A Domain Theory for Statistical Probabilistic Programming

MATTHIJS VÁKÁR, Columbia University, USA

OHAD KAMMAR, University of Oxford, UK

SAM STATON, University of Oxford, UK

We give an adequate denotational semantics for languages with recursive higher-order types, continuous probability distributions, and soft constraints. These are expressive languages for building Bayesian models of the kinds used in computational statistics and machine learning. Among them are untyped languages, similar to Church and WebPPL, because our semantics allows recursive mixed-variance datatypes. Our semantics justifies important program equivalences including commutativity.

Our new semantic model is based on ‘quasi-Borel predomains’. These are a mixture of chain-complete partial orders (cpo) and quasi-Borel spaces. Quasi-Borel spaces are a recent model of probability theory that focuses on sets of admissible random elements. Probability is traditionally treated in cpo models using probabilistic powerdomains, but these are not known to be commutative on any class of cpo with higher order functions. By contrast, quasi-Borel predomains do support both a commutative probabilistic powerdomain and higher-order functions. As we show, quasi-Borel predomains form both a model of Fiore’s axiomatic domain theory and a model of Kock’s synthetic measure theory.

CCS Concepts: • **Theory of computation** → **Probabilistic computation; Bayesian analysis; Denotational semantics**; • **Software and its engineering** → **Language types; Functional languages; Interpreters**; Domain specific languages; • **Computing methodologies** → *Machine learning*;

Additional Key Words and Phrases: denotational semantics, domain theory, probability, recursion, adequacy

ACM Reference Format:

Matthijs Vákár, Ohad Kammar, and Sam Staton. 2019. A Domain Theory for Statistical Probabilistic Programming. *Proc. ACM Program. Lang.* 3, POPL, Article 36 (January 2019), 29 pages. <https://doi.org/10.1145/3290349>

1 INTRODUCTION

The idea of statistical probabilistic programming is to use a programming language to specify statistical models and inference problems. It enables rapidly prototyping different models, because: (1) the model specification is separated from the technicalities of the inference/simulation algorithms; and (2) software engineering/programming techniques can be used to manage the complexity of statistical models. Here, we focus on the fundamental programming technique of *recursion*. We consider both recursive terms — looping — and recursive types, e.g., streams and untyped languages.

In a traditional programming language, recursion has long been analyzed using semantic domains based on ω -complete partial orders. The suitability of this approach is given by the adequacy theorem, which connects the compositional interpretation in domains with an operational interpretation. However, there are long standing open problems regarding using these domains with probability and measure. Here we sidestep these problems by introducing a new notion: *ω -quasi Borel spaces*. We look at a language with both statistical constructs and higher-order recursive types and terms, which is close to languages used in practice for statistical probabilistic programming, and we show the following adequacy theorem:

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

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/1-ART36

<https://doi.org/10.1145/3290349>

Fig. 1. Bayesian linear regression as (a) a first-order probabilistic program, (b) an informal specification, and (c) a plot of the prior and posterior distributions. Here $N(\mu, \sigma)$ is the normal (Gaussian) distribution with density $\text{normal-pdf}(x \mid \mu, \sigma) = (2\pi\sigma^2)^{-\frac{1}{2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$ and random number generator $\text{normal-rng}(\mu, \sigma)$.

THEOREM (6.9: ADEQUACY). *If two programs are equal in the ω qbs model then they are contextually equivalent with regard to a Monte-Carlo operational semantics.*

Domain theoretic semantics can verify compositional compiler optimizations. As an example, the following reordering transformation is valid, since it is readily verifiable in the ω qbses.

COROLLARY (6.11: COMMUTATIVITY). *The following two programs are contextually equivalent:*

$$\begin{array}{l} \text{let } x = t \text{ in} \\ \text{let } y = s \text{ in } r \end{array} \quad \approx \quad \begin{array}{l} \text{let } y = s \text{ in} \\ \text{let } x = t \text{ in } r \end{array}$$

This property says there is no implicit sequential state in the language. It is essential for [Shan and Ramsey](#)'s disintegration-based exact Bayesian inference technique [2017], implemented in the Hakaru system [Narayanan et al. 2016]. The corollary is related to Fubini's theorem for reordering integrals: informally, $\int dx \int dy r(x, y) = \int dy \int dx r(x, y)$. The important novelty here is that our semantic model extends this commutativity theorem to higher-order and recursive types, even if they do not fit easily into traditional measure theory.

1.1 Introduction to Statistical Probabilistic Programming

We introduce statistical probabilistic programming through a simple example of a regression problem, in Fig. 1. The problem is: supposing that there is a linear function $x \mapsto ax$ and three noisy measurements (1, 1.1), (2, 1.9) and (3, 2.7) of it with postulated noise scale 0.25, find a posterior distribution on the slope a . As indicated, first-order probabilistic programs can be thought of as a direct translation of a Bayesian statistical problem. The probabilistic program has an operational reading in terms of Monte-Carlo simulation: first use a Gaussian sampler/random number generator, using `normal-rng`, to draw from the prior with mean 0 and standard deviation 2; then weight, using `score`, the resulting samples with respect to the three data points according to the Gaussian likelihood, using `normal-pdf`, assuming noisy measurements with standard deviation $\frac{1}{4}$. The resulting program represents the *unnormalised* posterior distribution, which can then be passed to an inference algorithm to approximate its normalisation, e.g. through Monte-Carlo sampling.

A case for recursion: higher-order functions over infinite data structures. Many probabilistic programming languages also allow other programming language features, including recursion. When looking at a whole closed program of ground type, these extra features pose little conceptual problem because the entire program will reduce to a first-order program, albeit a very large or even

Fig. 2. A Gaussian random walk as (a) a higher-order recursive program and (b) an informal specification, with (c) some samples. The function *reduce* is sometimes called *fold left*.

infinite one. However, this paper is not concerned with the problem of interpreting whole closed programs, but rather interpreting individual aspects of a program in a compositional way.

For example, consider the program in Fig. 2, which takes snapshots of a Gaussian random walk at a stream of times (t) to produce a stream of (t, y) co-ordinate pairs. The meaning of the whole program is clear, and can be reduced to the first-order statistical model (b). But supposing this appears as part of a bigger model, we would like to understand each part separately. What is the mathematical meaning of *reduce* here? It takes a parameterized random operation, an initial value, and a stream, and produces a random stream. By providing mathematical objects that represent these recursive concepts, we can understand *reduce* as a first-class construct and reason about it.

A case for recursion: untyped programs. Rather than distinguish between ground types and higher-order recursive types, an alternative approach is to combine the full *untyped* lambda calculus with the constructions of the simple statistical programming language (Fig. 7). Many probabilistic programming languages take this approach (Church [Goodman et al. 2008], Anglican [Wood et al. 2014], WebPPL [Goodman and Stuhlmüller 2014], Venture [Mansinghka et al. 2014]). Recall that we can express untyped calculi using recursive types, using a single recursive type like $\Lambda = (\mathbb{1} \mid \Lambda * \Lambda \mid \dots \mid \Lambda \rightarrow \Lambda)$. Thus a language with recursive types can be thought of as generalising this untyped situation.

1.2 Summary of Semantics in ω -Quasi Borel Spaces

The usual method for interpreting a programming language is as follows:

- types denote spaces (in an untyped language, there is just one universal space);
- closed programs denote points in a space; and
- program phrases denote functions assigning a point to every valuation of their free variables.

Probabilistic programs, on the other hand, vary this by saying

- closed programs denote *measures* (or distributions) on a space; and
- program phrases denote kernels between spaces.

Recursive probabilistic programming has a tension between what a space is and what it is used for:

- In traditional domain theory, a type denotes a topological space or cpo in which continuity and convergence model recursion using fixed points.
- For first-order probabilistic programs, a type denotes a topological or measurable space whose purpose is to support measures and expectations.

It is tempting to use a single topological (or cpo) structure to interpret both the probability and recursion. By contrast with existing mainstream approaches [Jones and Plotkin 1989], however,

we choose to keep both structures separate but compatible, through the following observation. In probability theory it is widely acknowledged that topological and measurable structure are only a precursor to the notion of random element. Recall that a random element in a set P is a function $\alpha : \Omega \rightarrow P$ where Ω is some space of random seeds, e.g. $\Omega = \mathbb{R}$ with a Lebesgue measure. When $P = \mathbb{R}$, one would typically ask that α be measurable so that there is an expected (average, mean) value of α . But random elements are relevant beyond their expectation. A measure on P can be understood as a random element modulo being ‘equal in distribution’, but it is helpful to also keep the distinction between random element and measure. This leads us to the following definition.

DEFINITION (3.5). *An ω qbs comprises a set P together with the following structure:*

- *a partial order $\leq$ on P such that limits of ω -chains exist: to model recursion;*
- *a set M of functions $\alpha : \mathbb{R} \rightarrow P$: these are thought of as random elements, to interpret probability;*

all subject to some compatibility conditions.

For example:

- (1) Let P be the set of subsets of $\mathbb{N}$. This is a space to interpret deterministic programs with a natural number argument: a subset $X \subseteq \mathbb{N}$ represents the program that returns for each member in X . The order $\leq$ is the inclusion order, so program X is below Y if X diverges whenever Y diverges. The random elements M are the functions $\alpha : \mathbb{R} \rightarrow P$ such that $\alpha^{-1}[\{S \mid n \in S\}]$ is (Borel) measurable for all $n \in \mathbb{N}$.
- (2) Let $\mathbb{R}$ be the set of real *values*. When thought of as values, we use the discrete order. The random elements M are the measurable functions $\mathbb{R} \rightarrow \mathbb{R}$.
- (3) Let $\mathbb{W} = [0, \infty]$ be the non-negative extended reals. Its elements stand for computation *weights* with the linear order. The random elements M are the measurable functions $\mathbb{R} \rightarrow [0, \infty]$.

The semantics in ω qbses supports higher-order functions. Moreover, we can interpret recursive types by using the recipe of [Fiore and Plotkin](#)’s axiomatic domain theory [1994]. That is to say:

THEOREM (COR. 3.10, §5.3). *The category ω Qbs has products, sums, function spaces, and a bilimit compact expansion (sufficient structure to interpret recursive types).*

1.3 A Probabilistic Powerdomain

With higher-order functions and recursive types dealt with, the remaining ingredient is measures (and probabilistic programs that generate measures) as first-class constructions. To this end, for every ω qbs P we will associate an ω qbs $T(P)$ of measures on P . Following [Moggi](#) [1989], we turn T into a monad encapsulating the probabilistic aspects of the programming language.

Recall that an ω qbs P comes with a set M of random elements, viz. functions $\alpha : \mathbb{R} \rightarrow P$. Because a statistical probabilistic program naturally describes an unnormalized posterior measure, we consider the basic space $\mathbb{R}$ with the full Lebesgue measure, which has infinite total measure. To consider finite measures we consider partial random elements, which are given by pairs (α, D) where $\alpha \in M$ and $D \subseteq \mathbb{R}$ is Borel. Given any partial random element (α, D) and any morphism $f : P \rightarrow \mathbb{W}$ ($\mathbb{W} = [0, \infty]$), the composite $f \circ \alpha$ is measurable, so we can find an expectation for f :

$$\mathbb{E}_{(\alpha, D)}[f] := \int_D f(\alpha(x)) \, dx. \quad (1)$$

We say that (α, D) and (α', D') are *equivalent* when they give the same expectation operator:

$$\text{for all } f : P \rightarrow \mathbb{W}, \quad \int_D f(\alpha(x)) \, dx = \int_{D'} f(\alpha'(x)) \, dx. \quad (2)$$

DEFINITION. A measure on an ω qbs P is an equivalence class of a partial random element (α, D) , modulo the equivalence relation (2). Equivalently, a measure is a morphism $\mathbb{W}^P \rightarrow \mathbb{W}$ of the form $\mathbb{E}_{(\alpha, D)}$ for some partial random element (α, D) .

Here we run into a technical problem: the set of all measures has a natural (pointwise) partial order structure but this set might not be closed under suprema of ω -chains. On the other hand, we know that $J(P) := \mathbb{W}^{(\mathbb{W}^P)}$ is always closed, because J is the continuation monad which makes sense in any category with function spaces. Thus we take the closure $T(P)$ of the set of measures in $J(P)$ as our space of measures. In other words, $T(P) \subseteq J(P)$ contains those expectation operators that arise as iterated suprema of ω -chains on P .

Our approach follows existing continuation-passing-style techniques [e.g. Keimel and Plotkin 2009; Kiselyov and Shan 2009; Olmedo et al. 2016]. CPS semantics is analogous to working with the full continuation monad J or a fragment of it. This fragment must be chosen carefully, or else the commutativity property fails in the model. Indeed J also has constants such as $\text{exit}_r = \lambda k. r \in J(P)$ violating the commutativity equation: $\text{put } t_1 = \text{exit}_1, t_2 = \text{exit}_2$. As we have the commutativity property, these constants lie outside our monad $T(P)$, hence are not definable in the language.

Aside on the Jung-Tix problem. A long standing problem in traditional domain theory is to find a category of Scott domains that is closed under function spaces and a commutative probabilistic powerdomain [Jung and Tix 1998]. This remains an open problem. We side-step this problem by using ω qbses instead of Scott domains. They inherit many of the properties and intuitions of ω -cpo's, are closed under function spaces, and support a commutative probabilistic powerdomain. We summarize further work in this direction in Sec. 7.

1.4 Summary

We have provided a domain theory for recursion in statistical probabilistic programming. The main contributions of this work are the following novel constructions:

- (1) a Cartesian closed category of (pre)-domains (Sec. 3), that admits the solution of recursive domain equations (Sec. 5),
- (2) a commutative probabilistic power-domain (Sec. 4);
- (3) an adequate denotational model (Sec. 6) for probabilistic programming with recursive types (Sec. 2), and in consequence
 - an adequate denotational model for a higher-order language with sampling from continuous distributions, term recursion and soft constraints (§2.4);
 - an adequate denotational model for untyped probabilistic programming (§2.3).

2 CALCULI FOR STATISTICAL PROBABILISTIC PROGRAMMING

We consider three call-by-value calculi for statistical probabilistic programming. The main calculus, Statistical FPC (SFPC) is a statistical variant of Fiore and Plotkin's Fixed-Point Calculus (FPC) [1994]. SFPC has sum, product, function, and recursive types, as well as a ground type $\mathbb{R}$ of real numbers, constants $\underline{f}$ for all measurable functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$, a construct **match** – **with** $\{\underline{0} \rightarrow - \mid - \rightarrow -\}$ for testing real numbers for zero, a construct **sample** for drawing a random number using the uniform (Lebesgue) measure $\mathbb{U}_{[0,1]}$ on $[0, 1]$ and a construct **score** for reweighting program traces (to implement soft constraints). We express the other two calculi as fragments of SFPC. The first, Idealised Church, is an untyped λ -calculus for statistical probabilistic programming based on the Church modelling language [Wingate et al. 2011]. The second, the Call-by-Value Statistical PCF (CBV SPCF), is a call-by-value variant of Plotkin's [1977] and Scott's [1993] simply typed λ -calculus with higher-order term recursion, extended with statistical primitives.

2.1 Preliminaries: Borel Measurability

We need the following fundamentals of measure theory. The *Borel subsets* of the real line $\mathbb{R}$ are given inductively by taking every interval $[a, b]$ to be a Borel subset, and closing under complements and countable unions. More generally, for every natural number n , the Borel subsets of $\mathbb{R}^n$ are given inductively by taking every n -dimensional box $[a_1, b_1] \times \dots \times [a_n, b_n]$ to be a Borel subset and closing under complements and countable unions. A (partial) function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is *Borel-measurable* when its inverse image maps every Borel subset $B \subseteq \mathbb{R}$ to a Borel subset $f^{-1}[B] \subseteq \mathbb{R}^n$. The set of Borel-measurable functions contains, for example, all the elementary functions.

A *measure* μ on $\mathbb{R}^n$ is an assignment of possibly-infinite, non-negative real values $\mu(B)$ to every Borel subset $B \subseteq \mathbb{R}^n$, assigning 0 to the empty set $\mu(\emptyset) = 0$, linear on disjoint unions $\mu(B_1 \uplus B_2) = \mu(B_1) + \mu(B_2)$, and continuous with respect to countably increasing sequences: if for every n , $B_n \subseteq B_{n+1}$, then $\mu(\bigcup_{n=0}^{\infty} B_n) = \lim_{n \rightarrow \infty} \mu(B_n)$. The *Lebesgue measure* λ is the unique measure on $\mathbb{R}$ assigning to each interval its length $\lambda([a, b]) = b - a$. A *probability measure* on $\mathbb{R}^n$ is a measure μ whose *total measure* $\mu(\mathbb{R}^n)$ is 1. The *uniform probability measure* $\mathbb{U}_{[a, b]}$ on an interval $[a, b]$ assigns to each Borel set B the relative Lebesgue measure it occupies in the interval $[a, b]$: $\mathbb{U}_{[a, b]}(B) := \frac{1}{b-a} \lambda(B \cap [a, b])$. Measurable functions $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ let us transport measures μ on $\mathbb{R}^n$ to their *push-forward* measure $f_*\mu$ on $\mathbb{R}^m$, by setting $f_*\mu(B) := \mu(f^{-1}[B])$.

We think of whole statistical probabilistic programs of real type as a formalism for describing measures. To work compositionally, we need the following analogous concept for program fragments, i.e., terms with unbound variables. A *probability kernel* k from $\mathbb{R}^n$ to $\mathbb{R}$, written as $k : \mathbb{R}^n \rightsquigarrow \mathbb{R}$ is a function assigning to every $\vec{x} \in \mathbb{R}^n$ a probability measure $k(\vec{x}, -)$ on $\mathbb{R}$, such that, for every Borel set B , the function $k(-, B) : \mathbb{R}^n \rightarrow \mathbb{R}$ is measurable. We will use the following key result about probability kernels, the *randomisation lemma*, which says that a straightforward random number generator $\mathbb{U}_{[0, 1]}$ suffices to implement any of them.

LEMMA 2.1 ([KALLENBERG 2006, LEMMA 3.22]). *For every probability kernel $k : \mathbb{R}^n \rightsquigarrow \mathbb{R}$, there is a measurable function $\text{rand}_k : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$, such that $\text{rand}_k(x_1, \dots, x_n, -) \cdot \mathbb{U}_{[0, 1]} = k(x_1, \dots, x_n)$.*

2.2 SFPC: Bayesian Statistical Modelling with Recursive Types

Syntax. As recursive types contain type-variables, we use a kind system to ensure types are well-formed. Fig. 3 (top left) presents the kinds of our calculus, and Fig. 3 (bottom left) presents the types of SFPC. We include type variables, taken from a countable set ranged over by α, β, γ . We include simple types: unit, product, function, and variant types. Variant types use constructor labels taken from a countable set ranged over by $\ell, \ell_1, \ell_2, \dots$. In our abstract syntax, variant types are *dictionaries*, partial functions with a finite domain, from the set of constructor labels to the set of types. The recursive type former $\mu\alpha.\tau$ binds α in τ . In our abstract syntax, term variable contexts Γ are dictionaries from the countable set of variables, ranged over by $x, y, z, \dots$, to the set of types.

We desugar stand-alone labels in a variant type $\{\dots \mid \ell \mid \dots\}$ to the unit type $\{\dots \mid \ell \mid 1 \mid \dots\}$. We also desugar top-level-like recursive type declarations $\tau := \sigma[\alpha \mapsto \tau]$ to $\tau := \mu\alpha.\sigma$.

Example 2.2. The type of booleans is $\text{bool} := \{\text{True} \mid \text{False}\}$. The type of natural numbers is $\mathbb{N} := \{\text{Zero} \mid \text{Succ } \mathbb{N}\}$ desugaring to $\mathbb{N} := \mu\alpha.\{\text{Zero} \mid \text{Succ } \alpha\}$. The type of τ -lists is $\text{List } \tau := \{\text{Nil} \mid \text{Cons } \tau * \text{List } \tau\}$, desugaring to $\text{List } \tau := \mu\alpha.\{\text{Nil } 1 \mid \text{Cons } \tau * \alpha\}$. The type of (infinite) stochastic processes in τ is $\text{Stoch } \tau := 1 \rightarrow (\tau * \text{Stoch } \tau)$, desugaring to $\text{Stoch } \tau := \mu\alpha.1 \rightarrow (\tau * \alpha)$. The type of untyped λ -terms with values in τ is $\Lambda\text{-Term } \tau := \{\text{Val } \tau \mid \text{Fun } (\Lambda\text{-Term } \tau \rightarrow \Lambda\text{-Term } \tau)\}$, desugaring to $\Lambda\text{-Term } \tau := \mu\alpha.\{\text{Val } \tau \mid \text{Fun } (\alpha \rightarrow \alpha)\}$.

Fig. 3 (right) presents the terms of SFPC. Value variables are taken from a countable set ranged over by $x, y, z, \dots$. We include primitive function constants $\underline{f}(t_1, \dots, t_n)$ for every measurable

$k ::=$	kinds	$t, s, r ::=$	terms	function:
type	type	x	term variable	$ \lambda x : \tau. t$ abstraction
context	context	$ \underline{f}(t_1, \dots, t_n)$	primitive	$ \underline{t}$ application
$\tau, \sigma, \rho ::=$	types	match t with	conditional	pattern matching:
α	variable	$\{ \underline{0} \rightarrow s \mid _ \rightarrow r \}$		match t with unit
$ \underline{\mathbb{R}}$	reals	sample	sampling	$() \rightarrow s$
$ \mathbf{1}$	unit	score t	conditioning	match t with product
$ \tau * \sigma$	product		constructors:	$(x, y) \rightarrow s$
$ \tau \rightarrow \sigma$	function	$ \underline{()}$	unit	match t with variant
$ \{ \ell_1 \tau_1$	variant	$ \underline{(t, s)}$	pairing	$\{ \ell_1 x_1 \rightarrow s_1 \mid \dots$
$ \dots$		$ \underline{\tau. \ell t}$	variant	$ \ell_n x_n \rightarrow s_n \}$
$ \ell_n \tau_n \}$		$ \underline{\tau. \text{roll}}(t)$	iso-recursive	match t with recursive
$ \mu\alpha. \tau$	iso-recursive			roll $x \rightarrow s$
$\Gamma := x_1 : \tau_1, \dots, x_n : \tau_n$ variable contexts				

Fig. 3. SFPC kinds and types (left) and terms (right)

Sugar	Elaboration
• $_$	fresh variable
• let $(x : \tau) = t$ in s	$(\lambda x : \tau. s) t$
• letrec $(x : \rho) = t$ in s	$(\lambda x : \rho. s)(\mu x : \rho. t)$
• $t; s$	match t with $() \rightarrow s$
• unroll t	match t with roll $x \rightarrow x$
	• $\mu x : \tau. t$ let $(\text{body} : \sigma \rightarrow \tau) = \lambda y : \sigma.$ let $(x : \tau) = \lambda z : \tau_1. \text{unroll } y \text{ in } t$ in $\text{body}(\sigma. \text{roll } \text{body})$ where $\tau := \tau_1 \rightarrow \tau_2$ and $\sigma := \mu\alpha. \alpha \rightarrow \tau$ • $\perp_\rho$ $\mu x : \rho. x$

Fig. 4. Term-level syntactic sugar

(partial) function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. The conditional construct tests whether its argument of type $\underline{\mathbb{R}}$ evaluates to 0. We include an effect **sample** for sampling a real number uniformly from the interval $[0, 1]$, for defining the prior distribution of the program. We also include an effect **score** r for reweighting the posterior distribution by the non-negative factor $|r| \in [0, \infty)$. We include standard constructors and pattern-matching constructs for the simple types, and standard function abstraction and application constructs. Finally, we include the standard *iso-recursive* constructors and pattern matching, which require an explicit rolling and unrolling of the recursive definition such as $\underline{\mathbb{N}}.\text{roll}(\text{Zero}())$. Variant and iso-recursive constructor terms as well as function abstraction annotate their binding occurrences with the appropriate closed type τ to ensure unique typing.

To aid readability, we use the standard syntactic sugar of Fig. 4, e.g. **let** $x = t$ **in** s for $(\lambda x. t)s$, $\mu x : \tau \rightarrow \sigma. t$ for the usual encoding of term level recursion using type level recursion [Abadi and Fiore 1996; Fiore 1996]. When dealing with a recursive variant type $\sigma = \mu\alpha. \{ \dots \mid \ell \tau \mid \dots \}$, we write ℓt for the more cumbersome constructor $\sigma.\text{roll}(\{ \dots \mid \ell \tau \mid \dots \}. \ell t)$ as long as σ is clear from the context. We can encode constructs $\underline{k}$ in our calculus for drawing from arbitrary probability kernels $k : \mathbb{R}^n \rightsquigarrow \mathbb{R}$. Using the Randomisation Lemma 2.1, first find the appropriate measurable function $\text{rand}_k : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ and express $\underline{k}(t_1, \dots, t_n)$ by $\text{rand}_k(t_1, \dots, t_n, \text{sample})$.

Example 2.3. For $\text{Stoch } \underline{\mathbb{R}} = \mu\alpha. \mathbf{1} \rightarrow (\underline{\mathbb{R}} * \alpha)$, define **draw** $:= \lambda x : \text{Stoch } \underline{\mathbb{R}}. \text{unroll } (x) ()$, which draws a value from the process and moves the process to the corresponding new state. As an example stochastic process, writing $\text{normal_rng} : \mathbb{R} \times \mathbb{R} \rightsquigarrow \mathbb{R}$ for a Gaussian probability kernel taking the mean and standard deviation as arguments, we define an example Gaussian random

$$\begin{array}{c}
\frac{}{\Delta \vdash_k \alpha : \text{type}} (\alpha \in \Delta) \quad \frac{}{\Delta \vdash_k \underline{\mathbb{R}} : \text{type}} \quad \frac{}{\Delta \vdash_k 1 : \text{type}} \quad \frac{\Delta \vdash_k \tau : \text{type} \quad \Delta \vdash_k \sigma : \text{type}}{\Delta \vdash_k \tau * \sigma : \text{type}} \\
\frac{\Delta \vdash_k \tau : \text{type} \quad \Delta \vdash_k \sigma : \text{type}}{\Delta \vdash_k \tau \rightarrow \sigma : \text{type}} \quad \frac{\text{for all } 1 \leq i \leq n: \quad \Delta \vdash_k \tau_i : \text{type}}{\Delta \vdash_k \{\ell_1 \tau_1 \mid \dots \mid \ell_n \tau_n\} : \text{type}} \quad \frac{\Delta, \alpha \vdash_k \tau : \text{type}}{\Delta \vdash_k \mu\alpha.\tau : \text{type}} \quad \frac{\text{for all } (x : \tau) \in \Gamma: \quad \vdash_k \tau : \text{type}}{\vdash_k \Gamma : \text{context}}
\end{array}$$

Fig. 5. SFPC kind system

$$\begin{array}{c}
\frac{(x : \tau) \in \Gamma}{\Gamma \vdash x : \tau} \quad \frac{f : \mathbb{R}^n \rightarrow \mathbb{R} \quad \text{for all } 1 \leq i \leq n: \Gamma \vdash t_i : \underline{\mathbb{R}}}{\Gamma \vdash \underline{f}(t_1, \dots, t_n) : \underline{\mathbb{R}}} \quad \frac{\Gamma \vdash t : \underline{\mathbb{R}} \quad \Gamma \vdash s : \tau \quad \Gamma \vdash r : \tau}{\Gamma \vdash \text{match } t \text{ with } \{0 \rightarrow s \mid _ \rightarrow r\} : \tau} \\
\frac{}{\Gamma \vdash \text{sample} : \underline{\mathbb{R}}} \quad \frac{\Gamma \vdash t : \underline{\mathbb{R}}}{\Gamma \vdash \text{score } t : 1} \quad \frac{}{\Gamma \vdash () : 1} \quad \frac{\Gamma \vdash t : \tau \quad \Gamma \vdash s : \sigma}{\Gamma \vdash (t, s) : \tau * \sigma} \\
\frac{\Gamma \vdash t : \tau_i}{\Gamma \vdash \tau.\ell_i t : \tau} (\tau = \{\ell_1 \tau_1 \mid \dots \mid \ell_n \tau_n\}) \quad \frac{\Gamma \vdash t : \sigma[\alpha \mapsto \tau]}{\Gamma \vdash \tau.\text{roll}(t) : \tau} (\tau = \mu\alpha.\sigma) \quad \frac{\Gamma, x : \tau \vdash t : \sigma}{\Gamma \vdash \lambda x : \tau. t : \tau \rightarrow \sigma} \\
\frac{\Gamma \vdash t : \sigma \rightarrow \tau \quad \Gamma \vdash s : \sigma}{\Gamma \vdash ts : \tau} \quad \frac{\Gamma \vdash t : \sigma * \rho \quad \Gamma, x : \sigma, y : \rho \vdash s : \tau}{\Gamma \vdash \text{match } t \text{ with } (x, y) \rightarrow s : \tau} \quad \frac{\Gamma \vdash t : 1 \quad \Gamma \vdash s : \tau}{\Gamma \vdash \text{match } t \text{ with } () \rightarrow s : \tau} \\
\frac{\Gamma \vdash t : \{\ell_1 \tau_1 \mid \dots \mid \ell_n \tau_n\} \quad \text{for each } 1 \leq i \leq n: \Gamma, x_i : \tau_i \vdash s_i : \tau}{\Gamma \vdash \text{match } t \text{ with } \{\ell_1 x_1 \rightarrow s_1 \mid \dots \mid \ell_n x_n \rightarrow s_n\} : \tau} \\
\frac{\Gamma \vdash t : \mu\alpha.\sigma \quad \Gamma, x : \sigma[\alpha \mapsto \mu\alpha.\sigma] \vdash s : \tau}{\Gamma \vdash \text{match } t \text{ with roll } x \rightarrow s : \tau} \quad \frac{\Gamma, x : \tau \rightarrow \sigma \vdash t : \tau \rightarrow \sigma}{\Gamma \vdash \mu x : \tau \rightarrow \sigma. t : \tau \rightarrow \sigma}
\end{array}$$

Fig. 6. SFPC type system, including (in gray) the derivable typing rule for term recursion

walk $\text{RW } \underline{\sigma} \underline{\mu}$ with initial position $\underline{\mu}$ and standard deviation σ by setting:

$$\text{RW} = \lambda x : \underline{\mathbb{R}}. \mu y : \underline{\mathbb{R}} \rightarrow \text{Stoch } \underline{\mathbb{R}}. \lambda z : \underline{\mathbb{R}}. \text{Stoch } \underline{\mathbb{R}}. \text{roll}(\lambda _ : 1. (z, y(\text{normal_rng}(z, x))))$$

Kind and Type Systems. To ensure the well-formedness of types, which involve type variables, we use the kind system presented in Fig. 5. Each kinding judgement $\Delta \vdash_k \tau : \text{type}$ asserts that a given type τ is well-formed in the *type variable context* Δ , which is a finite set of type variables. The kinding judgements are standard. All type variables must be bound by the enclosing context, or by a recursive type binder. *Variable contexts* Γ must assign closed types. We treat α -conversion and capture avoiding substitution of variables as usual, possessing the standard structural properties.

Fig. 6 presents the resulting type system, including the derivable typing judgement for the sugar $\mu x : \tau \rightarrow \sigma. t$ for term recursion. Each typing judgement $\Gamma \vdash t : \tau$ asserts that the term t is well-typed with the well-formed closed type $\vdash_k \tau : \text{type}$ in the variable context $\vdash_k \Gamma : \text{context}$. The rules are standard. By design, every term has at most one type in a given context.

Example 2.4. The types from Ex. 2.2 are well-kinded: $\text{bool}, \mathbb{N}, \text{List } \underline{\mathbb{R}}, \text{Stoch } \underline{\mathbb{R}}, \Lambda\text{-Term } \underline{\mathbb{R}} : \text{type}$. The terms from Ex. 2.3 are well-typed: $\text{draw} : \text{Stoch } \underline{\mathbb{R}} \rightarrow (\underline{\mathbb{R}} * \text{Stoch } \underline{\mathbb{R}})$, $\text{RW} : \underline{\mathbb{R}} \rightarrow \underline{\mathbb{R}} \rightarrow \text{Stoch } \underline{\mathbb{R}}$.

$t, s, r ::=$	terms		v	values
x	variable	factor t	$::= x$	variable
$ \underline{k}(t_1, \dots, t_n)$	kernel	$\lambda x. t$	abstraction	$\underline{r}$ real number
$ \text{ifz } t \text{ then } s \text{ else } r$	conditional	$t \ s$	application	$\lambda x. t$ abstraction

Fig. 7. Idealised Church, terms and values

Values v	Auxiliary sequential unpacking
$x \mapsto x \quad \underline{r} \mapsto \text{Val } \underline{r} \quad \lambda x. t \mapsto \text{Fun } (\lambda x : \tau. t^\dagger)$	$(u_1 := t_1, \dots, u_n := t_n; x_1 := s_1, \dots, x_m := s_m) \bullet r$
Terms t, s, r	
$v \mapsto v^\dagger$	$(;) \bullet r \mapsto r$
$t \ s \mapsto (u := t^\dagger;) \bullet u \ s^\dagger$	$(; x_1 := s_1, \dots) \bullet r \mapsto \text{match unroll } s_1 \text{ with } \{$
$\underline{k}(t_1, \dots, t_n) \mapsto (; x_1 := t_1^\dagger, \dots, x_n := t_n^\dagger) \bullet \text{Val } \underline{k}(x_1, \dots, x_n)$	$\text{Val } x_1 \rightarrow (; \dots) \bullet r$
$\text{ifz } t \text{ then } r \text{ else } s \mapsto (; x_1 := t^\dagger) \bullet \text{match } x \text{ with } \{ \underline{0} \rightarrow r^\dagger \mid _ \rightarrow s^\dagger \}$	$\mid \text{Fun } _ \rightarrow \perp_\sigma \}$
$\text{factor}(t) \mapsto (; x := t) \bullet \text{score } x; \text{Val } x$	$(u_1 := t_1, \dots; \dots) \bullet r \mapsto \text{match unroll } t_1 \text{ with } \{$
	$\text{Fun } u_1 \rightarrow (\dots; \dots) \bullet r$
	$\mid \text{Val } _ \rightarrow \perp_\sigma \}$

Fig. 8. A faithful translation $(-)^{\dagger}$ of Idealised Church into SFPC using the type $\sigma := \Lambda\text{-Term } \underline{\mathbb{R}}$.

2.3 Idealised Church: Untyped Statistical Modelling

In applied probabilistic programming systems, there has been considerable interest in using an *untyped* λ -calculus as the basis for probabilistic programming. For instance, the Church modelling language [Wingate et al. 2011], in idealised form, is an untyped call-by-value λ -calculus over the real numbers with constructs $\underline{k}(x_1, \dots, x_n)$ for drawing from probability kernels $k : \mathbb{R}^n \rightsquigarrow \mathbb{R}$ (including all real numbers, measurable functions like a Gaussian density *normal_pdf*($y \mid \mu, \sigma$) and proper kernels like a Gaussian random number generator *normal_rng*(μ, σ)), and a construct **factor** for reweighting program traces, to enforce soft constraints [Borgström et al. 2016]. Fig. 7 presents the syntax of Idealised Church, and we desugar **let** $x = t$ **in** s to mean $(\lambda x.s) t$.

Idealised Church arises as a sublanguage of SFPC. We encode Idealised Church terms as SFPC terms of the type $\Lambda\text{-Term } \mathbb{R} := \{\text{Val } \tau \mid \text{Fun } (\Lambda\text{-Term } \mathbb{R} \rightarrow \Lambda\text{-Term } \mathbb{R})\}$ from Ex. 2.2, using the translation $(-)^{\dagger}$ in Fig. 8. The translation uses an auxiliary SFPC construct $(\dots) \bullet r$ for sequentially evaluating its arguments $x := t$, unpacking each term t , ensuring it is either a function or a real value, and binding its unpacked value to x in r . This translation is faithful (Lemma 6.2).

2.4 CBV SPCF: Simply Typed Recursive Modelling

We consider a simply typed sublanguage of SFPC, a call-by-value (CBV) probabilistic variant of Plotkin and Scott’s PCF [1977; 1993]. The types and terms are given by the following grammars:

$$\tau, \sigma, \rho \quad ::= \underline{\mathbb{R}} \mid \tau \rightarrow \sigma \qquad t, s, r \quad ::= x \mid \underline{f}(t_1, \dots, t_n) \mid \text{match } t \text{ with } \{ \underline{0} \rightarrow s \mid _ \rightarrow r \} \mid \text{sample} \mid \text{factor}(t) \mid \lambda x : \tau. t \mid t \ s \mid \mu x : \tau \rightarrow \sigma. t$$

SPCF is a fragment of SFPC: term recursion $\mu x : \tau \rightarrow \sigma.t$ is interpreted as in Fig. 4 and conditioning **factor** (t) as **let** ($x : \mathbb{R}$) = t **in score** x ; x . SPCF derives its kind and type systems from SFPC.

3 QUASI-BOREL PRE-DOMAINS

Previous works on quasi-Borel spaces (qbses) give a denotational semantics for higher-order probabilistic languages with a range of types, but crucially excludes higher-order term recursion and recursive types. To do that, we further equip a qbs with a compatible ωcpo structure. We call this new semantic structure a *quasi-Borel pre-domain* or an ωqbs .

3.1 Preliminaries

Category theory. We assume familiarity with categories $\mathcal{C}, \mathcal{D}$, functors $F, G : \mathcal{C} \rightarrow \mathcal{D}$, natural transformations $\alpha, \beta : F \rightarrow G$, and their theory of (co)limits and adjunctions. We write:

- unary, binary, and I -ary products as $\mathbb{1}, X_1 \times X_2$, and $\prod_{i \in I} X_i$, writing π_i for the projections and $\langle \rangle, \langle x_1, x_2 \rangle$, and $\langle x_i \rangle_{i \in I}$ for the tupling maps;
- unary, binary, and I -ary coproducts as $0, X_1 + X_2$, and $\sum_{i \in I} X_i$, writing ι_i for the injections and $[], [x_1, x_2]$, and $[x_i]_{i \in I}$ for the cotupling maps;
- exponentials as X^Y , writing Λ for the currying maps.

Domain theory. We recall some basic domain theory. Let $\omega = \{0 \leq 1 \leq \dots\}$ be the ordinary linear order on the naturals. An ω -chain in a poset $P = (|P|, \leq)$ is a monotone function $a_- : \omega \rightarrow P$. A poset P is an ωcpo when every ω -chain $\langle a_n \rangle_{n \in \mathbb{N}}$ has a least upper bound (lub) $\bigvee_{n \in \mathbb{N}} a_n$ in P .

Example 3.1. Each set X equipped with the discrete partial order forms an ωcpo $(X, =)$. E.g., the discrete ωcpo $\mathbb{R}$ over the real line. The non-negative extended reals equipped with the ordinary order, $\mathbb{W} := ([0, \infty], \leq)$, is an ωcpo . The Borel subsets of $\mathbb{R}^n$ ordered by inclusion form an ωcpo $\mathcal{B}_n$.

For every pair of ωcpo s P and Q , a *Scott-continuous* function $f : P \rightarrow Q$ is a monotone function $f : |P| \rightarrow |Q|$ such that for every ω -chain a_- , we have: $f(\bigvee_n a_n) = \bigvee_n f(a_n)$. A Scott-continuous function $f : P \rightarrow Q$ is a *full mono* when, for every $a, b \in |P|$, we have $f(a) \leq f(b) \implies a \leq b$. Recall that the category ωCpo of ωcpo s and Scott-continuous functions is Cartesian closed: products are taken componentwise and the exponential Q^P has carrier $\omega\text{Cpo}(P, Q)$ and order $f \leq_{Q^P} g$ iff $\forall p \in |P|. f(p) \leq_Q g(p)$. A *domain* is an ωcpo with a least element $\perp$. A *strict* function between domains is a Scott-continuous function that preserves their least elements.

Example 3.2. A measure on $\mathbb{R}^n$ is a strict continuous function $\mu : \mathcal{B}_n \rightarrow \mathbb{W}$ that is linear on disjoint subsets. The measurable functions $\mathcal{B}(\mathbb{R}, [0, \infty])$ ordered pointwise fully include into the ωcpo $\mathbb{W}^{\mathbb{R}}$. The *integral* is the unique Scott-continuous function $\int \mu : \mathcal{B}(\mathbb{R}, [0, \infty]) \rightarrow \mathbb{W}$ satisfying, for all measurable partitions $\mathbb{R} = \sum_{n \in \mathbb{N}} U_n$ and weights $\langle w_n \rangle$ in $\mathbb{W}$: $\int \mu [\lambda r : U_n \cdot w_n]_{n \in \mathbb{N}} = \sum_{n \in \mathbb{N}} x_n \cdot \mu U_n$.

An ωCpo -*enriched* category $\mathcal{C}$ consists of a locally-small category $\underline{\mathcal{C}}$ together with an assignment of an ωcpo $C(A, B)$ to every $A, B \in \text{Ob}(\underline{\mathcal{C}})$ whose carrier is the set $\underline{\mathcal{C}}(A, B)$ such that composition is Scott-continuous. An ωCpo -*functor*, a.k.a. a *locally-continuous functor*, $F : \mathcal{C} \rightarrow \mathcal{D}$ between two ωCpo -categories is an ordinary functor $\underline{F} : \underline{\mathcal{C}} \rightarrow \underline{\mathcal{D}}$ between the underlying ordinary categories, such that every morphism map $\underline{F}_{A, B} : \underline{\mathcal{C}}(A, B) \rightarrow \underline{\mathcal{D}}(FA, FB)$ is Scott-continuous.

Example 3.3. Every locally-small category is an ωCpo -category whose hom- ωcpo s are discrete. The category ωCpo itself is an ωCpo -category. If $\mathcal{C}$ is an ωCpo -category, its categorical dual $\mathcal{C}^{\text{op}}$ is an ωCpo -category. The category of locally-continuous functors $\mathcal{C}^{\text{op}} \rightarrow \omega\text{Cpo}$, with the order on natural transformations $\alpha : F \rightarrow G$ given componentwise, is an ωCpo -category when $\mathcal{C}$ is.

Measure theory. A measurable space $X = (|X|, \Sigma_X)$ consists of a carrier set $|X|$ and a set of subsets $\Sigma_X \subseteq \mathcal{P}(X)$, called its σ -*algebra*, containing the empty set, and closed under complements and countable unions, thus axiomatising the measurable subsets of $\mathbb{R}^n$. A measurable function $f : X \rightarrow Y$ is a function $f : |X| \rightarrow |Y|$ whose inverse image maps measurable subsets to measurable

subsets. Thus every n -dimensional Borel set, together with its Borel subsets, forms a measurable space. The measurable spaces that are measurably isomorphic to a Borel set are called *standard Borel spaces*. A fundamental result in descriptive set theory is that every standard Borel space is measurably isomorphic to $\{i \in \mathbb{N} \mid i < n\}$ for some $n = 0, 1, \dots, \omega$, or to $\mathbb{R}$ [Kechris 2012]. We write **Meas** and **Sbs** for the categories of measurable and standard Borel spaces and measurable functions between them.

Quasi-Borel spaces. As a Cartesian closed alternative to measure theory, Heunen et al. [2017] introduced the category **Qbs** of *quasi-Borel spaces* (qbses). Measure theory axiomatises measurable subsets of a space X , and deriving the *random elements*: measurable functions $\alpha : \mathbb{R} \rightarrow X$, for pushing measures forward. Qbses axiomatise random elements directly.

A *quasi-Borel space* (qbs) $X = (|X|, M_X)$ consists of a carrier set $|X|$ and a set of functions $M_X \subseteq |X|^\mathbb{R}$, called the *random elements*, such that (i) all the constant functions are in M_X , (ii) M_X is closed under precomposition with measurable functions on $\mathbb{R}$, and (iii) if $\mathbb{R} = \bigcup_{n \in \mathbb{N}} U_n$, where U_n are pairwise-disjoint and Borel measurable, and $\alpha_n \in M_X$ for all n , then the countable case-splitting $[\alpha_n|_{U_n}]_{n \in \mathbb{N}}$ is in M_X . A *morphism* $f : X \rightarrow Y$ is a structure-preserving function $f : |X| \rightarrow |Y|$, i.e. if $\alpha \in M_X$ then $(f \circ \alpha) \in M_Y$. Morphisms compose as functions, and we have a category **Qbs**.

Example 3.4. We turn the n -dimensional space $\mathbb{R}^n$ into a qbs by taking the random elements to be the measurable functions $M_{\mathbb{R}^n} := \mathbf{Meas}(\mathbb{R}, \mathbb{R}^n) \cong \prod_{i=1}^n \mathbf{Meas}(\mathbb{R}, \mathbb{R})$, i.e., n -tuples of correlated random variables. We also turn every set X into a qbs by taking the random elements to be measurably piece-wise constant functions, i.e., the *step functions*.

Both these examples are special cases of a more abstract situation. Every measurable space X can be turned into a qbs by setting $M_X := \mathbf{Meas}(\mathbb{R}, X)$. This defines a functor $M_- : \mathbf{Meas} \rightarrow \mathbf{Qbs}$. It has a left adjoint $\Sigma_- : \mathbf{Qbs} \rightarrow \mathbf{Meas}$ which equips a qbs X with the largest σ -algebra such that all random elements $\alpha \in M_X$ are measurable. This adjunction restricts to an adjoint embedding of the category of standard Borel spaces **Sbs** as a full subcategory of **Qbs**. This embedding makes **Qbs** a conservative extension of the well-behaved standard Borel spaces.

The category of qbses possesses substantial pleasant categorical properties: it has all (co)limits and is Cartesian closed, and so can interpret simple types, quotients, and refinements. In fact, **Qbs** is a Grothendieck quasi-topos, and so can interpret an expressive internal logic. The conservativity of the embedding **Sbs** $\hookrightarrow$ **Qbs** means that interpreting closed programs of ground type and reasoning about them in **Qbs** have standard measure-theoretic counterparts. The benefit comes from doing so *compositionally*: program fragments that are higher-order functions have a compositional interpretation and reasoning principles in **Qbs**, but not in **Sbs** nor in **Meas**.

3.2 Definition and Some Simple Examples

The difficulty inherent in combining domain and measure theory stems from the following considerations [Jung and Tix 1998]. Each (pre-)domain induces a topological space whose open subsets are the *Scott-open* subsets (see Ex. 4.7 and Ex. 5.1), from which one generates a measurable space structure by closing over countable unions and complements. Both the domain-theoretic structure and the induced measure-theoretic structure possess a cartesian product construction. However, without further assumptions, the two product structures may differ. To get a commutative probabilistic powerdomain, one requires conditions that ensure these two structures agree while maintaining, e.g., cartesian closure. The search for such a category is known as the *Jung-Tix problem*. We circumvent the Jung-Tix problem, without solving it, by keeping the two structures, the domain-theoretic and the measurable, separate but compatible. Doing so also allows us to replace the measure-theoretic

Carrier	Random elements	Partial order
$\mathbb{1}$	$\{r \mapsto \langle \rangle\}$	$=_{\mathbb{1}}$
$P \times Q$	$ P \times Q $	$\langle p, q \rangle \leq \langle p', q' \rangle : p \leq_P p', q \leq_Q q'$
$\sum_{i=1}^n P_i$	$\sum_{i=1}^n P_i $	$\langle j, p \rangle \leq \langle k, q \rangle : j = k, p \leq_{P_j} q$
Q^P	$\omega\mathbf{Qbs}(P, Q)$	$\{ \Lambda(\alpha) : \mathbb{R} \rightarrow Q^P \mid \alpha \in \omega\mathbf{Qbs}(\mathbb{R} \times P, Q) \}$
$P_{\perp}$	$ P + \{\perp\}$	$\perp \leq_{P_{\perp}} x, (\uparrow a \leq_{P_{\perp}} \uparrow b \iff a \leq_P b)$

Fig. 9. The simply-typed structure of $\omega\mathbf{Qbs}$

structure, which is usually incompatible with higher-order structure, with a quasi-Borel space structure. The result is the following definition:

Definition 3.5. An $\omega\mathbf{qbs}$ P consists of a triple $P = \langle |P|, M_P, \leq_P \rangle$ where: $(|P|, M_P)$ is a $\mathbf{qbs}$; $(|P|, \leq_P)$ is an $\omega\mathbf{cpo}$ over $|P|$; and M_P is closed under pointwise sups of ω -chains w.r.t. the pointwise order. A *morphism* between $\omega\mathbf{qbs}$ es $f : P \rightarrow Q$ is a Scott-continuous function between their underlying $\omega\mathbf{cpo}$ s that is also a $\mathbf{Qbs}$ -morphism between their underlying $\mathbf{qbs}$ es. We denote the category of $\omega\mathbf{qbs}$ es and their morphisms by $\omega\mathbf{Qbs}$.

Example 3.6 (Real Values). We have the $\omega\mathbf{qbs}$ $\mathbb{R} = (\mathbb{R}, \mathbf{Meas}(\mathbb{R}, \mathbb{R}), =_{\mathbb{R}})$ with the discrete order $=_{\mathbb{R}}$. This pre-domain represents a space of *values* returned by probabilistic computations.

Example 3.7 (Real Weights). Contrast this pre-domain with $\mathbb{W} = ([0, \infty], \mathbf{Meas}(\mathbb{R}, [0, \infty]), \leq_{[0, \infty]})$ with the linear order. This pre-domain represents a space of *weights* of computation traces. Compare this space to the Sierpiński space $\{0, 1\}_{\leq}$, the full subspace $\{0, 1\}$ in $\mathbb{W}$.

The category $\omega\mathbf{Qbs}$ is an $\omega\mathbf{Cpo}$ -category, with each homset in $\omega\mathbf{Qbs}$ ordered *pointwise* by setting, for every pair of morphisms $f, g \in \omega\mathbf{Qbs}(X, Y)$: $f \leq g$ when $\forall x \in |X|. f(x) \leq_Y g(x)$.

3.3 Interpreting Simple Types and Partiality

We turn every $\mathbf{qbs}$ into the *discrete* $\omega\mathbf{qbs}$ over it by taking the discrete $\omega\mathbf{cpo}$ structure, i.e., equality as an order. This construction is the left adjoint to the evident forgetful functor $|-| : \omega\mathbf{Qbs} \rightarrow \mathbf{Qbs}$. Similarly, we turn every $\omega\mathbf{cpo}$ into the *free* $\omega\mathbf{qbs}$ over it whose random elements are lub of step functions. This construction is the left adjoint to the evident forgetful functor $|-| : \omega\mathbf{Qbs} \rightarrow \omega\mathbf{Cpo}$.

Example 3.8. The discrete $\omega\mathbf{qbs}$ on the $\mathbf{qbs}$ structure of the real line is the pre-domain $\mathbb{R}$ of real values. The free $\omega\mathbf{qbs}$ on the $\omega\mathbf{cpo}$ of weights (Ex. 3.1) is the pre-domain $\mathbb{W}$ of real weights (Ex. 3.7).

These adjoints equip $\omega\mathbf{Qbs}$ with well-behaved limits and coproducts:

LEMMA 3.9. *The forgetful functors $|-| : \omega\mathbf{Qbs} \rightarrow \omega\mathbf{Cpo}$, $|-| : \omega\mathbf{Qbs} \rightarrow \mathbf{Qbs}$ preserve limits and coproducts. This uniquely determines the limits and coproducts of $\omega\mathbf{Qbs}$, which exist for small diagrams.*

In Sec. B we will see that $\omega\mathbf{Qbs}$ also has quotients, but their construction is more subtle.

The category $\omega\mathbf{Qbs}$ is bi-Cartesian closed, with the concrete structure given in Fig. 9. The structural maps, such as tupling, projections, and so forth, are given as for sets. The figure also depicts a locally-continuous *lifting* monad¹ $\langle \perp_{-}, \text{return}, \gg= \rangle$, for interpreting partiality, where:

$$\uparrow x := \iota_1 x \quad \perp := \iota_2 \perp \quad \text{return } x := \uparrow x \quad ((\uparrow x) \gg= f) := f(x) \quad (\perp \gg= f) := \perp.$$

¹See §4.1 for the definition of monads.

This monad jointly lifts the partiality monad $P \mapsto P_\perp$ over $\omega\mathbf{Cpo}$ and the exception monad $X \mapsto X + \mathbb{1}$ over $\mathbf{Qbs}$.

COROLLARY 3.10. *The functor $|-| : \omega\mathbf{Qbs} \rightarrow \mathbf{Set}$ preserves the bi-Cartesian and partiality structures.*

So simple types, when interpreted in $\omega\mathbf{Qbs}$, retain their natural set-theoretic interpretations.

4 A COMMUTATIVE STATISTICAL POWERDOMAIN

Our powerdomain construction combines two classical ideas in probability theory. The first idea is Schwartz's treatment of distributions as expectation operators. We construct the powerdomain monad T as a submonad of the continuation monad $J := \mathbb{W}^{\mathbb{W}^-}$, where $\mathbb{W}$ is the space of weights over $[0, \infty]$ from Ex. 3.7. The second idea is the Randomisation Lemma 2.1: kernels $X \rightarrow TY$ should arise by pushing forward the Lebesgue measure along a partial function $X \times \mathbb{R} \rightarrow Y$, i.e., a morphism $X \rightarrow (Y_\perp)^\mathbb{R}$. Each element of $(Y_\perp)^\mathbb{R}$ induces a *randomisable* expectation operator via the Lebesgue integral. Combining the two ideas, we take T to be the smallest full submonad of the Schwartz distribution monad that contains all the randomisable distributions.

Sampling and conditioning have natural interpretations as expectation operators:

$$\text{sample} : \mathbb{1} \rightarrow J\mathbb{R}, \text{sample}(\cdot)[w] := \int_{[0,1]} w(r) dr; \quad \text{score} : \mathbb{R} \rightarrow J\mathbb{1}, \text{score } r[w] := |r| \cdot w(\cdot)$$

We will see that T is also the smallest full submonad of J which is closed under sample and score.

4.1 Preliminaries

Monads. A *strong monad structure* $\underline{T}$ over a Cartesian closed category $\mathcal{C}$ is a triple $\langle T, \text{return}, \gg \rangle$ consisting of an assignment of an object TX and a morphism $\text{return}_X : X \rightarrow TX$ for every object X , and an assignment of a morphism $\gg_{X,Y} : TX \times (TY)^X \rightarrow TY$. A *strong monad* is a strong monad structure $\underline{T}$ satisfying the monad laws below, expressed in the Cartesian closed internal language:

$$((\text{return } x) \gg f) = f(x) \quad (a \gg \text{return}) = a \quad ((a \gg f) \gg g) = (a \gg \lambda x. f(x) \gg g)$$

Every monad yields an endofunctor T on $\mathcal{C}$ -morphisms: $T(f) := \text{id} \gg^T (\text{return}^T \circ f)$. The *Kleisli* category $\mathcal{C}_{\underline{T}}$ consists of the same objects as $\mathcal{C}$, but morphisms are given by $\mathcal{C}_{\underline{T}}(X, Y) := \mathcal{C}(X, TY)$. A strong monad $\underline{T}$ is *commutative* when, for every pair of objects X, Y :

$$a : TX, b : TY \vdash a \gg \lambda x. b \gg \lambda y. \text{return}(x, y) = b \gg \lambda y. a \gg \lambda x. \text{return}(x, y)$$

Factorisation systems. We use the following concepts to factorise our powerdomain as a submonad of the Schwartz distribution monad. Recall that a *orthogonal factorisation system* on a category $\mathcal{C}$ is a pair $(\mathcal{E}, \mathcal{M})$ consisting of two classes of morphisms of $\mathcal{C}$ such that:

- Both $\mathcal{E}$ and $\mathcal{M}$ are closed under composition, and contain all isomorphisms.
- Every morphism $f : X \rightarrow Y$ in $\mathcal{C}$ factors into $f = m \circ e$ for some $m \in \mathcal{M}$ and $e \in \mathcal{E}$.
- Functoriality: for each situation as on the left, there is a unique $h : A \rightarrow A'$ as on the right:

$$\begin{array}{ccc} X & \xrightarrow{e \in \mathcal{E}} & A \xrightarrow{m \in \mathcal{M}} Y \\ f \downarrow & = & \downarrow g \\ X' & \xrightarrow{e' \in \mathcal{E}} & A' \xrightarrow{m' \in \mathcal{M}} Y' \end{array} \implies \begin{array}{ccc} X & \xrightarrow{e} & A \xrightarrow{m} Y \\ f \downarrow & = & \downarrow h = \downarrow g \\ X' & \xrightarrow{e'} & A' \xrightarrow{m'} Y' \end{array}$$

S-finite measures and kernels. Let X be a measurable space. Define measures and kernels by direct analogy with their definition on $\mathbb{R}^n$. A measure μ is *finite* when $\mu(X) < \infty$, and a kernel $k : X \rightsquigarrow Y$ is *finite* when there is some bound $B \in [0, \infty)$ such that for all $x \in |X|$, $k(x, Y) < B$. We compare measures and kernels pointwise, and both collections are $\omega\mathbf{cpo}$ s. Measures and kernels are closed under countable pointwise sums given as the lubs of the finite partial sums. A measure μ is *s-finite*

when it is a lub of finite measures, equivalently $\mu = \sum_{n \in \mathbb{N}} \mu_n$ for some countable sequence of finite measures, and similarly a kernel k is *s-finite* when it is a lub of finite kernels, equivalently $k = \sum_{n \in \mathbb{N}} k_n$ for some countable sequence of finite kernels [Staton 2017]. The Randomisation Lemma 2.1 generalises to s-finite kernels [Vákár and Ong 2018, Theorem 15]: for every s-finite kernel $k : X \rightsquigarrow Y$, with Y standard Borel, there is a partial measurable function $f : X \times \mathbb{R} \rightarrow Y$ such that, for every $x \in |X|$, $k(x, -) = f(x, -)_* \lambda$. Recall the 2-dimensional Lebesgue measure $\lambda \otimes \lambda$, which assigns to each rectangle $[a_1, b_1] \times [a_2, b_2]$ its area $(b_1 - a_1) \times (b_2 - a_2)$. Applying the Randomisation Lemma to $\lambda \otimes \lambda$ yields the *transfer principle*: there is a measurable $\varphi : \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}$ such that $\varphi_* \lambda = \lambda \otimes \lambda$.

4.2 Randomisable Expectation Operators

A *randomisation* of an ω qbs X is a partial ω Qbs-morphism $\alpha : \mathbb{R} \rightarrow X$, equivalently a total ω Qbs morphism $\mathbb{R} \rightarrow X_\perp$. Thanks to the Cartesian closure, we have an ω qbs of randomisations $RX := (X_\perp)^\mathbb{R}$. A randomisation $\alpha \in RX$ represents an intensional description of a measure on X by pushing forward the Lebesgue measure λ . The undefined part of α shaves some of the measure leaving us with the restriction of λ to $\text{Dom}(\alpha) := \alpha^{-1}[x \in |X_\perp| \mid x \neq \perp]$. By construction, $\text{Dom}(\alpha)$ is a Borel set, and for every weighting function $w : X \rightarrow \mathbb{W}$, the composition $w \circ \alpha : \text{Dom}(\alpha) \rightarrow X \rightarrow \mathbb{W}$ is an ω qbs morphism. Underlying this composition is a qbs morphism $w \circ \alpha : \text{Dom}(\alpha) \rightarrow \mathbb{W}$. Because Qbs is a conservative extension of Sbs, this morphism is a Borel measurable function. Thus, every randomisation induces an expectation operator. Moreover, this assignment is an ω qbs morphism:

$$\mathbb{E} : RX \rightarrow JX; \quad \mathbb{E}_\alpha[w] := \int_{\text{Dom}(\alpha)} \lambda(dr) w(\alpha(r))$$

A *randomisable* expectation operator $\mu \in JX$ is one where $\mu = \mathbb{E}_\alpha$ for some randomisation $\alpha \in RX$. Let $|SX|$ be the set of randomisable operators $\mathbb{E}_{|X_\perp^\mathbb{R}|} \subseteq |JX|$. Similarly, consider the randomisable random operators $M_{SX} := \mathbb{E} \circ [M_{RX}]$. Let $|TX|$ be the ω -chain-lub-closure of $|SX| \subseteq |JX|$, and M_{TX} be the closure of $M_{SX} \subseteq M_{JX}$ under (pointwise) lub of ω -chains. The ω qbs TX is the smallest ω qbs that is a full sub- ω cpo of JX and containing the randomisable random operators.

Example 4.1. When X is a standard Borel space with the discrete order, each randomisable operator defines an s-finite measure, and each randomisable random operator defines an s-finite kernel. By the Randomisation Lemma for s-finite kernels, conversely, every s-finite measure/kernel arises from a randomisable (random) operator. Moreover, each s-finite measure/kernel is a lub of finite measures/kernels. So TX is the smallest sub- ω cpo of JX containing the finite measures as elements and kernels as random elements, and consists of the s-finite measures and kernels.

The randomisation functor $R : \omega\text{Qbs} \rightarrow \omega\text{Qbs}$ has the following monad structure. Using the transfer principle, fix any measurable $\varphi : \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}$ satisfying $\varphi_* \lambda = \lambda \otimes \lambda$, and define:

$$\text{return } x := \left[[0, 1]. \uparrow x, [0, 1]^\complement, \perp \right] \quad (\alpha \gg= f) : \mathbb{R} \xrightarrow{\varphi} \mathbb{R} \times \mathbb{R} \xrightarrow{\mathbb{R} \times \alpha} \mathbb{R} \times X \xrightarrow{\mathbb{R} \times f} \mathbb{R} \times (Y_\perp)^\mathbb{R} \xrightarrow{\text{eval}} Y$$

The unit is a randomisation of the Dirac distribution, shaving from Lebesgue all but a probability distribution concentrated on x . The monadic bind splits the source of randomness, using the transfer principle, into two independent sources of randomness, one for α and one for the kernel f . The expectation morphism $\mathbb{E} : RX \rightarrow JX$ preserves this monad structure. The monad structure R does *not* satisfy the monad laws. While RX depends on the choice of φ , TX is independent of the choice.

4.3 Factorising Monad Structure Morphisms

To show that the monad structure of J restricts to T , we rely on a general theory, recently developed by McDermott and Kammar [2018]. The full monos between ωCpo s form the $\mathcal{M}$ -class of an orthogonal factorisation system on ωCpo , where the $\mathcal{E}$ -class consists of the *dense epis*: Scott-continuous functions $e : P \rightarrow Q$ whose image is dense, i.e., the closure of $e[P]$ is Q . Kammar and Plotkin [2012] and McDermott and Kammar [2018] use this factorisation system to decompose a locally-continuous monad over ωCpo into appropriate sub-monads indexed by the sub-collection of effect operation subsets. We use this construction to carve a sub-monad for sampling and conditioning.

A *full mono* between ωqbs s is a full mono between them as ωCpo s, i.e., an order reflecting ωqbs morphism. A *densely strong epi* $e : X \rightarrow Y$ is an ωqbs morphism that maps the random elements M_X into a Scott dense subset of M_Y w.r.t. the pointwise order.

LEMMA 4.2. *Densely strong epis and full monos form an orthogonal factorisation system on ωQbs . Moreover, the densely strong epis are closed under countable products, exponentiation with standard Borel spaces, and the lifting monad $(-)_\perp$.*

Therefore, the densely strong epis are closed under the randomisation functor R . We can now directly apply McDermott and Kammar's construction [2018] to turn TX into a canonical monad:

THEOREM 4.3. *The unit and bind of J restrict to T . The (densely strong epi, full mono)-factorisation of the expectation operator $\mathbb{E} : RX \rightarrow TX \rightarrow JX$ preserves these monad structures.*

4.4 Sampling and Conditioning

In the [introduction](#) to this section, we defined sampling and conditioning as expectation operators. Both arise as expectation operators for the following randomisations:

$$\begin{aligned} \text{sample} : \mathbb{1} &\rightarrow R\mathbb{R} & \text{score} & : \mathbb{R} \rightarrow R\mathbb{1} \\ \text{sample} &:= [[0, 1].\uparrow, [0, 1]^{\mathbb{G}}.\perp]; & \text{score}(r) &:= [[0, |r|].\uparrow\langle \rangle, [0, |r|]^{\mathbb{G}}.\perp] \end{aligned}$$

Post-composing with $\mathbb{E} : R \rightarrow T$, we have analogous operations for T . Let FX be the *free monad* over ωQbs with Kleisli arrows $\text{sample} : \mathbb{1} \rightarrow F\mathbb{R}$ and $\text{score} : \mathbb{R} \rightarrow F\mathbb{1}$. It exists because, for example, ωQbs is locally presentable (see Appx. B). The monad morphism $m_J : F \rightarrow J$ preserving sample and score given by initiality of F factors through the full inclusion of T in J as $m_J : F \xrightarrow{m_T} T \xrightarrow{\mathbb{E}} J$, where $m_T : F \rightarrow T$ is the unique monad morphism given by initiality. The randomisable operators are *fully definable* by sample and score , lubs , and the monad operations in the following sense:

LEMMA 4.4. *The unique monad morphism from the free monad preserving sampling and conditioning is a component-wise densely strong epi $m_T : F \rightarrow T$.*

PROOF (SKETCH). Define the Lebesgue integral $\mathbb{E}_\uparrow(w) := \int_{\mathbb{R}} \lambda(dr)w(r)$ by rescaling the normal distribution. Let $\text{normal-rng} : [0, 1] \rightarrow \mathbb{R}$ be a randomisation of the gaussian with mean 0 and standard deviation 1, and let $\text{normal-pdf} : \mathbb{R} \rightarrow \mathbb{R}$ be the probability density function of this gaussian. Set $\xi \in F\mathbb{R}$ as on the left, and calculate as on the right:

$$\begin{aligned} \xi &:= \text{sample} \gg \lambda s. \\ \text{let } r &= \text{normal-rng}(s) \text{ in } & (m_T \xi)(w) &= \int_{[0, 1]} \lambda(ds) \frac{w(\text{normal-rng}(s))}{\text{normal-pdf} \circ \text{normal-rng } s} = \int_{\mathbb{R}} \lambda(dr)w(r) \\ \text{score } \frac{1}{\text{normal-pdf } r} & ; \text{return } r \end{aligned}$$

For each non-empty ωqbs X , choose $x_0 \in |X|$, and consider a randomisable random operator $\alpha = \mathbb{E} \circ \Lambda f$, for some uncurried $f : \mathbb{R} \times \mathbb{R} \rightarrow X$. Set $\zeta(s) := \xi \gg \lambda r. \text{handle } f(s, r) \text{ with } x_0$ where the auxiliary function **handle** – **with** $x_0 : X_\perp \rightarrow FX$; is given by: **handle** $\uparrow x$ **with** $x_0 := \text{return } x$ and **handle** $\perp$ **with** $x_0 := \text{score } 0 ; \text{return } x_0$. Then: $m_T \circ \zeta(s)(w) = \int_{\mathbb{R}} \lambda(dr)f(s, r) = \alpha(s, w)$. $\square$

Using the functoriality of the factorisation system, we deduce:

PROPOSITION 4.5. *The monad T is the minimal full submonad of J that contains `sample` and `score`.*

4.5 Synthetic Measure Theory

Synthetic mathematics identifies structure and axioms from which we can recover the main concepts and results of specific mathematical theories, and transport them to new settings. Kock [2012] demonstrates that some measure theory can be reproduced for any suitable monad on a suitable category. Ścibior et al. [2017] impose this categorical structure on $\mathbf{Qbs}$ and use it to verify implementations of Bayesian inference algorithms. Our statistical powerdomain T is a very well-behaved monad. On the full subcategory $\mathbf{Qbs} \subseteq \omega\mathbf{Qbs}$, it restricts to the distribution monad given by Ścibior et al. [2017]. Like there, T makes $\omega\mathbf{Qbs}$ into a model of synthetic measure theory, enabling the interpretation of a statistical calculus. We hope to replicate their proofs for the more expressive SFPC in the future, and only state that $\omega\mathbf{Qbs}$ has this structure.

Recall the central definition to synthetic measure theory. A *measure category* $(C, \underline{T})$ consists of a Cartesian closed category C with countable limits and coproducts; and a commutative monad $\underline{T}$ over C such that the morphisms $\langle \rangle_{T0} : T0 \rightarrow \mathbb{1}$ and $\gg = \left[\langle \delta_{i,j} \rangle_j \right]_i : T \sum_{i \in \mathbb{N}} X_i \rightarrow \prod_{j \in \mathbb{N}} TX_j$ are invertible, where $\delta_{i,i} = \text{return}_{X_i}^T$ and $\delta_{i \neq j} = x \mapsto (\langle \rangle_{T0}^{-1} \langle \rangle) \gg (\text{return}^T \circ [])$. The intuition is that elements in $T \sum_i X_i$, thought of as a measure on a countable coproduct spaces are in one-to-one correspondence with tuples of measures on the component spaces. Surprisingly, this short definition guarantees that $T\mathbb{1}$ is a countably-additive semi-ring whose elements resemble *scalars*, and measures have a countably-additive structure and scalar multiplication operations. We then also have a morphism $\text{total} := T\langle \rangle : TX \rightarrow T\mathbb{1}$, which we think of as assigning to each measure the scalar consisting of its *total measure*.

THEOREM 4.6. *The statistical powerdomain T equips $\omega\mathbf{Qbs}$ with a measure category structure. The countable semiring of scalars is given by weights $T\mathbb{1} \cong \mathbb{W}$ with addition and multiplication. In particular, elements of TX are linear and T is a commutative monad.*

Recall that we obtain a *probabilistic* powerdomain as a full submonad $\underline{P}$ of $\underline{T}$ as the equalizer of $1, \text{total} : TX \rightarrow T\mathbb{1}$, and a *subprobabilistic* powerdomain similarly. We make no use of these additional powerdomains in this work, except to note that there is a continuous normalization function $TX \rightarrow (PX)_\perp + \{\top\}$, where $\top$ is maximal, which returns $\perp$ or $\top$ if the overall measure is 0 or ∞ respectively, and a normalized probability measure otherwise. In probabilistic programming, this normalization operation is usually used at the top level to extract a posterior distribution.

4.6 Valuations on Borel-Scott Open Subsets

We compare our approach to traditional notions of probabilistic powerdomains using the following concepts. Let X be a set. A *lub-lattice of X -subsets* is a family of subsets of X that is closed under finite unions and intersections, and countable unions of ω -chains $B_0 \subseteq B_1 \subseteq \dots$ w.r.t. inclusion.

Example 4.7 (Scott open subsets). Let P be an ωcpo . A subset $U \subseteq |P|$ is *Scott open* when it is

- upwards closed: for all $x \in U$, and $y \in |Q|$, if $x \leq y$ then $y \in U$;
- ω -inaccessible: for every ω -lub $\bigvee_{n \in \mathbb{N}} y_n \in U$, there is some $n \in \mathbb{N}$ for which $y_n \in U$.

The Scott open subsets form a lub-lattice of subsets.

Example 4.8 (Borel subsets). Let X be a qbs. A subset $B \subseteq |X|$ is *Borel* when, for every random element α , the subset $\alpha^{-1}[B] \subseteq \mathbb{R}$ is Borel. The Borel subsets form a lub-lattice of subsets.

Example 4.9 (Borel-Scott open subsets). Let X be a qbs. A subset $D \subseteq |X|$ is *Borel-Scott open* when it is Borel w.r.t. the underlying qbs and a Scott open w.r.t. the underlying ω cpo. The Borel-Scott open subsets form a lub-lattice of subsets.

The measurable functions into a Borel space are characterised as approximated by step functions (see Ex. 3.4). The same is true for ω qbses. Let X be a ω qbs, and B a Borel-Scott open subset of X . The *characteristic function* of B is given by $[- \in B] := [B.1, B^c.0] : X \rightarrow \{0, 1\} \subseteq \mathbb{W}$. A morphism $f : X \rightarrow \mathbb{W}$ is a *step function* when it is a finite weighted sum of characteristic functions.

LEMMA 4.10 (APPROXIMATION BY SIMPLE FUNCTIONS). *Let $f : X \rightarrow \mathbb{W}$ be a ω qbs-morphism. Then, there is an ω -chain $f_n : X \rightarrow \mathbb{W}$ of step functions such that $\bigvee_{n \in \mathbb{N}} f_n = f$.*

PROOF (SKETCH). Define $f_n := \sum_{1 \leq i \leq 4^n} \frac{[- \in f^{-1}[(\frac{i}{2^n}, \infty]]]}{2^n}$. □

Let $O = (|O|, \subseteq)$ be a lub-lattice over a set X ordered w.r.t. inclusion. An O -valuation [Lawson 1982] is a strict Scott-continuous function $v : O \rightarrow \mathbb{W}$ satisfying the binary *inclusion-exclusion principle*: $v(A \cup B) + v(A \cap B) = v(A) + v(B)$ for all $A, B \in |O|$. Every lub-lattice of subsets on X induces a measurable space structure on X by closing under complements and countable unions. Every valuation induces a unique measure on this measurable space. A valuation v over X is *finite* when $v(X) < \infty$, and *s-finite* when it is the pointwise lub of finite valuations. The s-finite valuations form an ω -chain-closed subset of the valuations.

The linearity and Scott continuity of expectation operators in TX now gives us the following.

COROLLARY 4.11. *Expectation operators are determined by their value on characteristic functions.*

This corollary tells us that elements of TX can be thought of as certain (s-finite) valuations on the lub-lattice of Borel-Scott open subsets. If P is an ω cpo, the valuations over the free ω qbs generated by P (see §3.3) coincide with traditional valuations on the Scott opens of P . When X is a qbs with the discrete order, restricting the expectation operators in TX to characteristic functions yields s-finite measures on the underlying measurable space Σ_X , and by Cor. 4.11, each such s-finite measure uniquely determines the operator. Similarly, when X and Y are qbses with the discrete order, each Kleisli arrow $X \rightarrow TY$ determines at most one s-finite kernel $\Sigma_X \leadsto \Sigma_Y$ by restricting to characteristic functions. For an sbs X , TX consists of all s-finite measures on X and M_{TX} of all s-finite kernels $\mathbb{R} \leadsto X$, by the randomisation lemma for s-finite kernels [Vákár and Ong 2018].

5 AXIOMATIC DOMAIN THEORY

Domain theory develops order-theoretic techniques for solving recursive domain equations. The theorems guaranteeing such solutions exist are technically involved. Fiore and Plotkin's *axiomatic domain theory* [1996; 1994] axiomatises categorical structure sufficient for solving such equations. They aggregate axioms of different strengths dealing with the same domain-theoretic aspects of the category at hand; the strength of one axiom can compensate for the weakness of another. This theory allows us to treat recursive domain equations in ω Qbs methodically.

This section is technical, and the main result is that ω Qbs has an expansion to a category **pBS** of ω qbses and *partial maps* between them, supporting the solution of recursive domain equations. The type-formers of SFPC, that will denote locally continuous mixed-variance functors over ω Qbs, then have a locally continuous extension to **pBS**, allowing us to use the solutions of recursive domain equations as denotations of recursive types.

5.1 Axiomatic Structure

We begin by isolating the structure Fiore postulates as it applies to ω Qbs. For the full account, see Fiore's thesis. A (**Pos**)-domain structure $\mathcal{D}$ is a pair $\langle C_{\mathcal{D}}, \mathbb{M}_{\mathcal{D}} \rangle$ consisting of a **Pos**-enriched category

$C_{\mathfrak{D}}$, and a locally small full-on-objects subcategory $\mathfrak{M}_{\mathfrak{D}}$ consisting solely of monomorphisms such that for every $m : D \rightarrowtail Y$ in $\mathfrak{M}_{\mathfrak{D}}$ and $f : X \rightarrow Y$ in $C_{\mathfrak{D}}$:

- the pullback of m along f exists in $C_{\mathfrak{D}}$; and
- in every pullback diagram, the pulled back morphism $f^*m : f^*D \rightarrowtail X$ is in $\mathfrak{M}_{\mathfrak{D}}$.

We call $C_{\mathfrak{D}}$ the category of *total maps* and the monos in $\mathfrak{M}_{\mathfrak{D}}$ *admissible*.

Example 5.1 (Scott open monos). A *Scott open mono* $m : P \rightarrow Q$ between ωcpos is a Scott-continuous function that is a full mono such that $m[P] \subseteq Q$ is Scott open (see Ex. 4.7). Equivalently, m is mono and m -images of Scott open subsets are Scott open. Taking the Scott open monos as admissible monos yields a domain structure **Scott** over ωCpo .

To define partial maps for **Qbs** and ωQbs , we use the following two examples. First, define a *strong mono* $m : X \rightarrowtail Y$ between qbses to be an injective function $m : |X| \rightarrowtail |Y|$ such that $m \circ [M_X] = M_Y \cap [m[X]]^{\mathbb{R}}$. The strong monos in **Qbs** coincide with the regular monos, and so are closed under pullbacks. They are also closed under composition and include all isomorphisms.

Example 5.2 (Borel open monos). Let X, Y be qbses. A *Borel open mono* $m : X \rightarrowtail Y$ is a strong mono whose image $m[X] \subseteq Y$ is Borel (see Ex. 4.8). Equivalently, a strong mono such that m -images of Borel subsets are Borel. The Borel open monos are closed under pullbacks, composition, and contain all isos. By the completeness of **Qbs** we have all pullbacks, so taking the Borel open monos as admissible monos yields a domain structure **Borel** over **Qbs**.

Example 5.3 (Borel-Scott open monos). Taking the ωQbs -morphisms that are both Borel open and Scott open yields **BS**, our domain structure of interest over ωQbs . It is a domain structure as by Lemma 3.9 the pullbacks in ωQbs are computed using the pullbacks of ωCpo and **Qbs**.

To incorporate effects, we extend Fiore's [1996] account in the following special case. When the *representability* axiom ($\dashv$), which we define below, holds, there is a strong monad over C called the *lifting monad* $-\perp$. We define an *effectful* domain structure to be a triple $\langle \mathfrak{D}, T, m \rangle$ consisting of a representable domain structure $\mathfrak{D}$ with finite products; a strong monad T over the category of total maps; and a strong monad morphism $m : -\perp \rightarrow T$, thought of as *encoding* partiality using T 's effects.

Our effectful domain structure consists of ωQbs , together with the Borel-Scott open monos; the probabilistic power-domain of Sec. 4; and, as the resulting lifting monad is the partiality monad of §3.3, the monad morphism m is the function mapping $\perp$ to the zero measure and every other element x to the Dirac measure δ_x .

5.2 Axioms and Derived Structure

We develop the domain theory following Fiore's development and describe the axioms it validates, summarised in Fig. 10. While doing so, we recall the structure Fiore derives from these axioms.

Partial maps. Each domain structure $\mathfrak{D} = (C, \mathfrak{M})$ constructs a *category $\mathbf{pD}$ of partial maps*. Let X, Y be C -objects. A *partial map description* $u : X \rightarrow Y$ from X to Y , is a pair $u = (\partial_u, \bar{u})$ consisting of an admissible mono $\partial_u : D_u \rightarrowtail X$ and a C -morphism $\bar{u} : D_u \rightarrow Y$. Two descriptions u and v are *equivalent*, $u \equiv v$, when there is an isomorphism $i : D_u \xrightarrow{\cong} D_v$ satisfying: $\bar{v} \circ i = \bar{u}$ and $\partial_v \circ i = \partial_u$.

A *partial map* $u : X \rightarrow Y$ is the equivalence class of a description, bearing in mind that $\mathfrak{M}$ is locally small. E.g., for **BS**, choose inclusions as the canonical representatives $\partial_u : D_u \subseteq X$, which uniquely determine the description. Partial maps form a category $\mathbf{pD}$, with identities given by $[\text{id}, \text{id}]$, and the composition $v \circ u : X \xrightarrow{u} Y \xrightarrow{v} Z$ via pullback (see Fig. 11, left). We have an identity-on-objects, faithful functor $J : C \rightarrow \mathbf{pD}$ mapping each total map $f : X \rightarrow Y$ to $[\text{id}, f] : X \rightarrow Y$.

Structure		Axioms	
$C_{\mathfrak{D}}$	total map category	(+) every object has a partial map classifier $\uparrow_X : X \rightarrow X_{\perp}$	$(\rightarrow_{\leq})$ $C_{\mathfrak{D}}$ has locally monotone exponentials
$f \leq g$	Pos -enrichment pointwise order	(fup) every admissible mono is full and upper-closed	(+) locally continuous total coproducts
$\mathfrak{M}_{\mathfrak{D}}$	admissible monos	$(\dashv_{\leq})$ $\lrcorner_{-}$ is locally monotone	(?!) $0 \rightarrow 1$ is admissible
	Borel-Scott opens	$(C_{\vee})$ $C_{\mathfrak{D}}$ is $\omega\mathbf{Cpo}$ -enriched	$(\times_{\vee})$ $C_{\mathfrak{D}}$ has a locally continuous products
T	monad for effects power-domain	(U) ω -colimits behave uniformly (Lemma 5.8)	(CL) $C_{\mathfrak{D}}$ is cocomplete
m	partiality encoding $m : -_{\perp} \rightarrow T, \perp \mapsto 0$	(1) $C_{\mathfrak{D}}$ has a terminal object	$(T_{\vee})$ T is locally continuous
Derived axioms/structure			
$\mathbf{pD}$	partial map category	$(\otimes)$ $\mathbf{pD}$ has partial products	(pCL) $\mathbf{pD}$ is cocomplete
$-_{\perp}$	partiality monad	$(\otimes_{\vee})$ $(\otimes)$ is locally continuous	(p+ $\vee$) $\mathbf{pD}$ has locally continuous partial coproducts
$(\dashv_{\vee})$	the adjunction $J \dashv L$ is locally continuous	$(\rightarrow_{\vee})$ $C_{\mathfrak{D}}$ has locally continuous exponentials	(BC) $J : C \hookrightarrow \mathbf{pD}$ is a bilimit compact expansion
$(\mathbf{p}_{\vee})$	$\mathbf{pD}$ is $\omega\mathbf{Cpo}$ -enriched	$(\Rightarrow_{\vee})$ $\mathbf{pD}$ has locally continuous partial exponentials	
$(1_{\leq})$	$\mathbf{pD}$ has a partial terminal		

Fig. 10. The axiomatic domain theory of $\omega\mathbf{Qbs}$ and its probabilistic power-domain

Fig. 11. Partial identities and composition (left) and characteristic maps (right).

Representability. Given a domain structure, a *classifier of partial maps* is a collection of admissible monos $\langle \uparrow_X : X \rightarrow LX \rangle_{X \in C}$, indexed by the objects of C , such that for every Y , and every partial map $f : X \rightarrow Y$, there is a unique (total) map $\lrcorner f \lrcorner : X \rightarrow LY$ such that Fig. 11 (right) is a pullback square. We call $\lrcorner f \lrcorner$ the *total representation* of f . Each classifier $\uparrow$ of partial maps induces a right adjoint $J \dashv L : \mathbf{pD} \rightarrow C$ with $\uparrow$ as unit. We give two well-known and two novel examples.

Example 5.4 (partiality in Set). The collection $\langle \iota_1 : X \rightarrow X + \{\perp\} \rangle_{X \in \mathbf{Set}}$ classifies the partial maps of $\mathbf{Fn}$, the domain structure over $\mathbf{Set}$ with injections as admissible monos. The representation of a partial function maps the elements outside the domain to $\perp$. The induced adjunction is isomorphic to the Kleisli resolution of the partiality monad $(-) + \{\perp\}$.

Example 5.5 (partiality in $\omega\mathbf{Cpo}$). Let $X_{\perp}$ be the lifting of the $\omega\mathbf{cpo}$ X by adjoining a new bottom element. The collection $\langle \uparrow : X \rightarrow X_{\perp} \rangle_{X \in \omega\mathbf{Cpo}}$ classifies the **Scott**-partial maps with $\lrcorner_{-}$ as in $\mathbf{Fn}$. The induced adjunction is isomorphic to the Kleisli resolution of the lifting monad $(-)_{\perp}$.

Example 5.6 (partiality in $\mathbf{Qbs}$). The collection $\langle \iota_1 : X \rightarrow X + \{\perp\} \rangle_{X \in \mathbf{Qbs}}$ classifies the **Borel**-partial maps with $\lrcorner_{-}$ as in $\mathbf{Fn}$. The induced adjunction is isomorphic to the Kleisli resolution of the partiality monad $(-) + \{\perp\}$.

Example 5.7 (partiality in $\omega\mathbf{Qbs}$). The collection $\langle \uparrow : X \rightarrow X_\perp \rangle_{X \in \mathbf{Qbs}}$ classifies the **BS**-partial maps with $\lrcorner_\perp$ as in **Fn**. The induced adjunction is the Kleisli resolution of the monad $(-)_\perp$ of §3.3.

Enrichment. The order enrichment of $\mathbf{C}_{\mathbf{BS}} = \omega\mathbf{Qbs}$ is an $\omega\mathbf{Cpo}$ -enrichment. Moreover, the partial map category inherits a potential **Pos**-enrichment: for $u, v : X \rightarrow Y$, write $u \sqsubseteq v$ to mean that there is some $i : D_u \rightarrow D_v$ such that $\partial_u = \partial_v \circ i$ and $\bar{u} \leq \bar{v} \circ i$. The isomorphism $\mathbf{pD} \cong \omega\mathbf{Qbs}_\perp$ respects $\sqsubseteq$ and $\leq$, i.e., $u \sqsubseteq v : X \rightarrow Y$ iff $\lrcorner_\perp u \leq \lrcorner_\perp v$ as morphisms $X \rightarrow Y_\perp$ in $\omega\mathbf{Qbs}$, and so $\mathbf{pD}$ is $\omega\mathbf{Cpo}$ -enriched. We can also deduce this fact from more fundamental axioms. First, every admissible mono is full and its image is upper-closed. As a consequence of **Fiore's** Prop. 4.2.4, the order $\sqsubseteq$ is a partial order and **Pos**-enriches $\mathbf{pD}$. Denote the inverse map to representation by $\ulcorner - \urcorner : \mathbf{C}_{\mathbf{BS}}(X, Y_\perp) \rightarrow \mathbf{pBS}(X, Y)$. Because $\ulcorner - \urcorner$ is monotone, the adjunction $J \dashv L : \mathbf{pD}$ is locally monotone [*ibid.*, Prop. 4.5.4], and, as a consequence $\mathbf{pD}$ is $\omega\mathbf{Cpo}$ -enriched [*ibid.*, Prop. 4.5.3].

Uniformity. As $\omega\mathbf{Qbs}$ is cocomplete (Cor. 7.2 in Appx. B), it has local ω -lubs, but these lubs can behave pathologically [*ibid.*, Sec. 4.3.2]. The following *uniformity* axiom avoids such pathologies.

LEMMA 5.8 (UNIFORMITY). *Let $\langle i_{n+1} : D_n \rightarrowtail D_{n+1} \rangle_{n \in \mathbb{N}}$ be an ω -chain of Borel-Scott open monos.*

- *Every colimiting cocone $(D, \langle \mu_n : D_n \rightarrow D \rangle_{n \in \mathbb{N}})$ consists of Borel-Scott opens.*
- *The mediating morphism into any other cocone of Borel-Scott opens is also Borel-Scott open.*

Unit type. As $\omega\mathbf{Qbs}$ has a terminal object $\mathbb{1}$, the partial maps have $\mathbb{1}$ as a *partial* terminal object: every hom-poset $\mathbf{pBS}(A, \mathbb{1})$ has a unique maximal morphism [*ibid.*, p. 84].

Product types. As $\omega\mathbf{Qbs}$ has binary products $X \times Y$, we can extend the binary product functor $(\times) : \mathbf{C}_{\mathbb{D}} \times \mathbf{C}_{\mathbb{D}} \rightarrow \mathbf{C}_{\mathbb{D}}$ to a *partial-product* functor $\otimes : \mathbf{pD} \times \mathbf{pD} \rightarrow \mathbf{pD}$ [*ibid.*, Prop. 5.1.1]. Because $\omega\mathbf{Qbs}$ has exponentials, the functor $- \times X$ preserves colimits. As a consequence, as products are locally continuous, and as axioms (U) , $(\mathbf{p}_V)$ hold, $(\otimes)$ is locally continuous [*ibid.*, Prop. 5.1.3].

Variant types. Because $\omega\mathbf{Qbs}$ is cocomplete, using axioms $(+), (\mathbb{1})$ we deduce that $\mathbf{pD}$ is cocomplete [*ibid.*, Theorem 5.3.14], and colimiting cocones of total diagrams comprise of total maps and are colimiting in $\mathbf{C}_{\mathbb{D}}$ [*ibid.*, Prop. 5.2.4]. As the coproducts of $\omega\mathbf{Qbs}$ are locally continuous, so are the coproducts in $\mathbf{pBS}$ [*ibid.*, Prop. 5.3.13]. As a consequence, the locally continuous finite-coproduct functor $\sum_{i \in I} : \omega\mathbf{Qbs}^I \rightarrow \omega\mathbf{Qbs}$ extends to a locally continuous functor $\coprod_{i \in I} : \mathbf{pD}^I \rightarrow \mathbf{pD}$ on partial maps [*ibid.*, remark following Cor. 5.3.10].

Effectful function types. The following development is new, as **Fiore** only considered the partiality effect. As we now saw, when the representability axioms $(+), (\lrcorner_\perp)$, and the enrichment axioms $(\mathbf{p}_V), (\otimes_V)$ hold, we have a locally continuous strong lifting monad $-\perp$ over $\mathbf{C}$. Recall the additional structure given in an effectful domain structure $\langle \mathbb{D}, T, m \rangle$, namely the locally monotone strong monad T over $\mathbf{C}_{\mathbb{D}}$, and the strong monad morphism $m : -_\perp \rightarrow T$. We further assume that T is locally continuous, which we call axiom (T_V) . Thm. 4.3 validates it for the statistical powerdomain.

We have an identity-on-objects locally continuous functor $\lrcorner_{\perp T} : \mathbf{pD} \rightarrow \mathbf{C}_T$, from the category of partial maps into the Kleisli category for T , given for every partial map $u : X \rightarrow Y$ by:

$$\lrcorner_{\perp T} u : X \xrightarrow{\lrcorner_\perp u} Y_\perp \xrightarrow{m} TY$$

When axiom $(\rightarrow_\leq)$ holds, the composite functor $(\odot X) : J - \otimes X_{\perp T} : \mathbf{C} \rightarrow \mathbf{C}_T$ has a right adjoint:

$$(X \Rightarrow_T -) : \mathbf{C}_T \rightarrow \mathbf{C} \quad X \Rightarrow_T Y := (TY)^X \quad X \Rightarrow_T v := \lambda X. \left((TY)^X \times X \xrightarrow{\text{eval}} TY_1 \xrightarrow{\gg v} TY_2 \right)$$

for every $v : Y_1 \rightarrow TY_2$. Axiom $(\rightarrow_\leq)$ implies the exponential adjunction is locally continuous, and as a consequence, $(\Rightarrow_T)$ is locally continuous. We use it to extend Kleisli exponentiation

$(- \Rightarrow T-) : C^{\text{op}} \times C \rightarrow C$ to a locally continuous functor $(\multimap_T) : \mathbf{pD}^{\text{op}} \times \mathbf{pD} \rightarrow \mathbf{pD}$ by setting, for every $u : X_2 \rightarrow X_1$ and $v : Y_1 \rightarrow Y_2$:

$$u \multimap_T v := \lambda \lambda X. \left((TY_1)^{X_1} \times X_2 \xrightarrow{\text{id} \otimes u, T} T((TY_1)^{X_1} \times X_1) \xrightarrow{\gg\text{-eval}} TY_1 \xrightarrow{\gg\text{-}\iota, v, T} TY_2 \right)$$

Recursive types. To solve recursive domain equations, we synthesise axiomatic domain theory with Levy's more modern account [2004]. Recall that an embedding-projection-pair (ep-pair) $u : A \rightleftharpoons B$ in an $\omega\mathbf{Cpo}$ -enriched category $\mathcal{B}$ is a pair consisting of a $\mathcal{B}$ -morphism $u^e : A \rightarrow B$, the *embedding*, and a $\mathcal{B}$ -morphism $u^p : B \rightarrow A$, the *projection*, such that $e \circ p \leq \text{id}$ and $p \circ e = \text{id}$. An *embedding* $u : A \hookrightarrow B$ is the embedding part of some ep-pair $A \rightleftharpoons B$. Every embedding $u : A \hookrightarrow B$ in a partial map category with axiom $(\mathbf{p}_V)$ is a total map $u : A \rightarrow B$ [Fiore 1996, Prop. 5.4.2].

An ω -chain of ep-pairs $(\langle A_n \rangle_{n \in \mathbb{N}}, \langle a_n \rangle_{n \in \mathbb{N}})$ in $\mathcal{B}$ consists of a countable sequence of objects A_n and a countable sequence of ep-pairs $a_n : A_n \rightleftharpoons A_{n+1}$. A *bilimit* (D, d) of such an ω -chain consists of an object D and a countable sequence of ep-pairs $d_n : A_n \rightleftharpoons D$ such that, for all $n \in \mathbb{N}$, $d_{n+1} \circ a_n = d_n$, and $\bigvee_{n \in \mathbb{N}} d_n^e \circ d_n^p = \text{id}_D$. The celebrated *limit-colimit coincidence* [Smyth and Plotkin 1982] states that the bilimit structure is equivalent to a colimit structure (D, d^e) for $(\langle A_n \rangle, \langle a_n^e \rangle)$, in which case d_n^p are uniquely determined, and similarly equivalent to a limit structure (D, d^p) for $(\langle A_n \rangle, \langle a_n^p \rangle)$, in which case d_n^e are uniquely determined. As we saw, the partial map category $\mathbf{pBS}$ has all colimits (derived axiom $(\mathbf{pCL})$), and so $\mathbf{pBS}$ has bilimits of ω -chains of ep-pairs.

A *zero object* is an object that is both initial and terminal. An *ep-zero object* in an $\omega\mathbf{Cpo}$ -category is a zero object such that every morphism into it is an embedding and every morphism out of it is a projection. We say that axiom $(?)$ holds in a domain structure $\mathcal{D}$ in which C has both an initial object $\mathbb{0}$ and terminal object $\mathbb{1}$, when the unique morphism $\mathbb{0} \rightarrow \mathbb{1}$ is an admissible mono. The Borel-Scott domain structure satisfies $(?)$. When axioms $(+), (\mathbb{1}), (\times_V), (\rightarrow_{\leq}), (CL), (?), (+),$ and $(+_{\leq})$ hold, $\mathbf{pD}$ has $\mathbb{0}$, the initial object of C , as an ep-zero object. So in $\mathbf{pBS}$ the empty $\omega\mathbf{qbs}$ is an ep-zero.

A *bilimit compact category* is an $\omega\mathbf{Cpo}$ -category $\mathcal{B}$ with an ep-zero and ep-pair ω -chain bilimits. So $\mathbf{pBS}$ is bilimit compact. When $\mathcal{A}, \mathcal{B}$ are bilimit compact, every locally continuous, mixed-variance functor $F : \mathcal{A}^{\text{op}} \times \mathcal{B}^{\text{op}} \times \mathcal{A} \times \mathcal{B} \rightarrow \mathcal{B}$ has a parameterised solution to the recursive equation $\text{roll} : F(A, X, A', X) \xrightarrow{\cong} X$, for every A, A' in $\mathcal{A}$, qua the bilimit $(\mu B.F(A, B, A', B), d_{F, A, A'})$ of

$$\mathbb{0} \rightleftharpoons F(A, \mathbb{0}, A', \mathbb{0}) \xleftarrow{F^{\text{op}}(A, \multimap, A', \multimap)} F(A, F(A, \mathbb{0}, A', \mathbb{0}), A', F(A, \mathbb{0}, A', \mathbb{0})) \rightleftharpoons \dots \rightleftharpoons F^n(A, A') \rightleftharpoons \dots$$

The solution is minimal in the sense of Pitts [1996], and we denote the inverse to roll by unroll. The assignments $\mu B.F(A, B, A', B)$ extend to a mixed-variance functor $\mu B.F(-, B, -, B) : \mathcal{A}^{\text{op}} \times \mathcal{A} \rightarrow \mathcal{B}$ by $\mu B.F(f, B, g, B) := \bigvee_n d_n^e \circ F^n(f, g) \circ d_n^p$.

Finally, a *bilimit compact expansion* $J : C \hookrightarrow \mathcal{B}$ is a triple consisting of an $\omega\mathbf{Cpo}$ -category C , a bilimit compact category $\mathcal{B}$; and an identity-on-objects, locally continuous, order reflecting functor $J : C \rightarrow \mathcal{B}$ such that, for every ep-pair ω -chains $(A, a), (B, b)$ in $\mathcal{B}$, their bilimits $(D, d), (E, e)$, and countable collection of C -morphisms $\langle \alpha_n : A_n \rightarrow B_n \rangle_{n \in \mathbb{N}}$ such that for all n :

$$J\alpha_{n+1} \circ a_n^e = b_n^e \circ J\alpha_n, \quad J\alpha_n \circ a_n^p = b_n^p \circ J\alpha_{n+1}$$

(i.e., $J\alpha : (A, a^e) \rightarrow (B, b^e)$ and $J\alpha : (A, a^p) \rightarrow (B, b^p)$ are natural transformations), there is a C -morphism $f : A \rightarrow B$ such that $Jf = \bigvee_n e_n^e \circ \alpha_n \circ d_n^p$. The motivation for this definition: given two bilimit compact expansions $I : \mathcal{D} \hookrightarrow \mathcal{A}, J : C \hookrightarrow \mathcal{B}$, and two locally continuous functors

$$F : \mathcal{D}^{\text{op}} \times C^{\text{op}} \times \mathcal{D} \times C \rightarrow C \quad G : \mathcal{A}^{\text{op}} \times \mathcal{B}^{\text{op}} \times \mathcal{A} \times \mathcal{B} \rightarrow \mathcal{B} \quad \text{s.t.} \quad G \circ I^{\text{op}} \times J^{\text{op}} \times I \times J = J \circ F$$

the functor $\mu B.G(-, B, -, B) : \mathcal{A}^{\text{op}} \times \mathcal{A} \rightarrow \mathcal{B}$ restricts to $\mu B.G(J^{\text{op}}-, B, J-, B) : \mathcal{D}^{\text{op}} \times \mathcal{D} \rightarrow \mathcal{C}$. Bilimit compact expansions are closed under products and opposites [Levy 2004].

Returning to axiomatic domain theory, when axioms $(\dashv), (\dashv_{\leq}), (\mathbb{1}), (\times_{\vee}), (\rightarrow_{\leq}), (CL)$, and $(?)$ hold, the embedding $J : \mathcal{C} \hookrightarrow \mathbf{pD}$ is a bilimit compact expansion of total maps to partial maps. To show the third condition in the definition of expansion, note that embeddings in $\mathbf{pD}$ are total, and we can use the limit-colimit coincidence to reflect the two bilimits along J to the colimits in $\mathcal{C}$, and find a mediating morphism $f : D \rightarrow E$. This morphism maps to the mediating morphism in $\mathbf{pD}$, which is precisely $\bigvee_n e_n^e \circ \alpha_n \circ d_n^p$.

5.3 Summary: Semantics of Types

Using axiomatic domain theory, we constructed an expansion $J : \omega\mathbf{Qbs} \hookrightarrow \mathbf{pBS}$ to a suitable category of partial maps. We use it to interpret SFPC's type-system. We define two interpretations for type-variable contexts Δ as $\omega\mathbf{Cpo}$ -categories, and two locally continuous interpretations of well-kinded types $\Delta \vdash_k \tau : \text{type}$: a total map interpretation $\llbracket - \rrbracket$ and its extension $\{\!\{- \}\!\}$ to partial maps, i.e. $J \circ \llbracket - \rrbracket = \{\!\{- \}\!\} \circ J^{\text{op}} \times J$:

$$\begin{aligned} \llbracket \Delta \rrbracket &:= \prod_{\alpha \in \Delta} \omega\mathbf{Qbs} & \{\!\{\Delta\}\!\} &:= \prod_{\alpha \in \Delta} \mathbf{pBS} & \llbracket \tau \rrbracket : \llbracket \Delta \rrbracket^{\text{op}} \times \llbracket \Delta \rrbracket &\rightarrow \omega\mathbf{Qbs} & \{\!\{\tau\}\!\} : \{\!\{\Delta\}\!\}^{\text{op}} \times \{\!\{\Delta\}\!\} &\rightarrow \mathbf{pBS} \\ \llbracket \alpha \rrbracket &:= \pi_{\alpha} & \{\!\{\alpha\}\!\} &:= \pi_{\alpha} & \llbracket \tau * \sigma \rrbracket &:= \llbracket \tau \rrbracket \times \llbracket \sigma \rrbracket & \{\!\{\tau * \sigma\}\!\} &:= \{\!\{\tau\}\!\} \otimes \{\!\{\sigma\}\!\} \\ \llbracket \underline{\mathbb{R}} \rrbracket &:= \mathbb{R} & \{\!\{\underline{\mathbb{R}}\}\!\} &:= \mathbb{R} & \llbracket \tau \rightarrow \sigma \rrbracket &:= \llbracket \tau \rrbracket \Rightarrow T \llbracket \sigma \rrbracket & \{\!\{\tau \rightarrow \sigma\}\!\} &:= \{\!\{\tau\}\!\} \multimap_T \{\!\{\sigma\}\!\} \\ \llbracket \mathbb{1} \rrbracket &:= 1 & \{\!\{\mathbb{1}\}\!\} &:= 1 \\ \left\llbracket \begin{array}{c} \ell_1 \tau_1 \\ \vdots \\ \ell_n \tau_n \end{array} \right\rrbracket &:= \sum_{i=1}^n \llbracket \tau_i \rrbracket & \left\{ \begin{array}{c} \ell_1 \tau_1 \\ \vdots \\ \ell_n \tau_n \end{array} \right\} &:= \prod_{i=1}^n \{\!\{\tau_i\}\!\} & \llbracket \mu\alpha. \tau \rrbracket &:= \mu\alpha. & \{\!\{\mu\alpha. \tau\}\!\} &:= \mu\alpha. \\ & & & & \{\!\{\tau\}\!\}(J-, \alpha, J-, \alpha) & & \{\!\{\tau\}\!\}(-, \alpha, -, \alpha) \end{aligned}$$

6 SEMANTICS FOR SFPC

Operational semantics for statistical probabilistic calculi such as SFPC require a mathematically technical development. The difficulty comes from dealing with continuous distributions over programs that manipulate continuous data. We follow the recipe previously employed by Staton et al. [2016] (others have used broadly similar methods, e.g. Borgström et al. [2016]). Each program phrase involving the real numbers $r_1, \dots, r_n$ can be represented as an open term t with free variables $x_1, \dots, x_n$ each used linearly and in sequence, together with the tuple $\langle r_1, \dots, r_n \rangle \in \mathbb{R}^n$. This separation lets us equip the abstract syntax with a measurable space structure, and define a big-step operational semantics as kernels between these syntactic spaces [Staton et al. 2016].

6.1 Measure Theoretic Preliminaries

Constructing spaces. Every set is the carrier of a *discrete* measurable space, in which all subsets are measurable. We will typically only consider discrete spaces of countable carriers, as these are precisely the discrete standard Borel spaces. Let I be a set indexing some measurable spaces X_i . The measurable subsets of the disjoint union $\sum_{i \in I} X_i$ are the disjoint unions $\sum_{i \in I} A_i$ of measurable subsets A_i in X_i . When I is countable, the measurable subsets of the cartesian product $\prod_{i \in I} X_i$ are given by taking the boxes $\prod_{i \in I} A_i$, with each A_i measurable in X_i , and closing under countable unions and complements. These three constructions are the free measurable space over a set, categorical coproduct, and categorical countable product in the category of measurable spaces.

Constructing kernels. The *Dirac* kernel $\delta_- : X \rightsquigarrow X$ assigns to each $x \in |X|$ its *Dirac distribution* defined by $\delta_x(A) = 1$ for $x \in A$ and $\delta_x(A) = 0$ otherwise. Every measurable function $f : X \rightarrow Y$

induces a kernel $f : X \rightsquigarrow Y$ by setting $f(x, A) := \delta_{f(x)}(A)$. Each kernel $k : X \rightsquigarrow Y$ acts on measures μ over X yielding the unique measure $\mu \gg k$ over Y satisfying, for all measurable $\varphi : Y \rightarrow [0, \infty]$:

$$\int_Y (\mu \gg k)(dy) \varphi(y) = \int_X \mu(dx) \int_Y k(x, dy) \varphi(y)$$

Moreover, given $k : X \rightsquigarrow Y$ and $l : Y \rightsquigarrow Z$, the mapping $x \mapsto k(x) \gg l$ is a kernel $k \gg l : X \rightsquigarrow Z$. The s-finite kernels are closed under this composition [Kallenberg 2017; Staton 2017].

6.2 Syntax Spaces

Consider the set of terms of type τ in variable context Γ , $\text{Trm}^{\Gamma \vdash \tau} := \{t \mid \Gamma \vdash t : \tau\}$. The set of *values* of type τ in context Γ consists of the subset $\text{Val}^{\Gamma \vdash \tau} \subseteq \text{Trm}^{\Gamma \vdash \tau}$ given inductively by:

$$v, w, u ::= x \mid \underline{r} \mid () \mid (v, w) \mid \tau.\ell v \mid \tau.\text{roll}(v) \mid \lambda x : \tau. t \quad \text{values.}$$

We want to equip $\text{Trm}^{\Gamma \vdash \tau}$ and $\text{Val}^{\Gamma \vdash \tau}$ with a measurable space structure. Let Δ range over *non-symmetric linear* variable contexts of type $\mathbb{R}$, namely finite sequences of variables, and let *mixed* variable contexts $\Gamma; \Delta$ be pairs of variable context Γ and linear contexts Δ with disjoint sets of variables. Each such mixed context can be turned into an ordinary variable context $\Gamma, \Delta : \mathbb{R}$ by treating the identifiers in Δ as variables of type $\mathbb{R}$. We will use the identifiers $\square, \square_1, \square_2$ for the variables in Δ . A term *template* is a term that does not involve nullary primitives $\underline{r}$. We refine the SFPC type-system to judgements of the form $\Gamma; \Delta \vdash t : \tau$ by treating Δ non-symmetric linearly, e.g.:

$$\frac{}{\Gamma; \square \vdash \square : \mathbb{R}} \quad \frac{\Gamma; \Delta_1 \vdash t : \sigma \rightarrow \tau \quad \Gamma; \Delta_2 \vdash s : \sigma}{\Gamma; \Delta_1, \Delta_2 \vdash ts : \tau}$$

Let $\text{TTrm}^{\Gamma; \Delta \vdash \tau} \subseteq \text{Trm}^{\Gamma, \Delta \vdash \tau}$ be the set of well-formed templates $\Gamma; \Delta \vdash t : \tau$ and $\text{TVal}^{\Gamma; \Delta \vdash \tau} \subseteq \text{Val}^{\Gamma, \Delta \vdash \tau}$ be the set of well-formed template values. Given any $\Delta = \square_1, \dots, \square_k$, and $r_1, \dots, r_k \in \mathbb{R}$, we define the following function

$$-[r_1, \dots, r_n] : \text{TTrm}^{\Gamma; \Delta \vdash \tau} \rightarrow \text{Trm}^{\Gamma \vdash \tau} \quad t[r_1, \dots, r_n] := t[\square_1 \mapsto \underline{r}_1, \dots, \square_k \mapsto \underline{r}_k]$$

Combining these functions we have a bijection $\sum_{n \in \mathbb{N}, t \in \text{TTrm}^{\Gamma; \square_1, \dots, \square_{n+\tau}} \mathbb{R}^n \cong \text{Trm}^{\Gamma \vdash \tau}$.

Example 6.1. The Bayesian linear regression program from Fig. 1(a) is represented by:

$$\left(\begin{array}{ll} \text{let } a = \text{normal-rng}(\square_1, \square_2) \text{ in} & (0, 2, \\ 11, \text{score}(\text{normal-pdf}(\square_3 \mid a * \square_4, \square_5)); & 1.1, 1, 0.25, \\ \text{score}(\text{normal-pdf}(\square_6 \mid a * \square_7, \square_8)); & 1.9, 2, 0.25, \\ \text{score}(\text{normal-pdf}(\square_9 \mid a * \square_{10}, \square_{11})); a & 2.7, 3, 0.25) \end{array} \right)$$

The representation $\sum_{n \in \mathbb{N}, t \in \text{TTrm}^{\Gamma; \square_1, \dots, \square_{n+\tau}} \mathbb{R}^n$ has a canonical measurable space structure, combining the Borel σ -algebra on $\mathbb{R}^n$ and the coproduct σ -algebra. As a consequence, we equip $\text{Trm}^{\Gamma \vdash \tau}$ with the measurable space structure making the representation a measurable isomorphism. Similarly, we equip $\text{Val}^{\Gamma \vdash \tau}$ with a measurable space structure. Summing over all types, we get the measurable spaces of closed terms, Trm , and closed values, Val .

6.3 Operational Semantics

We define a structural operational semantics as a kernel $\Downarrow : \text{Trm} \rightsquigarrow \text{Val}$. As usual, the operational semantics is *not* compositional. E.g., the semantics of function application is not defined in terms of that of a subterm, but of a substituted subterm. As a consequence, we construct an increasing sequence of kernels $\Downarrow_n : \text{Trm} \rightsquigarrow \text{Val}$ indexed by the natural numbers and set $\Downarrow := \bigvee_n \Downarrow_n$. At level 0, the evaluation kernel is the least kernel, namely $t \Downarrow_0 A := 0$ for every closed term t and measurable

$$(\text{sample } \Downarrow_n) := \mathbb{U}_{[0,1]} \gg \lambda r. \delta_r \quad (\text{score}(t) \Downarrow_n) := (t \Downarrow_n) \gg \lambda r. |r| \cdot \delta_0$$

Fig. 12. big-step operational semantics for SFPC's statistical primitives ($n > 0$)

$A \subseteq \text{Val}$. For positive levels $n > 0$, we use the following three notational conventions that highlight that the semantics is a standard adaptation of the more familiar operational semantics. First:

$$\frac{k_1(t) w_1 \quad k_2(t, w_1) w_2 \quad \dots \quad k_n(t, w_1, \dots, w_n) v}{l(t) f(t, w_1, \dots, w_n, v)} \quad \text{means} \quad \begin{array}{ll} l(t) := k_1(t) & \gg \lambda w_1. \\ k_2(t, w_1) & \gg \lambda w_2. \dots \\ k_n(t, w_1, \dots, w_{n-1}) & \gg \lambda v. \end{array}$$

$$\text{Second: } l(t) := k_1(t) \quad l(t) := k_2(t) \text{ means } l(t) := k_1(t) + k_2(t) \quad \delta_{f(t, w_1, \dots, w_n, v)}$$

i.e., we sum overlapping definitions. Finally, if one of the intermediate kernels k_i is undefined (on a measurable subset), we define $l(t)$ to have measure 0 on this subset.

Fig. 12 shows the crucial definitions involved in the indexed evaluation kernels are sampling and conditioning, which evaluate using the interpreter's corresponding primitives (see Appx. A for the full semantics). The kernel $\Downarrow$ is s-finite, as the s-finite kernels are closed under composition and lubs [Staton 2017].

6.4 Contextual Equivalence

The operational semantics lets us compare the meaning of terms using the following standard notions. Given variable context Γ_1, Γ_2 , we say that Γ_2 *extends* Γ_1 , and write $\Gamma_2 \geq \Gamma_1$ when, for all $(x : \tau) \in \Gamma_1$, we have $(x : \tau) \in \Gamma_2$. *Program contexts of type σ with a hole – of type $\Gamma \vdash \tau$* are terms $C[\Gamma \vdash - : \tau]$ of type σ with a single variable of type τ , where this variable – always occurs inside the term in contexts $\Gamma' \geq \Gamma$. Write $C[t]$ for the *capturing* substitution $C[\Gamma \vdash - : \tau][- \mapsto t]$.

Two terms $t, s \in \text{Trm}^{\Gamma \vdash \tau}$ are in the *contextual preorder* $t \lesssim s$ when for all program contexts $C[\Gamma \vdash - : \tau]$ of type $\mathbb{R}$, we have that $C[t] \Downarrow \leq C[s] \Downarrow$ in the usual (pointwise) order of measures on $\text{Val}^{\mathbb{R}}$. We say that t and s are *contextually equivalent*, writing $t \approx s$, when $t \lesssim s$ and $s \lesssim t$.

The observational preorder and equivalence are the same even if we vary the definitions:

- (1) using contexts of type 1 and observing only the weight of convergence;
- (2) using contexts of ground type (i.e. iterated sums and products of primitive types) and observing the distribution over Val^{τ} ;
- (3) using contexts of arbitrary type τ and observing the induced distribution over $\text{Val}^{\tau} / \sim$, where $\sim$ is the smallest congruence identifying all λ -abstractions (as done in [Pitts 1996]).

Indeed, for the equivalence of 1., 2. and 3., use characteristic functions $\mathbb{R}^n \rightarrow \mathbb{R}$. For the equivalence with our notion, observe that 1 embeds into $\mathbb{R}$ and, conversely, distinguish any distribution on $\text{Val}^{\mathbb{R}}$ with contexts of type 1 by using **match** $[- \in A](\neg_{\mathbb{R}})$ **with** $\{0 \rightarrow \perp \mid _ \rightarrow ()\}$, for Borel $A \subseteq \mathbb{R}$.

6.5 Idealised Church and SPCF

Since both Idealised Church (§2.3) and SPCF (§2.4) are fragments of SFPC, they inherit an operational semantics and a notion of contextual equivalence w.r.t. contexts in the fragment. For Idealised Church, these contexts translate into certain SFPC program contexts of type Λ -Term $\mathbb{R}$. The induced notion of contextual preorder is now that for two Idealised Church terms we set $t \lesssim s$ when $C[t]^{\dagger} \Downarrow \leq C[s]^{\dagger} \Downarrow$ as distributions on $\text{Val}^{\Lambda\text{-Term } \mathbb{R}} / \sim$ for all Idealised Church program contexts $C[-]$, where $\sim$ identifies all values of the form $\text{Fun } (\lambda x. t)$. For CBV SPCF program contexts of type $\mathbb{R}$ induce an analogous notion of contextual preorder and equivalence.

$$\begin{array}{l}
r \leq_v^{\mathbb{R}} r \Leftrightarrow \top \quad \langle \rangle \leq_v^{\perp} () \Leftrightarrow \top \quad \langle a_1, a_2 \rangle \leq_v^{\tau_1 * \tau_2} (v_1, v_2) \Leftrightarrow a_1 \leq_v^{\tau_1} v_1 \wedge a_2 \leq_v^{\tau_2} v_2 \\
\iota_i(a) \leq_v^{\tau} \tau. \ell_j v \Leftrightarrow i = j \wedge a \leq_v^{\tau_i} v \quad \text{roll } a \leq_v^{\tau} \tau. \text{roll } v \Leftrightarrow \Lambda(k) \leq_v^{\tau \rightarrow \sigma} \lambda x : \tau. t \Leftrightarrow \\
(\tau = \{\ell_1 \tau_1 \mid \dots \mid \ell_n \tau_n\}) \quad a \leq_v^{\sigma[\alpha \mapsto \tau]} v \quad (\tau = \mu \alpha. \sigma) \quad k \leq \llbracket t \rrbracket \wedge k \leq_c^{x:\tau \vdash \sigma} t \\
\mu \leq_v^{\tau} t \Leftrightarrow b \leq_v^{\Gamma \vdash \tau} w \Leftrightarrow \forall \langle a_x \leq_v^{\Gamma \vdash \sigma} v_x \rangle_{(x:\sigma) \in \Gamma} . b(a_x)_{(x:\sigma) \in \Gamma} \leq_v^{\tau} w[x \mapsto v_x]_{(x:\sigma) \in \Gamma} \\
\mu \leq \llbracket t \Downarrow \rrbracket_T^v \quad k \leq_c^{\Gamma \vdash \tau} t \Leftrightarrow \forall \langle a_x \leq_v^{\Gamma \vdash \sigma} v_x \rangle_{(x:\sigma) \in \Gamma} . k(a_x)_{(x:\sigma) \in \Gamma} \leq_c^{\tau} t[x \mapsto v_x]_{(x:\sigma) \in \Gamma}
\end{array}$$

Fig. 14. The properties of the relational interpretation (logical relation) $\leq_v$ and $\leq_c$.

6.8 Adequacy

The key to relating the denotational and operational semantics is to establish that $\llbracket t \rrbracket = \llbracket t \Downarrow \rrbracket_T^v$ as expectation operators in $T[\tau]$. As usual, we prove one inequality by induction on the syntax.

LEMMA 6.6 (COMPUTATIONAL SOUNDNESS). *For every closed term $\vdash t : \tau$, we have $\llbracket t \rrbracket \geq \llbracket t \Downarrow \rrbracket_T^v$.*

For the other direction, we construct a logical relation, using a variation of Pitts' minimal invariant relations method in bilimit compact categories [Levy 2004; Pitts 1996].

LEMMA 6.7 (RELATIONAL INTERPRETATION). *For every context Γ and type τ of SFPC, there are relations $\leq_v^{\Gamma \vdash \tau} \subseteq \omega\mathbf{Qbs}(\llbracket \Gamma \rrbracket, \llbracket \tau \rrbracket) \times \text{Val}^{\Gamma \vdash \tau}$ and $\leq_c^{\Gamma \vdash \tau} \subseteq \omega\mathbf{Qbs}(\llbracket \Gamma \rrbracket, T[\tau]) \times \text{Trm}^{\Gamma \vdash \tau}$ satisfying Fig. 14.*

We have that $a \leq_v^{\Gamma \vdash \tau} v$ implies that $a \leq \llbracket v \rrbracket^v$. Therefore, using the Substitution Lemma 6.3 and the definitions of $\Downarrow$, $\llbracket - \rrbracket^v$ and $\llbracket - \rrbracket$, we establish the following fundamental lemma, by a lengthy mutual induction on the structure of w and s .

LEMMA 6.8 (FUNDAMENTAL). *For every $w \in \text{Val}^{\Gamma \vdash \tau}$ and $s \in \text{Trm}^{\Gamma \vdash \tau}$ $\llbracket w \rrbracket^v \leq_v^{\Gamma \vdash \tau} w$ and $\llbracket s \rrbracket \leq_c^{\Gamma \vdash \tau} s$.*

In particular, for closed terms $\vdash t : \tau$ we have $\llbracket t \rrbracket \leq \llbracket t \Downarrow \rrbracket_T^v$. Adequacy now follows:

THEOREM 6.9 (ADEQUACY). *For all types τ of SFPC and all closed terms $t, s \in \text{Trm}^{\vdash \tau}$, we have that $\llbracket t \rrbracket \leq \llbracket s \rrbracket$ implies that $t \lesssim s$. In particular, $\llbracket t \rrbracket = \llbracket s \rrbracket$ implies that $t \approx s$.*

PROOF. Assume $\llbracket t \rrbracket \leq \llbracket s \rrbracket$ and consider any any context $C[\Gamma \vdash - : \tau]$ of type $\mathbb{R}$. By the Compositionality Thm. 6.4, we have: $\llbracket C[t] \rrbracket \leq \llbracket C[s] \rrbracket$. Therefore: $\llbracket C[t] \Downarrow \rrbracket_T^v = \llbracket C[t] \rrbracket \leq \llbracket C[s] \rrbracket = \llbracket C[s] \Downarrow \rrbracket_T^v$. By Lemma 6.5, we deduce $C[t] \Downarrow \leq C[s] \Downarrow$. So $t \lesssim s$. $\square$

By Lemma 6.2, it now follows that the induced denotational semantics of Idealised Church at the $\omega\mathbf{qbs} \llbracket \Lambda\text{-Term } \mathbb{R} \rrbracket$ and the induced denotational semantics of CBV PPCF are adequate.

COROLLARY 6.10. *For Idealised Church or CBV PPCF terms t, s , $\llbracket t \rrbracket \leq \llbracket s \rrbracket$ implies that $t \lesssim s$.*

We conclude with two applications of the Adequacy Thm. 6.9. First, evaluation order in SFPC and its sub-fragments does not matter, as a consequence of the monad commutativity (Thm. 4.6):

COROLLARY 6.11. *For every $\Gamma \vdash t : \tau$, $\Gamma \vdash s : \sigma$, and $\Gamma, x : \tau, y : \sigma \vdash r : \rho$ we have:*

$$\begin{array}{lcl}
\text{let } (x : \tau) = t \text{ in} & \approx & \text{let } (y : \sigma) = s \text{ in} \\
\text{let } (y : \sigma) = s \text{ in } r & & \text{let } (x : \tau) = t \text{ in } r
\end{array}$$

Second, we show that our definable term-level recursion operator is indeed a fixed-point operator:

COROLLARY 6.12. *Let $\tau = \tau_1 \rightarrow \tau_2$ be a function type. For every $\Gamma, x : \tau \vdash t : \tau$ we have $\mu x : \tau. t \approx t[x \mapsto \mu x : \tau. t]$. Therefore, the following derivation rule is admissible:*

$$\frac{t[x \mapsto \mu x : \tau. t] \Downarrow v}{\mu x : \tau. t \Downarrow v}$$

7 RELATED WORK AND CONCLUDING REMARKS

There is an extensive body of literature on probabilistic power-domain constructions. We highlight the first denotational models of higher-order probabilistic languages with recursion [Saheb-Djahromi 1980] and the work of Jones and Plotkin [1989] who give an adequate model of FPC with *discrete* probabilistic choice based on a category of pre-domains and a probabilistic power-domain construction using valuations. Crucially, in this approach, commutativity of the probabilistic power-domain may fail unless one restricts to certain subcategories of continuous domains which are known to not be Cartesian closed. Jung and Tix [1998] survey the challenges in the search for such a sub-category of continuous domains that admits function spaces and probabilistic power-domains. One ingredient of this problem is that the measurable space structure used is generated from the Scott topology of the domain at hand. Subsequently Barker [2016]; Goubault-Larrecq and Varacca [2011]; Mislove [2016]; and Bacci et al. [2018] proposed variations on this in which random variables play a crucial role, as they do in qbses. We extend their reach by establishing commutativity.

Ehrhard et al. [2017] gave an adequate semantics for a *call-by-name* PCF with continuous probabilistic choice. The semantics is based on a variation of *stable* domain theory, using cones and stable functions, equipped with a notion of random elements analogous to a qbs-structure which they call *measurability tests*. We extend their reach by interpreting soft constraints, recursive types, and commutativity. A large technical difference with our work arises from their choice to work with cones rather than ω cpos, and with stable functions, rather than mere continuous functions, and from our choice to use monadic semantics.

Borgström et al. [2016] conducted a thorough operational treatment of an untyped probabilistic calculus, à la Idealised Church. Our work complements this analysis with a denotational counterpart.

Topological domain theory (TDT) [Battenfeld et al. 2007] accommodates most of the features that we have considered in this paper. In particular, it provides a cartesian closed category with an algebraically compact expansion and a commutative probabilistic powerdomain. Indeed, Huang et al. [2018] have already proposed it as a basis for statistical probabilistic programming. We leave the formal connections between topological domains and ω qbses for future investigation, but we compare briefly, following §1.2. In traditional domain theory, the Borel structure is derived from the order; in TDT, both order and Borel structure are derived from the topology; and in ω qbses, both order and Borel structure are independent (see §3.3). Concretely, TDT disallows topological discontinuities such as our zero-testing conditionals; moreover it makes a direct connection to computability [Ackerman et al. 2011].

To conclude, we developed pre-domains (§3) for statistical probabilistic recursive programs via:

- a *convenient* category ω Qbs, being Cartesian closed and (co)complete (Cor. 3.10);
- a commutative probabilistic power-domain modelling synthetic measure theory (§4);
- an interpretation of recursive types, through axiomatic domain theory (§5);
- adequate models of recursive languages with continuous probabilistic choice and soft constraints, of recursively typed, untyped and simply types varieties (§6);
- canonicity through four independent characterisations of ω Qbs (§B).

This semantics gives sound reasoning principles about recursive probabilistic programs.

ACKNOWLEDGMENTS

Research supported by Balliol College Oxford (Career Development Fellowship), EPSRC (grants EP/N007387/1, EP/M023974/1), and the Royal Society. We are grateful to the reviewers for their suggestions. It has been helpful to discuss this work with the Oxford PL and Foundations groups, at the Domains 2018 and HOPE 2018 workshops, and with Marcelo Fiore, Chris Heunen, Paul Levy, Carol Mak, Gordon Plotkin, Alex Simpson, Hongseok Yang, amongst others.

REFERENCES

- Martin Abadi and Marcelo P Fiore. 1996. Syntactic considerations on recursive types. In *Proc. LICS 1996*. IEEE, 242–252.
- Nathanael L Ackerman, Cameron E. Freer, and Daniel M. Roy. 2011. Noncomputable conditional distributions. In *Proc. LICS 2011*. 107–116.
- J. Adamek and J. Rosicky. 1994. *Locally Presentable and Accessible Categories*. Cambridge University Press.
- Giorgio Bacci, Robert Furber, Dexter Kozen, Radu Mardare, Prakash Panangaden, and Dana Scott. 2018. Boolean-valued semantics for the stochastic λ -calculus. In *Proc. LICS 2018*.
- Tyler Barker. 2016. A monad for randomized algorithms. In *Proc. MFPS 2016. Electronic Notes in Theoretical Computer Science* 325, 47–62.
- Ingo Battenfeld, Matthias Schröder, and Alex Simpson. 2007. A convenient category of domains. *Electronic Notes in Theoretical Computer Science* 172 (2007), 69–99.
- Johannes Borgström, Ugo Dal Lago, Andrew D. Gordon, and Marcin Szymczak. 2016. A lambda-calculus foundation for universal probabilistic programming. In *Proc. ICFP 2016*. 33–46. Full version at [arxiv:abs/1512.08990](https://arxiv.org/abs/1512.08990).
- Thomas Ehrhard, Michele Pagani, and Christine Tasson. 2017. Measurable cones and stable, measurable functions: a model for probabilistic higher-order programming. *Proceedings of the ACM on Programming Languages* 2, POPL (2017), 59.
- Marcelo P. Fiore. 1996. *Axiomatic Domain Theory in Categories of Partial Maps*. Cambridge University Press.
- Marcelo P Fiore and Gordon D Plotkin. 1994. An axiomatisation of computationally adequate domain theoretic models of FPC. In *Proc. LICS 1994*. 92–102.
- Noah Goodman, Vikash Mansinghka, Daniel M Roy, Keith Bonawitz, and Joshua B Tenenbaum. 2008. Church: a language for generative models. In *Proc. UAI 2008*.
- Noah Goodman and Andreas Stuhlmüller. 2014. Design and Implementation of Probabilistic Programming Languages. <http://dippl.org>. Online book.
- Jean Goubault-Larrecq and Daniele Varacca. 2011. Continuous random variables. In *Logic in Computer Science (LICS), 2011 26th Annual IEEE Symposium on*. IEEE, 97–106.
- Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. 2017. A convenient category for higher-order probability theory. In *Proc. LICS 2017*.
- Daniel Huang, Greg Morrisett, and Bas Spitters. 2018. An application of computable distributions to the semantics of probabilistic programs. (2018). [arxiv:1806.07966](https://arxiv.org/abs/1806.07966).
- Claire Jones and Gordon D Plotkin. 1989. A probabilistic powerdomain of evaluations. In *Proc. LICS 1989*. 186–195.
- Achim Jung and Regina Tix. 1998. The troublesome probabilistic powerdomain, In *Proc. Comprom 1998. Electronic Notes in Theoretical Computer Science* 13, 70–91.
- Olav Kallenberg. 2006. *Foundations of modern probability*. Springer.
- Olav Kallenberg. 2017. *Random Measures, Theory and Applications*. Probability Theory and Stochastic Modelling, Vol. 77. Springer.
- Ohad Kammar and Gordon D. Plotkin. 2012. Algebraic foundations for effect-dependent optimisations. In *Proc. POPL 2012*. 349–360.
- Alexander Kechris. 2012. *Classical descriptive set theory*. Vol. 156. Springer Science & Business Media.
- Klaus Keimel and Gordon D Plotkin. 2009. Predicate transformers for extended probability and non-determinism. *Mathematical Structures in Computer Science* 19, 3 (2009), 501–539.
- Oleg Kiselyov and Chung-Chieh Shan. 2009. Embedded probabilistic programming. In *Domain-Specific Languages*. Springer, 360–384.
- Anders Kock. 2012. Commutative monads as a theory of distributions. *Theory and Applications of Categories* 26, 4 (2012), 97–131.
- Jimmy D. Lawson. 1982. Valuations on continuous lattices. In *Continuous Lattices and Related Topics*, Rudolf-Eberhard Höffman (Ed.). Mathematik Arbeitspapiere, Vol. 27. Universitat Bremen, 204–225.
- Paul Blain Levy. 2004. *Call-By-Push-Value: A Functional/Imperative Synthesis*. Kluwer.
- Vikash K. Mansinghka, Daniel Selsam, and Yura N. Perov. 2014. Venture: a higher-order probabilistic programming platform with programmable inference. *arXiv:1404.0099* (2014). [http://arxiv.org/abs/1404.0099](https://arxiv.org/abs/1404.0099)
- Dylan McDermott and Ohad Kammar. 2018. Factorisation systems for logical relations and monadic lifting in type-and-effect system semantics. *Proc. MFPS 2018* (2018). To appear.
- Michael W Mislove. 2016. Domains and random variables. *arXiv preprint arXiv:1607.07698* (2016).
- Eugenio Moggi. 1989. Computational lambda-calculus and monads. In *Proc. LICS 1989*. 14–23.
- Praveen Narayanan, Jacques Carette, Wren Romano, Chung-chieh Shan, and Robert Zinkov. 2016. Probabilistic inference by program transformation in Hakaru (system description). In *Proc. FLOPS 2016*. Springer, 62–79.
- Federico Olmedo, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. 2016. Reasoning about recursive probabilistic programs. In *Proc. LICS 2016*. 672–681.
- Andrew M Pitts. 1996. Relational properties of domains. *Inform. Comput.* 127, 2 (1996), 66–90.

- Gordon D. Plotkin. 1977. LCF Considered as a Programming Language. *Theor. Comput. Sci.* 5, 3 (1977), 223–255.
- Nasser Saheb-Djahromi. 1980. CPO's of measures for nondeterminism. *Theoret. Comput. Sci.* 12, 1 (1980), 19–37.
- Adam Šcibior, Ohad Kammar, Matthijs Vákár, Sam Staton, Hongseok Yang, Yufei Cai, Klaus Ostermann, Sean K. Moss, Chris Heunen, and Zoubin Ghahramani. 2017. Denotational validation of higher-order Bayesian inference. *Proc. ACM Program. Lang.* 2, POPL, Article 60 (Dec. 2017), 29 pages.
- Dana S. Scott. 1993. A type-theoretical alternative to ISWIM, CUCH, OWHY. *Theoret. Comput. Sci.* 121, 1 (1993), 411–440.
- Chung-chieh Shan and Norman Ramsey. 2017. Exact Bayesian inference by symbolic disintegration. In *Proc. POPL 2017*. 130–144.
- Michael B. Smyth and Gordon D. Plotkin. 1982. The category-theoretic solution of recursive domain equations. *SIAM J. Comput.* 11, 4 (1982), 761–783.
- Sam Staton. 2017. Commutative semantics for probabilistic programming. In *Proc. ESOP 2017*.
- Sam Staton, Hongseok Yang, Frank Wood, Chris Heunen, and Ohad Kammar. 2016. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *Proc. LICS 2016*. 525–534.
- Matthijs Vákár and Luke Ong. 2018. On S-Finite Measures and Kernels. *arXiv preprint arXiv:1810.01837* (2018).
- David Wingate, Andreas Stuhlmüller, and Noah Goodman. 2011. Lightweight implementations of probabilistic programming languages via transformational compilation. In *Proc. AISTATS 2011*.
- Frank Wood, Jan Willem van de Meent, and Vikash Mansinghka. 2014. A new approach to probabilistic programming inference. In *Proc. AISTATS 2014*. 1024–1032.