



HAL
open science

On Bisimilarity in Lambda Calculi with Continuous Probabilistic Choice

Ugo Dal Lago, Francesco Gavazzo

► **To cite this version:**

Ugo Dal Lago, Francesco Gavazzo. On Bisimilarity in Lambda Calculi with Continuous Probabilistic Choice. MFPS 2019 - 35th Conference on the Mathematical Foundations of Programming Semantics, Jun 2019, London, United Kingdom. hal-02386083

HAL Id: hal-02386083

<https://inria.hal.science/hal-02386083>

Submitted on 8 Jan 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On Bisimilarity in Lambda Calculi with Continuous Probabilistic Choice

Ugo Dal Lago^{*1} and Francesco Gavazzo²

¹*University of Bologna & INRIA Sophia Antipolis*

²*IMDEA Software Institute*

Abstract

Applicative bisimilarity is a coinductively-defined program equivalence in which programs are tested as argument-passing processes. Starting with the seminal work by Abramsky, applicative bisimilarity has been proved to be a powerful technique for higher-order program equivalence. Recently, applicative bisimilarity has also been generalised to lambda calculi with algebraic effects, and with discrete probabilistic choice in particular. In this paper, we show that applicative bisimilarity behaves well in a lambda-calculus in which probabilistic choice is available in a more general form, namely through an operator for sampling of values from *continuous* distributions. Our main result shows that applicative bisimilarity is sound for contextual equivalence, hence providing a new reasoning principle for higher-order probabilistic languages.

1 Introduction

Program equivalence and refinement are central concepts in program semantics: giving meaning to programs has the positive effect of allowing one to *compare* programs and to dub them *equivalent* whenever having the same semantics. But what comes first, the chicken or the egg? After all, *any* equivalence relation between program phrases implicitly gives meaning to a program through their equivalence class.

Among the many possible notions of program equivalence, Morris' style contextual equivalence [Morris \(1969\)](#) has been considered as somehow canonical, being the largest, and thus coarsest, adequate and compatible relation between programs: *by construction*, one cannot do any better than that when trying to identify programs behaving the same. As such, contextual equivalence is accepted as the reference notion of equivalence in an higher-order scenario. When proving that two programs are actually equivalent, however, the universal quantification over all contexts (i.e. over all program environments) in the definition of contextual equivalence makes the proof burdensome, if not impossible.

This has stimulated the development of alternative techniques for program equivalence which, in one way or another, *get rid of* the aforementioned universal quantification. Examples are logical relations [G. Plotkin \(1973\)](#); [Reynolds \(1983\)](#); [Sieber \(1992\)](#), applicative bisimilarity [Abramsky \(1990\)](#), and denotational semantics [Scott \(1970\)](#); [Scott and Strachey \(1971\)](#). Traditionally, the first step in studying the nature of these notions of program equivalence consists, again, in proving compatibility and adequacy, from which inclusion in contextual equivalence easily follows. In some cases this comes more or less by definition, while in other cases, like the one of applicative bisimilarity, proofs of compatibility are nontrivial, and are typically carried out in indirect ways, by methodologies like the so-called Howe's method [Howe \(1996\)](#). Starting from Abramsky's pioneering work on applicative bisimilarity [Abramsky \(1990\)](#), coinduction has been proved to be a useful methodology for program equivalence in the context of pure λ -calculi.

*Thanks to the ANR projects 14CE250005 ELICA and 16CE250011 REPAS.

All these (and others) techniques, give purely functional programming languages a clean semantics which makes it relatively easy to reason about them. Injecting various forms of effects (such as I/O, exceptions, or nondeterminism) complicates the matter considerably. Some of the aforementioned techniques for program equivalence have been proved to be remarkably robust in this respect, with contributions ranging from extensions encompassing *specific* notions of effect [Bizjak and Birkedal \(2015\)](#); [Dal Lago, Sangiorgi, and Alberti \(2014\)](#); [De Liguoro and Piperno \(1995\)](#); [Jones \(1990\)](#); [S. Lassen \(1998\)](#) to more abstract notions in which effects *of a certain kind* (e.g. algebraic effects [G. D. Plotkin and Power \(2001, 2002\)](#)) are proved to work well in a given semantic framework [Biernacki, Piróg, Polesiuk, and Sieczkowski \(2018\)](#); [Dal Lago, Gavazzo, and Levy \(2017\)](#); [Goubault-Larrecq, Lasota, and Nowak \(2008\)](#); [Johann, Simpson, and Voigtländer \(2010\)](#); [Simpson and Voorneveld \(2018\)](#).

One kind of effect which has received a lot of attention recently [Staton \(2017,?\)](#), but whose impact to the traditional techniques for program equivalence has been studied only marginally, is *continuous* probabilistic choice, i.e. the operation of sampling from continuous distributions. This by itself poses a challenge to the definition of an operational semantics, and thus to operationally-based techniques. Concerning the latter, the only contributions the authors are aware of are those by Culpepper et al. [Culpepper and Cobb \(2017\)](#); [Wand, Culpepper, Giannakopoulos, and Cobb \(2018\)](#) on logical relations for a probabilistic λ -calculus with continuous distributions and scoring.

In this paper, we show that applicative (bi)simulation remains a sound methodology for program equivalence and refinement in presence of sampling from continuous distributions. We do it for a λ -calculus endowed with a very liberal recursive type system.

Noticeably, probabilistic bisimulation has been first defined in the abstract setting of discrete probabilistic transition systems, also known as labelled Markov chains [Larsen and Skou \(1989\)](#), and later generalised to Markov decision processes [Blute, Desharnais, Edalat, and Panangaden \(1997\)](#). The main technical difficulty one faces when trying to turn probabilistic bisimulation into a notion of equivalence for higher-order *programs* consists in coming up with a sensible notion of *lifting*, allowing to turn a relation between programs into a relation between program (sub)distributions [Panangaden \(2011\)](#): there is a fundamental tension between the properties one typically gets from the aforementioned abstract notions of bisimilarity, and those one *needs* in the proofs of congruence. This paper shows that there is indeed a point in which the two meet.

The rest of the paper is organised as follows. We begin our analysis with an example of a simple program transformation. Proving correctness of such a transformation, however, is nontrivial and requires the design of suitable operational techniques. We then introduce our main vehicle calculus, namely a probabilistic call-by-value λ -calculus with sum and recursive types (Section 2), as well as its operational semantics (Section 2). Section 4 introduces the relational calculus we will use to study notions of program equivalence and refinement. We instantiate such a framework to define contextual equivalence and approximation (Section 5), and applicative (bi)simulation (Section 6). We prove soundness of applicative (bi)similarity in Section 6. We conclude this paper with a short discussion on full abstraction (Section 7).

1.1 Warming Up

Before entering into the technicalities of our analysis, we discuss a simple, yet nontrivial example of program equivalence in presence of continuous distributions. This way, we also introduce some features of the languages we are going to study.

We consider higher-order functional languages enriched with primitives for real-valued (measurable) functions and for sampling from continuous (and discrete) distributions. For instance, the program **normal** $\mu \sigma$ samples a real number from the normal distribution with mean (encoded by the numeral) μ and standard deviation σ . Similarly, we can write a program **poisson** μ which samples a real number from the Poisson distribution with rate μ .

We can also take advantage of the higher-order nature of functional languages, and use arbitrary programs

e_μ, e_σ to compute the mean and standard deviation of the normal distribution:

$$\text{let } x_\mu = e_\mu \text{ in } (\text{let } x_\sigma = e_\sigma \text{ in } (\text{normal } x_\mu \ x_\sigma^2)).$$

Notice that e_μ, e_σ may be defined sampling from other distributions as in:

$$\text{let } x_\mu = (\text{poisson } e_{\lambda_1}) \text{ in } (\text{let } x_\sigma = (\text{poisson } e_{\lambda_2}) \text{ in } (\text{normal } x_\mu \ x_\sigma^2)).$$

This immediately raises the question of whether the order in which we evaluate (independent) expressions matter. That is, we ask whether the equivalence

$$\text{let } x_1 = e_1 \text{ in } (\text{let } x_2 = e_2 \text{ in } e_3) \equiv \text{let } x_2 = e_2 \text{ in } (\text{let } x_1 = e_1 \text{ in } e_3)$$

holds, assuming the variable x_1 not to be free in e_2 , and similarity for x_2 and e_1 . Such an equivalence is known as *commutativity*, and has been studied in [Staton \(2017\)](#); [Staton, Yang, Wood, Heunen, and Kammar \(2016\)](#) by means of denotational methods, and in [Culpepper and Cobb \(2017\)](#); [Wand et al. \(2018\)](#) by means of logical relations (in both cases, however, the target language also includes primitives for conditioning).

Let us now look at a more sophisticated example. Consider the following instances of the combinators **map append**, and **zipwith**:

$$\begin{aligned} \text{map} &: (\text{unit} \rightarrow \text{real}) \rightarrow \text{list}(\text{unit}) \rightarrow \text{list}(\text{real}) \\ \text{append} &: \text{list}(\text{real}) \rightarrow \text{list}(\text{real}) \rightarrow \text{list}(\text{real}) \\ \text{zipwith} &: (\text{real} \rightarrow \text{real}) \rightarrow \text{list}(\text{real}) \rightarrow \text{list}(\text{real}) \rightarrow \text{list}(\text{real}) \end{aligned}$$

where **map** is the usual combinator for mapping functions on lists, **append** is the usual function for appending lists, and **zipwith** is a combinator taking as input a binary operation and two lists (which, for the sake of the argument, we assume to have the same length), and returning the list obtained by the pointwise application of the operation to the elements occurring at the same position in both lists. Let now zs be a list of n elements of type **unit** (we use zs to generate lists of random numbers), and let us write $\lambda.e : \text{unit} \rightarrow \text{real}$ for the thunk of a program $e : \text{real}$. Is the following program transformation, which avoids a (possibly expensive) list traversing correct?

$$\lambda x. \left(\begin{array}{l} \text{let } xs = \text{map } (\lambda. \text{normal } e_{\mu_1} \ e_{\sigma_1}^2) \ zs \\ \quad ys = \text{map } (\lambda. \text{normal } e_{\mu_2} \ e_{\sigma_2}^2) \ zs \\ \text{in } \text{append } x \ (\text{zipwith } (+) \ xs \ ys) \end{array} \right) \equiv \lambda x. \left(\begin{array}{l} \text{let } xs = \text{map } (\lambda. \text{normal } (e_{\mu_1} + e_{\mu_2}) \ (e_{\sigma_1}^2 + e_{\sigma_2}^2)) \ zs \\ \text{in } \text{append } x \ xs \end{array} \right)$$

Proving the correctness of the above transformation requires to combine denotational reasoning about probability measures with operational reasoning about program behaviour. On the one hand, the two programs above are values, so that looking at their computational behaviour in isolation is useless: in order to understand what these programs do we have to test them against any possible input the environment can pass them (which might have a probabilistic behaviour itself). On the other hand, such programs are somehow related by basic results in statistics (notably, given random variables with normal probability distribution $X_1 \sim \mathcal{N}(\mu_1, \sigma_1^2)$ and $X_2 \sim \mathcal{N}(\mu_2, \sigma_2^2)$, we have $X_1 + X_2 \sim \mathcal{N}(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$). As we will see, the techniques we develop in this paper give a handy proof of correctness of the above program transformation.

Now that the reader has some insights on the kind of problems we aim to solve in this work, we can start our formal analysis of program equivalence and refinement.

2 A Calculus with Continuous Probabilities: Syntax

The target language of this work is a probabilistic fine-grain call-by-value [Levy, Power, and Thielecke \(2003\)](#) with finite sums and recursive types reminiscent of probabilistic PCF [Ehrhard, Pagani, and Tasson \(2018\)](#); [Park, Pfenning, and Thrun \(2005\)](#). The syntax and static semantics of Λ_P is defined in Figure 1.

Types of Λ_P are built starting from a countable set of type variables (denoted by the letter α in Figure 1) and the basic type **real** for real numbers, using finite sum, arrow, and recursive type constructors. In particular, in a type of the form $\sum_{i \in I} \tau_i$ we assume the letter I to stand for a finite set whose elements are denoted by \hat{i}, \hat{j}, \dots . We write **void** and **unit** for the empty sum type and **void** \rightarrow **void**, respectively.

Terms of Λ_P are divided in two classes: values and expressions (also called computations). Intuitively, a value is the result of a computation, whereas an expression is a value producer, i.e. a term that once evaluated may produce a value (the evaluation process might not terminate) as well as side effects (the latter being probabilistic). Accordingly, computations must be explicitly sequenced by means of the sequencing constructor **let** $x = -$ **in** $-$.

We use judgments of the form $\Gamma \vdash^A e : \tau$ for expressions, and $\Gamma \vdash^V v : \tau$ for values. In a judgment of the form $\Gamma \vdash^A e : \tau$ (resp. $\Gamma \vdash^V v : \tau$), Γ denotes an environment, i.e. a finite sequence $x_1 : \tau_1, \dots, x_n : \tau_n$ of distinct variables with associated *closed* types (we denote by \cdot the empty environment), τ is a *closed* type, and e (resp. v) is an expression (resp. value) with free variables among Γ . Notice that we work with closed types only. We use the letter e (possibly with sup- and subscripts) to denote expressions, and v (possibly with sup- and subscripts) to denote values. We also refer to sequents $\Gamma \vdash^A \tau$ as computation or expression sequents, and to $\Gamma \vdash^V \tau$ as value sequents. When the distinction between values and expressions is not relevant, we generically refer to terms.

Each real number r is represented in Λ_P by the value r . Additionally, Λ_P is parametric with respect to a \mathbb{N} -indexed family C of *countable* sets C_n , each of which contains (symbols for) measurable functions from \mathbb{R}^n to \mathbb{R} . In particular, we assume standard real-valued arithmetic operations to be in C . We use letters F, G, \dots to denote elements in C_n . Notice that a function $F \in C_n$ has values as arguments, rather than expressions. Indeed, we can encode the expression $F(e_1, \dots, e_n)$ as¹ **let** $x_1 = e_1$ **in** (**let** $x_2 = e_2$ **in** \dots (**let** $x_n = e_n$ **in** $F(x_1, \dots, x_n)$)).

Finally, the expression **sample** stands for the uniform distribution over the unit interval, which is nothing but the Lebesgue measure λ on $[0, 1]$. As it is shown in e.g. Ehrhard et al. (2018); Park et al. (2005), starting from **sample** it is possible to define several probabilistic measures (e.g. binomial, geometric, and exponential distribution) using functions in C .

Example 1. Given closed terms e_μ, e_σ of type **real** (encoding mean μ and standard deviation σ) we represent the normal distribution with mean μ and standard deviation σ as the Λ_P expression **normal** $e_\mu e_\sigma$, where the term **normal**_{std} encodes the normal distribution with mean 0 and standard deviation 1. Notice that the expressions e_μ, e_σ may be themselves defined in terms of other distributions:

$$\begin{aligned} \mathbf{normal}_{\text{std}} &\triangleq \mathbf{let} \ x = \mathbf{sample} \ \mathbf{in} \ (\mathbf{let} \ y = \mathbf{sample} \ \mathbf{in} \ (\sqrt{-2 \log(x)} \cos(2\pi y))) \\ \mathbf{normal} \ e_\mu e_\sigma &\triangleq \mathbf{let} \ x_\mu = e_\mu, \ x_\sigma = e_\sigma, \ y = \mathbf{normal}_{\text{std}} \ \mathbf{in} \ ((x_\sigma * y) + x_\mu). \end{aligned}$$

We adopt the standard syntactical conventions as in Barendregt (1984), notably the so-called variable convention. In particular, we denote by $FV(e)$ (resp. $FV(v)$) the collection of free variables of e (resp. v) (notice that, e.g., in a term of the form **case** v **of** $\{\mathbf{fold} \ x \rightarrow e\}$ the variable x is bound in e). We refer to closed expressions as *programs*. We denote by $e[v/x]$ (resp. $v'[v/x]$) the capture-free substitution of the value v for all free occurrences of x in e (resp. v') and identify terms up to renaming of bound variables. We extend the aforementioned conventions to types. For instance, we denote by $\tau_1[\tau_2/\alpha]$ the result of capture-avoiding substitution of the type τ_2 for the type variable α in τ_1 . Finally, we use the notation Λ and \mathcal{V} to denote the collections of all typable expressions and values, respectively.

3 A Calculus with Continuous Probabilities: Semantics

The dynamic semantics of Λ_P is given by a type-indexed family of evaluation functions $\llbracket - \rrbracket_\tau$ mapping programs of type τ to sub-probability measures over values of type τ . The collection of such measures can be

¹This encoding reflects the standard semantics of operations, where to evaluate an expression of the form $F(e_1, \dots, e_n)$ one first sequentially evaluates its arguments, proceeding from left to right.

$\frac{}{\Gamma, x : \tau \vdash^V x : \tau}$	$\frac{\Gamma \vdash^V v : \tau}{\Gamma \vdash^\Lambda v : \tau}$	$\frac{r \in \mathbb{R}}{\Gamma \vdash^V r : \mathbf{real}}$	$\frac{}{\Gamma \vdash^\Lambda \mathbf{sample} : \mathbf{real}}$	$\frac{\Gamma \vdash^V v_1 : \mathbf{real} \quad \dots \quad \Gamma \vdash^V v_n : \mathbf{real} \quad F \in C_n}{\Gamma \vdash^\Lambda F(v_1, \dots, v_n) : \mathbf{real}}$
$\frac{\Gamma, x : \tau_1 \vdash^\Lambda e : \tau_2}{\Gamma \vdash^V \lambda x. e : \tau_1 \rightarrow \tau_2}$	$\frac{\Gamma \vdash^V v_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash^V v_2 : \tau_2}{\Gamma \vdash^\Lambda v_1 v_2 : \tau_2}$	$\frac{\Gamma \vdash^\Lambda e_1 : \tau_1 \quad \Gamma, x : \tau_1 \vdash^\Lambda e_2 : \tau_2}{\Gamma \vdash^\Lambda \mathbf{let } x = e_1 \mathbf{ in } e_2 : \tau_2}$	$\frac{\Gamma \vdash^V v : \tau_1}{\Gamma \vdash^V \langle i, v \rangle : \sum_{i \in I} \tau_i}$	
$\frac{\Gamma \vdash^V v : \sum_{i \in I} \tau_i \quad \Gamma, x : \tau_i \vdash^\Lambda e_i : \tau \quad (\forall i \in I)}{\Gamma \vdash^\Lambda \mathbf{case } v \mathbf{ of } \{ \langle i, x \rangle \rightarrow e_i \} : \tau}$	$\frac{\Gamma \vdash^V v : \tau[\mu\alpha.\tau/\alpha]}{\Gamma \vdash^V \mathbf{fold } v : \mu\alpha.\tau}$	$\frac{\Gamma \vdash^V v : \mu\alpha.\tau_1 \quad \Gamma, x : \tau_1[\mu\alpha.\tau_1/\alpha] \vdash^\Lambda e : \tau_2}{\Gamma \vdash^\Lambda \mathbf{case } v \mathbf{ of } \{ \mathbf{fold } x \rightarrow e \} : \tau_2}$		

Figure 1: Static semantics of Λ_P .

abstractly described using the *sub-Giry monad* \mathcal{G} [Giry \(1982\)](#), so that the evaluation function $\llbracket - \rrbracket_\tau$ will map a program $e \in \Lambda_\tau$ to a sub-probability measure $\llbracket e \rrbracket_\tau \in \mathcal{G}(\mathcal{V}_\tau)$. Additionally, since $\mathcal{G}X$ carries a measurable space structure whenever X does, we would like $\llbracket - \rrbracket_\tau$ to be a measurable function. For that to make sense, however, we first need to make Λ a measurable space itself. We follow [Ehrhard et al. \(2018\)](#); [Sangiorgi, Kobayashi, and Sumii \(2011\)](#); [Staton et al. \(2016\)](#) and rely on the measurable space structure of \mathbb{R}^n and co-products in **Meas**, the category of measurable spaces². Before entering into the details we recall some useful mathematical preliminaries.

3.1 Mathematical Preliminaries: Stochastic Kernels and the Giry Monad

In this section we recall some mathematical preliminaries which are needed in order to give Λ_P dynamic semantics. Unfortunately, there is no hope to be comprehensive, and thus we assume the reader to be familiar with basic measure theory [Pollard, Gill, and Ripley \(2002\)](#), domain theory [Abramsky and Jung \(1994\)](#), and category theory [MacLane \(1971\)](#).

We denote by **Meas** the category of measurable spaces and measurable functions, by **Set** the category of sets and functions, and by $\omega\mathbf{Cppo}$ the category of ω -cpo and continuous functions. Given a measurable space (X, Σ_X) with a measure μ on it and a function $f : X \rightarrow [0, \infty]$, we denote by $\int_X f d\mu$ or $\int_X f(x)\mu(dx)$ the number in $[0, \infty]$ obtained by Lebesgue integrating f over X with respect to μ (when clear from the context, we will omit the subscript X).

In order to deal with nontermination, we work with sub-probability distributions rather than with full distributions. For instance, the measure $\llbracket \Omega \rrbracket$ associated to the purely divergent expression Ω will naturally assign 0 to any measurable sets of values. Mathematically, this means working with the sub-Giry monad rather than with the (full) Giry monad.

Definition 1. The sub-Giry functor $\mathcal{G} : \mathbf{Meas} \rightarrow \mathbf{Meas}$ maps a measurable space (X, Σ_X) to $(\mathcal{G}X, \Sigma_{\mathcal{G}X})$, where $\mathcal{G}X$ is the set of all sub-probability measures on X , and $\Sigma_{\mathcal{G}X}$ is the least σ -algebra making the Σ_X -indexed family of maps $\text{eval}_A : \mathcal{G}(X) \rightarrow [0, 1]$, defined by $\text{eval}_A(v) = v(A)$, measurable. Measurable functions $f : X \rightarrow Y$ are mapped to $\mathcal{G}f : \mathcal{G}X \rightarrow \mathcal{G}Y$, where $(\mathcal{G}f)(v) \in \mathcal{G}Y = v(f^{-1}(-))$.

The functor \mathcal{G} carries a monad structure, with unit $\eta_X : X \rightarrow \mathcal{G}X$ given by Dirac measures (recall that for any $x \in X$, the Dirac measure δ_x on X is defined as $\delta_x(A) = 1$ if $x \in A$, and $\delta_x(A) = 0$, otherwise), and multiplication $\mu_X : \mathcal{G}\mathcal{G}X \rightarrow \mathcal{G}X$ defined as $\mu_X(M)(A) = \int_{\mathcal{G}X} \text{eval}_A(x)M(dx)$, for $A \in \Sigma_X$. Equivalently, we can use \mathcal{G} to build a Kleisli triple [MacLane \(1971\)](#) defining³ $f^\dagger(v)(B) = \int f(x)(B)v(dx)$, for $f : X \rightarrow \mathcal{G}Y$. Actually, Fubini's Theorem implies that \mathcal{G} has the structure of a commutative monad, with double strength $\text{dst}_{X,Y} : \mathcal{G}X \times \mathcal{G}Y \rightarrow \mathcal{G}(X \times Y)$ mapping measures $\mu \in \mathcal{G}X$, $\nu \in \mathcal{G}Y$ to their product measure $\mu \otimes \nu$.

A measurable function $f : X \rightarrow \mathcal{G}Y$ is also called a (stochastic) *kernel* [Panangaden \(1999\)](#) from X to Y , as we can uncurry f obtaining a map $f' : X \times \Sigma_Y \rightarrow [0, 1]$. We write **Kern** for the Kleisli category

²Another approach is given in [Borgström, Lago, Gordon, and Szymczak \(2016\)](#), where the set of terms inherits a measurable structure from a syntax-oriented metric.

³Notice that the function mapping an element x in X to $f(x)(A) \in [0, 1]$ is integrable.

of \mathcal{G} . The category **Kern** enjoys several nice properties [Panangaden \(1999\)](#). In particular, the collection of measurable functions from X to $\mathcal{G}Y$ carries an ω -cpo structure under the pointwise order \sqsubseteq (i.e. $f \sqsubseteq g$ if and only if $\forall x \in X. \forall A \in \Sigma_Y. f(x)(A) \leq g(x)(A)$). The bottom element \perp is defined as the always zero distribution, whereas the least upper bound of an ω -chain $(f_n)_{n \geq 0}$ is defined as $\sup_n f_n$. Additionally, the monotone convergence theorem and the very definition of Lebesgue integration give the following identities: $\int \sup_n f_n dv = \sup_n \int f_n dv$, $\int f d(\sup_n \nu_n) = \sup_n \int f d\nu_n$. Abstractly, this means that **Kern** is ω **Cppo**-enriched [Kelly \(2005\)](#). Finally, we recall that **Meas** is cartesian (although not cartesian closed) and that **Kern** has countable coproducts. For a countable family (X_i, Σ_{X_i}) of measurable spaces we denote by $\coprod_i (X_i, \Sigma_{X_i})$ its coproduct.

3.2 Evaluation Semantics

We define a (type-indexed) map $\llbracket - \rrbracket_\tau : \Lambda_\tau \rightarrow \mathcal{G}(\mathcal{V}_\tau)$ associating to each program $e \in \Lambda_\tau$ a (sub)probability measure over values in \mathcal{V}_τ . We rely on the ω **Cppo**-enrichment of **Kern** to deal with non-termination.

Definition 2. Define the type-indexed family of (\mathbb{N} -indexed) maps $\llbracket - \rrbracket_\tau^{(n)} : \Lambda_\tau \rightarrow \mathcal{G}(\mathcal{V}_\tau)$ as follows (for readability, we omit type annotations in $\llbracket - \rrbracket^{(n)}$):

$$\begin{aligned} \llbracket e \rrbracket^{(0)}(A) &= 0 \\ \llbracket v \rrbracket^{(n+1)}(A) &= \delta_v(A) \\ \llbracket F(\mathbf{r}_1, \dots, \mathbf{r}_n) \rrbracket^{(n+1)}(A) &= \llbracket \mathbf{F}(\mathbf{r}_1, \dots, \mathbf{r}_n) \rrbracket^{(n)}(A) \\ \llbracket \text{sample} \rrbracket^{(n+1)}(A) &= \lambda\{r \mid \mathbf{r} \in A\} \\ \llbracket (\lambda x. e)v \rrbracket^{(n+1)}(A) &= \llbracket e[v/x] \rrbracket^{(n)}(A) \\ \llbracket \text{let } x = e_1 \text{ in } e_2 \rrbracket^{(n+1)}(A) &= \int \llbracket e_2[v/x] \rrbracket^{(n)} \llbracket e_1 \rrbracket^{(n)}(dv) \\ \llbracket \text{case } \langle i, v \rangle \text{ of } \{ \langle i, x \rangle \rightarrow e_i \} \rrbracket^{(n+1)}(A) &= \llbracket e_i[v/x] \rrbracket^{(n)}(A) \\ \llbracket \text{case (fold } v \text{) of {fold } x \rightarrow e} \rrbracket^{(n+1)}(A) &= \llbracket e[v/x] \rrbracket^{(n)}(A). \end{aligned}$$

Lemma 1. For any type τ and natural number n , the map $\llbracket - \rrbracket_\tau^{(n)}$ is a kernel.

Sketch. First of all, we notice that for the statement of Lemma 1 to make sense, we first need to make Λ_τ (as well as \mathcal{V}_τ) a measurable space. This can be done as in [Ehrhard et al. \(2018\)](#); [Staton et al. \(2016\)](#), taking advantages of coproducts in **Meas**. By observing that for any expression $\Gamma, x : \tau_1 \vdash^\Lambda e : \tau_2$, the map $e[-/x] : \mathcal{V}_{\Gamma+\tau_1} \rightarrow \Lambda_{\Gamma+\tau_2}$ is measurable, one can easily show $\llbracket - \rrbracket^{(n)}$ to be measurable. \square

Additionally, we see that $\llbracket e \rrbracket_\tau^{(n)}$ is actually an ω -chain in $\mathcal{G}(\mathcal{V}_\tau)$, so that we can define $\llbracket e \rrbracket_\tau$ as $\sup_n \llbracket e \rrbracket_\tau^{(n)}$. In particular, due to ω **Cppo**-enrichment of **Kern**, we can replace all the occurrences of $\llbracket - \rrbracket^{(n)}$ in Definition 2 with $\llbracket - \rrbracket$. As usual, where possible we omit type subscripts in $\llbracket - \rrbracket$.

Having defined the operational behaviour of programs, we can finally move to the study of notions of program equivalence and refinement.

4 Probabilistic Equivalence and Refinement

In the next sections we define notions of program equivalence and refinement for Λ_p . Specifically, we introduce (probabilistic) contextual equivalence and approximation, and (probabilistic) applicative (bi)similarity. Our main result states that applicative similarity (resp. bisimilarity) is a sound proof technique for contextual approximation (resp. equivalence).

Studying notions of program equivalence and refinement, it is useful to fix a proper notation and vocabulary for program relations. We do so following [Gordon \(1994\)](#); [S. Lassen \(1998\)](#).

4.1 Relational Calculus

We use letter R, S, \dots to range over relations and write $\text{Rel}(X, Y)$ for the set of relations over X and Y . We call elements of $\text{Rel}(X, X)$ *endorelations*. Given a relation $R \subseteq X \times Y$ and set $A \subseteq X$, we denote by $R[A]$ the R -image of A , i.e. $\{y \in Y \mid \exists x \in A. (x, y) \in R\}$. For relations $R \subseteq X \times Y, S \subseteq Y \times Z$, we write $S \circ R$ for their composition, and denote by $=_X$ the equality relation on X . As usual, an endorelation $R \in \text{Rel}(X, X)$ is reflexive if $=_X \subseteq R$, transitive if $R \circ R \subseteq R$, and symmetric if $R^\top \subseteq R$, where R^\top denotes the converse or transpose of R . We also recall that $\text{Rel}(X, X)$ carries a complete lattice structure with order given by set-theoretic inclusion.

Relational reasoning about programs oftentimes require to reason about open terms. It is thus useful to work with families of relations relating expressions typable within the same sequent. We refer to such (families of) relations as *term relations*.

Definition 3. A term relation R associates to each computation sequent $\Gamma \vdash^\Lambda \tau$ (resp. value sequent $\Gamma \vdash^\mathcal{V} \tau$) a relation $\Gamma \vdash^\Lambda - R - : \tau$ (resp. $\Gamma \vdash^\mathcal{V} - R - : \tau$) between its inhabitants.

Formally, a term relation is thus an element in the set $\text{Rel} = \prod_{\Gamma \vdash \tau} \text{Rel}(\Lambda_{\Gamma \vdash \tau}, \Lambda_{\Gamma \vdash \tau}) \times \text{Rel}(\mathcal{V}_{\Gamma \vdash \tau}, \mathcal{V}_{\Gamma \vdash \tau})$. The latter inherits a complete lattice structure from $\text{Rel}(\Lambda_{\Gamma \vdash \tau}, \Lambda_{\Gamma \vdash \tau})$ and $\text{Rel}(\mathcal{V}_{\Gamma \vdash \tau}, \mathcal{V}_{\Gamma \vdash \tau})$ pointwise, this way allowing one to define term relations both inductively and coinductively. Concerning notation, we write $R \subseteq S$ if:

$$\forall \Gamma, e_1, e_2. \Gamma \vdash^\Lambda e_1 R e_2 \implies \Gamma \vdash^\Lambda e_1 S e_2, \quad \forall \Gamma, v_1, v_2. \Gamma \vdash^\mathcal{V} v_1 R v_2 \implies \Gamma \vdash^\mathcal{V} v_1 S v_2.$$

Example 2. The identity term relation is defined by the (0-ary) rules: $\Gamma \vdash^\Lambda e =_\Lambda e : \tau$, $\Gamma \vdash^\mathcal{V} v =_\mathcal{V} v : \tau$. Given term relations R, S , we define $S \circ R$ as follows:

$$\frac{\Gamma \vdash^\Lambda e_1 R e_2 : \tau \quad \Gamma \vdash^\Lambda e_2 S e_3 : \tau}{\Gamma \vdash^\Lambda e_1 (S \circ R) e_3 : \tau} \quad \frac{\Gamma \vdash^\mathcal{V} v_1 R v_2 : \tau \quad \Gamma \vdash^\mathcal{V} v_2 S v_3 : \tau}{\Gamma \vdash^\mathcal{V} v_1 (S \circ R) v_3 : \tau}$$

The notions of a *preorder* and *equivalence* term relation are defined as usual. In fact, we say that a term relation R is reflexive if $= \subseteq R$, transitive if $R \circ R \subseteq R$, and symmetric if $R^\top \subseteq R$ (where R^\top is defined in the obvious way).

Oftentimes we will be interested in defining relations between programs only. We denote by Rel_0 the subspace $\prod_{\tau} \text{Rel}(\Lambda_\tau, \Lambda_\tau) \times \text{Rel}(\mathcal{V}_\tau, \mathcal{V}_\tau)$ of Rel , and refer to its elements as *closed term relations*. Every closed term relation $R \in \text{Rel}_0$ induces a term relation, called the *open extension* of R , as follows: given $\Gamma \vdash^\Lambda e : \tau$, let γ range over closed substitutions for Γ (i.e. for any variable $(x_i : \tau_i) \in \Gamma, \gamma(x_i) \in \mathcal{V}_{\tau_i}$). Define:

$$\Gamma \vdash^\Lambda e_1 R e_2 : \tau \iff \forall \gamma. \cdot \vdash^\Lambda e_1 \gamma R e_2 \gamma : \tau.$$

A similar definition can be given for values. Given a closed term relation R , we sometimes write $(R_\Lambda, R_\mathcal{V})$ to denote the (type-indexed families of) relations defining its action on expressions and values.

We tacitly require term relations to be closed under weakening, meaning that the following hold.

$$\frac{\Gamma \vdash^\Lambda e_1 R e_2 : \tau}{\Gamma, x : \tau' \vdash^\Lambda e_1 R e_2 : \tau} \quad \frac{\Gamma \vdash^\mathcal{V} v_1 R v_2 : \tau}{\Gamma, x : \tau' \vdash^\mathcal{V} v_1 R v_2 : \tau}$$

Notice that the open extension of a closed relation is indeed closed under weakening.

In order to formalise compositional reasoning about program behaviour, we introduce the notion of a compatible term relation, i.e. of a term relation closed under term constructors of Λ_p .

Definition 4. Given a term relation R , its compatible refinement \widehat{R} is defined by the rules in Figure 2. A term relation R is compatible if $\widehat{R} \subseteq R$, and a closed term relation is compatible if its open extension is.

$$\begin{array}{c}
\frac{\Gamma, x : \tau \vdash^V x \widehat{R} x : \tau}{\Gamma \vdash^A v_1 \widehat{R} v_2 : \tau} \quad \frac{\Gamma \vdash^V v_1 \widehat{R} v_2 : \tau}{\Gamma \vdash^V r \widehat{R} r : \mathbf{real}} \quad \frac{\Gamma \vdash^V v_1 \widehat{R} v_2 : \tau}{\Gamma \vdash^A \mathbf{sample} \widehat{R} \mathbf{sample} : \mathbf{real}} \\
\frac{\Gamma \vdash^V v_1 R v'_1 : \mathbf{real} \quad \cdots \quad \Gamma \vdash^V v_n R v'_n : \mathbf{real}}{\Gamma \vdash^A F(v_1, \dots, v_n) \widehat{R} F(v'_1, \dots, v'_n) : \mathbf{real}} \quad \frac{\Gamma, x : \tau_1 \vdash^A e_1 R e_2 : \tau_2}{\Gamma \vdash^V \lambda x. e_1 \widehat{R} \lambda x. e_2 : \tau_1 \rightarrow \tau_2} \quad \frac{\Gamma \vdash^V v_1 R v'_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash^V v_2 R v'_2 : \tau_1}{\Gamma \vdash^A v_1 v_2 \widehat{R} v'_1 v'_2 : \tau_2} \\
\frac{\Gamma \vdash^A e_1 R e'_1 : \tau_1 \quad \Gamma, x : \tau_1 \vdash^A e_2 R e'_2 : \tau_2}{\Gamma \vdash^A \mathbf{let} x = e_1 \mathbf{in} e_2 \widehat{R} \mathbf{let} x = e'_1 \mathbf{in} e'_2 : \tau_2} \quad \frac{\Gamma \vdash^V v_1 R v_2 : \tau_i}{\Gamma \vdash^V \langle i, v_1 \rangle \widehat{R} \langle i, v_2 \rangle : \sum_{i \in I} \tau_i} \quad \frac{\Gamma \vdash^V v_1 R v_2 : \tau[\mu\alpha.\tau/\alpha]}{\Gamma \vdash^V \mathbf{fold} v_1 \widehat{R} \mathbf{fold} v_2 : \mu\alpha.\tau} \\
\frac{\Gamma \vdash^V v R v' : \sum_{i \in I} \tau_i \quad \Gamma, x : \tau_i \vdash^A e_i R e'_i : \tau \quad (\forall i \in I)}{\Gamma \vdash^A \mathbf{case} v \mathbf{of} \{ \langle i, x \rangle \rightarrow e_i \} \widehat{R} \mathbf{case} v' \mathbf{of} \{ \langle i, x \rangle \rightarrow e'_i \} : \tau} \quad \frac{\Gamma \vdash^V v R v' : \mu\alpha.\tau_1 \quad \Gamma, x : \tau_1[\mu\alpha.\tau_1/\alpha] \vdash^A e R e' : \tau_2}{\Gamma \vdash^A \mathbf{case} v \mathbf{of} \{ \mathbf{fold} x \rightarrow e \} \widehat{R} \mathbf{case} v' \mathbf{of} \{ \mathbf{fold} x \rightarrow e' \} : \tau_2}
\end{array}$$

Figure 2: Compatible refinement.

Notice that \widehat{R} is indeed a term relation (notably, \widehat{R} is closed under weakening). Definition 4 induces a monotone endofunction $\widehat{\cdot}$ on Rel which distributes over composition ($\widehat{S \circ R} = \widehat{S} \circ \widehat{R}$) and transpose ($\widehat{R^\top} = (\widehat{R})^\top$). In particular, a term relation is compatible if and only if it is a pre-fixed point of $\widehat{\cdot}$. It is not hard to prove that the identity term relation of Example 2 is a pre-fixed point of $\widehat{\cdot}$, and actually the least such. As a consequence, any compatible relation is reflexive.

Lemma 2. *The collection of compatible λ -term relations ordered by \subseteq has a complete lattice.*

Lemma 2 allows us to define compatible λ -term relations both inductively and coinductively. Another central notion dealing with term relation is the one of *substitutivity*.

Definition 5. 1. A term relation R is *value-substitutive* if the following hold:

$$\frac{\Gamma, x : \tau_1 \vdash^A e_1 R e_2 : \tau_2 \quad \cdot \vdash^V v : \tau_1}{\Gamma \vdash^A e_1[v/x] R e_2[v/x] : \tau_2} \quad \frac{\Gamma, x : \tau_1 \vdash^V v_1 R v_2 \quad \cdot \vdash^V v : \tau_1}{\Gamma \vdash^V v_1[v/x] R v_2[v/x] : \tau_2}$$

2. A term relation R is *substitutive* if the following hold:

$$\frac{\Gamma, x : \tau_1 \vdash^A e_1 R e_2 : \tau_2 \quad \cdot \vdash^V v_1 R v_2 : \tau_1}{\Gamma \vdash^A e_1[v_1/x] R e_2[v_2/x] : \tau_2} \quad \frac{\Gamma, x : \tau_1 \vdash^V u_1 R u_2 : \tau_2 \quad \cdot \vdash^V v_1 R v_2 : \tau_1}{\Gamma \vdash^V u_1[v_1/x] R u_2[v_2/x] : \tau_2}$$

A closed relation is (value) substitutive if its open extension is. Moreover, we notice that the open extension of a closed term relation is trivially value-substitutive.

5 Contextual Approximation and Equivalence

Contextual equivalence [Morris \(1969\)](#) equates programs that behave the same in any possible environment. The notion of an environment is formalised by means of *contexts*, which are, roughly speaking, terms with a hole to be filled in with the program we aim to test. In order to avoid syntactic bureaucracy, we follow [Gordon \(1994\)](#); [S. Lassen \(1998\)](#) and give a property-based definition of contextual equivalence and approximation based on the notion of (*pre*)*adequacy*. For readability, from now on we write \mathbf{R} for $\mathcal{V}_{\mathbf{real}} \in \Sigma_{\mathbf{real}}$.

Definition 6. Given a term relation R , we say that R is *preadequate* (resp. *adequate*) if:

$$\cdot \vdash^A e_1 R e_2 : \mathbf{real} \implies \llbracket e_1 \rrbracket(\mathbf{R}) \leq \llbracket e_2 \rrbracket(\mathbf{R}); \quad (\text{resp. } \cdot \vdash^A e_1 R e_2 : \mathbf{real} \implies \llbracket e_1 \rrbracket(\mathbf{R}) = \llbracket e_2 \rrbracket(\mathbf{R})).$$

Definition 7. Contextual approximation \leq^{ctx} (resp. contextual equivalence \simeq^{ctx}) is the largest compatible and preadequate (resp. adequate) term relation.

Definition 7 is not well-posed, as it is not clear whether such largest term relations exist. In fact, (pre)adequacy is not a monotone property, meaning that we cannot define \leq^{ctx} (resp. \simeq^{ctx}) as the greatest fixed point of the (non-monotone) endofunction on Rel induced by Definition 7. Nonetheless, we see that the union of compatible (pre)adequate term relations is itself a compatible (pre)adequate term relation.

Additionally, Definition 7 gives a proof principle for contextual approximation (resp. equivalence) resembling a coinduction proof principle. In order to prove that a term relation R is included in \leq^{ctx} (resp. \simeq^{ctx}), it is sufficient to show that R is preadequate (resp. adequate) and compatible. Using such a proof principle it is a straightforward exercise to prove that \leq^{ctx} is a precongruence and that \simeq^{ctx} is a congruence term relation.

Example 3. Contextual equivalence gives the structural equalities in Figure 3, as well the quasi-denotational law: $\forall e_1, e_2 \in \Lambda_\tau. \llbracket e_1 \rrbracket = \llbracket e_2 \rrbracket \implies \cdot \vdash^\Lambda e_1 \simeq^{\text{ctx}} e_2 : \tau$. The quasi-denotational law can be used to export powerful results from measure theory, for instance (we slightly abuse notation by desequencing some computations):

$$\text{let } x_1 = (\text{normal } e_{\mu_1} e_{\sigma_1}^2) \text{ in } (\text{let } x_2 = (\text{normal } e_{\mu_2} e_{\sigma_2}^2) \text{ in } (x_1 + x_2)) \simeq^{\text{ctx}} \text{normal } (e_{\mu_1} + e_{\mu_2}) (e_{\sigma_1}^2 + e_{\sigma_2}^2)$$

Combining such laws with the structural equivalences in Figure 3 we can prove non-trivial program equivalences. For instance, recalling the program transformation studied in Section 1.1, we see that contextual equivalence gives correctness of the following program transformation, which avoid a list traversing.

$$\left(\begin{array}{l} \text{let } xs = \text{map } (\lambda. \text{normal } e_{\mu_1} e_{\sigma_1}^2) zs \\ \quad ys = \text{map } (\lambda. \text{normal } e_{\mu_2} e_{\sigma_2}^2) zs \\ \text{in } \text{append } x (\text{zipwith } (+) xs ys) \end{array} \right) \simeq^{\text{ctx}} \left(\begin{array}{l} \text{let } xs = \text{map } (\lambda. \text{normal } (e_{\mu_1} + e_{\mu_2}) (e_{\sigma_1}^2 + e_{\sigma_2}^2)) zs \\ \text{in } \text{append } x xs \end{array} \right)$$

Similarly, we can prove the correctness of the program transformation studied in Section 1.1. Proving such equivalences and laws from first principles, however, is highly non-trivial. We will obtain these results relying on our notion of applicative bisimulation.

$$\begin{array}{l} (\lambda x.e)v \simeq^{\text{ctx}} e[v/x] \\ \text{let } x = v \text{ in } e \simeq^{\text{ctx}} e[v/x] \\ \text{let } x = e_1 \text{ in } e_2 \simeq^{\text{ctx}} e_2 \qquad (x \notin FV(e_2)) \\ \text{let } x = e \text{ in } (\text{let } y = e' \text{ in } e'') \simeq^{\text{ctx}} \text{let } y = e' \text{ in } (\text{let } x = e \text{ in } e'') \qquad (x \notin FV(e'), y \notin FV(e)) \end{array}$$

Figure 3: Structural identities.

6 Applicative (Bi)simulation

Abramsky's applicative bisimilarity [Abramsky \(1990\)](#) tests programs as argument-passing processes, i.e. by allowing the environment to interact with a program by passing it arbitrary inputs only. In a probabilistic setting, however, the result of a program is not a value, but a subdistribution of values, meaning that we have to refine the conceptual apparatus behind applicative bisimilarity to take into account probabilistic behaviours.

Following [Hughes and Jacobs \(2004\)](#); [Levy \(2011\)](#); [Thijs \(1996\)](#), we look at notions of (bi)simulation in terms of relation lifting operations. Suppose we are given a term relation R such that $\vdash^\Lambda e_1 R e_2$. How should we relate $\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket \in \mathcal{G}\mathcal{V}_\tau$? Answering such a question means nothing more than finding a proper way to lift R_ν to $\mathcal{G}\mathcal{V}_\tau \times \mathcal{G}\mathcal{V}_\tau$.

6.1 Probabilistic Relation Lifting

From an abstract perspective, what we need in order to define applicative (bi)similarity is to find a lifting from $\text{Rel}(X, Y)$ to $\text{Rel}(\mathcal{G}X, \mathcal{G}Y)$. The literature on probabilistic equivalence abounds of such notions of lifting (the interested reader can consult [Panangaden \(2011\)](#) for an overview, and [Katsumata, Sato, and Uustalu \(2018\)](#) for a more technical treatment). For our purposes, the kind of lifting we are interested in is one that allows applicative (bi)similarity:

1. To be definable by coinduction;
2. To be a preorder (resp. equivalence) term relation;
3. To be a compatible and (pre)adequate term relation, and thus a sound proof technique for contextual approximation/equivalence.

The collection of notions of lifting satisfying these requirements has been identified in [Dal Lago et al. \(2017\)](#) with the one of *relators* or *lax extension* [Barr \(1970\)](#); [Thijs \(1996\)](#).

Unfortunately, finding non-trivial relators for the Giry monad has been proved to be extremely difficult, as the latter lacks some fundamental structural properties (notably, it does not preserve weak pullbacks [Panangaden \(2011\)](#)). Nonetheless, as first argued in [Van Breugel, Mislove, Ouaknine, and Worrell \(2005\)](#), restricting the attention from arbitrary binary relations to *reflexive endorelations* allows one to define well-behaved notions of relation lifting for \mathcal{G} .

At this point it is instructive to recall the standard notion of relation lifting proposed for \mathcal{G} (see, e.g., [Katsumata et al. \(2018\)](#)), and the kind of problems one faces when working with it. We also remark that the notions of lifting we are going to analyse are actually meant to model notions of program refinement, rather than equivalence. In fact, applicative bisimilarity can be defined in terms of applicative similarity, and compatibility of the former follows from compatibility of the latter.

Definition 8. Given measurable spaces (X, Σ_X) , (Y, Σ_Y) , and a relation $R \subseteq X \times Y$ between them, we define $\Gamma R \subseteq \mathcal{G}X \times \mathcal{G}Y$ as follows: $\Gamma R = \{(v_1, v_2) \mid \forall A \in \Sigma_X. \forall B \in \Sigma_Y. R[A] \subseteq B \implies v_1(A) \leq v_2(B)\}$.

The map $\Gamma : \text{Rel}(X, Y) \rightarrow \text{Rel}(\mathcal{G}X, \mathcal{G}Y)$ satisfies many interesting properties (see e.g. [Katsumata et al. \(2018\)](#)), but fails to satisfy the inclusion $\Gamma S \circ \Gamma R \subseteq \Gamma(S \circ R)$, which we refer to as *quasi transitivity*. Quasi transitivity ensures applicative similarity to be transitive. Most importantly, it is a central ingredient in Howe's method [Howe \(1996\)](#); [Pitts \(2011\)](#), the standard technique used to prove applicative (bi)similarity to be compatible.

A counterexample to quasi transitivity of Γ is given in [Katsumata et al. \(2018\)](#) (Example 4.15), taking advantage of the sets X and Y to be distinct (notably, taking X and Y to be the discrete and indiscrete space over the two-elements set, respectively). However, as we are interested in term relations, we can look at endorelations only. Moreover, as observed in [Van Breugel et al. \(2005\)](#), working with reflexive endorelations (which still form a complete lattice), it is possible to refine Γ in such a way to guarantee quasi transitivity.

Definition 9. Given a measurable space (X, Σ_X) , and a reflexive relation $R \subseteq X \times X$, define $\Gamma R \subseteq \mathcal{G}X \times \mathcal{G}X$ as follows⁴: $\Gamma R = \{(v_1, v_2) \mid \forall A \in \Sigma_X. R[A] \subseteq A \implies v_1(A) \leq v_2(A)\}$.

Lemma 3. Let $\text{Rel}_=(X, X)$ denote the complete lattice of endorelations on X . Then $\Gamma : \text{Rel}_=(X, X) \rightarrow \text{Rel}_=(\mathcal{G}X, \mathcal{G}X)$ satisfies the following properties, for all $R, S \in \text{Rel}_=(X, X)$.

$$\begin{aligned}
 &=_{\mathcal{G}X} \subseteq \Gamma(=_X) && \text{(quasi reflexivity)} \\
 \Gamma S \circ \Gamma R &\subseteq \Gamma(S \circ R). && \text{(quasi transitivity)} \\
 R \subseteq S &\implies \Gamma R \subseteq \Gamma S. && \text{(monotonicity)}
 \end{aligned}$$

⁴Notice that R being reflexive, $R[A] \subseteq A$ implies $R[A] = A$.

The proof of Lemma 3 is straightforward. Remarkably, we can give another characterisation of Γ .

Proposition 1. *Given $\mu, \nu \in \mathcal{G}X$ and $R \in \text{Rel}_=(X, X)$, we have:*

$$\mu \Gamma R \nu \iff \forall f : (X, R) \xrightarrow{1} ([0, 1], \leq). \int_X f d\mu \leq \int_X f d\nu,$$

where $f : (X, R) \xrightarrow{1} ([0, 1], \leq)$ means that f is a measurable function (notice that $[0, 1] \cong \mathcal{G}\{*\}$) such that $x R y$ implies $f(x) \leq f(y)$.

The proof of Proposition 1 goes essentially as in [Katsumata et al. \(2018\)](#) (Proposition 4.4). See also [Sato \(2016\)](#). Using Proposition 1 we can show that Γ is well-behaved with respect to the monad structure of \mathcal{G} , and thus that it nicely interacts with sequencing (observe that $\llbracket \text{let } x = e_1 \text{ in } e_2 \rrbracket(A) = \int \llbracket e_2[v/x] \rrbracket(A) \llbracket e_1 \rrbracket(dv) = \llbracket e_2[-/x] \rrbracket^\dagger \llbracket e_1 \rrbracket(A)$).

Lemma 4. *Given $R \in \text{Rel}_=(X, X)$, $S \in \text{Rel}_=(Y, Y)$, and measurable functions $f, g : X \rightarrow Y$, the following hold, for all $x, y \in X$ and $\mu, \nu \in \mathcal{G}X$*

$$x R y \implies \delta_x \Gamma R \delta_y \quad (\text{dirac})$$

$$(\forall x_1, x_2 \in X. x_1 R x_2 \implies f(x_1) \Gamma S g(x_2)) \implies (\mu \Gamma R \nu \implies f^\dagger(\mu) \Gamma S g^\dagger(\nu)) \quad (\text{bind})$$

Finally, Γ nicely interacts with the ωCppo -enrichment of \mathcal{G} , allowing the following induction principle.

Lemma 5. *For any relation $R \in \text{Rel}_=(X, X)$, any ω -chain $(\mu_n)_{n \geq 0}$ in $\mathcal{G}X$, and any $\nu \in \mathcal{G}X$, if $\mu_n \Gamma R \nu$ holds for all $n \geq 0$, then $\sup_n \mu_n \Gamma R \nu$.*

We can now rely on Γ to define a suitable notion of applicative (bi)simulation.

6.2 Applicative Similarly and Bisimilarity

We are now ready to formally define our notion of an applicative (bi)simulation.

Definition 10. *A reflexive closed term relation $R = (R_\wedge, R_\vee)$ is an applicative simulation if the following conditions hold:*

$$\begin{aligned} \cdot \vdash^\wedge e_1 R e_2 : \tau &\implies \llbracket e_1 \rrbracket \Gamma R_\vee \llbracket e_2 \rrbracket && (\text{app eval}) \\ \cdot \vdash^\vee \mathbf{r}_1 R \mathbf{r}_2 : \mathbf{real} &\implies \mathbf{r}_1 = \mathbf{r}_2 && (\text{app num}) \\ \cdot \vdash^\vee \lambda x. e_1 R \lambda x. e_2 : \tau_1 \rightarrow \tau_2 &\implies \forall v \in \mathcal{V}_\tau. \cdot \vdash^\wedge e_1[v/x] R e_2[v/x] : \tau_2 && (\text{app abs}) \\ \cdot \vdash^\vee \langle i, v_i \rangle R \langle j, v_j \rangle &\implies i = j \wedge \cdot \vdash^\vee v_i R v_j : \tau_i && (\text{app inj}) \\ \cdot \vdash^\vee \mathbf{fold} v_1 R \mathbf{fold} v_1 : \mu \alpha. \tau &\implies \cdot \vdash^\vee v_1 R v_2 : \tau[\mu \alpha. \tau / \alpha]. && (\text{app fold}) \end{aligned}$$

Applicative similarity \leq^\wedge is defined as the largest applicative simulation, whereas applicative bisimilarity \simeq^\wedge is the largest symmetric applicative simulation.

To see that \leq^\wedge indeed exists, we define the endofunction $[-]$ on the complete lattice Rel_0 as follows:

$$\begin{aligned} \cdot \vdash^\wedge e_1 [R] e_2 &\iff \llbracket e_1 \rrbracket \Gamma R_\vee \llbracket e_2 \rrbracket \\ \cdot \vdash^\vee \mathbf{r}_1 [R] \mathbf{r}_2 : \mathbf{real} &\iff \mathbf{r}_1 = \mathbf{r}_2 \\ \cdot \vdash^\vee \lambda x. e_1 [R] \lambda x. e_2 : \tau_1 \rightarrow \tau_2 &\iff \forall v \in \mathcal{V}_\tau. \cdot \vdash^\wedge e_1[v/x] R e_2[v/x] : \tau_2 \\ \cdot \vdash^\vee \langle i, v_i \rangle [R] \langle j, v_j \rangle &\iff i = j \wedge \cdot \vdash^\vee v_i R v_j : \tau_i \\ \cdot \vdash^\vee \mathbf{fold} v_1 [R] \mathbf{fold} v_1 : \mu \alpha. \tau &\iff \cdot \vdash^\vee v_1 R v_2 : \tau[\mu \alpha. \tau / \alpha]. \end{aligned}$$

We see that R is an applicative simulation if and only if $R \subseteq [R]$. Besides, Γ being monotone, $[-]$ is monotone as well, and thus it has a greatest fixed point which is \leq^\wedge .

Remark 1. Notice that we have not required $R_V \subseteq R_\Lambda$ for an applicative simulation $R = (R_\Lambda, R_V)$. However, it is straightforward to see that given such a simulation R , we can always extend R_Λ adding the pairs of values $(v_1, v_2) \in R_V$ still obtaining a simulation. This directly follows from condition (dirac). In particular, $\leq_V^\Lambda \subseteq \leq_\Lambda^\Lambda$.

Applicative similarity being defined coinductively, it comes with an associated coinduction proof principle, which can be formally given as follows:

$$\frac{\exists R. R \subseteq [R] \quad \cdot \vdash^\Lambda e_1 R e_2 : \tau}{\cdot \vdash^\Lambda e_1 \leq^\Lambda e_2 : \tau} \quad \frac{\exists R. R \subseteq [R] \quad \cdot \vdash^V v_1 R v_2 : \tau}{\cdot \vdash^V v_1 \leq^\Lambda v_2 : \tau}$$

An easy application of the coinduction proof principle allows us to prove \leq^Λ to be a preorder term relation.

Lemma 6. *Applicative similarity is reflexive and transitive.*

Sketch. By coinduction, relying on the properties (quasi reflexivity) and (quasi transitivity). \square

Additionally, we see that $\simeq^\Lambda = \leq^\Lambda \cap (\leq^\Lambda)^\top$, so that \simeq^Λ is an equivalence term relation. In order to ensure compositionality, we need to prove \leq^Λ and \simeq^Λ to be compatible term relations. We first prove that applicative similarity is compatible, i.e. $\widehat{\leq^\Lambda} \subseteq \leq^\Lambda$. To achieve such a goal, we use Howe's technique Howe (1996); Pitts (2011).

Definition 11. Given a closed term relation R , the Howe extension of R is the term relation R^H defined as the least solution to the equation:

$$\rho = R^o \circ \widehat{\rho},$$

where R^o denotes the open extension of R . Such a solution exists, since both $-^o$ and $\widehat{\cdot}$ are monotone endofunctions. More explicitly, we can inductively define R^H by:

$$\frac{\Gamma \vdash^\Lambda e_1 \widehat{R^H} e_3 : \tau \quad \Gamma \vdash^\Lambda e_3 R e_2 : \tau}{\Gamma \vdash^\Lambda e_1 R^H e_2 : \tau} \quad \frac{\Gamma \vdash^V v_1 \widehat{R^H} v_3 : \tau \quad \Gamma \vdash^V v_3 R v_2 : \tau}{\Gamma \vdash^V v_1 R^H v_2 : \tau}$$

Notice that R^H is not only the least solution to the equation $\rho = R^o \circ \widehat{\rho}$, but actually to the inclusion $R^o \circ \widehat{\rho} \subseteq \rho$ S. Lassen (1998); Levy (2006). The Howe extension of a reflexive and transitive (closed) term relation enjoys several interesting properties. In particular, it is a reflexive, compatible, and substitutive term relation.

Lemma 7. *Let R be reflexive and transitive closed term relation. Then the following hold:*

1. $R^o \subseteq R^H$.
2. $R^o \circ R^H \subseteq R^H$
3. R^H is compatible, and thus reflexive.
4. R^H is substitutive.

The proof of Lemma 7 is standard, and the reader is referred to S. Lassen (1998); Pitts (2011) for details. We are finally ready to prove our first main result, namely that applicative similarity is a precongruence term relation. To achieve such a goal, we prove that $(\leq^\Lambda)^H$ (restricted to closed terms) is an applicative simulation, and it is thus contained in \leq^Λ . Since by Lemma 7 we have the opposite inclusion, we will conclude $(\leq^\Lambda)^H = \leq^\Lambda$.

Lemma 8 (Key Lemma). *Given a reflexive and transitive closed term relation R , if R is an applicative simulation, then so is R^H (restricted to closed expressions and values).*

Sketch. By Lemma 7, R^H is reflexive, and thus a possible candidate simulation. It is easy to see that R_v^H (restricted to closed values) satisfies conditions (app num)-(app fold). It thus remain to prove (app eval), i.e. $\cdot \vdash^\wedge e_1 R e_2 : \tau \implies \llbracket e_1 \rrbracket \Gamma R_v^H \llbracket e_2 \rrbracket$. Since $\llbracket e_1 \rrbracket = \sup_n \llbracket e_1 \rrbracket^{(n)}$ we can appeal to Lemma 5, and prove the statement: $\cdot \vdash^\wedge e_1 R^H e_2 : \tau \implies \forall n \geq 0. \llbracket e_1 \rrbracket^{(n)} \Gamma R_v^H \llbracket e_2 \rrbracket$.

We proceed by induction on n . The base case is trivial. The inductive step is proved by case analysis on the derivation of $\cdot \vdash^\wedge e_1 R^H e_2 : \tau$, which in turn gives a case analysis on $\widehat{R^H}$. A central role is played by the inclusion $\Gamma R \circ \Gamma R^H \subseteq \Gamma R^H$ (which directly follows from Lemma 7 by monotonicity of Γ). Sequencing is handled as in Dal Lago et al. (2017), relying on property (bind), whereas the cases for measurable functions and numerals follows by condition (dirac). All other cases follow by induction hypothesis, relying on substitutivity of R^H . \square

Theorem 1. *Applicative similarity is a precongruence term relation and a sound proof technique for contextual approximation.*

Proof. It is sufficient to show that \leq^\wedge is compatible and preadequate. By Lemma 8 and Lemma 7 we know that $(\leq^\wedge)^H = \leq^\wedge$ when restricted to closed terms. This also holds on arbitrary terms, since $(\leq^\wedge)^H$ is substitutive and, dealing with closed values only, sequential and simultaneous substitution coincide. As a consequence, the open extension of applicative similarity coincides with $(\leq^\wedge)^H$, and thus it is compatible. To see that \leq^\wedge is preadequate, simply observe that $\leq_v^\wedge[\mathbf{R}] = \mathbf{R}$. \square

Finally, since $\simeq^\wedge = \leq^\wedge \cap (\leq^\wedge)^\top$, we see that \simeq^\wedge is compatible and adequate, and thus a sound proof technique for contextual equivalence.

Theorem 2. *Applicative bisimilarity \simeq^\wedge is a congruence term relation and a sound proof technique for contextual equivalence.*

Finally, we observe that we can use \simeq^\wedge to prove the equalities and laws in Example 3. In particular, the commutativity equation follows by commutativity of \mathcal{G} , whereas the soundness of the program transformation of Section 1.1 can be proved combining the quasi-denotational law with congruence properties of \simeq^\wedge , by passing the two programs an arbitrary value as input.

7 Digression: Towards Full Abstraction

Theorem 2 states that applicative bisimilarity is sound for contextual equivalence: is it also fully abstract (i.e. $\simeq^\wedge = \simeq^{\text{ctx}}$)? In Crubillé and Dal Lago (2014) this question is answered in the affirmative for a λ -calculus with sampling from discrete probability distributions. Such a result is proved going through testing equivalence De Nicola and Hennessy (1983); Larsen and Skou (1989), showing that tests can be implemented as λ -term contexts (meaning that testing equivalence includes contextual equivalence), on one hand, and that applicative bisimilarity coincides with testing equivalence, on the other hand.

The latter result has been proved in full generality in Van Breugel et al. (2005) relying on nontrivial results from domain theory, category theory, and measurable space theory. Remarkably, the equivalence between bisimilarity and testing equivalence holds not only for Markov chains, but also Markov processes, hence suggesting the possibility of proving full abstraction of applicative bisimilarity also in the continuous case. In fact, it is not hard to see that the operational semantics of Λ_P induces a (labelled) Markov process on terms, so that one naturally obtains notions of bisimilarity and testing equivalence for it. However, there is a fundamental difference between such a Markov process and the ones studied in Van Breugel et al. (2005). The former has *uncountably* many labels, whereas the latter requires the of labels to be *countable*.

Indeed, given a state $\lambda x.e$ of the Markov process, in order to model applicative bisimulation we need to consider transitions of the form $\lambda x.e \xrightarrow{v} e[v/x]$, meaning that the set of labels of the Markov process needs to contain (at least) all closed values, which are uncountably many.

A natural way to fix such a problem is to work with (uncountably many) forms of ‘countable’ bisimilarity relations \simeq_E^\wedge , where E is a *countable* set of labels of the Markov process, and to work with $\bigcap_E \simeq_E^\wedge$ in place of \simeq^\wedge . Showing $\bigcap_E \simeq_E^\wedge$ to be well-behaved, however, turns out to be highly nontrivial and to require the (abstract) identity $\Gamma(\bigcap_{i \in I} R_i) = \bigcap_{i \in I} \Gamma R_i$. Up to this point, the authors do not know whether such an identity holds, and leave further investigations on full abstraction as future work.

8 Conclusion

In this paper, we gave a notion of applicative bisimilarity for higher-order languages endowed with an operator performing sampling from *continuous* distributions. We proved that applicative bisimilarity is adequate and compatible, and thus sound for contextual equivalence. As far as the authors know, this is the first contribution on coinductive notions of equivalence for continuous probabilistic λ -calculi. We also gave evidence on the effectiveness of the introduced methodology by proving some nontrivial example equivalences. It is also interesting to notice that we can modify our framework to prove soundness of *open* (also known as *normal form*) bisimilarity S. B. Lassen (1999); Sangiorgi (1994), along the lines of Lago and Gavazzo (2019) (although we should remark that ‘standard’ open bisimilarity is not very well-suited for typed languages S. B. Lassen and Levy (2007)).

Related Work Higher-order programming languages featuring sampling from continuous distributions have received quite some attention in the last ten years, due to their use as idioms for bayesian programming. This has stimulated the study of operational Borgström et al. (2016); Park et al. (2005); Ramsey and Pfeffer (2002) and denotational Ehrhard et al. (2018); Staton (2017); Staton et al. (2016); Vákár, Kammar, and Staton (2019) kinds of semantics for these languages. Recently, contextual equivalence and logical relations have been introduced and proved to coincide both in presence and in absence of full recursion by Culpepper et al. Culpepper and Cobb (2017); Wand et al. (2018). Such logical relations require to test programs against any possible evaluation context, a feature reflected by their characterisation as CIU equivalence Mason and Talcott (1991). This makes reasoning with and about logical relations quite difficult. However, it is important to stress that such logical relations are defined for languages with conditioning (which makes program equivalence and semantics considerably harder), a feature not present in Λ_P .

Future Work A feature which is absent here, but which is desirable in fully-fledged probabilistic programming is conditioning, for example as expressed as a form of *scoring*. Coming up with an operational semantics for an extension of our λ -calculus with a scoring operator would be relatively easy, through the notion of *s*-finite kernel Staton (2017); Staton et al. (2016). The latter, however, seem not to carry a monad structure. A better choice might be to work with the monad $\mathcal{G}(\mathbb{R}_{\geq 0}^\infty \times -)$ associating to each expression a probability measures over (measurable sets of) pairs score-value. However, finding a way to compare such measures seems nontrivial. Of course, one can define a relation lifting for $\mathcal{G}(\mathbb{R}_{\geq 0}^\infty \times -)$ composing Γ with the ‘canonical’ relation lifting for $\mathbb{R}_{\geq 0}^\infty \times -$, the latter being nothing but an instance of the output monad. Such a lifting, however, is too fine-grained, as related values are required to have the same score. A better lifting of a relation R might be obtained requiring the expectations of the scores in R -closed measurable subsets of $\mathbb{R}_{\geq 0}^\infty \times -$ to be equal. Proving such a notion of lifting to be well-behaved is, however, nontrivial.

Finally, a further extension of the present work is the design of behavioural distances for probabilistic languages, possibly along the lines of de Amorim, Gaboardi, Hsu, Katsumata, and Cherigui (2017); Gavazzo (2018); Reed and Pierce (2010).

References

Abramsky, S. (1990). The lazy lambda calculus. In D. Turner (Ed.), *Research topics in functional programming* (pp. 65–117). Addison Wesley.

- Abramsky, S., & Jung, A. (1994). Domain theory. In *Handbook of logic in computer science* (pp. 1–168). Clarendon Press.
- Barendregt, H. (1984). *The lambda calculus: its syntax and semantics*. North-Holland.
- Barr, M. (1970). Relational algebras. *Lect. Notes Math.*, 137, 39–55.
- Biernacki, D., Piróg, M., Polesiuk, P., & Sieczkowski, F. (2018). Handle with care: relational interpretation of algebraic effects and handlers. *PACMPL*, 2(POPL), 8:1–8:30.
- Bizjak, A., & Birkedal, L. (2015). Step-indexed logical relations for probability. In *Proc. of FOSSACS 2015* (pp. 279–294).
- Blute, R., Desharnais, J., Edalat, A., & Panangaden, P. (1997). Bisimulation for labelled markov processes. In *Proc. of LICS* (pp. 149–158).
- Borgström, J., Lago, U. D., Gordon, A. D., & Szymczak, M. (2016). A lambda-calculus foundation for universal probabilistic programming. In *Proceedings of the 21st ACM SIGPLAN international conference on functional programming, ICFP 2016, nara, japan, september 18-22, 2016* (pp. 33–46).
- Crubillé, R., & Dal Lago, U. (2014). On probabilistic applicative bisimulation and call-by-value lambda-calculi. In *Proc. of ESOP 2014* (pp. 209–228).
- Culpepper, R., & Cobb, A. (2017). Contextual equivalence for probabilistic programs with continuous random variables and scoring. In *Proceedings of ESOP 2017* (pp. 368–392).
- Dal Lago, U., Gavazzo, F., & Levy, P. (2017). Effectful applicative bisimilarity: Monads, relators, and howe’s method. In *Proc. of LICS 2017* (pp. 1–12).
- Dal Lago, U., Sangiorgi, D., & Alberti, M. (2014). On coinductive equivalences for higher-order probabilistic functional programs. In *Proc. of POPL 2014* (pp. 297–308).
- De Nicola, R., & Hennessy, M. (1983). Testing equivalence for processes. In *Automata, languages and programming, 10th colloquium, barcelona, spain, july 18-22, 1983, proceedings* (pp. 548–560).
- de Amorim, A., Gaboardi, M., Hsu, J., Katsumata, S., & Cherigui, I. (2017). A semantic account of metric preservation. In *Proc. of POPL 2017* (pp. 545–556).
- De Liguoro, U., & Piperno, A. (1995). Non deterministic extensions of untyped lambda-calculus. *Inf. Comput.*, 122(2), 149–177.
- Ehrhard, T., Pagani, M., & Tasson, C. (2018). Measurable cones and stable, measurable functions: a model for probabilistic higher-order programming. *PACMPL*, 2(POPL), 59:1–59:28.
- Gavazzo, F. (2018). Quantitative behavioural reasoning for higher-order effectful programs: Applicative distances. In *Proceedings of the 33rd annual ACM/IEEE symposium on logic in computer science, LICS 2018, oxford, uk, july 09-12, 2018* (pp. 452–461).
- Giry, M. (1982). A categorical approach to probability theory. In B. Banaschewski (Ed.), *Categorical aspects of topology and analysis* (pp. 68–85). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Gordon, A. (1994, September). A tutorial on co-induction and functional programming. In *Workshops in computing* (p. 78-95). Springer London.
- Goubault-Larrecq, J., Lasota, S., & Nowak, D. (2008). Logical relations for monadic types. *Mathematical Structures in Computer Science*, 18(6), 1169–1217.
- Howe, D. (1996). Proving congruence of bisimulation in functional programming languages. *Inf. Comput.*, 124(2), 103–112.
- Hughes, J., & Jacobs, B. (2004). Simulations in coalgebra. *Theor. Comput. Sci.*, 327(1-2), 71–108.
- Johann, P., Simpson, A., & Voigtländer, J. (2010). A generic operational metatheory for algebraic effects. In *Proc. of LICS 2010* (pp. 209–218). IEEE Computer Society.
- Jones, C. (1990). *Probabilistic non-determinism* (Unpublished doctoral dissertation). University of Edinburgh, UK.
- Katsumata, S., Sato, T., & Uustalu, T. (2018). Codensity lifting of monads and its dual. *Logical Methods in Computer Science*, 14(4).
- Kelly, G. M. (2005). Basic concepts of enriched category theory. *Reprints in Theory and Applications of Categories*(10), 1–136.

- Lago, U. D., & Gavazzo, F. (2019). Effectful normal form bisimulation. In *Proc. of ESOP 2019*. (To appear)
- Larsen, K. G., & Skou, A. (1989). Bisimulation through probabilistic testing. In *Proceedings of POPL 1989* (pp. 344–352).
- Lassen, S. (1998). *Relational reasoning about functions and nondeterminism* (Unpublished doctoral dissertation). Dept. of Computer Science, University of Aarhus.
- Lassen, S. B. (1999). Bisimulation in untyped lambda calculus: Böhm trees and bisimulation up to context. *Electr. Notes Theor. Comput. Sci.*, 20, 346–374.
- Lassen, S. B., & Levy, P. B. (2007). Typed normal form bisimulation. In *Computer science logic, 21st international workshop, CSL 2007, 16th annual conference of the eacsl, lausanne, switzerland, september 11-15, 2007, proceedings* (pp. 283–297).
- Levy, P. (2006). Infinitary howe’s method. *Electr. Notes Theor. Comput. Sci.*, 164(1), 85–104.
- Levy, P. (2011). Similarity quotients as final coalgebras. In *Proc. of FOSSACS 2011* (Vol. 6604, pp. 27–41).
- Levy, P., Power, J., & Thielecke, H. (2003). Modelling environments in call-by-value programming languages. *Inf. Comput.*, 185(2), 182–210.
- MacLane, S. (1971). *Categories for the working mathematician*. Springer-Verlag.
- Mason, I. A., & Talcott, C. L. (1991). Equivalence in functional languages with effects. *J. Funct. Program.*, 1(3), 287–327.
- Morris, J. (1969). *Lambda calculus models of programming languages* (Unpublished doctoral dissertation). MIT.
- Panangaden, P. (1999). The category of markov kernels. *Electr. Notes Theor. Comput. Sci.*, 22, 171–187.
- Panangaden, P. (2011). Probabilistic bisimulation. In D. Sangiorgi & J. Rutten (Eds.), *Advanced topics in bisimulation and coinduction* (p. 290–326). Cambridge University Press.
- Park, S., Pfenning, F., & Thrun, S. (2005). A probabilistic language based upon sampling functions. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT symposium on principles of programming languages, POPL 2005, long beach, california, usa, january 12-14, 2005* (pp. 171–182).
- Pitts, A. (2011). Howe’s method for higher-order languages. In D. Sangiorgi & J. Rutten (Eds.), *Advanced topics in bisimulation and coinduction* (Vol. 52, pp. 197–232). Cambridge University Press.
- Plotkin, G. (1973). *Lambda-definability and logical relations*. (Technical Report SAI-RM-4, School of A.I., University of Edinburgh)
- Plotkin, G. D., & Power, J. (2001). Adequacy for algebraic effects. In *Proc. of FOSSACS 2001* (pp. 1–24).
- Plotkin, G. D., & Power, J. (2002). Notions of computation determine monads. In *Proc. of FOSSACS 2002* (pp. 342–356).
- Pollard, D., Gill, R., & Ripley, B. (2002). *A user’s guide to measure theoretic probability*. Cambridge University Press.
- Ramsey, N., & Pfeffer, A. (2002). Stochastic lambda calculus and monads of probability distributions. In *Proc. of POPL 2002* (pp. 154–165).
- Reed, J., & Pierce, B. (2010). Distance makes the types grow stronger: a calculus for differential privacy. In *Proc. of ICFP 2010* (pp. 157–168).
- Reynolds, J. (1983). Types, abstraction and parametric polymorphism. In *IFIP congress* (pp. 513–523).
- Sangiorgi, D. (1994). The lazy lambda calculus in a concurrency scenario. *Inf. Comput.*, 111(1), 120–153.
- Sangiorgi, D., Kobayashi, N., & Sumii, E. (2011). Environmental bisimulations for higher-order languages. *ACM Trans. Program. Lang. Syst.*, 33(1), 5:1–5:69.
- Sato, T. (2016). Approximate relational hoare logic for continuous random samplings. *Electr. Notes Theor. Comput. Sci.*, 325, 277–298.
- Scott, D. (1970, November). *Outline of a mathematical theory of computation* (Tech. Rep. No. PRG02). OUCL.
- Scott, D., & Strachey, C. (1971, August). *Toward a mathematical semantics for computer languages* (Tech. Rep. No. PRG06). OUCL.
- Sieber, K. (1992). Reasoning about sequential functions via logical relations. In J. P. T. Fourman M. P. & A. M. Pitts (Eds.), *Applications of categories in computer science* (Vol. 177, pp. 258–269). Cambridge

University Press.

- Simpson, A., & Voorneveld, N. (2018). Behavioural equivalence via modalities for algebraic effects. In *Proc. of ESOP 2018* (pp. 300–326).
- Staton, S. (2017). Commutative semantics for probabilistic programming. In *Programming languages and systems - 26th european symposium on programming, ESOP 2017, held as part of the european joint conferences on theory and practice of software, ETAPS 2017, uppsala, sweden, april 22-29, 2017, proceedings* (pp. 855–879).
- Staton, S., Yang, H., Wood, F. D., Heunen, C., & Kammar, O. (2016). Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *Proceedings of the 31st annual ACM/IEEE symposium on logic in computer science, LICS '16, new york, ny, usa, july 5-8, 2016* (pp. 525–534).
- Thijs, A. (1996). *Simulation and fixpoint semantics*. Rijksuniversiteit Groningen.
- Vákár, M., Kammar, O., & Staton, S. (2019). A domain theory for statistical probabilistic programming. *PACMPL*, 3(POPL), 36:1–36:29.
- Van Breugel, F., Mislove, M., Ouaknine, J., & Worrell, J. (2005). Domain theory, testing and simulation for labelled markov processes. *Theor. Comput. Sci.*, 333(1-2), 171–197.
- Wand, M., Culpepper, R., Giannakopoulos, T., & Cobb, A. (2018). Contextual equivalence for a probabilistic language with continuous random variables and recursion. *PACMPL*, 2(ICFP), 87:1–87:30.