Under consideration for publication in Math. Struct. in Comp. Science

Disintegration and Bayesian Inversion via String Diagrams[†]

KENTA CHO and BART JACOBS

Institute for Computing and Information Sciences, Radboud University
P.O.Box 9010, 6500 GL Nijmegen, the Netherlands
Email: {K.Cho,bart}@cs.ru.nl

Received 14 June 2018

The notions of disintegration and Bayesian inversion are fundamental in conditional probability theory. They produce channels, as conditional probabilities, from a joint state, or from an already given channel (in opposite direction). These notions exist in the literature, in concrete situations, but are presented here in abstract graphical formulations. The resulting abstract descriptions are used for proving basic results in conditional probability theory. The existence of disintegration and Bayesian inversion is discussed for discrete probability, and also for measure-theoretic probability — via standard Borel spaces and via likelihoods. Finally, the usefulness of disintegration and Bayesian inversion is illustrated in several examples.

1. Introduction

The essence of conditional probability can be summarised informally in the following equation about probability distributions:

$$\textit{joint} = \textit{conditional} \cdot \textit{marginal}.$$

A bit more precisely, when we have joint probabilities $\mathrm{P}(x, y)$ for elements x, y ranging over two sample spaces, the above equation splits into two equations,

$$\mathrm{P}(y \mid x) \cdot \mathrm{P}(x) = \mathrm{P}(x, y) = \mathrm{P}(x \mid y) \cdot \mathrm{P}(y), \tag{1}$$

where $\mathrm{P}(x)$ and $\mathrm{P}(y)$ describe the marginals, which are obtained by discarding variables. We see that conditional probabilities $\mathrm{P}(y \mid x)$ and $\mathrm{P}(x \mid y)$ can be constructed in two directions, namely y given x, and x given y. We also see that we need to copy variables: x on the left-hand-side of the equations (1), and y on the right-hand-side.

Conditional probabilities play a crucial role in Bayesian probability theory. They form

[†] The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement nr. 320571.

the nodes of Bayesian networks (Pearl, 1988; Bernardo and Smith, 2000; Barber, 2012), which reflect the conditional independencies of the underlying joint distribution via their graph structure. As part of our approach, we shall capture conditional independence in an abstract manner.

The main notion of this paper is *disintegration*. It is the process of extracting a conditional probability from a joint probability. Disintegration, as we shall formalise it here, gives a structural description of the above equation (1) in terms of *states* and *channels*. In general terms, a *state* is a probability distribution of some sort (discrete, measure-theoretic, or even quantum) and a *channel* is a map or morphism in a probabilistic setting, like $\mathrm{P}(y \mid x)$ and $\mathrm{P}(x \mid y)$ as used above. It can take the form of a stochastic matrix, probabilistic transition system, Markov kernel, conditional probability table (in a Bayesian network), or morphism in a Kleisli category of a 'probability monad' (Jacobs, 2017). A state is a special kind of channel, with trivial domain. Thus we can work in a monoidal category of channels, where we need discarding and copying — more formally, a comonoid structure on each object — in order to express the above conditional probability equations (1).

In this article we abstract away from interpretation details and will describe disintegration pictorially, in the language of string diagrams. This language can be seen as the internal language of symmetric monoidal categories (Selinger, 2010) — with comonoids in our case. The essence of disintegration becomes: extracting a conditional probability channel from a joint state.

Categorical approaches to Bayesian conditioning have appeared for instance in (Culbertson and Sturtz, 2014; Staton et al., 2016; Clerc et al., 2017) and in (Jacobs et al., 2015; Jacobs and Zanasi, 2016; Jacobs, 2017). The latter references use effectus theory (Jacobs, 2015; Cho et al., 2015), a new comprehensive approach aimed at covering the logic of both quantum theory and probability theory, supported by a Python-based tool *EfProb*, for 'effectus probability'. This tool is used for the (computationally extensive) examples in this paper.

Disintegration, also known as regular conditional probability, is a notoriously difficult operation in measure-theoretic probability, see *e.g.* (Pollard, 2002; Panangaden, 2009; Chang and Pollard, 1997): it may not exist (Stoyanov, 2014); even if it exists it may be determined only up to negligible sets; and it may not be continuous or computable (Ackerman et al., 2011). Disintegration has been studied using categorical language in (Culbertson and Sturtz, 2014), which focuses on a specific category of probabilistic mappings. Our approach here is more axiomatic.

We thus describe disintegration as going from a joint state to a channel. A closely related concept is *Bayesian inversion*: it turns a channel (with a state) into a channel in opposite direction. We show how Bayesian inversion can be understood and expressed easily in terms of disintegration — and also how, in the other direction, disintegration can be obtained from Bayesian inversion. Bayesian inversion is taken as primitive notion in (Clerc et al., 2017). Here we start from disintegration. The difference is a matter of choice.

Bayesian inversion is crucial for backward inference. We explain it informally: let σ be a state of a domain/type X, and $c\colon X \to Y$ be a channel; Bayesian inversion yields a

channel $d\colon Y \to X$. Informally, it produces for an element $y \in Y$, seen as singleton/point predicate $\{y\}$, the conditioning of the state σ with the pulled back evidence $c^{-1}(\{y\})$. A concrete example involving such 'point observations' will be described at the end of Section 8. More generally, disintegration and Bayesian inversion are used to structurally organise state updates in the presence of new evidence in probabilistic programming, see *e.g.* (Gordon et al., 2014; Borgström et al., 2013; Staton et al., 2016; Katoen et al., 2015). See also (Shan and Ramsey, 2017), where disintegration is handled via symbolic manipulation.

Disintegration and Bayesian inversion are relatively easy to define in discrete probability theory. The situation is much more difficult in measure-theoretic probability theory, first of all because point predicates $\{y\}$ do not make much sense there, see also (Chang and Pollard, 1997). A common solution to the problem of the existence of disintegration / Bayesian inversion is to restrict ourselves to standard Borel spaces, as in (Clerc et al., 2017). We take this approach too. There is still an issue that disintegration is determined only up to negligible sets. We address this by defining 'almost equality' in our abstract pictorial formulation. This allows us to present a fundamental result from (Clerc et al., 2017) abstractly in our setting, see Section 5.

Another common, more concrete solution is to assume a *likelihood*, that is, a probabilistic relation $X \times Y \to \mathbb{R}_{\geq 0}$. Such a likelihood gives rise to probability density function (pdf), providing a good handle on the situation, see (Pawitan, 2001). The technical core of Section 8 is a generalisation of this likelihood-based approach.

The paper is organised as follows. It starts with a brief introduction to the graphical language that we shall be using, and to the underlying monoidal categories with discarding and copying. Then, Section 3 introduces both disintegration and Bayesian inversion in this graphical language, and relates the two notions. Subsequently, Section 4 contains an elaborated example, namely of naive Bayesian classification. A standard example from the literature (Witten et al., 2011) is redescribed in the current setting: first, channels are extracted via disintegration from a table with given data; next, Bayesian inversion is applied to the combined extracted channels, giving the required classification. This is illustrated in both the discrete and the continuous version of the example.

Next, Section 5 is more technical and elaborates the standard equality notion of 'equal almost everywhere' in the current setting. This is used for describing Bayesian inversion in a more formal way, following (Clerc et al., 2017). Section 6 uses our graphical approach to review conditional independence and to prove at an abstract level several known results, namely the equivalence of various formulations of conditional independence, and the 'graphoid' axioms from (Verma and Pearl, 1988; Geiger et al., 1990). Section 7 relaxes the requirement that maps are causal, so that 'effects' can be used as the duals of states for validity and conditioning. The main result relates conditioning of joint states to forward and backward inference via the extracted channels, in the style of (Jacobs and Zanasi, 2016; Jacobs and Zanasi, 2018); it is illustrated in a concrete example, where a Bayesian network is seen as a graph in a Kleisli category — following (Fong, 2012). Finally, Section 8 gives the likelihood formulation of disintegration and inversion, as briefly described above.

2. Graphical language

The basic idea underlying this paper is to describe probability theory in terms of *channels*. A channel $f\colon X \to Y$ is a (stochastic) process from a system of type X into that of Y. Concretely, it may be a probability matrix or kernel. Our standing assumption is that types (as objects) and channels (as arrows) form a *symmetric monoidal category*. For the formal definition we refer to (Mac Lane, 1998). We informally summarise that we have the following constructions.

1 Sequential composition $g \circ f\colon X \to Z$ for appropriately typed channels $f\colon X \to Y$ and $g\colon Y \to Z$.

2 Parallel composition $f \otimes g\colon X \otimes Z \to Y \otimes W$ for $f\colon X \to Y$ and $g\colon Z \to W$. This involves composition of types $X \otimes Z$.

3 Identity channels $\mathrm{id}_X : X \to X$, which 'do nothing'. Thus $\mathrm{id} \circ f = f = \mathrm{id} \circ f$.

4 A unit type I, which represents 'no system'. Thus $I \otimes X \cong X \cong X \otimes I$.

5 Swap isomorphisms $X \otimes Y \cong Y \otimes X$ and associativity isomorphisms $(X \otimes Y) \otimes Z \cong X \otimes (Y \otimes Z)$, so the ordering in composed types does not matter.

The representation of such channels in the ordinary 'formula' notation easily becomes complex and thus reasoning becomes hard to follow. A graphical language known as *string diagrams* offers a more convenient and intuitive way of reasoning in a symmetric monoidal category.

In string diagrams, types/objects are represented as wires '|', with information flowing bottom to top. The composition of types is depicted by juxtaposition of wires, and the unit type is 'no diagram' as below.

Channels/arrows are represented by boxes with an input wire(s) and an output wire(s), in upward direction. When a box does not have input or output, we write it as a triangle or diamond. For example, $f\colon X \to Y \otimes Z$, $\omega\colon I \to X$, $p\colon X \to I$, and $s\colon I \to I$ are respectively depicted as:

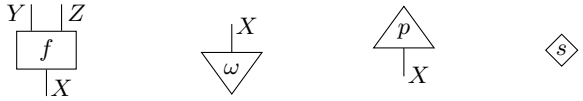

The identity channels are represented by 'no box', *i.e.* just wires, and the swap isomorphisms are represented by crossing of wires:

Finally, the sequential composition of channels is depicted by connecting the input and

output wires, and the parallel composition is given by juxtaposition, respectively as below:

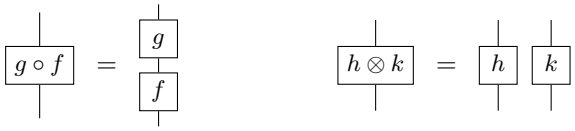

The use of string diagrams is justified by the following 'coherence' theorem; see (Joyal and Street, 1991; Selinger, 2010) for details.

Theorem 2.1. A well-formed equation between composites of arrows in a symmetric monoidal category follows from the axioms of symmetric monoidal categories if and only if the string diagrams of both sides are equal up to isomorphism of diagrams. □

We further assume the following structure in our category. For each type X there are a discarder $\overline{\top}_X \colon X \to I$ and a copier $\curlyvee_X \colon X \to X \otimes X$. They are required to satisfy the following equations:

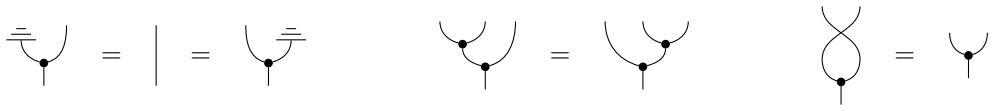

This says that $(\curlyvee_X, \overline{\top}_X)$ forms a commutative comonoid on X. By the associativity we may write:

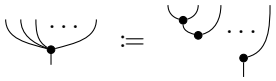

Moreover we assume that the comonoid structures $(\curlyvee_X, \overline{\top}_X)$ are compatible with the monoidal structure $(\otimes, I)$, in the sense that the following equations hold.

Note that we do *not* assume that these maps are natural. Explicitly, we do not necessarily have $\curlyvee \circ f = (f \otimes f) \circ \curlyvee$ or $\overline{\top} \circ f = \overline{\top}$.

We will use these symmetric monoidal categories throughout in the paper. For convenience, we introduce a term for them.

Definition 2.2. A *CD-category* is a symmetric monoidal category $(\mathbf{C}, \otimes, I)$ with a commutative comonoid $(\curlyvee_X, \overline{\top}_X)$ for each $X \in \mathbf{C}$, suitably compatible as described above.

Here 'CD' stands for Copy/Discard.

Definition 2.3. An arrow $f \colon X \to Y$ in a CD-category is said to be *causal* if

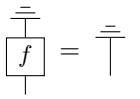

.

A CD-category is *affine* if all the arrows are causal, or equivalently, the tensor unit I is a final object.

The term 'causal' comes from (categorical) quantum foundation (Coecke and Kissinger, 2017; D'Ariano et al., 2017), and is related to relativistic causality, see *e.g.* (Coecke, 2016).

We reserve the term 'channel' for causal arrows. Explicitly, causal arrows $c\colon X \to Y$ in a CD-category are called *channels*. A channel $\omega\colon I \to X$ with input type I is called a *state* (on X). For the time being we will only use channels (and states), and thus only consider affine CD-categories. Non-causal arrows will not appear until Section 7.

Our main examples of affine CD-categories are two Kleisli categories $\mathcal{K}\ell(\mathcal{D})$ and $\mathcal{K}\ell(\mathcal{G})$, respectively, for discrete probability, and more general, measure-theoretic probability. We explain them in order below. There are more examples, like the Kleisli category of the non-empty powerset monad, or the (opposite of) the category of commutative C^*-algebras (with positive unital maps), but they are out of scope here.

Example 2.4. What we call a *distribution* or a *state* over a set X is a finite subset $\{x_1, x_2, \ldots, x_n\} \subseteq X$, called the support where each element x_i occurs with a multiplicity $r_i \in [0,1]$, such that $\sum_i r_i = 1$. Such a convex combination is often written as $r_1|x_1\rangle + \cdots + r_n|x_n\rangle$ with $r_i \in [0,1]$. The ket notation $|-\rangle$ is meaningless syntactic sugar that is used to distinguish elements $x \in X$ from occurrences in such formal sums. Notice that a distribution can also be written as a function $\omega\colon X \to [0,1]$ with finite support $\mathrm{supp}(\omega) = \{x \in X \mid \omega(x) \neq 0\}$. We shall write $\mathcal{D}(X)$ for the set of distributions over X. This $\mathcal{D}$ is a monad on the category **Set** of sets and functions.

A function $f\colon X \to \mathcal{D}(Y)$ is called a *Kleisli* map; it forms a channel $X \to Y$. Such maps can be composed as matrices, for which we use special notation $\circ$.

$$(g \circ f)(x)(z) = \textstyle\sum_{y \in Y} f(x)(y) \cdot g(y)(z) \qquad \text{for } g\colon Y \to \mathcal{D}(Z) \text{ with } x \in X, z \in Z.$$

We write $1 = \{*\}$ for a singleton set, and $2 = 1 + 1 = \{0,1\}$. Notice that $\mathcal{D}(1) \cong 1$ and $\mathcal{D}(2) \cong [0,1]$. We can identify a state on X with a channel $1 \to X$.

The monad $\mathcal{D}$ is known to be *commutative*. This implies that finite products of sets $X \times Y$ give rise to a symmetric monoidal structure on the Kleisli category $\mathcal{K}\ell(\mathcal{D})$. Specifically, for two maps $f\colon X \to \mathcal{D}(Y)$ and $g\colon Z \to \mathcal{D}(W)$, the tensor product / parallel composition $f \otimes g\colon X \times Z \to \mathcal{D}(Y \times W)$ is given by:

$$(f \otimes g)(x,z)(y,w) = f(x,y) \cdot g(z,w).$$

For each set X there are a copier $\curlyvee_X\colon X \to \mathcal{D}(X \times X)$ and a discarder $\overset{=}{\top}_X\colon X \to \mathcal{D}(1)$ given by $\curlyvee_X(x) = 1|x,x\rangle$ and $\overset{=}{\top}_X(x) = 1|*\rangle$, respectively. They come from the cartesian (finite product) structure of the base category **Set**, through the obvious functor $\mathbf{Set} \to \mathcal{K}\ell(\mathcal{D})$. Therefore $\mathcal{K}\ell(\mathcal{D})$ is a CD-category. It is moreover affine, since the monad is affine in the sense that $\mathcal{D}(1) \cong 1$.

Example 2.5. Let $X = (X, \Sigma_X)$ be a measurable space, where Σ_X is a σ-algebra on X. A *probability measure*, also called a *state*, on X is a function $\omega\colon \Sigma_X \to [0,1]$ which is countably additive and satisfies $\omega(X) = 1$. We write $\mathcal{G}(X)$ for the collection of all such probability measures on X. This set $\mathcal{G}(X)$ is itself a measurable space with the smallest σ-algebra such that for each $A \in \Sigma_X$ the 'evaluation' map $\mathrm{ev}_A\colon \mathcal{G}(X) \to [0,1]$, $\mathrm{ev}_A(\omega) = \omega(A)$, is measurable. Notice that $\mathcal{G}(X) \cong \mathcal{D}(X)$ when X is a finite set (as

discrete space). In particular, $\mathcal{G}(2) \cong \mathcal{D}(2) \cong [0,1]$. This $\mathcal{G}$ is a monad on the category **Meas** of measurable spaces, with measurable functions between them; it is called the Giry monad, after (Giry, 1982).

A Kleisli map, that is, a measurable function $f\colon X \to \mathcal{G}(Y)$ is a channel (or a *probability kernel*, see Example 7.2). These channels can be composed, via Kleisli composition $\odot$, using integration[†]:

$$(g \odot f)(x)(C) = \int_Y g(y)(C)\, f(x)(\mathrm{d}y) \quad \text{where } g\colon Y \to \mathcal{G}(Z) \text{ and } x \in X, C \in \Sigma_Z.$$

It is well-known that the monad $\mathcal{G}$ is commutative and affine, see also (Jacobs, 2017). Thus, in a similar manner to the previous example, the Kleisli category $\mathcal{K}\ell(\mathcal{G})$ is an affine CD-category. The parallel composition $f \otimes g\colon X \times Z \to \mathcal{G}(Y \times W)$ for $f\colon X \to \mathcal{G}(Y)$ and $g\colon Z \to \mathcal{G}(W)$ is given as:

$$(f \otimes g)(x,z)(B \times D) \;=\; f(x)(B) \cdot g(z)(D),$$

for $x \in X$, $z \in Z$, $B \in \Sigma_Y$, and $D \in \Sigma_W$. Since $f(x)$ and $g(z)$ are (σ-)finite measures, this indeed determines a unique measure $(f \otimes g)(x,z) \in \mathcal{G}(Y \times W)$, namely the unique product measure of $f(x)$ and $g(z)$. Explicitly, for $E \in \Sigma_{Y \times W}$,

$$(f \otimes g)(x,z)(E) = \int_W \left(\int_Y \mathbf{1}_E(y,w)\, f(x)(\mathrm{d}y) \right) g(z)(\mathrm{d}w)$$

where $\mathbf{1}_E$ is the indicator function.

3. Marginalisation, integration and disintegration

Let $\mathbf{C}$ be an affine CD-category. We think of states $\omega\colon I \to X$ in $\mathbf{C}$ as abstract (probability) distributions on type X. States of the form $\omega\colon I \to X \otimes Y$, often called (bipartite) joint states, are seen as joint distributions on X and Y. Later on we shall also consider n-partite joint states, but for the time being we restrict ourselves to bipartite ones. For a joint distribution $\mathrm{P}(x,y)$ in discrete probability, we can calculate the marginal distribution on X by summing (or marginalising) Y out, as $\mathrm{P}(x) = \sum_y \mathrm{P}(x,y)$. The marginal distribution on Y is also calculated by $\mathrm{P}(y) = \sum_x \mathrm{P}(x,y)$. In our abstract setting, given a joint state $\omega\colon I \to X \otimes Y$, we can obtain marginal states simply by discarding wires, as in:

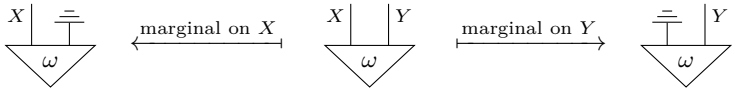

In other words, the marginal states are the state ω composed with the projection maps $\pi_1\colon X \otimes Y \to X$ and $\pi_2\colon X \otimes Y \to Y$, as below.

$\pi_1 :=$  $\qquad\qquad \pi_2 :=$ 

[†] We denote the integral of a function f with respect to a measure μ by $\int_X f(x)\, \mu(\mathrm{d}x)$.

Example 3.1. For a joint state $\omega \in \mathcal{D}(X \times Y)$ in $\mathcal{K}\ell(\mathcal{D})$, the first marginal $\omega_1 = \pi_1 \mathbin{\odot} \omega$ is given by $\omega_1(x) = \sum_{y \in Y} \omega(x, y)$, as expected. Similarly the second marginal is given by $\omega_2(y) = \sum_{x \in X} \omega(x, y)$.

Example 3.2. For a joint state $\omega \in \mathcal{G}(X \times Y)$ in $\mathcal{K}\ell(\mathcal{G})$, the first marginal is given by $\omega_1(A) = \omega(A \times Y)$ for $A \in \Sigma_X$, and the second marginal by $\omega_2(B) = \omega(X \times B)$ for $B \in \Sigma_Y$.

A channel $c \colon X \to Y$ is seen as an abstract *conditional* distribution $\mathrm{P}(y|x)$. In (discrete) probability theory, we can calculate a joint distribution $\mathrm{P}(x, y)$ from a distribution $\mathrm{P}(x)$ and a conditional distribution $\mathrm{P}(y|x)$ by the formula $\mathrm{P}(x, y) = \mathrm{P}(y|x) \cdot \mathrm{P}(x)$, which is often called the *product rule*. Similarly we have $\mathrm{P}(x, y) = \mathrm{P}(x|y) \cdot \mathrm{P}(y)$. In our setting, starting from a state $\sigma \colon I \to X$ and a channel $c \colon X \to Y$, or a state $\tau \colon I \to Y$ and a channel $d \colon Y \to X$, we can 'integrate' them into a joint state on $X \otimes Y$ as follows, respectively:

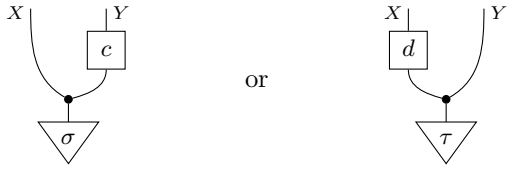

(2)

Example 3.3. Let $\sigma \in \mathcal{D}(X)$ and $c \colon X \to \mathcal{D}(Y)$ be a state and a channel in $\mathcal{K}\ell(\mathcal{D})$. An easy calculation verifies that $\omega = (\mathrm{id} \otimes c) \mathbin{\odot} \curlyvee \mathbin{\odot} \sigma$, the joint state on $X \times Y$ defined as in (2), satisfies $\omega(x, y) = c(x)(y) \cdot \sigma(x)$, as we expect from the product rule.

Example 3.4. For a state $\sigma \in \mathcal{G}(X)$ and a channel $c \colon X \to \mathcal{G}(Y)$ in $\mathcal{K}\ell(\mathcal{G})$, the joint state $\omega = (\mathrm{id} \otimes c) \mathbin{\odot} \curlyvee \mathbin{\odot} \sigma$ is given by $\omega(A \times B) = \int_A c(x)(B)\, \sigma(\mathrm{d}x)$ for $A \in \Sigma_X$ and $B \in \Sigma_Y$. This 'integration' construction of a joint probability measure is standard, see *e.g.* (Pollard, 2002; Panangaden, 2009).

Disintegration is an inverse operation of the 'integration' of a state and a channel into a joint state, as in (2). More specifically, it starts from a joint state $\omega \colon I \to X \otimes Y$ and extracts either a state $\omega_1 \colon I \to X$ and a channel $c_1 \colon X \to Y$, or a state $\omega_2 \colon I \to Y$ and a channel $c_2 \colon Y \to X$ as below,

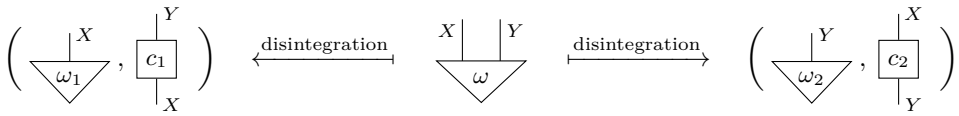

such that the equation on the left or right below holds, respectively.

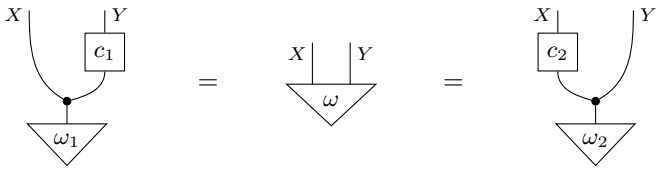

(3)

We immediately see from the equation that ω_1 and ω_2 must be marginals of ω:

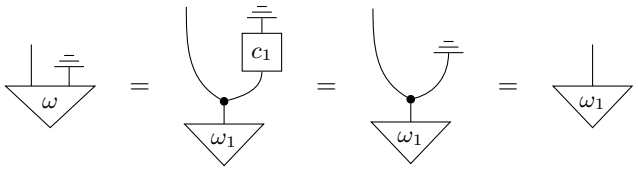

and similarly

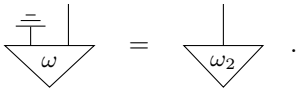

.

Therefore a disintegration of ω may be referred to by a channel c_i only, rather than a pair (ω_i, c_i). This leads to the following definition:

Definition 3.5. Let $\omega \colon I \to X \otimes Y$ be a joint state. A channel $c_1 \colon X \to Y$ (or $c_2 \colon Y \to X$) is called a *disintegration* of ω if it satisfies the equation (3) with ω_i the marginals of ω.

Let us look at concrete instances of disintegrations, in the Kleisli categories $\mathcal{K}\ell(\mathcal{D})$ and $\mathcal{K}\ell(\mathcal{G})$ from Examples 2.4 and 2.5.

Example 3.6. Let $\omega \in \mathcal{D}(X \times Y)$ be a joint state in $\mathcal{K}\ell(\mathcal{D})$. We write $\omega_1 \in \mathcal{D}(X)$ for the first marginal, given by $\omega_1(x) = \sum_y \omega(x, y)$. Then a channel $c \colon X \to \mathcal{D}(Y)$ is a disintegration of ω if and only if $\omega(x, y) = c(x)(y) \cdot \omega_1(x)$ for all $x \in X$ and $y \in Y$. It turns out that there is always such a channel c. We define a channel c by:

$$c(x)(y) \coloneqq \tfrac{\omega(x,y)}{\omega_1(x)} \qquad \text{if} \quad \omega_1(x) \neq 0 \ , \tag{4}$$

and $c(x) \coloneqq \tau$ if $\omega_1(x) = 0$, for an arbitrary state $\tau \in \mathcal{D}(Y)$. (Here $\mathcal{D}(Y)$ is nonempty, since so is $\mathcal{D}(X \times Y)$.) This indeed defines a channel c satisfying the required equation. Roughly speaking, disintegration in discrete probability is nothing but the 'definition' of conditional probability: $\mathrm{P}(y|x) = \mathrm{P}(x, y) / \mathrm{P}(x)$. There is still some subtlety — disintegrations need not be unique, when there are $x \in X$ with $\omega_1(x) = 0$. They are, nevertheless, 'almost surely' unique; see Section 5.

Example 3.7. Disintegrations in measure-theoretic probability, in $\mathcal{K}\ell(\mathcal{G})$, are far more difficult. Let $\omega \in \mathcal{G}(X \times Y)$ be a joint state, with $\omega_1 \in \mathcal{G}(X)$ the first marginal. A channel $c \colon X \to \mathcal{G}(Y)$ is a disintegration of ω if and only if

$$\omega(A \times B) = \int_A c(x)(B)\, \omega_1(\mathrm{d}x)$$

for all $A \in \Sigma_X$ and $B \in \Sigma_Y$. This is the ordinary notion of *disintegration* (of probability measures), also known as *regular conditional probability*; see *e.g.* (Faden, 1985; Pollard, 2002; Panangaden, 2009). We see that there is no obvious way to obtain a channel c here, unlike the discrete case. In fact, a disintegration may not exist (Stoyanov, 2014). There are, however, a number of results that guarantee the existence of a disintegration in certain situations. We will come back to this issue later in the section.

Bayesian inversion is a special form of disintegration, occurring frequently. We start

from a state $\sigma\colon I \to X$ and a channel $c\colon X \to Y$. We then integrate them into a joint state on X and Y, and disintegrate it in the other direction, as below.

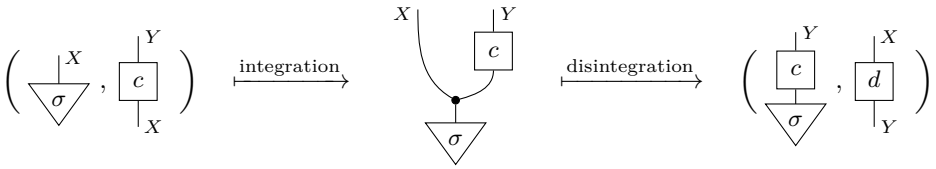

We call the disintegration $d\colon Y \to X$ a *Bayesian inversion* for $\sigma\colon I \to X$ along $c\colon X \to Y$. By unfolding the definitions, a channel $d\colon Y \to X$ is a Bayesian inversion if and only if

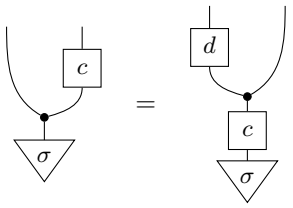

(5)

The composite $c \circ \sigma$ is also written as $c_*(\sigma)$. The operation $\sigma \mapsto c_*(\sigma)$, called *state transformation*, is used to explain forward inference in (Jacobs and Zanasi, 2016).

Example 3.8. Let $\sigma \in \mathcal{D}(X)$ and $c\colon X \to \mathcal{D}(Y)$ be a state and a channel in $\mathcal{K}\ell(\mathcal{D})$. Then a channel $d\colon Y \to \mathcal{D}(X)$ is a Bayesian inversion for σ along c if and only if $c(x)(y) \cdot \sigma(x) = d(y)(x) \cdot c_*(\sigma)(y)$, where $c_*(\sigma)(y) = \sum_{x'} c(x')(y) \cdot \sigma(x')$. In a similar manner to Example 3.6, we can obtain such a d by:

$$d(y)(x) \coloneqq \frac{c(x)(y) \cdot \sigma(x)}{c_*(\sigma)(y)} = \frac{c(x)(y) \cdot \sigma(x)}{\sum_{x'} c(x')(y) \cdot \sigma(x')}$$

for $y \in Y$ with $c_*(\sigma)(y) \neq 0$. For $y \in Y$ with $c_*(\sigma)(y) = 0$, we may define $d(y)$ to be an arbitrary state in $\mathcal{D}(X)$. We can recognise the above formula as the Bayes formula:

$$\mathrm{P}(x|y) = \frac{\mathrm{P}(y|x) \cdot \mathrm{P}(x)}{\mathrm{P}(y)} = \frac{\mathrm{P}(y|x) \cdot \mathrm{P}(x)}{\sum_{x'} \mathrm{P}(y|x') \cdot \mathrm{P}(x')} \ .$$

Example 3.9. Let $\sigma \in \mathcal{G}(X)$ and $c\colon X \to \mathcal{G}(Y)$ be a state and a channel in $\mathcal{K}\ell(\mathcal{G})$. A channel $d\colon Y \to \mathcal{G}(X)$ is a Bayesian inversion if and only if

$$\int_A c(x)(B)\, \sigma(\mathrm{d}x) = \int_B d(y)(A)\, c_*(\sigma)(\mathrm{d}y)$$

for all $A \in \Sigma_X$ and $B \in \Sigma_Y$. Here $c_*(\sigma) \in \mathcal{G}(Y)$ is the measure given by $c_*(\sigma)(B) = \int_X c(x)(B)\, \sigma(\mathrm{d}x)$. As we see below, Bayesian inversions are in some sense equivalent to disintegrations, and thus, they are as difficult as disintegrations. In particular, a Bayesian inversion need not exist.

In practice, however, the state σ and channel c are often given via density functions. This setting, so-called (absolutely) continuous probability, makes it easy to compute a Bayesian inversion. Suppose that X and Y are subspaces of $\mathbb{R}$, and that σ and c admit

density functions as

$$\sigma(A) = \int_A f(x)\,\mathrm{d}x \qquad\qquad c(x)(B) = \int_B \ell(x,y)\,\mathrm{d}y$$

for measurable functions $f\colon X \to \mathbb{R}_{\geq 0}$ and $\ell\colon X \times Y \to \mathbb{R}_{\geq 0}$. The conditional probability density $\ell(x,y)$ of y given x is often called the *likelihood* of x given y. By the familiar Bayes formula for densities — see *e.g.* (Bernardo and Smith, 2000) — the conditional density of x given y is:

$$k(y,x) \coloneqq \frac{\ell(x,y)\cdot f(x)}{\int_X \ell(x',y)\cdot f(x')\,\mathrm{d}x'} \quad .$$

This k then gives a channel $d\colon Y \to \mathcal{G}(X)$ by

$$d(y)(A) = \int_A k(y,x)\,\mathrm{d}x$$

for each $y \in Y$ such that $\int_X \ell(x',y)\cdot f(x')\,\mathrm{d}x' \neq 0$. For the other y's we define $d(y)$ to be some fixed state in $\mathcal{G}(X)$. An elementary calculation verifies that d is indeed a Bayesian inversion for σ along c. Later, in Section 8, we generalise this calculation into our abstract setting.

Although Bayesian inversions are a special case of disintegrations, we can conversely obtain disintegrations from Bayesian inversions, as in the proposition below. Therefore, in some sense the two notions are equivalent.

Proposition 3.10. Let ω be a state on $X \otimes Y$. Let $d\colon X \to X \otimes Y$ be a Bayesian inversion for ω along the first projection $\pi_1\colon X \otimes Y \to X$ on the left below.

$\pi_1 =$  $\qquad\qquad \pi_2 \circ d =$ 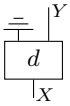

Then the composite $\pi_2 \circ d\colon X \to Y$ shown on the right above is a disintegration of ω.

Proof. We prove that the first equation in (3) holds for $c_1 = \pi_2 \circ d$, as follows.

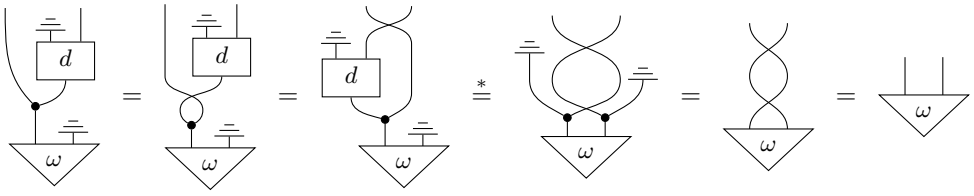

For the marked equality $\overset{*}{=}$ we used the equation (5) for the Bayesian inversion d. □

We say that an affine CD-category $\mathbf{C}$ *admits disintegration* if for every bipartite state $\omega\colon I \to X \otimes Y$ there exist a disintegration $c_1\colon X \to Y$ of ω. Note that in such categories there also exists a disintegration $c_2\colon Y \to X$ of ω in the other direction, since it can be

obtained as a disintegration of the following state:

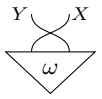

.

By Proposition 3.10, admitting disintegration is equivalent to admitting Bayesian inversion.

In Example 3.6, we have seen that $\mathcal{K}\ell(\mathcal{D})$ admits disintegration, but that in measure-theoretic probability, in $\mathcal{K}\ell(\mathcal{G})$, disintegrations may not exist. There are however a number of results that guarantee the existence of disintegrations in specific situations, see *e.g.* (Pachl, 1978; Faden, 1985). We here invoke one of these results and show that there is a subcategory of $\mathcal{K}\ell(\mathcal{G})$ that admits disintegration. A measurable space is called a *standard Borel space* if it is measurably isomorphic to a Polish space with its Borel σ-algebra, or equivalently, if it is measurably isomorphic to a Borel subspace of $\mathbb{R}$. Then the following theorem is standard, see *e.g.* (Pollard, 2002, §5.2) or (Faden, 1985, §5).

Theorem 3.11. Let X be any measurable space and Y be a standard Borel space. Then for any state (*i.e.* a probability measure) $\omega \in \mathcal{G}(X \times Y)$ in $\mathcal{K}\ell(\mathcal{G})$, there exists a disintegration $c_1 \colon X \to \mathcal{G}(Y)$ of ω. □

Let $\mathbf{pKrn}_{\mathrm{sb}}$ be the full subcategory of $\mathcal{K}\ell(\mathcal{G})$ consisting of standard Borel spaces as objects. It is not hard to see that the product of two standard Borel spaces is a standard Borel space. It follow that $\mathbf{pKrn}_{\mathrm{sb}}$ is an affine CD-category. Then the previous theorem immediately shows:

Corollary 3.12. The category $\mathbf{pKrn}_{\mathrm{sb}}$ admits disintegration. □

We note that $\mathbf{pKrn}_{\mathrm{sb}}$ can also be seen as the Kleisli category of the Giry monad restricted on the category of standard Borel spaces, see e.g. (Giry, 1982; Doberkat, 2009).

Since there are various 'existence' theorems like Theorem 3.11, there may be other subcategories of $\mathcal{K}\ell(\mathcal{G})$ that admit disintegration. A likely candidate is the category of perfect probabilistic mappings in (Culbertson and Sturtz, 2014). We do not go into this question here, since $\mathbf{pKrn}_{\mathrm{sb}}$ suffices for the present paper.

4. Example: naive Bayesian classifiers via inversion

Bayesian classification is a well-known technique in machine learning that produces a distribution over data classifications, given certain sample data. The distribution describes the probability, for each data (classification) category, that the sample data is in that category. Here we consider an example of 'naive' Bayesian classification, where the features are assumed to be independent. We consider a standard classification example from the literature which forms an ideal setting to illustrate the use of both disintegration and Bayesian inversion. Disintegration is used to extract channels from a given table, and subsequently Bayesian inversion is applied to (the tuple of) these channels to obtain the actual classification. The use of channels and disintegration/inversion in this classification setting is new, as far as we know.

For the description of the relevant operations in this example we use notation for

marginalisation and disintegration that we borrowed from the *EfProb* library (Cho and Jacobs, 2017). There are many ways to marginalise an n-partite state, namely one for each subset of the wires $\{1, 2, \ldots, n\}$. Such a subset can be described as a *mask*, consisting of a list of n zero's or one's, where a zero at position i means that the i-th wire/component is marginalised out, and a one at position i means that it remains. Such a mask $M = [b_1, \ldots, b_n]$ with $b_i \in \{0, 1\}$ is used as a post-fix selection operation in $\omega.M$ on an n-partite state ω. An example explains it all:

if $\omega =$ 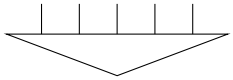 then $\omega.[1, 0, 1, 0, 0] =$ 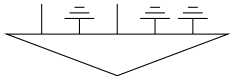

In a similar way one can disintegrate an n-partite state in 2^n may ways, where a mask of length n is now used to describe which wires are used as input to the extracted channel and which ones as output. We write $\omega /\!\!/ M$ for such a disintegration, where M is a mask, as above. A systematic description will be given in Section 6 below.

In practice it is often useful to be able to marginalise first, and disintegrate next. The general description in n-ary form is a bit complicated, so we use an example for $n = 5$. We shall label the wires with x_i, as on the left below. We seek the conditional probability written conventionally as $c = \omega[x_1, x_4 \mid x_2, x_5]$ on the right below.

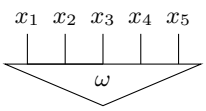

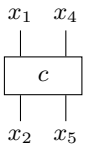

This channel c must satisfy:

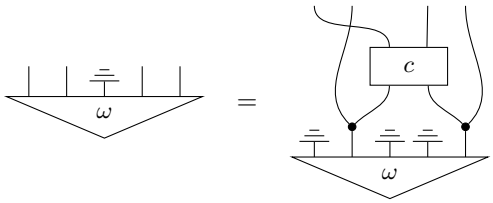

This picture shows how to obtain the channel c from ω: we first marginalise to restrict to the relevant wires x_1, x_2, x_4, x_5. This is written as $\omega.[1, 1, 0, 1, 1]$. Subsequently we disintegrate with x_1, x_4 as output and x_2, x_5 as input. Hence:

$$\begin{aligned} c &\coloneqq \omega.[1, 1, 0, 1, 1] /\!\!/ [0, 1, 0, 1] \\ &=: \omega.\big[\, [1, 0, 0, 1, 0] \;\big|\; [0, 1, 0, 0, 1] \,\big] \quad \text{as we shall write in the sequel.} \end{aligned}$$

We see that the latter post-fix $\big[[1, 0, 0, 1, 0] \,\big|\, [0, 1, 0, 0, 1]\big]$ is a 'variable free' version of the traditional notation $[x_1, x_4 \mid x_2, x_5]$, selecting the relevant positions.

We have now prepared the ground and can turn to the classification example that we announced. It involves the classification of 'playing' (yes or no) for certain weather data, used in (Witten et al., 2011). We shall first go through the discrete example in some detail. The relevant data are in the table in Figure 1. The question is: given this table, what can

Outlook	**Temperature**	**Humidity**	**Windy**	**Play**
Sunny	hot	high	false	no
Sunny	hot	high	true	no
Overcast	hot	high	false	yes
Rainy	mild	high	false	yes
Rainy	cool	normal	false	yes
Rainy	cool	normal	true	no
Overcast	cool	normal	true	yes
Sunny	mild	high	false	no
Sunny	cool	normal	false	yes
Rainy	mild	normal	false	yes
Sunny	mild	normal	true	yes
Overcast	mild	high	true	yes
Overcast	hot	normal	false	yes
Rainy	mild	high	true	no

Fig. 1. Weather and play data, copied from (Witten et al., 2011).

be said about the probability of playing if the outlook is *Sunny*, the temperature is *Cold*, the humidity is *High* and it is Windy?

Our plan is to first organise these table data into four channels d_O, d_T, d_H, d_W in a network of the form:

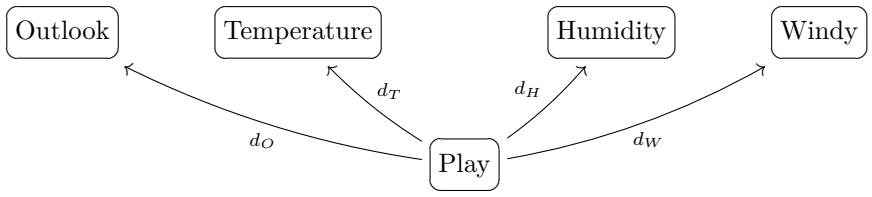

(6)

We start by extracting the underlying sets for for the categories in the table in Figure 1. We choose abbreviations for the entries in each of the categories.

$$O = \{s, o, r\} \qquad W = \{t, f\} \qquad T = \{h, m, c\} \qquad P = \{y, n\} \qquad H = \{h, n\}.$$

These sets are combined into a single product domain:

$$D = O \times T \times H \times W \times P.$$

It combines the five columns in Figure 1. The table itself in the figure is represented as a uniform distribution $\tau \in \mathcal{D}(D)$. This distribution has 14 entries — like in the table — and looks as follows.

$$\tau = \tfrac{1}{14}|s, h, h, f, n\rangle + \tfrac{1}{14}|s, h, h, t, n\rangle + \cdots + \tfrac{1}{14}|r, m, h, t, n\rangle.$$

We extract the four channels in Diagram (6) via appropriate disintegrations, from the Play column to the Outlook / Temperature / Humidity / Windy columns.

$$d_O = \tau.\big[\,[1,0,0,0,0] \;\big|\; [0,0,0,0,1]\,\big] \qquad d_T = \tau.\big[\,[0,1,0,0,0] \;\big|\; [0,0,0,0,1]\,\big]$$
$$d_H = \tau.\big[\,[0,0,1,0,0] \;\big|\; [0,0,0,0,1]\,\big] \qquad d_W = \tau.\big[\,[0,0,0,1,0] \;\big|\; [0,0,0,0,1]\,\big].$$

Thus, as described in the beginning of this section, the 'outlook' channel $d_O \colon P \to \mathcal{D}(O)$ is extracted by first marginalising the table τ to the relevant (first and last) wires, and then disintegrating. Explicitly, d_O is $\tau.[1,0,0,0,1] \mathbin{/\!\!/} [0,1]$ and satisfies:

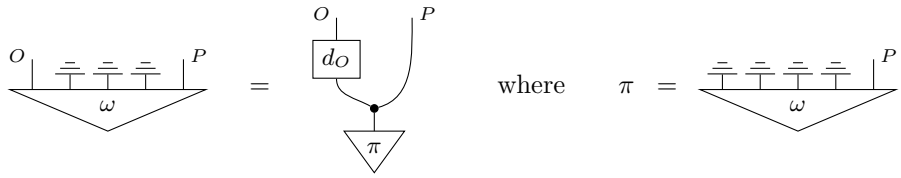

In a next step we combine these four channels into a single channel $d \colon P \to O \times T \times H \times W$ via tupling:

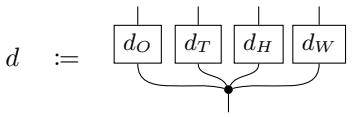

The answer that we are looking for will be obtained by Bayesian inversion of this channel d wrt. the above fifth marginal Play state $\pi = \tau.[0,0,0,0,1] = \frac{9}{14}|y\rangle + \frac{5}{14}|n\rangle \in \mathcal{D}(P)$. We write this inversion as a channel $e \colon O \times T \times H \times W \to P$. It satisfies, by construction, according to the pattern in (5):

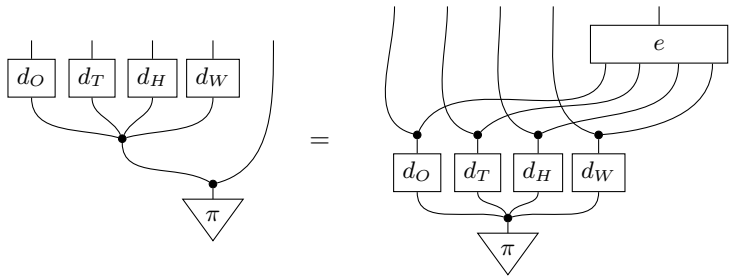

We can finally answer the original classification question. The assumptions — *Sunny* outlook, *Cold* temperature, *High* humidity, *true* windiness — are used as input to the inversion channel e. This yields the probability of play that we are looking for:

$$e(s, c, H, t) = 0.205|y\rangle + 0.795|n\rangle$$

The resulting classification probability[‡] of 0.205 coincides with the probability of 20.5% that is computed in (Witten et al., 2011) — without the above systematic approach.

One could complain that our approach is 'too' abstract, since it remains implicit what these extracted channels do. We elaborate the outlook channel $d_O \colon P \to O$. For the two elements in $P = \{y, n\}$ we have:

$$d_O(y) = \tfrac{2}{9}|s\rangle + \tfrac{4}{9}|o\rangle + \tfrac{3}{9}|r\rangle \qquad\qquad d_O(n) = \tfrac{3}{5}|s\rangle + 0|o\rangle + \tfrac{2}{5}|r\rangle.$$

These outcomes arise from the general formula (4). But we can also understand them at a

[‡] This outcome has been calculated with the tool *EfProb* (Cho and Jacobs, 2017) that can do both disintegration and inversion.

more concrete level: for the first distribution $d_O(y)$ we need to concentrate on the 9 lines in Figure 1 for which Play is *yes*; in these lines, in the first Outlook column, 2 out of 9 entries are *Sunny*, 4 out of 9 are *Overcast*, and 3 out of 9 are *Rainy*. This corresponds to the above first distribution $d_O(y)$. Similarly, the second distribution $d_O(n)$ captures the Outlook for the 5 lines where Play is *no*: 3 out of 5 are *Sunny* and 2 out of 5 are *Rainy*.

5. Almost equality of channels

This section explains how the standard notion of 'equal up to negligible sets' or 'equal almost everywhere' (with respect to a measure) can be expressed abstractly using string diagrams. Via this equality relation Bayesian inversion can be characterised very neatly, following (Clerc et al., 2017). We consider an affine CD-category, continuing in the setting of Section 3.

Definition 5.1. Let $c, d\colon X \to Y$ be two parallel channels, and $\sigma\colon I \to X$ be a state on their domain. We say that c is *σ-almost equal* to d, written as $c \overset{\sigma}{\sim} d$ if

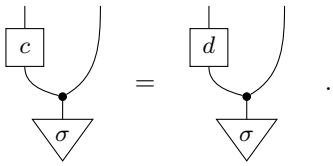

It is obvious that $\overset{\sigma}{\sim}$ is an equivalence relation on channels of type $X \to Y$. When S is a set of arrows of type $X \to Y$, we write S/σ for the quotient $S/\overset{\sigma}{\sim}$.

To put it more intuitively, we have $c \overset{\sigma}{\sim} d$ iff c and d can be identified whenever the input wires are connected to σ, possibly through copiers. For instance, using the associativity and commutativity of copiers, by $c \overset{\sigma}{\sim} d$ we may reason as:

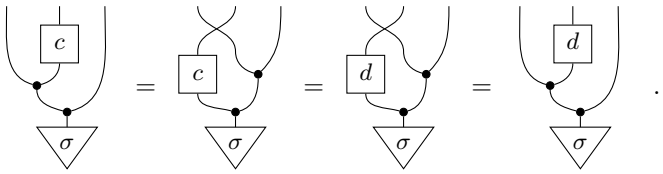

In particular, $c \overset{\sigma}{\sim} d$ if and only if

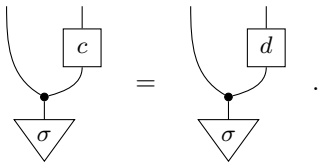

Now the following is an obvious consequence from the definition.

Proposition 5.2. If both $c, d\colon X \to Y$ are disintegrations of a joint state $\omega\colon I \to X \otimes Y$, then $c \overset{\omega_1}{\sim} d$, where $\omega_1\colon I \to X$ is the first marginal of ω. □

For channels $f, g\colon X \to \mathcal{D}(Y)$ and a state $\sigma \in \mathcal{D}(X)$ in $\mathcal{K}\ell(\mathcal{D})$, it is easy to see that

$f \overset{\sigma}{\sim} g$ if and only if $f(x)(y) \cdot \sigma(x) = g(x)(y) \cdot \sigma(x)$ for all $x \in X$ and $y \in Y$ if and only if $f(x) = g(x)$ for any $x \in X$ with $\sigma(x) \neq 0$. Almost equality in $\mathcal{K}\ell(\mathcal{G})$ is less trivial but characterised in an expected way.

Proposition 5.3. Let $f, g \colon X \to \mathcal{G}(Y)$ be channels and $\mu \in \mathcal{G}(X)$ a state in $\mathcal{K}\ell(\mathcal{G})$. Then $f \overset{\mu}{\sim} g$ if and only if for any $B \in \Sigma_Y$, $f(-)(B) = g(-)(B)$ μ-almost everywhere.

Proof. By expanding the definition, $f \overset{\mu}{\sim} g$ if and only if

$$\int_A f(x)(B)\, \mu(\mathrm{d}x) = \int_A g(x)(B)\, \mu(\mathrm{d}x)$$

for all $A \in \Sigma_X$ and $B \in \Sigma_Y$. This is equivalent to $f(-)(B) = g(-)(B)$ μ-almost everywhere for all $B \in \Sigma_Y$, see (Fremlin, 2000, 131H). □

Almost-everywhere equality of probability kernels $f, g \colon X \to \mathcal{G}(Y)$ is often formulated by the stronger condition that $f = g$ μ-almost everywhere. The next proposition shows that the stronger variant is equivalent under a reasonable assumption (any standard Borel space is countably generated, for example).

Proposition 5.4. In the setting of the previous proposition, additionally assume that the measurable space Y is countably generated. Then $f \overset{\mu}{\sim} g$ if and only if $f = g$ μ-almost everywhere.

Proof. Let the σ-algebra Σ_Y on Y be generated by a countable family $(B_n)_n$. We may assume that $(B_n)_n$ is a π-system, *i.e.* a family closed under binary intersections. Let $A_n = \{x \in X \mid f(x)(B_n) = g(x)(B_n)\}$, and $A = \bigcap_n A_n$. Each A_n is μ-conegligible, and thus A is μ-conegligible. For each $x \in A$, we have $f(x)(B_n) = g(x)(B_n)$ for all n. By application of the Dynkin π-λ theorem, it follows that $f(x) = g(x)$. Therefore $f = g$ μ-almost everywhere. □

We can now present a fundamental result from (Clerc et al., 2017, §3.3) in our abstract setting. Let $\mathbf{C}$ be an affine CD-category and $(I \downarrow \mathbf{C})$ be the comma (coslice) category. Objects in $(I \downarrow \mathbf{C})$ are states in $\mathbf{C}$, formally pairs (X, σ) of objects $X \in \mathbf{C}$ and states $\sigma \colon I \to X$. Arrows from (X, σ) to (Y, τ) are *state-preserving* channels, namely $c \colon X \to Y$ in $\mathbf{C}$ satisfying $c \circ \sigma = \tau$. A joint state $(X \otimes Y, \omega) \in (I \downarrow \mathbf{C})$ is called a *coupling* of two states $(X, \sigma), (Y, \tau) \in (I \downarrow \mathbf{C})$ if

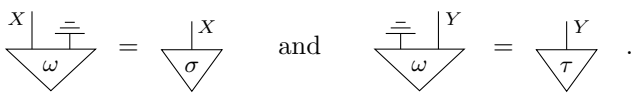

We write $\mathrm{Coupl}((X, \sigma), (Y, \tau))$ for the set of couplings of (X, σ) and (Y, τ).

Theorem 5.5. Let $\mathbf{C}$ be an affine CD-category that admits disintegration. For each pair of states $(X, \sigma), (Y, \tau) \in (I \downarrow \mathbf{C})$, there is the following bijection:

$$(I \downarrow \mathbf{C})\big((X, \sigma), (Y, \tau)\big)/\sigma \;\cong\; \mathrm{Coupl}((X, \sigma), (Y, \tau))$$

Proof. For each $c \in (I \downarrow \mathbf{C})\big((X, \sigma), (Y, \tau)\big)$, we define a joint state $I \to X \otimes Y$ to be

the 'integration' of σ and c as below, for which we use the following *ad hoc* notation:

$\sigma \circledcirc c \coloneqq$ 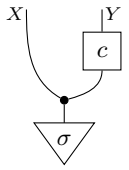

It is easy to check that $\sigma \circledcirc c$ is a coupling of σ and τ. For two channels $c, d\colon X \to Y$, we have $\sigma \circledcirc c = \sigma \circledcirc d$ if and only if $c \overset{\sigma}{\sim} d$, by the definition of $\overset{\sigma}{\sim}$. This means the mapping

$$c \longmapsto \sigma \circledcirc c\,, \quad (I \downarrow \mathbf{C})\big((X,\sigma),(Y,\tau)\big)/\sigma \longrightarrow \mathrm{Coupl}((X,\sigma),(Y,\tau))$$

is well-defined and injective. To prove the surjectivity let $(X \otimes Y, \omega) \in \mathrm{Coupl}((X,\sigma),(Y,\tau))$. Let $c\colon X \to Y$ be a disintegration of ω. Then c is state-preserving since

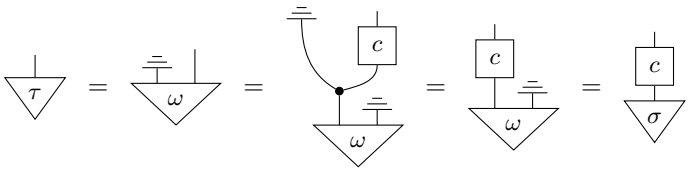

Moreover we have $\sigma \circledcirc c = \omega$, as desired. □

Via the symmetry $X \otimes Y \overset{\cong}{\to} Y \otimes X$ we have the obvious bijection $\mathrm{Coupl}((X,\sigma),(Y,\tau)) \cong \mathrm{Coupl}((Y,\tau),(X,\sigma))$. This immediately gives the following corollary.

Corollary 5.6. Let $\mathbf{C}$ be an affine CD-category that admits disintegration. For any states $(X,\sigma),(Y,\tau) \in (I \downarrow \mathbf{C})$ we have

$$(I \downarrow \mathbf{C})\big((X,\sigma),(Y,\tau)\big)/\sigma \;\cong\; (I \downarrow \mathbf{C})\big((Y,\tau),(X,\sigma)\big)/\tau$$

The bijection sends a channel $c\colon X \to Y$ to a Bayesian inversion $d\colon Y \to X$ for σ along c. □

Theorem 2 of (Clerc et al., 2017) is obtained as an instance, for the category $\mathbf{pKrn}_{\mathrm{sb}}$. This bijective correspondence yields a 'dagger' $(-)^\dagger$ functor on (a suitable quotient of) the comma category $(I \downarrow \mathbf{C})$ — as noted by the authors of (Clerc et al., 2017).

Strong almost equality

Unfortunately, the almost equality defined above is not the most useful notion for equational reasoning between string diagrams. Indeed, later in the proofs of Proposition 6.4 and Theorem 8.3 we encounter situations where we need a stronger notion of almost equality. We define the stronger one as follows.

Definition 5.7. Let $c, d\colon X \to Y$ be channels and $\sigma\colon I \to X$ be a state. We say that c is *strongly σ-almost equal* to d if

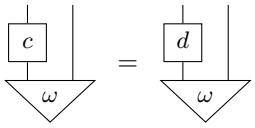

holds for any $\omega\colon I \to X \otimes Z$ such that σ is the first marginal of ω, that is:

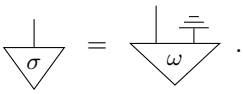

Notice that strong almost quality implies almost quality, via the following ω:

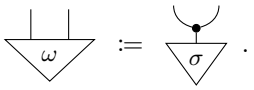

The good news is that the converse often holds too. We say that an affine CD-category admits *equality strengthening* if c is strongly σ-almost equal to d whenever $c \overset{\sigma}{\sim} d$, i.e. c is σ-almost equal to d. We present two propositions that guarantee this property.

Proposition 5.8. If an affine CD-category admits disintegration, then it admits equality strengthening.

Proof. Suppose that $c \overset{\sigma}{\sim} d$ holds for a state $\sigma\colon I \to X$ and for channels $c, d\colon X \to Y$. Let $\omega\colon I \to X \otimes Z$ be a state whose first marginal is σ. We can disintegrate ω as in:

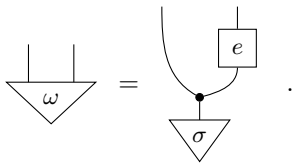

Then we have

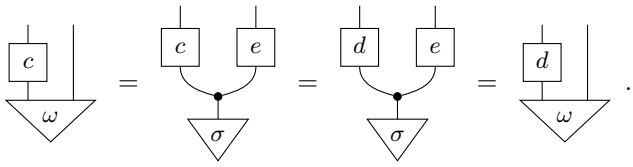

Therefore c is strongly σ-almost equal to d. □

Recall that $\mathcal{K}\ell(\mathcal{G})$ does not admit disintegration. Nevertheless, it admits equality strengthening.

Proposition 5.9. The category $\mathcal{K}\ell(\mathcal{G})$ admits equality strengthening.

Proof. Assume $c \overset{\sigma}{\sim} d$ for $\sigma \in \mathcal{G}(X)$ and $c, d\colon X \to \mathcal{G}(Y)$. Let $\omega \in \mathcal{G}(X \otimes Z)$ be a probability measure whose first marginal is σ, i.e. $(\pi_1)_*(\omega) = \sigma$. We need to prove

$(c \otimes \eta_Z) \circ\!\circ \omega = (d \otimes \eta_Z) \circ\!\circ \omega$, which is equivalent to:

$$\int_{X\times Z} c(x)(B)\mathbf{1}_C(z)\,\omega(\mathrm{d}(x,z)) = \int_{X\times Z} d(x)(B)\mathbf{1}_C(z)\,\omega(\mathrm{d}(x,z)) \tag{7}$$

for all $B \in \Sigma_Y$ and $C \in \Sigma_Z$. Using

$$\big|c(x)(B)\mathbf{1}_C(z) - d(x)(B)\mathbf{1}_C(z)\big| = \big|c(x)(B) - d(x)(B)\big|\mathbf{1}_C(z) \leq \big|c(x)(B) - d(x)(B)\big| \ ,$$

we have

$$\begin{aligned}
&\Big|\int_{X\times Z} \big(c(x)(B)\mathbf{1}_C(z) - d(x)(B)\mathbf{1}_C(z)\big)\,\omega(\mathrm{d}(x,z))\Big| \\
&\leq \int_{X\times Z} \big|c(x)(B)\mathbf{1}_C(z) - d(x)(B)\mathbf{1}_C(z)\big|\,\omega(\mathrm{d}(x,z)) \\
&\leq \int_{X\times Z} \big|c(x)(B) - d(x)(B)\big|\,\omega(\mathrm{d}(x,z)) \\
&= \int_X \big|c(x)(B) - d(x)(B)\big|\,(\pi_1)_*(\omega)(\mathrm{d}x) \\
&= \int_X \big|c(x)(B) - d(x)(B)\big|\,\sigma(\mathrm{d}x) \\
&= 0 \ ,
\end{aligned}$$

where the last equality holds since $c(-)(B) = d(-)(B)$ σ-almost everywhere by Proposition 5.3. Therefore the desired equality (7) holds. □

6. Conditional independence

Throughout this section, we consider an affine CD-category that admits disintegration.

6.1. *Disintegration of multipartite states*

So far we have concentrated on bipartite states — except in the classification example in Section 4. In order to deal with a general n-partite state $\omega\colon I \to X_1 \otimes \cdots \otimes X_n$, we will introduce several notations and conventions in Definitions 6.1, 6.2 and 6.3 below; they are in line with standard practice in probability theory.

In the conventions, an n-partite state, as below, is fixed, and used implicitly.

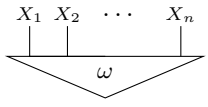

Definition 6.1. When we write

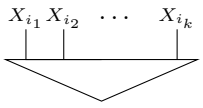

where $i_1, \ldots, i_k$ are distinct, it denotes the state $I \to X_{i_1} \otimes \cdots \otimes X_{i_k}$ obtained from ω by

marginalisation and permutation of wires (if necessary). Let us give a couple of examples, for $n = 5$.

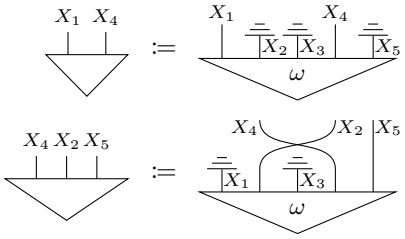

We permute wires via a combination of crossing. This is unambiguous by the coherence theorem.

Below we will use symbols $X, Y, Z, W, \ldots$ to denote not only a single wire X_i but also multiple wires $X_i \otimes X_j \otimes \cdots$. Disintegrations more general than in the bipartite case are now introduced as follows.

Definition 6.2. For $X = X_{i_1} \otimes \cdots \otimes X_{i_k}$ and $Y = X_{j_1} \otimes \cdots \otimes X_{j_l}$, with all $i_1, \ldots, i_k, j_1, \ldots, j_l$ distinct, a disintegration $X \to Y$ is defined to be a disintegration of

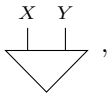

,

the marginal state given by the previous convention. We denote the disintegration simply as on the left below,

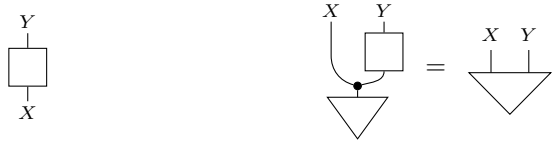

By definition, it must satisfy the equation on the right above. Let us give an example. The disintegration $X_5 \otimes X_2 \to X_1 \otimes X_4$ on the left below is defined by the equation on the right.

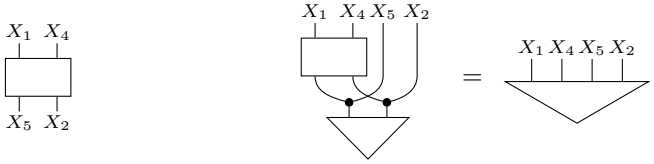

More specifically, assuming $n = 5$ and expanding the notation for marginals, the equation

is:

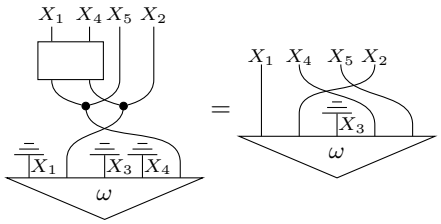

Note that disintegrations need not be unique. Thus when we write 

, we in fact *choose* one of them. Nevertheless, such disintegrations are unique up to almost-equality with respect to 

, which is good enough for our purpose.

Finally we make a convention about almost equality (Definition 5.1).

Definition 6.3. Let S and T be string diagrams of type $X \to Y$ that are made from marginals and disintegrations of ω as defined in Definitions 6.1 and 6.2. When we say S is almost equal to T (or write $S \sim T$) without reference to a state, it means that S is almost equal to T with respect to the state 

.

We shall make use of the following auxiliary equations involving discarding and composition of disintegrations.

Proposition 6.4. In the conventions and notations above, the following hold.

1 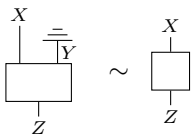

2 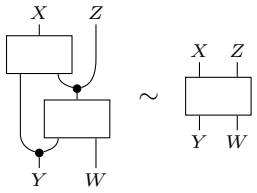

Proof. By the definition of almost equality, 1 is proved by:

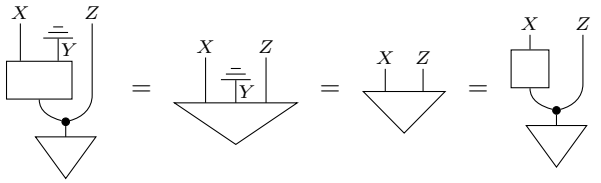

.

Similarly, we prove 2 as follows.

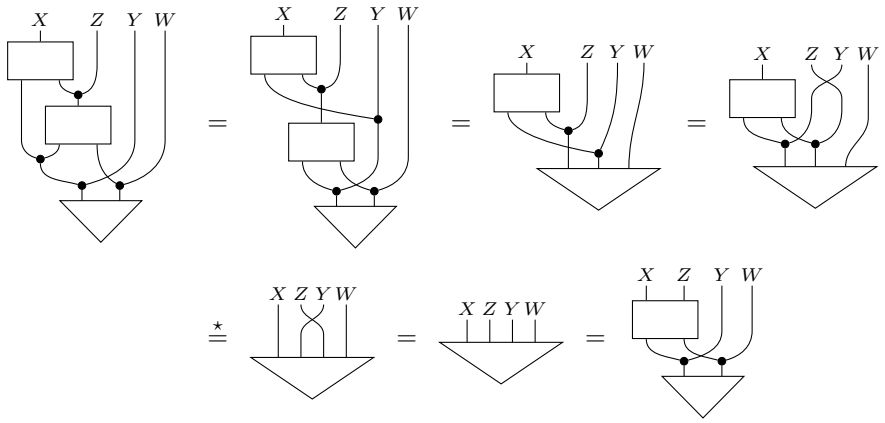

The marked equality $\stackrel{\star}{=}$ holds by strong almost equality, see Proposition 5.8. □

The equations correspond respectively to $\sum_y \mathrm{P}(x, y|z) = \mathrm{P}(x|z)$ and $\mathrm{P}(x|y, z) \cdot \mathrm{P}(z|y, w) = \mathrm{P}(x, z|y, w)$ in discrete probability.

Remark 6.5. We here keep our notation somewhat informal, e.g. using symbols $X, Y, Z, \ldots$ as a sort of meta-variables. We refer to (Joyal and Street, 1991; Selinger, 2010) for more formal aspects of string diagrams.

6.2. *Conditional independence*

We continue using the notations in the previous subsection. Recall that we fix an n-partite state

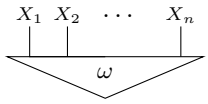

and use symbols $X, Y, Z, W, \ldots$ to denote a wire X_i or multiple wires $X_i \otimes X_j \otimes \cdots$.

We now introduce the notion of conditional independence. Although it is defined with respect to the underlying state ω, we leave the state ω implicit, like an underlying probability space Ω in conventional probability theory.

Definition 6.6. Let X, Y, Z denote distinct wires. Then we say X and Y are conditionally independent given Z, written as $X \perp\!\!\!\perp Y \mid Z$, if

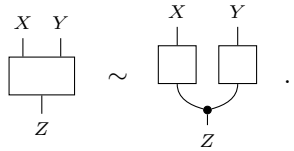

.

The definition is analogous to the condition $\mathrm{P}(x, y|z) = \mathrm{P}(x|z)\,\mathrm{P}(y|z)$ in discrete

probability. Indeed our definition coincides with the usual conditional independence, as explained below.

Example 6.7. In $\mathcal{K}\ell(\mathcal{D})$, let $c_{X|Z} \colon Z \to \mathcal{D}(X)$, $c_{Y|Z} \colon Z \to \mathcal{D}(Y)$, $c_{XY|Z} \colon Z \to \mathcal{D}(X \times Y)$ be disintegrations of some joint state, say $\omega \in \mathcal{D}(X \times Y \times Z)$. Let $\omega_Z \in \mathcal{D}(Z)$ be the marginal on Z. Then $X \perp\!\!\!\perp Y \mid Z$ if and only if

$$c_{XY|Z}(z)(x, y) = c_{X|Z}(z)(x) \cdot c_{Y|Z}(z)(y) \quad \text{whenever} \quad \omega_Z(z) \neq 0$$

for all $x \in X$, $y \in Y$ and $z \in Z$. If we write $\mathrm{P}(x, y|z) = c_{XY|Z}(z)(x, y)$, $\mathrm{P}(x|z) = c_{X|Z}(z)(x)$, $\mathrm{P}(y|z) = c_{Y|Z}(z)(y)$, and $\mathrm{P}(z) = \omega_Z(z)$, then the condition will look more familiar:

$$\mathrm{P}(x, y|z) = \mathrm{P}(x|z) \cdot \mathrm{P}(y|z) \quad \text{whenever} \quad \mathrm{P}(z) \neq 0 \ .$$

Example 6.8. Similarly, in $\mathcal{K}\ell(\mathcal{G})$, let $c_{X|Z} \colon Z \to \mathcal{G}(X)$, $c_{Y|Z} \colon Z \to \mathcal{G}(Y)$, $c_{XY|Z} \colon Z \to \mathcal{G}(X \times Y)$, and $\omega_Z \in \mathcal{G}(Z)$ be appropriate disintegrations and a marginal of some joint probability measure ω. Then $X \perp\!\!\!\perp Y \mid Z$ if and only if

$$c_{XY|Z}(z)(A \times B) = c_{X|Z}(z)(A) \cdot c_{Y|Z}(z)(B) \qquad \text{for } \omega_Z\text{-almost all } z \in Z$$

for all $A \in \Sigma_X$ and $B \in \Sigma_Y$.

The equivalences in the next result are well-known in conditional probability. Our contribution is that we formulate and prove them at an abstract, graphical level.

Proposition 6.9. The following are equivalent.

1 $X \perp\!\!\!\perp Y \mid Z$

2

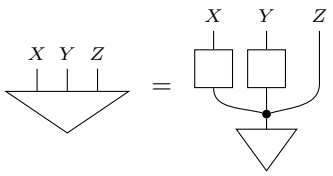

3

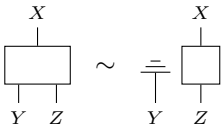

4

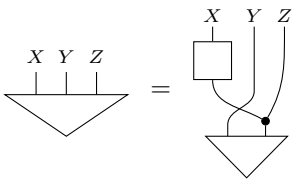

As we will see below, conditional independence $X \perp\!\!\!\perp Y \mid Z$ is symmetric in X and Y. Therefore the obvious symmetric counterparts of 3 and 4 are also equivalent to them.

Proof. By definition of almost equality, 1 is equivalent to

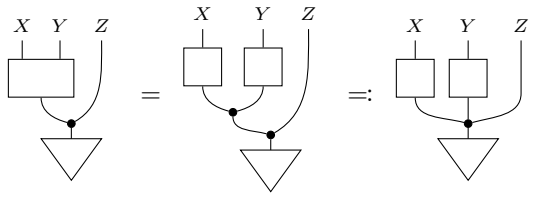

We then have $1 \Leftrightarrow 2$, since the identity below holds by the definition of disintegration.

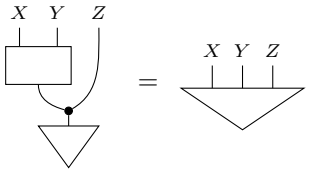

Next, assuming 3, we obtain 4 as follows.

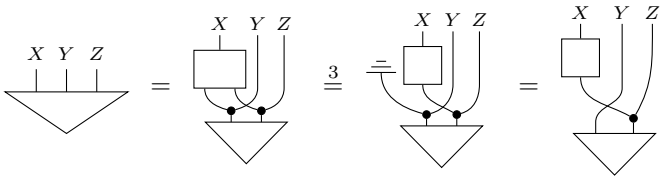

We prove $4 \Rightarrow 3$ similarly. Finally note that the following equation holds.

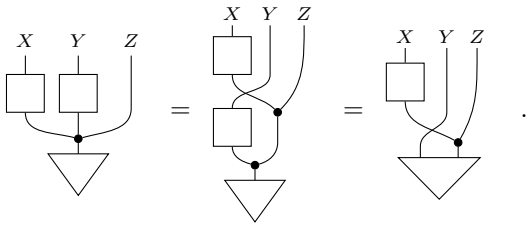

From this $2 \Leftrightarrow 4$ is immediate. □

Note that the condition 3 of the proposition is an analogue of $\mathrm{P}(x|y,z) = \mathrm{P}(x|z)$. The other conditions 2 and 4 say that the joint state can be factorised in certain ways, corresponding to the following equations:

$$\mathrm{P}(x,y,z) = \mathrm{P}(x|z)\,\mathrm{P}(y|z)\,\mathrm{P}(z) = \mathrm{P}(x|z)\,\mathrm{P}(y,z).$$

The proposition below shows that our abstract formulation of conditional independence does satisfy the basic 'rules' of conditional independence, which are known as *(semi-) graphoids axioms* (Verma and Pearl, 1988; Geiger et al., 1990).

Proposition 6.10. Conditional independence $(-) \perp\!\!\!\perp (-) \mid (-)$ satisfies:

1 (Symmetry) $X \perp\!\!\!\perp Y \mid Z$ if and only if $Y \perp\!\!\!\perp X \mid Z$.
2 (Decomposition) $X \perp\!\!\!\perp Y \otimes Z \mid W$ implies $X \perp\!\!\!\perp Y \mid W$ and $X \perp\!\!\!\perp Z \mid W$.
3 (Weak union) $X \perp\!\!\!\perp Y \otimes Z \mid W$ implies $X \perp\!\!\!\perp Y \mid Z \otimes W$.

4 (Contraction) $X \perp\!\!\!\perp Z \mid W$ and $X \perp\!\!\!\perp Y \mid Z \otimes W$ imply $X \perp\!\!\!\perp Y \otimes Z \mid W$.

Proof. We will freely use Proposition 6.9.
(1) Suppose $X \perp\!\!\!\perp Y \mid Z$. Then

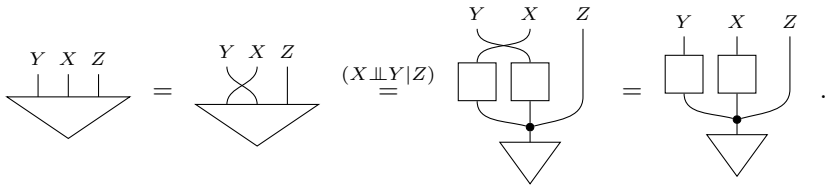

This means $Y \perp\!\!\!\perp X \mid Z$.
(2) Suppose $X \perp\!\!\!\perp Y \otimes Z \mid W$, namely:

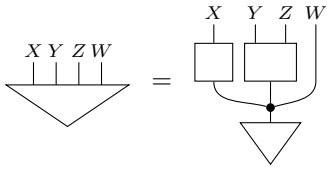

Marginalising Z, we obtain

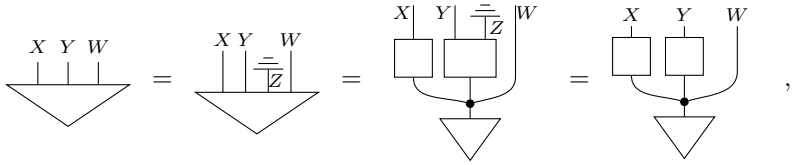

by Proposition 6.4.1. Thus $X \perp\!\!\!\perp Y \mid W$. Similarly we prove $X \perp\!\!\!\perp Z \mid W$.

Finally, we prove 3 and 4 at the same time. Note that $X \perp\!\!\!\perp Y \otimes Z \mid W$ implies $X \perp\!\!\!\perp Z \mid W$, as shown above. Therefore what we need to prove is that $X \perp\!\!\!\perp Y \otimes Z \mid W$ if and only if $X \perp\!\!\!\perp Y \mid Z \otimes W$, under $X \perp\!\!\!\perp Z \mid W$. Assume $X \perp\!\!\!\perp Z \mid W$, so we have

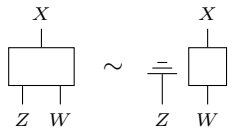

Then

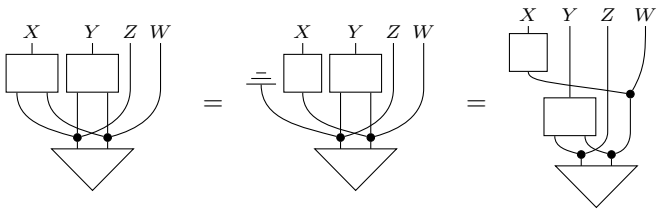

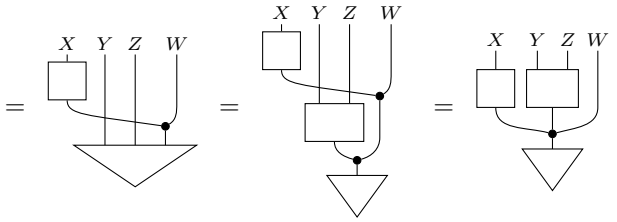

This proves $X \perp\!\!\!\perp Y \otimes Z \mid W$ if and only if $X \perp\!\!\!\perp Y \mid Z \otimes W$. □

The four properties from the graphoid axioms are essential in reasoning of conditional independence with DAGs or Bayesian networks (Verma and Pearl, 1988; Geiger et al., 1990).

Conditional independence has been studied categorically in (Simpson, 2017). There, a categorical notion of *(conditional) independence structure* is introduced, generalising algebraic axiomatisations of conditional independence, such as with graphoids and separoids (Dawid, 2001). We leave it to future work to precisely relate our approach to conditional probability to Simpson's categorical framework.

7. Beyond causal channels

All CD-categories $\mathbf{C}$ that we have considered so far are affine in the sense that all arrows $f\colon X \to Y$ are causal: $\overline{\top} \circ f = \overline{\top}$. We now drop the affineness, in order to enlarge our category to include 'non-causal' arrows, which enables us to have new notions such as scalars and effects. Essentially, we lose nothing by this change: all the arguments so far can still be applied to the subcategory $\mathrm{Caus}(\mathbf{C}) \subseteq \mathbf{C}$ containing all the objects and causal arrows. The category $\mathrm{Caus}(\mathbf{C})$ is an affine CD-category, inheriting the monoidal structure $(\otimes, I)$ and the comonoid structures $(\curlyvee, \overline{\top})$ from $\mathbf{C}$.

Recall that *channels* in $\mathbf{C}$ are causal arrows, *i.e.* arrows in $\mathrm{Caus}(\mathbf{C})$. *States* are channels of the form $\sigma\colon I \to X$. We call endomaps $I \to I$ on the tensor unit *scalars*. The set $\mathbf{C}(I, I)$ of scalars forms a monoid via the composition $s \cdot t = s \circ t$ and $1 = \mathrm{id}_I$. The monoid of scalars is always commutative — in fact, this is the case for any monoidal category, see *e.g.* (Abramsky and Coecke, 2009, §3.2). In string diagram scalars are written as $\langle s \rangle$ or simply as s. We can multiply scalars s to any arrows $f\colon X \to Y$ by the parallel composition, or diagrammatically by juxtaposition:

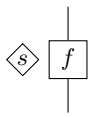

We call an arrow $\sigma\colon I \to X$ is *normalisable* if the scalar $\overline{\top} \circ \sigma\colon I \to I$ is (multiplicatively) invertible. In that case we can normalise σ into a proper state as follows.

$$\mathrm{nrm}(\sigma) \coloneqq$$

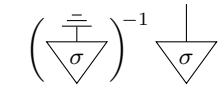

Effects in $\mathbf{C}$ are arrows of the form $p\colon X \to I$; they correspond to observables, with

predicates as special case. Diagrammatically they are written as on the left below.

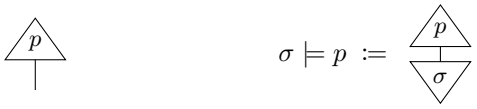

On the right the *validity* $\sigma \models p$ of a state $\sigma\colon I \to X$ and a effect $p\colon X \to I$ is defined. It is the scalar given by composition. Note that effects are not causal in general; by definition, only discarders $\overline{\top}$ are causal ones. States $\sigma\colon I \to X$ can be *conditioned* by effects $p\colon X \to I$ via normalisation, as follows.

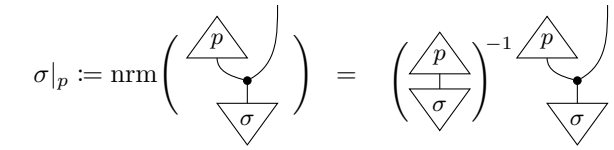

The conditional state $\sigma|_p$ is defined if the validity $\sigma \models p$ is invertible. Conditioning $\sigma|_p$ generalises conditional probability $\mathrm{P}(A|B)$ given *event* B, see Example 7.2 below. At the end of the section we will explain how conditioning and disintegration are related.

Recall, from Examples 2.4 and 2.5, that our previous examples $\mathcal{K}\ell(\mathcal{D})$ and $\mathcal{K}\ell(\mathcal{G})$ are both affine. We give two non-affine CD-categories that have $\mathcal{K}\ell(\mathcal{D})$ and $\mathcal{K}\ell(\mathcal{G})$ as subcategories, respectively.

Example 7.1. For discrete probability, we use *multisets* (or *unnormalised distributions*) over nonnegative real numbers $\mathbb{R}_{\geq 0} = [0, \infty)$, such as

$$1|x\rangle + 0.5|y\rangle + 3|z\rangle \qquad \text{on a set } X = \{x, y, z, \dots\}$$

We denote by $\mathcal{M}(X)$ the set of multisets over $\mathbb{R}_{\geq 0}$ on X. More formally:

$$\mathcal{M}(X) = \{\phi\colon X \to \mathbb{R}_{\geq 0} \mid \phi \text{ has finite support}\} \ .$$

It extends to a commutative monad $\mathcal{M}\colon \mathbf{Set} \to \mathbf{Set}$, see (Coumans and Jacobs, 2013). In a similar way to the distribution monad $\mathcal{D}$, we can check that the Kleisli category $\mathcal{K}\ell(\mathcal{M})$ is a CD-category. For a Kleisli map $f\colon X \to \mathcal{M}(Y)$, causality $\overline{\top} \circ f = \overline{\top}$ amounts to the condition $\sum_y f(x)(y) = 1$ for all $x \in X$. It is thus easy to see that $\mathrm{Caus}(\mathcal{K}\ell(\mathcal{M})) \cong \mathcal{K}\ell(\mathcal{D})$. In fact, the distribution monad $\mathcal{D}$ can be obtained from $\mathcal{M}$ as its *affine submonad*, see (Jacobs, 2017).

An effect $p\colon X \to 1$ in $\mathcal{K}\ell(\mathcal{M})$ is a function $p\colon X \to \mathbb{R}_{\geq 0}$. Its validity $\sigma \models p$ in a state ω is given by the expected value $\sum_x \sigma(x) \cdot p(x)$. The state $\sigma|_p$ updated with 'evidence' p is defined as $\sigma|_p(x) = \frac{\sigma(x) \cdot p(x)}{\sigma \models p}$.

Example 7.2. For general, measure-theoretic probability, we use *s-finite kernels* between measurable spaces (Kallenberg, 2017; Staton, 2017). Let X and Y be measurable spaces. A function $f\colon X \times \Sigma_Y \to [0, \infty]$ is called a *kernel* from X to Y if

— $f(x, -)\colon \Sigma_Y \to [0, \infty]$ is a measure for each X; and
— $f(-, B)\colon X \to [0, \infty]$ is measurable[§] for each $B \in \Sigma_Y$.

[§] The σ-algebra on $[0, \infty]$ is the standard one generated by $\{\infty\}$ and measurable subsets of $\mathbb{R}_{\geq 0}$.

We write $f\colon X \rightsquigarrow Y$ when f is a kernel from X to Y. A *probability kernel* is a kernel $f\colon X \rightsquigarrow Y$ with $f(x, Y) = 1$ for all $x \in X$. A kernel $f\colon X \rightsquigarrow Y$ is *finite* if there exists $r \in [0, \infty)$ such that for all $x \in X$, $f(x, Y) \leq r$. (Note that it must be 'uniformly' finite.) A kernel $f\colon X \rightsquigarrow Y$ is *s-finite* if $f = \sum_n f_n$ for some countable family $(f_n\colon X \to Y)_{n\in\mathbb{N}}$ of finite kernels.

For two s-finite kernels $f\colon X \rightsquigarrow Y$ and $g\colon Y \rightsquigarrow Z$, we define the (sequential) composite $g \circ f\colon X \rightsquigarrow Z$ by

$$(g \circ f)(x, C) = \int_Y g(y, C)\, f(x, \mathrm{d}y)$$

for $x \in X$ and $C \in \Sigma_Z$. There are identity kernels $\eta_X\colon X \rightsquigarrow X$ given by $\eta_X(x, A) = \mathbf{1}_A(x)$. With these data, measurable spaces and s-finite kernels form a category, which we denote by **sfKrn**. There is a monoidal structure on **sfKrn**. For measurable spaces X, Y we define the tensor product $X \otimes Y = X \times Y$ to be the cartesian product of measurable spaces. The tensor unit $I = 1$ is the singleton space. For s-finite kernels $f\colon X \rightsquigarrow Y$ and $g\colon Z \rightsquigarrow W$, we define $f \otimes g\colon X \times Z \rightsquigarrow Y \times W$ by

$$\begin{aligned}(f \otimes g)((x, z), E) &= \int_Y \Big(\int_W \mathbf{1}_E(y, w)\, g(z, \mathrm{d}w) \Big) f(x, \mathrm{d}y) \\ &= \int_W \Big(\int_Y \mathbf{1}_E(y, w)\, f(x, \mathrm{d}y) \Big) g(z, \mathrm{d}w)\end{aligned}$$

for $x \in X, z \in Z, E \in \Sigma_{Y\times W}$. The latter equality holds by the Fubini-Tonelli theorem for s-finite measures. These make the category **sfKrn** symmetric monoidal. Finally, for each measurable space X there is a 'copier' $\curlyvee\colon X \rightsquigarrow X \times X$ and a 'discarder' $\dot{\top}\colon X \rightsquigarrow 1$, given by $\curlyvee(x, E) = \mathbf{1}_E(x, x)$ and $\dot{\top}(x, 1) = 1$, so that **sfKrn** is a CD-category. For more technical details we refer to (Kallenberg, 2017; Staton, 2017).

Note that an s-finite kernel $f\colon X \rightsquigarrow Y$ is causal if and only if it is a probability kernel, which is nothing but a Kleisli map $X \to \mathcal{G}(Y)$ for the Giry monad. Therefore the causal subcategory of **sfKrn** is the Kleisli category of the Giry monad: $\mathrm{Caus}(\mathbf{sfKrn}) \cong \mathcal{K}\ell(\mathcal{G})$. In particular, states in **sfKrn** are probability measures $\sigma \in \mathcal{G}(X)$.

An effect $p\colon X \rightsquigarrow 1$ in **sfKrn**, *i.e.* an s-finite kernel $p\colon X \times \Sigma_1 \to [0, \infty]$, can be identified with a measurable function $p\colon X \to [0, \infty]$. The validity $\sigma \models p$ is then the integral $\int_X p(x)\, \sigma(\mathrm{d}x)$, defined in $[0, \infty]$. The conditional state $\sigma|_p \in \mathcal{G}(X)$ is defined by:

$$\sigma|_p(A) = \frac{\int_A p(x)\, \sigma(\mathrm{d}x)}{\sigma \models p}$$

for $A \in \Sigma_X$, when the validity $\sigma \models p$ is neither 0 nor ∞. In particular, for any 'event' $B \in \Sigma_X$, the obvious indicator function $\mathbf{1}_B\colon X \to [0, \infty]$ is an effect. Then conditioning yields a new state on X:

$$\sigma|_{\mathbf{1}_B}(A) = \frac{\int_A \mathbf{1}_B(x)\, \sigma(\mathrm{d}x)}{\sigma \models \mathbf{1}_B} = \frac{\sigma(A \cap B)}{\sigma(B)} \quad \text{for } A \in \Sigma_X.$$

This amounts to conditional probability $\mathrm{P}(A|B) = \mathrm{P}(A, B)/\mathrm{P}(B)$ given event B.

We need to generalise definitions from the previous sections in the non-affine/causal setting. Here we make only a minimal generalisation that is required in the next section.

For example, we still restrict ourselves to disintegration of (causal) states, although in the literature, the notion of disintegration exists also for non-probability (i.e. non-causal) measures and even for non-finite measures, see e.g. (Chang and Pollard, 1997, Definition 1). We leave such generalisations to future work.

Let $\sigma\colon I \to X$ be a state. We define σ-almost equality $f \overset{\sigma}{\sim} g$ between arbitrary arrows $f, g\colon X \to Y$ in the same way as Definition 5.1. Similarly, strong σ-almost equality between arbitrary arrows is defined as in Definition 5.7 (here ω still ranges over states).

Disintegrations of a joint state $\omega\colon I \to X \otimes Y$ are defined as in Definition 3.5, except that an arrow $c_1\colon X \to Y$ (or $c_2\colon Y \to X$) is not necessarily causal. Nevertheless, disintegrations of a state are *almost causal* in the following sense.

Definition 7.3. Let $\sigma\colon I \to X$ be a state. We say that an arrow $c\colon X \to Y$ is *σ-almost causal* if $\overline{\top} \circ c \overset{\sigma}{\sim} \overline{\top}$.

Then the following is immediate from the definition.

Proposition 7.4. Let $c_1\colon X \to Y$ be a disintegration of a state $\omega\colon I \to X \otimes Y$. Then c_1 is ω_1-almost causal, where $\omega_1\colon I \to X$ is the first marginal of ω. □

Example 7.5. In **sfKrn**, a s-finite kernel $f\colon X \rightsquigarrow Y$ is σ-almost causal if and only if $f(x, Y) = 1$ for σ-almost all $x \in X$. In that case, we can find a probability kernel $f'\colon X \rightsquigarrow Y$ such that $f \overset{\sigma}{\sim} f'$, by tweaking f in the obvious way. From this it follows that a disintegration of a joint probability measure $\omega \in \mathcal{G}(X \times Y)$ exists in **sfKrn** if and only if it exists in $\mathcal{K}\ell(\mathcal{G})$.

By Proposition 5.9, we can strengthen almost equality between channels in **sfKrn**. It is easy to see that equality strengthening is valid also for almost causal maps: if $f, g\colon X \rightsquigarrow Y$ are σ-almost causal maps with $f \overset{\sigma}{\sim} g$, then f, g are strongly σ-almost equal. (It is not clear whether equality strengthening is valid for arbitrary maps in **sfKrn**, but we will not need such a general result in this paper.)

In the remainder of the section, we present a basic relationship between disintegration and conditioning $\sigma|_p$, introduced above. We assume a joint state ω with its two disintegrations c_1 and c_2 in:

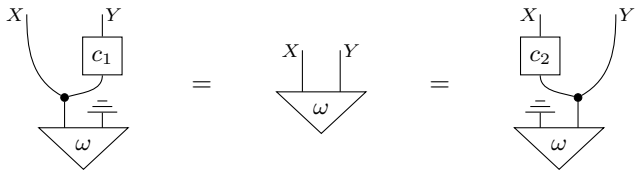

(8)

We write ω_1 and ω_2 for the first and second marginals of ω. (Thus the equations (8) above are the same as (3).)

Let q be an effect on Y. It can be extended to an effect $\mathbf{1} \otimes q$ on $X \otimes Y$, where:

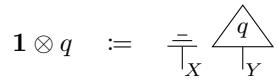

Then we can form the conditioned state $\omega|_{\mathbf{1}\otimes q}$. In a next step we take its first marginal,

written as $\big(\omega|_{\mathbf{1}\otimes q}\big)_1$. It turns out that, in general, this first marginal is different from the original first marginal ω_1, even though the effect q only applies to the second coordinate. This is called 'crossover influence' in (Jacobs and Zanasi, 2017). It happens when the state ω is 'entwined', that is, when its two coordinates are correlated.

A fundamental result in this context is that this crossover influence can also be captured via the channels c_1, c_2 that are extracted from ω via disintegrations. This works via effect transformation $c^*(p) \coloneqq p \circ c$ and state transformation $c_*(\sigma) \coloneqq c \circ \sigma$ along a channel.

Theorem 7.6. In the above setting, assuming that the relevant conditioned states exist, there are equalities of states:

$$\omega_1|_{c_1^*(q)} \;=\; \big(\omega|_{\mathbf{1}\otimes q}\big)_1 \;=\; \big(c_2\big)_*(\omega_2|_q). \tag{9}$$

Following (Jacobs and Zanasi, 2016) we can say that the expression on the left in (9) uses *backward* inference, and the one on the right uses *forward* inference.

Proof. We first note that the state in the middle of (9) is the first marginal of:

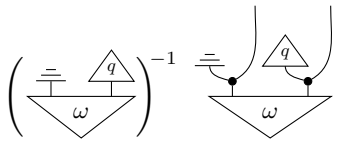

Hence:

$$\big(\omega|_{\mathbf{1}\otimes q}\big)_1 \quad = \quad \left(\begin{array}{c}\\ q\\ \omega\end{array}\right)^{-1} \begin{array}{c} q\\ \omega\end{array} \tag{10}$$

We note that the above scalar (that is inverted) can also be obtained as:

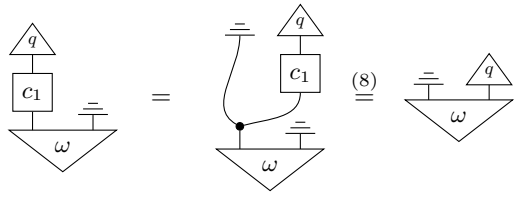

(11)

Hence we can prove the equation on the left in (9):

$$\omega_1|_{c_1^*(q)} \quad = \quad \left(\begin{array}{c} q\\ c_1\\ \omega\end{array}\right)^{-1} \begin{array}{c} q\\ c_1\\ \omega\end{array} \quad \overset{(8),(11)}{=} \quad \left(\begin{array}{c}\\ q\\ \omega\end{array}\right)^{-1} \begin{array}{c} q\\ \omega\end{array}$$

In a similar way we prove the equation on the right in (9), since $\big(c_2\big)_*(\omega_2|_q)$ equals:

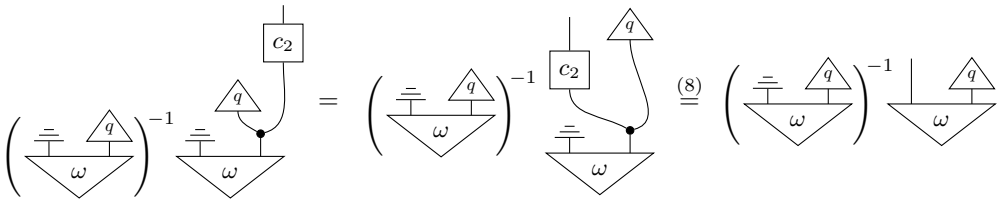

□

The two equations in Theorem 9 will be illustrated in the 'disease and mood' example below, where a particular state (probability) will be calculated in three different ways. For a gentle introduction to Bayesian inference with channels, and many more examples, we refer to (Jacobs and Zanasi, 2018).

Remark 7.7. Another way of relating conditioning to disintegration / Bayesian inversion was pointed out by a reviewer of the paper: for a state $\sigma\colon I \to X$ and an effect $p\colon X \to I$, a conditional state $\sigma|_p\colon I \to X$ is a Bayesian inversion for σ along p. This, however, involves disintegration of non-causal maps, which we leave to future work.

Disease and mood example

We describe an example of probabilistic (Bayesian) reasoning. The setting is the following. We consider a joint state about the occurrence and non-occurrence of a disease, written as a two-element set $D = \{d, d^\perp\}$, jointly with the occurrence and non-occurrence of a good mood, written as $M = \{m, m^\perp\}$, where $m^\perp$ stands for a bad (not good) mood. The joint distribution, called $\omega \in \mathcal{D}(M \times D)$, that we start from is of the form:

$$\omega \;=\; 0.05|m, d\rangle + 0.4|m, d^\perp\rangle + 0.5|m^\perp, d\rangle + 0.05|m^\perp, d^\perp\rangle. \tag{12}$$

Suppose there is a test for the disease, which is positive in 90% of all cases of people having the disease, and still 5% positive for people without the disease. Suppose the disease comes out positive. What is then the mood? It is expected that the mood will deteriorate, since the disease and the mood are 'entwined' (correlated) in the above joint state (12): a high likelihood of disease corresponds to a low mood.

The sensitivity is expressed as a channel $s\colon D \to 2$, where $2 = \{t, f\}$, with:

$$c(d) \;=\; \tfrac{9}{10}|t\rangle + \tfrac{1}{10}|f\rangle \qquad c(d^\perp) \;=\; \tfrac{1}{20}|t\rangle + \tfrac{19}{20}|f\rangle.$$

We calculate the first and second marginal of ω:

$$\omega_1 \coloneqq \omega.[1,0] \;=\; 0.45|m\rangle + 0.55|m^\perp\rangle \qquad \omega_2 \coloneqq \omega.[0,1] \;=\; 0.55|d\rangle + 0.45|d^\perp\rangle.$$

The a priori probability of a positive test is obtained by state transformation, applied to the second marginal:

$$s_*\big(\omega_2\big) \;=\; 0.518|t\rangle + 0.482|f\rangle.$$

As explained above, we are interested in the mood after a positive test. We write tt

for the predicate on $2 = \{t, f\}$ given by $tt(t) = 1$ and $tt(f) = 0$. We can transform it to a predicate $s^*(tt)$ on $D = \{d, d^\perp\}$, namely $s^*(tt)(d) = \frac{9}{10}$ and $s^*(tt)(d^\perp) = \frac{1}{20}$. We now describe three equivalent ways to calculate the posterior mood, as in Theorem 7.6, with predicate $q = s^*(tt)$.

1 First we use the truth predicate $\mathbf{1}$ on $M = \{m, m^\perp\}$, which is always 1, to extend (weaken) the predicate $s^*(tt)$ on D to $\mathbf{1} \otimes s^*(tt)$ on $M \times D$. The latter predicate can be used to update the joint state $\omega \in \mathcal{D}(M \times D)$. If we then take the first marginal we obtain the updated mood probability:

$$\big(\omega\big|_{\mathbf{1}\otimes s^*(tt)}\big).[1,0] \;=\; 0.126|m\rangle + 0.874|m^\perp\rangle.$$

Clearly, a positive test leads to a lower mood: a reduction from 0.45 to 0.126.

2 Next we extract channels $c_1 \colon M \to D$ and $c_2 \colon D \to M$ from ω via disintegration as in (8). For instance, c_1 is defined as:

$$\begin{aligned} c_1(m) \;&=\; \tfrac{\omega(m,d)}{\omega_1(m)}|d\rangle + \tfrac{\omega(m,d^\perp)}{\omega_1(m)}|d^\perp\rangle \;=\; \tfrac{1}{9}|d\rangle + \tfrac{8}{9}|d^\perp\rangle \\ c_1(m^\perp) \;&=\; \tfrac{\omega(m^\perp,d)}{\omega_1(m^\perp)}|d\rangle + \tfrac{\omega(m^\perp,d^\perp)}{\omega_1(m^\perp)}|d^\perp\rangle \;=\; \tfrac{10}{11}|d\rangle + \tfrac{1}{11}|d^\perp\rangle. \end{aligned}$$

We can now transform the predicate $s^*(tt)$ on D along c_1 to get the predicate $c_1^*(s^*(tt)) = (s_1 \mathbin{⊙} c_1)^*(tt))$ on M given by: $d \mapsto 13/90$ and $d^\perp \mapsto 181/220$. We can now use this predicate to update the a priori 'mood' marginal $\omega_1 \in \mathcal{D}(M)$. This gives the same a posteriori outcome as before:

$$\omega_1\big|_{c_1^*(s^*(tt))} \;=\; 0.126|m\rangle + 0.874|m^\perp\rangle.$$

3 We can also update the second 'disease' marginal $\omega_2 \in \mathcal{D}(D)$ directly with the predicate $s^*(tt)$ and then do state transformation along the channel $c_2 \colon D \to M$. This gives the same updated mood:

$$(c_2)_*\big(\omega_2|_{s^*(tt)}\big) \;=\; 0.126|m\rangle + 0.874|m^\perp\rangle.$$

8. Disintegration via likelihoods

We continue in the setting of Section 7 in a CD-category that is not necessarily affine. The goal of this section is to present Theorem 8.3, which generalises a construction of Bayesian inversions using densities/likelihoods shown in Example 3.9.

We first introduce 'likelihoods' in our setting.

Definition 8.1. We say a channel $c \colon X \to Y$ is represented by a effect ℓ on $X \otimes Y$ with respect to an arrow $\nu \colon I \to Y$ if

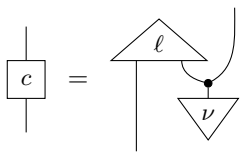

(13)

We call ℓ a *likelihood relation* for the channel c with respect to ν.

Interpreted in the category **sfKrn**, the definition says: a kernel $c\colon X \rightsquigarrow Y$ satisfies

$$c(x, B) = \int_B \ell(x, y)\, \nu(\mathrm{d}y)$$

for a kernel $\ell\colon X \times Y \rightsquigarrow 1$ (identified with a measurable function $\ell\colon X \times Y \to [0, \infty]$) and a measure $\nu\colon 1 \rightsquigarrow Y$. This is basically the same as what we have in Example 3.9, but here ν is not necessarily the Lebesgue measure.

Definition 8.2. Let $\sigma\colon I \to X$ be a state. An effect $p\colon X \to I$ is *σ-almost invertible* if there is an effect $q\colon X \to I$ such that:

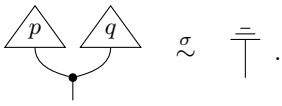

The definition allows us to normalise an arrow $f\colon X \to Y$ into an almost causal one, as follows. If an effect $\overline{\top} \circ f$ is σ-almost inverted by q, as on the left below,

then clearly the arrow $X \to Y$ on the right is σ-almost causal.

We can now formulate and prove our main technical result.

Theorem 8.3. Let σ be a state on X, and $c\colon X \to Y$ be a channel represented by a likelihood relation ℓ with respect to ν as in (13) above. Assume that the category admits equality strengthening for almost causal maps, and that the effect

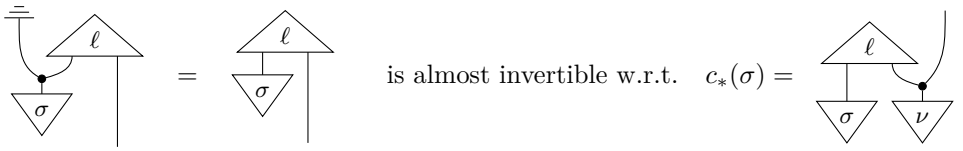

is almost invertible w.r.t. $c_*(\sigma) =$.

Then, writing $q\colon Y \to I$ for an almost inverse to the effect, the channel

$$d\colon Y \to X \;:=$$

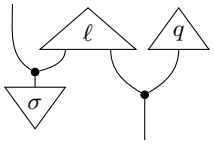

is a Bayesian inversion for σ along $c\colon X \to Y$. Namely, together they satisfy the equation (5).

Proof. We reason as follows.

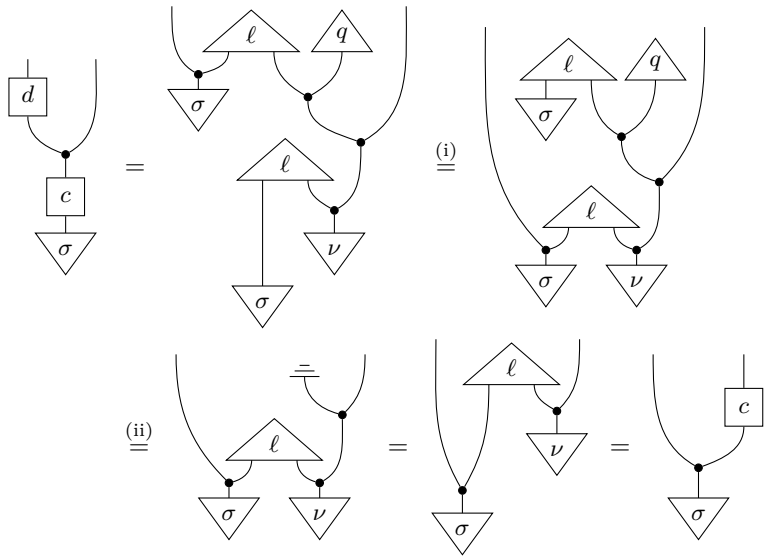

For the equality $\overset{\text{(i)}}{=}$ we use associativity and commutativity of copiers $\curlyvee$. The equality $\overset{\text{(ii)}}{=}$ follows by:

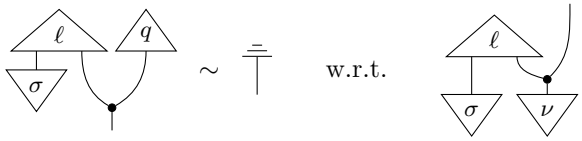

via equality strengthening. □

Example 8.4. We instantiate the Theorem 8.3 in **sfKrn**. Let $c\colon X \rightsquigarrow Y$ be a probability kernel represented by a likelihood relation $\ell\colon X \times Y \rightsquigarrow 1$ with respect to $\nu\colon 1 \rightsquigarrow Y$. The relation ℓ is identified with a measurable function $\ell\colon X \times Y \to [0,\infty]$ and ν with a measure $\nu\colon \Sigma_Y \to [0,\infty]$. The equation (13) amounts to

$$c(x,B) = \int_B \ell(x,y)\,\nu(\mathrm{d}y) \quad .$$

In particular, each $\ell(x,-)$ is a probability density function, satisfying $\int_Y \ell(x,y)\,\nu(\mathrm{d}y) = 1$. Typically, we use the Lebesgue measure as ν, with Y a subspace of $\mathbb{R}$. Let $\sigma\colon 1 \rightsquigarrow X$ be a probability measure. Then $c_*(\sigma)\colon 1 \rightsquigarrow Y$ is given as:

$$c_*(\sigma)(B) = \int_X c(x,B)\,\sigma(\mathrm{d}x) = \int_X \int_B \ell(x,y)\,\nu(\mathrm{d}y)\,\sigma(\mathrm{d}x)$$

The effect

$$p\colon Y \rightsquigarrow 1 \quad = \quad \text{ 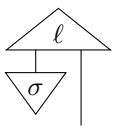 } \quad \text{is given as:} \qquad p(y) = \int_X \ell(x,y)\,\sigma(\mathrm{d}x) \quad .$$

To define an inverse of p, we claim that $0 < p < \infty$, $c_*(\sigma)$-almost everywhere. We prove

that $p^{-1}(\{0,\infty\}) = p^{-1}(0) \cup p^{-1}(\infty)$ is $c_*(\sigma)$-negligible, as:

$$\begin{aligned} c_*(\sigma)\big(p^{-1}(0)\big) &= \int_X \int_{p^{-1}(0)} \ell(x,y)\,\nu(\mathrm{d}y)\,\sigma(\mathrm{d}x) \\ &= \int_{p^{-1}(0)} p(x)\,\nu(\mathrm{d}y) \\ &= \int_{p^{-1}(0)} 0\,\nu(\mathrm{d}y) = 0 \end{aligned}$$

and, similarly we have

$$\int_{p^{-1}(\infty)} \infty\,\nu(\mathrm{d}y) = \int_{p^{-1}(\infty)} p(x)\,\nu(\mathrm{d}y) = c_*(\sigma)\big(p^{-1}(\infty)\big) \leq 1$$

but this is possible only when $\nu(p^{-1}(\infty)) = 0$, hence $c_*(\sigma)\big(p^{-1}(\infty)\big) = \int_{p^{-1}(\infty)} p(x)\,\nu(\mathrm{d}y) = 0$. Now define an effect $q\colon Y \rightsquigarrow 1$ by

$$q(y) = \begin{cases} p(y)^{-1} & \text{if } 0 < p(y) < \infty \\ 0 & \text{otherwise.} \end{cases}$$

Then p is $c_*(\sigma)$-almost inverted by q. By Theorem 8.3, the Bayesian inversion for σ along c is given by

$$d\colon Y \to X \quad := \quad$$

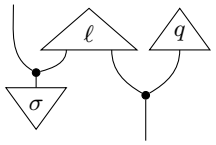

,

namely,

$$\begin{aligned} d(y,A) &= q(y) \int_A \ell(x,y)\,\sigma(\mathrm{d}x) \\ &= \frac{\int_A \ell(x,y)\,\sigma(\mathrm{d}x)}{\int_X \ell(x,y)\,\sigma(\mathrm{d}x)} \qquad \text{whenever} \quad 0 < \int_X \ell(x,y)\,\sigma(\mathrm{d}x) < \infty \end{aligned} \tag{14}$$

This may be seen as a variant of the Bayes formula. The calculation in Example 3.9 is reproduced when σ is also given via a density function.

In the end we reconsider the naive Bayesian classification example from Section 4. There we only considered the discrete version. The original source (Witten et al., 2011) also contains a 'hybrid' version, combining discrete and continuous probability.

In that hybrid form the Temperature and Humidity columns in Figure 1 are different, and are given by numerical values. What these values are does not matter too much here, since they are only used to calculate *mean* (μ) and *standard deviation* (σ) values. This is done separately for the cases where Play is *yes* or *no*. It results in the following table.

	Temperature	**Humidity**
Play = yes	$\mu = 73$, $\sigma = 6.2$	$\mu = 79.1$, $\sigma = 10.2$
Play = no	$\mu = 74.6$, $\sigma = 7.9$	$\mu = 86.2$, $\sigma = 9.7$

These μ, σ values are used to define two functions $c_T\colon P \to \mathcal{G}(\mathbb{R})$ and $c_H\colon P \to \mathcal{G}(\mathbb{R})$,

with normal distributions $\mathcal{N}$ as pdf. To be precise, these channels are defined as follows, where $A \subseteq \mathbb{R}$ is a measurable subset.

$$c_T(y)(A) = \int_A \mathcal{N}(73, 6.2)(x)\,\mathrm{d}x \qquad c_H(y)(A) = \int_A \mathcal{N}(79.1, 10.2)(x)\,\mathrm{d}x$$
$$c_T(n)(A) = \int_A \mathcal{N}(74.6, 7.9)(x)\,\mathrm{d}x \qquad c_H(y)(A) = \int_A \mathcal{N}(86.2, 9.7)(x)\,\mathrm{d}x.$$

These functions c_T and c_H form channels for the Giry monad $\mathcal{G}$.

The two columns Outlook and Windy in Figure 1 remain the same. The corresponding maps $d_O \colon P \to \mathcal{D}(O)$ and $d_W \colon P \to \mathcal{D}(W)$ for the discrete probability monad $\mathcal{D}$ — in Diagram (6) — are now written as maps $c_O \colon P \to \mathcal{G}(O)$ and $c_W \colon P \to \mathcal{G}(W)$ for the monad $\mathcal{G}$, using the obvious inclusion $\mathcal{D} \hookrightarrow \mathcal{G}$.

We now have four channels $c_O \colon P \to O$, $c_T \colon P \to \mathbb{R}$, $c_H \colon P \to \mathbb{R}$ and $c_W \colon P \to W$. We combine them, as before, into a single channel $c \colon P \to T \times \mathbb{R} \times \mathbb{R} \times W$. In combination with the Play marginal distribution π from Section 4, we can compute c's Bayesian inversion $f \colon T \times \mathbb{R} \times \mathbb{R} \times W \to \mathcal{G}(P)$. We apply it to the input data used in (Witten et al., 2011), and get the following Play distribution.

$$f(s, 66, 90, t) = 0.207|y\rangle + 0.793|n\rangle.$$

The latter inversion computation produces the probability of 0.207 for playing when the outlook is *Sunny*, the temperature is *66* (Fahrenheit), the humidity is *90%* and the windiness is *true*. The value computed in (Witten et al., 2011) is 20.8%. The minor difference of 0.001 with our outcome can be attributed to (intermediate) rounding errors.

Our computation is done via *EfProb*, using the formula (14).

References

Abramsky, S. and Coecke, B. (2009). Categorical quantum mechanics. In *Handbook of Quantum Logic and Quantum Structures: Quantum Logic*, pages 261–323. Elsevier.

Ackerman, N. L., Freer, C. E., and Roy, D. M. (2011). Noncomputable conditional distributions. In *Logic in Computer Science*, pages 107–116. IEEE.

Barber, D. (2012). *Bayesian Reasoning and Machine Learning*. Cambridge Univ. Press.

Bernardo, J. and Smith, A. (2000). *Bayesian Theory*. John Wiley & Sons.

Borgström, J., Gordon, A., Greenberg, M., Margetson, J., and Gael, J. V. (2013). Measure transformer semantics for Bayesian machine learning. *Logical Methods in Comp. Sci.*, 9(3):1–39.

Chang, J. T. and Pollard, D. (1997). Conditioning as disintegration. *Statistica Neerlandica*, 51(3):287–317.

Cho, K. and Jacobs, B. (2017). The EfProb library for probabilistic calculations. In *Conference on Algebra and Coalgebra in Computer Science*, volume 72 of *LIPIcs*. Schloss Dagstuhl.

Cho, K., Jacobs, B., Westerbaan, A., and Westerbaan, B. (2015). An introduction to effectus theory. Preprint. arXiv:1512.05813 [cs.LO].

Clerc, F., Danos, V., Dahlqvist, F., and Garnier, I. (2017). Pointless learning. In *Foundations of Software Science and Computation Structures*, volume 10203 of *Lect. Notes Comp. Sci.*, pages 355–369. Springer.

Coecke, B. (2016). Terminality implies no-signalling ...and much more than that. *New Generation Computing*, 34(1):69–85.

Coecke, B. and Kissinger, A. (2017). *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press.

Coumans, D. and Jacobs, B. (2013). Scalars, monads and categories. In Heunen, C., Sadrzadeh, M., and Grefenstette, E., editors, *Quantum Physics and Linguistics. A Compositional, Diagrammatic Discourse*, pages 184–216. Oxford Univ. Press.

Culbertson, J. and Sturtz, K. (2014). A categorical foundation for Bayesian probability. *Applied Categorical Structures*, 22(4):647–662.

D'Ariano, G. M., Chiribella, G., and Perinotti, P. (2017). *Quantum Theory from First Principles: An Informational Approach*. Cambridge University Press.

Dawid, A. (2001). Separoids: A mathematical framework for conditional independence and irrelevance. *Annals of Mathematics and Artificial Intelligence*, 32(1):335–372.

Doberkat, E.-E. (2009). *Stochastic Coalgebraic Logic*. Monographs in Theoretical Computer Science. Springer.

Faden, A. M. (1985). The existence of regular conditional probabilities: Necessary and sufficient conditions. *The Annals of Probability*, 13(1):288–298.

Fong, B. (2012). Causal theories: A categorical perspective on Bayesian networks. Master's thesis, Univ. of Oxford. arXiv:1301.6201 [math.PR].

Fremlin, D. H. (2000). *Measure Theory (5 volumes)*. Torres Fremlin.

Geiger, D., Verma, T., and Pearl, J. (1990). Identifying independence in Bayesian networks. *Networks*, 20:507–534.

Giry, M. (1982). A categorical approach to probability theory. In *Categorical Aspects of Topology and Analysis*, volume 915 of *Lecture Notes in Mathematics*, pages 68–85. Springer.

Gordon, A., Henzinger, T., Nori, A., and Rajamani, S. (2014). Probabilistic programming. In *Future of Software Engineering*, pages 167–181. ACM.

Jacobs, B. (2015). New directions in categorical logic, for classical, probabilistic and quantum logic. *Logical Methods in Comp. Sci.*, 11(3):1–76.

Jacobs, B. (2017). From probability monads to commutative effectuses. *Journ. of Logical and Algebraic Methods in Programming*, 156.

Jacobs, B., Westerbaan, B., and Westerbaan, A. (2015). States of convex sets. In *Foundations of Software Science and Computation Structures*, volume 9034 of *Lect. Notes Comp. Sci.*, pages 87–101. Springer.

Jacobs, B. and Zanasi, F. (2016). A predicate/state transformer semantics for Bayesian learning. In *Math. Found. of Programming Semantics*, volume 325 of *Elect. Notes in Theor. Comp. Sci.*, pages 185–200. Elsevier.

Jacobs, B. and Zanasi, F. (2017). A formal semantics of influence in Bayesian reasoning. In *Math. Found. of Computer Science*, volume 83 of *LIPIcs*. Schloss Dagstuhl.

Jacobs, B. and Zanasi, F. (2018). The logical essentials of Bayesian reasoning. See `arxiv.org/abs/1804.01193`.

Joyal, A. and Street, R. (1991). The geometry of tensor calculus, I. *Advances in Mathematics*, 88(1):55–112.

Kallenberg, O. (2017). *Random Measures, Theory and Applications*. Springer.

Katoen, J.-P., Gretz, F., Jansen, N., Lucien Kaminski, B., and Olmedo, F. (2015). Understanding probabilistic programs. In *Correct System Design*, volume 9360 of *Lect. Notes Comp. Sci.*, pages 15–32. Springer.

Mac Lane, S. (1998). *Categories for the Working Mathematician*. Springer, second edition.

Pachl, J. K. (1978). Disintegration and compact measures. *Mathematica Scandinavica*, 43:157–168.

Panangaden, P. (2009). *Labelled Markov Processes*. Imperial College Press.

Pawitan, Y. (2001). *In All Likelihood. Statistical Modelling and Inference Using Likelihood*. Clarendon Press.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.

Pollard, D. (2002). *A User's Guide to Measure Theoretic Probability*. Cambridge University Press.

Selinger, P. (2010). A survey of graphical languages for monoidal categories. In *New Structures for Physics*, volume 813 of *Lecture Notes in Physics*, pages 289–355. Springer.

Shan, C.-c. and Ramsey, N. (2017). Exact Bayesian inference by symbolic disintegration. In *Princ. of Programming Languages*, pages 130–144. ACM.

Simpson, A. (2017). Category-theoretic structure for independence and conditional independence. In *Math. Found. of Programming Semantics*. To appear.

Staton, S. (2017). Commutative semantics for probabilistic programming. In *European Symp. on Programming*, volume 10201 of *Lect. Notes Comp. Sci.*, pages 855–879. Springer.

Staton, S., Yang, H., Heunen, C., Kammar, O., and Wood, F. (2016). Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *Logic in Computer Science*, pages 525–534. ACM.

Stoyanov, J. M. (2014). *Counterexamples in Probability*. Dover, third edition.

Verma, T. and Pearl, J. (1988). Causal networks: Semantics and expressiveness. In *Uncertainty in Artificial Intelligence*.

Witten, I., Frank, E., and Hall, M. (2011). *Data Mining – Practical Machine Learning Tools and Techniques*. Elsevier, Amsterdam.