

A Convenient Category for Higher-Order Probability Theory

Chris Heunen Ohad Kammar Sam Staton Hongseok Yang
 University of Edinburgh, UK University of Oxford, UK University of Oxford, UK University of Oxford, UK

Abstract—Higher-order probabilistic programming languages allow programmers to write sophisticated models in machine learning and statistics in a succinct and structured way, but step outside the standard measure-theoretic formalization of probability theory. Programs may use both higher-order functions and continuous distributions, or even define a probability distribution on functions. But standard probability theory does not handle higher-order functions well: the category of measurable spaces is not cartesian closed.

Here we introduce quasi-Borel spaces. We show that these spaces: form a new formalization of probability theory replacing measurable spaces; form a cartesian closed category and so support higher-order functions; form a well-pointed category and so support good proof principles for equational reasoning; and support continuous probability distributions. We demonstrate the use of quasi-Borel spaces for higher-order functions and probability by: showing that a well-known construction of probability theory involving random functions gains a cleaner expression; and generalizing de Finetti’s theorem, that is a crucial theorem in probability theory, to quasi-Borel spaces.

I. INTRODUCTION

To express probabilistic models in machine learning and statistics in a succinct and structured way, it pays to use *higher-order* programming languages, such as Church [16], Venture [24], or Anglican [37]. These languages support advanced features from both programming language theory and probability theory, while providing generic inference algorithms for answering probabilistic queries, such as marginalization and posterior computation, for all models written in the language. As a result, the programmer can succinctly express a sophisticated probabilistic model and explore its properties while avoiding the nontrivial busywork of designing a custom inference algorithm.

This exciting development comes at a foundational price. Programs in these languages may combine higher-order functions and continuous distributions, or even define a probability distribution on functions. But the standard measure-theoretic formalization of probability theory does not handle higher-order functions well, as the category of measurable spaces is not cartesian closed [1]. For instance, the Anglican implementation of Bayesian linear regression in Figure 1 goes beyond the standard measure-theoretic foundation of probability theory, as it defines a probability distribution on functions $\mathbb{R} \rightarrow \mathbb{R}$.

We introduce a new formalization of probability theory that accommodates higher-order functions. The main notion

```

1  (defquery Bayesian-linear-regression
2    (let [f (let [s (sample (normal 0.0 3.0))
3                 b (sample (normal 0.0 3.0))]
4                   (fn [x] (+ (* s x) b)))]
5      (observe (normal (f 1.0) 0.5) 2.5)
6      (observe (normal (f 2.0) 0.5) 3.8)
7      (observe (normal (f 3.0) 0.5) 4.5)
8      (observe (normal (f 4.0) 0.5) 6.2)
9      (observe (normal (f 5.0) 0.5) 8.0)
10     (predict :f f)))

```




Fig. 1. Bayesian linear regression in Anglican. The program defines a probability distribution on functions $\mathbb{R} \rightarrow \mathbb{R}$. It first samples a random linear function f by randomly selecting slope s and intercept b . It then adjusts the probability distribution of the function to better describe five observations $(1.0, 2.5)$, $(2.0, 3.8)$, $(3.0, 4.5)$, $(4.0, 6.2)$ and $(5.0, 8.0)$ by posterior computation. In the graph, each line has been sampled from the posterior distribution over linear functions.

replacing a measurable space is a *quasi-Borel space*: a set X equipped with a collection of functions $M_X \subseteq [\mathbb{R} \rightarrow X]$ satisfying certain conditions (Def. 7). Intuitively, M_X is the set of random variables of type X . Here \mathbb{R} means that the randomness of random variables in M_X comes from (a probability distribution on) \mathbb{R} , one of the best behaving measurable spaces. Thus the primitive notion shifts from measurable subset to random variable, which is traditionally a derived notion. For related ideas see §IX.

Quasi-Borel spaces have good properties and structure.

- The category of quasi-Borel spaces is *well-pointed*, since a morphism is just a structure-preserving function (§III). (This is in contrast to [34, §8]).
- The category of quasi-Borel spaces is cartesian closed (§IV), so that it becomes a setting to study probability distributions on *higher-order* functions.

- There is a natural notion of probability measure on quasi-Borel spaces (Def. 10). The space of all probability measures is again a quasi-Borel space, and forms the basis for a commutative *monad* on the category of quasi-Borel spaces (§V). Thus quasi-Borel spaces form semantics for a probabilistic programming language in the monadic style [26].

We also illustrate the use of quasi-Borel spaces.

- *Bayesian regression* (§VI). Quasi-Borel spaces are a natural setting for understanding programs such as the one in Figure 1: the prior (Lines 2–4) defines a probability distribution over functions f , i.e. a measure on $\mathbb{R}^{\mathbb{R}}$, and the posterior (illustrated in the graph), is again a probability measure on $\mathbb{R}^{\mathbb{R}}$, conditioned by the observations (Lines 5–9).
- *Randomization* (§VII). A key idea of categorical logic is that $\forall\exists$ statements should become statements about quotients of objects. The structure of quasi-Borel spaces allows us to rephrase a crucial randomization lemma in this way. Classically, it says that every probability kernel arises from a random function. In the setting of quasi-Borel spaces, it says that the space of probability kernels $P(\mathbb{R})^X$ is a quotient of the space of random functions, $P(\mathbb{R}^X)$ (Theorem 26). Notice that the higher-order structure of quasi-Borel spaces allows us to succinctly state this result.
- *De Finetti’s theorem* (§VIII). Probability theorists often encounter problems when working with arbitrary probability measures on arbitrary measurable spaces. Quasi-Borel spaces allow us to better manage the source of randomness. For example, de Finetti’s theorem is a foundational result in Bayesian statistics which says that every exchangeable random sequence can be generated by randomly mixing multiple independent and identically distributed sequences. The theorem is known to hold for standard Borel spaces [8] or measurable spaces that arise from good topologies [18], but not for arbitrary measurable spaces [9]. We show that it holds for all quasi-Borel spaces (Theorem 29).

All of this is evidence that quasi-Borel spaces form a convenient category for higher-order probability theory.

II. PRELIMINARIES ON PROBABILITY MEASURES AND MEASURABLE SPACES

Definition 1. The Borel sets form the least collection $\Sigma_{\mathbb{R}}$ of subsets of \mathbb{R} that satisfies the following properties:

- intervals (a, b) are Borel sets;
- complements of Borel sets are Borel;
- countable unions of Borel sets are Borel.

The Borel sets play a crucial role in probability theory because of the tight connection between the notion of probability measure and the axiomatization of Borel sets.

Definition 2. A probability measure on \mathbb{R} is a function $\mu: \Sigma_{\mathbb{R}} \rightarrow [0, 1]$ satisfying $\mu(\mathbb{R}) = 1$ and $\mu(\biguplus S_i) = \sum \mu(S_i)$ for any countable sequence of disjoint Borel sets S_i .

The natural generalization gives measurable spaces.

Definition 3. A σ -algebra on a set X is a nonempty family of subsets of X that is closed under complements and countable unions. A measurable space is a pair (X, Σ_X) of a set X and a σ -algebra Σ_X on it. A probability measure on a measurable space X is a function $\mu: \Sigma_X \rightarrow [0, 1]$ satisfying $\mu(X) = 1$ and $\mu(\biguplus S_i) = \sum \mu(S_i)$ for any countable sequence of disjoint sets $S_i \in \Sigma_X$.

The Borel sets of the reals form a leading example of a σ -algebra. Other important examples are countable sets with their *discrete σ -algebra*, which contains all subsets. We can characterize these spaces as standard Borel spaces, but first introduce the appropriate structure-preserving maps.

Definition 4. Let (X, Σ_X) and (Y, Σ_Y) be measurable spaces. A measurable function $f: X \rightarrow Y$ is a function such that $f^{-1}(U) \in \Sigma_X$ when $U \in \Sigma_Y$.

Thus a measurable function $f: X \rightarrow Y$ lets us *push-forward* a probability measure μ on X to a probability measure $f_*\mu$ on Y by $(f_*\mu)(U) = \mu(f^{-1}(U))$. Measurable spaces and measurable functions form a category **Meas**.

Real-valued measurable functions $f: X \rightarrow \mathbb{R}$ can be integrated with respect to a probability measure μ on (X, Σ_X) . The *integral* of a nonnegative function f is

$$\int_X f \, d\mu \stackrel{\text{def}}{=} \sup_{\{U_i\}} \sum_i \left(\mu(U_i) \cdot \inf_{x \in U_i} f(x) \right),$$

where $\{U_i\}$ ranges over finite partitions of X into measurable subsets. When f may be negative, its integral is

$$\int_X f \, d\mu \stackrel{\text{def}}{=} \left(\int_X \max(0, f) \, d\mu \right) - \left(\int_X \max(0, -f) \, d\mu \right)$$

when those two integrals exist. When it is convenient to make the integrated variable explicit, we write $\int_{x \in U} f(x) \, d\mu$ for $\int_X (\lambda x. f(x) \cdot [x \in U]) \, d\mu$, where $U \in \Sigma_X$ is a measurable subset and $[\varphi]$ has the value 1 if φ holds and 0 otherwise.

A. Standard Borel spaces

Proposition 5 (e.g. [23], App. A1). For a measurable space (X, Σ_X) the following are equivalent:

- (X, Σ_X) is a retract of $(\mathbb{R}, \Sigma_{\mathbb{R}})$, that is, there exist measurable $X \xrightarrow{f} \mathbb{R} \xrightarrow{g} X$ such that $g \circ f = \text{id}_X$;
- (X, Σ_X) is either measurably isomorphic to $(\mathbb{R}, \Sigma_{\mathbb{R}})$ or countable and discrete;
- X has a complete metric with a countable dense subset and Σ_X is the least σ -algebra containing all open sets.

When (X, Σ_X) satisfies any of the above conditions, we call it *standard Borel space*. These spaces play an important role in probability theory because they enjoy properties that do not hold for general measurable spaces, such as the existence of conditional probability kernels [23], [28] and de Finetti’s theorem for exchangeable random processes [9].

Besides \mathbb{R} , another popular uncountable standard Borel space is $(0, 1)$ with the σ -algebra $\{U \cap (0, 1) \mid U \in \Sigma_{\mathbb{R}}\}$. As

the above proposition indicates, these spaces are isomorphic by, for instance, $\lambda r. \frac{1}{(1+e^{-r})} : \mathbb{R} \rightarrow (0, 1)$.

B. Failure of cartesian closure

Proposition 6 (Aumann, [1]). *The category \mathbf{Meas} is not cartesian closed: there is no space of functions $\mathbb{R} \rightarrow \mathbb{R}$.*

Specifically, the evaluation function

$$\varepsilon: \mathbf{Meas}(\mathbb{R}, \mathbb{R}) \times \mathbb{R} \rightarrow \mathbb{R} \quad \text{with} \quad \varepsilon(f, r) = f(r)$$

is never measurable ($\mathbf{Meas}(\mathbb{R}, \mathbb{R}) \times \mathbb{R}, \Sigma \otimes \Sigma_{\mathbb{R}} \rightarrow (\mathbb{R}, \Sigma_{\mathbb{R}})$) regardless of the choice of σ -algebra Σ on $\mathbf{Meas}(\mathbb{R}, \mathbb{R})$. Here, $\Sigma \otimes \Sigma_{\mathbb{R}}$ is the product σ -algebra, generated by rectangles $(U \times V)$ for $U \in \Sigma$ and $V \in \Sigma_{\mathbb{R}}$.

III. QUASI-BOREL SPACES

The typical situation in probability theory is that there is a fixed measurable space $(\Omega, \Sigma_{\Omega})$, called the *sample space*, from which all randomness originates, and that observations are made in terms of random variables, which are pairs (X, f) of a measurable space of observations (X, Σ_X) and a measurable function $f: \Omega \rightarrow X$. From this perspective, the notion of measurable function is more important than the notion of measurable space. In some ways, the σ -algebra Σ_X is only used as an intermediary to restrain the class of measurable functions $\Omega \rightarrow X$.

We now use this idea as a basis for our new notion of space. In doing so, we assume that our sample space Ω is the real numbers, which makes probabilities behave well.

Definition 7. A quasi-Borel space is a set X together with a set $M_X \subseteq [\mathbb{R} \rightarrow X]$ satisfying:

- $\alpha \circ f \in M_X$ if $\alpha \in M_X$ and $f: \mathbb{R} \rightarrow \mathbb{R}$ is measurable;
- $\alpha \in M_X$ if $\alpha: \mathbb{R} \rightarrow X$ is constant;
- if $\mathbb{R} = \biguplus_{i \in \mathbb{N}} S_i$, with each set S_i Borel, and $\alpha_1, \alpha_2, \dots \in M_X$, then β is in M_X , where $\beta(r) = \alpha_i(r)$ for $r \in S_i$.

The name ‘quasi-Borel space’ is motivated firstly by analogy to quasi-topological spaces (see §IX), and secondly in recognition of the intimate connection to the standard Borel space \mathbb{R} (see also Prop. 15(2)).

Example 8. For every measurable space (X, Σ_X) , let M_{Σ_X} be the set of measurable functions $\mathbb{R} \rightarrow X$. Thus M_{Σ_X} is the set of X -valued random variables. In particular: \mathbb{R} itself can be considered as a quasi-Borel space, with $M_{\mathbb{R}}$ the set of measurable functions $\mathbb{R} \rightarrow \mathbb{R}$; the two-element discrete space 2 can be considered as a quasi-Borel space, with M_2 the set of measurable functions $\mathbb{R} \rightarrow 2$, which are exactly the characteristic functions of the Borel sets (Def. 1).

Before we continue, we remark that the notion of quasi-Borel space is invariant under replacing \mathbb{R} with a different uncountable standard Borel space.

Proposition 9. For any measurable space $(\Omega, \Sigma_{\Omega})$, any measurable isomorphism $\iota: \mathbb{R} \rightarrow \Omega$, any set X , and any set N

of functions $\Omega \rightarrow X$, the pair $(X, \{\alpha \circ \iota \mid \alpha \in N\})$ is a quasi-Borel space if and only if:

- $\alpha \circ f \in N$ if $\alpha \in N$ and $f: \Omega \rightarrow \Omega$ is measurable;
- $\alpha \in N$ if $\alpha: \Omega \rightarrow X$ is constant;
- if $\Omega = \biguplus_{i \in \mathbb{N}} S_i$, with each set $S_i \in \Sigma_{\Omega}$, and $\alpha_1, \alpha_2, \dots \in N$, then β is in N , where $\beta(r) = \alpha_i(r)$ if $r \in S_i$.

By Prop. 5, the measurable spaces isomorphic to \mathbb{R} are the uncountable standard Borel spaces. Note that the choice of isomorphism ι is not important: it does not appear in the three conditions.

Probability theory typically considers a basic probability measure on the sample space Ω . Each random variable, that is each measurable function $\Omega \rightarrow X$, then induces a probability measure on X by pushing forward the basic measure. Quasi-Borel spaces take this idea as an axiomatic notion of probability measure.

Definition 10. A probability measure on a quasi-Borel space (X, M_X) is a pair (α, μ) of $\alpha \in M_X$ and a probability measure μ on \mathbb{R} (as in Def. 2).

A. Morphisms and integration

Definition 11. A morphism of quasi-Borel spaces $(X, M_X) \rightarrow (Y, M_Y)$ is a function $f: X \rightarrow Y$ such that $f \circ \alpha \in M_Y$ if $\alpha \in M_X$. Write $\mathbf{QBS}((X, M_X), (Y, M_Y))$ for the set of morphisms from (X, M_X) to (Y, M_Y) .

In particular, elements of M_X are precisely morphisms $(\mathbb{R}, M_{\mathbb{R}}) \rightarrow (X, M_X)$, so $M_X = \mathbf{QBS}((\mathbb{R}, M_{\mathbb{R}}), (X, M_X))$.

Morphisms compose as functions, and identity functions are morphisms, so quasi-Borel spaces form a category \mathbf{QBS} .

Example 12. There are two canonical ways to equip a set X with a quasi-Borel space structure. The first structure M_X^R consists of all functions $\mathbb{R} \rightarrow X$. The second structure M_X^L consists of all functions $\beta: \mathbb{R} \rightarrow X$ for which there exist: a countable subset $I \subseteq \mathbb{N}$; a measurable $f: \mathbb{R} \rightarrow \mathbb{R}$; a partition $\mathbb{R} = \biguplus_{i \in I} S_i$ with every S_i measurable; and a sequence $(x_i)_{i \in I}$ in X , such that $\beta(r) = x_i$ whenever $f(r) \in S_i$. These are the right and left adjoints, respectively, to the forgetful functor from \mathbf{QBS} to \mathbf{Set} .

Def. 11 is independent of \mathbb{R} : the sample space may be any uncountable standard Borel space.

Proposition 13. Consider a measurable space $(\Omega, \Sigma_{\Omega})$ with a measurable isomorphism $\iota: \mathbb{R} \rightarrow \Omega$. For $i \in 1, 2$, let X_i be a set and N_i a set of functions $\Omega \rightarrow X_i$ such that

$$M_i = (X_i, \{\alpha \circ \iota \mid \alpha \in N_i\})$$

are quasi-Borel spaces. A function $g: X_1 \rightarrow X_2$ is a morphism $(X_1, M_1) \rightarrow (X_2, M_2)$ if and only if $g \circ \alpha \in N_2$ for $\alpha \in N_1$.

Morphisms between quasi-Borel spaces are analogous to measurable functions between measurable spaces. The crucial properties of measurable functions are that they work well with (probability) measures: we can push-forward these measures,

and integrate over them. Morphisms of quasi-Borel spaces also support these constructions.

- *Pushing forward*: if $f: X \rightarrow Y$ is a morphism and (α, μ) is a probability measure on X then $f \circ \alpha$ is by definition in M_Y and so $(f \circ \alpha, \mu)$ is a probability measure on Y .
- *Integrating*: If $f: X \rightarrow \mathbb{R}$ is a morphism of quasi-Borel spaces and (α, μ) is a probability measure on X , the integral of f with respect to (α, μ) is

$$\int f d(\alpha, \mu) \stackrel{\text{def}}{=} \int_{\mathbb{R}} (f \circ \alpha) d\mu. \quad (1)$$

So integration formally reduces to integration on \mathbb{R} .

B. Relationship to measurable spaces

If we regard a subset $S \subseteq X$ as its characteristic function $\chi_S: X \rightarrow 2$, then we can regard a σ -algebra on a set X as a set of characteristic functions $F_X \subseteq [X \rightarrow 2]$ satisfying certain conditions. Thus a measurable space (Def. 3) could equivalently be described as a pair (X, F_X) of a set X and a collection $F_X \subseteq [X \rightarrow 2]$ of characteristic functions. Moreover, from this perspective, a measurable function $f: (X, F_X) \rightarrow (Y, F_Y)$ is simply a function $f: X \rightarrow Y$ such that $\chi \circ f \in F_X$ if $\chi \in F_Y$. Thus quasi-Borel spaces shift the emphasis from characteristic functions $X \rightarrow 2$ to random variables $\mathbb{R} \rightarrow X$.

1) *Quasi-Borel spaces as structured measurable spaces*: A subset $S \subseteq X$ is in the σ -algebra Σ_X of a measurable space (X, Σ_X) if and only if its characteristic function $X \rightarrow 2$ is measurable. With this in mind, we *define* a measurable subset of a quasi-Borel space (X, M_X) to be a subset $S \subseteq X$ such that the characteristic function $X \rightarrow 2$ is a morphism of quasi-Borel spaces.

Proposition 14. *The collection of all measurable subsets of a quasi-Borel space (X, M_X) is characterized as*

$$\Sigma_{M_X} \stackrel{\text{def}}{=} \{U \mid \forall \alpha \in M_X. \alpha^{-1}(U) \in \Sigma_{\mathbb{R}}\} \quad (2)$$

and forms a σ -algebra.

Thus we can understand a quasi-Borel space as a measurable space (X, Σ_X) equipped with a class of measurable functions $M_X \subseteq [\mathbb{R} \rightarrow X]$ determining the σ -algebra by $\Sigma_X = \Sigma_{M_X}$ as in (2).

Moreover, every morphism $(X, M_X) \rightarrow (Y, M_Y)$ is also a measurable function $(X, \Sigma_{M_X}) \rightarrow (Y, \Sigma_{M_Y})$ (but the converse does not hold in general).

A probability measure (α, μ) on a quasi-Borel space (X, M_X) induces a probability measure $\alpha_* \mu$ on the underlying measurable space. Integration as in (1) matches the standard definition for measurable spaces.

2) *An adjunction embedding standard Borel spaces*: Under some circumstances morphisms of quasi-Borel spaces coincide with measurable functions.

Proposition 15. *Let (Y, Σ_Y) be a measurable space.*

- 1) *If (X, M_X) is a quasi-Borel space, a function $X \rightarrow Y$ is a measurable function $(X, \Sigma_{M_X}) \rightarrow (Y, \Sigma_Y)$ if and only if it is a morphism $(X, M_X) \rightarrow (Y, M_{\Sigma_Y})$.*

- 2) *If (X, Σ_X) is a standard Borel space, a function $X \rightarrow Y$ is a morphism $(X, M_{\Sigma_X}) \rightarrow (Y, M_{\Sigma_Y})$ if and only if it is a measurable function $(X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$.*

Proposition 15(1) means there is an adjunction

$$\text{Meas} \begin{array}{c} \xleftarrow{L} \\[-1ex] \perp \\[-1ex] \xrightarrow{R} \end{array} \text{QBS}$$

where $L(X, M_X) = (X, \Sigma_{M_X})$ and $R(X, \Sigma_X) = (X, M_{\Sigma_X})$. Proposition 15(2) means that the functor R is full and faithful when restricted to standard Borel spaces. Equivalently, $L(R(X, \Sigma_X)) = (X, \Sigma_X)$, that is $\Sigma_X = \Sigma_{M_{\Sigma_X}}$ for standard Borel spaces (X, Σ_X) .

IV. PRODUCTS, COPRODUCTS AND FUNCTION SPACES

Quasi-Borel spaces support products, coproducts, and function spaces. These basic constructions form the basis for interpreting simple type theory in quasi-Borel spaces.

Proposition 16 (Products). *If $(X_i, M_{X_i})_{i \in I}$ is a family of quasi-Borel spaces indexed by a set I , then $(\prod_i X_i, M_{\prod_i X_i})$ is a quasi-Borel space, where $\prod_i X_i$ is the set product, and*

$$M_{\prod_i X_i} \stackrel{\text{def}}{=} \{f: \mathbb{R} \rightarrow \prod_i X_i \mid \forall i. (\pi_i \circ f) \in M_{X_i}\}.$$

The projections $\prod_i X_i \rightarrow X_i$ are morphisms, and provide the structure of a categorical product in **QBS**.

Proposition 17 (Coproducts). *If $(X_i, M_{X_i})_{i \in I}$ is a family of quasi-Borel spaces indexed by a countable set I , then $(\coprod_i X_i, M_{\coprod_i X_i})$ is a quasi-Borel space, where $\coprod_i X_i$ is the disjoint union of sets,*

$$M_{\coprod_i X_i} \stackrel{\text{def}}{=} \{\lambda r. (f(r), \alpha_{f(r)}(r)) \mid f: \mathbb{R} \rightarrow I \text{ is measurable}, (\alpha_i \in M_{X_i})_{i \in \text{image}(f)}\},$$

and I carries the discrete σ -algebra. This space has the universal property of a coproduct in the category **QBS**.

Proof notes. The third condition of quasi-Borel spaces is needed here. It is a crucial step in showing that for an I -indexed family of morphisms $(f_i: X_i \rightarrow Z)_{i \in I}$, the copairing $[f_i]_{i \in I}: \coprod_{i \in I} X_i \rightarrow Z$ is again a morphism. \square

Proposition 18 (Function spaces). *If (X, M_X) and (Y, M_Y) are quasi-Borel spaces, so is (Y^X, M_{Y^X}) , where $Y^X \stackrel{\text{def}}{=} \text{QBS}(X, Y)$ is the set of morphisms $X \rightarrow Y$, and*

$$M_{Y^X} \stackrel{\text{def}}{=} \{\alpha: \mathbb{R} \rightarrow Y^X \mid \text{uncurry}(\alpha) \in \text{QBS}(\mathbb{R} \times X, Y)\}.$$

The evaluation function $Y^X \times X \rightarrow Y$ is a morphism and has the universal property of the function space. Thus **QBS** is a cartesian closed category.

Proof notes. The only difficult part is showing that (Y^X, M_{Y^X}) satisfies the third condition of quasi-Borel spaces. Prop. 17 is useful here. \square

A. Relationship with standard Borel spaces

Recall that standard Borel spaces can be thought of as a full subcategory of the quasi-Borel spaces, that is, the functor $R: \mathbf{Meas} \rightarrow \mathbf{QBS}$ is full and faithful (Prop. 15(2)) when restricted to the standard Borel spaces. This full subcategory has the same countable products, coproducts and function spaces (whenever they exist). We may thus regard quasi-Borel spaces as a conservative extension of standard Borel spaces that supports simple type theory.

Proposition 19. *The functor $R(X, \Sigma_X) = (X, M_{\Sigma_X})$:*

- 1) *preserves products of standard Borel spaces:* $R(\prod_i X_i) = \prod_i R(X_i)$, where $(X_i, \Sigma_{X_i})_{i \in I}$ is a countable family of standard Borel spaces;
- 2) *preserves spaces of functions between standard Borel spaces whenever they exist:* if (Y, Σ_Y) is countable and discrete, and (X, Σ_X) is standard Borel, then $R(X^Y) = R(X)^{R(Y)}$;
- 3) *preserves countable coproducts of standard Borel spaces:* $R(\coprod_i X_i) = \coprod_i R(X_i)$, where $(X_i, \Sigma_{X_i})_{i \in I}$ is a countable family of standard Borel spaces.

Consequently, a standard programming language semantics in standard Borel spaces can be conservatively embedded in quasi-Borel spaces, allowing higher-order functions while preserving all the type theoretic structure.

We note, however, that in light of Prop. 6, the quasi-Borel space $\mathbb{R}^{\mathbb{R}}$ does not come from a standard Borel space. Moreover, the left adjoint $L: \mathbf{QBS} \rightarrow \mathbf{Meas}$ does not preserve products in general. For quasi-Borel spaces (X, M_X) and (Y, M_Y) , we always have $\Sigma_{M_X} \otimes \Sigma_{M_Y} \subseteq \Sigma_{M_{X \times Y}}$, but not always \supseteq . Indeed, $\Sigma_{M_{\mathbb{R}^{\mathbb{R}}}} \otimes \Sigma_{\mathbb{R}} \neq \Sigma_{M_{(\mathbb{R}^{\mathbb{R}} \times \mathbb{R})}}$, by Prop. 6.

V. A MONAD OF PROBABILITY MEASURES

In this section we will show that the probability measures on a quasi-Borel space form a quasi-Borel space again. This gives a commutative monad that generalizes the Giry monad for measurable spaces [15].

A. Monads

We use the Kleisli triple formulation of monads (see e.g. [26]). Recall that a monad on a category \mathcal{C} comprises

- for any object X , an object $T(X)$;
- for any object X , a morphism $\eta: X \rightarrow T(X)$;
- for any objects X, Y , a function

$$(\gg): \mathcal{C}(X, T(Y)) \rightarrow \mathcal{C}(T(X), T(Y)).$$

We write $(t\gg=f)$ for $(\gg)(f)(t)$.

This is subject to the conditions $(t\gg=\eta) = t$, $(\eta(x)\gg=f) = f(x)$, and $t\gg=(\lambda x. (f(x)\gg=g)) = (t\gg=f)\gg=g$.

The intuition is that $T(X)$ is an object of computations returning X , that η is the computation that returns immediately, and that $t\gg=f$ sequences computations, first running computation t and then calling f with the result.

When \mathcal{C} is cartesian closed, a monad is *strong* if $(\gg): (T(Y))^X \rightarrow (T(Y))^{T(X)}$, internalizes to an operation $(\gg): (T(Y))^X \rightarrow (T(Y))^{T(X)}$,

and then the conditions are understood as expressions in a cartesian closed category.

B. Kernels and the Giry monad

We recall the notion of probability kernel, which is a measurable family of probability measures.

Definition 20. *Let (X, Σ_X) and (Y, Σ_Y) be measurable spaces. A probability kernel from X to Y is a function $k: X \times \Sigma_Y \rightarrow [0, 1]$ such that $k(x, -)$ is a probability measure for all $x \in X$ (Def. 3), and $k(-, U)$ is a measurable function for all $U \in \Sigma_Y$ (Def. 4).*

We can classify probability kernels as follows. Let $G(X)$ be the set of probability measures on (X, Σ_X) . We can equip this set with the σ -algebra generated by $\{\mu \in G(X) \mid \mu(U) < r\}$, for $U \in \Sigma_X$ and $r \in [0, 1]$, to form a measurable space $(G(X), \Sigma_{G(X)})$. A measurable function $X \rightarrow G(Y)$ amounts to a probability kernel from X to Y .

The construction G has the structure of a monad, as first discussed by Giry [15]. A computational intuition is that $G(X)$ is a space of probabilistic computations over X , and this provides a semantic foundation for a first-order probabilistic programming language (see e.g. [34]). The unit $\eta: X \rightarrow G(X)$ lets $\eta(x)$ be the Dirac measure on x , with $\eta(x)(U) = 1$ if $x \in U$, and $\eta(x)(U) = 0$ if $x \notin U$. If $\mu \in G(X)$ and k is a measurable function $X \rightarrow G(Y)$, then $(\mu\gg=Gk)$ is the measure in $G(Y)$ with $(\mu\gg=Gk)(U) = \int_{x \in X} k(x)(U) d\mu$.

C. Equivalent measures on quasi-Borel spaces

Recall (Def. 10) that a probability measure (α, μ) on a quasi-Borel space (X, M_X) is a pair (α, μ) of a function $\alpha \in M_X$ and a probability measure μ on \mathbb{R} . Random variables are often equated when they describe the same distribution. Every probability measure (α, μ) determines a push-forward measure $\alpha_* \mu$ on the corresponding measurable space (X, Σ_{M_X}) , that assigns to $U \subseteq X$ the real number $\mu(\alpha^{-1}(U))$. We will identify two probability measures when they define the same push-forward measure, and write \sim for this equivalence relation.

This is a reasonable notion of equality even if we put aside the notion of measurable space, because two probability measures have the same push-forward measure precisely when they have the same integration operator: $(\alpha, \mu) \sim (\alpha', \mu')$ if and only if $\int f d(\alpha, \mu) = \int f d(\alpha', \mu')$ for all morphisms $f: (X, M_X) \rightarrow \mathbb{R}$. Nevertheless, other notions of equivalence could be used.

D. A probability monad

We now explain how to build a monad of probability measures on the category of quasi-Borel spaces, modulo this notion of equivalence. This monad P will inherit properties from the Giry monad. Technically, the functor $L: \mathbf{QBS} \rightarrow \mathbf{Meas}$ (Prop. 15) is a ‘monad opfunctor’ taking P to the Giry monad G , which means that it extends to a functor from the Kleisli category of P to the Kleisli category of G [36].

a) *On objects*: For a quasi-Borel space (X, M_X) , let

$$\begin{aligned} P(X) &= \{(\alpha, \mu) \text{ probability measure on } (X, M_X)\} / \sim, \\ M_{P(X)} &= \{\beta: \mathbb{R} \rightarrow P(X) \mid \exists \alpha \in M_X. \exists g \in \mathbf{Meas}(\mathbb{R}, G(\mathbb{R})). \\ &\quad \forall r \in \mathbb{R}. \beta(r) = [\alpha, g(r)]\}, \end{aligned}$$

where $[\alpha, \mu]$ denotes the equivalence class. Note that

$$P(X) \cong \{\alpha_* \mu \in G(X, \Sigma_{M_X}) \mid \alpha \in M_X, \mu \in G(\mathbb{R})\} \quad (3)$$

as sets, and $l_X([\alpha, \mu]) = \alpha_* \mu$ defines a measurable injection $l_X: L(P(X)) \rightarrow G(X, \Sigma_{M_X})$.

b) *Monad unit (return)*: Recall that the constant functions $(\lambda r. x)$ are all in M_X . For any probability measure μ on \mathbb{R} , the push-forward measure $(\lambda r. x)_* \mu$ on (X, Σ_{M_X}) is the Dirac measure on x , with $((\lambda r. x)_* \mu)(U) = 1$ if $x \in U$ and 0 otherwise. Thus $(\lambda r. x, \mu) \sim (\lambda r. x, \mu')$ for all measures μ, μ' on \mathbb{R} . The unit of P at (X, M_X) is the morphism $\eta: X \rightarrow P(X)$ given by

$$\eta_{(X, M_X)}(x) = [\lambda r. x, \mu] \quad (4)$$

for an arbitrary probability measure μ on \mathbb{R} .

c) *Bind*: To define $(\gg): P(Y)^X \rightarrow (P(Y))^{P(X)}$, suppose $f: X \rightarrow P(Y)$ is a morphism and $[\alpha, \mu]$ in $P(X)$. Since f is a morphism, there is a measurable $g: \mathbb{R} \rightarrow G(\mathbb{R})$ and a function $\beta \in M_Y$ such that $(f \circ \alpha)(r) = [\beta, g(r)]$. Set $([\alpha, \mu] \gg f) = [\beta, \mu \gg_G g]$, where $\mu \gg_G g$ is the bind of the Giry monad. This matches the bind of the Giry monad, since $((\alpha_* \mu) \gg_G (l_Y \circ f)) = \beta_*(\mu \gg_G g)$.

Theorem 21. *The data $(P, \eta, (\gg))$ above defines a strong monad on the category **QBS** of quasi-Borel spaces.*

Proof notes. The monad laws can be reduced to the laws for the monad G on **Meas** [15]. The monad on **QBS** is strong because $(\gg): P(Y)^X \rightarrow (P(Y))^{P(X)}$ is a morphism, which is shown by expanding the definitions. \square

Proposition 22. *The monad P satisfies these properties:*

- 1) *For $f: (X, M_X) \rightarrow (Y, M_Y)$, the functorial action $P(f): P(X) \rightarrow P(Y)$ is $[\alpha, \mu] \mapsto [f \circ \alpha, \mu]$.*
- 2) *It is a commutative monad, i.e. the order of sequencing doesn't matter: if $p \in P(X)$, $q \in P(Y)$, and f is a morphism $X \times Y \rightarrow P(Z)$, then $p \gg \lambda x. q \gg \lambda y. f(x, y)$ equals $q \gg \lambda y. p \gg \lambda x. f(x, y)$.*
- 3) *The faithful functor $L: \mathbf{QBS} \rightarrow \mathbf{Meas}$ with $L(X, M_X) = (X, \Sigma_{M_X})$ extends to a faithful functor $\text{Kleisli}(P) \rightarrow \text{Kleisli}(G)$, i.e. (L, l) is a monad opfuctor [36].*
- 4) *When (X, Σ_X) is a standard Borel space, the map l_X of Eq. (3) is a measurable isomorphism.*

VI. EXAMPLE: BAYESIAN REGRESSION

We are now in a position to explain the semantics of the Anglican program in Figure 1. The program can be split into three parts: a prior, a likelihood, and a posterior. Recall that Bayes' law says that the posterior is proportionate to the product of the prior and the likelihood.




Fig. 2. Illustration of 1000 sampled functions from the prior on $\mathbb{R}^{\mathbb{R}}$ for Bayesian linear regression (5).

a) *Prior*: Lines 2–4 define a prior measure on $\mathbb{R}^{\mathbb{R}}$:

```
prior  $\stackrel{\text{def}}{=} (\text{let } [s \text{ (sample (normal 0.0 3.0))} \\ \text{b (sample (normal 0.0 3.0))}] \\ (\text{fn } [x] (+ (* s x) b)))$ 
```

To describe this semantically, observe the following.

Proposition 23. *Let $(\Omega, \Sigma_{\Omega})$ be a standard Borel space, and (X, M_X) a quasi-Borel space. Let $\alpha: R(\Omega, \Sigma_{\Omega}) \rightarrow X$ be a morphism and μ a probability measure on $(\Omega, \Sigma_{\Omega})$. Any section-retraction pair $(\Omega \xrightarrow{\varsigma} \mathbb{R} \xrightarrow{\rho} \Omega) = \text{id}_{\Omega}$ has a probability measure $[\alpha \circ \rho, \varsigma_* \mu] \in P(X)$, that is independent of the choice of ς and ρ .*

Write $[\alpha, \mu]$ for the probability measure in this case.

Now, the program fragment *prior* describes the distribution $[\alpha, \nu \otimes \nu]$ in $P(\mathbb{R}^{\mathbb{R}})$ where ν is the normal distribution on \mathbb{R} with mean 0 and standard deviation 3, and where $\alpha: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^{\mathbb{R}}$ is given by $\alpha(s, b) \stackrel{\text{def}}{=} \lambda r. s \cdot r + b$. Informally,

$$\llbracket \text{prior} \rrbracket = [\alpha, \nu \otimes \nu] \in P(\mathbb{R}^{\mathbb{R}}). \quad (5)$$

Figure 2 illustrates this measure $[\alpha, \nu \otimes \nu]$. This denotational semantics can be made compositional, by using the commutative monad structure of P and the cartesian closed structure of the category **QBS** (following e.g. [26], [34]), but in this paper we focus on this example rather than spelling out the general case once again.

b) *Likelihood*: Lines 5–9 define the likelihood of the observations:

```
obs  $\stackrel{\text{def}}{=} (\text{observe (normal (f 1.0) 0.5) 2.5}) \\ \dots (\text{observe (normal (f 5.0) 0.5) 8.0})$ 
```

This program fragment has a free variable f of type $\mathbb{R}^{\mathbb{R}}$. Let us focus on line 5 for a moment:

```
obs1  $\stackrel{\text{def}}{=} (\text{observe (normal (f 1.0) 0.5) 2.5})$ 
```

Given a function $f: \mathbb{R} \rightarrow \mathbb{R}$, the likelihood of drawing 2.5 from a normal distribution with mean $f(1.0)$ and standard deviation 0.5 is

$$\llbracket f: \mathbb{R}^{\mathbb{R}} \vdash \text{obs1} \rrbracket = d(f(1.0), 2.5),$$

where $d: \mathbb{R}^2 \rightarrow [0, \infty)$ is the density of the normal distribution function with standard deviation 0.5:

$$d(\mu, x) = \sqrt{\frac{2}{\pi}} e^{-2(x-\mu)^2}.$$

Notice that we use a normal distribution to allow for some noise in the measurement. Informally, we are not recording an observation that $f(1.0)$ is *exactly* 2.5, since this would make regression impossible; rather, $f(1.0)$ is roughly 2.5.

Overall, lines 5–9 describe a likelihood weight which is the product of the likelihoods of the five data points, given $f: \mathbb{R}^{\mathbb{R}}$:

$$\begin{aligned} \llbracket f: \mathbb{R}^{\mathbb{R}} \vdash \text{obs} \rrbracket &= d(f(1), 2.5) \cdot d(f(2), 3.8) \cdot d(f(3), 4.5) \\ &\quad \cdot d(f(4), 6.2) \cdot d(f(5), 8.0). \end{aligned}$$

c) *Posterior*: We follow the recipe for a semantic posterior given in [34]. Putting the prior and likelihood together gives a probability measure in $P(\mathbb{R}^{\mathbb{R}} \times [0, \infty))$ which is found by pushing forward the measure $\llbracket \text{prior} \rrbracket \in P(\mathbb{R})$ along the function $(\text{id}, \llbracket \text{obs} \rrbracket): \mathbb{R}^{\mathbb{R}} \rightarrow \mathbb{R}^{\mathbb{R}} \times [0, \infty)$. This push-forward measure

$$P(\text{id}, \llbracket \text{obs} \rrbracket)(\llbracket \text{prior} \rrbracket) \in P(\mathbb{R}^{\mathbb{R}} \times [0, \infty))$$

is a measure over pairs (f, w) of functions together with their likelihood weight. We now find the posterior by multiplying the prior and the likelihood, and dividing by a normalizing constant. To do this we define a morphism

$$\text{norm}: P(X \times [0, \infty)) \rightarrow P(X) \uplus \{\text{error}\}$$

$$\text{norm}([\alpha, \beta], \nu) \stackrel{\text{def}}{=} \begin{cases} [\alpha, \nu_{\beta}/(\nu_{\beta}(\mathbb{R}))] & \text{if } 0 \neq \nu_{\beta}(\mathbb{R}) \neq \infty \\ \text{error} & \text{otherwise} \end{cases}$$

where $\nu_{\beta}: \Sigma_{\mathbb{R}} \rightarrow [0, \infty] \stackrel{\text{def}}{=} \lambda U. \int_{r \in U} (\beta(r)) d\nu$. The idea is that if $\beta: \mathbb{R} \rightarrow [0, \infty)$ and ν is a probability measure on \mathbb{R} then ν_{β} is always a posterior measure on \mathbb{R} , but it is typically not normalized, i.e. $\nu_{\beta}(\mathbb{R}) \neq 1$. We normalize it by dividing by the normalizing constant, as long as this division is well-defined.

Now, the semantics of the entire program in Figure 1 is $\text{norm}(P(\text{id}, \llbracket \text{obs} \rrbracket)(\llbracket \text{prior} \rrbracket))$, which is a measure in $P(\mathbb{R}^{\mathbb{R}})$. Calculating this posterior using Anglican’s inference algorithm 1mh gives the plot in the lower half of Figure 1.

d) *Defunctionalized regression and non-linear regression*: Of course, one can do regression without explicitly considering distributions over the space of all measurable functions, by instead directly calculating posterior distributions for the slope s and the intercept b . For example, one could defunctionalize the program in Fig. 1 in the style of Reynolds [29]. But defunctionalization is a whole-program transformation. By structuring the semantics using quasi-Borel spaces, we are able to work compositionally, without mentioning s and b explicitly on lines 5–10. The internal posterior calculations actually happen at the level of standard Borel spaces, and so a defunctionalized version would be in some sense equivalent, but from the programming perspective it helps to abstract away from this. The regression program in Fig. 1 is quickly adapted to fit other kinds of functions, e.g.

polynomials, or even programs from a small domain-specific language, simply by changing the prior in Lines 2–4.

VII. RANDOM FUNCTIONS

We discuss random variables and random functions, starting from the traditional setting. Let $(\Omega, \Sigma_{\Omega})$ be a measurable space with a probability measure. A *random variable* is a measurable function $(\Omega, \Sigma_{\Omega}) \rightarrow (X, \Sigma_X)$. A *random function* between measurable spaces (X, Σ_X) and (Y, Σ_Y) is a measurable function $(\Omega \times X, \Sigma_{\Omega} \otimes \Sigma_X) \rightarrow (Y, \Sigma_Y)$.

We can push forward a probability measure on Ω along a random variable $(\Omega, \Sigma_{\Omega}) \rightarrow (X, \Sigma_X)$ to get a probability measure on X , but in the traditional setting we cannot push forward a measure along a random function. Measurable spaces are not cartesian closed (Prop. 6), and so we cannot form a measurable space Y^X and we cannot curry a random function in general.

Now, if we revisit these definitions in the setting of quasi-Borel spaces, we *do* have function spaces, and so we can push forward along random functions. In fact, this is somewhat tautologous because a probability measure (Def. 10) on a function space is essentially the same thing as a random function: a probability measure on a function space $(Y, \Sigma_Y)^{(X, \Sigma_X)}$ is defined to be a pair (f, μ) of a probability measure μ on \mathbb{R} , our sample space, and a morphism $f: \mathbb{R} \rightarrow Y^X$; but to give a morphism $\mathbb{R} \rightarrow Y^X$ is to give a morphism $\mathbb{R} \times X \rightarrow Y$ (Prop. 18) as in the traditional definition of random function.

We have already encountered an example of a random function in Section VI: the prior for linear regression is a random function from \mathbb{R} to \mathbb{R} over the measurable space $(\mathbb{R} \times \mathbb{R}, \Sigma_{\mathbb{R}} \otimes \Sigma_{\mathbb{R}})$ with the measure $\nu \otimes \nu$. Random functions abound throughout probability theory and stochastic processes. The following section explores their use in the so-called randomization lemma, which is used throughout probability theory. By moving to quasi-Borel spaces, we can state this lemma succinctly (Theorem 26).

A. Randomization

An elementary but useful trick in probability theory is that every probability distribution on \mathbb{R} arises as a push-forward of the uniform distribution on $[0, 1]$. Even more useful is that this can be done in a parameterized way.

Proposition 24 ([23], Lem. 2.22). *Let (X, Σ_X) be a measurable space. For any kernel $k: X \times \Sigma_{\mathbb{R}} \rightarrow [0, 1]$ there is a measurable function $f: \mathbb{R} \times X \rightarrow \mathbb{R}$ such that $k(x, U) = v\{r \mid f(r, x) \in U\}$, where v is the uniform distribution on $[0, 1]$.*

For quasi-Borel spaces we can phrase this more succinctly: it is a result about a quotient of the space of random functions. We first define quotient spaces.

Proposition 25. *Let (X, M_X) be a quasi-Borel space, let Y be a set, and let $q: X \rightarrow Y$ be a surjection. Then (Y, M_Y) is a quasi-Borel space with $M_Y = \{q \circ \alpha \mid \alpha \in M_X\}$.*

We call such a space a *quotient space*.

Theorem 26. Let (X, M_X) be a quasi-Borel space. The space $(P(\mathbb{R}))^X$ of kernels is a quotient of the space $P(\mathbb{R}^X)$ of random functions.

Before proving this theorem, we use Prop. 24 to give an alternative characterization of our probability monad.

Lemma 27. Let (X, M_X) be a quasi-Borel space. The function $q: X^\mathbb{R} \rightarrow P(X)$ given by $q(\alpha) \stackrel{\text{def}}{=} [\alpha, v]$ is a surjection, with corresponding quotient space $(P(X), M_{P(X)})$:

$$M_{P(X)} = \{\lambda r \in \mathbb{R}. [\gamma(r), v] \mid \gamma \in M_{X^\mathbb{R}}\}, \quad (6)$$

where v is the uniform distribution on $[0, 1]$.

Proof notes. The direction (\subseteq) follows immediately from Prop. 24. For the direction (\supseteq) we must consider $\gamma \in M_{X^\mathbb{R}}$ and show that $(\lambda r \in \mathbb{R}. [\gamma(r), v])$ is in $M_{P(X)}$. This follows by considering the kernel $k: \mathbb{R} \rightarrow G(\mathbb{R} \times \mathbb{R})$ with $k(r) = v \otimes \delta_r$, so that $[\gamma(r), v] = [\text{uncurry}(\gamma), k(r)]$. Here we are using Prop. 23. \square

Proof of Theorem 26. Consider the evident morphism $q: P(\mathbb{R}^X) \rightarrow (P(\mathbb{R}))^X$ that comes from the monadic strength. That is, $(q([\alpha, \mu]))(x) = [\lambda r. \alpha(r)(x), \mu]$. We show that q is a quotient morphism.

We first show that q is surjective. To give a morphism $k: (X, M_X) \rightarrow P(\mathbb{R})$ is to give a measurable function $(X, \Sigma_{M_X}) \rightarrow G(\mathbb{R})$, since $(P(\mathbb{R}), M_{P(\mathbb{R})}) \cong (G(\mathbb{R}), M_{\Sigma_{G(\mathbb{R})}})$ (Prop. 22(4)) and by using the adjunction between measurable spaces and quasi-Borel spaces (Prop. 15(1)). Directly, we understand a morphism $k: (X, M_X) \rightarrow P(\mathbb{R})$ as the kernel $k^\sharp: X \times \Sigma_{\mathbb{R}} \rightarrow [0, 1]$ with $k^\sharp(x, U) \stackrel{\text{def}}{=} \mu_x(\alpha_x^{-1}(U))$ whenever $k(x) = [\alpha_x, \mu_x]$. The definition of k^\sharp does not depend on the choice of α_x, μ_x .

Now we can use the randomization lemma (Prop. 24) to find a measurable function $f_{k^\sharp}: \mathbb{R} \times X \rightarrow \mathbb{R}$ such that $k^\sharp(x, U) = v\{r \mid f_{k^\sharp}(r, x) \in U\}$. In general, if a function $Y \times X \rightarrow Z$ is jointly measurable then it is also a morphism from the product quasi-Borel space. So f_{k^\sharp} is a morphism, and we can form $(\text{curry } f_{k^\sharp}): \mathbb{R} \rightarrow \mathbb{R}^X$. So,

$$\begin{aligned} q([\text{curry } f_{k^\sharp}, v])(x) &= [\lambda r. \text{curry } f_{k^\sharp}(r)(x), v] \\ &= [\lambda r. f_{k^\sharp}(r, x), v] = k(x), \end{aligned}$$

and q is surjective, as required.

Finally we show that $M_{(P(\mathbb{R}))^X} = \{q \circ \alpha \mid \alpha \in M_{P(\mathbb{R}^X)}\}$. We have (\supseteq) since q is a morphism, so it remains to show (\subseteq) . Consider $\beta \in M_{(P(\mathbb{R}))^X}$. We must show that $\beta = q \circ \alpha$ for some $\alpha \in M_{P(\mathbb{R}^X)}$. By Prop. 18, $\beta \in M_{(P(\mathbb{R}))^X}$ means the uncurried function $(\text{uncurry } \beta): \mathbb{R} \times X \rightarrow P(\mathbb{R})$ is a morphism. As above, this morphism corresponds to a kernel $(\text{uncurry } \beta)^\sharp: (\mathbb{R} \times X) \times \Sigma_{\mathbb{R}} \rightarrow [0, 1]$. The randomization lemma (Prop. 24) gives a measurable function $f_\beta: \mathbb{R} \times (\mathbb{R} \times X) \rightarrow \mathbb{R}$ such that $(\text{uncurry } \beta)^\sharp((r, x), U) = v\{s \mid f_\beta(s, (r, x)) \in U\}$. By Prop. 15(1) and the fact that the σ -algebra of a product quasi-Borel space $\mathbb{R} \times (\mathbb{R} \times X)$ includes the product σ -algebras $\Sigma_{\mathbb{R}} \otimes \Sigma_{M_{\mathbb{R} \times X}}$, this function f_β is also a morphism. Define $\gamma: \mathbb{R} \rightarrow (\mathbb{R}^X)^\mathbb{R}$ by

$\gamma = \lambda r. \lambda s. \lambda x. f_\beta(s, (r, x))$. This is a morphism since we can interpret λ -calculus in a cartesian closed category. Define $\alpha: \mathbb{R} \rightarrow P(\mathbb{R}^X)$ by $\alpha(r) = [\gamma(r), v]$; this function is in $M_{P(\mathbb{R}^X)}$ by Lemma 27. A direct calculation now gives $\beta = q \circ \alpha$, as required. \square

VIII. DE FINETTI'S THEOREM

De Finetti's theorem [8] is one of the foundational results in Bayesian statistics. It says that every exchangeable sequence of random observations on \mathbb{R} or another well-behaved measurable space can be modeled accurately by the following two-step process: first choose a probability measure on \mathbb{R} randomly (according to some distribution on probability measures) and then generate a sequence with independent samples from this measure. Limiting observations to values in a well-behaved space like \mathbb{R} in the theorem is important: Dubins and Freedman proved that the theorem fails for a general measurable space [9].

In this section, we show that a version of de Finetti's theorem holds for all quasi-Borel spaces, not just \mathbb{R} . Our result does not contradict Dubins and Freedman's obstruction; probability measures on quasi-Borel spaces may only use \mathbb{R} as their source of randomness, whereas those on measurable spaces are allowed to use any measurable space for the same purpose. As we will show shortly, this careful choice of random source lets us generalize key arguments in a proof of de Finetti's theorem [2] to quasi-Borel spaces.

Let (X, M_X) be a quasi-Borel space and (X^n, M_{X^n}) the product quasi-Borel space $\prod_{i=1}^n X$ for each positive integer n . Recall that $P(X)$ consists of equivalence classes $[\beta, \nu]$ of probability measures (β, ν) on X . For $n \geq 1$, define a morphism $\text{iid}_n: P(X) \rightarrow P(X^n)$ by

$$\text{iid}_n([\beta, \nu]) = [(\prod_{i=1}^n \beta \circ \iota_n), ((\iota_n^{-1})_* \otimes_{i=1}^n \nu)]$$

where ι_n is a measurable isomorphism $\mathbb{R} \rightarrow \prod_{i=1}^n \mathbb{R}$, and $\otimes_{i=1}^n \nu$ is the product measure formed by n copies of ν . The name iid_n represents 'independent and identically distributed'. Indeed, iid_n transforms a probability measure (β, ν) on X to the measure of the random sequence in X^n that independently samples from (β, ν) . The function iid_n is a morphism $P(X) \rightarrow P(X^n)$ because it can also be written in terms of the strength of the monad P .

Write (X^ω, M_{X^ω}) for the countable product $\prod_{i=1}^\infty X$.

Definition 28. A probability measure (α, μ) on X^ω is exchangeable if for all permutations π on positive integers, $[\alpha, \mu] = [\alpha_\pi, \mu]$, where $\alpha_\pi(r)_i \stackrel{\text{def}}{=} \alpha(r)_{\pi(i)}$ for all r and i .

Theorem 29 (Weak de Finetti for quasi-Borel spaces). If (α, μ) is an exchangeable probability measure on X^ω , then there exists a probability measure (β, ν) in $P(P(X))$ such that for all $n \geq 1$, the measure $[(\beta, \nu) \gg= \text{iid}_n]$ on $P(X^n)$ equals $P((-)_{1\dots n})(\alpha, \mu)$ when considered as a measure on the product measurable space $(X^n, \bigotimes_{i=1}^n \Sigma_{M_X})$. (Here $(-)_{1\dots n}: X^\omega \rightarrow X^n$ is $(x)_{1\dots n} \stackrel{\text{def}}{=} (x_1, \dots, x_n)$.)

In the theorem, (β, ν) represents a random variable that has a probability measure on X as its value. The theorem says that (every finite prefix of) a sample sequence from (α, μ) can be generated by first sampling a probability measure on X according to (β, ν) , then generating independent X -valued samples from the measure, and finally forming a sequence with these samples.

We call the theorem *weak* for two reasons. First, the σ -algebra $\Sigma_{M_X^n}$ includes the product σ -algebra $\bigotimes_{i=1}^n \Sigma_{M_X}$, but we do not know that they are equal; two different probability measures in $P(X^n)$ may induce the same measure on $(X^n, \bigotimes_{i=1}^n \Sigma_{M_X})$, although they always induce different measures on $(X^n, \Sigma_{M_X^n})$. In the theorem, we equate such measures, which lets us use a standard technique for proving the equality of measures on product σ -algebras. Second, we are unable to construct a version of iid_n for infinite sequences, i.e. a morphism $P(X) \rightarrow P(X^\omega)$ implementing the independent identically-distributed random sequence. The theorem is stated only for finite prefixes.

The rest of this section provides an overview of our proof of Theorem 29. The starting point is to unpack definitions in the theorem, especially those related to quasi-Borel spaces, and to rewrite the statement of the theorem purely in terms of standard measure-theoretic notions.

Lemma 30. *Let (α, μ) be an exchangeable probability measure on X^ω . Then, the conclusion of Theorem 29 holds if and only if there exist a probability measure $\xi \in G(\mathbb{R})$, a measurable function $k: \mathbb{R} \rightarrow G(\mathbb{R})$, and $\gamma \in M_X$ such that for all $n \geq 1$ and all $U_1, \dots, U_n \in \Sigma_{M_X}$,*

$$\begin{aligned} \int_{r \in \mathbb{R}} \left(\prod_{i=1}^n [\alpha(r)_i \in U_i] \right) d\mu \\ = \int_{r \in \mathbb{R}} \prod_{i=1}^n \left(\int_{s \in \mathbb{R}} [\gamma(s) \in U_i] d(k(r)) \right) d\xi. \end{aligned}$$

Here we express the domain of integration and the integrated variable explicitly to avoid confusion.

Proof. Let (α, μ) be an exchangeable probability measure on X^ω . We unpack definitions in the conclusion of Theorem 29. The first definition to unpack is the notion of probability measure in $P(P(X))$. Here are the crucial facts that enable this unpacking. First, for every probability measure (β, ν) on $P(X)$, there exist a function $\gamma: \mathbb{R} \rightarrow X$ in M_X and a measurable $k: \mathbb{R} \rightarrow G(\mathbb{R})$ such that $\beta(r) = [\gamma, k(r)]$ for all $r \in \mathbb{R}$. Second, conversely, for a function $\gamma \in M_X$, a measurable $k: \mathbb{R} \rightarrow G(\mathbb{R})$, and a probability measure $\nu \in G(\mathbb{R})$, the function $(\lambda r. [\gamma, k(r)], \nu)$ is a probability measure in $P(P(X))$. Thus, we can look for (γ, k, ν) in the conclusion of the theorem instead of (β, ν) .

The second is the definition of $[\beta, \nu] \gg= \text{iid}_n$. Using (γ, k, ν) instead of (β, ν) , we find that $[\beta, \nu] \gg= \text{iid}_n$ is the measure $[(\prod_{i=1}^n \gamma) \circ \iota_n, (\iota_n^{-1})_* (\nu \gg= \lambda r. \bigotimes_{i=1}^n k(r))]$.

Recall that two measures p and q on the product space $(X^n, \bigotimes_{i=1}^n X)$ are equivalent when $p(U_1 \times \dots \times U_n)$ equals

$q(U_1 \times \dots \times U_n)$ for all $U_1, \dots, U_n \in \Sigma_{M_X}$. Thus we must show that $((-)_1 \dots n \circ \alpha)_* \mu(U_1 \times \dots \times U_n)$ is equal to $((\prod_{i=1}^n \gamma)_* (\nu \gg= \lambda r. \bigotimes_{i=1}^n k(r)))(U_1 \times \dots \times U_n)$. This equation is equivalent to the one in the statement of the lemma with $\xi = \nu$. \square

Thus we just need to show how to construct ξ , k and γ in Lemma 30 from a given exchangeable probability measure (α, μ) on X^ω . Constructing ξ and γ is easy:

$$\xi \stackrel{\text{def}}{=} \mu, \quad \gamma \stackrel{\text{def}}{=} \lambda r. \alpha(r)_1.$$

Note that these definitions type-check: $\xi = \mu \in G(\mathbb{R})$, and $\gamma \in M_X$ because $\alpha \in M_{X^\omega}$ and the first projection $(-)_1$ is a morphism $X^\omega \rightarrow X$.

Constructing k is not that easy. We need to use the fact that μ is defined over \mathbb{R} , a standard Borel space. This fact itself holds because all probability measures on quasi-Borel spaces use \mathbb{R} as their source of randomness. Define measurable functions $\alpha_e, \alpha_o: (\mathbb{R}, \Sigma_{\mathbb{R}}) \rightarrow (X^\omega, \Sigma_{M_{X^\omega}})$ by

$$\alpha_e(r)_i \stackrel{\text{def}}{=} \alpha(r)_{2i} \quad (\text{even}), \quad \alpha_o(r)_i \stackrel{\text{def}}{=} \alpha(r)_{2i-1} \quad (\text{odd}).$$

Since μ is a probability measure on \mathbb{R} , there exists a measurable function $k': (X^\omega, \Sigma_{M_{X^\omega}}) \rightarrow (G(\mathbb{R}), \Sigma_{G(\mathbb{R})})$, called a conditional probability kernel, such that for all measurable $f: \mathbb{R} \rightarrow \mathbb{R}$ and $U \in (\alpha_e)^{-1}(\Sigma_{M_{X^\omega}})$,

$$\int_{r \in U} f(r) d\mu = \int_{r \in U} \left(\int_{\mathbb{R}} f d((k' \circ \alpha_e)(r)) \right) d\mu. \quad (7)$$

Define $k \stackrel{\text{def}}{=} k' \circ \alpha_e$.

Our ξ , k and γ satisfy the requirement in Lemma 30 because of the following three properties, which follow from exchangeability of (α, μ) .

Lemma 31. *For all $n \geq 1$ and all $U_1, \dots, U_n \in \Sigma_{M_X}$,*

$$\int_{r \in \mathbb{R}} \left(\prod_{i=1}^n [\alpha(r)_i \in U_i] \right) d\mu = \int_{r \in \mathbb{R}} \left(\prod_{i=1}^n [\alpha_o(r)_i \in U_i] \right) d\mu.$$

Proof. Consider $n \geq 1$ and $U_1, \dots, U_n \in \Sigma_{M_X}$. Pick a permutation π on positive integers such that $\pi(i) = 2i - 1$ for all integers $1 \leq i \leq n$. Then, $[\alpha, \mu] = [\alpha_\pi, \mu]$ by the exchangeability of (α, μ) . Thus

$$\int_{r \in \mathbb{R}} \left(\prod_{i=1}^n [\alpha(r)_i \in U_i] \right) d\mu = \int_{r \in \mathbb{R}} \left(\prod_{i=1}^n [\alpha_\pi(r)_i \in U_i] \right) d\mu,$$

from which the statement follows. \square

Lemma 32. *For all $U \in \Sigma_{M_X}$ and all $i, j \geq 1$,*

$$\int_{s \in \mathbb{R}} [\alpha_o(s)_i \in U] d(k(r)) = \int_{s \in \mathbb{R}} [\alpha_o(s)_j \in U] d(k(r))$$

holds for μ -almost all $r \in \mathbb{R}$.

Proof. Consider a measurable set $U \in \Sigma_{M_X}$ and $i, j \geq 1$. The function $\lambda r. \int_{s \in \mathbb{R}} [\alpha_o(s)_i \in U] d(k(r)) : \mathbb{R} \rightarrow \mathbb{R}$ is a conditional expectation of the indicator function $\lambda s. [\alpha_o(s)_i \in U]$ with respect to the probability measure μ and the σ -algebra

generated by the measurable function $\alpha_e: \mathbb{R} \rightarrow (X^\omega, \Sigma_{M_{X^\omega}})$. By the almost-sure uniqueness of conditional expectation, it suffices to show that $\lambda r. \int_{s \in \mathbb{R}} [\alpha_o(s)_j \in U] d(k(r))$ is also a conditional expectation of $\lambda s. [\alpha_o(s)_i \in U]$ with respect to μ and α_e . Pick a measurable subset $V \in \Sigma_{M_{X^\omega}}$. Then:

$$\begin{aligned} & \int_{r \in \mathbb{R}} [\alpha_e(r) \in V] \cdot \left(\int_{s \in \mathbb{R}} [\alpha_o(s)_j \in U] d(k(r)) \right) d\mu \\ &= \int_{r \in \mathbb{R}} [\alpha_o(r)_j \in U \wedge \alpha_e(r) \in V] d\mu \\ &= \int_{r \in \mathbb{R}} [\alpha_o(r)_i \in U \wedge \alpha_e(r) \in V] d\mu. \end{aligned}$$

The first equation holds because the function $\lambda r. \int_{s \in \mathbb{R}} [\alpha_o(s)_j \in U] d(k(r))$ is a conditional expectation of $\lambda s. [\alpha_o(s)_j \in U]$ with respect to μ and α_e . The second equation follows from the exchangeability of (α, μ) . We have just shown that $\lambda r. \int_{s \in \mathbb{R}} [\alpha_o(s)_j \in U] d(k(r))$ is a conditional expectation of $\lambda s. [\alpha_o(s)_i \in U]$ with respect to μ and α_e . \square

Lemma 33. For all $n \geq 1$ and all $U_1, \dots, U_n \in \Sigma_{M_X}$,

$$\begin{aligned} & \int_{s \in \mathbb{R}} \left(\prod_{i=1}^n [\alpha_o(s)_i \in U_i] \right) d(k(r)) \\ &= \prod_{i=1}^n \int_{s \in \mathbb{R}} [\alpha_o(s)_i \in U_i] d(k(r)) \end{aligned}$$

holds for μ -almost all $r \in \mathbb{R}$.

Proof notes. Use induction on $n \geq 1$. There is nothing to prove for the base case $n = 1$. To handle the inductive case, assume that $n > 1$. Let U_1, \dots, U_n be subsets in Σ_{M_X} . Define a function $\alpha': \mathbb{R} \rightarrow X^\omega$ as follows:

$$\alpha'(r)_i = \begin{cases} \alpha_o(r)_i & \text{if } 1 \leq i \leq n-1 \\ \alpha_e(r)_{i-n+1} & \text{otherwise.} \end{cases}$$

Then, α' is in M_{X^ω} , so that α' is a measurable function $(\mathbb{R}, \Sigma_{\mathbb{R}}) \rightarrow (X^\omega, \Sigma_{M_{X^\omega}})$. Thus there exists a measurable $k'_0: (X^\omega, \Sigma_{M_{X^\omega}}) \rightarrow (G(\mathbb{R}), \Sigma_{G(\mathbb{R})})$, the conditional probability kernel, such that for all measurable functions $f: \mathbb{R} \rightarrow \mathbb{R}$, $\lambda r. \int_{\mathbb{R}} f d((k'_0 \circ \alpha')(r))$ is a conditional expectation of f with respect to μ and the σ -algebra generated by α' . Define $k': \mathbb{R} \rightarrow G(\mathbb{R}) = k'_0 \circ \alpha'$. Then k' is measurable because so are k'_0 and α' . More importantly, for μ -almost all $r \in \mathbb{R}$,

$$\int_{s \in \mathbb{R}} [\alpha_o(s)_n \in U_n] d(k(r)) = \int_{s \in \mathbb{R}} [\alpha_o(s)_n \in U_n] d(k'(r)). \quad (8)$$

The proof of this equality appears in the full version of this paper.

Recall that $k = k_0 \circ \alpha_e$ and $k' = k'_0 \circ \alpha'$ are defined in terms of conditional expectation. Thus, they inherit all the properties of conditional expectation. In particular, for μ -almost all $r \in \mathbb{R}$ and all measurable $h: \mathbb{R} \rightarrow \mathbb{R}$,

$$\begin{aligned} & \int_{s \in \mathbb{R}} \prod_{i=1}^n [\alpha_o(s)_i \in U_i] d(k(r)) \\ &= \int_{s \in \mathbb{R}} \left(\int_{t \in \mathbb{R}} \prod_{i=1}^n [\alpha_o(t)_i \in U_i] d(k'(s)) \right) d(k(r)), \end{aligned} \quad (9)$$

$$\begin{aligned} & \int_{s \in \mathbb{R}} \prod_{i=1}^n [\alpha_o(s)_i \in U_i] d(k'(r)) \\ &= \prod_{i=1}^{n-1} [\alpha_o(r)_i \in U_i] \cdot \int_{s \in \mathbb{R}} [\alpha_o(s)_n \in U_n] d(k'(r)), \end{aligned} \quad (10)$$

$$\begin{aligned} & \int_{s \in \mathbb{R}} \left(h(s) \cdot \int_{t \in \mathbb{R}} [\alpha_o(t)_n \in U_n] d(k(s)) \right) d(k(r)) \\ &= \left(\int_{t \in \mathbb{R}} [\alpha_o(t)_n \in U_n] d(k(r)) \right) \cdot \left(\int_{s \in \mathbb{R}} h(s) d(k(r)) \right). \end{aligned} \quad (11)$$

Using the assumption (8) and the properties (9), (10) and (11), we complete the proof of the inductive case as follows: for all subsets $V \in (\alpha_e)^{-1}(\Sigma_{M_{X^\omega}})$,

$$\begin{aligned} & \int_{r \in V} \int_{s \in \mathbb{R}} \prod_{i=1}^n [\alpha_o(s)_i \in U_i] d(k(r)) d\mu \\ &= \int_{r \in V} \int_{s \in \mathbb{R}} \int_{t \in \mathbb{R}} \prod_{i=1}^n [\alpha_o(t)_i \in U_i] d(k'(s)) d(k(r)) d\mu \\ &= \int_{r \in V} \int_{s \in \mathbb{R}} \prod_{i=1}^{n-1} [\alpha_o(s)_i \in U_i] \\ & \quad \cdot \int_{t \in \mathbb{R}} [\alpha_o(t)_n \in U_n] d(k'(s)) d(k(r)) d\mu \\ &= \int_{r \in V} \int_{s \in \mathbb{R}} \prod_{i=1}^{n-1} [\alpha_o(s)_i \in U_i] \\ & \quad \cdot \int_{t \in \mathbb{R}} [\alpha_o(t)_n \in U_n] d(k(s)) d(k(r)) d\mu \\ &= \int_{r \in V} \left(\int_{t \in \mathbb{R}} [\alpha_o(t)_n \in U_n] d(k(r)) \right) \\ & \quad \cdot \left(\int_{s \in \mathbb{R}} \prod_{i=1}^{n-1} [\alpha_o(s)_i \in U_i] d(k(r)) \right) d\mu \\ &= \int_{r \in V} \prod_{i=1}^n \int_{s \in \mathbb{R}} [\alpha_o(s)_i \in U_i] d(k(r)) d\mu. \end{aligned}$$

The first and the second equalities hold because of (9) and (10). The third equality uses (8), and the fourth the equality in (11). The fifth follows from the induction hypothesis. Our derivation implies that both $\lambda r. \int_{s \in \mathbb{R}} \prod_{i=1}^n [\alpha_o(s)_i \in U_i] d(k(r))$ and $\lambda r. \prod_{i=1}^n \int_{s \in \mathbb{R}} [\alpha_o(s)_i \in U_i] d(k(r))$ are conditional expectations of the same function with respect to μ and the same σ -algebra. So, they are equal for μ -almost all inputs r . \square

The following calculation combines these lemmas and shows that ξ , k and γ satisfy the requirement in Lemma 30:

$$\begin{aligned} & \int_{r \in \mathbb{R}} \prod_{i=1}^n [\alpha(r)_i \in U_i] d\mu \\ &= \int_{r \in \mathbb{R}} \prod_{i=1}^n [\alpha_o(r)_i \in U_i] d\mu \end{aligned} \quad \text{Lem. 31}$$

$$\begin{aligned}
&= \int_{r \in \mathbb{R}} \left(\int_{s \in \mathbb{R}} \prod_{i=1}^n [\alpha_o(s)_i \in U_i] d(k(r)) \right) d\mu && \text{Eq. (7)} \\
&= \int_{r \in \mathbb{R}} \prod_{i=1}^n \left(\int_{s \in \mathbb{R}} [\alpha_o(s)_i \in U_i] d(k(r)) \right) d\mu && \text{Lem. 33} \\
&= \int_{r \in \mathbb{R}} \prod_{i=1}^n \left(\int_{s \in \mathbb{R}} [\alpha_o(s)_1 \in U_i] d(k(r)) \right) d\mu && \text{Lem. 32} \\
&= \int_{r \in \mathbb{R}} \prod_{i=1}^n \left(\int_{s \in \mathbb{R}} [\gamma(s) \in U_i] d(k(r)) \right) d\xi && \text{Def. of } \gamma, \xi.
\end{aligned}$$

This concludes our proof outline for Theorem 29.

IX. RELATED WORK

A. Quasi-topological spaces and categories of functors

Our development of a cartesian closed category from measurable spaces mirrors the development of cartesian closed categories of topological spaces over the years.

For example, quasi-Borel spaces are reminiscent of *subsequential spaces* [20]: a set X together with a collection of functions $Q \subseteq [\mathbb{N} \cup \{\infty\} \rightarrow X]$ satisfying some conditions. The functions in Q are thought of as convergent sequences. Another notion of generalized topological space is *C-space* [38]: a set X together with a collection $Q \subseteq [2^{\mathbb{N}} \rightarrow X]$ of ‘probes’ satisfying some conditions; this is a variation on Spanier’s early notion of *quasi-topological space* [33]. Another reminiscent notion in the context of differential geometry is a *diffeological space* [3]: a set X together with a set $Q_U \subseteq [U \rightarrow X]$ of ‘plots’ for each open subset U of \mathbb{R}^n satisfying some conditions. These examples all form cartesian closed categories.

A common pattern is that these spaces can be understood as extensional (concrete) sheaves on an established category of spaces. Let \mathbf{SMeas} be the category of standard Borel spaces and measurable functions. There is a functor $J: \mathbf{QBS} \rightarrow [\mathbf{SMeas}^{\text{op}}, \mathbf{Set}]$ with $(J(X, M_X))(Y, \Sigma_Y) \stackrel{\text{def}}{=} \mathbf{QBS}((Y, M_{\Sigma_Y}), (X, M_X))$, which is full and faithful by Prop. 15(2). We can characterize those functors that arise in this way.

Proposition 34. *Let $F: \mathbf{SMeas}^{\text{op}} \rightarrow \mathbf{Set}$ be a functor. The following are equivalent:*

- F is naturally isomorphic to $J(X, M_X)$, for some quasi-Borel space (X, M_X) ;
- F preserves countable products and F is extensional: the functions $i_{(X, \Sigma_X)}: F(X, \Sigma_X) \rightarrow \mathbf{Set}(X, F(1))$ are injective, where $(i_{(X, \Sigma_X)}(\xi))(x) = (F(\lceil x \rceil))(\xi)$, and we consider $x \in X$ as a function $\lceil x \rceil: 1 \rightarrow X$.

There are similar characterizations of subsequential spaces [20], quasi-topological spaces [10] and diffeological spaces [3]. Prop. 34 is an instance of a general pattern (e.g. [3], [10]); but that is not to say that the definition of quasi-Borel space (Def. 7) arises automatically. The method of extensional presheaves also arises in other models of computation such as finiteness spaces [11] and realizability models [30]. This work appears to be the first application to probability theory,

although via Prop. 34 there are connections to Simpson’s probability sheaves [32].

The characterization of Prop. 34 gives a canonical categorical status to quasi-Borel spaces. It also connects with our earlier work [34], which used the cartesian closed category of countable-product-preserving functors in $[\mathbf{SMeas}^{\text{op}}, \mathbf{Set}]$. Quasi-Borel spaces have several advantages over this functor category. For one thing, they are more concrete, leading to better intuitions for their constructions. For example, measures in [34] are built abstractly from left Kan extensions, whereas for quasi-Borel spaces they have a straightforward concrete definition (Def. 10). For another thing, in contrast to the functor category in [34], quasi-Borel spaces form a well-pointed category: if two morphisms $(X, M_X) \rightarrow (Y, M_Y)$ are different then they disagree on some point in X . From the perspective of semantics of programming languages, where terms in context $\Gamma \vdash t : A$ are interpreted as morphisms $[\![t]\!] : [\![\Gamma]\!] \rightarrow [\![A]\!]$, well-pointedness is a crucial property. It says that if two open terms are different, $[\![t]\!] \neq [\![u]\!] : [\![\Gamma]\!] \rightarrow [\![A]\!]$, then there is a ground context $C: 1 \rightarrow [\![\Gamma]\!]$ that distinguishes them: $[\![C[t]]!] \neq [\![C[u]]!] : 1 \rightarrow [\![A]\!]$.

Quasi-Borel spaces add objects to make the category of measurable spaces cartesian closed. Another interesting future direction is to add morphisms to make more objects isomorphic, and so find a cartesian closed subcategory [35].

B. Domains and valuations

In this paper our starting point has been the standard foundation for probability theory, based on σ -algebras and probability measures. An alternative foundation for probability is based on topologies and valuations. An advantage of our starting point is that we can reference the canon of work on probability theory. Having said this, an advantage to the approach based on valuations is that it is related to domain theoretic methods, which have already been used to give semantics to programming languages.

Jones and Plotkin [21] showed that valuations form a monad which is analogous to our probability monad. However, there is considerable debate about which cartesian closed category this monad should be based on (e.g. [22], [17]). For a discussion of the concerns in the context of programming languages, see e.g. [13]. One recent proposal is to use Girard’s probabilistic coherence spaces [12]. Another is to use a topological domain theory as a cartesian closed category for analysis and probability ([5], [27], [19]).

Concerns about probabilistic powerdomains have led instead to domains of random variables (e.g. [25], [4], [31]). We cannot yet connect formally with this work, but there are many intuitive links. For example, our measures on quasi-Borel spaces (Def. 10) are reminiscent of continuous random variables on a dcpo.

An additional advantage of a domain theoretic approach is that it naturally supports recursion. We are currently investigating a notion of ‘ordered quasi-Borel space’, by enriching Prop. 34 over dcpo’s.

C. Other related work

Our work is related to two recent semantic studies on probabilistic programming languages. The first is Borgström et al.'s *operational* (not denotational as in this paper) semantics for a higher-order probabilistic programming language with continuous distributions [6], which has been used to justify a basic inference algorithm for the language. Recently, Culpepper and Cobb refined this operational approach using logical relations [7]. The second study is Freer and Roy's results on a computable variant of de Finetti's theorem and its implication on exchangeable random processes implemented in higher-order probabilistic programming languages [14]. One interesting future direction is to revisit the results about logical relations and computability in these studies with quasi-Borel spaces, and to see whether they can be extended to spaces other than standard Borel spaces.

X. CONCLUSION

We have shown that quasi-Borel spaces (§III) support higher-order functions (§IV) as well as spaces of probability measures (§V). We have illustrated the power of this new formalism by giving a semantic analysis of Bayesian regression (§VI), by rephrasing the randomization lemma as a quotient-space construction (§VII), and by showing that it supports de Finetti's theorem (§VIII).

ACKNOWLEDGMENT

We thank Radha Jagadeesan and Dexter Kozen for encouraging us to think about a well-pointed cartesian closed category for probability theory, Vincent Danos and Dan Roy for nudging us to work on de Finetti's theorem, Mike Mislove for discussions of quasi-Borel spaces, and Martin Escardó for explaining C-spaces, and Alex Simpson for detailed report with many suggestions. This research was supported by a Royal Society Research Fellowship and EPSRC grants EP/L002388/2 and EP/N007387/1, and also by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (No.R0190-16-2011, Development of Vulnerability Discovery Technologies for IoT Software Security).

REFERENCES

- [1] R. J. Aumann, "Borel structures for function spaces," *Illinois Journal of Mathematics*, vol. 5, pp. 614–630, 1961.
- [2] T. Austin, "Exchangeable random arrays," 2013. [Online]. Available: <https://cims.nyu.edu/~tim/ExchnotesforIISc.pdf>
- [3] J. C. Baez and A. E. Hoffnung, "Convenient categories of smooth spaces," *Trans. Amer. Math. Soc.*, vol. 363, 2011.
- [4] T. Barker, "A monad for randomized algorithms," in *Proc. MFPS*, 2016, pp. 47–62.
- [5] I. Battenfeld, M. Schröder, and A. Simpson, "A convenient category of domains," ser. ENTCS, vol. 172, 2007.
- [6] J. Borgström, U. Dal Lago, A. D. Gordon, and M. Szymczak, "A lambda-calculus foundation for universal probabilistic programming," in *Proc. ICFP*, 2016, pp. 33–46.
- [7] R. Culpepper and A. Cobb, "Contextual equivalence for probabilistic programs with continuous random variables and scoring," in *Proc. ESOP*, 2017.
- [8] B. de Finetti, "La prévision : ses lois logiques, ses sources subjectives," *Annales de l'institut Henri Poincaré*, vol. 7, 1937.
- [9] L. E. Dubins and D. A. Freedman, "Exchangeable processes need not be mixtures of independent, identically distributed random variables," *Z. Angew. Math. Mech.*, vol. 48, 1979.
- [10] E. J. Dubuc, "Concrete quasitopoi," in *Applications of sheaves*, ser. Lect. Notes Math. Springer, 1977, vol. 753, pp. 239–254.
- [11] T. Ehrhard, "On finiteness spaces and extensional presheaves over the Lawvere theory of polynomials," *J. Pure Appl. Algebra*, 2007, to appear.
- [12] T. Ehrhard, C. Tasson, and M. Pagani, "Probabilistic coherence spaces are fully abstract for probabilistic PCF," in *POPL 2014*.
- [13] M. H. Escardó, "Semi-decidability of may, must and probabilistic testing in a higher-type setting," in *Proc. MFPS*, 2009.
- [14] C. E. Freer and D. M. Roy, "Computable de Finetti measures," *Ann. Pure Appl. Logic*, vol. 163, no. 5, pp. 530–546, 2012.
- [15] M. Giry, "A categorical approach to probability theory," in *Categorical Aspects of Topology and Analysis*, 1982, pp. 68–85.
- [16] N. Goodman, V. Mansinghka, D. M. Roy, K. Bonawitz, and J. B. Tenenbaum, "Church: a language for generative models," in *UAI*, 2008.
- [17] J. Goubault-Larrecq, " ω QRB-domains and the probabilistic powerdomain," in *Proc. LICS*, 2010, pp. 352–361.
- [18] E. Hewitt and L. J. Savage, "Symmetric measures on cartesian products," *Trans. Amer. Math. Soc.*, vol. 80, pp. 470–501, 1955.
- [19] D. Huang and G. Morrisett, "An application of computable distributions to the semantics of probabilistic programs: part 2." [Online]. Available: <http://pps2017.soic.indiana.edu/files/2016/12/comp-dist-sem.pdf>
- [20] P. Johnstone, "On a topological topos," *Proc. London Math. Soc.*, vol. 3, no. 38, pp. 237–271, 1979.
- [21] C. Jones and G. D. Plotkin, "A probabilistic powerdomain of evaluations," in *Proc. LICS*, 1989, pp. 186–195.
- [22] A. Jung and R. Tix, "The troublesome probabilistic powerdomain," ser. ENTCS, vol. 13, 1998, pp. 70–91.
- [23] O. Kallenberg, *Foundations of Modern Probability*, 2nd ed. Springer, 2002.
- [24] V. K. Mansinghka, D. Selsam, and Y. N. Perov, "Venture: a higher-order probabilistic programming platform with programmable inference," *arXiv:1404.0099*, 2014.
- [25] M. W. Mislove, "Anatomy of a domain of continuous random variables I," *Theor. Comput. Sci.*, vol. 546, pp. 176–187, 2014.
- [26] E. Moggi, "Notions of computation and monads," *Inf. Comput.*, vol. 93, no. 1, pp. 55–92, 1991.
- [27] M. Pape and T. Streicher, "Computability in basic quantum mechanics," 2016. [Online]. Available: <https://arxiv.org/abs/1610.09209>
- [28] C. Preston, "Some notes on standard Borel and related spaces," 2008. [Online]. Available: <https://arxiv.org/abs/0809.3066>
- [29] J. C. Reynolds, "Definitional interpreters for higher-order programming languages," in *Proc. ACM Conference*, 1972.
- [30] G. Rosolini and T. Streicher, "Comparing models of higher type computation," ser. ENTCS, vol. 23, 1999.
- [31] D. S. Scott, "Stochastic λ -calculi: An extended abstract," *J. Applied Logic*, vol. 12, no. 3, pp. 369–376, 2014.
- [32] A. Simpson, "Probability sheaves," 2016, CIPPMI. [Online]. Available: <https://synapse.math.univ-toulouse.fr/index.php/s/QWrxKeXn31mN3gz>
- [33] E. Spanier, "Quasi-topologies," *Duke Math. J.*, pp. 1–14, 1963.
- [34] S. Staton, H. Yang, C. Heunen, O. Kammar, and F. Wood, "Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints," in *LICS*, 2016.
- [35] N. E. Steenrod, "A convenient category of topological spaces," *Michigan Mathematical Journal*, vol. 14, pp. 133–152, 1967.
- [36] R. Street, "The formal theory of monads," *J. Pure Appl. Algebra*, vol. 2, pp. 149–168, 1972.
- [37] F. Wood, J. W. van de Meent, and V. Mansinghka, "A new approach to probabilistic programming inference," in *Proc. AISTATS*, 2014.
- [38] C. Xu and M. Escardó, "A constructive model of uniform continuity," in *Proc. TLCA*, 2013.