Parametricity versus the Universal Type

DOMINIQUE DEVRIESE, KU Leuven, Belgium

MARCO PATRIGNANI*, MPI-SWS & CISPA, Germany

FRANK PIESENS, KU Leuven, Belgium

There has long been speculation in the scientific literature on how to dynamically enforce parametricity such as that yielded by System F. Almost 20 years ago, Sumii and Pierce proposed a formal compiler from System F into the cryptographic lambda calculus: an untyped lambda calculus extended with an idealised model of encryption. They conjectured that this compiler was fully abstract, i.e. that compiled terms are contextually equivalent if and only if the original terms were, a property that can be seen as a form of secure compilation. The conjecture has received attention in several other publications since then, but remains open to this day.

More recently, several researchers have been looking at gradually-typed languages that extend System F. In this setting it is natural to wonder whether embedding System F into these gradually-typed languages preserves contextual equivalence and thus parametricity.

In this paper, we answer both questions negatively. We provide a concrete counterexample: two System F terms whose contextual equivalence is not preserved by the Sumii-Pierce compiler, nor the embedding into the polymorphic blame calculus. This counterexample relies on the absence in System F of what we call a *universal type*, i.e., a type that all other types can be injected into and extracted from. As the languages in which System F is compiled have a universal type, the compilation cannot be fully abstract; this paper explains why.

We believe this paper thus sheds light on recent results in the field of gradually typed languages and it provides a perspective for further research into secure compilation of polymorphic languages.

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

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. 2018. Parametricity versus the Universal Type. *Proc. ACM Program. Lang.* 2, POPL, Article 38 (January 2018), 23 pages. <https://doi.org/10.1145/3158126>

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. Parametricity expresses the property that polymorphic functions in System F cannot behave differently when invoked at different types. It can be formalised as an (often relational) property that all inhabitants of a type

*This work was performed when the author was at MPI-SWS.

Authors' addresses: Dominique Devriese, imec-Distrinet, KU Leuven, Leuven, Belgium, name.surname@cs.kuleuven.be; Marco Patrignani, MPI-SWS & CISPA, Saarbrücken, Germany, marcopat@mpi-sws.org; Frank Piessens, imec-Distrinet, KU Leuven, Leuven, Belgium, name.surname@cs.kuleuven.be.

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

© 2018 Association for Computing Machinery.

2475-1421/2018/1-ART38

<https://doi.org/10.1145/3158126>

automatically satisfy [Reynolds 1983; Wadler 1989]. Such properties are automatically satisfied for any function without the need to verify their code. A typical example is the fact that any value of type $\forall X. X \rightarrow X$ must behave as the identity function (or possibly diverge in a non-terminating variant of System F).

The topic of this paper is the interaction of System F programs with programs in a weaker type-system. In such weaker systems, a System F component may receive values of a certain type from more weakly typed code. It is then not obvious that such values will satisfy the same properties as System F values, specifically parametricity. For example, if a System F component receives a function of type $\forall X. X \rightarrow X$ from a weakly-typed component, will it necessarily behave as the identity function (or diverge)?

It is already worth noting that the weaker type systems we consider are all non-terminating, and when interacting with them, the only properties that have a chance of being preserved, are those that hold in the presence of non-termination. This is the reason that the variant of System F that we consider is non-terminating, obtained by adding recursive types. We discuss in more detail why this is the right choice in Section 3.2.

We will look at the question of preserving properties of System F terms in two specific settings.

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 low-level components. The goal of secure compilation is to ensure that those low-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, security of a compiler has been mostly expressed as full abstraction: given two source-level contextually equivalent programs, their target-level compilation are also contextually equivalent (and vice versa) [Abadi 1998]. Compiler full abstraction has been proven for compilers that rely on address-space layout randomisation [Abadi and Plotkin 2012; Jagadeesan et al. 2011], or secure enclaves [Agten et al. 2012; Larmuseau et al. 2015, 2016; Patrignani et al. 2015, 2016], tagged architectures [Juglaret et al. 2016, 2015], JavaScript [Fournet et al. 2013], typed closure conversion [Ahmed and Blume 2008], cryptographic primitives [Abadi et al. 1998, 1999, 2000; Bugliesi and Giunti 2007] etc.

In a paper from the year 2000, Pierce and Sumii [2000] proposed a compiler from System F to a cryptographic lambda calculus, which enforces parametricity using a form of idealised encryption primitives (sealing) called lambda-seal (λ^{σ}). They conjectured that this compiler was fully abstract, a conjecture that has received further research attention, but remains open to this day [Siek and Wadler 2016; Sumii and Pierce 2004].

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

A natural way to formally express that a gradual language preserves the typed language's reasoning principles is to reuse the same notion of fully abstract compilation that we mentioned

above. Specifically, if the embedding of the typed language into the gradual language is fully abstract, then similarly to secure compilation, this expresses that untyped code can only interact with typed code in ways that are also possible using just typed code.

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. 2011a](#)], further developed by [Igarashi et al. \[2017\]](#) and [Ahmed et al. \[2017\]](#). [Ahmed et al.](#) also prove a form of parametricity in the polymorphic blame calculus, but their formulation is rather different than existing ones (e.g. [Dreyer et al. \[2009\]](#); [Reynolds \[1983\]](#)) because of its use of Kripke worlds to capture the interpretation of generated runtime types. As such, it remains unclear how this notion of parametricity is related to existing ones and whether the polymorphic blame calculus preserves all System F reasoning, particularly its contextual equivalences (i.e. whether it embeds System F in a fully abstract way).

Two questions, one answer. In this paper, we solve both open problems, answering the questions negatively. We prove that Sumii and Pierce’s compiler is not fully abstract and we show the same for the embedding of System F into the polymorphic blame calculus (i.e. equivalences are not preserved under the embedding).

Our counterexamples are essentially based on the observation that System F parametricity forbids the existence of a universal type: a type which any other type can be embedded into and extracted from.¹ More precisely, consider the following type:

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

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

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

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

Finally, we discuss a number of consequences and perspectives that follow from our results. First, we discuss some thoughts on how Sumii and Pierce’s compiler might be fixed (so that it does enforce full abstraction), and what could be modified in the polymorphic blame calculus to make it

¹Note that by this definition, the term *universal type* is broader than Abadi’s *dynamic type* [[Abadi et al. 1991](#)] or Siek and Taha’s *gradual type* [[Siek and Taha 2006](#)]: the former includes *any* type that arbitrary values can be embedded into and extracted from, while the latter are specific primitive types, specifically intended for representing untyped values in a typed language.

preserve all of System F's contextual equivalences. However, in neither case there appears to be a panacea solution: potential fixes all seem to come with certain downsides. Because of this, we also discuss whether we should not instead adjust our expectations and find a way to formalise the guarantees that we do get from both Sumii and Pierce's compiler and the polymorphic blame calculus. This exercise allows us to shed some new light on recent results from the literature.

Outline. We start our discussion in Section 2 by repeating the definition of System F and presenting the terms t_u and t_d . We prove they are contextually equivalent from the degeneracy of System F, relying on an existing formalisation of System F's parametricity [Dreyer et al. 2009]. Next, we disprove Sumii and Pierce's conjecture in Section 3 by explaining their compiler and how it treats t_d and t_u and fails to preserve their equivalence. In Section 4, we turn to gradual typing, introducing polymorphic blame calculi and demonstrating how t_u and t_d 's contextual equivalence is also lost in the presence of these calculi's universal type. Finally, we discuss perspectives and consequences of our results in Section 5, related work in Section 6 and we conclude in Section 7. For space constraints we omit most proofs and formalisation, which can be found in the supplementary material.

To make the distinction between source and target language visually apparent, we typeset elements of System F in a **blue, bold** font, elements of λ^σ in a **red, sans-serif** font and elements of the blame calculus in an *orange, italic* font.

2 SYSTEM F AND THE TYPE Univ

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

2.1 The Source Language λ^F

Figure 1 presents the variant of System F that we will be using in this paper, which we indicate with λ^F . In addition to standard polymorphic functions ($\forall X. \tau$) and existential packages ($\exists X. \tau$), the variant includes recursive types $\mu X. \tau$ (so there are no termination guarantees), as well as **Unit** and **Bool** and product $\tau_1 \times \tau_2$ types. In the figure, we present terms t , values v and types τ . We show the most important typing rules $\Delta; \Gamma \vdash t : \tau$ in terms of term and type variable contexts Γ and Δ (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 $\mathbb{E}$.

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

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

2.2 Parametricity

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

Example 2.2 ($\mathbb{Z}_n$ implementation in λ^F). We could for instance represent the type $\mathbb{Z}_n$ of integers modulo n as a tuple $\langle \text{zero}, \text{succ}, \text{zero?} \rangle$ of type $\exists X. (X \times (X \rightarrow X)) \times (X \rightarrow \text{Bool})$, and then either

Syntax:

$$\begin{aligned}
t &::= v \mid x \mid t \mid t.1 \mid t.2 \mid \langle t, t \rangle \mid \Lambda X. t \mid t \tau \mid \text{if } t \text{ then } t \text{ else } t \\
&\quad \mid \text{pack } \langle \tau, t \rangle \text{ as } \exists X. \tau \mid \text{unpack } t \text{ as } \langle X, x \rangle \text{ in } t \mid \text{roll } t \mid \text{unroll } t \\
v &::= \text{unit} \mid \text{true} \mid \text{false} \mid \lambda x : \tau. t \mid \langle v, v \rangle \mid \Lambda X. t \mid \text{pack } \langle \tau, v \rangle \text{ as } \exists X. \tau \mid \text{roll } v \\
\tau &::= \text{Unit} \mid \text{Bool} \mid \tau \rightarrow \tau \mid \tau \times \tau \mid X \mid \forall X. \tau \mid \exists X. \tau \mid \mu X. \tau \\
\Gamma &::= \emptyset \mid \Gamma, (x : \tau) \quad \Delta ::= \emptyset \mid \Delta, X \\
\mathbb{E} &::= [\cdot] \mid \mathbb{E} t \mid v \mathbb{E} \mid \mathbb{E}.1 \mid \mathbb{E}.2 \mid \langle \mathbb{E}, t \rangle \mid \langle v, \mathbb{E} \rangle \mid \mathbb{E} \tau \mid \mathbb{E}; t \mid \text{inl } \mathbb{E} \mid \text{inr } \mathbb{E} \\
&\quad \mid \text{case } \mathbb{E} \text{ of inl } x_1 \mapsto t \mid \text{inr } x_2 \mapsto t \mid \text{if } \mathbb{E} \text{ then } t \text{ else } t \mid \text{pack } \langle \tau, \mathbb{E} \rangle \text{ as } \exists X. \tau \\
&\quad \mid \text{unpack } \mathbb{E} \text{ as } \langle X, x \rangle \text{ in } t \mid \text{unpack } v \text{ as } \langle X, x \rangle \text{ in } \mathbb{E} \mid \text{roll } \mathbb{E} \mid \text{unroll } \mathbb{E}
\end{aligned}$$

Typing rules (excerpt):

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

Evaluation rules (excerpt):

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

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

implement this type as a k -tuple of booleans, or by means of a type of natural numbers defined using recursive types.

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 parametricity in the form of a theorem that all λ^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 τ_1, τ_2 , any relation R between values of types τ_1 and τ_2 , and any two closed values v_1, v_2 of type τ_1 and τ_2 respectively, if (v_1, v_1) is in R , then $f \tau_1 v_1$ and $f \tau_2 v_2$ will either both diverge or reduce to values $(v'_1, v'_2) \in R$.

The relational property is derived from the type using what is known as a logical relation.² For this paper, it is not necessary to understand how this works, so we refer the interested reader to other articles that explain this well [Dreyer et al. 2009; Wadler 1989].

2.3 The Type Univ

A central role in this paper is reserved for the type **Univ**:

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

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

In our non-terminating variant of λ^F , **Univ** is clearly inhabited, for example by this value:

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

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

A crucial observation for this paper is that *all* values of type **Univ** are degenerate in this sense. Intuitively, this is because a single type **Y** needs to be chosen, independently of the types **X** that will be embedded into it. Because nothing is known upfront about these **X**s and nothing can be learnt about them after invocation (because **X** must be treated parametrically), no viable choice for **Y** can be made.³

To understand more formally how the degeneracy of **Univ** follows from parametricity, it is instructive to consider the same type in a simpler setting. Let us consider the following lemma, in a terminating variant of λ^F , which has universal and existential types, the unit type but no recursive types. We indicate the terminating variant of λ^F as $\lambda_{\perp}^F$.

LEMMA 2.3 (**Univ** IS NOT INHABITED IN $\lambda_{\perp}^F$). *Assume that*

- A value **v** of type **Univ**: $\emptyset; \emptyset \vdash \mathbf{v} : \mathbf{Univ}$
- **v** satisfies unary parametricity, i.e. there exist $\tau_Y, P_Y \in \mathcal{P}(\llbracket \tau_Y \rrbracket)$ and **v'** such that we have $\mathbf{v} \equiv \text{pack } \langle \tau_Y, \mathbf{v}' \rangle \text{ as } \mathbf{Univ}$ and for all $\tau_X, P_X \in \mathcal{P}(\llbracket \tau_X \rrbracket)$, $\mathbf{v}_X \in P_X$ and $\mathbf{v}_Y \in P_Y$, there exist $\mathbf{v}'_X$ and $\mathbf{v}'_Y$ such that $(\mathbf{v}' \tau_X).1 \mathbf{v}_X \hookrightarrow^* \mathbf{v}'_Y$ and $(\mathbf{v}' \tau_X).2 \mathbf{v}_Y \hookrightarrow^* \mathbf{v}'_X$ and $\mathbf{v}'_Y \in P_Y$ and $\mathbf{v}'_X \in P_X$.

This is not possible.

In Lemma 2.3, we assume some value **v** of type **Univ**. Additionally, we avoid formulating a general parametricity result for this calculus but instead just postulate that **v** satisfies its unary parametricity property (assuming a set semantics for closed types $\llbracket \cdot \rrbracket$). The lemma claims that these two assumptions lead to a contradiction, so that (if we believe unary parametricity), we can conclude that **Univ** is empty.

To prove this lemma, we use the same trick as we do in the more complex proof of degeneracy of **Univ** in λ^F : we get a single P_Y , but we instantiate P_X twice with two different predicates:

- PROOF. • Take $\tau_X = \mathbf{Unit}$, $P_X = \{\mathbf{unit}\}$ and $\mathbf{v}_X = \mathbf{unit} \in P_X$. The property above gives us a $\mathbf{v}'_Y$ such that $(\mathbf{v}' \mathbf{Unit}).1 \mathbf{unit} \hookrightarrow^* \mathbf{v}'_Y$ and $\mathbf{v}'_Y \in P_Y$.
- Now take $\tau_X = \mathbf{Unit}$, $P'_X = \emptyset$ and take $\mathbf{v}_Y$ to be the $\mathbf{v}'_Y \in P_Y$ from the previous point. The property above gives us a $\mathbf{v}'_X$ such that $(\mathbf{v}' \mathbf{Unit}).2 \mathbf{v}_Y \hookrightarrow^* \mathbf{v}'_X$ and $\mathbf{v}'_X \in P'_X$.
 - Contradiction: $\mathbf{v}'_X \in P'_X$ but $P'_X = \emptyset$. □

²Unary and n-ary variants of parametricity also exist, though the relational one is the most commonly studied.

³Note that the situation is entirely different if we swap the quantifications in the type: $\mathbf{Triv} \stackrel{\text{def}}{=} \forall X. \exists Y. (X \rightarrow Y) \times (Y \rightarrow X)$.

Essentially, this proof exploits the intuitive argument that we outlined above: a value of type **Univ** needs to provide a single interpretation for the existentially quantified **Y** that must not depend on the choice of **X**. However, rather than applying this argument to the choice of types τ_Y and τ_X , it exploits parametricity by applying the argument to the choice of predicate interpretations P_Y and P_X .

2.4 Two Contextually Equivalent Terms

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

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

The reason that t_u and t'_d 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 discussed above, while for t'_d , the term ω_{Unit} in the body ensures divergence. Note that the degeneracy of **Univ** is essential: if the context were able to produce a non-degenerate value of type **Univ**, then t_u would not diverge when applied to it, so that the context could distinguish t_u from t'_d .

For technical reasons, our proofs will not actually use the term t'_d but rather the following term t_d , which first uses the value of type **Univ** in the same way as t_u , but then diverges afterwards.

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

The only reason we do this is because it is easier to prove that $t_u \simeq_{\text{ctx}} t_d$ than to do the same for t'_d (it seems like a proof for t'_d would require both a unary and binary logical relation).

We now have the following theorem. A full proof of this result can be found in the technical appendix ⁴, but we provide an overview here.

THEOREM 2.4 (t_u AND t_d ARE CONTEXTUALLY EQUIVALENT IN λ^F). $\emptyset; \emptyset \vdash t_u \simeq_{\text{ctx}} t_d : \text{Univ} \rightarrow \text{Unit}$.

Proof Sketch. This proof uses a pre-existing formulation of parametricity for λ^F by Dreyer et al. [2009]. The proof essentially applies the same idea as the above proof for Lemma 2.3: we use parametricity for the argument of type **Univ**, from which we obtain a single relational interpretation for the existentially quantified type **Y**. We can then use parametricity with two different relational interpretations for the universally quantified type **X**, and we end up with two terms related in the expression relation for an empty relation. From this, we can conclude that both must diverge. $\square$

3 DISPROVING THE SUMII-PIERCE CONJECTURE

The first technical result of this paper is that we disprove a conjecture by Pierce and Sumii [2000] about the dynamic enforcement of parametricity in a cryptographic calculus.

Sumii and Pierce's conjecture is about a compiler from λ^F to an untyped lambda calculus with sealing (idealised encryption) called λ^σ . In this section, we first introduce the target language λ^σ (Section 3.1) and the compiler from λ^F to λ^σ (Section 3.2). We then prove that the compiler is not fully-abstract: there exist two terms that are contextually-equivalent in λ^F but whose compilation is inequivalent in λ^σ (Section 3.3).

⁴The technical appendix is available as auxiliary material for this publication through the ACM Digital Library.

Syntax:

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

Evaluation rules (excerpt):

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

Fig. 2. λ^σ syntax and evaluation rules (excerpt). A λ^σ 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.

Remark. Sumii and Pierce have presented both a typed [Pierce and Sumii 2000] as well as an untyped [Sumii and Pierce 2004] version of λ^σ . We use the untyped version because it is quite a bit simpler.⁵ However, the typed version suffers from the same problem, as we show in detail in the technical report. Essentially, both settings break degeneracy of Univ because they feature a universal type: the untype of all values in the untyped target language and the type **bits** of ciphertexts produced by encryption (sealing) in the typed target language.

3.1 The Cryptographic Lambda Calculus λ^σ

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

Most syntactic constructs are standard for an untyped λ -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 σ 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 $\text{let } \{x\}_{t_1} = t_2 \text{ in } t_3 \text{ else } t_4$ is for unsealing or decrypting. It first evaluates t_1 to a seal σ_1 and t_2 to $\{v_2\}_{\sigma_2}$ (or produces a run-time error if either result is not of that form). If σ_1 and σ_2 are equal, t_3 is evaluated with x bound to the decrypted value v_2 , otherwise t_4 is evaluated.

Program contexts in λ^σ are defined as for λ^F and are denoted with $\mathbb{C}$. Contextual equivalence in λ^σ , indicated with $\approx_{ctx}$ is defined analogously to Definition 2.1. Note that the quantified contexts

⁵The 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.

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

Example 3.1 ($\mathbb{Z}_3$ in λ^σ). First, we introduce two helper functions:

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

Then, we can define a λ^σ correspondent of [z3](#) as:

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

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

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

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

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

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

We can implement a similar protection in λ^σ . To enforce parametric behavior of a polymorphic function received from the context, we create a new seal for every invocation, and we encrypt all parameters of the quantified type with that seal. For instance, the term above could be implemented in λ^σ as follows:

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

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

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

3.2 Sumii and Pierce's Compiler

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

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

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

We now present these steps and discuss their meaning. These definitions are taken from Sumii and Pierce [2004], except that we add support for recursive types (see below), and that we make some notational changes and some minor technical changes.

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.⁶

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

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 τ . 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 τ . Dually, applying $\text{confine}_{\eta;\tau}$ to a value v prevents v from behaving in a way that is not allowed by type τ . For any free type variables X in τ , η 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, η has the following syntax: $\eta ::= \emptyset \mid \eta, X \mapsto (t, t)$.

For defining **protect**, and **confine**, we rely on the following seal_σ and unseal_σ functions, which seal and unseal values with a given seal σ .

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

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 τ . 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 τ 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 σ . Similarly, confining

⁶We are assuming a standard desugaring of let expressions to function applications.

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

Fig. 3. The dynamic wrappers of Sumii and Pierce’s compiler. fix_2 is presented in Eq. (2) p. 12.

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 3.1.

Applying the compiler to $\mathbf{z3}$ from Example 3.1 results in the λ^σ term $\mathbf{z3}$ from Example 2.2 (modulo some additional β -reductions).

Technical remark. Sumii and Pierce’s conjectured compiler for untyped λ^σ already had a technical flaw, albeit unrelated to parametricity. In fact, they had a terminating source language $\lambda_{\mathbb{F}}$ and an untyped – thus possibly diverging – target language. With this discrepancy it is impossible to build a fully-abstract compiler for a simple reason. Consider these two terms, which are equivalent in $\lambda_{\mathbb{F}}$ (but not in λ^F):

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

Their compilation to λ^σ should intuitively be (modulo $\text{protect}_\cdot$, which we elide for simplicity):

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

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

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

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

We don't go into detail for this, but we can define **fix₂** as follows, using a definition of Plotkin's Z combinator⁷ as **fix**:

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

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

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

3.3 Disproving the Sumii-Pierce Conjecture

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

CONJECTURE 3.2 (SUMII AND PIERCE). $\emptyset; \emptyset \vdash t_1 \simeq_{ctx} t_2 : \tau$ if and only if $\emptyset \vdash \llbracket t_1 \rrbracket_{\lambda}^{\lambda^F} \simeq_{ctx} \llbracket t_2 \rrbracket_{\lambda}^{\lambda^F}$

However, we can now prove that this conjecture is false. A counterexample is given by the terms **t_u** and **t_d** which we proved contextually equivalent in Theorem 2.4. Compiling **t_u** and **t_d** does not produce contextually equivalent λ^σ terms:

THEOREM 3.3 (**t_u** AND **t_d** ARE NOT EQUIVALENT AFTER COMPILATION TO λ^σ). $\emptyset \vdash \llbracket t_u \rrbracket_{\lambda}^{\lambda^F} \not\simeq_{ctx} \llbracket t_d \rrbracket_{\lambda}^{\lambda^F}$

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

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

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

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 $\mathbb{C}$ 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 behavior.

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

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

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

THEOREM 3.4 ($\llbracket \cdot \rrbracket_{\lambda}^{\lambda^F}$ IS NOT FULLY ABSTRACT). *It is not true that*

$$\forall t_1, t_2. t_1 \simeq_{ctx} t_2 \iff \llbracket t_1 \rrbracket_{\lambda}^{\lambda^F} \simeq_{ctx} \llbracket t_2 \rrbracket_{\lambda}^{\lambda^F}$$

PROOF. Follows easily from Theorem 2.4 and the counterexample in Theorem 3.3. $\square$

⁷Plotkin's Z combinator is a variant of the more well-known Y combinator that works in a call-by-value setting, but is restricted to constructing fixpoints that are functions, see Pierce [2002, §5.2].

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

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

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

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

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

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

4 SYSTEM F EQUIVALENCES VS. GRADUAL TYPES

The second technical result in this paper is related to parametric polymorphism in gradually-typed languages. As explained in the introduction, gradually-typed languages provide a path for migrating codebases of untyped code to typed code. From this high-level idea, a number of natural design goals follow and the literature contains a number of correctness properties that formally express these objectives.

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

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

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

Syntax:

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

Fig. 4. 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.

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

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

In this section, we answer this question negatively for the polymorphic blame calculus, which extends System F into a gradually typed language. We first present Ahmed et al.'s λ^B , a gradually-typed, polymorphic lambda-calculus (Section 4.1). Next, we reconsider t_u and t_d from Theorem 2.4 in λ^B and present a λ^B context that differentiates them (Section 4.2). This shows that the embedding of λ^F into λ^B is not fully abstract.

4.1 The λ^B Calculus [Ahmed et al. 2017]

There exist several versions of the polymorphic blame calculus in the literature [Ahmed et al. 2017, 2011a; Igarashi et al. 2017]. We take λ^B to be the polymorphic blame calculus as described by Ahmed et al. [2017]. This version modifies certain behavior that was “topsy turvy” in the original version [Ahmed et al. 2011a]: 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 λ^B is presented in Fig. 4. The calculus contains all terms and types of λ^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]:

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

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

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

Then, λ^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 $\star$. λ^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 τ to τ' that can potentially fail at runtime (in which case *blame* p is raised). Casts can be used to inject types τ 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 τ and τ' 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, λ^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 γ into $\star$.

Polymorphic function application in λ^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 $(AX. \lambda x : X. x) \text{Unit}$ does not reduce to $\lambda x : \text{Unit}. x$ but to $\lambda x : \alpha. x$ for a fresh runtime type label α , where the assignment $\alpha := \text{Unit}$ is remembered in a type-name store Σ . Conversions ($t : \tau \xRightarrow{\phi} \tau'$) represent a notion of static casts for converting between such a runtime type label α and the type that is assigned to it in the store Σ . More generally, a convertibility judgement $\Sigma; \Delta \vdash \tau <^{\phi} \tau'$ defines when τ can be legally converted into τ' by expanding $(+\alpha \in \phi)$ or reducing $(-\alpha \in \phi)$ runtime type labels' definitions.

4.2 Embedding of λ^F into λ^B is not Fully Abstract

To embed λ^F (defined in Section 2.1) into λ^B , most constructs can simply be mapped to the corresponding construct in λ^B . However, there is a discrepancy between type abstractions in λ^B and λ^F . The difference is that λ^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 we use the logical relation by Dreyer et al. [2009] and the polymorphic blame calculus defined by Ahmed et al. [2017], which make different choices and would both be non-trivial to modify.

Therefore, we embed polymorphic functions from λ^F into λ^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 $AX. t$ are mapped to $[AX. t] \stackrel{\text{def}}{=} AX. \lambda_{-}. [t]$.

Our two contextually equivalent System F terms t_u and t_d embed into λ^B as $[t_u]$ and $[t_d]$. 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 λ^B context $\mathcal{C}^B$ that differentiates them. As before, the context simply applies the terms to a non-degenerate value of type Univ , which we can construct thanks to the existence of the universal type $\star$:

$$\mathcal{C}^B \stackrel{\text{def}}{=} [\cdot] (\text{pack } \langle \star, AX. \langle \lambda x : X. x : X \xRightarrow{p} \star, \lambda x : \star. x : \star \xRightarrow{p'} X \rangle \rangle \text{ as Univ})$$

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

THEOREM 4.1 ($\llbracket t_u \rrbracket$ AND $\llbracket t_d \rrbracket$ ARE NOT EQUIVALENT IN λ^B). $\llbracket t_u \rrbracket \not\approx_{ctx} \llbracket t_d \rrbracket$.

PROOF. We have that $\mathcal{C}^B[\llbracket t_u \rrbracket] \hookrightarrow^* unit$ while $\mathcal{C}^B[\llbracket t_d \rrbracket] \uparrow$. We have verified this on paper and using the interpreter provided by Jamner and Siek to support Ahmed et al. [2017]'s results.^{8 9} We provide the literal encoding of $\mathcal{C}^B[\llbracket t_u \rrbracket]$ and $\mathcal{C}^B[\llbracket t_d \rrbracket]$ for use in the interpreter in the technical appendix. $\square$

THEOREM 4.2 (EMBEDDING λ^F INTO λ^B IS NOT FULLY ABSTRACT). *It is not true that*

$$\forall t_1, t_2. t_1 \approx_{ctx} t_2 \iff \llbracket t_1 \rrbracket \approx_{ctx} \llbracket t_2 \rrbracket$$

PROOF. Follows directly from Theorem 4.1 and Theorem 2.4. $\square$

5 DISCUSSION

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

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

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

5.1 Fixing Sumii and Pierce's Compiler?

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

Imagine that a compiled System F term (e.g. t_u or t_d) is communicating with such an adversary over a communication channel and we want to enforce that the adversary respects the contract represented by System F type **Univ**. What happens is the following:

- (1) The compiled term transmits the value **unit** of type **Unit**, encrypted as $\{\text{unit}\}_\sigma$ of type X that is kept opaque from the adversary. The seal σ represents a fresh cryptographic key that we take care not to disclose to the adversary.
- (2) The adversary replies with a value of the unknown type Y (chosen by the adversary). The actual value transmitted back in our counterexample, is simply the encrypted value $\{\text{unit}\}_\sigma$ received in step 1.

⁸ Available at <http://www.ccs.neu.edu/home/dijamner/paramblame/artifact/>

⁹ Thanks to Jeremy Siek and others, for their kind support in doing this.

- (3) The compiled term does not inspect the received value but simply transmits it back to the adversary as a value of type $\mathbf{Y}$.
- (4) The adversary now takes the received value $\{\mathbf{unit}\}_\sigma$ and sends it back as a value of type $\mathbf{X}$.
- (5) The compiled term receives this value, decrypts it using the private cryptographic key σ and uses the result as a value of type $\mathbf{Unit}$.

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

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

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

To fix the compiler, it seems like we need to add some kind of feature that takes us beyond a pure-cryptography setting. We believe such a feature could take the form of a primitive that checks whether a value directly or indirectly contains values sealed with a certain seal. Such a primitive could perhaps allow to perform the required check on the value received from the adversary in step 2. Such a primitive does not correspond to a form of idealised encryption: noticing, for example, that a value like $\{\{v\}_{\sigma_1}\}_{\sigma_2}$ contains a value sealed with σ_1 would break the cryptographic interpretation of λ^σ , as it requires looking inside an encrypted value without access to the key (σ_2) and requires detecting, for example, arbitrarily XORed versions of a ciphertext. However, the primitive could still be implementable in non-cryptographic settings, like the hardware-enforced seals that are present in capability machines: a form of processor with native support for capabilities and sealing that has been developed recently [Watson et al. 2015]. Note that the attacker model in this setting is a bit different: we assume that the untrusted attacker is now running on trusted hardware.

5.2 Polymorphic Blame Calculus Without a Universal Type?

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

We believe it may not be needed to remove the type $\star$ entirely, but that we can instead replace it with a family of types $\star_{\tilde{X}}$, indexed by a set of type variables $\tilde{X}$. The type $\star_{\tilde{X}}$ would only be legal (i.e. well-formed) when the type variables $\tilde{X}$ are in scope. For values whose type is a type variable X' , injection into $\star_{\tilde{X}}$ would only be legal if X' is in the set $\tilde{X}$.

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

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

5.3 Adjusting our Expectations

The above suggested solutions seem to us like they might work, but they are not without downsides. The modified Sumii-Pierce compiler could no longer be used in a purely-cryptographic setting and the family of universal types $\star_{\{X\}}$ would be harder to use than the original $\star$. As such, we might consider an alternative way to address the lack of full abstraction. We could also decide to simply adjust our expectations: perhaps preserving all System F equivalences is overly ambitious and we should find a way to eat what is on the table instead. Both in the case of Sumii and Pierce's compiler and the gradual lambda calculus, it seems like something non-trivial is being enforced, even though it is not the preservation of arbitrary System F contextual equivalences. A way to formalise this is to recover full abstraction by modifying the source language System F: weakening its contextual equivalences in order to make them easier to preserve.

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

Without working this out in more detail, we find it plausible that we can recover full abstraction with such a modification, both for Sumii and Pierce's compiler and the embedding into the polymorphic blame calculus. The question is if this is a direction we want to take: is such a modified version of System F a source language that programmers would want to work in? To answer this question, we need to understand better what sort of language we end up with. For example, does such a language still satisfy a form of parametricity?

A recent result shows that the answer to this last question is positive. [Ahmed et al. \[2017\]](#) formulate and prove a form of parametricity for the polymorphic blame calculus. Their formulation of parametricity uses a logical relation with a Kripke world that assigns a relational interpretation to a fresh runtime type label when it is first introduced. This resembles the assignment of an invariant to a freshly allocated reference in a logical relation for a language with mutable references (see e.g. [Ahmed et al. \[2009\]](#)). Interestingly, this seemingly effectful formulation seems to preclude the trick that we have used in the proof of Theorem 2.4 of assigning two different relational interpretations to a single type variable: when the new runtime type label for the type application $x' \text{ Univ}$ (the subterm of t_u) is allocated, we need to choose a single relational interpretation for it and there seems to be no way to use two different ones instead. Intuitively, this suggests that the authors are allowing for the fresh runtime type label to play a global role in the application, rather than its role being restricted to the scope of the new type variable. This makes sense considering how a runtime type label can escape its scope when it is cast into the universal type $\star$. In such a setting, the degeneracy of **Univ** does not follow, but there still is a useful form of parametricity left, as [Ahmed et al.](#) demonstrate with some examples.

In other words, it appears that the degeneracy of type `Univ` does not follow from parametricity in general, but just from the particular kind of parametricity in typical formulations of System F and that [Ahmed et al.](#) have devised a weaker, alternative formulation that holds in a polymorphic blame calculus. Note that it must be weaker because it holds in a calculus where the type `Univ` is not degenerate, whereas the more standard formulation that we have used for System F implies this degeneracy. However, the weaker form of parametricity still implies useful results: [Ahmed et al.](#) prove free theorems for types $\forall X. \forall Y. X \rightarrow Y \rightarrow X$ and $\forall X. X \times X \rightarrow X \times X$.

6 RELATED WORK

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

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

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

Sealing was also used to enforce polymorphic contracts in PLT Scheme by [Guha et al. \[2007\]](#). While technically similar, the assumptions in that work are somewhat different than in the above, as contracts are about protecting the context from misbehaving terms, while Sumii and Pierce’s compiler protects trusted terms from a misbehaving context. Like other work discussed in this paper, [Guha et al.](#)’s polymorphic contracts fail to enforce the degeneracy of `Univ`, so if they satisfy a form of parametricity, it would have to be similar to [Ahmed et al. \[2017\]](#) in the sense that it does not imply this degeneracy.

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

Gradual typing and parametric polymorphism. As an instance of a multi-language (as proposed by Matthews and Findler [2009]), Matthews and Ahmed [2008] construct a language that can embed both Scheme (i.e. an untyped lambda calculus) and ML (i.e. a parametrically polymorphic typed lambda calculus) using a notion of dynamic casts from one to the other. They enforce parametric polymorphism using a notion of run-time type generation¹⁰ and prove a parametricity result. An error in the proof was later corrected [Ahmed et al. 2011c].

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

Very recently, Ahmed et al. [2017] present a new version of the polymorphic blame calculus rectifying certain “topsy-turvy” aspects of its operational semantics (see Section 4). Additionally, they prove a parametricity result for this calculus, which we have discussed from the perspective of our results in Section 5.3.

Igarashi et al. [2017] also present a new version of the polymorphic blame calculus, as an internal runtime representation for System F_G a new gradually-typed extension of System F. They make a more substantial change to the language by adding a distinction between static and gradual type variables. The former are used when embedding System F terms and they cannot be cast into or out of the gradual type $\star$. This prevents situations where a polymorphic type like $\forall X. X \rightarrow X$ is cast to $\forall X. \star \rightarrow X$ which would cause sealing or unsealing positions to be forgotten, breaking certain correctness properties. The addition allows them to prove properties like the gradual guarantee, but it does not seem to prevent our counterexample to the fully abstract embedding of System F.

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

Non-parametric parametricity. In a related setting, Neis et al. [2009, 2011] study a form of runtime type generation to protect parametrically polymorphic functions when interacting with code that can use a typecase primitive. They prove a parametricity result that applies to appropriately wrapped System F values. We suspect that our results may be adapted to this setting, as in such a non-parametrically polymorphic setting, the type $\forall X. X$ could be used as a universal type (that every other type can be embedded into or out from).

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

¹⁰They call this sealing, but since their seals have no run-time representation, we prefer the term run-time type generation.

7 CONCLUSION

This work started out as an effort to prove Sumii and Pierce’s conjecture, discussed in Section 3. However, rather than solving a decade-old open problem, this paper has turned out to point out new problems instead. We disprove the conjecture rather than proving it, and we identify what we see as an important problem in current parametrically polymorphic gradual languages: programmers reasoning about System F programs cannot trust contextual equivalences to remain valid in the gradually typed extended language. Solving those problems does not seem easy: it is not even clear in which direction such solutions had best be sought. We point out some directions in Section 5, but none seems perfect: they each come with downsides like reduced programmer convenience or the requirement of additional target language primitives. However, despite the absence of positive results, our results improve the understanding of the interaction between parametrically polymorphic languages and less strongly typed ones and may be useful to steer research into the concerned topics.

ACKNOWLEDGMENTS

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

REFERENCES

- Martín Abadi. 1998. Protection in Programming-Language Translations. In *ICALP’98*. 868–883.
- Martín Abadi, Luca Cardelli, Benjamin Pierce, and Gordon Plotkin. 1991. Dynamic Typing in a Statically Typed Language. *ACM Trans. Program. Lang. Syst.* 13, 2 (April 1991), 237–268. DOI: <http://dx.doi.org/10.1145/103135.103138>
- Martín Abadi, Cédric Fournet, and Georges Gonthier. 1998. Secure Implementation of Channel Abstractions. In *IEEE Symposium on Logic in Computer Science*. 105–116.
- Martín Abadi, Cédric Fournet, and Georges Gonthier. 1999. Secure Communications Processing for Distributed Languages. In *IEEE Symposium on Security and Privacy*. 74–88.
- Martín Abadi, Cédric Fournet, and Georges Gonthier. 2000. Authentication Primitives and their Compilation. In *Principles of Programming Languages*. ACM, 302–315. DOI: <http://dx.doi.org/10.1145/325694.325734>
- Martín Abadi and Gordon D. Plotkin. 2012. On Protection by Layout Randomization. *ACM Transactions on Information and System Security* 15, Article 8 (July 2012), 8:1–8:29 pages. DOI: <http://dx.doi.org/10.1145/2240276.2240279>
- Pieter Agten, Raoul Strackx, Bart Jacobs, and Frank Piessens. 2012. Secure Compilation to Modern Processors. In *Computer Security Foundations Symposium*. IEEE, 171–185. DOI: <http://dx.doi.org/10.1109/CSF.2012.12>
- Amal Ahmed and Matthias Blume. 2008. Typed Closure Conversion Preserves Observational Equivalence. In *International Conference on Functional Programming*. ACM, 157–168. DOI: <http://dx.doi.org/10.1145/1411204.1411227>
- Amal Ahmed, Justin Damner, Jeremy G. Siek, and Philip Wadler. 2017. Theorems for Free for Free. In *ICFP*.
- Amal Ahmed, Derek Dreyer, and Andreas Rossberg. 2009. State-dependent Representation Independence. In *Principles of Programming Languages*. ACM, 340–353. DOI: <http://dx.doi.org/10.1145/1480881.1480925>
- Amal Ahmed, Robert Bruce Findler, Jeremy G. Siek, and Philip Wadler. 2011a. Blame for All. In *Principles of Programming Languages*. 201–214.
- Amal Ahmed, Robert Bruce Findler, Jeremy G. Siek, and Philip Wadler. 2011b. *Blame for All*. Technical Report. <https://plt.eecs.northwestern.edu/blame-for-all/>
- Amal Ahmed, Lindsey Kuper, and Jacob Matthews. 2011c. Parametric polymorphism through run-time sealing, or, Theorems for low, low prices! (April 2011). <http://www.ccs.neu.edu/home/amal/papers/paramseal-tr.pdf>
- Amal Ahmed, Jacob Matthews, Robert Bruce Findler, and Philip Wadler. 2009. Blame for All. In *Workshop on Script-to-Program Evolution (STOP)*. 1–13.
- Robert Atkey, Neil Ghani, and Patricia Johann. 2014. A relationally parametric model of dependent type theory. In *Principles of Programming Languages*. ACM, 503–515. DOI: <http://dx.doi.org/10.1145/2535838.2535852>

- Michele Bugliesi and Marco Giunti. 2007. Secure Implementations of Typed Channel Abstractions. In *Principles of Programming Languages*. ACM, 251–262. DOI: <http://dx.doi.org/10.1145/1190216.1190253>
- Matteo Cimini and Jeremy G. Siek. 2016. The Gradualizer: A Methodology and Algorithm for Generating Gradual Type Systems. In *Principles of Programming Languages*. ACM. DOI: <http://dx.doi.org/10.1145/2837614.2837632>
- Dominique Devriese, Marco Patrignani, and Frank Piessens. 2016. Fully-abstract Compilation by Approximate Back-translation. In *Principles of Programming Languages*. 164–177.
- Dominique Devriese, Marco Patrignani, Frank Piessens, and Steven Keuchel. 2017. Modular, Fully-Abstract Compilation by Approximate Back-Translation. to appear in LMCS. (2017). <http://arxiv.org/abs/1703.09988>
- D. Dreyer, A. Ahmed, and L. Birkedal. 2009. Logical Step-Indexed Logical Relations. In *Logic In Computer Science*. 71–80.
- Robert Bruce Findler and Matthias Felleisen. 2002. Contracts for Higher-order Functions. In *International Conference on Functional Programming*. ACM. DOI: <http://dx.doi.org/10.1145/581478.581484>
- Cedric Fournet, Nikhil Swamy, Juan Chen, Pierre-Evariste Dagand, Pierre-Yves Strub, and Benjamin Livshits. 2013. Fully Abstract Compilation to JavaScript. In *Principles of Programming Languages*. ACM, 371–384. DOI: <http://dx.doi.org/10.1145/2429069.2429114>
- Ronald Garcia, Alison M. Clark, and Alric Tanter. 2016. Abstracting Gradual Typing. In *Principles of Programming Languages*. ACM. DOI: <http://dx.doi.org/10.1145/2837614.2837670>
- Jean-Yves Girard. 1972. *Interprétation Fonctionnelle et Élimination Des Coupures de l'arithmétique d'ordre Supérieur*. Ph.D. Dissertation. Université Paris VII.
- Arjun Guha, Jacob Matthews, Robert Bruce Findler, and Shriram Krishnamurthi. 2007. Relationally-parametric Polymorphic Contracts. In *Symposium on Dynamic Languages*. ACM. DOI: <http://dx.doi.org/10.1145/1297081.1297089>
- Yuu Igarashi, Taro Sekiyama, and Atsushi Igarashi. 2017. On Polymorphic Gradual Typing. In *International Conference on Functional Programming*. ACM.
- Radha Jagadeesan, Corin Pitcher, Julian Rathke, and James Riely. 2011. Local Memory via Layout Randomization. In *Computer Security Foundations Symposium*. IEEE Computer Society, 161–174. DOI: <http://dx.doi.org/10.1109/CSF.2011.18>
- Yannis Juglaret, Cătălin Hrițcu, Arthur Azevedo de Amorim, and Benjamin C. Pierce. 2016. Beyond Good and Evil: Formalizing the Security Guarantees of Compartmentalizing Compilation. In *Computer Security Foundations Symposium*.
- Yannis Juglaret, Catalin Hritcu, Arthur Azevedo de Amorim, Benjamin C. Pierce, Antal Spector-Zabusky, and Andrew Tolmach. 2015. Towards a Fully Abstract Compiler Using Micro-Policies: Secure Compilation for Mutually Distrustful Components. *CoRR* abs/1510.00697 (2015). <http://arxiv.org/abs/1510.00697>
- Adriaan Larmuseau, Marco Patrignani, and Dave Clarke. 2015. A Secure Compiler for ML Modules. In *Programming Languages and Systems - 13th Asian Symposium*. 29–48. DOI: http://dx.doi.org/10.1007/978-3-319-26529-2_3
- Adriaan Larmuseau, Marco Patrignani, and Dave Clarke. 2016. Implementing a Secure Abstract Machine. In *Symposium on Applied Computing*. ACM, 2041–2048. DOI: <http://dx.doi.org/10.1145/2851613.2851796>
- John R. Longley. 2003. Universal Types and What They are Good For. In *Domain Theory, Logic and Computation*. Springer, Dordrecht, 25–63. DOI: http://dx.doi.org/10.1007/978-94-017-1291-0_2
- Jacob Matthews and Amal Ahmed. 2008. Parametric Polymorphism through Run-Time Sealing or, Theorems for Low, Low Prices! *LNCS*, Vol. 4960. 16–31.
- Jacob Matthews and Robert Bruce Findler. 2009. Operational Semantics for Multi-language Programs. *ACM Trans. Program. Lang. Syst.* 31, 3 (April 2009), 12:1–12:44. DOI: <http://dx.doi.org/10.1145/1498926.1498930>
- James H. Morris, Jr. 1973a. Protection in Programming Languages. *Commun. ACM* 16, 1 (Jan. 1973), 15–21. DOI: <http://dx.doi.org/10.1145/361932.361937>
- James H. Morris, Jr. 1973b. Types Are Not Sets. In *Principles of Programming Languages*. 120–124.
- Georg Neis, Derek Dreyer, and Andreas Rossberg. 2009. Non-parametric Parametricity. In *ICFP '09*. 135–148.
- Georg Neis, Derek Dreyer, and Andreas Rossberg. 2011. Non-parametric parametricity. *Journal of Functional Programming* 21 (2011), 497–562.
- Marco Patrignani, Pieter Agten, Raoul Strackx, Bart Jacobs, Dave Clarke, and Frank Piessens. 2015. Secure Compilation to Protected Module Architectures. *ACM Trans. Program. Lang. Syst.* 37, Article 6 (April 2015), 6:1–6:50 pages. DOI: <http://dx.doi.org/10.1145/2699503>
- Marco Patrignani, Dominique Devriese, and Frank Piessens. 2016. On Modular and Fully Abstract Compilation. In *Computer Security Foundations Symposium*.
- Benjamin Pierce. 2002. *Types and Programming Languages*. MIT Press.
- Benjamin Pierce and Eijiro Sumii. 2000. Relating Cryptography and Polymorphism. (2000). <http://www.kb.ecei.tohoku.ac.jp/~sumii/pub/infohide.pdf> manuscript.
- Gordon D. Plotkin. 1977. LCF Considered as a Programming Language. *Theoretical Computer Science* 5 (1977), 223–255.
- John C. Reynolds. 1974. Towards a Theory of Type Structure. In *Programming Symposium*. Lecture Notes in Computer Science, Vol. 19. Springer-Verlag, 408–423.
- J. C. Reynolds. 1983. Types, Abstraction, and Parametric Polymorphism. In *Information Processing*. North Holland, 513–523.

- Jeremy Siek and Philip Wadler. 2016. The key to blame: Gradual typing meets cryptography. (2016). <http://homepages.inf.ed.ac.uk/wadler/papers/blame-key/blame-key.pdf> draft.
- Jeremy G. Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. In *SCHEME*. 81–92.
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015. Refined Criteria for Gradual Typing. In *Summit on Advances in Programming Languages (Leibniz International Proceedings in Informatics (LIPIcs))*, Vol. 32. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 274–293. DOI : <http://dx.doi.org/10.4230/LIPIcs.SNAPL.2015.274>
- Richard Statman. 1991. *A local translation of untyped [lambda] calculus into simply typed [lambda] calculus*. Technical Report. <http://repository.cmu.edu/cgi/viewcontent.cgi?article=1454&context=math>
- Christopher Strachey. 2000. Fundamental Concepts in Programming Languages. *Higher-Order and Symbolic Computation* 13, 1-2 (April 2000), 11–49. DOI : <http://dx.doi.org/10.1023/A:1010000313106>
- Eijiro Sumii and Benjamin C. Pierce. 2003. Logical Relations for Encryption. *J. Comput. Secur.* 11, 4 (July 2003), 521–554.
- Eijiro Sumii and Benjamin C. Pierce. 2004. A Bisimulation for Dynamic Sealing. In *Principles of Programming Languages*. 161–172.
- Asumu Takikawa, Daniel Feltey, Ben Greenman, Max S. New, Jan Vitek, and Matthias Felleisen. 2016. Is Sound Gradual Typing Dead?. In *Principles of Programming Languages*. ACM. DOI : <http://dx.doi.org/10.1145/2837614.2837630>
- Philip Wadler. 1989. Theorems for Free!. In *Functional Programming Languages and Computer Architecture*. ACM, 347–359.
- Philip Wadler. 2007. The Girard-Reynolds isomorphism (second edition). *Theoretical Computer Science* 375, 1 (2007). DOI : <http://dx.doi.org/10.1016/j.tcs.2006.12.042>
- Philip Wadler and Robert Bruce Findler. 2009. Well-Typed Programs Can’t Be Blamed. In *Programming Languages and Systems*. Springer, Berlin, Heidelberg, 1–16. DOI : http://dx.doi.org/10.1007/978-3-642-00590-9_1
- R. N. M. Watson, J. Woodruff, P. G. Neumann, S. W. Moore, J. Anderson, D. Chisnall, N. Dave, B. Davis, K. Gudka, B. Laurie, S. J. Murdoch, R. Norton, M. Roe, S. Son, and M. Vadera. 2015. CHERI: A Hybrid Capability-System Architecture for Scalable Software Compartmentalization. In *IEEE Symposium on Security and Privacy*. DOI : <http://dx.doi.org/10.1109/SP.2015.9>