

# Two Parametricities versus Three Universal Types\*

DOMINIQUE DEVRIESE, KU Leuven, Belgium  
MARCO PATRIGNANI, University of Trento, Italy  
FRANK PIESSENS, KU Leuven, Belgium

The formal calculus System F models the essence of polymorphism and abstract data types, features that exist in many programming languages. The calculus' core property is parametricity: a theorem expressing the language's abstractions and validating important principles like information hiding and modularity.

When System F is combined with features like recursive types, mutable state, continuations or exceptions, 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 clear how this formulation should change 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 satisfy forms of parametricity (with Kripke worlds carrying semantic interpretations of types). However, the relative power of the modified formulations of parametricity with respect to others and the relative expressiveness of System F with and without these extensions are unknown.

In this paper, we explain that the aforementioned different settings have a common characteristic: they do not enforce or preserve the lexical scope of System F's type variables. Formally, this results in the existence of a *universal type* (note: this is not the same as a *universally-quantified type*). We explain why standard parametricity is incompatible with such a type and how type worlds resolve this. 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. 2022. Two Parametricities versus Three Universal Types. *J. ACM* 0, 0, Article 0 (2022), 43 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

*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 [Patrignani 2020].*

\*A *universal type* is not the same as a *universally-quantified type*. Following Longley [2003] (and category theory intuition), we use *universal type* as a term for a type which any other type can be embedded into and extracted from.

Authors' addresses: Dominique Devriese, DistriNet, KU Leuven, Leuven, Belgium, [name.surname@kuleuven.be](mailto:name.surname@kuleuven.be); Marco Patrignani, University of Trento, Trento, Italy, [name.surname@unitn.it](mailto:name.surname@unitn.it); Frank Piessens, DistriNet, KU Leuven, Leuven, Belgium, [name.surname@cs.kuleuven.be](mailto: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).

© 2022 Association for Computing Machinery.

0004-5411/2022/0-ART0 \$15.00

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

## 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 (LR): an (often relational) property about values and terms, derived from their type [Reynolds 1983]. A Fundamental Theorem of Logical Relations or Abstraction Theorem 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 formulation of parametricity must be adapted when language 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. Adding mutable state changes the situation again: it is now possible for  $f$  to return its argument in some invocations and to diverge in others, even though the choice can still not depend on the argument. 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 behavioural differences are reflected in the different definition of the logical relation used to formalize parametricity. Thus, formulations of parametricity using logical relations capture semantic guarantees about the language as a whole.<sup>1</sup>

In many cases (like the addition of state mentioned above), it is well understood that a particular change in the formulation of parametricity introduces changes both in the set of equivalences it implies as well as in the logical relation used to state it (see e.g., [Dreyer et al. 2010]). However, there are also cases where a certain formulation of parametricity is motivated by technical considerations, but it is not clear what effects this change has on equivalence reasoning and whether the change could have been somehow avoided. 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], Toro et al. [2019], and New et al. [2019a] on parametricity in gradually-typed variants of System F. Although they work in different contexts, these authors prove (a form of) parametricity with a logical relation indexed by System F types (or a superset of these types), with the exception of Sumii and Pierce who index with simple types for cryptographic primitives. Additionally, the logical relation used by all these parametricity results (including Sumii and Pierce's) follow a particular pattern: they are indexed by a *type world*. For this reason, we refer to these LRs as type-world logical relations (TWLRs). TWLRs should not be regarded as a clearly delineated subclass of logical relations, but rather as a design pattern that can be applied and varied upon in the design of a logical relation.

The type worlds in TWLRs are a form of Kripke worlds: they capture a set of assumptions with respect to which the logical relation holds (or does not hold). There is also an order relation on worlds that expresses when one world (a “future world”) represents a stronger set of assumptions

<sup>1</sup> In this paper we strive to reserve the term parametricity for the informal property that polymorphic functions must preserve relatedness of values of parametric types, following Strachey's [Strachey 2000] and Reynolds' [Reynolds 1983] use of the term. However, it is sometimes hard to formally draw the line between this restricted interpretation of parametricity and other properties (like purity or forms of type safety) that happen to be formulated together with the intuitive notion of parametricity in the Fundamental Theorem of Logical Relations for a specific logical relation. We will refer to this formal theorem as “a formulation of parametricity”. Because the line between (informal) parametricity and a formal formulation of it is not very clear, we are not always very strict about this distinction.

than another and whenever the logical relation holds with respect to a world, it automatically also holds for future worlds. While traditional logical relations (e.g., Dreyer et al. [2011b]; Reynolds [1983]) keep track of semantic interpretations for type variables in a type environment (as we discuss later in Section 2.3), the logical relations used by Ahmed et al.; Neis et al.; New et al.; Sumii and Pierce; Toro et al. carry semantic interpretations for dynamically-allocated opaque type variables or seals in the Kripke world. Additionally, instantiating polymorphic functions (or allocating fresh seals in the case of Sumii and Pierce [2003]) results in terms that are only related in a world that stores the relation between the applied types as the type interpretation for the instantiated type (or seal).

So why is the same kind of TWLR used in these three different domains? In this paper, we suggest this is because they are used in settings where the lexical scope of type variables is not enforced. Thus, in types like  $\forall X. X \rightarrow A$  (with  $X$  not free in  $A$ ), it is possible for values of type  $X$  to escape from their scope, be stored as values of a type that does not mention  $X$  (such as  $A$ ) and somehow be recovered from there as a value of type  $X$  again.

Instead of TWLRs, other work that predates TWLRs used logical relations that enforced such lexical scoping, similar to Reynolds' original formulation [Reynolds 1983]. Thus, we refer to these traditional logical relations as Reynolds-style logical relations (RLRs).

Although some of our insights might appear obvious, the importance of understanding the pattern of type-world logical relations can be seen in the literature, in two ways. First, if we look at the history of one of the aforementioned results [Ahmed et al. 2017], then we see that an 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 our results make it clear that the RLR could not possibly have been compatible with the language's non-lexically scoped type variables, i.e., the theorem was wrong, not just the proof. The flaw was 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].

A second consequence is that researchers had wrong expectations of how program equivalence changes when not enforcing lexical scope of type variables in extensions of System F. Concretely, our insights allow us to disprove two long-standing 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, 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>2</sup> Compiler full abstraction has been proven for compilers that rely

<sup>2</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 9.

148 on address-space layout randomisation [Abadi and Plotkin 2012; Jagadeesan et al. 2011], or secure  
149 enclaves [Agten et al. 2012; Larmuseau et al. 2015, 2016; Patrignani et al. 2015, 2016], tagged  
150 architectures [Juglaret et al. 2015, 2016], dynamic type checking in JavaScript [Fournet et al. 2013],  
151 typed closure conversion [Ahmed and Blume 2008], cryptographic primitives [Abadi et al. 1998,  
152 1999, 2000; Bugliesi and Giunti 2007] etc, we refer the interested reader to the survey of Patrignani  
153 et al. [2019].

154 In a paper from the year 2000, Pierce and Sumii [2000] proposed a compiler from System F to a  
155 cryptographic lambda calculus, which enforces parametricity using a form of idealised encryption  
156 primitives (sealing) called lambda-seal ( $\lambda^\sigma$ ). They conjectured that this compiler was fully abstract,  
157 owing part of the complexity of such a proof to the target language being more expressive than the  
158 source. Their conjecture has received further research attention, but remains open to this day [Siek  
159 and Wadler 2016; Sumii and Pierce 2004]. In other work, the same authors proposed a TWLR for  
160 the cryptographic lambda calculus [Sumii and Pierce 2003].

161  
162  
163 *Non-parametric parametricity.* Some programming languages include a way to perform *intensional*  
164 *type analysis* through a type cast operator [Abadi et al. 1995; Rossberg 2003]. This appears to be in  
165 direct conflict with parametric polymorphism, possibly violating parametricity and representation  
166 independence guarantees [Mitchell 1986]. Researchers have proposed runtime type generation as  
167 a way to regain parametricity for languages with a type cast operator. Ideally, when an abstract  
168 type is defined, a fresh type name should also be generated at runtime. Such a name should then be  
169 used as a runtime representative of the abstract type for type analysis purposes.

170  
171 Runtime type generation has been proven to indeed provide parametricity guarantees for System F  
172 terms that interact with terms that can perform type casts [Neis et al. 2009]. This guarantee has  
173 been dubbed *non-parametric parametricity*.

174 Neis et al. [2009] also conjectured that their way of using runtime type generation preserves *any*  
175 System F abstraction in their type cast language. Although they do not explain the reasons behind  
176 this conjecture, we have an idea of the intuition behind it. Parametricity is the most powerful  
177 abstraction programmers think that System F has, so it seems logical to believe that once that  
178 has been preserved, all other abstractions will follow. These authors also use a TWLR for stating  
179 parametricity in their system, and our results let us conclude that their version is weaker than the  
180 parametricity one states with RLR.

181  
182  
183 *Gradual typing.* The final field where we contribute new insight is gradual typing. In order to allow  
184 programmers to incrementally migrate large, untyped code bases to a typed programming language,  
185 gradual programming languages allow for typed and untyped code to interact. Such languages  
186 generally strive to preserve the benefits of the statically-typed components of an application, even  
187 when interacting with untyped components. Such benefits include performance benefits, absence  
188 of runtime type errors but also benefits for reasoning [de Amorim et al. 2020; New et al. 2019b;  
189 Toro et al. 2018]. While the literature mentions several ways to formalise the former two properties,  
190 the latter has received less attention.

191 Based on a suggestion in the conference version of this paper [Devriese et al. 2018], Jacobs et al.  
192 [2021] have recently proposed to formally express that a gradual language preserves the reasoning  
193 principles of the typed language by reusing the same notion of fully abstract compilation that we  
194 mentioned above. Specifically, if the embedding of the typed language into the gradual language  
195  
196

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>3</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], Toro et al. [2019] and New et al. [2019a]. The latter two papers also formulate parametricity using a TWLR in the polymorphic blame calculus. In these cases it is also unclear what are the benefits of relying on TWLRs has as opposed to relying on RLR.

*Three questions, one answer.* To answer these questions, we have to define non-lexical scoping of type variables. Our best formal characterisation is based on the existence of what we call a *universal type*.<sup>4</sup> A universal type is a type which any other type can be embedded into and extracted from. Thus, by this definition, the term *universal type* is broader than Abadi et al. [1991]'s *dynamic type* or Siek and Taha [2006]'s *gradual type*: 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.<sup>5</sup> We point out that enforcing the lexical scope of type variables forbids the existence of a universal type such as the following one:

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

Type  $\mathbf{Univ}$  expresses the existence of a universal 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  $\mathbf{Univ}$ , but we prove that they are all degenerate in the sense that mapping a value into the universal type and back must necessarily diverge. Our key finding is that this degeneracy can be proven using RLRs (and this illustrates the incompatibility of RLRs with universal types) but it cannot be proven using TWLRs.

Thus, by clarifying the subtleties of TWLRs and RLRs, we solve these aforementioned open problems and answer these questions negatively. Concretely, 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 cast, or any of the published polymorphic blame calculi.

More in detail, Sumii and Pierce's compiler fails to enforce this degeneracy of  $\mathbf{Univ}$ . 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  $\mathbf{Univ}$ . 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  $\mathbf{Univ}$ , despite the presence of runtime type generation primitives.

<sup>3</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>4</sup> This seems to imply non-lexical scoping of type variables, so it suffices for our results. Ultimately, we think that the non-lexical scope of type variables is the more fundamental characteristic of the systems we are interested in.

<sup>5</sup> The term universal type was introduced by Longley [2003] based on a similarity to *universal objects* in category theory and related terms have been used in the literature [New et al. 2016a]. As mentioned before, the term universal type should be clearly distinguished from *universally-quantified* types, i.e., types of the form  $\forall X. \tau$ .

246 The type  $\forall Z. Z$ , which is normally only inhabited by diverging terms, becomes a universal type in  
 247 the presence of a dynamic type cast, as any value can be embedded in it and extracted from it. By  
 248 relying on this type, System F does not embed fully abstractly into this language, contradicting  
 249 Neis et al.'s conjecture.

250 Finally, in the field of gradual typing, we demonstrate that existing polymorphic blame calculi  
 251 also break the degeneracy of [Univ](#).<sup>6</sup> Like Sumii and Pierce's target language, they also provide  
 252 a universal type: the type of untyped values  $\star$ , common to most gradual languages (also often  
 253 indicated as  $?$ ). Exploiting the existence of this type, we demonstrate that the polymorphic blame  
 254 calculus does not embed System F in a fully abstract way and as a result, they do not preserve  
 255 System F's parametricity.

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

264 *Outline.* We start our discussion by repeating the definition of System F and defining two logical  
 265 relations for it (an RLR and a TWLR), which we will use in subsequent proofs (Section 2). Then  
 266 we present type [Univ](#) and two key terms:  $t_u$  and  $t_\omega$ , and we discuss their contextual equivalence  
 267 (Section 3). We then prove that [Univ](#) is degenerate and that these terms are contextually equivalent  
 268 using the previously-defined logical relations (Section 4). Next, we present Sumii and Pierce's  
 269 compiler and disprove their conjecture by explaining how it treats  $t_\omega$  and  $t_u$  and fails to preserve  
 270 their equivalence (Section 5). We then present  $G$ , an extension of System F with type casts and  
 271 runtime type generation and demonstrate how embedding  $t_u$  and  $t_\omega$  into  $G$  breaks contextual  
 272 equivalence (Section 6). Next, we turn to gradual typing, introducing polymorphic blame calculi  
 273 and demonstrating how the contextual equivalence of  $t_u$  and  $t_\omega$  is also lost in the presence of these  
 274 calculi's universal type (Section 7). Finally, we discuss perspectives and consequences of our results  
 275 (Section 8), related work (Section 9) and we conclude (Section 10). We omit only few auxiliary  
 276 lemmas, their proofs and tedious, long reductions used in the main proofs, all of this can be found  
 277 in the supplementary material.

278 *Relation with the Previous Version.* This paper extends a paper by the same authors that was  
 279 published at POPL 2018 [Devriese et al. 2018]. In this version, the main changes include (1) a  
 280 more detailed analysis of the type [Univ](#) and how its degeneracy varies in the presence of effects  
 281 and value polymorphism (2) a presentation of a Reynolds-style or lexically-scoped (Section 2.3),  
 282 Kripke (Section 2.4) and type-world logical relation (Section 2.5), (3) a more elegant proof of the  
 283 degeneracy of [Univ](#) (Section 4), (4) a detailed analysis of the relation between the different logical  
 284 relations, non-lexically-scoped type variables and degeneracy of the universal type (Section 4.2),  
 285 and (5) the disproof of the conjecture by Neis et al. [2011] in Section 6.

## 288 2 SYSTEM F

289 We now consider System F itself (Section 2.1), the standard formulation of parametricity for it  
 290 (Section 2.2) and the different kinds of logical relations for it: Reynolds-style or lexically-scoped  
 291

292  
 293 <sup>6</sup> Unlike the previous two, this was not an existing conjecture we debunk.  
 294

(Section 2.3) and type-world (Section 2.5). To introduce type-world logical relations, we first explain their more general variant: Kripke logical relations (Section 2.4)

*Note on divergence and recursive types.* Technically, the language we should be using is System F with recursive types since that is the language considered by the conjectures we examine [Pierce and Sumii 2000]. We choose against including recursive types in our technical development since their presence is not central to our argument for disproving existing conjectures. Recursive types are required only insofar as System F is allowed some way to diverge. Fortunately, there are simpler ways to allow a language to diverge, as for example the addition of a diverging term (a solution also proposed by Pierce and Sumii [2000] and that we also adopt). Moreover, well-founded logical relations for languages with recursive types require *step indices*, and the addition of steps makes the technical development noisy without adding particular insights. Thus, in this paper we remove recursive types from System F and instead add a diverging term  $\omega$ .

## 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 **Unit** and **Bool** and product  $\tau_1 \times \tau_2$  types. In the figure, we present types  $\tau$ , values  $v$  and terms  $t$ , here the most peculiar addition is  $\omega_\tau$ , which is a diverging term of type  $\tau$ . 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  $E$ . There, we indicate the usual capture-avoiding substitution of value  $v$  (or type  $\tau'$ ) for variable  $x$  (or type variable  $X$ ) in term  $t$  (or type  $\tau$ ) as  $t[v/x]$  (as  $\tau[\tau'/X]$ ). We indicate lists of such substitutions as  $\gamma$ .

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 contexts  $\Delta$  and  $\Gamma$ , 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].

*Value polymorphism.* An interesting aspect of this definition of System F is that the body of a polymorphic functions  $\lambda X. t$  is a general term  $t$ . This means that it is possible for polymorphic functions to diverge or perform effects when instantiated. This means, for example, that the polymorphic type  $\forall X. \perp$  (for an empty type  $\perp$ ) is inhabited by the function  $\lambda X. \omega$ . In effectful variants of System F, we could similarly write a function of type  $\forall X. \text{Ref (Maybe X)}: \lambda X. \text{ref nothing}$ . For this function, it is important that different applications of this function are evaluated separately and produce different mutable reference locations, as otherwise, we could use a single location to store a value of one type and read a value of another, leading to a type error. In particular, this means that polymorphic lambdas and type applications cannot be erased away in an untyped execution scheme.

Although the choice of allowing arbitrary terms in the body of a polymorphic function is standard in formal accounts of System F, imperative polymorphic languages like ML make a different choice. They use *value polymorphism*, originally proposed by Wright [1995], which restricts ML polymorphism to values. This choice can be modeled in System F by restricting the syntax of polymorphic functions from  $\lambda X. t$  to  $\lambda X. v$ , i.e., the body of the polymorphic function must be a

344

Syntax:

345

346  $t ::= v \mid x \mid tt \mid t.1 \mid t.2 \mid \langle t, t \rangle \mid t\tau \mid \text{if } t \text{ then } t \text{ else } t \mid \text{pack } \langle \tau, t \rangle \text{ as } \exists X. \tau$ 

347

348  $\mid \text{unpack } t \text{ as } \langle X, x \rangle \text{ in } t \mid \omega_\tau$ 

349

349  $\text{Val } \ni v ::= \text{unit} \mid \text{true} \mid \text{false} \mid \lambda x : \tau. t \mid \langle v, v \rangle \mid \Lambda X. t \mid \text{pack } \langle \tau, v \rangle \text{ as } \exists X. \tau$ 

350

350  $\tau ::= \text{Unit} \mid \text{Bool} \mid \tau \rightarrow \tau \mid \tau \times \tau \mid X \mid \forall X. \tau \mid \exists X. \tau$ 

351

351  $\Gamma ::= \emptyset \mid \Gamma, (x : \tau) \quad \Delta ::= \emptyset \mid \Delta, X$ 

352

352  $E ::= [\cdot] \mid E t \mid v E \mid E.1 \mid E.2 \mid \langle E, t \rangle \mid \langle v, E \rangle \mid E \tau \mid \text{if } E \text{ then } t \text{ else } t \mid \text{unpack } E \text{ as } \langle X, x \rangle \text{ in } t$ 

353

Typing rules (excerpts):

355

$$\frac{\Delta; \Gamma \vdash \diamond \quad (x : \tau) \in \Gamma}{\Delta; \Gamma \vdash x : \tau} \quad \frac{\Delta; \Gamma, x : \tau \vdash t : \tau'}{\Delta; \Gamma \vdash \lambda x : \tau. t : \tau \rightarrow \tau'} \quad \frac{\Delta, X; \Gamma \vdash t : \tau}{\Delta; \Gamma \vdash \Lambda X. t : \forall X. \tau} \quad \frac{\Delta; \Gamma \vdash t : \tau' \rightarrow \tau \quad \Delta; \Gamma \vdash t' : \tau'}{\Delta; \Gamma \vdash tt' : \tau}$$

359

$$\frac{\Delta \vdash \tau' \quad \Delta; \Gamma \vdash t : \forall X. \tau}{\Delta; \Gamma \vdash t\tau' : \tau[\tau'/X]} \quad \frac{\Delta \vdash \tau \quad \Delta; \Gamma \vdash t : \tau[\tau'/X]}{\Delta; \Gamma \vdash \text{pack } \langle \tau', t \rangle \text{ as } \exists X. \tau : \exists X. \tau}$$

360

361

362

363

$$\frac{\Delta; \Gamma \vdash t : \exists X. \tau \quad \Delta \vdash \tau' \quad \Delta, X; \Gamma, x : \tau \vdash t_1 : \tau'}{\Delta; \Gamma \vdash \text{unpack } t \text{ as } \langle X, x \rangle \text{ in } t_1 : \tau'} \quad \frac{}{\Delta; \Gamma \vdash \omega_\tau : \tau}$$

364

365

Evaluation rules (excerpts):

366

367

368

369

$$\frac{t \hookrightarrow_0 t'}{E[t] \hookrightarrow E[t']} \quad \frac{}{(\lambda x : \tau. t) v \hookrightarrow_0 t[v/x]} \quad \frac{}{(\Lambda X. t) \tau \hookrightarrow_0 t[\tau/X]}$$

370

371

372

$$\frac{}{\text{unpack } (\text{pack } \langle \tau', v \rangle \text{ as } \exists X. \tau) \text{ as } \langle X', x \rangle \text{ in } t \hookrightarrow_0 t[v/x][\tau'/X']} \quad \frac{}{\omega_\tau \hookrightarrow_0 \omega_\tau}$$

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

373

374

375

376

377

378

379

380

## 2.2 Parametricity

381

382

383

384

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.

385

386

387

388

*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 implement this type as a  $k$ -tuple of booleans.

389

390

391

392

*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  $(v_1, v_2)$  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.

### 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.<sup>7</sup>

The Reynolds-style or lexically-scoped 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  $R$  on values of type  $\tau$ . The value relation  $\mathcal{V}[\cdot]^\rho$  for type variables  $X$  defers to the appropriate entry in  $\rho$  for type variables  $X$ . 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. We apply the type environment to the type  $\tau$  (denoted with  $\rho(\tau)$ ) to indicate type  $\tau$  with any type variable  $X$  replaced with its binding in  $\rho$ , i.e., with  $\rho(X).1$ . Polymorphic functions are accepted when they can be applied to an arbitrary type  $\tau'$  and an arbitrary relation  $R$  on  $\tau'$  to obtain a term in the appropriate term relation, with  $\rho$  extended with  $R$ . Finally, existential packages must contain a type  $\tau'$  and a value related at the appropriate type, extending  $\rho$  with some relation  $R$  on  $\tau'$ . Values in the value relation are also required to be syntactically well-typed, as denoted by function of type  $(v, \tau)$ , which simply checks  $\emptyset; \emptyset \vdash v : \tau$ . The term relation accepts terms that diverge or produce a suitable value.

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

<sup>7</sup>We will continue to use words like relation, related etc. despite the fact that they are unary.

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.

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

With these ingredients, we can define the relation  $\Delta; \Gamma \Vdash t : \tau$  on open terms. This relation accepts terms  $t$  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).*

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

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*). We do not provide proofs of these lemmas as they are quite standard [Dreyer et al. 2011b].

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

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  $t' \in \mathcal{E}[\tau]^\rho$ , then  $t \in \mathcal{E}[\tau]^\rho$ .*

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).**

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

We also mention the Boring lemma (Lemma 2.9)<sup>8</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*

$$\mathcal{E}[\tau]^{\rho_1} = \mathcal{E}[\tau]^{\rho_2}$$

<sup>8</sup> For lack of a better name, we call this lemma as in Dreyer et al. [2011b]'s lecture notes. Neis et al. [2011] call this the *irrelevance* lemma.

Finally, it is worth mentioning that this logical relation can be easily adapted to the use of value polymorphism, discussed in Section 2.1. This requires only a change to the case for polymorphic types  $\forall X. \tau$ , which then looks as follows:

$$\mathcal{V} \llbracket \forall X. \tau \rrbracket^\rho = \left\{ \lambda X. v \mid \begin{array}{l} \text{of type } (\lambda X. v, \forall X. \tau) \text{ and} \\ \forall R \in \text{Re1}(\tau'). v[\tau'/X] \in \mathcal{V} \llbracket \tau \rrbracket^{\rho, X \mapsto (\tau', R)} \end{array} \right\}$$

The change is limited: polymorphic function bodies are now restricted to values and the instantiated bodies are now required to be in the value relation rather than the term relation, as one might expect.

## 2.4 Kripke Logical Relations

For languages with additional features, particularly effects like higher-order state, we can replace the basic logical relation from the previous section with a Kripke logical relation. The idea is to index the logical relation with a form of *possible worlds* or *Kripke worlds*  $\mathbf{W}$  which represent a set of shared assumptions, often about shared state like the heap. Possible worlds are partially ordered by a relation  $\sqsubseteq$ , where  $\mathbf{W}' \sqsupseteq \mathbf{W}$  expresses that  $\mathbf{W}'$  represents a stronger set of assumptions than  $\mathbf{W}$ . A Kripke version of the logical relation from before looks like the one below (as a notation convention, we typeset elements of Kripke LR with a grey background, to distinguish them from elements of the RLR, which have no background). Note that we do not give a concrete definition of worlds in this section, since the goal here is just to make the reader familiar with their treatment – the next section will provide more concrete details about worlds.

$$\mathcal{V} \llbracket X \rrbracket^\rho = \rho(X).R$$

$$\mathcal{V} \llbracket \text{Unit} \rrbracket^\rho = \{(\mathbf{W}, \text{unit})\}$$

$$\mathcal{V} \llbracket \text{Bool} \rrbracket^\rho = \{(\mathbf{W}, \text{true}), (\mathbf{W}, \text{false})\}$$

$$\mathcal{V} \llbracket \tau \rightarrow \tau' \rrbracket^\rho = \left\{ (\mathbf{W}, \lambda x : \rho(\tau). t) \mid \begin{array}{l} \text{of type } (\lambda x : \rho(\tau). t, \tau \rightarrow \tau') \text{ and } \forall \mathbf{W}' \sqsupseteq \mathbf{W}. \forall v. \\ \text{if } (\mathbf{W}', v) \in \mathcal{V} \llbracket \tau \rrbracket^\rho \text{ then } (\mathbf{W}', (\lambda x : \tau. t) v) \in \mathcal{E} \llbracket \tau' \rrbracket^\rho \end{array} \right\}$$

$$\mathcal{V} \llbracket \tau \times \tau' \rrbracket^\rho = \left\{ (\mathbf{W}, \langle v, v' \rangle) \mid \begin{array}{l} \text{of type } (\langle v, v' \rangle, \tau \times \tau') \text{ and} \\ (\mathbf{W}, v) \in \mathcal{V} \llbracket \tau \rrbracket^\rho \text{ and } (\mathbf{W}, v') \in \mathcal{V} \llbracket \tau' \rrbracket^\rho \end{array} \right\}$$

$$\mathcal{V} \llbracket \forall X. \tau \rrbracket^\rho = \left\{ (\mathbf{W}, \lambda X. t) \mid \begin{array}{l} \text{of type } (\lambda X. t, \forall X. \tau) \text{ and } \forall \mathbf{W}' \sqsupseteq \mathbf{W}. \forall \tau', R \in \text{Re1}[\tau']. \\ (\mathbf{W}', t[\tau'/X]) \in \mathcal{E} \llbracket \tau \rrbracket^{\rho, X \mapsto (\tau', R)} \end{array} \right\}$$

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

$$\mathcal{E} \llbracket \tau \rrbracket^\rho = \{(\mathbf{W}, t) \mid \forall v. \text{if } t \dot{\hookrightarrow} v \text{ then } \exists \mathbf{W}' \sqsupseteq \mathbf{W}. (\mathbf{W}', v) \in \mathcal{V} \llbracket \tau \rrbracket^\rho \}$$

$$\mathcal{G} \llbracket \emptyset \rrbracket^\rho = \{\emptyset\}$$

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

The differences with before are that all cases of the logical relation now mention the world  $\mathbf{W}$  with respect to which values are related. Additionally, the expression relation existentially

quantifies over a future world in which resulting values will be related. This generally models the fact that operational steps may introduce fresh assumptions, for example about freshly allocated mutable variables, which the resulting values' relation depends on. Finally, the value relation for function types  $\tau \rightarrow \tau'$  and  $\forall X. \tau$  is polymorphically quantified over an arbitrary world  $W' \sqsubseteq W$  that extends the assumptions of the current world  $W$ . This represents the fact that these values must be valid whenever they are invoked, including when additional assumptions may have been introduced.

Again, the logical relation is easy to adapt to a value-polymorphic variant of System F, by modifying only the case for polymorphic function types:

$$\mathcal{V}[\forall X. \tau]^\rho = \left\{ (W, \lambda X. v) \mid \begin{array}{l} \text{of type } (\lambda X. v, \forall X. \tau) \text{ and } \forall W' \sqsupseteq W. \forall \tau', R \in \text{Rel}[\tau'] \\ (W', v[\tau'/X]) \in \mathcal{V}[\tau]^\rho, X \mapsto (\tau', R) \end{array} \right\}$$

As before, the instantiated function body is now required to be in the value relation rather than the term relation.

## 2.5 Logical Relation with Type Worlds

In this paper, we are not so interested in general effects but we focus on non-lexically-scoped type variables. For this reason, we are interested in a pattern in logical relation definitions that we dub Type World Logical Relations (TWLRs) [Ahmed et al. 2017; Neis et al. 2009; New et al. 2019a; Sumii and Pierce 2003; Toro et al. 2019]. Note that the term is not intended to denote a clearly delineated subclass of logical relations but rather a design pattern that may be used and varied upon in the definition of logical relations. A TWLR is a form of Kripke logical relation which uses Kripke worlds  $W$  to represent the interpretations of type variables instead of (or in addition to) the semantic type relations  $\rho$ . Additionally, when related polymorphic functions are instantiated with related types, the resulting terms are only related in a world that stores the relation between the applied types as the type interpretation for the instantiated type. This different treatment of type environments removes the requirement to enforce the lexical scope of type quantifiers and makes the logical relation compatible with universal types, as we will explain in Section 3.

In a TWLR, 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). Here, we present a TWLR variant of the logical relation from Section 2.3 which uses worlds that contain types and relations for instantiated type variables. The future world relation  $\sqsubseteq$  enforces that once a binding is added to a world, it can never be removed (this is true for Kripke LRs too, but in Section 2.4 we did not have to define what constitutes worlds).

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

$$W' \sqsupseteq W = W' \supseteq W$$

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

$$R \in \text{Rel}[\tau]$$

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

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. Fortunately, this is a well-known problem and a standard solution exists: step-indexing [Ahmed 2004; Ahmed et al. 2009a; Dreyer et al. 2011a].

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  $\mathbf{W}$ . As before, replacing all type variables  $\mathbf{X}$  with their bindings in  $\mathbf{W}$  in type  $\tau$  is denoted as  $\mathbf{W}(\tau)$ .

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

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

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

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

$$\mathcal{V}[\tau \times \tau'] = \left\{ (\mathbf{W}, \langle v, v' \rangle) \mid \begin{array}{l} \text{of type } (\langle v, v' \rangle, \tau \times \tau') \text{ and} \\ (\mathbf{W}, v) \in \mathcal{V}[\tau] \text{ and } (\mathbf{W}, v') \in \mathcal{V}[\tau'] \end{array} \right\}$$

$$\mathcal{V}[\forall \mathbf{X}. \tau] = \left\{ (\mathbf{W}, \lambda \mathbf{X}. t) \mid \begin{array}{l} \text{of type } (\lambda \mathbf{X}. t, \forall \mathbf{X}. \tau) \text{ and } \forall \mathbf{W}' \supseteq \mathbf{W}. \forall \tau', \mathbf{R} \in \text{Rel}[\tau']. \\ (\mathbf{W}' + (\mathbf{X}, \tau', \mathbf{R}), t[\tau'/\mathbf{X}]) \in \mathcal{E}[\tau] \end{array} \right\}$$

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

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

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

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

It is crucial to compare the case for  $\mathcal{V}[\forall \mathbf{X}. \tau]$  against the one from Section 2.3. Contrary to there, the instantiated term  $t[\tau'/\mathbf{X}]$  is only required to be related in worlds  $\mathbf{W}' + (\mathbf{X}, \tau', \mathbf{R})$  that store the relation  $\mathbf{R}$  as the semantic interpretation for the instantiated type variable  $\mathbf{X}$ .

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.

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

$$639 \quad \Delta; \Gamma \# t : \tau \stackrel{\text{def}}{=} \forall \mathbf{W} \vdash \Delta, \forall \mathbf{y} \in \mathcal{G}[\Gamma], (\mathbf{W}, \mathbf{ty}) \in \mathcal{E}[\tau]$$

641 The fundamental property of TWLRs is analogous to that for RLR. As in the previous section,  
642 we do not provide proofs of these lemmas, which can be derived from similar (binary) statements  
643 in the works of e.g., Ahmed et al. [2017]; Neis et al. [2009].

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

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

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

$$652 \quad (\mathbf{W}_1, \mathbf{t}) \in \mathcal{V}[\tau] \iff (\mathbf{W}_2, \mathbf{t}) \in \mathcal{V}[\tau]$$

654 Contrary to the RLR, this lemma does not hold for our TWLR. If we were to try and prove it, we  
655 would quickly get stuck in the cases for lambdas and big lambdas where we get a future world  $\mathbf{W}'$   
656 of, for example,  $\mathbf{W}_1$  but we have no way to relate it to world  $\mathbf{W}_2$ .

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

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

$$661 \quad \text{if } (\mathbf{W}_1, \mathbf{v}) \in \mathcal{V}[\tau] \text{ then } (\mathbf{W}_2, \mathbf{v}) \in \mathcal{V}[\tau]$$

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

665 The difference between the RLR of Section 2.3 and the TWLR of this section may appear technical.  
666 However, we will see in the next section that this technical change from RLR to TWLR renders a  
667 formulation of parametricity compatible with universal types. This fact is witnessed by several  
668 TWLR-based parametricity proofs for languages with such types [Ahmed et al. 2011c, 2017; Neis  
669 et al. 2009, 2011; New et al. 2019a; Sumii and Pierce 2003; Toro et al. 2019]. However, none of those  
670 papers observed that using a TWLR reduced the set of equivalences that the LR implies and that  
671 this reduction was necessary in the presence of non-lexically-scoped type variables.

672 Before concluding, it is again interesting to consider what our TWLR would look like for a  
673 value-polymorphic language:

$$674 \quad \mathcal{V}[\forall X. \tau] = \left\{ (\mathbf{W}, \Lambda X. \mathbf{v}) \mid \begin{array}{l} \text{of type } (\Lambda X. \mathbf{v}, \forall X. \tau) \text{ and } \forall \mathbf{W}' \supseteq \mathbf{W}. \forall \tau', \mathbf{R} \in \text{Rel}[\tau']. \\ (\mathbf{W}' + (X, \tau', \mathbf{R}), \mathbf{v}[\tau'/X]) \in \mathcal{V}[\tau] \end{array} \right\}$$

675 Observe that here, the instantiated body is still required to be related by the value relation and thus  
676 not able to allocate additional assumptions. This makes sense, since the body cannot allocate, for  
677 example, fresh mutable heap locations. However, one particular assumption is still added to the  
678 world in which the body is related: the instantiated of  $X$  to type  $\tau'$  and relation  $\mathbf{R}$ .

681 In the next Sections, we provide further insight into how the use of a TWLR underlines a  
682 compatibility with the universal type in the language. We do this by demonstrating the degeneracy  
683 of the universal type with LRLS and explaining why such a proof fails using a TWLR. Before  
684 showing these technical proofs, we must define the universal type.

### 3 THE TYPE Univ

As mentioned, in this paper we rely heavily on the type Univ:

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

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 (recall that  $\omega_X$  is a diverging term of type  $X$ ):

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

However, this value is *degenerate* in the sense that injecting a value into the packaged  $Y$  and extracting it again diverges. Note that it is not necessarily the function of type  $Y \rightarrow X$  that diverges. For example, we can construct the following inhabitants:

$$\begin{aligned} & \text{pack } \langle (\forall Z. Z), \lambda X. \langle \lambda_ : X. \omega_{\forall Z. Z}, \lambda y : (\forall Z. Z). y X \rangle \rangle \text{ as } \underline{\text{Univ}} \\ & \text{pack } \langle (\forall Z. Z), \lambda X. \langle \lambda_ : X. \lambda Z. \omega_Z, \lambda y : (\forall Z. Z). y X \rangle \rangle \text{ as } \underline{\text{Univ}} \end{aligned}$$

These other inhabitants instantiate  $Y$  to  $\forall Z. Z$  and make the function of type  $Y \rightarrow X$  simply use the received value of type  $\forall Z. Z$  to obtain the required value of type  $X$ . It is now the function of type  $X \rightarrow Y$  that diverges, either directly or after being applied to a type  $Z$ .

A crucial observation for this paper is that *all* System F values of type Univ are degenerate in the above 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>9</sup>

From another perspective, the degeneracy results from the lexical scope of type variable  $X$  in the polymorphic function of type  $\forall X. (X \rightarrow Y) \times (Y \rightarrow X)$ . In  $\lambda^F$ , implementations of this function are required to respect this lexical scope and importantly, the variable is not in scope at the point where a choice for existential variable  $Y$  needs to be made. A non-degenerate implementation of Univ essentially must somehow pass a value of type  $X$  as the result of the function of type  $X \rightarrow Y$  and recover it from the argument to the function of type  $Y \rightarrow X$ . This would require that value to survive exiting and reentering the lexical scope of type variable  $X$ , and degeneracy expresses exactly the impossibility of this, i.e.  $\lambda^F$ 's respect for the lexical scope of type variables like  $X$ .

Note that in this paper, lexical scoping of type variables in System F is regarded as an informal property. Our best attempt to characterize it uses the degeneracy of Univ and the existence of a universal type as a sufficient criterion.

#### 3.1 Two Contextually Equivalent Terms

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

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

The reason that  $t_u$  and  $t_\omega$  are contextually equivalent is that both will diverge when applied to any argument of type Univ. For  $t_u$ , this follows from the degeneracy of the type Univ, as we

<sup>9</sup> Obviously, the situation is entirely different if we swap the quantifications in the type:  $\text{Triv} \stackrel{\text{def}}{=} \forall X. \exists Y. (X \rightarrow Y) \times (Y \rightarrow X)$ .

demonstrate below, while for  $t_\omega$ , the term  $\omega_{\text{Unit}}$  in the body ensures divergence. Note that the degeneracy of  $\text{Univ}$  is essential: if the context were able to produce a non-degenerate value of type  $\text{Univ}$ , 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  $\lambda^F$ ).  $\emptyset; \emptyset \vdash t_u \simeq t_\omega : \text{Univ} \rightarrow \text{Unit}$ .

Note that we have chosen  $t_u$  and  $t_\omega$  because their equivalence is relatively easy to prove. However, we could have taken many other equivalences which follow from the degeneracy of  $\text{Univ}$  and particularly, we could have taken example terms which will both terminate with different results in the presence of a non-degenerate universal type:

$$\begin{aligned} t_1 &\stackrel{\text{def}}{=} \lambda x : \text{Univ}. \text{unpack } x \text{ as } \langle Y, x' \rangle \text{ in} \\ &\quad \text{let } x'' : (\text{Unit} \rightarrow Y) \times (Y \rightarrow \text{Unit}) = x' \text{ Unit in } x''.2 (x''.1 \text{ unit}); 1 \\ t_2 &\stackrel{\text{def}}{=} \lambda x : \text{Univ}. \text{unpack } x \text{ as } \langle Y, x' \rangle \text{ in} \\ &\quad \text{let } x'' : (\text{Unit} \rightarrow Y) \times (Y \rightarrow \text{Unit}) = x' \text{ Unit in } x''.2 (x''.1 \text{ unit}); 2 \end{aligned}$$

### 3.2 $\text{Univ}$ in Other Settings

Before we look at proving Theorem 3.1, it is useful to build a better intuition of the meaning and the cause of the degeneracy of  $\text{Univ}$ . To this end, it is useful to consider the type's properties in variants of System F.

For example, if we imagine a variant of System F with errors or exceptions, converting a function of type  $X$  to  $Y$  and back, must no longer necessarily diverge, but can now also result in an error or an exception. Intuitively, this is perhaps still a form of degeneracy, but formally, it is no longer true that  $t_u \simeq t_\omega$  and thus, the two terms no longer form a counterexample to full abstraction. However, while the concrete counterexample is broken, the more general observation remains: System F implies properties about the type  $\text{Univ}$  that are no longer true in the languages we discuss in the next sections, even when both languages are extended with errors or exceptions. Concretely, we can easily adapt the counterexample by making  $t_\omega$  not always diverge, but diverge after invoking the two functions of the  $\text{Univ}$  value.

$$\begin{aligned} t_u &\stackrel{\text{def}}{=} \lambda x : \text{Univ}. \text{unpack } x \text{ as } \langle Y, x' \rangle \text{ in} \\ &\quad \text{let } x'' : (\text{Unit} \rightarrow Y) \times (Y \rightarrow \text{Unit}) = x' \text{ Unit in } x''.2 (x''.1 \text{ unit}) \\ t'_\omega &\stackrel{\text{def}}{=} \lambda x : \text{Univ}. \text{unpack } x \text{ as } \langle Y, x' \rangle \text{ in} \\ &\quad \text{let } x'' : (\text{Unit} \rightarrow Y) \times (Y \rightarrow \text{Unit}) = x' \text{ Unit in } x''.2 (x''.1 \text{ unit}); \omega_{\text{Unit}} \end{aligned}$$

Our results should not be interpreted as strictly related to the specific example terms  $t_u$  and  $t_\omega$  and not even to the type  $\text{Univ}$ . It appears clear at least that  $\text{Univ}$  is not the only type for which an RLR implies properties that do not follow from a TWLR (see Section 4). For example, it appears clear that similar properties will hold for a variant of  $\text{Univ}$  with one component duplicated:

$$\exists Y. \forall X. (X \rightarrow Y) \times (X \rightarrow Y) \times (Y \rightarrow X)$$

Our modified example shows that the phenomenon also manifests itself in the presence of effects like errors and exceptions.

It is hard to identify a root cause for the phenomenon: some of the settings we consider feature a form of type generativity, all feature a universal type, but it is hard to say whether these are symptom or disease. The use of TWLRs suggests (at least to us) that it is essentially the lexical scoping of type variables in System F which is not enforced in the different settings we consider.



885 However, as mentioned before, we don't know how to formalize this intuitive property or argue  
886 that it is the root cause.

887 Extending System F with ML-like references breaks the counterexample in a seemingly more  
888 severe way. Concretely, consider extending System F in a standard way with types `Ref  $\tau$`  for heap  
889 references and constructs for allocating (`ref $\tau$  v`), assigning (`x := v`) and dereferencing (`!x` heap  
890 variables), see, e.g., Pierce [2002]. In this case, the type `Univ` gains inhabitants like the following:

$$891 \text{pack } \left\langle \text{Unit}, \lambda X. \text{let } r : \text{Ref (Maybe X)} = \text{ref}_X \text{ None in} \right. \\ 892 \left. \langle \lambda x : X. r := \text{Some } x; \text{unit}, \lambda \_ : \text{Unit}. !r \rangle \right\rangle \text{ as } \text{Univ}$$

894 Rather than attempting to store values of type `X` in an appropriately chosen `Y` in some way, this  
895 inhabitant instantiates `Y` to the non-informative type `Unit`. Instead, it allocates a mutable variable  
896 of type `X`, shared between the two functions. The first function then stores its argument in the  
897 heap variable for the second function to retrieve.

898 The existence of this inhabitant shows that ML references break our results as well and in  
899 this case, it is not as obvious how to recover equivalences that rely on lexical scoping of type  
900 variables. One might interpret this to mean that the degeneracy of `Univ` is an artifact of the purity  
901 of System F, and that it only holds in the absence of effects like mutable state. Perhaps what we  
902 call non-lexically-scoped type variables should simply be interpreted as impurity of polymorphic  
903 function applications and as such, not different from other effects?

904 We think it is not so simple and one way to see this is to consider what happens with the above  
905 example in a value-polymorphic language. Interestingly, value polymorphism breaks the `Univ`  
906 inhabitant mentioned above: it is no longer possible to allocate a reference cell that is shared by the  
907 functions from `X → Y` and `Y → X`. In fact, we expect degeneracy of `Univ` to hold in value-polymorphic  
908 variants of System F, even in the presence of ML references or other effects.

909 Nevertheless, it will be clear in Sections 5 to 7 that value polymorphism does not similarly  
910 break the inhabitants of `Univ` we discuss there. This demonstrates that even if one interprets the  
911 non-lexical scope of type variables in those languages as a form of effect, it is at least a form of  
912 effect that behaves rather differently than other effects, as it does not disappear from polymorphic  
913 functions with the introduction of value-polymorphism. From that point of view, we think our  
914 results remain relevant in the presence of effects, although value polymorphism appears essential  
915 when those effects can be used to create a communication channel for transmitting values of type  
916 `X` between the two functions in `Univ` (as discussed for ML references).

#### 917 4 PROVING DEGENERACY OF `Univ`

918 This section attempts to prove Theorem 3.1 (the contextual equivalence of `tu` and `tω` which we  
919 recap below) using both the RLR and the TWLR and in doing so, it focusses on the key result for  
920 this proof: showing that `Univ` is degenerate. This result is doable with the RLR (from Section 2.3),  
921 highlighting why they are incompatible with non-lexically-scoped type variables but the same fact  
922 cannot be shown with the TWLR (from Section 2.5).

$$923 \text{t}_u \stackrel{\text{def}}{=} \lambda x : \text{Univ}. \text{unpack } x \text{ as } \langle Y, x' \rangle \text{ in} \\ 924 \text{let } x'' : (\text{Unit} \rightarrow Y) \times (Y \rightarrow \text{Unit}) = x' \text{Unit in } x''.2 (x''.1 \text{unit}) \\ 925 \text{t}_\omega \stackrel{\text{def}}{=} \lambda x : \text{Univ}. \omega_{\text{Unit}}$$

926 The first step of the proof of Theorem 3.1 is captured by Lemma 4.1 (Diverging functions are  
927 contextually equivalent to an omega function), which states that a function that diverges for every  
928 argument is equivalent to a function whose body is the omega term.

LEMMA 4.1 (DIVERGING FUNCTIONS ARE CONTEXTUALLY EQUIVALENT TO AN OMEGA FUNCTION).

If  $\emptyset; \emptyset \vdash \lambda x : \tau'. t : \tau' \rightarrow \text{Unit}$  and (for all  $\emptyset; \emptyset \vdash v : \tau'$  we have that  $(\lambda x : \tau'. t) v \Downarrow$ )

Then  $\emptyset; \emptyset \vdash \lambda x : \tau'. t \simeq \lambda x : \tau'. \omega_{\text{Unit}} : \tau' \rightarrow \text{Unit}$

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. [2011a]. To avoid distracting from the main point of our paper, we do not offer a proof here.

With Lemma 4.1, it suffices to prove that  $t_u$  always diverges when supplied with a value of type  $\text{Univ}$  in order to conclude Theorem 3.1. This result we derive from the “degeneracy” of  $\text{Univ}$ , i.e., from the fact that  $\text{Univ}$  is only inhabited by diverging terms. We show this can be proven with the aid of the RLR in Lemma 4.2 below (Section 4.1) and also show that this is not provable with the TWLR (Section 4.2). We then discuss how this proof is carried out with other kinds of LR (Section 4.3).

#### 4.1 $\text{Univ}$ Degeneracy Using an RLR

LEMMA 4.2 ( $\text{Univ}$  IS DEGENERATE). For all  $\emptyset; \emptyset \vdash v : \text{Univ}$ , we have that  $t_u v \Downarrow$ .

PROOF. The proof proceeds largely in a standard way: we unfold the definition of value relation for the value of universal type 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 inhabitants is empty. Specifically, we will prove that it belongs to  $\mathcal{E} \llbracket X \rrbracket^{\dots, X \mapsto (\text{Unit}, \emptyset)}$ . By definition this means that the term must diverge or reduce to a value in  $\mathcal{V} \llbracket X \rrbracket^{\dots, X \mapsto (\text{Unit}, \emptyset)}$ . 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  $\mathcal{E} \llbracket X \rrbracket^{\dots, X \mapsto (\text{Unit}, \emptyset)}$ , the main trick we use relies on having two relations for  $X$  to inhabit. We will then instantiate the quantification over  $X$  with two different semantic interpretations.

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

Now take  $\emptyset; \emptyset \vdash v : \text{Univ}$ . By Theorem 2.5 (Fundamental property for RLR) with  $\emptyset; \emptyset \vdash v : \text{Univ}$ , we have  $(\text{HLR})^{10} \emptyset; \emptyset \Vdash v : \text{Univ}$ . By Definition 2.4 (Reynolds-style Logical Relation) with HLR (taking  $\rho = \emptyset$  and  $\gamma = \emptyset$ ), we have  $v \in \mathcal{E} \llbracket \text{Univ} \rrbracket^{\emptyset}$ . Because  $v$  is a value, we have  $v \in \mathcal{V} \llbracket \text{Univ} \rrbracket^{\emptyset}$ .

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

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

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

$$\begin{aligned}
 & t_u v \\
 \equiv & \left| \begin{array}{l} (\lambda x : \text{Univ}. \text{unpack } x \text{ as } \langle Y, x' \rangle \text{ in} \\ \text{let } x'' : (\text{Unit} \rightarrow Y) \times (Y \rightarrow \text{Unit}) = x' \text{ Unit in} \\ x''.2 (x''.1 \text{ unit}) v \end{array} \right. \\
 \hookrightarrow & \text{unpack } v \text{ as } \langle Y, x' \rangle \text{ in let } x'' : (\text{Unit} \rightarrow Y) \times (Y \rightarrow \text{Unit}) = x' \text{ Unit in } x''.2 (x''.1 \text{ unit})
 \end{aligned}$$

<sup>10</sup>For ease of discussion, we use this notation to give names to hypotheses as they become available during the proof and use the name later in the proof when we use the hypothesis.

$$\begin{aligned}
& \equiv \left| \begin{array}{l} \text{unpack (pack } \langle \tau_Y, v' \rangle \text{ as } \underline{\text{Univ}} \text{ as } \langle Y, x' \rangle \text{ in} \\ \text{let } x'' : (\text{Unit} \rightarrow Y) \times (Y \rightarrow \text{Unit}) = x' \text{ Unit in} \\ \qquad \qquad \qquad x'' . 2 (x'' . 1 \text{ unit}) \end{array} \right. \\
& \hookrightarrow \text{let } x'' : (\text{Unit} \rightarrow \tau_Y) \times (\tau_Y \rightarrow \text{Unit}) = v' \text{ Unit in } x'' . 2 (x'' . 1 \text{ unit})
\end{aligned}$$

By Lemma 2.6 (Antireduction) and Lemma 2.9 (Boring lemma), to conclude our thesis ( $\mathbf{t}_u \mathbf{v} \in \mathcal{E} \llbracket \mathbf{X} \rrbracket^{0, X \mapsto (\text{Unit}, \text{RE})}$ ) it is sufficient to show that:

$$\text{let } x'' : (\text{Unit} \rightarrow \tau_Y) \times (\tau_Y \rightarrow \text{Unit}) = v' \text{ Unit in } x'' . 2 (x'' . 1 \text{ unit}) \in \mathcal{E} \llbracket \mathbf{X} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RE})}$$

Now assume that the term

$$\text{let } x'' : (\text{Unit} \rightarrow \tau_Y) \times (\tau_Y \rightarrow \text{Unit}) = v' \text{ Unit in } x'' . 2 (x'' . 1 \text{ unit})$$

terminates to a value  $v_r$ . By definition of the term relation, it suffices to prove that

$$v_r \in \mathcal{V} \llbracket \mathbf{X} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RE})}$$

It is easy to see that there must be an intermediate value  $v_U$  such that  $v' \text{ Unit} \hookrightarrow^* v_U$  and

$$\text{let } x'' : (\text{Unit} \rightarrow \tau_Y) \times (\tau_Y \rightarrow \text{Unit}) = v_U \text{ in } x'' . 2 (x'' . 1 \text{ unit}) \hookrightarrow^* v_r$$

From HPVP, we know that

$$v' \text{ Unit} \in \mathcal{E} \llbracket \mathbf{X} \rightarrow \mathbf{Y} \times \mathbf{Y} \rightarrow \mathbf{X} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RE})}$$

By definition of the term relation, we have that (HPVU)  $v_U \in \mathcal{V} \llbracket \mathbf{X} \rightarrow \mathbf{Y} \times \mathbf{Y} \rightarrow \mathbf{X} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RE})}$ . It will be useful later that the same property holds if we replace **RE** with **RU** (HPVU2).

By definition of the term relation, it suffices to prove that

$$v_U . 2 (v_U . 1 \text{ unit}) \in \mathcal{E} \llbracket \mathbf{X} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RE})}$$

This is a top-level application, so by Lemma 2.8 (Compatibility for applications), it suffices to prove the following (recall that **RE** =  $\emptyset$ ):

- (1)  $v_U . 2 \in \mathcal{E} \llbracket \mathbf{Y} \rightarrow \mathbf{X} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RE})}$
- (2)  $v_U . 1 \text{ unit} \in \mathcal{E} \llbracket \mathbf{Y} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RE})}$

The former follows easily from HPVU by unfolding the definition of the value relation for 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 **X** in the term relation for **Y** (since variable **X** is not mentioned in type **Y**), then add the second relation and all the term relations are equivalent:

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

So instead of proving:

$$v_U . 1 \text{ unit} \in \mathcal{E} \llbracket \mathbf{Y} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RE})}$$

we can just prove (notice the change in the semantic interpretation for **X**):

$$v_U . 1 \text{ unit} \in \mathcal{E} \llbracket \mathbf{Y} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RU})}$$

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:

- (3)  $v_U . 1 \in \mathcal{E} \llbracket \mathbf{X} \rightarrow \mathbf{Y} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RU})}$
- (4)  $\text{unit} \in \mathcal{E} \llbracket \mathbf{X} \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, \text{RU})}$

Similarly to how we derived HPVU from HPVP and the definition of the term relation, Item 3 follows from the fact that  $v_U \in \mathcal{V} \llbracket X \rightarrow Y \times Y \rightarrow X \rrbracket^{\emptyset, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, RU)}$  (HPVU2) and then unfolding the definition of the value relation for pairs.

Proving Item 4 is not hard either. The term relation includes the value relation, so it suffices to prove:

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

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

If we had not performed our main trick, all other proof obligations would hold, but proving Item 4 would not be possible. In fact, there we would have had to prove that:

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

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

## 4.2 Why Universal Types Require Type Worlds

The proof from Section 3.1 shows that Reynolds-style or lexically-scoped logical relations are incompatible with a universal type. But to thoroughly understand what is going on, it is useful to take a closer look at the counterexample.

Recall the definition of the type Univ.

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

According to the RLR, any value of this type is of the form **pack**  $\langle \tau_Y, v \rangle$  **as** Univ and comes with a predicate  $R_Y \in \text{Re1}(\tau_Y)$ , i.e., a predicate on values of type  $\tau_Y$ . Importantly, this predicate needs to be chosen independently of choices that are made later, particularly the choice of the type Unit for **X** and the predicate  $R_X \in \text{Re1}(\text{Unit})$  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 were a universal type **U** in System F with total injection and extraction functions  $\text{in}_Z : Z \rightarrow \text{U}$  and  $\text{out}_Z : \text{U} \rightarrow Z$  for arbitrary types **Z**. Then we could define a value of type Univ by taking  $Y = \text{U}$  and constructing a pair with  $\text{in}_X$  and  $\text{out}_X$ :

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

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  $\text{in}_X$  **unit** when **X** is later instantiated to Unit (as in the counterexample). However, whether or not this value can be legally embedded in **U** fundamentally depends on the choice for  $R_X \in \text{Re1}(\text{Unit})$ . Particularly, if  $R_X = \emptyset$ , then it should be rejected, but if  $R_X = \{\text{unit}\}$  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 2.5, 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:

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

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

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  $\mathbf{W}$ . This world  $\mathbf{W}$  will contain choices of predicates for other type variables, particularly the choice for  $\mathbf{R}_X$ . In each of the TWLRs where a universal type is at play, the logical relation will accept different values depending on the world  $\mathbf{W}$ , thus solving the conundrum outlined above.

It is also interesting to consider what this means in practice, when proving theorems using a formulation of 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 4.2 ( $\mathbf{Univ}$  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 2.5 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 4.2, we see that we were able to prove the following:

$$\mathbf{v}_{U.1} \text{ unit} \in \mathcal{E} \llbracket \mathbf{Y} \rrbracket^{\theta, Y \mapsto (\tau_Y, R_Y), X \mapsto (\mathbf{Unit}, \mathbf{RU})}$$

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

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

We could then conclude that

$$\mathbf{v}_{U.1} \text{ unit} \in \mathcal{E} \llbracket \mathbf{Y} \rrbracket^{\theta, Y \mapsto (\tau_Y, R_Y), X \mapsto (\mathbf{Unit}, \mathbf{RE})}$$

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

- $\mathbf{W}_{yx1} = ((Y, \tau_Y, R_Y); (X, \mathbf{Unit}, \mathbf{RU}))$
- $\mathbf{W}_{yx2} = ((Y, \tau_Y, R_Y); (X, \mathbf{Unit}, \mathbf{RE}))$

Clearly,  $\mathbf{W}_{yx2} \not\sqsupseteq \mathbf{W}_{yx1}$  (HPNF).

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

$$(\mathbf{W}_{yx1}, \mathbf{v}_U).1 \text{ unit} \in \mathcal{E} \llbracket \mathbf{Y} \rrbracket$$

but the step to  $\mathbf{W}_{yx2}$  no longer holds:

$$(\mathbf{W}_{yx2}, \mathbf{v}_U).1 \text{ unit} \notin \mathcal{E} \llbracket \mathbf{Y} \rrbracket$$

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  $\mathbf{Y}$  in world  $\mathbf{W}_{yx1}$  than in  $\mathbf{W}_{yx2}$ . This dependence of the relation for  $\mathbf{Y}$  on the world is precisely what a TWLR purposefully allows. As a result, it is impossible to transfer the result about  $\mathbf{v}'$   $\llbracket \mathbf{Unit} \rrbracket$  from the world  $\mathbf{W}_{yx1}$  (with the permissive predicate  $\mathbf{RU}$  for  $\mathbf{X}$ ) to the world  $\mathbf{W}_{yx2}$  (with the more restrictive predicate  $\mathbf{RE}$  for  $\mathbf{X}$ ).

### 4.3 Degeneracy and Other Variants of Logical Relations

If we consider the other variants of the logical relation discussed in Section 2.4, our expectations are essentially confirmed. We have seen in Section 3.2 that the degeneracy of  $\mathbf{Univ}$  does not hold in the presence of higher-order effects and indeed, the above proof fails if we use the Kripke logical relation from Section 2.4, for a similar reason as for the TWLR.

For some  $\mathbf{W}$ , we would be able to obtain the following two facts:

$$(\mathbf{W}, v' \text{Unit}) \in \mathcal{E} \llbracket X \rightarrow Y \times Y \rightarrow X \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, RE)}$$

$$(\mathbf{W}, v' \text{Unit}) \in \mathcal{E} \llbracket X \rightarrow Y \times Y \rightarrow X \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, RU)}$$

Combined with  $v' \text{Unit} \hookrightarrow^* v_U$ , this gives us some  $\mathbf{W}'$ , for which

$$(\mathbf{W}', v_U) \in \mathcal{V} \llbracket X \rightarrow Y \times Y \rightarrow X \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, RE)}$$

From the second expression relation above, we also get a  $\mathbf{W}''$  such that

$$(\mathbf{W}'', v_U) \in \mathcal{V} \llbracket X \rightarrow Y \times Y \rightarrow X \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, RU)}$$

However, these two worlds  $\mathbf{W}'$  and  $\mathbf{W}''$  are potentially distinct and one is not (necessarily) a future world of the other, which will later prevent us from combining both facts later in the proof.

In Section 3.2, we have also seen that the value-polymorphic version of the same calculus with higher-order state should preserve degeneracy of  $\text{Univ}$ . And indeed, using the value-polymorphic KLR, we can proceed differently and deduce that  $v' = \Lambda X. v''$  and

$$(\mathbf{W}, v''[\text{Unit}/X]) \in \mathcal{V} \llbracket X \rightarrow Y \times Y \rightarrow X \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, RE)}$$

$$(\mathbf{W}, v''[\text{Unit}/X]) \in \mathcal{V} \llbracket X \rightarrow Y \times Y \rightarrow X \rrbracket^{0, Y \mapsto (\tau_Y, R_Y), X \mapsto (\text{Unit}, RU)}$$

In other words, value polymorphism gives us these two facts immediately in the same world and as a result, contrary to above, we can combine them in the remainder of the proof.

Thus, the different logical relation variants that we have discussed in Section 2 either prevent or allow the proof of degeneracy of  $\text{Univ}$ , as appropriate for the language variant they are intended for. In addition to confirming our analysis of the universal type in Section 3, this also provides more insight into how the different logical relation variants influence the particular proof techniques used in the proof outlined here.

In addition to clarifying the relation between TWLRs and non-lexically-scoped type variables, 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].

## 5 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 5.1) and the compiler from  $\lambda^F$  to  $\lambda^\sigma$  (Section 5.2). 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 5.3).

*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>11</sup> However, the typed version suffers from the same problem, as we show in detail

<sup>11</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.

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.

### 5.1 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].

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):

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

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.

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  `$\{t_1\}_{t_2}$`  first evaluates  `$t_1$`  and  `$t_2$`  to values  `$v_1$`  and  `$v_2$`  and then creates the sealed value  `$\{v_1\}_{v_2}$`  (or leads to a run-time error if  `$v_2$`  is not a seal). One can think of such a sealed value as  `$v_1$`  encrypted under  `$v_2$` . The final construct `let  $\{x\}_{t_1} = t_2$  in  $t_3$  else  $t_4$`  is for unsealing or decrypting. It first evaluates  `$t_1$`  to a seal  `$\sigma_1$`  and  `$t_2$`  to  `$\{v_2\}_{\sigma_2}$`  (or produces a run-time error if either result is not of that form). If  `$\sigma_1$`  and  `$\sigma_2$`  are equal,  `$t_3$`  is evaluated with  `$x$`  bound to the decrypted value  `$v_2$` , otherwise  `$t_4$`  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

1128 implementation of  $\mathbb{Z}_3$  from Example 2.2 in terms of pairs of booleans, uses sealing to protect the  
 1129 abstraction.

1130 *Example 5.1* ( $\mathbb{Z}_3$  in  $\lambda^\sigma$ ). First, we introduce two helper functions:  
 1131

1132  $\text{seal}_\sigma \stackrel{\text{def}}{=} \lambda y. \{y\}_\sigma$                        $\text{unseal}_\sigma \stackrel{\text{def}}{=} \lambda y. \text{let } \{x\}_\sigma = y \text{ in } x \text{ else wrong}$   
 1133

1134 Then, we can define a  $\lambda^\sigma$  correspondent of **z3** as:

1135  $\text{z3} = \text{vs.} \langle \langle \text{zero}, \text{succ} \rangle, \text{zero?} \rangle$

1136 where  $\left\{ \begin{array}{l} \text{zero} = \text{seal}_s \langle \text{false}, \text{false} \rangle \\ \text{succ} = \lambda p. \text{if } (\text{unseal}_s p).2 \text{ then } \text{seal}_s \langle \text{false}, \text{false} \rangle \text{ else} \\ \quad \text{if } (\text{unseal}_s p).1 \text{ then } \text{seal}_s \langle \text{false}, \text{true} \rangle \text{ else } \text{seal}_s \langle \text{true}, \text{false} \rangle \\ \text{zero?} = \lambda p. \text{if } (\text{unseal}_s p).1 \text{ then false else if } (\text{unseal}_s p).2 \text{ then false else true} \end{array} \right.$   
 1137  
 1138  
 1139  
 1140  
 1141

1142 Whenever values of the abstract type leave the scope of the abstract type definition, they are  
 1143 encrypted, and when they are passed back in they are decrypted before use. Intuitively, one can see  
 1144 that this protects against a context looking into or tampering with representation values, similarly  
 1145 to the protection offered by type checking of parametric polymorphism.

1146 Also the protection required for the dual case, where a term calls a universally quantified function  
 1147 provided by the context, can be implemented in  $\lambda^\sigma$ . Consider for instance the  $\lambda^F$  term:

1148  $\lambda f : \forall X. X \times X \rightarrow X. \text{if } f \text{ Bool } \langle \text{true}, \text{true} \rangle \text{ then true else false}$   
 1149

1150 The polymorphic function **f** that will be passed in by the context, can only return one of its  
 1151 arguments (or diverge), and hence the invocation in the term above will necessarily return **true** (or  
 1152 diverge).

1153 We can implement a similar protection in  $\lambda^\sigma$ . To enforce parametric behaviour of a polymorphic  
 1154 function received from the context, we create a new seal for every invocation, and we encrypt all  
 1155 parameters of the quantified type with that seal. For instance, the term above could be implemented  
 1156 in  $\lambda^\sigma$  as follows:

1157  $\lambda f. \text{if } \text{vs.} \text{unseal}_s (f \langle \text{seal}_s \text{ true}, \text{seal}_s \text{ true} \rangle) \text{ then true else false}$   
 1158

1159 Before calling **f**, its arguments are encrypted with a new seal, and the return value gets decrypted  
 1160 with that seal, hence all that **f** can do is either return one of its arguments or diverge (doing anything  
 1161 else would lead to a run-time error).

1162 Again, this gives us a *dynamic* guarantee that a function has to treat its arguments opaquely,  
 1163 where parametric polymorphism gives us this guarantee *statically* by type checking. This technique  
 1164 of dynamically protecting values with appropriately scoped seals is the essence of the idea behind  
 1165 Sumii and Pierce's compiler for  $\lambda^F$ .

## 1166 5.2 Sumii and Pierce's Compiler

1167 The compiler  $\llbracket \cdot \rrbracket_{\lambda^\sigma}^{\lambda^F}$  first performs standard type erasure (function erase  $(\cdot)$ ) and then wraps it with  
 1168 a dynamic check ( $\text{protect}_{\eta; \tau}$ ) that will insert dynamic applications of sealing and unsealing:

1169 if  $\emptyset; \emptyset \vdash t : \tau$ , then  $\llbracket t \rrbracket_{\lambda^\sigma}^{\lambda^F} \stackrel{\text{def}}{=} \text{let } x = \text{erase } (t) \text{ in } \text{protect}_{\emptyset; \tau} x$   
 1170  
 1171

1172 Note that we restrict the compiler to closed terms here. Lifting this limitation is quite straightforward,  
 1173 as presented by Devriese et al. [2017].

1174 We now present these steps and discuss their meaning. These definitions are taken from Sumii  
 1175 and Pierce [2004], except that we make some notational changes and some minor technical changes.  
 1176



*Type erasure.* This pass is mostly straightforward, the only non-trivial aspect is how to deal with the quantified types. Type abstraction and application are erased to a dummy lambda abstraction and an application to a **unit** parameter, and unpacking is erased to a **let**-binding.<sup>12</sup>

$$\begin{aligned}
 \text{erase } (x) &\stackrel{\text{def}}{=} x & \text{erase } (t \ \tau') &\stackrel{\text{def}}{=} \text{erase } (t) \ \text{unit} \\
 \text{erase } (\lambda X. t) &\stackrel{\text{def}}{=} \lambda \_ . \text{erase } (t) & \text{erase } (\text{pack } \langle \tau', t \rangle \text{ as } \exists X. \tau) &\stackrel{\text{def}}{=} \text{erase } (t) \\
 \text{erase } (\lambda x : \tau. t) &\stackrel{\text{def}}{=} \lambda x. \text{erase } (t) & \text{erase } (\text{unpack } t \text{ as } \langle X, x \rangle \text{ in } t') &\stackrel{\text{def}}{=} \text{let } x = \text{erase } (t) \text{ in } \text{erase } (t')
 \end{aligned}$$

*Dynamic wrappers.* The second phase of the compiler wraps compiled terms with dynamic wrappers. These are formalised as the function  $\text{protect}_{\eta;\tau}$ , which is defined in Figure 3, together with its dual  $\text{confine}_{\eta;\tau}$  by mutual induction on  $\tau$ . We use the names **protect** and **confine** (following Devriese et al. [2016]) to refer to the wrappers that Sumii and Pierce call  $\mathcal{E}^+$  and  $\mathcal{E}^-$  respectively [Sumii and Pierce 2004].

Intuitively, applying  $\text{protect}_{\eta;\tau}$  to a value  $v$  ensures that  $v$  cannot be used in ways that are not allowed by type  $\tau$ . Dually, applying  $\text{confine}_{\eta;\tau}$  to a value  $v$  prevents  $v$  from behaving in a way that is not allowed by type  $\tau$ . For any free type variables  $X$  in  $\tau$ ,  $\eta$  tells us how to protect/confine values of type  $X$ . Concretely,  $\eta(X) = (t_p, t_c)$  where  $t_p$  and  $t_c$  are untyped terms that should be applied to protect/confine (respectively) values of type  $X$ . Formally,  $\eta$  has the following syntax:  $\eta ::= \emptyset \mid \eta, X \mapsto (t, t)$ .

For defining **protect**, and **confine**, we rely on the following  $\text{seal}_\sigma$  and  $\text{unseal}_\sigma$  functions, which seal and unseal values with a given seal  $\sigma$ .

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

Protecting at a function type confines the argument and protects the result with the same polarity, and similarly for confining at a function type. Confining at ground type inserts a dynamic check 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 5.1.

Applying the compiler to **z3** from Example 5.1 results in the  $\lambda^\sigma$  term **z3** from Example 2.2 (modulo some additional  $\beta$ -reductions).

### 5.3 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 5.2 (SUMII AND PIERCE).  $\emptyset; \emptyset \vdash t_1 \simeq t_2 : \tau$  if and only if  $\emptyset \vdash \llbracket t_1 \rrbracket_{\lambda^\sigma}^{\lambda^F} \simeq \llbracket t_2 \rrbracket_{\lambda^\sigma}^{\lambda^F}$

<sup>12</sup>We are assuming a standard desugaring of let expressions to function applications.

1226  $\text{protect}_{\eta;\text{Unit}} x \stackrel{\text{def}}{=} x$   
 1227  $\text{protect}_{\eta;\text{Bool}} x \stackrel{\text{def}}{=} x$   
 1228  $\text{protect}_{\eta;\tau_1 \times \tau_2} x \stackrel{\text{def}}{=} \text{let } x_1 = x.1 \text{ in let } x_2 = x.2 \text{ in } \langle \text{protect}_{\eta;\tau_1} x_1, \text{protect}_{\eta;\tau_2} x_2 \rangle$   
 1229  $\text{protect}_{\eta;\tau_1 \rightarrow \tau_2} x \stackrel{\text{def}}{=} \lambda y. \text{let } z = x \text{ (confine}_{\eta;\tau_1} y \text{) in protect}_{\eta;\tau_2} z$   
 1230  $\text{protect}_{\eta;\forall X.\tau} x \stackrel{\text{def}}{=} \lambda \_ . \text{let } y = x \text{ unit in protect}_{\eta, X \mapsto (\lambda x.x, \lambda x.x); \tau} y$   
 1231  $\text{protect}_{\eta;\exists X.\tau} x \stackrel{\text{def}}{=} \text{vs. } \text{protect}_{\eta, X \mapsto (\text{seals}, \text{unseals}); \tau} x$   
 1232  $\text{protect}_{\eta;X} x \stackrel{\text{def}}{=} t_p x \quad \text{where } \eta(X) = (t_p, \_)$   
 1233  $\text{confine}_{\eta;\text{Unit}} x \stackrel{\text{def}}{=} x; \text{unit}$   
 1234  $\text{confine}_{\eta;\text{Bool}} x \stackrel{\text{def}}{=} \text{if } x \text{ then true else false}$   
 1235  $\text{confine}_{\eta;\tau_1 \times \tau_2} x \stackrel{\text{def}}{=} \text{let } x_1 = x.1 \text{ in let } x_2 = x.2 \text{ in } \langle \text{confine}_{\eta;\tau_1} x_1, \text{confine}_{\eta;\tau_2} x_2 \rangle$   
 1236  $\text{confine}_{\eta;\tau_1 \rightarrow \tau_2} x \stackrel{\text{def}}{=} \lambda y. \text{let } z = x \text{ (protect}_{\eta;\tau_1} y \text{) in confine}_{\eta;\tau_2} z$   
 1237  $\text{confine}_{\eta;\forall X.\tau} x \stackrel{\text{def}}{=} \lambda \_ . \text{vs. let } x' = x \text{ unit in confine}_{\eta, X \mapsto (\text{seals}, \text{unseals}); \tau} x'$   
 1238  $\text{confine}_{\eta;\exists X.\tau} x \stackrel{\text{def}}{=} \text{confine}_{\eta, X \mapsto (\lambda x.x, \lambda x.x); \tau} x$   
 1239  $\text{confine}_{\eta;X} x \stackrel{\text{def}}{=} t_c x \quad \text{where } \eta(X) = (\_, t_c)$

Fig. 3. The dynamic wrappers of Sumii and Pierce's compiler.

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 5.3** ( $t_u$  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 supplementary material.

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.

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

$$C \left[ \llbracket t_u \rrbracket_{\lambda^\sigma}^{\lambda^F} \right] \hookrightarrow^* \text{unit} \quad C \left[ \llbracket t_\omega \rrbracket_{\lambda^\sigma}^{\lambda^F} \right] \hookrightarrow^* (\lambda r. r) \omega \hookrightarrow^* (\lambda r. r) \omega \hookrightarrow^* \dots \uparrow$$

We spell out the reductions in full detail in the supplementary material. 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.

1275 THEOREM 5.4 ( $\llbracket \cdot \rrbracket_{\lambda^F}^{\sigma}$  IS NOT FULLY ABSTRACT). *It is not true that*

$$1276 \quad \forall \mathbf{t}_1, \mathbf{t}_2. \mathbf{t}_1 \simeq \mathbf{t}_2 \iff \llbracket \mathbf{t}_1 \rrbracket_{\lambda^F}^{\sigma} \simeq \llbracket \mathbf{t}_2 \rrbracket_{\lambda^F}^{\sigma}$$

1277  
1278 PROOF. Follows easily from Theorem 3.1 and the counterexample in Theorem 5.3.  $\square$

1280 *The problem is in the easy case.* It is worth noticing that most of the work in Sumii and Pierce's  
1281 compiler (see Fig. 3) is in enforcing that existentially quantified types passed to the context are  
1282 treated opaquely and, dually, that polymorphic functions received from the context are forced  
1283 to treat their argument type opaquely. However, it is not in these cases that the counterexample  
1284 highlights a problem.

1285 Instead, it goes wrong in the cases where we receive an existential type from the context (and  
1286 dually when we pass a polymorphic function to the context). For these seemingly simple cases,  
1287 the dynamic wrappers from Fig. 3 do not perform any specific kind of enforcement (except for  
1288 recursing on their body type). However, it is there that our counterexample uncovers a problem:  
1289 the value that it provides as a value of  $\mathbf{Univ} = \exists \mathbf{Y}. \forall \mathbf{X}. (\mathbf{X} \rightarrow \mathbf{Y}) \times (\mathbf{Y} \rightarrow \mathbf{X})$  does not correspond to any  
1290 legal choice of  $\mathbf{Y}$ , but the dynamic type wrappers have no way to detect this. In fact, an alternative  
1291 way to understand what goes wrong is that the value  $\lambda_. \langle \lambda x. x, \lambda x. x \rangle$  provided by the context,  
1292 behaves as if it can choose  $\mathbf{Y}$  equal to  $\mathbf{X}$ . However, this is not possible in  $\lambda^F$  because  $\mathbf{X}$  is not in  
1293 scope at the moment when  $\mathbf{Y}$  needs to be chosen.

1294 In other words, what seems to be missing in Sumii and Pierce's dynamic enforcement is an  
1295 enforcement of the type variable scope of existentially quantified types. To see this, consider the  
1296 following type  $\mathbf{Triv}$ , which is identical to  $\mathbf{Univ}$ , except for the order of the quantifiers:

$$1297 \quad \mathbf{Univ} \stackrel{\text{def}}{=} \exists \mathbf{Y}. \forall \mathbf{X}. (\mathbf{X} \rightarrow \mathbf{Y}) \times (\mathbf{Y} \rightarrow \mathbf{X}) \quad \mathbf{Triv} \stackrel{\text{def}}{=} \forall \mathbf{X}. \exists \mathbf{Y}. (\mathbf{X} \rightarrow \mathbf{Y}) \times (\mathbf{Y} \rightarrow \mathbf{X})$$

1299 Even though the type  $\mathbf{Triv}$  is more liberal, as it allows to instantiate  $\mathbf{Y}$  with  $\mathbf{X}$ , the wrappers of  
1300 Fig. 3 treat the types essentially the same. From this perspective, it is no surprise that the value  
1301  $\lambda_. \langle \lambda x. x, \lambda x. x \rangle$  is accepted as a value of type  $\mathbf{Univ}$ , because it does in fact correspond to a legal  $\lambda^F$   
1302 value of type  $\mathbf{Triv}$ :

$$1303 \quad \Delta \mathbf{X}. \text{pack } \langle \mathbf{X}, \langle \lambda x : \mathbf{X}. x, \lambda x : \mathbf{X}. x \rangle \rangle \text{ as } \exists \mathbf{Y}. (\mathbf{X} \rightarrow \mathbf{Y}) \times (\mathbf{Y} \rightarrow \mathbf{X})$$

## 1304 6 ENFORCING PARAMETRICITY IN THE PRESENCE OF TYPE CASTS

1307 The second conjecture we disprove is by Neis et al. [2009, 2011]. These authors study a form  
1308 of runtime type generation to protect parametrically polymorphic functions when interacting  
1309 with code that can use a type cast primitive. They prove a parametricity result that applies to  
1310 appropriately wrapped System F values once embedded in  $\mathbf{G}$ , a language with a type cast primitive.  
1311 They conjecture [Neis et al. 2009, section 10] that this wrapping is fully abstract, and while they prove  
1312 equivalence reflection, they only conjecture equivalence preservation due to a lack of sophistication  
1313 of the proof techniques available at the time.<sup>13</sup>

1314 It turns out that our results may be adapted to this setting, as in such a non-parametrically  
1315 polymorphic setting, the type  $\forall \mathbf{X}. \mathbf{X}$  can be used as a universal type (that every other type can be  
1316 embedded into or out from). In other words, we disprove this conjecture too: embedding System F  
1317 into  $\mathbf{G}$  à la Neis-Dreyer-Rossberg (NDR, in the sequel) is not fully abstract.

1318 We first present  $\mathbf{G}$  (Section 6.1) and the wrapper for protecting polymorphic functions from  
1319 interacting with  $\mathbf{G}$  terms (Section 6.2). Then we demonstrate how  $\forall \mathbf{Z}. \mathbf{Z}$  is a universal type in  $\mathbf{G}$

1320 <sup>13</sup> A claim that was indeed true, since much of the research in proof techniques for fully abstract compilation postdates that  
1321 paper [Abate et al. 2019; Ahmed and Blume 2011; Devriese et al. 2016; Fournet et al. 2013; New et al. 2016b; Patrignani and  
1322 Garg 2019; Schmidt-Schauß et al. 2015]. An exception is the work by Ahmed and Blume [2008] on typed closure conversion.

1324 (Section 6.3) and prove that due to that type, the wrapper does not preserve contextual equivalence  
 1325 (Section 6.4).

1326

1327

## 6.1 The $\mathbb{G}$ Language

1328  $\mathbb{G}$  extends System F with two primitives (Figure 4). The first one casts values of type  $\tau_1$  to values  
 1329 of type  $\tau_2$ . The second one generates a fresh type name  $X$  that is registered to be an equivalent  
 1330 classifier to a type  $\tau$  although data of the two types is not equivalent (so casting between  $X$  and  $\tau$   
 1331 will not succeed).

1332

1333

1334

Syntax:

1335

$$t ::= \dots \mid \text{cast } \tau_1 \tau_2 \mid \text{new } X \approx \tau \text{ in } t$$

1336

$$\Delta ::= \emptyset \mid \Delta, X \mid \Delta, X \approx \tau$$

1337

$$\Sigma ::= \emptyset \mid \Sigma, X \approx \tau$$

1338

1339

Typing rules (excerpts):

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

1350

1351

Evaluation rules (excerpts):

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

$$\frac{\Delta \vdash \tau_1 \quad \Delta \vdash \tau_2}{\Delta; \Gamma \vdash \text{cast } \tau_1 \tau_2 : \tau_1 \rightarrow \tau_2 \rightarrow \tau_2}$$

FRANK PIESSENS, KU Leuven, Belgium

$$\frac{\Delta, X \approx \tau; \Gamma \vdash t : \tau'}{\Delta; \Gamma \vdash \text{new } X \approx \tau \text{ in } t : \tau'}$$

$$\frac{\Delta; \Gamma \vdash t : \tau \quad \Delta \vdash \tau \approx \tau'}{\Delta; \Gamma \vdash t : \tau'}$$

$$\frac{\tau_1 = \tau_2}{\Sigma \triangleright \text{cast } \tau_1 \tau_2 \leftrightarrow_0 \Sigma \triangleright \lambda x_1 : \tau_1. \lambda x_2 : \tau_2. x_1}$$

FRANK PIESSENS, KU Leuven, Belgium

$$\frac{\tau_1 \neq \tau_2}{\Sigma \triangleright \text{cast } \tau_1 \tau_2 \leftrightarrow_0 \Sigma \triangleright \lambda x_1 : \tau_1. \lambda x_2 : \tau_2. x_2}$$

$$\Sigma \triangleright \text{new } X \approx \tau \text{ in } t \leftrightarrow_0 \Sigma, (X \approx \tau) \triangleright t$$

Fig. 4. The  $\mathbb{G}$  language: syntax, typing rules and evaluation rules (excerpts). The semantics relation  $\leftrightarrow$  relies on the primitive reductions indicated as  $\leftrightarrow_0$  and it relates configurations of the form  $\Sigma \triangleright t$

1364

1365

1366

1367

1368

1369

1370

1371

1372

## 6.2 Enforcing Parametricity in $\mathbb{G}$

To avoid much code repetition, we use the same notation for the wrapper as that used by Neis et al. [2009]. 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 (i.e., from  $\pm$  to  $\mp$ ).

$$W_{\tau}^{\pm}(t) = \text{let } x = t \text{ in Wrap}_{\tau}^{\pm}(x)$$

1373  $\text{Wrap}_X^\pm(v) = v$   
 1374  $\text{Wrap}_{\text{Bool}}^\pm(v) = v$   
 1375  $\text{Wrap}_{\tau \rightarrow \tau'}^\pm(v) = \lambda x : \tau. W_{\tau'}^\pm((v \text{ Wrap}_\tau^\mp(x)))$   
 1376  $\text{Wrap}_{\tau \times \tau'}^\pm(v) = \langle W_\tau^\pm(v.1), W_{\tau'}^\pm(v.2) \rangle$   
 1377  $\text{Wrap}_{\forall X. \tau}^\pm(v) = \langle W_\tau^\pm(v.1), W_{\tau'}^\pm(v.2) \rangle$   
 1378  $\text{Wrap}_{\exists X. \tau}^\pm(v) = \Delta X. \text{NewTy}^\mp X \text{ in } W_\tau^\pm((v X))$   
 1379  $\text{NewTy}^+ X \text{ in } t = \text{new } Y \approx X \text{ in } t[Y/X]$   
 1380  $\text{NewTy}^- X \text{ in } t = t$

1384 The wrapper consists of three parts. The first one,  $W_\tau^\pm(t)$ , is the term wrapper, it reduces a term  $t$   
 1385 of source type  $\tau$  to a value and applies the value wrapper. The second one,  $\text{Wrap}_\tau^\pm(v)$ , is the value  
 1386 wrapper, it is responsible for generating the fresh type variables for outgoing universal values and  
 1387 incoming existential packages. The third one,  $\text{NewTy}^\pm X \text{ in } t$  is the code that effectively generates  
 1388 the fresh type variables, depending on the polarity of the invocation. As the authors themselves  
 1389 note, the wrapper functions in a way analogous to the dynamic checks inserted by the Sumii-Pierce  
 1390 compiler.

1391 The compiler from  $\lambda^F$  to  $\mathbb{G}$ , denoted with  $\wr$ , leaves the term untouched (which we indicate with  
 1392  $\equiv$ ) and wraps it with a wrapper of the appropriate type.

$$1393 \quad \wr t \stackrel{\text{def}}{=} W_\tau^+(t) \quad \text{if } \emptyset; \emptyset \vdash t : \tau \text{ and } t \equiv t$$

1395 The NDR conjecture (Conjecture 6.1) states that this wrapper is fully abstract. Contextual  
 1396 equivalence in  $\mathbb{G}$ , indicated with  $\simeq$ , is defined analogously to Definition 2.1.

1397 CONJECTURE 6.1 (NDR CONJECTURE).  $\forall t_1, t_2. t_1 \simeq t_2 \iff \wr t_1 \simeq \wr t_2$

1399 We will prove that this statement is not true (Theorem 6.3).

### 1401 6.3 $\mathbb{G}$ has an Instance of an Universal Type: $\forall Z. Z$

1402 To explain how we disprove the NDR conjecture, we need to explain that  $\mathbb{G}$  indeed has a universal  
 1403 type (unlike System F). This type is  $\forall Z. Z$  and it functions as a universal type because we can take  
 1404 any other type  $X$  and (i) embed a value  $v$  of type  $X$  into  $\forall Z. Z$ , (ii) extract the same value  $v$  from  $\forall Z. Z$   
 1405 at type  $X$ . It is possible to do this while remaining parametric in the type  $X$  that these functions  
 1406 work with, i.e., keeping type variable  $X$  free in these terms and binding it in a larger term using  
 1407 both these functions.

1408 These two functionalities will be key for building a distinguishing context for  $\mathbb{G}$  which is analogous  
 1409 to the distinguishing contexts of for  $\lambda^\sigma$  and  $\lambda^B$  (the latter will be presented later). A value  $v$  of an  
 1410 arbitrary type  $\tau$  is represented as a value of type  $\forall Z. Z$ . The cast operator in  $\mathbb{G}$  allows us to make  
 1411 this function behave differently when it is applied to type  $\tau$  than when it is applied to other types.  
 1412 In the former case, we will make the function return  $v$  itself, while in the latter case, we simply  
 1413 make the function diverge.

1414 Concretely, to extract a value from  $\forall Z. Z$  into  $X$ , we simply apply the polymorphic function to  
 1415 type  $X$ :

$$1416 \quad \lambda z : (\forall Z. Z). z X$$

1417 Injecting a value of type from  $X$  into  $\forall Z. Z$  is a bit more complex. What we would like to write is  
 1418 the following:

$$1420 \quad \lambda x : X. \Lambda Z. \text{cast } X Z \times \omega_Z$$

1421



perform a series of unpacking and type renaming which are intended to preserve parametricity.

$$\lambda x_1 : \underline{\text{Univ}}. \left( \begin{array}{l} (\lambda x : \underline{\text{Univ}}. \text{unpack } x \text{ as } \langle Y, x' \rangle \text{ in let } x'' = x' \text{ Unit in } x''.2 (x''.1 \text{ unit})) \\ \text{unpack } x_1 \text{ as } \langle X_1, x_4 \rangle \text{ in} \\ \text{pack } \left\langle X_1, \Lambda X_2. \text{new } X_3 \approx X_2 \text{ in } \left\langle \begin{array}{l} \lambda x_5 : X_3.(x_4 X_3).1 x_5, \\ \lambda x_6 : X_1.(x_4 X_3).2 x_6 \end{array} \right\rangle \right\rangle \text{ as } \underline{\text{Univ}} \end{array} \right)$$

The distinguishing context for  $\mathbb{G}$  passes a parameter of type  $\underline{\text{Univ}}$  to the term in the hole (which will be either  $\{t_u\}$  or  $\{t_\omega\}$ ). The goal of this parameter is to make  $\{t_u\}$  terminate and  $\{t_\omega\}$  diverge. Let us discuss the passed parameter. Since  $\underline{\text{Univ}}$  is an existential type at the top level, the parameter is a  $\text{pack } \langle \cdot \rangle \text{ as } \cdot$ . 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  $\Lambda Z. \cdot$ . Then comes the pair of functions for projecting into and extracting from the universal type  $\forall Z. Z$ , as explained in Section 6.3. As a convention, we name variables and type variables from the context with  $z$ :

$$\mathbb{C}^{\mathbb{G}} \stackrel{\text{def}}{=} [\cdot] \left( \text{pack } \left\langle \forall Z. Z, \Lambda Z1. \left\langle \begin{array}{l} \lambda z1 : Z1. \Lambda Z2. \left( \text{cast } (\text{Unit} \rightarrow Z1) (\text{Unit} \rightarrow Z2) \right) \text{unit}, \\ \lambda z2 : (\forall Z. Z). z2 Z1 \end{array} \right\rangle \right\rangle \text{ as } \underline{\text{Univ}} \right)$$

THEOREM 6.2 ( $\{t_u\}$  AND  $\{t_\omega\}$  ARE NOT EQUIVALENT IN  $\mathbb{G}$ ).  $\{t_u\} \neq \{t_\omega\}$

PROOF. We have that  $\mathbb{C}^{\mathbb{G}} [\{t_u\}] \hookrightarrow^* \text{unit}$  while  $\mathbb{C}^{\mathbb{G}} [\{t_\omega\}] \uparrow$ . □

THEOREM 6.3 (EMBEDDING  $\lambda^{\text{F}}$  INTO  $\mathbb{G}$  IS NOT FULLY ABSTRACT). *It is not true that*

$$\forall t_1, t_2. t_1 \approx t_2 \iff \{t_1\} \approx \{t_2\}$$

PROOF. Follows directly from Theorem 6.2 and Theorem 3.1. □

*Additional instances of universal types in  $\mathbb{G}$ .*  $\forall Z. Z$  is not the only universal type in  $\mathbb{G}$ . The type  $\exists X. X$  works just as well and might be more intuitive to some readers. The context below is analogous to the one above and can also differentiate between the compilation of  $t_u$  and of  $t_\omega$ .

$$\mathbb{C}^{\mathbb{G}} \stackrel{\text{def}}{=} [\cdot] \left( \text{pack } \langle \exists X. X, \Lambda Z. \langle \text{inj}, \text{ext} \rangle \rangle \text{ as } \underline{\text{Univ}} \right)$$

$$\text{inj} \stackrel{\text{def}}{=} \lambda z : Z. \text{pack } \langle Z, z \rangle \text{ as } \exists X. X$$

$$\text{ext} \stackrel{\text{def}}{=} \lambda z : (\exists X. X). \text{unpack } z \text{ as } \langle U, u \rangle \text{ in } \left( \begin{array}{l} \text{cast } (\text{Unit} \rightarrow U) (\text{Unit} \rightarrow Z) \\ (\lambda_ : \text{Unit}. u) (\lambda_ : \text{Unit}. \omega_Z) \end{array} \right) \text{unit}$$

## 7 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, which include the *gradual guarantee*.

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.

Jacobs et al. [2021] have recently argued the importance of on another benefit of well-typedness that has received less attention so far: 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 the terms' types. Formulations of System F parametricity similarly allows us to deduce properties from programs' types.

Jacobs et al. propose a criterion that requires gradual type systems to preserve the validity of such reasoning in the original typed language. Specifically, the criterion requires that 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  $[t]$ . The criterion then becomes: if  $t_1 \simeq t_2$ , do we have that  $[t_1] \simeq [t_2]$ , or, in other words, is the embedding of the typed language into the gradual language fully abstract? Based on this view, Jacobs et al. refer to their criterion as the Fully Abstract Embedding (FAE) criterion.

In this section, we show that the criterion fails 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 7.1). 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 7.2). This shows that the embedding of  $\lambda^F$  into  $\lambda^B$  is not fully abstract.<sup>14</sup>

## 7.1 The $\lambda^B$ Calculus

There exist several versions of polymorphic gradual languages in the literature [Ahmed et al. 2011b, 2017; Igarashi et al. 2017; New et al. 2019a; 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.

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]:

$$\begin{aligned} \exists X. \tau &\stackrel{\text{def}}{=} \forall Y. (\forall X. \tau \rightarrow Y) \rightarrow Y \\ \text{pack } \langle \tau', t \rangle \text{ as } \exists X. \tau &\stackrel{\text{def}}{=} \Lambda Y. \lambda f : (\forall X. \tau \rightarrow Y). f [\tau'] t \\ \text{unpack } t_1 \text{ as } \langle X, x \rangle \text{ in } t_2 &\stackrel{\text{def}}{=} t_1 [\tau_2] (\Lambda X. \lambda x : \tau_1. t_2) \quad \text{where } t_1 : \exists X. \tau_1 \text{ and } t_2 : \tau_2 \end{aligned}$$

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.

<sup>14</sup> As mentioned, to the best of our knowledge this is not an existing conjecture that we disprove.



1569	Syntax:	
1570		
1571	Terms	$t ::= v \mid \text{if } t \text{ then } t \text{ else } t \mid x \mid tt \mid t\tau \mid t.1 \mid t.2 \mid \langle t, t \rangle$
1572		
1573		$\mid t : \tau \xRightarrow{p} \tau \mid t : \tau \xRightarrow{\phi} \tau \mid \text{blame } p$
1574	Values	$v ::= \text{unit} \mid \text{true} \mid \text{false} \mid \lambda x : \tau. t \mid \Lambda X. v \mid \langle v, v \rangle$
1575		
1576		$\mid v : \tau \rightarrow \tau \xRightarrow{\phi} \tau \rightarrow \tau \mid v : \forall X. \tau \xRightarrow{\phi} \forall X. \tau \mid v : \tau \xRightarrow{\neg\alpha} \alpha$
1577		
1578		$\mid v : \tau \rightarrow \tau \xRightarrow{p} \tau \rightarrow \tau \mid v : \tau \xRightarrow{p} \forall X. \tau \mid v : \gamma \xRightarrow{p} \star$
1579	Types	$\tau ::= \text{Unit} \mid \text{Bool} \mid \tau \times \tau \mid \tau \rightarrow \tau \mid \forall X. \tau \mid X \mid \alpha \mid \star$
1580	Ground types	$\gamma ::= \text{Unit} \mid \text{Bool} \mid \alpha \mid \star \times \star \mid \star \rightarrow \star$
1581	Convertibility labels	$\phi ::= \alpha \mid \neg\alpha$
1582	Compatibility labels	$p ::= l \mid \neg l$
1583		
1584		

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.

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  $(\Lambda X. \lambda x : X. x) \text{Unit}$  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.

## 7.2 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  $\text{Univ}$ , which we can construct thanks to the existence of the universal type  $\star$ :

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

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.

In the following, contextual equivalence in  $\lambda^B$ , indicated with  $\simeq$ , is defined analogously to Definition 2.1.

**THEOREM 7.1** ( $[t_u]$  AND  $[t_\omega]$  ARE NOT EQUIVALENT IN  $\lambda^B$ ).  $[t_u] \not\simeq [t_\omega]$ .

**PROOF.** We have that  $C^B[t_u] \hookrightarrow^* \text{unit}$  while  $C^B[t_\omega] \uparrow$ . We have verified this on paper and using the interpreter provided by Jamner and Siek to support Ahmed et al. [2017]'s results.<sup>15</sup>  
<sup>16</sup> We provide the literal encoding of  $C^B[t_u]$  and  $C^B[t_\omega]$  for use in the interpreter in the supplementary material.  $\square$

**THEOREM 7.2** (EMBEDDING  $\lambda^F$  INTO  $\lambda^B$  IS NOT FULLY ABSTRACT). *It is not true that*

$$\forall t_1, t_2. t_1 \simeq t_2 \iff [t_1] \simeq [t_2]$$

**PROOF.** Follows directly from Theorem 7.1 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], New et al. [2019a] and Xie et al. [2018].

## 8 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

<sup>15</sup>Available at <http://www.ccs.neu.edu/home/dijamner/paramblame/artifact/>

<sup>16</sup>Thanks to Jeremy Siek and others, for their kind support in doing this.

1667 Sumii-Pierce compiler is a useful secure compiler, even if it is not fully abstract. Hence we discuss  
 1668 how to weaken the requirements that full abstraction imposes on a translation from System F, by  
 1669 modifying System F in a way that weakens its contextual equivalences.

1670

### 1671 8.1 Fixing Sumii and Pierce's Compiler?

1672 One of the attractive features of Sumii and Pierce's original conjecture is that it relies *only* on  
 1673 encryption (or at least, an idealised version of encryption in the form of seals). Because it required  
 1674 only encryption, this suggested that System F types could even be enforced as the contract for  
 1675 an untrusted adversary, running at the other end of a communication channel, on an untrusted  
 1676 computer. However, if we analyse the counterexample, this ambition of using just encryption seems  
 1677 hard to maintain.

1678 Imagine that a compiled System F term (e.g.,  $t_u$  or  $t_\omega$ ) is communicating with such an adversary  
 1679 over a communication channel and we want to enforce that the adversary respects the contract  
 1680 represented by System F type `Univ`. What happens is the following:

- 1681 (1) The compiled term transmits the value `unit` of type `Unit`, encrypted as  $\{\text{unit}\}_\sigma$  of type `X` that  
 1682 is kept opaque from the adversary. The seal  $\sigma$  represents a fresh cryptographic key that we  
 1683 take care not to disclose to the adversary.
- 1684 (2) The adversary replies with a value of the unknown type `Y` (chosen by the adversary). The  
 1685 actual value transmitted back in our counterexample, is simply the encrypted value  $\{\text{unit}\}_\sigma$   
 1686 received in step 1.
- 1687 (3) The compiled term does not inspect the received value but simply transmits it back to the  
 1688 adversary as a value of type `Y`.
- 1689 (4) The adversary now takes the received value  $\{\text{unit}\}_\sigma$  and sends it back as a value of type `X`.
- 1690 (5) The compiled term receives this value, decrypts it using the private cryptographic key  $\sigma$  and  
 1691 uses the result as a value of type `Unit`.

1692  
 1693 What goes wrong in the above communication is that the value  $\{\text{unit}\}_\sigma$  sent back by the adversary  
 1694 in step 2 is essentially illegal. The adversary should have chosen a `Y` independently of `X` and values  
 1695 of such a type should intuitively not be able to contain values of type `X` (unless they are themselves  
 1696 packed in an existential package somehow). Let us try to think of what we could change in the  
 1697 communication protocol to enforce this.

1698 In a pure cryptographic setting, we believe there is little hope to fix the compiler. To understand  
 1699 this, consider how, in the cryptography setting, the value  $\{\text{unit}\}_\sigma$ , that we send to the adversary in  
 1700 step 1, simply represents a sequence of bits that is the result of some encryption algorithm applied  
 1701 to value `unit`. On the other hand, the value sent back by the adversary in step 2 is another sequence  
 1702 of bits that represents a value of an unknown type `Y`. We want to prevent this second value from  
 1703 somehow including the first value  $\{\text{unit}\}_\sigma$ , but since the type `Y` is unknown and the adversary is  
 1704 running on an untrusted computer, there seems to be little the compiler can check. Any sequence  
 1705 of bits received from the adversary could in principle be a cleverly-encoded version of  $\{\text{unit}\}_\sigma$ : they  
 1706 could have XORed the value with an arbitrary other bitsequence and still be able to retrieve the  
 1707 original afterwards.

1708 This (informal) argument suggests that the Sumii-Pierce compiler cannot be fixed in any way, as  
 1709 long as the target language contains only features that can be interpreted as a form of (idealised)  
 1710 cryptography. This would include the original sealing primitives (whatever way they are used), but  
 1711 also possible extensions that model a form of idealised signing (rather than encryption) (a track we  
 1712 were initially exploring).

1713 To fix the compiler, it seems like we need to add some kind of feature that takes us beyond a  
 1714 pure-cryptography setting. We believe such a feature could take the form of a primitive that checks

1715

1716 whether a value directly or indirectly contains values sealed with a certain seal. Such a primitive  
 1717 could perhaps allow to perform the required check on the value received from the adversary in step  
 1718 2. Such a primitive does not correspond to a form of idealised encryption: noticing, for example, that  
 1719 a value like  $\{\{v\}_{\sigma_1}\}_{\sigma_2}$  contains a value sealed with  $\sigma_1$  would break the cryptographic interpretation  
 1720 of  $\lambda^\sigma$ , as it requires looking inside an encrypted value without access to the key ( $\sigma_2$ ) and requires  
 1721 detecting, for example, arbitrarily XORed versions of a ciphertext. However, the primitive could  
 1722 still be implementable in non-cryptographic settings, like the hardware-enforced seals that are  
 1723 present in capability machines: a form of processor with native support for capabilities and sealing  
 1724 that has been developed recently [Watson et al. 2015]. Note that the attacker model in this setting  
 1725 is a bit different: we assume that the untrusted attacker is now running on trusted hardware.  
 1726

## 1727 8.2 Polymorphic Blame Calculus Without a Universal Type?

1728 Whether or not it is feasible to construct a gradual polymorphic language that fully abstractly  
 1729 embeds System F is unclear. In fact, even simply reconciling type safety results like the Blame  
 1730 Theorem (see Section 7) or the Dynamic Gradual Guarantee (which expresses that removing type  
 1731 annotations from a term should never cause the term to fail at runtime) with parametricity is the  
 1732 topic of ongoing research. In fact, in the design of their GSF calculus, Toro et al. [2019] explicitly  
 1733 abandon the dynamic gradual guarantee in order to salvage parametricity, formulated using a  
 1734 TWLR. New et al. [2019a] have recently reconciled (a TWLR-based form of) parametricity with the  
 1735 dynamic gradual guarantee, by requiring explicit sealing annotations in the source.

1736 All published gradual polymorphic calculi make use of dynamic sealing and as such, do not satisfy  
 1737 the FAE criterion, except perhaps for one. Labrada et al. [2022] have investigated an alternative  
 1738 design, based on a form of lexically-scoped sealing (in addition to some other new ideas like  
 1739 plausible sealing). It would lead us too far to explain the details here, but the design replaces the  
 1740 type  $\star$  with a family of types  $\star_{\{\bar{X}\}}$ , indexed by a set of type variables  $\bar{X}$ . The type  $\star_{\bar{X}}$  is only legal  
 1741 (i.e., well-formed) when the type variables  $\bar{X}$  are in scope. Injecting values of type  $X'$  into  $\star_{\bar{X}}$   
 1742 is only legal if  $X'$  is in the set  $\bar{X}$ . For our FAE counterexample, the effect is that the context can no  
 1743 longer produce a non-degenerate value of type [Univ](#), as there would no longer be a viable choice  
 1744 for  $Y$ . Any instance  $\star_{\{\bar{X}\}}$  of the universal type we could choose, could not have  $X$  as part of  $\{\bar{X}\}$ , as  
 1745  $X$  is not yet in scope at the moment where  $Y$  needs to be produced.

1746 Although the new design still has some restrictions, it satisfies a form of RLR-based parametricity.  
 1747 As such, the jury is still out, but we believe the FAE criterion may not be out of reach in the context  
 1748 of gradual parametricity.  
 1749

## 1750 8.3 Adjusting our Expectations

1751 For us, the solutions suggested above look like they might work, but they are not without downsides.  
 1752 The modified Sumii-Pierce compiler could no longer be used in a purely-cryptographic setting and  
 1753 the family of universal types  $\star_{\{\bar{X}\}}$  might be harder to use than the original  $\star$ . As such, we might  
 1754 consider alternative ways to address the lack of full abstraction.  
 1755

1756 One alternative is to abandon the choice of fully abstract compilation and rely on other notions.  
 1757 More concretely, perhaps a secure compiler or gradually typed language should not preserve  
 1758 arbitrary System F equivalences but only some of them, namely those that follow from a TWLR-  
 1759 based formulation of parametricity?

1760 Another alternative is to keep relying on fully abstract compilation and decide to simply adjust  
 1761 our expectations: perhaps preserving all System F equivalences is overly ambitious and we should  
 1762 find a way to eat what is on the table instead. Both in the case of Sumii and Pierce's compiler and  
 1763 the gradual lambda calculus, it seems like something non-trivial is being enforced, even though it  
 1764

1765 is not the preservation of arbitrary System F contextual equivalences. A way to formalise this is  
1766 to recover full abstraction by modifying the source language System F: weakening its contextual  
1767 equivalences in order to make them easier to preserve.

1768 In fact, our counterexample suggests a way to accomplish this: the problem is essentially that the  
1769 type [Univ](#) is degenerate in System F, but not in the target language. So what if we modify System F  
1770 to remove that degeneracy in the source language too? Specifically, what if we add a primitive type  
1771 that all other types can be embedded into and extracted from? Interestingly, it seems like what  
1772 we end up here is a simple version of the gradual type  $\star$ , together with injection and extraction  
1773 functions.

1774 Without working this out in more detail, we find it plausible that we can recover full abstraction  
1775 with such a modification, both for Sumii and Pierce’s compiler and the embedding into the poly-  
1776 morphic blame calculus. Ahmed et al. [2017], Toro et al. [2019] and New et al. [2019a] have shown  
1777 that such a variant of System F still satisfies useful (TWLR-based) parametricity results and that  
1778 useful free theorems follow from it, suggesting it is a suitable language for programmers to work  
1779 in.

1780

## 1781 9 RELATED WORK

1782 *System F and parametricity.* Parametric polymorphism was first introduced 50 years ago as an  
1783 informal concept by Strachey [2000]. A few years later, System F was independently discovered by  
1784 Reynolds [1974] and Girard [1972]. Reynolds [1983] later formalised Strachey’s informal concept  
1785 of parametricity using a logical relation for System F, as explained in Example 2.2 and Wadler  
1786 [1989] popularised the property using the slogan “Theorems for Free”. The property was further  
1787 developed by many different researchers over the years but for space reasons, we refer to Atkey  
1788 et al. [2014]; Wadler [2007] for an overview of related work.

1789 *Enforcing parametricity using dynamic sealing.* Dynamic sealing/unsealing was (informally)  
1790 proposed more than 40 years ago by Morris [1973a,b] as a way for dynamically enforcing type  
1791 abstractions. Much later, Pierce and Sumii [2000] developed this idea into a compiler that uses sealing  
1792 to enforce System F’s parametricity and conjectured full abstraction of the proposed compiler. The  
1793 target language in this work is a simply typed cryptographic  $\lambda$ -calculus. They already mention that  
1794 the dynamic enforcement of parametricity may also be useful to combine parametric polymorphism  
1795 and untyped languages, foreshadowing the work on parametrically polymorphic gradual type  
1796 systems that we discuss next.

1797 A few years later, the same authors report further on the simply typed cryptographic  $\lambda$ -calculus  
1798 and construct a logical relation for proving contextual equivalences for it [Sumii and Pierce 2003].  
1799 Another year later, they report on a bisimulation that can be used to prove contextual equivalence in  
1800  $\lambda^\sigma$ , which they originally constructed for proving the conjectured full abstraction [Sumii and Pierce  
1801 2004]. In this last paper, the compiler is presented for an untyped version of the cryptographic  
1802  $\lambda$ -calculus (as in this text), which renders it significantly simpler.

1803 Sealing was also used to enforce polymorphic contracts in PLT Scheme by Guha et al. [2007].  
1804 While technically similar, the assumptions in that work are somewhat different than in the above,  
1805 as contracts are about protecting the context from misbehaving terms, while Sumii and Pierce’s  
1806 compiler protects trusted terms from a misbehaving context. Like other work discussed in this  
1807 paper, Guha et al.’s polymorphic contracts fail to enforce the degeneracy of [Univ](#), so if they satisfy  
1808 a form of parametricity, it would have to be TWLR-based.

1809  
1810 *Gradual typing.* Gradual typing is a specific way of combining dynamic and static typing in a  
1811 single language, intended to create a gradual migration path from untyped to typed codebases.  
1812 Since it was originally proposed by Siek and Taha [2006] and Tobin-Hochstadt and Felleisen [2006],  
1813

1814 gradual typing has received a lot of attention: gradual extensions have been constructed for many  
 1815 different type systems, the notion of blame was adapted from the world of contracts [Findler and  
 1816 Felleisen 2002] to track the origin of a dynamic cast failure [Wadler and Findler 2009], correctness  
 1817 criteria were studied [Siek et al. 2015], the process of constructing a gradually-typed version of  
 1818 a pre-existing type system was to some extent automated [Cimini and Siek 2016; Garcia et al.  
 1819 2016] and last but not least, the entire approach has also been declared dead because of severe  
 1820 performance issues [Takikawa et al. 2016].

1821 Prior to the work on gradual typing, calculi which combined statically-typed languages with  
 1822 a universal type for interacting with untyped code, have been proposed by Henglein [1994] and  
 1823 Abadi et al. [1991].

1824 As mentioned in Section 7, formal criteria for gradual type systems have been proposed by Siek  
 1825 et al. [2015], including the gradual guarantee. Garcia and Tanter [2020] have later argued that  
 1826 gradually typed languages should additionally preserve reasoning principles of the static type  
 1827 system. Jacobs et al. [2021] have proposed that this can be formulated in a standard way in the  
 1828 form of the Fully Abstract Embedding (FAE) criterion and have established the criterion for the  
 1829 GTLC, a basic gradual type system.

1830 *Gradual typing and parametric polymorphism.* As an instance of a multi-language (as proposed by  
 1831 Matthews and Findler [2009]), Matthews and Ahmed [2008] construct a language that can embed  
 1832 both Scheme (i.e., an untyped lambda calculus) and ML (i.e., a parametrically polymorphic typed  
 1833 lambda calculus) using a notion of dynamic casts from one to the other. They enforce parametric  
 1834 polymorphism using a notion of run-time type generation<sup>17</sup> and prove a parametricity result. An  
 1835 error in the proof was later corrected [Ahmed et al. 2011c].

1836 Next, Ahmed et al. [2009b, 2011b] present the first version of the polymorphic blame calculus, a  
 1837 gradually-typed extension of System F that includes a notion of blame. The authors prove a number  
 1838 of correctness results, but Perconti found an intractable error in one of the proofs later [Ahmed  
 1839 et al. 2011a]. However, they do not consider parametricity or preservation of System F contextual  
 1840 equivalences (i.e., fully abstract embedding).

1841 Ahmed et al. [2017] present  $\lambda_B$ : a new version of the polymorphic blame calculus rectifying  
 1842 certain “topsy-turvy” aspects of its operational semantics (see Section 7). Additionally, they prove a  
 1843 parametricity result for this calculus, which we have discussed from the perspective of our results  
 1844 in Section 8.3.

1845 Igarashi et al. [2017] also present a new version of the polymorphic blame calculus, as an internal  
 1846 runtime representation for System  $F_G$  a new gradually-typed extension of System F. They make  
 1847 certain special restrictions to the consistency and precision relation of the system (later criticised by  
 1848 others [Toro et al. 2019]), which allow them to prove a version of the gradual guarantee. However,  
 1849 it does not seem to prevent our counterexample to the fully abstract embedding of System F.

1850 In an unpublished draft, Siek and Wadler [2016] study and interrelate three ways to achieve  
 1851 relational parametricity: universal types, runtime type generation and cryptographic sealing.  
 1852 They show translations from the polymorphic blame calculus from [Ahmed et al. 2017] into the  
 1853 cryptographic lambda calculus and back that they show to be simulations. They also study a calculus  
 1854  $\lambda_G$ , obtained by removing the universal type  $\star$  and casts from the polymorphic blame calculus,  
 1855 but keeping runtime type generation. They show that embedding System F into  $\lambda_G$  is also fully  
 1856 abstract. Both of these full abstraction results tally with our observations: both System F and  $\lambda_G$   
 1857 lack a universal type (making [Univ](#) degenerate) while both the polymorphic blame calculus and the  
 1858 cryptographic lambda calculus feature a universal type (making [Univ](#) non-degenerate). As such,  
 1859

1860  
 1861 <sup>17</sup>They call this sealing, but since their seals have no run-time representation, we prefer the term run-time type generation.  
 1862

1863 the embedding of  $\lambda_G$  into the polymorphic blame calculus will not be fully abstract, supplementing  
1864 their results.

1865 Xie et al. [2018] present a gradually typed language with parametric polymorphism, focusing on  
1866 the interplay of gradual typing with the subtyping relation in the presence of implicit polymorphism.  
1867 The result of their work is a source language that elaborates to  $\lambda_B$ , which we discussed before. As  
1868 such, the language provides the same form of parametricity than  $\lambda_B$  and we think our results apply  
1869 to it as well.

1870 More recently, Toro et al. [2019] proposed a new gradually typed calculus with explicit polymor-  
1871 phism, based on the AGT methodology by Garcia et al. [2016]. The methodology allows them to  
1872 construct a system that satisfies the refined criteria by Siek et al. [2015], except for the Dynamic  
1873 Gradual Guarantee. They show that this property is in conflict with parametricity in their system,  
1874 but they demonstrate a weaker property which they do satisfy. The gradual type is also a universal  
1875 type in their system and they prove a TWLR-based parametricity.

1876 Finally, New et al. [2019a] have developed PolyG<sup>v</sup>, the first gradual language to support both  
1877 parametricity and the dynamic gradual guarantee (which they refer to as graduality) at the same  
1878 time. The syntax of PolyG<sup>v</sup> departs from System F, requiring programmers to write explicit sealing  
1879 and unsealing annotations. Like previous work, New et al. use dynamic sealing and they prove a  
1880 form of parametricity based on a TWLR logical relation, although it is closer to standard System F  
1881 logical relations in some other respects.

1882  
1883 *Universal types.* Universal types have been studied by Longley [2003]. Our observation and proof  
1884 that System F's parametricity excludes a universal type is, to the best of our knowledge, novel.  
1885

1886  
1887 *Alternatives for full abstraction.* The property of full abstraction was proposed by Abadi [1998].  
1888 Abate et al. [2019] provide a lattice of secure compilation criteria (dubbed robust compilation) that  
1889 preserve classes of hyperproperties [Clarkson and Schneider 2010], i.e., arbitrary behaviours. The  
1890 general nature of the framework makes it hard to make precise statements, but we believe our  
1891 counterexamples would also disprove most of these alternative properties.

1892

## 1893 10 CONCLUSION

1894 This work started out years ago as an effort to prove Sumii and Pierce's conjecture, discussed in  
1895 Section 5. Our failure to do so has perhaps proven more interesting than a success might have been.  
1896 Specifically, we disprove the conjecture rather than proving it, we disprove another conjecture for  
1897 languages with non-parametric parametricity and we identify what we see as an important problem  
1898 in current parametrically polymorphic gradual languages: programmers reasoning about System F  
1899 programs cannot trust contextual equivalences to remain valid in the gradually typed extended  
1900 language. In addition to highlighting these problems, we discuss in Section 8 some ideas about how  
1901 the issues might be solved. None of them seem easy to solve and the solutions we propose have  
1902 downsides of their own, but we do believe they might be worth exploring further.

1903 During the work on the Sumii-Pierce conjecture, we also gained some more high-level insights  
1904 about variations of parametricity (type-world LRs versus lexically-scoped LRs) and their relation to  
1905 the lexical scope of type variables and the existence of universal types. These insights were not  
1906 mentioned explicitly in the previous version of this paper [Devriese et al. 2018] and in discussions  
1907 with experts in the field, we found that these insights were unclear. For this reason, this paper  
1908 focuses very clearly on these more high-level insights and we hope they may improve understanding  
1909 of the issues involved.

1910

1911

1912 **ACKNOWLEDGMENTS**

1913 This research is partially funded by project grants from the Research Fund KU Leuven, and from  
 1914 the Research Foundation Flanders (FWO). This work was partially supported by the Air Force  
 1915 Office of Scientific Research under award number FA9550-21-1-0054. This work was partially  
 1916 supported by the German Federal Ministry of Education and Research (BMBF) through funding for  
 1917 the CISPA-Stanford Center for Cybersecurity (FKZ: 13N1S0762). This work was partially supported  
 1918 by the Italian Ministry of Education through funding for the Rita Levi Montalcini grant (call of  
 1919 2019). The authors thank Phil Wadler for interesting comments and suggestions.

1922 **REFERENCES**

- 1923 Martín Abadi. Protection in programming-language translations. In *ICALP'98*, pages 868–883, 1998.
- 1924 Martín Abadi and Gordon D. Plotkin. On protection by layout randomization. *ACM Transactions on Information and System*  
 1925 *Security*, 15:8:1–8:29, July 2012. ISSN 1094-9224. doi: 10.1145/2240276.2240279.
- 1926 Martín Abadi, Luca Cardelli, Benjamin Pierce, and Gordon Plotkin. Dynamic Typing in a Statically Typed Language.  
 1927 *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>.
- 1928 Martín Abadi, Luca Cardelli, Benjamin Pierce, and Didier Rémy. Dynamic typing in polymorphic languages. *JOURNAL OF*  
 1929 *FUNCTIONAL PROGRAMMING*, 5:92–103, 1995.
- 1930 Martín Abadi, Cédric Fournet, and Georges Gonthier. Secure implementation of channel abstractions. In *IEEE Symposium*  
 1931 *on Logic in Computer Science*, pages 105–116, 1998.
- 1932 Martín Abadi, Cédric Fournet, and Georges Gonthier. Secure communications processing for distributed languages. In *IEEE*  
 1933 *Symposium on Security and Privacy*, pages 74–88, 1999.
- 1934 Martín Abadi, Cédric Fournet, and Georges Gonthier. Authentication primitives and their compilation. In *Principles of*  
 1935 *Programming Languages*, pages 302–315. ACM, 2000. doi: 10.1145/325694.325734.
- 1936 Carmine Abate, Arthur Azevedo de Amorim, Roberto Blanco, Ana Nora Evans, Guglielmo Fachini, Catalin Hritcu, Théo  
 1937 Laurent, Benjamin C. Pierce, Marco Sironi, and Andrew Tolmach. When good components go bad: Formally secure  
 1938 compilation despite dynamic compromise. CCS '18, 2018.
- 1939 Carmine Abate, Roberto Blanco, Deepak Garg, Marco Hrițcu, Cătălin Patrignani, and Jérémy Thibault. Journey beyond  
 1940 full abstraction: Exploring robust property preservation for secure compilation. In *2019 IEEE 32th Computer Security*  
 1941 *Foundations Symposium*, CSF 2019, June 2019.
- 1942 Pieter Agten, Raoul Strackx, Bart Jacobs, and Frank Piessens. Secure compilation to modern processors. In *Computer*  
 1943 *Security Foundations Symposium*, pages 171–185. IEEE, 2012. doi: 10.1109/CSF.2012.12.
- 1944 Amal Ahmed. *Semantics of Types for Mutable State*. PhD thesis, USA, 2004. AAI3136691.
- 1945 Amal Ahmed and Matthias Blume. Typed closure conversion preserves observational equivalence. In *International Conference*  
 1946 *on Functional Programming*, pages 157–168. ACM, 2008. ISBN 978-1-59593-919-7. doi: 10.1145/1411204.1411227.
- 1947 Amal Ahmed and Matthias Blume. An equivalence-preserving CPS translation via multi-language semantics. In *Proceedings*  
 1948 *of the 16th ACM SIGPLAN International Conference on Functional Programming*, ICFP '11, pages 431–444. ACM, 2011.  
 1949 ISBN 978-1-4503-0865-6. doi: 10.1145/2034773.2034830. URL <http://doi.acm.org/10.1145/2034773.2034830>.
- 1950 Amal Ahmed, Derek Dreyer, and Andreas Rossberg. State-dependent Representation Independence. In *Principles of*  
 1951 *Programming Languages*, pages 340–353. ACM, 2009a. doi: 10.1145/1480881.1480925.
- 1952 Amal Ahmed, Jacob Matthews, Robert Bruce Findler, and Philip Wadler. Blame for all. In *Workshop on Script-to-Program*  
 1953 *Evolution (STOP)*, pages 1–13, 2009b.
- 1954 Amal Ahmed, Robert Bruce Findler, Jeremy G. Siek, and Philip Wadler. Blame for all. Technical report, 2011a. URL  
 1955 <https://plt.eecs.northwestern.edu/blame-for-all/>.
- 1956 Amal Ahmed, Robert Bruce Findler, Jeremy G. Siek, and Philip Wadler. Blame for all. In *Principles of Programming Languages*,  
 1957 pages 201–214, 2011b.
- 1958 Amal Ahmed, Lindsey Kuper, and Jacob Matthews. Parametric polymorphism through run-time sealing, or, Theorems for  
 1959 low, low prices!, April 2011c. URL <http://www.ccs.neu.edu/home/amal/papers/paramseal-tr.pdf>.
- 1960 Amal Ahmed, Dustin Jamner, Jeremy G. Siek, and Philip Wadler. Theorems for free for free: Parametricity, with and without  
 types. *Proceedings of the ACM on Programming Languages*, 1(ICFP):39:1–39:28, August 2017. doi: 10.1145/3110283.
- 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.



- 1961 Matteo Cimini and Jeremy G. Siek. The Gradualizer: A Methodology and Algorithm for Generating Gradual Type Systems.  
1962 In *Principles of Programming Languages*. ACM, 2016. doi: 10.1145/2837614.2837632.
- 1963 Michael R. Clarkson and Fred B. Schneider. Hyperproperties. *J. Comput. Secur.*, 18(6):1157–1210, September 2010.
- 1964 Arthur Azevedo de Amorim, Matt Fredrikson, and Limin Jia. Reconciling noninterference and gradual typing. In Holger  
1965 Hermanns, Lijun Zhang, Naoki Kobayashi, and Dale Miller, editors, *LICS '20: 35th Annual ACM/IEEE Symposium on Logic  
1966 in Computer Science, Saarbrücken, Germany, July 8–11, 2020*, pages 116–129. ACM, 2020. doi: 10.1145/3373718.3394778.  
URL <https://doi.org/10.1145/3373718.3394778>.
- 1967 Dominique Devriese, Marco Patrignani, and Frank Piessens. Fully-abstract compilation by approximate back-translation. In  
1968 *Principles of Programming Languages*, pages 164–177, 2016.
- 1969 Dominique Devriese, Marco Patrignani, Frank Piessens, and Steven Keuchel. Modular, Fully-abstract Compilation by  
1970 Approximate Back-translation. *Logical Methods in Computer Science*, 13(4 lmcsc:4011), October 2017. doi: 10.23638/  
LMCS-13(4:2)2017.
- 1971 Dominique Devriese, Marco Patrignani, and Frank Piessens. Parametricity versus the universal type. In *Proceedings of the  
1972 45th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2018, Los Angeles, CA,  
1973 USA, 2016*, 2018.
- 1974 Derek Dreyer, Georg Neis, and Lars Birkedal. The Impact of Higher-order State and Control Effects on Local Relational  
1975 Reasoning. In *International Conference on Functional Programming*, pages 143–156. ACM, 2010. doi: 10.1145/1863543.  
1863566.
- 1976 Derek Dreyer, Amal Ahmed, and Lars Birkedal. Logical step-indexed logical relations. *Logical Methods in Computer Science*,  
1977 7(2), 2011a.
- 1978 Derek Dreyer, Ralf Jung, Jan-Oliver Kaiser, Hoang-Hai Dang, and David Swasey. Semantics of type systems lecture notes,  
2011b. Available at: <https://plv.mpi-sws.org/semantics/2017/lecturenotes.pdf>.
- 1979 Matthias Felleisen. On the expressive power of programming languages. In *Selected Papers from the Symposium on 3rd  
1980 European Symposium on Programming, ESOP '90*, pages 35–75, New York, NY, USA, 1991. Elsevier North-Holland, Inc.
- 1981 Robert Bruce Findler and Matthias Felleisen. Contracts for Higher-order Functions. In *International Conference on Functional  
1982 Programming*. ACM, 2002. doi: 10.1145/581478.581484.
- 1983 Cedric Fournet, Nikhil Swamy, Juan Chen, Pierre-Evariste Dagand, Pierre-Yves Strub, and Benjamin Livshits. Fully abstract  
1984 compilation to JavaScript. In *Principles of Programming Languages*, pages 371–384. ACM, 2013. doi: 10.1145/2429069.  
2429114.
- 1985 Ronald Garcia and Éric Tanter. Gradual typing as if types mattered. Workshop on Gradual Typing, 2020. URL <https://wgt20.irif.fr/wgt20-final28-acmpaginated.pdf>.
- 1986 Ronald Garcia, Alison M. Clark, and Éric Tanter. Abstracting Gradual Typing. In *Principles of Programming Languages*.  
1987 ACM, 2016. doi: 10.1145/2837614.2837670.
- 1988 Jean-Yves Girard. *Interprétation Fonctionnelle et Élimination Des Coupures de l'arithmétique d'ordre Supérieur*. PhD thesis,  
1989 Université Paris VII, 1972.
- 1990 Arjun Guha, Jacob Matthews, Robert Bruce Findler, and Shriram Krishnamurthi. Relationally-parametric Polymorphic  
1991 Contracts. In *Symposium on Dynamic Languages*. ACM, 2007. doi: 10.1145/1297081.1297089.
- 1992 Fritz Henglein. Dynamic typing: Syntax and proof theory. *Science of Computer Programming*, 22(3):197–230, June 1994.  
1993 ISSN 0167-6423. doi: 10.1016/0167-6423(94)00004-2.
- 1994 Yuu Igarashi, Taro Sekiyama, and Atsushi Igarashi. On polymorphic gradual typing. In *International Conference on Functional  
1995 Programming*. ACM, 2017.
- 1996 Koen Jacobs, Amin Timany, and Dominique Devriese. Fully abstract from static to gradual. *Proceedings of the ACM on  
1997 Programming Languages*, 5(POPL):7:1–7:30, January 2021. doi: 10.1145/3434288.
- 1998 Radha Jagadeesan, Corin Pitcher, Julian Rathke, and James Riely. Local memory via layout randomization. In *Computer  
1999 Security Foundations Symposium*, pages 161–174. IEEE Computer Society, 2011. doi: 10.1109/CSF.2011.18.
- 2000 Yannis Juglaret, Catalin Hritcu, Arthur Azevedo de Amorim, Benjamin C. Pierce, Antal Spector-Zabusky, and Andrew  
2001 Tolmach. Towards a fully abstract compiler using micro-policies: Secure compilation for mutually distrustful components.  
2002 *CoRR*, abs/1510.00697, 2015. URL <http://arxiv.org/abs/1510.00697>.
- 2003 Yannis Juglaret, Cătălin Hrițcu, Arthur Azevedo de Amorim, and Benjamin C. Pierce. Beyond good and evil: Formalizing  
2004 the security guarantees of compartmentalizing compilation. In *Computer Security Foundations Symposium*, 2016.
- 2005 Elizabeth Labrada, MatĀnjas Toro, ĀĀric Tanter, and Dominique Devriese. Plausible sealing for gradual parametricity.  
2006 *Proceedings of the ACM on Programming Languages*, 7(OOPSLA), 2022. To appear.
- 2007 Adriaan Larmuseau, Marco Patrignani, and Dave Clarke. A secure compiler for ML modules. In *Programming Languages  
2008 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).
- 2009 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://>

- 2010 //doi.acm.org/10.1145/2851613.2851796.
- 2011 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.
- 2012 Jacob Matthews and Amal Ahmed. Parametric polymorphism through run-time sealing or, theorems for low, low prices! volume 4960 of *LNC3*, pages 16–31. 2008.
- 2013 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.
- 2014 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>.
- 2015 John C. Mitchell. On abstraction and the expressive power of programming languages. *Science of Computer Programming*, 21(2):141 – 163, 1993. ISSN 0167-6423.
- 2016 James H. Morris, Jr. Protection in programming languages. *Commun. ACM*, 16(1):15–21, January 1973a. ISSN 0001-0782. doi: 10.1145/361932.361937.
- 2017 James H. Morris, Jr. Types are not sets. In *Principles of Programming Languages*, pages 120–124, 1973b.
- 2018 Georg Neis, Derek Dreyer, and Andreas Rossberg. Non-parametric parametricity. In *ICFP '09*, pages 135–148, 2009.
- 2019 Georg Neis, Derek Dreyer, and Andreas Rossberg. Non-parametric parametricity. *Journal of Functional Programming*, 21:497–562, 2011. ISSN 1469-7653.
- 2020 Max New, William J. Bowman, and Amal Ahmed. Fully abstract compilation via universal embedding. In *ICFP'16*, 2016a.
- 2021 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, 2016b. doi: 10.1145/2951913.2951941.
- 2022 Max S. New, Dustin Jamner, and Amal Ahmed. Graduality and parametricity: Together again for the first time. *Proceedings of the ACM on Programming Languages*, 4(POPL):46:1–46:32, December 2019a. doi: 10.1145/3371114.
- 2023 Max S. New, Daniel R. Licata, and Amal Ahmed. Gradual type theory. *Proc. ACM Program. Lang.*, 3(POPL):15:1–15:31, 2019b. doi: 10.1145/3290328. URL <https://doi.org/10.1145/3290328>.
- 2024 Joachim Parrow. Expressiveness of process algebras. *Elec. Not. Theo. Comp. Sci.*, 209(0):173 – 186, 2008.
- 2025 Marco Patrignani. Why should anyone use colours? or, syntax highlighting beyond code snippets. CoRR abs/2001.11334, 2020.
- 2026 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*.
- 2027 Marco Patrignani and Deepak Garg. Robustly safe compilation. In *Programming Languages and Systems - 28th European Symposium on Programming, ESOP 2019, ESOP'19, 2019*.
- 2028 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.
- 2029 Marco Patrignani, Dominique Devriese, and Frank Piessens. On Modular and Fully Abstract Compilation. In *Computer Security Foundations Symposium*, 2016.
- 2030 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.
- 2031 Benjamin Pierce. *Types and Programming Languages*. MIT Press, 2002.
- 2032 Benjamin Pierce and Eijiro Sumii. Relating cryptography and polymorphism. manuscript, 2000. URL <http://www.kb.ecei.tohoku.ac.jp/~sumii/pub/infohide.pdf>.
- 2033 Andrew M. Pitts. Existential types: Logical relations and operational equivalence. In Kim G. Larsen, Sven Skyum, and Glynn Winskel, editors, *Automata, Languages and Programming*, Lecture Notes in Computer Science, pages 309–326. Springer Berlin, Heidelberg, 1998. doi: 10.1007/BFb0055063.
- 2034 Gordon D. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, 5:223–255, 1977.
- 2035 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.
- 2036 John C. Reynolds. Types, abstraction, and parametric polymorphism. In *Information Processing*, pages 513–523. North Holland, 1983.
- 2037 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>.
- 2038 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>.
- 2039 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>.

- 2059 Jeremy G. Siek and Walid Taha. Gradual typing for functional languages. In *SCHEME*, pages 81–92, 2006.
- 2060 Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. Refined criteria for gradual typing. In *Summit on*  
2061 *Advances in Programming Languages*, volume 32 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 274–293,  
2062 Dagstuhl, Germany, 2015. Schloss Dagstuhl-Leibniz-Zentrum für Informatik. doi: 10.4230/LIPIcs.SNAPL.2015.274.
- 2063 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>.
- 2064 Christopher Strachey. Fundamental Concepts in Programming Languages. *Higher-Order and Symbolic Computation*, 13(1-2):  
2065 11–49, April 2000. ISSN 1388-3690, 1573-0557. doi: 10.1023/A:1010000313106.
- 2066 Eijiro Sumii and Benjamin C. Pierce. Logical Relations for Encryption. *J. Comput. Secur.*, 11(4):521–554, July 2003. ISSN  
2067 0926-227X.
- 2068 Eijiro Sumii and Benjamin C. Pierce. A bisimulation for dynamic sealing. In *Principles of Programming Languages*, pages  
161–172, 2004.
- 2069 Asumu Takikawa, Daniel Feltey, Ben Greenman, Max S. New, Jan Vitek, and Matthias Felleisen. Is Sound Gradual Typing  
2070 Dead? In *Principles of Programming Languages*. ACM, 2016. doi: 10.1145/2837614.2837630.
- 2071 Sam Tobin-Hochstadt and Matthias Felleisen. Interlanguage migration: From scripts to programs. In *Companion to the 21st*  
2072 *ACM SIGPLAN Symposium on Object-Oriented Programming Systems, Languages, and Applications*, OOPSLA '06, page  
964–974, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 159593491X. doi: 10.1145/1176617.  
2073 1176755. URL <https://doi.org/10.1145/1176617.1176755>.
- 2074 Matias Toro, Ronald Garcia, and Éric Tanter. Type-driven gradual security with references. *ACM Trans. Program. Lang.*  
2075 *Syst.*, 40(4):16:1–16:55, 2018. doi: 10.1145/3229061. URL <https://doi.org/10.1145/3229061>.
- 2076 Matias Toro, Elizabeth Labrada, and Éric Tanter. Gradual Parametricity, Revisited. *Proc. ACM Program. Lang.*, 3(POPL):  
2077 17:1–17:30, January 2019. ISSN 2475-1421. doi: 10.1145/3290330.
- 2078 Philip Wadler. Theorems for free! In *Functional Programming Languages and Computer Architecture*, pages 347–359. ACM,  
1989.
- 2079 Philip Wadler. The Girard-Reynolds isomorphism (second edition). *Theoretical Computer Science*, 375(1), 2007. ISSN  
2080 0304-3975. doi: 10.1016/j.tcs.2006.12.042.
- 2081 Philip Wadler and Robert Bruce Findler. Well-typed programs can't be blamed. In *Programming Languages and Systems*,  
2082 pages 1–16. Springer, Berlin, Heidelberg, March 2009. doi: 10.1007/978-3-642-00590-9\_1.
- 2083 Robert N. M. Watson, Jonathan Woodruff, Peter G. Neumann, Simon W. Moore, Jonathan Anderson, David Chisnall, Nirav H.  
2084 Dave, Brooks Davis, Khilan Gudka, Ben Laurie, Steven J. Murdoch, Robert M. Norton, Michael Roe, Stacey D. Son, and  
2085 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.
- 2086 Andrew K. Wright. Simple imperative polymorphism. *LISP and Symbolic Computation*, 8(4):343–355, December 1995. ISSN  
2087 1573-0557. doi: 10.1007/BF01018828.
- 2088 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.
- 2089
- 2090
- 2091
- 2092
- 2093
- 2094
- 2095
- 2096
- 2097
- 2098
- 2099
- 2100
- 2101
- 2102
- 2103
- 2104
- 2105
- 2106
- 2107