unit** application and recursively protect at type  $\tau$ . This is because there is nothing to protect: intuitively, the context cannot use a term of type  $\forall X. \tau$  in a way that is not allowed by the type. Similarly, confining at an existential type  $\exists X. \tau$  just recurses over type  $\tau$  without doing anything special for values of type  $X$ , because there is intuitively no way for a value of type  $\exists X. \tau$  to behave that is not allowed by the type.

Finally, when protecting a term of type  $\exists X. \tau$ , we want to make sure that the context treats the type  $X$  opaquely, so values of type  $X$  are sealed (encrypted) with a fresh seal  $\sigma$ . Similarly, confining at type  $\forall X. \tau$  generates a fresh seal for every invocation to protect values of type  $X$  with. This is the idea we have explained before in Section 4.1.

Applying the compiler to **z3** from Example 4.1 results in the  $\lambda^\sigma$  term **z3** from Example 2.2 (modulo some additional  $\beta$ -reductions).

*Technical remark.* Sumii and Pierce’s conjectured compiler for untyped  $\lambda^\sigma$  already had a technical flaw, albeit unrelated to parametricity. In fact, they had a terminating source language  $\lambda_{\text{F}}$  and an untyped – thus possibly diverging – target language. With this discrepancy it is impossible to build

a fully-abstract compiler for a simple reason. Consider these two terms, which are equivalent in  $\lambda_{\perp}^F$  (but not in  $\lambda^F$ ):

$$\lambda x : \text{Unit} \rightarrow \text{Unit}. \text{let } \_ = x \text{ unit in unit} \qquad \lambda x : \text{Unit} \rightarrow \text{Unit}. \text{unit}$$

Their compilation to  $\lambda^\sigma$  should intuitively be (modulo `protect`, which we elide for simplicity):

$$\lambda x. \text{let } \_ = x \text{ unit in unit} \qquad \lambda x. \text{unit}$$

However, a target context can distinguish them by passing them a term that once applied, diverges:

$$C \stackrel{\text{def}}{=} [\cdot] (\lambda y. \omega)$$

*Dealing with recursive types.* We avoid this problem by extending the compiler to the recursive types that we have in our non-terminating version of  $\lambda^F$ . Essentially, for recursive types  $\mu X. \tau$ , `protect` and `confine` can naturally be defined as a mutually recursive pair of functions, i.e., recursively in terms of themselves *and* each other. The definition in Fig. 3 achieves this using a `fix2` function, which constructs the fixpoint of a function mapping a pair of functions to a pair of functions.

We don't go into detail for this, but we can define `fix2` as follows, using a definition of Plotkin's Z combinator<sup>9</sup> as `fix`:

$$\text{fix} \stackrel{\text{def}}{=} \lambda f. (\lambda x. f (\lambda y. x \ x \ y)) (\lambda x. f (\lambda y. x \ x \ y)) \quad (1)$$

$$\text{fix}_2 \stackrel{\text{def}}{=} \lambda f. \text{fix} (\lambda \text{rec}. f \langle \lambda x. (\text{rec unit}).1 \ x, \lambda x. (\text{rec unit}).2 \ x \rangle) \text{unit} \quad (2)$$

It is worth noting at this point that recursive types are not essential for our counterexample. In fact, instead of recursive types, we could just as well have added a fixpoint primitive to make System F non-terminating and our counterexample would continue to apply without modification.

### 4.3 Disproving the Sumii-Pierce Conjecture

Sumii and Pierce conjectured that their compiler  $\llbracket \cdot \rrbracket_{\lambda^\sigma}^{\lambda^F}$  is fully abstract. In other words, two  $\lambda^F$  terms are contextually equivalent if and only if they are compiled to equivalent  $\lambda^\sigma$  terms.

CONJECTURE 4.2 (SUMII AND PIERCE).  $\emptyset; \emptyset \vdash t_1 \simeq t_2 : \tau$  if and only if  $\emptyset \vdash \llbracket t_1 \rrbracket_{\lambda^\sigma}^{\lambda^F} \simeq \llbracket t_2 \rrbracket_{\lambda^\sigma}^{\lambda^F}$

However, we can now prove that this conjecture is false. A counterexample is given by the terms  $t_u$  and  $t_\omega$  which we proved contextually equivalent in Theorem 3.1. Compiling  $t_u$  and  $t_\omega$  does not produce contextually equivalent  $\lambda^\sigma$  terms:

THEOREM 4.3 ( $t_u$  AND  $t_\omega$  ARE NOT EQUIVALENT AFTER COMPILATION TO  $\lambda^\sigma$ ).  $\emptyset \vdash \llbracket t_u \rrbracket_{\lambda^\sigma}^{\lambda^F} \neq \llbracket t_\omega \rrbracket_{\lambda^\sigma}^{\lambda^F}$

*Proof Sketch.* A full proof with all details can be found in the technical report.

The terms  $\llbracket t_u \rrbracket_{\lambda^\sigma}^{\lambda^F}$  and  $\llbracket t_\omega \rrbracket_{\lambda^\sigma}^{\lambda^F}$  can be discriminated by the following context:

$$C \stackrel{\text{def}}{=} [\cdot] (\lambda \_ . \langle \lambda x. x, \lambda x. x \rangle)$$

To understand the role of this context, recall that the contextual equivalence of  $t_u$  and  $t_d$  relies on the degeneracy of type `Univ`. The context `C` breaks this assumption by invoking the terms with a non-degenerate value of type `Univ`, which is constructed by using the untype of all untyped values as a universal type.

<sup>9</sup>Plotkin's Z combinator is a variant of the more well-known Y combinator that works in a call-by-value setting, but is restricted to constructing fixpoints that are functions, see Pierce [2002, §5.2].

By unfolding definitions and executing the operational semantics, it is easy to check that we get the following behaviour.

$$C \left[ \llbracket \mathbf{t}_u \rrbracket_{\lambda}^{\lambda^F} \right] \hookrightarrow^* \text{unit} \qquad C \left[ \llbracket \mathbf{t}_\omega \rrbracket_{\lambda}^{\lambda^F} \right] \hookrightarrow^* (\lambda r. r) \omega \hookrightarrow^* (\lambda r. r) \omega \hookrightarrow^* \dots \uparrow$$

We spell out the reductions in full detail in the technical appendix. From this behaviour, it follows immediately that the terms are not contextually equivalent.  $\square$

We can thus prove that Sumii and Pierce's conjecture is false as follows.

**THEOREM 4.4** ( $\llbracket \cdot \rrbracket_{\lambda}^{\lambda^F}$  IS NOT FULLY ABSTRACT). *It is not true that*

$$\forall \mathbf{t}_1, \mathbf{t}_2. \mathbf{t}_1 \simeq \mathbf{t}_2 \iff \llbracket \mathbf{t}_1 \rrbracket_{\lambda}^{\lambda^F} \simeq \llbracket \mathbf{t}_2 \rrbracket_{\lambda}^{\lambda^F}$$

**PROOF.** Follows easily from Theorem 3.1 and the counterexample in Theorem 4.3.  $\square$

*The problem is in the easy case.* It is worth noticing that most of the work in Sumii and Pierce's compiler (see Fig. 3) is in enforcing that existentially quantified types passed to the context are treated opaquely and, dually, that polymorphic functions received from the context are forced to treat their argument type opaquely. However, it is not in these cases that the counterexample highlights a problem.

Instead, it goes wrong in the cases where we receive an existential type from the context (and dually when we pass a polymorphic function to the context). For these seemingly simple cases, the dynamic wrappers from Fig. 3 do not perform any specific kind of enforcement (except for recursing on their body type). However, it is there that our counterexample uncovers a problem: the value that it provides as a value of  $\text{Univ} = \exists Y. \forall X. (X \rightarrow Y) \times (Y \rightarrow X)$  does not correspond to any legal choice of  $Y$ , but the dynamic type wrappers have no way to detect this. In fact, an alternative way to understand what goes wrong is that the value  $\lambda_. \langle \lambda x. x, \lambda x. x \rangle$  provided by the context, behaves as if it can choose  $Y$  equal to  $X$ . However, this is not possible in  $\lambda^F$  because  $X$  is not in scope at the moment when  $Y$  needs to be chosen.

In other words, what seems to be missing in Sumii and Pierce's dynamic enforcement is an enforcement of the type variable scope of existentially quantified types. To see this, consider the following type  $\text{Triv}$ , which is identical to  $\text{Univ}$ , except for the order of the quantifiers:

$$\text{Univ} \stackrel{\text{def}}{=} \exists Y. \forall X. (X \rightarrow Y) \times (Y \rightarrow X) \qquad \text{Triv} \stackrel{\text{def}}{=} \forall X. \exists Y. (X \rightarrow Y) \times (Y \rightarrow X)$$

Even though the type  $\text{Triv}$  is more liberal, as it allows to instantiate  $Y$  with  $X$ , the wrappers of Fig. 3 treat the types essentially the same. From this perspective, it is no surprise that the value  $\lambda_. \langle \lambda x. x, \lambda x. x \rangle$  is accepted as a value of type  $\text{Univ}$ , because it does in fact correspond to a legal  $\lambda^F$  value of type  $\text{Triv}$ :

$$\Delta X. \text{pack } \langle X, \langle \lambda x : X. x, \lambda x : X. x \rangle \rangle \text{ as } \exists Y. (X \rightarrow Y) \times (Y \rightarrow X)$$

## 5 ENFORCING PARAMETRICITY IN THE PRESENCE OF TYPE CASTS

The second conjecture we disprove is by Neis et al. [2009, 2011]. These authors study a form of runtime type generation to protect parametrically polymorphic functions when interacting with code that can use a type cast primitive. They prove a parametricity result that applies to appropriately wrapped System F values once embedded in  $\mathbb{G}$ , a language with a type cast primitive. They conjecture [Neis et al. 2009, section 10] that this wrapping is fully abstract, and while they prove

equivalence reflection, they only conjecture equivalence preservation due to a lack of sophistication of the proof techniques available at the time.<sup>10</sup>

It turns out that our results may be adapted to this setting, as in such a non-parametrically polymorphic setting, the type  $\forall X. X$  can be used as a universal type (that every other type can be embedded into or out from). In other words, we disprove this conjecture too: embedding System F into  $G$  à la Neis-Dreyer-Rossberg (NDR, in the sequel) is not fully abstract.

We first present  $G$  (Section 5.1) and the wrapper for protecting polymorphic functions from interacting with  $G$  terms (Section 5.2). Then we demonstrate how  $\forall Z. Z$  is a universal type in  $G$  (Section 5.3) and prove that due to that type, the wrapper does not preserve contextual equivalence (Section 5.4).

## 5.1 The $G$ Language

$G$  extends System F with two primitives (Figure 4). The first one casts between values of type  $\tau_1$  to values of type  $\tau_2$ . The second one generates a fresh type name  $X$  that is registered to be an equivalent classifier to a type  $\tau$  although data of the two types is not equivalent (so casting between  $X$  and  $\tau$  will not succeed).

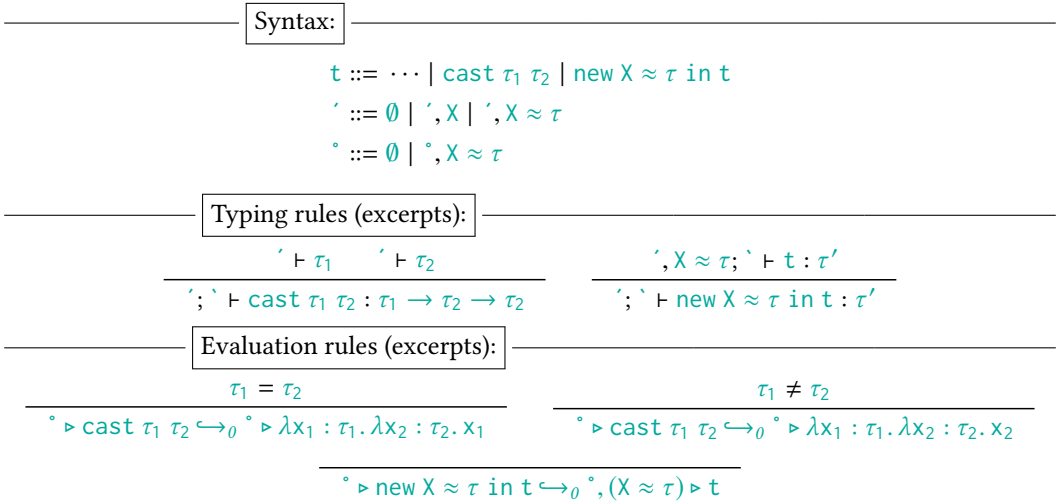


Fig. 4. The  $G$  language: syntax, typing rules and evaluation rules (excerpts). The semantics relation  $\hookrightarrow$  relies on the primitive reductions indicated as  $\hookrightarrow_0$  and it relates configurations of the form  $\circ \triangleright t$

## 5.2 Enforcing Parametricity in $G$

To avoid much code repetition, we use the same notation for the wrapper as that used by Neis et al. [2009]. For simplicity, we omit the case for recursive types and report only the wrapper cases that we rely only later on. Instead of writing two recursive functions such as `protect`, and `confine`, we annotate the wrapper with a *polarity*: positive polarity (+) is analogous to `protect`, negative polarity (−) is analogous to `confine`. When the *other* function is invoked, the polarity is switched

<sup>10</sup>A claim that was indeed true, since much of the research in proof techniques for fully abstract compilation postdates that paper [Abate et al. 2019; Ahmed and Blume 2011; Devriese et al. 2016; Fournet et al. 2013; New et al. 2016; Patrignani and Garg 2019; Schmidt-Schauß et al. 2015]. An exception is the work by Ahmed and Blume [2008] on typed closure conversion.

(i.e., from  $\pm$  to  $\mp$ ).

$$\begin{aligned}
W_{\tau}^{\pm}(t) &= \text{let } x = t \text{ in Wrap}_{\tau}^{\pm}(x) \\
\text{Wrap}_X^{\pm}(v) &= v \\
\text{Wrap}_{\text{Bool}}^{\pm}(v) &= v \\
\text{Wrap}_{\tau \rightarrow \tau'}^{\pm}(v) &= \lambda x : \tau. W_{\tau'}^{\pm}((v \text{ Wrap}_{\tau}^{\mp}(x))) \\
\text{Wrap}_{\tau \times \tau'}^{\pm}(v) &= \langle W_{\tau}^{\pm}(v.1), W_{\tau'}^{\pm}(v.2) \rangle \\
\text{Wrap}_{\forall X. \tau}^{\pm}(v) &= \sim X. \text{NewTy}^{\mp} X \text{ in } W_{\tau}^{\pm}(v X) \\
\text{Wrap}_{\exists X. \tau}^{\pm}(v) &= \text{unpack } v \text{ as } \langle X, x \rangle \text{ in } \text{NewTy}^{\pm} X \text{ in } \text{pack } \langle X, \text{Wrap}_{\tau}^{\pm}(x) \rangle \text{ as } \exists X. \tau \\
\text{NewTy}^+ X \text{ in } t &= \text{new } Y \approx X \text{ in } t[Y/X] \\
\text{NewTy}^- X \text{ in } t &= t
\end{aligned}$$

The wrapper consists of three parts. The first one,  $W_{\tau}^{\pm}(t)$ , is the term wrapper, it reduces a term  $t$  of source type  $\tau$  to a value and applies the value wrapper. The second one,  $\text{Wrap}_X^{\pm}(v)$ , is the value wrapper, it is responsible for generating the fresh type variables for outgoing universal values and incoming existential packages. The third one,  $\text{NewTy}^{\pm} X \text{ in } t$  is the code that effectively generates the fresh type variables, depending on the polarity of the invocation. As the authors themselves note, the wrapper functions in a way analogous to the dynamic checks inserted by the Sumii-Pierce compiler.

The compiler from  $\lambda^F$  to  $\mathbb{G}$ , denoted with  $\wr \cdot$ , leaves the term untouched and wraps it with a wrapper of the appropriate type.

$$\wr t \stackrel{\text{def}}{=} W_{\tau}^+(t) \quad \text{if } \emptyset; \emptyset \vdash t : \tau$$

The NDR conjecture (Theorem 5.1) states that this wrapper is fully abstract:

$$\text{THEOREM 5.1 (NDR CONJECTURE). } \forall t_1, t_2. t_1 \simeq t_2 \iff \wr t_1 \simeq \wr t_2$$

We will prove that this statement is not true (Theorem 5.3).

### 5.3 $\mathbb{G}$ has a Universal Type: $\forall Z. Z$

To explain how we disprove the NDR conjecture, we need to explain that  $\mathbb{G}$  indeed has a universal type (unlike System F). This type is  $\forall Z. Z$  and it functions as a universal type because we can take any other type  $X$  and (i) embed a value  $v$  of type  $X$  into  $\forall Z. Z$ , (ii) extract the same value  $v$  from  $\forall Z. Z$  at type  $X$ . It is possible to do this while remaining parametric in the type  $X$  that these functions work with, i.e., keeping type variable  $X$  free in these terms and binding it in a larger term using both these functions.

These two functionalities will be key for building a distinguishing context for  $\mathbb{G}$  which is analogous to the distinguishing contexts of for  $\lambda^{\sigma}$  and  $\lambda^B$  (the latter will be presented later). A value  $v$  of an arbitrary type  $\tau$  is represented as a value of type  $\forall Z. Z$ . The cast operator in  $\mathbb{G}$  allows us to make this function behave differently when it is applied to type  $\tau$  than when it is applied to other types. In the former case, we will make the function return  $v$  itself, while in the latter case, we simply make the function diverge.

Concretely, to extract a value from  $\forall Z. Z$  into  $X$ , we simply apply the polymorphic function to type  $X$ :

$$\lambda z : (\forall Z. Z). z X$$

$$\begin{aligned}
& \left( \left( \left( \left( \lambda x_2 : \forall X. X. x_2 X \right) \right) \right) \left( \left( \left( \lambda x : X. \sim X. \left( \text{cast } (\text{Unit} \rightarrow X) (\text{Unit} \rightarrow X) \right) \text{unit} \right) x \right) \right) \right) \text{Bool} \right) \text{true} \\
& \hookrightarrow \left( \left( \left( \left( \lambda x : \text{Bool}. \left( \left( \left( \lambda x_2 : \forall X. X. x_2 \text{Bool} \right) \right) \right) \left( \left( \lambda x : \text{Bool}. \sim X. \left( \text{cast } (\text{Unit} \rightarrow \text{Bool}) (\text{Unit} \rightarrow X) \right) \text{unit} \right) x \right) \right) \right) \right) \right) \text{true} \\
& \hookrightarrow \left( \left( \left( \left( \lambda x_2 : \forall X. X. x_2 \text{Bool} \right) \right) \right) \left( \left( \left( \lambda x : \text{Bool}. \sim X. \left( \text{cast } (\text{Unit} \rightarrow \text{Bool}) (\text{Unit} \rightarrow X) \right) \text{unit} \right) \text{true} \right) \right) \right) \\
& \hookrightarrow \left( \left( \left( \lambda x_2 : \forall X. X. x_2 \text{Bool} \right) \right) \right) \left( \sim X. \left( \text{cast } (\text{Unit} \rightarrow \text{Bool}) (\text{Unit} \rightarrow X) (\lambda_ : \text{Unit}. \text{true}) (\lambda_ : \text{Unit}. \omega_X) \right) \text{unit} \right) \\
& \hookrightarrow \left( \sim X. \left( \text{cast } (\text{Unit} \rightarrow \text{Bool}) (\text{Unit} \rightarrow X) (\lambda_ : \text{Unit}. \text{true}) (\lambda_ : \text{Unit}. \omega_X) \right) \text{unit} \right) \text{Bool} \\
& \hookrightarrow \left( \text{cast } (\text{Unit} \rightarrow \text{Bool}) (\text{Unit} \rightarrow \text{Bool}) (\lambda_ : \text{Unit}. \text{true}) (\lambda_ : \text{Unit}. \omega_X) \right) \text{unit} \\
& \hookrightarrow \left( (\lambda y_1 : (\text{Unit} \rightarrow \text{Bool}). \lambda y_2 : (\text{Unit} \rightarrow \text{Bool}). y_1) (\lambda_ : \text{Unit}. \text{true}) (\lambda_ : \text{Unit}. \omega_X) \right) \text{unit} \\
& \hookrightarrow^* (\lambda_ : \text{Unit}. \text{true}) \text{unit} \\
& \hookrightarrow \text{true}
\end{aligned}$$

Fig. 5. Reductions of the combination of the two functionalities.

Injecting a value of type from  $X$  into  $\forall Z. Z$  is a bit more complex. What we would like to write is the following:

$$\lambda x : X. \sim Z. \text{cast } X Z x \omega_Z$$

This term uses the `cast` primitive to cast a value of type  $X$  into  $\forall Z. Z$ . If everything goes well, the requested type  $Z$  is the same as the type  $X$  of the value contained, and the constructed function returns  $x$ . Otherwise, it diverges by calling  $\omega_Z$ .

Unfortunately, the above doesn't work as intended because we are in a call-by-value setting, so  $\omega_Z$  is evaluated before checking the type equality and the whole term always diverges. Fortunately, this can be resolved by the standard trick of "thunking" the term and the expected value  $x$  into lambdas that expect a `Unit` value, and pass `unit` to them after the cast:

$$\lambda x : X. \sim Z. (\text{cast } (\text{Unit} \rightarrow X) (\text{Unit} \rightarrow Z) (\lambda_ : \text{Unit}. x) (\lambda_ : \text{Unit}. \omega_Z)) \text{unit}$$

To see how these functions work in practice in our counterexample, we build a term that uses them in order to inject `true` of type `Bool` into  $\forall Z. Z$  and extract `true` back. The term and its reductions are in Figure 5.

## 5.4 Proving the NDR Conjecture

We now have all the technical machinery to disprove the NDR conjecture.

Compiling  $t_u$  to  $G$  results in the following term (the appendix contains the full expansion). As a convention we name variables and type variables from the wrapper with  $y$ . The first line is the  $t_u$  term ported to  $G$ , the second and third lines (as well as the wrapping lambda) are wrapper-generated code. This code will take the parameter that needs to be passed to  $t_u$  and perform a series of

unpacking and type renaming which are intended to preserve parametricity.

$$\lambda y_1 : \underline{\text{Univ}}. \left( \begin{array}{l} (\lambda x : \underline{\text{Univ}}. \text{unpack } x \text{ as } \langle Y, x' \rangle \text{ in let } x'' = x' \text{ Unit in } x''.2 (x''.1 \text{ unit})) \\ \text{unpack } y_1 \text{ as } \langle Y_1, y_4 \rangle \text{ in new } Y_2 \approx Y_1 \text{ in} \\ \text{pack } \langle Y_2, \Lambda Y_3. \langle (\lambda y_5 : Y_3. ((y_4 Y_3).1 y_5)), (\lambda y_6 : Y_2. ((y_4 Y_3).2 y_6)) \rangle \rangle \text{ as } \underline{\text{Univ}} \end{array} \right)$$

The distinguishing context for  $\mathbb{G}$  passes a parameter of type  $\underline{\text{Univ}}$  to the term in the hole (which will be either  $\wr \mathbf{t}_u \wr$  or  $\wr \mathbf{t}_\omega \wr$ ). The goal of this parameter is to make  $\wr \mathbf{t}_u \wr$  terminate and  $\wr \mathbf{t}_\omega \wr$  diverge. Let us discuss the passed parameter. Since  $\underline{\text{Univ}}$  is an existential type at the top level, the parameter is a `pack <.> as .`. The packed type is the universal type that exists in  $\mathbb{G}$ :  $\forall Z. Z$ . After the existential,  $\underline{\text{Univ}}$  has a universal quantification, and thus the body of the existential package contains a  $\sim Z. .$ . Then comes the pair of functions for projecting into and extracting from the universal type  $\forall Z. Z$ , as explained in Section 5.3. As a convention, we name variables and type variables from the context with  $z$ :

$$\mathbb{C}^{\mathbb{G}} \stackrel{\text{def}}{=} [\cdot] \left( \text{pack } \left\langle \forall Z. Z, \sim Z.1. \left( \lambda z1 : Z1. \Lambda Z2. \left( \text{cast } (\text{Unit} \rightarrow Z1) (\text{Unit} \rightarrow Z2) \right) \text{unit}, \right) \right\rangle \text{ as } \underline{\text{Univ}} \right)$$

**THEOREM 5.2** ( $\wr \mathbf{t}_u \wr$  AND  $\wr \mathbf{t}_\omega \wr$  ARE NOT EQUIVALENT IN  $\mathbb{G}$ ).  $\wr \mathbf{t}_u \wr \neq \wr \mathbf{t}_\omega \wr$

**PROOF.** We have that  $\mathbb{C}^{\mathbb{G}} [\wr \mathbf{t}_u \wr] \hookrightarrow^* \text{unit}$  while  $\mathbb{C}^{\mathbb{G}} [\wr \mathbf{t}_\omega \wr] \uparrow$ . □

**THEOREM 5.3** (EMBEDDING  $\lambda^{\text{F}}$  INTO  $\mathbb{G}$  IS NOT FULLY ABSTRACT). *It is not true that*

$$\forall \mathbf{t}_1, \mathbf{t}_2. \mathbf{t}_1 \simeq \mathbf{t}_2 \iff \wr \mathbf{t}_1 \wr \simeq \wr \mathbf{t}_2 \wr$$

**PROOF.** Follows directly from Theorem 5.2 and Theorem 3.1. □

## 6 SYSTEM F EQUIVALENCES VS. GRADUAL TYPES

After discussing these conjectures, we turn our attention to polymorphic gradual calculi: gradually typed languages featuring parametric polymorphism. As explained in the introduction, gradually-typed languages provide a path for migrating codebases of untyped code to typed code. From this high-level idea, a number of natural design goals follow and the literature contains a number of correctness properties that formally express these objectives.

First, gradual languages are intended to preserve the semantics of existing typed and untyped code. Additionally, turning untyped programs into typed ones by adding correct (!) type signatures should not modify the semantics of programs. Without going into detail (because it is not relevant for our discussion), [Siek et al. \[2015\]](#) formalise these objectives as a number of formal criteria for gradually-typed languages.

Another high-level design goal of gradually-typed languages is that the typed components continue to enjoy the benefits of well-typedness in the presence of untyped other components. [Wadler and Findler \[2009\]](#) have proposed the Blame Theorem that expresses this property when it comes to one such benefit: the absence of type errors at runtime. Gradual languages rely on dynamic casts (which can fail at runtime) to coerce untyped values into typed ones. [Wadler and Findler](#) add a notion of blame, which is essentially a way to identify the type cast that caused a runtime type error. The Blame Theorem then expresses that well-typed components are never to blame for such failures, i.e., the property that well-typed components never cause runtime type errors remains valid in the gradual language.

For our purposes, we are interested in another benefit of well-typedness that has received less attention in the literature on gradual types: the benefits that well-typedness provides in terms of

reasoning. Many type systems allow us to deduce properties of components from their types, for example, the function  $\lambda x : \text{Unit}. x$  is easily seen to be equivalent to  $\lambda x : \text{Unit}. \text{unit}$ , based on their types. System F parametricity similarly allows us to deduce properties from programs' types.

It is natural to wonder whether gradual type systems preserve the validity of such reasoning in the original typed language. For this paper, specifically, we are interested in the question whether contextual equivalences between typed terms continue to hold when these terms are considered in the gradual language. Formally, there is an embedding of  $\lambda^F$  terms  $t$  into  $\lambda^B$  terms that we will denote as  $\lfloor t \rfloor$ . The question then becomes: if  $t_1 \simeq t_2$ , do we have that  $\lfloor t_1 \rfloor \simeq \lfloor t_2 \rfloor$ , or, in other words, is the embedding of the typed language into the gradual language fully abstract?

In this section, we answer this question negatively for the polymorphic blame calculus, which extends System F into a gradually typed language. We first present Ahmed et al.'s  $\lambda^B$ , a gradually-typed, polymorphic lambda-calculus (Section 6.1). Next, we reconsider  $t_u$  and  $t_\omega$  from Theorem 3.1 in  $\lambda^B$  and present a  $\lambda^B$  context that differentiates them (Section 6.2). This shows that the embedding of  $\lambda^F$  into  $\lambda^B$  is not fully abstract.

## 6.1 The $\lambda^B$ Calculus

There exist several versions of the polymorphic gradual languages in the literature [Ahmed et al. 2011b, 2017; Igarashi et al. 2017; Toro et al. 2019; Xie et al. 2018]. We take  $\lambda^B$  to be the polymorphic blame calculus as described by Ahmed et al. [2017]. This version modifies certain behaviour that was “topsy turvy” in the original version [Ahmed et al. 2011b]: a peculiar operational semantics that performs evaluation under type and value abstractions and an ad hoc postponing of run-time type generation in certain situations.

Syntax:	
Terms	$t ::= v \mid \text{if } t \text{ then } t \text{ else } t \mid x \mid tt \mid t\tau \mid t.1 \mid t.2 \mid \langle t, t \rangle$ $\mid t : \tau \xRightarrow{p} \tau \mid t : \tau \xRightarrow{\phi} \tau \mid \text{blame } p \mid t; t$
Values	$v ::= \text{unit} \mid \text{true} \mid \text{false} \mid \lambda x : \tau. t \mid \Lambda X. v \mid \langle v, v \rangle$ $\mid v : \tau \rightarrow \tau \xRightarrow{\phi} \tau \rightarrow \tau \mid v : \forall X. \tau \xRightarrow{\phi} \forall X. \tau \mid v : \tau \xRightarrow{\neg\alpha} \alpha$ $\mid v : \tau \rightarrow \tau \xRightarrow{p} \tau \rightarrow \tau \mid v : \tau \xRightarrow{p} \forall X. \tau \mid v : \gamma \xRightarrow{p} \star$
Types	$\tau ::= \text{Unit} \mid \text{Bool} \mid \tau \times \tau \mid \tau \rightarrow \tau \mid \forall X. \tau \mid X \mid \alpha \mid \star$
Ground types	$\gamma ::= \text{Unit} \mid \text{Bool} \mid \alpha \mid \star \times \star \mid \star \rightarrow \star$
Convertibility labels	$\phi ::= \alpha \mid \neg\alpha$
Compatibility labels	$p ::= l \mid \neg l$

Fig. 6. The polymorphic blame calculus [Ahmed et al. 2017]. Note: we have adapted notations, added the *Unit* type and removed the *int* type to align more closely with other calculi in this paper.

The syntax of  $\lambda^B$  is presented in Fig. 6. The calculus contains all terms and types of  $\lambda^F$  except for existentials. However, we can use a standard encoding of existentials in terms of universals as follows [Pierce 2002, §24.3, pp. 377-379]:

$$\exists X. \tau \stackrel{\text{def}}{=} \forall Y. (\forall X. \tau \rightarrow Y) \rightarrow Y$$

$$\text{pack } \langle \tau', t \rangle \text{ as } \exists X. \tau \stackrel{\text{def}}{=} \Lambda Y. \lambda f : (\forall X. \tau \rightarrow Y). f [\tau'] t$$

*unpack*  $t_1$  as  $\langle X, x \rangle$  in  $t_2 \stackrel{\text{def}}{=} t_1 [\tau_2]$  ( $\Lambda X. \lambda x : \tau_1. t_2$ )      where  $t_1 : \exists X. \tau_1$  and  $t_2 : \tau_2$

Even though existentials under this encoding are not entirely equivalent to regular existentials in our non-terminating calculus (because they contain a kind of bottom value), we can adapt our counterexample without trouble.

Then,  $\lambda^B$  contains a number of constructs to let typed and untyped code coexist. First we have the type  $\star$ : the type of untyped values. Untyped code can be typed with respect to the single, universal type  $\star$ .  $\lambda^B$  provides both a notion of casts ( $t : \tau \xRightarrow{p} \tau'$ ) and a notion of conversions ( $t : \tau \xRightarrow{\phi} \tau'$ ).

Casts ( $t : \tau \xRightarrow{p} \tau'$ ) represent a form of dynamic casts from  $\tau$  to  $\tau'$  that can potentially fail at runtime (in which case *blame*  $p$  is raised). Casts can be used to inject types  $\tau$  into the universal type  $\star$  and also to extract those types out of  $\star$  again. More generally, they can be used to convert between any types  $\tau$  and  $\tau'$  that satisfy a compatibility judgement  $\Sigma; \Delta \vdash \tau < \tau'$ , but we omit details for this because they are not relevant for our discussion. As part of the design that solves certain topsy-turvy aspects of the operational semantics in previous versions,  $\lambda^B$  carefully defines some casts as values: those where the term being cast is a value and the cast is either (1) between function types, (2) towards a polymorphic function type or (3) from a ground type  $\gamma$  into  $\star$ .

Polymorphic function application in  $\lambda^B$  does not use type substitution like in System F, but uses a notion of runtime type generation instead. Full details are not relevant for our discussion, but essentially, a polymorphic function application ( $\Lambda X. \lambda x : X. x$ ) *Unit* does not reduce to  $\lambda x : \text{Unit}. x$  but to  $\lambda x : \alpha. x$  for a fresh runtime type label  $\alpha$ , where the assignment  $\alpha := \text{Unit}$  is remembered in a type-name store  $\Sigma$ . Conversions ( $t : \tau \xRightarrow{\phi} \tau'$ ) represent a notion of static casts for converting between such a runtime type label  $\alpha$  and the type that is assigned to it in the store  $\Sigma$ . More generally, a convertibility judgement  $\Sigma; \Delta \vdash \tau <^{\phi} \tau'$  defines when  $\tau$  can be legally converted into  $\tau'$  by expanding ( $+\alpha \in \phi$ ) or reducing ( $-\alpha \in \phi$ ) runtime type labels' definitions.

## 6.2 Embedding of $\lambda^F$ into $\lambda^B$ is not Fully Abstract

To embed  $\lambda^F$  (defined in Section 2.1) into  $\lambda^B$ , most constructs can simply be mapped to the corresponding construct in  $\lambda^B$ . However, there is a discrepancy between type abstractions in  $\lambda^B$  and  $\lambda^F$ . The difference is that  $\lambda^B$  uses value polymorphism: the bodies of type abstractions are required to be values. This choice has some desirable consequences, particularly that type abstractions and applications can be removed entirely during type erasure.

The difference is essentially orthogonal to the topics of this paper, but we are unable to standardise on using value polymorphism or not, because the non-parametrically polymorphic language  $\mathbb{G}$  from the previous section and the polymorphic blame calculus  $\lambda^B$  make different choices and would both be non-trivial to modify.

Therefore, we embed polymorphic functions from  $\lambda^F$  into  $\lambda^B$  by introducing a form of thunking for polymorphic functions: the type  $\forall X. \tau$  is mapped to  $[\forall X. \tau] \stackrel{\text{def}}{=} \forall X. \text{Unit} \rightarrow [\tau]$  and type abstractions  $\Lambda X. t$  are mapped to  $[\Lambda X. t] \stackrel{\text{def}}{=} \Lambda X. \lambda \_ . [t]$ .

Our two contextually equivalent System F terms  $t_u$  and  $t_\omega$  embed into  $\lambda^B$  as  $[t_u]$  and  $[t_\omega]$ . In this section, we show that they do not remain contextually equivalent, showing that the embedding is not fully abstract. Similarly to before, we can construct a  $\lambda^B$  context  $C^B$  that differentiates them. As before, the context simply applies the terms to a non-degenerate value of type *Univ*, which we can construct thanks to the existence of the universal type  $\star$ :

$$C^B \stackrel{\text{def}}{=} [\_ ] \left( \text{pack} \left\langle \star, \Lambda X. \left\langle \lambda x : X. x : X \xRightarrow{p} \star, \lambda x : \star. x : \star \xRightarrow{p'} X \right\rangle \right\rangle \text{ as } \underline{\text{Univ}} \right)$$

The constructed  $\text{Univ}$  value simply takes  $\star$  as the existentially quantified universal type and uses casts to implement the functions from an arbitrary  $X$  into  $\star$  and back.

**THEOREM 6.1** ( $[\mathbf{t}_u]$  AND  $[\mathbf{t}_\omega]$  ARE NOT EQUIVALENT IN  $\lambda^B$ ).  $[\mathbf{t}_u] \neq [\mathbf{t}_\omega]$ .

**PROOF.** We have that  $C^B[\mathbf{t}_u] \hookrightarrow^* \text{unit}$  while  $C^B[\mathbf{t}_\omega] \not\Downarrow$ . We have verified this on paper and using the interpreter provided by Jamner and Siek to support Ahmed et al. [2017]’s results.<sup>11 12</sup> We provide the literal encoding of  $C^B[\mathbf{t}_u]$  and  $C^B[\mathbf{t}_\omega]$  for use in the interpreter in the technical appendix.  $\square$

**THEOREM 6.2** (EMBEDDING  $\lambda^F$  INTO  $\lambda^B$  IS NOT FULLY ABSTRACT). *It is not true that*

$$\forall \mathbf{t}_1, \mathbf{t}_2. \mathbf{t}_1 \simeq \mathbf{t}_2 \iff [\mathbf{t}_1] \simeq [\mathbf{t}_2]$$

**PROOF.** Follows directly from Theorem 6.1 and Theorem 3.1.  $\square$

Although the above discussion looks just at  $\lambda^B$  by Ahmed et al. [2017], our results also apply to the more recent polymorphic gradual languages proposed by Toro et al. [2019] and Xie et al. [2018].

TODO:  
And  
Igarashi  
et al.  
[2017]?

## 7 DISCUSSION

So Sumii and Pierce’s compiler is not fully abstract, the polymorphic blame calculus breaks contextual equivalences in System F and enforcing parametricity in a non-parametric calculus also breaks contextual equivalence. But what to conclude from this? Should we start looking for alternative ways to dynamically enforce parametricity or were we wrong to hope for these properties to hold in the first place? In this section we present some thoughts on the different possible options.

First, it is interesting to investigate whether full abstraction could be recovered by fixing Sumii and Pierce’s compiler or the polymorphic blame calculus. We discuss in the sections below some possible paths to explore, but we believe there are no straightforward solutions.

Second, another possible conclusion from our negative results is that full abstraction is perhaps too strong a property to aim for when doing secure compilation or gradual typing. We still think the Sumii-Pierce compiler is a useful secure compiler, even if it is not fully abstract. Hence we discuss how to weaken the requirements that full abstraction imposes on a translation from System F, by modifying System F in a way that weakens its contextual equivalences.

### 7.1 Fixing Sumii and Pierce’s Compiler?

One of the attractive features of Sumii and Pierce’s original conjecture is that it relies *only* on encryption (or at least, an idealised version of encryption in the form of seals). Because it required only encryption, this suggested that System F types could even be enforced as the contract for an untrusted adversary, running at the other end of a communication channel, on an untrusted computer. However, if we analyse the counterexample, this ambition of using just encryption seems hard to maintain.

Imagine that a compiled System F term (e.g.,  $\mathbf{t}_u$  or  $\mathbf{t}_\omega$ ) is communicating with such an adversary over a communication channel and we want to enforce that the adversary respects the contract represented by System F type  $\text{Univ}$ . What happens is the following:

- (1) The compiled term transmits the value  $\text{unit}$  of type  $\text{Unit}$ , encrypted as  $\{\text{unit}\}_\sigma$  of type  $X$  that is kept opaque from the adversary. The seal  $\sigma$  represents a fresh cryptographic key that we take care not to disclose to the adversary.

<sup>11</sup>Available at <http://www.ccs.neu.edu/home/dijamner/paramblame/artifact/>

<sup>12</sup>Thanks to Jeremy Siek and others, for their kind support in doing this.

- (2) The adversary replies with a value of the unknown type  $Y$  (chosen by the adversary). The actual value transmitted back in our counterexample, is simply the encrypted value  $\{\text{unit}\}_\sigma$  received in step 1.
- (3) The compiled term does not inspect the received value but simply transmits it back to the adversary as a value of type  $Y$ .
- (4) The adversary now takes the received value  $\{\text{unit}\}_\sigma$  and sends it back as a value of type  $X$ .
- (5) The compiled term receives this value, decrypts it using the private cryptographic key  $\sigma$  and uses the result as a value of type  $\text{Unit}$ .

What goes wrong in the above communication is that the value  $\{\text{unit}\}_\sigma$  sent back by the adversary in step 2 is essentially illegal. The adversary should have chosen a  $Y$  independently of  $X$  and values of such a type should intuitively not be able to contain values of type  $X$  (unless they are themselves packed in an existential package somehow). Let us try to think of what we could change in the communication protocol to enforce this.

In a pure cryptographic setting, we believe there is little hope to fix the compiler. To understand this, consider how, in the cryptography setting, the value  $\{\text{unit}\}_\sigma$ , that we send to the adversary in step 1, simply represents a sequence of bits that is the result of some encryption algorithm applied to value  $\text{unit}$ . On the other hand, the value sent back by the adversary in step 2 is another sequence of bits that represents a value of an unknown type  $Y$ . We want to prevent this second value from somehow including the first value  $\{\text{unit}\}_\sigma$ , but since the type  $Y$  is unknown and the adversary is running on an untrusted computer, there seems to be little the compiler can check. Any sequence of bits received from the adversary could in principle be a cleverly-encoded version of  $\{\text{unit}\}_\sigma$ : they could have XORed the value with an arbitrary other bitsequence and still be able to retrieve the original afterwards.

This (informal) argument suggests that the Sumii-Pierce compiler cannot be fixed in any way, as long as the target language contains only features that can be interpreted as a form of (idealised) cryptography. This would include the original sealing primitives (whatever way they are used), but also possible extensions that model a form of idealised signing (rather than encryption) (a track we were initially exploring).

To fix the compiler, it seems like we need to add some kind of feature that takes us beyond a pure-cryptography setting. We believe such a feature could take the form of a primitive that checks whether a value directly or indirectly contains values sealed with a certain seal. Such a primitive could perhaps allow to perform the required check on the value received from the adversary in step 2. Such a primitive does not correspond to a form of idealised encryption: noticing, for example, that a value like  $\{\{v\}_{\sigma_1}\}_{\sigma_2}$  contains a value sealed with  $\sigma_1$  would break the cryptographic interpretation of  $\lambda^\sigma$ , as it requires looking inside an encrypted value without access to the key ( $\sigma_2$ ) and requires detecting, for example, arbitrarily XORed versions of a ciphertext. However, the primitive could still be implementable in non-cryptographic settings, like the hardware-enforced seals that are present in capability machines: a form of processor with native support for capabilities and sealing that has been developed recently [Watson et al. 2015]. Note that the attacker model in this setting is a bit different: we assume that the untrusted attacker is now running on trusted hardware.

## 7.2 Polymorphic Blame Calculus Without a Universal Type?

When it comes to the polymorphic blame calculus, it is the existence of the universal type  $\star$  that breaks the fully abstract embedding of System F in our counterexample. So perhaps we should remove or modify the type  $\star$  in order to solve the problem?

We believe it may not be needed to remove the type  $\star$  entirely, but that we can instead replace it with a family of types  $\star_{\{\bar{X}\}}$ , indexed by a set of type variables  $\bar{X}$ . The type  $\star_{\bar{X}}$  would only be legal

(i.e., well-formed) when the type variables  $\tilde{X}$  are in scope. For values whose type is a type variable  $X'$ , injection into  $\star_{\tilde{X}}$  would only be legal if  $X'$  is in the set  $\tilde{X}$ .

For our counterexample, the effect would be that the context could no longer produce a non-degenerate value of type `Univ`, as there would no longer be a viable choice for  $Y$ . Any instance  $\star_{\{X\}}$  of the universal type we could choose, could not have  $X$  as part of  $\{\tilde{X}\}$ , as  $X$  is not yet in scope at the moment where  $Y$  needs to be produced. More generally, we conjecture that such a modified polymorphic blame calculus would embed System F in a fully abstract way.

An interesting question at this stage is whether this calculus would still adequately embed the untyped lambda calculus. Perhaps untyped code would be embeddable for any choice of  $\tilde{X}$ ? Would such untyped code still be able to interact as we would like with typed code?

### 7.3 Adjusting our Expectations

For us, the solutions suggested above look like they might work, but they are not without downsides. The modified Sumii-Pierce compiler could no longer be used in a purely-cryptographic setting and the family of universal types  $\star_{\{X\}}$  might be harder to use than the original  $\star$ . As such, we might consider alternative ways to address the lack of full abstraction.

One alternative is to abandon the choice of fully abstract compilation and rely on other notions. More concretely, perhaps a secure compiler or gradually typed language should not preserve arbitrary System F equivalences but only some of them, namely those that follow from a TWLR-parametricity?

Another alternative is to keep relying on fully abstract compilation and decide to simply adjust our expectations: perhaps preserving all System F equivalences is overly ambitious and we should find a way to eat what is on the table instead. Both in the case of Sumii and Pierce's compiler and the gradual lambda calculus, it seems like something non-trivial is being enforced, even though it is not the preservation of arbitrary System F contextual equivalences. A way to formalise this is to recover full abstraction by modifying the source language System F: weakening its contextual equivalences in order to make them easier to preserve.

In fact, our counterexample suggests a way to accomplish this: the problem is essentially that the type `Univ` is degenerate in System F, but not in the target language. So what if we modify System F to remove that degeneracy in the source language too? Specifically, what if we add a primitive type that all other types can be embedded into and extracted from? Interestingly, it seems like what we end up here is a simple version of the gradual type  $\star$ , together with injection and extraction functions.

Without working this out in more detail, we find it plausible that we can recover full abstraction with such a modification, both for Sumii and Pierce's compiler and the embedding into the polymorphic blame calculus. [Ahmed et al. \[2017\]](#) and [Toro et al. \[2019\]](#) have shown that such a variant of System F still satisfies useful (TWLR-based) parametricity results and that useful free theorems follow from it, suggesting it is a suitable language for programmers to work in.

## 8 RELATED WORK

*System F and parametricity.* Parametric polymorphism was first introduced 50 years ago as an informal concept by [Strachey \[2000\]](#). A few years later, System F was independently discovered by [Reynolds \[1974\]](#) and [Girard \[1972\]](#). [Reynolds \[1983\]](#) later formalised Strachey's informal concept of parametricity using a logical relation for System F, as explained in [Example 2.2](#) and [Wadler \[1989\]](#) popularised the property using the slogan "Theorems for Free". The property was further developed by many different researchers over the years but for space reasons, we refer to [Atkey et al. \[2014\]](#); [Wadler \[2007\]](#) for an overview of related work.

*Enforcing parametricity using dynamic sealing.* Dynamic sealing/unsealing was (informally) proposed more than 40 years ago by Morris [1973a,b] as a way for dynamically enforcing type abstractions. Much later, Pierce and Sumii [2000] developed this idea into a compiler that uses sealing to enforce System F’s parametricity and conjectured full abstraction of the proposed compiler. The target language in this work is a simply typed cryptographic  $\lambda$ -calculus. They already mention that the dynamic enforcement of parametricity may also be useful to combine parametric polymorphism and untyped languages, foreshadowing the work on parametrically polymorphic gradual type systems that we discuss next.

A few years later, the same authors report further on the simply typed cryptographic  $\lambda$ -calculus and construct a logical relation for proving contextual equivalences for it [Sumii and Pierce 2003]. Another year later, they report on a bisimulation that can be used to prove contextual equivalence in  $\lambda^\sigma$ , which they originally constructed for proving the conjectured full abstraction [Sumii and Pierce 2004]. In this last paper, the compiler is presented for an untyped version of the cryptographic  $\lambda$ -calculus (as in this text), which renders it significantly simpler.

Sealing was also used to enforce polymorphic contracts in PLT Scheme by Guha et al. [2007]. While technically similar, the assumptions in that work are somewhat different than in the above, as contracts are about protecting the context from misbehaving terms, while Sumii and Pierce’s compiler protects trusted terms from a misbehaving context. Like other work discussed in this paper, Guha et al.’s polymorphic contracts fail to enforce the degeneracy of Univ, so if they satisfy a form of parametricity, it would have to be TWLR-based.

*Gradual typing.* Gradual typing is a specific way of combining dynamic and static typing in a single language, intended to create a gradual migration path from untyped to typed codebases. Since it was originally proposed by Siek and Taha [2006], gradual typing has received a lot of attention: gradual extensions have been constructed for many different type systems, the notion of blame was adapted from the world of contracts [Findler and Felleisen 2002] to track the origin of a dynamic cast failure [Wadler and Findler 2009], correctness criteria were studied [Siek et al. 2015], the process of constructing a gradually-typed version of a pre-existing type system was to some extent automated [Cimini and Siek 2016; Garcia et al. 2016] and last but not least, the entire approach has also been declared dead because of severe performance issues [Takikawa et al. 2016].

Prior to the work on gradual typing, calculi which combined statically-typed languages with a universal type for interacting with untyped code, have been proposed by Henglein [1994] and Abadi et al. [1991].

*Gradual typing and parametric polymorphism.* As an instance of a multi-language (as proposed by Matthews and Findler [2009]), Matthews and Ahmed [2008] construct a language that can embed both Scheme (i.e., an untyped lambda calculus) and ML (i.e., a parametrically polymorphic typed lambda calculus) using a notion of dynamic casts from one to the other. They enforce parametric polymorphism using a notion of run-time type generation<sup>13</sup> and prove a parametricity result. An error in the proof was later corrected [Ahmed et al. 2011c].

Next, Ahmed et al. [2009b, 2011b] present the first version of the polymorphic blame calculus, a gradually-typed extension of System F that includes a notion of blame. The authors prove a number of correctness results, but Perconti found an intractable error in one of the proofs later [Ahmed et al. 2011a]. However, they do not consider parametricity or preservation of System F contextual equivalences (i.e., fully abstract embedding).

Ahmed et al. [2017] present  $\lambda_B$ : a new version of the polymorphic blame calculus rectifying certain “topsy-turvy” aspects of its operational semantics (see Section 6). Additionally, they prove a

<sup>13</sup>They call this sealing, but since their seals have no run-time representation, we prefer the term run-time type generation.

parametricity result for this calculus, which we have discussed from the perspective of our results in Section 7.3.

Igarashi et al. [2017] also present a new version of the polymorphic blame calculus, as an internal runtime representation for System  $F_G$  a new gradually-typed extension of System F. They make certain special restrictions to the consistency and precision relation of the system (later criticised by others [Toro et al. 2019]), which allow them to prove a version of the gradual guarantee. However, it does not seem to prevent our counterexample to the fully abstract embedding of System F.

In an unpublished draft, Siek and Wadler [2016] study and interrelate three ways to achieve relational parametricity: universal types, runtime type generation and cryptographic sealing. They show translations from the polymorphic blame calculus from [Ahmed et al. 2017] into the cryptographic lambda calculus and back that they show to be simulations. They also study a calculus  $\lambda_G$ , obtained by removing the universal type  $\star$  and casts from the polymorphic blame calculus, but keeping runtime type generation. They show that embedding System F into  $\lambda_G$  is also fully abstract. Both of these full abstraction results tally with our observations: both System F and  $\lambda_G$  lack a universal type (making `Univ` degenerate) while both the polymorphic blame calculus and the cryptographic lambda calculus feature a universal type (making `Univ` non-degenerate). As such, the embedding of  $\lambda_G$  into the polymorphic blame calculus will not be fully abstract, supplementing their results.

Xie et al. [2018] present a gradually typed language with parametric polymorphism, focusing on the interplay of gradual typing with the subtyping relation in the presence of implicit polymorphism. The result of their work is a source language that elaborates to  $\lambda_B$ , which we discussed before. As such, the language provides the same form of parametricity than  $\lambda_B$  and we think our results apply to it as well.

Most recently, Toro et al. [2019] proposed a new gradually typed calculus with explicit polymorphism, based on the AGT methodology by Garcia et al. [2016]. The methodology allows them to construct a system that satisfies the refined criteria by Siek et al. [2015], except for the Dynamic Gradual Guarantee. They show that this property is in conflict with parametricity in their system, but they demonstrate a weaker property which they do satisfy. The gradual type is also a universal type in their system and they prove a TWLR-based parametricity.

*Universal types.* Universal types have been studied by Longley [2003]. Our observation and proof that System F's parametricity excludes a universal type is, to the best of our knowledge, novel.

*Alternatives for Full Abstraction.* The property of full abstraction was proposed by Abadi [1998]. Abate et al. [2019] provide a lattice of secure compilation criteria (dubbed robust compilation) that preserve classes of hyperproperties [Clarkson and Schneider 2010], i.e., arbitrary behaviours. The general nature of the framework makes it hard to make precise statements, but we believe our counterexamples would also disprove most of these alternative properties.

## 9 CONCLUSION

This work started out as an effort to prove Sumii and Pierce's conjecture, discussed in Section 4. Our failure to do so has perhaps proven more interesting than a success might have been. Specifically, we disprove the conjecture rather than proving it, we disprove another conjecture for languages with non-parametric parametricity and we identify what we see as an important problem in current parametrically polymorphic gradual languages: programmers reasoning about System F programs cannot trust contextual equivalences to remain valid in the gradually typed extended language. In addition to highlighting these problems, we discuss in Section 7 some ideas about how the issues might be solved. None of them seem easy to solve and the solutions we propose have downsides of their own, but we do believe they might be worth exploring further.

During the work on the Sumii-Pierce conjecture, we also gained some more high-level insights about variations of parametricity (type-world LR versus Reynolds-style LR) and their relation to the existence of universal types. These insights were not mentioned explicitly in the previous version of this paper [Devriese et al. 2018] and in discussions with experts in the field, we found that these insights were not very clearly understood. For this reason, this paper focuses very clearly on these more high-level insights and we hope they may improve understanding of the issues involved.

## ACKNOWLEDGMENTS

This research is partially funded by project grants from the Research Fund KU Leuven, and from the Research Foundation Flanders (FWO). This work was partially supported by the German Federal Ministry of Education and Research (BMBF) through funding for the CISPA-Stanford Center for Cybersecurity (FKZ: 13N1S0762). The authors thank Phil Wadler for interesting comments and suggestions.

## REFERENCES

- Martín Abadi. Protection in programming-language translations. In *ICALP'98*, pages 868–883, 1998.
- Martín Abadi and Gordon D. Plotkin. On protection by layout randomization. *ACM Transactions on Information and System Security*, 15:8:1–8:29, July 2012. ISSN 1094-9224. doi: 10.1145/2240276.2240279.
- Martín Abadi, Luca Cardelli, Benjamin Pierce, and Gordon Plotkin. Dynamic Typing in a Statically Typed Language. *ACM Trans. Program. Lang. Syst.*, 13(2):237–268, April 1991. ISSN 0164-0925. doi: 10.1145/103135.103138. URL <http://doi.acm.org/10.1145/103135.103138>.
- Martín Abadi, Luca Cardelli, Benjamin Pierce, and Didier Rémy. Dynamic typing in polymorphic languages. *JOURNAL OF FUNCTIONAL PROGRAMMING*, 5:92–103, 1995.
- Martín Abadi, Cédric Fournet, and Georges Gonthier. Secure implementation of channel abstractions. In *IEEE Symposium on Logic in Computer Science*, pages 105–116, 1998.
- Martín Abadi, Cédric Fournet, and Georges Gonthier. Secure communications processing for distributed languages. In *IEEE Symposium on Security and Privacy*, pages 74–88, 1999.
- Martín Abadi, Cédric Fournet, and Georges Gonthier. Authentication primitives and their compilation. In *Principles of Programming Languages*, pages 302–315. ACM, 2000. doi: 10.1145/325694.325734.
- Carmine Abate, Arthur Azevedo de Amorim, Roberto Blanco, Ana Nora Evans, Guglielmo Fachini, Catalin Hritcu, Théo Laurent, Benjamin C. Pierce, Marco Stronati, and Andrew Tolmach. When good components go bad: Formally secure compilation despite dynamic compromise. *CCS '18*, 2018.
- Carmine Abate, Roberto Blanco, Deepak Garg, Marco Hrițcu, Cătălin and Patrignani, and Jérémy Thibault. Journey beyond full abstraction: Exploring robust property preservation for secure compilation. In *2019 IEEE 32th Computer Security Foundations Symposium*, CSF 2019, June 2019.
- Pieter Agten, Raoul Strackx, Bart Jacobs, and Frank Piessens. Secure compilation to modern processors. In *Computer Security Foundations Symposium*, pages 171–185. IEEE, 2012. doi: 10.1109/CSF.2012.12.
- Amal Ahmed and Matthias Blume. Typed closure conversion preserves observational equivalence. In *International Conference on Functional Programming*, pages 157–168. ACM, 2008. ISBN 978-1-59593-919-7. doi: 10.1145/1411204.1411227.
- Amal Ahmed and Matthias Blume. An equivalence-preserving CPS translation via multi-language semantics. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming*, ICFP '11, pages 431–444. ACM, 2011. ISBN 978-1-4503-0865-6. doi: 10.1145/2034773.2034830. URL <http://doi.acm.org/10.1145/2034773.2034830>.
- Amal Ahmed, Derek Dreyer, and Andreas Rossberg. State-dependent Representation Independence. In *Principles of Programming Languages*, pages 340–353. ACM, 2009a. doi: 10.1145/1480881.1480925.
- Amal Ahmed, Jacob Matthews, Robert Bruce Findler, and Philip Wadler. Blame for all. In *Workshop on Script-to-Program Evolution (STOP)*, pages 1–13, 2009b.
- Amal Ahmed, Robert Bruce Findler, Jeremy G. Siek, and Philip Wadler. Blame for all. Technical report, 2011a. URL <https://plt.eecs.northwestern.edu/blame-for-all/>.
- Amal Ahmed, Robert Bruce Findler, Jeremy G. Siek, and Philip Wadler. Blame for all. In *Principles of Programming Languages*, pages 201–214, 2011b.
- Amal Ahmed, Lindsey Kuper, and Jacob Matthews. Parametric polymorphism through run-time sealing, or, Theorems for low, low prices!, April 2011c. URL <http://www.ccs.neu.edu/home/amal/papers/paramseal-tr.pdf>.
- Amal Ahmed, Justin Damner, Jeremy G. Siek, and Philip Wadler. Theorems for free for free. In *ICFP*, 2017.

- Andrew W. Appel and David McAllester. An Indexed Model of Recursive Types for Foundational Proof-carrying Code. *ACM Trans. Program. Lang. Syst.*, 23(5):657–683, September 2001. ISSN 0164-0925. doi: 10.1145/504709.504712.
- Robert Atkey, Neil Ghani, and Patricia Johann. A relationally parametric model of dependent type theory. In *Principles of Programming Languages*, pages 503–515. ACM, 2014. doi: 10.1145/2535838.2535852.
- Michele Bugliesi and Marco Giunti. Secure implementations of typed channel abstractions. In *Principles of Programming Languages*, pages 251–262. ACM, 2007. doi: 10.1145/1190216.1190253.
- Matteo Cimini and Jeremy G. Siek. The Gradualizer: A Methodology and Algorithm for Generating Gradual Type Systems. In *Principles of Programming Languages*. ACM, 2016. doi: 10.1145/2837614.2837632.
- Michael R. Clarkson and Fred B. Schneider. Hyperproperties. *J. Comput. Secur.*, 18(6):1157–1210, September 2010.
- Dominique Devriese, Marco Patrignani, and Frank Piessens. Fully-abstract compilation by approximate back-translation. In *Principles of Programming Languages*, pages 164–177, 2016.
- Dominique Devriese, Marco Patrignani, Frank Piessens, and Steven Keuchel. Modular, Fully-abstract Compilation by Approximate Back-translation. *Logical Methods in Computer Science*, 13(4 lmc:4011), October 2017. doi: 10.23638/LMCS-13(4:2)2017.
- Dominique Devriese, Marco Patrignani, and Frank Piessens. Parametricity versus the universal type. In *Proceedings of the 45th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2018, Los Angeles, CA, USA, 2016, 2018*.
- Derek Dreyer, Amal Ahmed, and Lars Birkedal. Logical step-indexed logical relations. *Logical Methods in Computer Science*, 7(2), 2011.
- Derek Dreyer, Ralf Jung, Jan-Oliver Kaiser, Hoang-Hai Dang, and David Swasey. Semantics of type systems lecture notes, 2018. Available at: <https://plv.mpi-sws.org/semantics/2017/lecturenotes.pdf>.
- Matthias Felleisen. On the expressive power of programming languages. In *Selected Papers from the Symposium on 3rd European Symposium on Programming, ESOP '90*, pages 35–75, New York, NY, USA, 1991. Elsevier North-Holland, Inc.
- Robert Bruce Findler and Matthias Felleisen. Contracts for Higher-order Functions. In *International Conference on Functional Programming*. ACM, 2002. doi: 10.1145/581478.581484.
- Cedric Fournet, Nikhil Swamy, Juan Chen, Pierre-Evariste Dagand, Pierre-Yves Strub, and Benjamin Livshits. Fully abstract compilation to JavaScript. In *Principles of Programming Languages*, pages 371–384. ACM, 2013. doi: 10.1145/2429069.2429114.
- Ronald Garcia, Alison M. Clark, and Éric Tanter. Abstracting Gradual Typing. In *Principles of Programming Languages*. ACM, 2016. doi: 10.1145/2837614.2837670.
- Jean-Yves Girard. *Interprétation Fonctionnelle et Élimination Des Coupures de l'arithmétique d'ordre Supérieur*. PhD thesis, Université Paris VII, 1972.
- Arjun Guha, Jacob Matthews, Robert Bruce Findler, and Shriram Krishnamurthi. Relationally-parametric Polymorphic Contracts. In *Symposium on Dynamic Languages*. ACM, 2007. doi: 10.1145/1297081.1297089.
- Fritz Henglein. Dynamic typing: Syntax and proof theory. *Science of Computer Programming*, 22(3):197–230, June 1994. ISSN 0167-6423. doi: 10.1016/0167-6423(94)00004-2.
- Yuu Igarashi, Taro Sekiyama, and Atsushi Igarashi. On polymorphic gradual typing. In *International Conference on Functional Programming*. ACM, 2017.
- Radha Jagadeesan, Corin Pitcher, Julian Rathke, and James Riely. Local memory via layout randomization. In *Computer Security Foundations Symposium*, pages 161–174. IEEE Computer Society, 2011. doi: 10.1109/CSF.2011.18.
- Yannis Juglaret, Catalin Hritcu, Arthur Azevedo de Amorim, Benjamin C. Pierce, Antal Spector-Zabusky, and Andrew Tolmach. Towards a fully abstract compiler using micro-policies: Secure compilation for mutually distrustful components. *CoRR*, abs/1510.00697, 2015. URL <http://arxiv.org/abs/1510.00697>.
- Yannis Juglaret, Cătălin Hrițcu, Arthur Azevedo de Amorim, and Benjamin C. Pierce. Beyond good and evil: Formalizing the security guarantees of compartmentalizing compilation. In *Computer Security Foundations Symposium*, 2016.
- Adriaan Larmuseau, Marco Patrignani, and Dave Clarke. A secure compiler for ML modules. In *Programming Languages and Systems - 13th Asian Symposium*, pages 29–48, 2015. doi: 10.1007/978-3-319-26529-2\_3. URL [http://dx.doi.org/10.1007/978-3-319-26529-2\\_3](http://dx.doi.org/10.1007/978-3-319-26529-2_3).
- Adriaan Larmuseau, Marco Patrignani, and Dave Clarke. Implementing a secure abstract machine. In *Symposium on Applied Computing*, pages 2041–2048. ACM, 2016. ISBN 978-1-4503-3739-7. doi: 10.1145/2851613.2851796. URL <http://doi.acm.org/10.1145/2851613.2851796>.
- John R. Longley. Universal types and what they are good for. In *Domain Theory, Logic and Computation*, Semantic Structures in Computation, pages 25–63. Springer, Dordrecht, 2003. doi: 10.1007/978-94-017-1291-0\_2.
- Jacob Matthews and Amal Ahmed. Parametric polymorphism through run-time sealing or, theorems for low, low prices! volume 4960 of *LNCS*, pages 16–31. 2008.
- Jacob Matthews and Robert Bruce Findler. Operational Semantics for Multi-language Programs. *ACM Trans. Program. Lang. Syst.*, 31(3):12:1–12:44, April 2009. ISSN 0164-0925. doi: 10.1145/1498926.1498930.

- John C. Mitchell. Representation independence and data abstraction. In *Proceedings of the 13th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, POPL '86, pages 263–276, New York, NY, USA, 1986. ACM. doi: 10.1145/512644.512669. URL <http://doi.acm.org/10.1145/512644.512669>.
- John C. Mitchell. On abstraction and the expressive power of programming languages. *Science of Computer Programming*, 21(2):141 – 163, 1993. ISSN 0167-6423.
- James H. Morris, Jr. Protection in programming languages. *Commun. ACM*, 16(1):15–21, January 1973a. ISSN 0001-0782. doi: 10.1145/361932.361937.
- James H. Morris, Jr. Types are not sets. In *Principles of Programming Languages*, pages 120–124, 1973b.
- Georg Neis, Derek Dreyer, and Andreas Rossberg. Non-parametric parametricity. In *ICFP '09*, pages 135–148, 2009.
- Georg Neis, Derek Dreyer, and Andreas Rossberg. Non-parametric parametricity. *Journal of Functional Programming*, 21: 497–562, 2011. ISSN 1469-7653.
- Max S. New, William J. Bowman, and Amal Ahmed. Fully abstract compilation via universal embedding. In *International Conference on Functional Programming*, pages 103–116. ACM, 2016. doi: 10.1145/2951913.2951941.
- Joachim Parrow. Expressiveness of process algebras. *Elec. Not. Theo. Comp. Sci.*, 209(0):173 – 186, 2008.
- Marco Patrignani and Deepak Garg. Secure Compilation and Hyperproperties Preservation. In *Proceedings of the 30th IEEE Computer Security Foundations Symposium CSF 2017, Santa Barbara, USA*, CSF 2017, 2017.
- Marco Patrignani and Deepak Garg. Robustly safe compilation. In *Programming Languages and Systems - 28th European Symposium on Programming, ESOP 2019, ESOP'19*, 2019.
- Marco Patrignani, Pieter Agten, Raoul Strackx, Bart Jacobs, Dave Clarke, and Frank Piessens. Secure compilation to protected module architectures. *ACM Trans. Program. Lang. Syst.*, 37:6:1–6:50, April 2015. ISSN 0164-0925. doi: 10.1145/2699503.
- Marco Patrignani, Dominique Devriese, and Frank Piessens. On Modular and Fully Abstract Compilation. In *Computer Security Foundations Symposium*, 2016.
- Marco Patrignani, Amal Ahmed, and Dave Clarke. Formal approaches to secure compilation a survey of fully abstract compilation and related work. *ACM Comput. Surv.*, 51(6):125:1–125:36, January 2019.
- Benjamin Pierce. *Types and Programming Languages*. MIT Press, 2002.
- Benjamin Pierce and Eijiro Sumii. Relating cryptography and polymorphism. manuscript, 2000. URL <http://www.kb.ecei.tohoku.ac.jp/~sumii/pub/infohide.pdf>.
- Gordon D. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, 5:223–255, 1977.
- John C. Reynolds. Towards a theory of type structure. In *Programming Symposium*, volume 19 of *Lecture Notes in Computer Science*, pages 408–423. Springer-Verlag, 1974.
- John C. Reynolds. Types, abstraction, and parametric polymorphism. In *Information Processing*, pages 513–523. North Holland, 1983.
- Andreas Rossberg. Generativity and dynamic opacity for abstract types. In *Proceedings of the 5th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, PPDP '03, pages 241–252, New York, NY, USA, 2003. ACM. ISBN 1-58113-705-2. doi: 10.1145/888251.888274. URL <http://doi.acm.org/10.1145/888251.888274>.
- Manfred Schmidt-Schauß, David Sabel, Joachim Niehren, and Jan Schwinghammer. Observational program calculi and the correctness of translations. *Theoretical Computer Science*, 577:98 – 124, 2015. ISSN 0304-3975. doi: <http://dx.doi.org/10.1016/j.tcs.2015.02.027>.
- Jeremy Siek and Philip Wadler. The key to blame: Gradual typing meets cryptography. draft, 2016. URL <http://homepages.inf.ed.ac.uk/wadler/papers/blame-key/blame-key.pdf>.
- Jeremy G. Siek and Walid Taha. Gradual typing for functional languages. In *SCHEME*, pages 81–92, 2006.
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. Refined criteria for gradual typing. In *Summit on Advances in Programming Languages*, volume 32 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 274–293, Dagstuhl, Germany, 2015. Schloss Dagstuhl-Leibniz-Zentrum für Informatik. doi: 10.4230/LIPIcs.SNAPL.2015.274.
- Richard Statman. A local translation of untyped  $[\lambda]$  calculus into simply typed  $[\lambda]$  calculus. Technical report, 1991. URL <http://repository.cmu.edu/cgi/viewcontent.cgi?article=1454&context=math>.
- Christopher Strachey. Fundamental Concepts in Programming Languages. *Higher-Order and Symbolic Computation*, 13(1-2): 11–49, April 2000. ISSN 1388-3690, 1573-0557. doi: 10.1023/A:1010000313106.
- Eijiro Sumii and Benjamin C. Pierce. Logical Relations for Encryption. *J. Comput. Secur.*, 11(4):521–554, July 2003. ISSN 0926-227X.
- Eijiro Sumii and Benjamin C. Pierce. A bisimulation for dynamic sealing. In *Principles of Programming Languages*, pages 161–172, 2004.
- Asumu Takikawa, Daniel Feltey, Ben Greenman, Max S. New, Jan Vitek, and Matthias Felleisen. Is Sound Gradual Typing Dead? In *Principles of Programming Languages*. ACM, 2016. doi: 10.1145/2837614.2837630.
- Matias Toro, Elizabeth Labrada, and Éric Tanter. Gradual Parametricity, Revisited. *Proc. ACM Program. Lang.*, 3(POPL): 17:1–17:30, January 2019. ISSN 2475-1421. doi: 10.1145/3290330.

- Philip Wadler. Theorems for free! In *Functional Programming Languages and Computer Architecture*, pages 347–359. ACM, 1989.
- Philip Wadler. The Girard-Reynolds isomorphism (second edition). *Theoretical Computer Science*, 375(1), 2007. ISSN 0304-3975. doi: 10.1016/j.tcs.2006.12.042.
- Philip Wadler and Robert Bruce Findler. Well-typed programs can't be blamed. In *Programming Languages and Systems*, pages 1–16. Springer, Berlin, Heidelberg, March 2009. doi: 10.1007/978-3-642-00590-9\_1.
- Robert N. M. Watson, Jonathan Woodruff, Peter G. Neumann, Simon W. Moore, Jonathan Anderson, David Chisnall, Nirav H. Dave, Brooks Davis, Khilan Gudka, Ben Laurie, Steven J. Murdoch, Robert M. Norton, Michael Roe, Stacey D. Son, and Munraj Vadera. CHERI: A hybrid capability-system architecture for scalable software compartmentalization. In *IEEE Symposium on Security and Privacy*, 2015. doi: 10.1109/SP.2015.9.
- Ningning Xie, Xuan Bi, and Bruno C. d. S. Oliveira. Consistent Subtyping for All. In Amal Ahmed, editor, *Programming Languages and Systems*, Lecture Notes in Computer Science, pages 3–30. Springer International Publishing, 2018.