

Joint Distributions in Probabilistic Semantics

Dexter Kozen Alexandra Silva

^a *Computer Science Department
Cornell University
Ithaca, New York 14853-7501, USA*

Erik Voogd

^b *Department of Informatics
University of Oslo
Gaustadalléen 23B, 0373 Oslo, Norway*

Abstract

Various categories have been proposed as targets for the denotational semantics of higher-order probabilistic programming languages. One such proposal involves joint probability distributions (couplings) used in Bayesian statistical models with conditioning. In previous treatments, composition of joint measures was performed by disintegrating to obtain Markov kernels, composing the kernels, then reintegrating to obtain a joint measure. Disintegrations exist only under certain restrictions on the underlying spaces. In this paper we propose a category whose morphisms are joint finite measures in which composition is defined without reference to disintegration, allowing its application to a broader class of spaces. The category is symmetric monoidal with a pleasing symmetry in which the dagger structure is a simple transpose.

Keywords: probabilistic programming, disintegration, Bayesian inference, conditioning

1 Introduction

Bayesian inference and *conditioning* are important tools in probabilistic programming. Modern probabilistic languages for machine learning, e.g. Church [8] and Anglican [14], generally incorporate these tools in some form. To formalize the semantics of such languages, several authors have proposed categories for modeling Bayesian inference and conditioning via *disintegration*.

Approaches based on disintegration have the disadvantage that they can only be applied to certain spaces. At the core of the matter is the *disintegration theorem* which only holds under certain assumptions—usually that the spaces are standard Borel, as in e.g. [10,2,1,6,11,5]. This theorem connects joint distributions and Markov Kernels. A Markov kernel P gives rise to a joint distribution $\mathbf{J}P$ on $X \times Y$ with marginals μ and ν in a natural way. The disintegration theorem says that this construction can be inverted: for any joint distribution θ on $X \times Y$ with marginals μ and ν , there is a kernel P , unique up to a μ -nullset, such that $\theta = \mathbf{J}P$. The inversion can be explored in defining composition and proving properties of the underlying category; we will give two examples below.

In this paper, we want to show that the category \mathbf{JDist} whose objects are measure spaces (X, \mathcal{A}, μ) and whose morphisms $(X, \mathcal{A}, \mu) \rightarrow (Y, \mathcal{B}, \nu)$ are joint distributions θ on $X \times Y$ with marginals μ and ν can be taken as first-class citizen to formalize semantics of probabilistic languages. In particular, we will show that composition in \mathbf{JDist} can be defined without having to use disintegration, thereby making this category more generally applicable than previous approaches.

One recent proposal by Dahlqvist et al. [5] was the category \mathbf{Krn} , whose objects are probability spaces (X, \mathcal{A}, μ) with \mathcal{A} a standard Borel space. The morphisms $(X, \mathcal{A}, \mu) \rightarrow (Y, \mathcal{B}, \nu)$ are equivalence classes modulo μ -nullsets of Markov kernels $P: X \rightarrow Y$ such that

$$\nu(B) = \int_{s \in X} P(s, B) d\mu(s).$$

Dahlqvist et al. [5] use disintegration to show that \mathbf{Krn} is a dagger category with involutive functor $\dagger: \mathbf{Krn} \rightarrow \mathbf{Krn}^{\text{op}}$. Thus $P^\dagger: Y \rightarrow X$ is a kernel in the opposite direction that models Bayesian inference of an input distribution conditioned on output samples. This works even for continuous distributions in which outcomes can occur with zero probability. We will show that the category \mathbf{Krn} can be embedded faithfully in the category \mathbf{JDist} , and this is a symmetric monoidal category with joint distributions as tensors and transpose as symmetry: $\theta^\dagger(A \times B) = \theta(B \times A)$.

Another category was proposed by Abramsky et al. [1] who studied a subcategory of \mathbf{JDist} on Polish spaces called \mathbf{PRel} , using disintegration to define composition: to compose two joint distributions, one disintegrates to obtain Markov kernels, composes the kernels as in \mathbf{Krn} , then reintegrates to obtain a joint measure. This construction can be performed only under various restrictions admitting disintegration [1,6,11,4,3]. The most common assumption is that the spaces are standard Borel. The most general result of this type seems to be that of Culbertson and Sturtz [4] based on work of Faden [7], who assume countably generated spaces but that the measures are perfect.

In this paper, we show that composition in \mathbf{JDist} can be defined independently without reference to disintegration and without any restriction on the underlying spaces. This allows probabilistic programs to be interpreted in a more general class of spaces with a pleasing symmetry in which the notion of equivalence up to a nullset is built in. We will illustrate our approach in defining the composition in \mathbf{JDist} through an example in Section 2. We will then define the composition in \mathbf{JDist} in 3 steps. We will first present the definition in the discrete setting (Section 3), as this is a simpler case which does not require extra measure infrastructure. Second, we will define and prove some results akin to Radon-Nikodým approximants (Section 5). This will provide the necessary results to define composition in \mathbf{JDist} and a functor $\mathbf{J}: \mathbf{Krn} \rightarrow \mathbf{JDist}$ that provides a faithful embedding of categories (Section 6).

2 Illustrative Example

To explain how composition in \mathbf{JDist} circumvents disintegration, we consider a modified version of the example from Dahlqvist et al. [5], pseudocode presented on the right: the goal is to measure the probability that $x > 1$ for $x \sim \mathcal{N}(0, 1)$ given an observation $z > 0.5$ for $z \sim \mathcal{N}(y, 1)$, and $y \sim \mathcal{N}(x, 1)$. That is, x is a standard Gaussian sample (Line 1) and y and z are normally distributed with means x and y respectively, and both with variance one (Lines 2 and 3). Intuitively, Line 4 conditions on the observation that z is greater than 0.5, and Line 5 returns the probability of x being greater than one given this observation.

```

1  x := normal(0, 1);
2  y := normal(x, 1);
3  z := normal(y, 1);
4  observe(z > 0.5);
5  return(x > 1);


```

To illustrate the nature of the category \mathbf{JDist} , we somewhat informally explain the semantics of this program: first in terms of Markov kernels (\mathbf{Krn}), and after that in terms of joint distributions (\mathbf{JDist}).

- In \mathbf{Krn} , the state of the program after Line 1 can be thought of as a measure space (X, \mathcal{A}, μ) with $(X, \mathcal{A}) = (\mathbb{R}, \mathcal{B}_{\mathbb{R}})$ the Borel σ -algebra on \mathbb{R} , and μ the standard Gaussian measure.
- Line 2 constructs a Markov kernel $P: (X, \mathcal{A}) \rightarrow (Y, \mathcal{B})$, where $(Y, \mathcal{B}) = (\mathbb{R}, \mathcal{B}_{\mathbb{R}})$, that, for every x , yields the Gaussian measure with variance one centered around x . The state of y in the program after Line 2 is then the measure space (Y, \mathcal{B}, ν) where $\nu(B) := \int_{x \in X} P(x, B) d\mu$.
- Analogously, Line 3 constructs a kernel $Q: (Y, \mathcal{B}) \rightarrow (Z, \mathcal{C})$ of Gaussian measures $Q(y, -)$ on $(Z, \mathcal{C}) = (\mathbb{R}, \mathcal{B}_{\mathbb{R}})$ centered around y , and the state of z is the measure space (Z, \mathcal{C}, ρ) with $\rho(C) := \int_{y \in Y} Q(y, C) d\nu$. Equivalently, ρ is obtained by integration of the kernel composition $P; Q$ w.r.t. μ .
- Lines (4-5) compute an inverse kernel of $P; Q$ as follows: the composition is integrated w.r.t. μ to obtain a joint measure ζ on $X \times Z$. This joint measure is transposed to ζ^\dagger on $Z \times X$ and then disintegrated

to the inverse kernel R . Then $\int_{z \in (0.5, \infty)} R(z, -) d\rho$ evaluated on the interval $(1, \infty)$ yields the desired result.

Lines (1-3) can be pictorially represented using three concrete samples x, y, z as follows:




Interpreting the program in terms of joint distributions can be done as follows:

- Line 1 constructs the measure space (X, \mathcal{A}, μ) with μ standard Gaussian as above.
- Line 2 constructs a joint measure as follows: let (Y, \mathcal{B}) and P be as above. Define the joint measure θ on $X \times Y$ on its rectangles as

$$\theta(A \times B) := \int_{x \in A} P(x, B) d\mu$$

Then the left marginal $\theta(- \times Y)$ of θ is μ and the right marginal $\theta(X \times -)$ is ν . We consider the state of the program after Line 2 to be the joint distribution θ . Intuitively, this is a way to *remember the input distribution*, and this is, computationally speaking, crucial for Bayesian inference.

- Similarly, Line 3 can be thought of as constructing a joint η on $Y \times Z$ using $(Z, \mathcal{C}) = (\mathbb{R}, \mathcal{B}_{\mathbb{R}})$ and Q as above. This joint will have marginals ν and ρ . The joint measures obtained from Lines (1-3) can be visualized as follows:



In order to say anything about the state after Line 3, we have to be able to compose the joint measures θ and η to $\zeta = \theta ; \eta$. How to define this composition without assumptions on the underlying space is the central contribution of this paper, explained in the following sections.

- Lines (4,5) transpose the joint measure ζ on $X \times Z$ to ζ^\dagger on $Z \times X$ and the output of the program is the ζ^\dagger -measure of the rectangle $(0.5, \infty) \times (1, \infty)$.¹

Our goal now is to define how to compose joint measures θ and η (as above) to a joint measure on X and Z such that

- (i) the marginals of $\theta ; \eta$ are μ and ρ , and
- (ii) integration of the kernel composition $P; Q$ with respect to μ yields $\theta ; \eta$.

Before our general treatment of this problem with continuous measures, we describe the problem and its solution in a discrete setting for instructive purposes.

¹ Computing the result with observe statements on *zero measure events*, e.g. $z = 0.5$, still requires disintegration: the joint ζ^\dagger is disintegrated at $z = 0.5$ and then evaluated on $(1, \infty)$.

3 Discrete Approach

We consider a finite state space; with the appropriate convergence assumptions, this generalizes straightforwardly to countable state spaces. Through the below development, we will have forward pointers to the analogous definitions of the general case which we will present in Section 4, we mark this as “c.f. (n) below”.

3.1 Preliminaries

Formally, a *probability measure* on a finite set $\{1, \dots, n\}$ (from now on overloadingly written n) is a finite additive set function $2^n \rightarrow [0, 1]$ defined uniquely on its points by $i \mapsto x_i$ (sometimes written $x(i)$) such that $\sum_{i=1}^n x_i = 1$. Equivalently, it is a vector $x \in [0, 1]^n$ such that $x^T \cdot \mathbf{1} = 1$, where $\mathbf{1}$ denotes the n -dimensional vector of ones. Integrating a function $f : n \rightarrow \mathbb{R}$ (always measurable) w.r.t. x is done by weighting each element $f(i)$ with its corresponding probability measure x_i . Thus, integration in the discrete world is just matrix multiplication $x^T f$.

A *Markov kernel* is a map $A : n \rightarrow [0, 1]^n$ (always measurable) such that $A(i)$ is a probability measure for every $i \in n$. Equivalently, it is a matrix $A \in [0, 1]^{n \times n}$ such that $A\mathbf{1} = \mathbf{1}$. Integrating a Markov kernel A w.r.t. a probability measure x yields a new probability measure $y = x^T A$. This procedure is given componentwise by $y(j) = \sum_{i=1}^n A(i, j) \cdot x(i)$ (c.f. (2) below).

A (joint) probability measure on the product space $n \times n$ is a matrix $\alpha \in [0, 1]^{n \times n}$ whose entries all sum to one. We have $1 = \sum_{i,j} \alpha(i, j) = \mathbf{1}^T \alpha \mathbf{1} = \mathbf{1}^T \alpha^T \mathbf{1}$, so that the row-sums $\alpha \mathbf{1}$ and the column-sums $\alpha^T \mathbf{1}$ are probability measures; they are the *left* and *right marginal*, respectively. A joint measure α on $n \times n$ defines unique marginals by definition, but the converse is not true in general: there can be many joint measures for two given marginals. (In this setting it is solving $2n$ linear equations of n^2 unknowns.)

3.2 Discrete Markov Kernels to Joint Distributions

We now describe what the functor $\mathbf{J} : \mathbf{Krn} \rightarrow \mathbf{JDist}$ does to *discrete* Markov kernels. For a Markov kernel A , first define the map $A' : n \rightarrow [0, 1]^{n \times n}$ by (c.f. (4) below):

$$A'(i, j, k) = \begin{cases} A(i, k) & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

Equivalently, A' is a $n \times n \times n$ matrix where each page (ranged over by k) is a diagonal $n \times n$ matrix, and the diagonal of the k -th page is given by the k -th column of A .

Then, we have that $(x^T \cdot A')(i, j) = A(i, j) \cdot x_i$ for every $i, j \in n$ and measure $x \in [0, 1]^n$ (c.f. (5) below). So, each i -th row-sum of $x^T \cdot A'$ is just x_i (because the row-sums of A are one) and the column-sums are just entry-wise computations of $A^T x$. This means that $x^T \cdot A'$ is a measure on $n \times n$ with left marginal x and right marginal $y = A^T x$. The functor \mathbf{J} maps the kernel A to the joint measure $x^T A'$. It implicitly depends on x .

Disintegrating the joint measure $\mathbf{J}A$ back to a kernel is always possible for discrete kernels, but does this give us back A ? In the finite setting, to disintegrate a joint measure $\alpha \in [0, 1]^{n \times n}$ on the product space $n_1 \times n_2$ (where $n_1 = n_2 = n$), let $\pi : n_1 \times n_2 \rightarrow n_1$ be the first projection, and $x = \alpha \circ \pi^{-1}$ be the pushforward measure of α through this projection. The choice of name for x is not a coincidence; it is the left marginal of α :

$$x(i) = \sum_{j=1}^n \alpha(i, j).$$

A *disintegration* of α along π is defined as a finite set $\{a^{(i)}\}_{i \in n_1}$ of measures on n_2 such that

$$\forall (i, j) \in n_1 \times n_2 : \quad \alpha(i, j) = a^{(i)}(j) \cdot x(i)$$

Equivalently, this is a kernel $n_1 \rightarrow [0, 1]^{n_2}$, or a matrix in $[0, 1]^{n_1 \times n_2}$. The condition for disintegration in a non-discrete setting contains an integral amounting to a finite sum here, but the measures are uniquely

defined by their point masses, so this condition suffices. In addition, the usual definition involves measurability conditions, but these are trivially satisfied here. The kernel A such that $x^T A = y^T$ is a disintegration of the joint measure $\mathbf{J}A = x^T A'$ defined in the previous paragraph.

The question is now whether A is the only disintegration of $\mathbf{J}A$. The answer is yes, *but only up to negligible events*. More precisely, given α , the disintegration $a^{(i)}$ at i – the i -th row of the kernel A – can be anything if $x(i) = 0$. It is then natural to consider an equivalence class \equiv_x on kernels $A, B \in [0, 1]^{n \times n}$ defined by

$$A \equiv_x B : \iff \forall i \in n \ (x_i > 0 \implies A(i) = B(i)) \quad (\text{c.f. (6)})$$

Recall that, for all i, j , $\mathbf{J}A(i, j) = A(i, j) \cdot x_i$ and $\mathbf{J}B(i, j) = B(i, j) \cdot x_i$. So, we will have that

$$\mathbf{J}A = \mathbf{J}B \text{ if and only if } A \equiv_x B \quad (1)$$

This is exactly the analogue of Lemma 4.1 below in the discrete case.

The goal of our work is to compose joint measures $\alpha, \beta \in [0, 1]^{n \times n}$ with marginals $x, y \in [0, 1]^n$ for α and marginals $y, z \in [0, 1]^n$ for β , to a joint measure $\gamma \in [0, 1]^{n \times n}$, without disintegrating them to kernels first. We want this composition to satisfy some sensible properties: if A and B are kernels such that $\alpha = \mathbf{J}A = x^T A'$ and $\beta = \mathbf{J}B = y^T B'$, then

- (i) the marginals of γ are x and z ;
- (ii) $\gamma = \mathbf{J}(A \cdot B) = x^T (AB)'$ (c.f. Theorem 6.3, \mathbf{J} is a functor); and
- (iii) if C is such that $\gamma = \mathbf{J}C$ then $AB \equiv_x C$ (faithfulness of \mathbf{J} , c.f. Lemma 4.1).

In the discrete case, the composition γ can be defined entry-wise by

$$\gamma(i, k) = \sum_{\substack{j=1 \\ y_j > 0}}^n \frac{\alpha(i, j) \cdot \beta(j, k)}{y_j} \quad (\text{c.f. (16)})$$

We can then verify the above properties. First, the left marginal of γ is x :

$$(\gamma \mathbf{1})(i) = \sum_{k=1}^n \sum_{\substack{j=1 \\ y_j > 0}}^n \frac{\alpha(i, j) \cdot \beta(j, k)}{y_j} = \sum_{\substack{j=1 \\ y_j > 0}}^n \left(\sum_{k=1}^n \beta(j, k) \right) / y_j \cdot \alpha(i, j) = \sum_{j=1}^n \alpha(i, j) = (\alpha \mathbf{1})(i) = x_i$$

Here, we have used that $\sum_{k=1}^n \beta(j, k) = y_j$ and, if, for some j , $y_j = 0$, then $\sum_{i=1}^n \alpha(i, j) = 0$, so $\alpha(i, j) = 0$ (which is why we can leave out the subscript $y_j > 0$ from the sum). A similar calculation shows that its right marginal is z , so property (i) is verified. For property (ii), we have

$$\mathbf{J}(A \cdot B)(i, k) = (A \cdot B)(i, k) \cdot x_i = \sum_{j=1}^n A(i, j) \cdot B(j, k) \cdot x_i = \sum_{\substack{j=1 \\ y_j > 0}}^n \alpha(i, j) \cdot B(j, k) \cdot \frac{y_j}{y_j} = \sum_{\substack{j=1 \\ y_j > 0}}^n \frac{\alpha(i, j) \cdot \beta(j, k)}{y_j}$$

identifying $\mathbf{J}(A \cdot B)$ with γ . Property (iii) is immediate from (1) and property (ii).

4 Krn and JDist

We will define the basics to introduce the categories **Krn** and **JDist**, though the composition of the latter will be defined in Section 6 as we will need a few more results to provide it without resorting to disintegration. However, we can already give the mapping $\mathbf{J}: \mathbf{Krn} \rightarrow \mathbf{JDist}$ and prove that it is well-defined (Lemma 4.1).

4.1 Preliminaries

Let (X, \mathcal{A}) and (Y, \mathcal{B}) be measurable spaces. We abbreviate (X, \mathcal{A}) by X when \mathcal{A} is understood. The letters A, B, C, D denote measurable sets. The space (X, \mathcal{A}) is a *standard Borel space* if \mathcal{A} is the set of Borel sets of a Polish space (separable and completely metrizable). The space (X, \mathcal{A}) is *countably generated* if there exists a countable set $\mathcal{A}_0 \subseteq \mathcal{A}$ such that \mathcal{A} is the smallest σ -algebra containing \mathcal{A}_0 . All standard Borel spaces are countably generated. For an in-depth treatment of measure theory, see, e.g., [6,9,12].

A *Markov kernel* $P : X \rightarrow Y$ is a map $P : X \times \mathcal{B} \rightarrow [0, 1]$ such that

- for fixed $s \in X$, $P(s, -)$ is a probability measure on Y ,
- for fixed $B \in \mathcal{B}$, $P(-, B)$ is a measurable function on X .

These properties allow kernels to be sequentially composed by Lebesgue integration. The class of measurable spaces and Markov kernels forms a category **SRel** [11,6], which is isomorphic to the Kleisli category of the Giry monad. We write $P : X \rightarrow Y$ for the kernel P regarded as a morphism in this category.

For $P : X \rightarrow Y$ a Markov kernel and μ a finite measure on X , write $\mu ; P$ for the measure on Y such that

$$(\mu ; P)(B) = \int_X P(s, B) d\mu(s). \quad (2)$$

This gives a bounded linear map $(- ; P) : \mathbf{Meas} X \rightarrow \mathbf{Meas} Y$ that is monotone and continuous in both the metric and Scott topologies [13].

For $A \in \mathcal{A}$, let A also denote the subidentity kernel

$$A(s, B) = \mathbf{1}_X(s, A \cap B) = \begin{cases} 1, & s \in A \cap B, \\ 0, & s \notin A \cap B. \end{cases}$$

Then for all $A \in \mathcal{A}$ and $B \in \mathcal{B}$,

$$(\mu ; A ; P)(B) = \int_A P(s, B) d\mu(s). \quad (3)$$

The category **Krn**² has as objects probability spaces (X, \mathcal{A}, μ) . The morphisms $(X, \mathcal{A}, \mu) \rightarrow (Y, \mathcal{B}, \nu)$ are equivalence classes modulo μ -nullsets of Markov kernels $P : X \rightarrow Y$ such that

$$\nu(B) = \int_{s \in X} P(s, B) d\mu(s).$$

4.2 From **Krn** to **JDist**

The category **JDist** is the category whose objects are probability spaces (X, \mathcal{A}, μ) and whose morphisms $(X, \mathcal{A}, \mu) \rightarrow (Y, \mathcal{B}, \nu)$ are joint distributions or *couplings* on $X \times Y$ with marginals μ and ν . We will define composition in **JDist** formally in Section 6, but we can already define the embedding functor $\mathbf{J} : \mathbf{Krn} \rightarrow \mathbf{JDist}$. It is the identity on objects, and for morphisms $P : X \rightarrow Y$, let $P' : X \rightarrow X \times Y$ be the kernel that behaves like P , but also remembers its input state: on measurable rectangles $A \times B$,

$$P'(s, A \times B) = \begin{cases} P(s, B), & \text{if } s \in A, \\ 0, & \text{if } s \notin A. \end{cases} \quad (4)$$

² It is worth noting that the category **Krn** mentioned in the introduction, considered by Dahlqvist et al. [5], restricts \mathcal{A} to be standard Borel; we do not impose this restriction.

The kernel P' together with the probability measure μ on X induce a joint measure $\mathbf{J}P$ on $X \times Y$ whose value on measurable rectangles $A \times B$ is

$$\mathbf{J}P(A \times B) = \int_X P'(s, A \times B) d\mu(s) = \int_A P(s, B) d\mu(s). \quad (5)$$

Since morphisms in \mathbf{Krn} are equivalence classes of kernels, there is at this point an obligation for proof of well-definedness. For $P, Q : X \rightarrow Y$ Markov kernels and μ a measure on X , the equivalence is defined as

$$P \equiv_\mu Q \Leftrightarrow \forall B \in \mathcal{B} \quad \mu(\{s \mid P(s, B) \neq Q(s, B)\}) = 0. \quad (6)$$

The following lemma now shows that \mathbf{J} is well-defined:

Lemma 4.1 *Let (X, \mathcal{A}) and (Y, \mathcal{B}) be measurable spaces, \mathcal{B} countably generated. Let $P, Q : X \rightarrow Y$ be Markov kernels and μ a probability measure on X . The following are equivalent:*

- (i) $P \equiv_\mu Q$
- (ii) $P \equiv_\xi Q$ for all ξ absolutely continuous with respect to μ (notation: $\xi \ll \mu$)
- (iii) for all $A \in \mathcal{A}$, $\mu ; A ; P = \mu ; A ; Q$
- (iv) $\mathbf{J}P = \mathbf{J}Q$, considering P and Q as \mathbf{Krn} -morphisms $(X, \mathcal{A}, \mu) \rightarrow (Y, \mathcal{B}, \nu)$, where $\nu = \mu ; P = \mu ; Q$.

Proof. The equivalence of (i) and (ii) is clear from the definition (6).

For (i) \Rightarrow (iii), suppose that $P \equiv_\mu Q$. Let $E_B = \{s \mid P(s, B) \neq Q(s, B)\}$. By definition, $\mu(X \setminus E_B) = 0$. For all $A \in \mathcal{A}$ and $B \in \mathcal{B}$,

$$\begin{aligned} (\mu ; A ; P)(B) &= \int_A P(s, B) d\mu(s) \\ &= \int_{A \cap E_B} P(s, B) d\mu(s) + \int_{A \setminus E_B} P(s, B) d\mu(s) \end{aligned} \quad (7)$$

$$\begin{aligned} &= \int_{A \cap E_B} Q(s, B) d\mu(s) + \int_{A \setminus E_B} Q(s, B) d\mu(s) \\ &= \int_A Q(s, B) d\mu(s) = (\mu ; A ; Q)(B). \end{aligned} \quad (8)$$

The left-hand terms of (7) and (8) agree because the integral is restricted to E_B , and the right-hand terms are 0 because $A \setminus E_B$ is a μ -nullset. As B was arbitrary, $\mu ; A ; P = \mu ; A ; Q$.

Conversely, for (iii) \Rightarrow (i), if $P \not\equiv_\mu Q$, then $\mu(\{s \mid |P(s, B) - Q(s, B)| \geq 1/n\}) > 0$ for some $B \in \mathcal{B}$ and $n > 0$. Letting A be this set, we have

$$\int_A |P(s, B) - Q(s, B)| d\mu(s) \geq \frac{1}{n} \mu(A) > 0,$$

so

$$(\mu ; A ; P)(B) = \int_A P(s, B) d\mu(s) \neq \int_A Q(s, B) d\mu(s) = (\mu ; A ; Q)(B).$$

The equivalence of (iii) and (iv) follows from (3) and (5). \square

Let (X, \mathcal{A}) be a measurable space. The countable measurable partitions \mathcal{D} of X form an upper semi-lattice ordered by refinement, denoted \sqsubseteq , with least common refinement as join, denoted \sqcup . We will often consider the limiting behavior of functions defined on increasingly finer such partitions. For any such map ϕ taking values in a topological space, if $\phi(\mathcal{D}_n)$ converges to the same value for all chains $\mathcal{D}_0 \sqsubseteq \mathcal{D}_1 \sqsubseteq \dots$ that eventually become sufficiently fine, we write $\lim_{\mathcal{D}} \phi(\mathcal{D})$ for this value.

5 Radon-Nikodým Approximants

Integration and Radon-Nikodým derivatives are typically defined in terms of approximants. We will use a similar technique to define the composition of joint measures in **JDist**. In this section we develop the necessary infrastructure.

Let ν and μ be finite measures on (X, \mathcal{A}) . For any $B \in \mathcal{A}$, consider the set

$$\left\{ \frac{\nu(C)}{\mu(C)} \mid C \subseteq B, \mu(C) > 0 \right\} \subseteq \mathbb{R}. \quad (9)$$

This set is nonempty iff $\mu(B) > 0$. In that case, the set has a finite infimum, since $\nu(B)/\mu(B)$ is a member, but it may be unbounded above.

Lemma 5.1 *Let ν and μ be finite measures on (X, \mathcal{A}) . For any $\varepsilon > 0$, there exists a countable measurable partition \mathcal{D} of X such that for all $B \in \mathcal{D}$ with $\mu(B) > 0$, the set (9) is bounded above and*

$$\sup_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} - \inf_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \leq \varepsilon \quad \left(\sup_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} - \inf_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \right) \left(\sup_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \right) \leq \varepsilon^2. \quad (10)$$

Moreover, these properties are preserved under refinement; that is, if $\mathcal{D} \sqsubseteq \mathcal{D}'$ and \mathcal{D} satisfies the inequalities (10) for all $B \in \mathcal{D}$ with $\mu(B) > 0$, then the same is true of \mathcal{D}' .

Proof. For $k \geq 1$, consider the signed measure³ $\nu - (\varepsilon \ln k)\mu$. By the Hahn decomposition theorem, there exists a family of measurable partitions $\{A_k^+, A_k^-\}$ of X , one for each natural number $k \geq 1$, such that $\nu - (\varepsilon \ln k)\mu$ is purely nonnegative on A_k^+ and purely nonpositive on A_k^- . That is, for all measurable $C \subseteq A_k^+$, $\nu(C) \geq (\varepsilon \ln k)\mu(C)$ and for all measurable $C \subseteq A_k^-$, $\nu(C) \leq (\varepsilon \ln k)\mu(C)$. Without loss of generality we can assume $A_1^+ = X$ and $A_1^- = \emptyset$. Let the partition \mathcal{D} consist of the pairwise disjoint sets $\bigcap_{i=1}^k A_i^+ \cap A_{k+1}^-$, $k \geq 1$, along with $\bigcap_{i=1}^{\infty} A_i^+$. The last is a μ -nullset, since if $\mu(C) > 0$, then $C \not\subseteq A_k^+$ for any $k > \exp(\frac{\nu(C)}{\varepsilon\mu(C)})$.

For any measurable $C \subseteq \bigcap_{i=1}^k A_i^+ \cap A_{k+1}^-$, we have $(\varepsilon \ln k)\mu(C) \leq \nu(C) \leq (\varepsilon \ln(k+1))\mu(C)$, so if $\mu(C) > 0$, then $\nu(C)/\mu(C)$ exists and lies in the interval $[\varepsilon \ln k, \varepsilon \ln(k+1)]$. Since $\ln(1+x) \leq x$ for $x \geq 1$,

$$\begin{aligned} \varepsilon \ln(k+1) - \varepsilon \ln k &= \varepsilon \ln\left(\frac{k+1}{k}\right) = \varepsilon \ln\left(1 + \frac{1}{k}\right) \leq \frac{\varepsilon}{k} \leq \varepsilon \\ (\varepsilon \ln(k+1) - \varepsilon \ln k)\varepsilon \ln(k+1) &= \varepsilon \ln\left(1 + \frac{1}{k}\right)\varepsilon \ln(k+1) \leq \frac{\varepsilon}{k} \cdot \varepsilon k = \varepsilon^2, \end{aligned}$$

from which the bounds (10) follow. \square

Let \mathcal{D} be a countable measurable partition of \mathcal{A} . For $s \in X$, define

$$F_{\mathcal{D}} = \bigcup \{B \in \mathcal{D} \mid \mu(B) > 0\} \quad (11)$$

$$\left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^+(s) = \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \sup_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \cdot \mathbf{1}_X(s, B) \quad \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^-(s) = \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \inf_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \cdot \mathbf{1}_X(s, B). \quad (12)$$

The set $F_{\mathcal{D}}$ is measurable with $\mu(X \setminus F_{\mathcal{D}}) = 0$, and the functions $(d\nu/d\mu)_{\mathcal{D}}^+$ and $(d\nu/d\mu)_{\mathcal{D}}^-$ are measurable

³ The signed measure $\nu - \varepsilon k \mu$ already implies the left-hand inequality of (10), which suffices for most purposes, and standard treatments use this. But $\nu - (\varepsilon \ln k)\mu$ gives the right-hand inequality as well, which will be useful in §6 in the definition of composition in **JDist**.

step functions that vanish outside $F_{\mathcal{D}}$. From Lemma 5.1, we have that for sufficiently fine \mathcal{D} ,

$$\left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^{+} - \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^{-} \leq \varepsilon \quad \left(\left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^{+} - \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^{-}\right) \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^{+} \leq \varepsilon^2.$$

These definitions lead to an enhanced version of the Radon-Nikodým theorem.

Theorem 5.2 (Lebesgue-Radon-Nikodým) *Let ν and μ be finite measures on (X, \mathcal{A}) . There exist measures ν_0, ν_1 , a measurable set $F \in \mathcal{A}$, a measurable real-valued function f defined on X , and a countable \sqsubseteq -chain $\mathcal{D}_0 \sqsubseteq \mathcal{D}_1 \sqsubseteq \dots$ such that*

- (i) *(Lebesgue decomposition) ν_0 and ν_1 form a Lebesgue decomposition of ν on F with respect to μ ; that is,*

$$\nu = \nu_0 + \nu_1 \quad \nu_0 \ll \mu \quad \nu_1(F) = 0 \quad \mu(X \setminus F) = 0;$$

- (ii) *(Radon-Nikodým theorem) $f(s) = 0$ for all $s \notin F$ and*

$$\int_A f(s) d\mu(s) = \nu_0(A), \quad A \in \mathcal{A}; \quad (13)$$

- (iii) *(Uniform approximation) The sequence $(d\nu/d\mu)_n^- = (d\nu/d\mu)_{\mathcal{D}_n}^-$ is monotone nondecreasing on F , the sequence $(d\nu/d\mu)_n^+ = (d\nu/d\mu)_{\mathcal{D}_n}^+$ is monotone nonincreasing everywhere, and both sequences converge pointwise to f and converge uniformly on F .*

If $\nu \ll \mu$, we can take $\nu_0 = \nu$ and $\nu_1 = 0$ in (i), in which case (ii) gives $\int_A f(s) d\mu(s) = \nu(A)$. In this case, f is the standard Radon-Nikodým derivative $d\nu/d\mu$.

The version [12, Theorem 3, p. 258] asserts (i) and (ii) only without reference to the approximants $(d\nu/d\mu)_n^+$ and $(d\nu/d\mu)_n^-$, but (iii) is essentially how it is proved. In fact, [12] gives three proofs. We give a fourth here for completeness.

Proof. Let $(d\nu/d\mu)_{\mathcal{D}}^+$, $(d\nu/d\mu)_{\mathcal{D}}^-$, and $F_{\mathcal{D}}$ be defined as in (11) and (12). By definition,

$$\left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^{-}(s) \leq \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \frac{\nu(B)}{\mu(B)} \cdot \mathbf{1}_X(s, B) \leq \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^{+}(s).$$

If $\mathcal{D} \sqsubseteq \mathcal{D}'$, then $F_{\mathcal{D}'} \subseteq F_{\mathcal{D}}$, $(d\nu/d\mu)_{\mathcal{D}'}^- \leq (d\nu/d\mu)_{\mathcal{D}}^-$ pointwise on $F_{\mathcal{D}'}$, and $(d\nu/d\mu)_{\mathcal{D}'}^- \leq (d\nu/d\mu)_{\mathcal{D}'}^+ \leq (d\nu/d\mu)_{\mathcal{D}}^+$ pointwise everywhere. Moreover, by Lemma 5.1, for all $\varepsilon > 0$ there exists a sufficiently fine \mathcal{D} that $(d\nu/d\mu)_{\mathcal{D}}^+ - (d\nu/d\mu)_{\mathcal{D}}^- < \varepsilon$ pointwise. It follows that for any countable chain $\mathcal{D}_0 \sqsubseteq \mathcal{D}_1 \sqsubseteq \dots$ of sufficiently fine countable measurable partitions, $(d\nu/d\mu)_n^+$ and $(d\nu/d\mu)_n^-$ converge pointwise to a measurable function $f = \inf_n (d\nu/d\mu)_n^+$ and converge uniformly on $F = \bigcap_n F_{\mathcal{D}_n}$. In addition, the region of uniform convergence F is of full μ -measure, as $\mu(F) = \inf_n \mu(F_{\mathcal{D}_n}) = \mu(X)$, and f vanishes outside F . This establishes (iii).

For (i), if $\mu(C) = 0$, then $\nu(C \cap A_k^-) \leq (\varepsilon \ln k) \mu(C \cap A_k^-) = 0$ for all k , where the A_k^- are the Hahn decomposition sets constructed in the proof of Lemma 5.1. Assuming \mathcal{D}_n refines one of the partitions defined in that lemma, we have $F \subseteq F_{\mathcal{D}_n} \subseteq \bigcup_k A_k^-$, so $\nu(C \cap F) = 0$. Taking $\nu_0(C) = \nu(C \cap F)$ and $\nu_1(C) = \nu(C \setminus F)$ give a Lebesgue decomposition of ν on F with respect to μ ; in particular,

$$\mu(C) = 0 \Rightarrow \nu(C) = \nu(C \setminus F) = \nu_1(C). \quad (14)$$

For (ii), for sufficiently fine \mathcal{D} ,

$$\begin{aligned} \int_A \left(\frac{d\nu}{d\mu} \right)_D^+ (s) d\mu(s) &= \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \sup_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \cdot \mu(A \cap B) \leq \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \left(\inf_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} + \varepsilon \right) \cdot \mu(A \cap B) \\ &\leq \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \left(\frac{\nu(B)}{\mu(B)} + \varepsilon \right) \cdot \mu(B) \leq \nu(X) + \varepsilon \mu(X), \end{aligned}$$

thus the integral exists and is finite. Moreover,

$$\int_A \left(\frac{d\nu}{d\mu} \right)_n^+ (s) d\mu(s) = \sum_{\substack{B \in \mathcal{D}_n \\ \mu(B) > 0}} \sup_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \cdot \mu(A \cap B) = \sum_{\substack{B \in \mathcal{D}_n \\ \mu(A \cap B) > 0}} \sup_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \cdot \mu(A \cap B). \quad (15)$$

The right-hand equality in (15) follows from the observation that all summands corresponding to $B \in \mathcal{D}_n$ with $\mu(A \cap B) = 0$ vanish, whereas for those with $\mu(A \cap B) > 0$, the test $\mu(B) > 0$ is redundant. Specializing the supremum in (15) at $C = A \cap B$,

$$\begin{aligned} \int_A \left(\frac{d\nu}{d\mu} \right)_n^+ (s) d\mu(s) &\geq \sum_{\substack{B \in \mathcal{D}_n \\ \mu(A \cap B) > 0}} \nu(A \cap B) = \nu(A) - \sum_{\substack{B \in \mathcal{D}_n \\ \mu(A \cap B) = 0}} \nu(A \cap B) \\ &= \nu(A) - \sum_{\substack{B \in \mathcal{D}_n \\ \mu(A \cap B) = 0}} \nu_1(A \cap B) && \text{by (14)} \\ &\geq \nu(A) - \sum_{B \in \mathcal{D}_n} \nu_1(A \cap B) = \nu(A) - \nu_1(A) = \nu_0(A). \end{aligned}$$

Similarly,

$$\int_A \left(\frac{d\nu}{d\mu} \right)_n^- (s) d\mu(s) = \sum_{\substack{B \in \mathcal{D}_n \\ \mu(A \cap B) > 0}} \inf_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \cdot \mu(A \cap B) = \sum_{\substack{B \in \mathcal{D}_n \\ \mu(A \cap B \cap F) > 0}} \inf_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \cdot \mu(A \cap B \cap F).$$

The former equality follows from an argument similar to (15), the latter from the fact that μ vanishes outside F . Specializing the infimum at $C = A \cap B \cap F$,

$$\int_A \left(\frac{d\nu}{d\mu} \right)_n^- (s) d\mu(s) \leq \sum_{\substack{B \in \mathcal{D}_n \\ \mu(A \cap B \cap F) > 0}} \nu(A \cap B \cap F) \leq \sum_{B \in \mathcal{D}_n} \nu_0(A \cap B) = \nu_0(A).$$

Thus $\nu_0(A)$ is the limit (13). \square

The value of the integral (13) is independent of the choice of the \sqsubseteq -chain \mathcal{D}_n , but F and f are not, and there is no one choice that works uniformly for all \sqsubseteq -chains. That is why $d\nu/d\mu$ is only defined up to a μ -nullset. However, by taking least common refinements, one can find F and f that work for several \sqsubseteq -chains at once. That is one reason for making (iii) explicit.

6 Composition in \mathbf{JDist}

To show that \mathbf{JDist} is a category and that $\mathbf{J} : \mathbf{Krn} \rightarrow \mathbf{JDist}$ is a functor, we must define composition and the identity morphisms in \mathbf{JDist} and show that they are preserved by \mathbf{J} .

Composition in **JDist** is defined as follows. For $\theta : (X, \mathcal{A}, \mu) \rightarrow (Y, \mathcal{B}, \nu)$ and $\eta : (Y, \mathcal{B}, \nu) \rightarrow (Z, \mathcal{C}, \xi)$, define

$$(\theta ; \eta)(A \times C) = \lim_{\mathcal{D}} \sum_{\substack{B \in \mathcal{D} \\ \nu(B) > 0}} \frac{\theta(A \times B) \cdot \eta(B \times C)}{\nu(B)}, \quad (16)$$

where the limit is taken over countable measurable partitions \mathcal{D} of the mediating space Y . We argue below (Theorem 6.1) that the limit exists.

The identity morphisms are the joint distributions $\mathbf{J}1_X$ obtained from the identity kernels $1_X : X \rightarrow X$. Checking associativity of the composition is a mechanical exercise, where one employs commutativity of the countable limit sums over the partitions of the two intermedating spaces: for θ and η as above, and $\zeta : (Z, \mathcal{C}, \xi) \rightarrow (W, \mathcal{F}, \rho)$, one has

$$(\theta ; \eta ; \zeta)(A \times F) = \lim_{\mathcal{D}} \sum_{\substack{B \in \mathcal{D} \\ \nu(B) > 0}} \lim_{\mathcal{E}} \sum_{\substack{C \in \mathcal{E} \\ \xi(C) > 0}} \frac{\theta(A \times B) \cdot \eta(B \times C) \cdot \zeta(C \times F)}{\nu(B) \cdot \xi(C)}$$

where \mathcal{D} and \mathcal{E} range over countable measurable partitions of Y and Z , respectively.

Note that the definition of composition is completely symmetric in the input and output space. The category **JDist** is thus a dagger category whose involution \dagger is composition with transpose: $\theta^\dagger(C \times A) = \theta(A \times C)$.

Theorem 6.1 *Let μ, ν, ξ be finite measures on Y . Let $(d\nu/d\mu)_{\mathcal{D}}^+$ and $(d\xi/d\mu)_{\mathcal{D}}^+$ be the approximants defined in (12). If there exists a countable measurable partition \mathcal{D} such that*

$$\int_Y \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^+(t) \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^+(t) d\mu(t) < \infty,$$

then the limit

$$\lim_{\mathcal{D}} \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \frac{\nu(B) \cdot \xi(B)}{\mu(B)} \quad (17)$$

exists and is equal to

$$\int_Y f(t) g(t) d\mu(t) = \inf_n \int_Y \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}_n}^+(t) \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}_n}^+(t) d\mu(t) \quad (18)$$

for any sufficiently fine countable \sqsubseteq -chain $\mathcal{D}_0 \sqsubseteq \mathcal{D}_1 \sqsubseteq \dots$, where $f = \inf_n (d\nu/d\mu)_{\mathcal{D}_n}^+$ and $g = \inf_n (d\xi/d\mu)_{\mathcal{D}_n}^+$.

Proof. By definition of $(d\nu/d\mu)_{\mathcal{D}}^+$ and $(d\xi/d\mu)_{\mathcal{D}}^+$,

$$\begin{aligned} \int_Y \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^+(t) \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^+(t) d\mu(t) &= \int_Y \left(\sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \sup_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} 1_Y(t, B) \right) \left(\sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \sup_{\substack{D \subseteq B \\ \mu(D) > 0}} \frac{\xi(D)}{\mu(D)} 1_Y(t, B) \right) d\mu(t) \\ &= \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \left(\sup_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} \right) \left(\sup_{\substack{D \subseteq B \\ \mu(D) > 0}} \frac{\xi(D)}{\mu(D)} \right) \mu(B). \end{aligned}$$

To show (17) and (18), for arbitrarily small positive ε , we have by Lemma 5.1

$$\begin{aligned}
 \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \frac{\nu(B)\xi(B)}{\mu(B)} &= \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \frac{\nu(B)\xi(B)}{\mu(B)\mu(B)} \mu(B) \leq \int_Y \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^+(t) \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^+(t) d\mu(t) \\
 &\leq \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \left(\inf_{\substack{C \subseteq B \\ \mu(C) > 0}} \frac{\nu(C)}{\mu(C)} + \varepsilon \right) \left(\inf_{\substack{D \subseteq B \\ \mu(D) > 0}} \frac{\xi(D)}{\mu(D)} + \varepsilon \right) \mu(B) \\
 &\leq \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \left(\frac{\nu(B)}{\mu(B)} + \varepsilon \right) \left(\frac{\xi(B)}{\mu(B)} + \varepsilon \right) \mu(B) \\
 &\leq \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \left(\frac{\nu(B)\xi(B)}{\mu(B)} + \varepsilon\xi(B) + \varepsilon\nu(B) + \varepsilon^2\mu(B) \right) \\
 &\leq \left(\sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \frac{\nu(B)\xi(B)}{\mu(B)} \right) + \varepsilon\xi(Y) + \varepsilon\nu(Y) + \varepsilon^2\mu(Y).
 \end{aligned}$$

For the remaining statement (18), we use the stronger claim (10) of Lemma 5.1. Since

$$\left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^-(t) \leq f(t) \leq \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^+(t) \quad \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^-(t) \leq g(t) \leq \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^+(t),$$

we have

$$\left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^-(t) \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^-(t) \leq f(t)g(t) \leq \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^+(t) \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^+(t).$$

We must choose \mathcal{D}_n so that

$$\left(\frac{d\nu}{d\mu}\right)_n^+(t) \left(\frac{d\xi}{d\mu}\right)_n^+(t) - \left(\frac{d\nu}{d\mu}\right)_n^-(t) \left(\frac{d\xi}{d\mu}\right)_n^-(t)$$

becomes arbitrarily small. By (10) of Lemma 5.1, \mathcal{D} can be chosen so that

$$\begin{aligned}
 &\left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^+(t) \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^+(t) - \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^-(t) \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^-(t) \\
 &\leq \left(\left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^+(t) - \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^-(t) \right) \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^+(t) + \left(\frac{d\nu}{d\mu}\right)_{\mathcal{D}}^+(t) \left(\left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^+(t) - \left(\frac{d\xi}{d\mu}\right)_{\mathcal{D}}^-(t) \right) \leq 2\varepsilon^2.
 \end{aligned}$$

□

Corollary 6.2 *The map $\theta ; \eta$ as defined in (16) on measurable rectangles extends to a joint probability measure on $X \times Z$.*

Proof. This can be done using the Carathéodory-Hahn-Kolmogorov extension theorem⁴. It suffices to verify the premises of that theorem, namely

- (i) $\theta ; \eta$ is finitely additive on measurable rectangles: if $\{A_n \times C_n\}_n$ is a finite set of pairwise disjoint measurable rectangles whose union is a measurable rectangle, then

$$(\theta ; \eta)\left(\bigcup_n (A_n \times C_n)\right) = \sum_n (\theta ; \eta)(A_n \times C_n).$$

⁴ This result is sometimes referred to simply as Carathéodory's extension theorem and Hahn-Kolmogorov theorem, among other names.

- (ii) If for each $i \geq 0$, $\{A_n^i \times C_n^i\}_n$ is a finite collection of pairwise disjoint measurable rectangles with $\bigcup_n (A_n^{i+1} \times C_n^{i+1}) \subseteq \bigcup_n (A_n^i \times C_n^i)$, and if $\bigcap_i \bigcup_n (A_n^i \times C_n^i) = \emptyset$, then

$$\inf_i (\theta ; \eta) \left(\bigcup_n (A_n^i \times C_n^i) \right) = 0.$$

For (i), we can assume without loss of generality that the A_n are pairwise disjoint and the C_m are pairwise disjoint, and we are to show

$$(\theta ; \eta) \left(\bigcup_n A_n \times \bigcup_m C_m \right) = \sum_n \sum_m (\theta ; \eta) (A_n \times C_m).$$

Since limits commute with finite sums,

$$\begin{aligned} (\theta ; \eta) \left(\bigcup_n A_n \times \bigcup_m C_m \right) &= \lim_{\mathcal{D}} \sum_{B \in \mathcal{D}} \frac{\theta(\bigcup_n A_n \times B) \cdot \eta(B \times \bigcup_m C_m)}{\nu(B)} = \lim_{\mathcal{D}} \sum_{B \in \mathcal{D}} \sum_n \sum_m \frac{\theta(A_n \times B) \cdot \eta(B \times C_m)}{\nu(B)} \\ &= \sum_n \sum_m \lim_{\mathcal{D}} \sum_{B \in \mathcal{D}} \frac{\theta(A_n \times B) \cdot \eta(B \times C_m)}{\nu(B)} = \sum_n \sum_m (\theta ; \eta) (A_n \times C_m). \end{aligned}$$

For (ii),

$$\begin{aligned} \inf_i (\theta ; \eta) \left(\bigcup_n (A_n^i \times C_n^i) \right) &= \inf_i \sum_n (\theta ; \eta) (A_n^i \times C_n^i) \\ &= \inf_i \sum_n \lim_{\mathcal{D}} \sum_{B \in \mathcal{D}} \frac{\theta(A_n^i \times B) \cdot \eta(B \times C_n^i)}{\nu(B)} = \inf_i \lim_{\mathcal{D}} \sum_{B \in \mathcal{D}} \sum_n \frac{\theta(A_n^i \times B) \cdot \eta(B \times C_n^i)}{\nu(B)}. \end{aligned}$$

We argue that if $\bigcap_i \bigcup_n (A_n^i \times C_n^i) = \emptyset$, then either $\bigcap_i \bigcup_n A_n^i = \emptyset$ or $\bigcap_i \bigcup_n C_n^i = \emptyset$. Suppose not. Let $s \in \bigcap_i \bigcup_n A_n^i$ and $t \in \bigcap_i \bigcup_n C_n^i$. Then for all i there exists n such that $s \in A_n^i$ and there exists m such that $t \in C_m^i$. By renumbering if necessary, we can assume that for all i , $s \in A_1^i$ and $t \in C_1^i$. Then $(s, t) \in \bigcap_i (A_1^i \times C_1^i) \subseteq \bigcap_i \bigcup_n (A_n^i \times C_n^i)$.

By symmetry, assume without loss of generality that $\bigcap_i \bigcup_n A_n^i = \emptyset$. Since

$$\eta(B \times C_n^i) / \nu(B) = \eta(B \times C_n^i) / \eta(B \times Z) \leq 1,$$

we have

$$\inf_i \lim_{\mathcal{D}} \sum_{B \in \mathcal{D}} \sum_n \frac{\theta(A_n^i \times B) \cdot \eta(B \times C_n^i)}{\nu(B)} \leq \inf_i \lim_{\mathcal{D}} \sum_{B \in \mathcal{D}} \sum_n \theta(A_n^i \times B) = \inf_i \theta \left(\bigcup_n A_n^i \times Y \right) = 0.$$

□

Theorem 6.3 (Faithfulness) *The map \mathbf{J} constitutes a faithful embedding $\mathbf{J} : \mathbf{Krn} \rightarrow \mathbf{JDist}$.*

Proof. Composition in \mathbf{Krn} is defined by $[P]_\mu ; [Q]_\nu = [P ; Q]_\mu$. This is well defined by Lemma 4.1. To confirm functoriality, we show that $\mathbf{J}(P ; Q) = \mathbf{J}P ; \mathbf{J}Q$ and that the $\mathbf{J}1_X$ are identities for composition in \mathbf{JDist} .

For composition, using the fact that for any P , $\mathbf{J}P(A \times B) = (\mu ; A ; P)(B)$, the left-hand side gives

$$\mathbf{J}(P ; Q)(A \times C) = (\mu ; A ; P ; Q)(C).$$

For the right-hand side, we observe that

$$\inf_{t \in B} Q(t, C) \int_B \nu(dt) \leq \int_B Q(t, C) \nu(dt) \leq \sup_{t \in B} Q(t, C) \int_B \nu(dt),$$

or in other words,

$$\inf_{t \in B} Q(t, C) \nu(B) \leq (\nu ; B ; Q)(C) \leq \sup_{t \in B} Q(t, C) \nu(B),$$

so for $\nu(B) > 0$,

$$\inf_{t \in B} Q(t, C) \leq \frac{(\nu ; B ; Q)(C)}{\nu(B)} \leq \sup_{t \in B} Q(t, C).$$

Thus

$$\begin{aligned} (\mathbf{J}P ; \mathbf{J}Q)(A \times C) &= \lim_{\mathcal{D}} \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \frac{\mathbf{J}P(A \times B) \cdot \mathbf{J}Q(B \times C)}{\nu(B)} = \lim_{\mathcal{D}} \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} \frac{(\mu; A; P)(B) \cdot (\nu; B; Q)(C)}{\nu(B)} \\ &= \lim_{\mathcal{D}} \sum_{\substack{B \in \mathcal{D} \\ \mu(B) > 0}} (\mu ; A ; P)(B) \cdot \sup_{t \in B} Q(t, C) = \int_Y (\mu ; A ; P)(dt) \cdot Q(t, C) = (\mu ; A ; P ; Q)(C). \end{aligned}$$

For the identities, $\mathbf{J}1_X ; \mathbf{J}P = \mathbf{J}(1_X ; P) = \mathbf{J}P$ and $\mathbf{J}P ; \mathbf{J}1 = \mathbf{J}(P ; 1_Y) = \mathbf{J}P$. The equivalence of Lemma 4.1(i) and (iv) establishes the faithfulness of the embedding. \square

The embedding is also full when \mathbf{Krn} is restricted to standard Borel spaces—hence, the category considered by Dahlqvist et al.[5] can be fully and faithfully embedded in \mathbf{JDist} .

7 Conclusion

We have presented a category of joint distributions \mathbf{JDist} in which composition can be defined without reference to disintegration. To define this composition we explored approximants and an enhanced version of the Radon-Nikodým theorem. Our motivation to define this new category is to provide a setting in which semantics of probabilistic programs can be done under weaker assumptions than previous work (where e.g. standard Borel spaces were considered). We showed that the category \mathbf{Krn} of Markov kernels on countably generated spaces can be faithfully embedded to \mathbf{JDist} and, when \mathbf{Krn} is restricted to standard Borel spaces, the embedding is also full.

Although the assumption of spaces being standard Borel may not seem to be a very restrictive one, we believe our new way of composing joint measures will be profitable nevertheless. The category \mathbf{JDist} , with its definition of composition, is a natural domain to contrive and prove correctness of approximation schemes of joint distributions for probabilistic semantics. This is something we already hinted at with the discrete approach in Section 3.

A natural direction for future work is to explore the formulation of a point-free treatment based on the observation that the objects (X, \mathcal{A}, μ) of \mathbf{Krn} and \mathbf{JDist} do not really depend on the measure μ , but only its σ -ideal of nullsets.

Acknowledgments

Thanks to Fredrik Dahlqvist, Vincent Danos, Nate Foster, Justin Hsu, Bart Jacobs, Bobby Kleinberg, Radu Mardare, Prakash Panangaden, Daniel Roy, and Steffen Smolka. Special thanks to Sam Staton for catching a serious error in an earlier draft. Thanks to the Bellairs Research Institute of McGill University for providing a wonderful research environment. The support of the National Science Foundation under grant CCF-2008083 is gratefully acknowledged.

References

- [1] Abramsky, S., R. Blute and P. Panangaden, *Nuclear and trace ideals in tensored *-categories*, Journal of Pure and Applied Algebra **143**, pages 3–47 (1999), ISSN 0022-4049.
[https://doi.org/https://doi.org/10.1016/S0022-4049\(98\)00106-6](https://doi.org/https://doi.org/10.1016/S0022-4049(98)00106-6)

- [2] Chang, J. T. and D. Pollard, *Conditioning as disintegration*, *Statistica Neerlandica* **51**, pages 287–317 (1997).
- [3] Cho, K. and B. Jacobs, *Disintegration and Bayesian inversion, both abstractly and concretely*, *CoRR* **abs/1709.00322** (2017). [1709.00322](https://arxiv.org/abs/1709.00322).
<http://arxiv.org/abs/1709.00322>
- [4] Culbertson, J. and K. Sturtz, *A categorical foundation for Bayesian probability*, *Applied Categorical Structures* **22**, page 647–662 (2014).
<https://doi.org/10.1007/s10485-013-9324-9>
- [5] Dahlqvist, F., V. Danos, I. Garnier and A. Silva, *Borel kernels and their approximation, categorically* (2018).
<https://arxiv.org/abs/1803.02651v2>
- [6] Doberkat, E.-E., *Stochastic Relations: Foundations for Markov Transition Systems*, *Studies in Informatics*, Chapman Hall (2007).
- [7] Faden, A. M., *The existence of regular conditional probabilities: necessary and sufficient conditions*, *Ann. Prob.* **13**, pages 288–298 (1985).
https://projecteuclid.org/download/pdf_1/euclid.aop/1176993081
- [8] Goodman, N. D., V. K. Mansinghka, D. Roy, K. Bonawitz and J. B. Tenenbaum, *Church: a language for generative models*, in: *Proc. 24th Conf. Uncertainty in Artificial Intelligence (UAI'08)*, page 220–229 (2008).
- [9] Halmos, P. R., *Measure theory*, Springer (2013).
- [10] Pachl, J. K., *Disintegration and compact measures*, *Math. Scand.* **43**, pages 157–168 (1978).
- [11] Panangaden, P., *Labelled Markov Processes*, Imperial College Press (2009).
- [12] Rao, M. M., *Measure Theory and Integration*, Wiley-Interscience (1987).
- [13] Smolka, S., P. Kumar, N. Foster, D. Kozen and A. Silva, *Cantor meets scott: semantic foundations for probabilistic networks*, in: G. Castagna and A. D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, pages 557–571, ACM (2017).
<https://doi.org/10.1145/3009837.3009843>
- [14] Tolpin, D., J.-W. van de Meent and F. Wood, *Probabilistic programming in Anglican*, in: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECMLKDD '15)*, pages 308–311 (2015).