



Plausible Sealing for Gradual Parametricity

ELIZABETH LABRADA*, University of Chile, Chile

MATÍAS TORO, University of Chile, Chile

ÉRIC TANTER, University of Chile, Chile

DOMINIQUE DEVRIESE, KU Leuven, Belgium

Graduality and parametricity have proven to be extremely challenging notions to bring together. Intuitively, enforcing parametricity gradually requires possibly sealing values in order to detect violations of uniform behavior. Toro et al. (2019) argue that the two notions are incompatible in the context of System F, where sealing is transparently driven by potentially imprecise type information, while New et al. (2020) reconcile both properties at the cost of abandoning the syntax of System F and requiring user-provided sealing annotations that are not subject to graduality guarantees. Furthermore, all current proposals rely on a global form of dynamic sealing in order to enforce parametric behavior at runtime, which weakens parametric reasoning and breaks equivalences in the static language. Based on the observation that the tension between graduality and parametricity comes from the early commitment to seal values based on type information, we propose *plausible sealing* as a new intermediate language mechanism that allows postponing such decisions to runtime. We propose an intermediate language for gradual parametricity, Funky, which supports plausible sealing in a simplified setting where polymorphism is restricted to instantiations with base and variable types. We prove that Funky satisfies both parametricity and graduality, mechanizing key lemmas in Agda. Additionally, we avoid global dynamic sealing and instead propose a novel lexically-scoped form of sealing realized using a representation of evidence inspired by the category of spans. As a consequence, Funky satisfies a standard formulation of parametricity that does not break System F equivalences. In order to show the practicality of plausible sealing, we describe a translation from Funk, a source language without explicit sealing, to Funky, that takes care of inserting plausible sealing forms. We establish graduality of Funk, subject to a restriction on type applications, and explain the source-level parametric reasoning it supports. Finally, we provide an interactive prototype along with illustrative examples both novel and from the literature.

CCS Concepts: • **Theory of computation** → **Operational semantics**.

Additional Key Words and Phrases: Gradual typing, polymorphism, parametricity

ACM Reference Format:

Elizabeth Labrada, Matías Toro, Éric Tanter, and Dominique Devriese. 2022. Plausible Sealing for Gradual Parametricity. *Proc. ACM Program. Lang.* 6, OOPSLA1, Article 70 (April 2022), 28 pages. <https://doi.org/10.1145/3527314>

*Research conducted while Elizabeth was affiliated with Vrije Universiteit Brussel and KU Leuven, Belgium. This work is partially funded by ANID FONDECYT projects 1190058 and 3200583, Chile.

Authors' addresses: Elizabeth Labrada, University of Chile, PLEIAD Lab, Computer Science Department (DCC), Beauchef 851, Santiago, Chile, elabrada@dcc.uchile.cl; Matías Toro, University of Chile, PLEIAD Lab, Computer Science Department (DCC), Beauchef 851, Santiago, Chile, mtoro@dcc.uchile.cl; Éric Tanter, University of Chile, PLEIAD Lab, Computer Science Department (DCC), Beauchef 851, Santiago, Chile, etanter@dcc.uchile.cl; Dominique Devriese, KU Leuven, imec - DistriNet, Leuven, Belgium, dominique.devriese@kuleuven.be.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2022 Copyright held by the owner/author(s).

2475-1421/2022/4-ART70

<https://doi.org/10.1145/3527314>

1 INTRODUCTION

Parametric polymorphism enables the generic definition of functions and types, providing as benefits code reusability and representation independence. System F [Girard 1972; Reynolds 1974] is the standard language to formalize this notion of parametric polymorphism. Relational parametricity stipulates that the behavior of polymorphic functions must be independent of the specific types they are instantiated with. For instance, the behavior of a function f of type $\forall X. X \rightarrow X$ should not depend on the type X it is instantiated with and, consequently, should treat the argument of type X opaquely. Hence, f [Int] 42 should never return a different value than 42.

Also, in recent years, gradual typing has become a relevant feature for programming languages because it combines the best of static and dynamic type checking. Central to gradual typing is the notion of *precision* between types, which can range from fully-precise static types to the fully-imprecise unknown type (hereafter written $?$), with partially specified types in between, such as $\text{Int} \rightarrow ?$. Among the expected properties of gradual languages [Siek et al. 2015], a particularly challenging one is the *dynamic gradual guarantee* (DGG), also called *graduality* [New and Ahmed 2018]. Informally, graduality is a monotonicity property of reduction with respect to precision: introducing imprecision in a program ought not change its behavior. For instance, the function $\lambda x : ?. x + 1$ of type $? \rightarrow \text{Int}$ should be transparently usable in place of $\lambda x : \text{Int}. x + 1$.

It turns that out that integrating parametric polymorphism and gradual typing into a language while preserving parametricity and graduality is extremely challenging [Ahmed et al. 2009, 2017; Igarashi et al. 2017; Matthews and Ahmed 2008; New et al. 2020; Toro et al. 2019]. The difficulty observed in these efforts is a strong tension between the two desirable properties, considering that functions may attempt to use gradual typing to bypass parametricity. For example, the following function of type $\forall X. X \rightarrow X$ should not be allowed to treat the value x as an integer, even when X happens to be instantiated to Int :

$$(\lambda X. \lambda x : X. ((x :: ?) + 1) :: X) [\text{Int}] 42 \quad (\text{we write } t :: T \text{ for type ascriptions})$$

To prevent this application from reducing to $((42 :: ?) + 1) :: \text{Int}$, gradual polymorphic languages have generally relied on a form of *dynamic sealing* [Matthews and Ahmed 2008]. Essentially, the function $(\lambda X. \lambda x : X. ((x :: ?) + 1) :: X)$ is not applied to type Int and value 42; instead, the language generates a fresh seal α and applies the function to α and a sealed version of the value 42, and unseals the result. This approach ensures that effectively-parametric code behaves as usual, but that the above example fails (because addition fails on sealed values):

$$\begin{aligned} (\lambda X. \lambda x : X. x) [\text{Int}] 42 &\rightarrow^* \text{unseal}_\alpha(\text{seal}_\alpha(42)) \rightarrow^* 42 \\ (\lambda X. \lambda x : X. ((x :: ?) + 1) :: X) [\text{Int}] 42 &\rightarrow^* \text{unseal}_\alpha(\boxed{(\text{seal}_\alpha(42) :: ?) + 1}) :: \alpha \rightarrow^* \text{error} \end{aligned}$$

Unfortunately, when applying a polymorphic function with an imprecise type, the decision of whether arguments should be sealed or not is not so clear-cut. Consider, for example, the functions $f_1 = \lambda X. \lambda x : ?. x :: X$ and $f_2 = \lambda X. \lambda x : ?. x :: \text{Int}$. By graduality, the two functions should behave like their more precisely-typed versions $\lambda X. \lambda x : X. x :: X$ and $\lambda X. \lambda x : \text{Int}. x :: \text{Int}$, respectively. However, this means that applying both functions to type Int and value 42 should treat their arguments differently even though they have the same parameter type. Applying f_1 [Int] 42 should seal the argument 42, while f_2 [Int] 42 should not. Most proposed gradual parametric languages decide whether to seal or not based on the type of the argument and the function being applied, i.e. they apply *type-driven* sealing. However, there is no way to make this choice *a priori* and modularly, without breaking graduality. For example, GSF [Toro et al. 2019] does not seal the argument here, breaking graduality for f_1 .

A recent proposal by New et al. [2020] side-steps the conundrum by shifting the burden of choice to programmers using *term-driven* sealing. PolyG^v requires programmers to specify whether arguments should be sealed or not, by writing for example $f_1 [X=\text{Int}]$ ($\text{seal}_X(42)$) and $f_2 [X=\text{Int}]$ 42, respectively.¹ While this strategy has produced the first parametric gradual calculus, it is important to realize that this calculus does not solve the same problem as the one tackled by other proposals like GSF, λB [Ahmed et al. 2017] or System F_G [Igarashi et al. 2017]. Gradual languages are intended to smoothly support the static-to-dynamic checking spectrum, but as noted by New et al., PolyG^v supports this only when the untyped code already contains the right sealing annotations. In other words, in PolyG^v, sealing annotations are not subject to graduality guarantees, *i.e.* $f [Int]$ 42 and $f [Int]$ $\text{seal}_X(42)$ are unrelated by precision, and therefore graduality does not relate their respective behavior.

In this paper, we revisit the original problem: gradual parametricity with type-driven sealing. Consider again the applications $f_1 [Int]$ 42 and $f_2 [Int]$ 42. Instead of making an arbitrary choice between sealing or not sealing, we propose to keep both options open, so the decision can be made when the value 42 is actually used. This novel technique, called *plausible sealing*, essentially allows our calculus to treat the applications as $f_i [Int]$ ($\text{maybeSeal}_X(42)$). The maybe-sealed value 42 embeds the fact that it may be both sealed at X and unsealed, which makes the two applications successfully reduce to 42. To study plausible sealing, we propose an intermediate gradual parametric language, Funky (F_ϵ^2), which can be used as the elaboration target of different gradual source languages; we describe one such source language, Funk (F^2), with the familiar syntax of System F.² Note that for simplicity, we formalize the approach in a setting where polymorphism is limited to instantiations with base and variable types.

The key novelty of the intermediate language F_ϵ^2 is that it introduces maybe-sealing forms, which are interpreted thanks to an innovative runtime tracking technique. Additionally, F_ϵ^2 avoids the use of dynamically-generated *global* seals. In previous calculi, a seal α can continue to exist when the type variable X for which it was created goes out of scope: $(\Lambda X. \lambda x : X. x \ :: ?)$ [Int] 42 \rightarrow^* $\text{seal}_\alpha(42)$. In fact, seals in these calculi behave as a form of symbolic cryptography, which makes it possible to embed languages with runtime sealing [Pierce and Sumii 2000; Sumii and Pierce 2004]. But at the same time, global seals have been shown to break equivalences that hold in System F [Devriese et al. 2018]. This global nature of seals is also the reason that parametricity theorems for gradual calculi so far have used formulations based on Kripke worlds containing semantic types for dynamically-allocated seals. F_ϵ^2 features *lexically-scoped sealing*, and it is the first to support a stronger formulation of parametricity where semantic types are tracked in a lexical environment, similar to traditional formulations of parametricity [Reynolds 1983]. As such, F_ϵ^2 could perhaps satisfy the ambitious criterion for gradual languages recently proposed by Jacobs et al. [2021]: fully abstract embedding of the statically-typed language into the gradually-typed language. This has been disproved by Devriese et al. [2018] for λB , but their counterexample, which essentially relies on the global nature of seals in λB and GSF, does not apply to F_ϵ^2 .³ Finally, we prove both graduality and parametricity for the intermediate language F_ϵ^2 .

The elaboration of the source language F^2 to F_ϵ^2 is in charge of introducing maybe-sealing forms when imprecise types occur in type applications. For F^2 , we establish graduality, currently subject to a restriction on type applications. Specifically, reasoning about graduality requires users to verify that the types of polymorphic functions being applied have the same shape; for instance, graduality holds between functions of types $\forall X. X \rightarrow X$ and $\forall X. ? \rightarrow ?$, but not between $\forall X. X \rightarrow X$ and

¹As the syntax suggests, type variables in PolyG^v are introduced at instantiation time, with outward scoping; this requires linear typing environments and a mechanism to limit their propagation to the current lambda abstraction [New et al. 2020].

²Funk is for **F**-unknown (F^2), and Funky is for Funk with evidence (F_ϵ^2).

³See the technical report for a proof sketch that their counterexample does not apply to F_ϵ^2 .

$\forall X.?$. Except for this technical restriction, type and term precision is standard and graduality in $F^?$ allows programmers to reason in much the same way as they would with the natural notion of term precision.

In addition to graduality, we explain the source-level parametric reasoning that $F^?$ offers. It is worth noting that parametric reasoning at the source level of a gradual language is subtle because of another point of tension between parametricity and gradual typing that was pointed out by [New et al. \[2020\]](#). Consider the two applications: $(\lambda X. \lambda x : ?.x :: X) [\text{Int}] 42$ and $(\lambda X. \lambda x : ?.x :: X) [\text{Bool}] 42$. Since the behavior of the polymorphic function $\lambda X. \lambda x : ?.x :: X$ should not depend on the type it is applied to, a strict interpretation of parametricity dictates that both applications should behave the same. At the same time, by graduality, the first application should behave equivalently to the following more precisely typed version, which reduces to 42: $(\lambda X. \lambda x : X.x :: X) [\text{Int}] 42 \rightarrow^* 42$. However, the second application is of type `Bool` and there is no reasonable way to come up with a boolean value to return. Even worse, because parametricity implies preservation of relatedness of values, successfully returning a boolean in the second application would imply a contradiction, because that boolean would have to be related to 42 in an arbitrary, caller-chosen relation, even when that relation is empty. In other words, this strict interpretation of source-level parametricity is incompatible with graduality. However, that is not the end of the story.

In $F^?$, the second application fails at runtime: the value 42 does not have the right type to be sealed at type X , so it is not maybe-sealed, and we simply report an error when it is treated as a value of type X . This means that some polymorphic $F^?$ terms may behave differently depending on the type they are applied to, as we have $(\lambda X. \lambda x : ?.x :: X) [\text{Int}] 42 \mapsto^* 42$ and $(\lambda X. \lambda x : ?.x :: X) [\text{Bool}] 42 \mapsto^* \text{error}$. It would however be incorrect to conclude that $F^?$ is not parametrically polymorphic. First, uniformity of behavior is satisfied for polymorphic functions of fully precise types,⁴ even if they internally use type applications that do (!). In these cases, the definition of parametricity coincides with the standard definition for System F—except that related terms may also simultaneously fail with a runtime type error. In other words, the differences in behavior can only occur for imprecise types (and can therefore be avoided using ascriptions to precise types). Intuitively, these differences are a consequence of $F^?$ applying plausible sealing in an attempt to infer whether the programmer intended to treat arguments (or results) as values of the quantified type X , in a maximally permissive way. However, the behavior of plausible sealing is entirely predictable based on type information available *statically* at the call site, and does not depend on runtime type information. When one takes this behavior into account, gradual parametricity in $F^?$ still implies useful free theorems. For example, for any $f : \forall X. ? \rightarrow X$, $f [\text{Bool}] \text{true}$ may diverge, fail or return the value `true`, but it can never return `false`.

Contributions. We develop a novel approach to gradual parametricity based on plausible sealing. Technically, we use lexically-scoped rather than global sealing, and a novel runtime tracking mechanism based on proof-relevant precision to account for postponing sealing decisions. This is achieved using a representation of evidence inspired by the category of spans. Focusing on the new ideas, we formally develop our approach in a simplified setting where polymorphism is restricted to instantiations with base and variable types. We prove that the proposed intermediate language $F_\epsilon^?$ satisfies both parametricity and graduality, and mechanize the two key lemmas in Agda needed to prove these properties. We illustrate the practicality of $F_\epsilon^?$ by providing a translation from the source gradual language $F^?$. For $F^?$, we establish graduality, subject to a restriction on type applications, and explain the source-level parametric reasoning it offers.

⁴Later on, we introduce a mechanism to annotate occurrences of the unknown type with the subset of type variables in scope that it might denote, and explain the impact of this feature on parametric reasoning for imprecise types.

Overview. In Section 2, we illustrate the behavior of $F_\epsilon^?$ programs by starting from their $F^?$ source counterparts, and compare to other approaches. We then formalize the core calculus $F_\epsilon^?$ (Section 3), describe its novel form of runtime tracking mechanism for plausible sealing (Section 4), and prove parametricity (Section 5) and the gradual guarantees (Section 6). We then formalize the source language $F^?$ and its elaboration to $F_\epsilon^?$. We discuss the parametric reasoning enjoyed by $F^?$, and the gradual guarantees, subject to a technical restriction on type applications (Section 7). We discuss the lifting of the technical restrictions of this work in Section 8. Section 9 discusses related work and Section 10 concludes.

Full definitions and proofs of the main results can be found in the companion technical report, provided as supplementary material. Mechanized proofs of two key technical results in Agda (Lemmas 4.6 and 6.2, marked with \checkmark) are also included as supplementary material. The implementation (<https://doi.org/10.5281/zenodo.6341550>) exhibits typing derivations, the translation from $F^?$ to $F_\epsilon^?$, and reduction traces, including all the examples mentioned in this paper and of the related literature.

2 BACKGROUND AND OVERVIEW OF $F_\epsilon^?$

This section recalls the basics of gradual typing, emphasizing the Abstracting Gradual Typing methodology [Garcia et al. 2016], which inspired this work. It also outlines the behavior of $F_\epsilon^?$ with specific source program examples in $F^?$ from the current state of the art of gradual parametricity, informally shedding light on how plausible sealing is realized and compares to other approaches.

2.1 Background on (Abstracting) Gradual Typing

Basics of gradual typing. Gradual typing smoothly supports the range from static to dynamic type checking by introducing the unknown type (here denoted $?$) and allowing types to be partially specified [Siek and Taha 2006]. For instance, $? \rightarrow \text{Int}$ is the gradual type of functions whose domain is statically unknown $?$ and whose codomain is Int . This type is said to be less *precise* (or more imprecise) than static types such as $\text{Bool} \rightarrow \text{Int}$, and more precise than both $? \rightarrow ?$ and $?$ [Siek et al. 2015]; this is noted $\text{Bool} \rightarrow \text{Int} \sqsubseteq ? \rightarrow \text{Int} \sqsubseteq ? \rightarrow ? \sqsubseteq ?$. Optimistically, a variable of type $?$ can be used at any type statically; at runtime, some mechanism ensures that a runtime type error is raised before any unsafe operation is performed. For instance, the application $(\lambda x : ?. x + 1)$ `false` is well-typed, but results in a runtime error before addition is performed.

The flexibility of gradual typing is achieved by relaxing type predicates (such as type equality, subtyping, etc.) to optimistically account for imprecision. For example, type *consistency* (denoted \sim) is the relaxation of type equality [Garcia et al. 2016; Siek and Taha 2006]. The application example above is well typed because $\text{Bool} \sim ?$ and $? \sim \text{Int}$ (but $\text{Bool} \not\sim \text{Int}!$). The dynamic semantics of a gradual source language is typically given by elaboration to a cast calculus [Siek and Taha 2006]. The elaboration inserts casts to guarantee that violations of static assumptions are detected, triggering type errors at runtime. For instance, the source term $\lambda x : ?. x + 1$ would typically be elaborated to the target term $\lambda x : ?. \langle \text{Int} \leftarrow ? \rangle x + 1$, where the cast $\langle \text{Int} \leftarrow ? \rangle$ ensures that the argument given at runtime is indeed an **Int** value, otherwise an error is raised.⁵

Abstracting gradual typing. The Abstracting Gradual Typing framework (AGT) [Garcia et al. 2016] derives the static and dynamic semantics of a gradual language starting from a static language and its type safety argument. The static semantics exploit a Galois connection between gradual types and the sets of static types they denote: predicates on gradual types are obtained by existential lifting of static predicates. For instance, consistency is the lifting of equality: two gradual types

⁵We use the blue color and sans serif fonts for source languages and the red color and bold fonts for target languages.

are consistent iff there exist two static types in their denotations (a.k.a. concretizations) that are equal. More advanced predicates and functions for gradual types, such as consistent subtyping and consistent join, can also be derived following this approach.

In AGT, the dynamic semantics is defined by reduction of gradual typing derivations augmented with *evidence* for consistent judgments. Equivalently, one can understand this approach as defining an elaboration to an *evidence-based* target language, by analogy with cast calculi. The key notion here is that of evidence, which tracks the most precise information regarding a consistent judgment at runtime. During reduction, evidences are combined, just like casts, and this combination may fail with a runtime error, whenever the resulting consistent judgment is not justified anymore. This mechanism at least ensures type safety, and can be adjusted to ensure other properties (e.g. noninterference [Toro et al. 2018], parametricity [Toro et al. 2019]).

While the concept of evidence is very general and applies to a variety of typing disciplines, Garcia et al. [2016] observe that for a language with only type consistency, evidence coincides with the middle type of threesomes [Siek and Wadler 2010]. A threesome is a three-place cast, $\langle G_2 \stackrel{G}{\Leftarrow} G_1 \rangle$, representing a downcast from the source type G_1 to the middle type G , followed by an upcast from the middle type to the target type G_2 . This representation allows for space efficiency of cast calculi: when combining two threesomes, it is sufficient to retain the outermost types and keep the *meet* \sqcap (according to the precision partial order) of the middle types. If such a meet is not defined, the combination of threesomes fails with a cast error. For instance, the combination $\langle \text{Int} \stackrel{\text{Int}}{\Leftarrow} ? \rangle \langle ? \stackrel{\text{Bool}}{\Leftarrow} \text{Bool} \rangle$ fails because $\text{Int} \sqcap \text{Bool}$ is undefined.

Likewise, an evidence ϵ for a consistency judgment, noted $\epsilon : G_1 \sim G_2$, is naturally represented by a common more precise type G such that $G \sqsubseteq G_1$ and $G \sqsubseteq G_2$; for instance $\text{Int} : \text{Int} \sim ?$. At runtime, reduction proceeds by combining evidences through *consistent transitivity* (\circ), which is, like for threesomes, the precision meet of the evidences. For example, term $\epsilon_2 (\epsilon_1 x :: \text{Int}) :: ?$ (with $\epsilon_1 = \text{Int}$) reduces to $(\epsilon_1 \circ \epsilon_2) x :: ?$, where $\epsilon_1 \circ \epsilon_2 = \text{Int} \sqcap \text{Int} = \text{Int}$.

The elaboration from the source language to the evidence-based target language simply inserts the *initial evidence* of all consistent judgments used in the gradual typing derivation of the term. For example, if $\vdash t : G$ then the source term $t :: G'$ would elaborate to $\epsilon t :: G'$, where t is the elaboration of the subterm t , and ϵ is the initial evidence between the type G and the ascribed type G' , i.e. $G \sqcap G'$. In an elimination form such as a function application, elaboration introduces an ascription to ensure that the top-level type constructor matches.

2.2 Evidence for Plausible Sealing

In this work, we adopt AGT for deriving the static semantics of $F^?$ and $F_\epsilon^?$, and define the dynamic semantics of $F^?$ by elaboration to the evidence-based target language $F_\epsilon^?$. Prior work using AGT for gradual parametricity (GSF [Toro et al. 2019]) has shown that the semantics obtained blindly with AGT only ensure type safety, but not parametricity. Ensuring parametricity requires a refined representation of evidence and consistent transitivity. In GSF, evidence is represented not as a single type, but as a pair of types (extended with type names tracked globally), in order to capture the directionality of consistent judgments, which can intuitively denote either *sealing* or *unsealing*. Consistent transitivity is refined to forbid unsound unsealing and hence enforce parametricity. In order to address the limitations discussed in the introduction, we design a novel representation of evidence in $F_\epsilon^?$, to realize plausible sealing. The rest of this section informally describes this novel representation of evidence and the achieved behavior.

Let us focus on the two terms (1) $f_1 [\text{Int}]$ 42 and (2) $f_2 [\text{Int}]$ 42 used in the introduction. As explained, these are key illustrations of the challenge of type-driven sealing: any early decision to either seal or not seal the argument would make one of these examples fail, thereby breaking

graduality. Our approach consists of capturing the different possibilities regarding sealing, and postponing the choice to seal or not to seal until a value is used; as a consequence, both programs successfully reduce to 42. This is achieved by a novel representation of evidence, which accommodates the different valid usages of an argument of unknown type, whenever the unknown type is in scope of some type variables. The first step consists of decorating the unknown type with the type variables that are in scope. So the type of both elaborated polymorphic functions in F_ϵ^2 are $\forall X. ?_X \rightarrow X$ and $\forall X. ?_X \rightarrow \text{Int}$, respectively, since X is the only type variable in the scope of the unknown type. The argument 42 is of type Int , so upon elaboration an ascription to $?_0$ is introduced—there are no type variables in scope at that point. Hence, the elaboration of both examples (where G stands for either X (1) or Int (2)) is:

$$(\epsilon_1 ((AX.\lambda x : ?_X. \epsilon \ x :: G) [\text{Int}]) :: ?_0 \rightarrow \text{Int}) (\epsilon_2 \ 42 :: ?_0) \quad (1)$$

When these polymorphic functions are instantiated at type Int , the decorations of unknown types are enriched with the instantiation information, so the lambda-abstractions both take an argument of type $?_{X:\text{Int}}$. To proceed with the beta reduction, the argument $\epsilon_2 42 :: ?_0$ is ascribed to the expected argument type of the lambda, yielding the value $v = \epsilon' 42 :: ?_{X:\text{Int}}$. This value is the $\text{maybeSeal}_X(42)$ used in the introduction. Observe that there are *two* ways in which the type of 42, Int , is consistent with $?_{X:\text{Int}}$: either because $?_{X:\text{Int}}$ stands for Int , or because it stands for X (which happens to be instantiated with Int). So it is *plausible* that the value be sealed at type X , though not mandatory. In order to account for this multiplicity of possibilities, we let ϵ' be a *set* of justifications, rather than a single justification as is standard in AGT (and in GSF). Both justifications support the same consistency judgment $\text{Int} \sim ?_{X:\text{Int}}$, so using just the meet is insufficient. Instead, we represent a justification of a consistent judgment between types G_1 and G_2 as a triple (G, c_1, c_2) , where G is the meet, and c_1 (resp. c_2) is a proof term that characterizes *how* G is more precise than G_1 (resp. G_2). Hence, precision in F_ϵ^2 is a *proof-relevant* notion, and evidences carry these proofs. In the example, the precision judgments are $\text{inj}_X : \text{Int} \sqsubseteq ?_{X:\text{Int}}$ and $\text{inj}_{\text{Int}} : \text{Int} \sqsubseteq ?_{X:\text{Int}}$, where the proof terms inj_X and inj_{Int} denote the two possible injections of imprecision. We write refl_{Int} for the proof term of $\text{Int} \sqsubseteq \text{Int}$. So we have:

$$\epsilon' = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_X), (\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}})\}$$

When v is substituted in the body, reduction proceeds by combining ϵ' with ϵ , the evidence inserted by the elaboration of the ascription in the body (Equation 1), using consistent transitivity. Importantly, in Example (1), ϵ justifies that the unknown type is consistent with X via inj_X , and when Int is substituted for X , the proof term inj_X in ϵ does not change (although it now justifies the judgment $\text{Int} \sqsubseteq ?_{X:\text{Int}}$ rather than $X \sqsubseteq ?_X$). Then reduction proceeds by checking that there is at least one justification in ϵ' that is compatible with ϵ ; otherwise an error is raised. Because such a justification exists in both examples, they both successfully reduce to 42.

In essence, we treat type precision \sqsubseteq in F_ϵ^2 not simply as a preorder, but as a category, and we construct evidence as a variant of the category of spans. Spans are the triples (G, c_1, c_2) , and evidences are sets of spans. Composition of evidence through consistent transitivity can then be defined in terms of a category-theoretic pullback operation, again generalizing the order-theoretic meet that is used in regular threesomes and AGT.

2.3 Comparing Plausible Sealing and Prior Approaches

We now outline the behavior of F_ϵ^2 , informally shedding light on how plausible sealing is realized and compares to other approaches. For the sake of simplicity and understanding, we use source F^2 programs for the comparison. Note that, in order to be well typed, source terms in F^2 need to be augmented with evidence in F_ϵ^2 , casts in λB , and seal/unseal terms in PolyG^V (possibly yielding

Table 1. Comparisons of gradual parametricity approaches.

	Source term in $F^?$	$F_\epsilon^?$	λB	System F_G	GSF	PolyG ^v
1	$(\lambda X.\lambda x:?.x :: X) \text{ [Int] } 42$	42	error	error	error	error / 42
2	$(\lambda X.\lambda x:?.x :: \text{Int}) \text{ [Int] } 42$	42	42	42	42	42 / error
3	$(\lambda X.\lambda x:?.x :: X) \text{ [Bool] } 42$	error	error	error	error	error
4	$((\lambda X.\lambda x:X.x :: ?) \text{ [Int] } 42) + 1$	43	error	error	error	error / 43
5	$(\lambda X.\lambda x:X.(x :: ?) + 1) \text{ [Int] } 3$	error	error	error	error	error
6	$(\lambda X.\lambda Y.\lambda x:?.(x, x) :: X \times Y) \text{ [Int] [Int] } 42$	$\langle 42, 42 \rangle$	error	error	error	error / error

two possible well-typed variants), in addition to superficial syntactic differences. Table 1 compares $F_\epsilon^?$ with prior approaches using a number of key examples from the literature—except Example (6)—either adapted or verbatim. Additional examples are provided in the technical report.

Examples (1) and (2) are the key examples discussed in Section 2.2. In GSF, λB and System F_G , Example (1) fails with an error, and Example (2) yields 42, because these systems eagerly choose not to seal the argument when it has the unknown type. In PolyG^v, programmers have to use explicit sealing to decide to seal or not, but this cannot be done modularly; one can obtain different behaviors accordingly. Example (3) raises a runtime error at the ascription to X , as the type variable is instantiated to `Bool` but a value of type `Int` is provided; together with Example (1), it illustrates the shallow non-parametric behavior of $F^?$ discussed in the introduction. Note that other approaches also raise an error in this example because the argument is not sealed, which implies that Example (1) fails as well. Example (4) illustrates that, contrary to other approaches that use global type names as a runtime sealing mechanism, sealing in $F_\epsilon^?$ is lexically scoped: seals cannot outlive the lexical boundary of a type abstraction. In the example, when 42 is returned by the function, it is automatically unsealed and usable as a regular integer. In PolyG^v, an explicit unseal is needed to avoid failure. Example (5) illustrates the prevention of a violation of parametricity at runtime. Example (6) illustrates yet another flexibility of plausible sealing that makes it more expressive than prior approaches: evidence as sets of spans can support multiple sealing behaviors. In this example, the argument of the function, 42, is treated as plausibly sealed to both X and Y at the same time. This example fails in GSF. In PolyG^v, programmers have to pick in advance whether to seal with X or Y , and the example fails in both cases. Observe that this program does not have a fully statically-typed counterpart, and therefore showcases an expressiveness gain of the gradual language, which compromises neither graduality nor parametricity.

These examples illustrate the flexibility afforded by plausible sealing, as a novel point in the design space of gradual parametricity.

3 THE EVIDENCE-BASED LANGUAGE $F_\epsilon^?$

Now that we have informally explained our representation of evidence and the obtained behavior, we turn to the formalization of $F_\epsilon^?$ and its properties: parametricity and graduality. This section centers on presenting the language without entering into the details of evidence: evidence and its operators are treated abstractly. We provide the full details of evidence for $F_\epsilon^?$ in Section 4. Sections 5 and 6 establish parametricity and graduality of $F_\epsilon^?$, respectively. Section 7 then studies the source language $F^?$, its elaboration to $F_\epsilon^?$, and its properties.

Syntax and static semantics. Figure 1 presents the syntax and semantics of $F_\epsilon^?$. A type G can be either a base type, a type variable, a function type, a polymorphic type, or the unknown

$$\begin{aligned}
& \mathbf{X} \in \text{TYPEVAR}, \quad \mathbf{G} \in \text{GTYPE}, \quad \epsilon \in \text{EVIDENCE}, \quad \mathbf{t} \in \text{TERM}, \quad \Delta \subset \text{TYPEVAR}, \quad \Gamma \in \text{VAR} \xrightarrow{\text{fin}} \text{GTYPE} \\
& \mathbf{F} ::= \mathbf{B} \mid \mathbf{X} \quad \mathbf{G} ::= \mathbf{B} \mid \mathbf{X} \mid \mathbf{G} \rightarrow \mathbf{G} \mid \forall \mathbf{X}. \mathbf{G} \mid ?_{\delta} \quad \delta ::= \delta, \mathbf{X} : \mathbf{F} \mid \emptyset \\
& \mathbf{u} ::= \mathbf{b} \mid \lambda \mathbf{x} : \mathbf{G}. \mathbf{t} \mid \Lambda \mathbf{X}. \mathbf{t} \quad \mathbf{v} ::= \epsilon \mathbf{u} :: \mathbf{G} \quad \mathbf{t} ::= \mathbf{v} \mid \mathbf{x} \mid \mathbf{t} \mathbf{t} \mid \mathbf{t} [\mathbf{F}] \mid \epsilon \mathbf{t} :: \mathbf{G} \quad \mathbf{s} ::= \mathbf{u} \mid \mathbf{t}
\end{aligned}$$

$\Delta; \Gamma \vdash \mathbf{s} : \mathbf{G}$

Term typing

$$\begin{array}{c}
\Delta; \Gamma \vdash \mathbf{s} : \mathbf{G}' \quad \Delta \vdash \mathbf{G} \quad \epsilon : \mathbf{G}' \sim \mathbf{G} \\
\text{Gasc} \frac{}{\Delta; \Gamma \vdash \epsilon \mathbf{s} :: \mathbf{G} : \mathbf{G}} \\
\text{GappG} \frac{\Delta; \Gamma \vdash \mathbf{t} : \forall \mathbf{X}. \mathbf{G} \quad \Delta \vdash \mathbf{F}}{\Delta; \Gamma \vdash \mathbf{t} [\mathbf{F}] : \mathbf{G} [\mathbf{F}/\mathbf{X}]}
\end{array}$$

$\mathbf{t} \longrightarrow \mathbf{t}$ or **error**

Notion of reduction

$$\begin{array}{l}
(\text{Rasc}) \quad \epsilon_2 (\epsilon_1 \mathbf{u} :: \mathbf{G}_1) :: \mathbf{G}_2 \longrightarrow \begin{cases} \epsilon_1 \mathbf{u} :: \mathbf{G}_2 & \text{if } \epsilon = \epsilon_1 \circ \epsilon_2 \\ \text{error} & \text{otherwise} \end{cases} \\
(\text{Rapp}) \quad \begin{array}{l} (\epsilon_1 (\lambda \mathbf{x} : \mathbf{G}_{11}. \mathbf{t}) :: \mathbf{G}_1 \rightarrow \mathbf{G}_2) \\ (\epsilon_2 \mathbf{u} :: \mathbf{G}_1) \end{array} \longrightarrow \begin{cases} \text{cod}(\epsilon_1) (\mathbf{t} [(\epsilon_2 \mathbf{u} :: \mathbf{G}_{11})/\mathbf{x}]) :: \mathbf{G}_2 & \text{if } \epsilon = \epsilon_2 \circ \text{dom}(\epsilon_1) \\ \text{error} & \text{otherwise} \end{cases} \\
(\text{RappG}) \quad (\epsilon (\Lambda \mathbf{X}. \mathbf{t}) :: \forall \mathbf{X}. \mathbf{G}) [\mathbf{F}] \longrightarrow (\text{schm}(\epsilon) \mathbf{t} :: \mathbf{G}) [\mathbf{F}/\mathbf{X}]
\end{array}$$

Fig. 1. F_{ϵ}^2 : Syntax, Static and Dynamic Semantics (fragment).

type.⁶ Observe that static types from System F are syntactically included in gradual types \mathbf{G} . In F_{ϵ}^2 , polymorphic types can only be instantiated with base types and type variables, called *instantiation types*, and denoted by metavariable \mathbf{F} . As mentioned in the introduction, this restriction on polymorphism simplifies the already-dense technical development while still manifesting all the subtleties of gradual parametricity identified in prior work. Another distinctive feature of F_{ϵ}^2 is that it avoids the use of a global typename store as used in all prior work on gradual parametricity thanks to the fact that the unknown type is indexed by an environment δ . This instantiation environment keeps track of the static and dynamic information related to type variables in scope: $?_{\mathbf{X}:\text{Int}}$ expresses that type variable \mathbf{X} is in scope and instantiated to Int . Uninstantiated type variables are associated with themselves $\mathbf{X} : \mathbf{X}$, which for brevity we simply write as \mathbf{X} . It is worth noting that in the type $?_{\mathbf{X}:\mathbf{X}}$, the two occurrences of \mathbf{X} play a different role: the first is merely a label, while the second is an actual occurrence of the type variable \mathbf{X} .

A term \mathbf{t} can be a value \mathbf{v} , a variable, a term application, a type application (to an instantiation type), or an ascription. Note the presence of an evidence ϵ in an ascription, to justify the fact that the underlying term is of a type consistent with the ascribed type. Values \mathbf{v} are ascribed raw values $\epsilon \mathbf{u} :: \mathbf{G}$, where ϵ justifies that the type of \mathbf{u} is consistent with \mathbf{G} . A raw value \mathbf{u} can be a base value \mathbf{b} , a function, or a type abstraction. To avoid duplication of typing rules, we use metavariable \mathbf{s} to denote both raw values \mathbf{u} and terms \mathbf{t} .

The typing judgment $\Delta; \Gamma \vdash \mathbf{s} : \mathbf{G}$ establishes that \mathbf{s} has type \mathbf{G} , under type variable environment Δ , and type environment Γ . Δ is used to track type variables in scope, and Γ to map variables to their types. Most of the type rules are standard, closely following System F. Note that rule (Gasc) is the only rule that uses the consistency relation; all other elimination rules require types to match exactly. Elaboration from the source language F^2 is in charge of introducing the necessary ascriptions to safely support the flexibility of gradual typing. Rule (GappG) is almost standard, save for the fact that it restricts instantiations to instantiation types \mathbf{F} . The type substitution operator $\mathbf{G} [\mathbf{F}/\mathbf{X}]$

⁶We omit formal definitions for pairs, which are unsurprising and found in the technical report.

is standard, except for occurrences of unknown types, for which type substitution is applied to their instantiation environments δ : $(\delta, X : F')[F/X] = \delta[F/X], X : F'[F/X]$. For instance, $?_{Y:X,X:X}[Int/X] = ?_{Y:Int,X:Int}$. Notice that type substitution on an instantiation environment δ only affects type variable occurrences, not labels.

Dynamic semantics. The dynamic semantics of F_ϵ^2 are usual for an evidence-based reduction semantics [Garcia et al. 2016; Toro et al. 2019], using reduction frames and notions of reduction. Reduction uses the consistent transitivity operator \circ to combine evidence and justify transitive judgments. If $\epsilon_1 \circ \epsilon_2$ is defined then it yields a more precise evidence, otherwise an error is raised. For example, rule (Rasc) reduces nested ascriptions, such as value $\epsilon_1 u :: G_1$ ascribed to G_2 using evidence ϵ_2 . Recall that ϵ_1 justifies that G_u , the type of u , is consistent with G_1 , noted $\epsilon_1 : G_u \sim G_1$, and likewise, $\epsilon_2 : G_1 \sim G_2$. Therefore, if $\epsilon_1 \circ \epsilon_2$ is defined, then the resulting evidence justifies the transitive judgment between G_u and G_2 , i.e. $\epsilon_1 \circ \epsilon_2 : G_u \sim G_2$. Rule (Rapp) reduces a term application substituting the argument in the body of the function. It first ascribes the argument to G_{11} , the type of x . To justify the transitive judgment of the beta reduction, it combines ϵ_2 , with $dom(\epsilon_1)$. Evidence $dom(\epsilon_1)$ and $cod(\epsilon_1)$ can be extracted from ϵ_1 by reasoning about inversion on consistency ($\epsilon_1 : G_{11} \rightarrow G_{12} \sim G_1 \rightarrow G_2$). Evidence ϵ_1 justifies that $G_{11} \rightarrow G_{12}$, the underlying type of the function, is consistent with $G_1 \rightarrow G_2$. Thus, evidence $dom(\epsilon_1)$ justifies that G_1 is consistent with G_{11} , and therefore $\epsilon \circ dom(\epsilon_1)$, if defined, justifies that the type of the raw value u is consistent with G_{11} . The output of the function is ascribed to the expected return type G_2 using the co-domain evidence $cod(\epsilon_1)$. Rule (RappG) reduces type application by substituting type F in the schema evidence $schm(\epsilon)$, in the body of the type abstraction t and in the scheme type G . By inversion on consistency, if $\epsilon : \forall X.G' \sim \forall X.G$, then $schm(\epsilon) : G' \sim G$. Substitution on evidence $\epsilon[F/X]$ is defined using substitution on types, for all type information that appears in the evidence (Section 4). Substitution on terms $t[F/X]$ is recursively defined over subterms, evidences, and types.

It is worth noting that rule (RappG) is remarkably standard unlike other gradual polymorphic calculi where dynamic type generation happens in this rule, being stored in a global store. This is made possible thanks to the use of the annotated unknown type $?_\delta$. Another point that has relevance in the reduction of type applications is that the type of the redex can contain instantiated type variables in scope. For example, term $(\epsilon_2 (\lambda x.\lambda x : ?_X.\epsilon_1 x :: X) :: \forall X.?_X \rightarrow X) [Int]$ has type $?_{X:Int} \rightarrow Int$ with X in the instantiation environment of the unknown type. To obtain a term that can be applied to an argument of type $?_0$, an external evidence to the application is necessary to justify that $?_{X:Int} \rightarrow Int$ is consistent with $?_0 \rightarrow Int$. As we saw in Section 2 and will explain in detail in Section 7, this evidence is inserted by the elaboration from F^2 to F_ϵ^2 .

Properties. As expected from any language, F_ϵ^2 is type safe (i.e. well-typed F_ϵ^2 terms do not get stuck). Thus, a well-typed program either evaluates to a value, a runtime error, or diverges. In order to prove type safety, it is necessary to have some properties about the evidence, such as the resulting evidence from consistent transitivity supports the transitive consistency judgment and type substitution over the evidence supports the substitution over the judgment (Section 4).

LEMMA 3.1 (TYPE SAFETY). *If $\vdash t : G$ then either $t \xrightarrow{*} v$ with $\vdash v : G$, $t \xrightarrow{*} \mathbf{error}$, or t diverges.*

Of course, the most interesting properties of F_ϵ^2 are parametricity and graduality. We dive into the details of these properties in Section 5 and Section 6 respectively, after giving a detailed account of evidence, including its representation, operations, and properties thereof, in particular *associativity* and *monotonicity* of consistent transitivity.

$$c ::= \text{refl}_F \mid c \rightarrow c \mid \forall X.c \mid \text{inj}_F \mid \text{inj}_{\rightarrow}(c) \mid \text{inj}_{\vee}(c) \mid \text{inj}_{?}$$

$c : G \sqsubseteq G$ **Proof-relevant precision**

$$\frac{}{\text{refl}_B : B \sqsubseteq B} \quad \frac{}{\text{refl}_X : X \sqsubseteq X} \quad \frac{\delta \sqsubseteq \delta'}{\text{inj}_{?} : ?_{\delta} \sqsubseteq ?_{\delta'}}$$

$$\frac{c : G_1 \sqsubseteq G_2 \quad c' : G'_1 \sqsubseteq G'_2}{c \rightarrow c' : G_1 \rightarrow G'_1 \sqsubseteq G_2 \rightarrow G'_2} \quad \frac{c : G_1 \sqsubseteq G_2}{\forall X.c : \forall X.G_1 \sqsubseteq \forall X.G_2} \quad \frac{X : F \in \delta \quad \delta \vdash F}{\text{inj}_X : F \sqsubseteq ?_{\delta}}$$

$$\frac{}{\text{inj}_B : B \sqsubseteq ?_{\delta}} \quad \frac{c : G \sqsubseteq ?_{\delta} \rightarrow ?_{\delta}}{\text{inj}_{\rightarrow}(c) : G \sqsubseteq ?_{\delta}} \quad \frac{c : G \sqsubseteq \forall X.?_{\delta}, X}{\text{inj}_{\vee}(c) : G \sqsubseteq ?_{\delta}}$$

$c;c = c$ **Composition of precision proof terms**

$$\text{refl}_B; \text{refl}_B = \text{refl}_B \quad \text{refl}_X; \text{refl}_X = \text{refl}_X \quad (c_1 \rightarrow c_2); (c'_1 \rightarrow c'_2) = (c_1; c'_1) \rightarrow (c_2; c'_2)$$

$$(\forall X.c); (\forall X.c') = \forall X.(c; c') \quad \text{refl}_B; \text{inj}_B = \text{inj}_B \quad \text{refl}_F; \text{inj}_X = \text{inj}_X$$

$$c_1; \text{inj}_{\rightarrow}(c_2) = \text{inj}_{\rightarrow}(c_1; c_2) \quad c; \text{inj}_{?} = c$$




Fig. 2. Proof-relevant type precision, composition, and examples.

4 EVIDENCE FOR PLAUSIBLE SEALING IN F_{ϵ}^2

We now turn to the key technical innovation that makes F_{ϵ}^2 (and by extension, F^2) able to address the dilemma presented in Section 1: plausible sealing, implemented via a novel representation of evidence based on a *proof-relevant* notion of gradual type precision. As explained in Section 2, for a consistency judgment $G_1 \sim G_2$, instead of having evidence only track a common more precise type G , evidence is a set of spans, where each span includes a common more precise type G and two proof terms that describe *how* $G \sqsubseteq G_1$ and $G \sqsubseteq G_2$ hold, respectively.

Proof-relevant precision. As mentioned in Section 2.2, there can be multiple ways of satisfying a precision relation $G \sqsubseteq G'$. To differentiate them, we extend the precision relation between types with a proof term c that expresses how G is more precise than G' .

Proof-relevant precision is presented in Figure 2. The proof relevant judgment $c : G \sqsubseteq G'$ denotes that *proof term* c justifies that G is more precise than G' . A reflexive proof term refl_F justifies that F is more precise than F . A function proof term $c \rightarrow c'$ witnesses that a function type is more precise than another function type if their domains and codomains are related; likewise for polymorphic proof terms $(\forall X.c)$. Proof term inj_X represents an injection from X into $?_{\delta}$, and witnesses that if X is associated to F in δ and F is well-formed with respect to δ , then F is more precise than $?_{\delta}$. We say that a type F is well-formed with respect to δ ($\delta \vdash F$) if F is a base type B or is a type variable X and $X : X \in \delta$. For example, if X is not yet instantiated and $X : X \in \delta$, then $\text{inj}_X : X \sqsubseteq ?_{X.X}$. If $X : \text{Int} \in \delta$, then $\text{inj}_X : \text{Int} \sqsubseteq ?_{X:\text{Int}}$. Injection inj_B witnesses that a base type B is more precise than any unknown type. Proof term sequence $\text{inj}_{\rightarrow}(c)$ justifies that function types are more precise than unknown: if c witnesses that a type G is more precise than

$?_8 \rightarrow ?_8$, then $\text{inj}_{\rightarrow}(\mathbf{c})$ justifies that \mathbf{G} is more precise than $?_8$. Similarly, $\text{inj}_{\vee}(\mathbf{c})$ witnesses that polymorphic types are more precise than unknown respectively. $\text{inj}_{?}$ justifies that an unknown type is more precise than another if the environment of the former is contained in the environment of the latter.

To support transitive judgments of precision, we define the composition of proof terms in Figure 2. A reflexive proof term combined with itself yields the same proof term. The combinations of function, and abstraction proof terms are defined inductively. The combination of a reflexive base type proof term with an injection from that type yields the latter. Similarly, the combination of a reflexive refl_F proof term with an injection from \mathbf{X} , yields just the injection from \mathbf{X} . Finally, the combination of an $\text{inj}_{?}$ from the right can always be dropped. Figure 2 illustrates graphically the composition function along some examples. If $\mathbf{c}_1 : \mathbf{G}_1 \sqsubseteq \mathbf{G}$, and $\mathbf{c} : \mathbf{G} \sqsubseteq \mathbf{G}_2$, then $\mathbf{c}_1; \mathbf{c}_2 : \mathbf{G}_1 \sqsubseteq \mathbf{G}_2$. If $\text{refl}_{\text{Int}} : \text{Int} \sqsubseteq \text{Int}$, and $\text{inj}_{\text{Int}} : \text{Int} \sqsubseteq ?_{\mathbf{X}:\text{Int}}$, then as $\text{refl}_{\text{Int}}; \text{inj}_{\text{Int}} = \text{inj}_{\text{Int}}$, then $\text{inj}_{\text{Int}} : \text{Int} \sqsubseteq ?_{\mathbf{X}:\text{Int}}$. Finally, if $\text{inj}_{\text{Int}} : \text{Int} \sqsubseteq ?_{\mathbf{X}:\text{Int}}$, and $\text{inj}_{?} : ?_{\mathbf{X}:\text{Int}} \sqsubseteq ?_{\mathbf{X}:\text{Int}; \mathbf{Y}:\text{Int}}$, then $\text{inj}_{\text{Int}} : \text{Int} \sqsubseteq ?_{\mathbf{X}:\text{Int}; \mathbf{Y}:\text{Int}}$. With these reflexivity and composition operators, gradual types and type precision proof terms can be seen as the objects and morphisms of a category, which will be useful in Section 4.

Evidence and consistent transitivity. As mentioned before, evidence is defined as a set of justifications, accounting for the multiple possibilities in which two types can be consistent. Using proof-relevant type precision, evidence ϵ is defined as a non-empty set of spans $\{\mathbf{S}, \dots\}$, where a span \mathbf{S} is a tuple $(\mathbf{G}, \mathbf{c}_1, \mathbf{c}_2)$ such that if $\epsilon : \mathbf{G}_1 \sim \mathbf{G}_2$, then \mathbf{G} is a common more precise type than \mathbf{G}_1 and \mathbf{G}_2 , and \mathbf{c}_1 and \mathbf{c}_2 justify “how”, respectively.

Definition 4.1. $\epsilon : \mathbf{G}_1 \sim \mathbf{G}_2$ iff $\forall (\mathbf{G}, \mathbf{c}_1, \mathbf{c}_2) \in \epsilon, \mathbf{c}_1 : \mathbf{G} \sqsubseteq \mathbf{G}_1 \wedge \mathbf{c}_2 : \mathbf{G} \sqsubseteq \mathbf{G}_2$.

For example, $\epsilon = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}}), (\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\mathbf{X}})\}$ justifies that $\text{Int} \sim ?_{\mathbf{X}:\text{Int}}$, because $\text{refl}_{\text{Int}} : \text{Int} \sqsubseteq \text{Int}$, $\text{inj}_{\text{Int}} : \text{Int} \sqsubseteq ?_{\mathbf{X}:\text{Int}}$, and $\text{inj}_{\mathbf{X}} : \text{Int} \sqsubseteq ?_{\mathbf{X}:\text{Int}}$. Therefore, as explained in Section 2, the term $\epsilon 42 :: ?_{\mathbf{X}:\text{Int}}$ corresponds exactly to the maybe-sealed value $\text{maybeSeal}_{\mathbf{X}}(42)$ from the introduction: the evidence holds *both* possible justifications.


The type substitution definition on evidence, and more precisely on proof terms, is fundamental for the plausible sealing mechanism to preserve parametricity. Type substitution for evidence is defined as the type substitution for each of its spans. Type substitution for a span is defined as the type substitution of its components. For example, $(\mathbf{X}, \text{refl}_{\mathbf{X}}, \text{refl}_{\mathbf{X}})[\mathbf{F}/\mathbf{X}] = (\mathbf{F}, \text{refl}_{\mathbf{F}}, \text{refl}_{\mathbf{F}})$ and $(\mathbf{X}, \text{inj}_{\mathbf{X}}, \text{refl}_{\mathbf{X}})[\mathbf{F}/\mathbf{X}] = (\mathbf{F}, \text{inj}_{\mathbf{X}}, \text{refl}_{\mathbf{F}})$. Note that $\text{inj}_{\mathbf{X}}[\mathbf{F}/\mathbf{X}] = \text{inj}_{\mathbf{X}}$ is essential to preserve sealed values; otherwise, we would forget the sealing if we apply the substitution.

To define consistent transitivity for this representation of evidence, we first define the *pullback* operator between proof terms.

LEMMA 4.2 (PULLBACK OPERATOR AND ITS UNIVERSAL PROPERTY).

There exists a partial pullback operator such that if $\mathbf{c}_1 : \mathbf{G}_1 \sqsubseteq \mathbf{G}$ and $\mathbf{c}_2 : \mathbf{G}_2 \sqsubseteq \mathbf{G}$, and $\text{pullback}(\mathbf{G}, (\mathbf{G}_1, \mathbf{c}_1), (\mathbf{G}_2, \mathbf{c}_2)) = (\mathbf{G}_3, \mathbf{c}_3, \mathbf{c}_4, \mathbf{c}_5)$, then $\mathbf{c}_3 : \mathbf{G}_3 \sqsubseteq \mathbf{G}_1$, $\mathbf{c}_4 : \mathbf{G}_3 \sqsubseteq \mathbf{G}_2$, $\mathbf{c}_3; \mathbf{c}_1 = \mathbf{c}_5$ and $\mathbf{c}_4; \mathbf{c}_2 = \mathbf{c}_5$. The pullback operator is universal in the following sense. If there exists $\mathbf{G}'_3, \mathbf{c}'_3, \mathbf{c}'_4$ and \mathbf{c}'_5 such that $\mathbf{c}'_3; \mathbf{c}_1 = \mathbf{c}'_5$ and $\mathbf{c}'_4; \mathbf{c}_2 = \mathbf{c}'_5$, then $\text{pullback}(\mathbf{G}, (\mathbf{G}_1, \mathbf{c}_1), (\mathbf{G}_2, \mathbf{c}_2)) = (\mathbf{G}_3, \mathbf{c}_3, \mathbf{c}_4, \mathbf{c}_5)$ and there exists a unique $\mathbf{c} : \mathbf{G}'_3 \sqsubseteq \mathbf{G}_3$ such that $\mathbf{c}; \mathbf{c}_3 = \mathbf{c}'_3$, $\mathbf{c}; \mathbf{c}_4 = \mathbf{c}'_4$ and $\mathbf{c}; \mathbf{c}_5 = \mathbf{c}'_5$.

Our pullback operator and its universal property are a mild adaptation of the standard definition in category theory [nLab contributors 2021a]. Figure 3 illustrates our *pullback* operator, along with two examples. The first example (second diagram) calculates $\text{pullback}(?_{\mathbf{X}:\text{Int}}, (\text{Int}, \text{inj}_{\mathbf{X}}), (\text{Int}, \text{inj}_{\mathbf{X}}))$, returning $(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}}, \text{inj}_{\mathbf{X}})$. Note that the diamond diagram commutes, obtaining $\text{inj}_{\mathbf{X}}$ both on the left and on the right. The second example (third diagram) tries to calculate






Fig. 3. Pullback and examples.




Fig. 4. Consistent transitivity for spans and example.

$\text{pullback}(\text{?X:Int}, (\text{Int}, \text{inj}_X), (\text{Int}, \text{inj}_{\text{Int}}))$, but it is undefined since there is no gradual type and proof terms such that the diagram commutes. The definition of the pullback operator is algorithmic, proceeding in most cases congruently.

The universal property in Lemma 4.2 establishes that if G'_3 (with proof terms c'_3 and c'_4) makes the pullback diagram commute, then the pullback is defined, resulting in type G_3 (with proof terms c_3 and c_4), the less precise type that makes the diagram commute. Additionally, there exists a proof term c such that $c : G_3 \sqsubseteq G'_3$ and all sub-diagrams commute. For example, suppose $c_1 = c_2 = \text{inj}_?$ and $G = G_1 = G_2 = \text{?X:Int}$. Since $G'_3 = \text{Int}$, with $c'_3 = c'_4 = \text{inj}_X$, makes the diagram commute, we know that the pullback exists. In this case, we know that $\text{pullback}(\text{?X:Int}, (\text{?X:Int}, \text{inj}_?), (\text{?X:Int}, \text{inj}_?)) = (\text{?X:Int}, \text{inj}_?, \text{inj}_?, \text{inj}_?)$. Observe that there exist two proof terms c such that $c : G_3 \sqsubseteq G'_3$, $c = \text{inj}_{\text{Int}}$ and $c = \text{inj}_X$. However, only $c = \text{inj}_X$ satisfies $c;c_3 = c'_3$, $c;c_4 = c'_4$ and $c;c_5 = c'_5$. Consistent transitivity between spans is then defined as follows:

Definition 4.3 (Consistent transitivity for spans). Let $c_1 : G_1 \sqsubseteq G$, $c'_1 : G_1 \sqsubseteq G'_1$, $c_2 : G_2 \sqsubseteq G$ and $c'_2 : G_2 \sqsubseteq G'_2$. We pose $(G_1, c'_1, c_1) \circ (G_2, c_2, c'_2) = \{(G_3, c_6, c_7) \mid \text{pullback}(G, (G_1, c_1), (G_2, c_2)) = (G_3, c_3, c_4, c_5) \wedge c_3; c'_1 = c_6 \wedge c_4; c'_2 = c_7\}$.

The definition is very close to the standard definition of composition of spans [nLab contributors 2021b], except that we are dealing with a *partial* pullback and a *set of spans* rather than a single span. Figure 4 graphically supports the definition of consistent transitivity, along with an example. First the pullback of c_1 and c_2 is computed. If the pullback is defined, then the new evidence type is computed using the common gradual type from the pullback G_3 , and new proofs that G_3 is more precise than G'_1 and G'_2 using the proof-relevant composition operator. Note that the result of consistent transitivity for spans is either a singleton set or the empty set. In the example, we have that $(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_X) : \text{Int} \sim \text{?X:Int}$ (representing a seal at type X), and $(\text{Int}, \text{inj}_X, \text{inj}_Y) : \text{?X:Int} \sim \text{?Y:Int}$ (an unseal at type X , followed by a seal at type Y). Consistent transitivity $(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_X) \circ (\text{Int}, \text{inj}_X, \text{inj}_Y)$ is computed by first computing $\text{pullback}(\text{?X:Int}, (\text{Int}, \text{inj}_X), (\text{Int}, \text{inj}_X)) = (\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}}, \text{inj}_X)$. As $\text{refl}_{\text{Int}}; \text{refl}_{\text{Int}} = \text{refl}_{\text{Int}}$, and $\text{refl}_{\text{Int}}; \text{inj}_Y = \text{inj}_Y$, the result is $(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_Y)$.

Finally, consistent transitivity between evidences is just defined as the natural lifting of consistent transitivity of spans to sets of spans.

Definition 4.4 (Consistent transitivity for evidence). Let $\epsilon_1 : G_1 \sim G$, and $\epsilon_2 : G \sim G_2$.

$$\epsilon_1 \circ \epsilon_2 ::= \begin{cases} \epsilon & \text{if } \epsilon = \{S \mid S \in S_1 \circ S_2, S_1 \in \epsilon_1, S_2 \in \epsilon_2\} \neq \emptyset \\ \mathbf{error} & \text{otherwise} \end{cases}$$

Note that if the resulting set is empty, then consistent transitivity is undefined, representing a runtime type error because plausibility of well-typedness has been refuted. Otherwise, if consistent transitivity is defined, then the obtained evidence justifies the transitive judgment.

LEMMA 4.5. *Let $\epsilon_1 : G_1 \sim G$, and $\epsilon_2 : G \sim G_2$. If $\epsilon_1 \circ \epsilon_2$ is defined, then $\epsilon_1 \circ \epsilon_2 : G_1 \sim G_2$.*

Associativity of consistent transitivity. Associativity of consistent transitivity is a key property in evidence-based semantics, used to establish type soundness as well as space efficiency optimizations [Bañados Schwerter et al. 2021; Toro and Tanter 2020]. In particular, in this work the associativity lemma is used extensively in the proof of parametricity of $F_\epsilon^?$ (Section 5).

LEMMA 4.6 (✓ ASSOCIATIVITY OF EVIDENCE COMPOSITION). $(\epsilon_1 \circ \epsilon_2) \circ \epsilon_3 = \epsilon_1 \circ (\epsilon_2 \circ \epsilon_3)$.

The proof of the associativity lemma relies on the universal property of the pullback (Lemma 4.2).

Examples of reduction. Armed with the dynamic semantics of $F_\epsilon^?$ and the concrete representation of evidence, we first illustrate the reduction of Example (1) from Section 2.3: $(\lambda X.\lambda x : ?x :: X) [\text{Int}] 42$, which reduces to 42. Its elaboration (omitting some trivial evidence for conciseness) and reduction proceed as follows:

$$\begin{aligned}
& (\epsilon_2 (\lambda X.\lambda x : ?X.\epsilon_1 x :: X [\text{Int}]) :: ? \rightarrow \text{Int}) (\epsilon_3 42 :: ?) \quad \text{where } \epsilon_1 = \{(X, \text{inj}_X, \text{refl}_X)\} \text{ and} \\
& \epsilon_2 = \{(\text{Int} \rightarrow \text{Int}, \text{inj}_X \rightarrow \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}} \rightarrow \text{refl}_{\text{Int}}), \\
& \quad (? \rightarrow \text{Int}, \text{inj}_? \rightarrow \text{refl}_{\text{Int}}, \text{inj}_? \rightarrow \text{refl}_{\text{Int}})\} \text{ and } \epsilon_3 = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}})\} \\
(\text{RappG}) \quad & \mapsto (\epsilon_2 (\lambda x : ?X:\text{Int}.\epsilon'_1 x :: \text{Int}) :: ? \rightarrow \text{Int}) (\epsilon_3 42 :: ?) \quad \text{where } \epsilon'_1 = \{(\text{Int}, \text{inj}_X, \text{refl}_{\text{Int}})\} \\
(\text{Rapp}) \quad & \mapsto \text{cod}(\epsilon_2) (\epsilon'_1 (\epsilon'_3 42 :: ?) :: \text{Int}) :: \text{Int} \quad \text{where } \epsilon'_3 = \epsilon_3 \circ \text{dom}(\epsilon_2) = \\
& \quad \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_X), (\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}})\} \\
(\text{Rasc}) \quad & \mapsto \text{cod}(\epsilon_2) (\epsilon_4 42 :: \text{Int}) :: \text{Int} \quad \text{where } \epsilon_4 = \epsilon'_3 \circ \epsilon'_1 = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\} \\
(\text{Rasc}) \quad & \mapsto \epsilon_5 42 :: \text{Int} \quad \text{where } \epsilon_5 = \epsilon_4 \circ \text{cod}(\epsilon_2) = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\}
\end{aligned}$$

We now illustrate the reduction of Example (3) from Section 2.3: $(\lambda X.\lambda x : ?x :: X) [\text{Bool}] 42$. The elaboration and reduction are as follows:

$$\begin{aligned}
& (\epsilon_2 (\lambda X.\lambda x : ?X.\epsilon_1 x :: X [\text{Bool}]) :: ? \rightarrow \text{Bool}) (\epsilon_3 42 :: ?) \quad \text{where } \epsilon_1 = \{(X, \text{inj}_X, \text{refl}_X)\} \text{ and} \\
& \epsilon_2 = \{(\text{Bool} \rightarrow \text{Bool}, \text{inj}_X \rightarrow \text{refl}_{\text{Bool}}, \text{inj}_{\text{Bool}} \rightarrow \text{refl}_{\text{Bool}}), \\
& \quad (? \rightarrow \text{Bool}, \text{inj}_? \rightarrow \text{refl}_{\text{Bool}}, \text{inj}_? \rightarrow \text{refl}_{\text{Bool}})\} \text{ and } \epsilon_3 = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}})\} \\
(\text{RappG}) \quad & \mapsto (\epsilon_2 (\lambda x : ?X:\text{Bool}.\epsilon'_1 x :: \text{Bool}) :: ? \rightarrow \text{Bool}) (\epsilon_3 42 :: ?) \quad \text{where } \epsilon'_1 = \{(\text{Bool}, \text{inj}_X, \text{refl}_{\text{Bool}})\} \\
(\text{Rapp}) \quad & \mapsto \text{cod}(\epsilon_2) (\epsilon'_1 (\epsilon'_3 42 :: ?) :: \text{Bool}) :: \text{Bool} \quad \text{where } \epsilon'_3 = \epsilon_3 \circ \text{dom}(\epsilon_2) = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}})\} \\
(\text{Rasc}) \quad & \mapsto \mathbf{error} \quad \text{because } \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}})\} \circ \{(\text{Bool}, \text{inj}_X, \text{refl}_{\text{Bool}})\} \text{ is undefined}
\end{aligned}$$

Finally, we show the reduction of Example (5) from Section 2.3, which illustrates the prevention of a violation of parametricity at runtime: $(\lambda X.\lambda x : X.(x :: ?) + 1) [\text{Int}] 3$. The elaboration and reduction are as follows:

$$\begin{aligned}
& (\epsilon_3 ((\lambda X.\lambda x : X.(\epsilon_2 (\epsilon_1 x :: ?X) :: \text{Int}) + (\epsilon_{\text{Int}} 1 :: \text{Int})) [\text{Int}]) :: \text{Int} \rightarrow \text{Int}) (\epsilon_{\text{Int}} 3 :: \text{Int}) \\
& \text{where } \epsilon_1 = \{(X, \text{refl}_X, \text{inj}_X)\}, \epsilon_2 = \{(\text{Int}, \text{inj}_{\text{Int}}, \text{refl}_{\text{Int}})\}, \epsilon_{\text{Int}} = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\} \\
& \text{and } \epsilon_3 = \{(\text{Int} \rightarrow \text{Int}, \text{refl}_{\text{Int}} \rightarrow \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}} \rightarrow \text{refl}_{\text{Int}})\} \\
(\text{RappG}) \quad & \mapsto (\epsilon_3 (\lambda x : \text{Int}.\epsilon_2 (\epsilon'_1 x :: ?X:\text{Int}) :: \text{Int}) + (\epsilon_{\text{Int}} 1 :: \text{Int})) :: \text{Int} \rightarrow \text{Int}) (\epsilon_{\text{Int}} 3 :: \text{Int}) \\
& \text{where } \epsilon'_1 = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_X)\} \\
(\text{Rapp}) \quad & \mapsto \text{cod}(\epsilon_3) ((\epsilon_2 (\epsilon'_1 (\epsilon_{\text{Int}} 3 :: \text{Int}) :: ?X:\text{Int}) :: \text{Int}) + (\epsilon_{\text{Int}} 1 :: \text{Int})) :: \text{Int} \text{ where } \epsilon_{\text{Int}} \circ \text{dom}(\epsilon_3) = \epsilon_{\text{Int}} \\
(\text{Rasc}) \quad & \mapsto \text{cod}(\epsilon_3) ((\epsilon_2 (\epsilon'_1 3 :: ?X:\text{Int}) :: \text{Int}) + (\epsilon_{\text{Int}} 1 :: \text{Int})) :: \text{Int} \text{ where } \epsilon_{\text{Int}} \circ \epsilon'_1 = \epsilon'_1 \\
(\text{Rasc}) \quad & \mapsto \mathbf{error} \quad \text{because } \epsilon'_1 \circ \epsilon_2 = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_X)\} \circ \{(\text{Int}, \text{inj}_{\text{Int}}, \text{refl}_{\text{Int}})\} \text{ is undefined}
\end{aligned}$$

5 $F_\epsilon^?$: GRADUAL PARAMETRICITY

In this section, we present parametricity for $F_\epsilon^?$. Since gradual types support non-terminating terms, we use a standard technique for establishing this result: *step-indexed* logical relations [Ahmed 2006;

$\mathcal{V}_\rho[[B]]$	$= \{(n, \mathbf{v}, \mathbf{v}) \in \text{ATOM}_\rho[B]\}$
$\mathcal{V}_\rho[[G_1 \rightarrow G_2]]$	$= \{(n, \mathbf{v}_1, \mathbf{v}_2) \in \text{ATOM}_\rho[G_1 \rightarrow G_2] \mid \forall n' \leq n, \mathbf{v}'_1, \mathbf{v}'_2. \\ \triangleright (n', \mathbf{v}'_1, \mathbf{v}'_2) \in \mathcal{V}_\rho[[G_1]] \Rightarrow (n', \mathbf{v}_1 \mathbf{v}'_1, \mathbf{v}_2 \mathbf{v}'_2) \in \mathcal{J}_\rho[[G_2]]\}$
$\mathcal{V}_\rho[[\forall X.G]]$	$= \{(n, \mathbf{v}_1, \mathbf{v}_2) \in \text{ATOM}_\rho[\forall X.G] \mid \forall \vdash \mathbf{B}_1, \vdash \mathbf{B}_2, R \in \text{REL}[\mathbf{B}_1, \mathbf{B}_2]. \\ (n, \mathbf{v}_1 [\mathbf{B}_1], \mathbf{v}_2 [\mathbf{B}_2]) \in \mathcal{J}_\rho; X \mapsto (\mathbf{B}_1, \mathbf{B}_2, R)[[G]]\}$
$\mathcal{V}_\rho[[X]]$	$= \rho.R(X)$
$\mathcal{V}_\rho[[?_\delta]]$	$= \{(n, \mathbf{v}_1, \mathbf{v}_2) \in \text{ATOM}_\rho[?_\delta] \mid \forall G_R, \epsilon, \vdash \epsilon : \delta \rightarrow G_R. \\ (n, \rho_1(\epsilon) \mathbf{v}_1 :: \rho_1(G_R), \rho_2(\epsilon) \mathbf{v}_2 :: \rho_2(G_R)) \in \mathcal{J}_\rho[[G_R]]\}$
$\mathcal{J}_\rho[[G]]$	$= \{(n, \mathbf{t}_1, \mathbf{t}_2) \in \text{ATOM}_\rho[G] \mid \forall i < n. \\ (\forall \mathbf{v}_1. \mathbf{t}_1 \mapsto^i \mathbf{v}_1 \Rightarrow \exists \mathbf{v}_2. \mathbf{t}_2 \mapsto^* \mathbf{v}_2 \wedge \triangleright^i (n, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}_\rho[[G]]) \wedge \\ (\mathbf{t}_1 \mapsto^i \mathbf{error} \Rightarrow \mathbf{t}_2 \mapsto^* \mathbf{error})\}$
$\mathcal{D}[[\emptyset]]$	$= \{(n, \emptyset)\}$
$\mathcal{D}[[\Delta, X]]$	$= \{(n, \rho[X \mapsto (\mathbf{B}_1, \mathbf{B}_2, R)]) \mid (n, \rho) \in \mathcal{D}[[\Delta]] \wedge R \in \text{REL}[\mathbf{B}_1, \mathbf{B}_2]\}$
$\mathcal{G}_\rho[[\emptyset]]$	$= \{(n, \emptyset)\}$
$\mathcal{G}_\rho[[\Gamma, \mathbf{x} : G]]$	$= \{(n, \gamma[\mathbf{x} \mapsto (\mathbf{v}_1, \mathbf{v}_2)]) \mid (n, \gamma) \in \mathcal{G}_\rho[[\Gamma]] \wedge (n, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}_\rho[[G]]\}$
$\Delta; \Gamma \vdash \mathbf{t}_1 \leq \mathbf{t}_2 : G \triangleq \Delta; \Gamma \vdash \mathbf{t}_1 : G \wedge \Delta; \Gamma \vdash \mathbf{t}_2 : G \wedge \forall n, \rho, \gamma. ((n, \rho) \in \mathcal{D}[[\Delta]] \wedge (n, \gamma) \in \mathcal{G}_\rho[[\Gamma]]) \Rightarrow \\ (n, \rho_1(\gamma_1(\mathbf{t}_1)), \rho_2(\gamma_2(\mathbf{t}_2))) \in \mathcal{J}_\rho[[G]]$	
$\Delta; \Gamma \vdash \mathbf{t}_1 \approx \mathbf{t}_2 : G \triangleq \Delta; \Gamma \vdash \mathbf{t}_1 \leq \mathbf{t}_2 : G \wedge \Delta; \Gamma \vdash \mathbf{t}_2 \leq \mathbf{t}_1 : G$	
$\text{ATOM}_\rho[G] = \{(n, \mathbf{t}_1, \mathbf{t}_2) \in \text{ATOM}[\rho.1(G), \rho.2(G)]\}$	
$\text{ATOM}[G_1, G_2] = \{(n, \mathbf{t}_1, \mathbf{t}_2) \mid \vdash \mathbf{t}_1 : G_1 \wedge \vdash \mathbf{t}_2 : G_2\} \quad \text{ATOM}^{\text{val}}[G_1, G_2] = \{(n, \mathbf{v}_1, \mathbf{v}_2) \in \text{ATOM}[G_1, G_2]\}$	
$\text{REL}[G_1, G_2] = \{R \subseteq \text{ATOM}^{\text{val}}[G_1, G_2] \mid \forall n' \leq n, \mathbf{v}_1, \mathbf{v}_2. (n, \mathbf{v}_1, \mathbf{v}_2) \in R \Rightarrow (n', \mathbf{v}_1, \mathbf{v}_2) \in R\}$	

Fig. 5. Gradual logical relation and auxiliary definitions.

[Appel and McAllester 2001]. Step indexing ensures the well-foundedness of the logical relation. We start by defining the logical relation for values and terms, and then we establish the fundamental property or parametricity. Our proposal is the first gradual polymorphic language to support a formulation of parametricity where semantic types are tracked in a lexical environment, similar to traditional formulations of parametricity [Reynolds 1983].

Logical relations. Figure 5 presents the logical relation for parametricity along with some auxiliary definitions. The relational interpretation is presented using *atoms* of the form $(n, \mathbf{t}_1, \mathbf{t}_2) \in \text{ATOM}[G_1, G_2]$, where n denotes the step index, and \mathbf{t}_1 and \mathbf{t}_2 denote closed well-typed terms at types G_1 and G_2 respectively. The logical relation is defined using two mutually-defined interpretations: one for values $\mathcal{V}_\rho[[G]]$ and one for computations $\mathcal{J}_\rho[[G]]$. Both interpretations are indexed by a type G , and an environment ρ , which maps type variables to two types G_1 and G_2 and a relation $R \in \text{REL}[G_1, G_2]$. $\text{REL}[G_1, G_2]$ defines the set of all admissible relations R such that $R \subseteq \text{ATOM}^{\text{val}}[G_1, G_2]$ (the subset of atoms where terms are values). For convenience, if $\rho = \{\mathbf{X}_i \mapsto (G_{i1}, G_{i2}, R_i)\}$, then $\rho.1, \rho.2$, and $\rho.R$ are abbreviations for $\{\mathbf{X}_i \mapsto G_{i1}\}$ and $\{\mathbf{X}_i \mapsto G_{i2}\}$, and $\{\mathbf{X}_i \mapsto R_i\}$ respectively. Thus $\rho.j(G)$ is an abbreviation for multiple substitutions $G[\mathbf{X}_i \mapsto G_{ij}]$. Finally, $\text{ATOM}_\rho[G]$ denotes the set of atoms $\text{ATOM}[\rho.1(G), \rho.2(G)]$.

Logical relation for values. The definition of related values is standard except for the unknown type. Two base values of type B are related if they are the same. Two functions are related at type $G_1 \rightarrow G_2$, if given two related arguments at type G_1 (and a strictly smaller index), the application

yields related computations at type \mathbf{G}_2 . We use notation $\triangleright^i (n, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}_\rho \llbracket \mathbf{G} \rrbracket$ as an abbreviation for $(n - i, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}_\rho \llbracket \mathbf{G} \rrbracket$, and $\triangleright (n, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}_\rho \llbracket \mathbf{G} \rrbracket$ for $\triangleright^1 (n, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}_\rho \llbracket \mathbf{G} \rrbracket$. Two type abstractions are related if their instantiations to two arbitrary base types yields related computations for any given relation between the instantiated types. Two values are related at an abstract type \mathbf{X} , if they are contained in the relation for \mathbf{X} . Two values are related at the unknown type $?_\delta$, if given any evidence ϵ that justifies that any \mathbf{G}_R is *ground* with respect to δ , notation $\vdash \epsilon : \delta \rightarrow \mathbf{G}_R$, then both values ascribed to $\rho.1(\mathbf{G}_R)$ and $\rho.2(\mathbf{G}_R)$, using $\rho.1(\epsilon)$ and $\rho.2(\epsilon)$ respectively, are related computations at type \mathbf{G}_R . This definition captures the fact that if two values are related at $?_\delta$, they are also related at some more precise ground type, either \mathbf{X} , \mathbf{B} , $?_\delta \rightarrow ?_\delta$, or $\forall \mathbf{X}.?_{\delta, \mathbf{X}}$, after removing the respective injections to the unknown type with the evidences $\rho.1(\epsilon)$ and $\rho.2(\epsilon)$. Both $\lambda\mathbf{B}$ and GSF use similar approaches for defining the logical relation for values of type unknown but are formalized differently, according to the syntax of the considered languages. Relation $\vdash \epsilon : \delta \rightarrow \mathbf{G}_R$ is defined such that \mathbf{G}_R is a *ground type* restricted to δ , and $\epsilon : ?_\delta \sim \mathbf{G}_R$:

$$\frac{\vdash \{(\mathbf{B}, \text{inj}_{\mathbf{B}}, \text{refl}_{\mathbf{B}})\} : \delta \rightarrow \mathbf{B} \quad \mathbf{X} : \mathbf{F} \in \delta \quad \delta \vdash \mathbf{F}}{\vdash \{(\mathbf{F}, \text{inj}_{\mathbf{X}}, \text{refl}_{\mathbf{F}})\} : \delta \rightarrow \mathbf{F}} \quad \frac{\vdash \{(?_\delta \rightarrow ?_\delta, \text{inj}_{\rightarrow}, \text{inj}_{? \rightarrow \text{inj}_{?}})\} : \delta \rightarrow ?_\delta \rightarrow ?_\delta}{\vdash \{(\forall \mathbf{X}.?_{\delta, \mathbf{X}}, \text{inj}_{\forall}, \forall \mathbf{X}. \text{inj}_{?})\} : \delta \rightarrow \forall \mathbf{X}.?_{\delta, \mathbf{X}}}$$

Let us illustrate the interpretation of the unknown type with examples. Consider the evidences: $\epsilon_{\text{Int}^?} = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}})\}$, $\epsilon_{\text{Int}} = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\}$ and $\epsilon_{\mathbf{X}^?} = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\mathbf{X}})\}$, where $\delta = \mathbf{X} : \mathbf{X}$, and $\delta' = \mathbf{X} : \text{Int}$.

- $(n, \epsilon_{\text{Int}^?42 :: ?_{\delta'}}, \epsilon_{\text{Int}^?42 :: ?_{\delta'}}) \in \mathcal{V}_\rho \llbracket ?_\delta \rrbracket$ because for $\epsilon = \{(\text{Int}, \text{inj}_{\text{Int}}, \text{refl}_{\text{Int}})\}$ and $\vdash \epsilon : \delta \rightarrow \text{Int}$, as $\epsilon_{\text{Int}^?} \circ \rho.i(\epsilon) = \epsilon_{\text{Int}}$, then $\triangleright (n, \epsilon_{\text{Int}^?42 :: \text{Int}}, \epsilon_{\text{Int}^?42 :: \text{Int}}) \in \mathcal{V}_\rho \llbracket \text{Int} \rrbracket$ (and for every other evidence ϵ and \mathbf{G}_R , such that $\vdash \epsilon : \delta \rightarrow \mathbf{G}_R$, consistent transitivity is not defined).
- $(n, \epsilon_{\text{Int}^?42 :: ?_{\delta'}}, \epsilon_{\text{Int}^?43 :: ?_{\delta'}}) \notin \mathcal{V}_\rho \llbracket ?_\delta \rrbracket$ because for $\vdash \{(\text{Int}, \text{inj}_{\text{Int}}, \text{refl}_{\text{Int}})\} : \delta \rightarrow \text{Int}$, as $\triangleright (n, \epsilon_{\text{Int}^?42 :: \text{Int}}, \epsilon_{\text{Int}^?43 :: \text{Int}}) \notin \mathcal{V}_\rho \llbracket \text{Int} \rrbracket$.
- Suppose $\triangleright (n, \epsilon_{\text{Int}^?42 :: \text{Int}}, \epsilon_{\text{Int}^?43 :: \text{Int}}) \in \rho.R(\mathbf{X})$. Then $(n, \epsilon_{\mathbf{X}^?42 :: ?_{\delta'}}, \epsilon_{\mathbf{X}^?43 :: ?_{\delta'}}) \in \mathcal{V}_\rho \llbracket ?_\delta \rrbracket$ because for $\epsilon = \{(\mathbf{X}, \text{inj}_{\mathbf{X}}, \text{refl}_{\mathbf{X}})\}$ and $\vdash \epsilon : \delta \rightarrow \mathbf{X}$, as $\epsilon_{\mathbf{X}^?} \circ \rho.i(\epsilon) = \epsilon_{\text{Int}}$, then $\triangleright (n, \epsilon_{\text{Int}^?42 :: \text{Int}}, \epsilon_{\text{Int}^?43 :: \text{Int}}) \in \mathcal{V}_\rho \llbracket \mathbf{X} \rrbracket = \rho.R(\mathbf{X})$ (and for every other evidence ϵ and \mathbf{G}_R , such that $\vdash \epsilon : \delta \rightarrow \mathbf{G}_R$, consistent transitivity is not defined).
- But $(n, \epsilon_{\mathbf{X}^?42 :: ?_{\delta'}}, \epsilon_{\mathbf{X}^?43 :: ?_{\delta'}}) \notin \mathcal{V}_\rho \llbracket ?_{\delta'} \rrbracket$ because for $\epsilon = \{(\text{Int}, \text{inj}_{\mathbf{X}}, \text{refl}_{\text{Int}})\}$ and $\vdash \epsilon : \delta' \rightarrow \text{Int}$, as $\epsilon_{\mathbf{X}^?} \circ \rho.i(\epsilon) = \epsilon_{\text{Int}}$, but $\triangleright (n, \epsilon_{\text{Int}^?42 :: \text{Int}}, \epsilon_{\text{Int}^?43 :: \text{Int}}) \notin \mathcal{V}_\rho \llbracket \text{Int} \rrbracket$.
- Suppose $\rho.R(\mathbf{X}) = \mathcal{V}_\rho \llbracket \text{Int} \rrbracket$, and $\epsilon_{\mathbf{m}} = \epsilon_{\text{Int}^?} \cup \epsilon_{\mathbf{X}^?}$. Then $(n, \epsilon_{\mathbf{m}42 :: ?_{\delta'}}, \epsilon_{\mathbf{m}42 :: ?_{\delta'}}) \in \mathcal{V}_\rho \llbracket ?_\delta \rrbracket$ because (1) for $\epsilon = \{(\mathbf{X}, \text{inj}_{\mathbf{X}}, \text{refl}_{\mathbf{X}})\}$ and $\vdash \epsilon : \delta \rightarrow \mathbf{X}$, as $\epsilon_{\mathbf{m}} \circ \rho.i(\epsilon) = \epsilon_{\mathbf{X}^?} \circ \rho.i(\epsilon) = \epsilon_{\text{Int}}$, then $\triangleright (n, \epsilon_{\text{Int}^?42 :: \text{Int}}, \epsilon_{\text{Int}^?42 :: \text{Int}}) \in \mathcal{V}_\rho \llbracket \mathbf{X} \rrbracket = \rho.R(\mathbf{X}) = \mathcal{V}_\rho \llbracket \text{Int} \rrbracket$; and (2) for $\epsilon = \{(\text{Int}, \text{inj}_{\text{Int}}, \text{refl}_{\text{Int}})\}$, $\vdash \epsilon : \delta \rightarrow \text{Int}$, as $\epsilon_{\mathbf{m}} \circ \rho.i(\epsilon) = \epsilon_{\text{Int}^?} \circ \rho.i(\epsilon) = \epsilon_{\text{Int}}$, then $\triangleright (n, \epsilon_{\text{Int}^?42 :: \text{Int}}, \epsilon_{\text{Int}^?42 :: \text{Int}}) \in \mathcal{V}_\rho \llbracket \text{Int} \rrbracket$.

Note that for the case of functions, type applications, and the unknown type, although the same step index is used in every recursive reasoning, the relations are well-formed as in each case a single step of reduction is always taken, lowering the index by one.

Logical relation for terms. Two computations are related at n steps if the first term yields a value in $i < n$ reduction steps, then the second must produce a value related at that type at $n - i$ steps; and if the first term fails, then the second also fails.

Logical relation for environments. The interpretation of environment Δ , specifies all type substitutions ρ , such that all type variables in Δ are mapped to a pair of base types and a relation

at those types. The interpretation of environment Γ , specifies all value substitution γ , such that every variable of type G is mapped to a pair of related values at that type.

Parametricity. The logical approximation $\Delta; \Gamma \vdash t_1 \leq t_2 : G$ states that given any step index, any environments ρ and γ that satisfy Δ and Γ respectively, the substituted terms are related computations. Similarly to ρ , for convenience if $\gamma = \{\mathbf{x} \mapsto (\mathbf{v}_{i1}, \mathbf{v}_{i2})\}$, then $\gamma_j = \{\mathbf{x} \mapsto \mathbf{v}_{ij}\}$. Finally, the fundamental property states that any well-typed term logically approximates itself.

THEOREM 5.1 (FUNDAMENTAL PROPERTY). *If $\Delta; \Gamma \vdash t : G$ then $\Delta; \Gamma \vdash t \leq t : G$.*

As standard [Ahmed 2004], the proofs of the fundamental property depends on numerous compatibility lemmas for each term constructor and the compositionality lemma, which in this work resembles compositionality for System F.

LEMMA 5.2 (COMPOSITIONALITY). *Let $\Delta \vdash F, \Delta, X \vdash G, (n, \rho) \in \mathcal{D}[\Delta]$, and $R = \nu_\rho[F]$, then $\nu_\rho[G[F/X]] = \nu_{\rho, X \mapsto (\rho_1(F), \rho_2(F), R)}[G]$.*

The most important lemma, used by almost all compatibility lemmas and compositionality is the ascription lemma, which says that the ascription of two related values yields related computations.

LEMMA 5.3 (ASCRPTION LEMMA). *If $(n, \mathbf{v}_1, \mathbf{v}_2) \in \nu_\rho[G]$, $(n, \rho) \in \mathcal{D}[\Delta]$, $\Delta \vdash G'$ and $\epsilon : G \sim G'$, then $(n, \rho_1(\epsilon)\mathbf{v}_1 :: \rho_1(G'), \rho_2(\epsilon)\mathbf{v}_2 :: \rho_2(G')) \in \mathcal{F}_\rho[G']$.*

We finalize this section by emphasizing that most of the logical relations and main lemmas are standard and defined just as in System F. In particular, and contrary to other gradual parametricity formulations, the definition of related values at polymorphic types is defined just as in System F, without the need for special notations and cases. The only unusual case is the definition of related values at the unknown type, but that is expected for any gradual language.

6 $F_\epsilon^?$: GRADUAL GUARANTEES

This section presents graduality for $F_\epsilon^?$. We start by presenting the definition of evidence and term precision. Similar to type precision, these definitions are also proof-relevant. Then, we show two of the main challenges of proving graduality: monotonicity of consistent transitivity and monotonicity of type substitution over evidence. We end this section by establishing graduality, more specifically, the static and dynamic gradual guarantees [Siek et al. 2015].

Evidence precision. To define precision between evidence we start by stating two intuitive requirements. Suppose $\epsilon_1 : G_1 \sim G'_1$ and $\epsilon_2 : G_2 \sim G'_2$. We say that ϵ_1 is more precise than ϵ_2 , if first, the types involved in the judgments are related by precision, i.e. $\mathbf{c} : G_1 \sqsubseteq G_2$ and $\mathbf{c}' : G'_1 \sqsubseteq G'_2$ for some \mathbf{c} and \mathbf{c}' ; and second, we require that for all $S_1 \in \epsilon_1$ there exists some $S_2 \in \epsilon_2$, such that S_1 is more precise than S_2 . Note that there may be some $S \in \epsilon_2$ not in precision with any element of ϵ_1 . This is intuitively expected by graduality, as it may cause ϵ_2 to “fail less” than ϵ_1 when combined with other evidence. Precision between spans $(G_{t1}, \mathbf{c}_1, \mathbf{c}'_1) \sqsubseteq (G_{t2}, \mathbf{c}_2, \mathbf{c}'_2)$ could be naively defined if there exists some proof term \mathbf{c}_t that justifies that G_{t1} is more precise than G_{t2} , i.e. $\mathbf{c}_t : G_{t1} \sqsubseteq G_{t2}$.

However, the above requirements are not sufficient to define precision among evidences. Suppose that we have $\epsilon_1 = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_X)\}$ and $\epsilon_2 = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}})\}$, where $\epsilon_1 : \text{Int} \sim ?_{X:\text{Int}}$ (Figure 6 supports this example). These two evidences meet all the above requirements: there exist $\mathbf{c} = \text{refl}_{\text{Int}}$, $\mathbf{c}' = \text{inj}_?$, and $\mathbf{c}_t = \text{refl}_{\text{Int}}$ such that $\mathbf{c} : \text{Int} \sqsubseteq \text{Int}$, $\mathbf{c}' : ?_{X:\text{Int}} \sqsubseteq ?_{X:\text{Int}}$ and $\mathbf{c}_t : \text{Int} \sqsubseteq \text{Int}$. We may be tempted to say that $\epsilon_1 \sqsubseteq \epsilon_2$ (or vice versa), but then graduality would not hold. In particular, monotonicity of consistent transitivity (MCT), a key lemma used to prove graduality, would be broken. MCT states that given two pairs of evidence related by precision $\epsilon_1 \sqsubseteq \epsilon_2$ and $\epsilon'_1 \sqsubseteq \epsilon'_2$, if $\epsilon_1 \circ \epsilon'_1$ is defined, then $\epsilon_1 \circ \epsilon'_1 \sqsubseteq \epsilon_2 \circ \epsilon'_2$. In the example, if we take evidence $\epsilon'_1 = \epsilon'_2 = \{(\text{Int},$





Fig. 6. Evidence precision auxiliary example.

$\text{inj}_{\text{Int}}, \text{refl}_{\text{Int}}\}$ that justifies $?X:\text{Int} \sim \text{Int}$ (and $\epsilon'_1 \sqsubseteq \epsilon'_2$), then $\epsilon_1 \circ \epsilon'_1 = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\}$ is defined but $\epsilon_2 \circ \epsilon'_2$ is not. We can use an analogous reasoning when assuming $\epsilon_2 \sqsubseteq \epsilon_1$, using $\epsilon'_1 = \epsilon'_2 = \{(\text{Int}, \text{inj}_X, \text{refl}_{\text{Int}})\}$. These two evidences should not be related by precision; we miss a connection between c_t and both c' and c as described next.

Definition 6.1 (Evidence Precision). If $\epsilon_1 : G_1 \sim G'_1$, $\epsilon_2 : G_2 \sim G'_2$, $c : G_1 \sqsubseteq G_2$ and $c' : G'_1 \sqsubseteq G'_2$, then we say that $[c]\epsilon_1 \sqsubseteq \epsilon_2[c']$ iff for all $(G_{t1}, c_1, c'_1) \in \epsilon_1$ there exists a $(G_{t2}, c_2, c'_2) \in \epsilon_2, c'_1, c'_2$ and c_t such that $c_t : G_{t1} \sqsubseteq G_{t2}$, $c_t; c_2 = c'_1$, $c_1; c = c'_1$, $c_t; c'_2 = c'_2$ and $c'_1; c' = c'_2$.



In addition to the requirements described above, evidence precision also requires that the combination of c_t and c_2 must commute with the combination of c_1 and c ; similarly, the combination of c_t and c'_2 must commute with the combination of c'_1 and c' . Going back to the example, ϵ_1 and ϵ_2 are not related by precision as the diagram does not commute: $\text{refl}_{\text{Int}}; \text{inj}_{\text{Int}} = \text{inj}_{\text{Int}}$ ($c_t; c'_2 = c'_2$) and $\text{inj}_X; \text{inj}_? = \text{inj}_X$ ($c'_1; c' = c'_2$), but $\text{inj}_{\text{Int}} \neq \text{inj}_X$. On the other hand, evidence $\{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\}$ is more precise than $\{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_X)\}$, because we can choose $c = c_t = \text{refl}_{\text{Int}}$, and $c' = \text{inj}_X$, such that the diagram commutes: $\text{refl}_{\text{Int}}; \text{inj}_X = \text{inj}_X$ ($c_t; c'_2 = c'_2$), and $\text{refl}_{\text{Int}}; \text{inj}_X = \text{inj}_X$ ($c'_1; c' = c'_2$).

Note that the evidence precision judgment $[c]\epsilon_1 \sqsubseteq \epsilon_2[c']$ explicitly tracks proof terms c and c' (we will refer to them as *boundary proofs*). The reason is that monotonicity of consistent transitivity only holds when adjacent boundary proofs match up.

LEMMA 6.2 (✓ MONOTONICITY OF CONSISTENT TRANSITIVITY). *If $[c]\epsilon_1 \sqsubseteq \epsilon_2[c']$, $[c']\epsilon'_1 \sqsubseteq \epsilon'_2[c'']$ and $(\epsilon_1 \circ \epsilon'_1)$ is defined, then $[c](\epsilon_1 \circ \epsilon'_1) \sqsubseteq (\epsilon_2 \circ \epsilon'_2)[c'']$.*

Let us consider the following example to understand why the “middle” boundary proof terms must match. We have that $[\text{refl}_{\text{Int}}]\{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\} \sqsubseteq \{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_X)\}[\text{inj}_X]$ and $[\text{inj}_{\text{Int}}]\{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\} \sqsubseteq \{(\text{Int}, \text{inj}_{\text{Int}}, \text{refl}_{\text{Int}})\}[\text{refl}_{\text{Int}}]$. The precision proofs do not match ($\text{inj}_X \neq \text{inj}_{\text{Int}}$), and even though $\{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\} \circ \{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\}$ is defined, $\{(\text{Int}, \text{refl}_{\text{Int}}, \text{inj}_X)\} \circ \{(\text{Int}, \text{inj}_{\text{Int}}, \text{refl}_{\text{Int}})\}$ is not. Similar to consistent transitivity, type substitution over evidence is also monotonous concerning evidence precision (two evidences related by precision remain related after type substitution).

LEMMA 6.3 (MONOTONICITY OF TYPE SUBSTITUTION). *If $[\forall X.c']\epsilon_1 \sqsubseteq \epsilon_2[\forall X.c]$, then $[c'](schm(\epsilon_1)) \sqsubseteq (schm(\epsilon_2))[c]$ and $[c'[F/X]](schm(\epsilon_1)[F/X]) \sqsubseteq (schm(\epsilon_2)[F/X])[c[F/X]]$.*

Term precision. Term precision is the natural lifting of type and evidence precision to terms, and is presented in Figure 7. Judgment $\Omega \vdash c : s_1 \sqsubseteq s_2$ denotes that term s_1 is more precise than s_2 justified by proof term c , under precision relation environment Ω . Boundary proof terms c are propagated for types, contexts, evidence, and subterms, justifying that the type of the less precise term is less precise than the type of the more precise term. Ω binds a term variable x to a type precision judgment $c : G_1 \sqsubseteq G_2$. Rule (\sqsubseteq_x) establishes that a term variable is related with itself

$\Omega \vdash c : s \sqsubseteq s$

Term precision

$$\begin{array}{c}
\text{Eq}_b \frac{}{\Omega \vdash \text{refl}_B : b \sqsubseteq b} \quad \text{Eq}_x \frac{\Omega(x) = c : G_1 \sqsubseteq G_2}{\Omega \vdash c : x \sqsubseteq x} \quad \text{Eq}_\lambda \frac{\Omega, x \mapsto c : G'_1 \sqsubseteq G'_2 \vdash c' : t_1 \sqsubseteq t_2}{\Omega \vdash c \rightarrow c' : \lambda x : G'_1.t_1 \sqsubseteq \lambda x : G'_2.t_2} \\
\text{Eq}_\Lambda \frac{\Omega \vdash c : t_1 \sqsubseteq t_2}{\Omega \vdash \forall X.c : \Lambda X.t_1 \sqsubseteq \Lambda X.t_2} \quad \text{Eq}_{\text{app}} \frac{\Omega \vdash c' \rightarrow c : t_1 \sqsubseteq t_2 \quad \Omega \vdash c' : t'_1 \sqsubseteq t'_2}{\Omega \vdash c : t_1 t'_1 \sqsubseteq t_2 t'_2} \\
\text{Eq}_{\text{appG}} \frac{\Omega \vdash \forall X.c : t_1 \sqsubseteq t_2}{\Omega \vdash c[F/X] : t_1[F] \sqsubseteq t_2[F]} \quad \text{Eq}_{\text{asc}} \frac{\Omega \vdash c' : s_1 \sqsubseteq s_2 \quad c : G_1 \sqsubseteq G_2 \quad [c']_{\epsilon_1} \sqsubseteq \epsilon_2[c]}{\Omega \vdash c : \epsilon_1 s_1 :: G_1 \sqsubseteq \epsilon_2 s_2 :: G_2}
\end{array}$$

Fig. 7. F_ϵ^2 : Term Precision (fragment).

along boundary proof term c if $x : (c : G_1 \sqsubseteq G_2) \in \Omega$, and Rule (Eq_λ) extends Ω with the judgment that justifies that the argument types are in precision. Analogous to MCT, rule (Eq_{app}) requires that the domain proof term of the function matches with the proof term of the arguments. Rule (Eq_{asc}) establishes that two ascriptions are related if the sub-terms, s_1 and s_2 , are in precision with proof term c' , the ascribed types, G_1 and G_2 , are in precision with the proof c , and evidences are in precision with the boundary proof terms c' and c .

Gradual guarantees. Armed with the definition of term precision, we now establish the graduality of F_ϵ^2 with the gradual guarantees [Siek et al. 2015].

THEOREM 6.4. *Suppose $\vdash t_1 : G_1$ and $\vdash c : t_1 \sqsubseteq t_2$. Then,*

- $\vdash t_2 : G_2$ and $c : G_1 \sqsubseteq G_2$.
- $t_1 \xrightarrow{*} v_1$ implies $t_2 \xrightarrow{*} v_2$ and $\vdash c : v_1 \sqsubseteq v_2$.
- t_1 diverges implies t_2 diverges.

The only peculiarity of this result compared to others in the literature is that the type and term precision judgments are proof-relevant. The static part of graduality (the static gradual guarantee) ensures that if t_1 with type G_1 is more precise than t_2 , justified by proof c , then t_2 has a less precise type G_2 justified by c . The dynamic part of graduality (the dynamic gradual guarantee) establishes that if the more precise term reduces to a value, then the less precise term also does, resulting in values in precision with the same type proof term c . The key lemmas to prove graduality, are MCT (Lemma 6.2), and monotonicity of type substitution (evidence precision is monotonous with respect to type substitution) (Lemma 6.3).

7 THE GRADUAL SOURCE LANGUAGE $F^?$

Having formalized the key technical innovation of this work, plausible sealing, and established both graduality and parametricity for the intermediate language F_ϵ^2 , we now turn to the source language $F^?$. This section presents the static semantics of $F^?$ and its translation to F_ϵ^2 . The static semantics of $F^?$ is derived systematically by applying AGT to System F_1 , which is a variation of System F where type instantiations are restricted to instantiation types (*i.e.* base types and type variables). The novel translation to F_ϵ^2 plays a crucial role since it is in charge of statically generating the maybe-sealing evidence for type applications. We study the gradual guarantees for $F^?$ and the resulting source-level parametric reasoning.

$F^?$: Statics. In order to apply AGT to obtain the static semantics (*i.e.* lifting functions and predicates), we use explicit type equalities and partial type functions in the typing rules of System F_1 . These partial functions also allow capturing elimination forms in a single rule that accounts for both precise and imprecise type information [Garcia et al. 2016]. For instance, functions *dom* and *cod* extract the

$F ::= B \mid X \quad G ::= F \mid G \rightarrow G \mid \forall X.G \mid ?_\delta \quad \delta ::= \delta, X : X \mid \emptyset \quad t ::= b \mid \lambda x : G.t \mid \Lambda X.t \mid x \mid t t \mid t [F] \mid t :: G$

$$\boxed{\Delta; \Gamma \vdash t : G} \text{ Term typing}$$

$$\text{Gasc} \frac{\Delta; \Gamma \vdash t : G' \quad \Delta \vdash G \quad G' \sim G}{\Delta; \Gamma \vdash t :: G : G} \quad \text{GappG} \frac{\Delta; \Gamma \vdash t : G \quad \Delta \vdash F}{\Delta; \Gamma \vdash t [F] : \text{inst}^\#(G, F)}$$

Fig. 8. $F^?$: Syntax and Static Semantics (fragment).

domain and codomain types, and *inst* instantiates a polymorphic type using a substitution function $T[T'/X]$, which replaces T' for X in T . For example, the typing rule for type applications is:

$$(\text{TappT}) \frac{\Delta; \Gamma \vdash t : T \quad \Delta \vdash F}{\Delta; \Gamma \vdash t [F] : \text{inst}(T, F)}$$

Figure 8 presents the syntax of source gradual types G . Source gradual types are syntactically contained in the gradual types of $F_\varepsilon^?$, and for simplicity throughout this section, we write G as the $F_\varepsilon^?$ counterpart of G . Source gradual types restricts the scope of unknown types to not-instantiated variables only. To represent this, we index unknown types with the δ meta-variable (included in δ); and as every type variable in $?_\delta$ is not instantiated (i.e. of the form $X : X$), for simplicity we just write δ as a set of type variables. One final note is that $?$ (without a scope δ , as used in previous sections) is syntactic sugar for $?_\Delta$, where Δ is the set of all variables in scope at that point. A straightforward and simple translation can insert these annotations before typing.

We give meaning to gradual types G through the concretization function $\gamma(\cdot)$ (omitted for space). The meaning of the unknown type $?_\delta$ is the set of all well-formed static types with respect to δ (i.e. $\gamma(?_\delta) = \{T \mid \delta \vdash T\}$). The concretization function helps us define precision ($G_1 \sqsubseteq G_2$ if and only if $\gamma(G_1) \subseteq \gamma(G_2)$) and consistency ($G_1 \sim G_2$ if and only if there exists T_1 and T_2 such that $T_1 = T_2$, $T_1 \in \gamma(G_1)$ and $T_2 \in \gamma(G_2)$). Precision and consistency resemble their $F_\varepsilon^?$ counterpart and can also be inductively defined. For instance, $X \sim ?_X$, but $X \not\sim ?_Y$ (for $X \neq Y$). However, precision in $F^?$ is no longer proof relevant: $F^?$ contains only unknown types $?_\delta$ with uninstantiated type variables, so that the precision relation from $F_\varepsilon^?$ (which had the structure of a category) reduces to a proof-irrelevant order relation in $F^?$.

Figure 8 presents a fragment of the term typing rules for $F^?$, which are obtained by replacing type predicates and functions with their corresponding liftings. The lifting is straightforward and uses the corresponding abstraction function $\alpha(\cdot)$ of $\gamma(\cdot)$, forming a Galois connection. For example, rule (Gasc) uses type consistency instead of type equality, and rule (GappG) uses the lifting of the function *inst*, defined for polymorphic types and the unknown type (i.e. $\text{inst}^\#(\forall X.G, G') = G[G'/X] \setminus X$, $\text{inst}^\#(?_\delta, G') = ?_\delta$ and undefined for other cases). Note that $\text{inst}^\#$ uses the scope removal function $G \setminus X$, which removes X from the scopes of unknown types in G : $?_{\delta_1, X, \delta_2} \setminus X = ?_{\delta_1, \delta_2}$. For instance, $(X \rightarrow ?_{X, Y})[Int/X] \setminus X = Int \rightarrow ?_Y$.

$F^?$: Elaboration to $F_\varepsilon^?$. The dynamic semantics of a $F^?$ program is given by a type-directed translation to $F_\varepsilon^?$. The rules are mostly standard save for the elaboration rule for type application. Judgment $\Delta; \Gamma \vdash t : G \rightsquigarrow t'$ expresses that term t is elaborated to t' , under type variable environment Δ , and type environment Γ . The elaboration rules use the function $\text{initEv}(G_1, G_2)$, which stands for the *initial evidence* between G_1 and G_2 . It computes the least precise evidence that justifies consistency between the types, and is defined as follows:

$$\begin{aligned}
\text{instEv}(\mathbf{B}, \mathbf{X}, \mathbf{F}) &= \text{reflEv}(\mathbf{B}) & \text{instEv}(\mathbf{Y}, \mathbf{X}, \mathbf{F}) &= \text{reflEv}(\mathbf{Y}) \text{ if } \mathbf{X} \neq \mathbf{Y} \\
\text{instEv}(\mathbf{X}, \mathbf{X}, \mathbf{F}) &= \text{reflEv}(\mathbf{F}) & \text{instEv}(\mathbf{G}_1 \rightarrow \mathbf{G}_2, \mathbf{X}, \mathbf{F}) &= \text{instEv}(\mathbf{G}_1, \mathbf{X}, \mathbf{F}) \rightarrow \text{instEv}(\mathbf{G}_2, \mathbf{X}, \mathbf{F}) \\
\text{instEv}(\delta, \mathbf{X}, \mathbf{F}) &= \text{reflEv}(\delta) \text{ if } \mathbf{X} \notin \delta & \text{instEv}(\forall \mathbf{Y}. \mathbf{G}, \mathbf{X}, \mathbf{F}) &= \forall \mathbf{Y}. \text{instEv}(\mathbf{G}, \mathbf{X}, \mathbf{F}) \\
\text{instEv}(\delta_1, \mathbf{X}, \delta_2, \mathbf{X}, \mathbf{F}) &= \{(\delta_1, \delta_2, \text{inj}_?, \text{inj}_?), (\mathbf{F}, \text{inj}_\mathbf{X}, \text{inj}_\mathbf{F})\}
\end{aligned}$$

Fig. 9. Instantiation evidence function.

Definition 7.1 (Initial Evidence). If $\mathbf{G}_1 \sim \mathbf{G}_2$ then $\mathbf{G} = \mathbf{G}_1 \sqcap \mathbf{G}_2$ and $\text{initEv}(\mathbf{G}_1, \mathbf{G}_2) = \{(\mathbf{G}, \text{initPT}(\mathbf{G}, \mathbf{G}_1), \text{initPT}(\mathbf{G}, \mathbf{G}_2))\}$.

This evidence consists of a single span, where the first component is the *meet* (greatest lower bound with respect to precision) $\mathbf{G}_1 \sqcap \mathbf{G}_2$ between \mathbf{G}_1 and \mathbf{G}_2 , and the second and third components are the *initial proof terms* between the meet and \mathbf{G}_1 and \mathbf{G}_2 , respectively. The meet $\mathbf{G}_1 \sqcap \mathbf{G}_2$ is a partial function and corresponds formally to $\alpha(\gamma(\mathbf{G}_1) \cap \gamma(\mathbf{G}_2))$. As for the definition of precision and congruence, we also define an inductive definition for the meet. Note that from AGT, $\mathbf{G} \sim \mathbf{G}'$ holds if $\gamma(\mathbf{G}_1) \cap \gamma(\mathbf{G}_2)$ is not empty, then if $\mathbf{G} \sim \mathbf{G}'$ then $\mathbf{G} \sqcap \mathbf{G}'$ will always be defined. The initial proof term between two types in precision is computed using the $\text{initPT}(\mathbf{G}, \mathbf{G}')$ function such that $\text{initPT}(\mathbf{G}, \mathbf{G}') : \mathbf{G} \sqsubseteq \mathbf{G}'$. It is important to note that the initial proof term between two types is unique since the type variables within the unknown type scope are not instantiated. Its definition is unsurprising and can be derived from the type precision judgment from Figure 2. For example, $\text{initPT}(\text{Int}, ?_X) = \text{inj}_{\text{Int}}$ and $\text{initPT}(X, ?_X) = \text{inj}_X$.

As $F_\epsilon^?$ requires all values to be ascribed, the elaboration rules ascribe base values, functions and type abstractions to their own type using the *reflexive evidence* operator $\text{reflEv}(\mathbf{G}) \triangleq \text{initEv}(\mathbf{G}, \mathbf{G})$. For instance, term 42 is elaborated to $\text{reflEv}(\text{Int}) \ 42 :: \text{Int}$, where $\text{reflEv}(\text{Int}) = \{(\text{Int}, \text{refl}_{\text{Int}}, \text{refl}_{\text{Int}})\}$. The elaboration process also inserts ascriptions to equate types in elimination forms. In particular, the translation of a term application (Eapp) ascribes the function term t_1 to a function type that matches [Cimini and Siek 2016] with its own type ($\mathbf{G}_1 \rightarrow \mathbf{G}_{11} \rightarrow \mathbf{G}_{12}$): its own type if t_1 has a function type; otherwise $\delta \rightarrow \delta$ (if the type is δ). Also, the argument term t_2 is ascribed to the argument type of the ascribed function.

$$\text{Eapp} \frac{\Delta; \Gamma \vdash t_1 : \mathbf{G}_1 \rightsquigarrow t'_1 \quad \Delta; \Gamma \vdash t_2 : \mathbf{G}_2 \rightsquigarrow t'_2 \quad \mathbf{G}_1 \rightarrow \mathbf{G}_{11} \rightarrow \mathbf{G}_{12} \quad \epsilon_1 = \text{initEv}(\mathbf{G}_1, \mathbf{G}_{11} \rightarrow \mathbf{G}_{12}) \quad \epsilon_2 = \text{initEv}(\mathbf{G}_2, \mathbf{G}_{11})}{\Delta; \Gamma \vdash t_1 \ t_2 : \mathbf{G}_{12} \rightsquigarrow (\epsilon_1 t'_1 :: \mathbf{G}_{11} \rightarrow \mathbf{G}_{12}) (\epsilon_2 t'_2 :: \mathbf{G}_{11})}$$

Rule (EappG) elaborates type applications, and is responsible of inserting “maybe-seal” evidence.

$$\text{EappG} \frac{\Delta; \Gamma \vdash t : \mathbf{G} \rightsquigarrow t' \quad \Delta \vdash \mathbf{F} \quad \mathbf{G} \rightarrow \forall \mathbf{X}. \mathbf{G}' \quad \epsilon_1 = \text{initEv}(\mathbf{G}, \forall \mathbf{X}. \mathbf{G}') \quad \epsilon_2 = \text{instEv}(\mathbf{G}', \mathbf{X}, \mathbf{F})}{\Delta; \Gamma \vdash t [\mathbf{F}] : \mathbf{G}'[\mathbf{F}/\mathbf{X}] \setminus \mathbf{X} \rightsquigarrow \epsilon_2((\epsilon_1 t' :: \forall \mathbf{X}. \mathbf{G}') [\mathbf{F}]) :: \mathbf{G}'[\mathbf{F}/\mathbf{X}] \setminus \mathbf{X}}$$

The elaborated type application is ascribed to the instantiated scheme type $\mathbf{G}'[\mathbf{F}/\mathbf{X}] \setminus \mathbf{X}$ (removing \mathbf{X} from the environments of unknown types), using a special evidence that justifies that $\mathbf{G}'[\mathbf{F}/\mathbf{X}]$ is consistent with $\mathbf{G}'[\mathbf{F}/\mathbf{X}] \setminus \mathbf{X}$. This evidence is computed using the *instantiation evidence* function instEv defined in Figure 9. Function $\text{instEv}(\mathbf{G}, \mathbf{X}, \mathbf{F}) : \mathbf{G}[\mathbf{F}/\mathbf{X}] \sim \mathbf{G}'[\mathbf{F}/\mathbf{X}] \setminus \mathbf{X}$ is defined (inductively) almost as the reflexive evidence operator, save for the case when type \mathbf{G} is δ and \mathbf{X} is in scope δ . Then, instEv generates an evidence which consist of two spans: one span that “seals to \mathbf{X} ” and another one that does not. More in detail, the first span $(\mathbf{F}, \text{inj}_\mathbf{X}, \text{inj}_\mathbf{F})$ represents that the unknown type should behave polymorphically in \mathbf{X} . On the contrary, the second span $(\delta_1, \delta_2, \text{inj}_?, \text{inj}_?)$ does not acknowledge the existence of variable \mathbf{X} . For example, we have that $(?_X \rightarrow \mathbf{X})[\text{Int}/\mathbf{X}] =$

$G \leq G$

Shape-restricted type precision

$$\begin{array}{c} \leq_B \frac{}{B \leq B} \quad \leq_X \frac{}{X \leq X} \quad \leq_{\rightarrow} \frac{G_1 \leq G_2 \quad G'_1 \leq G'_2}{G_1 \rightarrow G'_1 \leq G_2 \rightarrow G'_2} \quad \leq_{\forall} \frac{G_1 \leq G_2}{\forall X. G_1 \leq \forall X. G_2} \\ \leq_{B?} \frac{}{B \leq ?_B} \quad \leq_{X?} \frac{X \in \delta}{X \leq ?_B} \quad \leq_{?} \frac{\delta \subseteq \delta'}{?_B \leq ?_{B'}} \end{array}$$

$\Omega \vdash t : G \sqsubseteq t : G$

Term precision

$$\begin{array}{c} \sqsubseteq_{\text{asc}} \frac{\Omega \vdash t_1 : G'_1 \sqsubseteq t_2 : G'_2 \quad G_1 \sqsubseteq G_2}{\Omega \vdash t_1 :: G_1 : G_1 \sqsubseteq t_2 :: G_2 : G_2} \quad \sqsubseteq_{\text{app}G} \frac{\Omega \vdash t_1 : G_1 \sqsubseteq t_2 : G_2 \quad G_1 \leq G_2 \quad G_1 \rightarrow \forall X. G'_1 \quad G_2 \rightarrow \forall X. G'_2}{\Omega \vdash t_1 [F] : G'_1[F/X] \setminus X \sqsubseteq t_2 [F] : G'_2[F/X] \setminus X} \end{array}$$

Fig. 10. F^2 : Shape-restricted type precision and term precision (fragment).

$?_{X:\text{Int}} \rightarrow \text{Int}$ and $(?_X \rightarrow X)[\text{Int}/X] \setminus X = ? \rightarrow \text{Int}$. Therefore $\text{instEv}(?_X \rightarrow X, X, \text{Int}) : ?_{X:\text{Int}} \rightarrow \text{Int} \sim ? \rightarrow \text{Int}$, where $\text{instEv}(?_X \rightarrow X, X, \text{Int}) = \{(\text{Int} \rightarrow \text{Int}, \text{inj}_X \rightarrow \text{refl}_{\text{Int}}, \text{inj}_{\text{Int}} \rightarrow \text{refl}_{\text{Int}}), (? \rightarrow \text{Int}, \text{inj}_? \rightarrow \text{refl}_{\text{Int}}, \text{inj}_? \rightarrow \text{refl}_{\text{Int}})\}$. Note that this evidence makes it impossible that a sealed value leaks out of a polymorphic function application. The scope of a type variable (label) is limited to the type application in which it appears. For example, consider the type application $t [\text{Int}]$ in F^2 , where $\vdash t : \forall X. ?_X \rightarrow ?_X$. This term is elaborated to $\text{instEv}(?_X \rightarrow ?_X, X, \text{Int}) (t [\text{Int}]) :: ? \rightarrow ?$, where $t [\text{Int}]$ has type $?_{X:\text{Int}} \rightarrow ?_{X:\text{Int}}$; the generated instEv evidence is used to coerce this type to $? \rightarrow ?$. The label X may appear in unknown types that are used inside $t [\text{Int}]$, but the instEv evidence casts $t [\text{Int}]$ to a type not mentioning the label X , effectively restricting the scope of X to inside the term $t [\text{Int}]$. Even when the resulting term computes further and $t [\text{Int}]$ is reduced and combined with values from the context (for example, in a function application), the instEv evidence protects the scope of X , preventing it from interfering with possible other occurrences of the same name X introduced by other type applications. Because of this, there is no need for alpha-renaming.

It is important to clarify one important limitation: the instantiation evidence $\text{instEv}(G, X, F)$ is not general enough when $G = ?_{\delta_1, X, \delta_2}$. With the current definition, $?_{\delta_1, X, \delta_2}$ intuitively only represents something of type X or other well-formed static type with respect to δ_1, δ_2 . This means that at runtime, if $?_{\delta_1, X, \delta_2}$ is used in a consistent judgment with a function such as $X \rightarrow X$, the program could fail. For instance, the program $(\lambda X. \lambda x : X.x) :: \forall X. ?_X [\text{Int}] 1$ generates the instantiation evidence $\text{instEv}(?_X, X, \text{Int}) = \{(? , \text{inj}_?, \text{inj}_?), (\text{Int}, \text{inj}_X, \text{inj}_{\text{Int}})\}$, which will not seal argument 1 , making this program fail at runtime. To fix the program, and generate appropriate sealing, the type of the ascription had to be changed as follows: $(\lambda X. \lambda x : X.x) :: \forall X. ?_X \rightarrow ?_X [\text{Int}] 1$. The spans generated now include $(\text{Int} \rightarrow \text{Int}, \text{inj}_X \rightarrow \text{inj}_X, \text{inj}_{\text{Int}} \rightarrow \text{inj}_{\text{Int}})$ which makes the program run without errors. However, statically there is no way to know a priori the exact shape of the evidence needed when imprecise information is involved. Consequently, a more general mechanism for the generation of the instantiation evidence is needed (see Section 8).

Finally, we prove that the elaboration preserves typing:

THEOREM 7.2 (ELABORATION PRESERVES TYPING). *If $\Delta; \Gamma \vdash t : G$, then $\Delta; \Gamma \vdash t : G \rightsquigarrow t'$ and $\Delta; \Gamma \vdash t' : G$.*

Source-level graduality. Under the natural notion of type precision (Figure 8), some F^2 terms related by precision elaborate to F^2_ε terms that are *not* related by precision. Consider program $(\lambda X. \lambda x : X.x) :: \forall X. ?_X \rightarrow ?_X [\text{Int}] 1$ more precise than $(\lambda X. \lambda x : X.x) :: \forall X. ?_X [\text{Int}] 1$ (note that

$\forall X. ?_X \rightarrow ?_X \sqsubseteq \forall X. ?_X$). The first program elaborates to a program that reduces correctly, but the second to a program that fails. As explained in the previous section, this is because $\text{instEv}(?_X, X, \text{Int})$ does not generate evidence that contains function spans that seal the argument. This does not mean that there is no source-level graduality in F^2 at all; as first explored by Igarashi et al. [2017], the fact that the gradual guarantees are stated relative to a notion of precision means that we may be able to characterize source-level graduality via a restricted notion of precision.

To characterize the F^2 programs for which we can guarantee graduality, it is enough to restrict term precision *only* for type applications, enforcing that for such expressions, type precision be restricted to types of the same shape. Notice how rule ($\sqsubseteq_{\text{appG}}$) in Figure 10 uses *shape-restricted* type precision \sqsubseteq in its premise, while other rules, such as rule (\sqsubseteq_{asc}), use the natural type precision \sqsubseteq . Shape-restricted type precision \sqsubseteq is defined similarly to \sqsubseteq , but in the case of polymorphic and function types, the type constructors have to match. For example, $\forall X. X \rightarrow X \sqsubseteq \forall X. ?_X \rightarrow ?_X$ but $\forall X. (X \rightarrow X) \rightarrow X \not\sqsubseteq \forall X. ?_X \rightarrow ?_X$ and $\forall X. X \rightarrow X \not\sqsubseteq \forall X. ?_X$. This means that if $G_1 \sqsubseteq G_2$, then all sealing spans included in applying instEv to G_1 will be included in the application of instEv to G_2 . The other cases of term precision are derived just as the natural lifting of type precision to terms.

With this notion of term precision, two source terms related by precision elaborate to F^2_ϵ terms that are also related by precision:

LEMMA 7.3. *If $\vdash t_1 : G_1 \sqsubseteq t_2 : G_2$, $\vdash t_1 : G_1 \rightsquigarrow t_1$ and $\vdash t_2 : G_2 \rightsquigarrow t_2$, then $\vdash \text{initPT}(G_1, G_2) : t_1 \sqsubseteq t_2$.*

Note that precision in F^2_ϵ is proof-relevant, therefore we have to provide a proof term that justifies “how” two terms are related. We do that by using the initial proof term function between the types of the related terms.

The dynamic semantics of a F^2 term are given by first elaborating the term to F^2_ϵ and then reducing the F^2_ϵ term. Hence, for establishing the gradual guarantees in F^2 , we write $t \Downarrow v$ if $\vdash t : G \rightsquigarrow t$ and $t \mapsto^* v$. Similarly, we write $t \Uparrow$ if the elaboration of t diverges. Then, using Lemmas 7.2, 7.3 and 6.4 we can prove the gradual guarantees for F^2 :

THEOREM 7.4 (GRADUAL GUARANTEES). *Suppose $\vdash t_1 : G_1 \sqsubseteq t_2 : G_2$ and $\vdash t_1 : G_1$.*

- (1) $\vdash t_2 : G_2$ and $G_1 \sqsubseteq G_2$.
- (2) If $t_1 \Downarrow v_1$, then $t_2 \Downarrow v_2$ and $\vdash \text{initPT}(G_1, G_2) : v_1 \sqsubseteq v_2$.
If $t_1 \Uparrow$ then $t_2 \Uparrow$.

Source-level parametric reasoning. As a first form of parametric reasoning for F^2 , the elaborations of well-typed F^2 terms produces F^2_ϵ terms that are also well typed (by Theorem 7.2), and hence related to themselves (by Theorem 5.1):

COROLLARY 7.5. *If $\Delta; \Gamma \vdash t : G \rightsquigarrow t$ then $\Delta; \Gamma \vdash t \leq t : G$.*

This lemma is powerful, but it is not immediately clear what it means for concrete example terms in F^2 . We make this clearer as follows: a type abstraction f of type $\forall X. G$ applied to related types, produces related terms whenever X does not occur in the scopes of unknown types in G (a condition written $G \setminus X = G$ below):

LEMMA 7.6. $\forall n, \rho$

$$\frac{\Delta; \Gamma \vdash f : \forall X. G \quad \forall B_1, B_2, R \in \text{REL}[B_1, B_2] \quad ((n, \rho) \in \mathcal{D}[\Delta] \wedge (n, \gamma) \in \mathcal{C}_\rho[\Gamma]) \quad G \setminus X = G \quad \Delta; \Gamma \vdash f [B_i] : G [B_i/X] \setminus X \rightsquigarrow t_i}{(n, \rho_1(\gamma_1(t_1)), \rho_2(\gamma_2(t_2))) \in \mathcal{J}_{\rho, X \mapsto (B_1, B_2, R)}[\mathbb{G}]}$$

As a direct consequence of Lemma 7.6, every F^2 program ascribed to a static type behaves parametrically, even if it internally uses the unknown type. Since the logical relation for F^2_ϵ coincides

almost exactly with traditional formulations for System F, this means we get the same reasoning about such applications than in System F itself, save for the possibility that two related terms both raise a runtime error. From Lemma 7.6, we can also derive free theorems involving imprecise types. For instance, given a function of type $\forall X. ?_0 \rightarrow X$, then the application of the function either fails or diverges.

LEMMA 7.7. *If $\vdash f : \forall X. ?_0 \rightarrow X$, $\vdash v : B$ and $t = f [B] v$, then $t \Downarrow \mathbf{error}$, or $t \Uparrow$.*

This behavior is expected because $?_0$ can only stand for a type that does *not* involve X . Intuitively, this means that this gradual polymorphic function type denotes types such as $\forall X. \mathbf{Int} \rightarrow X$ and $\forall X. \mathbf{Bool} \rightarrow X$, but not $\forall X. X \rightarrow X$. Therefore, the function cannot create a value of type X out of thin air, and the argument v cannot possibly be sealed as a value of type X , so the function necessarily fails if it tries to return any value.

Lemma 7.6 does not apply to polymorphic functions whose type mentions unknown types with the quantified variable in scope. To understand why such types require more nuance, remember the example function $f = \Lambda X. \lambda x : ?_X :: X$ discussed in the introduction. When applied to type \mathbf{Int} and value 42, this function produces the value 42, but it throws a runtime type error when applied to \mathbf{Bool} and value 42. In such examples, $F^?$ inserts plausible sealing evidence, in an effort to guess whether the programmer intended 42 to be treated as a value of type X or not, in a maximally permissive way. However, this does not mean we do not get any form of parametricity for such examples, but rather, we need to keep in mind the intuitive effect of plausible sealing. In other words, $F^?$ supports more parametric reasoning than just what Lemma 7.6 expresses. Particularly, when a type variable is in the scope of an unknown type within a function type, we can also derive some free theorems using Corollary 7.5. For instance, if a function f has type $\forall X. ?_X \rightarrow X$, then by parametricity we can deduce that f behaves either as the identity function or fails or diverges:

LEMMA 7.8. *If $\vdash f : \forall X. ?_X \rightarrow X$, $\vdash v : B$ and $t = f [B] v$, then $t \Downarrow v$ with $\vdash v : B \rightsquigarrow v$, or $t \Downarrow \mathbf{error}$, or $t \Uparrow$.*

Contrary to GSF [Toro et al. 2019], in $F^?$ this result holds just by looking at the type of f , without the need to unfold its definition. Intuitively, this lemma takes into account that $F^?$ applies plausible sealing to the argument v , so f might return it as the result of type X . The function f can also diverge or fail, but parametricity for $F^?$ still implies that f cannot return any value other than v .

8 LIMITATIONS AND PERSPECTIVES

The technical development of plausible sealing in this article suffers from two technical limitations. The first is that we only formalize a simplified form of polymorphism with instantiation types. The second is that graduality for $F^?$ is restricted for type applications, since two type applications are only related when the polymorphic types have the same shape (Section 7).

Both limitations manifest in the definition of instantiation evidence for the unknown type (Figure 9): $\text{instEv}(?_{\delta_1; X; \delta_2}, X, F) = \{(?_{\delta_1; \delta_2}, \text{inj}_?, \text{inj}_?), (F, \text{inj}_X, \text{inj}_F)\}$. Recall that the role of this instantiation evidence is to cast, for example, a function application of type $?_{X; \mathbf{Int}} \rightarrow \mathbf{Int}$ to type $?_0 \rightarrow \mathbf{Int}$. The restricted form of polymorphism is apparent because instEv 's third argument F is restricted to an instantiation type (a base type or a type variable), and we simply use $\text{inj}_F : F \sqsubseteq ?_{\delta_1; \delta_2}$ to inject F into type $?_{\delta_1; \delta_2}$. Generalizing to full polymorphism would require replacing F and inj_F in the above definition with an arbitrary type G and a proof term $c : G \sqsubseteq ?_{\delta_1; \delta_2}$. This requires to extend the syntax of δ to allow for any type. An initial exploration suggests this is largely unproblematic and in fact, our Agda proofs of consistent transitivity associativity and monotonicity already support such a richer syntax of δ . A problem turns up when one of the other types in δ_1 or δ_2 mentions $?_{\delta'}$ with $X : X \in \delta'$. In that case, it appears we additionally need a proof term that

expresses precision between $?_{\delta}$ and $?_{\delta'}$ when δ is not just a subset of δ' but some of the types in δ are themselves strictly more precise than corresponding types in δ' .

Allowing a type abstraction to be instantiated at any type G would require saving the information of the instantiated type in the proof term inj_X . This information could then be refined through composition. Therefore, we should transform inj_X to the proof term sequence $\text{inj}_X(c)$, where $\text{inj}_X(c) : G' \sqsubseteq ?_{\delta}$, $X : G \in \delta$, $c : G' \sqsubseteq G$ and $\delta \vdash G'$. In addition, the composition of proof terms for this case would slightly change: $c_1; \text{inj}_X(c_2) = \text{inj}_X(c_1; c_2)$. Note that applying the type abstraction to any type G' with rule RappG would not change; the instantiated type G' would continue to be substituted in the body of the type abstraction and the evidence.

The other main limitation of $F^?$ is caused by the right-hand-side of the above definition. Problematically, it only mentions two cases: a value of type $?_{\delta_1; \delta_2}$ will be converted into type $?_{\delta_1; X; \delta_2}$ either (1) by not sealing at all and simply extending the scope of the unknown type, or (2) by sealing, converting a value of type F into a value sealed at type X in $?_{\delta_1; X; \delta_2}$. What is missing is a recursive case that would treat, for example, a value of type $F \rightarrow F$ as a value of type $X \rightarrow X$ and recursively seal it accordingly. In fact, this could be accommodated easily by extending the right-hand-side with an additional case: $(F \rightarrow F, \text{inj}_{\rightarrow}(\text{inj}_X \rightarrow \text{inj}_X), \text{inj}_{\rightarrow}(\text{inj}_F \rightarrow \text{inj}_F))$. We conjecture that a solution is to introduce a syntax of recursive evidence that would allow us to define $\text{instEv}(?_{\delta_1; X; \delta_2}, X, F)$ as:

$\mu \epsilon. \text{inj}_{\rightarrow}(\epsilon \rightarrow \epsilon) \uplus \text{inj}_X(\epsilon \times \epsilon) \uplus \text{inj}_V(\forall Y. \epsilon) \uplus \{(B, \text{inj}_B, \text{inj}_B)\} \uplus \{(F, \text{inj}_X, \text{inj}_F)\}$. In this notation, we construct an evidence $\text{inj}_{\rightarrow}(\epsilon \rightarrow \epsilon) : ?_{\delta} \sim ?_{\delta'}$ from $\epsilon : ?_{\delta} \sim ?_{\delta'}$ by combining two spans (G_1, c_1, c_2) and (G_2, c_3, c_4) from ϵ into the span $(G_1 \rightarrow G_2, \text{inj}_{\rightarrow}(c_1 \rightarrow c_3), \text{inj}_{\rightarrow}(c_2 \rightarrow c_4))$ and similarly for inj_X and inj_V . We leave the definition of the operational behavior of such recursive evidence and the proofs of its properties to future work.

9 DISCUSSION AND RELATED WORK

Gradual parametricity has been intensively studied [Ahmed et al. 2009, 2017; Igarashi et al. 2017; Matthews and Ahmed 2008; New et al. 2020; Toro et al. 2019; Xie et al. 2018]. We have already discussed in detail related work, emphasizing the most recent proposals. Recall that Section 2.3 compares existing languages via examples (additional examples are included in the technical report).

Sealing for Parametricity. Dynamic sealing, originally proposed by Morris [1973] to dynamically enforce type abstraction, has been widely used to guarantee parametricity in gradual languages. The notion of dynamic sealing combined with global runtime type name generation has driven the dynamic semantics of polymorphic gradual languages such as λB [Ahmed et al. 2017], CSA [Xie et al. 2018], GSF [Toro et al. 2019] and PolyG^V [New et al. 2020]. Type names are dynamically generated in each type application and are kept in a global store, making the dynamic semantics and the definitions and proof of parametricity less standard and more complex. $F_e^?$ avoids using type names generation and, therefore, a global store thanks to the fact that the unknown type is decorated by an environment and the sealing/unsealing mechanism is generated statically. It is worth noting that, unlike other developed gradual polymorphic languages, PolyG^V also includes explicit seal and unseal terms in its syntax. In this sense, we can say that $F_e^?$ also includes in its syntax explicit forms of sealing and unsealing, since for a program with an imprecise type to behave in a parametric way, it is necessary to introduce the evidence of sealing and unsealing statically. The key novelty of $F_e^?$ is to support evidence with *multiple* sealing justifications, which makes it possible to avoid to eagerly choose a sealing strategy when interacting with the unknown type.

Strict Precision. Igarashi et al. [2017] first proposed using a non-standard notion of precision in System F_G to address some problems with the dynamic gradual guarantee when the unknown type is allowed to stand for a type variable. Consequently, in System F_G the unknown type is not

consistent with any type variable. Here, the source language F^2 also restricts precision, but in a much less drastic manner: type precision is the standard precision, and only term precision is restricted, in order to only relate type applications to types of the same shape. We explain in the previous section how this restriction could be lifted. Finally, F_ε^2 has no such restriction, and satisfies the gradual guarantees with respect to the standard notion of precision, for both types and terms.

Proof Terms for Type Precision. PolyG^v , inspired by previous work [New and Ahmed 2018], adds a proof term to the type precision relation as a technical intermediate representation for the translation from PolyG^v to PolyC^v , a cast calculus that gives meaning to PolyG^v programs. It is important to note that proof terms in the type precision relation are canonical, *i.e.* there is at most one proof term that proves any given type precision judgment. Likewise, F_ε^2 language indexes the type precision relation with proof terms, but contrary to PolyG^v , proof terms are relevant, *i.e.* there can be multiple proof terms for the same precision judgment. Also, as a difference with PolyG^v , term precision in F_ε^2 is indexed by a relevant proof term.

Family of Unknown Types. Devriese et al. [2018] proposed to decorate the unknown type with the set of type variables in scope, as we do here, thus limiting the expressiveness of the unknown type by forming different families. This proposal aims to potentially reestablish the fully abstract embedding property of System F into λB . Devriese et al. [2018] and more recently Jacobs et al. [2021] proposed a new criterion for gradual typing named the fully abstract embedding (FAE) property: the embedding from the static to the gradual language should be fully abstract in order to preserve the semantics properties of the static languages. We conjecture that this criterion holds for the language F^2 , being a gradual version of System F that preserves its main semantic property (*i.e.* parametricity), although we leave a proof of FAE as future work.

Implicit Polymorphism. Several polymorphic gradual languages have explored implicit polymorphism present in languages such as Haskell. Xie et al. [2018] developed a gradual source language with implicit polymorphism, where the runtime semantics are given by compilation to λB . λB and System F_G , in turn, are languages with explicit polymorphism that accommodate some form of implicit polymorphism. F^2 and F_ε^2 , as well as PolyG^v and GSF, only support explicit polymorphism; exploring implicit polymorphism for F^2 is an interesting venue for future work, in order to enhance interoperability between typed and untyped code.

Evidence Representation. Evidence has been used in different scenarios, varying its representation according to the semantic properties to be preserved in the gradual language. For instance, Lehmann and Tanter [2017] develop a gradual language with refinement types, allowing smooth evolution and interoperability between simple types and logically refined types. In this case, the evidence for consistent subtyping is represented by a triple, where the first component accounts for the logical environment, and the second and third are types. Toro et al. [2018] develop a gradual language with security types and references, indexing types with gradual security labels. Driven by noninterference, types in evidence are indexed with *intervals* of security labels, representing (bounded) ranges of possible static types. Likewise, F_ε^2 represents evidence in a novel way. First, it enriches evidence with proof terms relevant that we call *span* and then generalizes the evidence to a set of *spans*, building evidence with the expressiveness to ensure both graduality and parametricity. This theory may be applicable in other complex settings as well. We conjecture, and leave as future work, that by representing the evidence of instantiation recursively when imprecise types occur, we can lift the restriction on the term precision relation, thus correctly running all programs that by graduality must end in a value. In addition, it would be interesting to explore if there is a way to systematically derive proof-relevant consistency with AGT.

Performance. Gradual parametricity is a very challenging topic at the theoretical level, with all current efforts trying to figure out how to achieve a good design backed by a strong metatheory. This work likewise focuses on the theory of gradual parametricity, contributing a novel approach

and technique. We leave the study of the performance and efficiency of a practical implementation as an open question to be addressed. Nevertheless, it is worth mentioning that the proposed language design satisfies a relevant criterion for space efficiency [Herman et al. 2010], namely associativity of evidence composition, which is known to allow for space efficiency in evidence-based semantics [Bañados Schwerter et al. 2021; Toro and Tanter 2020]. Whether the algorithmic definition of consistent transitivity can be efficiently implemented depends on whether evidence can be represented in memory in a form that uses space efficiently and allows an efficient implementation of evidence composition. All these are open research questions.

10 CONCLUSION

Previous work on gradual parametricity has had to compromise on important design goals like graduality [Toro et al. 2019] or type-driven sealing [New et al. 2020]. Rather than accepting these compromises, this paper attempts to revisit accepted wisdom like the use of globally scoped sealing and contribute new ideas like plausible sealing and the set-of-spans representation of evidence for proof-relevant precision. Although the results presented here still have some restrictions, they open a new path towards the goal of reconciling parametricity, graduality and type-driven sealing. Additionally, some of our novel techniques are potentially reusable in other settings. Finally, our use of lexically scoped sealing invalidates the counterexample of fully-abstract embedding put forth by Devriese et al. [2018], and thereby offers new hope of constructing a gradual language that satisfies the ambitious goal of embedding System F fully abstractly [Jacobs et al. 2021].

ACKNOWLEDGMENTS

This work was partially supported by the United States Air Force Office of Scientific Research under award number FA9550-21-1-0054, by the Flemish Research Programme Cybersecurity, and by ANID FONDECYT projects 1190058 and 3200583, Chile.

REFERENCES

- Amal Ahmed. 2004. *Semantics of Types for Mutable State*. Ph.D. Dissertation. Princeton University.
- Amal Ahmed. 2006. Step-Indexed Syntactic Logical Relations for Recursive and Quantified Types. In *Proceedings of the 15th European Symposium on Programming Languages and Systems (ESOP 2006) (Lecture Notes in Computer Science, Vol. 3924)*, Peter Sestoft (Ed.). Springer-Verlag, Vienna, Austria, 69–83.
- Amal Ahmed, Robert Bruce Findler, Jacob Matthews, and Philip Wadler. 2009. Blame for All. In *Workshop on Script to Program Evolution (STOP)*. Genova, Italy.
- Amal Ahmed, Dustin Jamner, Jeremy G. Siek, and Philip Wadler. 2017. Theorems for Free for Free: Parametricity, with and Without Types. See [ICFP 2017 2017], 39:1–39:28.
- Andrew W. Appel and David McAllester. 2001. An Indexed Model of Recursive Types for Foundational Proof-carrying Code. *ACM Transactions on Programming Languages and Systems* 23, 5 (Sept. 2001), 657–683.
- Felipe Bañados Schwerter, Alison M. Clark, and Jafery. 2021. Abstracting Gradual Typing Moving Forward: Precise and Space-Efficient. See [POPL 2021 2021], 61:1–61:28.
- Rastislav Bodík and Rupak Majumdar (Eds.). 2016. *Proceedings of the 43rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2016)*. ACM Press, St Petersburg, FL, USA.
- Matteo Cimini and Jeremy Siek. 2016. The gradualizer: a methodology and algorithm for generating gradual type systems. See [Bodík and Majumdar 2016], 443–455.
- Dominique Devriese, Marco Patrignani, and Frank Piessens. 2018. Parametricity versus the universal type. *Proceedings of the ACM on Programming Languages* 2, POPL (Jan. 2018), 38:1–38:23.
- Ronald Garcia, Alison M. Clark, and Éric Tanter. 2016. Abstracting Gradual Typing. See [Bodík and Majumdar 2016], 429–442. See erratum: <https://www.cs.ubc.ca/~rxg/agt-erratum.pdf>.
- Jean-Yves Girard. 1972. *Interprétation Fonctionnelle et Élimination des Coupures de l'Arithmétique d'Ordre Supérieur*. Ph.D. Dissertation. Université de Paris VII, Paris, France.
- David Herman, Aaron Tomb, and Cormac Flanagan. 2010. Space-efficient gradual typing. *Higher-Order and Symbolic Computation* 23, 2 (June 2010), 167–189.
- ICFP 2017 2017.

- Yuu Igarashi, Taro Sekiyama, and Atsushi Igarashi. 2017. On Polymorphic Gradual Typing. See [ICFP 2017 2017], 40:1–40:29.
- Koen Jacobs, Amin Timany, and Dominique Devriese. 2021. Fully abstract from static to gradual. See [POPL 2021 2021], 7:1–7:30.
- Nico Lehmann and Éric Tanter. 2017. Gradual Refinement Types. In *Proceedings of the 44th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2017)*. ACM Press, Paris, France, 775–788.
- Jacob Matthews and Amal Ahmed. 2008. Parametric Polymorphism Through Run-Time Sealing, or, Theorems for Low, Low Prices!. In *Proceedings of the 17th European Symposium on Programming Languages and Systems (ESOP 2008) (Lecture Notes in Computer Science, Vol. 4960)*, Sophia Drossopoulou (Ed.). Springer-Verlag, Budapest, Hungary, 16–31.
- James H. Morris. 1973. Protection in Programming Languages. *Commun. ACM* 16, 1 (Jan. 1973), 15–21.
- Max S. New and Amal Ahmed. 2018. Graduality from Embedding-Projection Pairs. *Proceedings of the ACM on Programming Languages* 2, ICFP (Sept. 2018), 73:1–73:30.
- Max S. New, Dustin Janner, and Amal Ahmed. 2020. Graduality and Parametricity: Together Again for the First Time. *Proceedings of the ACM on Programming Languages* 4, POPL (Jan. 2020), 46:1–46:32.
- nLab contributors. 2021a. pullback. <https://ncatlab.org/nlab/show/pullback>
- nLab contributors. 2021b. span. <https://ncatlab.org/nlab/show/span>
- Benjamin Pierce and Eijiro Sumii. 2000. Relating Cryptography and Polymorphism. Manuscript. POPL 2021 2021.
- John C. Reynolds. 1974. Towards a Theory of Type Structure. In *Proceedings of the Programming Symposium (Lecture Notes in Computer Science, Vol. 19)*. Springer-Verlag, 408–423.
- John C. Reynolds. 1983. Types, abstraction, and parametric polymorphism. In *Information Processing 83*, R. E. A. Mason (Ed.). Elsevier, 513–523.
- Jeremy Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. In *Proceedings of the Scheme and Functional Programming Workshop*. 81–92.
- Jeremy Siek and Philip Wadler. 2010. Threesomes, with and without blame. In *Proceedings of the 37th annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2010)*. ACM Press, Madrid, Spain, 365–376.
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015. Refined Criteria for Gradual Typing. In *1st Summit on Advances in Programming Languages (SNAPL 2015) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 32)*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Asilomar, California, USA, 274–293.
- Eijiro Sumii and Benjamin C. Pierce. 2004. A Bisimulation for Dynamic Sealing. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2005)*. ACM Press, Venice, Italy, 161–172.
- Matías Toro, Ronald Garcia, and Éric Tanter. 2018. Type-Driven Gradual Security with References. *ACM Transactions on Programming Languages and Systems* 40, 4 (Nov. 2018), 16:1–16:55.
- Matías Toro, Elizabeth Labrada, and Éric Tanter. 2019. Gradual Parametricity, Revisited. *Proceedings of the ACM on Programming Languages* 3, POPL (Jan. 2019), 17:1–17:30.
- Matías Toro and Éric Tanter. 2020. Abstracting Gradual References. *Science of Computer Programming* 197 (Oct. 2020), 1–65.
- Ningning Xie, Xuan Bi, and Bruno C. d. S. Oliveira. 2018. Consistent Subtyping for All. In *Proceedings of the 27th European Symposium on Programming Languages and Systems (ESOP 2018) (Lecture Notes in Computer Science, Vol. 10801)*, Amal Ahmed (Ed.). Springer-Verlag, Thessaloniki, Greece, 3–30.