



A Gradual Probabilistic Lambda Calculus*

WENJIA YE, The University of Hong Kong, China

MATÍAS TORO, PLEIAD Lab, Computer Science Department (DCC), University of Chile, Chile

FEDERICO OLMEDO, PLEIAD Lab, Computer Science Department (DCC), University of Chile, Chile

Probabilistic programming languages have recently gained a lot of attention, in particular due to their applications in domains such as machine learning and differential privacy. To establish invariants of interest, many such languages include some form of static checking in the form of type systems. However, adopting such a type discipline can be cumbersome or overly conservative.

Gradual typing addresses this problem by supporting a smooth transition between static and dynamic checking, and has been successfully applied for languages with different constructs and type abstractions. Nevertheless, its benefits have never been explored in the context of probabilistic languages.

In this work, we present and formalize GPLC, a gradual source probabilistic lambda calculus. GPLC includes a binary probabilistic choice operator and allows programmers to gradually introduce/remove static type –and probability– annotations. The static semantics of GPLC heavily relies on the notion of probabilistic couplings, as required for defining several relations, such as consistency, precision, and consistent transitivity. The dynamic semantics of GPLC is given via elaboration to the target language TPLC, which features a distribution-based semantics interpreting programs as probability distributions over final values. Regarding the language metatheory, we establish that TPLC –and therefore also GPLC– is *type safe* and satisfies two of the so-called *refined criteria* for gradual languages, namely, that it is a conservative extension of a fully static variant and that it satisfies the gradual guarantee, behaving smoothly with respect to type precision.

CCS Concepts: • **Theory of computation** → **Semantics and reasoning; Type structures; Operational semantics; Program reasoning; Probabilistic computation.**

Additional Key Words and Phrases: Type Systems, Gradual Typing, Probabilistic Lambda Calculus

ACM Reference Format:

Wenjia Ye, Matías Toro, and Federico Olmedo. 2023. A Gradual Probabilistic Lambda Calculus. *Proc. ACM Program. Lang.* 7, OOPSLA1, Article 84 (April 2023), 30 pages. <https://doi.org/10.1145/3586036>

1 INTRODUCTION

In a nutshell, *probabilistic programming languages* are traditional programming languages that, on top of their regular constructs, offer the possibility of sampling values from probability distributions [Gordon et al. 2014; van de Meent et al. 2018]. They find applications in a wealth of different areas, ranging from more traditional application domains such randomized algorithms [Motwani and Raghavan 1995] and cryptography [Goldwasser and Micali 1984] to more novel application domains such as differential privacy [Dwork and Roth 2014] and machine learnings [Claret et al. 2013;

*This work has been partially sponsored by Hong Kong Research Grants Council projects number 17209520 and 17209821, and ANID FONDECYT project 3200583, Chile.

Authors' addresses: Wenjia Ye, The University of Hong Kong, Hong Kong, China, yewenjia@connect.hku.hk; Matías Toro, PLEIAD Lab, Computer Science Department (DCC), University of Chile, Beauchef 851, Santiago, Chile, mtoro@dcc.uchile.cl; Federico Olmedo, PLEIAD Lab, Computer Science Department (DCC), University of Chile, Beauchef 851, Santiago, Chile, folmedo@dcc.uchile.cl.



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

© 2023 Copyright held by the owner/author(s).

2475-1421/2023/4-ART84

<https://doi.org/10.1145/3586036>

Ghahramani 2015]. These latter have led to a remarkable resurgence of probabilistic programming in the past years, with the development of a growing number of new probabilistic programming systems [Goodman et al. 2008; Goodman and Stuhlmüller 2014; Kiselyov 2016; Le et al. 2017; Pfeffer 2010; Tran et al. 2017].

To establish certain invariants of interest, programming languages traditionally incorporate some form of *typing*, backed up by a type checking phase. Depending on the moment in which type checking occurs, it is classified either as *static*—when taking place during compilation—, or as *dynamic*—when it takes place during runtime—, each of them having their own strengths and weaknesses. Concretely, programming languages with static typing allows detecting errors (i.e. invariant violations) at an early stage, but are not flexible enough for rapid prototyping. On the other hand, programming languages with dynamic typing accommodate better to changes, but present slower runtimes.

Gradual typing [Siek and Taha 2006] represents an effective alternative for integrating the benefits of static and dynamic typing at the same time, by allowing a smooth transition all along the spectrum. To do so, it introduces *imprecise* (a.k.a. *gradual*) types, which represent types possibly partially known at compile time. Imprecise types can range from fully precise static types (such as $\text{Real} \rightarrow \text{Bool}$), to the fully unknown (or imprecise) type, written $?$, with partially precise types (such as $\text{Real} \rightarrow ?$) in-between. At compile time, a gradual language typechecks programs optimistically, based on the notion of type consistency, (e.g., accepting the application of a function expecting an argument of type $\text{Real} \rightarrow ?$ and receiving an argument of type $? \rightarrow \text{Bool}$.) while the *runtime* is responsible for detecting (and reporting) any violation of such assumptions (e.g., if the received argument happens to have concrete type $\text{Bool} \rightarrow \text{Bool}$).

Gradual typing has been successfully applied to programming languages with diverse constructs and typing disciplines. Some relevant features include first-class classes [Takikawa et al. 2012], mutable references [Herman et al. 2010; Siek and Taha 2006; Siek et al. 2015c; Toro and Tanter 2020], effects as primitives [Bañados Schwerter et al. 2016], tagged and untagged unions [Toro and Tanter 2017], labeling operations (for reasoning about information flow) [Azevedo de Amorim et al. 2020; Disney and Flanagan 2011; Fennell and Thiemann 2013; Toro et al. 2018], and algebraic data types [Malewski et al. 2021]. However, it is an open question whether the benefits of gradual typing carry over to *probabilistic* programming languages.

In this work, we give a positive answer to this question by, on the one hand, designing, to the best of our knowledge, the first gradual probabilistic language and, on the other hand, establishing a set of metatheoretic results, natural to all gradual languages.

First, we present SPLC, a probabilistic λ -calculus that extends ordinary λ -calculus with a (binary) probabilistic choice operator and acts as the static end of our gradual language. It features a big-step semantics relating programs to the probability distribution of final values and to better accommodate the derivation of the gradual variant, its type system presents some distinguished features such as the presence of ascriptions, partial functions *dom* and *cod* over types, and explicit type equality. Furthermore, equality over types is semantic (instead of syntactic).

Second, we introduce GPLC, our source gradual language, whose derivation from SPLC is justified using the Abstracting Gradual Typing (AGT) methodology, a systematic approach for deriving gradual languages based on abstract interpretation [Garcia et al. 2016]. For the so-derived notions of type consistency and type and term precision, we also provide alternative—more amenable to automation—characterizations, based on the notion of probabilistic couplings [Deng and Du 2011]. In effect, probabilistic couplings are a fundamental ingredient, behind all our technical development.

Notably, GPLC allows unknown probabilities not only at the type level, but also at the term level (in probabilistic choices). This yields an increased expressivity and flexibility—characteristic of all gradual languages— and also the opportunity of leveraging the language for program refinement.

Third, we define the dynamic semantics of our gradual language by translating GPLC into the target gradual language TPLC. The runtime semantics of TPLC incorporates the required evidence to confirm or discard the optimistic assumptions made by GPLC type system. In turn, this requires adapting gradual types to encode unknown probabilities through symbolic variables, which are constrained by well-formedness conditions.

Finally, to formally validate our language design, we establish three fundamental properties of GPLC. First, we prove that it is a conservative extension of SPLC. Second, we show that it satisfies type safety. Lastly, we show that it behaves smoothly with respect to precision, a property known as the *gradual guarantee* [Siek et al. 2015b].

Altogether, this provides the first steps into the theoretical foundations of gradual probabilistic programming, and serves as starting point for developing gradual variants of more specialized, domain specific, probabilistic languages, e.g., as used for differential privacy [Reed and Pierce 2010].

Paper Organization. The rest of the paper is organized as follows. Section 2 discusses the motivation and some key design decisions and challenges behind our probabilistic gradual language. Section 3 presents the probabilistic lambda calculus (SPLC) acting as the static end of the gradualization. Section 4 develops the source gradual language (GPLC) and Section 5 the target language (TPLC), together with the metatheory. Section 6 overviews the related work and Section 7 concludes. Full definitions and proofs of the main results can be found in the supplementary material.

2 OVERVIEW

We next discuss the motivation behind a *gradual* probabilistic language through a concrete use case and summarize some key aspects and challenges behind the design of our gradual probabilistic language.

2.1 A Gradual Probabilistic Language: Why?

Assume we must develop a web application for a company, in particular, the login endpoint. To authenticate a user, we must verify that the user remains active in the company, information that is provided by an external web service (exposed by a foreign library). As usual, we support both production and development modes, where in development mode we replace the external web service with a local function, conveniently defined for developing and testing purposes.

Under these requirements, we quickly prototype the following (untyped) program, written in a SCALA-like language:

```
1 def isActive(user) = if (prod) externalCall(user) else localCall(user)
2 def login(user, pass) = if (isActive(user)) /* test password */ else false
```

Probabilistic modeling. The company now requires that the login endpoint have a 95% uptime (availability). However, after some research, we learn that the external web service `externalCall` has only a 90% uptime, returning a 503 Service Unavailable error when down. Written in such an untyped language, the above program is unable to capture this uptime information, let alone detect the impossibility to comply with the login requirements.

As a first step to address this problem, we can adopt a typed language that includes *distribution types*. Loosely speaking, distribution types represent probability distributions of “simpler” types, and can crisply model uptime information. For instance, the return type of `externalCall` shall now be represented by $\{\text{Bool}_{\frac{90}{100}}, \text{Error503}_{\frac{10}{100}}\}$, and the return type of `login` by $\{\text{Bool}_{\frac{95}{100}}, \text{Error503}_{\frac{5}{100}}\}$. Furthermore, we can implement a `localCall` function compatible with the uptime requirement of the login endpoint, as follows:

```
3 def localCall(user: String):  $\{\text{Bool}_{\frac{95}{100}}, \text{Error503}_{\frac{5}{100}}\}$  = true  $\oplus_{\frac{95}{100}}$  new Error503()
```

A program of the form $m \oplus_p n$ is known as a *probabilistic choice* between m and n , and behaves like m with probability p and like n with probability $1-p$.¹

Limitations of static typing. Adopting a static typing for our probabilistic language would be cumbersome, as it would require inserting type annotations everywhere, or else extending the language with a type inference mechanism. In either case, the static typing can be overly conservative, rejecting (at compile time) programs that may indeed go right at runtime. For example, declaring the return types of functions `externalCall` and `localCall` as argued above ($\{\text{Bool}_{\frac{90}{100}}, \text{Error503}_{\frac{10}{100}}\}$ and $\{\text{Bool}_{\frac{95}{100}}, \text{Error503}_{\frac{5}{100}}\}$, respectively), would render function `isUserActive` ill-typed as the two branches of the conditional in the function body would have different types. Note that, even though the uptime of the external web service is incompatible with the uptime requirements of the login endpoint, we still would like to have a program able to execute in development mode (and in production mode, with minor modifications in typing annotation, if uptime requirements are reconciled).

Gradual typing at rescue. Gradual typing addresses this problem by supporting a smooth transition between static and dynamic typing, introducing imprecision on static types via the unknown annotation $?$.² Intuitively, an unknown type (resp. probability) $?$ represents any type (resp. probability). For example, using gradual (distribution) types we can partially annotate the program to assert only a subset of function uptimes:

```

4 val externalCall: ? -> {Bool90/100, Error50310/100} = ...
5 def isUserActive(user: ?): ? = if (prod) externalCall(user) :: ? else
    ↪ localCall(user) :: ?
6 def login(user: ?, pass: ?): {Bool95/100, Error5035/100} = if (isUserActive(user)) ...

```

To render `isUserActive` well-typed, we also had to ascribe both its conditional branches to the unknown type (written $:: ?$), since the conditional branches have different (fully static) types.

The type checker of a gradual language treats type equality optimistically, through the notion of *consistency*. Consistency between gradual types tests the plausibility of equality between any of the static types they represent. For instance, gradual type $? \rightarrow \text{Bool}$ is consistent with $\text{Int} \rightarrow ?$, written $? \rightarrow \text{Bool} \sim \text{Int} \rightarrow ?$, because (during runtime) they can both represent, e.g., the fully static type $\text{Int} \rightarrow \text{Bool}$.

In view of this optimistic treatment of equality, the above program is accepted statically as the unknown type $?$ is (trivially) consistent with every other type. If the application is in development mode, then the `login` endpoint runs successfully. On the contrary, if the application is in production mode, a runtime error is raised. This is because, even though $\{\text{Bool}_{\frac{90}{100}}, \text{Error503}_{\frac{10}{100}}\} \sim ?$ and $? \sim \{\text{Bool}_{\frac{95}{100}}, \text{Error503}_{\frac{5}{100}}\}$, $\{\text{Bool}_{\frac{90}{100}}, \text{Error503}_{\frac{10}{100}}\} \not\sim \{\text{Bool}_{\frac{95}{100}}, \text{Error503}_{\frac{5}{100}}\}$. Said otherwise, consistency is not transitive. Therefore, gradual languages incorporate runtime checks to detect any potential violation of the optimistic assumptions performed statically during type checking.

Finally, note that we can increase the program precision, e.g., changing the return type of `isUserActive` from $?$ to $\{\text{Bool}_{\frac{95}{100}}, \text{Error503}_{\frac{5}{100}}\}$, which would make the program ill-typed, failing thus at compile time.

Besides for this enhanced expressivity, one can also employ our *gradual* probabilistic language for program refinement purposes. To illustrate this application, assume that the external service `externalCall` is now required to have an uptime of *at least* 95%. This can be modelled by declaring

¹A probabilistic choice $m \oplus_p n$ can be readily simulated (in an approximate manner) by all programming languages that include the commonplace primitive `random()`, returning an (approximately) uniform value in the $[0, 1]$ interval. It suffices to take program `if (random() <= p) m else n`.

²A fully untyped program is considered to have unknown annotations everywhere.

its return type as $\{\text{Bool}^{\frac{95}{100}}, \text{Bool}^?, \text{Error503}^?\}$. In contrast to the above example where $?$ represented unknown *types*, here, both occurrences of $?$ represent unknown (possibly different) *probabilities*, which together with $\frac{95}{100}$ must sum up to 1.

Furthermore, assume that the external service originally relied on a single server of 90% uptime (server1) to keep track of active users. To reach the desired uptime of (at least) 95%, the service provider decides to buy a new—very costly—server of 98% uptime (server2). A naive implementation of the service would simply dispense with server1 and rely only on server2 to respond queries. However, this would negatively impact on server2 lifetime, diminishing the return of the performed investment. To avoid this problem and still benefit from server1, a possible solution consists in, upon each query, *probabilistically* choosing either server to respond the query. The fundamental question left to answer is whether this design would result in an overall (expected) uptime of at least 95%. To answer this question, we can consider the following program:

```

7  val server1: ? -> {Bool90/100, Error50310/100} = ...
8  val server2: ? -> {Bool98/100, Error5032/100} = ...
9  def externalCall(user: ?): {Bool95/100, Bool?, Error503?} = server1(user) ⊕?
    ↪ server2(user)

```

where symbol $?$ in the probabilistic choice $\oplus_?$ also represents an unknown probability. Our gradual language correctly typechecks this program and its runtime informs us about the feasibility of the proposed `externalCall` design. For concreteness, assume that for any user, server1 (resp. server2) responds `true` with probability $\frac{90}{100}$ (resp. $\frac{98}{100}$) and `error503` with the complementary probability. The instrumentation of the language runtime introduces symbolic variables to represent unknown probabilities; say ω_{Bool} , ω_{Error503} and ω_{choice} represent the unknown probabilities encoded by $?$ respectively in $\text{Bool}^?$, $\text{Error503}^?$ and $\oplus_?$. The runtime semantics tells us that invoking `externalCall` with any user returns `true` with probability $\frac{95}{100} + \omega_{\text{Bool}}$ and `error503` with probability $\frac{2}{100} + (\frac{98}{100} - \frac{90}{100})\omega_{\text{choice}}$, where the symbolic variables are constrained by formulas $\frac{95}{100} + \omega_{\text{Bool}} + \omega_{\text{Error503}} = 1$ and $\frac{90}{100}\omega_{\text{choice}} + \frac{98}{100}(1 - \omega_{\text{choice}}) = \frac{95}{100} + \omega_{\text{Bool}}$.³ Any probability $\omega_{\text{choice}} \leq \frac{37.5}{100}$ yields a valid solution of the equation system, yielding a valid refinement of `externalCall` and validating the proposed design.

2.2 Design Decisions and Challenges

When designing our source (GPLC) and target (TPLC) gradual probabilistic languages, we faced several design decisions and challenges.

Where to introduce imprecision. In most traditional gradual typing calculi, imprecision is introduced via the unknown type $?$. To gain expressivity and flexibility, in this work we allow imprecision at the type level as well as the probability level. For example, given the fully static distribution type $\{\text{Bool}^{\frac{9}{10}}, \text{Error503}^{\frac{1}{10}}\}$, we can introduce imprecision either in probabilities, e.g. $\{\text{Bool}^?, \text{Error503}^{\frac{1}{10}}\}$, in the underlying types, e.g. $\{?^{\frac{9}{10}}, \text{Error503}^{\frac{1}{10}}\}$, or in both. Note that there is no need to introduce the unknown distribution as it can be represented by the gradual distribution type $\{?\}$. Interestingly, unknown probabilities are particularly useful for expressing *probability bounds*. As hinted above, we can use type $\{\text{Bool}^{\frac{95}{100}}, \text{Bool}^?, \text{Error503}^?\}$ to represent a service with an uptime of *at least* 95%. On the other hand, type $\{\text{Bool}^?, \text{Error503}^{\frac{5}{100}}, \text{Error503}^?\}$ models an uptime of *at most* 95%.

Tracking dependencies of probability annotations. When dealing with unknown probabilities, as in type $\{\text{Bool}^{\frac{9}{10}}, \text{Bool}^?, \text{Error503}^?\}$, the gradual language must ensure that the concrete probabilities

³Formally, the instrumentation of the runtime semantics yields a handful of further constraints, but altogether they are equivalent to the considered subset.

$r \in \mathbb{R}, \quad b \in \mathbb{B}, \quad x \in \text{Var}, \quad p \in [0, 1], \quad \tau \in \text{TYPE}, \quad T \in \text{DTYPE}$	
$\tau ::= \text{Real} \mid \text{Bool} \mid \tau \rightarrow T$	(simple types)
$T ::= \{\{\tau_i^{p_i} \mid i \in \mathcal{I}\}\}$	(distribution types)
$m, n ::= v \mid v \ w \mid \text{let } x = m \text{ in } n \mid m \oplus_p n$	(terms)
$m ::= T \mid v :: \tau \mid \text{if } v \text{ then } m \text{ else } n \mid v + w$	
$v, w ::= x \mid r \mid b \mid (\lambda x : \tau. m)$	(values)

Fig. 1. Syntax of SPLC.

they represent induce only well-defined static distribution types, with a total probability of 1. This requirement induces implicit dependencies and gradual probabilities are thus elaborated to fresh variables (ω) constrained by formulas, e.g. of the form $\frac{9}{10} + \omega_1 + \omega_2 = 1$.

Ascribing to distribution types. One of the fundamental features of GPLC is the possibility of ascribing programs to distribution types. For example, a program $f = (\lambda x : ?.x) :: \{(\text{Real} \rightarrow ?)^{\frac{1}{2}}, (? \rightarrow \text{Bool})^{\frac{1}{2}}\}$ behaves as a function that takes a number as argument with probability $\frac{1}{2}$, and as a function that returns a boolean also with probability $\frac{1}{2}$. Reducing an application to f and correctly propagating the respective type information is not a trivial task. Intuitively, our approach consists in “pushing” the real argument into each (compatible) type in the distribution. For instance, the reduction of program f 1 proceeds, informally, as follows:

$$f \ 1 \mapsto^* \{(\lambda x : ?.x) :: (\text{Real} \rightarrow ?)^{\frac{1}{2}} \ 1, (\lambda x : ?.x) :: (? \rightarrow \text{Bool})^{\frac{1}{2}} \ 1\} \mapsto^* \{1 :: ?^{\frac{1}{2}}, \mathbf{error}^{\frac{1}{2}}\}$$

Couplings as a central tool. Defining some key relations between distribution types is another technical challenge. For instance, should we consider distribution type $\{(\text{Real} \rightarrow ?)^{\frac{1}{2}}, (? \rightarrow \text{Real})^{\frac{1}{2}}\}$ consistent with $\{(? \rightarrow \text{Bool})^{\frac{1}{3}}, (\text{Real} \rightarrow ?)^{\frac{2}{3}}\}$? Is $\{\text{Real}^{\frac{1}{2}}, ?^{\frac{1}{2}}\}$ more precise than $\{\text{Bool}^{\frac{2}{3}}, ?^{\frac{1}{3}}\}$? To define these (and other) relations over distribution types we heavily rely on the notion of probabilistic coupling, which yields a canonical lifting from relations over pair of sets to probability distributions over the sets.

3 SPLC: STATIC LANGUAGE

In this section, we present SPLC, a statically-typed lambda calculus, extended with a probabilistic choice operator, which is the starting point –static end– of our gradualization effort. The static semantics of SPLC is based on that of λ_{\oplus} from [Lago and Grellois 2017], with two major differences: SPLC features a semantic (rather than syntactic) equality between types and also allows type ascriptions. As for the dynamic semantics, programs are interpreted as probability distributions over final values.

3.1 Syntax

The syntax of SPLC is presented in Figure 1, comprising its type and term languages.

Type language. The type language contains two (mutually defined) syntactic categories: simple types and distributions types. A *simple type*, ranged over by τ , can be the type Real of real numbers, the type Bool of Boolean values, or a function type of the form $\tau \rightarrow T$, where T is a distribution type. A *distribution type*, ranged over by T , is a multi-set of pairs comprised by a simple type τ and a probability p in the interval $[0, 1]$. Intuitively, we use $\{\{\tau_i^{p_i} \mid i \in \mathcal{I}\}\}$ to denote a distribution type in which simple type τ_i occurs with probability p_i , for each i in the (non-empty and finite) subset \mathcal{I} of the natural numbers. For instance, distribution type $\{\{\text{Real}^{\frac{1}{4}}, \text{Real}^{\frac{1}{4}}, \text{Bool}^{\frac{1}{2}}\}\}$ represents Real with probability $\frac{1}{4} + \frac{1}{4} = \frac{1}{2}$ and Bool with probability $\frac{1}{2}$. Notationwise, we sometimes omit the index set

$\Gamma \vdash_s v : \tau, \quad \Gamma \vdash_s m : T$

$\text{(Tr)} \frac{}{\Gamma \vdash_s r : \text{Real}}$	$\text{(Tb)} \frac{}{\Gamma \vdash_s b : \text{Bool}}$	$\text{(Tx)} \frac{\Gamma(x) = \tau}{\Gamma \vdash_s x : \tau}$
$\text{(Tl)} \frac{\Gamma, x : \tau \vdash_s m : T \quad \vdash \tau}{\Gamma \vdash_s \lambda x : \tau. m : \tau \rightarrow T}$	$\text{(Tv)} \frac{\Gamma \vdash_s v : \tau}{\Gamma \vdash_s v : \{\tau^1\}}$	$\text{(T:: } \tau) \frac{\Gamma \vdash_s v : \tau' \quad \tau' =_s \tau \quad \vdash \tau}{\Gamma \vdash_s v : \tau : \{\tau^1\}}$
$\text{(Tapp)} \frac{\Gamma \vdash_s v : \tau_1 \quad \Gamma \vdash_s w : \tau_2 \quad \text{dom}(\tau_1) =_s \tau_2}{\Gamma \vdash_s v w : \text{cod}(\tau_1)}$	$\text{(T}\oplus) \frac{\Gamma \vdash_s m : T_1 \quad \Gamma \vdash_s n : T_2}{\Gamma \vdash_s m \oplus_p n : p \cdot T_1 + (1-p) \cdot T_2}$	$\text{(T:: } T) \frac{\Gamma \vdash_s m : T' \quad T' =_s T \quad \vdash T}{\Gamma \vdash_s m : T : T}$
$\text{(Tlet)} \frac{\Gamma \vdash_s m : \{\tau_i^{p_i} \mid i \in \mathcal{I}\} \quad \forall i \in \mathcal{I}. \Gamma, x : \tau_i \vdash_s n : T_i}{\Gamma \vdash_s \text{let } x = m \text{ in } n : \sum_{i \in \mathcal{I}} p_i \cdot T_i}$	$p \cdot \{\tau_i^{p_i} \mid i \in \mathcal{I}\} = \{\tau_i^{p \cdot p_i} \mid i \in \mathcal{I}\}$	
$\text{dom}(\tau \rightarrow T) = \tau$	$\text{cod}(\tau \rightarrow T) = T$	$p \cdot \{\tau_i^{p_i} \mid i \in \mathcal{I}\} = \{\tau_i^{p \cdot p_i} \mid i \in \mathcal{I}\}$
$\text{dom}(\tau) \text{ undef. otherwise}$	$\text{cod}(\tau) \text{ undef. otherwise}$	
$\{\tau_i^{p_i} \mid i \in \mathcal{I}\} + \{\tau_j^{p_j} \mid j \in \mathcal{J}\} = \{\tau_i^{p_i} \mid i \in \mathcal{I}\} \cup \{\tau_j^{p_j} \mid j \in \mathcal{J}\} \quad \text{if } \sum_{i \in \mathcal{I}} p_i + \sum_{j \in \mathcal{J}} p_j \leq 1$		

Fig. 2. Type system of SPLC (excerpt).

\mathcal{I} and simply write $\{\tau_i^{p_i}\}$. Finally, note that distribution types—as the name suggests—represent *probability distributions* (over simple types) and therefore, well-typed programs are associated distribution types whose probabilities sum up to 1 (this restriction is formally captured by the notion of *type well-formedness* defined in Section 3.1).

Term language. Terms, ranged over by m, n , and values, ranged over by v, w , are mutually defined. A *term* can be a value v , an application $v w$ between two values, a let expression $\text{let } x = m \text{ in } n$, a probabilistic choice $m \oplus_p n$, a term ascription $m :: T$, a value ascription $v :: \tau$, a conditional if v then m else n , or an addition $v + w$ between two values. Note that terms are defined in A-normal form [Sabry and Felleisen 1993], which pushes all the reasoning about probabilities to the let construct. Randomization is introduced through probabilistic choices: program $m \oplus_p n$ behaves like (i.e. reduce to) m with probability p and like n with probability $1 - p$.

3.2 Type System

Figure 2 presents the type system of SPLC. Type rules are defined using a pair of mutually-defined judgments: one for values and another for computations. Judgment $\Gamma \vdash_s v : \tau$ (resp. $\Gamma \vdash_s m : T$) for values (resp. computations) denotes that value v (resp. term m) has simple type τ (resp. distribution type T) under type environment Γ , which maps variables to simple types.

Type rules for values are standard, only a few rules deserving special attention. For example, rule (Tv) allows assigning a value of simple type τ also distribution type $\{\tau^1\}$ (e.g. program 1 can be typed as $\{\text{Int}^1\}$). Also, note that rules (T λ), (T:: τ) and (T:: T) require all program type annotations to be well-formed. We say that a distribution type $\{\tau_i^{p_i} \mid i \in \mathcal{I}\}$ is *well-formed*, written $\vdash \{\tau_i^{p_i} \mid i \in \mathcal{I}\}$, if $\sum_{i \in \mathcal{I}} p_i = 1$ and simple type τ_i is well-formed for every $i \in \mathcal{I}$. A simple type τ is *well-formed*, written $\vdash \tau$, if it is either a base type (Real or Bool) or a function type $\tau \rightarrow T$, where τ and T are well-formed.

A particularity of SPLC's type system is that it relies on a semantic—rather than syntactic— notion of type equality ($=_s$), as used in rules (T:: τ), (Tapp) and (T:: T). For example, $\{\text{Real}^{\frac{1}{2}}, \text{Bool}^{\frac{1}{4}}, \text{Bool}^{\frac{1}{4}}\} =_s \{\text{Real}^{\frac{1}{3}}, \text{Real}^{\frac{1}{6}}, \text{Bool}^{\frac{1}{2}}\}$ because $\frac{1}{2} = \frac{1}{3} + \frac{1}{6}$ and $\frac{1}{4} + \frac{1}{4} = \frac{1}{2}$. Formally, type equality is given by rules:

$$\boxed{m \Downarrow_s \mathcal{V}} \quad \mathcal{V} ::= \{\{v_i^{p_i} \mid i \in \mathcal{I}\}\} \text{ (distribution values)}$$

$$\frac{}{v \Downarrow_s \{\{v^1\}\}} \quad \frac{m[v/x] \Downarrow_s \mathcal{V}}{(\lambda x : \tau. m) v \Downarrow_s \mathcal{V}} \quad \frac{m \Downarrow_s \mathcal{V}_1 \quad m \Downarrow_s \mathcal{V}_2}{m \oplus_p n \Downarrow_s p \cdot \mathcal{V}_1 + (1-p) \cdot \mathcal{V}_2}$$

$$\frac{m \Downarrow_s \{\{v_i^{p_i} \mid i \in \mathcal{I}\}\} \quad \forall i \in \mathcal{I}. n[v_i/x] \Downarrow_s \mathcal{V}_i}{\text{let } x = m \text{ in } n \Downarrow_s \sum_{i \in \mathcal{I}} p_i \cdot \mathcal{V}_i} \quad \frac{}{v :: \tau \Downarrow_s \{\{v^1\}\}} \quad \frac{m \Downarrow_s \mathcal{V}}{m :: T \Downarrow_s \mathcal{V}}$$

Fig. 3. Runtime semantics of SPLC (excerpt).

$$\frac{}{\text{Real} =_s \text{Real}} \quad \frac{}{\text{Bool} =_s \text{Bool}} \quad \frac{\tau_1 =_s \tau_2 \quad T_1 =_s T_2}{\tau_1 \rightarrow T_1 =_s \tau_2 \rightarrow T_2} \quad \frac{\forall \tau \in \text{supp}(T_1). T_1(\tau) = T_2(\tau)}{T_1 =_s T_2}$$

where $\text{supp}(T)$ represents the *support* of distribution type T defined by $\text{supp}(T) = \{\tau \mid \tau^p \in T \wedge p > 0\}$ and $T(\tau)$ represents the *probability* that T assigns to a simple type τ , defined by $\{\{v_i^{p_i} \mid i \in \mathcal{I}\}\}(\tau) = \sum_{i \in \mathcal{I} \mid \tau_i =_s \tau} p_i$.

Following the approach of Garcia et al. [2016], to ease the gradualization process we make all type relations and type functions explicit. For (Tapp) rule, we use partial functions dom and cod to extract the domain and codomain of a function type, respectively. Also we make explicit the fact that the type of the argument should be equal to the domain type of the function. Rule (T \oplus) combines the distribution types T_1 and T_2 of sub-expressions m and n , by first scaling T_1 by p and T_2 by $1-p$, and then adding the resulting scaled distributions types together. Scaling $p \cdot T$ is defined pointwise, i.e. by scaling all the probabilities in the distribution type by p ; the addition of two (sub) distribution types $T_1 + T_2$ is defined as the union of the two multi-sets, provided that the sum of the resulting probabilities do not exceed 1. For instance, consider program $1 \oplus_{\frac{1}{3}} \text{true}$. Expression 1 is typed as $\{\{\text{Real}^1\}\}$ and true as $\{\{\text{Bool}^1\}\}$. After scaling both distribution types and adding them together, the resulting distribution type is $\frac{1}{3} \cdot \{\{\text{Real}^1\}\} + \frac{2}{3} \cdot \{\{\text{Bool}^1\}\} = \{\{\text{Real}^{\frac{1}{3}}, \text{Bool}^{\frac{2}{3}}\}\}$. Rule (Tlet) propagates the type of m to n as follows. If m has a distribution type T , then for each type and probability $\tau^p \in T$, n is type checked under an extended environment where x is typed as τ . The resulting type of the let expression is computed by adding each distribution type of n scaled by its corresponding p .

Note that, as expected, well-typed terms are assigned well-formed types, only.

LEMMA 3.1 (TYPE WELL-FORMEDNESS). *For every value v , every term m , every simple type $\tau \in \text{TYPE}$, every distribution type $T \in \text{DTYPE}$ and every environment Γ ,*

- (1) If $\Gamma \vdash_s v : \tau$, then $\vdash \tau$
- (2) If $\Gamma \vdash_s m : T$, then $\vdash T$

3.3 Dynamic Semantics

We endow SPLC with a big-step *distribution-based* semantics that relates programs to probability distributions over final values [Lago and Zorzi 2012a], following a call-by-value reduction strategy. Concretely, judgment $m \Downarrow_s \mathcal{V}$ denotes that expression m reduces to a *distribution value* \mathcal{V} , i.e. a probability distribution over values. The reduction relation is formally defined in Figure 3.

A value v reduces to a Dirac distribution, i.e. a distribution that assigns probability 1 to v (and 0 to any other value). A function application reduces by substituting the argument for the variable binder in the function body. A probabilistic choice first reduces its pair of branches and then returns the weighted sum of the so obtained distribution values. The scaling and addition operators for

$r \in \mathbb{R}, \quad b \in \mathbb{B}, \quad x \in \text{Var}, \quad p \in [0, 1], \quad \rho \in \text{GPROB}, \quad \sigma \in \text{GTYPE}, \quad \mu \in \text{GDTYPE}$	
$\rho ::= p \mid ?$	(gradual probabilities)
$\sigma, \delta ::= \text{Real} \mid \text{Bool} \mid \sigma \rightarrow \mu \mid ?$	(gradual simple types)
$\mu, \nu ::= \{\{\sigma_i^{p_i} \mid i \in \mathcal{I}\}\}$	(gradual distribution types)
$m, n ::= v \mid v \ w \mid \text{let } x = m \text{ in } n \mid m \oplus_\rho n$	(terms)
$m ::= \mu \mid v ::= \sigma \mid \text{if } v \text{ then } m \text{ else } n \mid v + w$	
$v, w ::= x \mid r \mid b \mid \lambda x : \sigma. m$	(values)

Fig. 4. Syntax of GPLC.

distribution values are defined analogously to those for type distributions (Fig. 2). For example, program $(1 \oplus_{\frac{1}{2}} 2) \oplus_{\frac{2}{3}} \text{true}$ reduces to distribution value $\{\{1^{\frac{1}{3}}, 2^{\frac{1}{3}}, \text{true}^{\frac{1}{3}}\}\}$. The reduction of a let-expression $\text{let } x = m \text{ in } n$ is more involved and proceeds as follows. First, subterm m is reduced to a distribution value $\{\{v_i^{p_i} \mid i \in \mathcal{I}\}\}$. Second, subterm n is reduced by substituting each v_i (i.e. each possible outcome of m) for x , resulting in distribution values \mathcal{V}_i . The entire let-expression then reduces to the weighted sum $\sum_{i \in \mathcal{I}} p_i \cdot \mathcal{V}_i$. Finally, ascribed terms reduce by removing type ascription.

SPLC is *type safe*, meaning that every well-typed closed expression reduces to a distribution value. Formally, this follows from three results of GPLC that we establish in Section 4 (Theorem 4.13) and Section 5 (Theorems 5.11 and 5.15).

4 GPLC: GRADUAL SOURCE LANGUAGE

We now present GPLC, our gradual source probabilistic language. GPLC is derived from SPLC and its design is justified by the Abstracting Gradual Typing (AGT) methodology [Garcia et al. 2016]. The section is structured as follows. First, we introduce GPLC syntax, specifying, in particular, where we support (im)precision. Second, we present GPLC type system and define consistency, discussing why a naive approach to consistency is bound to fail. Third, we define type and term precision, proving that GPLC satisfies the gradual guarantee and that its type system conservatively extends that of SPLC. The dynamic semantics of GPLC is defined through an elaboration to a target language, introduced in Section 5.⁴

4.1 Syntax

The syntax of GPLC is presented in Figure 4. We introduce imprecision in the language by extending probabilities and simple types with the unknown annotation $?$. The unknown probability $?$ represents any probability in the interval $[0, 1]$, and similarly, the unknown simple type $?$ represents any static simple type. We do not need an (explicit) unknown distribution type as it can already be encoded by the singleton distribution type $\{\{?\}\}$ (of unknown simple type, with unknown probability). Notationwise, we use ρ to range over gradual probabilities (GPROB), σ, δ to range over simple gradual types (GTYPE), and μ, ν to range over gradual distribution types (GDTYPE).

Design driven by AGT. To justify some of the design decisions behind GPLC, we follow, in parallel, the Abstracting Gradual Typing (AGT) methodology [Garcia et al. 2016]. In short, the idea behind AGT is that starting from a specification of the meaning of gradual types in terms of sets of static types, we can systematically derive all relevant notions of the gradual language, which by construction, will enjoy a set of desired properties (to be discussed later). Unfortunately, some of the so-obtained definitions turn out not to be very amenable to implementation. To address this limitation, we also derive alternative (equivalent) definitions, with a more operational nature.

⁴We use the [blue color](#) for source languages and the [red color](#) for target languages.

$$\boxed{\Gamma \vdash v : \sigma, \quad \Gamma \vdash m : \mu}$$

$$\frac{\Gamma \vdash v : \sigma}{\Gamma \vdash v : \{\{\sigma^1\}\}} \quad \frac{\Gamma, x : \sigma \vdash m : \mu \quad \vdash \sigma}{\Gamma \vdash \lambda x : \sigma. m : \sigma \rightarrow \mu} \quad \frac{\Gamma \vdash v : \sigma \quad \sigma \sim \delta \quad \vdash \delta}{\Gamma \vdash v :: \delta : \{\{\delta^1\}\}}$$

$$\frac{\Gamma \vdash v : \sigma \quad \Gamma \vdash w : \delta \quad \delta \sim \widetilde{dom}(\sigma)}{\Gamma \vdash v w : \widetilde{cod}(\sigma)} \quad \frac{\Gamma \vdash m : \mu \quad \Gamma \vdash n : v}{\Gamma \vdash m \oplus_{\rho} n : \rho \cdot \mu + (1-\rho) \cdot v}$$

$$\frac{\Gamma \vdash m : \{\{\sigma_i^{\rho_i} \mid i \in \mathcal{I}\}\} \quad \forall i \in \mathcal{I}. \Gamma, x : \sigma_i \vdash n : \mu_i}{\Gamma \vdash \text{let } x = m \text{ in } n : \sum_{i \in \mathcal{I}} \rho_i \cdot \mu_i} \quad \frac{\Gamma \vdash m : \mu \quad \mu \sim v \quad \vdash v}{\Gamma \vdash m :: v : v}$$

Fig. 5. Type system of GPLC (excerpt).

As just hinted, we start providing the meaning of gradual types and probabilities via concretization functions that map *gradual* simple types, distribution types and probabilities to non-empty sets of *static* simple types, distribution types and probabilities, respectively.

$$\begin{aligned}
\gamma_p : [0, 1] \cup \{?\} &\rightarrow \mathcal{P}([0, 1]) & \gamma_\tau : \text{GTYPE} &\rightarrow \mathcal{P}(\text{TYPE}) \\
\gamma_p(?) = [0, 1] & \quad \gamma_p(p) = \{p\} & \gamma_\tau(\sigma \rightarrow \mu) &= \{\tau \rightarrow T \mid \tau \in \gamma_\tau(\sigma) \wedge T \in \gamma_T(\mu)\} \\
\gamma_T : \text{GDTYPE} &\rightarrow \mathcal{P}(\text{DTTYPE}) & \gamma_\tau(?) = \text{TYPE} & \quad \gamma_\tau(\text{Real}) = \{\text{Real}\} \quad \gamma_\tau(\text{Bool}) = \{\text{Bool}\} \\
\gamma_T(\{\{\sigma_i^{\rho_i} \mid i \in \mathcal{I}\}\}) &= \{\{\tau_i^{\rho_i} \mid i \in \mathcal{I}\} \mid \forall i \in \mathcal{I}. \tau_i \in \gamma_\tau(\sigma_i) \wedge \rho_i \in \gamma_p(\rho_i)\}
\end{aligned}$$

The concretization functions crisply captures the intuition behind imprecision: The meaning of the unknown gradual probability is any probability in the interval $[0, 1]$ and the meaning of the unknown gradual simple type is any static simple type. The meaning of a gradual distribution type is computed inductively, by computing the meaning of both gradual simple types and gradual probabilities.

4.2 Type System

Figure 5 shows the type system of GPLC, which is obtained from the type system of the static language (Figure 2) by replacing static elements with their gradual counterpart. Let us briefly describe these liftings. The lifting $\widetilde{dom} : \text{GTYPE} \rightarrow \text{GTYPE}$ (resp. $\widetilde{cod} : \text{GTYPE} \rightarrow \text{GDTYPE}$) of type function dom is standard: $\widetilde{dom}(\sigma \rightarrow \mu) = \sigma$, $\widetilde{dom}(?) = ?$, and \widetilde{dom} is undefined elsewhere. Function \widetilde{cod} is defined analogously. The lifting of the minus (resp. product) operation between probabilities, also denoted by $-$ (resp. \cdot), returns $?$ if either of the operands is $?$:

$$\rho_1 \text{ op } \rho_2 = \begin{cases} \rho_1 \text{ op } \rho_2 & \text{if } \rho_1, \rho_2 \in [0, 1] \\ ? & \text{otherwise} \end{cases} \quad \text{op} \in \{\cdot, -\}$$

The lifting of the scaling of distribution types, also denoted by \cdot , is defined pointwise, in terms of the lifting of the product between probabilities: $\rho \cdot \{\{\sigma_i^{\rho_i} \mid i \in \mathcal{I}\}\} = \{\{\sigma_i^{\rho \cdot \rho_i} \mid i \in \mathcal{I}\}\}$. The lifting of the sum between distribution types, also denoted by $+$, coincides with the original operation (see Fig. 2).⁵

The lifting of type equality, called type *consistency* and denoted by \sim in GPLC, plays a fundamental role in gradual languages. It allows soundly handling the notion of (im)precision, which is conveniently introduced via type ascriptions. For example, program $1 \oplus_{\frac{1}{2}} \text{true} :: \{\{?\}\} :: \{\{\text{Real}^{\frac{2}{3}}, \text{Bool}^{\frac{1}{3}}\}\}$ is (optimistically) accepted by the gradual type system of GPLC because $\{\{\text{Real}^{\frac{1}{2}}, \text{Bool}^{\frac{1}{2}}\}\} \sim \{\{?\}\}$

⁵Formally, side condition $\sum_{i \in \mathcal{I}} \rho_i + \sum_{j \in \mathcal{J}} \rho_j \leq 1$ is defined following the AGT approach, i.e. it holds if there exist concretizations $p_i \in \gamma_p(\rho_i)$ and $p_j \in \gamma_p(\rho_j)$ such that $\sum_{i \in \mathcal{I}} p_i + \sum_{j \in \mathcal{J}} p_j \leq 1$.




Fig. 6. Probability splitting to justify consistency between gradual distribution types.

and $\{\{?\}\} \sim \{\{\text{Real}^{\frac{2}{3}}, \text{Bool}^{\frac{1}{3}}\}\}$. Following AGT, we define type consistency by the existential lifting of type equality $=_s$:

Definition 4.1 (Type consistency, by AGT). For any pair of gradual simple types $\sigma, \delta \in \text{GTYPE}$ and any pair of gradual distribution types $\mu, \nu \in \text{GDTYPE}$, we define:

$$\begin{aligned} \sigma \sim_{\text{AGT}} \delta & \text{ if and only if } \exists \tau_1 \in \gamma_\tau(\sigma), \tau_2 \in \gamma_\tau(\delta). \tau_1 =_s \tau_2 \\ \mu \sim_{\text{AGT}} \nu & \text{ if and only if } \exists T_1 \in \gamma_T(\mu), T_2 \in \gamma_T(\nu). T_1 =_s T_2 \end{aligned}$$

In words, two gradual types are consistent if there exist static simple types in their concretizations that are equal. The problem with this definition is that is not practical, as it can depend on sets of infinitely many types. For gradual simple types, this can be partially addressed by stating that $?$ is consistent with every other gradual simple type, but for gradual distribution types the problem is more challenging as probabilities must also be taken into account.

4.3 Consistency, Refined

We are thus interested in an *inductive* definition of consistency. To illustrate the main idea behind our alternative characterization, consider the pair of gradual distribution types

$$\mu = \{\{(\text{Real} \rightarrow \{\{?^1\}\})^{\frac{1}{2}}, (? \rightarrow \{\{\text{Real}^1\}\})^{\frac{1}{2}}\}\} \quad \text{and} \quad \nu = \{\{(? \rightarrow \{\{\text{Bool}^1\}\})^{\frac{1}{3}}, (\text{Real} \rightarrow \{\{?^1\}\})^{\frac{2}{3}}\}\},$$

represented on the left and right hand side of Figure 6 (for concreteness, we assume that the elements of μ and ν are enumerated by index set $\mathcal{S} = \{1, 2\}$, thus, e.g., $(\text{Real} \rightarrow \{\{?^1\}\})^{\frac{1}{2}}$ corresponds to the simple type of index 1 in μ). Intuitively, μ and ν will be consistent if (and only if) there exists a *splitting* of the probabilities $\frac{1}{2}, \frac{1}{2}$ from μ and $\frac{1}{3}, \frac{2}{3}$ from ν that relates the simple types in μ with the simple types in ν as follows:

- (1) Type $\text{Real} \rightarrow \{\{?^1\}\}$ in μ is consistent with both $? \rightarrow \{\{\text{Bool}^1\}\}$ and $\text{Real} \rightarrow \{\{?^1\}\}$ in ν . This means that $\frac{1}{2}$, the probability of $\text{Real} \rightarrow \{\{?^1\}\}$ in μ , must be split into two, i.e. $\frac{1}{2} = \alpha_{11} + \alpha_{12}$, where α_{11} (resp. α_{12}) represents the probability of relating $\text{Real} \rightarrow \{\{?^1\}\}$, the first simple type in μ , with $? \rightarrow \{\{\text{Bool}^1\}\}$ (resp. $\text{Real} \rightarrow \{\{?^1\}\}$), the first (resp. second) simple type in ν .
- (2) Type $? \rightarrow \{\{\text{Real}^1\}\}$ in μ is consistent only with $\text{Real} \rightarrow \{\{?^1\}\}$ in ν . Therefore, the probability of $? \rightarrow \{\{\text{Real}^1\}\}$ in μ need not be split, leading to $\frac{1}{2} = \alpha_{22}$.
- (3) Similarly, type $? \rightarrow \{\{\text{Bool}^1\}\}$ in ν is consistent only with $\text{Real} \rightarrow \{\{?^1\}\}$ in μ , so $\frac{1}{3} = \alpha_{11}$.
- (4) Finally, type $\text{Real} \rightarrow \{\{?^1\}\}$ in ν is consistent with $\text{Real} \rightarrow \{\{?^1\}\}$ and $? \rightarrow \{\{\text{Real}^1\}\}$ in μ , resulting in $\frac{2}{3} = \alpha_{12} + \alpha_{22}$.

Since the system of four equations so derived is feasible, witnessed e.g. by solution $\alpha_{11} = \frac{1}{3}$, $\alpha_{12} = \frac{1}{6}$ and $\alpha_{22} = \frac{1}{2}$, we can conclude that μ and ν are consistent. This thought process constitutes a lifting of the consistency relation between simple types to distribution types through *couplings* [Deng and Du 2011], tool that has already been exploited e.g. in the context of probabilistic bisimulation [Segala and Lynch 1995] and verification of cryptographic properties [Barthe et al. 2009].

Definition 4.2. (Relation lifting) Assume that $\mathcal{A} = \{\{a_i^{p_i} \mid i \in \mathcal{I}\}\}$ and $\mathcal{B} = \{\{b_j^{q_j} \mid j \in \mathcal{J}\}\}$ are multi-set representations of discrete probability distributions over sets A and B , respectively (that is,

$a_i \in A$ and $p_i \in [0, 1]$ for all $i \in \mathcal{I}$, $b_j \in B$ and $q_j \in [0, 1]$ for all $j \in \mathcal{J}$, and $\sum_{i \in \mathcal{I}} p_i = \sum_{j \in \mathcal{J}} q_j = 1$. Moreover, let $R \subseteq A \times B$ be a relation between A and B . We say that \mathcal{A} and \mathcal{B} are related by the *lifting* of R , written $L_R(\mathcal{A}, \mathcal{B})$, iff there exist $\mathcal{C} = \{\{\alpha_{ij} \in [0, 1] \mid i \in \mathcal{I} \wedge j \in \mathcal{J}\}\}$ such that for all $i \in \mathcal{I}$ and all $j \in \mathcal{J}$,

$$(1) p_i = \sum_{j \in \mathcal{J}} \alpha_{ij} \wedge p_j = \sum_{i \in \mathcal{I}} \alpha_{ij}, \quad \text{and} \quad (2) \alpha_{ij} > 0 \Rightarrow a_i R b_j$$

We write $\mathcal{C} \vdash \mathcal{A} R \mathcal{B}$ to denote that \mathcal{C} is a witness of the relation $L_R(\mathcal{A}, \mathcal{B})$, i.e. to denote the conjunction between conditions 1 and 2 above. Moreover, any \mathcal{C} satisfying (only) condition 1 is called a *coupling* between \mathcal{A} and \mathcal{B} .

Type equality via couplings. To make a uniform treatment of type equality ($=_s$) in SPLC and type consistency (\sim) in GPLC, we start by redefining the type equality in SPLC in terms of couplings, via relation $=$:

$$\frac{}{\text{Real} = \text{Real}} \quad \frac{}{\text{Bool} = \text{Bool}} \quad \frac{\tau_1 = \tau_2 \quad T_1 = T_2}{\tau_1 \rightarrow T_1 = \tau_2 \rightarrow T_2} \quad \frac{L=(T_1, T_2)}{T_1 = T_2}$$

The last rule above says that two distribution types are equal if there exists a coupling that justifies the lifting of equality on simple types to distribution types (note that the $=$ symbol in the rule premise refers to equality over simple types, while the $=$ symbol in the conclusion refers to equality over distribution types). Using this rule we can, e.g., derive that $\{\{\text{Real}^{\frac{1}{2}}, \text{Bool}^{\frac{1}{2}}\}\} = \{\{\text{Real}^{\frac{1}{3}}, \text{Real}^{\frac{1}{6}}, \text{Bool}^{\frac{1}{2}}\}\}$ because the set of formulas $\frac{1}{2} = \alpha_{11} + \alpha_{12}$, $\frac{1}{2} = \alpha_{23}$, $\frac{1}{3} = \alpha_{11}$, $\frac{1}{6} = \alpha_{12}$ and $\frac{1}{2} = \alpha_{23}$ is satisfiable (by solution $\alpha_{11} = \frac{1}{3}$, $\alpha_{12} = \frac{1}{6}$, $\alpha_{23} = \frac{1}{2}$).

As expected, this alternative definition of equality is equivalent to the original from Section 3.

LEMMA 4.3 (ALT. CHARACTERIZATION OF EQUALITY). *For all pairs of simple types $\tau_1, \tau_2 \in \text{TYPE}$ and distributions $T_1, T_2 \in \text{DTYPE}$,*

$$\tau_1 =_s \tau_2 \text{ iff } \tau_1 = \tau_2 \quad \text{and} \quad T_1 =_s T_2 \text{ iff } T_1 = T_2$$

Armed with this new definition of equality based on couplings we proceed to define consistency.

Type consistency, a straightforward approach. A straightforward approach to define consistency in GPLC consists in (rule-wise) lifting the definition of type equality $=$ in SPLC, and extending it with rules stating that $?$ is consistent with any gradual simple type. An excerpt of the resulting set of rules would be:

$$\frac{}{? \sim \sigma} \quad \frac{}{\text{Real} \sim \text{Real}} \quad \frac{\sigma_1 \sim \sigma_2 \quad \mu_1 \sim \mu_2}{\sigma_1 \rightarrow \mu_1 \sim \sigma_2 \rightarrow \mu_2} \quad \frac{L=(\mu_1, \mu_2)}{\mu_1 \sim \mu_2}$$

According to this definition (in particular, by the last rule), establishing the consistency, e.g., between gradual distribution types $\{\{\sigma_i^{\rho_i} \mid i \in \mathcal{I}\}\}$ and $\{\{\sigma_j^{\rho_j} \mid j \in \mathcal{J}\}\}$ requires exhibiting a coupling between them. The problem here is that probabilities in either of the distribution types can be only partially known, that is, ρ_i could be $?$ for some $i \in \mathcal{I}$, rendering formula $\sum_{j \in \mathcal{J}} \alpha_{ij} = ?$ even ill-defined. A first approach to tackle this problem consists in lifting these formulas (from the static setting) to the gradual setting. Note, however, that this lifting cannot be done for each formula independently because the “same” $?$ will probably occur in multiple formulas. We should then lift all related formulas at the same time, but this still suffers from scoping problems because unknown probabilities (represented by $?$) must remain visible outside the formulas lifting: At runtime, we need to carry witness information about consistency, where gradual probabilities “flow” across reductions (see the let and probabilistic choice reduction rules in Fig. 3). To tackle this problem, we introduce (fresh) symbolic variables representing unknown probabilities, and analyze the existence of couplings for this symbolic representation of distribution types.

$\sigma, \delta \in \text{FSTYPE}, \quad \mu, \nu \in \text{FDTYPE}, \quad \omega \in \text{TVAR}, \quad \Phi \in \text{FORMULA}$	
$\omega ::= \langle \alpha, \ell, \mathcal{r} \rangle$	(tagged variables)
$\varrho ::= \omega \mid r$	(symbolic probabilities)
$\sigma, \delta ::= \text{Real} \mid \text{Bool} \mid \sigma \rightarrow \mu \mid ?$	(formula simple types)
$\mu, \nu ::= \Phi \triangleright \{\{\sigma_i^{\varrho_i} \mid i \in \mathcal{I}\}\}$	(formula distribution types)
$\Phi ::= \varphi = \varphi \mid \varphi \leq \varphi \mid \Phi \wedge \Phi$	(formulas)
$\varphi ::= \varrho \mid \varphi + \varphi \mid \varphi - \varphi \mid \varphi \cdot \varphi \mid \varphi / \varphi$	(expressions)

Fig. 7. Formula types.

Type consistency via symbolic liftings. To represent unknown probabilities as (free) variables in the set of formulas defining the lifting of consistency (from gradual simple types to gradual distribution types), we extend the syntax of GPLC with *formula simple types* (FSTYPE) and *formula distribution types* (FDTYPE) as shown in Figure 7. Intuitively, a formula type is the same as an ordinary gradual type, except that (1) unknown probabilities are replaced by variables, and (2) distribution types are guarded by formulas (like in refinement types). Formally, a *symbolic probability* ϱ is either a constant r or a tagged variable. A tagged variable ω represents a symbolic variable α (to be interpreted over the $[0, 1]$ interval), which for convenience is tagged by a pair of natural numbers ℓ and \mathcal{r} . For simplicity, we adopt the following notation conventions. First, we use $\omega.\alpha, \omega.\ell$ and $\omega.\mathcal{r}$ to access the first, second and third component of ω , respectively. Second, given $\omega = \langle \alpha, i, j \rangle$, we use $\omega(i, j)$ as a shorthand for α . Third, in order not to clutter formulas, we sometime write ω for $\omega.\alpha$ (e.g., $\omega_1 + \omega_2 = 1$ for $\omega_1.\alpha + \omega_2.\alpha = 1$). Finally, when clear from the context, we refer to tagged variables simply as variables. An *expression* φ represents either a symbolic probability ϱ or algebraic operations (addition, subtraction, multiplication or division) between symbolic probabilities. A *formula* Φ is either a comparison between two expressions or the conjunction of other two formulas. *Formula simple types* are defined similarly to gradual simple types (including the ? type), except that the codomain of function types are formula distribution types. Formula distribution types are now multi-sets of pairs of formula simple types and symbolic probabilities, closed under a formula Φ .

Let us introduce some handy notation for the rest of the presentation. First, given formula Φ , we use $FV(\Phi)$ to denote the set of (free) tagged variables occurring in Φ . Second, given a set of symbolic probabilities $\{\varrho_i \mid i \in \mathcal{I}\}$, we use $TV(\{\varrho_i \mid i \in \mathcal{I}\})$ to denote the subset of tagged variables, only (i.e. the result of filtering out static probabilities). Lastly, given formula Φ over tagged variables $\omega_1, \dots, \omega_n$, we use $\text{sat}(\Phi)$ to denote that Φ is satisfiable, i.e. as a shorthand for $\exists \omega_1, \dots, \exists \omega_n. \Phi$.

There is a canonical lifting from gradual types to formula gradual types. For example, the gradual distribution type $\{\{\text{Int}^{\frac{1}{3}}, \text{Bool}^?, ?\}\}$ can be represented by the formula distribution type $\Phi \triangleright \{\{\text{Int}^{\omega_1}, \text{Bool}^{\omega_2}, ?^{\omega_3}\}\}$, where $\Phi = \omega_1 = \frac{1}{3} \wedge \omega_2 \in [0, 1] \wedge \omega_3 \in [0, 1] \wedge \omega_1 + \omega_2 + \omega_3 = 1$, and $\varrho \in [r_1, r_2]$ is syntactic sugar for $r_1 \leq \varrho \wedge \varrho \leq r_2$. Formally, the lifting is captured by three mutually recursive functions that act over gradual simple types, gradual probabilities and gradual distribution types respectively as follows:

$$\begin{aligned}
[\cdot] &: \text{GTYPE} \rightarrow \text{FSTYPE} \\
[\text{Real}] &= \text{Real} \quad [\text{Bool}] = \text{Bool} \quad [?] = ? \quad [\sigma \rightarrow \mu] = [\sigma] \rightarrow [\mu] \\
[\cdot] &: \text{GPROB} \times \text{TVAR} \rightarrow \text{FORMULA} \\
[p]_\omega &= (\omega = p) \quad [?]_\omega = (\omega \in [0, 1]) \\
[\cdot] &: \text{GDTYPE} \rightarrow \text{FDTYPE} \\
[\{\{\sigma_i^{\rho_i} \mid i \in \mathcal{I}\}\}] &= (\bigwedge_{i \in \mathcal{I}} [\rho_i]_{\omega_i} \wedge \sum_{i \in \mathcal{I}} \omega_i = 1) \triangleright \{\{[\sigma_i]_{\omega_i}^{\omega_i} \mid i \in \mathcal{I}\}\} \quad \omega_i = \langle \alpha_i, i, i \rangle, \alpha_i \text{ is fresh}
\end{aligned}$$

To give the inductive definition of type consistency (and other forthcoming notions), we require a variant of the traditional notion of coupling. This variant differs from the traditional definition (see Def. 4.2) in that it operates over *symbolic* probability distributions, where probabilities are given by logical variables rather than concrete numbers, and these variables are subject to given constraints.

Definition 4.4. (Coupling over symbolic distributions) Assume that $\mathcal{A} = \{\{a_i^{p_i} \mid i \in \mathcal{I}\}\}$ and $\mathcal{B} = \{\{b_j^{q_j} \mid j \in \mathcal{J}\}\}$ are (multi-set representations of) symbolic discrete probability distributions over sets A and B . Moreover, let $R \subseteq A \times B$ be a relation between A and B . Given $\mathcal{C} = \{\{\alpha_{ij} \mid i \in \mathcal{I} \wedge j \in \mathcal{J}\}\}$, constraint ψ_1 over p_i , and constraint ψ_2 over q_j , we use $\mathcal{C} \vdash \mathcal{A}^{\psi_1} R \mathcal{B}^{\psi_2}$ to denote the conjunction of conditions 1 and 2 from Def. 4.2 (i.e. that \mathcal{C} is a “traditional” coupling between \mathcal{A} and \mathcal{B}), together with $\psi_1 \wedge \psi_2$. We also use $L_R(\mathcal{A}^{\psi_1}, \mathcal{B}^{\psi_2})$ to denote formula $\exists\{p_i \mid i \in \mathcal{I}\} \cup \{q_j \mid j \in \mathcal{J}\} \cup \{\alpha_{ij} \mid i \in \mathcal{I} \wedge j \in \mathcal{J}\}$. $\mathcal{C} \vdash \mathcal{A}^{\psi_1} R \mathcal{B}^{\psi_2}$.

Note that $L_R(\mathcal{A}^{\psi_1}, \mathcal{B}^{\psi_2})$ requires the existence not only of a coupling \mathcal{C} , but also of concretizations of (symbolic distributions) \mathcal{A} and \mathcal{B} , respectively satisfying ψ_1 and ψ_2 . Typically, ψ_1 and ψ_2 will require that probabilities sum up to 1. Like in Definition 4.4, in the rest of the presentation we allow ourselves some abuse of notation and write $\exists\{x_1, x_2, \dots, x_n\}$ as a shorthand for $\exists x_1. \exists x_2. \dots \exists x_n$, and similarly for $\forall\{x_1, x_2, \dots, x_n\}$.

Armed with the above notion of relation lifting, we can define the lifting of any relation R over gradual simple types to gradual distribution types as:

$$L_R\left(\Phi_1 \triangleright \{\{\sigma_i^{o_i} \mid i \in \mathcal{I}\}\}, \Phi_2 \triangleright \{\{\sigma_j^{o_j} \mid j \in \mathcal{J}\}\}\right) \quad \text{iff} \quad L_R\left(\{\{\sigma_i^{o_i} \mid i \in \mathcal{I}\}\}^{\Phi_1}, \{\{\sigma_j^{o_j} \mid j \in \mathcal{J}\}\}^{\Phi_2}\right)$$

Now, we can readily provide an inductive characterization of consistency, by simply lifting the definition of equality:

Definition 4.5 (Type consistency, inductively). The consistency relation \sim between gradual types and formula distribution types ($\mu, \nu \in \text{FDTYPE}$) is defined as follows:

$$\begin{array}{c} \text{Real} \sim \text{Real} \qquad \text{Bool} \sim \text{Bool} \qquad \sigma \sim ? \qquad ? \sim \sigma \\ \hline \frac{\sigma_1 \sim \sigma_2 \quad \mu_1 \sim \mu_2}{\sigma_1 \rightarrow \mu_1 \sim \sigma_2 \rightarrow \mu_2} \qquad \frac{[\mu_1] \sim [\mu_2]}{\mu_1 \sim \mu_2} \qquad \frac{L(\mu, \nu)}{\mu \sim \nu} \end{array}$$

As expected, the inductive definition of consistency (Def. 4.5) coincides with the one yielded by AGT (Def. 4.1):

LEMMA 4.6 (EQUIVALENCE OF CONSISTENCIES). *For any pair of gradual simple types $\sigma, \delta \in \text{GTYPE}$ and any pair of gradual distribution types $\mu, \nu \in \text{GDTYPE}$,*

$$\sigma \sim_{\text{AGT}} \delta \quad \text{iff} \quad \sigma \sim \delta \qquad \text{and} \qquad \mu \sim_{\text{AGT}} \nu \quad \text{iff} \quad \mu \sim \nu$$

Type well-formedness. Another relevant aspect of GPLC type system is that, like SPLC type system, programs are assigned well-formed types, only. The definition of well-formedness for gradual types is similar to that of static types, except that a gradual distribution type is well-formed iff it is *plausible* (rather than certain) that its underlying probabilities sum up to 1, and moreover, all its tagged variables occur in the closing formula:

Definition 4.7 (Type well-formedness). The well-formedness of gradual and formula types (denoted by symbol \vdash) is defined as follows:

$$\frac{}{\vdash \text{Real}} \quad \frac{}{\vdash \text{Bool}} \quad \frac{}{\vdash ?} \quad \frac{\vdash \sigma \quad \vdash \mu}{\vdash \sigma \rightarrow \mu} \quad \frac{\vdash [\mu]}{\vdash \mu}$$

$$\frac{TV(\{q_i \mid i \in \mathcal{I}\}) \subseteq FV(\Phi) \quad \text{sat}(\Phi \wedge \sum_{i \in \mathcal{I}} q_i = 1) \quad \forall i \in \mathcal{I}. \vdash \sigma_i}{\vdash \Phi \triangleright \{\{\sigma_i^{q_i} \mid i \in \mathcal{I}\}\}}$$

Note that while the first line of rules defines well-formedness for both gradual simple and gradual distribution types, the second line defines well-formedness for formula distribution types, only. Well-formedness for formula simple types follows the same rules as for gradual simple types (first four rules above).

LEMMA 4.8 (TYPE WELL-FORMEDNESS). *For any value v , any term m , any gradual simple type $\sigma \in \text{GTYPE}$ and gradual distribution type $\mu \in \text{GDTYPE}$ from GPLC, and any environment Γ ,*

(1) *If $\Gamma \vdash v : \sigma$, then $\vdash \sigma$*

(2) *If $\Gamma \vdash m : \mu$, then $\vdash \mu$*

An appealing property of the operator $[\cdot]$ lifting gradual distribution types to formula distribution types is that it preserves well-formedness:

LEMMA 4.9 (PRESERVATION OF TYPE WELL-FORMEDNESS). *For any gradual simple type $\sigma \in \text{GTYPE}$, and any gradual distribution type $\mu \in \text{GDTYPE}$,*

(1) *If $\vdash \sigma$, then $\vdash [\sigma]$*

(2) *If $\vdash \mu$, then $\vdash [\mu]$*

4.4 Refined Criteria

The refined criteria for gradual languages [Siek et al. 2015b] establish a set of distinguishing properties for such class of languages, where (only) two such properties are related to the static semantics: the *static gradual guarantee*, which guarantees that typing is monotone with respect to imprecision, and the *conservative extension of the static discipline*, which guarantees that every fully-statically-annotated well-typed term in the gradual language is also typeable in the static language (and vice versa). To establish the first property, the static gradual guarantee for GPLC, we first need to define a notion of precision between types, and subsequently between terms.

Type precision. AGT casts the definition of type precision in terms of set containment on the concretization of the gradual types, i.e. $G_1 \sqsubseteq G_2$ (meaning that gradual type G_1 is at least as precise as gradual type G_2) if and only if $\gamma(G_1) \subseteq \gamma(G_2)$. Nevertheless, in the presence of gradual distribution types, the definition based on set containment is not satisfactory as it assumes a *syntactic* equality between set elements. For instance, while $\{\{\text{Real}^1\}\} \sqsubseteq \{\{\text{Real}^{\frac{1}{2}}, \text{Real}^{\frac{1}{2}}\}\}$ is expected to hold since the involved pair of types are equal (under our semantic view of equality), a naive definition of precision would reject this relation. Therefore, we adopt an alternative definition of precision by [Lennon-Bertrand et al. 2022], which can be successfully applied when equality is not syntactic.

Definition 4.10 (Type precision). For any pair of gradual simple types $\sigma, \delta \in \text{GTYPE}$ and any pair of gradual distribution types $\mu, \nu \in \text{GDTYPE}$,

(1) $\sigma \sqsubseteq_{\text{AGT}} \delta$ if and only if $\forall \tau_1 \in \gamma_\tau(\sigma). \exists \tau_2 \in \gamma_\tau(\delta). \tau_1 = \tau_2$.

(2) $\mu \sqsubseteq_{\text{AGT}} \nu$ if and only if $\forall T_1 \in \gamma_T(\mu). \exists T_2 \in \gamma_T(\nu). T_1 = T_2$.

Like for consistency, the definition of precision above, despite being sound, is unpractical. We thus present an alternative, inductive characterization. This inductive characterization is rather standard, only the case of (gradual and formula) distribution types deserving special attention;

$$\frac{[\mu_1] \sqsubseteq [\mu_2]}{\mu_1 \sqsubseteq \mu_2} \quad \frac{\forall FV(\Phi_1). \Phi_1 \implies \exists FV(\Phi_2) \cup \{\omega_{ij} \mid i \in \mathcal{I} \wedge j \in \mathcal{J}\}. \quad \{\omega_{ij} \mid i \in \mathcal{I} \wedge j \in \mathcal{J}\} \vdash \{\sigma_i^{oi} \mid i \in \mathcal{I}\}^{\Phi_1} \sqsubseteq \{\sigma_j^{oj} \mid j \in \mathcal{J}\}^{\Phi_2}}{\Phi_1 \triangleright \{\sigma_i^{oi} \mid i \in \mathcal{I}\} \sqsubseteq \Phi_2 \triangleright \{\sigma_j^{oj} \mid j \in \mathcal{J}\}}$$

Fig. 8. Type precision in GPLC (excerpt).

$$\frac{m \sqsubseteq n \quad \mu \sqsubseteq \nu}{m :: \mu \sqsubseteq n :: \nu} \quad \frac{}{p \sqsubseteq p} \quad \frac{}{\rho \sqsubseteq ?} \quad \frac{m \sqsubseteq m' \quad n \sqsubseteq n' \quad \rho \sqsubseteq \rho'}{m \oplus_\rho n \sqsubseteq m' \oplus_{\rho'} n'}$$

Fig. 9. Term precision in GPLC (excerpt).

see Figure 8. Two gradual distribution types are in precision if their lifting to formula distribution types are in precision. Precision for formula distribution types is slightly different from consistency. Loosely speaking, formula distribution types μ_1 and μ_2 are related by precision iff every solution that makes the probabilities of μ_1 sum up to 1 can be “completed” to form a coupling between μ_1 and μ_2 that witnesses the lifting of precision. Intuitively, the definition is designed to reproduce the quantifier structure of Definition 4.10.

To illustrate how this new definition of type precision works, consider the following examples:

- $\{\{\text{Real}^{\frac{1}{2}}, ?^{\frac{1}{2}}\} \sqsubseteq \{\{\text{Real}^{\frac{1}{3}}, \text{Real}^{\frac{1}{6}}, ?^{\frac{1}{2}}\}\}$ holds because $\omega_{11} + \omega_{12} + \omega_{13} = \frac{1}{2} \wedge \omega_{21} + \omega_{22} + \omega_{23} = \frac{1}{2} \wedge \omega_{11} + \omega_{21} = \frac{1}{3} \wedge \omega_{12} + \omega_{22} = \frac{1}{6} \wedge \omega_{13} + \omega_{23} = \frac{1}{2}$ is satisfiable by the solution set $\{0 \leq \omega_{11} \leq \frac{1}{3}, 0 \leq \omega_{12} \leq \frac{1}{6}, 0 \leq \omega_{13} \leq \frac{1}{2}, 0 \leq \omega_{21} \leq \frac{1}{3}, 0 \leq \omega_{22} \leq \frac{1}{6}\}$.
- $\{\{\text{Real}^{\frac{1}{2}}, ?^{\frac{1}{2}}\} \not\sqsubseteq \{\{\text{Bool}^{\frac{2}{3}}, ?^{\frac{1}{3}}\}\}$ does not hold because $\omega_{12} = \frac{1}{2} \wedge \omega_{21} + \omega_{22} = \frac{1}{2} \wedge \omega_{21} = \frac{2}{3} \wedge \omega_{12} + \omega_{22} = \frac{1}{3}$ is not satisfiable.

As already hinted, this inductive definition of precision is equivalent to Definition 4.10:

LEMMA 4.11 (EQUIVALENCE OF TYPE PRECISION). *For any pair of gradual simple types $\sigma, \delta \in \text{GTYPE}$ and any pair of gradual distribution types $\mu, \nu \in \text{GDTYPE}$,*

$$\sigma \sqsubseteq_{\text{AGT}} \delta \text{ iff } \sigma \sqsubseteq \delta \quad \text{and} \quad \mu \sqsubseteq_{\text{AGT}} \nu \text{ iff } \mu \sqsubseteq \nu$$

Term precision. Term precision is the natural lifting of type precision to the space of terms. Its definition is rather standard, by induction in the term structure, as presented in Figure 9.

Metatheory. Armed with the definition of precision, we can now state the two fundamental properties that hold for the static semantics of GPLC. First, typeability is monotone w.r.t. imprecision:

THEOREM 4.12 (STATIC GRADUAL GUARANTEE FOR GPLC). *For every value v , every term m , every gradual simple type σ and every gradual distribution type μ from GPLC,*

- (1) *If $\vdash v : \sigma$ and $m \sqsubseteq n$, then there exists δ such that $\vdash n : \delta$ and $\sigma \sqsubseteq \delta$.*
- (2) *If $\vdash m : \mu$ and $m \sqsubseteq n$, then there exists ν such that $\vdash n : \nu$ and $\mu \sqsubseteq \nu$.*

Second, the static semantics of SPLC and GPLC are equivalent for fully-statically-annotated terms:

THEOREM 4.13 (CONSERVATIVE EXTENSION OF THE STATIC SEMANTIC). *For every value v , every term m , every simple type τ and every distribution type T from SPLC,*

$$\vdash_s v : \tau \text{ iff } \vdash m : \tau \quad \text{and} \quad \vdash_s m : T \text{ iff } \vdash m : T$$

$$\begin{aligned}
& r \in \mathbb{R}, \quad b \in \mathbb{B}, \quad x \in \text{Var}, \quad \sigma \in \text{FSTYPE}, \quad \mu \in \text{FDTYPE} \\
m, n ::= & v \mid v w \mid \text{let } x = m \text{ in } n \mid m_{\varrho_1} \oplus_{\varrho_2}^{\Phi} n \mid \xi m :: \mu & (\text{terms}) \\
\varepsilon v :: \sigma ::= & \text{if } v \text{ then } m \text{ else } n \mid v + w \mid \mathbf{error}_{\mu} \\
u ::= & r \mid b \mid (\lambda x : \sigma. m) & (\text{raw values}) \\
v, w ::= & x \mid \varepsilon u :: \sigma \mid \mathbf{error}_{\sigma} & (\text{values}) \\
\mathcal{V} ::= & \{v_i^{\varrho_i} \mid i \in \mathcal{I}\} & (\text{distribution values})
\end{aligned}$$

Fig. 10. Syntax of TPLC (excerpt).

4.5 Dynamic Semantics

Traditionally, when designing gradual languages, the runtime semantics are not defined directly over the gradual source language. The program is translated or elaborated into a *cast calculus* program, inserting casts at the boundaries between static and dynamic typing, ensuring at runtime that no static assumptions are violated. If a static assumption is violated, then a runtime error is raised. This cast calculus is usually called the gradual target language. The dynamic semantics of GPLC are no exception: taking inspiration from AGT, we elaborate GPLC into an evidence-based gradual target language, where evidence play the role of casts that justify consistency judgments. The gradual target language for GPLC, dubbed TPLC, is presented next.

5 TPLC: GRADUAL TARGET LANGUAGE

In this section, we introduce TPLC, an evidence-based target language for GPLC. We start by presenting the static semantics, followed by the dynamic semantics, which relates programs to probability distributions over values. Finally, we establish type safety and two refine criteria for TPLC (dynamic counterparts of the refined criteria already established for GPLC): the gradual guarantee, and that the language is a conservative extension of SPLC, its static counterpart.

5.1 Static Semantics

The static semantics of TPLC differs from GPLC in five key aspects: (1) we use formula distribution types from the beginning, (2) consistency judgment are augmented with concrete type information (called *evidence*) that justifies judgment validity, (3) explicit ascriptions are incorporated along type derivations to push all consistency judgments to the ascription type rules, (4) ascriptions carry their underlying evidence to justify consistency transitivity at runtime, and (5) to simplify the reduction rules and proofs, all values are ascribed.

Syntax. Figure 10 presents the syntax of TPLC. Types are the formula types from GPLC (see Fig. 7). Terms are now annotated with formula simple types and formula distribution types introduced in previous section. The probabilistic choice operator $m_{\varrho_1} \oplus_{\varrho_2}^{\Phi} n$ is now annotated with variables ϱ_1 and ϱ_2 closed by formula Φ , corresponding to the probability of taking the left or right branch respectively. Ascriptions $\varepsilon v :: \sigma$ and $\xi m :: \mu$ are augmented with evidences, where ε is an evidence for a formula simple type consistency judgment, and ξ for a formula distribution type consistency judgment (both kind of evidences, to be defined later). A raw value is either a real number r , a constant b or a lambda abstraction $\lambda x : \sigma. m$. As previously mentioned, all values in GPLC become ascribed values in TPLC. Therefore, a value v is either a variable x , an ascribed raw value u , or a tagged error \mathbf{error}_{σ} . Note that in contrast to classical gradual approaches, in TPLC error is also a term, and can be either a redex (\mathbf{error}_{μ}) or a value (\mathbf{error}_{σ}). The main reason for this is to simplify the metatheory when accounting for probabilistic branches that may fail during runtime. Errors also carry type information related to the expected type of the expression in order to establish type

$$\boxed{\Gamma \vdash v : \sigma, \quad \Gamma \vdash m : \mu, \quad \Gamma \vdash \Phi \triangleright \mathcal{V} : \mu}$$

$$\begin{array}{c}
\text{(Gerr}_\sigma\text{)} \frac{\vdash \sigma}{\Gamma \vdash \mathbf{error}_\sigma : \sigma} \quad \text{(Gerr}_\mu\text{)} \frac{\vdash \mu}{\Gamma \vdash \mathbf{error}_\mu : \mu} \quad \text{(G::}\sigma\text{)} \frac{\Gamma \vdash v : \sigma \quad \varepsilon \vdash \sigma \sim \delta \quad \vdash \delta}{\Gamma \vdash \varepsilon v :: \delta : \{\delta^1\}} \\
\text{(Gapp)} \frac{\Gamma \vdash v : \sigma \rightarrow \mu \quad \Gamma \vdash w : \sigma}{\Gamma \vdash v w : \mu} \quad \text{(Glet)} \frac{\Gamma \vdash m : \Phi \triangleright \{\sigma_i^{Q_i} \mid i \in \mathcal{I}\} \quad \forall i \in \mathcal{I}. \Gamma, x : \sigma_i \vdash n : \mu_i}{\Gamma \vdash \text{let } x = m \text{ in } n : \Phi \triangleright \sum_{i \in \mathcal{I}} Q_i \cdot \mu_i} \\
\text{(G::}\mu\text{)} \frac{\Gamma \vdash m : \mu \quad \xi \vdash \mu \sim v \quad \vdash v}{\Gamma \vdash \xi m :: v : v} \quad \text{(G}\oplus\text{)} \frac{\Gamma \vdash m : \mu \quad \Gamma \vdash n : v \quad \text{sat}(\Phi \Rightarrow Q_1 + Q_2 = 1)}{\Gamma \vdash m \quad Q_1 \oplus_{Q_2} n : \Phi \triangleright Q_1 \cdot \mu + Q_2 \cdot v} \\
\text{(GV)} \frac{\forall i \in \mathcal{I}. \vdash v_i : \sigma_i}{\Gamma \vdash \Phi \triangleright \{v_i^{Q_i} \mid i \in \mathcal{I}\} : \Phi \triangleright \{\sigma_i^{Q_i} \mid i \in \mathcal{I}\}}
\end{array}$$

$$\begin{aligned}
Q \cdot \Phi \triangleright \{\sigma_i^{Q_i} \mid i \in \mathcal{I}\} &= \Phi \wedge (\bigwedge_{i \in \mathcal{I}} \omega_i = Q \cdot Q_i) \triangleright \{\sigma_i^{Q_i} \mid i \in \mathcal{I}\} & \omega_i = \langle \alpha_i, \omega_i.\ell, \omega_i.\mathcal{F} \rangle, \alpha_i \text{ fresh} \\
\Phi \triangleright \sum_{i \in \mathcal{I}} Q_i \triangleright \{\sigma_j^{Q_j} \mid j \in \mathcal{J}_i\} &= \Phi \wedge (\bigwedge_{i \in \mathcal{I}} \Phi_i) \wedge (\sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}_i} Q_j = 1) \triangleright \bigcup_{i \in \mathcal{I}} \{\sigma_j^{Q_j} \mid j \in \mathcal{J}_i\}
\end{aligned}$$

Fig. 11. Type system of TPLC (excerpt).

safety, and can be removed in a real implementation. Finally, a *distribution value* \mathcal{V} stands for a distribution over values v .

Type System. The type system of TPLC is presented in Figure 11. Compared to GPLC, the only rules that use consistency are the ascription rules, making all top-level constructors match in the remaining type rules. Rule (G \oplus) requires that the fact that formula Φ entails that probabilities Q_1 and Q_2 sum up to 1 be plausible. Note that formula Φ is also pushed as part of the constraints of the resulting type—the weighted sum between the branch types. Similarly, rule (Glet) scales each μ_i with variable Q_i , so Φ is pushed to the resulting distribution type to close the type. Consistency judgments are now justified by some evidence, written $\varepsilon \vdash \sigma \sim \delta$ (for simple types) and $\xi \vdash \mu \sim v$ (for distribution types). Intuitively, evidences ε and ξ correspond to the most precise type information that support the respective consistency judgment; we elaborate on this in Section 5.2.

The notion of consistency of formula types is defined in the same way as in GPLC:

Definition 5.1 (Type consistency of formula types). Type consistency over simple (FSTYPE) and distribution (FDTYPE) formula types is defined as follows:

$$\frac{}{\text{Real} \sim \text{Real}} \quad \frac{}{\text{Bool} \sim \text{Bool}} \quad \frac{}{\sigma \sim ?} \quad \frac{}{? \sim \sigma} \quad \frac{\sigma_1 \sim \sigma_2 \quad \mu_1 \sim \mu_2}{\sigma_1 \rightarrow \mu_1 \sim \sigma_2 \rightarrow \mu_2} \quad \frac{L_-(\mu, v)}{\mu \sim v}$$

The definition of well-formedness is defined identical to Def. 4.7, and omitted for brevity. Like in SPLC and GPLC, all well-typed terms type check to well-formed formula types:

LEMMA 5.2. *For any value v , any term m , any formula simple type $\sigma \in \text{FSTYPE}$ and formula distribution type $\mu \in \text{FDTYPE}$ from TPLC, and any environment Γ ,*

- (1) *If $\Gamma \vdash v : \sigma$, then $\vdash \sigma$*
- (2) *If $\Gamma \vdash m : \mu$, then $\vdash \mu$*

5.2 Evidence

Following AGT, evidences are encoded as pairs of (gradual) types of the form $\langle G_1, G_2 \rangle$. Intuitively, each type of the pair corresponds to a type in the consistent judgment that the evidence shall

justify, e.g. in $\langle G_1, G_2 \rangle \vdash G'_1 \sim G'_2$, G_1 corresponds to G'_1 and G_2 to G'_2 . Furthermore, each type in the evidence is at least as precise as its corresponding type in the consistent judgment, i.e. $G_1 \sqsubseteq G'_1$ and $G_2 \sqsubseteq G'_2$. When dealing with consistency (the gradual counterpart of equality), both types in evidence coincide, and therefore evidence is represented by single types, namely

$$\varepsilon ::= \sigma \quad (\text{simple evidences}) \qquad \xi ::= \mu \quad (\text{distribution evidences})$$

A simple evidence ε (resp. distribution evidence ξ) is just a formula simple type (resp. formula distribution type) that justifies a consistency judgment between two formula simple types (resp. two formula distribution types). Formally, an evidence justifies a consistency judgment if and only if the evidence is more precise than both types:

Definition 5.3 (Evidence). For all formula simple types $\varepsilon, \sigma, \delta \in \text{FSTYPE}$ and all formula distribution types $\xi, \mu, \nu \in \text{FDTYPE}$, we define

$$\varepsilon \vdash \sigma \sim \delta \iff \varepsilon \sqsubseteq \sigma \wedge \varepsilon \sqsubseteq \delta \qquad \text{and} \qquad \xi \vdash \mu \sim \nu \iff \xi \sqsubseteq \mu \wedge \xi \sqsubseteq \nu$$

For instance, $\text{Real} \rightarrow \text{Bool} \vdash \text{Real} \rightarrow ? \sim ? \rightarrow \text{Bool}$. Now we can justify the role of tags in tagged variables. Consider judgment $\Phi \triangleright \{\{\sigma_k^{\omega_k}\}\} \vdash \Phi_1 \triangleright \{\{\sigma_i^{\omega_i}\}\} \sim \Phi_2 \triangleright \{\{\delta_j^{\omega_j}\}\}$. Tagged variables connect evidence with their underlying types, i.e. type σ_k justifies that σ_i is consistent with δ_j , because $\omega_k.\mathcal{L} = i$ and $\omega_k.\mathcal{R} = j$, where $\omega_k.\alpha$ is the *weight* of the connection between σ_i and δ_j . Note that a pair of simple types can be connected through multiple evidences.

As usual in gradual languages, consistency is not transitive, e.g. $\{\{\text{Real}^{\frac{2}{3}}, \text{Bool}^{\frac{1}{3}}\}\} \sim \{\{?\}\}$ and $\{\{?\}\} \sim \{\{\text{Real}^{\frac{1}{3}}, \text{Bool}^{\frac{2}{3}}\}\}$, but $\{\{\text{Real}^{\frac{2}{3}}, \text{Bool}^{\frac{1}{3}}\}\} \not\sim \{\{\text{Real}^{\frac{1}{3}}, \text{Bool}^{\frac{2}{3}}\}\}$. Therefore, during runtime, evidence is combined to try to justify transitivity. If the combination succeeds, the resulting (and possible more precise) evidence justifies the resulting judgment from transitivity, otherwise a runtime error is raised. The combination of evidence is formalized using the *consistent transitivity* operator, which coincides with the meet (least upper bound) operator w.r.t. the precision order, i.e.

$$\varepsilon_1 \circ \varepsilon_2 = \varepsilon_1 \sqcap \varepsilon_2 \qquad \text{and} \qquad \xi_1 \circ \xi_2 = \xi_1 \sqcap \xi_2$$

The meet operator is partial. For simple types, it is defined by the following clauses:

$$\text{Real} \sqcap \text{Real} = \text{Real} \qquad \text{Bool} \sqcap \text{Bool} = \text{Bool} \qquad ? \sqcap \sigma = \sigma \qquad \sigma \sqcap ? = \sigma$$

$$\sigma_1 \rightarrow \mu_1 \sqcap \sigma_2 \rightarrow \mu_2 = \sigma_1 \sqcap \sigma_2 \rightarrow \mu_1 \sqcap \mu_2$$

For distribution types, the definition follows the same approach as used for defining consistency and precision, in terms of the existence of couplings that justifies the lifting, but explicitly capturing all witness couplings. Formally,

$$\mu_1 \sqcap \mu_2 = W_{\sqcap}(\mu_1, \mu_2)$$

where $W: (\text{FSTYPE} \times \text{FSTYPE} \rightarrow \text{FSTYPE}) \times (\text{FDTYPE} \times \text{FDTYPE} \rightarrow \text{FDTYPE}) \rightarrow \text{FDTYPE}$ returns (a characterization of) all couplings that witness the lifting, and is defined as:

$$W_f(\Phi_1 \triangleright \{\{\sigma_i^{\omega_i} \mid i \in \mathcal{I}\}\}, \Phi_2 \triangleright \{\{\delta_j^{\omega_j} \mid j \in \mathcal{J}\}\}) = \\ \Phi \triangleright \{\{f(\sigma_i, \delta_j)^{\omega_{ij}} \mid (i, j) \in \mathcal{I} \times \mathcal{J} \wedge (\sigma_i, \delta_j) \in \text{dom}(f)\}\} \qquad \omega_{ij} \text{ fresh}$$

provided $\exists FV(\Phi_1) \cup FV(\Phi_2) \cup \{\omega_{ij} \mid (i, j) \in \mathcal{I} \times \mathcal{J}\}$. $\Phi' \wedge \Phi$, where $\Phi' = \forall i \in \mathcal{I}. \forall j \in \mathcal{J}. \omega_{ij}.\mathcal{L} = \omega_i.\mathcal{L} \wedge \omega_{ij}.\mathcal{R} = \omega_j.\mathcal{R}$, $\Phi = \{\{\omega_{ij} \mid (i, j) \in \mathcal{I} \times \mathcal{J}\}\} \vdash \{\{\sigma_i^{\omega_i} \mid i \in \mathcal{I}\}\}^{\Phi_1} R \{\{\delta_j^{\omega_j} \mid j \in \mathcal{J}\}\}^{\Phi_2}$ and $\sigma_i R \delta_j$ iff $(\sigma_i, \delta_j) \in \text{dom}(f)$.⁶

⁶Note that Φ can be cast as a FORMULA by taking $\mathcal{X} = \{(i, j) \mid (\sigma_i, \delta_j) \in \text{dom}(f)\}$ as the index set of the witness couplings.

As an example of the meet between formula distribution types, observe that $(\omega_1 = \frac{1}{2} \wedge \omega_2 = \frac{1}{2}) \triangleright \{ \{ (\text{Real} \rightarrow ?)^{\omega_1}, (? \rightarrow \text{Real})^{\omega_2} \} \sqcap (\omega_3 = \frac{1}{3} \wedge \omega_4 = \frac{2}{3}) \triangleright \{ \{ (\text{Real} \rightarrow ?)^{\omega_3}, (? \rightarrow \text{Real})^{\omega_4} \} = (\omega_{11} + \omega_{21} = \omega_1 \wedge \omega_{22} = \omega_2 \wedge \omega_{11} = \omega_3 \wedge \omega_{21} + \omega_{22} = \omega_4) \triangleright \{ \{ (\text{Real} \rightarrow ?)^{\omega_{11}}, (\text{Real} \rightarrow \text{Real})^{\omega_{12}}, (\text{Real} \rightarrow \text{Real})^{\omega_{21}}, (? \rightarrow \text{Real})^{\omega_{22}} \} \}$ because $\omega_{11} + \omega_{21} = \omega_1 \wedge \omega_{22} = \omega_2 \wedge \omega_{11} = \omega_3 \wedge \omega_{21} + \omega_{22} = \omega_4$ is satisfiable by the solution set $\{\omega_{11} = \frac{1}{2}, \omega_{21} = \frac{1}{6}, \omega_{22} = \frac{1}{3}\}$.

Importantly, the meet between a pair of types is at least as precise as either of them (and therefore, a valid evidence for their consistency).

LEMMA 5.4 (MONOTONICITY OF THE MEET OPERATOR). *For all formula simple types $\sigma_1, \sigma_2, \sigma_3 \in \text{FSTYPE}$ and all formula distribution types $\mu_1, \mu_2, \mu_3 \in \text{FDTYPE}$,*

- (1) If $\sigma_3 = \sigma_1 \sqcap \sigma_2$, then $\sigma_3 \sqsubseteq \sigma_1 \wedge \sigma_3 \sqsubseteq \sigma_2$ (2) If $\mu_3 = \mu_1 \sqcap \mu_2$, then $\mu_3 \sqsubseteq \mu_1 \wedge \mu_3 \sqsubseteq \mu_2$

Armed with the definition of evidence and the meet operator, we can now state the following invariant for distribution evidences that crisply captures when an evidence is well-defined with respect to a judgment.

$$\begin{array}{c} \text{(wdB)} \frac{\varepsilon \in \{\text{Real}, \text{Bool}, ?\} \quad \varepsilon \sqsubseteq \sigma \quad \varepsilon \sqsubseteq \delta}{\varepsilon \vdash \sigma \sim \delta} \qquad \text{(wd}\rightarrow\text{)} \frac{\varepsilon \vdash \sigma \sim \delta \quad \xi \vdash \mu \sim \nu}{\varepsilon \rightarrow \xi \vdash \sigma \rightarrow \mu \sim \delta \rightarrow \nu} \\ \\ \text{(wd}\xi\text{)} \frac{\begin{array}{c} \exists FV(\Phi) \cup FV(\Phi_1) \cup FV(\Phi_2). \Phi \wedge \Phi_1 \wedge \Phi_2 \wedge \Phi_L \wedge \Phi_R \wedge \Phi_{\sim} \\ \Phi_L = \forall i \in \mathcal{I}. \sum_{k|\omega_k.\ell=i} \omega_k = \varrho_i \quad \Phi_R = \forall j \in \mathcal{J}. \sum_{k|\omega_k.r=j} \omega_k = \varrho_j \\ \Phi_{\sim} = \forall k \in \mathcal{K}. \omega_k > 0 \Rightarrow \sigma_k \vdash \sigma_{\omega_k.\ell} \sim \sigma_{\omega_k.r} \end{array}}{\Phi \triangleright \{ \{ \sigma_k^{\omega_k} \mid k \in \mathcal{K} \} \vdash \Phi_1 \triangleright \{ \{ \sigma_i^{\varrho_i} \mid i \in \mathcal{I} \} \sim \Phi_2 \triangleright \{ \{ \sigma_j^{\varrho_j} \mid j \in \mathcal{J} \} \}} \end{array}$$

The invariant for distribution evidences ensures that (1) the sum of all the weights connected to a simple type must be equal to the probability of that type, and (2) each evidence in the distribution evidence of weight larger than zero must be well-defined (and thus more precise than both types involved in the consistency judgment). Finally, we use this invariant, to validate that the consistent transitivity operator preserves the invariant.

LEMMA 5.5 (INVARIANT PRESERVATION). *For all formula simple types $\varepsilon_1, \varepsilon_2, \sigma_1, \sigma_2, \delta \in \text{FSTYPE}$ and all formula distribution types $\xi_1, \xi_2, \mu_1, \mu_2, \nu \in \text{FDTYPE}$,*

- (1) Let $\varepsilon_1 \vdash \sigma_1 \sim \delta$ and $\varepsilon_2 \vdash \delta \sim \sigma_2$. If $\varepsilon_1 \circ \varepsilon_2$ is defined, then $\varepsilon_1 \circ \varepsilon_2 \vdash \sigma_1 \sim \sigma_2$
(2) Let $\xi_1 \vdash \mu_1 \sim \nu$ and $\xi_2 \vdash \nu \sim \mu_2$. If $\xi_1 \circ \xi_2$ is defined, then $\xi_1 \circ \xi_2 \vdash \mu_1 \sim \mu_2$.

5.3 Dynamic Semantics

We now present the dynamic semantics for TPLC, which relates programs to probability distributions over final values, through a big-step reduction relation (like in SPLC). The adoption of a distribution semantics is key to establish the dynamic gradual guarantee (DGG), which requires that reduction be monotone with respect to imprecision. To see this, assume that we adopt a (e.g. sampling) semantics that relates programs to *individual* final values. Under this semantics, program $(\lambda x : \text{Bool}. (x :: ? + 1) \oplus_{\frac{1}{2}} \text{true})$ false reduces to true (with probability $\frac{1}{2}$), while the less precise program $(\lambda x : ? . (x :: ? + 1) \oplus_{\frac{1}{2}} \text{true})$ false reduces to an error (also with probability $\frac{1}{2}$), which contradicts the DGG. Nevertheless, we can recover the DGG by considering all possible program outcomes at the same time, via a distribution semantics.

The distribution semantics is presented in Figure 12. Reduction judgment $m \Downarrow_k \Phi \triangleright \mathcal{V}$ denotes that term m reduces to *distribution configuration* $\Phi \triangleright \mathcal{V}$ in k steps. The number of steps of reduction judgments are only required to establish the metatheory, and can be removed in a real implementation. Several rules are defined similarly to SPLC, but accounting for the fact that values

$$\boxed{m \Downarrow_k \Phi \triangleright \mathcal{V}} \quad \mathcal{V} ::= \{\{v_i^{\mathcal{O}i} \mid i \in \mathcal{I}\}\} \text{ (distribution values)}$$

$$\begin{array}{c}
\text{(Dapp)} \frac{\widetilde{\text{dom}}(\varepsilon)v :: \delta \Downarrow_1 \cdot \triangleright \{\{w^1\}\}}{\text{sub}(\widetilde{\text{cod}}(\varepsilon)m :: \mu), w, x \Downarrow_k \Phi \triangleright \mathcal{V}} \quad \text{(D}\oplus\text{)} \frac{m \Downarrow_{k_1} \Phi_1 \triangleright \mathcal{V}_1 \quad n \Downarrow_{k_2} \Phi_2 \triangleright \mathcal{V}_2}{\Phi' = \Phi_1 \wedge \Phi_2 \wedge \Phi} \\
\text{(Dlet)} \frac{m \Downarrow_{k_1} \{\{v_i^{\mathcal{O}i} \mid i \in \mathcal{I}\}\} \quad \forall i \in \mathcal{I}. \text{sub}(n, v_i, x) \Downarrow_{k_2} \Phi_i \triangleright \mathcal{V}_i}{\text{let } x = m \text{ in } n \Downarrow_{k_1+k_2+1} (\bigwedge_{i \in \mathcal{I}} \Phi_i) \triangleright \sum_{i \in \mathcal{I}} \varrho_i \cdot \mathcal{V}_i} \quad \text{(Dv)} \frac{}{v \Downarrow_1 \cdot \triangleright \{\{v^1\}\}} \\
\text{(Derr)} \frac{\mu = \Phi \triangleright \{\{\sigma_i^{\mathcal{O}i} \mid i \in \mathcal{I}\}\}}{\mathbf{error}_\mu \Downarrow_1 \Phi \triangleright \{\{\mathbf{error}_{\sigma_i}^{\mathcal{O}i} \mid i \in \mathcal{I}\}\}} \quad \text{(Dmon)} \frac{m \Downarrow_k \mathcal{V}}{m \Downarrow_{k+1} \mathcal{V}} \\
\text{(D::}\sigma\text{)} \frac{}{\varepsilon_2(\varepsilon_1 u :: \sigma) :: \delta \Downarrow_1 \cdot \triangleright \begin{cases} \{\{\varepsilon_3 u :: \delta\}^1\} & \text{If } \varepsilon_1 \circ \varepsilon_2 = \varepsilon_3 \\ \{\{\mathbf{error}_\sigma^1\}\} & \text{otherwise} \end{cases}} \\
\text{(D::}\mu\text{)} \frac{m \Downarrow_{k'} \Phi_1 \triangleright \{\{v_i^{\mathcal{O}i} \mid i \in \mathcal{I}\}\} \quad \vdash \Phi_1 \triangleright \{\{v_i^{\mathcal{O}i} \mid i \in \mathcal{I}\}\} : \mu' \quad \xi \vdash \mu \sim v \quad v = \Phi_3 \triangleright \{\{\delta_j^{\mathcal{O}j} \mid j \in \mathcal{J}\}\}}{\text{(}\xi m :: v\text{)} \Downarrow_{k'+1} \begin{cases} \Phi_2 \triangleright \sum_{k \in \mathcal{K}} \omega_k \cdot \mathcal{V}_k & \text{If } (\mu' \parallel \mu) \circ \xi = \Phi_2 \triangleright \{\{\varepsilon_k^{\omega_k} \mid k \in \mathcal{K}\}\}, \text{ where } \forall k \in \mathcal{K}, \\ i = \omega_k \cdot \ell \wedge j = \omega_k \cdot \mathcal{r} \implies (\varepsilon_k v_i :: \delta_j) \Downarrow_1 \cdot \triangleright \mathcal{V}_k \\ \{\{\mathbf{error}_v^1\}\} & \text{otherwise} \end{cases}} \\
m[v/x] : \text{TERM} \times \text{VALUE} \times \text{VAR} \rightarrow \text{TERM} \\
\text{sub}(m, \varepsilon u :: \sigma, x) = m[\varepsilon u :: \sigma/x] \quad \text{sub}(m, \mathbf{error}_\sigma, x) = \mathbf{error}_\mu \quad \text{where } x : \sigma \vdash m : \mu
\end{array}$$

Fig. 12. Distribution semantics of TPLC (excerpt).

are always ascribed. Rule *(Dapp)* first ascribes argument v to δ , appealing to transitivity with the domain of ε as evidence. After its reduction, the obtained value w is substituted for x in the body of the function using the auxiliary function sub . If during v reduction transitivity does not hold (and w is thus \mathbf{error}_δ), sub yields term \mathbf{error}_μ , μ being the expected distribution type of the application. Rule *(Dlet)* reduces subterm n by substituting x by all the possible outcomes of m , using also function sub to properly handle the case where one such outcome is an error. The so obtained distribution configurations are combined (distribution values via their weighted sum and formulas via their conjunction) to form the final outcome of the let-expression. Rule *(D \oplus)* reduces the pair of branches and combines their results like the *(Dlet)* rule, the major difference being that formula Φ is also included in the resulting distribution configuration. Rule *(Dv)* lifts values to (Dirac) distribution values. Rule *(Derr)* reduces an error over distribution type μ to a distribution of errors over simple types σ_i . Rule *(Dmon)* establishes the monotonicity of the reduction relation with respect to the step index. Rule *(D:: σ)* analyzes whether type of u is consistent with δ , combining the respective evidences through the consistent transitivity operator. The rule yields either a Dirac distribution of a newly ascribed value (if consistent transitivity succeeds), or (else) an error.

Rule *(D:: μ)* is the most challenging. Intuitively, it “pushes” simple evidences within ξ into the outcomes of m . However, note that pushing every simple evidence within ξ into every possible outcome of m is not what we want. For instance, given the (informal) program $\{\{\text{Real}^{\frac{1}{2}}, \text{Bool}^{\frac{1}{2}}\}\}(1_{\frac{1}{2}} \oplus_{\frac{1}{2}} \text{true}) :: \{\{\text{Real}^{\frac{1}{2}}, \text{Bool}^{\frac{1}{2}}\}\}$, it would be futile to push evidence Real into true , or Bool into 1 . Here, we have two

problems to address. First, to determine what evidences must be pushed into what values. Second, to determine the probability of each such combination. We address both problems simultaneously, taking advantage of the consistent transitivity operator, and a subsidiary relation over formula types we introduce next.

Observe that rule $(D::\mu)$ proceeds by first reducing m to a value distribution of type μ' . However, this μ' can be (syntactically) different from μ , the actual type of m . What we require here is that μ' be a *reordering* of μ . We thus introduce the reordering relation $\stackrel{\Gamma}{=}$ over formula types, which (for distribution types) is nothing more than the coupling lifting of the syntactic equality.

Definition 5.6 (Reordering). The reordering relation $\stackrel{\Gamma}{=}$ over formula simple and distribution types is defined by the following clauses:

$$\frac{}{\text{Real} \stackrel{\Gamma}{=} \text{Real}} \quad \frac{}{\text{Bool} \stackrel{\Gamma}{=} \text{Bool}} \quad \frac{}{? \stackrel{\Gamma}{=} ?} \quad \frac{\sigma_2 \stackrel{\Gamma}{=} \sigma_1 \quad \mu_1 \stackrel{\Gamma}{=} \mu_2}{\sigma_1 \rightarrow \mu_1 \stackrel{\Gamma}{=} \sigma_2 \rightarrow \mu_2} \quad \frac{L_{\stackrel{\Gamma}{=}}(\mu, v)}{\mu \stackrel{\Gamma}{=} v}$$

We can construct an initial evidence for reordering judgments similarly to the meet operator:

Definition 5.7 (Reordering initial evidence). The partial operator \parallel over formula simple and distribution types is defined as follows:

$$\begin{aligned} \text{Real} \parallel \text{Real} &= \text{Real} & \text{Bool} \parallel \text{Bool} &= \text{Bool} & ? \parallel ? &= ? \\ (\sigma_1 \rightarrow \mu_1 \parallel \sigma_2 \rightarrow \mu_2) &= (\sigma_1 \parallel \sigma_2) \rightarrow (\mu_1 \parallel \mu_2) & \mu_1 \parallel \mu_2 &= W_{\parallel}(\mu_1, \mu_2) \end{aligned}$$

LEMMA 5.8. For all formula simple types $\sigma_1, \sigma_2 \in \text{FSTYPE}$ and all formula distribution types $\mu_1, \mu_2 \in \text{FDTYPE}$,

- (1) If $\sigma_1 \stackrel{\Gamma}{=} \sigma_2$, then $\sigma_1 \parallel \sigma_2$ is defined, and $\sigma_1 \parallel \sigma_2 \vdash \sigma_1 \stackrel{\Gamma}{=} \sigma_2$.
- (2) If $\mu_1 \stackrel{\Gamma}{=} \mu_2$, then $\mu_1 \parallel \mu_2$ is defined, and $\mu_1 \parallel \mu_2 \vdash \mu_1 \stackrel{\Gamma}{=} \mu_2$.

Here, $\varepsilon \vdash \sigma_1 \stackrel{\Gamma}{=} \sigma_2$ means that evidence ε justifies reordering $\sigma_1 \stackrel{\Gamma}{=} \sigma_2$ and holds if $\varepsilon \sqsubseteq \sigma_1 \parallel \sigma_2$. The notion of evidence for the reordering between distribution types is defined analogously.

Reordering and consistency interact nicely, in that their evidences can be soundly combined:

LEMMA 5.9. For all formula simple types $\varepsilon, \varepsilon', \sigma, \sigma', \delta \in \text{FSTYPE}$ and all formula distribution types $\xi, \xi', \mu, \mu', v \in \text{FDTYPE}$,

- (1) If $\varepsilon \vdash \sigma \stackrel{\Gamma}{=} \sigma', \varepsilon' \vdash \sigma' \sim \delta$ and $\varepsilon \circ \varepsilon'$ is defined, then $\varepsilon \circ \varepsilon' \vdash \sigma \sim \delta$.
- (2) If $\xi \vdash \mu \stackrel{\Gamma}{=} \mu', \xi' \vdash \mu' \sim v$ and $\xi \circ \xi'$ is defined, then $\xi \circ \xi' \vdash \mu \sim v$.

Returning to rule $(D::\mu)$, observe that evidence $(\mu' \parallel \mu) \circ \xi$ addresses both of the mentioned problems: every simple evidence ε_k in $(\mu' \parallel \mu) \circ \xi$ connects a simple type from μ' with a simple type from v , via its corresponding weight ω_k . To form the final distribution value, evidence ε_k is pushed to the ascription of value $v_{\omega_k \cdot \ell}$ with type $\delta_{\omega_k \cdot \nu}$. After reducing these terms, the obtained distribution values are combined (through a weighted sum) to yield the final distribution value.

To illustrate this process, consider expression $\xi m :: \{\{?^{\frac{2}{3}}, \text{Real}^{\frac{1}{3}}\}\}$, where $\xi = (\omega_1(1, 1) = \frac{1}{6} \wedge \omega_2(1, 2) = \frac{1}{3} \wedge \omega_3(2, 1) = \frac{1}{2}) \triangleright \{\{\text{Real}^{\omega_1}, \text{Real}^{\omega_2}, \text{Bool}^{\omega_3}\}\}$ and $\xi \vdash \{\{\text{Real}^{\frac{1}{2}}, \text{Bool}^{\frac{1}{2}}\}\} \sim \{\{?^{\frac{2}{3}}, \text{Real}^{\frac{1}{3}}\}\}$.⁷ If $m \Downarrow_1 \{\{((\text{Bool})\text{true} :: \text{Bool})^{\frac{1}{2}}, ((\text{Real})1 :: \text{Real})^{\frac{1}{2}}\}\}$, then the $\xi' = (\omega'_1(1, 2) = \frac{1}{2} \wedge \omega'_2(2, 1) = \frac{1}{2}) \triangleright \{\{\text{Real}^{\omega'_1}, \text{Bool}^{\omega'_2}\}\}$. Notice that $\xi' \circ \xi = (\omega''_1(1, 1) = \frac{1}{2} \wedge \omega''_2(2, 1) = \frac{1}{6} \wedge \omega''_3(2, 2) = \frac{1}{3}) \triangleright$

⁷We have simplified the notation by using real numbers instead of variables (thus omitting formulas), and omitting some trivial ascriptions.

$$\begin{array}{c}
\boxed{\Gamma \vdash v : \sigma \rightsquigarrow v} \\
\boxed{\Gamma \vdash m : \mu \rightsquigarrow m} \\
\text{(Eapp)} \frac{\Gamma \vdash v : \sigma \rightsquigarrow v \quad \Gamma \vdash w : \delta \rightsquigarrow w \quad \delta \sim \widetilde{\text{dom}}(\sigma)}{\varepsilon_1 = [\delta] \sqcap [\widetilde{\text{dom}}(\sigma)] \quad \varepsilon_2 = [\sigma] \sqcap [\widetilde{\text{dom}}(\sigma) \rightarrow \text{cod}(\sigma)]} \\
\Gamma \vdash v w : \widetilde{\text{cod}}(\sigma) \rightsquigarrow \text{let } x = \varepsilon_1 w :: [\widetilde{\text{dom}}(\sigma)] \text{ in let } y = \varepsilon_2 v :: [\widetilde{\text{dom}}(\sigma) \rightarrow \text{cod}(\sigma)] \text{ in } y x \\
\text{(E}\oplus\text{)} \frac{\Gamma \vdash m : \mu \rightsquigarrow m \quad \Gamma \vdash n : \nu \rightsquigarrow n \quad \xi_1 = [\mu] \quad \xi_2 = [\nu]}{\Phi = \Phi_1 \wedge \Phi_2 \wedge (\omega_1 + \omega_2 = 1) \quad \xi = \Phi \vdash (\omega_1 \cdot \xi_1 + \omega_2 \cdot \xi_2) \sqcap [\rho \cdot \mu + (1 - \rho) \cdot \nu]} \\
\omega_1, \omega_2 \text{ fresh} \quad [\rho]_{\omega_1} = \Phi_1 \quad [(1 - \rho)]_{\omega_2} = \Phi_2 \\
\Gamma \vdash m \oplus_{\rho} n : \rho \cdot \mu + (1 - \rho) \cdot \nu \rightsquigarrow \xi m_{\omega_1} \oplus_{\omega_2}^{\Phi} n :: [\rho \cdot \mu + (1 - \rho) \cdot \nu] \\
\text{(E}\lambda\text{)} \frac{\Gamma, x : \sigma \vdash m : \mu \rightsquigarrow m \quad \varepsilon = [\sigma \rightarrow \mu] \quad \vdash \sigma}{\Gamma \vdash \lambda x : \sigma. m : \sigma \rightsquigarrow \varepsilon \lambda x : [\sigma]. m :: [\sigma \rightarrow \mu]} \quad \text{(E::}\mu\text{)} \frac{\Gamma \vdash m : \mu \rightsquigarrow m \quad \mu \sim \nu \quad \xi = [\mu] \sqcap [\nu] \quad \vdash \nu}{\Gamma \vdash m :: \nu : \nu \rightsquigarrow \xi m :: [\nu]}
\end{array}$$

Fig. 13. Elaboration from GPLC to TPLC (excerpt).

$\{\{\text{Bool}^{\omega_1''}, \text{Real}^{\omega_2''}, \text{Real}^{\omega_3''}\}\}$. Finally, the whole expression reduces to $\{\{((\text{Bool})\text{true} :: ?)^{\frac{1}{2}}, ((\text{Real})1 :: \text{Real})^{\frac{1}{6}}, ((\text{Real})1 :: \text{Real})^{\frac{1}{3}}\}\}$.

5.4 Elaboration

As previously mentioned, the runtime semantics of GPLC is given via translation to the TPLC target language. Figure 13 presents the type-driven elaboration rules from GPLC to TPLC. Judgment $\Gamma \vdash v : \sigma \rightsquigarrow v$ (resp. $\Gamma \vdash m : \mu \rightsquigarrow m$), denotes the elaboration of value v (resp. term m) from value v (resp. term m), where v (resp. m) is typed σ (resp. μ) under environment Γ . For simplicity, we write $v : \sigma \rightsquigarrow v$ (resp. $m : \mu \rightsquigarrow m$) as a shorthand for $\cdot \vdash v : \sigma \rightsquigarrow v$ (resp. $\cdot \vdash m : \mu \rightsquigarrow m$). Rule (E λ) and the ones for elaborating other values, elaborate by inserting ascriptions to their types. The initial evidence between two gradual types is computed using the meet of the lifted types. Rule (Eapp) insert ascriptions in both the function and argument to make top-level constructor match using A-normal form. Rule (E:: μ) produces initial evidence to justify the consistency judgment between μ and ν . Rule (E \oplus) is designed to carefully deal with the probability annotations. First, it generates two fresh variables: ω_1 for annotation ρ , and ω_2 for the complement $1 - \rho$.⁸ Second, these two fresh variables are used to generate two formulas Φ_1 and Φ_2 by lifting ρ and $1 - \rho$. Third, the annotation formula Φ is computed by combining Φ_1 and Φ_2 , with the extra requirement that the sum of the probability variables must be one (in case $\rho = ?$). Finally, we insert an extra ascription to relate variables ω_1 and ω_2 with the fresh variables obtained from lifting the type of the expression (as they will be different).

To conclude, we establish that the elaboration rules preserve typing:

THEOREM 5.10 (ELABORATION PRESERVE TYPES). *For every value v , term m , simple type σ and distribution type μ from GPLC,*

- (1) *If $\Gamma \vdash v : \sigma$, then there exists value v in TPLC such that $\Gamma \vdash v : \sigma \rightsquigarrow v$ and $[\Gamma] \vdash v : [\sigma]$.*
- (2) *If $\Gamma \vdash m : \mu$, then there exists term m in TPLC such that $\Gamma \vdash v : \mu \rightsquigarrow m$ and $[\Gamma] \vdash m : [\mu]$.*

Here, $[\Gamma]$ is the pointwise lifting of Γ .

⁸For ω_1 and ω_2 we do not care about the indexes because they do not flow into evidences. For this reason we can arbitrary choose 0 and 0 as default values.

$$\begin{array}{c}
\forall FV(\Phi_1). \Phi_1 \implies \exists FV(\Phi_2) \cup \{\omega_{ij} \mid i \in \mathcal{I} \wedge j \in \mathcal{J}\}. \\
\{\{\omega_{ij} \mid i \in \mathcal{I} \wedge j \in \mathcal{J}\} \vdash \{\{v_i^{\omega_i} \mid i \in \mathcal{I}\}\}^{\Phi_1} \sqsubseteq \{\{v_j^{\omega_j} \mid j \in \mathcal{J}\}\}^{\Phi_2}\} \\
(\sqsubseteq^{\mathcal{V}}) \hline
\Phi_1 \triangleright \{\{v_i^{\omega_i} \mid i \in \mathcal{I}\}\} \sqsubseteq \Phi_2 \triangleright \{\{v_j^{\omega_j} \mid j \in \mathcal{J}\}\} \\
\hline
\frac{\vdash m : \delta \quad \sigma \sqsubseteq \delta}{\mathbf{error}_\sigma \sqsubseteq m} \quad (\sqsubseteq \oplus) \frac{m \sqsubseteq m' \quad n \sqsubseteq n' \quad \forall FV(\Phi_1). \Phi_1 \implies \Phi_2}{m_{\varrho_1} \oplus_{\varrho_2}^{\Phi_1} n \sqsubseteq m'_{\varrho_1} \oplus_{\varrho_2}^{\Phi_2} n'}
\end{array}$$

Fig. 14. Term precision of TPLC (excerpt).

5.5 Type Safety and Gradual Guarantee

We can now establish several properties about GPLC, based on the elaboration to TPLC. To this end, we start highlighting that even though our static language (SPLC) is terminating, when introducing unknown types, one can encode statically well-typed programs that diverge, rendering our gradual languages (GPLC and TPLC) non-terminating. The emblematic program illustrating this phenomenon is the omega term $\Omega = (\lambda x : ?.x x)(\lambda x : ?.x x)$.

In the remaining of this section, for simplicity, given m we write $m \Downarrow \Phi \triangleright \mathcal{V}$ if $\vdash m : \mu \rightsquigarrow m$ (for some μ and m) and there exists k such that $m \Downarrow_k \Phi \triangleright \mathcal{V}$ (for some Φ and \mathcal{V}), and $m \Uparrow$ if $\vdash m : \mu \rightsquigarrow m$ (for some μ and m) and there exists no such k . In the former case we say that m (similarly, m) *terminates*, reducing to a distribution value, while in the latter case we say that m (similarly, m) *diverges* (also denoted by $m \Uparrow$). Formally, this terminology corresponds to the notion of *certain* termination, where, intuitively, a program is considered terminating if there is a bound on the length of all its executions.

We note that probabilistic programs also support more general notions of termination, such as *almost-sure* termination, which intuitively allows diverging executions, but requires them to have an overall null probability. Support for reasoning about this class of programs is left as future work.

Type Safety. Following [Lennon-Bertrand et al. 2022], we model errors (\mathbf{error}_σ and \mathbf{error}_μ) as expressions, simplifying this way the statement of type safety, as we do not need to reason about error separately. Type safety for GPLC then states that if a term m is well-typed, then it either reduces to a distribution value of an equivalent type, or diverges:

THEOREM 5.11 (TYPE SAFETY FOR GPLC). *For every term m and gradual distribution type μ from GPLC, if $\vdash m : \mu$ then either*

$$(1) m \Downarrow \Phi \triangleright \mathcal{V}, \vdash \Phi \triangleright \mathcal{V} : \mu \text{ and } \mu \stackrel{\tau}{=} [\mu] \text{ for some } \Phi, \mathcal{V} \text{ and } \mu, \text{ or} \quad (2) m \Uparrow.$$

Dynamic Gradual Guarantee. To establish the dynamic gradual guarantee (DGG) for GPLC, we start by establishing the DGG for TPLC. This requires defining the notion of type and term precision for TPLC. Type precision is defined in the same way as for GPLC (see Figure. 8). Term precision is the natural lifting of type precision to terms, and is defined in Figure. 14. Rule $(\sqsubseteq^{\mathcal{V}})$ relates two value distributions by lifting the precision relation on values to distributions via couplings, similarly to type precision. Like in [Lennon-Bertrand et al. 2022; New et al. 2019], an \mathbf{error}_σ is more precise than any term provided σ is more precise than the term type. Rule $(\sqsubseteq \oplus)$ relates two probabilistic choices if the corresponding subterms are in precision relation, and most importantly, the formula in the more precise probabilistic choice entails the formula in the less precise probabilistic choice. The notion of entailment over FORMULA is rather standard, and thus omitted. Finally, note that the pair of probabilistic choices share the same variable names after alpha renaming. To illustrate this rule, assume that we want to show that a probabilistic choice with a static probability

$$\begin{aligned}
\mathcal{V}[\text{Real}] &= \{(r, \varepsilon r :: \text{Real}) \in \text{Atom}[\text{Real}] \mid r = r\} \\
\mathcal{V}[\tau \rightarrow T] &= \{(v_1, v_2) \in \text{Atom}[\tau \rightarrow T] \mid \forall (v'_1, v'_2) \in \mathcal{V}[\tau]. (v_1 v'_1, v_2 v'_2) \in \mathcal{T}[[T]]\} \\
\mathcal{V}[T] &= \{(\mathcal{V}, \mathcal{V}) \mid \mathcal{V} = \{v_i^{p_i} \mid i \in \mathcal{I}\} \wedge \mathcal{V} = \{v_j^{p_j} \mid j \in \mathcal{J}\} \wedge \exists \xi = \{\tau_k^{\omega_k} \mid k \in \mathcal{K}\}. \\
&\quad (\xi, \mathcal{V}, \mathcal{V}) \in \text{Atom}[T] \wedge \forall \omega_k > 0, i = \omega_k \cdot \ell, j = \omega_k \cdot r. (v_i, v_j) \in \mathcal{V}[[\tau_k]]\} \\
\mathcal{T}[[T]] &= \{(m_1, m_2) \mid m_1 \Downarrow_s^* \mathcal{V}_1 \wedge m_2 \Downarrow_s^* \mathcal{V}_2 \wedge (\mathcal{V}_1, \mathcal{V}_2) \in \mathcal{V}[[T]]\} \\
\mathcal{G}[\Gamma, x : \tau] &= \{\gamma[(v, v')/x] \mid \gamma \in \mathcal{G}[\Gamma] \wedge (v, v') \in \mathcal{V}[[\tau]]\} \quad \mathcal{G}[\cdot] = \{\emptyset\} \\
\text{Atom}[\tau] &= \{(v, v') \mid \vdash_s v : \tau \wedge \vdash_s v' : \tau\} \\
\text{Atom}[T] &= \{(\xi, \mathcal{V}, \mathcal{V}) \mid \vdash_s \mathcal{V} : T_1 \wedge \vdash_s \mathcal{V} : T_2 \wedge \xi \vdash T_1 \stackrel{\dagger}{=} T_2 \wedge T \stackrel{\dagger}{=} T_1 \wedge T \stackrel{\dagger}{=} T_2\} \\
\Gamma \vdash m_1 \approx m_2 : T &\iff \forall (\gamma_1, \gamma_2) \in \mathcal{G}[\Gamma]. (\gamma_1(m_1), \gamma_2(m_2)) \in \mathcal{T}[[T]]
\end{aligned}$$

Fig. 15. Logical relation between SPLC and TPLC (excerpt).

$\frac{1}{2}$ is more precise than the one that we obtain replacing the static probability with ?. The rule application would then generate, as premise, the entailment $(\omega_1 = \frac{1}{2} \wedge \omega_2 = \frac{1}{2} \wedge \omega_1 + \omega_2 = 1) \Rightarrow (\omega_1 \in [0, 1] \wedge \omega_2 \in [0, 1] \wedge \omega_1 + \omega_2 = 1)$, with ω_1, ω_2 universally quantified. The remaining rules are standard.

Having defined the notion of precision, the major pending challenge to establish the DGG is to prove that evidence combination is monotone with respect to imprecision:

LEMMA 5.12 (MONOTONICITY OF EVIDENCE COMBINATION). *For all formula simple types $\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4 \in \text{FSTYPE}$ and all formula distribution types $\xi_1, \xi_2, \xi_3, \xi_4 \in \text{FDTYPE}$,*

- (1) *If $\varepsilon_1 \sqsubseteq \varepsilon_2, \varepsilon_3 \sqsubseteq \varepsilon_4$ and $\varepsilon_1 \circ \varepsilon_3$ is defined then $\varepsilon_1 \circ \varepsilon_3 \sqsubseteq \varepsilon_2 \circ \varepsilon_4$*
- (2) *If $\xi_1 \sqsubseteq \xi_2, \xi_3 \sqsubseteq \xi_4$ and $\xi_1 \circ \xi_3$ is defined then $\xi_1 \circ \xi_3 \sqsubseteq \xi_2 \circ \xi_4$*

Proof sketch. For simple evidences, the proof proceeds by routine induction. For distribution evidences, the proof requires some coupling combinations. Assume $\mathcal{C}_{12}^{\sqsubseteq} \vdash \xi_1 \sqsubseteq \xi_2, \mathcal{C}_{34}^{\sqsubseteq} \vdash \xi_3 \sqsubseteq \xi_4, \mathcal{C}_{13}^{\sqsubseteq} \vdash \xi_1 \circ \xi_3$. To prove that $\xi_2 \circ \xi_4$ is defined, from $\mathcal{C}_{12}^{\sqsubseteq}$ we build a coupling $\mathcal{C}_{21}^{\sqsupset} \vdash \xi_2 \sqsupset \xi_1$ and then use $\mathcal{C}_{21}^{\sqsupset} \circ \mathcal{C}_{13}^{\sqsubseteq} \circ \mathcal{C}_{34}^{\sqsubseteq}$ as witness coupling, where the coupling composition operator \circ is defined as:

$$\mathcal{C}_1 \circ \mathcal{C}_2 = \left\{ \omega_{k_1 k_2}(i, j) \mid \omega_{k_1 k_2}(i, j) = \frac{\omega_{k_1}(i, h) \omega_{k_2}(h, j)}{\sum_{i, k_1 \mid \omega_{k_1}(i, h)} \omega_{k_1}(i, h)}, \omega_{k_1}(i, h) \in \mathcal{C}_1, \omega_{k_2}(h, j) \in \mathcal{C}_2 \right\}$$

To justify that $\xi_1 \circ \xi_3 \sqsubseteq \xi_2 \circ \xi_4$, we start from coupling \mathcal{C}_{13}° and transform its first dimension (the one typically iterated by index \mathcal{I}) following the associations stated by coupling $\mathcal{C}_{12}^{\sqsubseteq}$ and its second dimension (typically iterated by index \mathcal{J}) following the associations stated by coupling $\mathcal{C}_{34}^{\sqsubseteq}$.

Now we can establish the DGG for TPLC: reduction is monotone with respect to imprecision.

THEOREM 5.13 (DYNAMIC GRADUAL GUARANTEE FOR TPLC). *Suppose $m \sqsubseteq n, \vdash m : \mu$ and $\vdash n : \nu$.*

- (1) *If $m \Downarrow_{k_1} \Phi_1 \triangleright \mathcal{V}_1$, then $n \Downarrow_{k_2} \Phi_2 \triangleright \mathcal{V}_2$, and $\Phi_1 \triangleright \mathcal{V}_1 \sqsubseteq \Phi_2 \triangleright \mathcal{V}_2$.*
- (2) *If $m \Uparrow$, then $n \Uparrow$.*

The DGG for GPLC is given by first elaborating the source terms to TPLC and then reducing the TPLC terms.

THEOREM 5.14 (DYNAMIC GRADUAL GUARANTEE FOR GPLC). *Suppose $m \sqsubseteq n, \vdash m : \mu$ and $\vdash n : \nu$.*

- (1) *If $m \Downarrow \Phi_1 \triangleright \mathcal{V}_1$, then $n \Downarrow \Phi_2 \triangleright \mathcal{V}_2$ and $\Phi_1 \triangleright \mathcal{V}_1 \sqsubseteq \Phi_2 \triangleright \mathcal{V}_2$.*
- (2) *If $m \Uparrow$, then $n \Uparrow$.*

Conservative extension of the dynamic semantics. In Section 4, we establish the equivalence between the static semantics of SPLC and GPLC for fully-statically-annotated terms. Here —due to the syntactic differences between both languages— to establish the equivalence between the dynamic semantics, we use logical relations. The logical relation between SPLC and TPLC is presented in Figure 15, and states that two related terms reduce to related distributions. Formally, it is defined using three mutually-defined interpretations: one for values ($\mathcal{V}[[\tau]]$), one for distribution values ($\mathcal{V}[[T]]$), and another one for terms or computations ($\mathcal{T}[[T]]$).

We write $(v_1, v_2) \in \mathcal{V}[[\tau]]$ to denote that values v_1 and v_2 are related at simple type τ . Two values are related at type τ if, first, they type check to τ , written $(v, v) \in \text{Atom}[\tau]$. Two booleans (resp. real numbers) are related when the underlying values are the same. Two functions are related if their application to related argument yields related computations.

Two distribution values $\mathcal{V}, \mathcal{V}'$ are related at a distribution type T if, first, there exists a distribution evidence (of fully-static types) ξ that justifies that their types⁹ are equivalent to (i.e. a reorder of) T , written $(\xi, \mathcal{V}, \mathcal{V}') \in \text{Atom}[T]$.¹⁰ Second, for all positive probabilities in the evidence (the coupling), the corresponding values must be related at the corresponding type.

Two computations are related if both reduce to related distribution values. Two value substitutions are related at some type environment, if every variable in the domain of the environments is bound to related values. Finally, two open terms are related if the substitution to any two related value environment yield related computations.

We can now establish the conservative extension of the dynamic semantics of TPLC with respect to SPLC for fully-annotated terms.

THEOREM 5.15 (DYNAMIC CONSERVATIVE EXTENSION OF TPLC W.R.T. SPLC).

$$(1) \vdash_s m : \tau, m \rightsquigarrow m' : \tau, \text{ then } \vdash m \approx m' : \tau \quad (2) \vdash_s m : \tau, m \rightsquigarrow m' : T, \text{ then } \vdash m \approx m' : T$$

The proof of Theorem 5.15 relies on the fact that the composition of static evidences is always defined:

LEMMA 5.16.

- (1) If $\varepsilon_1 \vdash \tau_1 \sim \tau_2$ and $\varepsilon_2 \vdash \tau_2 \sim \tau_3$, then $\varepsilon_1 \circ \varepsilon_2$ is defined, and $\varepsilon_1 \circ \varepsilon_2 \vdash \tau_1 \sim \tau_3$.
- (2) If $\xi_1 \vdash T_1 \sim T_2$ and $\xi_2 \vdash T_2 \sim T_3$, then $\xi_1 \circ \xi_2$ is defined, and $\xi_1 \circ \xi_2 \vdash T_1 \sim T_3$.

6 RELATED WORK

Gradual typing. As previously mentioned, gradual typing has been applied to many type discipline and language constructs. To the best of our knowledge, gradual typing has not been applied to probabilistic languages, neither to non-deterministic languages.

Lehmann and Tanter [2017] presented gradual refinement types, which allow the smoothly transition —and interoperability— between simple types and logically-refined types. In this work, we use statically-typed refinement types to implement cast/evidence, but do not support for gradual refinement types at the source level. Phipps-Costin et al. [2021] present TYPEWHICH, an approach for automatic type migration, which tries to infer additional or improved type annotations in gradually typed languages. Similarly to this work, TYPEWHICH also generates constrains (formulas) during type checking, and relies in an SMT solver to find solutions to their objectives.

There exist many flavors to define the runtime semantics of gradual languages. The classical approach is via a translation to a cast calculus [Garcia 2013; Herman et al. 2007, 2010; Siek et al.

⁹The type rule for \mathcal{V} is defined analogously to \mathcal{V}' , and can be found in the supplementary material.

¹⁰In TPLC, as the lifting of static types T always yields distribution types with one-to-one equality formulas, e.g. $\{\{\text{Real}^{\frac{1}{2}}, \text{Bool}^{\frac{1}{2}}\}\}$ is lifted as $(\omega_1 = \frac{1}{2} \wedge \omega_2 = \frac{1}{2}) \vdash \{\{\text{Real}^{\omega_1}, \text{Bool}^{\omega_2}\}\}$, for simplicity, to avoid writing variables, in the rule definition we annotate probabilities as numbers instead (i.e. $v_i^{p_i}$ and $v_j^{p_j}$).

2015a; Siek and Wadler 2010; Siek et al. 2009; Wadler and Findler 2009]; Garcia et al. [2016] defined the runtime semantics directly in the source language, by mimicking the proof normalization steps done in type safety; and recently, Ye et al. [2021] also presented direct dynamic semantics by using type-directed operational semantics (TDOS) [Huang and Oliveira 2020]. In this work, we follow the classical approach –defining a source and target language (cast calculus)–, where casts are implemented by using evidences from the AGT methodology.

There has been active work on designing gradual languages that allows the combination/collection of types. Castagna and Lanvin [2017]; Castagna et al. [2019] proposed a theory for gradual set-theoretic types, supporting union, intersection and the unknown type. In parallel, Toro and Tanter [2017] explored tagged and untagged union types, and Jafery and Dunfield [2017] sums types. Beside many fundamental differences, this work could be seen as a generalization of gradual union types with gradual weights.

Probabilistic λ -calculus. We can trace the origin of probabilistic λ -calculus to the work of [Saheb-Djahromi 1978], who present a typed, higher-order calculus. They develop a denotational semantics based on Plotkin’s probabilistic powerdomain [Jones and Plotkin 1989] and an operational semantics in terms of Markov chains. [Lago and Zorzi 2012b] surveys a variety of operational semantics for a λ -calculus with a probabilistic choice operator including small/big-step, inductive/coinductive and call-by-value/name variants. All fell under the category of distribution-based semantics, relating programs to probability distribution of values. [Ramsey and Pfeffer 2002] develop a denotational semantics for a stochastic λ -calculus exploiting the monadic structure of probability distributions. More recently, [Danos and Ehrhard 2011] and [Ehrhard et al. 2017] study a denotational semantics for higher-order programs in terms of coherence spaces.

Different type systems have been developed for probabilistic λ -calculi, aimed at establishing different program invariants. [Lago and Grellois 2017] use sized types to reason about almost-sure termination of higher-order programs. [Avanzini et al. 2021] develop a type system based on refinement types, to perform complexity analysis of higher-order functional programs. [Reed and Pierce 2010] (and many subsequent extensions) present a type system for reasoning about program sensitivity, used for established differential privacy properties of programs.

7 CONCLUSION

In this work, we provide a first step into the theoretical foundation of gradual probabilistic programming. We develop GPLC, to the best of our knowledge, the first gradual probabilistic language. The language enables an increased flexibility and expressivity, allowing some form of probabilistic specifications and also of program refinement via unknown probabilities in probabilistic choices. The development of GPLC is justified using the AGT methodology. The dynamic semantics of GPLC is given via translation to an evidence-based calculus, called TPLC, which features a distribution-based dynamic semantics. The development of GPLC and TPLC heavily relies in the notion of probabilistic coupling, as required for defining several relations and functions, such as type consistency, precision and consistent transitivity. As for the metatheory, GPLC satisfies type safety as well as the refined criteria for gradual languages.

As future work, we plan to explore the addition of more features to the language, such as subtyping and polymorphism. Introducing subtyping may bring several challenges, such as the use of sub-distributions. Another possible line of future work are the practical aspects of the gradual language, such as efficient handling of evidences in runtime, and space-efficient reduction rules.

REFERENCES

- Martin Avanzini, Ugo Dal Lago, and Alexis Ghyselen. 2021. Type-Based Complexity Analysis of Probabilistic Functional Programs. In *Proceedings of the 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '19)*. IEEE Press, Article 41, 13 pages.
- Arthur Azevedo de Amorim, Matt Fredrikson, and Limin Jia. 2020. Reconciling Noninterference and Gradual Typing. In *Proceedings of the 2020 Symposium on Logic in Computer Science (LICS 2020)*.
- Felipe Bañados Schwerter, Ronald Garcia, and Éric Tanter. 2016. Gradual Type-and-Effect Systems. *Journal of Functional Programming* 26 (Sept. 2016), 19:1–19:69.
- Gilles Barthe, Benjamin Grégoire, and Santiago Zanella-Béguelin. 2009. Formal Certification of Code-Based Cryptographic Proofs. In *36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'09)*. ACM, New York, 90–101.
- Giuseppe Castagna and Victor Lanvin. 2017. Gradual Typing with Union and Intersection Types. *Proceedings of the ACM on Programming Languages* 1, ICFP (Sept. 2017), 41:1–41:28.
- Giuseppe Castagna, Victor Lanvin, Tommaso Petrucci, and Jeremy G. Siek. 2019. Gradual typing: a new perspective. See [POPL 2019], 16:1–16:32.
- Guillaume Claret, Sriram K. Rajamani, Aditya V. Nori, Andrew D. Gordon, and Johannes Borgström. 2013. Bayesian Inference using Data Flow Analysis. In *Proceedings of the 9th Joint Meeting on Foundations of Software Engineering (ESEC/FSE 2013)*. ACM, 92–102.
- Vincent Danos and Thomas Ehrhard. 2011. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Information and Computation* 209, 6 (2011), 966–991. <https://doi.org/10.1016/j.ic.2011.02.001>
- Yuxin Deng and Wenjie Du. 2011. Logical, Metric, and Algorithmic Characterisations of Probabilistic Bisimulation. *CoRR* abs/1103.4577 (2011). arXiv:1103.4577 <http://arxiv.org/abs/1103.4577>
- Tim Disney and Cormac Flanagan. 2011. Gradual information flow typing. In *International Workshop on Scripts to Programs*.
- Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.* 9, 3-4 (2014), 211–407. <https://doi.org/10.1561/04000000042>
- Thomas Ehrhard, Michele Pagani, and Christine Tasson. 2017. Measurable Cones and Stable, Measurable Functions: A Model for Probabilistic Higher-Order Programming. *Proc. ACM Program. Lang.* 2, POPL, Article 59 (dec 2017), 28 pages. <https://doi.org/10.1145/3158147>
- Luminous Fennell and Peter Thiemann. 2013. Gradual Security Typing with References. In *Proceedings of the 26th Computer Security Foundations Symposium (CSF)*. 224–239.
- Ronald Garcia. 2013. Calculating threesomes, with blame. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional programming*. 417–428.
- Ronald Garcia, Alison M. Clark, and Éric Tanter. 2016. Abstracting Gradual Typing. In *Proceedings of the 43rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2016)*, Rastislav Bodík and Rupak Majumdar (Eds.). ACM Press, St Petersburg, FL, USA, 429–442. See erratum: <https://www.cs.ubc.ca/~rxg/agt-erratum.pdf>.
- Zoubin Ghahramani. 2015. Probabilistic Machine Learning and Artificial Intelligence. *Nature* 521, 7553 (2015), 452–459.
- Shafi Goldwasser and Silvio Micali. 1984. Probabilistic Encryption. *J. Comput. Sys. Sci.* 28, 2 (1984), 270–299.
- Noah D. Goodman, Vikash K. Mansinghka, Daniel Roy, Keith Bonawitz, and Joshua B. Tenenbaum. 2008. Church: A Language for Generative Models. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI'08)*. AUAI Press, 220–229.
- Noah D Goodman and Andreas Stuhlmüller. 2014. The Design and Implementation of Probabilistic Programming Languages. <http://dippl.org>. Accessed: 2022-10-17.
- Andrew D. Gordon, Thomas A. Henzinger, Aditya V. Nori, and Sriram K. Rajamani. 2014. Probabilistic programming. In *Proceedings of the on Future of Software Engineering, FOSE 2014*. ACM, 167–181.
- David Herman, Aaron Tomb, and Cormac Flanagan. 2007. Space-efficient gradual typing. In *In Trends in Functional Programming (TFP)*.
- David Herman, Aaron Tomb, and Cormac Flanagan. 2010. Space-efficient gradual typing. *Higher-Order and Sympolic Computation* 23, 2 (June 2010), 167–189.
- Xuejing Huang and Bruno C. d. S. Oliveira. 2020. A Type-Directed Operational Semantics For a Calculus with a Merge Operator. In *ECOOP*.
- Khurram A. Jafery and Jana Dunfield. 2017. Sums of Uncertainty: Refinements Go Gradual. See [POPL 2017], 804–817.
- C. Jones and Gordon D. Plotkin. 1989. A probabilistic powerdomain of evaluations. [1989] *Proceedings. Fourth Annual Symposium on Logic in Computer Science* (1989), 186–195.
- Oleg Kiselyov. 2016. Probabilistic Programming Language and its Incremental Evaluation. In *Programming Languages and Systems - 14th Asian Symposium, APLAS 2016, Hanoi, Vietnam, November 21-23, 2016, Proceedings (Lecture Notes in Computer Science)*, Atsushi Igarashi (Ed.), Vol. 10017. 357–376. https://doi.org/10.1007/978-3-319-47958-3_19

- Ugo Dal Lago and Charles Grellois. 2017. Probabilistic Termination by Monadic Affine Sized Typing. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 41 (2017), 1 – 65.
- Ugo Dal Lago and Margherita Zorzi. 2012a. Probabilistic operational semantics for the lambda calculus. *RAIRO Theor. Informatics Appl.* 46, 3 (2012), 413–450. <https://doi.org/10.1051/ita/2012012>
- Ugo Dal Lago and Margherita Zorzi. 2012b. Probabilistic operational semantics for the lambda calculus. *ArXiv abs/1104.0195* (2012).
- Tuan Anh Le, Atilim Gunes Baydin, and Frank D. Wood. 2017. Inference Compilation and Universal Probabilistic Programming. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA (Proceedings of Machine Learning Research)*, Aarti Singh and Xiaojin (Jerry) Zhu (Eds.), Vol. 54. PMLR, 1338–1348. <http://proceedings.mlr.press/v54/le17a.html>
- Nico Lehmann and Éric Tanter. 2017. Gradual Refinement Types, See [POPL 2017], 775–788.
- Meven Lennon-Bertrand, Kenji Maillard, Nicolas Tabareau, and Éric Tanter. 2022. Gradualizing the Calculus of Inductive Constructions. *ACM Transactions on Programming Languages and Systems* (2022). To appear. To be presented at POPL'22.
- Stefan Malewski, Michael Greenberg, and Éric Tanter. 2021. Gradually Structured Data. *Proceedings of the ACM on Programming Languages* 5, OOPSLA (Nov. 2021), 126:1–126:28.
- Rajeev Motwani and Prabhakar Raghavan. 1995. *Randomized Algorithms*. Cambridge University Press.
- Max S. New, Daniel R. Licata, and Amal Ahmed. 2019. Gradual Type Theory. See [POPL 2019], 15:1–15:31.
- Avi Pfeffer. 2010. Practical Probabilistic Programming. In *Inductive Logic Programming - 20th International Conference, ILP 2010, Florence, Italy, June 27-30, 2010. Revised Papers (Lecture Notes in Computer Science)*, Paolo Frasconi and Francesca A. Lisi (Eds.), Vol. 6489. Springer, 2–3.
- Luna Phipps-Costin, Carolyn Jane Anderson, Michael Greenberg, and Arjun Guha. 2021. Solver-based gradual type migration. *Proc. ACM Program. Lang.* 5, OOPSLA (2021), 1–27. <https://doi.org/10.1145/3485488>
- POPL 2017. *Proceedings of the 44th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2017)*. ACM Press, Paris, France.
- POPL 2019. *Proceedings of the 46th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2019)*.
- Norman Ramsey and Avi Pfeffer. 2002. Stochastic lambda calculus and monads of probability distributions. In *POPL '02*.
- Jason Reed and Benjamin C. Pierce. 2010. Distance Makes the Types Grow Stronger: A Calculus for Differential Privacy. In *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming (ICFP'10)*. ACM, 157–168.
- Amr Sabry and Matthias Felleisen. 1993. Reasoning about Programs in Continuation-Passing Style. In *LISP AND SYMBOLIC COMPUTATION*. 288–298.
- Nasser Saheb-Djahromi. 1978. Probabilistic LCF. In *MFCS*.
- Roberto Segala and Nancy A. Lynch. 1995. Probabilistic Simulations for Probabilistic Processes. *Nord. J. Comput.* 2, 2 (1995), 250–273.
- Jeremy Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. In *Proceedings of the Scheme and Functional Programming Workshop*. 81–92.
- Jeremy Siek, Peter Thiemann, and Phil Wadler. 2015a. Blame and Coercion: Together Again for the First Time. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2015)*. ACM Press, Portland, OR, USA, 425–435.
- 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, Ronald Garcia, and Walid Taha. 2009. Exploring the Design Space of Higher-Order Casts. In *ESOP*.
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015b. 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.
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, Sam Tobin-Hochstadt, and Ronald Garcia. 2015c. Monotonic References for Efficient Gradual Typing. In *ESOP*.
- Asumu Takikawa, T. Stephen Strickland, Christos Dimoulas, Sam Tobin-Hochstadt, and Matthias Felleisen. 2012. Gradual Typing for First-Class Classes. In *Proceedings of the 27th ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA 2012)*. ACM Press, Tucson, AZ, USA, 793–810.
- Matias 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.
- Matias Toro and Éric Tanter. 2017. A Gradual Interpretation of Union Types. In *Proceedings of the 24th Static Analysis Symposium (SAS 2017) (Lecture Notes in Computer Science)*, Vol. 10422. Springer-Verlag, New York City, NY, USA, 382–404.
- Matias Toro and Éric Tanter. 2020. Abstracting Gradual References. *Science of Computer Programming* 197 (Oct. 2020), 1–65.
- Dustin Tran, Matthew D. Hoffman, Rif A. Saurous, Eugene Brevdo, Kevin Murphy, and David M. Blei. 2017. Deep Probabilistic Programming. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017*,

Conference Track Proceedings. OpenReview.net. <https://openreview.net/forum?id=Hy6b4Pqee>

Jan-Willem van de Meent, Brooks Paige, Hongseok Yang, and Frank Wood. 2018. An Introduction to Probabilistic Programming. *CoRR* abs/1809.10756 (2018). arXiv:1809.10756 <http://arxiv.org/abs/1809.10756>

Philip Wadler and Robert Bruce Findler. 2009. Well-Typed Programs Can't Be Blamed. In *Proceedings of the 18th European Symposium on Programming Languages and Systems (ESOP 2009) (Lecture Notes in Computer Science)*, Giuseppe Castagna (Ed.), Vol. 5502. Springer-Verlag, York, UK, 1–16.

Wenjia Ye, Bruno C. d. S. Oliveira, and Xuejing Huang. 2021. Type-Directed Operational Semantics for Gradual Typing. In *ECOOP*.

Received 2022-10-28; accepted 2023-02-25